1-30-2017

Analytic Methods for Geometric Modeling

Morad Behandish
Graduate Student, m.behandish@uconn.edu

Analytic Methods for Geometric Modeling

Morad Behandish, M.Sc.

University of Connecticut, 2017

Current practice in mechanical product description and computational processing is rooted in the theory of solid modeling whose principles were established almost four decades ago. Despite providing a concrete foundation for computerizing the product development process at the time, the limited capacity of the traditional models to answer the growing design and manufacturing needs is only recently being recognized in the scientific and engineering communities.

I propose an alternative approach that employs spatial functions (i.e., 3D signals) instead of (or in addition to) topological pointsets in the Euclidean 3−space to abstract physical objects and their various geometric, physical, and material properties. The central benefit of this new approach is its great promise and potential extensibility in response to the emerging needs for more meaningful form-function correlations (e.g., for conceptual design) and more powerful tools for modeling new materials (e.g., knitted composites) and processes (e.g., additive manufacturing). This requires a paradigm shift from 'explicit' descriptions equipped with 'combinatorial' methods—e.g., combinatorial intersection test between objects modeled as r-sets, approximated by sphere trees, and tested using set-theoretic operations that exploit efficient tree traversal—to 'implicit' descriptions equipped with 'analytic' methods—e.g., analytic intersection test between objects modeled as density functions, approximated by frequency domain samples, and tested using measure-theoretic operations that are streamlined via fast Fourier transforms (FFT).

The results obtained so far suggest that the proposed approach is a powerful unifying alternative to the conventional approaches to geometric computing and overcomes some of the key shortcomings of the traditional models. It opens up new promising theoretical and computational directions for future research in the nascent but emerging field of analytic solid geometry. I hope that the early-stage results in this thesis will inspire other researchers to develop more rigorous formal models and that it will eventually encourage broad industrial adoption of the analytic techniques into future generations of the product life-cycle management (PLM) software.

Analytic Methods for Geometric Modeling

Morad Behandish

M.Sc., University of Connecticut, 2016

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

At the

University of Connecticut

2017

APPROVAL PAGE

Doctor of Philosophy Dissertation

Analytic Methods for Geometric Modeling

Presented by
Morad Behandish, M.Sc.,

Major Advisor _____
Horea T. Ilieş

Associate Advisor _____
Kazem Kazerouninan

Associate Advisor _____
Sanguthevar Rajasekaran

Associate Advisor _____
Donald R. Sheehy

Associate Advisor _____
Julian A. Norato

University of Connecticut
2017

Acknowledgements

I am very grateful to my advisor Horea T. Ilieş, Professor of Mechanical Engineering at UConn, for years of selfless support and patient mentorship besides insightful research advisorship. Horea has been more than a great academic advisor by providing a truly exceptional example of *character* on top of academic and professional excellence.

Special Thanks to Vadim Shapiro, Professor of Mechanical Engineering at UW-Madison, my advisor's Ph.D. advisor, and my post-doc supervisor at ICSI, for all the stimulating and instructive discussions within the past year. I have already learned a lot (and have a lot more to learn) from him in a quest to "make an impact."[1]

I would like to thank my associate advisor Kazem Kazerounian, Dean of Engineering and Professor of Mechanical Engineering at UConn, for his continuous support of my academic and professional development throughout my years at UConn.

I would like to thank my associate advisor Sanguthevar Rajasekaran, Director of BECAT and Professor of Computer Science and Engineering at UConn, for introducing me to the realm of computer algorithms—an exceptional classroom experience.

Special thanks to my associate advisor Donald R. Sheehy, Assistant Professor of Computer Science at UConn, for the long and stimulating (and often unplanned) discussions that we have had in his office. Those discussions greatly helped me wrap my head around the ideas that were central to the message of this thesis. Don has also been immensely helpful in pointing me to the right directions as I have been trying to educate myself in computational geometry and topology.

I am also grateful to my associate advisor Julian A. Norato, Professor of Mechanical Engineering at UConn for his valuable feedback on the thesis and presentation.

Thanks to Saigopal Nelaturi, Research Scientist at PARC for numerous discussions on configuration modeling. Sai's Ph.D. research work at UW-Madison and the follow-up applications to manufacturing process planning at PARC have inspired me a lot as to the intellectual and commercial promise of the developed methods in this thesis.

Thanks to Mikola Lysenko, Software Engineer (and a truly exceptional one) at

[1] I am also thankful to Vadim for arranging the post-doc position, and to Horea for his recommendation, that spared me the additional stress of applying for jobs while writing this thesis!

Flexport for instructive discussions on measure-theoretic interpretations of geometric operations, and for giving me his 'blessings' to push these ideas forward following his decision to change career directions.

It has been a great pleasure working at the Computational Design Lab (CDL) at UConn where I had the chance to interact with bright individuals including Hüseyin Erdim, Ata Eftekharian, Radu Corcodel, Denis Dorozhkin, Frol Periverzov, Pouya Tavousi, Reed Williams, and Mariangela Trotta.

I am especially grateful to my UConn Outing Club comrades, particularly Jonathan Le May to whom and whose free spirit I owe some of the most adventurous caving, climbing, rafting, hiking, and camping experiences that kept me sane throughout this journey, and also to Ali Phillips, Derek Tata, Calvin Ng, and many others.

A special tip of the hat goes to Kamyar Momeni, an exceptional (ex-)housemate for about a half decade and a true friend for life.

Thanks to the dearest Raana Kashfi for being such an amazing companion especially throughout the past few stressful months.

Most importantly, I would like to thank my parents, to whom this thesis is dedicated. My late father Morteza Behandish has been and continues to be—more than a decade after his passing—my number one role-model having taught me principles I aspire to maintain throughout my life. My beloved mother Nasrin Tabrizi has continuously supported and trusted me during every step of my personal development. She deserves the biggest gratitude for embracing my decision of pursuing graduate education in the US, which meant a (seemingly endless) separation from his only son over thousands of miles for several years.

Finally, I express my appreciation for support from the National Science Foundation (NSF),[2] without which this research would not have been remotely possible.

Disclaimer: The responsibility for any errors and omissions lies solely with the author. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the NSF.

The author would also like to emphasize that the identification, exemplification, or choice of any commercial product or trade name throughout this thesis does not imply endorsement or recommendation.

Morad Behandish
June 2016
Storrs, CT, USA

[2]Grant numbers CMMI-1200089, CMMI-0927105, and CNS-0927105.

Contents

Preface

> "The future of model-based engineering hinges on the ability to support the ever increasing diversity of abstractions [of physical artifacts] in a systematic and computationally efficient manner."
>
> Vadim Shapiro, 2015 [319]

Research in geometric modeling, computing, and reasoning has matured significantly over the past four or five decades alongside the widespread application of computers in product design and development, with applications in various scientific, engineering, architectural, biological, consumer, and other domains. The applied field of geometric modeling, as practiced today, extensively overlaps with and draws upon multiple areas of mechanical engineering, computer science, and applied mathematics—including linkage kinematics [140, 293], multibody dynamics [14, 309], computational geometry [89, 91, 228, 259], computational topology [107], graphics [56, 158], algorithms [80, 151], and mathematical analysis [15].

This Ph.D. dissertation is specifically concerned with a set of applied methods in computer-aided geometric design (CAGD) [115] with a particular interest in consolidating solid and configuration modeling in a measure-theoretic context, and providing a set of power tools for uniformly formulating and solving shape and motion related problems—collectively referred to as 'analytic methods' throughout this thesis.

In its original form, solid modeling [275, 277, 278, 351, 352] was conceived in the 1970s and 80s as a consistent set of principles for mathematical modeling [275] and computer representation [278] of 3D solid objects [316]. In spite of its huge success, as a result of its emphasis on informational completeness, physical fidelity, and universality [316], in implementing a wide class of shape-related computations,[3] a number of critical challenges remained to be addressed, for example, in approaching several motion-related problems.[4] These difficulties could be traced in part to the lack of modeling techniques, as complete and universal as those available for general shapes, to reason about problems involving general motions—except for simplified cases, e.g., one-parametric continuous motions (i.e., sweep/unsweep) [161–163]—in spite of their ubiquity in most mechanical systems and applications.

More recently, configuration modeling [231, 303] was introduced to CAGD applications in order to solve the aforementioned problems by taking advantage of the kinematic descriptions developed in 1980s and 90s for motion planning in robotics [199, 210]. In a similar fashion to modeling shapes as topological pointsets in the Euclidean 3D space, configurations were also viewed as topological pointsets in the generally non-Euclidean and higher-dimensional spaces (e.g., the 6D motion group SE(3) for general rigid motions). Thus a variety of direct and inverse problems prevalent in design and manufacturing[5] were uniformly formulated in terms of interactions of shapes and motions (e.g., using Minkowski products and quotients) [235,291]. How-

[3] e.g., finite sequences of material addition or removal (i.e., Boolean constructions), instantaneous interference tests (i.e., collision detection), local surface modifications (i.e., Euler operators), bulk property computations (i.e., surface or volume integrals), etc.

[4] e.g., continuous sweeps of material addition or removal (i.e., morphological dilation), maximal or minimal noninterference objects under continuous motion (i.e., morphological erosion), generalized offset constructions (i.e., Minkowski operations), etc.

[5] Many of these problems can be recast in terms of the operations exemplified in the previous footnote. The artificial distinction between direct and inverse problems can also be eliminated by observing the duality between them [231, 235].

ever, many theoretical and computational problems remain to be solved, including:

1. handling dimensional heterogeneity of spatial planning problems, e.g., when subsets of the configuration space accessible to a moving robot degenerates to a lower-dimensional sub-space, or when studying sweeps of a 3D shape along 1D or 2D parametric trajectories for fold regions and other singularities [111, 113];

2. developing efficient and flexible representation schemes [36] or reconciling with the existing and extensively studied/optimized schemes of traditional solid modeling [277, 278], to go beyond point cloud sampling strategies [236];

3. extending the methods to broader applications such as automatic or semi-automatic assembly [35, 37], protein docking [25, 69], bin packing [216, 217], and other scoring/optimization problems.

In this thesis I will try to address examples from all of the above problems, unified under one elegant theoretical framework.

This project started from an earlier attempt to systematically address a specific configuration space problem; namely, formalizing the notion of a 'good fit' (from a purely geometric standpoint) between pairs of rigid objects of arbitrary shape. Such an attempt was motivated by a specific application in mind; namely, quantifying and optimizing a generic measure of *shape complementarity*—an early version of which I reported in [33]—motivated by the importance and difficulty of protein docking [288].[6]

Throughout the course of my research on different approaches to related problems, I came across a crucial observation that would eventually form the conceptual

[6]In fact, when I presented my literature review as part of the Ph.D. Qualification Exams at UConn back in 2012, it was titled "Studies and Directions in Protein Docking." As with most open-ended research endeavors, a bigger picture started to crystalize over time due to observing relations between similar techniques from different areas and I ended up with a completely different title.

framework of this thesis. Comparing the critical ingredients of a range of recently emerging techniques—e.g., analytic interference and contact measures [214,217], generalized morphological operations [215,218], FFT-based path planning [88,174], and (primarily) FFT-based protein docking [25,60,61,69]—I noticed a common theme, to which I refer as 'analytic methods' in the project title and throughout the thesis.[7]

The common denominator to all of these methods was the reinterpretation of families of pointsets (for shapes and configurations alike) as sub-/super-level sets of real-valued functions over their subsequent spaces to open up the possibility of applying familiar concepts from mathematical analysis—particularly measure-theoretic operations as fundamental as integration, e.g., computing norms, inner products, etc.—to reformulate important geometric computations. For instance, one of the fundamental ingredients is the group-theoretic formalism of the invertible motions which allows interpreting their cumulative application to shapes as *convolutions*. Another example is the interpretation of intersections as inner products, and can also be recast as convolutions when formulated as a function of relative motions. Such formulations in turn allow for approximate computations 'on a budget' by using Fourier transforms, anti-aliasing, and low-pass filtering. The crucial observation is that geometric problems can be recast in the language of functions and can tremendously benefit from established ideas in digital signal processing (DSP)—noting that one in this case deals with signals over higher-dimensional domains that carry geometric information.

As I continued looking into other important developments such as fuzzy solid

[7]The credit for this denomination belongs to Mikola Lysenko and Vadim Shapiro in the more restricted context of FFT-based collision detection [214] to whom I am indebted for fruitful discussions on the measure-theoretic interpretation of morphological operators [215,218] in a more general context. However, my first exposure to the idea of convolution-based shape analysis was through the works of Chandrajit Bajaj's group on FFT-based protein docking implemented into $\mathsf{F}^2\mathsf{Dock}$ [25,69] later formalized for general rigid body correlations [23].

modeling [219,369], tolerant solid modeling [260–262], CAD data exchange [270,331], and the more general problems of geometric interoperability [147,148], I realized that the analytic approach also lends itself more naturally to some of the most successful proposals for solving problems in those areas, which led me to the following premise:

Central Insight of the Thesis

Geometric and spatial information often categorized arbitrarily into solids and configurations, shapes and motions, structures and materials, and other physical notions at different scales and application contexts, can be uniformly conceptualized in terms of spatial distribution fields (i.e., signals). From that perspective, geometric operations central to CAD can be re-conceptualized and studied in terms of an *algebra of functions*. I show that traditional geometric modeling operations are subsumed and extended by a strikingly simple algebraic structure, formed by basic functions—namely, *inclusion* and *distance* queries—endowed with familiar measure-theoretic operators—namely, *convolution* and *composition* with Dirac $\delta-$functions of various orders—thus can be computed and reasoned about using higher-dimensional generalizations of signal processing techniques.

Although there has been numerous scattered efforts to introduce measure-theoretic methods to different research areas where their benefits are most vividly observed—a few of which were enumerated earlier—this thesis is an effort to unify and extend them into a well-structured and generic framework for geometric modeling and spatial reasoning. The research in this area is still at a stage of infancy and many problems remain to be solved. Nevertheless, there is a case to be made for the tremendous theoretical and computational advantages of these power tools, which is the main purpose of this effort. I hope that this thesis will motivate researchers, developers, and entrepreneurs to explore more opportunities in this direction.

Chapter 1

Introduction

> "Mechanical design automation and, more fundamentally, the understanding of mechanical design in a scientific sense are progressing slowly if at all. Thus, we have a growing technological imbalance, with manufacturing striding ahead of design in terms of both scientific understanding and automation."
>
> Herbert Voelcker, 1988 [350]

1.1 Historical Perspective

The objective of computer-aided design (CAD) is to provide computational aids to facilitate the transformation of a set of behavioral specifications (i.e., function) to a proper structural description (i.e., form) of a product—including geometric, physical, material, and manufacturing process specifications—that meets those functional requirements [164]. Such a transformation typically involves exploring the design space to produce candidate solutions (i.e., synthesis), evaluate them against the design spec-

ifications (i.e., analysis), inspect their manufacturability with the available processes (i.e., planning), and iterate over this cycle to improve the performance-to-cost ratio (i.e., optimization). Every component in this so-called 'digital thread' depends heavily on common mathematical abstractions used to model, represent, and process the product data (i.e., 'digital mock-ups' (DMU) [356]).

In spite of the substantial amount of research in computer-aided design, analysis, and manufacturing (CAD/CAE/CAM), the observation made by Herbert Voelcker[1] in 1988 (quoted above) remains to be valid, with the complex engineering systems and liberating manufacturing processes growing faster than the advancements in the research on the engineering abstractions for design.

The theory of solid modeling [275,277,278,351] was originally developed to support the computerized automation of the production processes in 1970s and 80s [353]—particularly numerically controlled (NC) machine tool operations. After about four decades today, the same theoretical and computational models continue to govern the product life-cycle management (PLM) infrastructure. The traditional solid modeling environments still serve as the primary medium for mechanical product description to support most engineering activities including conceptual design, detailed design, manufacturing (both subtractive and additive), inspection, assembly, and so on [160], many of which nowadays need to answer to the kinds of practical needs that have dramatically evolved ahead of their original intentions and intrinsic limitation, des-

[1]Herbert Voelcker is a professor emeritus of mechanical engineering at Cornell University and one of the founders of the modern theory and practice of solid modeling [351,352]. He founded the production automation project (PAP) [353] at the University of Rochester in 1972 aiming to develop the solid modeler called part and assembly description language (PADL) [243]. He was joined by Ari Requicha in 1973 who further contributed to the mathematical foundations [276], modeling [275], and representation [278] of mechanical parts and assemblies, followed by much of the pioneering work on what is now called solid modeling [202, 203, 277, 279, 285, 296, 301, 343]. A historical review of PAP can be found in [353].

perately calling for a major reconsideration of the foundations [319, 348].

The traditional solid modeling tools are simply insufficient to address the modern engineering problems faced with increasing complexity of CAD/CAE/CAM needs for emerging manufacturing technologies. Some of the fundamental challenges with crippling effects on the state-of-the-art are elaborated in Section 1.3.

1.2 Basic Terminology

The abstraction paradigm that Voelcker and Requicha [351] introduced for geometric product specification is illustrated in Fig. 1.2.1. Such an information model typically comes with three elements referred to as

- the 'referent' (i.e., physical object, system, or process) that is being modeled;
- the 'model' (i.e., mathematical abstraction) which captures some of the properties of an idealized physical referent, which enables posing and answering questions (i.e., 'queries') about the referent; and
- the 'representation' (i.e., symbol structure) which is a *finite* description of the model for digital computing purposes, which allows for algorithmic implementations of those answers.

Similar proposals with different terminology were given by others in related areas of computing, e.g., by Denning [90] who proposed that in general, information in is essentially a triplet ⟨referent, interpreter, sign⟩ whose elements correspond to the above items, respectively.[2]

[2]The credit for this observation and its relevance to CAD/CAE/CAM belongs to Vadim Shaprio.

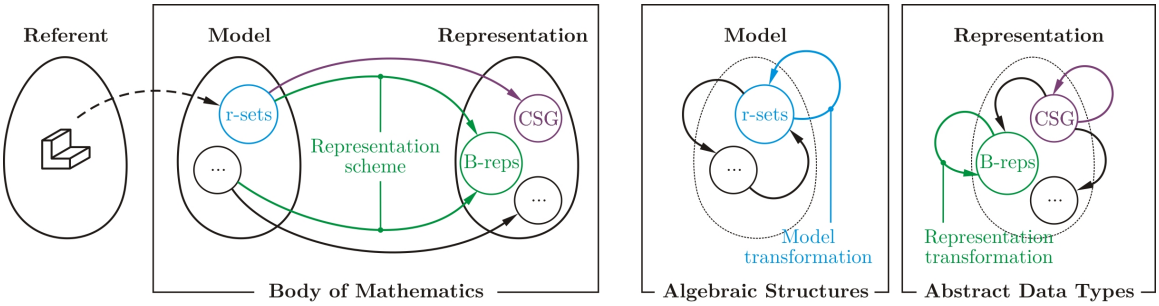

FIGURE 1.2.1: An abstract view of geometric specification (recreated from [351]) (left); and its endowment with algebraic tools (transformations and conversions) (right).

Computer representations are largely symbolic structures that can be thought of as compositions of computational primitives by rules specific to a given representation scheme [277]. A representation scheme, as illustrated in Fig. 1.2.1, is a conceptual mapping that defines the representation semantics, for which the notions of informational 'completeness', 'validity', 'consistency', and 'uniqueness' are formally defined and desired [278]. Among the common examples are

- constructive solid geometry (CSG) interpreted by pointset topological models (e.g., closed regular subsets of the Euclidean 3−space);
- boundary representations (B-reps) interpreted by algebraic topological models (e.g., orientable embedded manifold cell complexes); and
- (most commonly) hybrid consistent representations interpreted by multiple consistent models.

Convention 1.2.1. Hereafter I use the term 'description' to refer generically to both mathematical 'models' and computational 'representations' as originally defined by Voelcker and Requicha [351]. In particular, I will make a careful distinction between explicit and implicit descriptions:[3]

[3] I (ab)use the term 'description' extensively, for the lack of a better term, throughout this doc-

- 'Explicit' descriptions indicate models and representations of *pointsets*, which abstract and store an object in its entirety. It includes the traditional models such as r-sets, manifold cell complexes, semianalytic or semialgebraic sets, etc. and the most prevalent 'enumerative' and 'combinatorial' representations such as CSG, B-reps, cell decompositions, spatial enumerations, and groupings [277].

- 'Implicit' descriptions refer to models and representations of *functions*, from which the aforementioned pointset are implicitly inferred. It includes super-level functions such as characteristic (i.e., indicator) functions, membership classification or grade functions, R-functions, etc. and their 'constructive' and 'procedural' representations such as CSG,[4] F-reps, algebraic systems, etc. [252].

Both classes of representations are computational symbol structures with different primitives and composition rules. Enumerative representations are made of geometric primitives organized into combinatorial structures, while constructive representations are composed of functional primitives procedurally combined into more intricate functions. The distinction will become clearer in Section 1.4.

It is worthwhile emphasizing that this thesis is primarily concerned with modeling and the fundamental distinction between explicit and implicit models (i.e., pointsets versus functions). A discussion of enumerative versus constructive computer representations is out of scope. The reader is referred to [316] for an insightful discussion of fundamental classifications of representations.

---

ument when referring to a generic abstraction layer that collectively refers to models and representations.

[4]Constructive solid geometry (CSG) is special in the sense that it can be viewed both as a combinatorial representation that puts together geometric primitives into more complex constructions via regularized Boolean operations [282] and as a procedural representation that composes PMC functions of primitives into the PMC function of the object [316].

1.3 Research Challenges

Historically, the advent of new manufacturing technologies that surpass the modeling capabilities of existing computational tools has been the driving force for innovation. Every breakthrough in fabrication possibilities opens up a new space of previously unrealizable artifacts with additional dimensions of complexity and behavior, creating a need (and an opportunity) to revisit the foundations of modeling, representation, and processing the essential information. This evolutionary process inevitably keeps the computational models a step behind what would ideally leverage the full potential of the manufacturing enterprise most of the time, until a new breakthrough enables modeling tools to catch up.

The status quo in industrial PLM systems today is no exception to this trend.[5] The traditional solid modeling theory and practice is simply insufficient to manage the increasing complexity of computational design with the availability of new functional materials and advanced fabrication processes.[6] The evidence for this is ample, a number of which are given below, starting from purely geometric or kinematic problems (more relevant to this thesis) and continuing to those that also involve physical and material aspects (for future studies).

[5] This is aggravated by the overhead of introducing new tools to an existing PLM infrastructure, which is, unfortunately, a slow-paced process in an industrial setup due to the need for short-term economical incentives, extreme reliability guarantees, and other practical necessities that are difficult to fulfill simultaneously.

[6] "We don't really do "solid" modeling anymore!" –Jan H. Vandenbrande, Program Manager at DARPA, in his keynote address [348] at the SIAM GD/SPM'2015 conference.

1.3.1 Primary Challenges

Set-theoretic abstractions are suboptimal for geometric reasoning:

A typical computer representation of a geometric set that models an object of realistic details is a gigantic data structure with millions of metric (e.g., vertex coordinates) and combinatorial (e.g., incidence pointers) elements. In order to make sense of this enormous pile of information and to perform even the most basic geometric operations (e.g., constructions and queries) on them, 'smart' algorithms must be designed that make efficient use of the invariant patterns and composition rules in the representation scheme.[7] As a result, we have ended up with a bewildering variety of geometric representations, each of which are designed to efficiently handle some tasks that are clearly compatible with their internal structures, while they fall short at others.

For example, constructive solid geometry (CSG) trees lend themselves best to finite Boolean operations and point membership classification (PMC) queries but are suboptimal for operations on the boundary. Boundary representation (B-reps), on the other hand, are harder to use for performing Boolean operations and PMC tests, but are more suitable for visualization and freeform surface editions. To make the matters more involved, as soon as continuous *motion* is introduced—which is an indispensable functional component of most mechanical systems—almost all set-theoretic representations fail to provide straightforward means for efficient implementation of the most basic operations. Tasks as intuitive as computing the sweep of a shape under a given motion [5, 112, 145, 187, 254, 354, 372], finding maximal shapes and motions subject to containment constraints [162–164, 235], Minkowski operations [8, 30, 119, 201, 366], collision detection [169,188,208,339], and other direct and inverse configuration space

[7] See the sort-and-merge algorithm for the 1D example at the end of Section 1.2.

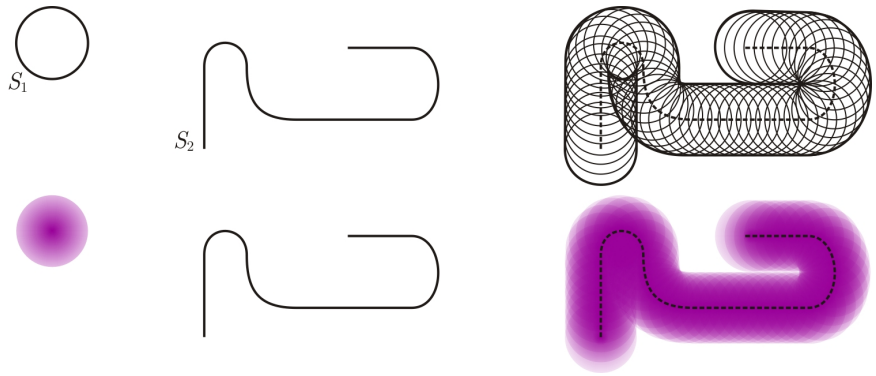

FIGURE 1.3.1: Set-theoretic versus measure-theoretic sweep of a disk along a planar curve.

problems prove to be more challenging than they appear to our intuition.

Perhaps the number one motivation for the adoption of the analytic approach is how it naturally and elegantly lends itself to a uniform formulation of shape and motion related problems with basic measure-theoretic constructs (i.e., integrals). For example, all of the kinematic problems mentioned above are uniformly posed in terms of convolutions of defining functions, which can be computed rapidly using FFTs [76, 258] regardless of the input shape or size complexity (e.g., polygon count).

Example 1.3.1. Figure 1.3.1 compares a set-theoretic versus measure-theoretic approach to computing the sweep of a disk along a 2D trajectory in the plane. Intuitively speaking, the former simulates the sweep as an infinite number of union operations that is hard to grasp coming from a classical CSG or B-rep mindset, while the latter simply resembles the accumulation of color on the paper as a paintbrush is stroke along a given path. From a theoretical perspective, the former involves somewhat complicated concepts (e.g., using differential geometry and envelope theory [40,41,136]) to establish rigorous relations between the geometry of the generator disk and its sweep, while the latter already has a simple mechanism in place for it: integration. From a computational perspective, the former suffers from robustness and complexity issues—particularly for identifying and trimming the so-called 'fold regions' (i.e., undercuts) [42, 111]—while the latter only requires standard sampling that applies to arbitrary geometry.

Generic shape analysis algorithms must replace ad hoc recipes:

In addition to the well-defined problems mentioned above that are readily formulated as morphological construction (e.g., sweep/unsweep) and configuration space queries (e.g., collision predicates), there are arguably harder problems that involve essentially ill-defined and somewhat fuzzy perceptions such as a notion of geometric alignment (e.g., approximate fit/contact measures) or rely on harder computational devices (e.g., a computable approximation of the medial axes (MA)). Consequently, a myriad of ad hoc solutions have been attempted at solving a number of classical problems in shape analysis. For example, *comparative* (i.e., similarity and complementarity) shape analysis problems are formulated in different ways with dissimilar implicit assumptions. A common theme in most of these methods is to extract a more compact description of the shape from the explicit representation (called shape 'descriptors' or 'signatures') that are quantitative measures of certain qualitative properties. It is not always clear what features or characteristics are being compared and what are the nature of transformations under which their invariance is desired. For instance, most shape matching methods that rely on MA [128, 131, 307, 329, 337] implicitly assume an interest in global topological properties (e.g., homotopy equivalence [207]), while partial shape retrieval and matching methods that rely on heat kernel signatures (HKS) [51, 92, 152, 250, 271] look at differential properties that are persistent under isometric transformations of the surfaces prevalent in biological forms. Other applications such as protein docking [25, 61, 110, 172, 190] require global measures of shape complementarity that are invariant under isometric transformations of the entire $3-$space (i.e., rigid body motions).

Once again, the analytic framework offers unique advantages and more flexibil-

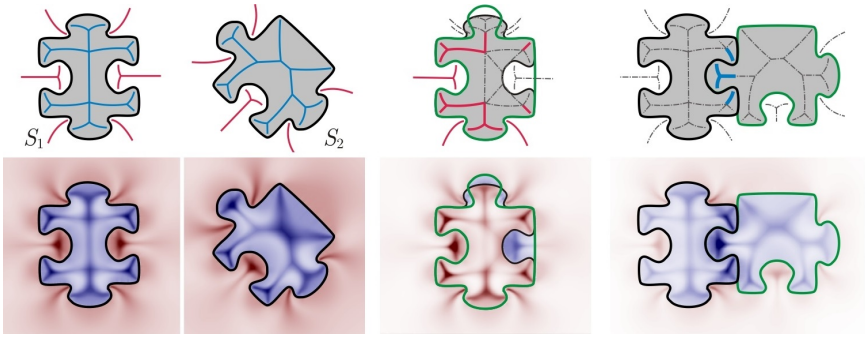

FIGURE 1.3.2: Set-theoretic versus measure-theoretic matching of planar jigsaw puzzles.

ity in the proper choice of shape and motion descriptors. For example, I show that the analytic correlation methods [23] that are popular in protein docking [25, 57, 69] can be generalized to subsume other shape matching techniques with the differences being captured by the choice of shape and motion descriptor functions. Such functions subsume and extend traditional descriptors such as spherical clusters [326, 327], double-skin layers [168, 357], fuzzy surfaces [216, 217], and skeletal densities [35, 37].

Example 1.3.2. Figure 1.3.2 compares a set-theoretic versus measure-theoretic approach to evaluating a score for shape similarity and complementarity between two 2D puzzle pieces across different rigid motions in the plane. Using an MA-based feature extraction approach, the former requires heuristic rules to measure partial matching between MA branches corresponding to different features, while the latter simply overlaps the two skeletal density function (SDF) [33] templates and aggregates their pointwise amplifications. From a theoretical perspective, the semantics used in the former are somewhat arbitrary and difficult to formalize (e.g., using methods such as 'shock graph' matching [307,328,329]), while the arbitrariness in the latter is summarized into a choice of a kernel for the SDF computation. From a computational perspective, the former suffers from robustness and complexity issues—particularly for stable MA approximating and pruning [17, 310]—while the latter is insensitive to noise (e.g., jagged bitmap boundary) and indifferent to geometric, topological, and syntactic complexity.

Traditional models are not computable in the presence of errors:

The traditional model of solidity [275] assumes access to computers with *exact* real number arithmetic capabilities and fails to provide formal principles to support robust algorithms for finite floating-point precision. For example, constructs as basic as the intersection surface for solids in contact [217] or the MA for piecewise smooth manifolds [59] are not 'computable' even for polyhedral shapes in the sense that a small perturbation in the shape data (e.g., noise in B-rep vertex coordinates) can cause dramatic changes in the output.

Edalat and Lieutier [104–106] proposed a domain-theoretic approach to a computable solid modeling framework that is inclusive of irregular sets and is closed under Boolean operations. In spite of its elegance, their approach misses clear links to the existing representations and algorithms based on traditional models. Qi and Shapiro [260–262] proposed $\epsilon-$solidity and $\epsilon-$regularity—which subsume Requicha's exact solidity and regularity as $\epsilon \to 0^+$—to establish formal principles for tolerant modeling. The major drawbacks of this approach are that it specifies a global single tolerance value and it does not guarantee topological consistency across different interchangeable shapes within the tolerance zones.[8]

I propose to replace this 'zone-based' approach to tolerant modeling with concepts from fuzzy modeling [219, 369], which are essentially analytic methods. In particular, the PMC is to be replaced with a membership density function (MDF) which, depending on the application, can be interpreted as the membership grade, mass density, or probability distribution. In particular, Boolean and Minkowski operations naturally generalize from real-valued arithmetics and convolution, respectively, on PMC for

[8]Moreover, the $\epsilon-$neighborhoods used to define the tolerance zones are prone to the same (even intensified) difficulties in terms of finite describability that infinitesimal neighborhoods face.

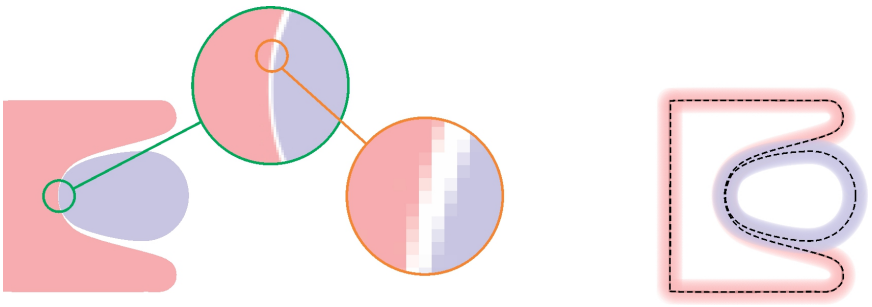

FIGURE 1.3.3: Set-theoretic versus measure-theoretic boundary contact characterization.

'crisp' sets to those on the MDF for 'fuzzy' sets with meaningful tolerant modeling interpretations. In addition, operations that are not computable on crisp sets, such as measuring surface contact and constructing the MA, due to instability with respect to small perturbations, turn into computable operations over fuzzy sets who boundaries are dissipated in a meaningful trade-off between precision and resilience.

Example 1.3.3. Figure 1.3.3 compares a set-theoretic versus measure-theoretic approach to measuring effective contact between two non-colliding 2D objects along their shared 1D interface. Although the notion of "contact curve length" is mathematically well-defined, it s computationally ill-posed for several reasons—that despite manifesting differently, are all traceable to a fundamental problem in formulation. First, the result can change from finite curve length (as seemingly the case in the figure) to zero curve length by moving one of the shapes by a small displacement. There is no lowerbound to how small a change can be to create such a large effect. Additionally, exact contact or overlap is almost never the case in practice, and one must specify a finite tolerance for lack of contact (e.g., separation) or noisy boundary (e.g., jagging) for the question to make sense. A sensible approach is to formulate 'effective' contact measures [217] in which the fuzziness of the boundary is captured a priori in the models, rather that being introduced after-the-fact based on arbitrary heuristics. The same fundamental problem is observed with many other computations that involve interactions of lower-dimensional features (e.g., curves and surfaces) embedded in the 3D or higher-D space.

Neighborhoods may be the right answer, but to the wrong question:

The traditional model of solidity [275] requires dimensional homogeneity (i.e., topological regularity) as a necessary condition all along to guarantee physically realizable shapes. However, regular sets are not closed under set-theoretic Boolean operations $(\cup, \cap, -)$ [276] which has led to the introduction of 'regularized' Boolean operations denoted by asterisks $(\cup^*, \cap^*, -^*)$ [343]. The problem is that unlike their set-theoretic counterparts, the implementation of r-operations on r-sets is more than that of simple logical combinations (i.e., table look-up) and appears quite unnatural to combinatorial methods. In particular, it is not possible to compute even the simplest finite r-operations—let alone infinite ones, e.g., in problems involving motions—without examining the *neighborhoods* of individual surface cells. Unfortunately, neighborhoods are quite difficult to model, represent, and compute.[9] Their current implementations are limited to case-by-case simplification recipes that are highly dependent on the representation scheme and easily breakable by introducing aberrations.

Shapiro [312] addressed this problem by providing a precise definitions of 'well-formed' set representations as well as establishing necessary and sufficient conditions for well-formedness. It was shown that well-formed representations do exist for semi-algebraic r-sets, but their existence for semianalytic r-sets is not clear.

I propose to abandon our long-lasting faith in regularity at the modeling phase and define solids as equivalence classes of irregular sets that are indistinguishable when looking through a measure-theoretic lens. The modeling space is then reduced to a

[9] It is easy to prove that the neighborhood of a 'candidate' set (e.g., a point) against a 'reference' set (e.g., a solid) is not finitely describable for general semianalytic r-sets except when the dimension of the two sets differs by one unit, e.g., testing a line against a polygon or a plane against a polyhedron. This is an important detail omitted in the set-theoretic foundations [342] leading to the false promise of informational completeness.

quotient space of the power set $\mathscr{P}(\mathbb{R}^3)$ that is isomorphic to the class of regular sets. I will clarify what I mean by this statement in Chapter 2 with more details to follow in Chapter 3. From this viewpoint, all regularization tasks, rather than burdening the combinatorial algorithms to maintain regularity at every step—or repair/heal irregularities—are delegated to the analytic integrals whose eventual finite-difference approximations will automatically eliminate irregularities.

'Informational completeness' is a myth:

Not only practically, as demonstrated in the above two challenge items, but also from a theoretical perspective. Requicha defines completeness (i.e., unambiguity) of a representation as containing "enough information to distinguish a single entity from all other entities in the modeling domain" [277] which eventually boils down to the ability to respond to PMC. However, a ternary PMC test: $\mathbb{R}^3 \to \{\text{“in”}, \text{“on”}, \text{“out”}\}$ does not provide any additional information (in principle) beyond the binary inclusion test: $\mathbb{R}^3 \to \{\text{“in/on”}, \text{“out”}\}$. It might appear that discriminating between interior and boundary points is helpful for resolving regularized Boolean operations. However, an "on" response would be useless without a description of the neighborhood—whose complete representation for semianalytic sets would essentially be $\infty-$dimensional—see the illustration in Fig. 1.3.4.

On the bright side, an "on" response along with the neighborhood properties are accessible from the differential properties of defining functions (e.g., indicator function or MDF). I posit that the smooth MDF descriptors that generalize the discontinuous indicator functions are the proper apparatus to achieve a more realistic form of informational completeness. Interesting problems to address in this area are

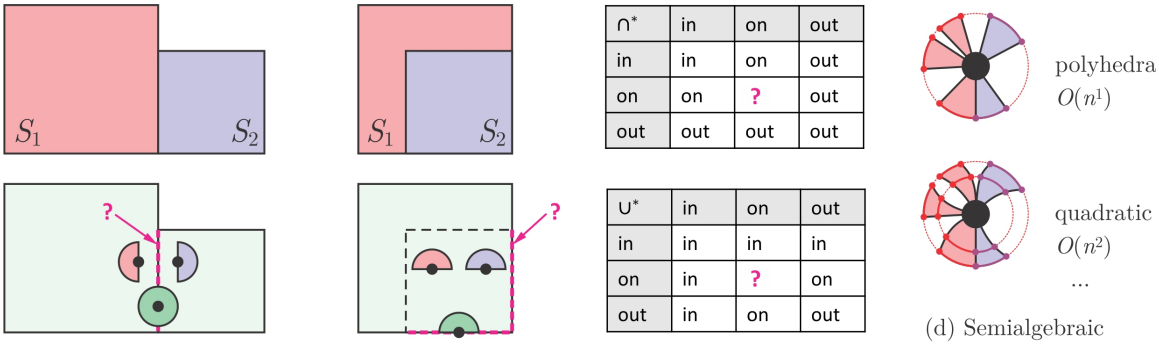

(a) Regularized union (b) Regularized intersection (c) Look-up table (d) Semialgebraic neighborhood

FIGURE 1.3.4: Performing implicit regularized Boolean operations via PMC requires explicit knowledge of the $\epsilon-$neighborhood geometry as $\epsilon \to 0^+$ to resolve ambiguities.

obtaining bounds on the level set variations (e.g., in terms of Hausdorff distance) and/or conditions for topological consistency (e.g., in terms of persistent homology) to reconcile with the zone-based tolerant modeling.

Algorithms that scale with structural complexity are bottlenecks:

The combinatorial methods are inevitably dependent on the topological complexity (e.g., connectivity and number of holes), geometric complexity (e.g., convexity and type of surfaces), and syntactic complexity (e.g., number of triangles or voxels) built into the structure of explicit representations. For instance, few practical algorithms are available for Minkowski operations on arbitrary shapes (even in 3D), and the input size can quickly make such computations impractical even in the case of polyhedral B-reps [8, 30, 119, 201, 366]. These difficulties can be overcome by sampling-based methods [88, 174, 205, 235, 349] that are becoming increasingly popular in robotics and spatial planning [199, 210]. In particular, the convolution+FFT-based methods [88,174] for motion planning can be viewed as a special class of analytic methods using indicator functions and uniform samples [214, 215, 218] a nonuniform and rotation-

invariant extension of which I presented recently in [36].

Most importantly, I intend to show that the analytic approach liberates one from the combinatorial blow-up of complexity and allows for flexible and meaningful trade-off between desired accuracy and available computational resources (i.e., time and memory) by leveraging techniques from harmonic analysis and signal processing.

1.3.2 Ambitious Horizons

*** More explicit form-function correlations are needed:**

The primary role of geometry to build a language for describing mechanical products and processes is historically well-justified, since shape is relatively easier to understand, visualize, and describe compared to mathematical models of mechanical function and physical processes [160]. However, traditional geometric modeling systems do not contain the key functional and behavioral information[10] needed for many engineering activities [160,251,321]. This, in turn, tends to instigate costly iterations in a product development cycle [160] that typically involve (numerous and blindfolded) repetitions of physics-based analysis on (excessive and unnecessary) levels of detail on geometry-based design trials. The need for and lack of richer models and representations that more explicitly link geometric shape (i.e., form) to mechanical behavior (i.e., function) have been constantly pronounced for a long time [9, 160, 321].

Over the years, a variety of solutions have been proposed to establish such explicit correlations in restricted contexts such as kinematic function [164] and physical behavior [251]. However, they have not yet received the academic attention or in-

[10]Except for simplistic symbolic and textual information to accessorize B-rep faces with (e.g., tolerances, material properties, process characteristics, and so on).

dustrial recognition that they deserve. The functional (i.e., analytic) approach to product description appear to be the proper setup to support systematic form synthesis with respect to both kinematic and physical aspects of function, and in fact brings together some of the existing promising approaches within a unified framework. In particular, configuration modeling [231, 303] for kinematic function and the use of discrete forms [344, 345] for physical behavior have their roots in the premise of unifying shapes, motions, and physical quantities in (continuous or discrete forms) of space-time tensor fields that generalize membership functions.

*** Interoperability and integration are open problems:**

A bewildering variety of shape representation schemes and software services have emerged in the past few decades in response to the growing application-specific needs. These schemes range from point clouds, parametric (e.g., nonuniform rational B-spline (NURBS)) or tessellated (e.g., triangular mesh) surfaces, and volumetric enumeration (e.g., voxelization) to fully parameterized feature-based structures, each appealing to the specific needs at a certain stage of the PLM chain. One of the big challenges today is of 'interoperability' of PLM services [147, 148]; namely, reliable exchange and translation of geometric data with minimal information loss and least possible compromise in validity and consistency of the translated representations. Such problems are inevitably caused by different precision measures of the sending and receiving systems, different semantics of the evaluation algorithms, and the ad hoc nature of representation healers and mesh repair algorithms [262].

A notorious example is the "CAD-CAE integration" problem [320] which impairs the iterative design+analysis cycle due to the necessity of manual interventions by

domain experts, e.g., for geometric simplification, quality mesh generation, and consistent mesh repair operations on exported data in neutral file formats such as the "standard for the exchange of product model data" (STEP) [1, 2].

Although a variety of approaches have been attempted to solve the data translation problem [260–262], it appears that this 'data-centric' approach using one-to-one customized translators—i.e., an attempt to passing over a complete description in a single shot—is part of the problem rather than a solution [319]. In contrast, a 'query-based' approach [147, 148] that enables dynamic and adaptive information exchange is more promising. Research in this area is still premature, but the functional (i.e., analytic) approach and its algebraic topological extensions (e.g., à la Tonti [344,345]) appear to be the proper abstractions to support it.

*** Multiscale and mutliphysics models are imperative:**

As the complexity of engineering systems grow, from self-assembling nano- and micro-structures functioning in multiphase environments to enormous aircraft or ship assemblies that are structured at multiple scales, more versatile design tools are demanded. Current theory of solid modeling assumes a homogeneous and isotropic continuum as the underlying space [275] and chooses to formulate queries (e.g., PMC or distance functions) at different points as the primitive elements of the space-time. However, a notion as simple and fundamental as a point becomes ambiguous in the face of the new modeling requirements. For example, when dealing with new materials (e.g., carbon fiber composites) and novel fabrication processes (e.g., 3D printing), different geometric primitives should be used at multiple scales—e.g., continuum models at macro-scale, texture models at meso-scale, and crystal models at micro-scale. The

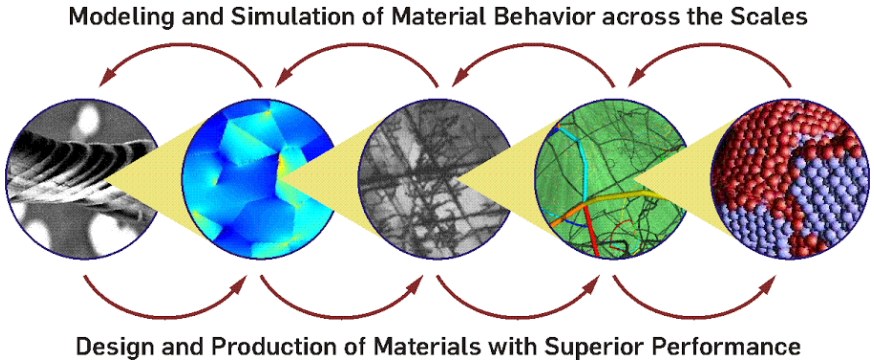

FIGURE 1.3.5: More advanced models and representations are needed to bridge across length and time scales for computational design of functional material structures. Figure courtesy of Dennis M. Kochmann (private communication) www.kochmann.caltech.edu.

traditional membership queries are no longer sufficient for informational completeness and more flexible functional (i.e., analytic) descriptions ranging from continuous forms (i.e., density fields over the continuum) or discrete forms (i.e., co-chains defined over cell decompositions) are demanded for representing geometric, physical, and material distributions at different space-time scales (Fig. 1.3.5).

I anticipate a surge of research during the remaining half of the decade in the development of multiscale and mutliphysics abstractions for model-based systems engineering (MBSE) [319] supported by federal funding (e.g., DARPA [244–246]). One way or another, it is expected that concepts from algebraic topology (e.g., simplicial homology) and category theory (e.g., sheaf theory) will be indispensable role players in the next generation of geometric modeling advancements. The proposed functional (i.e., analytic) approach is, in principle, extensible for the development of new modeling techniques using these power tools.

The solution of each and every one of the problems described above is fairly ambitious, and I am not likely to be able to address a number of them. I intend to

demonstrate the key advantages of the analytic approach to the first six challenges in this Ph.D. thesis and to suggest guidelines for future work on the last three problems marked by an asterisk (*) which appear significantly more challenging.[11]

1.4 Proposed Alternative

The commutative diagram on the front page, repeated below in Section 1.4.1, summarizes the proposed paradigm shift to support the next generation of geometric modeling and computing tools in addressing some of the aforementioned challenges. Section 1.4.2 briefly reviews examples from the relevant literature on scattered efforts that adopted a similar approach in different applications. Section 1.4.3 presents the main claims of the proposed framework in its ability to address classical and new geometric problems more effectively and efficiently.

1.4.1 A 'Detour' (or not!)

The traditional approach in many engineering problems is to abstract physical objects as explicit *pointsets* represented via enumerative representations that contain topological and geometric information in its entirety, in accordance with the requirement of 'informational completeness' [351]. Accordingly, the geometric operations can be viewed as mappings of pointsets-to-pointsets, which are implemented as representation transformations using combinatorial methods:

[11] Further research in the last two areas is the candidate's next project for a post-doctoral fellowship at the University of California at Berkeley after defending this Ph.D. dissertation!

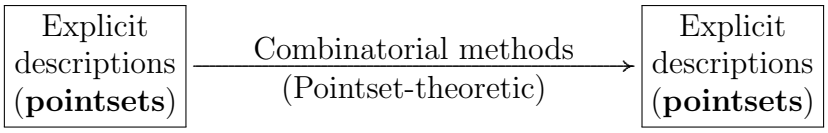

(1.4.1)

These pointset models and the operations applied to them form algebraic structures, whose implementations as enumerated finite constructs (i.e., data structures) and finite programs (i.e., algorithms) form the so-called abstract data types (ADT) [75].

The proposed paradigm shift is to think of physical objects as being implicitly (and partially) defined in terms of *functions* (e.g., spatial fields or signals) that describe the topological and geometric properties as spatial distributions. In other words, each object is 'characterized' by a set of 'queries' that need not reveal complete information about the object—but just enough to answer the same queries about the result of the geometric transformation by computing some analytic correlations between them. Thus the transformations are conceptualized as mappings from functions-to-functions:

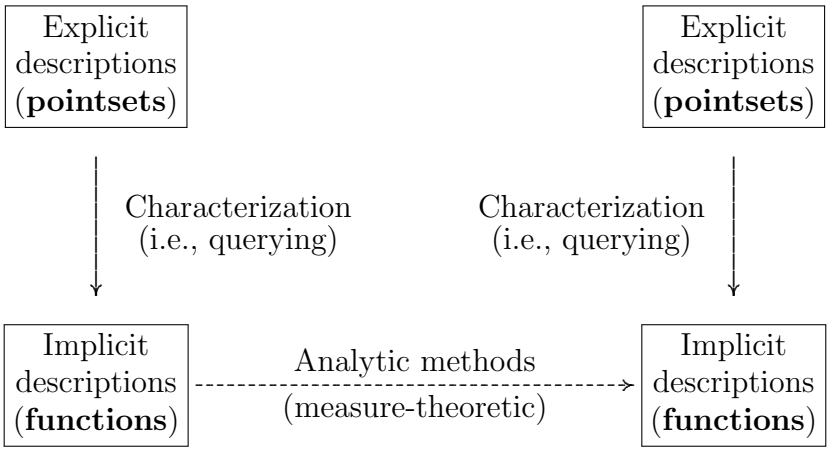

(1.4.2)

These functions can be viewed as *abilities to respond* to queries about local properties (e.g., set membership, distance, density, etc.), in contrast to a global account captured by explicit descriptions. The geometric operators (i.e., algebra of pointsets) are then replaced with corresponding functional operators (i.e., algebra of functions), which

can be represented using a variety of meshing and sampling techniques.

Although it is not strictly required by this paradigm, most cases studied in this thesis shall use informationally complete functional descriptions from which a point membership classification (PMC) can be retrieved by checking logical conditions on the value of the function (e.g., as its sub-/super-level set) and/or its coordinate system independent differential properties that capture neighborhood information. In this case, the vertical downward arrow in (1.4.2) correspond to invertible maps, whose inversion 'reconstructs' the explicit pointsets from the shape descriptor functions. Thus the complete diagram looks like this:

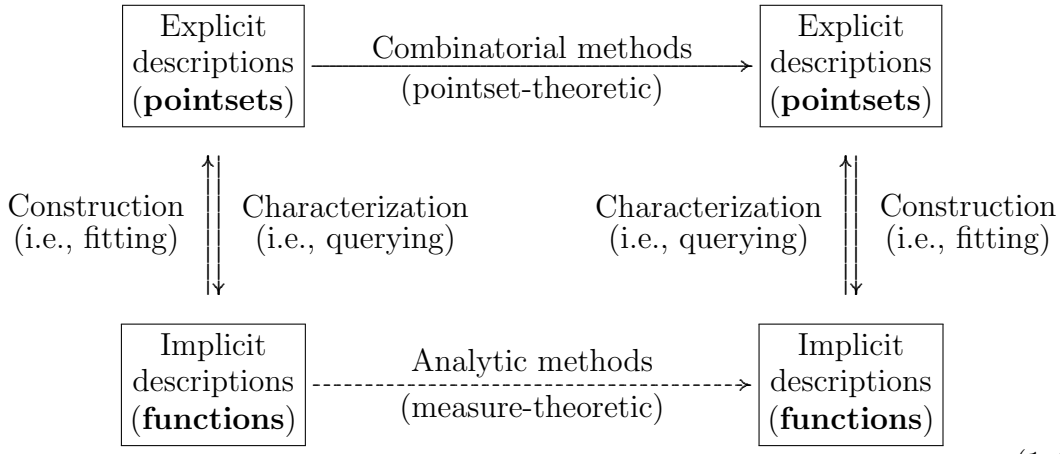

(1.4.3)

Example 1.4.1. Figure 1.4.1 illustrates the distinction with a simple 1D example. On the top panel, we have two 1D objects modeled as finite unions of intervals and represented as sequences of sorted endpoints with binary neighborhood membership information; the overlap is then modeled as a set intersection operation and computed by merging the representation and resolving the membership in between them. On the bottom panel, the same 1D objects are modeled as Boolean functions—i.e., 'black boxes' that are able to answer consistently to the "in or out?" question for a given point)—

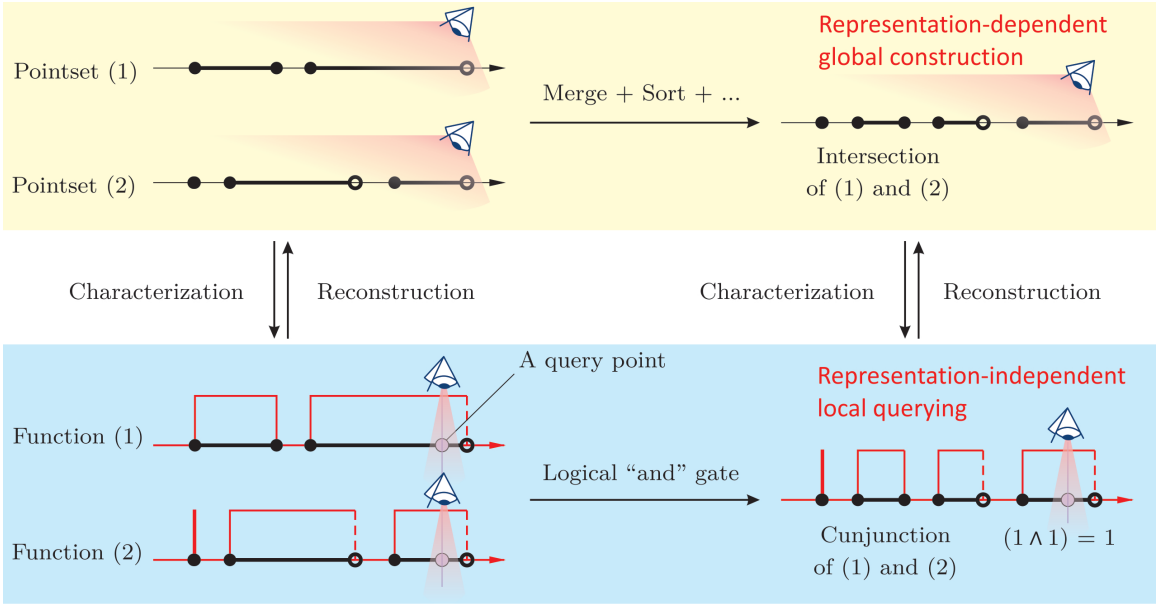

FIGURE 1.4.1: A one-dimensional example of the fundamental difference. Using the indicator functions of each set, evaluating the indicator of intersection at every query point amounts to evaluating the conjunction of individual indicators—i.e., inclusion test in the intersection is the same as an "and" combination of individual inclusion tests.

the overlap is then modeled as a logical "and" (i.e., conjunction) operation. One can in principle convert between the two models and descriptions by point membership classification (PMC) and shape reconstruction.

A more technical description of the above example would look like this: given two finitely describable 1D sets $S_1, S_2 \subseteq \mathbb{R}$, i.e., each decomposable to a finite union of open, closed, or mixed intervals (including singular points),

$$\mathbf{1}_{S_1 \cap S_2} = (\mathbf{1}_{S_1} \wedge \mathbf{1}_{S_2}), \quad \text{where } \mathbf{1}_S : \mathbb{R} \to \{0, 1\}, \text{ and } \mathbf{1}_S(x) = (x \overset{?}{\in} S). \tag{1.4.4}$$

As with every implicit description, the corresponding functions are obviously non-unique. Although the choice of indicator (i.e., characteristic) or PMC functions is obvious for simple Boolean operations—and as we shall see, also for Minkowski operations, sweeps/unsweeps, etc.—it is less so for more complex classes of problems that have to do with measuring contact, shape similarity/complementary, etc. In general, the main challenge is to define the appropriate set of rules to construct effective and

well-defined 'shape descriptor' functions such that the transformations depicted by a solid horizontal arrow in (1.4.1) convert to meaningful operations depicted by a dashed horizontal arrow in (1.4.2) whose implementation on the space of functions exhibit *substantial theoretical and computational advantages*—e.g., can be computed faster, support more general cases, are more robust to errors, etc. If this challenge is somehow overcome, a case can be made for the benefit of taking a detour like this:

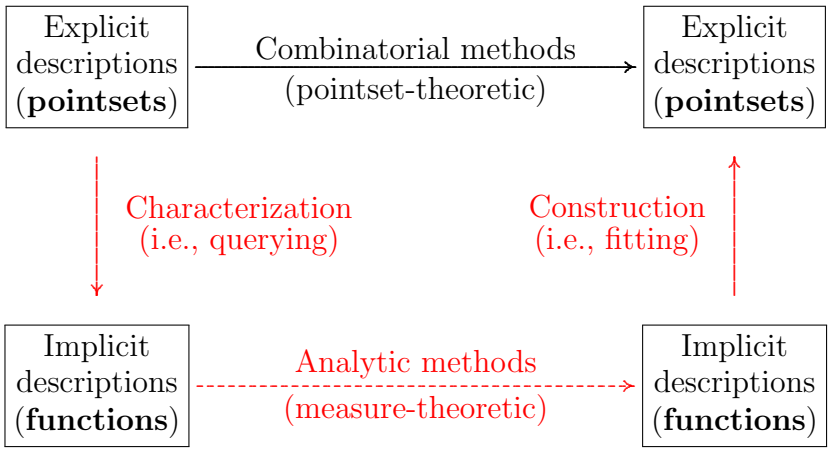

(1.4.5)

if one can show that it preserves the integrity of the results—i.e., the above diagram 'commutes' up to reasonable accuracy, at least for an important subclass of operations. The 'characterization' and 'construction' maps[12] are pointset $\rightleftarrows$ function convertors. The former can be as simple as a function generator that lifts only the included points in the pointset to a nonzero field (e.g., characteristic or distance function), for which the inverse map is one that takes the support (i.e., $0-$superlevel set) of given function, as in the example of Fig. 1.4.1, which are sufficient for a variety of fundamental problems [215, 218, 231]. More complex maps are needs, such as skeletal density functions (SDF) [33] for more complex problems [32, 35, 37].

[12]Or more technically, 'functors' in the language of category theory [332].

"But is it really a detour?"[13] Closer inspection of the problems in which analytic methods have been more popular than others reveals an interesting insight; namely, that the functional descriptors are more fundamental than geometric pointsets to the physical observations being modeled. In such cases the traditional approach is truly the circuitous one that should be carrying the burden of practical justification:

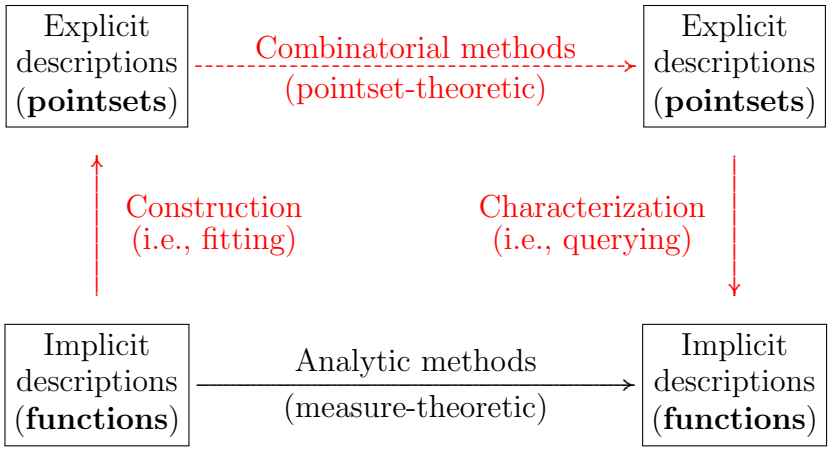

(1.4.6)

For the same reason that implicit characterizations are nonunique, the explicit classification entails a loss of information that could otherwise be useful in the following transformations (i.e., horizontal arrows). After all, explicit geometric constructions are made-up incomplete abstractions—partly because shapes are relatively easier to imagine and visualize using graphical drawings [160]—while functions (i.e., queries about shape properties) more directly correspond to physical measurements to which the shape is fitted according to often arbitrary rules.

[13] Almost verbatim from a comment from Don R. Sheehy in my Ph.D. defense, which turned into a long fruitful discussion leading to this paragraph.

1.4.2 A Premier Example

My favorite example to support the proposed paradigm shift, and to demonstrate the claim that it lends itself more naturally to geometric and physical modeling, is 'protein docking'. It refers to the task of computing the 3D structure of a protein complex starting from the individual structures of the constituent molecules [288].

Protein structures are experimentally determined using methods such as X-ray crystallography or nuclear magnetic resonance (NMR) techniques. The simplified geometric molecular models that we would like to obtain—such as collections of balls (for individual atoms) or combinations of $\alpha-$helices and $\beta-$strands (for secondary structure)—are indirectly inferred from these measurements. For instance, X-ray diffraction data yields a Fourier representation of the electron density function. After reconstructing the density function in the physical domain—which is subject to significant uncertainty due to the very notorious 'phase problem' [263]—a variety of computational techniques subject to a range of different assumptions and arbitrary parameters (thus more uncertainty) are used to fit spherical atomic models to density data. This process can be viewed as the characterization step, i.e., vertical downward arrow in (1.4.6) which suffers from loss of information and arbitrariness of geometric form fitting. The constructed structural representation is then stored in the protein data bank (PDB).[14] Figure 1.4.2 illustrated a high-level schematic of this process.

Many protein docking algorithms begin with a simplified rigid docking in which conformational the constituent proteins (called 'receptor' and 'ligand'), which has been justified as a reasonable starting point due to the observed similarity of the crystallographic structures between the bound and unbound conformations [63, 166].

[14]The PDB database is updated weekly, and contains more than 100,000 X-ray and more than 10,000 NMR structures at the time of this writing: www.rcsb.org/pdb.

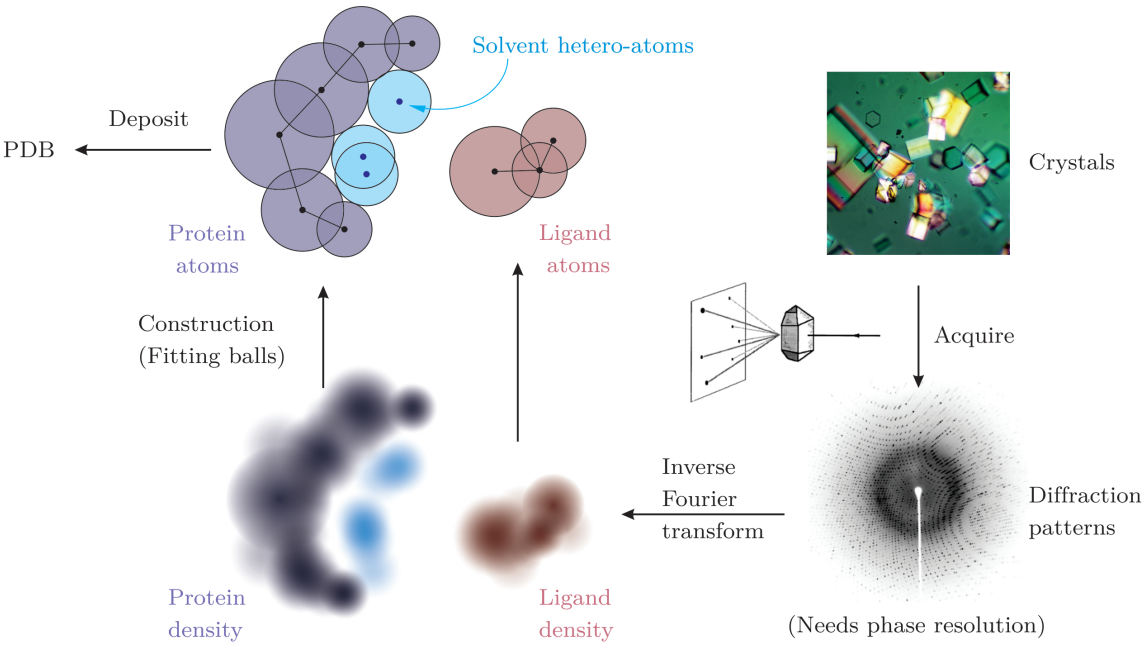

FIGURE 1.4.2: The most common approach to protein structure determination involves collecting X-ray diffraction data, resolving the phase, and geometric parameterization.

There is a widely accepted consensus among the computational biologists,[15] growing only stronger since the early days of protein structure determination in 1950s [177], that *shape complementarity* is a major factor in molecular recognition and binding. Most geometric docking algorithms use atomic coordinates from PDB as input and reproduce molecular surface representations. For surface matching, combinatorial algorithms are devised to search for docking configurations (i.e., 'poses') of high shape complementarity, defined in terms of how well the surface protrusions and depressions,[16] are aligned with each other. The alignment is characterized by matching the so-called 'shape descriptors' across the two molecules. Different shape descriptors and matching criteria have been used, for example:

[15]This can be traced back to Emil Fischer's 'lock-and-key' rigid binding model proposed 120+ years ago (1895 to be exact), extended by Daniel E. Koshland, Jr. in 1958 to 'induced fit' binding model taking conformational changes into account. See [189] for more details.

[16]Also called ridges'n grooves [194], knobs'n holes [74], bumps'n clefts [239], hills'n pits [240], etc.

- aligning spherical clusters packed inside the ligand against those packed into the receptor's binding pocket, using consistency of pairwise internal distances as the matching criteria [194, 326, 327];

- aligning critical points and normals selected on the ligand's surface against those of the receptor's surface with matching convexity quantifiers [239–241];

- aligning surface skin layer atoms of the ligand against an extra grown skin layer of pseudo-atoms swept over the receptor's surface [60, 61, 168, 172, 357];

and so on. The combinatorial matching is performed using methods ranging from graph-theoretic search [114, 225, 327] (e.g., using clique-finding and bin-matching algorithms) to geometric hashing [102, 117, 118, 204, 239–241] adapted from computer vision [196–198] to the molecular recognition problem. Disregarding the differences in combinatorial strategies used to speed up the search for the best docking poses (i.e., the optimization task), the common attribute among all of these methods is an attempt to quantify the geometric alignment between surface features of the receptor's exterior and the ligand's interior in the vicinity of the common interface (i.e., the scoring task). Aligning features are detected by comparing transformation-invariant signatures such as pairwise internal distances and angles between normals and center-lines across the two molecules, and the best alignment pose is found by minimizing root-mean-squared deviation (RMSD) of those features. Figure 1.4.3 schematically illustrates the combinatorial protein docking process.[17]

More recently, the so-called 'correlation techniques' [172] have found enormous

[17]For simplicity, the solvent hetero-atoms deposited alongside molecule's own atoms in the PDB in Fig. 1.4.2 are used to exemplify surface shape descriptors in Fig. 1.4.3. However, protein docking packages usually ignore the solvent hetero-atom position obtained from crystallography and repopulate the surface with pseudo-atoms (of the same size $1.2-1.4$ $\mathring{A}$).

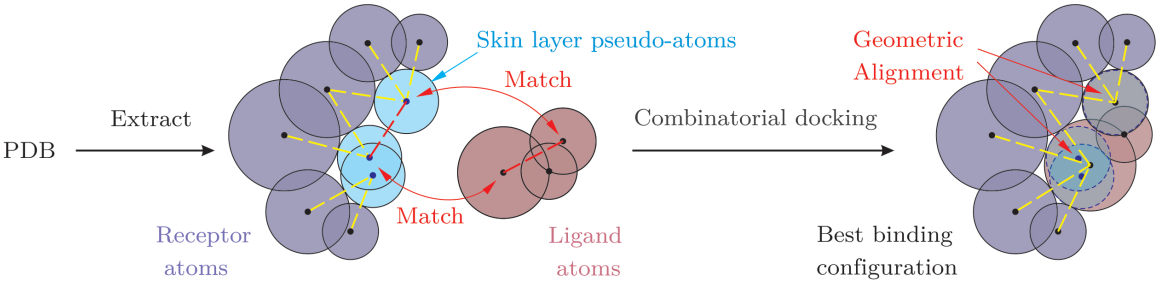

FIGURE 1.4.3: Protein docking via combinatorial methods that match the geometric shape descriptors obtained from the structural representations extracted form the PDB.

popularity in protein docking, which are the prefect example of the analytic detour expressed in (1.4.2). I this method, the molecules are characterized via their density functions reproduced from the geometric model, e.g., by assigning Gaussian radial kernels to each atom whose coordinates and radii are extracted from the PDB [127]. The geometric alignment for a given pose is in turn measured by computing the cumulative overlap of the two functions (i.e., their inner product), which, unlike the discrete combinatorial matching, yields a continuous scoring scheme over the configuration space of relative motions—referred to as the 'correlation function' in this context. The key observation is that inner products of the same pair of overlapped functions computed in different relative translations (i.e., shifts) is the classical definition of a convolution [173] which can be extended in principle to general rigid and affine transformations [65]. As a result of the well-known convolution theorem, the score function can be computed in the frequency domain as a simple pairwise multiplication. A variety of protein docking packages take advantage of this technique by rasterizing the volumetric occupancy of the molecules into uninform Cartesian grids to compute the translational correlations for sampled rotations using 3D fast Fourier transforms (FFT) [60, 61, 110, 126, 172, 191]. Others use polar grids to compute the translational and rotational correlations for sample intramolecular distance using 5D

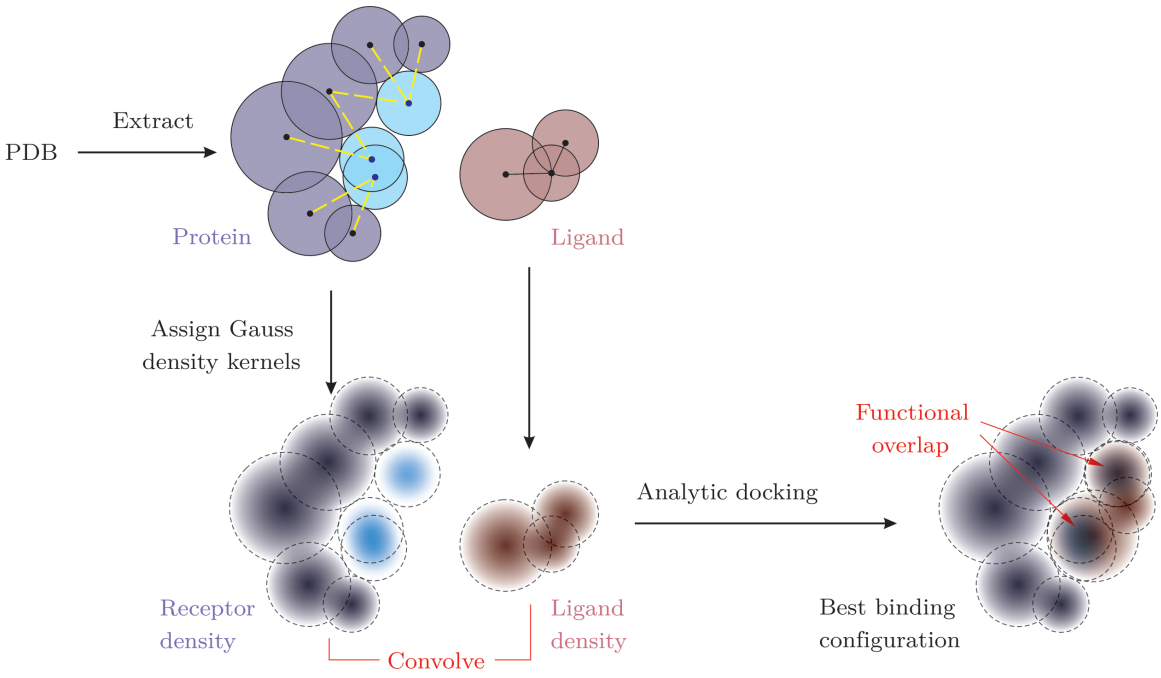

FIGURE 1.4.4: Protein docking via analytic methods that match the functional shape descriptors obtained by assigning Gaussian kernels to geometric units (e.g., atoms).

spherical harmonic Fourier transforms (SHFT) [190, 220, 289, 290]. More recently, grid-free generalizations have also been developed [25, 57, 69], where Gaussian kern taking advantage of nonequispaced FFT (NFFT) algorithms [258].

The main differences across these methods are in how they narrow and focus the search space (i.e., optimization task) and/or how they distinguish each grid cell based on its proximity to the protein surface and its neighborhood geometry (i.e., scoring task). As far as the scoring scheme goes, the horizontal arrow on the left-hand side of (1.4.2) corresponds to the different characterizations of the molecular densities (e.g., as Gaussian radial basis functions (RBF)) while the horizontal arrow is the combination of computing the convolution and selecting the set of 'good fit' configurations. There are a few benefits to using analytic methods for protein docking:

1. *Robustness*: The combinatorial matching is sensitive to small changes in geometric features due to their discrete operations (e.g., distance hashing and clique-finding), whereas the analytic methods are much more robust to variations and more capable of capturing approximate overlaps.

2. *Efficiency*: The combinatorial search dramatically slows down with the number of atoms, as the discrete search complexity grows exponentially, whereas the analytic method's complexity is independently adjusted by the choice of the grid resolution, allowing meaningful cost-to-benefit trade-offs.

3. *Effectiveness*: The combinatorial methods often have difficulties in dealing with molecular surfaces with extensive flat regions and few protrusions and depressions, as they operate over a space of shape descriptors enumerated by such notable features. The analytic methods, on the other hand, do not suffer from such limitations as they work with functional shape descriptors, which can be quired at all spatial points and can be modified to reflect different features.

Figure 1.4.4 schematically illustrates the analytic protein docking process.

The protein docking example is one in which the analytic correlation method appears more natural to the problem, as the functional descriptors are what one measures in the first place, while geometric models are error-prone arbitrary shape fitting attempts that have their roots in dated atomic models. The process of geometric reconstruction from X-ray data followed by functional characterization via Gaussian models appears to be redundant—ignoring the fact that a geometry-centric PDB database is valuable in many other aspects. Nevertheless, it is a natural question to ask whether this process can be eliminated altogether for protein docking (at least in principle) and if the FFT-based correlation techniques can be used directly on

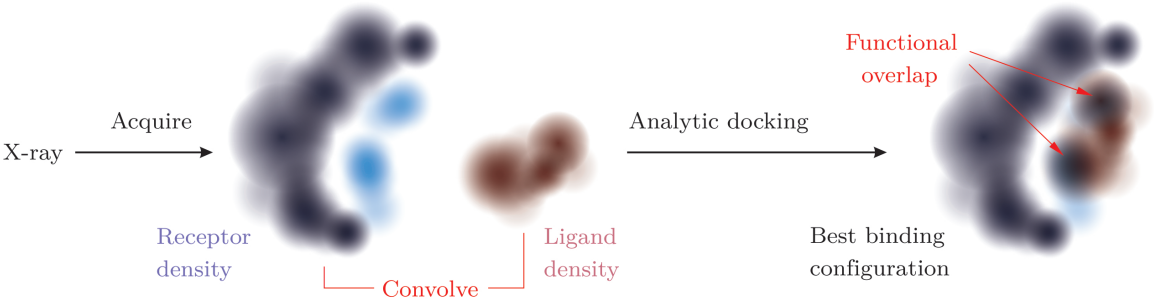

FIGURE 1.4.5: Protein docking via analytic methods that match the functional shape descriptors obtained directly from the X-ray diffraction data.

the original X-ray diffraction data measurements obtained in the frequency domain.[18] Such an exercise is beyond the scope of this thesis, but it bears value as to find out if (and to what extent) it would affect the reliability of docking predictions.

For more information on different aspects of protein docking, the reader is referred to several surveys [47, 288, 338, 346].

1.4.3 Central Claims

Based on the example problems reviewed in Section 1.4.2, I posit that a general theme is common to a large number of shape and motion related problems:

Remark 1.4.2. The 'direct' problems involving interactions of shapes and motions rely on the ability to create arrangements of a primitive 'template' repeated according to a (finite or infinite) 'pattern'. The goal is to reason (i.e., answer queries) about the resulting complex structure in terms of the queries of those simpler ingredient. The examples include but are not limited to

- applying a discrete set of motions to another set of shapes or motions to populate (intact or modified) copies of a unit element (i.e., 'building block') in order to

[18] The 'phase problem' remains a challenge as X-ray data reveals amplitudes only [263].

create an organized structure, which underlies numerous enumerative geometric representations;[19] and

- applying a continuous set of motions to another set of shapes or motions to create swept volumes or mixed motions, which underlies morphological shape synthesis, mechanism workspace analysis, etc.

Remark 1.4.3. The 'inverse' problems involving interactions of shapes and motions rely on the ability to characterize the 'cumulative overlaps' between features of given objects over their configuration space. The goal is to reason (i.e., answer queries) about the resulting overlap in terms of the queries of the participating objects. The examples include but are not limited to

- detecting and measuring interference (i.e., set intersection)[20] between different features of given objects which underlies collision detection, contact measurement, motion and path planning, etc.

- detecting and measuring alignments of shape descriptors inferred from given objects which underlies comparative shape analysis (e.g., similarity, complementarity, and symmetry measurement), etc.

The first important observation is that the two classes of problems are dual to each other, as long as the configurations are invertible, and the distinction made above is merely an artificial one motivated by different applications.

The second important observation is that in the implicit domain, direct problems can be informally viewed as computing some accumulation of functional characterization of the template, which is an integral in measure-theoretic terms. On the other

[19] I conjecture that this can be extended to develop multi-scale material representation schemes.
[20] This can be extended to unions by duality through De Morgan's laws.

hand, inverse problems are about computing the overlap of functional characterization of the objects, which is an inner product (i.e., also an integral). Both of these integrations are carried out for different configurations of the same objects, which unifies them under a single operation; namely, convolution.

Claim 1.4.4. In the functional (i.e., analytic) domain, *both direct and inverse problems can be formulated as convolutions*, as long as the pattern motions are invertible and the involved shape and motion descriptors are measurable. More precisely, the queries about the produced structure can be answered by convolving the queries about the template and the pattern.

A caveat is that convolutions—and all measure-theoretic tools in general—are unable to capture overlaps of lower-dimensional features. For example, volumetric convolutions cannot distinguish between collision-free surface or curve contact, and no contact/collision at all. To enable handling features of heterogeneous dimensions, an additional tool needs to be introduced:

Claim 1.4.5. In the functional (i.e., analytic) domain, *interactions of heterogeneous dimensions can be captured via Dirac* $\delta-$*calculus*. More precisely, lower-dimensional features are made artificially measurable by incorporating Dirac $\delta-$functions into the shape and motion descriptors.

It is important to note that the Dirac $\delta-$function is not computable; nevertheless, is provides an abstraction device for theoretical developments, which, in turn, represent a convergence 'limit' to which computational algorithms can approach to any desired precision, depending on the available resolution. For example, a lower-dimensional contact is theoretically measured as the overlap of $\delta-$singular surfaces or

curves, which is made computable in practice by allowing Gaussian dissipations into the surrounding volume. This is significantly more appealing that a combinatorial attempt to characterize contact measures, which is doomed to fail because of the fact that lower-dimensional contact is ill-defined and non-computable.[21]

As with most computational procedures, there is an inevitable trade-off to be made between accuracy and computational resources (i.e., time and memory). Unlike combinatorial methods that usually dwell on the edge of computability, as with the example of lower-dimensional contact, analytic methods are robust and computable, with the additional benefit of providing a meaningful trade-off mechanism.

1.5 Application Domains

The proposed approach transforms how we approach new and classical problems involved in numerous engineering activities. The following are among the main application areas, some of which go beyond traditional mechanical product development:

- **Next Generation of PLM Tools.** In spite of the constant improvements in user interface, agile modeling, collaborative design, and other software capabilities, there are fundamental issues that remain to obstruct the progress of computer-aided design, manufacturing, and engineering (CAD/CAM/CAE) tools in the marketplace, some of the main ones I shall cite in this proposal. The theoretical limitations of the conventional theory of solid modeling—e.g.,

[21] Informally, a function define over a continuum is computable by discrete approximations—which is inevitable as it underlies digital computers and computing models [347]—if small perturbations in the input do not dramatically change the output [105, 106]. If this does not hold even in the theoretical formulation of the problem, the computational algorithms are doomed to fail—as is the case with many combinatorial methods attempting to characterize notions such as lower-dimensional intersection [217] and medial axis (MA) [59].

related to regularization, neighborhoods, computability, etc.—have forced the software companies into appealing to provisional remedies rather than generic solutions. A major contribution of this thesis is to propose an alternative approach that can resolve many of these problems by leveraging existing (but underused) tools from well-established areas of mathematics.

- **Robotics and Autonomous Navigation.** The true power of analytic modeling is revealed when geometry meets kinematics, i.e., when concurrently analyzing or synthesizing shapes and motions. In the past few decades, configuration space modeling has become extremely popular in robotics and spatial planning. Among the most effective and efficient approaches are those which, one way or another, are analytic methods in disguise. The primary setup proposed in this thesis enables a uniform and elegant formulation of many classical motion and path planning problems. The secondary tools presented for comparative analysis are particularly useful for navigation in narrow or crowded environments where many other planners are prone to failure.

- **Design for Additive Manufacturing.** Additive manufacturing (AM) is democratizing fabrication in a variety of domains from rapid prototyping and circuit printing to building customized prosthetics. AM opens up completely new design spaces and material distributions that were not accessible to computerized NC (CNC) machining while it also introduces a whole new set of design constraints, all of which call for revisiting the foundations of our design and modeling apparatus. AM can greatly benefit from the shape analysis methods to be presented in this thesis to incorporate material utilization constraints as early as possible in the computational design process.

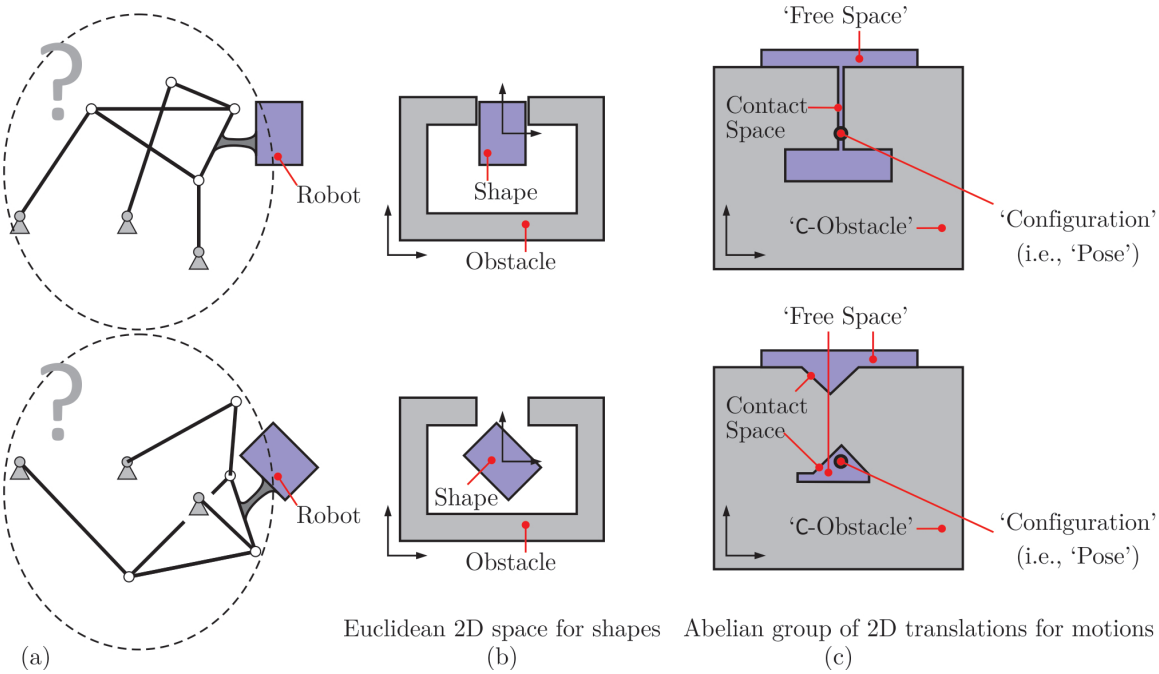

FIGURE 1.5.1: Designing a mechanism (a) in the presence of obstacles requires satisfying non-collision conditions at certain 'precision points' (i.e., configurations or poses) (b) which can be viewed as an inclusion query on the configuration space obstacle (c). The figures in (b, c) are reproduced from [170].

- **Shape and Topology Optimization.** The ability to exploit the efficiency of digital computing technology to explore the design space beyond what was possible before the digital era is to perform numerous iterations in a short amount of time. In doing so, design synthesis and optimization play a significant role, which require machine-readable parameterization of the design space. Whether by using voxelizations, polyhedral mesh, or spline level-set parameterizations, synthesis and optimization techniques are largely reliant on functional representations. The analytic descriptions presented in this thesis provide a sound basis for synthesis and optimization and their query-based (i.e., functional) interoperability with physical analysis, manufacture planning, and other tasks in a modern computational design workflow.

- **Shape Retrieval and Matching.** A large class of problems in computer graphics and shape recognition are attempts to extract compact descriptions of shapes (complete vs. incomplete, exact vs. approximate, rigid vs. flexible, etc.) that could serve as signatures for comparative analysis. The correlations that I shall present in this thesis provide reliable, robust, and fast measures for shape matching that apply to arbitrary shapes, are not impaired by topological, geometric, or syntactic complexity of the problem, and are persistent in the presence of noise and errors.

- **Automatic Assembly Planning.** An automatic identification of mating features in assemblies is one of the important challenges in the grand scheme of the design automation project. It is an important component of virtual reality and digital prototyping, which results in a significant reduction of time and cost associated with physical prototyping and facilitates the elimination of a large

subset of design problems in the earlier stages of the process. The problem can be recast in its most general form as an optimization of shape complementarity measures to be formulated in this thesis.

- **Pharmaceutical Developments.** The use of computational tools has significantly accelerated the expensive and time-consuming drug development process. A large class of drugs are small molecules or medium-sized proteins that bind to the active site of enzymes and inhibit the (mal)function of cellular processes. Geometric protein docking and surface analysis are key to the 'lead compound' development as a scaffold before adding biochemical functional groups. The tools being developed as part of this thesis can enrich the back-end of the computer-aided drug design software in the ultimate war against HIV, cancer, Alzheimer's, Parkinson's, and other deadly diseases.

Figure 1.5.1 demonstrates how the configuration space obstacle (or $\mathsf{C}-$obstacle for short) are generated and used in robot motion planning and mechanism design. $\mathsf{C}-$obstacles will be studied and generalized in depth throughout this thesis.

Figure 1.5.2 illustrates the last two example applications; namely, determination of spatial relations for assembly planning and prediction of molecular binding configurations for rational drug design.

1.6 Summary of Contributions

The paradigm shift advocated in this thesis offers several important advantages that I shall describe in three different categories; namely, theoretical, computational, and practical contributions that follow in Sections 1.6.1, 1.6.2, and 1.6.3, respectively.

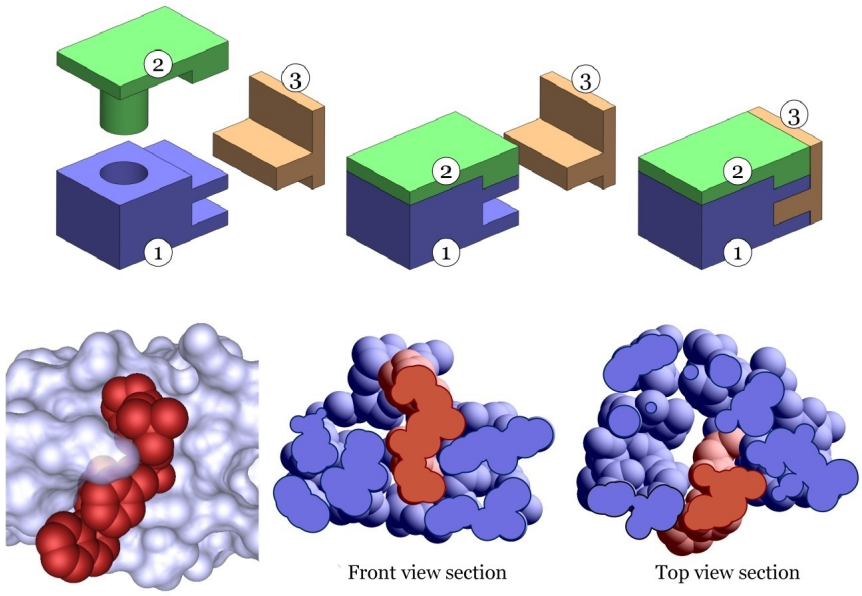

FIGURE 1.5.2: Example applications of analytic methods to geometric modeling and spatial reasoning: assembly planning (recreated from [199]) (top); and protein docking—e.g., HIV protease (enzyme) and saquinavir (drug) complex (PDB Code: 1FB7) [149] (bottom).

1.6.1 Theoretical Benefits

The analytic paradigm described in this thesis provides an elegant reformulation of the most fundamental geometric operations involving shapes and motions that enables:

- *Unification*: It allows unifying seemingly different computational tasks and reveals their dualities—e.g., by showing that numerous operations ranging from topological regularization to morphological operations such as Minkowski operations, sweep and unsweep, configuration space obstacle computation, and alike are uniformly describable in terms of convolutions of functions.

- *Generalization*: It provides insight into how these concepts can be extended to formulate even more powerful tools—e.g., by further customizing the profile of the shape descriptor functions whose sub-/super-level sets correspond to the

traditional set-theoretic descriptions of interacting shapes and motions.

- *Abstraction*: It offers a powerful framework for a theoretical study of the traditionally non-computable concepts by introducing artificial notions (e.g., Dirac $\delta-$calculus), with a natural mechanism to approximate them in practice—e.g., using Gaussian mollifiers to approximate the $\delta-$function.

To elaborate on the first point above, I show that seemingly different geometric operations that have been traditionally addressed by numerous researchers using various ad hoc algorithms—e.g., exploiting certain geometric properties such as polyhedral shape, dimensionality, convexity, etc.—can be unified in an algebra of functions. These operations are formulated in terms of a few geometric queries (namely, inclusion and distance) and their functional combinations (namely, arithmetics, convolution, and composition with the $\delta-$function) in a uniform fashion, regardless of combinatorial, topological, and geometric complexity, and without appealing to the properties of any particular computational representation. This has tremendous implications in terms of interoperability of computational systems and representations [147, 148], if the shape and motion descriptor functions are viewed as an ability of a representation device to respond to spatial queries.

To elaborate on the second point above, the uninform formulation also enables one to systematically extend and generalize the aforementioned operations to obtain more flexible computational tools. For example, characteristic functions that conceptualize inclusion queries can be replaced with general nonnegative real-valued defining functions with customized profiles to explore new geometric operations that subsume the known ones as special cases, or to come up with a new class of operations that traditional abstractions could not express. This has significant implications to

transform the design process from one that deals with solid objects (using simple in/on/out queries) to those that can characterize material distributions, probabilistic or deterministic micro-structure, shape and material uncertainty, and other spatial distributions pertaining to geometric and physical modeling.

Although measure-theoretic tools have been used for geometric modeling in recent works [215, 217, 218], a major limitation of such an approach has been recognized as its inability to capture lower-dimensional geometric entities with a single Lebesgue 3−measure. A major contribution of this thesis is to introduce a Dirac δ−calculus to implicitly describe lower-dimensional features (e.g., discrete points, curve segments, and surface patches) and assign them with artificial 'lumped' volumetric densities (e.g., per point, per unit length, and per unit area, respectively). To the best of my knowledge, this thesis is the first attempt in which lower-dimensional spatial elements are brought into the scene by adding convolution and composition with the δ−function among the algebraic operators on functional representations. The major advantage of this abstraction is the route it provides to turning non-computable tasks involving lower-dimensional features into computable approximations (in a meaningful trade-off) by dissipating the crisp geometric notions (e.g., a 2D surface in 3D, viewed as a δ−singularity) into fuzzy distribution forms (e.g., a Gaussian distribution).

1.6.2 Numerical Benefits

The advantages of the proposed paradigm shift go beyond theoretical elegance. Among the main contributions is providing a systematic framework to develop accurate and efficient algorithms whose effectiveness are independent of the combinatorial, topological, and geometric properties of the particular shapes and motions. Moreover, such

a framework provides several powerful approximation mechanisms—arising naturally from the formulation—for meaningful and controllable trade-offs between robustness and/or efficiency on the one hand, and the accuracy of the results, on the other hand:

- *Robustness*: The analytic (i.e., functional) paradigm provides computable approximations to the traditionally non-computable geometric constructs in the presence of errors/noise using several powerful mechanisms; for instance,

    1. *Mollification* can be viewed as approximating the $\delta-$function, used to characterize lower-dimensional spatial elements as volume singularities, using one of its limit representations (e.g., Gaussian kernel). It can be thought of as a 'dissipation' of lower-dimensional manifolds to impart robustness in the presence of geometric errors.

    2. *Anti-aliasing* can be achieved by convolution with a smooth filter (e.g., Gaussian kernel) followed by low-pass filtering—i.e., truncating the band-unlimited Fourier representations—of the functional descriptors of shapes and motions, eliminating their high-frequency fluctuations pertaining to geometric features of negligible size.

- *Efficiency*: The analytic (i.e., functional) paradigm supports fast (e.g., time-optimal) numerical computations by leveraging Fourier transforms when computing the fundamental measure-theoretic operators (e.g., inner products, convolutions, etc.) from which the rest of the operations are composed; namely,

    1. *Equispaced* (*i.e., ordinary/uniform*) fast Fourier transforms (FFT) can be used to efficiently implement convolutions in real-time on highly paral-

lel many-core GPUs, seamlessly leveraging high-performance computing (HPC) technologies (i.e., by 'plug-and-play') which are becoming more powerful and affordable every day.

2. *Nonequispaced* (*i.e., nonuniform*) fast Fourier transforms (NFFT) can be justified for problems in which uniform sampling (required for the ordinary FFTs) entails a nonoptimal use of computer memory, while nonuniform sampling could reduce the sample size at a smaller cost of slowing down the processing per unit sample size.

3. *Low-pass filtering* can also be used as an effective speed-up mechanism with a small and controllable compromise of accuracy, which enables one to choose the number of dominant modes of the shape and motion signals in the frequency domain depending on the available time budget.

Recall from Section 1.3 that traditional methods are intrinsically limited in addressing robustness issues when dealing with interactions of lower-dimensional manifolds. Some of the simplest operations are non-computable—e.g., computing the surface contact area of non-colliding objects, computing the medial axis of a general shape, etc.—as a small geometric perturbation due to errors/noise can lead to prohibitive changes in the outcome. The analytic approach overcomes these problems using the aforementioned mechanisms that trade off the problem's sensitivity to such errors with the confidence with which the results are computed.

On the other hand, the modularization of computing in terms of a few measure-theoretic tools such as inner products and convolutions allows one to break down algorithm design into a composition of such operations, which, in turn, can leverage (equispaced or nonequispaced) FFTs that are widely available as optimized CPU-

and GPU-based libraries. Although equispaced (i.e., ordinary) FFTs [76] have been used extensively in robotics [88, 174] and more recently in geometric modeling [214], nonequispaced FFTs [258] have found few applications in applications such as protein docking [25, 57]—where the advantages of nonuniform sampling (e.g., at molecule's atom centers) are more obvious than the case of arbitrary shapes.

To exploit "the best of both worlds", a spherical decomposition scheme is presented in Chapter 4 to represent functional descriptions in discrete forms (via radial basis functions (RBF)) and formalize their convolution algebra and harmonic analysis in terms of the RBF 'knots'. This approach enables tremendous benefits by allocating the time and memory resources more effectively to capture the geometric details near the boundary, as opposed to uniform grid-based sampling that essentially wastes numerous sample points in the interior. Thus a major contribution of this thesis is to reconcile the FFT-based analytic (i.e., functional) methods—which have been viewed for a long time to have the disadvantage of having to deal with uniform grids or voxmaps to leverage ordinary FFTs–with efficient hierarchical combinatorial representations at multiple levels-of-detail (LOD) (e.g., sphere-trees).

1.6.3 Pragmatic Benefits

A large fraction of the upcoming chapters is devoted to demonstrating how classical shape and motion related problems can be uniformly posed and efficiently solved in the analytic (i.e., functional) paradigm—e.g., for integration, topological regularization, boundary evaluation, morphological combinations, and skeletonization discussed in Chapter 3. These operations are mathematically well-defined and well-understood, even though their computational implementation varies in difficulty when using tra-

ditional explicit descriptions and combinatorial algorithms. In addition to demonstrating how implicit descriptions and analytic techniques provide more effective and efficient solutions to this class of problems, another main contribution of this thesis is the systematic approach it provides for formulating and solving harder problems that are not as well-posed. For example, on the one hand, many problems in the first class can be broken down into detecting volumetric *collision* and surface *contact* between objects of arbitrary shape, which are mathematically well-defined even though their accurate and efficient implementations are nontrivial and subject to extensive ongoing research. These problems range from morphological synthesis (i.e., shaping with motion) via sweeps and unsweeps [159, 163], or more generally, configuration products and quotients [235] to spatial planning [210] and evaluating shape skeletons [17], all of which can be addressed in the analytic (i.e., functional) paradigm using a few geometric queries (namely, inclusion and distance) and their functional combinations (namely, arithmetics, convolution, and composition with the $\delta-$function). On the other hand, there are more difficult problems (e.g., in comparative shape analysis) that deal with detecting and quantifying approximate shape *similarity*, *complementarity*, and *symmetry* for objects of arbitrary shape, for which there is no single rigorous mathematical formulation. My goal is to demonstrate that the proposed paradigm is also powerful in quantifying these notions, which are central to formulating and solving problems such as assembly planning [35, 37] and protein docking [25, 69]. In essence, I show that the heuristics used to define such measures can be summarized into the choice of more sophisticated shape descriptors—i.e., functions that locally capture more shape information than inclusion and distance queries—while the fundamental algebraic operations combining the descriptors (e.g., convolutions) have fairly universal and non-controversial semantics.

1.7 Outline of the Thesis

1.7.1 Organization

This thesis is organized into 5 chapters and an appendix, overviewed below:

- Chapter 1 (the current chapter) serves to introduce the reader with the existing challenges and future directions, and how this thesis positions itself in contributing to the creation of knowledge and advancement of the field. In Section 1.1 I overviewed the historical motivation for the development of the foundations of CAD/CAE/CAM as practiced today. Section 1.2 provided the basic terminology (e.g., referent, model, and representation) and the distinction between explicit and implicit descriptions. Section 1.3 described some of the most important classical challenges with crippling effects on the advancement of the field as well as newer challenges arising with the ever growing complexity of design for modern materials and fabrication processes, and described how the alternative paradigm outlined in this thesis provides a range of possibilities to address them. In Section 1.4, I described the proposed paradigm in detail from both philosophical and practical perspectives and demonstrated its benefits in the context of an important exemplar problem (namely, protein docking). In Section 1.5, I provided more CAD/CAE/CAM application areas that can benefits from the proposed approach from a higher-level view. The main contributions of the thesis were enumerated in Section 1.6.

- Chapter 2 elaborates the paradigm shift from the traditional set-theoretic (i.e., explicit) to a more powerful measure-theoretic (i.e., implicit) approach for mathematical abstraction of engineering problems (Section 2.1). In Section 2.2 I

present the preliminaries (postulates and definitions) including pointset-topological and group-theoretic foundations of solid and configuration modeling as well as basic concepts from harmonic analysis that will be useful in the subsequent chapters. In Section 2.3 I present the analytic formulation in terms of characteristic functions, nonnegative real-valued defining functions, and their extension to special functions that included singularities of different orders (to model 'lumped' volume along lower-dimensional spatial elements) using Dirac $\delta-$function. A modified $\delta-$calculus is presented with an emphasis on the sifting properties of the $\delta-$function to provide the building blocks of geometric operations to be presented in the next chapter.

- Chapter 3 presents some of the most important fundamental geometric operations reformulated in measure-theoretic terms using the analytic building blocks developed in the previous chapter (Section 3.1). These operations include computing integrals over domains of various dimensionalities (Section 3.2), topological operations such as regularization and boundary evaluation (Section 3.3), morphological operations such as Minkowski operations, dilation and erosion, closing and opening, offsetting and blending, sweeping and unsweeping, configuration space obstacle computations, and alike (Section 3.4), and last but not least, skeletonisation (Section 3.5), whose implicit formulation was the most technically challenging contribution of this thesis. All operations are first presented in set-theoretic (i.e., explicit) formulation, along with some literature review on the state-of-the-art and its modes of failure, followed by a measure-theoretic (i.e., implicit) formulation (e.g., using convolution algebra and harmonic analysis) along with its offered computational advantages.

- Chapter 4 presents a discretization scheme for representing and computing on the functional descriptions of shapes and motions (Section 4.1) using an efficient spherical decomposition that lends itself naturally to the convolution algebra and harmonic analysis (Section 4.2). In Section 4.3 I present the formalism for discretization of an arbitrary shape as a countable union of balls, its interpretation as a Minkowski sum, and its analytic description as a convolution. In Section 4.4 I formulate correlation predicates in terms of Minkowski sums across multiple shapes together with their convolution forms, and use the results from the previous section to carry the discretization into the configuration space. In Section 4.5 the results of effcient GPU implementations are demonstrated to outperform some of the state-of-the-art methods in those areas. Although the development is independent of the sampling algorithm, a particular such algorithm is presented in Section 4.6 whose performance is shown to outperform an existing popular 'sphere-packing' algorithm.

- Chapter 5 serves to showcase some of the engineering applications of the developed methods. It starts off with a discussion of fuzzy modeling and density functions in Section 5.1, viewed as the dissipated approximations to the singular shape descriptors formulated earlier. These density functions are obtained by substituting the Dirac $\delta-$singularities with their limit representations to make them computable for a particular choice of geometric resolution. The application of these density functions in solving fundamental problems are demonstrated in the following sections, along with their unique benefits in contrast to the traditional methods, the tremendous new horizons they enable for CAD/CAE/CAM advancement, and their current limitations that open up new

research opportunities. In particular, Section 5.2 provides insight into the development of more efficient collision detection and contact measure algorithms for real-time (e.g., virtual reality) applications on a limited rendering budget. Section 5.3 extends those ideas to quantifying shape similarity and complementarity of arbitrary shapes with numerous applications in assembly planning and protein docking. A particularly important application in haptic-assisted virtual prototyping and assembly is presented in Section 5.4. Section 5.5 presents how powerful synthesis tools can be developed to explore the design space using rich functional shape descriptors. Section 5.6 shows how versatile planning tools can be created based on assessing manufacturability of complex shapes (using additive or subtractive fabrication processes) and providing feedback to the design process. Last but not least, the tremendous benefits of the analytic (i.e., functional) paradigm in making systems of different representations with hidden proprietary internal structures (i.e., 'black boxes') work together in the context of query-based interoperability are briefly hinted in Section 5.7.

- The final chapter closes with the concluding remarks and future directions.

- Appendix A presents a quick overview of the group-theoretic and metric properties of the configuration space of rigid motions (i.e., rotations and translations) that are fundamental to the concepts presented (particularly in Chapter 3).

1.7.2 Prerequisites

Before moving on, it is helpful to clarify what this thesis is and is not meant to be.

The target audience for this document are engineers and researchers, particularly in mechanical and industrial engineering, with a keen interest in CAD/CAE/CAM, model-based systems engineering (MBSE), geometric modeling, computing, and reasoning, and more generally, professionals who are interested in applying computational tools in mechanical systems design.

I will assume the reader is comfortable with basic concepts from set theory, function analysis, and calculus, and has minimal familiarity of pointset and algebraic topology, differential geometry, measure theory, convolution algebra, and harmonic analysis, without necessarily having a deep or detailed proficiency. Although a quick skimming through the pages will give away the ubiquitous use of commutative diagrams, I am not using (or assuming the reader's familiarity) with category theory and related abstractions.

I have done my best to find a proper balance between rigor and clarity. Compared to most engineering thesis, this one may appear more "contaminated" with mathematical formalism, which I find inevitable in my quest to emphasize the importance of the shift in formulating, abstracting, and thinking about geometric problems in functional terms. However, this is *not* a thesis in mathematics and should not be judged to a level of scrutiny that a mathematical technical article should sustain. It is very likely that a keen mathematician can find errors of rigor in the presentation—some intentional, in the interest of simplicity, others not. Fixing all of them would make the document prohibitively long[22] and dilute the practical message.

[22]Not to mention that when writing this sentence for the first time, my estimate of the total page number was nearly half of what I ended up with!

1.7.3 Diagram Notations

The commutative diagrams of the form presented in Section 1.4 are central to the presentation of this thesis. The claim that a given diagram commutes implies that all paths in the diagram are essentially equivalent, i.e., for all objects in a source space, for which two different paths can be completed to arrive at two objects in a destination space, the resulting objects must be the same. For example, the diagram:

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\qquad f_1 \qquad} & C_1 \\ \Big\downarrow g & & \Big\downarrow h \\ A_2 \times B_2 & \xrightarrow{\qquad f_2 \qquad} & C_2 \end{array} \tag{1.7.1}$$

means that the specific functions $f_1 : (A_1 \times B_1) \rightarrow C_1$, $f_2 : (A_2 \times B_2) \rightarrow C_2$, $g : (A_1 \times B_1) \rightarrow (A_2 \times B_2)$, and $h : C_1 \rightarrow C_2$ are related via $(h \circ f_1) = (f_2 \circ g)$ where "$\circ$" is the function 'composition' operator. In other words, for every pair of elements $(a_1, b_1) \in (A_1 \times B_1)$, if $c_2, c_2' \in C_2$ such that $c_2 = h(f_1(a_1, a_2))$ and $c_2' = f_2(g(a_1, a_2))$ then $c_2 = c_2'$. The diagram may carry even more information, for example:

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\qquad f_1 \qquad} & C_1 \\ g^{-1} \Big\uparrow \Big\downarrow g & & \Big\downarrow h \\ A_2 \times B_2 & \xrightarrow{\qquad f_2 \qquad} & C_2 \end{array} \tag{1.7.2}$$

also has two other paths; namely, in addition to what was implied by (1.7.1):

- $(g^{-1} \circ g) = \mathsf{id}_{A_1 \times B_1}$, i.e., $(g^{-1}(g(a_1, b_1)) = (a_1, b_1)$ for all $(a_1, b_1) \in (A_1 \times B_1)$;
- $(g \circ g^{-1}) = \mathsf{id}_{A_2 \times B_2}$, i.e., $(g(g^{-1}(a_2, b_2)) = (a_2, b_2)$ for all $(a_2, b_2) \in (A_2 \times B_2)$;

- $(h \circ f_1 \circ g^{-1}) = f_2$, i.e., $h(f_1(g^{-1}(a_1, b_2))) = f_2(a_1, b_2)$ for all $(a_2, b_2) \in (A_2 \times B_2)$;

Two of which are sufficient to conclude the other two. At times, I will make use of special arrows that mean different things depending on the context, which are clarified on a case-by-case basis. For example, something that looks like this:

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\qquad\qquad f_1 \qquad\qquad} & C_1 \\ \vdots & & \| \\ \vdots\ g & & \| \equiv \\ \downarrow & & \| \\ A_2 \times B_2 & \xrightarrow{\qquad\qquad f_2 \qquad\qquad} & C_2 \end{array} \tag{1.7.3}$$

may imply by the dashed arrow that $g : (A_1 \times B_1) \to (A_2 \times B_2)$ is unknown, under investigation, or loosely implied—e.g., when using limit representations $g(a_1, b_1) := \lim_{(\epsilon_1, \epsilon_2) \to (0,0)} (a_1, b_1)$ with loosely specified semantics—or that its commutativity with the rest of the diagram is a conjecture, etc. The equality-like arrow implies equivalence, i.e., $f_1(a_1, b_1) \equiv f_2(g(a_1, b_1))$ for all $(a_1, b_1) \in (A_1 \times B_1)$, whatever the binary equivalence relation "$\equiv$" $\in (C_1 \times C_2)$ is supposed to mean.

There will be times at which I will use the notation $f(\cdot)$ and f interchangeably, where $(\cdot)$ implies a placeholder for an argument. For example, $g = (\cdot + 5)$ and $g(\cdot) = (\cdot + 5)$ both imply that $g(x) = (x + 5)$ for all $x \in \mathbb{R}$ when $g : \mathbb{R} \to \mathbb{R}$ is defined in terms of "$+$" $\in (\mathbb{R} \times \mathbb{R})$.[23] Thus the following diagram:

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\qquad\qquad f_1 \qquad\qquad} & C_1 \\ \vdots & & \vdots \\ \vdots\ g(\cdot) & & \vdots\ h(\cdot) \\ \downarrow & & \downarrow \\ A_2 \times B_2 & \xrightarrow{\qquad\qquad f_2 \qquad\qquad} & C_2 \end{array} \tag{1.7.4}$$

[23] If the reader is familiar with Church's $\lambda-$calculus [71], this is equivalent to the $\lambda-$notation: $g := \lambda x.(x + 5)$, meaning: "substitute all symbols "x" with expression "$(x + 5)$."

has the exact same meaning as (1.7.1). To extend this, $f(\cdot_1, \cdot_2, \cdot_3, \cdots)$ can be used for functions with multiple arguments, e.g., $f = \frac{(\cdot_1)+1}{(\cdot_2)-1}$ and $f(\cdot_1, \cdot_2) = \frac{(\cdot_1)+1}{(\cdot_2)-1}$ both mean $f(x_1, x_2) = \frac{x_1+1}{x_2-1}$ for all $(x_1, x_2) \in (\mathbb{R} \times \mathbb{R})$.[24] Thus the following diagram:

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\quad f_1(\cdot_1, \cdot_2) \quad} & C_1 \\ \Big\downarrow g & & \Big\downarrow h \\ A_2 \times B_2 & \xrightarrow{\quad f_2(\cdot_1, \cdot_2) \quad} & C_2 \end{array} \tag{1.7.5}$$

also has the same meaning as (1.7.1) and (1.7.4). However, the diagram

$$\begin{array}{ccc} A_1 \times B_1 & \xrightarrow{\quad f_1(\cdot_2, \cdot_1) \quad} & C_1 \\ \Big\downarrow g & & \Big\downarrow h \\ A_2 \times B_2 & \xrightarrow{\quad f_2(\cdot_1, \cdot_2) \quad} & C_2 \end{array} \tag{1.7.6}$$

has a completely different meaning, since it implies an order swap of the tuple $(a_1, b_1) \in (A_1 \times B_1)$ before the function $f_1 : (A_1 \times B_1) \to C_1$ is applied, i.e., $(g \circ f_1)(\cdot_2, \cdot_1) = (f_2 \circ h)(\cdot_1, \cdot_2)$ which is the same as $(g \circ f_1)(\cdot_2, \cdot_1) = (f_2 \circ h)$, meaning $(g \circ f_1)(b_1, a_1) = (f_2 \circ h)(a_1, b_2)$ for all $(a_1, b_1) \in (A_1 \times B_1)$. The notation with explicit arguments $(\cdot_1, \cdot_2, \cdot_3, \cdots)$ is only used when specifying the order becomes necessary to avoid potential confusions, while it is avoided as much as possible when the order is inferrable from the context.

[24] Once again, this is similar to the multivariate $\lambda-$notation $f := \lambda x_1 x_2 . \frac{x_1+1}{x_2-1}$.

Chapter 2

A Paradigm Shift

> "The way to improve this is to make a further abstract picture of our first picture of the problem, which eradicates its bias and retains only its abstract structural features; this second picture may then be examined according to precisely defined operations, in a way not subject to the bias of language and experience."
>
> Christopher Alexander, 1964 [9]

2.1 Engineering Abstractions

The demonstrative "this" at the beginning of the above epigraph refers to the conceptual abstraction of the engineering problems in general—and the engineering design problems in particular within the discussion scope of the cited source [9]. Such a mental picture, often vague and biased from the very moment of its conception, needs to be framed precisely into mathematical formalism. Once articulated in a clear and reliable mathematical language, computational models and representations are

developed to take advantage of computers for efficient processing. In his landmark technical series on principles of solid modeling [275, 278], Ari Requicha[1] opened with the following definition of computational modeling in general—with a subjective interest in the modeling processes for mechanical parts and assemblies:

"Computational models of real-world entities seek to capture abstractly certain aspects of the real entities in a form that it is suitable for computing a class of properties (i.e., for answering a class of questions about) the entities being modeled. The ultimate validity of such models must be ascertained experimentally by comparing the results of measurements performed in the real world." [275]

Mathematical abstraction (i.e., modeling) lies at the center of all scientific and engineering activities—and at the focal point of this thesis. Modeling can be viewed in general as a mapping from the universe of physical entities—or more precisely, from the mental picture of the class of entities in the scientist's or engineer's mind—to an appropriate mathematical structure that captures their relevant properties. The goal is to support subsequent mathematical operations on the latter to accurately and reliably predict the physical phenomena on the former—or more precisely, to match the interpretation of the experimentalist observing the phenomena. Obviously, every such mapping between the physical and mathematical realms exhibits two fundamental properties as a direct consequence of its subjective formation process; namely, only

[1] Aristides A. G. Requicha is a Professor of Computer Science and Electrical Engineering at the University of Southern California (USC), one of the pioneers of solid modeling as we know and practice it today [351, 352], and the first recipient of the first Bézier Award in 2007. He joined the production automation project (PAP) [353] at the University of Rochester in 1973 to develop the mathematical foundations for modeling and representation of mechanical parts and assemblies [202, 203, 275–282, 284, 285, 301, 343], making him one of the 282 highly cited researchers in computer science during 1980-1999 according to ISI Web of Knowledge. A historical review of PAP can be found in [353].

a limited subset of the attributes of the original entity is covered by the model (i.e., reduction) which is intended for a particular purpose that is in demand at a specific time and under specific conditions (i.e. pragmatism) [333].

As the demands on the engineering practice grow over time, the models need to be enriched to cover a larger collection of attributes and with better precision. The natural way to do this—and a more economic one, at least in the short-run—is to improve the existing model and augment it with additional (e.g., generalized or extended) properties. However, the initial bias built into the model due to the arbitrary decisions, limited perceptions, and specific requirements at the time persist to cripple the effectiveness of the model in a subtle fashion. Disruptive and bold paradigm shifts, despite experiencing strong academic and industrial resistance (e.g., due to economic, safety, and reliability concerns), have a chance to offer a leap of performance that is hard to achieve with incremental adjustment.

The rest of this chapter is dedicated to formalizing a paradigm shift in geometric modeling.[2] A brief review of the fundamental postulates and definitions from classical solid and configuration modeling is given in Section 2.2. I present an analytic reformulation in Section 2.3 along with the mathematical ingredients to support shifting to an analytic abstraction paradigm, before diving into the specific geometric and spatial problems in Chapter 3.

[2]I am indebted to Vadim Shapiro for the long discussions on the current state of engineering abstractions and future directions in geometric and physical modeling, which have influenced my thoughts when developing the rest of this chapter. The title of this section was inspired by his 2015 Bézier Award recipient keynote address titled "Engineering through Abstractions" [319] in the GD/SPM'2015 conference.

2.2 Preliminaries

In this section I present the preliminary concepts which are imperative to the development of the remainder of this chapter. Section 2.2.1 reviews some of the fundamental ideas that underly the theories of solid and configuration modeling as we know them today, and iterates over the basic requirements and alternative approaches to fulfill them. Section 2.2.2 presents basic definitions based on concepts from semianalytic geometry, Lie groups and Lie algebras, and other ingredients of classical solid and configuration modeling.

2.2.1 Basic Postulates

The early focus in solid modeling was on geometric specification of mechanical parts and assemblies to support the computerized automation of manufacturing at the time. In addition to the basic requirements for physical realizability—e.g., bounded size, well-defined 3D interior, well-behaved 2D boundary, etc.—a few additional assumptions were made to simplify the computational modeling—e.g., rigidity, homogeneity, isotropy, etc.—which are, as argued earlier in Section 1.3, no longer sufficient to keep up with the advances in manufacturing technology. These basic assumptions led to the assertion of a few postulates [316] that I describe below in informal terms before presenting precise definitions in the upcoming sections:

Review: Basic Postulates of Traditional Solid Modeling:

The following are typically asserted as the basic postulates of solid modeling [316]:

- *Boundedness* is imposed by requiring the solid to be contained in some bounding volume (e.g., ball or box) of finite dimensions in the Euclidean metric space.
- *Rigidity* is formalized in terms of *congruence* relation; i.e., a rigid body is defined as an equivalence class of congruent solids that can be mapped to one another by a transformation that preserves lengths and signed angles (i.e., a Euclidean isometry). This allows one to conceptualize an object's identity independently from the choice of a particular observer (e.g., affine coordinate system).
- *Structural homogeneity* is implicitly incorporated by assigning no distinctive identity or additional properties (e.g., density, orientation, etc.) to the individual points of the continuum solid, which implies material homogeneity and isotropy. Every material point has only one property—as far as solid modeling is concerned—which is its inclusion in the set, distinguishing it from the non-material points outside.
- *Dimensional homogeneity* is guaranteed by enforcing the topological notion of *regularity*. Informally, this means that the model of a solid object comprises a 3D open interior—with all points being surrounded by full neighborhoods—covered with a tight 2D 'skin', and cannot have artifacts such as 'dangling' faces or edges, singular (inclusion or exclusion) points, and other physically unrealizable features [275].
- *Finite triangulability*—and consequently *finite describability*—is ensured by requiring a solid to be a curved subpolyhedron of the Euclidean space, i.e., built up from a finite selection of curved tetrahedra in an embedded simplicial decomposition of the entire space. This eliminates the possibility of pathological behavior at the solid's boundary (e.g., 'wild' embeddings) [275].
- *Algebraic closure* of the modeling space under an appropriate set of mathematical operations that characterize the basic building blocks for modeling desired processes. In particular, operations similar to the set-theoretic Boolean combinations are important to abstract the manufacturing processes involving addition and removal of material as well as interference tests [275].

Some of the foregoing requirements perhaps deserve some explanation.

The boundedness condition is sometimes removed—at least for modeling, to be brought back when speaking of computer implementations—to enable closure under set complement operation, hence the formation of Boolean algebras of geometric sets (detailed in Section 3.4).[3] This in turn has significant implications, e.g., to establish dualities between important operations characterizing direct and inverse problems.[4]

The continuity assumption is implicit in the very choice of Euclidean topology as the modeling medium, where only macroscopic geometric information is being captured.[5] Consequently, 'informational completeness' can be achieved (at least in theory) by holding the computer representation to answer a single question about inclusion or exclusion of a given spatial point against a solid [275].[6]

The rigidity assumption can be relaxed by including more general invertible transformations to define the equivalence classes of solids that are closed under their actions (detailed in Section 3.4).[7] One way or another, such classifications follow from a group-theoretic formalism of configuration modeling (i.e., motion and kinematics)

[3]The exclusion of unbounded sets from the modeling space brings with it the exclusion of the set complement operation $\mathbf{c}(S) = (\mathbb{R}^3 - S)$ from the algebraic structure it admits; namely, turning the 'Boolean algebra' into the (less powerful) 'Boolean ring' [276].

[4]For example, the Minkowski sum and difference are related via $S_1 \ominus S_2 = \mathbf{c}((\mathbf{c}S_1) \oplus S_2)$.

[5]Microscopic geometric information (e.g., material microstructure, porosity, crystal structure, etc.) are typically assigned as annotations to the different cell decompositions of the solid, which are symbolic accessories outside the scope of traditional solid modeling.

[6]This is the most fundamental postulate whose violation for a greater good is the main theme of this chapter. It is desirable to develop models of geometric objects (i.e., shapes and motions) whose elements (i.e., points and configurations) have more properties than the basic binary-valued inclusion, making those objects richer in information than what is traditionally captured by the set-theoretic abstractions.

[7]The idea is to assign a conceptual identity to different instances of the same geometric 'form' in different 'configurations'. For example, rather than using the elements of SE(3) to define an equivalence relation (i.e., congruence) on $\mathbb{R}^3$, one can obtain a more 'flexible' equivalence relation using the elements of GA(3) to allow more (though still finite) degrees of freedom (DOF) for modeling global deformations. However, general models of local deformations require allowing for spatial distributions of transformations (i.e., infinite DOF), which are beyond the scope of this thesis.

that has been considered somewhat peripheral to the core idea of solid modeling (i.e., shape and structure).

The next two properties, namely, dimensional homogeneity and finite triangulability are often collectively referred to as the *solidity* requirement and form the basis for what we typically refer to as 'solid modeling' [279, 280, 283, 284, 351, 352], which is discussed below at length. In addition to eliminating pathological behavior at the boundary and ensuring that there exists a physically realizable referent for the modeled object, the finite describability has another important implication, as its name implies, to certify the feasibility of constructing finite symbol structures for computer representation.[8] Triangulability also ensures *measurability* which is of particular relevance to the methods developed in this thesis—meaning that curve, surface, and volume integrals (i.e., Hausdorff d−integrals) are well-defined. In principle, they can be computed by decomposing their domains into the curvilinear simplicial cells, each of which is in correspondence (via the same tame embedding) with multivariate parametric integrals (i.e., Riemann d−integrals) over Euclidean d−spaces for $0 \leq d < 4$.

Finally, algebraic closure refers to the requirement that whatever abstraction is used to model solids (e.g., as semianalytic sets, polyhedra, r-sets, etc.) it must be closed under algebraic operations that model geometric modifications and transformations that model design and manufacturing processes such as material addition/removal, deformations, and alike, to which I shall return in Chapter 3.

[8] Although finite triangulability implies finite describability in an almost trivial sense—namely, by explicitly representing the model via its native simplicial complex representation—the converse is not necessarily true in the most strict sense of the word; i.e., one might be able to come up with another consistent set of rules to construct a representation scheme for restricted classes of nontriangulable sets—see pp. 8, 9 of [275].

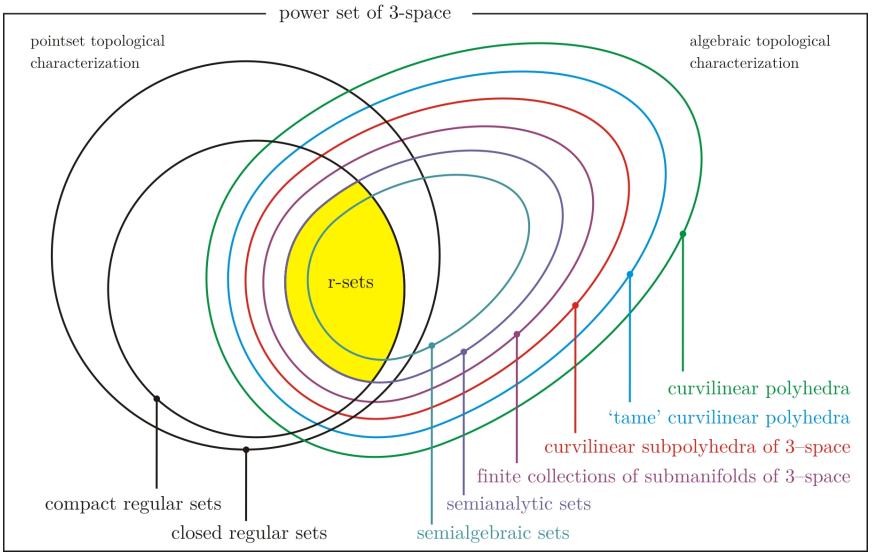

FIGURE 2.2.1: Requicha defined r-sets [275] as compact (i.e., bounded and closed) regular and semianalytic subsets of the Euclidean 3−space endowed with the usual topology.

2.2.1.1 Characterizing the Notion of 'Solidity' for Solid Modeling:

The notion of solidity can be characterized mathematically in more than one way, a few of which are described below and illustrated schematically in Fig. 2.2.1.

On the one hand, the classes of closed or open regular sets[9]—defined using the pointset topological notion of regularity, without necessarily introducing a metric or specifying a particular analytic parameterization of the shape profile—are too broad to capture the notion of solidity. Regular sets are guaranteed, by definition, to be homogeneously 3D and rule out the nonphysical artifacts such as 'dangling' faces or edges, singular points, etc. However, they are not closed under set-theoretic Boolean operations, which has led to the need for the introduction of 'regularized' Boolean

[9]There are two principal ways in which regular sets can be defined, which are dual to each other; namely, *closed regular* and *open regular* sets (Section 3.3.1). Either model can be used to construct a theory of solid modeling; while the former was chosen to develop solid modeling as we know it today, the latter (proposed in [16] for an alternative solid modeling) did not catch on. Hereafter, when I refer to regular sets, I means closed regular sets unless otherwise specified.

operations [343]. Moreover, regular sets are not necessarily finitely describable in some suitable sense and their boundaries can be ill-behaved.

The first difficulty can be resolved by further restricting them to topological polyhedra.[10] Polyhedral objects can also be open, closed, or neither, and can be irregular in general. However, the notion of regularity—which is hard to put into computable form in its pointset topological definition—can be captured exactly via combinatorial reasoning on the underlying cell decomposition of the polyhedra; namely, by ensuring that every lower-dimensional cell (i.e., face, edge, or vertex) of the cell complex is adjacent to a 3D cell (i.e., polyhedron) [275].[11] However, compact (i.e., bounded and closed) polyhedra are not closed under Boolean operations, regardless of whether they are regular and/or the set operations are regularized.

The second difficulty arises due to the so-called 'wild' embeddings that can give rise to nonphysical pathologies of a different kind than the one eliminated by regularity.[12] Although characterizing the suitable conditions on embeddings is nontrivial, an adequately 'tame' embeddings is one that is a restriction of a homeomorphism of the *entire* 3−space onto itself. In other words, a solid is a so-called (curvilinear) 'subpolyhedron' of $\mathbb{R}^3$, i.e., a (curvilinear) polyhedron constructed from a subcomplex of a (curvilinear) cell decomposition of the 3−space.

[10]Topological polyhedra (also known as triangulated spaces) are typically viewed as curvilinear polyhedra specified by an embedding of rectilinear polyhedra using structure-preserving elastic deformations (i.e., homeomorphisms).

[11]It is easy to show that a compact polyhedron is a regular set (in pointset topological terms) iff it is a homogeneously 3D polyhedron in embedded in the 3−space. In other words, in this restricted sense the pointset topological and algebraic topological notions of regularity are equivalent, with the latter having the extra benefit of computability.

[12]A famous example is 'Alexander's horned sphere' [11], which is a topological sphere that bounds a topological ball (i.e., homeomorphic to a tetrahedron) composed of countably infinite union of compact and simply connected sets. Although the embedding is a homeomorphism, the resulting shape is pathological and does not point to a physically realizable object, hence should not be classified as a solid [221].

On the other hand, the class of semianalytic sets[13]—whose definition requires going beyond pointset topology and appealing to a parameterization of the space using an affine coordinate system—satisfies all of the postulated requirements except the regularity condition. By definition, semianalytic sets are closed under the (set-theoretic or regularized) Boolean operations. They include the subclass of semialgebraic sets[14] as well as all sets represented by polynomial and rational equalities and inequalities [316]—which are versatile enough to include simple objects with machinable faces (e.g., planar, cylindrical, spherical, conical, etc.) as well as sculptured objects with higher-dimensional polynomial or rational surface patches (e.g., Hermite, Bézier, B-splines, NURBS, etc.) [275]. In addition, regularization, projection, and connected components of a semianalytic set are all semianalytic, and bounded semianalytic sets are also subpolyhedra of $\mathbb{R}^3$ (thus triangulable) [316]. However, semianalytic sets can be open, closed, or neither; and can be heterogeneous in dimensions of their different pieces [316].

Another approach is to define solids as finite collections of 3D submanifolds of the 3−space,[15] which are regular as a result of definition. It is known that compact submanifolds of $\mathbb{R}^3$ are also finite polyhedra, i.e., are finitely triangulable. However, they are not closed under the (set-theoretic or regularized) Boolean operations.

Putting these all together, Requicha [275] defined an 'r-set'—used synonymously

[13]Semianalytic sets are defined to include all sets that can be constructed by a finite sequence of Boolean combinations of equalities or inequalities of analytic functions. Here, an analytic function refers to a real function that admits a (possibly infinite) Taylor series expansion at all points in the 3−space. Note that the term 'analytic' in this context has almost nothing to do with its more prevalent usage throughout the thesis to classify the proposed methods.

[14]Semialgebraic sets are defined to include all sets that can be constructed by a finite sequence of Boolean combinations of equalities or inequalities of polynomial functions.

[15]A submanifold of $\mathbb{R}^3$ (or a 3D manifold embedded in $\mathbb{R}^3$) is defined as a subset of $\mathbb{R}^3$ whose points have neighborhoods that are either homeomorphic to an open ball (i.e., intrinsic interior) or a closed planar halfspace (i.e., intrinsic boundary)—which coincide with pointset topological notions of interior and boundary, respectively.

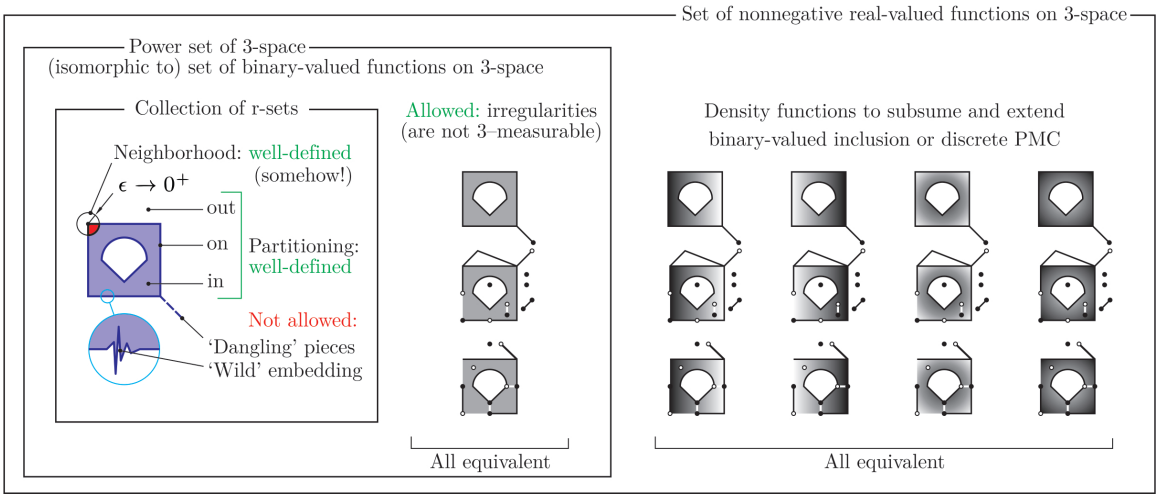

FIGURE 2.2.2: The collection of r-sets à la Requicha [275] (inner box) is to be subsumed and extended to include heterogenous dimensions (middle box) and heterogeneous distributions (outer box) in the remainder of this Chapter. See also Figs. 2.3.3 and 2.3.13.

with 'solid' ever since in most subsequent studies—as a compact (i.e., bounded and closed) regular and semianalytic subset of $\mathbb{R}^3$, which is automatically decomposable into a finite number of submanifolds (called 'strata' [209]). Hence r-sets appeared to be specific enough to capture all of the postulated requirements and, at the same time, broad enough to cover virtually all useful objects in engineering practice—at least the ones that were manufacturable by the available fabrication processes at the time (e.g., casting, molding, stamping, machining, etc.). As argued in Section 1.3, this is no longer the case; as neither the postulates of solid modeling nor the r-set abstraction are sufficient to capture the rich geometric, physical, and material properties of mechanical products that are manufacturable by the modern processes (e.g., 3D printing, composite weaving/knitting/braiding, etc.).

Figure 2.2.1 schematically summarizes the different approached enumerated above to characterize the notion of solidity. The implications on a simple planar shape are illustrated in simple terms in Fig. 2.2.2, where the idea of extending the mathematical

space to allow dimensional and structural heterogeneity that will be introduced in the rest of this chapter are also illustrated.

2.2.1.2 Spatial Modeling Relies More on Heterogenous Dimensions:

In contrast to the (more or less) common set of postulates for solid modeling (i.e., restricting 'shapes') that are still sufficient to tackle a large set of engineering applications, there are no such common set of strict postulates pronounced as specifically for spatial modeling (i.e., restricting 'motions'). Although the dimensional homogeneity and piecewise-manifoldness, discussed above at length, are universally accepted requirements for solids—except for a few application, giving rise to more generalized shape modeling proposals such as [137,299,300,359]—similar assumptions on motions can be too restrictive. Depending on the problem at hand, motions can be

- discretely instantiated (e.g., a finite sets of points in SE(3)) as applied in constructive solid representation [276, 282];
- one-parametric (e.g., piecewise continuous trajectories in SE(3)) as applied in one-parametric sweeps [5, 254];
- multi-parametric (e.g., piecewise continuous subspaces in SE(3)) as applied in workspace enumeration [79, 235];
- regular continuum (e.g., subsets of homogeneous dimensions in SE(3)) as assumed in formulating regularized morphology [215, 218]; or
- hybrid (e.g., subsets of heterogeneous dimensions in SE(3)) as in assembly planning [84] and other problems involving kinematic constraints of different DOFs between different features of objects.

In addition to (ordinary or regularized) Boolean operations that form the basis for constructive geometry and dominated the early focus in geometric modeling [276, 282, 343], there are other practically significant algebraic operations such as (ordinary or regularized) Minkowski operations that form the basis for morphological reasoning [139, 308]. These operations can be generalized to study the interactions of shapes and motions—or in morphological language, templates and filters—in a unified setup of 'group morphology' [291, 292]. I show in Section 3.4 that the range of problems exemplified above can be uniformly formulated as group morphological operations, which was first shown by Lysenko et al. [215, 218] in the context of geometric modeling applications. To enable an elegant analytic (i.e., measure-theoretic) interpretation, they assumed a regularized morphology formed with regularized Minkowski operations on regular (i.e., full-dimensional measurable) sets and established their implicit descriptions in terms of convolutions of indicator functions. However, regularization in the configuration space restricts the applicability to a very confined subset of problems in the above categories. In particular, problems involving heterogeneous dimensions such as computing one-parametric sweeps/unsweeps of 3D shapes along 1D trajectories in the configuration space cannot be formulated immediately as convolutions using a single (Lebesgue or Haar) measure.

To allow the so-obtained theory of implicit 'configuration modeling' [231] reach its full potential, multiple measures should be used to handle interactions between sets (i.e., shape and motion) of various dimensions. In the following sections, I show that Dirac $\delta-$calculus can be incorporated effectively into implicit descriptions to enrich the analytic (i.e., measure-theoretic) toolbox to uniformly formulate and solve more general shape and motion related problems.

2.2.2 Basic Definitions

Before presenting the specifics of analytic methods in Section 2.3, this section presents the basic definitions and terminology in detail, which are imperative to an effective presentation of the material in subsequent sections.

2.2.2.1 Basic Assumptions: Choosing Shape and Motion Spaces:

It is important for our purposes to clearly distinguish between shapes and motions and their distinctive properties that contributed to solid and configuration modeling as practiced today, before we can unify them back together.

The 'Shape' Space. The Euclidean 3−space $\mathsf{E} := \mathbb{R}^3$ endowed with its usual topology based on the standard L^2-metric is considered as the underlying space for shape description. The topological structure allows one to replace the set-theoretic Boolean algebra with a regularized Boolean algebra [343]. In addition to being a metric and topological space, E is also a vector space—thus an Abelian group, allowing one to define a group-theoretic Minkowski algebra—with continuous linear algebraic operations—thus a topological group, allowing one to define a regularized Minkowski algebra. The choice of the L^2-metric $\mathbb{d} : \mathbb{R}^3 \times \mathbb{R}^3 \to [0, +\infty)$, which is the same as the L^2-norm $\|\cdot\|_2 : \mathbb{R}^3 \to [0, +\infty)$ of the difference vector $\mathbf{x} := (\mathbf{x}_2 - \mathbf{x}_1)$, i.e., $\mathbb{d}(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_2 - \mathbf{x}_1\|_2$ is partly motivated by the fact that it is the only $L^{\mathrm{p}}-$norm that is independent of the choice of a Cartesian coordinate system, i.e., is invariant under rigid body motions (also called Euclidean isometries).

Definition 2.2.1. (Explicit Shapes) 'Explicit' descriptions of 'shapes' refer to the subsets of E. Thus $S \subseteq \mathsf{E}$ and $S \in \mathscr{P}(\mathsf{E})$ are equivalent expressions for the statement

"S is an (at most) 3D set of points." Shapes include isolated points, curve segments, surface patches, volume chunks, or heterogeneous combinations of them. For now, we assume no additional conditions (e.g., boundedness, regularity, etc.) on the sets.

Definition 2.2.2. (Implicit Shapes) 'Implicit' descriptions of 'shapes' refer to the functions defined over E. Thus $f : \mathsf{E} \to \mathbb{Q}$ and $f \in \mathbb{Q}^{\mathsf{E}}$ are equivalent expressions for the statement "$f(\mathbf{x}) \in \mathbb{Q}$ is a unique value assigned to $\mathbf{x} \in \mathsf{E}$." For now, $\mathbb{Q}$ can be any tensor space (most frequently $\mathbb{Q} := \mathbb{R}$ or $\mathbb{C} = \mathbb{R} + \mathbf{i}\mathbb{R} \cong \mathbb{R}^2$), and we assume no additional conditions (e.g., boundedness, continuity, etc.) on the function itself or its domain/codomain.

Often times, the explicit and implicit descriptions are associated syntactically by adding a subscript as $f_S \in \mathbb{Q}^{\mathsf{E}}$, whose relationship is determined by a logical condition on the function's value:

$$S = \big\{\mathbf{x} \in \mathsf{E} \mid \mathbb{E}[f_S(\mathbf{x})] = \mathtt{1}\big\}, \quad \text{where} \quad \mathbb{E}[\cdot] : \mathbb{Q} \to \{\mathtt{0}, \mathtt{1}\}. \tag{2.2.1}$$

The 'Motion' Space. In accordance with the rigidity assumption given in Section 2.2.1 (at least for the moment), the special Euclidean group $\mathsf{C} := \mathrm{SE}(3)$, representing the space of all rigid body motions, endowed with its usual topology based on the geodesic L^2-metric is considered as the underlying space for motion description, also known as the 'configuration space' in robotics literature [199]. Once again, both set-theoretic and regularized Boolean algebras can be defined for sets of configurations. In addition to being a metric and topological space, C is also a Lie group—i.e., a topological group and a Riemannian manifold, locally parameterized with its Lie algebra $\mathfrak{se}(3) \cong \mathbb{R}^6$, which, in turn, is a vector space—the two of which are related by the

exponential map $\exp|_{\mathfrak{se}(3)} : \mathfrak{se}(3) \to \mathrm{SE}(3)$. Thus both group-theoretic and regularized Minkowski algebras can be defined for sets of configurations. The geodesic L^2-metric $\mathbb{d} : \mathrm{SE}(3)\times\mathrm{SE}(3) \to [0, +\infty)$—that is the distance along the shortest curve connecting the two configurations—is obtained as the L^2-norm $\|\cdot\|_2 : \mathfrak{se}(3) \to [0, +\infty)$—called the Frobenius norm in this context—of the natural logarithm of the relative configuration $\mathbb{c} := \mathbb{c}_1^{-1}\mathbb{c}_2$, i.e., $\mathbb{d}(\mathbb{c}_1, \mathbb{c}_2) = \|\log(\mathbb{c}_1^{-1}\mathbb{c}_2)\|_2$—which subsumes the Euclidean metric over $\mathbb{R}^3$ as a special case.

Definition 2.2.3. (Explicit Motions) 'Explicit' descriptions of 'motions' refer to the subsets of C. Thus $M \subseteq \mathsf{C}$ and $M \in \mathscr{P}(\mathsf{C})$ are equivalent expressions for the statement "M is an (at most) 6D set of configurations." Motions include isolated configurations, various multi-parameteric trajectories, or heterogeneous combinations of them. For now, we assume no conditions (e.g., boundedness, regularity, etc.) on the sets.

Definition 2.2.4. (Implicit Motions) 'Implicit' descriptions of 'motions' refer to the functions defined over C. Thus $f : \mathsf{C} \to \mathbb{Q}$ and $f \in \mathbb{Q}^{\mathsf{C}}$ are equivalent expressions for the statement "$f(\mathbb{c}) \in \mathbb{Q}$ is a unique value assigned to $\mathbb{c} \in \mathsf{C}$." For now, $\mathbb{Q}$ can be any tensor space (most frequently $\mathbb{Q} := \mathbb{R}$ or $\mathbb{C} = \mathbb{R} + \mathbf{i}\mathbb{R} \cong \mathbb{R}^2$), and we assume no conditions (e.g., boundedness, continuity, etc.) on the function itself or its domain/codomain.

Once again, the explicit and implicit descriptions are associated syntactically by adding a subscript as $f_M \in \mathbb{Q}^{\mathsf{C}}$, whose relationship is determined by a logical condition on the function's value:

$$M = \big\{\mathbb{c} \in \mathsf{C} \mid \mathbb{E}[f_M(\mathbb{c})] = \mathtt{1}\big\}, \quad \text{where} \quad \mathbb{E}[\cdot] : \mathbb{Q} \to \{\mathtt{0}, \mathtt{1}\}. \tag{2.2.2}$$

Homogeneous Space. It can be shown that the shape space $\mathsf{E} = \mathbb{R}^3$ is in fact a 'homogeneous space' (in group-theoretic sense)[16] of the configuration space $\mathsf{C} = \mathrm{SE}(3)$ under the transitive group action of the elements in the latter (i.e., applying motions to shapes).[17] Indeed, upon choosing an arbitrary origin $\mathbf{0} \in \mathbb{R}^3$, one can express the motion group as the semidirect product $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$. The special orthogonal group $\mathrm{SO}(3) \cong \mathsf{stb}(\mathbf{0})$ (denoting 'stabilizer' of $\mathbf{0}$) is a Lie subgroup of SE(3) whose configuration elements stabilize/fix (i.e., do not move) the origin $\mathbf{0} \in \mathbb{R}^3$, representing all proper (i.e., reflection-free) rotations. The corresponding Lie algebra $\mathfrak{so}(3) \cong \mathbb{R}^3$ is a vector subspace of $\mathfrak{se}(3) \cong \mathbb{R}^6$ that can be viewed as an axis-angle parameterization of rotations in SO(3) via the (further restricted) exponential map $\exp|_{\mathfrak{so}(3)} : \mathfrak{so}(3) \to \mathrm{SO}(3)$. The translation group $\mathrm{T}(3) \cong \mathrm{SE}(3)/\mathrm{SO}(3)$, on the other hand, is the subgroup of SE(3) that is fixed-point free, and is in a one-to-one correspondence with the shape space, i.e., $\mathrm{T}(3) \cong \mathbb{R}^3$. Thus every rigid transformation can be represented by a tuple $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ of a rotation and translation—where $(\overset{0}{=})$ reads "is represented, with respect to the choice of origin $\mathbf{0} \in \mathbb{R}^3$, via the tuple"—applied in a noncommutative order—e.g., first a rotation $\mathbb{r} \in \mathrm{SO}(3)$, then a translation $\mathbf{t} \in \mathbb{R}^3$, as an immediate result of the semidirect product order. See Appendix A for more details on the motion group and its semidirect product structure.

Convention 2.2.5. (Measurable Sets) We are only interested in *measurable* sets, i.e., sets for which the notions of Lebesgue/Haar d−measures make sense. This includes

[16] In simple terms, this means that the entire Euclidean 3−space is symmetrical with respect to all rigid transformations (i.e., Euclidean isometries), and roughly corresponds to structural homogeneity and isotropy postulate given in Section 2.2.1.

[17] It is worthwhile mentioning that the terms 'configuration' or 'motion' are often used interchangeably in the literature—including some of my own past publications [35–37]—to refer to a single transformation $\mathbb{c} \in \mathsf{C}$ describing a particular 'pose' of the shape. However, for the lack of a better option, I use the term 'motion' to refer to a collection of configurations $M \in \mathsf{M}$.

infinite measures such as $(\mathrm{d}-1)-$measure of a $\mathrm{d}-$dimensional or $\mathrm{d}-$measure of an unbounded/fractal $\mathrm{d}-$dimensional shape, but excludes 'non-measurable' pathological sets.[18] Thus the reader should hereafter think of measurability as an implicit assumption in virtually all statements, equations, and diagrams.

Particularly, hereafter $\mathscr{P}(\mathsf{E})$ and $\mathscr{P}(\mathsf{C})$ stand for the collection of Lebesgue/Haar measurable subsets of the powerset—i.e., collection of all subsets—of E and C, respectively. Accordingly, hereafter $\mathbb{Q}^{\mathsf{E}}$ and $\mathbb{Q}^{\mathsf{C}}$ stand for the set of Lebesgue/Haar integrable $\mathbb{Q}-$valued functions over E and C, respectively.

The first distinction to note right away is that conventional solid modeling, despite being heavily dependent on the topological and metric structures of the Euclidean space, misses an opportunity to exploit its group structure.[19] On the other hand, conventional spatial planing is centered around the group structure of the configuration space, but is less explicit about the implications of its topological and metric properties. In the rest of this chapter, I will show that shapes and motions can leverage the properties on both ends, which are captured more naturally by analytic methods.

Figure 1.5.1 illustrates the distinction between shapes and motions with a simple planar motion planning problem (e.g., for mechanism design [140, 293]) in which a robot's motion in a space with obstacles is analyzed. If the motion is restricted

[18]For example, there are sets whose volume cannot satisfy basic properties (e.g., rotation or subdivision invariance)! Such sets are proven to exist in Zermelo-Fraenkel+choice (ZFC) set theory due to the axiom of choice [330]. We choose to ignore this technicality throughout this thesis, as such sets do not correspond to any shape of practical significance. A more rigorous mathematical treatment would require using concepts from descriptive set theory [175]

[19]Of course the Abelian group structure of the Euclidean vector space—implied by vector summations, subtractions, etc. from linear algebra—and the fact that it is acted on by different groups of linear transformations—applied as matrix multiplications—is extensively exploited in solid modeling, among many other important vector space properties of E. What I intend to point out here is the benefits of the realization that E is (isomorphic to) a Lie subgroup of the Lie groups that model its continuous transformations, providing an opportunity to unify shapes and motions and apply the established powertools of spatial planning and kinematic synthesis to solid modeling [231].

to the 2D translations in the plane, the motion space $\mathsf{C} := \mathrm{T}(3)$ is the same (up to isomorphism) with the shape space $\mathsf{E} := \mathbb{R}^3$, i.e., $\mathrm{T}(3) \cong \mathbb{R}^3$ upon fixing an origin. Thus for every fixed orientation of the shape, checking its collision against the obstacle can be viewed as an inclusion query on the configuration space obstacle—defined precisely in the subsequent sections. The goal in this case is to formulate that query in the motion space in terms of inclusion queries on the robot's end effector and obstacles in the shape space. I will show how this "queries-to-queries" mapping can be formulated analytically in terms of inner products and convolutions.

Basic Assumptions: Choosing Shape and Motion Subsets:

As depicted earlier in Section 2.2.1, it is often desirable to allow for heterogeneous dimensions in models for shapes and motions—which has been proposed in the past multiple times [137, 299, 300, 359]—without violating the algebraic closure postulate with respect to important operations.

Assumption 2.2.6. (Semianalytic Shapes) I choose to restrict all shapes to semianalytic subsets of the Euclidean space $\mathsf{E} = \mathbb{R}^3$, i.e., $\mathscr{P}(\mathsf{E})$ and $\mathbb{Q}^{\mathsf{E}}$ are both further restricted to include shapes that can be locally described via convergent power series expansions. This is motivated by the following properties of semianalytic sets:

- Semianalytic sets can be irregular, which allows modeling shapes and motions of heterogeneous dimensions. However, they are tamely embedded curvilinear subpolyhedra as depicted earlier, ensuring well-behavior, finite describability, and (most relevantly to the analytic approach) measurability.

- Semianalytic sets are decomposable into coarse disjoint manifold subsets called 'strata' while the corresponding decomposition is referred to as 'stratification'

[364]. Alternatively, semianalytic sets can be triangulated into a finite number of curved d−polytopes ($0 \leq d < 4$)—e.g., points, curve segments, curved polygons, and curved polyhedra in 3D [209]. This allows for formulating well-defined integrals (required for analytic modeling) over each stratum.

- Semianalytic sets are closed under topological operations of interest such as closure and boundary evaluation and regularization. In particular, boundaries of at most 3D semianalytic sets are at most 2D semianalytic sets [209].

- Semianalytic sets are closed under morphological operations of interest such as Boolean and Minkowski operations and skeletonization. In particular, medial axes (MA) of at most 3D semianalytic sets are at most 2D semianalytic sets [59].

For motions, similar restrictions can be imposed to generalize the power series expansions to locally describing motions (e.g., subanalytic subsets of $\mathsf{C} = \mathrm{SE}(3)$), to which I shall not attend in this thesis.[20]

The latter two properties enable interpreting geometric modeling and spatial reasoning operations uniformly as closed (ordinary or regularized) algebras over explicit or implicit shapes and motions. Therefore, rather than confining geometric operations to a closed subalgebra on regular semianalytic sets (e.g., r-sets joined with their complements), regularization is viewed as an operator that can be applied to the shapes and motions on demand, either explicitly or implicitly.

[20]This is motivated by the fact that the 6D Lie group $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$ is in a 1:2 diffeomorphism with the Cartesian product of the 3−sphere embedded in $\mathbb{R}^4$ (e.g., the space of all unit quaternions) and $\mathbb{R}^3$. Thus it can be embedded in the obvious fashion in a higher-dimensional Euclidean space. In addition, it can be locally parameterized using its Lie algebra $\mathfrak{se}(3) \cong \mathfrak{so}(3) \oplus \mathrm{T}(3)$.

2.2.2.2 Basic Operations: Inner Products and Convolutions:

Throughout this thesis, I shall make extensive use of nonegative real-valued functions of the form $f : \mathsf{E} \to \mathbb{R}_+$ (i.e., $f \in \mathbb{R}_+^{\mathsf{E}}$) and $f : \mathsf{C} \to \mathbb{R}_+$ (i.e., $f \in \mathbb{R}_+^{\mathsf{C}}$) defined over the shape and motion spaces, in which the notation $\mathbb{R}_+ := [0, +\infty)$ is used.[21] In a more general setting of complex-valued functions where the range is $\mathbb{C} = \mathbb{R} + \mathbf{i}\mathbb{R} \cong \mathbb{R}^2$ (where $\mathbf{i}^2 = -1$) is the complex plane—$\mathbb{R}_+$ being a restriction to half of the real-line—the space of all absolute-value integrable functions over shape and configuration spaces are dented via $L^1(\mathsf{E}) \subset \mathbb{C}^{\mathsf{E}}$ and $L^1(\mathsf{C}) \subset \mathbb{C}^{\mathsf{C}}$, respectively:

$$f \in L^1(\mathsf{E}) \quad \text{iff} \quad \int_{\mathsf{E}} |f(\mathbf{x})| \; d\mu^3[\mathbf{x}] < \infty \text{ (i.e., exists)}, \tag{2.2.3}$$

$$f \in L^1(\mathsf{C}) \quad \text{iff} \quad \int_{\mathsf{C}} |f(\mathbb{c})| \; d\mu^6[\mathbb{c}] < \infty \text{ (i.e., exists)}, \tag{2.2.4}$$

where $|\cdot| : \mathbb{C} \to \mathbb{R}_+$ is the absolute value. $\mu^3[\cdot]$ and $\mu^6[\cdot]$ represent Lebesgue and Haar measures over the shape and configuration spaces, respectively.

Definition 2.2.7. The 'inner product' $\langle \cdot_1, \cdot_2 \rangle : L^1(\mathsf{E}) \times L^1(\mathsf{E}) \to \mathbb{R}_+$ is defined as

$$\langle f_1, f_2 \rangle := \int_{\mathbb{R}^3} (f_1 \bullet f_2)(\mathbf{x}) \; d\mu^3[\mathbf{x}] = \int_{\mathbb{R}^3} f_1(\mathbf{x}) f_2(\mathbf{x}) \; d\mu^3[\mathbf{x}], \tag{2.2.5}$$

while the $L^{\mathrm{p}}-$norm of a function $\|\cdot\|_{\mathrm{p}} : L^1(\mathsf{E}) \to \mathbb{R}_+$ is defined in general as

$$\|f\|_{\mathrm{p}}^{\mathrm{p}} := \int_{\mathbb{R}^3} f^{\mathrm{p}}(\mathbf{x}) \; d\mu^3[\mathbf{x}] \quad \Rightarrow \quad \langle f_1, f_2 \rangle = \|f_1 \bullet f_2\|_1. \tag{2.2.6}$$

Thus the inner product can be viewed as an integration of the pointwise multiplication

[21] I use the notation $\mathbb{R}_+ := [0, +\infty)$ instead of (the more popular) $\mathbb{R}_+ := (0, +\infty)$ to avoid notational complexity that would result via $\mathbb{R}_+ \cup \{0\} = [0, +\infty)$.

of two functions over the entire 3−space.

The definition can be extended to $\langle \cdot_1, \cdot_2 \rangle : L^1(\mathsf{E}) \times L^1(\mathsf{E}) \to [0, \infty] := \mathbb{R}_+ \cup \{\text{“}\infty\text{”}\}$ in which the nominal infinity “∞” is included in the codomain to allow defining the above notions for the functions that are not strictly integrable—e.g., to account for the integration over input functions that are neither compactly supported nor decaying with a rapid rate for the integral to converge. However, if the functions are shape descriptors of bounded sets (e.g., r-sets)—in compliance with the boundedness postulate presented in Section 2.2.1—then they are compactly supported, hence are L^{p}−norm integrable (i.e., $f \in L^{\mathrm{p}}(\mathsf{E})$) for all $\mathrm{p} > 0$.[22]

Definition 2.2.8. The ‘convolution’ $* : L^1(\mathsf{E}) \times L^1(\mathsf{E}) \to L^1(\mathsf{E})$ is defined as

$$(f_1 * f_2)(\mathbf{x}) := \big\langle f_1, \big(\tilde{f}_2 \circ \mathsf{act}(\mathbb{r}_{\mathsf{id}}, -\mathbf{x})\big)\big\rangle = \int_{\mathsf{E}} f_1(\mathbf{x}')\ f_2(\mathbf{x} - \mathbf{x}')\ d\mu^3[\mathbf{x}'] \tag{2.2.7}$$

$$= \big\langle \big(\tilde{f}_1 \circ \mathsf{act}(\mathbb{r}_{\mathsf{id}}, -\mathbf{x})\big), f_2\big\rangle = \int_{\mathsf{E}} f_1(\mathbf{x} - \mathbf{x}')\ f_2(\mathbf{x}')\ d\mu^3[\mathbf{x}'], \tag{2.2.8}$$

where $\tilde{f}(\mathbf{x}) := -f(\mathbf{x}) := f(-\mathbf{x})$ is a reflection with respect to the origin $\mathbf{0} \in \mathsf{E}$.

Once again, the generalization of $L^1(\mathsf{E})$ to functions that are not strictly integrable, i.e., those with non-nconvergent integrals is desirable. Examples are bounded well-behaved functions over unbounded but measurable supports.

The convolution function is an ensemble of inner products of two functions, one of which is composed with the action of $\mathbb{c} :\overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{x})$ for different translations $\mathbf{x} \in \mathsf{E}$; i.e., can be viewed as shifting one function’s domain before integrating their pointwise

[22] As mentioned briefly earlier, allowing for unbounded sets is sometimes desirable to allow for closure under set complement operation, which, in turn, enables exploiting the dualities (in the form of De Morgan’s laws) between Boolean unions and intersections as well as between Minkowski sums and differences—or more generally, Minkowski products and quotients.

multiplication over the entire 3−space. The convolution over the vector space (i.e., Abelian group) $\mathsf{E} = \mathbb{R}^3$ is commutative as a direct result of the commutativity of vector sums (i.e., translations), which can be easily verified by a change of variables $\mathbf{x}' \mapsto (\mathbf{x} - \mathbf{x}')$ from (2.2.7) to (2.2.8) or vice versa.

The vector space (i.e., Abelian) convolution can be generalized to noncommutative group convolution $* : L^1(\mathsf{C}) \times L^1(\mathsf{C}) \to L^1(\mathsf{C})$ defined as

$$(f_1 * f_2)(\mathbb{c}) := \big\langle f_1, \big(\tilde{f}_2 \circ \mathsf{act}(\mathbb{c}^{-1})\big)\big\rangle = \int_{\mathsf{C}} f_1(\mathbb{c}')\ f_2\big((\mathbb{c}')^{-1}\mathbb{c}\big)\ d\mu^6[\mathbb{c}'] \tag{2.2.9}$$

$$\neq (f_2 * f_1)(\mathbb{c}) = \big\langle \big(\tilde{f}_1 \circ \mathsf{act}(\mathbb{c}^{-1})\big), f_2\big\rangle = \int_{\mathbb{c}} f_1\big((\mathbb{c}')^{-1}\mathbb{c}\big)\ f_2(\mathbb{c}')\ d\mu^6[\mathbb{c}'], \tag{2.2.10}$$

where $\tilde{f}(\mathbb{c}) := \boldsymbol{-}f(\mathbb{c}) := f(\mathbb{c}^{-1})$ is an inversion with respect to the identity element represented as $\mathbb{c}_{\mathsf{id}} \overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{0})$. See [215, 218] for more details.

Basic Operations: Fourier Transforms and Harmonic Analysis:

Definition 2.2.9. Using the orthonormal basis of the form $e^{\pm 2\pi \mathbf{i}(\boldsymbol{\omega}\cdot\mathbf{x})}$, the Fourier transform $\mathscr{F} : L^1(\mathsf{E}) \to L^1(\hat{\mathsf{E}})$ is defined as $\hat{f} := \mathscr{F}\{f\}$ where

$$\hat{f}(\boldsymbol{\omega}) := \big\langle f, e^{+2\pi \mathbf{i}(\boldsymbol{\omega}\cdot(\cdot))}\big\rangle = \int_{\mathsf{E}} f(\mathbf{x})\ e^{-2\pi \mathbf{i}(\boldsymbol{\omega}\cdot\mathbf{x})}\ d\mu^3[\mathbf{x}], \tag{2.2.11}$$

which can be shown to be invertible to retrieve the original function, and the inverse Fourier transform $\mathscr{F}^{-1} : L^1(\hat{\mathsf{E}}) \to L^1(\mathsf{E})$ is obtained as $f = \mathscr{F}^{-1}\{\hat{f}\}$ where

$$f(\mathbf{x}) := \big\langle \hat{f}, e^{-2\pi \mathbf{i}((\cdot)\cdot\mathbf{x})}\big\rangle = \int_{\hat{\mathsf{E}}} \hat{f}(\boldsymbol{\omega})\ e^{+2\pi \mathbf{i}(\boldsymbol{\omega}\cdot\mathbf{x})}\ d\mu^3[\boldsymbol{\omega}], \tag{2.2.12}$$

where E is viewed as the 'physical' domain,[23] using the notation $\mathbf{x} \in \mathsf{E}$ for 3D points, while $\hat{\mathsf{E}}$ is viewed as the 'frequency' domain, using the notation $\boldsymbol{\omega} \in \hat{\mathsf{E}}$ for frequencies (which are also 3D). Nevertheless, $\hat{\mathsf{E}} \cong \mathsf{E}$ and the two can be used interchangeably.

The primary motivation behind using inner products and convolutions in geometric modeling and spatial reasoning from a computational point of view, are the result of the following two well-known theorems [173]:

Theorem 2.2.10. (Parseval's Theorem) *Inner products are invariant under Fourier transform, i.e., for all* $f_1, f_2 \in L^1(\mathsf{E})$, $\langle \hat{f}_1, \hat{f}_2 \rangle = \langle f_1, f_2 \rangle$.

Proof. The proof is straightforward by algebraic manipulations of the Fourier transform of the inner product, noting the linearity of the Fourier transform. See a standard textbook on harmonic analysis [173]. □

In other word, (forward and inverse) Fourier transforms preserve the inner product structure of the space. Parseval's theorem can be illustrated effectively using the following commutative diagram:

$$\begin{array}{ccc} L^1(\mathsf{E}) \times L^1(\mathsf{E}) & & \\ \Big\uparrow \Big\downarrow (\mathscr{F}^{-1}, \mathscr{F}^{-1}) \quad (\mathscr{F}, \mathscr{F}) & \xrightarrow{\langle \cdot_1, \cdot_2 \rangle} & \mathbb{C} \\ L^1(\hat{\mathsf{E}}) \times L^1(\hat{\mathsf{E}}) & \xrightarrow{\langle \cdot_1, \cdot_2 \rangle} & \end{array} \tag{2.2.13}$$

Theorem 2.2.11. (Convolution Theorem) *The Fourier transform of convolution is*

[23] In the context of signal processing dealing with 1D functions representing time-variant signals, it is referred to as the 'time' domain.

the same as pointwise multiplication of Fourier transforms, i.e., for all $f_1, f_2 \in L^1(\mathsf{E})$*:*

$$\mathscr{F}\{f_1 * f_2\} = (\hat{f}_1 \bullet \hat{f}_1) \quad \text{i.e.,} \quad (f_1 * f_2) = \mathscr{F}^{-1}\{\hat{f}_1 \bullet \hat{f}_1\}. \tag{2.2.14}$$

Proof. The proof is straightforward by algebraic manipulations of the Fourier transform of the convolution, noting the linearity of the Fourier transform. See a standard textbook on harmonic analysis [173]. □

In other word, although convolution is a complex 'global' operation (i.e., integration) over the entire 3−space in the physical domain as defined in (2.2.7) and (2.2.8), it reduces to a simple 'local' operation, (i.e., pointwise multiplication) in the frequency domain, which can be computed much more efficiently.[24] The convolution theorem can be illustrated effectively using the following commutative diagram:

$$\begin{array}{ccc} L^1(\mathsf{E}) \times L^1(\mathsf{E}) & \xrightarrow{\quad * \quad} & L^1(\mathsf{E}) \\ \Big\uparrow\Big\downarrow (\mathscr{F}^{-1}, \mathscr{F}^{-1}) \quad (\mathscr{F}, \mathscr{F}) & & \mathscr{F}^{-1} \Big\uparrow\Big\downarrow \mathscr{F} \\ L^1(\hat{\mathsf{E}}) \times L^1(\hat{\mathsf{E}}) & \xrightarrow{\quad \bullet \quad} & L^1(\hat{\mathsf{E}}) \end{array} \tag{2.2.15}$$

Fourier transforms have other useful properties, a few of which are given below:

- Translation in the physical domain converts to a multiplier in the frequency domain, i.e., for a given $\mathbb{t} \in \mathsf{C}$ such that $\mathbb{t} \overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{t})$ with a 'shift' vector $\mathbf{t} \in \mathbb{R}^3$,

[24] Of course using Fourier transform one needs to compute two integrations for forward transform in (2.2.11) and one integration for inverse transform in (2.2.12). However, this is computed ones for a fixed pair of participating functions, rather than computing it for every relative shift as in the original formulation of convolution in (2.2.7) and (2.2.8). Moreover, forward and inverse Fourier transforms can be computed rapidly using the fast Fourier transform (FFT) algorithm [76].

such that $(\mathsf{act}(\mathbb{t}))(\mathbf{x}) = (\mathbf{x} + \mathbf{t})$ for all $\mathbf{x} \in \mathsf{E}$:

$$\mathscr{F}\{f \circ \mathsf{act}(\mathbb{t}^{-1})\} = (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{f}), \quad \text{i.e.,} \quad (f \circ \mathsf{act}(\mathbb{t}^{-1})) = \mathscr{F}^{-1}\{\hat{\varsigma}_{\mathbf{t}} \bullet \hat{f}\}, \tag{2.2.16}$$

where $\varsigma_{\mathbf{t}}(\mathbf{x}) = (\delta^3 \circ \mathsf{act}(\mathbb{t}^{-1}))(\mathbf{x}) = \delta^3(\mathbf{x} - \mathbf{t})$ is a shifted Dirac $\delta-$function (introduced in Section 2.3.2) whose Fourier transform is the sinusoidal basis function $\hat{\varsigma}_{\mathbf{t}}(\boldsymbol{\omega}) = \mathscr{F}\{\varsigma_{\mathbf{t}}\}(\boldsymbol{\omega}) = e^{-2\pi\mathbf{i}(\boldsymbol{\omega}\cdot\mathbf{t})}$ introduced in Definition 2.2.9. This property can be illustrated with the following commutative diagram:

$$\begin{array}{ccc} L^1(\mathsf{E}) & \underset{(\cdot)\circ\mathsf{act}(\mathbb{t})}{\overset{(\cdot)\circ\mathsf{act}(\mathbb{t}^{-1})}{\rightleftarrows}} & L^1(\mathsf{E}) \\ \mathscr{F}^{-1}\uparrow \;\downarrow\mathscr{F} & & \mathscr{F}^{-1}\uparrow \;\downarrow\mathscr{F} \\ L^1(\hat{\mathsf{E}}) & \underset{\hat{\varsigma}_{-\mathbf{t}}\bullet(\cdot)}{\overset{\hat{\varsigma}_{+\mathbf{t}}\bullet(\cdot)}{\rightleftarrows}} & L^1(\hat{\mathsf{E}}) \end{array} \tag{2.2.17}$$

- As a result of the linearity of (2.2.11), linear maps are preserved under the Fourier transform. For example, for a given $\mathbb{r} \in \mathrm{SO}(3)$, the rotations of Fourier transform is the same as Fourier transform of rotation:

$$\mathscr{F}\{f \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})\} = (\hat{f} \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})), \quad \text{i.e.,} \quad (f \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})) = \mathscr{F}^{-1}\{\hat{f} \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})\}, \tag{2.2.18}$$

which can also be illustrated via a commutative diagram:

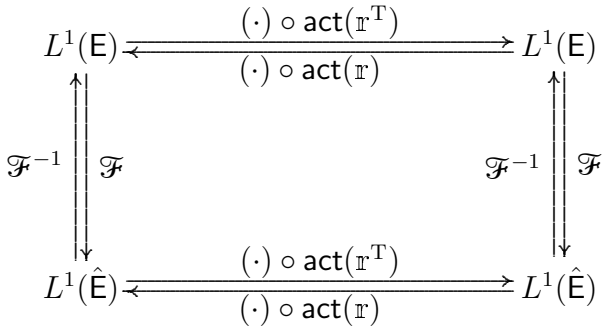

(2.2.19)

- Similarly, the reflection of Fourier transform is the same as the Fourier transform of reflection: $\mathscr{F}\{-f\} = -\mathscr{F}\{f\}$, (i.e., $\hat{\tilde{f}} = \tilde{\hat{f}}$), which can be illustrated as:

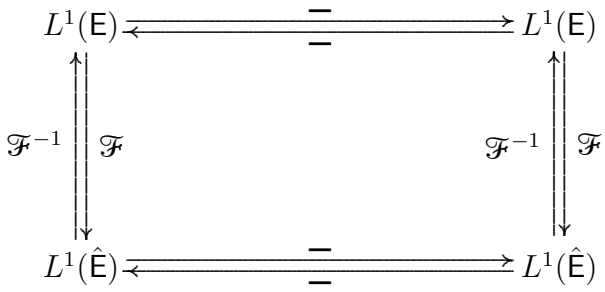

(2.2.20)

- For the special case when $f \in (L^1(\mathsf{E}) \cap \mathbb{R}^{\mathsf{E}})$ is real-valued, one has $f = \bar{f}$ hence the Fourier transform converts reflection to conjugation, i.e., $\mathscr{F}\{-f\} = \overline{\mathscr{F}\{f\}}$, (i.e., $\hat{\tilde{f}} = \bar{\hat{f}}$). Substituting the previous property then yields $-\mathscr{F}\{f\} = \overline{\mathscr{F}\{f\}}$, (i.e., $\tilde{\hat{f}} = \bar{\hat{f}}$), a property known as the 'Hermitian symmetry'.

Once again, the above results can be extended from the commutative Euclidean vector space $\mathsf{E} = \mathbb{R}^3$ to noncommutative group of rigid motions $\mathsf{C} = \mathrm{SE}(3)$. However, the convolution theorem becomes extremely messy due to the appearance of infinite dimensional matrices in the Fourier domain [65], which extremely limits their practical applicability. See [215, 218] for more details.

2.2.2.3 Topological Regularization: Explicit and Implicit Approaches:

Explicit Regularization. As mentioned in Section 2.2.1, classical solid modeling deals with algebraic operations on subsets of the Euclidean 3−space $\mathsf{E} = \mathbb{R}^3$ that are dimensionally homogeneous (i.e., regular)—informally, sets that contain only interior points and (optionally) all those limit points that have some interior points nearly; all other points with eroded lower-dimensional neighborhoods such as 'dangling' faces and edges or (positive or negative) singular points are not allowed.

Convention 2.2.12. (Basic Explicit Operators) To facilitate the upcoming discussions, the following notations for standard operators are used throughout this thesis:

- Set-theoretic complement, denoted $\mathscr{c} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the set of all points that are excluded from it, i.e., $\mathbf{x} \in \mathscr{c}(S)$ iff $\mathbf{x} \notin S$. Note that it is purely set-theoretic [153], and unlike the following operators, does not require appealing to a topological or metric structure.

- Topological 'interior', denoted $\mathscr{i} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the largest open set contained in it [229].

- Topological 'closure', denoted $\mathscr{k} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the smallest closed set containing it [229].[25]

The interior and closure are dual to each other with respect to set complement, i.e.,

- Topological 'exterior', denoted $\mathscr{e} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the interior of its complement, which is the same as the complement of its closure, i.e., $\mathscr{e}(S) = (\mathscr{i} \circ \mathscr{c})(S) = (\mathscr{c} \circ \mathscr{k})(S)$.

[25] Closure can be equivalently defined as the union of the set and its 'limit points', i.e., points (not necessarily included in the set) to which convergent sequences of included points converge.

- Similarly, the (nameless) operator $\hbar : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the closure of its complement, which is the same as the complement of its interior, i.e., $\hbar(S) = (\hbar \circ \mathit{c})(S) = (\mathit{c} \circ \mathit{i}\,)(S)$.

The above dualities can be illustrated via the following commutative diagrams:

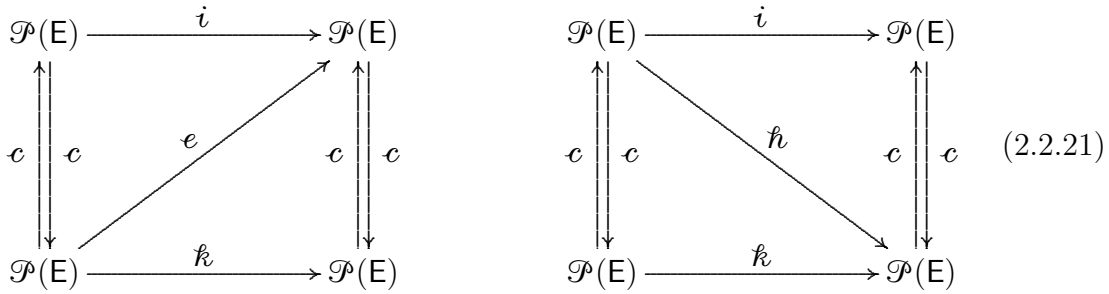

(2.2.21)

Based on the above operators, two dual definitions are possible for regularization:

- 'Closed regularization', denoted $\mathit{r}_{\mathrm{c}} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the closure of its interior, i.e., $\mathit{r}_{\mathrm{c}}(S) = (\hbar \circ \mathit{i}\,)(S)$.

- 'Open regularization', denoted $\mathit{r}_{\mathrm{o}} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the interior of its closure, i.e., $\mathit{r}_{\mathrm{o}}(S) = (\mathit{i} \circ \hbar)(S)$.

The duality between the above two can be added to the diagram in (2.2.21) as

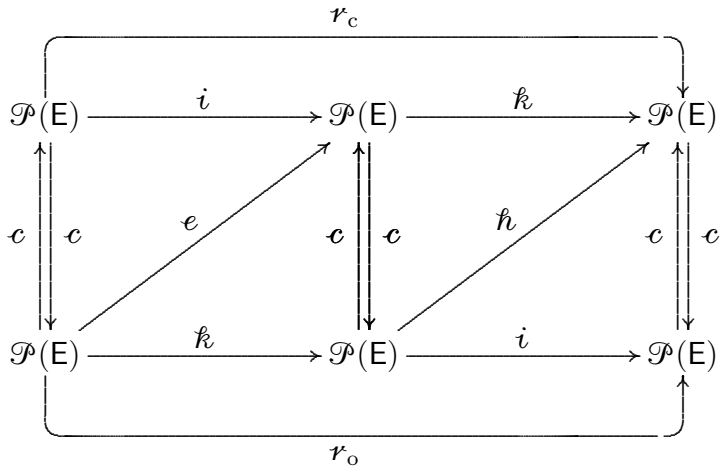

(2.2.22)

Note also that $\mathscr{r}_{\mathrm{c}}(S) = (\mathscr{h} \circ \mathscr{h})(S)$ and $\mathscr{r}_{\mathrm{o}}(S) = (\mathscr{e} \circ \mathscr{e})(S)$, not shown above to prevent cluttering.

- Topological 'boundary', denoted $\mathscr{b} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, converts a shape into the set-theoretic intersection of its closure and its complement's closure, i.e., $\mathscr{b}(S) = (\mathscr{k}(S) \cap \mathscr{h}(S))$. It includes every point, inside the set or otherwise, that has both an internal and external point nearby.

Figure 2.2.3 illustrates the above concepts for a simple planar shape.

Either closed or open regularization can be used, denoted $\mathscr{r} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ to intentionally keep the choice unspecified. In either case, it can viewed as a many-to-one map that eliminates the unrealizable features and returns a regular shape, with or without a tight 'skin' (i.e., boundary) around it being included in the set. The operator's image space $\mathscr{P}^*(\mathsf{E}) := \mathscr{r}(\mathscr{P}(\mathsf{E}))$ is called the 'regular powerset' of the $3-$space—corresponding to the larger circle on the left in Fig. 2.2.1 when closed regularization is used as in conventional solid modeling [276].

The fact that one of the two dual Boolean operations requires explicit regularization in each case results in at least two major difficulties; namely,

- To establish a CSG scheme [282] that is 'well-formed' [315] (i.e., guaranteed for correctness), one needs to explicitly introduce regularization to at least one of the two Boolean operations [343]—e.g., by somehow capturing $\epsilon-$neighborhood differential geometry as $\epsilon \to 0^+$, which is not finitely describable for general semianalytic sets, as illustrated in Fig. 1.3.4.

- In each case, the two operations can become non-computable in special configurations due to instability of the output with respect to infinitesimal changes in

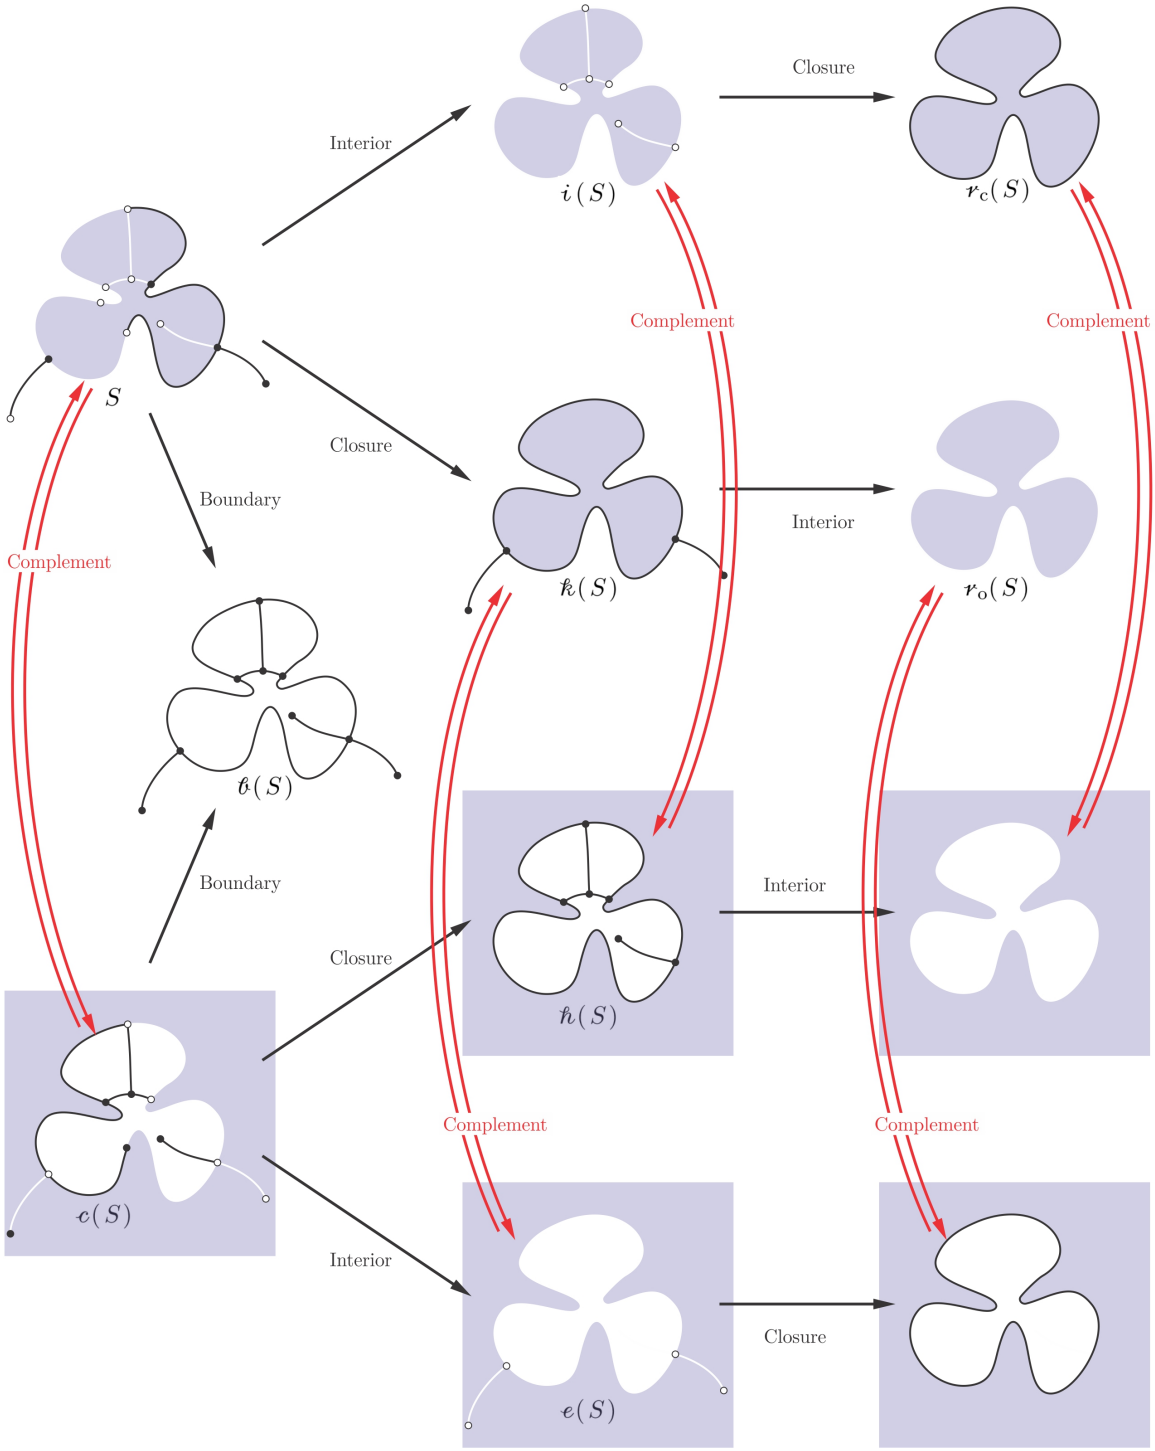

FIGURE 2.2.3: The pointset-topological operators introduced in Convention 2.2.12, their notations, and duality relationships.

the input [105, 106]. When the shapes are in degenerate lower-dimensional contact, a small perturbation in their shape (e.g., in Hausdorff topology) or configurations can cause large changes in the result with the possibility of points, curve segments, or surface patches of finite size suddenly appearing/disappearing.

Both problems can be avoided by replacing the explicit pointset-topological regularization with an implicit measure-theoretic regularization described below.

Implicit Regularization. A few concepts from measure theory will extremely advantageous. In the following, the Hausdorff d−measure—defined as the usual external measure on d−dimensional subsets of the 3−space based on the Euclidean L^2−metric—is denoted by $\mu^{\mathrm{d}} : \mathscr{P}(\mathsf{E}) \to \mathbb{R}_+ \cup \{\text{“}\infty\text{”}\}$, which corresponds to counting a finite number of points, and measuring curve length, surface area, and interior volume (i.e., Lebesgue 3−measure) for $0 \leq \mathrm{d} < 4$, respectively. The element “∞” is included to make the unbounded sets ‘measurable’.

Convention 2.2.13. A property is said to hold ‘almost everywhere’ (abbreviated ae) over the Euclidean 3−space $\mathsf{E} = \mathbb{R}^3$ if it holds everywhere except over a set with zero Lebesgue 3−measure (i.e., of null 3D volume).[26]

Definition 2.2.14. (Equal−ae Shapes) Two shapes $S_1, S_2 \in \mathscr{P}(\mathsf{E})$ are equal−ae (denoted $S_1 \overset{\mathsf{ae}}{=} S_2$) if their difference is a lower-dimensional set (i.e., one of zero 3−measure) $\mu^3[S_1 - S_2] = 0$. This means that the 3−measure of both sets’ intersections with every other measurable domain agree, i.e.,

$$S_1 \overset{\mathsf{ae}}{=} S_2 \quad \text{iff} \quad \mu^3[S_1 \cap \Omega] = \mu^3[S_2 \cap \Omega] \quad \text{for all } \Omega \in \mathscr{P}(\mathsf{E}), \tag{2.2.23}$$

[26] In probability theory, the equivalent notion is an event that is said to happen ‘almost surely’ if it happens with a probability of unity.

It is easy to verify that "$\overset{\mathsf{ae}}{=}$" $\subset \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E})$ is an equivalence relation over the Lebesgue 3−measurable subsets of the 3−space.[27]

As mentioned in Convention 2.2.5, it is important to emphasize the restriction of the definition of equality−ae to the measurable subsets of E, including sets of zero or infinite measures. The following is an important proposition that enables a measure-theoretic conceptualization of topological regularization, which is central to the development of implicit geometric operations:

Proposition 2.2.15. *Given Lebesgue-measurable sets $S_1, S_2 \in \mathscr{P}(\mathsf{E})$, their topological regularization is identical iff they are equal*−ae*—i.e., $\mathscr{r}(S_1) = \mathscr{r}(S_2)$ iff $S_1 \overset{\mathsf{ae}}{=} S_2$.*

Proof. This can be verified by noting that if the two sets share the same (closed or open) regularization $\mathscr{r}(S_1) = \mathscr{r}(S_2)$, then the set of irregularities obtained as their difference $(S_1 - S_2)$ is nonwhere dense, thus has a zero 3−measure.[28] □

Definition 2.2.16. The equivalence class $[S]^\dagger \in \mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}$ includes all Lebesgue 3−measurable sets that are equal−ae, i.e., share the same regularization:

$$[S]^\dagger := \left\{ S' \in \mathscr{P}(\mathsf{E}) \mid S \overset{\mathsf{ae}}{=} S' \text{ i.e., } \mathscr{r}(S) = \mathscr{r}(S') \right\}. \tag{2.2.24}$$

Disregarding the equivalence class of non-measurable sets, the quotient space $\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}$ is in a one-to-one correspondence with the regular power set $\mathscr{P}^*(\mathsf{E})$,[29] which can be

[27] From a probabilistic view, two sets are equal−ae iff the probability of any point included in one set being excluded from the other is zero.

[28] This is a simplified explanation rather than a rigorous proof, which would be beyond the scope of this thesis, as with many other technicalities.

[29] To avoid making significant changes to the notations, simply think of all non-measurable sets being dumped into a single outlier equivalence class. The rest of the quotient space is in a one-to-one correspondence with the class of regular sets $\mathscr{P}^*(\mathsf{E})$.

illustrated via the following commutative diagram (to be extensively cited later):

$$\begin{array}{ccccc} & & [\cdot]^\dagger & & \\ & \ulcorner & \rightarrow & \urcorner & \\ \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad r(\cdot) \quad} & \mathscr{P}^*(\mathsf{E}) & = & \mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=} \end{array} \tag{2.2.25}$$

where the horizonal equality arrow (=) implies isomorphism rather than strict equality. Such an implicit conceptualization of regularization can thus replace the explicit approach to avoid mathematical complications and computational difficulties associated with the latter altogether.

As an example, all nine shapes in Fig. 2.2.3 (except the boundary) are equal$-\mathsf{ae}$, as operations that depend on 3−measurable local and global properties would not be affected by their lower-dimensional differences.

The idea is to develop a language in which equality is replaced with equality$-\mathsf{ae}$ when developing algebraic structures. Going back to the aforementioned difficulties with explicit regularization, the first problem can be solved by noting that even though (one of the two) regularized Boolean operations on regular sets are not equal to the set-theoretic Boolean operations, they are equal$-\mathsf{ae}$ for general sets. The same can be said for set complement and difference operations, i.e., for all $S, S_1, S_2 \in \mathscr{P}(\mathsf{E})$, the following equivalence relations hold:

$$r((rS_1) \cup (rS_2)) \overset{\mathsf{ae}}{=} ((rS_1) \cup (rS_2)) \overset{\mathsf{ae}}{=} r(S_1 \cup S_2) \overset{\mathsf{ae}}{=} (S_1 \cup S_2), \tag{2.2.26}$$

$$r((rS_1) \cap (rS_2)) \overset{\mathsf{ae}}{=} ((rS_1) \cap (rS_2)) \overset{\mathsf{ae}}{=} r(S_1 \cap S_2) \overset{\mathsf{ae}}{=} (S_1 \cap S_2), \tag{2.2.27}$$

This means that irregularities become irrelevant up to equality$-\mathsf{ae}$ as far as finite sequences of Boolean operations are concerned. This is simply due to the fortunate

fact that eroded lower-dimensional features (i.e., irregularities) of nonzero $0-$, $1-$, or $2-$measures cannot contribute to 3D features with nonzero $3-$measure after a finite number of Boolean combinations.[30] On the other hand, the second problem is also resolved when assessing robustness and stability up to equality$-\mathsf{ae}$, since infinitesimal changes to the sets can only produce changes in the constructive geometry with large $0-$, $1-$, or $2-$measures, while the variations of the $3-$measure are still infinitesimal.

To summarize, by working in a space where strict equalities and exact geometric informational completeness are replaced with equalities$-\mathsf{ae}$ and completeness up to 'measurable' properties, a significant subset of problems in geometric reasoning will disappear. I find this approach more practical than the traditional one with excessive obsession with exact geometry, because exact geometric completeness is not achievable in practice on digital computers with finite-precision floating point arithmetic capabilities, leading to numerous data exchange and interoperability problems between solid modeling systems in the presence of inaccuracies [260–262].

The regularized Boolean operations and set complement are defined over the collection of (closed or open) regular sets as the (closed or open) regularization applied to the set-theoretic counterparts [343], i.e., $\cup^*, \cap^*, -^* : \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) \to \mathscr{P}^*(\mathsf{E})$ and $\boldsymbol{c}^* : \mathscr{P}^*(\mathsf{E}) \to \mathscr{P}^*(\mathsf{E})$, defined as $\odot^* := (\boldsymbol{r} \circ \odot)$ in which $\odot \in \{\cup, \cap, -\}$ and $\boldsymbol{c}^* := (\boldsymbol{r} \circ \boldsymbol{c})$. It is easy to verify that they form a Boolean algebra $\langle \mathscr{P}^*(\mathsf{E}); \cup^*, \cap^*, \boldsymbol{c}^* \rangle$. To avoid explicit post-regularization and the problems that come with it, I propose substituting these definitions with the following isomorphic forms defined for the equivalence classes of shapes (denoted by $[\cdot]^\dagger$) via the equivalence relation $\overset{\mathsf{ae}}{=}$, i.e., as

[30] The term "finite" is the key here, as it is not true for infinite Boolean operations, e.g., arising when defining Minkowski and sweep/unsweep operations.

$\cup^\dagger, \cap^\dagger, -^\dagger : (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}) \times (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}) \to (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=})$ and $\boldsymbol{c}^\dagger : (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}) \to (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=})$:

$$[S_1]^\dagger \cup^\dagger [S_2]^\dagger := \left\{ S_1' \cup S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\} = [S_1 \cup S_2]^\dagger, \tag{2.2.28}$$

$$[S_1]^\dagger \cap^\dagger [S_2]^\dagger := \left\{ S_1' \cap S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\} = [S_1 \cap S_2]^\dagger, \tag{2.2.29}$$

The fortunate observation that the equivalence relation commutes with Boolean operations—i.e., identities of the form $[S_1] \odot^\dagger [S_2] = [S_1 \odot S_2]$ for $\odot \in \{\cap, \cup, -\}$—ensures that the quotient space $\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}$ is closed under these generalized operations, in a similar way that the space of regular sets $\mathscr{P}^*(\mathsf{E})$ is closed under Tilove and Requicha's regularized operations [343]. Therefore, $\langle \mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}; \cup^\dagger, \cap^\dagger, \boldsymbol{c}^\dagger \rangle$ forms a Boolean algebra that is isomorphic to the traditional regularized Boolean algebra.

The above concepts can be readily extended from shapes to motions, with the main distinction being due to the use of the Haar measure over the Lie group $\mathsf{C} = \mathrm{SE}(3)$ to define the equivalence relation $\overset{\mathsf{ae}}{=} \subset \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C})$ between measurable motions, as a natural extension of the Lebesgue measure over the Euclidean space $\mathsf{E} = \mathbb{R}^3$.

Unfortunately, it is not as trivial to use a similar trick as the one used for regularized Boolean algebra to redefine regularized Minkowski algebra in terms of equivalence classes via $\overset{\mathsf{ae}}{=}$. This is due to the fact that unlike finite sequences of Boolean operations, infinite Boolean combinations—e.g., as in Minkowski operations over continuum sets $|S_1| = |S_2| = |\mathbb{R}|$—can produce 3D artifacts from cross-combinations of points, curves, and surfaces. In other words, eroded lower-dimensional features (i.e., irregularities) of nonzero $0-$, $1-$, or $2-$measures are unfortunately able to 'sweep' 3D features with nonzero $3-$measure through Minkowski combinations. For example, one might be tempted to redefine the regularized Minkowski sum and difference

$\oplus^\dagger, \ominus^\dagger : (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}) \times (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}) \to (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=})$ as

$$[S_1]^\dagger \oplus^\dagger [S_2]^\dagger \overset{!}{:=} \left\{ S_1' \oplus S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\} \neq [S_1 \oplus S_2]^\dagger, \tag{2.2.30}$$

$$[S_1]^\dagger \ominus^\dagger [S_2]^\dagger \overset{!}{:=} \left\{ S_1' \ominus S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\} \neq [S_1 \ominus S_2]^\dagger, \tag{2.2.31}$$

which are incorrect, simply because pairwise Minkowski combinations of shapes picked from the input equivalence classes do not output equivalence classes. In fact, one can check that such definitions yield $\mathsf{E} \in [S_1]^\dagger \oplus^\dagger [S_2]^\dagger$ and $\emptyset \in [S_1]^\dagger \ominus^\dagger [S_2]^\dagger$—e.g., think of S_1' and S_2' as sets with space-filling (positive or negative) curves among their irregularities—implying that $\mathsf{E} \overset{\mathsf{ae}}{=} (S_1 \oplus S_2)$ and $\emptyset \overset{\mathsf{ae}}{=} (S_1 \ominus S_2)$, which is a contradiction. The alternative redefinitions that immediately comes to mind is

$$[S_1]^\dagger \oplus^\dagger [S_2]^\dagger \overset{!}{:=} [S_1 \oplus S_2]^\dagger \neq \left\{ S_1' \oplus S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\}, \tag{2.2.32}$$

$$[S_1]^\dagger \ominus^\dagger [S_2]^\dagger \overset{!}{:=} [S_1 \ominus S_2]^\dagger \neq \left\{ S_1' \ominus S_2' \mid S_1' \in [S_1]^\dagger \text{ and } S_2' \in [S_2]^\dagger \right\}, \tag{2.2.33}$$

which are also incorrect, because the right-hand sides of $\overset{!}{:=}$ in the above definitions are dependent on the particular choices of S_1 and S_2 while the equivalence classes on the left-hand sides are meant not to; meaning that the above operations do not guarantee algebraic closure.

It turns out that unlike the case with the Boolean algebra where we avoided explicit regularization using an 'irregularity-tolerant' language of equality$-\mathsf{ae}$, it is not as trivial to form a closed Minkowski algebra without an explicit appeal to regularization of the operand sets. This is also the case with sweep/unsweep construction, configuration obstacle generation, and similar motion-related operations. On the other hand, the existence of an algebraic structure over $\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}$ which is isomorphic

to $\langle \mathscr{P}^*(\mathsf{E}); \oplus^*, \ominus^* \rangle$ is guaranteed due to the one-to-one correspondence (i.e., bijection) between $\mathscr{P}^*(\mathsf{E})$ and $\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=}$ established by $[\cdot]^*|_{\mathscr{P}^*(\mathsf{E})} : \mathscr{P}^*(\mathsf{E}) \to (\mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=})$. The question is, what is the right measure-theoretic abstraction to extend the equivalence via $\overset{\mathsf{ae}}{=}$ to handle irregularities implicitly for such operations? I show in Section 2.3.2 that Dirac $\delta-$functions provide the right toolset to overcome these difficulties.

In what follows, I present implicit (i.e., functional) descriptions of semianalytic sets in such a way that morphological (i.e., Boolean and Minkowski) operations can be interpreted as meaningful measure-theoretic operations on those functions.

2.3 Analytic Reformulation

Having specified the underlying spaces for shapes and motions, I dedicate this section to develop the foundations for analytic reformulation in great detail. Section 2.3.1 presents the alternative modeling paradigm by conceptualizing geometric entities as functions instead of sets, along with measure-theoretic characterization of solidity. The reconciliation with the classical set-theoretic and pointset-topological approach is enabled by equivalence classes of functions to which the algebraic operations are extended. To further expand the types of geometric entities that are not describable with the classical approach, a Dirac $\delta-$calculus for geometric modeling is presented in Section 2.3.2 to enable graceful transitions between subsets of heterogeneous dimensions and operating on them using a single (Lebesgue or Haar) measure.

2.3.1 Defining Functions

The ideas of implicit description, analytic modeling, and functional geometry are perhaps as dated as geometric modeling itself. For example, the idea of using functions to represent geometry for intersection tests between 3D shapes was introduced by Comba [73] about a half century ago. It was later adopted by Ricci [286] to establish a min/max Boolean algebra for constructive geometry with smooth blending capabilities. However, Rvachev's theory of R-functions [302] presents perhaps the most well-studied and rigorously developed mathematical foundations for implicit algebras of sets, which was further developed and applied to semianalytic solid geometry by Shapiro [312, 313, 317]. However, R-functions did not receive the deserved attention as a powerful modeling tool in solid modeling applications and systems, mostly due to the algebraic and analytic complexities that hindered their applicability to large models with numerous shape features.

The analytic methods and algebra of functions that I present in this section are much simpler to work with and provide tremendous computational benefits over other explicit and implicit methods. The idea is to exploit powerful tools from basic measure theory to reinterpret fundamental operations in terms of maps between functions, from which complete implicit morphologies can be developed—i.e., not only equipped with Boolean operations for constructive shape modeling as with R-functions, but also capable of morphing shapes with motions via Minkowski operations. In addition, they subsume several other techniques that use real-valued defining functions to solve geometric problems, including Gaussian functions for protein electron density modeling [25, 28, 127] and bump functions for blobby object modeling [43, 44, 214].

2.3.1.1 Characteristic Functions and Level Set Defining Functions:

Let us start from analytic models for describing shapes and algebraic operations on them to present a few concepts, which are easily extensible to motions.

Definition 2.3.1. (Characteristic Functions) From basic set theory [153] there is a one-to-one correspondence (i.e., bijection) between the following two sets:

$$\mathscr{P}(\mathsf{E}) = \big\{S \mid S \subseteq \mathsf{E}\big\} \quad \cong \quad 2^{\mathsf{E}} = \{0,1\}^{\mathsf{E}} = \big\{\mathbf{1} \mid \mathbf{1} : \mathsf{E} \to \{0,1\}\big\}, \tag{2.3.1}$$

where the bijection is provided by the indicator map $\mathsf{indc} : \mathscr{P}(\mathsf{E}) \to 2^{\mathsf{E}}$ that maps a shape $S \subseteq \mathsf{E}$ to a binary-valued function $\mathbf{1}_S : \mathsf{E} \to \{0,1\}$ defined as

$$\mathbf{1}_S := \mathsf{indc}(S) \in 2^{\mathsf{E}}, \quad \text{where} \quad \mathbf{1}_S(\mathbf{x}) = (\mathbf{x} \overset{?}{\in} S) = \begin{cases} 1 & \text{if } \mathbf{x} \in S, \\ 0 & \text{if } \mathbf{x} \notin S. \end{cases} \tag{2.3.2}$$

where $\mathbf{1}_{(\cdot)} = \mathsf{indc}(\cdot)$ are interchangeable notations for the so-called 'indicator function' or 'characteristic function' of a shape. The inverse is given by the support map $\mathsf{supp} : 2^{\mathsf{E}} \to \mathscr{P}(\mathsf{E})$ that converts a function $\mathbf{1} : \mathsf{E} \to \{0,1\}$ to its support:

$$\mathsf{supp}(\mathbf{1}_S) = S, \quad \text{where} \quad \mathsf{supp}(\mathbf{1}) = \big\{\mathbf{x} \in \mathsf{E} \mid \mathbf{1}(\mathbf{x}) \neq 0\big\}. \tag{2.3.3}$$

The first relationship to notice is that indicator function of set complement is the same as negation of indicator function: $\complement S = \mathsf{supp}(\neg \mathbf{1}_S) = \mathsf{supp}(1 - \mathbf{1}_S)$, i.e.,

$$(\mathsf{indc} \circ \complement)(S) = (\neg \circ \mathsf{indc})(S) \;\rightleftharpoons\; (\mathsf{supp} \circ \neg)(\mathbf{1}_S) = (\complement \circ \mathsf{supp})(\mathbf{1}_S), \tag{2.3.4}$$

which can be illustrated using a fully commutative diagram of the form:

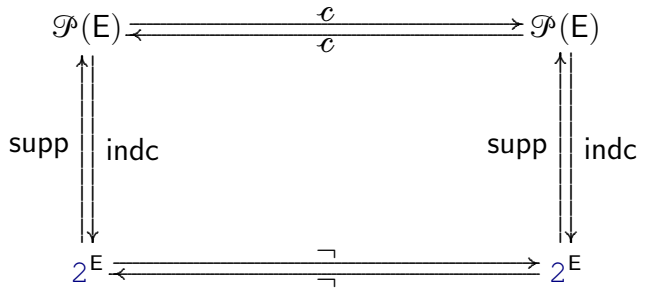

(2.3.5)

Another observation is that indicator functions are covariant with rigid (and other invertible) maps, meaning that indicator function of moved set is the same as indicator function of the original set composed from right with the action of inverted rigid map, i.e., $\mathbf{1}_{\mathbb{c}S} = \big(\mathbf{1}_S \circ \mathsf{act}(\mathbb{c}^{-1})\big)$ for all $\mathbb{c} \in \mathsf{C}$, i.e.,

$$\big(\mathsf{indc} \circ \mathsf{act}(\mathbb{c})\big)(\cdot) = \big(\mathbf{1}_{(\cdot)} \circ \mathsf{act}(\mathbb{c}^{-1})\big) \;\rightleftharpoons\; \mathsf{supp}\big((\cdot) \circ \mathsf{act}(\mathbb{c}^{-1})\big) = \big(\mathsf{act}(\mathbb{c}) \circ \mathsf{supp}\big)(\cdot), \tag{2.3.6}$$

which can be illustrated using a fully commutative diagram of the form:

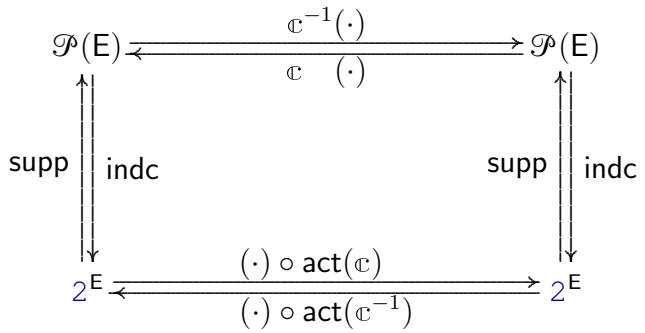

(2.3.7)

The relationships expressed via diagrams (2.3.5) and (2.3.7) will be significant in establishing fundamental dualities; namely

1. dualities of Boolean union/intersection via De Morgan's laws which will extend into dualities in Minkowski sum/different, product/quotient, sweep/unsweep,

and so on, that are carried to the implicit domain via (2.3.5); and

2. dualities of explicit and implicit definitions of the aforementioned operations by inverting the motions via (2.3.7).

Combining the two allows formulating most shape and motion related problems in terms of intersections alone—e.g., both sweep/unsweep to be defined as conditions on intersection of a shape with an inverted motion trajectory, and similarly for dual Minkowski operations.

The notion of a 'measure' is central to the analytic methods in this thesis. In particular, the Lebesgue 3−measure of a set (e.g., volume for 3D shapes) is computable as the Lebesgue 3−integral of its indicator function over the entire 3−space (i.e., L^1-norm):

$$\|\mathbf{1}_S\|_1 = \int_{\mathsf{E}} \mathbf{1}_S(\mathbf{x})\ d\mu^3[\mathbf{x}] = \int_S\ d\mu^3[\mathbf{x}] = \mu^3[S], \tag{2.3.8}$$

which can be illustrated using another simple commutative diagram:

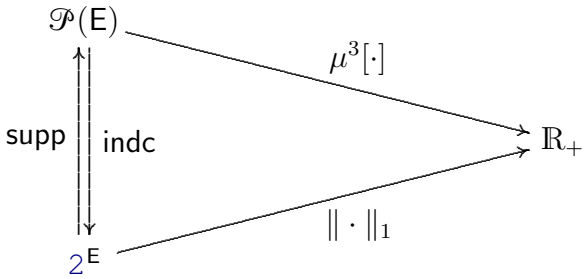

(2.3.9)

Next, set-theoretic Boolean operations $\cup, \cap : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ on explicit shapes (i.e., sets) are expressed via logical "or"/"and" gates $\vee, \wedge : 2^{\mathsf{E}} \times 2^{\mathsf{E}} \to 2^{\mathsf{E}}$ on

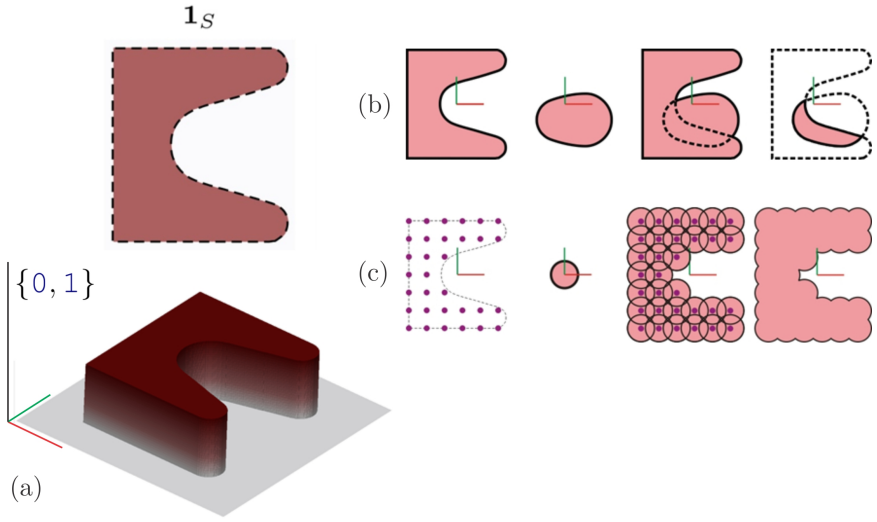

FIGURE 2.3.1: Describing a 2D shape with its indicator function (a) allows an implicit formulation of Boolean union/intersection with logical disjunction/conjunction, respectively (b). The idea can be extended to the Minkowski sum of discrete samples with a given continuum shape, viewed as a finite union of displaced instances of the later—thus finite disjunction of displaced indicator functions (c) but does not go further beyond that—e.g., for the general Minkowski sums. Compare with Fig. 2.3.2.

their indicator functions (i.e., implicit shapes) ia the following relationships:

$$(S_1 \cup S_2) = \mathsf{supp}(\mathbf{1}_{S_1} \vee \mathbf{1}_{S_2}), \quad \text{i.e.,} \quad \mathsf{indc}(S_1 \cup S_2) = (\mathbf{1}_{S_1} \vee \mathbf{1}_{S_2}), \tag{2.3.10}$$

$$(S_1 \cap S_2) = \mathsf{supp}(\mathbf{1}_{S_1} \wedge \mathbf{1}_{S_2}), \quad \text{i.e.,} \quad \mathsf{indc}(S_1 \cap S_2) = (\mathbf{1}_{S_1} \wedge \mathbf{1}_{S_2}), \tag{2.3.11}$$

which can be illustrated more effectively with the commutative diagrams:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\cup} & \mathscr{P}(\mathsf{E}) \\ \mathsf{supp} \uparrow \quad \downarrow \mathsf{indc} & & \mathsf{supp} \uparrow \quad \downarrow \mathsf{indc} \\ 2^{\mathsf{E}} \times 2^{\mathsf{E}} & \xrightarrow{\vee} & 2^{\mathsf{E}} \end{array} \qquad \begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\cap} & \mathscr{P}(\mathsf{E}) \\ \mathsf{supp} \uparrow \quad \downarrow \mathsf{indc} & & \mathsf{supp} \uparrow \quad \downarrow \mathsf{indc} \\ 2^{\mathsf{E}} \times 2^{\mathsf{E}} & \xrightarrow{\wedge} & 2^{\mathsf{E}} \end{array} \tag{2.3.12}$$

Figure 2.3.1 illustrates the indicator function for simple planar shapes, the implicit formulation of Boolean union/intersection via logical disjunction/conjunction, respectively, and its extension to the simplest form of Minkowski sum when one of the two shapes is sampled into a discrete set of points.

The strict isomorphism between Boolean algebra of sets and first-order logic over their indicator functions does not provide additional flexibility—e.g., to extend the formulation to infinite Boolean operations to formulate Minkowski operations, sweep/unsweep, etc. Alternative implicit formulations such as min/max-algebras [286] and R-functions [312] are subject to similar limitations, and do not extend beyond regularized Boolean operations. To overcome this, I shall introduce an $\mathbb{R}_+-$algebra in which logical gates are replaced with real number arithmetics, finite set unions turn into finite sums, which are generalized for infinite unions to integrals which, in turn, are coupled with intersections via De Morgan's laws and motion inversion, as described above.

As a first step, the binary indicator functions $\mathbf{1}_S : \mathsf{E} \to \{0, 1\}$ are embedded into the space of nonnegative real-valued functions via $\mathsf{emb} : \{0, 1\}^{\mathsf{E}} \hookrightarrow \mathbb{R}_+^{\mathsf{E}}$ in the obvious way:, i.e., $\mathsf{emb}(\mathbf{x} \to 0) = (\mathbf{x} \to 0)$ and $\mathsf{emb}(\mathbf{x} \to 1) = (\mathbf{x} \to 1)$.[31] The support map is accordingly replaced with the general $t-$superlevel set generator $U_t : \mathbb{R}_+^{\mathsf{E}} \to \mathscr{P}(\mathsf{E})$ whose restriction to the domain of binary-valued functions is sufficient when working with indicator functions:

$$0 \leq t < 1 \;\Rightarrow\; U_t(\mathbf{1}_S) = S, \quad \text{where} \quad U_t(f) = \{\mathbf{x} \in \mathsf{E} \mid f(\mathbf{x}) > t\}. \tag{2.3.13}$$

[31] The embedding maps the logical "false"/"true" denoted by 0/1 to arithmetic 0/1. The conceptual difference between the two are emphasized to a lesser extent hereafter to simply the notations.

Definition 2.3.2. (Defining Functions) To enable an "analytic" formulation of the geometric modeling operations, each shape $S \in \mathscr{P}(\mathsf{E})$ can be implicitly described as the $0-$superlevel set of a nonnegative real-valued function $f_S : \mathsf{E} \to \mathbb{R}_+$ where $\mathbb{R}_+ := [0, +\infty)$ is the set of all nonnegative real numbers:

$$f_S = \mathsf{desc}(S) \in \mathbb{R}_+^{\mathsf{E}}, \quad \text{where} \quad S = U_0(f_S), \quad \text{i.e.,} \quad (U_0 \circ \mathsf{desc}) = \mathsf{id}_{\mathscr{P}(\mathsf{E})}, \tag{2.3.14}$$

where $f_{(\cdot)} = \mathsf{desc}(\cdot)$ are interchangeable notations for the so-called 'defining function' of a shape (or 'shape descriptor' for short), conceptualized by the injective (but not surjective) mapping $\mathsf{desc} : \mathscr{P}(\mathsf{E}) \to \mathbb{R}_+^{\mathsf{E}}$, which is not unique as with most implicit descriptions. Thus the indicator function can be viewed as the simplest form of defining functions with a particular choice of $\mathsf{desc} := \mathsf{indc}$ within the pool of all possible injections.[32]

This embedding can be visualized via the following commutative diagram:

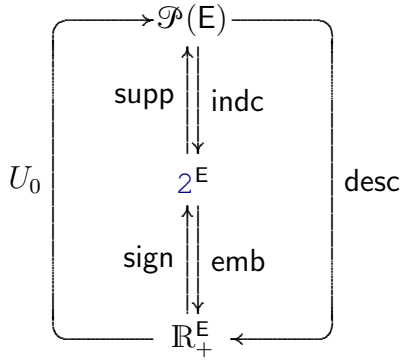

(2.3.15)

in which the sign function $\mathsf{sign} : \mathbb{R}_+^{\mathsf{E}} \to \{0, 1\}^{\mathsf{E}}$ is defined as the projection such that $(\mathsf{sign} \circ \mathsf{emb}) = \mathsf{id}_{2^{\mathsf{E}}}$, i.e., $\mathsf{sign}(\mathbf{x} \mapsto 0) = (\mathbf{x} \mapsto 0)$ and $\mathsf{sign}(\mathbf{x} \mapsto c) = (\mathbf{x} \mapsto 1)$ if $c > 0$.

[32] To be more accurate, the embedding of the characteristic functions from $2^{\mathsf{E}} = \{0, 1\}^{\mathsf{E}}$ into $\mathbb{R}_+^{\mathsf{E}}$—mapping the logical "false"/"true" denoted by $0/1$ to arithmetic $0/1$—is of concern. The conceptual difference between the two are hereafter ignored to simply the notations.

The logical negation on indicator functions can be extended to a pushed-negation on real-valued defining functions $\neg : \mathbb{R}_+^{\mathsf{E}} \to \mathbb{R}_+^{\mathsf{E}}$ obtained as $\neg f := (\neg \circ \mathsf{sign})(f)$ which maps every defining function of a given shape to the indicator function (among other possible choices of defining functions) of its complement:

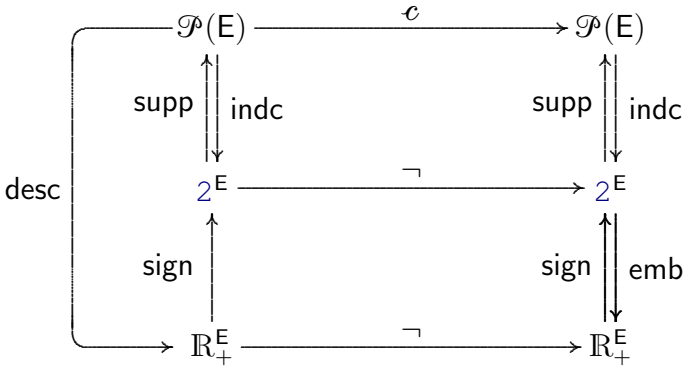

(2.3.16)

As shown for indicator functions in (2.3.6), defining functions are also covariant with rigid (and other invertible) maps, i.e., $f_{\mathbb{c}S} = \big(f_S \circ \mathsf{act}(\mathbb{c}^{-1})\big)$ for all $\mathbb{c} \in \mathsf{C}$, i.e.,

$$\big(\mathsf{desc} \circ \mathsf{act}(\mathbb{c})\big)(\cdot) = \big(f_{(\cdot)} \circ \mathsf{act}(\mathbb{c}^{-1})\big) \;\rightleftharpoons\; U_0\big((\cdot) \circ \mathsf{act}(\mathbb{c}^{-1})\big) = \big(\mathsf{act}(\mathbb{c}) \circ U_0\big)(\cdot), \quad (2.3.17)$$

which can be illustrated using a fully commutative diagram of the form:

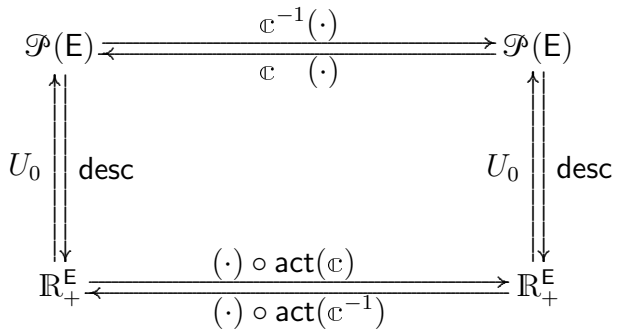

(2.3.18)

The most important of reasons behind the choice of nonnegative codomain and $0-$superlevel sets is their ability to describe set-theoretic Boolean operations $\cup, \cap$:

$\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ on sets (i.e., explicit shapes) in terms of $\mathbb{R}_+-$arithmetic operations $+, \bullet \colon \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} \to \mathbb{R}_+^{\mathsf{E}}$ on their defining functions (i.e., implicit shapes) as

$$(S_1 \cup S_2) = U_0(f_{S_1} + f_{S_2}), \quad \text{and} \quad (S_1 \cap S_2) = U_0(f_{S_1} \bullet f_{S_2}), \tag{2.3.19}$$

where the argument of U_0 is not an indicator function in general, even if the input defining functions are chosen to be indicator functions:[33]

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad\cup\quad} & \mathscr{P}(\mathsf{E}) \\ \Big\uparrow U_0 \times U_0 \Big\downarrow & & \Big\uparrow U_0 \\ \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & \xrightarrow{\quad+\quad} & \mathbb{R}_+^{\mathsf{E}} \end{array} \qquad \begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad\cap\quad} & \mathscr{P}(\mathsf{E}) \\ \Big\uparrow U_0 \times U_0 \Big\downarrow & & \Big\uparrow U_0 \\ \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & \xrightarrow{\quad\bullet\quad} & \mathbb{R}_+^{\mathsf{E}} \end{array} \tag{2.3.20}$$

Figure 2.3.2 illustrates the indicator function for simple planar shapes, the implicit formulation of Boolean union/intersection via logical disjunction/conjunction, respectively, and its extension to the simplest form of Minkowski sum when one of the two shapes is sampled into a discrete set of points.

If the shapes are restricted to (closed or open) regular sets, as required by the dimensional homogeneity postulate of solid modeling in Section 2.2.1, one can describe regularized Boolean operations $\cup^*, \cap^* : \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) \to \mathscr{P}^*(\mathsf{E})$ on regular shapes as

$$(S_1 \cup^* S_2) = U_0^*(f_{S_1} + f_{S_2}), \quad \text{and} \quad (S_1 \cap^* S_2) = U_0^*(f_{S_1} \bullet f_{S_2}), \tag{2.3.21}$$

[33] This is more of a problem for the summation but not for multiplication—since $1 + 1 = 2$ but $1 \bullet 1 = 1$. In other words, summations of real numbers 0/1 is not idempotent while their multiplications are. An obvious way to overcome this is to alter the algebraic foundations by positing $1+1 := 1$ which is equivalent to using logical "or" with bits 0/1 instead of arithmetic sum with real numbers 0/1. However, $\mathbb{R}_+-$algebra and the fact that summation keeps record of some measure of accumulation proves to be advantageous when formulating Minkowski operations as convolutions.

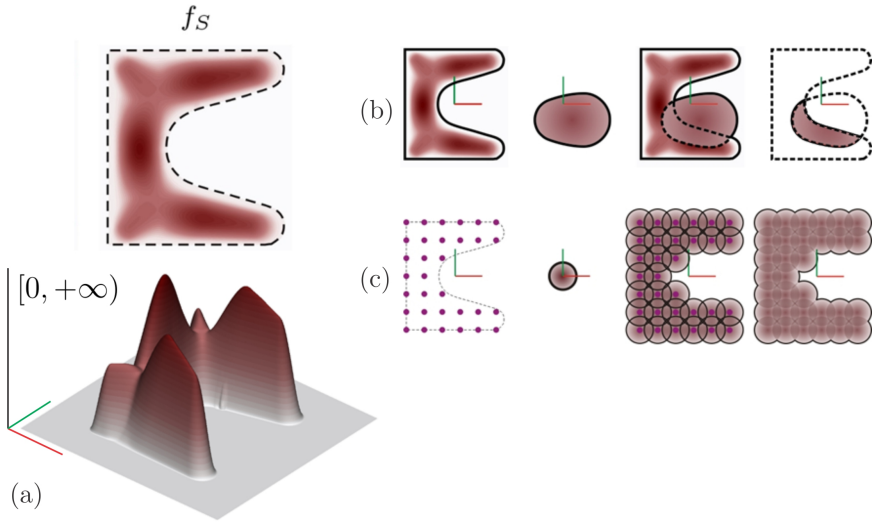

FIGURE 2.3.2: Describing a 2D shape with its defining function (a) allows an implicit formulation of Boolean union/intersection with arithmetic sum/product, respectively (b). The idea can be extended to the Minkowski sum of discrete samples with a given continuum shape, viewed as a finite union of displaced instances of the later—thus finite summation of displaced defining functions (c). The benefit of $\mathbb{R}_+-$arithmetic is that finite sums can be generalized to infinite sums (i.e., integrals). Compare with Fig. 2.3.1.

where $U_0^* = (r \circ U_0) : \mathbb{R}_+^{\mathsf{E}} \to \mathscr{P}^*(\mathsf{E})$, i.e., $U_0^*(S) = rU_0(S)$ denotes the regularized $0-$superlevel set. This is illustrated via commutative diagrams:

$$\begin{array}{ccc} \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cup^*} & \mathscr{P}^*(\mathsf{E}) \\ \Big\uparrow U_0^* \times U_0^* & & \Big\uparrow U_0^* \\ \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & \xrightarrow{+} & \mathbb{R}_+^{\mathsf{E}} \end{array} \qquad \begin{array}{ccc} \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cap^*} & \mathscr{P}^*(\mathsf{E}) \\ \Big\uparrow U_0^* \times U_0^* & & \Big\uparrow U_0^* \\ \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & \xrightarrow{\bullet} & \mathbb{R}_+^{\mathsf{E}} \end{array} \tag{2.3.22}$$

It is extremely helpful if such an explicit regularization could be avoided altogether, giving rise to the following important question: If one works entirely with defining functions in the implicit realm, what is the proper formalism to replace the

concept of regularization?[34] Starting from the observation that set-theoretic regularization corresponds to the elimination of lower-dimensional features that are not *measurable* via Lebesgue 3−measure (i.e., have zero volume), I aim to show that working with *equivalence classes* of defining functions, which cannot be distinguished using a measure-theoretic lens, is the proper setting to formulate an implicit regularization.

In spite of the one-to-one correspondence established via (2.3.2) between the sets and their characteristic functions, there are at least two limitations when using characteristic functions to establish analytic formulations of the fundamental algebraic (e.g., Boolean, Minkowski, etc.) operations for geometric modeling:

1. Characteristic functions are not closed under arithmetic summation employed in (2.3.19) or (2.3.21), which is an essential component of generalizing finite unions (described as sums) to infinite unions (described as integrals). More specifically, $(\mathbf{1}_{S_1} + \mathbf{1}_{S_2}) = f_{(S_1 \cup S_2)} \neq \mathbf{1}_{(S_1 \cup S_2)}$. Note that no such difficulty arises when dealing with intersections, i.e., $(\mathbf{1}_{S_1} \bullet \mathbf{1}_{S_2}) = f_{(S_1 \cap S_2)} = \mathbf{1}_{(S_1 \cap S_2)}$.

2. Characteristic functions are closed under arithmetic multiplication; however, when shapes are restricted to closed regular sets, multiplication (corresponding to set-theoretic intersection) does not guarantee the same algebraic closure to maintain regularity, as it might produce lower-dimensional features. More specifically, $(\mathbf{1}_{S_1} \bullet \mathbf{1}_{S_2}) = \mathbf{1}_{(S_1 \cap S_2)} \neq \mathbf{1}_{(S_1 \cap^* S_2)}$. Note that unlike intersections, regularization does not affect unions, i.e., $(\mathbf{1}_{S_1} + \mathbf{1}_{S_2}) = f_{(S_1 \cup S_2)} = f_{(S_1 \cup^* S_2)}$.

Both of the above problems can be solved by using equivalence classes of defining

[34]Note that if the input sets are restricted to regular sets, depending on the chosen regularization convention, only one of the two Boolean operations in (2.3.21) require explicit regularization. In particular, if closed-regularization is used, then regular sets are closed under unions (but not under intersections) thus U_0^* can be replaced safely with U_0 for unions (but not intersections) in (2.3.21). If open-regularization is used, the same statement applies by swapping unions and intersections.

functions—represented by characteristic functions, but including a wide range of other functions (called level set defining classes) elaborated below.

2.3.1.2 Equivalence Classes of Ordinary Defining Functions:

To solve the first problem, the key observation is that as far as Boolean operations are concerned, the actual value of the function $f_S(\mathbf{x})$ at a given query point $\mathbf{x} \in \mathsf{E}$ is of little relevance for $S := (S_1 \cup S_2)$, $S := (S_1 \cap S_2)$, etc., while its sign (i.e., zero or positive) determines the outcome—i.e., $\mathbf{x} \in S$ iff $f_S(\mathbf{x}) > 0$. Thus all defining functions that describe the same $0-$superlevel set should be deemed equivalent to enable mapping the idempotence of Boolean operations—i.e., $(S \cup S) = (S \cap S) = S$ for all $S \in \mathscr{P}(\mathsf{E})$—to that of arithmetic operations up to $0-$superlevel equivalence—i.e., $(f + f) \equiv (f \bullet f) \equiv f$, and so one for all $f \in \mathbb{R}_+^{\mathsf{E}}$. This correspondence is made precise by the following definition:

Definition 2.3.3. The equivalence class $[\mathbf{1}_S] \in \mathbb{R}_+^{\mathsf{E}}/\equiv$ includes all nonnegative real-valued functions over $\mathsf{E} = \mathbb{R}^3$ that share the common $0-$superlevel set $S \in \mathscr{P}(\mathsf{E})$:

$$[f] := \big\{ f' : \mathsf{E} \to \mathbb{R}_+ \mid U_0(f') = U_0(f) \big\}. \tag{2.3.23}$$

The corresponding equivalence relation "$\equiv$" $\subseteq \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}}$ is thus defined as

$$f_1 \equiv f_2 \quad \text{iff} \quad f_1(\mathbf{x}) > 0 \;\rightleftharpoons\; f_2(\mathbf{x}) > 0, \text{ for all } \mathbf{x} \in \mathsf{E}. \tag{2.3.24}$$

Thus one obtains $(\mathbf{1}_{S_1} + \mathbf{1}_{S_2}) \equiv \mathbf{1}_{(S_1 \cup S_2)}$ even though $(\mathbf{1}_{S_1} + \mathbf{1}_{S_2}) \neq \mathbf{1}_{(S_1 \cup S_2)}$.

The above definition allows to construct a closed set-theoretic Boolean algebra over the $0-$superlevel defining classes of shapes, to which I shall return in Chapter 3.

To solve the second problem, let us revisit the concept of properties that hold 'almost everywhere' presented in Convention 2.2.13. In particular, the equality$-\mathsf{ae}$ defined in Definition 2.2.16 can be restated in terms of characteristic functions as

$$S_1 \overset{\mathsf{ae}}{=} S_2 \quad \text{iff} \quad \langle \mathbf{1}_{S_1}, \mathbf{1}_\Omega \rangle > 0 \;\rightleftharpoons\; \langle \mathbf{1}_{S_2}, \mathbf{1}_\Omega \rangle > 0 \quad \text{for all } \Omega \in \mathscr{P}(\mathsf{E}), \tag{2.3.25}$$

noting that the two inner products are in face the same as the measure of intersection $\mu^3[S_1 \cap \Omega] = \mu^3[S_2 \cap \Omega]$. One can also use the following (more general) definition of equality$-\mathsf{ae}$ in terms of general nonnegative real-valued defining functions:

$$S_1 \overset{\mathsf{ae}}{=} S_2 \quad \text{iff} \quad \langle f_{S_1}, f_\Omega \rangle > 0 \;\rightleftharpoons\; \langle f_{S_2}, f_\Omega \rangle > 0 \quad \text{for all } \Omega \in \mathscr{P}(\mathsf{E}). \tag{2.3.26}$$

Recalling that given two 3$-$measurable sets $S_1, S_2 \in \mathscr{P}(\mathsf{E})$, $\mathscr{r}(S_1) = \mathscr{r}(S_2)$ iff $S_1 \overset{\mathsf{ae}}{=} S_2$, the following combines equality$-\mathsf{ae}$ and 0$-$superlevel set equivalence via $\equiv$ to define regularized 0$-$superlevel set equivalence via $\overset{*}{\equiv}$:

Definition 2.3.4. The equivalence class $[\mathbf{1}_S]^* \in \mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}$ includes all nonnegative real-valued functions over $\mathsf{E} = \mathbb{R}^3$ that share the common regularized 0$-$superlevel set $S \in \mathscr{P}(\mathsf{E})$, i.e., whose 0$-$superlevel sets are equal$-\mathsf{ae}$:[35]

$$[f]^* := \big\{ f' : \mathsf{E} \to \mathbb{R}_+ \mid U_0^*(f') = U_0^*(f) \big\}, \tag{2.3.27}$$

whose alternative definition, according to Proposition 2.2.15, is the following:

$$[f]^* := \big\{ f' : \mathsf{E} \to \mathbb{R}_+ \mid U_0(f') \overset{\mathsf{ae}}{=} U_0(f) \big\}. \tag{2.3.28}$$

[35]To deal with the defining functions of sets that are not Lebesgue 3$-$measurable at all, let us think of them all being dumped into a single outlier equivalence class. The rest of the quotient space is in a one-to-one correspondence with the class of regular sets $\mathscr{P}^*(\mathsf{E})$.

The corresponding equivalence relation "$\overset{*}{\equiv}$" $\subseteq \mathbb{R}^{\mathsf{E}}_{+} \times \mathbb{R}^{\mathsf{E}}_{+}$ is thus defined as

$$f_1 \overset{*}{\equiv} f_2 \quad \text{iff} \quad f_1(\mathbf{x}) > 0 \overset{\mathsf{ae}}{\rightleftharpoons} f_2(\mathbf{x}) > 0, \text{ for all } \mathbf{x} \in \mathsf{E}. \tag{2.3.29}$$

Thus one obtains $(\mathbf{1}_{S_1} \bullet \mathbf{1}_{S_2}) \overset{*}{\equiv} \mathbf{1}_{(S_1 \cap^* S_2)}$ even though $(\mathbf{1}_{S_1} \bullet \mathbf{1}_{S_2}) \neq \mathbf{1}_{(S_1 \cap^* S_2)}$.

The above results can be summarized into the following commutative diagram:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad [\cdot]^{\dagger} \quad} & \mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=} \\ \Big\uparrow U_0 \quad \Big\downarrow \mathsf{desc} & & \Big\uparrow U_0 \quad \Big\downarrow \mathsf{desc} \\ \mathbb{R}^{\mathsf{E}}_{+} & \xrightarrow{\quad [\cdot]^{*} \quad} & \mathbb{R}^{\mathsf{E}}_{+}/\overset{*}{\equiv} \end{array} \tag{2.3.30}$$

Composing the implicit regularization formulae in (2.3.30) and (2.2.25) yields:

$$\begin{array}{ccccc} \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad \mathscr{r}(\cdot) \quad} & \mathscr{P}^{*}(\mathsf{E}) & = & \mathscr{P}(\mathsf{E})/\overset{\mathsf{ae}}{=} \\ \Big\uparrow U_0 \quad \Big\downarrow \mathsf{desc} & & & & \Big\uparrow U_0 \quad \Big\downarrow \mathsf{desc} \\ \mathbb{R}^{\mathsf{E}}_{+} & & \xrightarrow{\qquad [\cdot]^{*} \qquad} & & \mathbb{R}^{\mathsf{E}}_{+}/\overset{*}{\equiv} \end{array} \tag{2.3.31}$$

The above definition allows to construct a closed regularized Boolean algebra over the 0−superlevel defining classes of shapes, to which I shall return in Chapter 3.

Note also that the equality−ae is a weaker condition than the strict equality of 0−superlevel sets, implying that $\mathbb{R}^{\mathsf{E}}_{+}/\overset{*}{\equiv}$ is a coarser partitioning than $\mathbb{R}^{\mathsf{E}}_{+}/\equiv$, i.e., $[\mathbf{1}_S] \subset [\mathbf{1}_S]^*$ thus $f_1 \equiv f_2$ implies $f_1 \overset{*}{\equiv} f_2$ while the converse may not be true.

Profile-Indifference. The ideas of the equivalence classes in Definitions 2.3.3 and 2.3.4 is illustrated in Fig. 2.3.3 for simple planar shapes with $0-$, $1-$, and $2-$dimensional features. For every shape, the defining functions that are supported by the same exact (regular or irregular) set are equivalent via $\equiv$ regardless of the function profile, to which I refer as a 'profile-indifferent' or 'profile-oblivious' shape description scheme. This is only useful because the equivalence classes themselves form a Boolean ring by extending the definitions of $+, \bullet : \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} \to \mathbb{R}_+^{\mathsf{E}}$ to $+, \bullet : (\mathbb{R}_+^{\mathsf{E}}/\equiv) \times (\mathbb{R}_+^{\mathsf{E}}/\equiv) \to (\mathbb{R}_+^{\mathsf{E}}/\equiv)$ as

$$[f_1] + [f_2] := \big\{ f_1 + f_2 \mid f_1 \in [f_1] \text{ and } f_2 \in [f_1] \big\} = [f_1 + f_2], \tag{2.3.32}$$

$$[f_1] \bullet [f_2] := \big\{ f_1 \bullet f_2 \mid f_1 \in [f_1] \text{ and } f_2 \in [f_1] \big\} = [f_1 \bullet f_2], \tag{2.3.33}$$

which is made possible by the fact that equivalence classes are closed under arithmetic operations on their member functions, thus $[\cdot]$ commutes with $+/\bullet$. This is extremely important, as it makes the above operations computable between pairs of equivalence classes by arbitrarily selecting representative functions (e.g., indicator functions) in each class and computing arithmetic operations on those representatives. In other words, the above relationships enable a profile-oblivious Boolean algebra:

$$(S_1 \cup S_2) = U_0(f_{S_1} + f_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(S_1 \cup S_2) \equiv (f_{S_1} + f_{S_2}), \tag{2.3.34}$$

$$(S_1 \cap S_2) = U_0(f_{S_1} \bullet f_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(S_1 \cap S_2) \equiv (f_{S_1} \bullet f_{S_2}). \tag{2.3.35}$$

Thus the previous commutative diagrams in (2.3.20) are modified as follows:

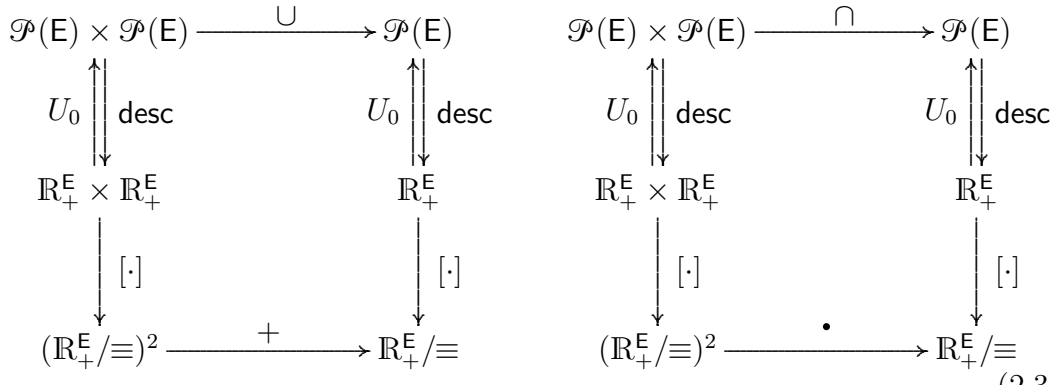

(2.3.36)

Irregularity Tolerance. On the other hand, the defining functions that are supported by the same set 'almost everywhere' are equivalent via $\overset{*}{\equiv}$ regardless of the function profile and also irrespective of the lower-dimensional features. In contrast to the pointset-topological approach to regularization that requires explicit removal of irregularities (e.g., dangling edges and singular points), this measure-theoretic notion of regularization can be viewed as an implicit elimination of their effects by thinking of all instances of deviation from the regular base shape collectively as a single mathematical entity (namely, the defining class $[\mathbf{1}_S]^* \in \mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}$), to which I refer as an 'irregularity-tolerant' shape description scheme. Once again, this is only useful because the equivalence classes themselves form a Boolean ring by extending the definitions of $+, \bullet : \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} \to \mathbb{R}_+^{\mathsf{E}}$ to $+, \bullet : (\mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}) \times (\mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}) \to (\mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv})$ as

$$[f_1]^* + [f_2]^* := \big\{ f_1 + f_2 \mid f_1 \in [f_1]^* \text{ and } f_2 \in [f_1]^* \big\} = [f_1 + f_2]^*, \tag{2.3.37}$$

$$[f_1]^* \bullet [f_2]^* := \big\{ f_1 \bullet f_2 \mid f_1 \in [f_1]^* \text{ and } f_2 \in [f_1]^* \big\} = [f_1 \bullet f_2]^*, \tag{2.3.38}$$

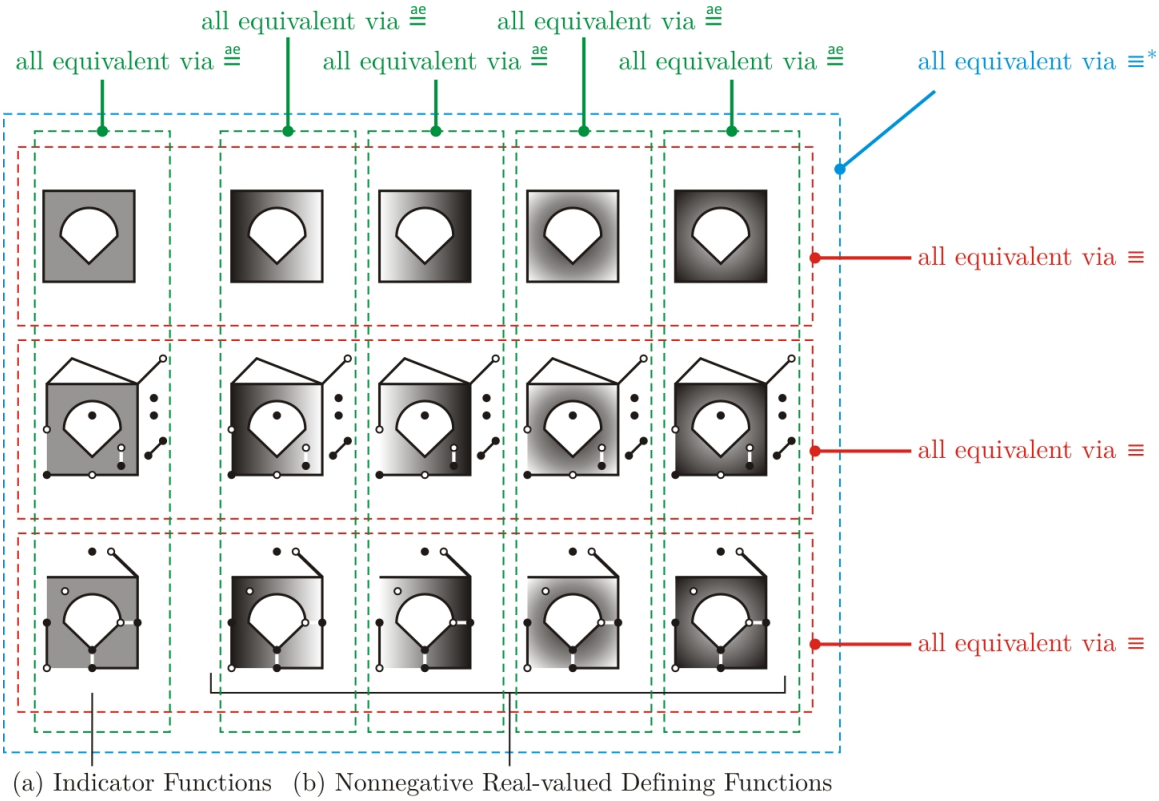

FIGURE 2.3.3: Each binary-valued indicator functions in (a) is equivalent via $\equiv$ to all non-negative real-valued defining functions that are supported by the same exact shape—i.e., have the same 0−superlevel set—exemplified by the functions in the same row in (b). Note that lower-dimensional features (e.g., dangling edges or singular points) must also match for equivalence via $\equiv$. However, all defining functions that are supported by the same shape 'almost everywhere'—i.e., have the same regularized 0-superlevel set—exemplified across different rows in (b) are equivalent via $\overset{*}{\equiv}$, making lower-dimensional features *irrelevant*.

which is made possible by the fact that equivalence classes are closed under arithmetic operations on their member functions, thus $[\cdot]^*$ commutes with $+/\bullet$—once again, making the above operations computable for higher-level abstractions such as equivalence classes, enabling an irregularity-tolerant Boolean algebra:

$$(S_1 \cup^* S_2) = U_0^*(f_{S_1} + f_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(S_1 \cup^* S_2) \overset{*}{\equiv} (f_{S_1} + f_{S_2}), \tag{2.3.39}$$

$$(S_1 \cap^* S_2) = U_0^*(f_{S_1} \bullet f_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(S_1 \cap^* S_2) \overset{*}{\equiv} (f_{S_1} \bullet f_{S_2}). \tag{2.3.40}$$

Thus the previous commutative diagrams in (2.3.22) are modified as follows:

$$\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cup^*} & \mathscr{P}^*(\mathsf{E}) \\
{\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} \times {\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} & & {\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} \\
\mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & & \mathbb{R}_+^{\mathsf{E}} \\
\downarrow{\scriptstyle [\cdot]^*} \times \downarrow{\scriptstyle [\cdot]^*} & & \downarrow{\scriptstyle [\cdot]^*} \\
(\mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv})^2 & \xrightarrow{+} & \mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}
\end{array}
\qquad
\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cap^*} & \mathscr{P}^*(\mathsf{E}) \\
{\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} \times {\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} & & {\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} \\
\mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} & & \mathbb{R}_+^{\mathsf{E}} \\
\downarrow{\scriptstyle [\cdot]^*} \times \downarrow{\scriptstyle [\cdot]^*} & & \downarrow{\scriptstyle [\cdot]^*} \\
(\mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv})^2 & \xrightarrow{\bullet} & \mathbb{R}_+^{\mathsf{E}}/\overset{*}{\equiv}
\end{array} \tag{2.3.41}$$

A natural question is whether more complex geometric modeling operations can be abstracted in terms of familiar analytic operations on these defining functions, and whether the aforementioned equivalence classes of functions enjoy a similar algebraic closure as the one established for (ordinary and regularized) Boolean operations? I show in Chapter 3 that this is in fact the case for an important class of morphological operations that are extremely helpful in uniformly formulating and solving a variety of geometric and spatial computing problems, exemplified in Chapter 5.

A significant challenge to overcome is that, although lower-dimensional features cannot create full-dimensional features after a finite sequence of Boolean operations,

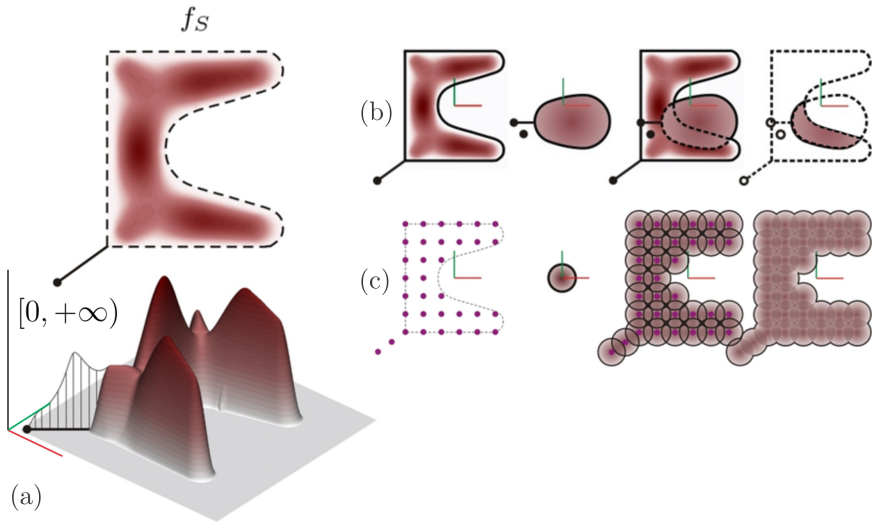

FIGURE 2.3.4: Unlike finite Boolean operations (b) that are indifferent to irregularities, as they endure as lower-dimensional features but cannot create full-dimensional features—the same is not true for Minkowski operations.

this is not true for infinite number of such operations (e.g., Minkowski operations, sweeps/unsweeps, etc.). In other words, irregularity tolerance is limited to finite sequences of Boolean operations, through which the irregularities endure to exist as they are passed downstream but their collective 3−measure remains zero. However, Minkowski operations and other kinds of infinite Boolean combinations can create and destroy full-dimensional features.[36] See Fig. 2.3.4 for a simple example. I will address this challenge by introducing 3−measurable singularities in Section 2.3.2.

To summarize, the equivalence relation $\overset{*}{\equiv}$ establishes a one-to-one correspondence (i.e., isomorphism) between set-theoretic abstraction of traditional 'solids' $S \in \mathscr{P}^*(\mathsf{E})$ and algebras on them on the one hand, and the corresponding defining class abstraction $[\mathbf{1}_S]^* \in \mathbb{R}^{\mathsf{E}}_+/\overset{*}{\equiv}$ and algebras on them on the other hand. However, by acknowl-

[36]For example, the Minkowski sum of a d_1-dimensional shape with a d_2-dimensional shape is at most $(\mathrm{d}_1+\mathrm{d}_2)-$dimensional [235]—e.g., a curve segment swept along another curve segment can create a surface. When $\mathrm{d}_1 = 0$ or $\mathrm{d}_2 = 0$, the dimension cannot increase, thus makes the operation indifferent (i.e., tolerant) to the irregularities.

edging the distinct identity of each function $f \in [f]^*$, one deals with more complex abstractions of geometric objects that cannot be captured by the traditional models. This is clearly depicted by the variety of objects in Fig. 2.3.3 that can represent 'fuzzy' shapes—e.g., think of proteins with nonuniform electron distributions.

Indicator Retrieval. It is important to note that the main point of using real-valued functions and fuzzy sets instead of binary-valued functions and crisp sets is to step beyond the profile-oblivious abstraction—otherwise the reformulation would be theoretically limited to the same capabilities of traditional models by strict isomorphism. In other words, I show in Chapter 5 that a range of important engineering problems can be solved by choosing an appropriate profile for the defining function that capture certain combinatorial, topological, and geometric properties that are relevant to the problem. Here it suffices to say that for computing a large subset of morphological operations, indicator function alone suffices as an input, even though a different defining function will be returned as the output [215, 218]. Moreover, indicator functions uniquely and unambiguously describes their equivalence classes via $\equiv$ and is the most memory-efficient representation (for fixed $3-$space decomposition) among other members of the class using bits alone. Thus retrieving $\mathbf{1}_S \in [\mathbf{1}_S]$ from the class $[\mathbf{1}_S] \in \mathbb{R}_+^{\mathsf{E}}$ is extremely important from both theoretical and computational perspectives, which is as simple as computing the sign of an arbitrary member:

$$\mathbf{1}_S := (\mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S], \quad \text{i.e., } \mathbf{1}_S(\mathbf{x}) = \begin{cases} 1 & \text{if } \big(\mathsf{choose}[\mathbf{1}_S]\big)(\mathbf{x}) > 0, \\ 0 & \text{if } \big(\mathsf{choose}[\mathbf{1}_S]\big)(\mathbf{x}) = 0, \end{cases} \tag{2.3.42}$$

where $\mathsf{choose} : [f] \mapsto f$ is a function that chooses an arbitrary function with the same $0-$superlevel set, i.e., $U_0(f) = U_0(f')$ for all $f' \in [f]$. This completes the diagrams in (2.3.36) to allow vertical transforms in both directions:

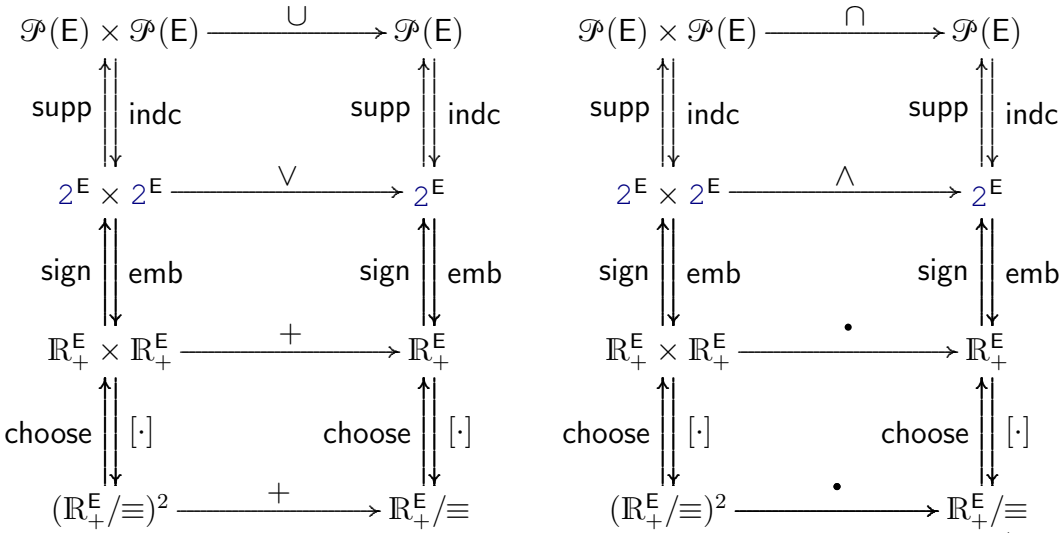

(2.3.43)

On the other hand, indicator functions do not uniquely describe their equivalence classes via $\overset{*}{\equiv}$, meaning that $(\mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S]^*$ returns a set of infinitely many different (though equal$-\mathsf{ae}$) indicator functions that look like the left-most column of Fig. 2.3.3, one of which needs to be selected according to some deterministic set of rule. For computational purposes, choosing the indicator function of the regularized $0-$superlevel set (e.g., the top-left instance in Fig. 2.3.3) appears to be the most appealing rule to follow, by defining $\mathsf{choose}^* : [f]^* \mapsto f$ such that $U_0(f) = U_0^*(f')$ for all $f' \in [f]^*$. However, it requires explicit pointset-topological regularization that we tried to avoid by introducing the notion of equality$-\mathsf{ae}$ in the first place. I show in Section 2.3.3 that pointset-topological regularization is also possible via implicit methods by applying a convolution-based filter to an arbitrary $\mathbf{1}_S := (\mathsf{choose} \circ \mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S]^*$.

The above notions developed for implicit operations on shapes $S \subseteq \mathsf{E}$ can be readily extended to implicit operations on motions $M \subseteq \mathsf{C}$. The most important distinctive technicalities are due to the generalization of measure-theoretic concepts to the Lie group $\mathsf{C} = \mathrm{SE}(3)$ by introducing Haar measures as the natural extensions of the Lebesgue measure. Once again, noting that $\mathsf{E} = \mathbb{R}^3$ is (isomorphic to) a subgroup of $\mathsf{C} = \mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$, the relationships defined over the former can be viewed as restrictions of those defined over the latter to the translation subgroup $\mathrm{T}(3) \cong \mathbb{R}^3$.

2.3.1.3 Ordinary Interference Test and Configuration Space Obstacles:

As I shall show in this chapter and the next, unions/intersections play a critical role in explicit/implicit formulation of shape and motion related problems. By dualities established via diagrams (2.3.5) and (2.3.7), having a gadget to test for one of them is enough, defined next.

Definition 2.3.5. (Collision Predicate) The function $\mathsf{col} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \{0, 1\}$ decides if two given shapes $S_1, S_2 \in \mathscr{P}(\mathsf{E})$ are 'colliding', defined in terms of regularized intersection as

$$\mathsf{col}(S_1, S_2) = \begin{cases} 1 & \text{if } (S_1 \cap^* S_2) \neq \emptyset, \\ 0 & \text{if } (S_1 \cap^* S_2) = \emptyset. \end{cases} \tag{2.3.44}$$

Alternatively, a collision can be defined as set-theoretic intersection over a region of nonzero 3−measure (i.e., non-null volume), i.e.,

$$\mathsf{col}(S_1, S_2) = \begin{cases} 1 & \text{if } \mu^3[S_1 \cap S_2] > 0, \\ 0 & \text{if } \mu^3[S_1 \cap S_2] = 0. \end{cases} \tag{2.3.45}$$

The collision predicate is covariant under rigid transformations, i.e., $\mathsf{col}(S_1, S_2) =$

$\mathsf{col}(\mathbb{c}S_1, \mathbb{c}S_2)$ for all $\mathbb{c} \in \mathrm{SE}(3)$, meaning that the following diagram commutes:

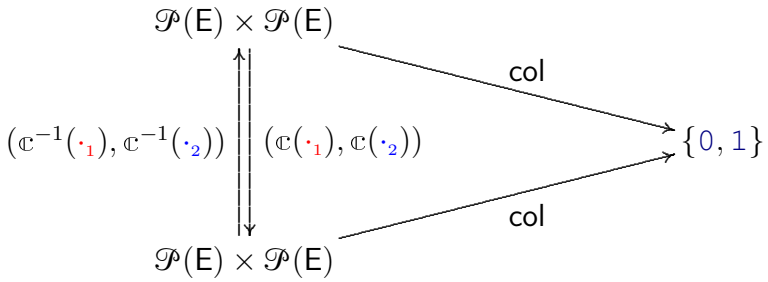

(2.3.46)

Lemma 2.3.6. (Null-Volume Lemma) *Given $S_1, S_2 \in \mathscr{P}(\mathsf{E})$, the inner product of their real-valued defining functions $f_{S_1}, f_{S_2} : \mathsf{E} \to \mathbb{R}_+$ gives a collision predicate, i.e.,*

$$\langle f_{S_1}, f_{S_2} \rangle > 0 \;\rightleftharpoons\; \mu^3[S_1 \cap S_2] > 0, \quad \text{i.e.,} \quad \mathsf{col} = \mathsf{sign} \circ \langle \cdot_1, \cdot_2 \rangle \circ \mathsf{desc}. \tag{2.3.47}$$

Furthermore, if characteristic functions are used, the inner product precisely yields the interference 3−measure (i.e., volume):

$$\langle \mathbf{1}_{S_1}, \mathbf{1}_{S_2} \rangle = \int_{\mathsf{E}} \mathbf{1}_{S_1}(\mathbf{x}) \mathbf{1}_{S_2}(\mathbf{x}) \; d\mu^3[\mathbf{x}] = \int_{S_1 \cap S_2} d\mu^3[\mathbf{x}] = \mu^3[S_1 \cap S_2]. \tag{2.3.48}$$

The parity between the pointset-topological formulation in (2.3.44) and the measure-theoretic (2.3.45) leading to (2.3.47) as a result of Lemma 2.3.6 can be illustrated via the following commutative diagram, which works regardless of regularization:

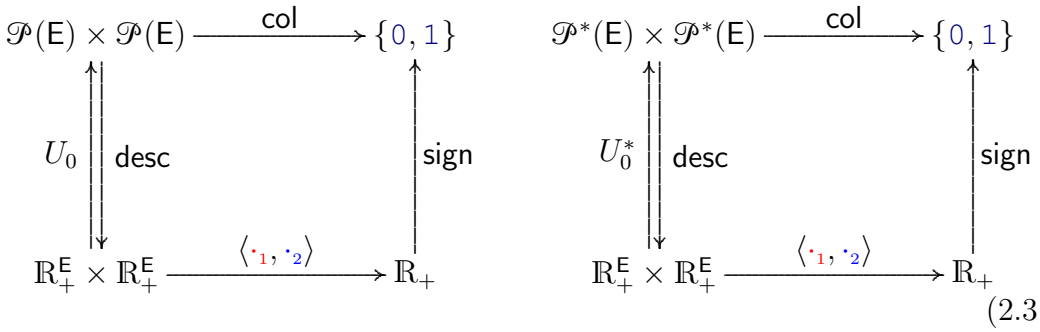

(2.3.49)

Now let the two shapes move via $\mathbb{c}_1, \mathbb{c}_2 \in \mathsf{C}$, respectively. As a result of the definition,

$$\mathsf{col}(\mathbb{c}_1 S_1, \mathbb{c}_2 S_2) = \mathsf{col}(S_1, (\mathbb{c}_1^{-1}\mathbb{c}_2)S_2) = \mathsf{col}((\mathbb{c}_2^{-1}\mathbb{c}_1)S_1, S_2), \tag{2.3.50}$$

Let $O_{S_1,S_2} \subseteq \mathsf{C}$ be the set of relative configurations $\mathbb{c} = (\mathbb{c}_1^{-1}\mathbb{c}_2) \in \mathsf{C}$ such that there is a collision, i.e., $\mathsf{col}(S_1, (\mathbb{c}S_2))$ evaluates "true":

$$O_{S_1,S_2} := \big\{\mathbb{c} \in \mathsf{C} \mid \mathsf{col}(S_1, (\mathbb{c}S_2))\big\} = \big\{\mathbb{c} \in \mathsf{C} \mid S_1 \cap^* (\mathbb{c}S_2) \neq \emptyset\big\}, \tag{2.3.51}$$

which is typically referred to as the 'configuration space obstacle' (or $\mathsf{C}-$obstacle for short). Lemma 2.3.6 implies that the inner product $\langle f_{S_1}, f_{\mathbb{c}S_2}\rangle$ gives a nonnegative real-valued defining function of the $\mathsf{C}-$obstacle:

$$O_{S_1,S_2} := \big\{\mathbb{c} \in \mathsf{C} \mid \mu^3[S_1, (\mathbb{c}S_2)] > 0\big\} = \big\{\mathbb{c} \in \mathsf{C} \mid \langle f_{S_1}, f_{\mathbb{c}S_2}\rangle > 0\big\}, \tag{2.3.52}$$

Recalling from (2.2.9) or (2.2.10), and noting also that $f_{\mathbb{c}S_2} = (f_{S_2} \circ \mathsf{act}(\mathbb{c}^{-1}))$ due to the covariance property in (2.3.17), the above inner product for different values of $\mathbb{c} \in \mathsf{C}$ gives a convolution: $\langle f_{S_1}, f_{\mathbb{c}S_2}\rangle = (f_{S_1} * \tilde{f}_{S_2})(\mathbb{c})$, therefore,

$$O_{S_1,S_2} = U_0(f_{S_1} * \tilde{f}_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(O_{S_1,S_2}) \equiv (f_{S_1} * \tilde{f}_{S_2}), \tag{2.3.53}$$

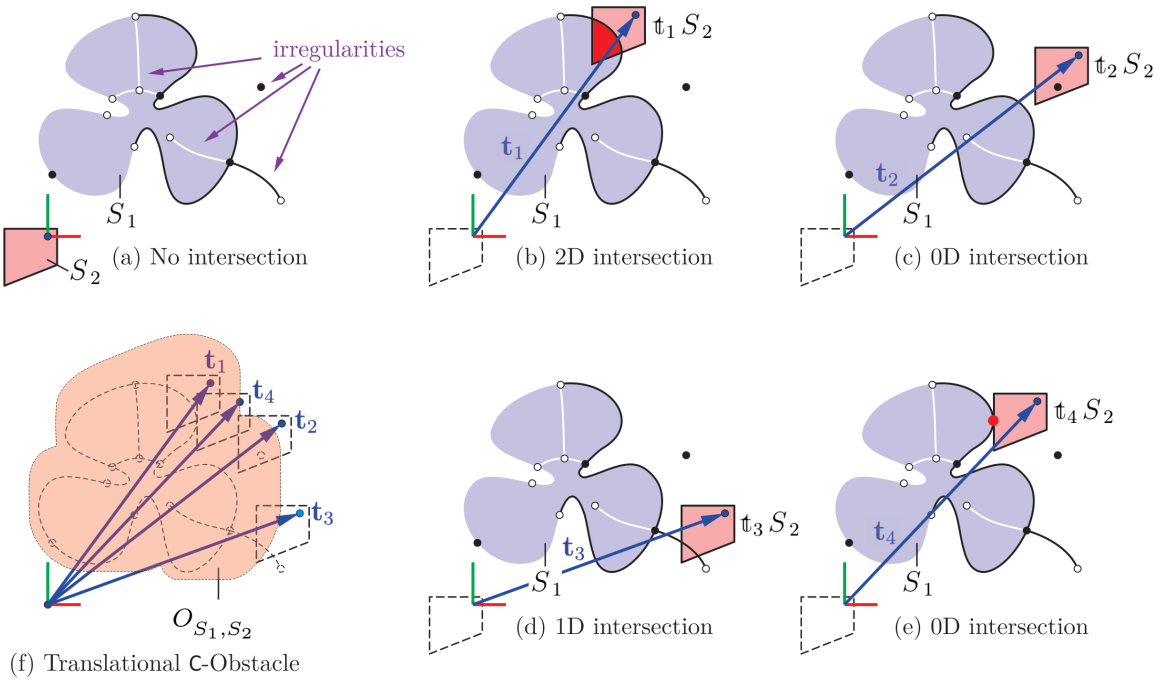

FIGURE 2.3.5: The translational C−obstacle in (f) for a pair of shapes on the 2D plane is the subset of all relative displacements at which there is an intersection of nonzero 2−measure. The nonempty 0D or 1D intersections do not count as 2−collisions.

which can be illustrated using the following commutative diagram:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{O_{(\cdot_1,\cdot_2)}} & \mathscr{P}(\mathsf{C}) \\ \Big\uparrow U_0 \Big\downarrow \mathsf{desc} & & \Big\uparrow U_0 \\ \mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & \xrightarrow{(\cdot_1) * (- \cdot_2)} & \mathbb{R}^{\mathsf{C}}_+ \end{array} \qquad \begin{array}{ccc} \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{O_{(\cdot_1,\cdot_2)}} & \mathscr{P}(\mathsf{C}) \\ \Big\uparrow U_0^* \Big\downarrow \mathsf{desc} & & \Big\uparrow U_0 \\ \mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & \xrightarrow{(\cdot_1) * (- \cdot_2)} & \mathbb{R}^{\mathsf{C}}_+ \end{array} \tag{2.3.54}$$

Figure 2.3.5 illustrates the C−obstacle for a pair of planar shapes, by restricting the motion to translation only—which can be viewed as a constant rotation slice through the full C−obstacle. Note that irregularities (i.e., lower-dimensional elements) do not affect the shape of the C−obstacle, since intersections of zero 3−measure in the

3−space do *not* qualify as 3−collisions. A generalization of the concept to capture various d−collisions for all $0 \leq \mathrm{d} < 4$ is presented in the end of Section 2.3.2. See Fig. 2.3.15 for an illustration.

A major benefit of the measure-theoretic formulation of collision predicates and C−obstacles is that one can use Fourier transforms to compute them efficiently. Using Parseval's theorem (Theorem 2.2.10), the inner product in (2.3.47) can be computed in the Fourier domain. Thus Lemma 2.3.6 can be restated via:

$$\langle \hat{f}_{S_1}, \hat{f}_{S_2} \rangle > 0 \;\rightleftharpoons\; \mu^3[S_1 \cap S_2] > 0, \quad \text{i.e.,} \quad \mathsf{col} = \mathsf{sign} \circ \langle \cdot_1, \cdot_2 \rangle \circ \mathfrak{F} \circ \mathsf{desc}, \tag{2.3.55}$$

where $\hat{f}_{S_1} = \mathfrak{F}\{f_{S_1}\}$ and $\hat{f}_{S_2} = \mathfrak{F}\{f_{S_2}\}$, noting that $\langle \cdot_1, \cdot_2 \rangle \circ \mathfrak{F} = \langle \cdot_1, \cdot_2 \rangle$ due to (2.2.13). Applying (2.2.13) to (2.3.49) extends the diagrams to

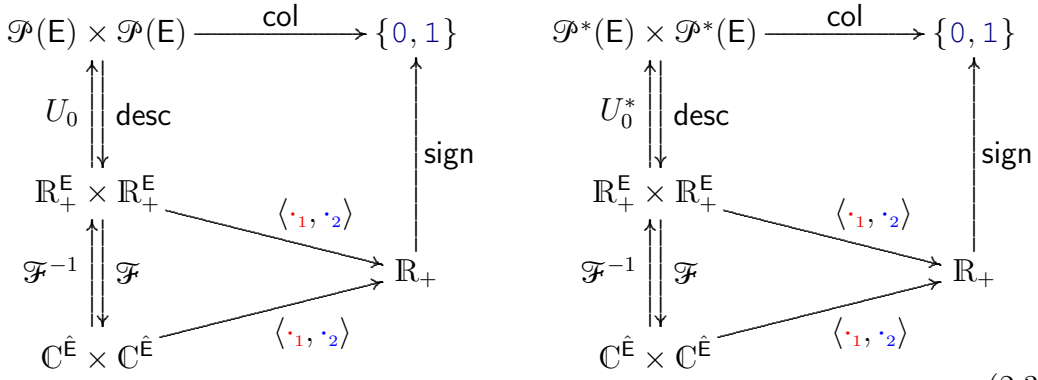

(2.3.56)

One can obtain an approximate collision predicate by truncating the Fourier expansions, i.e., low-pass filtering the Fourier descriptors $\hat{f}_{S_1}, \hat{f}_{S_2} \in \mathbb{C}^{\hat{\mathsf{E}}}$. For instance, multiplying both descriptors with the indicator function of a compact window $\hat{W}_r \subset \hat{\mathsf{E}}$ around the origin—e.g., a 3−ball in the frequency domain $\hat{W}_r := \hat{B}^3_r(\hat{\mathbf{0}})$—denoted by $\mathbf{1}_{\hat{W}_r} : \hat{\mathsf{E}} \to \{0, 1\}$ yields a low-pass filtered approximation of the inner product that retains the dominant low-frequency modes within an arbitrary distance of $r > 0$ to

the origin $\hat{\mathbf{0}} \in \hat{\mathsf{E}}$ and neglecting the high-frequency modes:

$$\langle \hat{f}_{S_1}, \hat{f}_{S_2} \rangle \approx \big\langle (\hat{f}_{S_1} \bullet \mathbf{1}_{\hat{W}_r}), (\hat{f}_{S_2} \bullet \mathbf{1}_{\hat{W}_r}) \big\rangle = \int_{\hat{W}_r} \hat{f}_{S_1}(\boldsymbol{\omega}) \hat{f}_{S_2}(\boldsymbol{\omega}) \, d\mu^3[\boldsymbol{\omega}]. \tag{2.3.57}$$

Noting that $\lim_{r \to +\infty} \hat{W}_r = \hat{\mathsf{E}}$ and assuming that shapes are compact, the frequency domain representations have decaying high-frequency amplitudes as $r \to +\infty$, thus the above approximation approaches the exact collision predicate as $r \to +\infty$.

Accordingly, the convolution theorem (Theorem 2.2.11) allows one to convert the convolution of defining functions in (2.3.53) to a pointwise multiplication of their Fourier transforms. In other words, the Fourier descriptor of the $\mathsf{C}-$obstacles can be obtained from pointwise multiplication of Fourier descriptors of the shapes—with an additional reflection $\boldsymbol{-} : \mathbb{R}^{\hat{\mathsf{C}}}_+ \to \mathbb{R}^{\hat{\mathsf{C}}}_+$:

$$O_{S_1,S_2} = (U_0 \circ \mathscr{F}^{-1})(\hat{f}_{S_1} \bullet \check{f}_{S_2}), \quad \text{i.e.,} \quad (\mathscr{F} \circ \mathsf{desc})(O_{S_1,S_2}) \equiv (\hat{f}_{S_1} \bullet \check{f}_{S_2}), \tag{2.3.58}$$

where $\check{f}_S := \tilde{\hat{f}}_S = \hat{\tilde{f}}_S$ is the reflection (i.e., inversion) of the Fourier transform, which is the same as the Fourier transform of reflection due to linearity. Accordingly, the diagram in (2.2.15) extends the diagram in (2.3.54) to:

$$\begin{array}{ccc}
\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{O_{(\cdot_1,\cdot_2)}} & \mathscr{P}(\mathsf{C}) \\
{\scriptstyle U_0}\uparrow\downarrow{\scriptstyle \mathsf{desc}} & & \uparrow{\scriptstyle U_0} \\
\mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & \xrightarrow{(\cdot_1) * (\boldsymbol{-}\, \cdot_2)} & \mathbb{R}^{\mathsf{C}}_+ \\
{\scriptstyle \mathscr{F}^{-1}}\uparrow\downarrow{\scriptstyle \mathscr{F}} & & {\scriptstyle \mathscr{F}^{-1}}\uparrow\downarrow{\scriptstyle \mathscr{F}} \\
\mathbb{C}^{\hat{\mathsf{E}}} \times \mathbb{C}^{\hat{\mathsf{E}}} & \xrightarrow{(\cdot_1) \bullet (\boldsymbol{-}\, \cdot_2)} & \mathbb{R}^{\hat{\mathsf{C}}}_+
\end{array}
\qquad
\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{O_{(\cdot_1,\cdot_2)}} & \mathscr{P}(\mathsf{C}) \\
{\scriptstyle U_0^*}\uparrow\downarrow{\scriptstyle \mathsf{desc}} & & \uparrow{\scriptstyle U_0} \\
\mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & \xrightarrow{(\cdot_1) * (\boldsymbol{-}\, \cdot_2)} & \mathbb{R}^{\mathsf{C}}_+ \\
{\scriptstyle \mathscr{F}^{-1}}\uparrow\downarrow{\scriptstyle \mathscr{F}} & & {\scriptstyle \mathscr{F}^{-1}}\uparrow\downarrow{\scriptstyle \mathscr{F}} \\
\mathbb{C}^{\hat{\mathsf{E}}} \times \mathbb{C}^{\hat{\mathsf{E}}} & \xrightarrow{(\cdot_1) \bullet (\boldsymbol{-}\, \cdot_2)} & \mathbb{R}^{\hat{\mathsf{C}}}_+
\end{array}
\tag{2.3.59}$$

Once again, the low-pass filtering trick can be used to obtain an approximation of the $\mathsf{C}-$obstacle's Fourier descriptor:

$$\mathsf{desc}(O_{S_1,S_2}) \equiv (f_{S_1} * \tilde{f}_{S_2}) = \mathscr{F}^{-1}\big\{\hat{f}_{S_1} \bullet \check{f}_{S_2}\big\} \approx \mathscr{F}^{-1}\big\{\hat{f}_{S_1} \bullet \check{f}_{S_2} \bullet \mathbf{1}_{\hat{W}_r}\big\}. \qquad (2.3.60)$$

And as before, the $\mathsf{C}-$obstacle descriptor approximation approaches the exact values at all configurations as $r \to +\infty$. However, the caveat is that this does not mean the same limit behavior for the $0-$superlevel set (i.e, $\mathsf{C}-$obstacle itself) in the Hausdorff topology: This is because of the 'rippling' effect cause by truncating the Fourier expansions. Compact shapes and motions do not have band-limited Fourier descriptors, and approximating them as such will create sinusoidal (thus non-compact) artifacts, as illustrated in Fig. 5.2.1 of Chapter 5 for a simple planar example with translational motion. Motivated by signal processing techniques, a possible solution to explore is to use smooth (e.g., Gaussian) kernels in lieu of $\mathbf{1}_{\hat{W}_r}$ to anti-alias the convolution, and use $\epsilon-$superlevel sets (for $\epsilon > 0$) to obtain a compact approximation of the $\mathsf{C}-$obstacle in the Hausdorff topology, which requires further research.[37]

It must be clear at this point that the measure-theoretic tools are restricted by their inability to capture anything that is lower-dimensional than the dimension of the space—e.g., contact over surface patches, curve segments, or isolated points. Although this was leveraged to liberate us from the daunting task of explicit regularization, it significantly restricts the method in formulating geometric operations that involve heterogeneous dimensions. Next section is dedicated to introduce new tools to mitigate such limitations.

[37]This is only an issue when one desires to go back to the explicit description of the $\mathsf{C}-$obstacle. If one can remain in the implicit domain for subsequent downstream operations—which is the ultimate goal advocated in this thesis—the approximation in (2.3.60) is likely to be sufficient.

2.3.2 Delta Singularities

As explained in Section 2.2.1, a number of important problems in spatial reasoning require working with shapes and motions with heterogeneous dimensions. A prime example is the fundamental task of computing sweep and unsweep of a solid shape, which is a 3D regular set in $\mathsf{E} = \mathbb{R}^3$, along a one-parametric motion, which is a 1D trajectory in $\mathsf{C} = \mathrm{SE}(3)$, requiring the computation of morphological operations (namely, dilation and erosion) in the 6D configuration space without regularization. However, the state-of-the-art proposals for the analytic formulation of such morphological operations in terms of convolutions [215, 218] are applicable to shapes and motions that can be embedded as 6D regular sets in $\mathsf{C} = \mathrm{SE}(3)$. This is due to the fact that convolutions in particular, and measure-theoretic integrals in general, implicitly regularize the outcomes simply because lower-dimensional features cannot contribute to the volumetric integration.

In order to tackle this problem, one needs a device to deliberately make the lower-dimensional features measurable. I propose to model the defining functions of embedded lower-dimensional sets in terms of the Dirac $\delta-$function, originally introduced by the 20th century theoretical physicist Paul Dirac [99] to formulate *singularities* in quantum mechanics and later applied extensively to the solution of ordinary and partial differential equations, control theory, fluid mechanics, statistics, etc.

2.3.2.1 Dirac $\delta-$Function for Dealing with Heterogeneous Dimensions:

The Dirac $\delta-$function is an "improper" function as the limit of a sequence of ordinary (i.e., "proper") functions—whose use alongside ordinary functions in mathematical analysis "must be confined to certain simple types of expression for which it is obvious

that no inconsistency can arise" [99]. Since its introduction by Paul Dirac[38] to quantum mechanics, it has found applications in may other areas and has been redefined in a variety of different ways in efforts to rigorously formalize its calculus, including definitions via the theory of distributions using Schwartz and Sobolev's concepts of 'generalized' function/calculus [306], and more recently using Cortizo's concepts of 'virtual' functions/calculus [81], both of which extend the classical notions of functions and integration. Here I resist the temptation to get into such details and restrict ourselves to a more simplistic definition using a set of axioms that are the extensions to Dirac's original definition and, if manipulated consistently, provide a sufficiently complete definition for our purposes.

Definition 2.3.7. (Dirac $\delta-$Function) The $3^{\text{rd}}-$order $\delta-$function over the d$-$space denoted by $\delta^3 : \mathsf{E} \to \{0, \text{“}\infty^{\text{d}}\text{”}\}$ is axiomatically defined as

$$\delta^3(\mathbf{x}) := \begin{cases} \text{“}\infty^3\text{”} & \text{if } \mathbf{x} = \mathbf{0}, \\ 0 & \text{if } \mathbf{x} \neq \mathbf{0}. \end{cases} \quad \text{such that} \quad \int_S \delta^3(\mathbf{x})\, d\mu^3[\mathbf{x}] = \begin{cases} 1 & \text{if } \mathbf{0} \in \dot{\imath}\,(S), \\ 0 & \text{if } \mathbf{0} \in \dot{e}(S), \end{cases} \tag{2.3.61}$$

for every domain $S \in \mathscr{P}(\mathsf{E})$ whose topological interior and exterior (i.e., interior of complement) are denoted by $\dot{\imath}\,(S), \dot{e}(S) \in \mathscr{P}(\mathsf{E})$, respectively. This is an incomplete definition in the sense that the value of the integral is left ambiguous if the domain's boundary passes through the origin (i.e., $\mathbf{0} \in \dot{b}(S)$). One way of resolving this ambiguity is to define the $\delta-$function as the 'limit representation' of a proper integrable real-valued function (sometimes called a 'nascent' $\delta-$function) $g^3_\sigma \in L^1(\mathsf{E})$ as $\sigma \to 0^+$,

[38] Paul Dirac was an English theoretical physicist and mathematician, and a professor of Mathematics at the University of Cambridge. He has made significant contributions to the field of quantum mechanics and shared the Nobel Prize in physics in 1933 with Erwin Schrödinger for his contributions to new productive forms of 'atomic theory'.

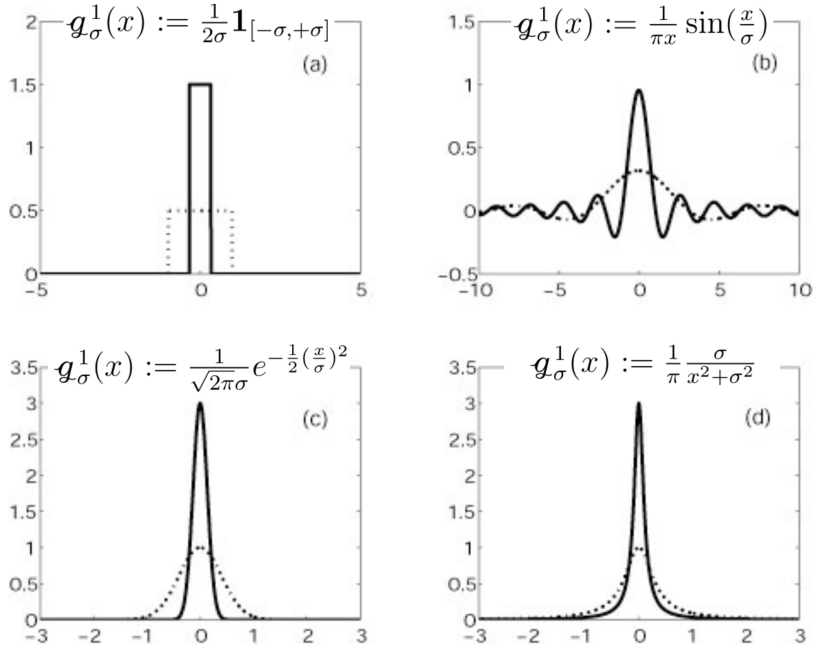

Figure adopted from dmpeli.math.mcmaster.ca

FIGURE 2.3.6: A number of limit representations for $\delta^1(x)$.

TABLE 2.3.1: A number of limit representations for $\delta^1(x)$.

$g^1_\sigma(x) := \frac{1}{\pi}\frac{\sigma}{x^2+\sigma^2}$	$g^1_\sigma(x) := \frac{1}{2\sigma}\mathbf{1}_{[-\sigma,+\sigma]}$	$g^1_\sigma(x) := \frac{1}{2}\mathbf{1}_{[-1,+1]}\lvert x\rvert^{\sigma-1}$
$g^1_\sigma(x) := \frac{1}{2\sigma}e^{-\lvert x\rvert/\sigma}$	$g^1_\sigma(x) := \frac{1}{\pi x}\sin(\frac{x}{\sigma})$	$g^1_\sigma(x) := \log\left[\coth(\frac{x}{\sigma})\right]$
$g^1_\sigma(x) := \frac{1}{2\sigma}\mathrm{sech}^2(\frac{x}{\sigma})$	$g^1_\sigma(x) := \frac{1}{\sigma}\mathrm{csch}(\frac{x}{\sigma})$	$g^1_\sigma(x) := \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{1}{2}(\frac{x}{\sigma})^2}$

i.e., $g_\sigma : \mathsf{E} \to \mathbb{R}$ such that the $\delta-$properties in (2.3.61) hold in the limit:

$$\lim_{\sigma\to 0^+} g^{\mathrm{d}}_\sigma(\mathbf{x}) = \begin{cases} \infty & \text{if } \mathbf{x} = \mathbf{0}, \\ 0 & \text{if } \mathbf{x} \neq \mathbf{0}. \end{cases} \quad \text{and} \quad \lim_{\sigma\to 0^+}\int_S g^3(\mathbf{x})\, d\mu^3[\mathbf{x}] = \begin{cases} 1 & \text{if } \mathbf{0} \in \imath\,(S), \\ 0 & \text{if } \mathbf{0} \in \epsilon(S). \end{cases} \tag{2.3.62}$$

A few examples of one-dimensional nascent $\delta-$functions are given in Table 2.3.1. For the moment, I will use the axiomatic definition in (2.3.61) and present the propositions without appealing to a particular $g_\sigma-$function. Nevertheless, the following assumption is necessary:

Assumption 2.3.8. Integrating a $\delta-$function over a domain $S \in \mathscr{P}(\mathsf{E})$ whose boundary passes through the origin yields an unknown value in the range $[0,1]$ which is a function of the domain's shape $\lambda : \mathscr{P}(\mathsf{E}) \to [0,1]$, where

$$\lambda(S) := \langle \delta^3, \mathbf{1}_S \rangle = \int_S \delta^3(\mathbf{x}) \; d\mu^3[\mathbf{x}] = \lim_{\sigma \to 0^+} \int_S \mathscr{g}^3(\mathbf{x}) \; d\mu^3[\mathbf{x}]. \tag{2.3.63}$$

The precise value of $0 \leq \lambda(S) \leq 1$ will depend on the chose of the limit representation $\mathscr{g}^3_\sigma : \mathsf{E} \to \mathbb{R}$ and local differential geometry of the domain. In particular,

- If $\mathbf{0} \in \mathscr{i}\,(S)$ then $\lambda = 1$ as a direct result of the definition in (2.3.61), but the converse is not necessarily true, i.e., if $\lambda = 1$ then $\mathbf{0} \in \mathscr{i}\,(S)$ or $\mathbf{0} \in \mathscr{b}(S)$:

$$\mathbf{0} \in \mathscr{i}\,(S) \;\Rightarrow\; \lambda = 1 \quad \text{but} \quad \lambda = 1 \;\Rightarrow\; \mathbf{0} \in (\mathscr{i}\,(S) \cup \mathscr{b}(S)) = \mathscr{h}(S). \tag{2.3.64}$$

- If $\mathbf{0} \in \mathscr{e}(S)$ then $\lambda = 0$ as a direct result of the definition in (2.3.61), but the converse is not necessarily true, i.e., if $\lambda = 0$ then $\mathbf{0} \in \mathscr{e}(S)$ or $\mathbf{0} \in \mathscr{b}(S)$:

$$\mathbf{0} \in \mathscr{e}(S) \;\Rightarrow\; \lambda = 0 \quad \text{but} \quad \lambda = 0 \;\Rightarrow\; \mathbf{0} \in (\mathscr{e}(S) \cup \mathscr{b}(S)) = \mathscr{h}(S). \tag{2.3.65}$$

- If $\mathbf{0} \in \mathscr{b}(S)$ then $\lambda \in [0,1]$ (note: closed interval) as a result of the above assumption. Obviously, $\lambda \in [0,1]$ is not conclusive at all; however, $\lambda \in (0,1)$ can only happen if $\mathbf{0} \in \mathscr{b}(S)$.

In other words, the definition given in (2.3.61) requires an additional clause:

$$\langle \delta^3, \mathbf{1}_S \rangle = \int_S \delta^3(\mathbf{x})\ d\mu^3[\mathbf{x}] = \begin{cases} 1 & \text{if } \mathbf{0} \in \dot{\iota}\,(S), \\ \lambda & \text{if } \mathbf{0} \in \partial(S), \quad \text{for some } \lambda \in [0,1], \\ 0 & \text{if } \mathbf{0} \in \dot{e}(S). \end{cases} \tag{2.3.66}$$

The importance of the above assumption is that it allows us to use the $\delta-$function as a means to detect if a boundary point lies on a lower-dimensional segment that would be eliminated upon regularization. I shall use this in Section 2.3.3 to define pointset-topological superlevel set regularization of the defining functions.

Lemma 2.3.9. (Volume Fraction Lemma) *If the* $\delta-$*function is assumed to be a radial, i.e., rotationally symmetric* $\delta^3 = \delta^3 \circ \mathsf{act}^{-1}(\mathbb{r})$ *for all* $\mathbb{r} \in \mathrm{SO}(3)$*, the integral in* (2.3.66) *returns the fraction of the volume of an infinitesimal* $3-$*ball neighborhood around the origin occupied by the domain, i.e.,*

$$\lambda(S) = \lim_{\epsilon \to 0^+} \frac{\mu^3[S \cap B_\epsilon(\mathbf{0})]}{\mu^3[B_\epsilon(\mathbf{0})]}, \quad \text{where } B_r(\mathbf{x}) = \big\{ \mathbf{x}' \in \mathsf{E} \mid \|\mathbf{x} - \mathbf{x}'\|_2 \leq r \big\}. \tag{2.3.67}$$

Proof. The proof follows from the radial symmetry of the integral around the origin, also noting that the integral's value is completely determined by the $\epsilon-$neighborhood of the origin, as it vanishes everywhere else due to (2.3.66). □

The beauty of the analytic approach is that such a computation will never be necessary, and it will carry itself through implicitly as part of the chosen $g-$kernel.[39]

The $\delta-$function can be thought of as the defining function of a 'lumped' material of

[39] I shall return to the practical aspects of relaxing the introduced $\delta-$singularities with $\sigma \to 0^+$ for theoretical abstraction to the so-called 'geometric densities' and letting them 'dissipate' in space by using finite $\sigma > 0$ for practical computation in Chapter 5.

unit 3−measure concentrated at the origin $\mathbf{0} \in \mathsf{E}$, making a singular point *artificially* Lebesgue 3−measurable. Note that $\delta^{\mathrm{d}}(\mathbf{x}) = \sum_{0 \leq j < \mathrm{d}} \delta^1(x_j)$ if $\mathbf{x} := \sum_{0 \leq j < \mathrm{d}} x_j \mathbf{e}_j$ where $x_j := (\mathbf{x} \cdot \mathbf{e}_j) \in \mathbb{R}$ are orthogonal coordinates for $0 \leq j < \mathrm{d} \leq 3$—implying a natural extension of the $\mathbb{R}_+$−algebra to the nominal infinities—e.g., $\infty^2 = \infty^1 \cdot \infty^1$, $\infty^3 = \infty^2 \cdot \infty^1$, etc. Note also that by definition, $\delta^0(\mathbf{x}) = \mathbf{1}_{\{\mathbf{0}\}}(\mathbf{x})$ which returns $\infty^0 := 1$ if $\mathbf{x} = \mathbf{0}$ and zero otherwise.[40]

2.3.2.2 Modeling Artificially Measurable Geometric Singularities:

To model singularities of different orders embedded in $\mathsf{E} = \mathbb{R}^3$ around a given point $\mathbf{x}_0 \in \mathsf{E}$, one needs a local parameterization $\Gamma = (\gamma_1, \gamma_2, \gamma_3) : (-\epsilon, +\epsilon)^3 \to \mathsf{E}$ for which $\mathbf{x}_0 = \Gamma(\mathbf{0})$ which is a C^∞−diffeomorphism with nonzero Jacobian determinant. To simplify the development and presentation, I make the following stronger assumption that is in line with Assumption 2.2.6 where I assumed the shapes to be semianalytic subsets of the 3−space:

Assumption 2.3.10. Each $0-, 1-, 2-$, or $3-$dimensioal analytic stratum in a given shape $S \subseteq \mathsf{E}$ (in accordance with Assumption 2.2.6 of Section 2.2.2) is an embedding of a cell (i.e., point, interval, polygon, or polyhedron) via a nowhere-singular analytic automorphism $\Gamma : \mathsf{E} \to \mathsf{E}$—i.e., a conformal (i.e., angle-preserving) and homeomorphic (i.e., continuous and invertible) map that tamely embeds the entire E onto itself. Informally, this ensures the existence of a single *orthogonal* curvilinear coordinate chart per stratum, whose origin, one of 3 axes, one of 3 planes, or entire grid is aligned with the carrier geometry of the stratum.

[40] One must be very careful when working with these new elements to prevent inconsistencies further down the line—"never fool around with infinities." –Javad Feiz

Convention 2.3.11. Thus every shape $S \in \mathscr{P}(\mathsf{E})$ can be decomposed into four groups of strata as $S := \bigcup_{0\leq \mathrm{d}<4} S^{\mathrm{d}}$. For $0 \leq \mathrm{d} < 4$, each S^{d} is a d−measurable semianalytic set. Thus each shape is a finite collection of n_{d} distinct open d−dimensional C^{ω}−manifolds $S^{\mathrm{d}} = \bigcup_{0\leq i<n_{\mathrm{d}}} \Gamma_i(\Delta^{\mathrm{d}})$ indexed via $0 \leq i < n_{\mathrm{d}}$, with each analytic submanifold being parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^{\mathrm{d}} \subseteq \mathbb{R}^3$ is an open d−cell.

Parameterized Analytic Strata of Heterogeneous Dimensions

1. $S^0 \subseteq S$ is a 0−measurable semianalytic set, i.e., a finite collection of n_0 distinct points $S^0 = \{\mathbf{x}_i\}_{0\leq i<n_0}$ indexed via $0 \leq i < n_0$ corresponding to an indexed collection of embeddings (i.e., instantiations) via $\Gamma_i : \mathsf{E} \to \mathsf{E}$ of the origin of the parameter space $\mathbf{0} \in \mathbb{R}^3$ to the shape space, i.e., $\mathbf{x}_i = \Gamma_i(\mathbf{0})$ for $0 \leq i < n_0$.
2. $S^1 \subseteq S$ is a 1−measurable semianalytic set, i.e., a finite collection of n_1 distinct open C^{ω}−curve segments $S^1 = \bigcup_{0\leq i<n_1} \Gamma_i(\Delta^1)$ indexed via $0 \leq i < n_1$, each being parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^1 \subseteq \mathbb{R}$ is an open 1−cell (i.e., interval).
3. $S^2 \subseteq S$ is a 2−measurable semianalytic set, i.e., a finite collection of n_2 distinct open C^{ω}−surface patches $S^2 = \bigcup_{0\leq i<n_2} \Gamma_i(\Delta^2)$ indexed via $0 \leq i < n_2$, each being parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^2 \subseteq \mathbb{R}^2$ is an open 2−cell (i.e., polygon).
4. $S^3 \subseteq S$ is a 3−measurable semianalytic set, i.e., a finite collection of n_3 distinct open C^{ω}−solid volumes $S^3 = \bigcup_{0\leq i<n_3} \Gamma_i(\Delta^3)$ indexed via $0 \leq i < n_3$, each being parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^3 \subseteq \mathbb{R}^3$ is an open 3−cell (i.e., polyhedron).

Figure 2.3.7 illustrates the stratification for a semianalanytic planar set.

Convention 2.3.12. For a given embedding $\Gamma : \mathsf{E} \to \mathsf{E}$ that satisfies the conditions in Assumption 2.3.10, let $\gamma_j := (\Gamma \cdot \mathbf{e}_j) : \mathbb{R}^3 \to \mathbb{R}$ be the individual transformations that map the conforming curvilinear coordinates $\mathbf{t} \in \mathbb{R}^3$ to the (fixed) world Cartesian coordinates $\mathbf{x} \in \mathsf{E}$ as

$$(x_1, x_2, x_3) = \Gamma(t_1, t_2, t_3) = \Big(\gamma_1(t_1, t_2, t_3), \gamma_2(t_1, t_2, t_3), \gamma_3(t_1, t_2, t_3)\Big). \tag{2.3.68}$$

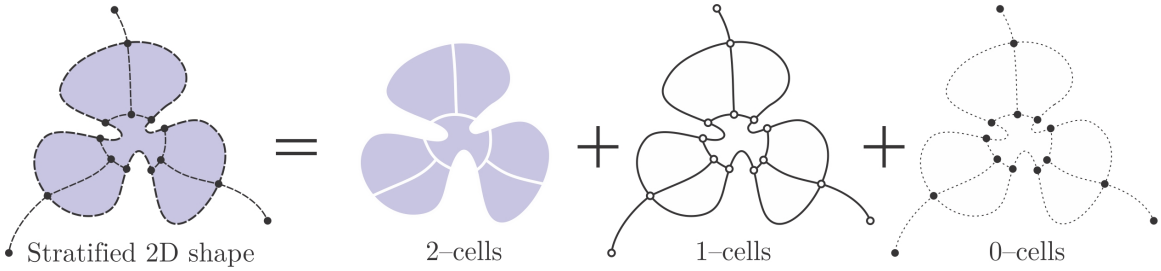

FIGURE 2.3.7: Every semianalytic subset of the d−space (here, d := 2) can be stratified into $0-, 1-, \cdots,$ d−simensional strata. See also Fig. 2.3.12.

where $\mathbf{t} = (t_1, t_2, t_3) = \sum_{1 \le j \le 3} t_j \mathbf{e}_j$ and $\mathbf{x} = (x_1, x_2, x_3) = \sum_{1 \le j \le 3} x_j \mathbf{e}_j$. The Jacobian determinant is defined and denoted by

$$|\mathfrak{J}(t_1, t_2, t_3)| = \begin{vmatrix} \frac{\partial \gamma_1}{\partial t_1}(t_1, t_2, t_3) & \frac{\partial \gamma_1}{\partial t_2}(t_1, t_2, t_3) & \frac{\partial \gamma_1}{\partial t_3}(t_1, t_2, t_3) \\ \frac{\partial \gamma_2}{\partial t_1}(t_1, t_2, t_3) & \frac{\partial \gamma_2}{\partial t_2}(t_1, t_2, t_3) & \frac{\partial \gamma_2}{\partial t_3}(t_1, t_2, t_3) \\ \frac{\partial \gamma_3}{\partial t_1}(t_1, t_2, t_3) & \frac{\partial \gamma_3}{\partial t_2}(t_1, t_2, t_3) & \frac{\partial \gamma_3}{\partial t_3}(t_1, t_2, t_3) \end{vmatrix}. \tag{2.3.69}$$

The 1^{st} and 2^{nd} minors of the Jacobian will be useful when assigning 1−singualrities and 2−singualrities to parameterized surface patches and curve segments, respectively, obtained via applying the restricted mappings

- $\Gamma|_{\mathbb{R}^2} = (\gamma_1, \gamma_2) : \mathbb{R}^2 \to \mathbb{R}^3$ to a 2−cell $\Delta^2 \subset \mathbb{R}^2$ to obtain a curvilinear surface patch $\Gamma(\Delta^2) \subset \mathsf{E}$, parameterized via $(t_1, t_2) \in \mathbb{R}^2$ for $t_3 := 0$, whose normal vectors are given by $\frac{\partial \gamma_3}{\partial t_3}\big|_{t_3=0}$ over the surface patch; and

- $\Gamma|_{\mathbb{R}} = \gamma_1 : \mathbb{R} \to \mathbb{R}^3$ to an 1−cell $\Delta^1 \subset \mathbb{R}$ to obtain a curvilinear curve segment $\Gamma(\Delta^1) \subset \mathsf{E}$, parameterized via $t_1 \in \mathbb{R}$ for $(t_2, t_3) := (0, 0)$, whose normal planes are given by $\frac{\partial \gamma_1}{\partial t_1}\big|_{t_1=0}$ and $\frac{\partial \gamma_2}{\partial t_2}\big|_{t_2=0}$ along the curve segment.

These minors are defined and denoted by $|\mathfrak{J}'(\mathbf{t})|$ and $|\mathfrak{J}''(\mathbf{t})|$, respectively:

$$|\mathfrak{J}'(t_1,t_2,0)| = \begin{vmatrix} \frac{\partial\gamma_1}{\partial t_1}(t_1,t_2,0) & \frac{\partial\gamma_1}{\partial t_2}(t_1,t_2,0) \\ \frac{\partial\gamma_2}{\partial t_1}(t_1,t_2,0) & \frac{\partial\gamma_2}{\partial t_2}(t_1,t_2,0) \end{vmatrix}, \quad \text{and} \quad |\mathfrak{J}''(t_1,t_2,t_3)| = \left|\frac{\partial\gamma_1}{\partial t_1}(t_1,0,0)\right|. \tag{2.3.70}$$

Note that even though all three parameters $t_1, t_2, t_3 \in \mathbb{R}$ appear on the left-hand side of the above two definitions—chosen as such to enable short notations $|\mathfrak{J}'(\mathbf{t})|$ and $|\mathfrak{J}''(\mathbf{t})|$—the right-hand side implies that the former is independent of t_3 and the latter does not depend on t_2 and t_3 (i.e., substitutes them with 0 anyway).

Therefore, for a given semianalytic set and an a priori knowledge of its stratification, it can be decomposed into four groups of $0-$, $1-$, $2-$, and $3-$ dimensional strata as $S := \sum_{0\leq \mathrm{d}<4} S^{\mathrm{d}} = S^0 \cup S^1 \cup S^2 \cup S^3$ illustrated via

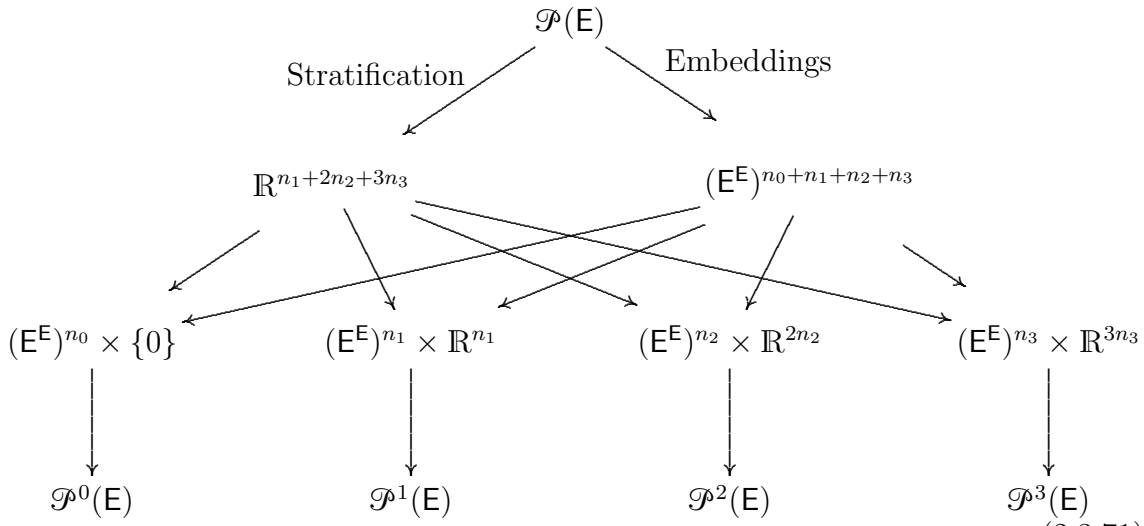

(2.3.71)

in which $\mathscr{P}^{\mathrm{d}}(\mathsf{E}) \subset \mathscr{P}(\mathsf{E})$ for $0 \leq \mathrm{d} < 4$ denote the collections of all shapes that can be decomposed into a finite number of disjoint open $\mathrm{d}-$manifolds.

It is desirable to define embedded $1^{\text{st}}-$, $2^{\text{nd}}-$, and $3^{\text{rd}}-$order $\delta-$functions over $\mathsf{E} = \mathbb{R}^3$ at $\mathbf{x} := \Gamma(\mathbf{t}) \in \mathsf{E}$ with respect to the geometry-aligned curvilinear coordinate system to define singularities over lower-dimensional strata and make them *artificially* $3-$measurable whenever desired—e.g., to allow them to participate in inner products or convolutions between defining functions and other volumetric integrals, rather than rendering them irrelevant up to equality$-\mathsf{ae}$.

In the following, I demonstrate how artificial volumes can be assigned to collections of $\mathrm{d}-$ elements for $0 \leq \mathrm{d} < 4$ when explicit parametrizations are known a priori. I shall come back to the question of how to do it when that is not the case (e.g., starting from implicit descriptions) in Chapter 3.

Discrete Pointsets. Take for example the finite pointset $S^0 = \{\mathbf{x}_i\}_{0 \leq i < n_0} \subset \mathsf{E}$ with a cardinality (i.e., Hausdorff $0-$measure) $|S^0| = \mu^0[S^0] = n_0$, whose defining function $f_{S^0} : \mathsf{E} \to \mathbb{R}_+$ can be constructed as

$$f_{S^0}(\mathbf{x}) := \sum_{0 \leq i < n_0} c_i \delta^0(\mathbf{x} - \mathbf{x}_i), \quad \text{where } c_i \in \mathbb{R}_+ \text{ for } 0 \leq i < n_0. \tag{2.3.72}$$

The 'weight coefficients' $c_i > 0$ are arbitrary (for now) as a result of the profile-oblivious abstraction, and $\delta^0(\mathbf{x} - \mathbf{x}_i) = \mathbf{1}_{\{\mathbf{0}\}}(\mathbf{x} - \mathbf{x}_i) = \mathbf{1}_{\{\mathbf{x}_i\}}(\mathbf{x})$, thus $f_{S^0} \equiv \mathbf{1}_{S^0}$ with strict equality when all $c_i = 1$. Obviously, the $3-$measure (i.e., volume) of the set is zero. Now, let us replace the above real-valued function $f_{S^0} : \mathsf{E} \to \mathbb{R}_+$ with an "improper" 'extended'-real-valued function $f^3_{S^0} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^3\text{”}\}$ as

$$f^3_{S^0}(\mathbf{x}) := \sum_{0 \leq i < n_0} c_i \delta^3(\mathbf{x} - \mathbf{x}_i), \quad \text{where } c_i \in \mathbb{R}_+ \text{ for } 0 \leq i < n_0, \tag{2.3.73}$$

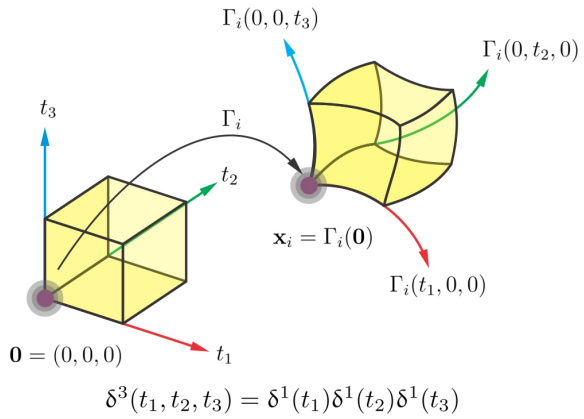

FIGURE 2.3.8: A $3^{\text{rd}}-$order $\delta-$singular isolated point is obtained by mapping the origin of the parameter space, at which the $3^{\text{rd}}-$order singularity $\delta^3(t_1,t_2,t_3)$ is concentrated.

conceptualized as putting a 'lumped' volume (i.e., $3^{\text{rd}}-$order volume singularity) of intensity $c_i > 0$ concentrated at each point. Consequently, the artificial $3-$measure of the set becomes nonzero:

$$\int_{\mathsf{E}} f^3_{S^0}(\mathbf{x})\; d\mu^3[\mathbf{x}] = \sum_{0\le i<n_0} c_i \int_{\mathsf{E}} \delta^3(\mathbf{x}-\mathbf{x}_i)\; d\mu^3[\mathbf{x}] = \sum_{0\le i<n_0} c_i. \tag{2.3.74}$$

In general, one can choose to have both ordinary points and lumped volume points in a collection $S^0 = \{\mathbf{x}_i\}_{0\le i<n_0}$. But of course adding ordinary points does not affect the set up to equality$-$ae.

For the special case with $c_i := 1$ the singular defining function in (2.3.73) turns into a $3^{\text{rd}}-$order $\delta-$singular characteristic function $\mathbb{1}^3_{S^0} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^3\text{”}\}$ for which the L^1-norm gives the precise point count $\|\mathbb{1}^3_{S^0}\|_1 = \sum_{0\le i<n_0} 1 = \mu^0[S^0] = |S^0|$.

Before presenting the extension of the above ideas to higher-dimensional analytic submanifold strata, note that each point $\mathbf{x}_i \in S_0$ can be viewed as the origin of a curvilinear coordinate system, defined via a mapping $\Gamma_i : \mathsf{E} \to \mathsf{E}$. Adopting Convention

2.3.12 with $\Gamma_i : (\gamma_{i,1}, \gamma_{i,2}, \gamma_{i,3}) : \mathbb{R}^3 \to \mathbb{R}$, for every $0 \leq i < n_0$ one obtains

$$(x_1, x_2, x_3) := \Big(\gamma_{i,1}(0,0,0), \gamma_{i,2}(0,0,0), \gamma_{i,2}(0,0,0)\Big), \tag{2.3.75}$$

i.e., $\mathbf{x}_i = \Gamma_i(\mathbf{0}) \in \mathsf{E}$. Thus (2.3.73) can be rewritten in parametric terms as

$$f^3_{S^0}(\mathbf{x}) := \sum_{0\leq i<n_0} c_i \ |\mathfrak{J}_i(0,0,0)|^{-1} \ \delta^3(t_1,t_2,t_3) \ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.76}$$

The Jacobian of the parameterization is included in (2.3.76) to account for the effects of the non-isometric mapping on the artificial $3-$measure of the set:

$$\int_{\mathsf{E}} f^3_{S^0}(\mathbf{x}) \ d\mu^3[\mathbf{x}] = \iiint_{\mathbb{R}^3} (f^3_{S^0} \circ \Gamma_i)(t_1,t_2,t_3) \ |\mathfrak{J}_i(\mathbf{t})| \ dt_1dt_2dt_3 \tag{2.3.77}$$

$$= \sum_{0\leq i<n_0} \iiint_{\mathbb{R}^3} c_i \frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(0,\ 0,\ 0)|} \ \delta^3(t_1,t_2,t_3) \ dt_1dt_2dt_3. \tag{2.3.78}$$

Using Lemma 2.3.16 with d := 3 and the function $(t_1,t_2,t_3) \mapsto \rho_i(t_1) \ \frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(0,\ 0,\ 0)|}$ yields

$$\int_{\mathsf{E}} f^3_{S^0}(\mathbf{x}) \ d\mu^3[\mathbf{x}] = \sum_{0\leq i<n_0} \int_{\mathbb{R}} c_i \ \frac{|\mathfrak{J}_i(0,0,0)|}{|\mathfrak{J}_i(0,0,0)|} \ dt_1 = \sum_{0\leq i<n_1} c_i. \tag{2.3.79}$$

which is the expected result as obtained in (2.3.74). Figure 2.3.8 illustrates the idea of constructing $3^{\text{rd}}-$order $\delta-$singular isolated points via parameterization.

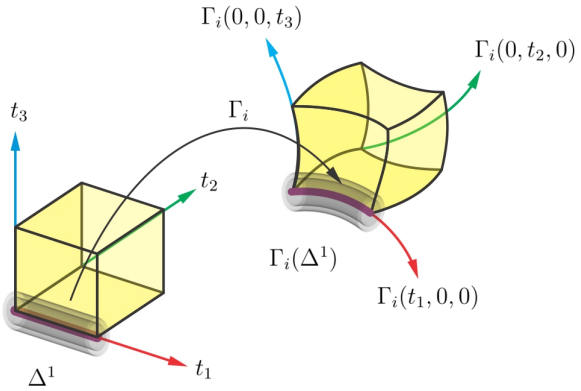

$$\mathbf{1}_{\Delta^1}(t_1)\delta^2(t_2,t_3) = \mathbf{1}_{\Delta^1}(t_1)\delta^1(t_2)\delta^1(t_3)$$

FIGURE 2.3.9: A $2^{\text{nd}}-$order $\delta-$singular curve segment is obtained by mapping a parameterized $1-$cell $\Delta^1 \subset \mathbb{R}$, at which the $2^{\text{rd}}-$order singularity $\delta^2(t_2,t_3)$ is concentrated.

Curve Segments. A Hausdorff $1-$measurable union of disjoint $C^\omega-$curve segments $S^1 = \bigcup_{0\le i<n_1} \Gamma_i(\Delta^1)$—with $\mu^1[S^1]$ equal to the total curve length of the n_1 pieces—can be made into an artificially $3-$measurable set. Adopting Convention 2.3.12 with $\Gamma_i : (\gamma_{i,1},\gamma_{i,2},\gamma_{i,3}) : \mathbb{R}^3 \to \mathbb{R}$, for every $0 \le i < n_1$ one obtains

$$(x_1,x_2,x_3) := \Big(\gamma_{i,1}(t_1,0,0),\gamma_{i,2}(t_1,0,0),\gamma_{i,2}(t_1,0,0)\Big), \quad t_1 \in \Delta^1, \tag{2.3.80}$$

which parameterizes the curve segment spanned by $\mathbf{x} = \Gamma_i(\mathbf{t}) \in \mathsf{E}$ via the $1-$cell (i.e., interval) $\Delta^1 \subset \mathbb{R}$, spanned by $\mathbf{t} = \Gamma_i^{-1}(\mathbf{t}) \in \mathbb{R}\times\{(0,0)\}$ for $(t_2,t_3) := (0,0)$. The original (i.e., unextended) defining function $f_{S^1} : \mathsf{E} \to \mathbb{R}_+$ can thus be constructed as

$$f_{S^1}(\mathbf{x}) := \sum_{0\le i<n_1} \rho_i'(t_1)\ \delta^0(t_2,t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.81}$$

where $\delta^0(t_2,t_3)$ is the indicator function of the origin $(t_2,t_3) := (0,0)$ of the 2D parametric t_2t_3-plane that is normal to the $2-$singular $1-$cell $\Delta^1 \subset \mathbb{R}$.

The (presumably nonzero) Jacobian determinant $|\mathfrak{J}_i(\mathbf{t})|$ can be included to nor-

malize the metrics after coordinate transformation,[41] noting, however, that it does not affect the defining function up to equivalence via $\equiv$. Thus letting the density function be $\rho_i(t_1) := \rho_i'(t_1)|\mathfrak{J}_i(t_1,0,0)|$,

$$f_{S^1}(\mathbf{x}) := \sum_{0\le i<n_1} \rho_i(t_1)\ |\mathfrak{J}_i(t_1,0,0)|^{-1}\ \delta^0(t_2,t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.82}$$

The functions $\rho_i : \mathbb{R} \to \mathbb{R}_+$ are called 'curve density' distributions that are nonzero along the parameterization $1-$cell $\Delta^1 \subset \mathbb{R}$ and zero elsewhere, i.e., $\rho_i \equiv \mathbf{1}_{\Delta^1}$, where $\mathbf{1}_{\Delta^1} : \mathbb{R} \to \{0,1\}$ is the cell's indicator function, thus the basis function $\rho_i(t_1)\ |\mathfrak{J}_i(t_1,0,0)|^{-1}\ \delta^0(t_2,t_3)$ is a defining function of each instance of embedding the $1-$cell to individual curve segments. For unit constant density $\rho_i(t_1) := 1$, one obtains the indicator function of the curve segments:

$$\mathbf{1}_{S^1}(\mathbf{x}) := \sum_{0\le i<n_1} |\mathfrak{J}_i(t_1,0,0)|^{-1}\ \delta^0(t_2,t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.83}$$

Next, let us replace the above real-valued function $f_{S^1} : \mathsf{E} \to \mathbb{R}_+$ with an "improper" 'extended'-real-valued function $f^2_{S^1} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^2\text{”}\}$ as

$$f^2_{S^1}(\mathbf{x}) := \sum_{0\le i<n_1} \rho_i(t_1)\ |\mathfrak{J}_i(t_1,0,0)|^{-1}\ \delta^2(t_2,t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.84}$$

conceptualized as putting a $2^{\text{nd}}-$order volume singularity of intensity $\rho_i(t_1) > 0$ distributed along each curve segment. Consequently, the artificial $3-$measure of the

[41] The Jacobian is included to ensure that $\rho_i : \mathbb{R} \to \mathsf{E}$ represents distribution per unit physical length along the curve $\Gamma_i(\Delta^1) \subset \mathsf{E}$ rather than per unit parametric increment along the $1-$cell $\Delta^1 \subset \mathbb{R}$ regardless of the (possibly non-isometric) parameterization.

set, obtained by taking the L^1-norm of (2.3.84), becomes nonzero:

$$\int_{\mathsf{E}} f^2_{S^1}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \iiint_{\mathbb{R}^3} (f^2_{S^1} \circ \Gamma_i)(t_1,t_2,t_3)\ |\mathfrak{J}_i(\mathbf{t})|\ dt_1 dt_2 dt_3 \tag{2.3.85}$$

$$= \sum_{0\le i<n_1} \iiint_{\mathbb{R}^3} \rho_i(t_1)\ \frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(t_1,\ 0,\ 0)|}\ \delta^2(t_2,t_3)\ dt_1 dt_2 dt_3. \tag{2.3.86}$$

Using Lemma 2.3.16 with $\mathrm{d} := 2$ and the function $(t_2,t_3) \mapsto \rho_i(t_1)\ \frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(t_1,\ 0,\ 0)|}$ yields

$$\int_{\mathsf{E}} f^2_{S^1}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \sum_{0\le i<n_1} \int_{\mathbb{R}} \rho_i(t_1)\ \frac{|\mathfrak{J}_i(t_1,0,0)|}{|\mathfrak{J}_i(t_1,0,0)|}\ dt_1. \tag{2.3.87}$$

Letting $c_i := \|\rho_i\|_1$ be the total artificial volume of each curve segment,

$$\|f^2_{S^1}\|_1 = \sum_{0\le i<n_1} \int_{\mathbb{R}} \rho_i(t_1)\ dt_1 = \sum_{0\le i<n_1} \|\rho_i\|_1 = \sum_{0\le i<n_1} c_i. \tag{2.3.88}$$

For the special case with $\rho_i(t_1) := 1$ the singular defining function in (2.3.84) turns into a $2^{\text{nd}}-$order $\delta-$singular characteristic function $\mathbb{1}^2_{S^1} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^2\text{”}\}$ as

$$\mathbb{1}^2_{S^1}(\mathbf{x}) := \sum_{0\le i<n_1} |\mathfrak{J}_i(t_1,0,0)|^{-1}\ \delta^2(t_2,t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.89}$$

for which the L^1-norm gives the precise total curve length:

$$\|\mathbb{1}^2_{S^1}\|_1 = \int_{\mathsf{E}} \mathbb{1}^2_{S^1}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \sum_{0\le i<n_1} \mu^1[\Gamma_i(\Delta^1)] = \mu^1[S^1]. \tag{2.3.90}$$

Once again, one can choose to have both ordinary- and singular-volume curve segments in a collection $S^1 = \bigcup_{0\le i<n_1} \Gamma_i(\Delta^1)$, while the former does not affect the implicit set being described up to equality$-\mathsf{ae}$. See Fig. 2.3.9 for an illustration.

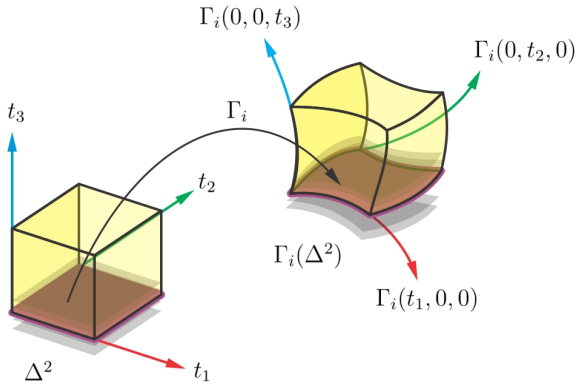

$$\mathbf{1}_{\Delta^2}(t_1,t_2)\delta^1(t_3) = \mathbf{1}_{\Delta^1}(t_1)\mathbf{1}_{\Delta^1}(t_2)\delta^1(t_3)$$

FIGURE 2.3.10: A $1^{\text{st}}-$order $\delta-$singular surface patch is obtained by mapping a parameterized 2−cell $\Delta^2 \subset \mathbb{R}$, at which the $1^{\text{st}}-$order singularity $\delta^1(t_3)$ is concentrated.

Surface Patches. A Hausdorff 2−measurable union of disjoint $C^\omega-$surface patches $S^2 = \bigcup_{0\le i<n_2} \Gamma_i(\Delta^2)$—with $\mu^2[S^2]$ equal to the total surface area of the n_2 pieces—can be made into an artificially 3−measurable set. Adopting Convention 2.3.12 with $\Gamma_i : (\gamma_{i,1}, \gamma_{i,2}, \gamma_{i,3}) : \mathbb{R}^3 \to \mathbb{R}$, for every $0 \le i < n_2$ one obtains

$$(x_1, x_2, x_3) := \Big(\gamma_{i,1}(t_1,t_2,0), \gamma_{i,2}(t_1,t_2,0), \gamma_{i,2}(t_1,t_2,0)\Big), \quad (t_1,t_2) \in \Delta^2, \qquad (2.3.91)$$

which parameterizes the surface patch spanned by $\mathbf{x} = \Gamma_i(\mathbf{t}) \in \mathsf{E}$ via the 2−cell (i.e., rectangle) $\Delta^2 \subset \mathbb{R}^2$, spanned by $\mathbf{t} = \Gamma_i^{-1}(\mathbf{t}) \in \mathbb{R}^2 \times \{0\}$ for $t_3 := 0$. The original (i.e., unextended) defining function $f_{S^2} : \mathsf{E} \to \mathbb{R}_+$ can thus be constructed as

$$f_{S^2}(\mathbf{x}) := \sum_{0\le i<n_2} \rho_i'(t_1,t_2)\ \delta^0(t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \qquad (2.3.92)$$

where $\delta^1(t_3)$ is the indicator function of the origin $t_3 := 0$ of the 1D parametric t_3-line that is normal to the 1−singular 2−cell $\Delta^2 \subset \mathbb{R}^2$.

The (presumably nonzero) Jacobian determinant $|\mathfrak{J}_i(\mathbf{t})|$ can be included to nor-

malize the metrics after coordinate transformation.[42] noting, however, that it does not affect the defining function up to equivalence via $\equiv$. Thus letting the density function be $\rho_i(t_1,t_2) := \rho'_i(t_1,t_2)|\mathfrak{J}_i(t_1,t_2,0)|$,

$$f_{S^2}(\mathbf{x}) := \sum_{0\leq i<n_2} \rho_i(t_1,t_2)\ |\mathfrak{J}_i(t_1,t_2,0)|^{-1}\ \delta^0(t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.93}$$

The functions $\rho_i : \mathbb{R}^2 \to \mathbb{R}_+$ are called 'surface density' distributions that are nonzero over the parameterization 2−cell $\Delta^2 \subset \mathbb{R}^2$ and zero elsewhere, i.e., $\rho_i \equiv \mathbf{1}_{\Delta^2}$, where $\mathbf{1}_{\Delta^2} : \mathbb{R}^2 \to \{0,1\}$ is the cell's indicator function, thus the basis function $\rho_i(t_1,t_2)\ |J_i(t_1,t_2,0)|^{-1}\ \delta^0(t_3)$ is a defining function of each instance of embedding the 2−cell to individual surface patches. For unit constant density $\rho_i(t_1,t_2) := 1$, one obtains the indicator function of the surface patches:

$$\mathbf{1}_{S^2}(\mathbf{x}) := \sum_{0\leq i<n_2} |\mathfrak{J}_i(t_1,t_2,0)|^{-1}\ \delta^0(t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.94}$$

Next, let us replace the above real-valued function $f_{S^2} : \mathsf{E} \to \mathbb{R}_+$ with an "improper" 'extended'-real-valued function $f^1_{S^2} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^1\text{”}\}$ as

$$f^1_{S^2}(\mathbf{x}) := \sum_{0\leq i<n_2} \rho_i(t_1,t_2)\ |\mathfrak{J}_i(t_1,t_2,0)|^{-1}\ \delta^1(t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.95}$$

conceptualized as putting a 1^{st}−order volume singularity of intensity $\rho_i(t_1,t_2) > 0$ distributed over each surface patch. Consequently, the artificial 3−measure of the

[42]The Jacobian is included to ensure that $\rho_i : \mathbb{R}^2 \to \mathsf{E}$ represents distribution per unit physical area over the surface $\Gamma_i(\Delta^2) \subset \mathsf{E}$ rather than per unit parametric elements over the 2−cell $\Delta^2 \subset \mathbb{R}^2$ regardless of the (possibly non-isometric) parameterization.

set, obtained by taking the L^1-norm of (2.3.95), becomes nonzero:

$$\int_{\mathsf{E}} f^1_{S^2}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \iiint_{\mathbb{R}^3} (f^1_{S^2} \circ \Gamma_i)(t_1,t_2,t_3)\ |\mathfrak{J}_i(\mathbf{t})|\ dt_1dt_2dt_3 \tag{2.3.96}$$

$$= \sum_{0\leq i<n_2} \iiint_{\mathbb{R}^3} \rho_i(t_1,t_2) \frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(t_1,t_2,0)|}\ \delta^1(t_3)\ dt_1dt_2dt_3. \tag{2.3.97}$$

Using Lemma 2.3.16 with $\mathrm{d} := 1$ and the function $t_3 \mapsto \rho_i(t_1,t_2)\frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(t_1,t_2,\,0)|}$ yields

$$\int_{\mathsf{E}} f^1_{S^2}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \sum_{0\leq i<n_2} \iint_{\mathbb{R}^2} \rho_i(t_1,t_2)\ \frac{|\mathfrak{J}_i(t_1,t_2,0)|}{|\mathfrak{J}_i(t_1,t_2,0)|}\ dt_1dt_2. \tag{2.3.98}$$

Letting $c_i := \|\rho_i\|_1$ be the total artificial volume of each surface patch,

$$\|f^1_{S^2}\|_1 = \sum_{0\leq i<n_2} \iint_{\mathbb{R}^2} \rho_i(t_1,t_2)\ dt_1dt_2 = \sum_{0\leq i<n_2} \|\rho_i\|_1 = \sum_{0\leq i<n_2} c_i. \tag{2.3.99}$$

For the special case with $\rho_i(t_1,t_2) := 1$ the singular defining function in (2.3.95) turns into a $1^{\text{st}}-$order $\delta-$singular characteristic function $\mathbb{1}^1_{S^2} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^1\text{”}\}$ as

$$\mathbb{1}^1_{S^2}(\mathbf{x}) := \sum_{0\leq i<n_2} |\mathfrak{J}_i(t_1,t_2,0)|^{-1}\ \delta^2(t_3)\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.100}$$

for which the L^1-norm gives the precise total surface area:

$$\|\mathbb{1}^1_{S^2}\|_1 = \int_{\mathsf{E}} \mathbb{1}^1_{S^2}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \sum_{0\leq i<n_2} \mu^2[\Gamma_i(\Delta^1)] = \mu^2[S^2]. \tag{2.3.101}$$

Once again, one can choose to have both ordinary- and singular-volume surface patches in a collection $S_2 = \bigcup_{0\leq i<n_2} \Gamma_i(\Delta^2)$, while the former does not affect the implicit set being described up to equality$-\mathsf{ae}$. See Fig. 2.3.10 for an illustration.

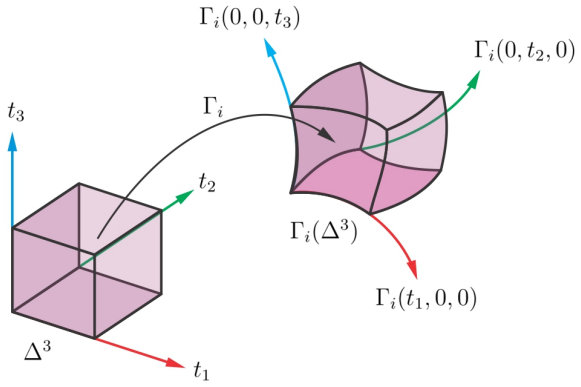

$$\mathbf{1}_{\Delta^3}(t_1,t_2,t_3)=\mathbf{1}_{\Delta^1}(t_1)\mathbf{1}_{\Delta^1}(t_2)\mathbf{1}_{\Delta^1}(t_3)$$

FIGURE 2.3.11: A 3D volume is obtained by mapping a parameterized $3-$cell $\Delta^3\subset\mathbb{R}^3$.

Solid Volumes. Finally, a Hausdorff $3-$measurable union of disjoint $C^\omega-$solid submanifolds $S^3=\bigcup_{0\leq i<n_3}\Gamma_i(\Delta^3)$—with $\mu^3[S^3]$ equal to the total volumes of the n_3 pieces—need not any $\delta-$singularities. Once again, adopting Convention 2.3.12 with $\Gamma_i:(\gamma_{i,1},\gamma_{i,2},\gamma_{i,3}):\mathbb{R}^3\to\mathbb{R}$, for every $0\leq i<n_3$ one obtains

$$(x_1,x_2,x_3):=\Big(\gamma_{i,1}(t_1,t_2,t_3),\gamma_{i,2}(t_1,t_2,t_3),\gamma_{i,2}(t_1,t_2,t_3)\Big),\quad (t_1,t_2,t_3)\in\Delta^3, \tag{2.3.102}$$

which parameterizes the submanifold spanned by $\mathbf{x}=\Gamma_i(\mathbf{t})\in\mathsf{E}$ via the $3-$cell (i.e., cube) $\Delta^3\subset\mathbb{R}^3$, spanned by $\mathbf{t}=\Gamma_i^{-1}(\mathbf{t})\in\mathbb{R}^3$. The original defining function $f_{S^3}:\mathsf{E}\to\mathbb{R}_+$ can thus be constructed as

$$f_{S^3}(\mathbf{x}):=\sum_{0\leq i<n_2}\rho_i'(t_1,t_2,t_3)\;\big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.103}$$

Once again, the Jacobian determinant $|\mathfrak{J}_i(\mathbf{t})|$ can be included to normalize the metrics after coordinate transformation.[43] noting, however, that it does not affect the defining

[43]The Jacobian is included to ensure that $\rho_i:\mathbb{R}^3\to\mathsf{E}$ represents distribution per unit physical area over the volume $\Gamma_i(\Delta^3)\subset\mathsf{E}$ rather than per unit parametric elements over the $3-$cell $\Delta^3\subset\mathbb{R}^3$

function up to equivalence via $\equiv$. Thus letting $\rho_i(t_1,t_2,t_3) := \rho'_i(t_1,t_2,t_3)|\mathfrak{J}_i(t_1,t_2,t_3)|$,

$$f_{S^3}(\mathbf{x}) := \sum_{0\le i<n_3} \rho_i(t_1,t_2,t_3)\ |\mathfrak{J}_i(t_1,t_2,t_3)|^{-1}\big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.104}$$

The functions $\rho_i : \mathbb{R}^3 \to \mathbb{R}_+$ are called 'volume density' distributions that are nonzero over the parameterization tetrahedron $\Delta^3 \subset \mathbb{R}^3$ and zero elsewhere, i.e., $\rho_i \equiv \mathbf{1}_{\Delta^3}$, where $\mathbf{1}_{\Delta^3} : \mathbb{R}^3 \to \{0,1\}$ is the tetrahedron indicator function, thus the basis function $\rho_i(t_1,t_2,t_3)\ |J_i(t_1,t_2,t_3)|^{-1}$ is a defining function of each instance of embedding the tetrahedron to individual volumetric pieces. For unit constant density $\rho_i(t_1,t_2,t_3) := 1$, one obtains the indicator function of the solid volumes:

$$\mathbf{1}_{S^3}(\mathbf{x}) := \sum_{0\le i<n_3} |\mathfrak{J}_i(t_1,t_2,t_3)|^{-1}\ \big|_{(t_1,t_2,t_3):=(\Gamma_i)^{-1}(\mathbf{x})}. \tag{2.3.105}$$

To reconcile with the previous discussion of lower-dimensional strata, (2.3.104) can be viewed as a special case of an "improper" 'extended'-real-valued function $f^0_{S^3} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^0\text{”} = 1\}$:

$$f^0_{S^3}(\mathbf{x}) := \sum_{0\le i<n_3} \rho_i(t_1,t_2,t_3)\ |\mathfrak{J}_i(t_1,t_2,t_3)|^{-1}\ \big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}. \tag{2.3.106}$$

which can be used to obtain the $3-$measure of the set as:

$$\int_{\mathsf{E}} f^0_{S^3}(\mathbf{x})\ d\mu^3[\mathbf{x}] = \iiint_{\mathbb{R}^3} (f^0_{S^3}\circ\Gamma_i)(t_1,t_2,t_3)\ |\mathfrak{J}_i(\mathbf{t})|\ dt_1dt_2dt_3 \tag{2.3.107}$$

$$= \sum_{0\le i<n_3} \iiint_{\mathbb{R}^3} \rho_i(t_1,t_2,t_3)\frac{|\mathfrak{J}_i(t_1,t_2,t_3)|}{|\mathfrak{J}_i(t_1,t_2,t_3)|}\ dt_1dt_2dt_3. \tag{2.3.108}$$

regardless of the (possibly non-isometric) parameterization.

Letting $c_i := \|\rho_i\|_1$ be the total volume of each solid submanifold,

$$\|f^0_{S^3}\|_1 = \sum_{0\le i<n_3} \iint_{\mathbb{R}^2} \rho_i(t_1,t_2,t_3)\; dt_1dt_2dt_3 = \sum_{0\le i<n_3} \|\rho_i\|_1 = \sum_{0\le i<n_3} c_i. \tag{2.3.109}$$

Once again, for the special case with $\rho_i(t_1,t_2,t_3) := 1$ the characteristic function in (2.3.105) can be restated as $\mathbb{1}^0_{S^3} : \mathsf{E} \to \mathbb{R}_+ \otimes \{\text{“}\infty^0\text{”} = 1\}$:

$$\mathbb{1}^0_{S^3}(\mathbf{x}) := \sum_{0\le i<n_3} |\mathfrak{J}_i(t_1,t_2,t_3)|^{-1} \,\big|_{(t_1,t_2,t_3):=\Gamma_i^{-1}(\mathbf{x})}, \tag{2.3.110}$$

for which the L^1-norm gives the precise total solid volume, as expected:

$$\|\mathbb{1}^0_{S^3}\|_1 = \int_{\mathsf{E}} \mathbb{1}^0_{S^3}(\mathbf{x})\; d\mu^3[\mathbf{x}] = \sum_{0\le i<n_3} \mu^3[\Gamma_i(\Delta^3)] = \mu^3[S^3]. \tag{2.3.111}$$

See Fig. 2.3.11 for an illustration, and compare with Figs. 2.3.8 through 2.3.10.

Singular Strata (Summary). The extended-real-valued defining functions of the $\delta-$singular $\mathrm{d}-$polytopes ($0 \le \mathrm{d} < 4$) given in (2.3.76), (2.3.84), (2.3.95), and (2.3.106) are repeated below for $\mathbf{x} = \Gamma_i(\mathbf{t}) = \Gamma_i(t_1,t_2,t_3)$:

$$f^3_{S^0}(\mathbf{x}) = \sum_{0\le i<n_0} \frac{c_i\;\delta^3(t_1,t_2,t_3)}{|\mathfrak{J}_i(0,\;0,\;0)|}, \quad \text{and} \quad \mathbb{1}^3_{S^0}(\mathbf{x}) = \sum_{0\le i<n_0} \frac{\delta^3(t_1,t_2,t_3)}{|\mathfrak{J}_i(0,\;0,\;0)|}, \tag{2.3.112}$$

$$f^2_{S^1}(\mathbf{x}) = \sum_{0\le i<n_1} \frac{\rho_i(t_1)\;\delta^2(t_2,t_3)}{|\mathfrak{J}_i(t_1,\;0,\;0)|}, \quad \text{and} \quad \mathbb{1}^2_{S^1}(\mathbf{x}) = \sum_{0\le i<n_1} \frac{\delta^2(t_2,t_3)}{|\mathfrak{J}_i(t_1,\;0,\;0)|}, \tag{2.3.113}$$

$$f^1_{S^2}(\mathbf{x}) = \sum_{0\le i<n_2} \frac{\rho_i(t_1,t_2)\;\delta^1(t_3)}{|\mathfrak{J}_i(t_1,t_2,\;0)|}, \quad \text{and} \quad \mathbb{1}^1_{S^2}(\mathbf{x}) = \sum_{0\le i<n_2} \frac{\delta^1(t_3)}{|\mathfrak{J}_i(t_1,t_2,\;0)|}, \tag{2.3.114}$$

$$f^0_{S^3}(\mathbf{x}) = \sum_{0\le i<n_3} \frac{\rho_i(t_1,t_2,t_3)}{|\mathfrak{J}_i(t_1,t_2,t_3)|}, \quad \text{and} \quad \mathbb{1}^0_{S^3}(\mathbf{x}) = \sum_{0\le i<n_3} \frac{\rho_i(t_1,t_2,t_3)}{|\mathfrak{J}_i(t_1,t_2,t_3)|}, \tag{2.3.115}$$

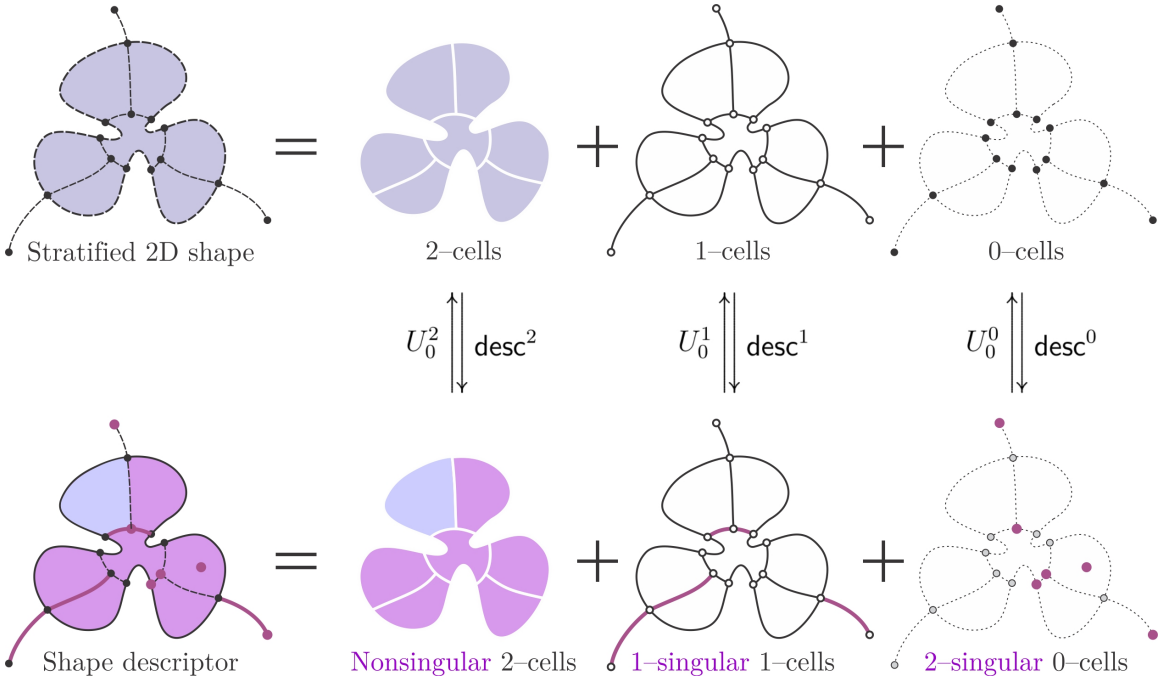

FIGURE 2.3.12: Assigning artificial $\delta-$singular $\mathrm{d}-$measures (here, $\mathrm{d} := 2$) concentrated along some of the lower dimensional strata in Fig. 2.3.7.

the last row being the same as the ordinary defining and characteristic functions. Figure 2.3.12 illustrates the idea for a planar semianalytic shape.

The above derivations can be summarized into the following formal definition:

Definition 2.3.13. (Extended Defining Functions) To enable geometric modeling operations for heterogeneous features, for each shape $S \in \mathscr{P}(\mathsf{E})$ with a priori stratification due to Assumption 2.3.10, each $S^{\mathrm{d}} \in \mathscr{P}^{\mathrm{d}}(\mathsf{E})$ for $0 \leq \mathrm{d} < 4$ can be implicitly described as the $0-$superlevel set of an 'extended'-nonnegative real-valued function $f_S^{\mathrm{d}} : \mathsf{E} \to \mathbb{R}_\star^{\mathrm{d}}$ where $\mathbb{R}_\star^{\mathrm{d}}$ is the set of 'extended'-nonnegative real numbers:

$$\mathbb{R}_\star^{\mathrm{d}} := \ \mathbb{R}_+ \otimes \{1, \text{“}\infty^{3-\mathrm{d}}\text{”}\} := \mathbb{R}_+ \cup \left\{c \times \text{“}\infty^{3-\mathrm{d}}\text{”} \mid c \in \mathbb{R}_+\right\}, \tag{2.3.116}$$

in which the coefficients $c \in \mathbb{R}_+$ can be thought of as intensities of the artificial

d−infinity, added to accommodate the d−dimensional δ−function. Letting $\mathbb{R}_\star := \bigcup_{0\le \mathrm{d}<4} \mathbb{R}^{\mathrm{d}}_\star$, the 'extended'-real-valued defining function for a shape of heterogeneous dimensions $f_S : \mathsf{E} \to \mathbb{R}_\star$ is conceptualized as a formal sum $f_S = \sum_{0\le \mathrm{d}<4} f^{3-\mathrm{d}}_{S^{\mathrm{d}}}$ where $f^{3-\mathrm{d}}_{S^{\mathrm{d}}} = \mathsf{desc}^{\mathrm{d}}(S^{\mathrm{d}}) : \mathsf{E} \to \mathbb{R}^{\mathrm{d}}_\star$ are related to the shape decomposition via

$$f^{3-\mathrm{d}}_{S^{\mathrm{d}}} = \mathsf{desc}^{\mathrm{d}}(S^{\mathrm{d}}) \in (\mathbb{R}^{\mathrm{d}}_\star)^{\mathsf{E}}, \qquad \text{where} \quad S^{\mathrm{d}} = U^{\mathrm{d}}_0(f^{3-\mathrm{d}}_{S^{\mathrm{d}}}). \tag{2.3.117}$$

If the nonnegative real-valued density functions $\rho_i : \mathbb{R}^{\mathrm{d}} \to \mathbb{R}_+$ defined earlier over the parameterized d−strata are chosen to be unity, the so-obtained extended characteristic function is denoted by $\mathbb{1}^{3-\mathrm{d}}_{S^{\mathrm{d}}} = \mathsf{indc}^{\mathrm{d}}(S^{\mathrm{d}}) : \mathsf{E} \to \{0, \text{“}\infty^{3-\mathrm{d}}\text{”}\}$:

$$\mathbb{1}^{3-\mathrm{d}}_{S^{\mathrm{d}}} = \mathsf{indc}^{\mathrm{d}}(S^{\mathrm{d}}) \in \{0, \text{“}\infty^{3-\mathrm{d}}\text{”}\}^{\mathsf{E}}, \quad \text{where} \quad S^{\mathrm{d}} = U^{\mathrm{d}}_0(\mathbb{1}^{3-\mathrm{d}}_{S^{\mathrm{d}}}), \tag{2.3.118}$$

For $\mathrm{d} := 3$, $\mathbb{1}^0_{S^3} = \mathbf{1}_{S^3}$ and $f^0_{S^3} = f_{S^3}$ are the original (i.e., nonsingular) characteristic and defining functions, as in Definitions 2.3.1 and 2.3.2, respectively.

Using the δ−singular defining functions of various orders, the bottom row of the diagram in (2.3.71) can be extended as follows:

$$\begin{array}{cccc}
\cdots & \cdots & \cdots & \cdots \\
\downarrow & \downarrow & \downarrow & \downarrow \\
\mathscr{P}^0(\mathsf{E}) & \mathscr{P}^1(\mathsf{E}) & \mathscr{P}^2(\mathsf{E}) & \mathscr{P}^3(\mathsf{E}) \\
U^0_0 \uparrow\!\downarrow \mathsf{desc}^0 & U^1_0 \uparrow\!\downarrow \mathsf{desc}^1 & U^2_0 \uparrow\!\downarrow \mathsf{desc}^2 & U^3_0 \uparrow\!\downarrow \mathsf{desc}^3 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} & (\mathbb{R}^1_\star)^{\mathsf{E}} & (\mathbb{R}^2_\star)^{\mathsf{E}} & (\mathbb{R}^3_\star)^{\mathsf{E}}
\end{array} \tag{2.3.119}$$

in which the right-most branch is identical to the outer cycle of the diagram in (2.3.36). Using unit densities, one obtains

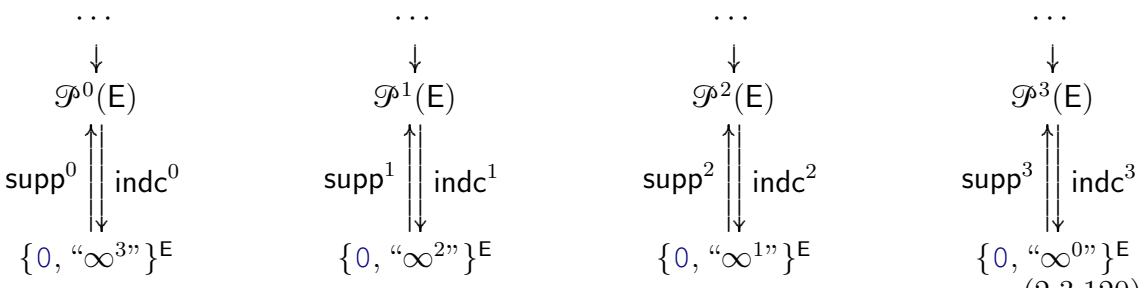

(2.3.120)

Most of the relationships obtained earlier for characteristic and defining functions generalize (with care) to the extended defining functions. For example, the measure relation in (2.3.8) illustrated in (2.3.9) naturally extends to:

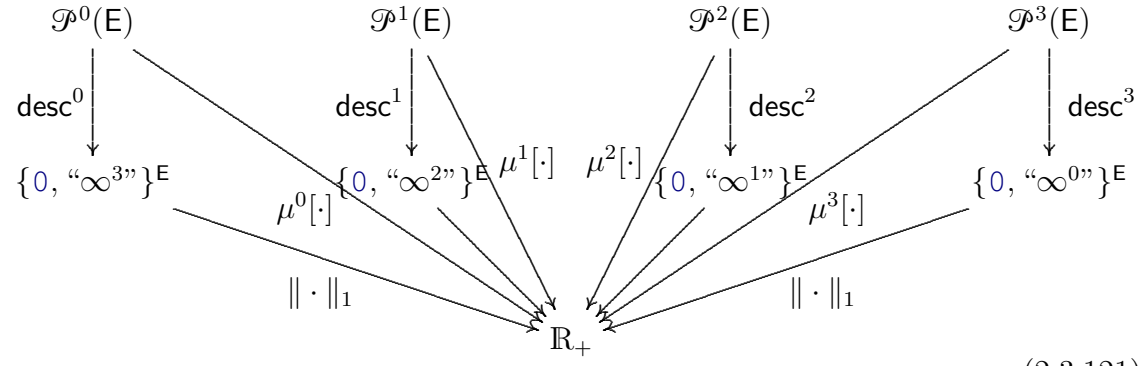

(2.3.121)

serving to emphasize that the L^1-norm based on a single measure uniformly quantifies various lower-dimensional features, provided that $\delta-$singularities are properly assigned. In fact, using the formal sum $\mathbb{1}_S := \sum_{0\le \mathrm{d}<4} \mathbb{1}_{S^{\mathrm{d}}}^{3-\mathrm{d}}$, the arithmetic sum of $0-$, $1-$, $2-$, and $3-$measures is computed in one shot using a single L^1-norm:

$$\|\mathbb{1}_S\|_1 = \|\sum_{0\le \mathrm{d}<4} \mathbb{1}_{S^{\mathrm{d}}}^{3-\mathrm{d}}\|_1 = \sum_{0\le \mathrm{d}<4} \|\mathbb{1}_{S^{\mathrm{d}}}^{3-\mathrm{d}}\|_1 = \sum_{0\le \mathrm{d}<4} \mu^{\mathrm{d}}[S^{\mathrm{d}}]. \tag{2.3.122}$$

Once again, the explicit sets extended with the notion of singular lumped volume

along some of their lower-dimensional features are in one-to-one correspondence with implicit extended-real-valued defining functions modulo function profile, i.e.,

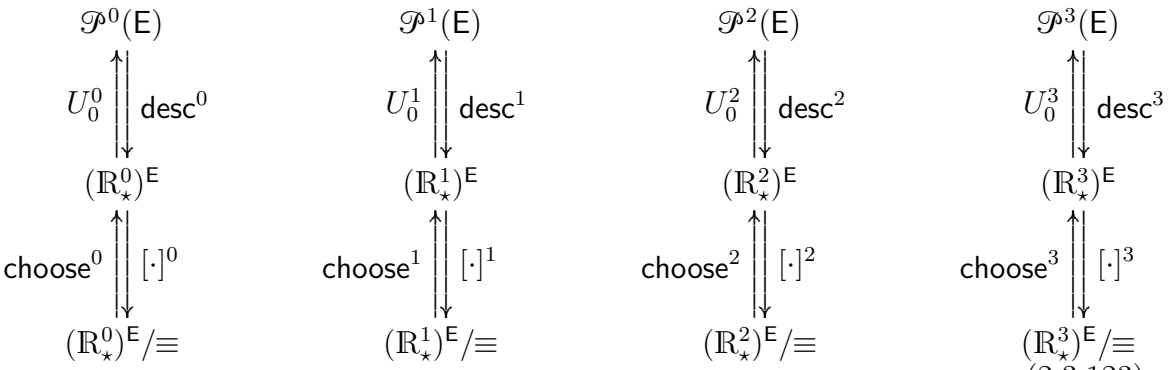

(2.3.123)

in which $[\cdot]^{\mathrm{d}} : (\mathbb{R}_\star^{\mathrm{d}})^{\mathsf{E}} \to (\mathbb{R}_\star^{\mathrm{d}})^{\mathsf{E}}/\equiv$ and $\mathsf{choose}^{\mathrm{d}} : (\mathbb{R}_\star^{\mathrm{d}})^{\mathsf{E}}/\equiv \to (\mathbb{R}_\star^{\mathrm{d}})^{\mathsf{E}}$ are natural generalizations of the equivalence class in (2.3.23) and choice function in (2.3.42) based on the extended definitions. Note also that nominal infinities of different orders are formally summed together into the single extended defining function $f_S : \mathsf{E} \to \mathbb{R}_\star$ defined above, making them theoretically separable:

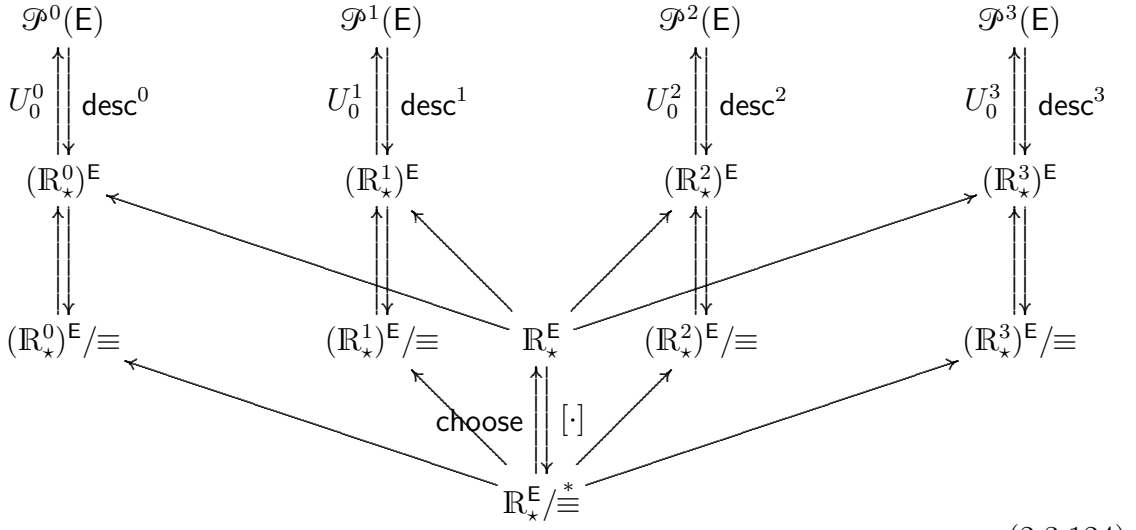

(2.3.124)

In practice, none of the defining functions $f_{S^{\mathrm{d}}}^{3-\mathrm{d}} : \mathsf{E} \to \mathbb{R}_\star^{\mathrm{d}}$ is computable (except for $\mathrm{d} := 3$ when $\mathbb{R}_\star^3 = \mathbb{R}_+$). The trick is to think of them as the limit of real-valued functions, in which the $\delta-$functions of various orders in (2.3.73), (2.3.84), and (2.3.95) are expressed in terms of their limit representations (e.g., selected from Table 2.3.1):

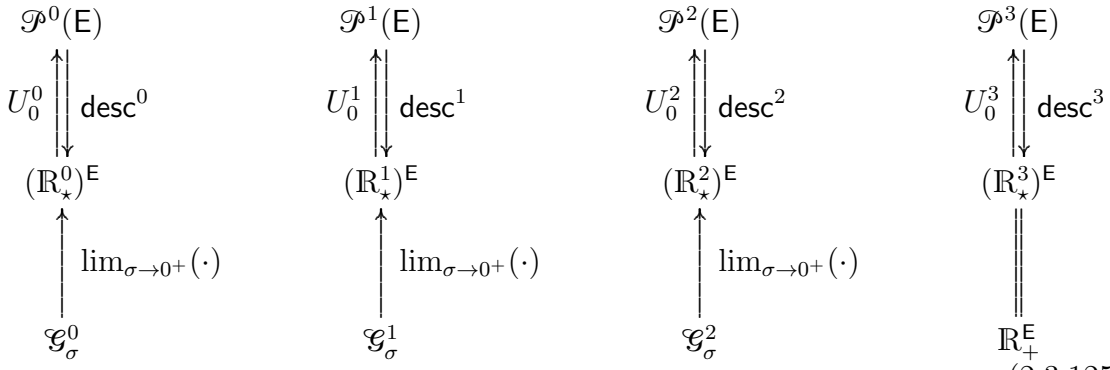

(2.3.125)

where $\mathscr{G}^{3-\mathrm{d}} \subset \mathbb{R}_+^{\mathsf{E}}$ is the set of all functions parameterized by a real variable $\sigma \in (0, \infty)$ such that for all $g_\sigma^{\mathrm{d}} \in \mathscr{G}_\sigma^{\mathrm{d}}$, $\lim_{\sigma\to 0^+} g_\sigma^{\mathrm{d}}(\mathbf{x}) = \delta^{\mathrm{d}}(\mathbf{x})$ for all $\mathbf{x} \in \mathsf{E}$ and $\mathrm{d} = 1, 2$, and 3. Similarly, the formal sum $f_S = \sum_{0\le \mathrm{d}<4} f_{S^{\mathrm{d}}}^{3-\mathrm{d}}$ can be expressed as a single limit:

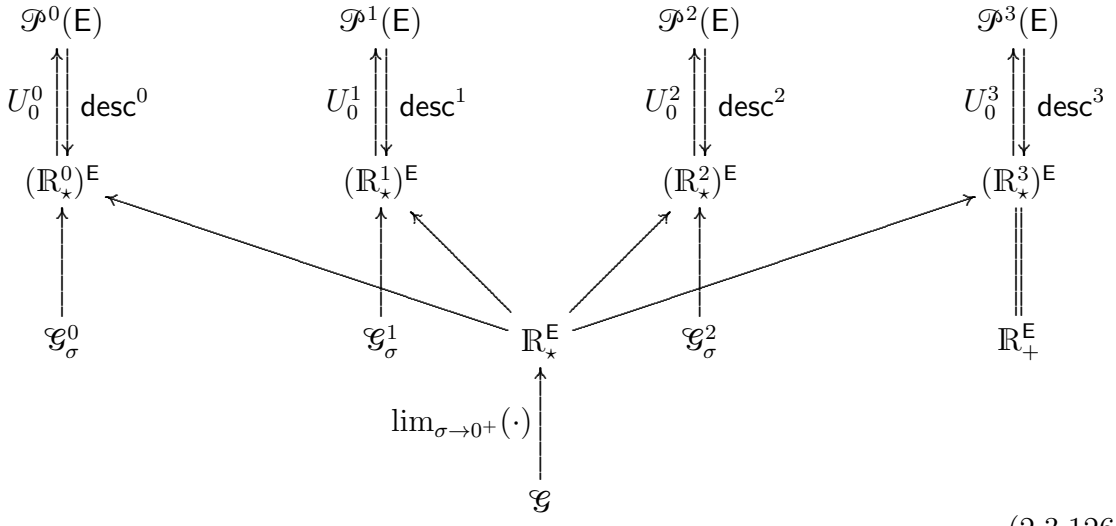

(2.3.126)

where $\mathcal{G} \subset \mathbb{R}_+^{\mathsf{E}}$ is obtained as the direct sum $\mathcal{G} = \mathcal{G}_\sigma^0 \oplus \mathcal{G}_\sigma^1 \oplus \mathcal{G}_\sigma^2$, using a single parameter $\sigma \in (0, \infty)$ for all $\delta-$singularities of various dimensions. One can in principle choose different parameters for each category—or even separately for each stratum—to add flexibility when working with approximations of the $\delta-$functions using finite $\sigma > 0$.

The above diagram provides a natural approximation mechanism by choosing a single finite $\sigma > 0$ depending on the available resources, allowing for a graceful-tradeoff between required computational resources (e.g., time and memory) for resolution and desired precision, to which I shall return in Chapter 5.

2.3.2.3 Equivalence Classes of Extended Defining Functions:

Figure 2.3.13 repeats the defining functions illustrated in Fig. 2.3.3 for 2D shapes with the addition of a few more examples that have lumped densities distributed over arbitrarily chosen edges and vertices. It is important to note that if the extended real semiline $\mathbb{R}_\star$ defined in (2.3.116) is linearly ordered in the natural way, Definition 2.3.3 can be readily extended to redefine the equivalence relation "$\equiv$" $\subset \mathbb{R}_\star^{\mathsf{E}} \times \mathbb{R}_\star^{\mathsf{E}}$ and the equivalence classes $[f_S] \in \mathbb{R}_\star^{\mathsf{E}}/\equiv$ are still closed under the generalized $\mathbb{R}_\star-$arithmetic operations defined similarly to those of (2.3.32) and (2.3.33). In fact, the addition of lumped densities to an implicit shape descriptor does not affect the profile-oblivious abstraction scheme—i.e., the corresponding $0-$superlevel set is still the same.

In a similar fashion, Definition 2.3.4 can also be extended to redefine the equivalence relation "$\overset{*}{\equiv}$" $\subset \mathbb{R}_\star^{\mathsf{E}} \times \mathbb{R}_\star^{\mathsf{E}}$ and the equivalence class $[f_S]^* \in \mathbb{R}_\star^{\mathsf{E}}/\overset{*}{\equiv}$, which are still closed under the generalized $\mathbb{R}_\star-$arithmetic operations defined similarly to those of (2.3.37) and (2.3.38). However, this time the addition of lumped densities

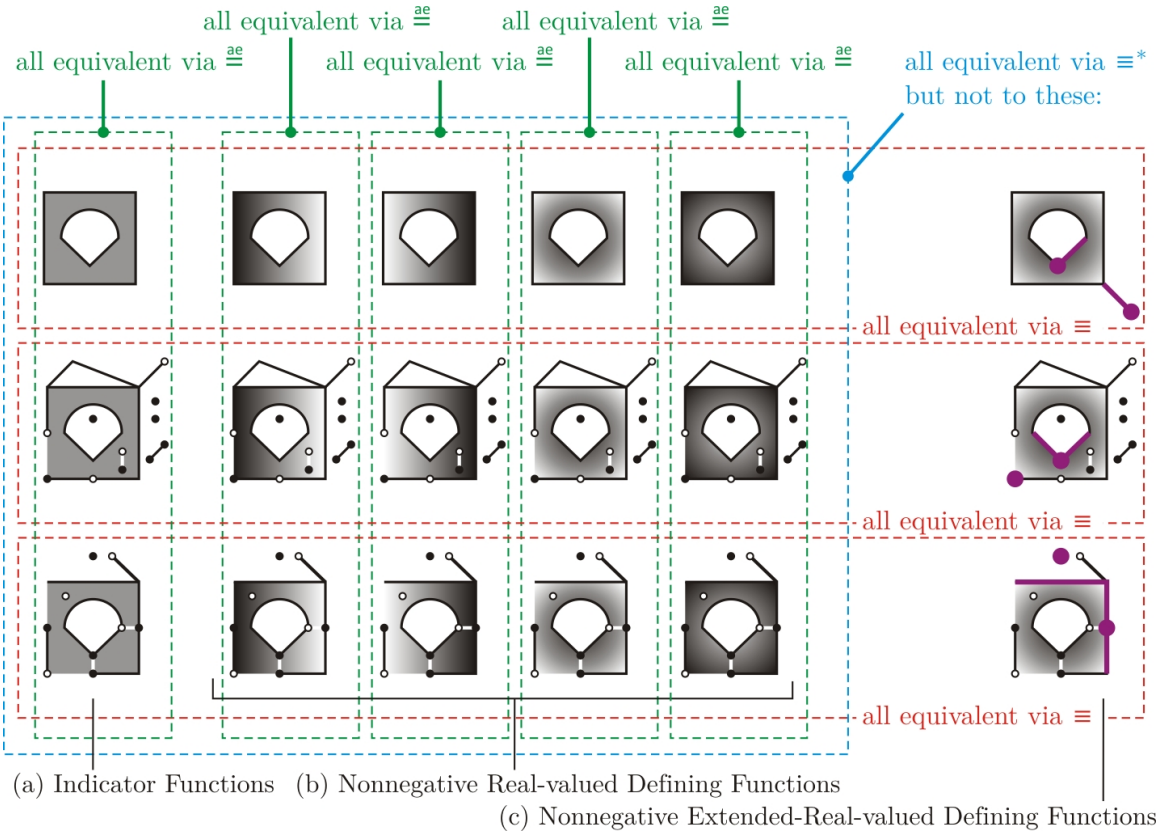

FIGURE 2.3.13: The addition of $\delta-$function-based singularities (shown with thicker bullets and lines) does not affect the $0-$superlevel set of the defining functions, thus keeps them in the same equivalence class via $\equiv$. However, even though the regularized $0-$superlevel sets are also unaltered by the addition of lumped material—e.g., lumped surface areas concentrated into edges and vertices, making them artificially $2-$measurable in this planar example—the defining function no longer have the same measure-theoretic properties 'almost everywhere', thus the new functions belong to different equivalence classes via $\overset{*}{\equiv}$.

to an implicit shape descriptor changes the measure-theoretic properties of the abstracted entities. Therefore, even though the irregularity-tolerant abstraction scheme is expected to remain unaffected—i.e., the corresponding regularized $0-$superlevel set is still the same—the implicit shape descriptor is no longer a member of the same equivalence class—i.e., the function signs (i.e., zero or positive) no longer agrees$-\mathsf{ae}$—as depicted by Fig. 2.3.13. In other words, $0-$superlevel equality$-\mathsf{ae}$ and function sign agreement$-\mathsf{ae}$ do not represent the same semantics anymore. In other words, the equality of regularized $0-$superlevel sets in (2.3.27), which is equivalent to equality$-\mathsf{ae}$ of $0-$superlevel sets in (2.3.28), no longer corresponds to the function sign agreement$-\mathsf{ae}$ in (2.3.29) in Definition (2.3.4). Hereafter, I shall use the latter as the criteria for equivalence via $\overset{*}{\equiv}$ between functions whose codomain is extended to $\mathbb{R}_\star$, i.e., the extended equivalence relation "$\overset{*}{\equiv}$" $\subseteq \mathbb{R}_\star^{\mathsf{E}} \times \mathbb{R}_\star^{\mathsf{E}}$ is redefined as

$$f_1 \overset{*}{\equiv} f_2 \quad \text{iff} \quad f_1(\mathbf{x}) > 0 \overset{\mathsf{ae}}{\rightleftharpoons} f_2(\mathbf{x}) > 0, \text{ for all } \mathbf{x} \in \mathsf{E}, \tag{2.3.127}$$

where the sign agreement$-\mathsf{ae}$ condition on the right-hand side is equivalent to

$$\int_\Omega f_1(\mathbf{x})\ d\mu^3[\mathbf{x}] > 0 \ \rightleftharpoons\ \int_\Omega f_2(\mathbf{x})\ d\mu^3[\mathbf{x}] > 0, \quad \text{for all } \Omega \in \mathscr{P}(\mathsf{E}). \tag{2.3.128}$$

Therefore, the equivalence classes $[f]^* \in \mathbb{R}_\star^{\mathsf{E}}/\overset{*}{\equiv}$ are redefined accordingly as

$$[f]^* := \left\{ f' : \mathsf{E} \to \mathbb{R}_+ \mid (\mathsf{sign} \circ f) \overset{\mathsf{ae}}{=} (\mathsf{sign} \circ f') \right\}, \tag{2.3.129}$$

where the sign function $\mathsf{sign} : \mathbb{R}_\star \to \{0, +1\}$ is extended in the obvious way, including positive multiples of different orders of nominal infinity.

To summarize, the extension of the shape description scheme to include lumped

densities and singularities with the aid of Dirac $\delta-$functions allows for another level of generalization of geometric models. Once again, the individual functions $f \in [f]^*$ are capable of representing more complex geometric objects that carry more information than captured by traditional 'solids'—i.e., their regularized $0-$superlevel sets.

It is important to understand the implications of the $\delta-$singularities when used in Boolean operations. In particular, it is easy to verify that the profile-oblivious Boolean algebra enabled by the $\mathbb{R}_+-$arithmetic operations in (2.3.34) and (2.3.35) extends naturally to $\mathbb{R}_\star-$arithmetic operations.[44] However, the regularized Boolean algebra in (2.3.39) and (2.3.40) will be affected, as the $\delta-$singularities make lower-dimensional features $3-$measurable—hence immune to implicit regularization characterized via equality-ae. However, this does not affect the irregularities that appear and disappear in the intermediate stages.

Finally, it is conceivable that the same concepts can be applied to extend the implicitly regularized motion descriptor classes using Haar measures and tangent spaces (i.e., Lie algebra) for the Lie group $\mathsf{C} = \mathrm{SE}(3)$.

2.3.2.4 Extended Interference Test and Configuration Space Obstacles:

The collision predicate defined in Definition 2.3.5 to capture $3-$measurable interference can be naturally extended to compute lower-dimensional contact:

Definition 2.3.14. (Overlap Predicate) The function $\mathsf{col}^{\mathrm{d}} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \{0, 1\}$ decides if two shapes $S_1, S_2 \in \mathscr{P}^*(\mathsf{E})$ have an intersection with nonzero $\mathrm{d}-$measure,

[44] Taking for granted simple algebraic rules such as $(0 \bullet \infty^{\mathrm{d}}) = 0$, $(\infty^{\mathrm{d}_1} \bullet \infty^{\mathrm{d}_2}) = \infty^{\mathrm{d}_1+\mathrm{d}_2}$, etc.

i.e.,

$$\mathsf{col}^{\mathrm{d}}(S_1, S_2) = \begin{cases} 1 & \text{if } \mu^{\mathrm{d}}[S_1 \cap S_2] > 0, \\ 0 & \text{if } \mu^{\mathrm{d}}[S_1 \cap S_2] = 0. \end{cases} \tag{2.3.130}$$

The d−overlap predicate is covariant under rigid transformations, i.e., $\mathsf{col}^{\mathrm{d}}(S_1, S_2) = \mathsf{col}^{\mathrm{d}}(\mathbb{c}S_1, \mathbb{c}S_2)$ for all $\mathbb{c} \in \mathsf{C}$, meaning that the following diagram commutes:

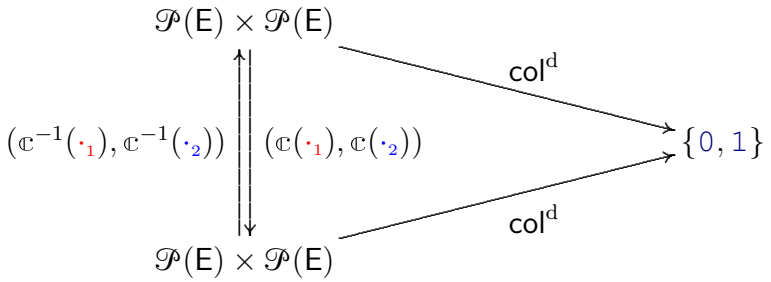

(2.3.131)

The d−overlap subsumes and extends the collision predicate defined earlier in Definition 2.3.5 (in the special case of $\mathrm{d} = 3$):

- For $\mathrm{d} = 0$, it decides if there is an intersection, i.e., at least contact over some points, i.e., $\mathsf{col}^0(S_1, S_2) = 1$ iff $(S_1 \cap S_2) \neq \emptyset$.
- For $\mathrm{d} = 1$, it decides if there is at least contact over a curve of nonzero length;
- For $\mathrm{d} = 2$, it decides if there is at least contact over a surface of nonzero area;
- For $\mathrm{d} = 3$, it decides if there is an interference, i.e., collision over a nonzero volume, i.e., $\mathsf{col}^3(S_1, S_2) = 1$ iff $(S_1 \cap^* S_2) \neq \emptyset$.

Thus a total order $\mathsf{col}^3(S_1, S_2) \preccurlyeq \mathsf{col}^2(S_1, S_2) \preccurlyeq \mathsf{col}^1(S_1, S_2) \preccurlyeq \mathsf{col}^0(S_1, S_2)$ exists, where the minimal overlap corresponds to set-theoretic intersection and the maximal overlap corresponds to regularized intersection.

Given two shapes that are stratified as $S_1 := \bigcup_{0 \leq \mathrm{d} < 4} S_1^{\mathrm{d}}$ and $S_2 := \bigcup_{0 \leq \mathrm{d} < 4} S_2^{\mathrm{d}}$, an overlap between their features of heterogeneous dimensions can be detected using an

inner product of their extended-real-valued defining functions in a similar fashion to Lemma 2.3.6, only if the said features are properly assigned with the correct order of $\delta-$singularities. In particular, if a d_1-dimensional feature of one shape $S_1^{\mathrm{d}_1} \subseteq S_1$ and a d_2-dimensional feature of the other shape $S_2^{\mathrm{d}_2} \subseteq S_2$ are each separately assigned with $(3-\mathrm{d}_1)-$ and $(3-\mathrm{d}_2)-$singularities, respectively, their intersection over a $\mathrm{d}-$dimensional region (i.e., $\mathrm{d}-$overlap) is detected only if it is a *transverse* intersection,[45] in which case

$$(3-\mathrm{d}_1)+(3-\mathrm{d}_2) \geq (3-\mathrm{d}) \quad \text{i.e.,} \quad \mathrm{d} \geq (\mathrm{d}_1+\mathrm{d}_2)-3. \tag{2.3.132}$$

For example, if two 2D manifold surfaces embedded in the 3−space are each assigned with 1−singularities, i.e., $\mathrm{d}_1 = \mathrm{d}_2 = 2$ thus $(\mathrm{d}_1+\mathrm{d}_2)-3 = 1$, their contact over surfaces ($\mathrm{d} = 2$) or curves ($\mathrm{d} = 1$) are captured by a nonzero inner product (since $\mathrm{d} \geq 1$) whereas their contact over points ($\mathrm{d} = 0$) cannot be detected. Figure 2.3.14 illustrates different possible scenarios for simple planar shapes overlapped along their lower-dimensional features, for which 3 must be replaced with 2 in (2.3.132).

Lemma 2.3.15. (Null-Overlap Lemma) *Given* $S_1^{\mathrm{d}_1} \in \mathscr{P}^{\mathrm{d}_1}(\mathsf{E})$ *and* $S_2^{\mathrm{d}_2} \in \mathscr{P}^{\mathrm{d}_2}(\mathsf{E})$*, the inner product of their extended-real-valued defining functions* $f_{S_1^{\mathrm{d}}}^{3-\mathrm{d}_1}, f_{S_2^{\mathrm{d}}}^{3-\mathrm{d}_1} : \mathsf{E} \to \mathbb{R}_\star$ *gives a* $\mathrm{d}-$*overlap predicate if the intersection* $S^{\mathrm{d}} := (S_1^{\mathrm{d}_1} \cap S_2^{\mathrm{d}_2})$ *is transverse:*

$$\langle f_{S_1^{\mathrm{d}_1}}^{3-\mathrm{d}_1}, f_{S_2^{\mathrm{d}_2}}^{3-\mathrm{d}_2} \rangle > 0 \;\rightleftharpoons\; \mu^{\mathrm{d}}[S_1^{\mathrm{d}_1} \cap S_2^{\mathrm{d}_2}] > 0, \quad \text{i.e.,} \quad \mathsf{col}^{\mathrm{d}} = \mathsf{sign} \circ \langle \cdot_1, \cdot_2 \rangle \circ \mathsf{desc}^{\mathrm{d}_{1,2}}, \tag{2.3.133}$$

where $\mathrm{d} = (\mathrm{d}_1+\mathrm{d}_2)-3$ *as a result of transversality assumption.*

[45] An intersection of two manifolds is transverse in E if the direct sum of the tangents to the individual manifolds spans the entire E.

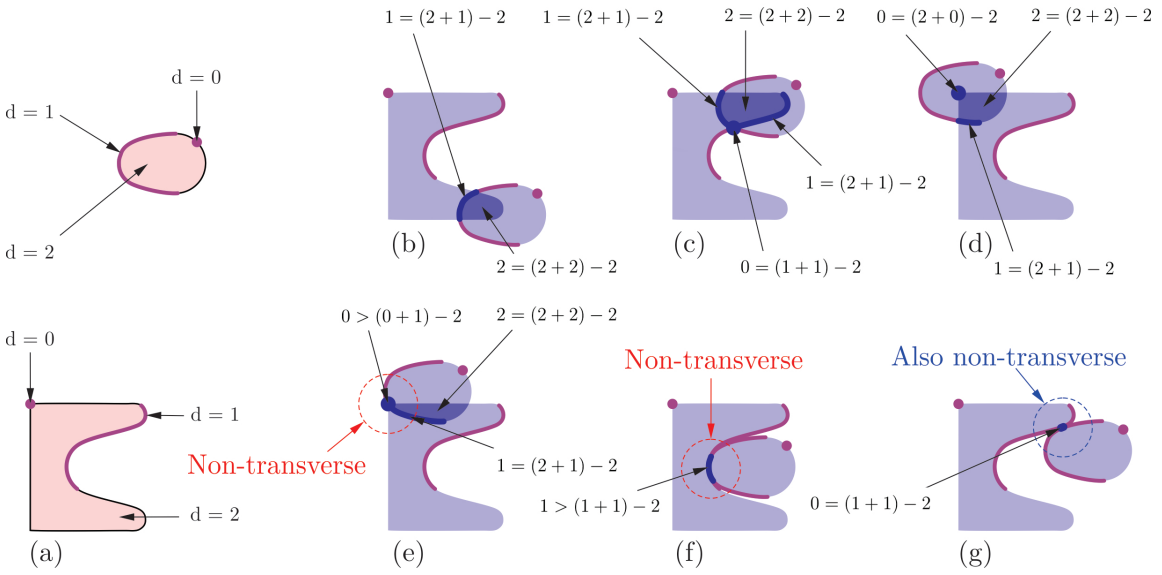

FIGURE 2.3.14: Two 2D shapes in (a) assigned with $(2-\mathrm{d})-$singularities assigned to some of their curve segments $(\mathrm{d}_1, \mathrm{d}_2 = 1)$ and points $(\mathrm{d}_1, \mathrm{d}_2 = 0)$ are tested for d−overlaps in (b–g). In each case, if $\mathrm{d} \geq (\mathrm{d}_1 + \mathrm{d}_2) - 2$ the overlap contributes a nonzero value to the inner product in (2.3.133). The limit behavior at the tangent contact curves is unknown.

Once again, if proper δ−singular characteristic functions are used, the inner product precisely yields the interference d−measure for $\mathrm{d} = (\mathrm{d}_1 + \mathrm{d}_2) - 3$:

$$\langle \mathbb{1}_{S_1^{\mathrm{d}_1}}^{3-\mathrm{d}_1}, \mathbb{1}_{S_2^{\mathrm{d}_2}}^{3-\mathrm{d}_2} \rangle = \int_{\mathsf{E}} \mathbb{1}_{S_1^{\mathrm{d}_1}}^{3-\mathrm{d}_1}(\mathbf{x}) \mathbb{1}_{S_2^{\mathrm{d}_2}}^{3-\mathrm{d}_2}(\mathbf{x})\, d\mu^3[\mathbf{x}] = \int_{S_1^{\mathrm{d}_1} \cap S_2^{\mathrm{d}_2}} d\mu^{\mathrm{d}}[\mathbf{x}] = \mu^{\mathrm{d}}[S_1^{\mathrm{d}_1} \cap S_2^{\mathrm{d}_2}]. \tag{2.3.134}$$

Thus if for two stratified shapes $S_1 = \bigcup_{0 \leq \mathrm{d}_1 < 4} S_1^{\mathrm{d}_1}$ and $S_2 = \bigcup_{0 \leq \mathrm{d}_2 < 4} S_2^{\mathrm{d}_2}$, the δ−singular characteristic functions of different strata are put together in formal sums $\mathbb{1}_{S_1} := \sum_{0 \leq \mathrm{d}_1 < 4} \mathbb{1}_{S_1^{\mathrm{d}_1}}^{3-\mathrm{d}_1}$ and $\mathbb{1}_{S_2} := \sum_{0 \leq \mathrm{d}_2 < 4} \mathbb{1}_{S_2^{\mathrm{d}_2}}^{3-\mathrm{d}_2}$, respectively, the sum of measures of (presumably transverse) intersections is given in a single shot by:

$$\sum_{0 \leq \mathrm{d}_1, \mathrm{d}_2 < 4} \mu^{(\mathrm{d}_1 + \mathrm{d}_2) - 3}[S_1^{\mathrm{d}_1} \pitchfork S_2^{\mathrm{d}_2}] = \sum_{0 \leq \mathrm{d}_1, \mathrm{d}_2 < 4} \langle \mathbb{1}_{S_1^{\mathrm{d}_1}}^{3-\mathrm{d}_1}, \mathbb{1}_{S_2^{\mathrm{d}_2}}^{3-\mathrm{d}_2} \rangle = \langle \mathbb{1}_{S_1}, \mathbb{1}_{S_2} \rangle, \tag{2.3.135}$$

where "⫛" stands for an intersection (i.e., "∩") that is transverse. The non-transverse

intersections (e.g., tangencies) lead to infinite d−measures for $\mathrm{d} = (\mathrm{d}_1 + \mathrm{d}_2) - 3$. Nevertheless, the sign of the inner product remains to be an intersection predicate.

The diagram in (2.3.49) is extended accordingly to measure lower-dimensional overlaps (e.g., surfaces, curves, and points) via inner products, assuming that the features whose overlaps we wish to measure are properly assigned with δ−singularities:

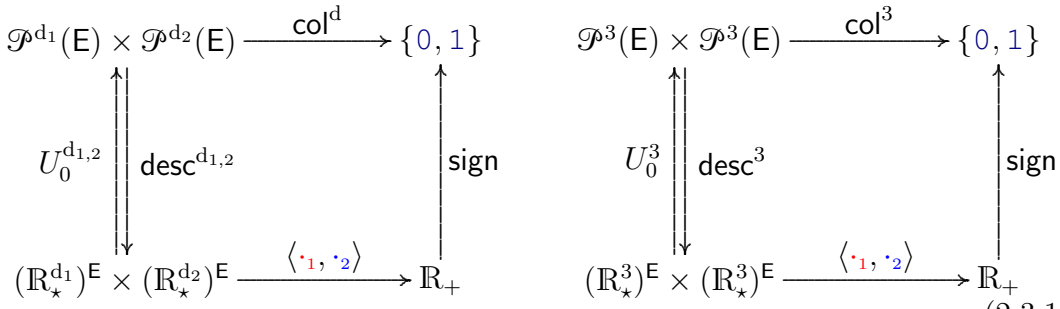

(2.3.136)

where the right-hand side diagram is a special case of the left-hand side diagram for $\mathrm{d}_1 = \mathrm{d}_2 = 3$ (i.e., open 3−manifolds) whose intersection is always transverse in 3D with $\mathrm{d} = (3 + 3) - 3 = 3$. It roughly corresponds to the right-hand side diagram of (2.3.49) for collision test between regular solids.

Once again, if the shapes move via $\mathbb{c}_1, \mathbb{c}_2 \in \mathsf{C}$, respectively, one obtains:

$$\mathsf{col}^{\mathrm{d}}(\mathbb{c}_1 S_1, \mathbb{c}_2 S_2) = \mathsf{col}^{\mathrm{d}}(S_1, (\mathbb{c}_1^{-1}\mathbb{c}_2) S_2) = \mathsf{col}^{\mathrm{d}}((\mathbb{c}_2^{-1}\mathbb{c}_1) S_1, S_2), \tag{2.3.137}$$

Let $O^{\mathrm{d}}_{S_1,S_2} \subseteq \mathsf{C}$ be the set of relative configurations $\mathbb{c} = (\mathbb{c}_1^{-1}\mathbb{c}_2) \in \mathsf{C}$ such that there is a d−overlap, i.e., $\mathsf{col}^{\mathrm{d}}(S_1, (\mathbb{c} S_2))$ evaluates "true". Lemma 2.3.15 implies that the inner product $\langle f_{S_1}, f_{\mathbb{c} S_2} \rangle$—in which the extended defining functions are assigned properly as formal sums $f_{S_1} := \sum_{0 \le \mathrm{d} < 4} f_{S_1^{\mathrm{d}}}^{3-\mathrm{d}_1}$ and $f_{S_2}^{3-\mathrm{d}_2} := \sum_{0 \le \mathrm{d} < 4} f_{S_2^{\mathrm{d}}}$—gives a nonnegative

real-valued defining function of the so-defined 'extended' $\mathsf{C}-$obstacle:

$$O^{\mathrm{d}}_{S_1,S_2} := \big\{\mathbb{c} \in \mathsf{C} \mid \mu^{\mathrm{d}}[S_1, (\mathbb{c}S_2)] > 0\big\} = \big\{\mathbb{c} \in \mathsf{C} \mid \langle f_{S_1}, f_{\mathbb{c}S_2}\rangle > 0\big\}, \tag{2.3.138}$$

The ordering relationship over the overlap predicates is preserved via containment $O^3_{S_1,S_2} \subseteq O^2_{S_1,S_2} \subseteq O^1_{S_1,S_2} \subseteq O^0_{S_1,S_2}$. One again, the above inner product for different values of $\mathbb{c} \in \mathsf{C}$ gives a convolution: $\langle f_{S_1}, f_{\mathbb{c}S_2}\rangle = (f_{S_1} * \tilde{f}_{S_2})(\mathbb{c})$, therefore,

$$O^{\mathrm{d}}_{S_1,S_2} = U_0(f_{S_1} * \tilde{f}_{S_2}), \quad \text{i.e.,} \quad \mathsf{desc}(O^{\mathrm{d}}_{S_1,S_2}) \equiv (f_{S_1} * \tilde{f}_{S_2}), \tag{2.3.139}$$

which can be illustrated using the following commutative diagram:

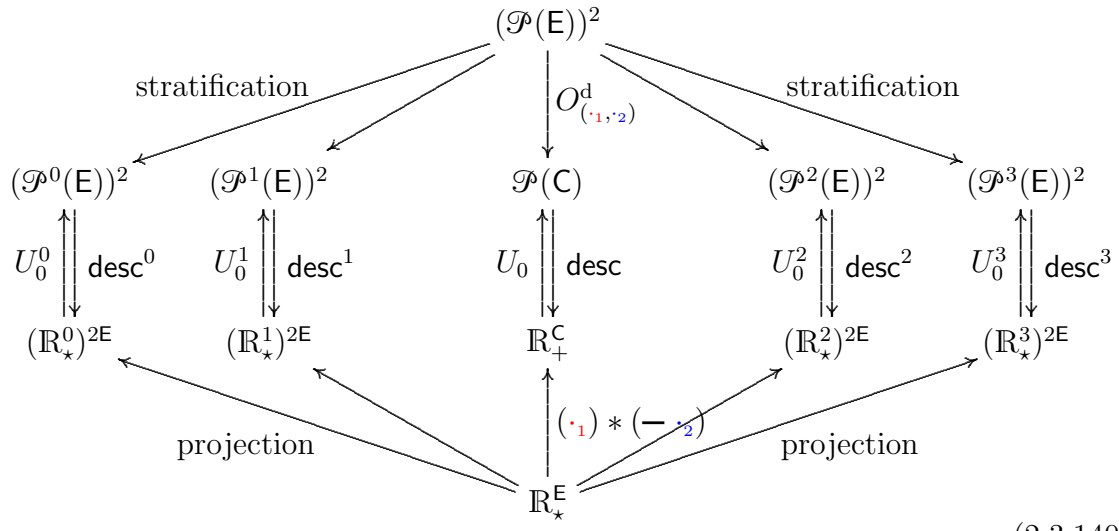

(2.3.140)

Figure 2.3.15 redraws Fig. 2.3.5 of the $\mathsf{C}-$obstacles for a pair of planar shapes, this time taking into account lower-dimensional intersections as well. Once again, the motion is restricted to translation only for illustrative purposes. Note that configurations that result in d$-$ and higher-dimensional intersections are included in the $O^{\mathrm{d}}_{S_1,S_2}$, which is shown for $\mathrm{d} := 0, 2$ in the figure.

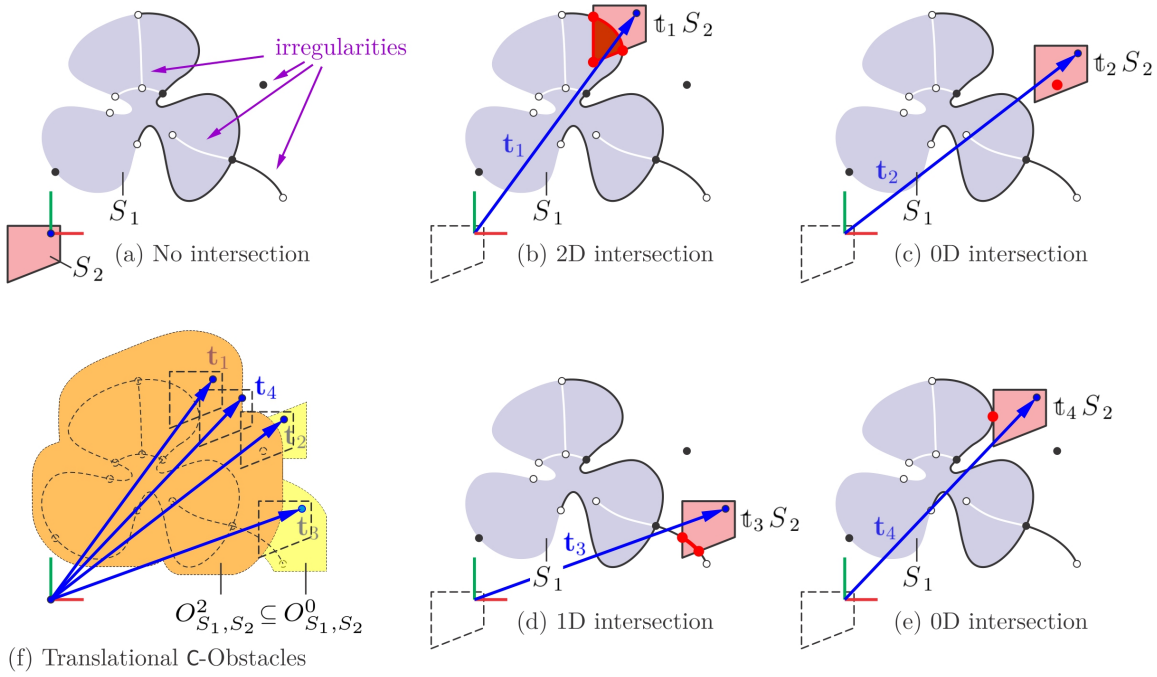

FIGURE 2.3.15: The generalized translational C−obstacles in (f) for a pair of shapes on the 2D plane are the subset of all relative displacements at which there are intersections of nonzero d−measure for d := 0, 1, and 2. Compare to Fig. 2.3.5.

As before Fourier transforms can be leveraged for efficient computation. Using Parseval's theorem (Theorem 2.2.10), the inner product in (2.3.133) can be computed in the Fourier domain. Thus Lemma 2.3.15 can be restated via:

$$\langle \hat{f}^{3-\mathrm{d}_1}_{S_1^{\mathrm{d}_1}}, \hat{f}^{3-\mathrm{d}_2}_{S_2^{\mathrm{d}_2}} \rangle > 0 \;\rightleftharpoons\; \mu^{\mathrm{d}}[S_1^{\mathrm{d}_1} \cap S_2^{\mathrm{d}_2}] > 0, \quad \text{i.e.,} \quad \mathsf{col}^{\mathrm{d}} = \mathsf{sign} \circ \langle \cdot_1, \cdot_2 \rangle \circ \mathscr{F} \circ \mathsf{desc}^{\mathrm{d}_{1,2}}, \tag{2.3.141}$$

where $\hat{f}^{3-\mathrm{d}_1}_{S_1^{\mathrm{d}_1}} = \mathscr{F}\{f^{3-\mathrm{d}_1}_{S_1^{\mathrm{d}_1}}\}$ and $\hat{f}^{3-\mathrm{d}_2}_{S_2^{\mathrm{d}_2}} = \mathscr{F}\{f^{3-\mathrm{d}_2}_{S_2^{\mathrm{d}_2}}\}$, noting that $\langle \cdot_1, \cdot_2 \rangle \circ \mathscr{F} = \langle \cdot_1, \cdot_2 \rangle$ due to (2.2.13). Applying (2.2.13) to (2.3.136) extends the diagrams to

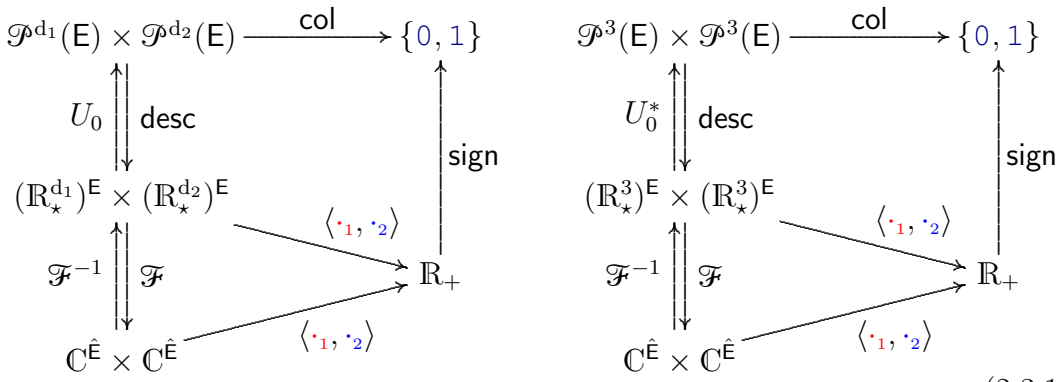

(2.3.142)

Similar to the case with the collision predicate in (2.3.56), one can obtain an approximate d−overlap predicate by truncating the Fourier expansions, i.e., low-pass filtering the Fourier descriptors $\hat{f}^{3-\mathrm{d}_1}_{S_1^{\mathrm{d}_1}} \in \mathbb{R}^{\hat{\mathsf{E}}}_\star$ and $\hat{f}^{3-\mathrm{d}_1}_{S_2^{\mathrm{d}_2}} \in \mathbb{R}^{\hat{\mathsf{E}}}_\star$. The development is similar to that of (2.3.57) and will not be repeated.

Once again, the convolution theorem (Theorem 2.2.11) allows one to convert the convolution of shape descriptor functions in (2.3.139) to a pointwise multiplication of their Fourier transforms. In other words, the Fourier descriptor of the C−obstacles can be obtained from pointwise multiplication of Fourier descriptors of the shapes—

with an additional reflection $\mathsf{-} : \mathbb{R}^{\hat{\mathsf{C}}}_{+} \to \mathbb{R}^{\hat{\mathsf{C}}}_{+}$:

$$O^{\mathrm{d}}_{S_1,S_2} = (U_0 \circ \mathscr{F}^{-1})(\hat{f}_{S_1} \bullet \check{f}_{S_2}), \quad \text{i.e.,} \quad (\mathscr{F} \circ \mathsf{desc})(O_{S_1,S_2}) \equiv (\hat{f}_{S_1} \bullet \check{f}_{S_2}), \quad (2.3.143)$$

where $\check{f}_S := \tilde{\hat{f}}_S = \hat{\tilde{f}}_S$ is the reflection (i.e., inversion) of the Fourier transform, which is the same as the Fourier transform of reflection due to linearity. Once again, the low-pass filtering trick can be used to obtain an approximation of the generalized $\mathsf{C}-$obstacle's Fourier descriptor, whose development is similar to that of (2.3.60) and will not be repeated here.

2.3.3 Delta Calculus

2.3.3.1 $\delta-$Sifting: Measure-Theoretic Tool for Support Regularization:

The following is the fundamental property of the $\delta-$function which allows one to replicate integrals over lower-dimensional manifolds with uniform Lebesgue 3−integrals:

Lemma 2.3.16. (Sifting Property) *Given any real-valued function $f : \mathsf{E} \to \mathbb{R}$ that is C^1-differentiable at the origin $\mathbf{0} \in \mathsf{E}$ and the 3−measurable domain $S \in \mathscr{P}(\mathsf{E})$,*

$$\int_S \delta^3(\mathbf{x}') f(\mathbf{x}') \, d\mu^3[\mathbf{x}'] = \begin{cases} f(\mathbf{0}) & \text{if } \mathbf{0} \in \mathscr{i}\,(S), \\ \lambda f(\mathbf{0}) & \text{if } \mathbf{0} \in \mathscr{b}(S), \quad \text{for some } \lambda \in [0,1], \\ 0 & \text{if } \mathbf{0} \in \mathscr{e}(S). \end{cases} \quad (2.3.144)$$

which satisfies the conditions in Assumption 2.3.8.

The above equation is a generalization of the integral equation in (2.3.66) corresponding to the special case $f(\mathbf{x}') = 1$ for all $\mathbf{x}' \in S$. Now, let $\varsigma^3_{\mathbf{x}}(\mathbf{x}') = \delta^3(\mathbf{x}' - \mathbf{x})$ denote a

'shifted' $\delta-$function, where $\mathbb{c} \stackrel{0}{:=} (\mathbb{r}_{\mathsf{id}}, \mathbf{x})$ is a translation. Thus (2.3.146) yields:

Corollary 2.3.17. (Sifting Property) *Given any real-valued function* $f : \mathsf{E} \to \mathbb{R}$ *that is* C^1-*differentiable at the point* $\mathbf{x} \in \mathsf{E}$ *and the* $3-$*measurable domain* $S \in \mathscr{P}(\mathsf{E})$,

$$\int_S \varsigma^3_{\mathbf{x}}(\mathbf{x}') f(\mathbf{x}') \; d\mu^3[\mathbf{x}'] = \begin{cases} f(\mathbf{x}) & \text{if } \mathbf{x} \in \mathscr{i}\,(S), \\ \lambda f(\mathbf{x}) & \text{if } \mathbf{x} \in \mathscr{b}(S), \quad \text{for some } \lambda \in [0,1], \\ 0 & \text{if } \mathbf{x} \in \mathscr{e}(S). \end{cases} \tag{2.3.145}$$

which satisfies the conditions in Assumption 2.3.8.

The left-hand side of (2.3.145) can be rewritten in terms of the indicator functions as

$$\left\langle (f \bullet \mathbf{1}_S), \varsigma^3_{\mathbf{x}} \right\rangle = \left\| (f \bullet \mathbf{1}_S) \bullet \varsigma^3_{\mathbf{x}} \right\|_1 = \begin{cases} \lambda' f(\mathbf{x}) & \text{if } \mathbf{x} \in \mathscr{k}(S), \\ 0 & \text{if } \mathbf{x} \in \mathscr{e}(S). \end{cases} \tag{2.3.146}$$

where $\lambda' = 1$ if $\mathbf{x} \in \mathscr{i}\,(S)$ and $\lambda' = \lambda$ if $\mathbf{x} \in \mathscr{b}(S)$ which satisfies Assumption 2.3.8, i.e., $\lambda > 0$ iff the query point $\mathbf{x} \in \mathsf{E}$ has a full-dimensional (i.e., non-zero volume) neighborhood in S, i.e., for all $r > 0$:

$$\left\langle (f \bullet \mathbf{1}_S), \varsigma^3_{\mathbf{x}} \right\rangle > 0 \quad \rightleftharpoons \quad \mu^3[S \cap B_r(\mathbf{x})] > 0 \text{ and } f(\mathbf{x}) > 0. \tag{2.3.147}$$

Observe that if $f(\mathbf{x}) > 0$ for all $\mathbf{x} \in S$, then $(f \bullet \mathbf{1}_S) \equiv \mathbf{1}_S$ is a nonnegative real-valued defining function of the domain. In that case, the above inner product gives a predicate to test if the query point belongs to the regularization of the given domain, if the following assumption is made:

Assumption 2.3.18. (Cusp Free Shapes) Hereafter, we assume the shapes $S \in \mathscr{P}(\mathsf{E})$ to be 'cusp-free' everywhere on their regularized boundaries, meaning that the exists

$\theta > 0$ such that at every point $\mathbf{x} \in \partial(\mathscr{r}(S))$, the neighborhood $(S \cap B_r(\mathbf{x}))$ can be swept rotationally by an open half-cone of half-angle $\theta > 0$ around its apex.[46] Subsequently, it can be verified that such shapes have full-dimensional neighborhoods at all points after regularization, i.e., $\mathbf{x} \in \mathscr{r}(S)$ iff $\mu^3[S \cap B_\epsilon(\mathbf{x})] = O(\epsilon^3)$ as $\epsilon \to 0^+$.

Proposition 2.3.19. *For a given shape $S \in \mathscr{P}(\mathsf{E})$, whose nonnegative real-valued defining function is $f_S : \mathsf{E} \to \mathbb{R}_+$, $\langle f_S, \varsigma^3_{\mathbf{x}} \rangle > 0$ iff $f_{\mathscr{r}(S)}(\mathbf{x}) > 0$, i.e., the inner product with $\varsigma^3_{\mathbf{x}}(\cdot) = \delta^3(\mathbf{x} - \cdot)$ yields the defining function of the regularized set $\mathscr{r}(S) \in \mathscr{P}^*(\mathsf{E})$, in accordance with* Assumptions 2.3.8 and 2.3.18.

Proof. Lemma 2.3.9 implies that $\lambda(S) = \langle \delta^3, \mathbf{1}_S \rangle > 0$, or equivalently, $\langle \delta^3, f_S \rangle > 0$ iff $\mu^3[S \cap B_\epsilon(\mathbf{0})] = O(\epsilon^3)$—noting that $\mu^3[B_\epsilon(\mathbf{0})] = \frac{4}{3}\pi\epsilon^3 = O(\epsilon^3)$—as $\epsilon \to 0^+$. In turn, Assumption 2.3.18 posits that this is the case iff $\mathbf{0} \in \mathscr{r}(S)$. By shifting the origin to the arbitrary query point, $\langle \varsigma^3_{\mathbf{x}}, f_S \rangle > 0$ iff $\mathbf{x} \in \mathscr{r}(S)$, where $f_S := (f \bullet \mathbf{1}_S)$. □

On the other hand, inner products of the form $\langle f_1(\cdot), f_2(\mathbf{x} - \cdot) \rangle$ can be interpreted as instantiations of a convolution $(f_1 * f_2)(\cdot)$, as demonstrated in Definition 2.2.8. Therefore, the $\delta-$function serves as the identity element for the convolution algebra on $L^1(\mathsf{E})$ up to equality$-\mathsf{ae}$. In other words, for every integrable function $f \in L^1(\mathsf{E})$, its convolution with the $\delta-$functon gives another function that is equal$-\mathsf{ae}$ to the original function (i.e., $f \overset{\mathsf{ae}}{=} f'$). Furthermore, the output function's $0-$superlevel set is the closure of the input function's $0-$superlevel set:

Theorem 2.3.20. ($\delta-$Regularization) *Convolution with the $\delta-$function is the implicit analogue of topological regularization for every $S \in \mathscr{P}(\mathsf{E})$ and $f_S \in \mathbb{R}_+^{\mathsf{E}}$:*

$$(f_S * \delta^3) \equiv f_{\mathscr{r}(S)} \quad \rightleftharpoons \quad U_0(f_S * \delta^3) = U_0^*(f_S) = \mathscr{r}(S), \tag{2.3.148}$$

[46] This is a reasonable assumption for shapes of practical usefulness.

in accordance with Assumptions 2.3.8 and 2.3.18.

In other words, $\delta-$convolution of defining function is equivalent via $\equiv$ to the defining function of (closed) regularization, i.e., $(\mathsf{desc} * \delta^3) \equiv (\mathscr{r} \circ \mathsf{desc})$. This can be illustrated by the following commutative diagram—remembering the notation $U_0^* = (\mathscr{r} \circ U_0)$:

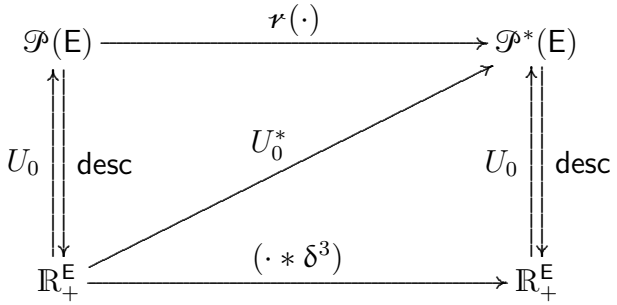

(2.3.149)

Recalling the implicit regularization from Section 2.2.2 where regularization was viewed as a quotient operation—i.e., noting that regular sets are isomorphic to general sets modulo all possible irregularities—the above diagram can be extended as

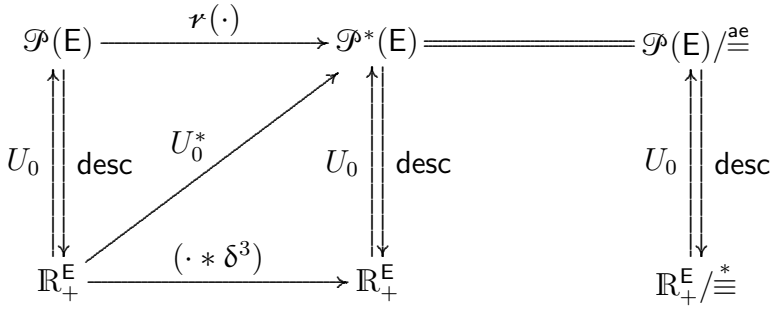

(2.3.150)

To complete the picture, the above diagram and (2.3.31) can be composed to give

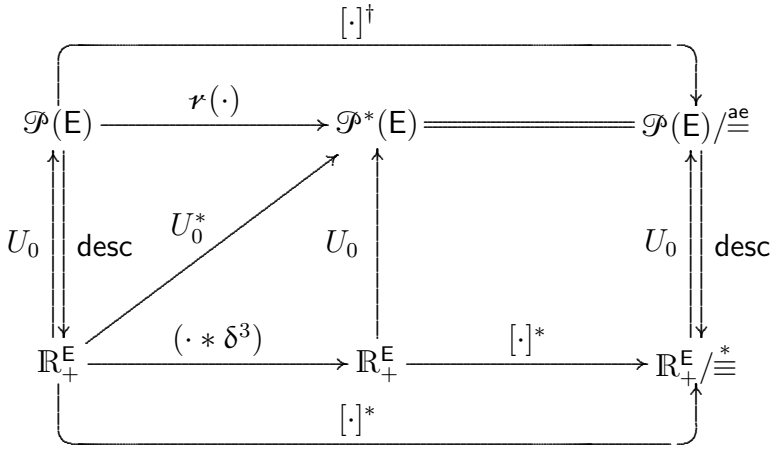

(2.3.151)

Therefore, the $\delta-$function can be used as a generic topological operator for the implicit paradigm. To summarize, two different views are possible to interpret (2.3.159):

Point Membership Classification. From a local perspective, (2.3.159) provides a means for point membership classification (PMC) queries on the implicit regularization. In particular, given a candidate point $\mathbf{x} \in \mathsf{E}$ against a reference set $S \in \mathscr{P}(\mathsf{E})$ described via its defining function $f_S : \mathsf{E} \to \mathbb{R}_+$ the $\delta-$convolution function $(f_S * \delta^3) : \mathsf{E} \to \mathbb{R}_+$ classifies the point into four different regions:

- If $\mathbf{x} \in i(S)$, $\lambda' = 1$ in (2.3.146) thus $(f_S * \delta^3)(\mathbf{x}) = f_S(\mathbf{x})$, i.e., $(f_S * \delta^3) = f_S$.
- If $\mathbf{x} \in b(S)$ and it has a 3−neighborhood in S, $\lambda' = \lambda > 0$ in (2.3.146), thus $(f_S * \delta^3)(\mathbf{x}) > 0$ iff $f_S(\mathbf{x}) > 0$, i.e., $(f_S * \delta^3) \equiv f_S$.
- If $\mathbf{x} \in b(S)$ but it does not have a 3−neighborhood in S, $\lambda' = \lambda = 0$ in (2.3.146), thus $(f_S * \delta^3)(\mathbf{x}) = 0$ regardless of $f_S(\mathbf{x})$.
- If $\mathbf{x} \in e(S)$, $(f_S * \delta^3)(\mathbf{x}) = 0$ due to the second clause of (2.3.146).

Importantly, this also provides a natural mechanism to extend the non-computable PMC to a membership density function (MDF), suggested in the past by other researchers [219, 369] for fuzzy solid modeling. In fact, if the $\delta-$function is viewed as the limit $\delta^3(\mathbf{x}) = \lim_{\sigma\to 0^+} \not{g}^3_\sigma(\mathbf{x})$ (e.g., selected from Table 2.3.1), the MDF returns an intermediate value between 0 and $\lambda' f_S(\mathbf{x})$ if the query point is "close enough" to the boundary, depending on the choice of $\sigma > 0$. This is referred to as 'mollification', to which I shall return in Chapter 5. [47]

Analytic Support Regularization. From a global perspective, the sifting property is extremely powerful as it provides an analytic mechanism to perform pointset-topological regularization on implicit descriptions. In particular, given the defining function of an arbitrary set as $f_S : \mathsf{E} \to \mathbb{R}_+$,

- the defining function of its closed regularization $f_{r_c(S)} : \mathsf{E} \to \mathbb{R}_+$ amounts to computing a $\delta-$convolution; whereas
- the defining function of its open regularization $f_{r_o(S)} : \mathsf{E} \to \mathbb{R}_+$ amounts to computing a $\delta-$convolution and two negations, exploiting its duality with the former via generalized negation.

Once again, if the $\delta-$function is is viewed as the limit $\delta^3(\mathbf{x}) = \lim_{\sigma\to 0^+} \not{g}^3_\sigma(\mathbf{x})$ (e.g., selected from Table 2.3.1), these topological transformations can be viewed as the limit of 'mollification' processes via convolution algebra and the mollifier $\not{g}_\sigma-$functions.

[47] The computational properties of the $\delta-$singularities requires more research. Qualitatively, one can approximate the $\delta-$function by its limit representations and approach to the exact results as much as desired by decreasing the finite precision $\sigma > 0$. This gives the opportunity to formalize a computational theory of implicit geometric modeling with finite precision computing capabilities, whose rigorous formalization is beyond the scope of this thesis.

These transformations smooth out the geometry and allow formulating operations such as offsetting, blending, etc., which are the subject matter of Chapter 5.

The above notions can be generalized (in principle) from shapes to motions by defining (left and right) $\delta-$functions in terms of (left and right) Haar measures. The vector space (i.e., Abelian) convolution is also replaced with noncommutative group convolution operator $* : L^1(\mathsf{C}) \times L^1(\mathsf{C}) \rightarrow L^1(\mathsf{C})$ defined in (2.2.9) and (2.2.10). Conceivably, (left and right) sifting can be obtained by convolving (left and right) $\delta-$functions with defining functions of motion over the configuration space. The detailed treatment of $\delta-$calculus for the configuration space $\mathsf{C} = \mathrm{SE}(3)$—and for and Lie groups and Lie algebras in general—is an interesting problem to study, but lies beyond the scope of this thesis.

2.3.3.2 $\delta-$Composition: Direct Generation of $\delta-$Singular Level Sets

The assignment of $\delta-$singularities of various orders to lower-dimensional features presented after Convention 2.3.11 requires a priori knowledge of an explicit stratification of the shape. However, the whole point of this thesis is constructing an algebra of functions that is expressive enough to replicate important geometric operations on implicit descriptions without a need to reproduce explicit descriptions in intermediate states. If a given shape is obtained from upsteam implicit methods, it is desirable to obtain the $\delta-$singular defining functions of its lower-dimensional features directly from its nonsingular defining function without appealing to explicit stratification.

The following generalization of Lemma 2.3.16 obtained from [150] is the single most important behavioral characteristic of the $\delta-$function that makes it a powerful tool to link measure-theoretic properties to topology, particularly for navigating across

heterogeneous dimensions, using a single Lebesgue measure:

Lemma 2.3.21. (Generalized Root Sum Integral) *Given a C^1-differentiable real-valued function $f : \mathsf{E} \to \mathbb{R}$ whose $0-$isolevel set (i.e., 'root set') $Z := f^{-1}(0)$ is an integrable surface with $\nabla f(\mathbf{x}') \neq \mathbf{0}$ for all $\mathbf{x}' \in Z$, the composition $(\delta^1 \circ f)$ yields a $\delta-$singularity distribution over the surface:*

$$(\delta^1 \circ f)(\mathbf{x}) = \int_Z \frac{\varsigma^2_{\mathbf{x}}(\mathbf{x}')}{\|\nabla f(\mathbf{x}')\|_2} \, d\mu^2_Z[\mathbf{x}'] = \int_Z \frac{\delta^2(\mathbf{x} - \mathbf{x}')}{\|\nabla f(\mathbf{x}')\|_2} \, d\mu^2_Z[\mathbf{x}'], \tag{2.3.152}$$

where $\mu^2_Z[\cdot]$ is the Hausdorff $2-$measure (i.e., surface area) over the root set.

Proof. The proof is beyond the scope of this thesis. A generalization of the above result can be found in distribution theory, referred to as the 'simple layer' integral [150]. See also [129, 130] for a 2D version (therein called 'sweet' $\delta-$formulae). □

The above result is extremely useful in implicit modeling, especially when computing lower-dimensional features of shape—e.g., computing a defining function for the boundary and the medial axis (MA) from the defining function of the given shape. The following is an important corollary that will be leveraged in Chapter 3:

Corollary 2.3.22. (Generalized Sifting Property) *Given a C^1-differentiable real-valued function $f : \mathsf{E} \to \mathbb{R}$ whose $0-$isolevel set (i.e., 'root set') $Z := f^{-1}(0)$ is an integrable surface with $\nabla f(\mathbf{x}') \neq \mathbf{0}$ for all $\mathbf{x}' \in Z$, the L^1-norm of the composition $(\delta^1 \circ f)$ yields a weighted $2-$measure of the root set:*

$$\|\delta^1 \circ f\|_1 = \int_{\mathsf{E}} (\delta^1 \circ f)(\mathbf{x}) \, d\mu^3[\mathbf{x}] = \int_Z \frac{1}{\|\nabla f(\mathbf{x}')\|_2} \, d\mu^2[\mathbf{x}']. \tag{2.3.153}$$

The above result simplifies to the (unweighted) 2−measure (i.e., surface area) $\mu^{\mathrm{d}-1}[Z]$ if the gradient's norm happens to have unit norm everywhere over Z:

$$\|\nabla f(\mathbf{x}')\|_2\Big|_{\mathbf{x}'\in Z} = 1 \;\Rightarrow\; \|\delta^1 \circ f\|_1 = \int_Z d\mu^2[\mathbf{x}'] = \mu^2[Z]. \tag{2.3.154}$$

Figure 2.3.16 illustrates the idea for the simplest case of a 1D (i.e., univariate) function $f : \mathbb{R} \to \mathbb{R}$ whose root set $Z = f^{-1}(0)$ is a finite number of isolated points with nonzero gradient $\frac{df}{dx}(x') \neq 0$ for all $x' \in Z$, meaning that none of the roots have 'multiplicity' greater than one. A lower-dimensional version of (2.3.152)—sometimes referred to as the original 'root sum' lemma—yields the following for every $x \in \mathbb{R}$:

$$(\delta^1 \circ f)(x) = \sum_{x'\in Z} \frac{\varsigma^1_x(x')}{|\frac{df}{dx}(x')|} = \sum_{x'\in Z} \frac{\delta^1(x-x')}{|\frac{df}{dx}(x')|} = \sum_{0\leq i<n_0} c_i \delta^1(x - x'_i). \tag{2.3.155}$$

where $Z = \{x'_i\}_{0\leq i<n_0}$ with a cardinality $|Z| = n_0$, and $c_i := \left(\frac{df}{dx}(x'_i)\right)^{-1}$. Thus $(\delta^1 \circ f)$ assigns 1^{st}−order δ−singularities to each of the isolated single roots whose intensities are inversely proportional to the derivative of the function at the roots. Integrating this function over the entire real-line (i.e., L^1−norm) returns the sum of c_i coefficients:

$$\|\delta^1 \circ f\|_1 = \int_{\mathbb{R}} (\delta^1 \circ f)(x)\; dx = \int_{\mathbb{R}} \sum_{0\leq i<n_0} c_i \delta^1(x - x'_i)\; dx = \sum_{0\leq i<n_0} c_i. \tag{2.3.156}$$

In addition, if the function is chosen carefully such that $\frac{df}{dx}(x'_i) = \pm 1$ at every root $x'_i \in Z$ for $0 \leq i < n_0$, the above sum precisely *counts* the number of roots:

$$\left|\frac{df}{dx}(x'_i)\right| = 1 \quad \Rightarrow \quad \|\delta^1 \circ f\|_1 = \sum_{0\leq i<n_0} \left|\frac{df}{dx}(x'_i)\right|^{-1} = \sum_{0\leq i<n_0} 1 = |Z|. \tag{2.3.157}$$

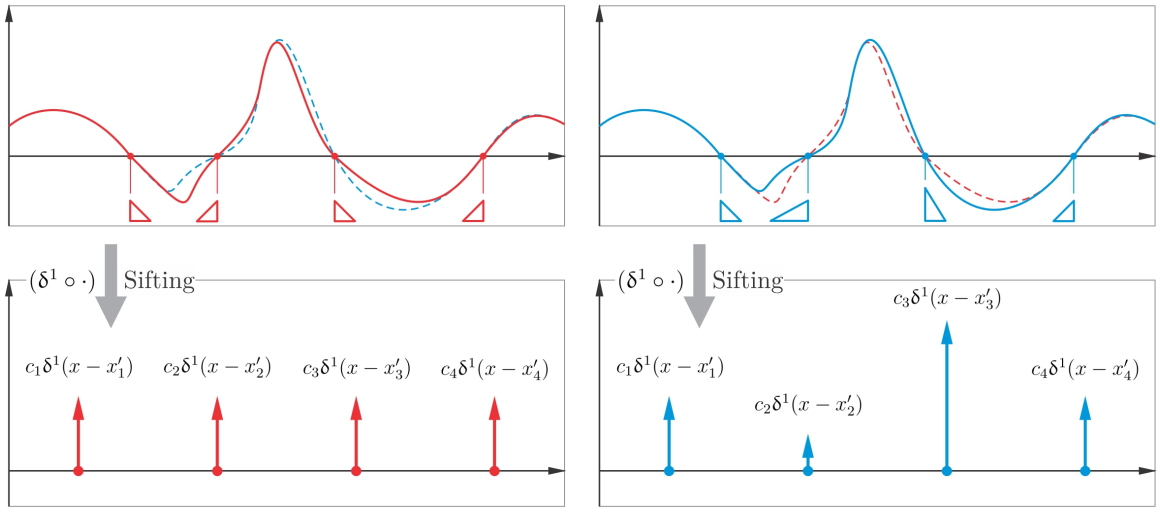

(a) $1 + 1 + 1 + 1 = 4$ (b) $1 + \frac{1}{2} + 2 + 1 = 4\frac{1}{2}$

FIGURE 2.3.16: Given a differentiable function $f : \mathbb{R} \to \mathbb{R}$ with isolated simple roots, the composition $(\delta^1 \circ f)$ puts a $\delta-$singularity at each root $x'_i \in f^{-1}(0)$ with an intensity proportional to the function's derivative at the root (i.e., $c_i := |\frac{df}{dx}(x'_i)|^{-1}$). Integrating the singularities over the real line (i.e., computing its L^1-norm) aggregates the intensities. If $c_i = 1$ for all roots, the integral simply counts the number of roots.

In other words, the composition of an ordinary function with the $\delta-$function is equivalent to assigning $\delta-$singularities to its roots, making them artificially measurable.

Going back to the 3D case in (2.3.152), if the function $f : \mathsf{E} \to \mathbb{R}$, has a non-degenerate $0-$isolevel set Z (i.e., 'isosurface' for $f(\mathbf{x}) = 0$), then $(\delta^1 \circ f)$ distributes $3^{\text{rd}}-$order $\delta-$singularities over that isosurface whose intensities are inversely proportional to the norm of the function's gradient. Integrating this function over the entire $3-$space (i.e., L^1-norm) measures the weighted area of the isosurface with a weight distribution inversely proportional to the gradient over the surface in (2.3.153). In addition, if the function is chosen carefully such that $\|\nabla f(\mathbf{x}')\|_2 = 1$ everywhere over the isosurface, the integral returns the exact $2-$measure (i.e., surface area) of the isosurface in (2.3.154). Importantly, this means that the composition of an ordinary function with the $\delta-$function yields a $1^{\text{st}}-$order $\delta-$singular descriptor of the

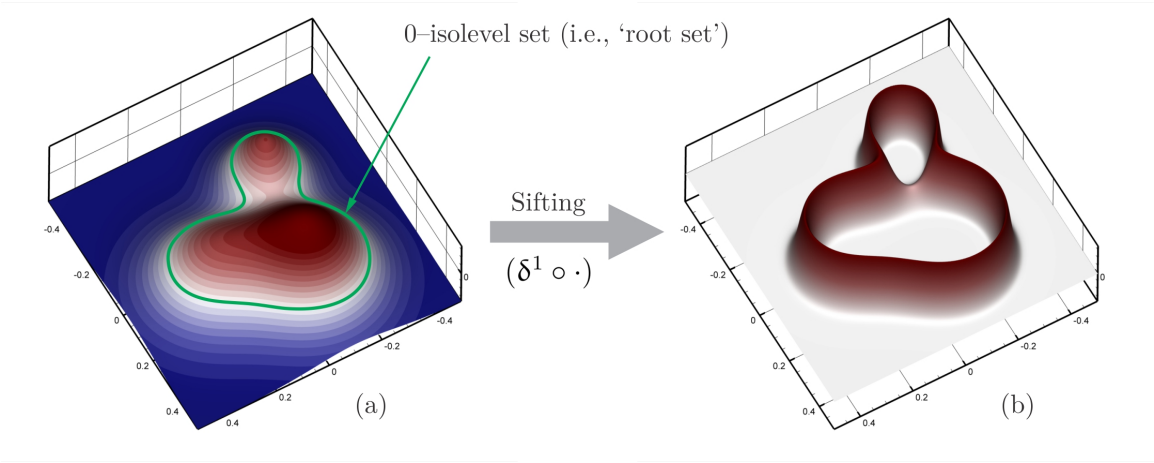

FIGURE 2.3.17: Given a differentiable function $f : \mathbb{R}^2 \to \mathbb{R}$ in (a), its root set is $\delta-$sifted by composition $(\delta^1 \circ f)$, approximated by $(g^1_\sigma \circ f)$ using a Gaussian $g^1_\sigma(x) := \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}(\frac{x}{\sigma})^2}$ with $\sigma > 0$. The $\delta-$singular descriptor of the $0-$isolevel curve is obtained as $\sigma \to 0^+$.

isosurface, making it artificially $3-$measurable.

The essential benefit is that we no longer need to know an a priori stratification and/or a specific parameterization of this surface to distribute $\delta-$singularities; all we need is an implicit description in terms of a function that has a non-degenerate isosurface with nonsingular gradients.

Note that for 3D shapes (e.g., solids), it is not possible to use the root sum lemma with nonnegative real-valued defining functions, since their $0-$level set contains the $3-$measurable set complement, rather than a $2-$measurable surface. However, one can devise different shape descriptors over the $3-$space that are C^1-differentiable, whose $2-$measurable $0-$isolevel set correspond to a collection of 2D surface patches, and whose gradient evaluates to a unit vector over that isosurface. Let us denote the space of such functions with $\mathbb{R}^{\mathsf{E}}_\circ$:

$$\mathbb{R}^{\mathsf{E}}_\circ := \left\{ f \in (\mathbb{R}^{\mathsf{E}} \cap C^1(\mathsf{E})) \mid \mu^2[f^{-1}(0)] < \infty \text{ and } \|\nabla f(f^{-1}(0))\|_2 = 1 \right\}. \quad (2.3.158)$$

Theorem 2.3.23. (δ−Composition) *composition with the δ−function generates uniform 2^{nd}−order δ−singular 0−isosurfaces for every $S \in \mathscr{P}(\mathsf{E})$ and $f_S \in \mathbb{R}_\circ^{\mathsf{E}}$:*

$$(\delta^1 \circ f_S) \equiv \mathbb{1}_{S^2} \quad \text{and} \quad \|\delta^1 \circ f_S\|_1 = \|f_{S^2}\|_1 = \mu^2[S^2], \tag{2.3.159}$$

where $S^2 := f_S^{-1}(0)$ and $\mathbb{1}_{S^2} = \mathsf{indc}^2(S_2) \in (\mathbb{R}_\star^2)^{\mathsf{E}}$ is the δ−singular characteristic function of the 0−isosurface.

In other words, to make a 2D feature $S^2 \in \mathscr{P}^2(\mathsf{E})$ of a shape $S \in \mathscr{P}(\mathsf{E})$ artificially 3−measurable, all we need is to express a shape descriptor $f_S := \mathsf{desc}^\circ(S) \in \mathbb{R}_\circ^{\mathsf{E}}$ that implicitly defines the feature as an isosurface, i.e., $S^2 = f_S^{-1}(0)$. As long as the function satisfies the conditions expressed in (2.3.158), its composition with the 1^{st}−order δ−function distributes a uniform 1^{st}−order δ−singularity over the surface. In other words, the following diagram commutes:

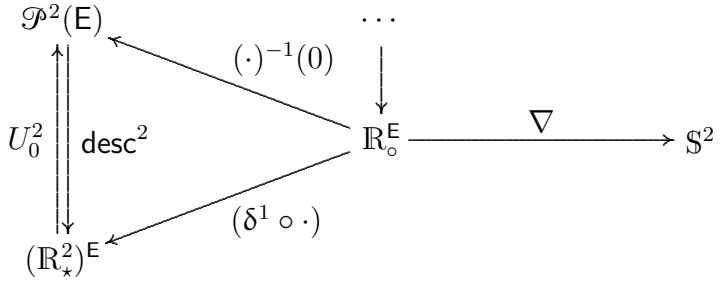

(2.3.160)

in which the vertical downward arrow in the middle implies that (for now) we do not know where the shape descriptor that satisfies (2.3.158) comes from, and $\mathbb{S}^2$ is the unit 2−sphere centered at the origin (i.e., the set of all unit vectors). The fact that unit gradient results in a uniform singularity distribution, whose L^1−norm gives the

exact 2−measure (i.e., surface area) of the 0−isosurface can be shown by:

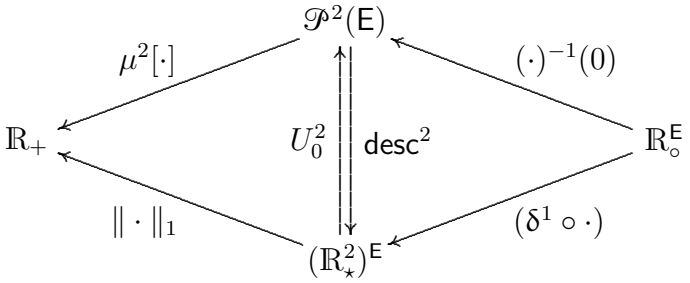

(2.3.161)

In Chapter 3, I will present two important examples of shape descriptors that satisfy (2.3.158) which are useful in computing important constructs implicitly; namely

- the signed distance function $\mathsf{dist} : \mathscr{P}^*(\mathsf{E}) \to \mathbb{R}$ used for boundary evaluation:

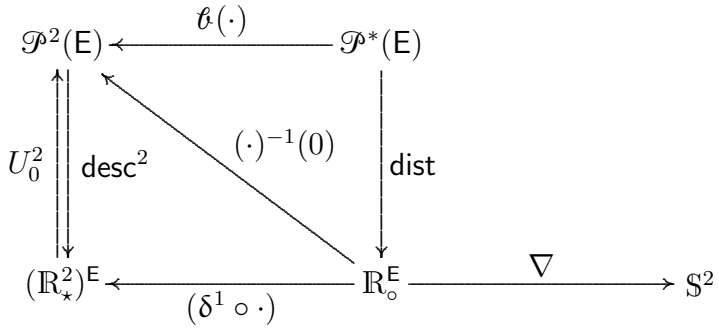

(2.3.162)

- a particular radial basis descriptor $\mathsf{rad} : \mathscr{P}^*(\mathsf{E}) \to \mathbb{R}$ used for skeletonization:

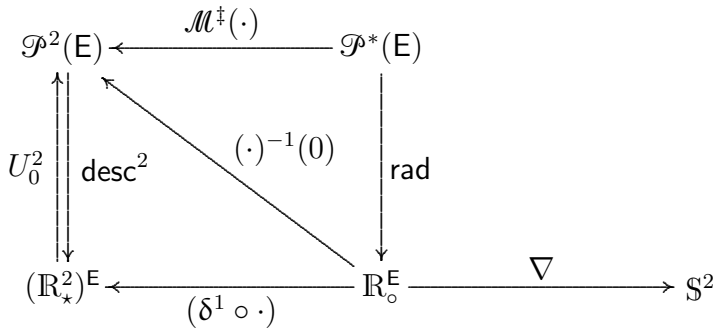

(2.3.163)

To summarize, the root sum lemma and δ−sifting gives the mechanism to link

full-dimensional integration with a *single* Lebesgue measure—e.g., volume integrals using the 3−measure in 3D—to lower-dimensional integration—e.g., surface integrals, curve integrals, and finite summations in 3D. In particular, it enables an analytic (i.e., measure-theoretic) device for generalized *counting*—e.g., counting the number of closest boundary points for implicit characterization of the medial axis (MA).

Once again, the above notions can be generalized (in principle) with some extra care to the 6D configuration space $\mathsf{C} = \mathrm{SE}(3)$ using Haar measures to define the $\delta-$function, and using Lie algebras and exponential map to formulate the differentiation over the Riemannian manifold $\mathrm{SE}(3) = \exp(\mathfrak{se}(3))$.

Dimension Discovery. As mentioned earlier, the set of singular defining functions given in (2.3.73), (2.3.84), and (2.3.95) for separated collections of isolated points, curve segments, and surface patches, respectively, are useful for conversions from explicit-nonsingular to implicit-singular descriptions, assuming that a parameterized stratification for the former is known a priori. This is not always the case with real-world representations of explicit models.[48] Moreover, a true paradigm shift to implicit modeling is not possible until the need for explicit stratification is eliminated.

The first tool that we need is for discovering the local topological properties of the original defining function. At the very least, given a query point $\mathbf{x} \in \mathsf{E}$, one needs to find out the maximum dimension of the strata that are incident to (i.e., pass through or are bounded by) the point from a knowledge of $f_S : \mathsf{E} \to \mathbb{R}_+$ alone. This turns out to be a difficult problem, which cannot be solved using basic measure-theoretic tools presented so far alone. However, $\delta-$regularization (Theorem 2.3.20) makes it

[48] In fact, most existing solid modeling systems use hybrid representations, from which retrieving a parameterized stratification is difficult for objects of complex geometry.

possible to do so if *some* auxiliary information are provided; namely,

- If we know that the set is (closed or open) regular and cusp-free (as per Assumption 2.3.18), $\delta-$sifting provides the answer, i.e., the neighborhood of $\mathbf{x} \in S$ is full if $(f_S * \delta^3)(\mathbf{x}) = f_S(\mathbf{x})$ and partial if $(f_S * \delta^3)(\mathbf{x}) < f_S(\mathbf{x})$.

- If we do not know whether the set is regular or not, $(f_S * \delta^3) = 0$ can mean $0-$, $1-$, or $2-$dimensional neighborhood. There is no general solution around this using a single Lebesgue measure, until some additional knowledge of the local differential properties of the $0-$superlevel set is made explicit.

The good news is that in almost all practical applications dealing with irregularities, the $0-, 1-$, and $2-$dimensional subsets of shapes and/or motions that one deals with are either specified explicitly—e.g., when computing the sweep or unsweep of a body under a given one- or two-parametric motion—or are produced from operations on regular sets—e.g., when computing the boundary or skeleton of a body or Minkowski differences of bodies in low-DOF contact. For the first case, the parametric integrals (2.3.73), (2.3.84), and (2.3.95) are sufficient to produce implicit singularities from an explicit knowledge of lower-dimensional features. For the rest of the problems, I shall present different techniques on a case-by-case basis in Chapter 5. Once the $\delta-$functions are populated, analytic methods can be used in a variety of downstream operations to transform, combine, or compute properties of shapes and motions.

Regularized Indicator Retrieval. Following up on the discussion in Section 2.3.1, I argue that the $\delta-$function provides a filter for the retrieval of the regularized shape's indicator function from the irregularity-tolerant description scheme. Earlier,

$(\mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S] = \mathbf{1}_S$ was presented as a tool in (2.3.42) to retrieve the unique indicator function from the level set equivalence class $[\mathbf{1}_S]$. It was also discussed that $(\mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S]^* \overset{*}{\equiv} \mathbf{1}_S$ returns a subclass of indicator functions that are equal$-$ae (i.e., the left-most column in Fig. 2.3.3) from which one indicator function can be selected arbitrarily (e.g., randomly) as $\mathbf{1}_S := (\mathsf{choose} \circ \mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S]^*$ due to a lack of proper rules to choose one over another. Now it is clear that the sifting property of the $\delta-$function in (2.3.147) can be used to retrieve the unique indicator function of the regularized set $\mathbf{1}_{r(S)} \in [\mathbf{1}_S]^*$ from the class $[\mathbf{1}_S]^* \in \mathbb{R}^{\mathsf{E}}_+$ in a deterministic fashion:

$$\mathbf{1}_{r(S)} = \mathsf{choose}^*[\mathbf{1}_S]^* := \big((\mathsf{choose} \circ \mathsf{sign} \circ \mathsf{choose})[\mathbf{1}_S]^* * \delta^3\big), \tag{2.3.164}$$

This completes the diagrams in (2.3.41) as:

$$\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cup^*} & \mathscr{P}^*(\mathsf{E}) \\
\mathsf{supp} \uparrow\downarrow \mathsf{indc} & & \mathsf{supp} \uparrow\downarrow \mathsf{indc} \\
2^{\mathsf{E}} \times 2^{\mathsf{E}} & & 2^{\mathsf{E}} \\
\mathsf{sign} \uparrow\downarrow \mathsf{emb} & & \mathsf{sign} \uparrow\downarrow \mathsf{emb} \\
\mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & & \mathbb{R}^{\mathsf{E}}_+ \\
\mathsf{choose} \uparrow\downarrow [\cdot]^* & & \mathsf{choose}^* \uparrow\downarrow [\cdot]^* \\
(\mathbb{R}^{\mathsf{E}}_+/\overset{*}{\equiv})^2 & \xrightarrow{+} & \mathbb{R}^{\mathsf{E}}_+/\overset{*}{\equiv}
\end{array}
\qquad
\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\cap} & \mathscr{P}^*(\mathsf{E}) \\
\mathsf{supp} \uparrow\downarrow \mathsf{indc} & & \mathsf{supp} \uparrow\downarrow \mathsf{indc} \\
2^{\mathsf{E}} \times 2^{\mathsf{E}} & & 2^{\mathsf{E}} \\
\mathsf{sign} \uparrow\downarrow \mathsf{emb} & & \mathsf{sign} \uparrow\downarrow \mathsf{emb} \\
\mathbb{R}^{\mathsf{E}}_+ \times \mathbb{R}^{\mathsf{E}}_+ & & \mathbb{R}^{\mathsf{E}}_+ \\
\mathsf{choose}^* \uparrow\downarrow [\cdot]^* & & \mathsf{choose} \uparrow\downarrow [\cdot]^* \\
(\mathbb{R}^{\mathsf{E}}_+/\overset{*}{\equiv})^2 & \xrightarrow{\bullet} & \mathbb{R}^{\mathsf{E}}_+/\overset{*}{\equiv}
\end{array} \tag{2.3.165}$$

Finally, the application of Dirac $\delta-$claculus to shape descriptors can be extended in principle to motion descriptors. The most important distinction to note is that local parameterization of lower-dimensional motion strata in the configuration space—e.g.,

curves and surfaces embedded in the Lie group $\mathsf{C} = \mathrm{SE}(3)$ to model one- and two-parametric sweeps of a shape—needed for defining singularities of the forms (2.3.84) or (2.3.95) are made possible by the Lie algebra $\mathfrak{se}(3) \cong (\mathfrak{so}(3) \oplus \mathbb{R}^3) \cong \mathbb{R}^6$, which is a 6D vector space, and the exponential map $\exp|_{\mathfrak{se}(3)} : \mathfrak{se}(3) \to \mathrm{SE}(3)$.

Chapter 3

Geometric Operations

> "And if you open your eyes again, you wouldn't know that they'd moved. But it's the motion that really characterizes the symmetry inside the Alhambra. But it's also about producing a language to describe this, and the power of mathematics is often to change one thing into another, to change *geometry* into *language*."
>
> Marcus du Sautoy, 2009 [318]

3.1 The Analytic 'Toolbox'

In Chapter 2 I presented a collection of measure-theoretic tools that can be viewed as the 'building blocks' of an analytic approach to geometric modeling and spatial reasoning. The purpose of this chapter is to apply those tools to a variety of important problems that are seemingly very different, but can be unified by systematic composition of a handful of fundamental operations presented on basic 'queries' (modeled as functions) on geometric entities (i.e., shapes and motions). Although it is not pos-

sible to cover every important geometric computation, this chapter will demonstrate a methodology to extract commonalities between shape and motion related problems in terms of the basic computations involved, and to observe the general patterns of transition from explicit to implicit techniques.

Each section will have two subsections, one for the explicit methods that are more popular in the traditional set-theoretic perspective, and one for the implicit methods to apply the alternative measure-theoretic techniques. For the former, high-level procedural reviews of the existing "recipes" for each problem class will be given in enumerated form (step 1, step 2, $\cdots$). For the latter, the measure-theoretic building blocks developed in Chapter 2 and summarized in Table 3.1.1 will be composed to construct new operations. These operations include:

1. Computing integral properties of shapes, which amounts to computing integrals of scalar, vector, or (more generally) tensor functions over a geometric domain. This is the subject of Section 3.2.

2. Computing topological operations such as boundary evaluation and regularization, which amount to explicitly mapping a shape to its topological boundary or regular counterpart, or to implicitly mapping the functional shape descriptors. This is the subject of Section 3.3.

3. Computing morphological operations across shapes and motions, which include Minkowski sum and difference, Minkowski product and quotient, dilation and erosion, sweep and unsweep, closing and opening, configuration space obstacle generation, and their special cases that give offsetting, blending, grouping, etc. This is the subject of Section 3.4.

4. Skeletonization, which amounts to explicitly mapping a shape to its medial axis (MA) or skeleton, or to implicitly mapping the functional shape descriptors. This is the subject of Section 3.5.

In each case, the shapes and/or motions are modeled and represented either explicitly as a set or implicitly as one or more functions, which abstract the *ability to respond* to basic queries. I will show that for the above purposes, being able to answer two basic queries will be sufficient; namely:

- The *set inclusion queries* that are modeled as indicator (i.e., characteristic) functions $\mathbf{1}_S : \mathsf{E} \to \{0, 1\}$ defined as

$$\mathbf{1}_S := \mathsf{indc}(S) \in 2^{\mathsf{E}}, \quad \text{where} \quad \mathbf{1}_S(\mathbf{x}) = (\mathbf{x} \overset{?}{\in} S) = \begin{cases} 1 & \text{if } \mathbf{x} \in S, \\ 0 & \text{if } \mathbf{x} \notin S. \end{cases} \tag{3.1.1}$$

- The *boundary distance queries* that are modeled as signed infimum distance functions $\mathbb{d}_S : \mathsf{E} \to \mathbb{R}$ defined as

$$\mathbb{d}_S := \mathsf{dist}(S) \in \mathbb{R}^{\mathsf{E}}, \quad \text{where} \quad \mathbb{d}_S(\mathbf{x}) = \big[1 - 2(\mathbf{1}_S(\mathbf{x}))\big] \inf_{\mathbf{x}' \in \partial(S)} \|\mathbf{x} - \mathbf{x}'\|_2, \tag{3.1.2}$$

Note that the infimum distance computation in (3.1.2) is independent (i.e., cannot be constructed implicitly) from an ability to query inclusion modeled via characteristic (or other level set defining) functions, and requires explicit enumeration of the boundary. Therefore, there is no other choice than taking it as a separate query that is made available a priori by an explicit representation.[1] Once the two implicit descriptors

[1] It is worthwhile noting that the signed distance function is obtained from the unsigned distance function $|\mathbb{d}_S(\mathbf{x})| = \inf_{\mathbf{x}' \in \partial S} \|\mathbf{x} - \mathbf{x}'\|_2$ and the indicator function by adjusting the sign by multiplying

$\mathsf{desc}_1 := \mathsf{indc}$ and $\mathsf{desc}_2 := \mathsf{dist}$ are taken for granted, the downstream operations—resulting in implicit descriptors of new shapes and motions—will be composed from them using the following analytic operations, without a need for appealing to explicit information again:

- *arithmetics* of real-valued functions, including extended-real-valued defining functions (subsuming indicator functions) and distance functions;
- *convolutions* including convolutions among defining functions themselves and between defining functions and $\delta-$function (i.e., $\delta-$regularization); and
- *compositions* with the $\delta-$function (i.e., $\delta-$sifting).

In other words, the two fundamental queries span the space of all important geometric computations via implicit descriptions—i.e., viewed as queries about resulting shapes and motions—through the algebraic application of the above operations and guided by the set of rules prescribed by Table 3.1.1 to tie the them back to explicit descriptions.

Interestingly, inclusion queries equipped with arithmetics and convolution algebra alone suffice for integral and morphological computations. However, when addressing more difficult problems such as boundary and skeleton evaluations, the need for distance queries and Dirac $\delta-$sifting will become self-evident for transcending measure-theoretic limitations when dealing with features of heterogeneous dimensions.

Most of the concepts developed for shapes in the Euclidean space $\mathsf{E} = \mathbb{R}^3$ equipped with Euclidean distance (based on Euclidean L^2-norm) and Lebesgue 3$-$measure can be extended to the configuration space $\mathsf{C} = \mathrm{SE}(3)$ equipped with Riemannian metric

the coefficient $\left[1 - 2(\mathbf{1}_S(\mathbf{x}))\right] = \pm 1$. Thus it seems more appropriate to take $|\mathbb{d}_S|$ and $\mathbf{1}_S$ as the independent pair of fundamental queries, and view $\mathbb{d}_S$ as a dependent query computed from the knowledge of both. However, I choose the signed distance function as a basic query, without loss of generality, to simplify the notations for resulting implicit compositions.

TABLE 3.1.1: The analytic "toolbox"—Most geometric modeling and spatial reasoning operations can be implicitly formulated by composing inclusion and distance queries via arithmetics, convolutions, and $\delta-$sifting, and by leveraging the following relationships:

Rule Diagram	Rule Nickname	Rule Description
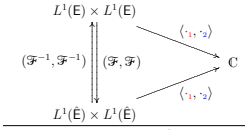	Parseval's Theorem (2.2.13)	• Fourier transforms preserve inner product structure. • Allows for efficient low-pass filtered approximation.
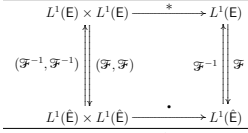	Convolution Theorem (2.2.15)	Fourier transforms convert convolution to multiplication. Allows for efficient FFT-based cumulative computations. Allows for efficient low-pass filtered approximation.
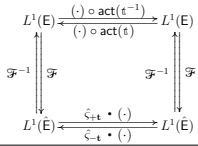	Translation Multiplier (2.2.17)	Fourier transforms convert translation to multiplier. Allows for efficient incorporation of translational motion.
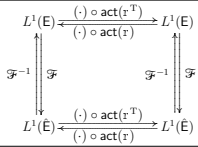	Rotation Invariance (2.2.19)	Fourier transforms are commutative with rotations. Allows for efficient incorporation of rotational motion.
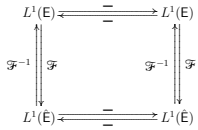	Reflection Invariance (2.2.20)	Fourier transforms are commutative with reflections. Allows for efficient incorporation of reflection/inversion.
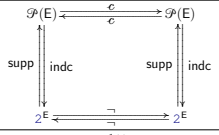	Implicit Complement (2.3.5)	Set complement corresponds to indicator function negation. Allows for implicit formulation of De Morgan dualities.
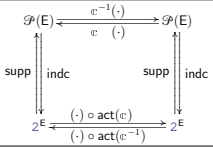	Motion Covariance (2.3.7)	Indicator functions are covariant under rigid motions. Allows for implicit formulation in terms of inverse motions.
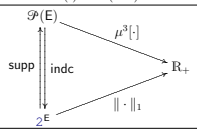	3−Measure via L^1-Norm (2.3.9)	3−Measures of sets are L^1-norms of indicator functions. Allows for implicit formulation in terms of full-D integrals.
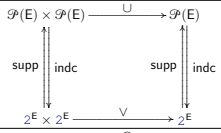	Set-Theoretic Union (2.3.12)	Set union corresponds to indicator function disjunction. Allows for implicit Boolean algebra via binary logic.
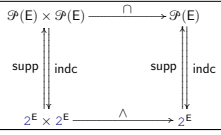	Set-Theoretic Intersection (2.3.12)	Set Intersection corresponds to indicator function conjunction. Allows for implicit Boolean algebra via binary logic.

Rule Diagram	Rule Nickname	Rule Description
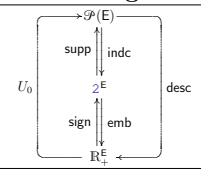	Defining Functions (2.3.15)	Nonnegative real-valued functions to replace binary indicator. Allows extending pure logic-based algebra to $\mathbb{R}_+-$algebra. Allows a profile-oblivious view of implicit methods. Allows a mechanism to model infinite unions as integrals.
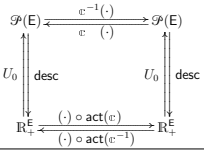	Motion Covariance (2.3.18)	Defining functions are covariant under rigid motions. Allows for implicit formulation in terms of inverse motions.
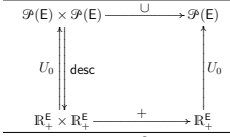	Set-Theoretic Union (2.3.20)	Set union is implicitly formulated via summation. Allows for implicit Boolean algebra via $\mathbb{R}_+-$arithmetics.
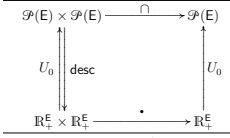	Set-Theoretic Intersection (2.3.20)	Set Intersection is implicitly formulated via multiplication. Allows for implicit Boolean algebra via $\mathbb{R}_+-$arithmetics.
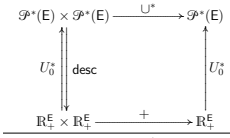	Regularized Union (2.3.22)	Set union is implicitly formulated via summation. Allows for implicit Boolean algebra via $\mathbb{R}_+-$arithmetics.
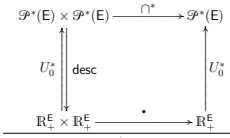	Regularized Intersection (2.3.22)	Set Intersection is implicitly formulated via multiplication. Allows for implicit Boolean algebra via $\mathbb{R}_+-$arithmetics.
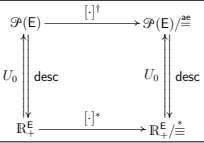	Implicit Regularization (2.3.30)	Regularization viewed as equivalence via equality$-\mathsf{ae}$. Allows for implicit measure-theoretic regularization.
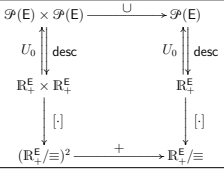	Set-Theoretic Union (2.3.36)	$\mathbb{R}_+-$summation is extended to equivalence via $\equiv$. Allows concept of profile-oblivious Boolean algebra.
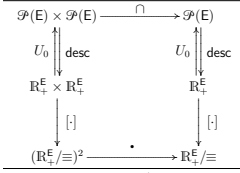	Set-Theoretic Intersection (2.3.36)	$\mathbb{R}_+-$multiplication is extended to equivalence via $\equiv$. Allows concept of profile-oblivious Boolean algebra.
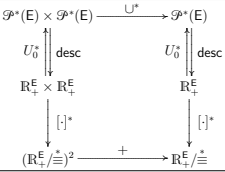	Regularized Union (2.3.41)	$\mathbb{R}_+-$summation is extended to equivalence via $\overset{*}{\equiv}$. Allows concept of irregularity-tolerant Boolean algebra.
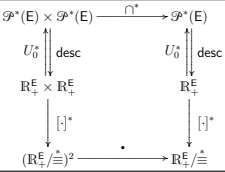	Regularized Intersection (2.3.41)	$\mathbb{R}_+-$multiplication is extended to equivalence via $\overset{*}{\equiv}$. Allows concept of irregularity-tolerant Boolean algebra.

Rule Diagram	Rule Nickname	Rule Description
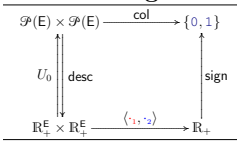	3−Collision Predicate (2.3.49)	Inner product of descriptors gives a collision predicate. Allows interference test between general sets.
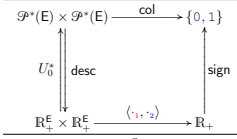	3−Collision Predicate (2.3.49)	Inner product of descriptors gives a collision predicate. Allows interference test between regular sets.
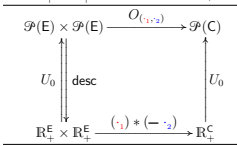	C−Space 3−Obstacle (2.3.54)	Convolution of descriptors is a 3−obstacle descriptor. Allows cumulative interference test between general sets.
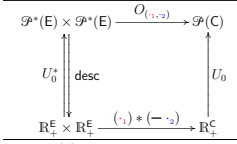	C−Space 3−Obstacle (2.3.54)	Convolution of descriptors is a 3−obstacle descriptor. Allows cumulative interference test between regular sets.
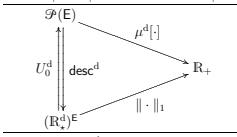	d−Measure via L^1−Norm (2.3.121)	d−Measures of sets are L^1−norms of δ−singular descriptors. Allows for implicit formulation in terms of lower-D integrals.
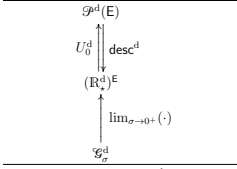	Limit Representation (2.3.125)	δ−functions can be viewed in terms of limit representations. Allows for construction of computable descriptors. Allows for a resolution/precision trade-off mechanism.
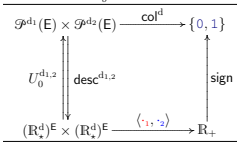	d−Overlap Predicate (2.3.136)	Inner product of descriptors give a d−overlap predicate. Allows d−overlap test between general sets.
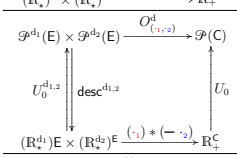	C−Space d−Obstacle (2.3.140)	Convolution of descriptors is a d−obstacle descriptor. Allows cumulative d−overlap test between general sets.
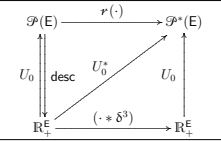	δ−Regularization (2.3.149)	δ−Convolution of a descriptor regularizes its support. Allows topological regularization on implicit descriptions. Allows a natural mechanism for mollification into MDF.
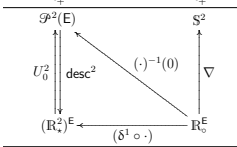	δ−Sifting (2.3.160)	δ−Composition of a descriptor is a δ−singular isosurface. Allows expressing 2D integrals via 3D L^1−norms. Allows a device to transcend single-measure limits.

(based on Frobenius L^2-norm) and Haar 6$-$measure. In particular, some of the most important extensions in the context of group morphology are discussed in Section 3.4 along with fundamental lifting/projection maps that removes the artificial distinction between shapes and motions [215, 218, 235].

3.2 Integral Functions

Let us start with the problems of computing integral properties for shapes restricted to the shape class of semianalytic sets that are guaranteed to be stratifiable into integrable analytic pieces of various dimensions—including points, curve segments, surface patches, and solid volumes parameterized with an atlas of curvilinear coordinate charts that map them to points, line segments, polygonal patches patches, and polyhedral volumes in the parametric domain, respectively.

3.2.1 Explicit Integration

3.2.1.1 Pointset-Theoretic Perspective for Evaluating Shape Integrals:

Among the two popular approaches to numerical integration are dimensional reduction and sampling. The former is reliant on Gauss-Ostrogradsky (i.e., divergence) and generalized Stoke's theorems to reformulate an integral over a higher-dimensional domain to one over its lower-dimensional boundary by exploiting the integrand's continuity and conservation properties (if applicable) [212]. Sooner or later, numerical integration reduces to breaking the domain apart into relatively simpler pieces—e.g., cubes, tetrahedra, etc.

Integration of L^1-Functions over Analytic Strata

Given a tensor-valued function $g : \mathsf{E} \to \mathbb{Q}$, the problem is to compute its integral restricted to fully specified analytic spatial elements, e.g.,

- Over a finite set of isolated points $S^0 := \{\mathbf{x}_i\}_{0\leq i\leq n_0}$ with $|S^0| = n_0$:

$$\mathcal{I}^0(g; S^0) := \sum_{\mathbf{x}_i\in S^0} g(\mathbf{x}_i) = \sum_{0\leq i<n_0} g(\mathbf{x}_i). \tag{3.2.1}$$

- Over a finite union of disjoint $C^\omega-$curve segments $S^1 := \bigcup_{0\leq i\leq n_1} \Gamma_i(\Delta^1)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^1 \subseteq \mathbb{R}$ is an open 1−cell:

$$\begin{aligned}\mathcal{I}^1(g; S^1) &:= \int_{S^1} g(\mathbf{x})\ d\mu^1[\mathbf{x}] = \sum_{0\leq i<n_1} \int_{\Gamma_i(\Delta^1)} g(\mathbf{x})\ d\mu^1[\mathbf{x}] \\ &= \sum_{0\leq i<n_1} \int_{\Delta^1} (g\circ\Gamma_i)(t_1, 0, 0)\ |\mathfrak{J}_i''(t_1, 0, 0)|\ dt_1,\end{aligned} \tag{3.2.2}$$

  where $|\mathfrak{J}_i''(\mathbf{t})|$ is the Jacobian minor defined in (2.3.70).

- Over a finite union of disjoint $C^\omega-$surface patches $S^2 := \bigcup_{0\leq i\leq n_2} \Gamma_i(\Delta^2)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^2 \subseteq \mathbb{R}^2$ is an open 2−cell:

$$\begin{aligned}\mathcal{I}^2(g; S^2) &:= \int_{S^2} g(\mathbf{x})\ d\mu^2[\mathbf{x}] = \sum_{0\leq i<n_2} \int_{\Gamma_i(\Delta^2)} g(\mathbf{x})\ d\mu^2[\mathbf{x}] \\ &= \sum_{0\leq i<n_2} \iint_{\Delta^2} (g\circ\Gamma_i)(t_1, t_2, 0)\ |\mathfrak{J}_i'(t_1, t_2, 0)|\ dt_1 dt_2,\end{aligned} \tag{3.2.3}$$

  where $|\mathfrak{J}_i'(\mathbf{t})|$ is the Jacobian minor defined in (2.3.70).

- Over a finite union of disjoint $C^\omega-$solid volumes $S^3 := \bigcup_{0\leq i\leq n_3} \Gamma_i(\Delta^3)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^3 \subseteq \mathbb{R}^3$ is an open 3−cell:

$$\begin{aligned}\mathcal{I}^3(g; S^3) &:= \int_{S_3} g(\mathbf{x})\ d\mu^3[\mathbf{x}] = \sum_{0\leq i<n_3} \int_{\Gamma_i(\Delta^3)} g(\mathbf{x})\ d\mu^3[\mathbf{x}] \\ &= \sum_{0\leq i<n_3} \iiint_{\Delta^3} (g\circ\Gamma_i)(t_1, t_2, t_3)\ |\mathfrak{J}_i(t_1, t_2, t_3)|\ dt_1 dt_2 dt_3.\end{aligned} \tag{3.2.4}$$

  where $|\mathfrak{J}_i(\mathbf{t})|$ is the Jacobian determinant defined in (2.3.69).

The notations follow those of Convention 2.3.11. Note that in each case, the parameterization $\Gamma_i = (\gamma_{i,1}, \gamma_{i,2}, \gamma_{i,3}) : \mathbb{R}^3 \to \mathsf{E}$ describes a conformal curvilinear coordinate system due to Assumption 2.3.10 even though for integrals over lower-dimensional strata, only the first d components of $\mathbf{t} = (t_1, t_2, t_3) \in \mathbb{R}^3$ are used to parameterize the geometry, where $0 \leq \mathrm{d} < 4$ in (3.2.1) through (3.2.4), respectively. The integrand can be a tensor function in general, which can be integrated componentwise, i.e., by integrating a finite number of real- or scalar-valued components. To make the notations simpler, I assume an extension of the function's domain from S^{d} to the entire 3−space by zeropadding, without loss of generality.

Explicit methods exist for exact evaluation of polynomial functions over linear polyhedral cells (within the machine precision) [39, 58, 206, 226]. These methods typically follow a set of steps that look like these:[2]

1. breaking the integral down into a sum of integrals over simpler subdomains (in shape and/or dimensionality) using conforming or non-conforming and often hierarchical domain decompositions;

2. transforming the integrals on subdomains to integrals over even simple shapes, if possible (e.g., parameterized simplices) by algebraic manipulations and differential transformations (i.e., using Jacobian determinants);

3. using explicit algebraic formulae or well-established cubature rules [77, 78, 335] to compute the simplicial integrals;[3] and

[2]The author is thankful to Vaidyanathan Thiagarajan and Vadim Shapiro [341] for providing the herein cited references to integration methods. That being said, the responsibility of potential misrepresentation or oversimplification of the literature lies with the author.

[3]See nines.cs.kuleuven.be/ecf for an encyclopedia of cubature formulae.

4. adding them up to obtain the aggregate result.

However, for more general functions and/or cells for engineering applications, the evaluation is performed only approximately based on the function values at carefully chosen (e.g. Gauss) quadrature points in the cell, hierarchical or geometrical adaptive methods, etc. [211, 340]. These sampling and decomposition techniques can be classified under implicit approaches, which are the subject of Section 3.2.2.1.

The analogues of the above Lebesgue integrals over shapes (i.e., sets of points) can be formulated in principle as Haar integrals over motions (i.e., sets of configurations) of $\mathrm{d} = 0, 1, \cdots, 6$, representing kinematic constraints allowing $\mathrm{d} \leq 6$ degrees of freedom (DOF) for a given object. However, in practice, one is predominantly concerned with integrating differential equations over some subspace of the 3D physical space (or 4D space-time, at most) for solving boundary value problems, and the need to integrate over the 6D configuration space rarely occurs—except along lower-dimensional (e.g., 1D or 2D) motion trajectories.

3.2.2 Implicit Integration

3.2.2.1 Measure-Theoretic Perspective for Evaluating Shape Integrals:

Advocating for implicit methods for integration requires the least of efforts, as most effective techniques that apply to arbitrarily complex shapes and distribution functions are intrinsically implicit. As explained in Section 3.2.1.1, except for the simplest shapes and integrands—e.g., polynomials integrated over polyhedra [39,58,206,226]—no explicit method can outperform implicit sampling-based techniques for general semianalytic shapes and integrable functions.

In the following, I show how the integrals in (3.2.1) through (3.2.4) can be expressed in measure-theoretic language in terms of (an ability to compute) inner products between the integrand and (and ability to respond to) weighted inclusion queries about geometry modeled as $0-$superlevel nonnegative extended-real-valued defining functions of the form $f_S : \mathsf{E} \to \mathbb{R}_\star$.

Finite Summations. Let us start with computing the finite sum of a given tensor-valued function $g : \mathsf{E} \to \mathbb{Q}$ over the finite set of points $S^0 := \{\mathbf{x}_i\}_{0 \leq i \leq n_0}$ with $|S^0| = n_0$ in (3.2.1), which can be rearranged as

$$\mathscr{I}^0(g; S^0) = \sum_{0 \leq i < n_0} g(\mathbf{x}_i) = \int_{\mathsf{E}} \sum_{0 \leq i < n_0} \delta^3(\mathbf{x} - \mathbf{x}_i) g(\mathbf{x}_i) \; d\mu^3[\mathbf{x}], \tag{3.2.5}$$

as a direct result of the definition of the Dirac $\delta-$function in (2.3.66). Comparing (3.2.5) with the defining function $f^3_{S^0} : \mathsf{E} \to \mathbb{R}^0_\star$ presented in (2.3.73) yields:

$$\mathscr{I}^0(g; S^0) = \langle \mathbb{1}^3_{S^0}, g \rangle = \| \mathbb{1}^3_{S^0} \bullet g \|_1, \quad \text{where} \quad c_i := 1 \; (0 \leq i < n_0), \tag{3.2.6}$$

i.e., if the weight coefficients are distributed uniformly among all points. However, one can choose to assign different real-valued weights or integer-valued 'multiplicities' to the points—e.g., if each point represents multiple distinct physical singularities whose positions are coincident at a moment of time. If that is the case, choosing different $c_i \neq 1$ allows to account for this multiplicity when integrating an 'extensive' property function (i.e., a quantity that is defined "per [physical] point"):

$$\mathscr{I}^{0'}(g; S^0) := \langle f^3_{S^0}, g \rangle = \| f^3_{S^0} \bullet g \|_1, \quad \text{where} \quad c_i > 0 \; (0 \leq i < n_0), \tag{3.2.7}$$

Curve Integrals. Given a tensor-valued function $g : \mathsf{E} \to \mathbb{Q}$ restricted to a finite union of disjoint $C^{\omega}-$curve segments $S^1 := \bigcup_{0\leq i\leq n_1} \Gamma_i(\Delta^1)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^1 \subseteq \mathbb{R}$ is an open $1-$cell, in accordance with Convention 2.3.11, the integral in (3.2.2) can be rearranged as

$$\mathcal{I}^1(g; S^1) = \sum_{0\leq i<n_1} \int_{\mathbb{R}} \mathbf{1}_{\Delta^1}(t_1)\ (g \circ \Gamma_i)(t_1, 0, 0)\ |\mathfrak{J}_i''(t_1, 0, 0)|\ dt_1, \tag{3.2.8}$$

where $\mathbf{t} := (t_1, t_2, t_3) \in (\mathbb{R} \times \mathbb{R}^2) \cong \mathsf{E}$ are coordinates in the parametric space and $\mathbf{1}_{\Delta^1} : \mathbb{R} \to \{0, 1\}$ is the indicator function of the $1-$cell. The single (i.e., 1D) integral can be converted to a triple (i.e., 3D) integral by using a $2^{\text{nd}}-$order $\delta-$function in the parametric domain, i.e., by noting that

$$(g \circ \Gamma_i)(t_1, 0, 0)\ |\mathfrak{J}_i''(t_1, 0, 0)| = \iint_{\mathbb{R}^2} \delta^2(t_2, t_3)\ (g \circ \Gamma_i)(\mathbf{t})\ |\mathfrak{J}_i''(\mathbf{t})|\ dt_2 dt_3, \tag{3.2.9}$$

where the function $(t_2, t_3) \mapsto (g \circ \Gamma_i)(t_1, t_2, t_3)\ |\mathfrak{J}_i''(t_1, t_2, t_3)|$ was used to apply Lemma 2.3.16 with $\mathrm{d} := 2$. Substituting the latter into (3.2.8) yields

$$\mathcal{I}^1(g; S^1) = \sum_{0\leq i<n_1} \int_{\mathbb{R}} \iint_{\mathbb{R}^2} \mathbf{1}_{\Delta^1}(t_1)\ \delta^2(t_2, t_3)\ (g \circ \Gamma_i)(\mathbf{t})\ |\mathfrak{J}_i''(\mathbf{t})|\ dt_1\ dt_2 dt_3. \tag{3.2.10}$$

The above Riemann integral can be converted to a Lebesgue integral as

$$\mathcal{I}^1(g; S^1) = \int_{\mathsf{E}} \sum_{0\leq i<n_1} \mathbf{1}_{\Delta^1}(t_1)\ \delta^2(t_2, t_3)\ g(\mathbf{x})\ |\mathfrak{J}_i''(\mathbf{t})|\ |\mathfrak{J}_i(\mathbf{t})|^{-1} \Big|_{\mathbf{t}:=\Gamma_i^{-1}(\mathbf{x})} d\mu^3[\mathbf{x}]. \tag{3.2.11}$$

Now let the defining function $f^2_{S^1} : \mathsf{E} \to \mathbb{R}^1_\star$ of $S^1 \in \mathscr{P}^1(\mathsf{E})$ in (2.3.84) be

$$f^2_{S^1}(\mathbf{x}) := \sum_{0 \le i < n_1} \mathbf{1}_{\Delta^1}(t_1)\ \delta^2(t_2, t_3)\ |\mathfrak{J}''_i(\mathbf{t})|\ |\mathfrak{J}_i(\mathbf{t})|^{-1} \Big|_{\mathbf{t} := \Gamma_i^{-1}(\mathbf{x})}. \tag{3.2.12}$$

Comparing the above definition with (2.3.84) and noting that the function vanishes at all $(t_2, t_3) \neq (0, 0)$ thus $\delta^2(t_2, t_3)\ |\mathfrak{J}_i(t_1, t_2, t_3)|^{-1} = \delta^2(t_2, t_3)\ |\mathfrak{J}_i(t_1, 0, 0)|^{-1}$, and that $\mathbf{1}_{\Delta^1}(t_1) = \mathbf{1}_{\Delta^1 \times \mathbb{R}^2}(t_1, t_2, t_3)$, the curve density function is obtained as

$$\rho_i(t_1) := \mathbf{1}_{\Delta^1}(t_1)\ |\mathfrak{J}''_i(t_1, 0, 0)| = \big(\mathbf{1}_{\Delta^1 \times \mathbb{R}^2}\ \bullet\ |\mathfrak{J}''_i|\big) \circ \Gamma_i^{-1}(\mathbf{x})\ \Big|_{\mathbf{x} := \Gamma_i(t_1, 0, 0)}, \tag{3.2.13}$$

from which the total artificial volume of each curve segment is obtained as

$$c_i = \|\rho_i\|_1 = \int_{\mathbb{R}} \rho_i(t_1)\ dt_1 = \int_{\Delta^1} |\mathfrak{J}''_i(\cdot, 0, 0)|\ dt_1 = \langle \mathbf{1}_{\Delta^1}, |\mathfrak{J}''_i(\cdot, 0, 0)| \rangle. \tag{3.2.14}$$

Substituting for the density function from (3.2.13) back to (3.2.11) yields

$$\mathscr{I}^1(g; S^1) = \langle \mathbb{1}^2_{S^1}, g \rangle = \|\mathbb{1}^2_{S^1}\ \bullet\ g\|_1, \quad \text{where} \quad c_i := 1\ (0 \le i < n_1), \tag{3.2.15}$$

However, one can choose to assign different density distributions when integrating an 'extensive' property function that needs to be scaled with the artificial "heaviness" assigned to each infinitesimal curve element:

$$\mathscr{I}^1(g; S^1) := \langle f^2_{S^1}, g \rangle = \|f^2_{S^1}\ \bullet\ g\|_1, \quad \text{where} \quad c_i > 0\ (0 \le i < n_1). \tag{3.2.16}$$

Surface Integrals. Given a tensor-valued function $g : \mathsf{E} \to \mathbb{Q}$ restricted to a finite union of disjoint $C^{\omega}-$surface patches $S^2 := \bigcup_{0\leq i\leq n_2} \Gamma_i(\Delta^2)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^2 \subseteq \mathbb{R}^2$ is an open 2−cell, in accordance with Convention 2.3.11, the integral in (3.2.3) can be rearranged as

$$\mathscr{I}^2(g; S^2) = \sum_{0\leq i<n_2} \iint_{\mathbb{R}^2} \mathbf{1}_{\Delta^2}(t_1, t_2)\ (g \circ \Gamma_i)(t_1, t_2, 0)\ |\mathfrak{J}_i'(t_1, t_2, 0)|\ dt_1 dt_2, \tag{3.2.17}$$

where $\mathbf{t} := (t_1, t_2, t_3) \in (\mathbb{R}^2 \times \mathbb{R}) \cong \mathsf{E}$ are coordinates in the parametric space and $\mathbf{1}_{\Delta^2} : \mathbb{R}^2 \to \{0, 1\}$ is the indicator function of the 2−cell. The double (i.e., 2D) integral can be converted to a triple (i.e., 3D) integral by using a $1^{\text{st}}-$order $\delta-$function in the parametric domain, i.e., by noting that

$$(g \circ \Gamma_i)(t_1, t_2, 0)\ |\mathfrak{J}_i'(t_1, t_2, 0)| = \int_{\mathbb{R}} \delta^1(t_3)\ (g \circ \Gamma_i)(\mathbf{t})\ |\mathfrak{J}_i'(\mathbf{t})|\ dt_3, \tag{3.2.18}$$

where the function $t_3 \mapsto (g \circ \Gamma_i)(t_1, t_2, t_3)\ |\mathfrak{J}_i'(t_1, t_2, t_3)|$ was used to apply Lemma 2.3.16 with $\mathrm{d} := 1$. Substituting the latter into (3.2.17) yields

$$\mathscr{I}^2(g; S^2) = \sum_{0\leq i<n_2} \iint_{\mathbb{R}^2} \int_{\mathbb{R}} \mathbf{1}_{\Delta^2}(t_1, t_2)\ \delta^1(t_3)\ (g \circ \Gamma_i)(\mathbf{t})\ |\mathfrak{J}_i'(\mathbf{t})|\ dt_1 dt_2\ dt_3. \tag{3.2.19}$$

The above Riemann integral can be converted to a Lebesgue integral as

$$\mathscr{I}^2(g; S^2) = \int_{\mathsf{E}} \sum_{0\leq i<n_2} \mathbf{1}_{\Delta^2}(t_1, t_2)\ \delta^1(t_3)\ g(\mathbf{x})\ |\mathfrak{J}_i'(\mathbf{t})|\ |\mathfrak{J}_i(\mathbf{t})|^{-1}\ \Big|_{\mathbf{t}:=(\Gamma_i^2)^{-1}(\mathbf{x})}\ d\mu^3[\mathbf{x}]. \tag{3.2.20}$$

Now let the defining function $f^1_{S^2} : \mathsf{E} \to \mathbb{R}^2_\star$ of $S^2 \in \mathscr{P}^2(\mathsf{E})$ in (2.3.95) be

$$f^1_{S^2}(\mathbf{x}) := \sum_{0 \le i < n_2} \mathbf{1}_{\Delta^2}(t_1, t_2)\ \delta^1(t_3)\ |\mathfrak{J}'_i(\mathbf{t})|\ |\mathfrak{J}_i(\mathbf{t})|^{-1}\Big|_{\mathbf{t}:=\Gamma_i^{-1}(\mathbf{x})}. \tag{3.2.21}$$

Comparing the above definition with (2.3.95) and noting that the function vanishes at all $t_3 \neq 0$ thus $\delta^1(t_3)\ |\mathfrak{J}_i(t_1, t_2, t_3)|^{-1} = \delta^1(t_3)\ |\mathfrak{J}_i(t_1, t_2, 0)|^{-1}$, and that $\mathbf{1}_{\Delta^2}(t_1, t_2) = \mathbf{1}_{\Delta^2\times\mathbb{R}}(t_1, t_2, 0)$, the surface density function is obtained as

$$\rho_i(t_1, t_2) := \mathbf{1}_{\Delta^2}(t_1, t_2)\ |\mathfrak{J}'_i(t_1, t_2, 0)| = \big(\mathbf{1}_{\Delta^2\times\mathbb{R}}\ \bullet\ |\mathfrak{J}'_i|\big) \circ \Gamma_i^{-1}(\mathbf{x})\ \Big|_{\mathbf{x}:=\Gamma_i(t_1,t_2,0)}, \tag{3.2.22}$$

from which the total artificial volume of each surface patch is obtained as

$$c_i = \|\rho_i\|_1 = \iint_{\mathbb{R}^2} \rho_i(t_1, t_2)\ dt_1 dt_2 = \iint_{\Delta^2} |\mathfrak{J}'_i(t_1, t_2, 0)|\ dt_1 dt_2 = \langle \mathbf{1}_{\Delta^2}, |\mathfrak{J}'_i(\cdot, \cdot, 0)| \rangle. \tag{3.2.23}$$

Substituting for the density function from (3.2.22) back to (3.2.20) yields

$$\mathcal{F}^2(g; S^2) = \langle \mathbb{1}^1_{S^2}, g \rangle = \|\mathbb{1}^1_{S^2}\ \bullet\ g\|_1, \quad \text{where} \quad c_i := 1\ (0 \le i < n_2), \tag{3.2.24}$$

However, one can choose to assign different density distributions when integrating an 'extensive' property function that needs to be scaled with the artificial "heaviness" assigned to each infinitesimal surface element:

$$\mathcal{F}^2(g; S^2) := \langle f^1_{S^2}, g \rangle = \|f^1_{S^2}\ \bullet\ g\|_1, \quad \text{where} \quad c_i > 0\ (0 \le i < n_2). \tag{3.2.25}$$

Volume Integrals. Given a tensor-valued function $g : \mathsf{E} \to \mathbb{Q}$ restricted to a finite union of disjoint $C^{\omega}-$solid volumes $S^3 := \bigcup_{0\leq i\leq n_3} \Gamma_i(\Delta^3)$ parameterized via $\Gamma_i : \mathsf{E} \to \mathsf{E}$, where $\Delta^3 \subseteq \mathbb{R}^3$ is an open 3−cell, in accordance with Convention 2.3.11, the integral in (3.2.4) directly yields

$$\mathcal{I}^3(g; S^3) = \langle \mathbf{1}_{S^3}, g\rangle = \|\mathbf{1}_{S^3} \bullet g\|_1. \tag{3.2.26}$$

Once again, one can choose to assign different density distributions when integrating an 'extensive' property function that needs to be scaled with the density values assigned to each infinitesimal volume element. This is captured by replacing the indicator function in (3.2.26) with some other nonnegative real-valued defining function:

$$\mathcal{I}^3(g; S^3) := \langle f_{S^3}, g\rangle = \|f_{S^3} \bullet g\|_1. \tag{3.2.27}$$

The above results can be summarized into the following commutative diagram:

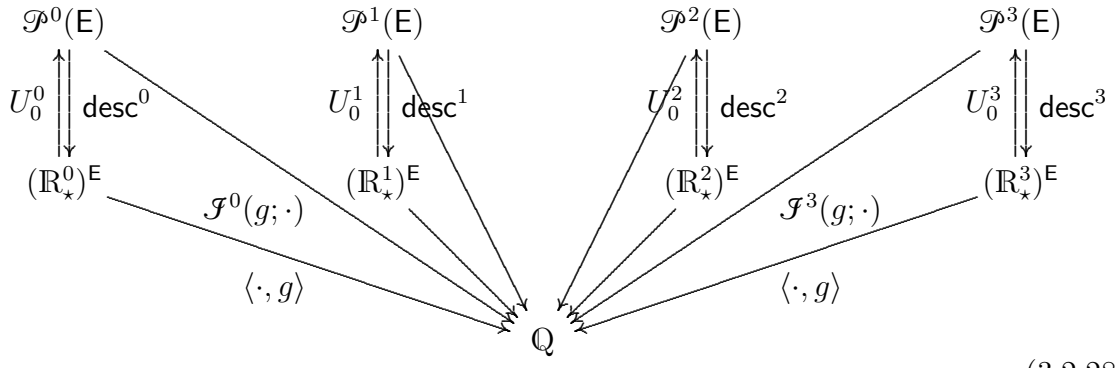

(3.2.28)

in which the labels for some arrows $\mathscr{P}^{\mathrm{d}}(\mathsf{E}) \xrightarrow{\mathcal{I}^{\mathrm{d}}(g;\cdot)} \mathbb{Q} \xleftarrow{\langle\cdot,g\rangle} (\mathbb{R}^{\mathrm{d}}_\star)^{\mathsf{E}}$ for $\mathrm{d} = 1$ and 2 are removed to avoid a cluttered picture. This diagram subsumes (2.3.121) in Section 2.3.2 in which $g := \mathbf{1}_{\mathsf{E}}$ is a unit constant. Once again, integration over features of

heterogeneous dimensions are unified as an inner product based on a single Lebesgue 3−measure, provided that δ−singularities are properly assigned. In fact, using the formal sum $f_S = \sum_{0\le \mathrm{d}<4} f_{S^{\mathrm{d}}}$, the tensor sum of 0−, 1−, 2−, and 3−integrals is computed in one shot using a single inner product:

$$\langle f_S, g\rangle = \langle \sum_{0\le \mathrm{d}<4} f_{S^{\mathrm{d}}}, g\rangle = \sum_{0\le \mathrm{d}<4} \mathcal{I}^{\mathrm{d}}(g; S^{\mathrm{d}}) = \sum_{0\le \mathrm{d}<4} \mathcal{I}^{\mathrm{d}}(g; U_0^{\mathrm{d}}(f_S)). \tag{3.2.29}$$

Figures 3.2.1 and 3.2.2 illustrate the implicit view of 2− and 1−integration of an arbitrary property for a simple planar shape and its boundary, using nonsingular and 1st−order δ−singular shape descriptors, respectively.

To make the above expression computable, recall that each δ−function can be expressed in terms of its limit representation (e.g., chosen from Table 2.3.1) as $\sigma \to 0^+$. Composing the diagrams in (2.3.125) and (3.2.28) one obtains:

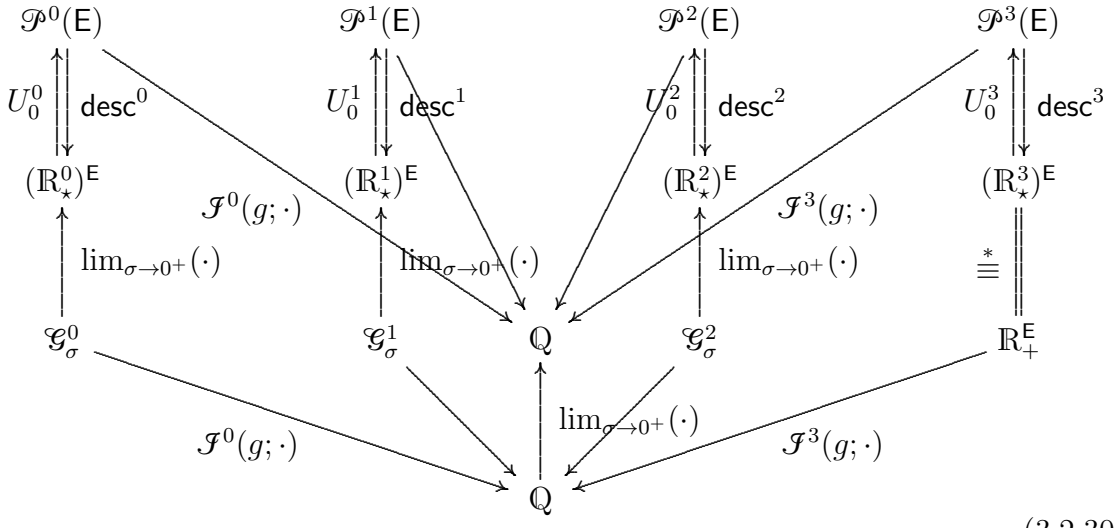

(3.2.30)

The vertical arrow $\mathbb{Q} \xrightarrow{\lim_{\sigma\to 0^+}(\cdot)} \mathbb{Q}$ implies that the result of the inner products for $\sigma > 0$ approximates the exact lower-dimensional integrals with 3D integrals over

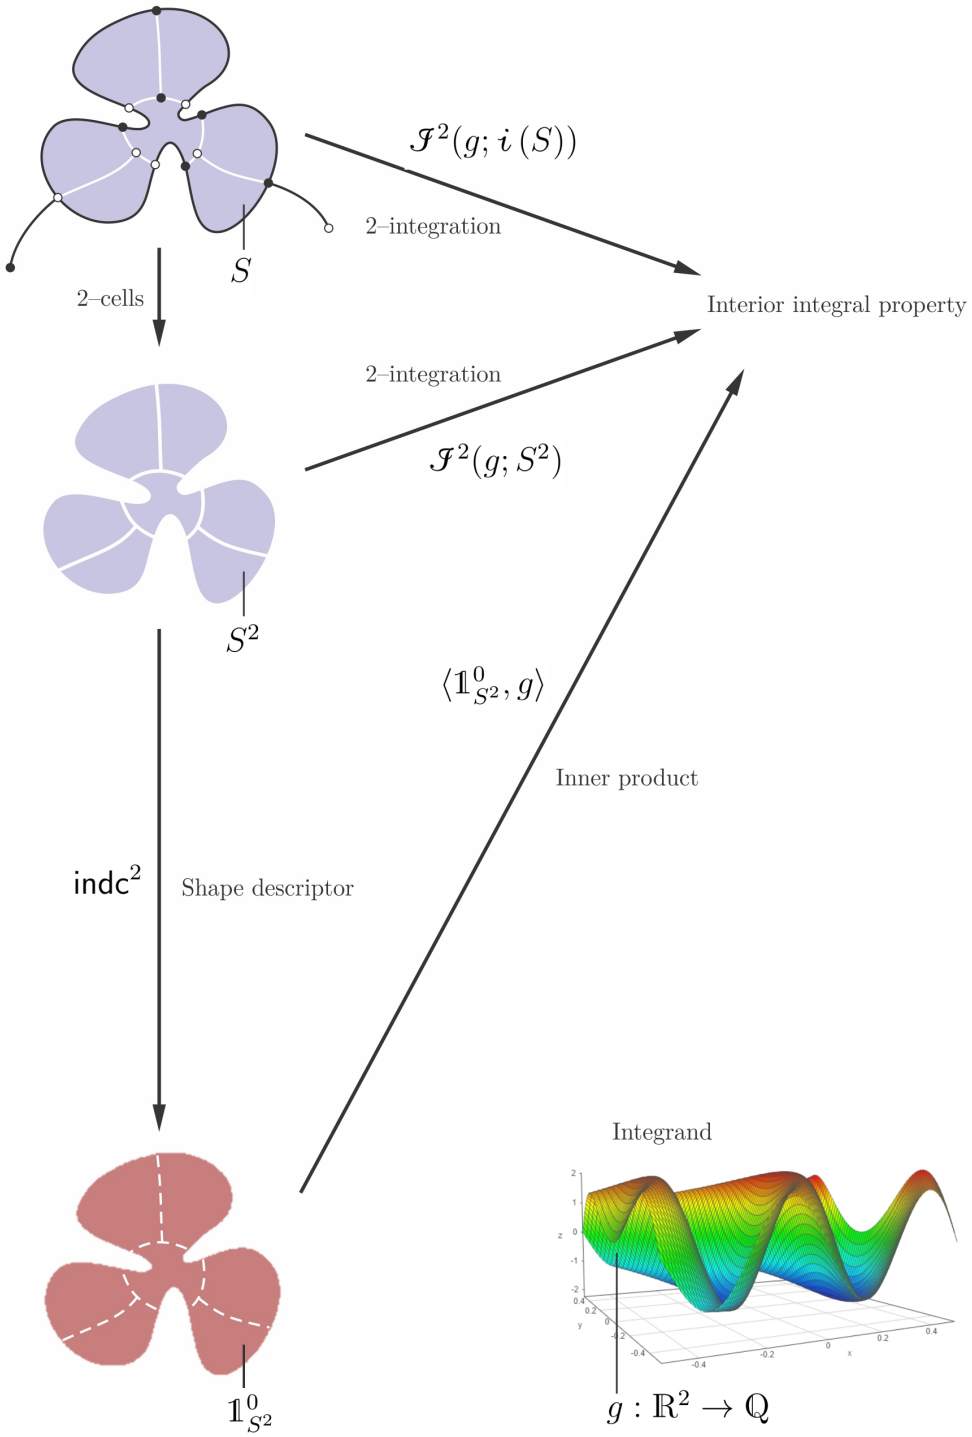

FIGURE 3.2.1: Integrating $g : \mathbb{R}^2 \to \mathbb{Q}$ over the interior of a planar shape amounts to an inner product with the nonsingular descriptor of $S^2 \in \mathcal{P}^2(\mathbb{R}^2)$ (here $S^2 = i(S)$). Note that $\mathbf{1}_S \overset{*}{\equiv} \mathbf{1}_{S^2} = \mathbb{1}^0_{S^2}$ thus $\langle \mathbf{1}_S, g \rangle = \langle \mathbf{1}_{S^2}, g \rangle = \langle \mathbb{1}^0_{S^2}, g \rangle$ which yield $\mathcal{J}^2(g; S) = \mathcal{J}^2(g; S^2)$.

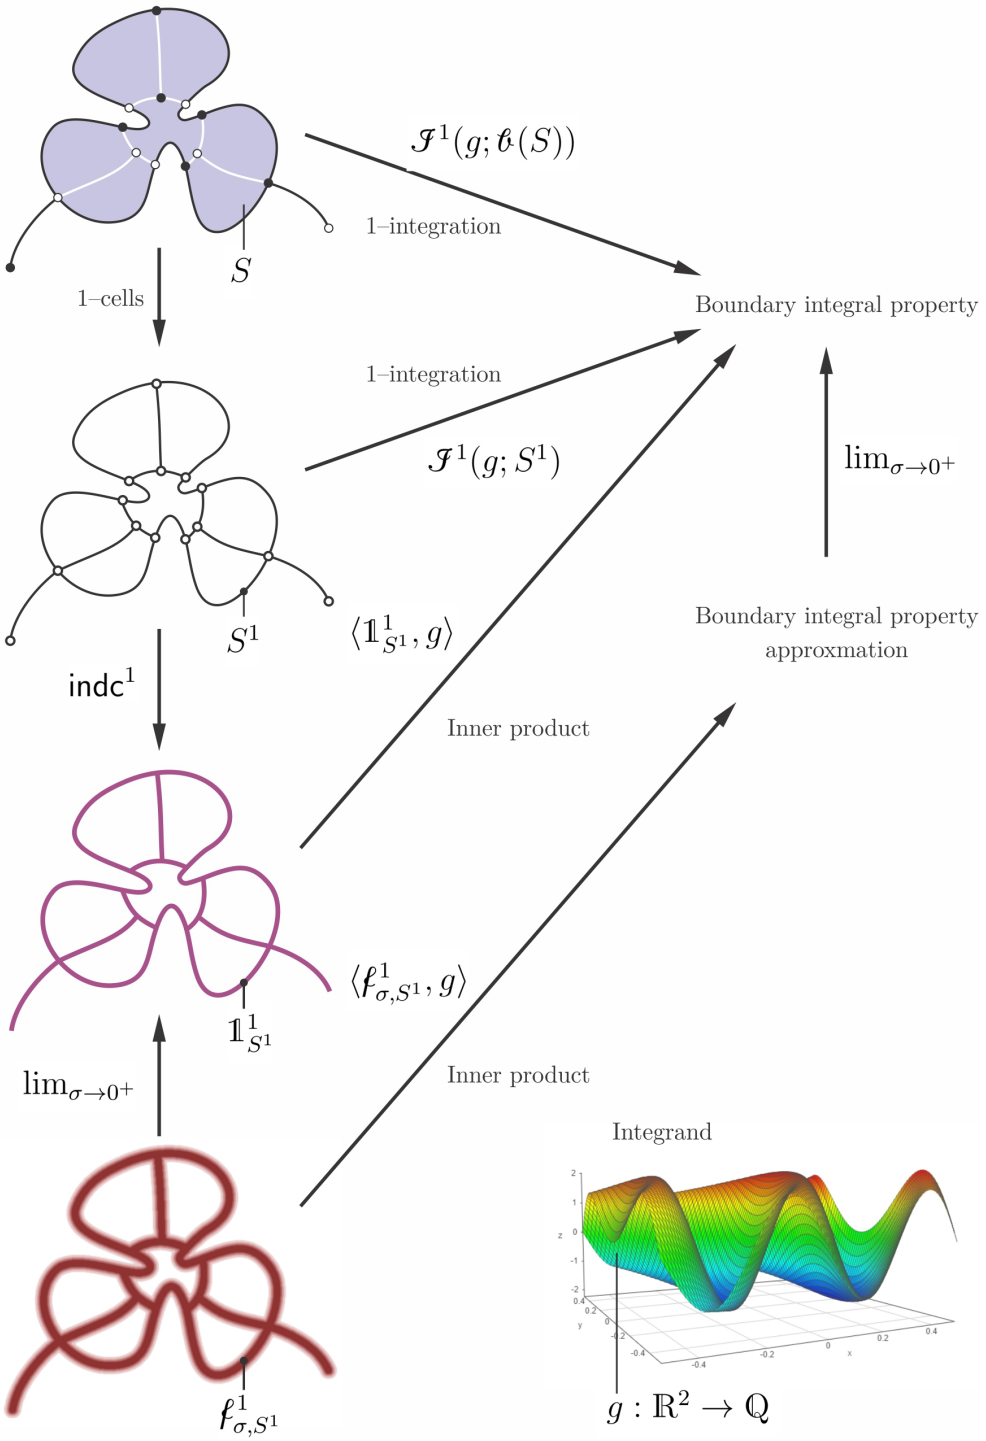

FIGURE 3.2.2: Integrating $g:\mathbb{R}^2\to\mathbb{Q}$ over the boundary of a planar shape amounts to an inner product with the $\delta-$singular descriptor of $S^1\in\mathcal{P}^1(\mathbb{R}^2)$ (here $S^1\subseteq\mathcal{C}(S)$). Note that $\mathbb{1}^1_{S^1}=\lim_{\sigma\to 0^+}\ell_{\sigma,S^1}$ thus $\langle\mathbb{1}^1_{S^1},g\rangle=\lim_{\sigma\to 0^+}\langle\ell_{\sigma,S^1},g\rangle$ where $\ell_{\sigma,S^1}\in\mathcal{G}^1_\sigma$.

the 'dissipated' geometry. The makes the analytic method significantly more robust than the traditional approach to integration, especially if the integrand function's value is sensitive to geometric errors in the integration domain. The choice of finite $\sigma > 0$ provides a natural trade-off mechanism between a reliable resolution for describing lower-dimensional geometric features embedded in the $3-$space and the desired precision of computational results. However, more research is needed to study the convergence rate and error bounds.

Another approach to making the inner product of singular functions computable is to apply Fourier transforms, noting that the $\delta-$singularities, despite being non-computable in the physical domain, convert to ordinary sinusoidal functions in the frequency domain. In particular, Parseval's theorem (Theorem 2.2.10) enables one to rewrite the inner product in (3.2.29) as

$$\mathcal{I}(g;S) := \langle f_S, g\rangle = \langle \hat{f}_S, \hat{g}\rangle = \int_{\hat{\mathsf{E}}} (\hat{f}_S \bullet \hat{g})(\boldsymbol{\omega})\, d\mu^3[\boldsymbol{\omega}], \tag{3.2.31}$$

which can be illustrated by composing the diagrams in (2.2.13) and (3.2.28):

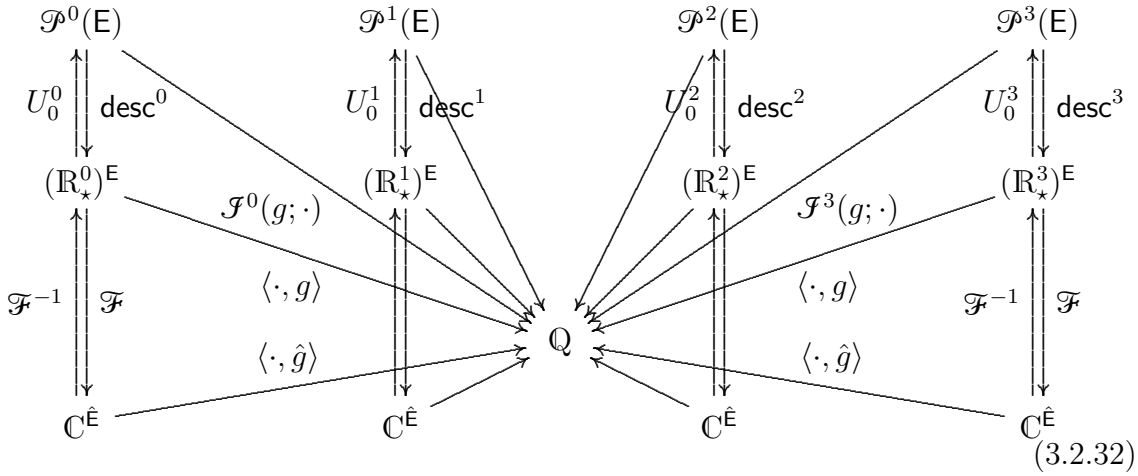

(3.2.32)

In addition, working in the Fourier space provides access to a variety of signal processing techniques such as low-pass filtering. For instance, multiplying the integrand in (3.2.31) with the indicator function of a compact window $\hat{W}_r \subset \hat{\mathsf{E}}$ around the origin—e.g., a 3−ball in the frequency domain $\hat{W}_r := \hat{B}^3_r(\hat{\mathbf{0}})$—denoted by $\mathbf{1}_{\hat{W}_r} : \hat{\mathsf{E}} \to \{0, 1\}$ yields a low-pass filtered approximation of the integral that retains the dominant low-frequency modes within an arbitrary distance of $r > 0$ to the origin $\hat{\mathbf{0}} \in \hat{\mathsf{E}}$ and neglecting the high-frequency modes:

$$\mathscr{I}(g; S) \approx \bar{\mathscr{I}}_r(g; S) := \int_{\hat{\mathsf{E}}} \big((\hat{f}_S \bullet \hat{g}) \bullet \mathbf{1}_{\hat{W}_r}\big)(\boldsymbol{\omega})\, d\mu^3[\boldsymbol{\omega}], \tag{3.2.33}$$

Noting that $\lim_{r\to+\infty} \hat{W}_r = \hat{\mathsf{E}}$ and assuming that the integrand is absolute value integrable (i.e., $|g| \in L^1(\mathsf{E})$) while the integration domain is compact (i.e., $|f_S| \in L^1(\mathsf{E})$) the frequency domain representations have decaying high-frequency amplitudes as $r \to +\infty$, thus $\lim_{r\to+\infty} \bar{\mathscr{I}}_r(g; S) = \mathscr{I}(g; S)$, illustrated via:

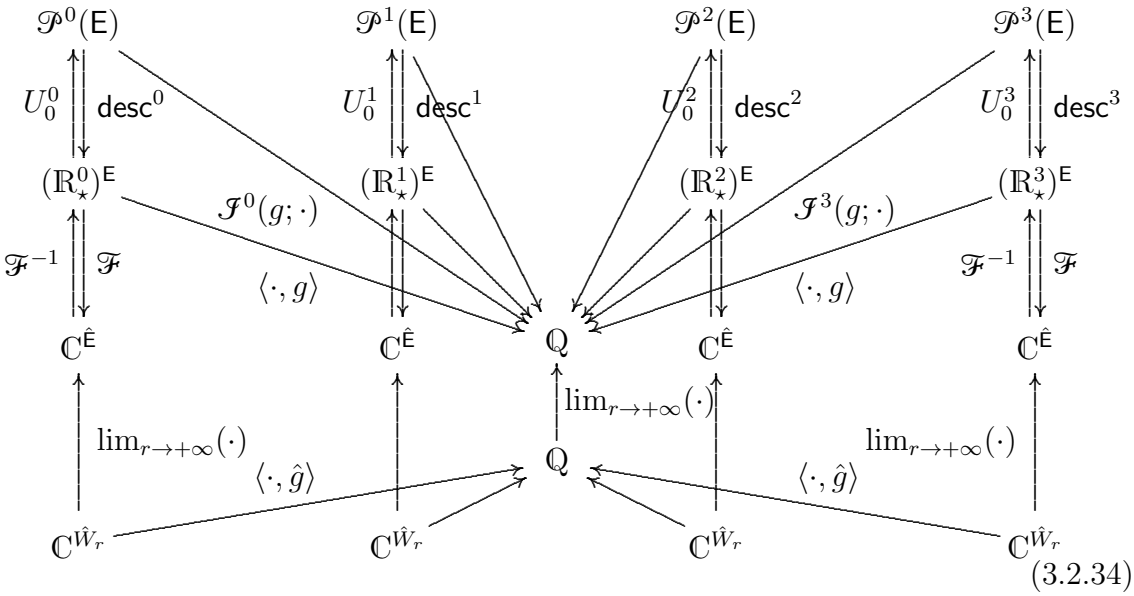

(3.2.34)

If the integrand function is sufficiently well-behaved (e.g., having small high-frequency fluctuations) the approximation $\hat{\mathcal{J}}_r(g;\cdot)$ rapidly converges to the exact value $\mathcal{J}(g;\cdot)$ regardless of the shape complexity of the integration domain. A good approximation can be obtained by computing the integral in (3.2.33) much more efficiently over the integration window $\hat{W}_r \subseteq \hat{\mathsf{E}}$ in the frequency domain, compared to the support $S \subseteq \mathsf{E}$ in the physical domain. In fact, if the integral is computed sequentially, the window size can be increased incrementally, starting from $r = 0$ and growing the window until either a convergence criteria is satisfied or the allocated time budget runs out: One way or another, the choice of finite $r > 0$ provides a natural trade-off mechanism between the available computational resources (e.g., time and memory) and the accuracy of computational results.

The analytic method of integration over shapes can be extended to integration over motions in a straightforward fashion, by replacing the Lebesgue $3-$measure on $\mathsf{E} = \mathbb{R}^3$ with Haar $6-$measure on $\mathsf{C} = \mathrm{SE}(3)$.[4]

3.3 Topological Operations

The two most common approaches to characterizing physical objects (e.g., shapes and motions) in geometric modeling rely on pointset (i.e., general) topology and algebraic topology, both of which are important, giving rise to complementary mathematical models and computer representations [316]. As far as geometric modeling is concerned, the continuum models based on pointset (i.e., general) topology provide the foundations to specify objects and processes in terms of their *local* properties—using

[4]Note that even though the configuration space itself is not a linear space, the space of defining functions is a linear space for both shapes and motions alike, allowing a straightforward generalization of the inner product interpretation of implicit integration to the latter.

either explicit forms in terms of pointsets, or their implicit counterparts as real-valued continuum functions proposed in this thesis—while the combinatorial models based on algebraic topology provide the tools to extend them to *global* constructions—whose explicit forms are orientable manifold cell complexes whose implicit counterparts would perhaps entail using 'co-chain' models to discretize the aforementioned continuum functions integrated over spatial elements (i.e., cells) of heterogeneous dimensions.[5]

The topological operations that I shall discuss in this section are regularization and boundary evaluation, both of which are central to solid modeling, with extensive benefits upon extension to configuration modeling.

3.3.1 Explicit Topology

3.3.1.1 Pointset-Theoretic Regularization via Neighborhoods:

As discussed earlier in Section 2.2.2, there are two dual models of pointset-topological regularization (illustrated in Fig. 3.3.1), namely:

- closed-regularization $r_c := (k \circ i) : \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$, which amounts to computing the closure of interior of a set, i.e., $r_c(S) := k(i(S))$; and
- open-regularization $r_o := (i \circ k) : \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$, which amounts to computing the interior of closure of a set, i.e., $r_o(S) := i(k(S))$;

which are isomorphic by set-theoretic complement $c : \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$ or the pair of topological closure/interior maps[6] $k, i : \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$ defined respectively as the

[5] This is one of my ongoing research directions, and is beyond the scope of this thesis.

[6] Given a topology $\mathbb{T} \subset \mathcal{P}(\mathsf{E})$ that defines open sets—e.g., the 'natural' or 'usual' topology on $\mathsf{E} = \mathbb{R}^3$ based on the L^2-metric—one can define a dual topology in at least two different ways; namely, $\mathbb{T}' := c(\mathbb{T}) \rightleftharpoons \mathbb{T} = c(\mathbb{T}')$ or alternatively $\mathbb{T}' := k(\mathbb{T}) \rightleftharpoons \mathbb{T} := i(\mathbb{T}')$ in which all open and

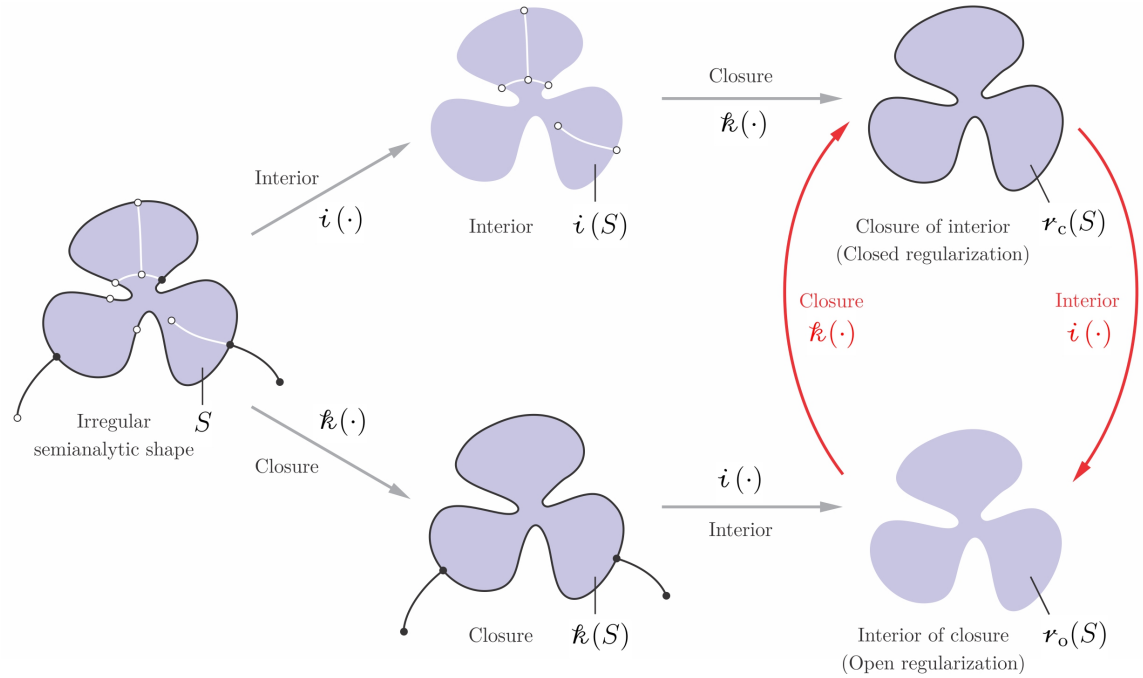

FIGURE 3.3.1: The two types of regularization (i.e., closed and open) are dual to each other. The difference is that one contains the boundary and the other does not.

smallest closed set that contains the set and the largest open set that is contained in the set [276]. In each case, the image space $\mathcal{P}^*(\mathsf{E}) := r(\mathcal{P}(\mathsf{E}))$ is referred to as the (closed or open) 'regular power set' of the 3−space.

For mostly arbitrary reasons, the closed-regularization model gained more popularity as the underlying abstraction for solid modeling developed by Requicha [275, 277, 278], and has been used extensively ever since as the field of solid modeling grew more mature. The open-regularization model, despite offering the same amount of advantages and drawbacks, was advocated in fewer occasions—e.g., by Arbab [16] to allow overlapping boundaries for assembly modeling—but did not catch on [320].

Here are the most basic pros and cons to each choice:

closed sets are swapped. The isomorphism in both cases can be provided either by the set-theoretic complement, or by the pair of topological closure/interior operators.

- closed-regular sets $\mathscr{P}^*(\mathsf{E}) := \boldsymbol{\kappa}_{\mathrm{c}}(\mathscr{P}(\mathsf{E}))$ are closed under regularized union (but not under regularized intersection), i.e.,

$$S_1, S_2 \in \mathscr{P}^*(\mathsf{E}) \;\Rightarrow\; (S_1 \cup S_2) \in \mathscr{P}^*(\mathsf{E}),\ \text{i.e., } (S_1 \cup^* S_2) = (S_1 \cup S_2), \quad (3.3.1)$$

  where $\cup^* := (\boldsymbol{\kappa} \circ \cup)$ thus $\cup^*|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})} = \cup|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})}$.

- open-regular sets $\mathscr{P}^*(\mathsf{E}) := \boldsymbol{\kappa}_{\mathrm{o}}(\mathscr{P}(\mathsf{E}))$ are closed under regularized intersection (but not under regularized union), i.e.,

$$S_1, S_2 \in \mathscr{P}^*(\mathsf{E}) \;\Rightarrow\; (S_1 \cap S_2) \in \mathscr{P}^*(\mathsf{E}),\ \text{i.e., } (S_1 \cap^* S_2) = (S_1 \cap S_2), \quad (3.3.2)$$

  where $\cap^* := (\boldsymbol{\kappa} \circ \cap)$ thus $\cap^*|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})} = \cap|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})}$.

The fact that one of the two dual Boolean operations requires explicit regularization in each case results in at least two major difficulties; namely,

- To establish a CSG scheme [276,282] that is 'well-formed' [315] (i.e., guaranteed for correctness), one needs to explicitly introduce regularization to at least one of the two Boolean operations [343]—e.g., by somehow capturing $\epsilon-$neighborhood differential geometry as $\epsilon \to 0^+$, which is not finitely describable for general semianalytic sets, as illustrated in Fig. 1.3.4.

- In each case, one of the two operations becomes non-computable in particular configurations due to instability of the output with respect to infinitesimal changes in the input [105, 106], meaning that a small perturbation in shapes (e.g., in Hausdorff topology) can cause large changes in the constructive geometry—e.g., due to degenerate conditions with the possibility of points,

curve segments, or surface patches of finite size suddenly appearing or disappearing upon perturbation when two solids are in collision-free contact, i.e., intersect over a subset of their boundaries.

Both problems can be solved by replacing the explicit pointset-topological regularization with an implicit measure-theoretic regularization conceptualized below.

As long as solid models and operations are restricted to the closed Boolean algebra of regular sets [276], regularization does not appear as an explicit operation, but is ensured by maintaining regularity step-by-step—e.g., via neighborhood analysis and maintenance. This is the ideal scenario which underlies 'well-formed' solid representation schemes [315]. However, most solid modeling systems do not strictly satisfy well-formedness due to the computational difficulties with representing and combining neighborhoods with finite precision computing capabilities. A major consequence is that data exchange and translation across different solid modelers can result in invalid or inconsistent representations, e.g., syntactic structures that either map to incorrect semantics (e.g., non-solids) or correct semantics that are different than those of the original model (e.g., erroneous solids). For the former cases, algebraic closure is violated and explicit regularization is sometimes incorporated as heuristic techniques to 'heal' or 'repair' the translated representations [29,46,230,358,367], for which there is neither a guarantee nor a formal understanding of the limitations.

Even in the dream world of well-formed representations and accurate conversions, when solids are modeled as more general sets with heterogeneous dimensions (e.g., as in [137,299,300,359]) due to Assumption 2.2.6, explicit regularization becomes a tool that can be called on demand when (the properties of) the model's full-dimensional subset is being queried. Although combinatorial modeling provides simple and direct

means for this task, it is difficult to implement explicit regularization with continuum modeling; because pointset-topological notion of regularity is a local property defined in terms of infinitesimal neighborhoods and is not computable on finite-precision machines. Current solid modelers generate reliable results for semialgebraic sets arising from CSG with simple lower-order primitives (e.g., planar, cylindrical, spherical, conical, etc.) or free-form surfaces (e.g., Hermite, Bézier, B-splines, NURBS, etc.) by using ad hoc methods to handle neighborhood information in one way or another. However, there are no guarantees and it is not particularly hard to make them fail as soon as their (only qualitatively understood) limitations are pushed.

3.3.1.2 Pointset-Theoretic Boundary Evaluation via Decomposition:

The 'boundary' of a solid is defined as the set of points that have both interior and exterior points nearby; or more precisely put, the set of (included or excluded) points in the 3−space whose every open ϵ−neighborhood in $\mathsf{E} = \mathbb{R}^3$ intersects both interior and exterior (i.e., interior of complement) of the solid. The pointset-topological boundary operator $\mathscr{b} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is thus defined as

$$\mathscr{b}(S) := \mathscr{k}(S) \cap \mathscr{h}(S) = \mathscr{k}(S) \cap (\mathscr{k} \circ \mathscr{c})(S) = \mathscr{k}(S) \cap (\mathscr{c} \circ \mathscr{i})(S), \tag{3.3.3}$$

i.e., as the intersection of two closed sets, one being the closure of the set itself, and the other being the closure of its complement, which is the same as the complement of its interior, as illustrated by Fig. 3.3.2.

Boundaries of (closed or open) regular solids are finite collections of homogeneously 2D orientable manifolds [275]. The Jordan-Brouwer separation theorem [10]

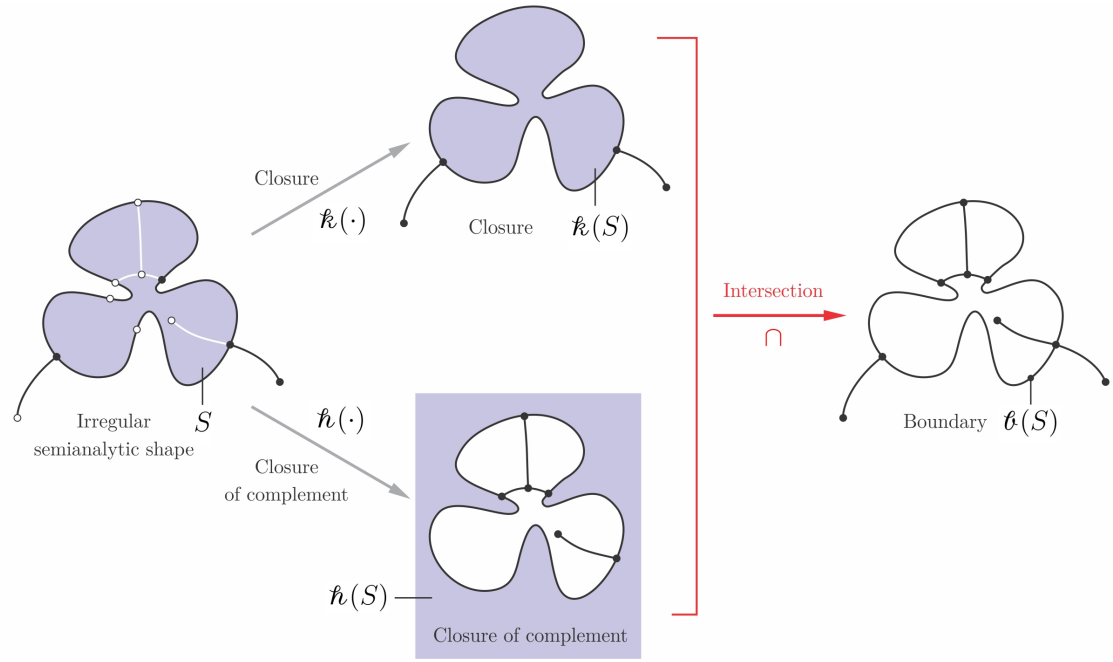

FIGURE 3.3.2: The boundary of a given shape is the set of points that are shared between its closure and closure of complement (i.e., complement of interior).

guarantees that each solid can be uniquely represented by its boundary,[7] meaning that the boundary operator is a one-to-one correspondence between 3D solids and 2D piecewise-manifold surfaces. Therefore, the boundary is a lower-dimensional set that *locally* captures the most important topological and differential properties of the shape.[8] Furthermore, the boundary determines all physical interactions with the environment (e.g., mechanical contact/friction, heat transfer, fluid flow, etc.) ex-

[7]The Jordan-Brouwer separation theorem—which is a generalization of the well-known planar Jordan curve theorem—states that every compact, connected surface in $\mathbb{R}^3$ divides the space into two connected regions; namely, the 'inside' and the 'outside', the former being a compact manifold whose boundary is the said surface [10], implying a one-to-one correspondence between manifold solids and manifold surfaces in the 3−space.

[8]Once again, this stems from the postulate of structural homogeneity and the conceptualization of solids as crisp sets whose interior point behave "all the same" in a sense that they each have a full neighborhood and no relevant property or identity other than being included in the interior—thus no particularly interesting *local* topological or differential properties. This of course changes once a functional approach takes over to model objects as fuzzy sets, assigning grade functions or density functions with meaningful local topological and differential properties to each point.

pressed as boundary conditions—another conceptual interaction being its appearance for computer graphics and haptics rendering.

Although solid modeling was centered around CSG [282] during its early conception in the 1970s and 80s—mostly for its intuitive way of capturing how engineers designed parts in their heads and its compatibility with data transfer and documentation for subtractive manufacturing (i.e., machining)—B-reps took over in the following years for their compactness (i.e., memory efficiency), and amenability to free-form modifications. The 'boundary evaluation' problem was originally formulated as the CSG $\rightarrow$ B-rep conversion [285]—or more generally, conversion between any other 'finite set-theoretic representation' (FSR) [311] to B-reps. It was recognized as one of the most fundamental procedures from the earliest years of solid modeling.[9]

Recalling that boundaries of semianalytic sets are nowhere dense (i.e., at most 2D) semianalytic sets [209]—thus finitely stratifiable into vertices, curved edges, and curved faces, and the fact that boundaries of Boolean constructions are subsets of boundaries of contributing elements, an explicit approach to boundary evaluation [82, 83, 192, 285], typically requires the following set of steps:

1. obtaining a partitioning of the $3-$space into the canonical intersection terms (i.e., 'atoms') [314] induced by some constructive representation (e.g., CSG);

2. testing each of the semianalytic strata (i.e., vertices, edges, and faces) in the boundaries of the atoms against the given representation's set membership classifier (SMC) [342]—which can be reduced, if necessary, to PMC of sampled points on them[10]—retaining those strata that classify as "on the boundary"

[9] To emphasize on its significance, Shapiro [316] called it "the queen of all representation conversion procedures—both in complexity and its importance in solid modeling."

[10] Although one might expect otherwise, SMC is in general more manageable than PMC due to the

and discarding the rest that are now part of resultant interior or exterior; and

3. extracting the combinatorial information to stitch the included strata back together and appropriately orient the surface.

The main problem with this approach is that slight numerical errors in representation and/or computation can cause invalid or incorrect results, which is further magnified when dealing with higher-degree polynomial curves and surfaces [249]. Put differently, neither PMC nor SMC are strictly computable [105, 106] in general in the sense that they are not stable with respect to small perturbations in geometry, especially when classifying boundary elements. In particular, a small numerical imprecision can change whether a particular stratum on the boundary is identified as such or not. Appealing to finite numerical thresholds and other heuristics, on the other hand, does not guarantee validity or consistency.

Exact methods have been proposed for boundary evaluation of Boolean constructions (i.e. CSG) for polyhedral and semialgebraic sets [120, 178, 179, 182–184] which draw on the extensive body of research on computing geometric predicates (particularly intersections) for algebraic curves and surfaces [123, 180, 181] using exact algebraic numbers [368]. The problem with this approach is that it gets computationally expensive when numerous primitives and boundary elements are involved in large geometric models, thus its implementation has been limited to lower-degree polynomials.

---

simplicity of neighborhood geometry when the dimensionality of the candidate set and the reference set are closer to one another. In fact, the neighborhoods become subsets of d−dimensional balls where d is the dimensionality difference—e.g., a binary orientation when testing a surface against a solid, a curve against a surface, or a point against a curve (i.e., $d = 1$), wedges of a 2D disk when testing a curve against a solid or a point against a curve (i.e., $d = 2$), and wedges of a 3D ball when testing a point against a solid (i.e., $d = 3$). Only the first case is finitely describable for general semianalytic sets that are infinitely differentiable, thus need infinitely many parameters to describe their differential properties of the wedges for $d > 1$.

The inverse problem of obtaining a solid from its boundary corresponds to an inverse representation conversion from a given B-rep to another constructive representation (e.g., CSG) which is a nontrivial (though well-studied) problem [100, 322–324]. The inverse task typically follows a similar decompose-and-test process as described above except for the fact that the halfspaces of boundary element carriers are not sufficient as the primitives of the canonical decomposition for the inverse problem, requiring one to add a set of 'separating halfspaces' [324].

Although the Jordan-Brouwer separation theorem [10] guarantees a one-to-one correspondence between models of solids and their boundaries, the FSR construction are clearly nonunique, thus both CSG $\rightarrow$ B-rep and B-rep $\rightarrow$ CSG conversions are technically optimization problems. However, there is little explicit optimization involved in the former as the canonical decomposition typically yields coarse boundary elements, while the need for additional halfspaces in the latter further complicates the problem and might lead to an excessive number of CSG primitives.

3.3.2 Implicit Topology

3.3.2.1 Measure-Theoretic Regularization—Alternative Approaches:

In Sections 2.3.2 and 2.3.3 I presented two different approaches to performing regularization on implicit representation.

The first method, to which I referred as 'implicit regularization' approach or 'irregularity-tolerant' modeling entails a change of perspective—i.e., viewing regularization as a conceptual classification of shapes modulo all of their possible (undesirable) irregularities.[11] What makes this possible in implicit modeling is the obser-

[11] In other words, the first solution to regularization is... Not regularizing at all!

vation that irregularities form lower-dimensional features that are not measurable. From a practical perspective, looking at the functional description as a 'black box' with the *ability to respond* to queries, the probability of hitting a query point right on such features is precisely zero.[12] Thus one should not worry about formation of irregularities along the process as long as measurable properties are of concern.

The second approach takes a more proactive (i.e., 'pseudo-explicit' for the lack of a better term) approach to regularization, in which, a defining function with irregular $0-$superlevel set is converted to another defining function with a regular $0-$superlevel set that is equal$-\mathsf{ae}$ to that of the former, by computing its convolution with the $\delta-$function—at least in theory, whose practical realization amounts to using limit representations (e.g., chosen from Table 2.3.1).

For example, if two shapes $S_1, S_2 \in \mathscr{P}(\mathsf{E})$ are abstracted by their defining functions $f_{S_1}, f_{S_2} : \mathsf{E} \to \mathbb{R}_+$, and their intersect over a 2D surface $(S_1 \cap S_2)$ with no $3-$collision $\mu^3[S_1 \cap S_2] = 0$, the first approach relies on the fact that the defining function of intersection is equal$-\mathsf{ae}$ to a constant zero function (i.e., indicator of empty set), denoted via $(S_1 \cap S_2) \overset{\mathsf{ae}}{\equiv} (S_1 \cap^* S_2)$ because $(S_1 \cap^* S_2) = \emptyset$. Thus both functions return the same output for the same inputs (i.e., query points) except for the case when the query point is on the 2D intersection, which is probabilistically impossible:

$$f_{S_1 \cap^* S_2} \overset{*}{\equiv} f_{S_1 \cap S_2} \quad \text{i.e.,} \quad \mathsf{prob}_{\mathbf{x} \in \mathsf{E}}\big[\mathsf{sign}(f_{S_1 \cap S_2}(\mathbf{x})) \neq \mathsf{sign}(f_{S_1 \cap^* S_2}(\mathbf{x}))\big] = 0,$$

even though there exist (uncountably many, but not measurably much) $\mathbf{x} \in \mathsf{E}$ for which the inequality can occur. The second approach, on the other hand, provides a

[12] Or very small in practice, if finite-precision floating-point arithmetics is used to implement the internal structure of the so-called 'black box' representations.

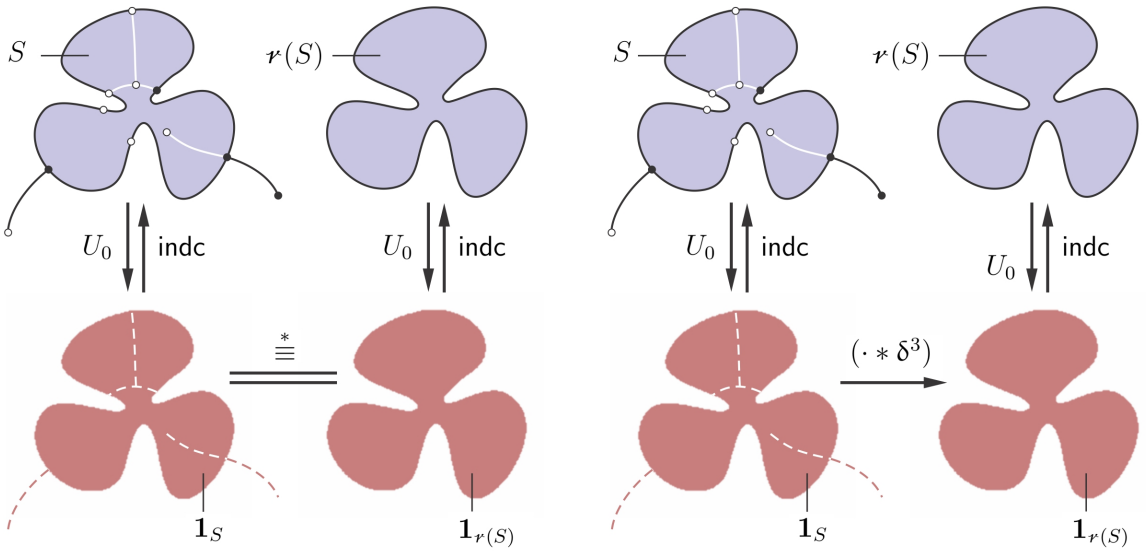

FIGURE 3.3.3: The two different approaches to implicit regularization: (a) one can adopt an irregularity tolerant approach, in which irregularities are simply ignored due to their lack of participation in measure-theoretic operations; or (b) one can convolve the defining function with the $\delta-$function to prune the irregularities from its $0-$superlevel set.

mechanism that would convert $f_{S_1 \cap S_2}$ to $f_{S_1 \cap^* S_2}$ precisely (not almost) everywhere:

$$f_{S_1 \cap^* S_2} \equiv (f_{S_1 \cap S_2} * \delta^3) \quad \text{i.e.,} \quad \mathsf{sign}(f_{S_1 \cap^* S_2}(\mathbf{x})) = \mathsf{sign}((f_{S_1 \cap S_2} * \delta^3)(\mathbf{x})),$$

for all $\mathbf{x} \in \mathsf{E}$ without any exception. The difference is illustrated in Fig. 3.3.3.

Note that the first approach is indifferent to closed versus open regularization, as they are equal$-$$\mathsf{ae}$, which can be expressed as:

$$f_S \overset{*}{\equiv} f_{r_c(S)} \overset{*}{\equiv} f_{r_o(S)} \quad \text{i.e.,} \quad U_0(f_S) \overset{\mathsf{ae}}{=} U_0(f_{r_c(S)}) \overset{\mathsf{ae}}{=} U_0(f_{r_o(S)}). \tag{3.3.4}$$

This can be illustrated with the following commutative diagram:

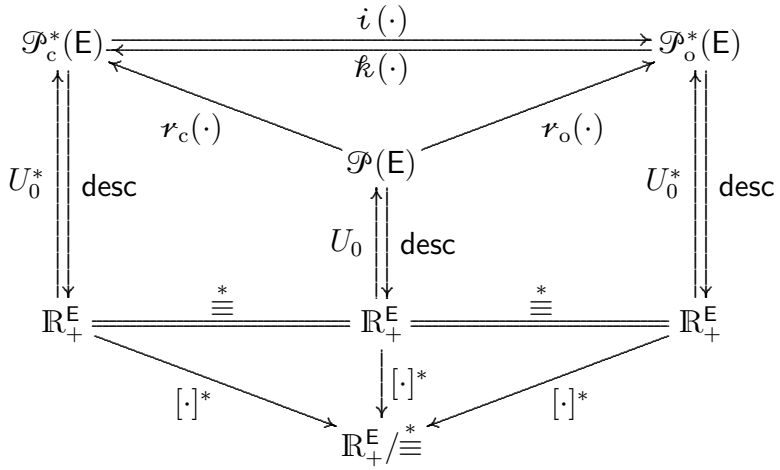

(3.3.5)

in which $\mathscr{P}^*_{\mathrm{c}}(\mathsf{E}) = \mathscr{r}_{\mathrm{c}}(\mathscr{P}(\mathsf{E}))$ and $\mathscr{P}^*_{\mathrm{o}}(\mathsf{E}) = \mathscr{r}_{\mathrm{o}}(\mathscr{P}(\mathsf{E}))$ are dual to each other. On the other hand, the second approach allows one to obtain closed regularization via $\delta-$convolution, from which open regularization is obtain via De Morgan duality (i.e., using two negations), which can be illustrated via:

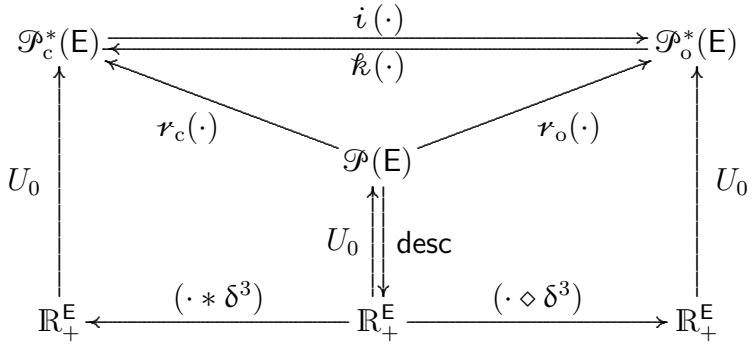

(3.3.6)

where $*, \diamond : L^1(\mathsf{E}) \times L^1(\mathsf{E}) \to L^1(\mathsf{E})$ are the convolution and its dual operator defined via $(f_1 \diamond f_2) := \neg((\neg f_1) * f_2)$, respectively. Figure 3.3.4 gives an illustration.

The reason closed regularization is favored by the $\delta-$calculus over open regularization in spite of the fundamental symmetry between them is a direct consequence of a choice that I made in defining the $\delta-$function in Assumption 2.3.8 of Section

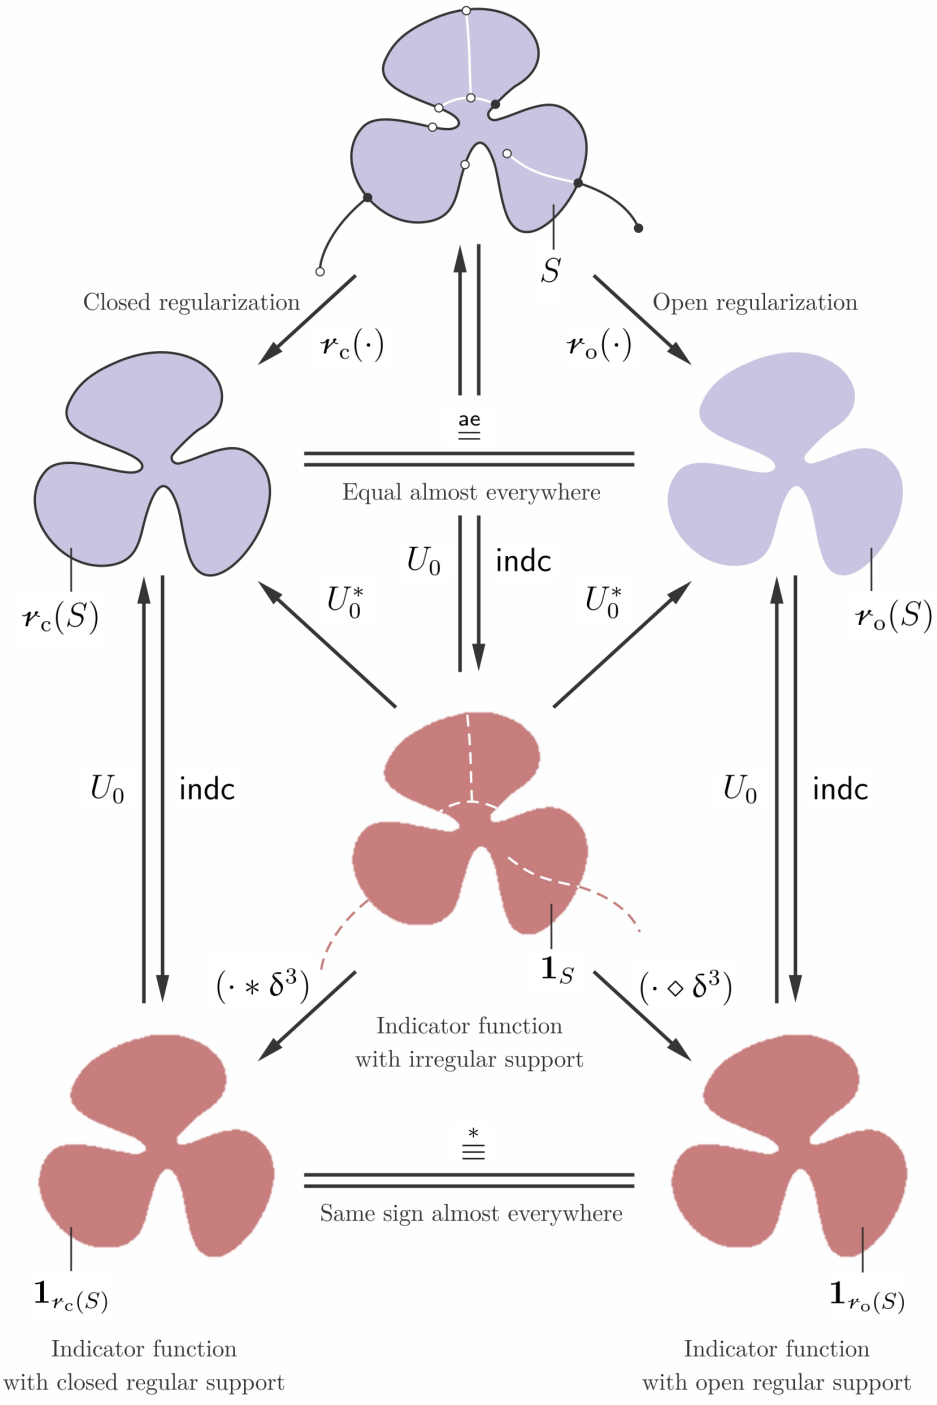

FIGURE 3.3.4: Closed and open regularization in the explicit realm (top) correspond to convolution and its dual operation with the $\delta-$function, respectively, in the implicit realm (bottom). The resulting functions are equivalent to the original function modulo $\overset{*}{\equiv}$.

2.3.2; namely, the way in which the neighborhood ambiguities were resolved when integrating over a region with a $\delta-$singularity on its boundary. More precisely, as long as the $\delta-$function is viewed as the limit of a proper function (e.g., Gaussian dissipation of the lumped singularity) as $\sigma \to 0^+$, its effects on its neighborhood are positive contributions to the defining function being regularized, which, in turn, when used with the convention of strict $0-$superlevel set shape description, results in an infinitesimal *growth* of the shape upon $\delta-$convolution, rather than *shrinking* it—thus closed rather than open regularization.

The two approaches to measure-theoretic regularization and their relationship were summarized in (2.3.151) of Section 2.3.3.1, repeated below:

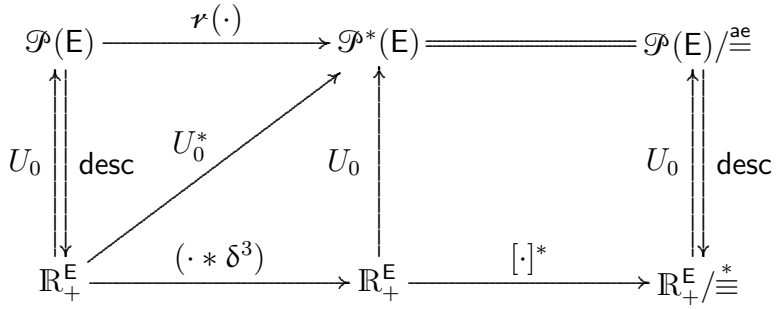

(3.3.7)

Once again, the $\delta-$function is non-computable in practice, which requires approximation using a nascent $\delta-$function (e.g., selected from Table 2.3.1), which essentially results in 'mollification' of 'blurring' of the indicator function. In particular, if a compactly supported nascent $\delta-$function (e.g., a mollifier) is used, the convolution $(f_S * \mathscr{g}_\sigma^3) : \mathsf{E} \to \mathbb{R}_+$ returns a nonnegative real-valued defining function of a grown shape that contains the regularized shape $\mathscr{r}(S) \in \mathscr{P}^*(\mathsf{E})$. For example, it can be shown that if $\mathscr{g}_\sigma^3 := \mathbf{1}_{B_\sigma(\mathbf{0})}/\mu^3[B_\sigma(\mathbf{0})]$ is chosen to be the normalized indicator function of a 3D $\sigma-$ball of volume $\mu^3[B_\sigma(\mathbf{0})] = \frac{4}{3}\pi\sigma^3$, the result is the defining function of a $+\sigma-$offset of the regularized shape: $(f_S * \mathscr{g}_\sigma^3) \equiv f_{(\mathbb{O}_r^{\uparrow\ddagger} \circ \mathscr{r})(S)} = f_{\mathscr{r}(S) \oplus^\ddagger B_\sigma(S)}$. This

can be expressed by the following extension to the left-hand side of (3.3.7):

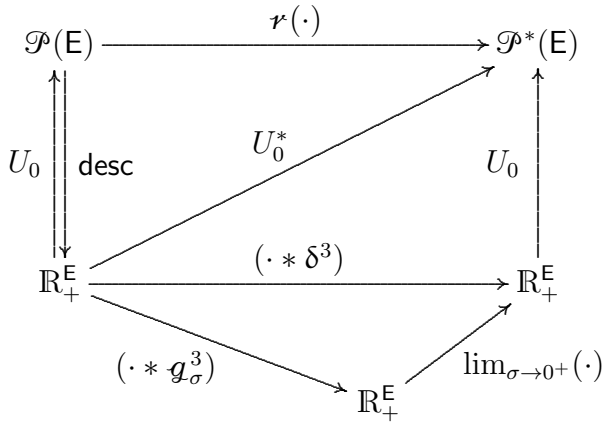

(3.3.8)

The convolution with the nascent $\delta-$function can in turn be converted to a pointwise multiplication in the Fourier domain using the convolution theorem (Theorem 2.2.11):

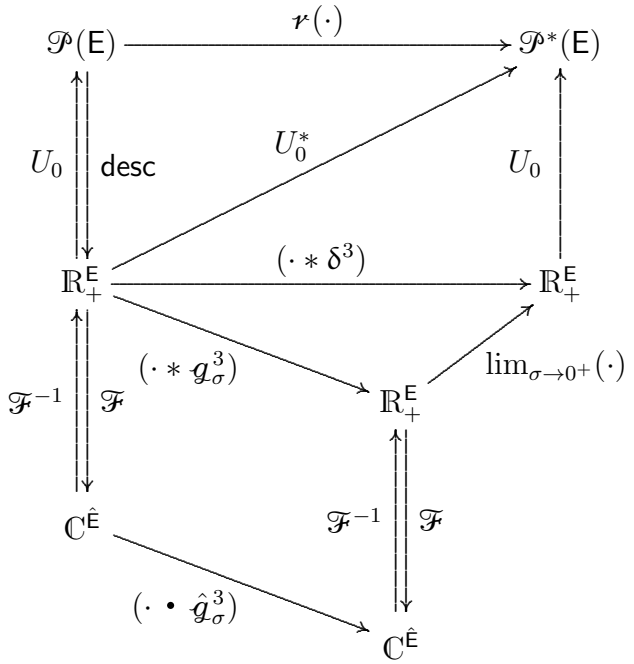

(3.3.9)

In essence, regularization in the implicit realm is the limit representation of 'anti-aliasing' of shape descriptors, using an arbitrary filter $\boldsymbol{g}_\sigma^3 : \mathsf{E} \to \mathbb{R}_+$. In particular,

if a Gaussian kernel $g^3_\sigma(\mathbf{x}) := \frac{1}{2(\sqrt{\pi}\sigma)^3} e^{-\|\mathbf{x}\|_2^2/\sigma^2}$ is chosen—whose Fourier transform happens to be another Gaussian, over the frequency domain—the anti-aliasing can be viewed as a Gaussian amplification of the low-frequency modes and weakening of the high-frequency terms. The high-frequency terms are not completely zeroed out, unlike the case with band-limited filters whose truncation of the high-frequency tail creates 'rippling' effects in the physical domain, and possibly an infinitely long wave with unbounded support. However, this desirable property of the Gaussian filter comes at a cost; namely, the convolution function becomes strictly positive with $U_0(f_S * g^3_\sigma) = \mathsf{E}$ for all $\sigma > 0$—but the value decays double-exponentially with distance from the original shape. To obtain a bounded approximation to the regularized shape, one needs to take an $\epsilon-$superlevel set for $\epsilon > 0$, which I shall elaborate shortly. See Fig. 3.3.5 for an illustration of $\delta-$regularization and its limit representation as anti-aliasing for a simple planar shape.

Figure 3.3.6 illustrates the local meaning of $\delta-$regularization (the second method) on a planar irregular shape's nonnegative real-valued defining function $f_S : \mathsf{E} \to \mathbb{R}_+$. The evaluation of $(f_S * \delta^3) : \mathsf{E} \to \mathbb{R}_+$ at a query point $\mathbf{x} \in \mathsf{E}$ amounts to placing a small $\sigma-$neighborhood $B_\sigma(\mathbf{x})$ at the query point and testing if its intersection with the original shape has a nonzero $3-$measure:

$$f_{r(S)}(\mathbf{x}) := (f_S * \delta^3)(\mathbf{x}) > 0 \quad \text{iff} \quad \exists \sigma > 0 : \mu^3[S \cap B_\sigma(\mathbf{x})] > 0. \tag{3.3.10}$$

To understand this, note that the $\delta-$function can be viewed as the limit of the defining function of a $\sigma-$neigborhood at the origin:

$$\delta^3(\mathbf{x}) = \lim_{\sigma \to 0^+} \frac{\mathbf{1}_{B_\sigma(\mathbf{0})}(\mathbf{x})}{\mu^3[B_\sigma(\mathbf{0})]} \quad \Rightarrow \quad \delta^3(\mathbf{x} - \mathbf{x}') = \lim_{\sigma \to 0^+} \frac{\mathbf{1}_{B_\sigma(\mathbf{x}')}(\mathbf{x})}{\mu^3[B_\sigma(\mathbf{x}')]}, \tag{3.3.11}$$

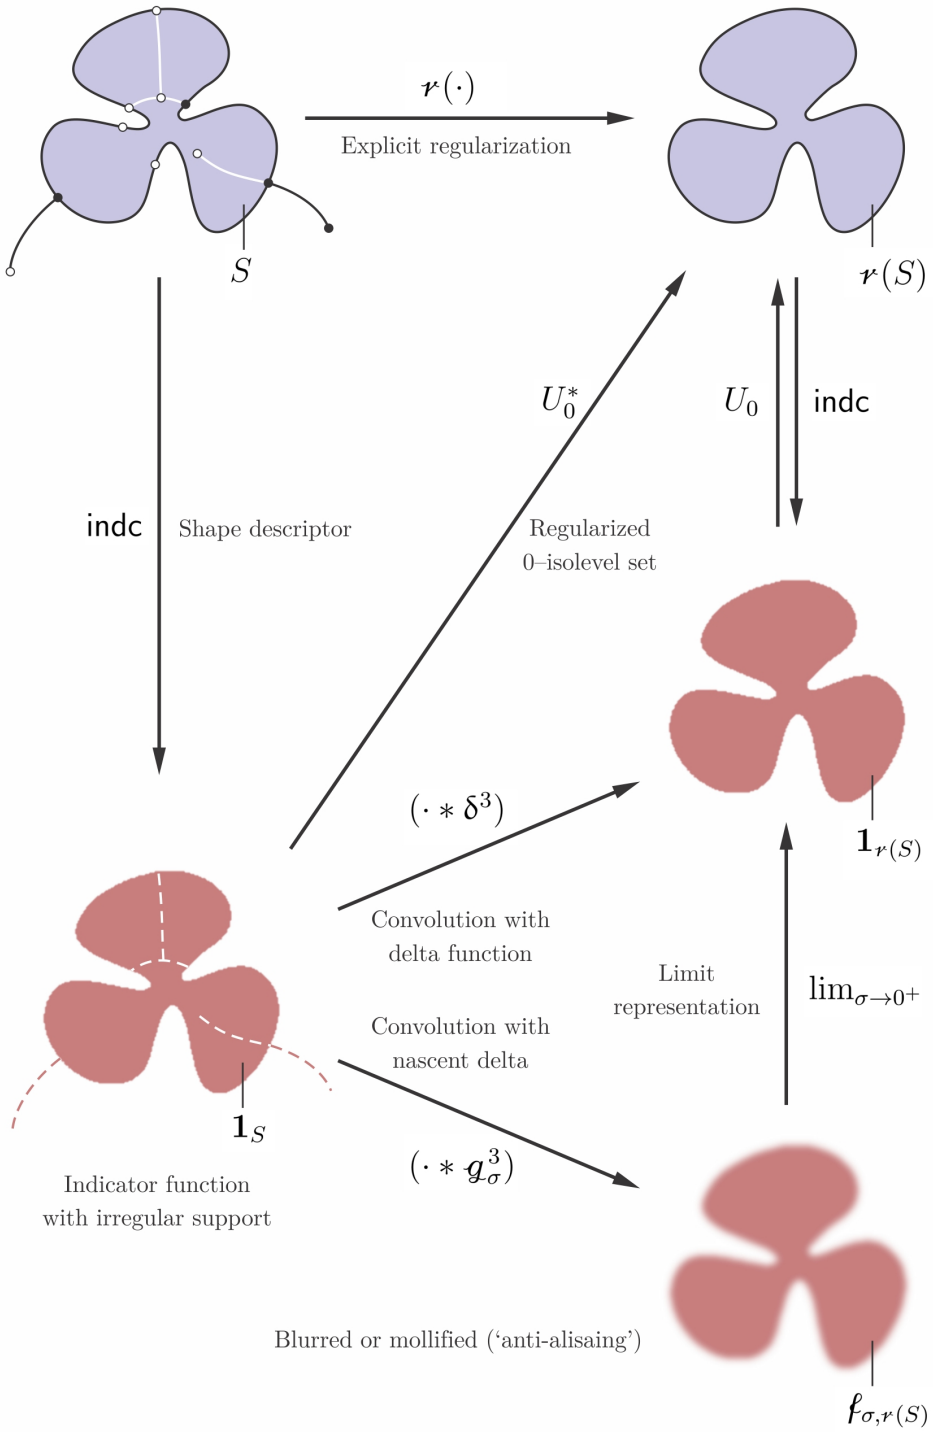

FIGURE 3.3.5: The convolution of a nonnegative real-valued defining function of a 2D regular shape with the $3^{\text{rd}}-$order $\delta-$functions yields a nonnegative real-valued defining function of its closed regularization.

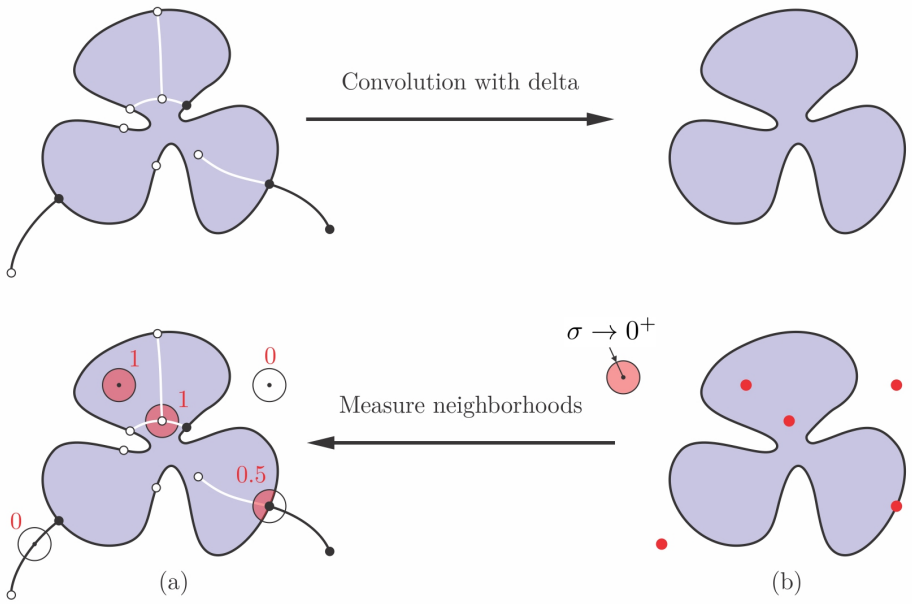

FIGURE 3.3.6: Given $f_S : \mathsf{E} \to \mathbb{R}_+$, evaluating $(f_S * \delta^3)(\mathbf{x})$ at a given $\mathbf{x} \in \mathsf{E}$ in (b) can be viewed as testing if $\mu^3[S \cap B_\sigma(\mathbf{x})]$ remains nonzero as $\sigma \to 0^+$ in (a).

in which the indicator function is normalized by $\mu^3[B_\sigma(\mathbf{0})] = \mu^3[B_\sigma(\mathbf{x}')] = \frac{4}{3}\pi\sigma^3$ (in case of a 3−ball neighborhood) to obtain a unit L^1−norm as required by Definition 2.3.7, which does not alter the function up to equivalence via $\equiv$. As a result, regularization can be viewed as a morphological operation in the limit—namely, a measurable Minkowski sum (defined in Section 3.4) with a σ−ball as $\sigma \to 0^+$, i.e.,

$$\mathscr{r}(S) \stackrel{?}{=} \lim_{\sigma \to 0^+} (S \oplus^{\ddagger} B_\sigma(\mathbf{0})) \quad \text{i.e.,} \quad f_{\mathscr{r}(S)} \stackrel{?}{\equiv} \lim_{\sigma \to 0^+} (f_S * \mathbf{1}_{B_\sigma(\mathbf{0})}) \equiv (f_S * \delta^3). \tag{3.3.12}$$

The caveat is that guaranteeing the convergence of the shape on the left-hand side (e.g., in Hausdorff topology) appears highly nontrivial.[13] Nevertheless, the limit representation provides a way to approximate the non-computable convolution with

[13] See a related discussion in [260–262] of ϵ−topological operators and how they extend the traditional notions of regularity and solidity.

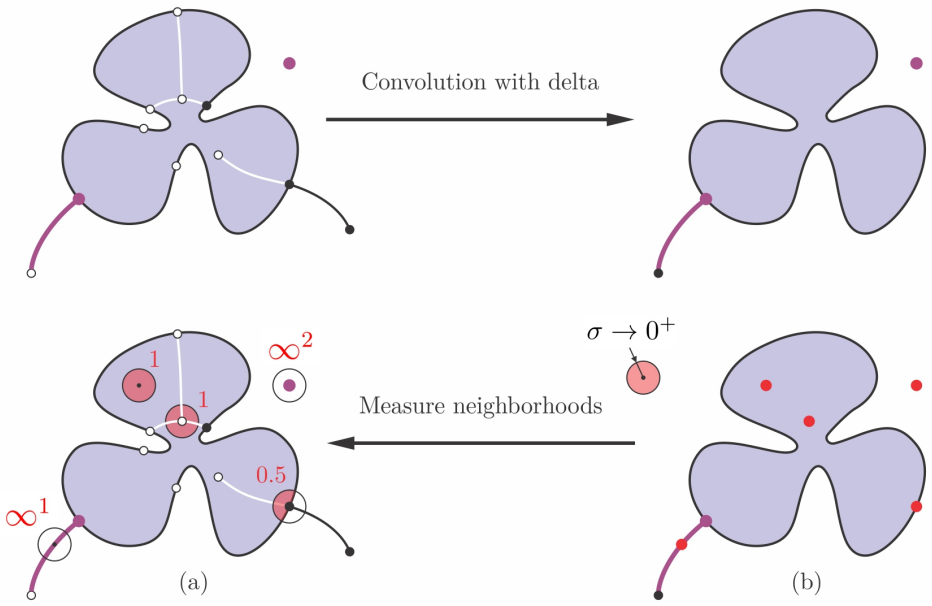

FIGURE 3.3.7: Given $f_S : \mathsf{E} \to \mathbb{R}_\star$, evaluating $(f_S * \delta^3)(\mathbf{x})$ at a given $\mathbf{x} \in \mathsf{E}$ in (b) can be viewed as characterizing how $\mu^3[S \cap B_\sigma(\mathbf{x})]$ behaves as $\sigma \to 0^+$ in (a).

$\delta-$function with a computable convolution with a mollifier $g^3_\sigma : \mathsf{E} \to \mathbb{R}$ (e.g., Gaussian or otherwise selected from Table 2.3.1).

Now consider the case when extended-real-valued defining functions are used, introducing artificial 3−measure to the lower-dimensional features. As illustrated in Fig. 3.3.7, convolution with the $\delta-$function cannot regularize such features that are assigned with $\delta-$singularities of proper orders that make them 3−measurable—namely, d−dimensional features with $(3 - \mathrm{d})-$singularities—as the intersections of the $\sigma-$neighborhoods with such features are artificially 3−measurable.

I conjecture that the regularization can be obtained in the limit (e.g., in Hausdorff topology) as the $\epsilon-$superlevel set of the convolution of a defining function with the

mollifier, which, in turn, is a limit representation of the $\delta-$function, i.e.,

$$\mathscr{r}(S) = U_0(f_S * \delta^3) = U_0(f_S * (\lim_{\sigma \to 0^+} \mathscr{g}_\sigma^3)) \stackrel{?}{=} \lim_{\epsilon,\sigma \to 0^+} U_\epsilon(f_S * \mathscr{g}_\sigma^3), \tag{3.3.13}$$

where the relationship between $\sigma, \epsilon > 0$ is unknown. More research is required to understand the behavior of the approximate regularization for objects of arbitrary shape. In other words, I conjecture that $\mathscr{r}(S) = \lim_{\sigma \to 0^+} U_{\mathscr{f}(\sigma)}(f_S * \mathscr{g}_\sigma^3)$ holds under meaningful practical conditions, in which the relationship between the relative rate of convergence of the parameters $\sigma, \epsilon \to 0^+$ may depend on the choice of the $\mathscr{g}_\sigma-$kernel.

3.3.2.2 Measure-Theoretic Boundary Evaluation via $\delta-$Sifting:

Let us start from the definition of the boundary in (3.3.3), which can be converted to the following implicit definition:

$$\mathbf{x} \in \mathscr{b}(S) \quad \text{iff} \quad \forall \epsilon > 0 : B_\epsilon(\mathbf{x}) \cap \mathscr{k}(S) \neq \emptyset \quad \text{and} \quad B_\epsilon(\mathbf{x}) \cap \mathscr{h}(S) \neq \emptyset. \tag{3.3.14}$$

where $B_\epsilon(\mathbf{x}) = \left\{\mathbf{x}' \in \mathsf{E} \mid \|\mathbf{x} - \mathbf{x}'\|_2 < r\right\}$ is an open 3−ball of radius $r > 0$ centered at the query point $\mathbf{x} \in \mathsf{E}$, and as in (3.3.3), $\mathscr{h}(S) = (\mathscr{k} \circ \mathscr{c})(S) = (\mathscr{c} \circ \mathscr{i})(S)$. The above conditions mean that for every $\epsilon > 0$ and a given query point $\mathbf{x} \in \mathsf{E}$, the point belongs to the boundary (i.e., $f_{\mathscr{b}(S)}(\mathbf{x}) > 0$) iff there exist at least two distinct points $\mathbf{x}_1 \neq \mathbf{x}_2$, one almost internal $\mathbf{x}_1 \in \mathscr{k}(S)$ and one almost external $\mathbf{x}_2 \in \mathscr{h}(S)$, whose distances to the query point are bounded by $\epsilon > 0$:

$$\mathbf{x} \in \mathscr{b}(S) \quad \text{iff} \quad \forall \epsilon > 0 : \qquad \exists \mathbf{x}_1 \in \mathscr{k}(S) : \|\mathbf{x} - \mathbf{x}_1\|_2 < \epsilon, \tag{3.3.15}$$

$$\text{and} \quad \exists \mathbf{x}_2 \in \mathscr{h}(S) : \|\mathbf{x} - \mathbf{x}_2\|_2 < \epsilon, \tag{3.3.16}$$

i.e., $\mathbf{x} \in \mathscr{b}(S)$ iff $\mathrm{d}(\mathscr{k}(S), \mathbf{x}) = \mathrm{d}(\mathscr{h}(S), \mathbf{x}) = 0$ where the distance between from a query point to a set is defined as $\mathrm{d}(S, \mathbf{x}) := \inf_{\mathbf{x}' \in S} \|\mathbf{x} - \mathbf{x}'\|_2$. Alternatively, the boundary can be implicitly defined as the $0-$isolevel set of the C^0-continuous signed distance function defined in (3.1.2), which can be rewritten as

$$\mathrm{d}_S(\mathbf{x}) = \begin{cases} -\mathrm{d}(\mathscr{h}(S), \mathbf{x}) = -\mathrm{d}(\mathscr{b}(S), \mathbf{x}) & \text{if } \mathbf{x} \in \mathscr{i}(S), \\ 0 & \text{if } \mathbf{x} \in \mathscr{b}(S), \\ +\mathrm{d}(\mathscr{k}(S), \mathbf{x}) = +\mathrm{d}(\mathscr{b}(S), \mathbf{x}) & \text{if } \mathbf{x} \in \mathscr{e}(S), \end{cases} \tag{3.3.17}$$

thus $\mathbf{x} \in \mathscr{b}(S)$ iff $\mathrm{d}_S(\mathbf{x}) = \mathrm{d}(\mathscr{b}(S), \mathbf{x}) = 0$. Consequently, the defining function $f^0_{\mathscr{b}(S)}(\mathbf{x}) : \mathsf{E} \to \mathbb{R}_+$ can be constructed in terms of the signed distance function as

$$f^0_{\mathscr{b}(S)}(\mathbf{x}) := \big(\rho \bullet (\delta^0 \circ \mathrm{d}_S)\big)(\mathbf{x}) = \rho(\mathbf{x})\ \delta^0\big(\mathrm{d}_S(\mathbf{x})\big), \tag{3.3.18}$$

where $\delta^0(r) = \mathbf{1}_{\{0\}}(r)$, i.e., $\delta^0(r) = 1$ if $r := \mathrm{d}_S(\mathbf{x}) = 0$ and $\delta^0(r) = 0$ otherwise. The function $\rho : \mathsf{E} \to \mathbb{R}_+$ is an arbitrary density function that is nonzero over the boundary, whose profile does not change the fact that $U_0(f_{\mathscr{b}(S)}) = \mathscr{b}(S)$. However, this function is not very useful in the sense that it is not artificially $3-$measurable, thus cannot be used in downstream processes that are all formulated in terms of Lebesgue $3-$integrals—e.g., L^1-norms, inner products, and more generally, convolutions. The idea is simple—and the above expression was deliberately written in a way to give it away: replace δ^0 with δ^1 in (3.3.18)to impart an artificial $3-$measure, obtaining the extended-real-valued defining function $f^1_{\mathscr{b}(S)}(\mathbf{x}) : \mathsf{E} \to \mathbb{R}_\star$:

$$f^1_{\mathscr{b}(S)}(\mathbf{x}) := \big(\rho \bullet (\delta^1 \circ \mathrm{d}_S)\big)(\mathbf{x}) = \rho(\mathbf{x})\ \delta^1\big(\mathrm{d}_S(\mathbf{x})\big), \tag{3.3.19}$$

Recall from Section 2.3.2 that the L^1-norm of the weighted (via $\rho-$function) composition of the $\delta-$function with every differentiable function is obtained from the generalized sifting formulae given by Corollary 2.3.22. Substituting for the distance function whose 0−level set is the boundary itself (i.e., $Z := \mathbb{d}_S^{-1}(0) = \mathscr{b}(S)$) yields:

$$\|f^1_{\mathscr{b}(S)}\|_1 = \|\rho \bullet (\delta^1 \circ \mathbb{d}_S)\|_1 = \int_{\mathscr{b}S} \frac{\rho(\mathbf{x}')}{\|\nabla \mathbb{d}_S(\mathbf{x}')\|_2} \, d\mu^2[\mathbf{x}'], \tag{3.3.20}$$

A caveat is that the distance function, despite being C^0-continuous everywhere, may not necessarily be C^1-continuous everywhere. In fact, the distance function is not differentiable along the cut locus (i.e., closure of medial axis (MA)) of the interior and exterior, whose intersection with the boundary are exactly at the "sharp" (i.e., convex) or "dull" (i.e., concave) edges and vertices where the surface strata are stitched together. Nevertheless, if the shape is (closed or open) regular, in addition to being semianalytic due to Assumption 2.2.6, the boundary is a surface that admits an analytic stratification and the distance function is smooth in the interior of the surface patches. In other words, it is differentiable 'almost everywhere' (denoted $C^1-\mathsf{ae}$) with respect to the Hausdorff 2−measure on the surface, and the integration in (3.3.20) is not affected by those sharp or dull corners. Furthermore, *the distance function has a unit gradient norm* $\|\nabla \mathbb{d}_S(\mathbf{x}')\|_2 = 1$ everywhere else, i.e.,

$$\|f^1_{\mathscr{b}(S)}\|_1 = \|\rho \bullet (\delta^1 \circ \mathbb{d}_S)\|_1 = \int_{\mathscr{b}(S)} \rho(\mathbf{x}') \, d\mu^2[\mathbf{x}'], \tag{3.3.21}$$

Letting $\rho(\mathbf{x}') := 1$ for all $\mathbf{x}' \in \mathscr{b}(S)$ leads to the $\delta-$singular characteristic function:

$$\|\mathbb{1}^1_{\mathscr{b}(S)}\|_1 = \|\delta^1 \circ \mathbb{d}_S\|_1 = \int_{\mathscr{b}(S)} d\mu^2[\mathbf{x}'] = \mu^2[\mathscr{b}(S)]. \tag{3.3.22}$$

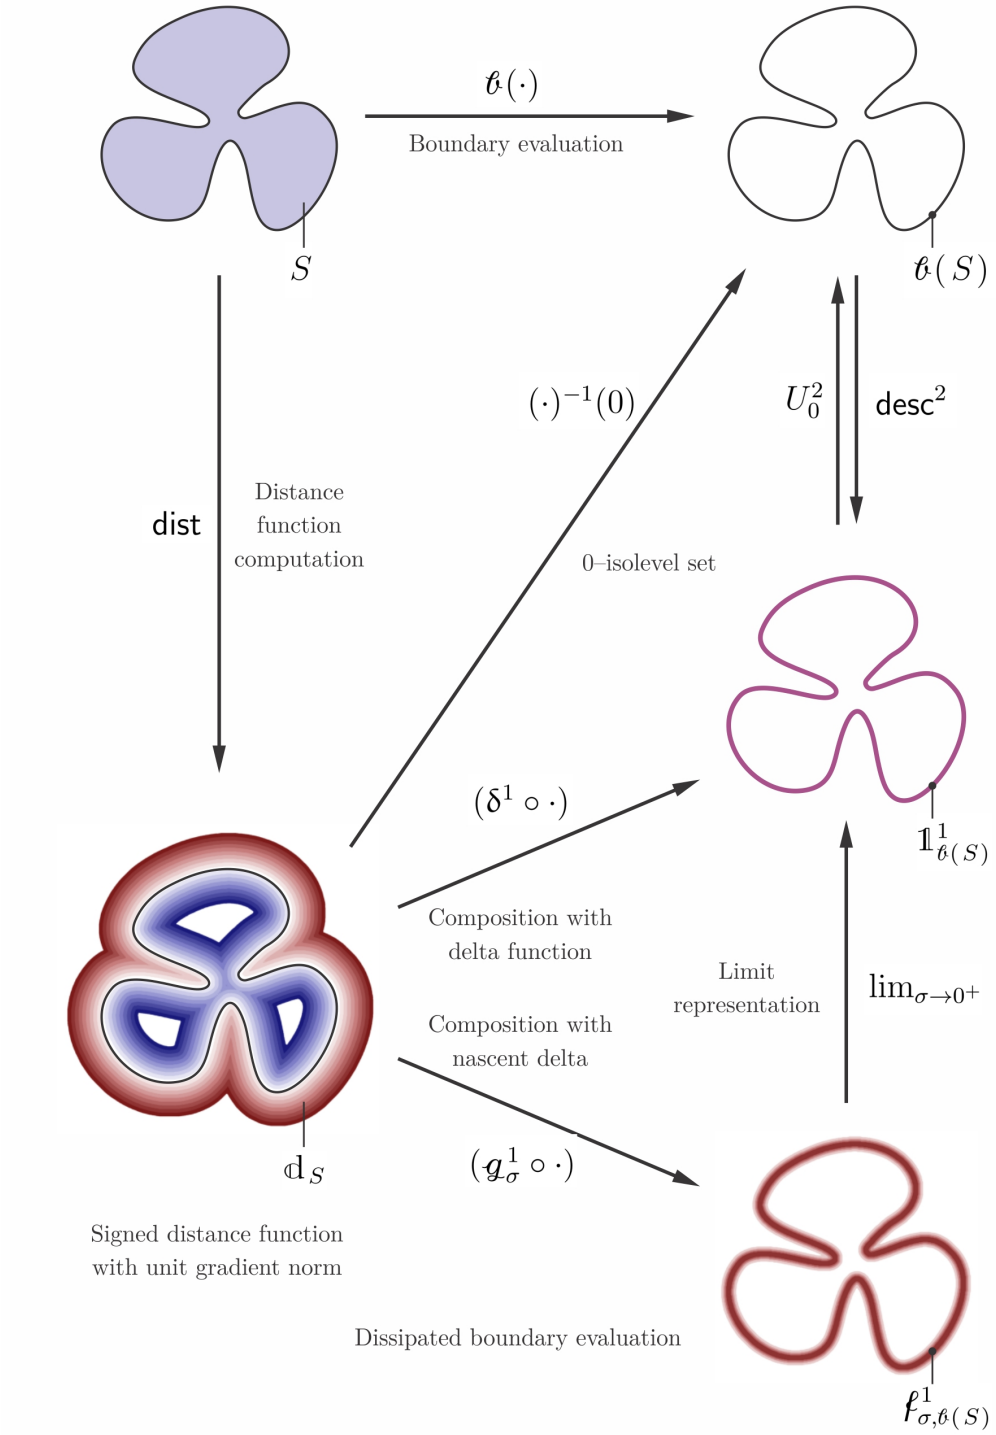

FIGURE 3.3.8: The composition of the signed distance function of a 2D regular shape with the $1^{\text{st}}-$order $\delta-$function yields a $1^{\text{st}}-$order $\delta-$singular boundary descriptor in 2D. Similarly, the same composition yields a $2^{\text{nd}}-$order $\delta-$singular boundary descriptor in 3D.

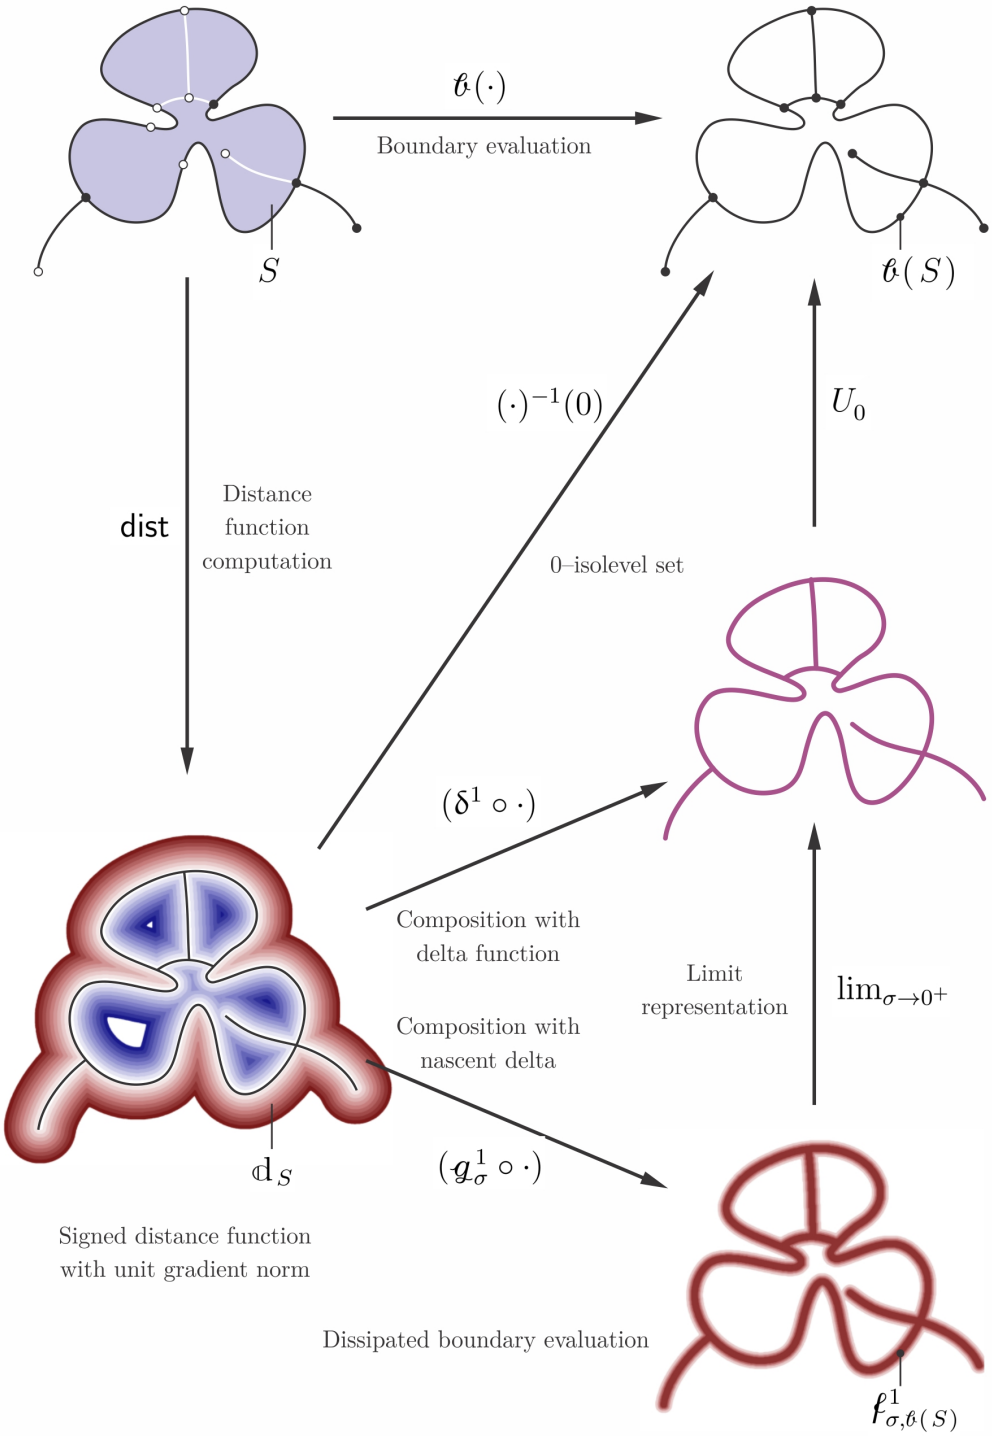

FIGURE 3.3.9: It is conceivable that under certain conditions, the diagram in Fig. 3.3.8 can be extended to general (e.g., irregular) sets with well-defined distance functions, leading to descriptors of different orders of singularity.

Therefore, we have obtained a $\delta-$singular defining function for the boundary that has made it artificially $3-$measurable, and can be used in downstream applications.

The above $\delta-$sifting process of the distance function in (3.3.17) to obtain a $2^{\text{nd}}-$order $\delta-$singular shape descriptor of the boundary (illustrated in Fig. 3.3.8) can be illustrated using the following commutative diagram:

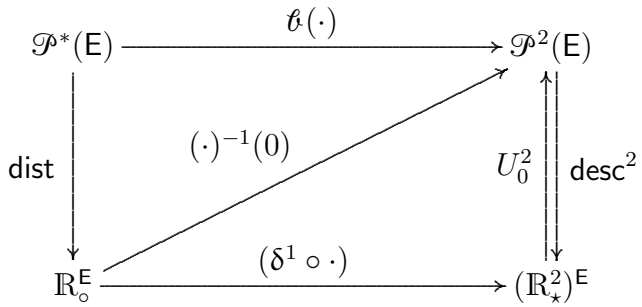

(3.3.23)

Once again, we encounter a situation in which the exact geometry becomes non-computable. Although this may seem like a drawback at the first glance, it demonstrates a fundamental strength of the analytic method in revealing theoretical limitations, and proves to be useful in practice. It is a well-known problem that computations that depend on the boundary of a shape, such as evaluating lower-dimensional intersections, contact surface and length measure, and medial axis transform (MAT) are nonrobust with respect to small geometric perturbations. These problems are mitigated once we reject the idea of crisp lower-dimensional manifolds and work with their dissipated approximations. The approximation can be obtained by using a nascent $\delta-$function (e.g., selected from Table 2.3.1), which can be thought of 'thickening' or 'dissipation' of the boundary into the adjacent volume. In particular, if a compactly supported nascent $\delta-$function (e.g., a mollifier) is used, the composition $(\mathscr{g}^1_\sigma \circ \mathbb{d}_S) : \mathsf{E} \to \mathbb{R}_+$ returns a nonegative real-valued defining function of a thickened boundary that contains the exact boundary $\mathscr{b}(S) \in \mathscr{P}^2(\mathsf{E})$. For example, it can be

shown that if $\mathscr{g}^1_\sigma := \mathbf{1}_{(-\sigma,+\sigma)}/\mu^1[-\sigma,+\sigma]$ is chosen to be the normalized indicator function of a 1D $\sigma-$interval of volume $\mu^1[-\sigma,+\sigma] = 2\sigma$, the result is the defining function of a $\pm\sigma-$offset of the exact boundary: $(\mathscr{g}^1_\sigma \circ \mathbb{d}_S) \equiv f_{(\mathbb{6}^{\uparrow *}_r \circ \mathscr{b})(S)} = f_{\mathscr{b}(S) \oplus^* B_\sigma(S)}$. This can be expressed by the following extension to (3.3.23):

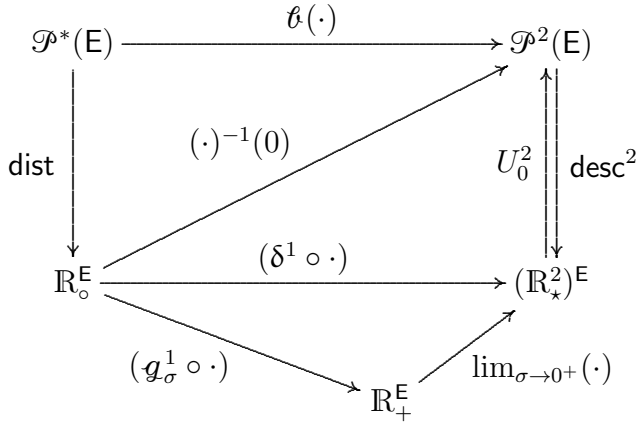

(3.3.24)

In essence, boundary evaluation in the implicit realm is the limit representation of the thin shell obtained by magnifying near-zero distance via $\mathscr{g}^1_\sigma : \mathbb{R} \to \mathbb{R}_+$. In particular, if a simple normalized indicator function $\mathscr{g}^1_\sigma(r) := \frac{1}{2\sigma}\mathbf{1}_{(\sigma,+\sigma)}(r)$ is chosen, the shape descriptor $(\mathscr{g}^1_\sigma \circ \mathbb{d}_S)$ can be viewed as a uniform probability distribution within a $\pm r-$offset tolerance zone around the exact boundary—i.e., $\mathbf{x} \in U_0(\mathscr{g}^1_\sigma \circ \mathbb{d}_S)$ iff $-r < \mathbb{d}_S(\mathbf{x}) < +r$. If a Gaussian kernel $\mathscr{g}^1_\sigma(r) := \frac{1}{(\sqrt{\pi}\sigma)} e^{-r^2/\sigma^2}$ is chosen, it can be thought of as a normal probability distribution, and so on.[14]

I conjecture that (3.3.24) can be extended to boundaries of general semianalytic

[14]This conceptualization of defining functions as probability distributions is particularly attractive as it leads to meaningful interpretations of downstream applications. Particularly, the addition and multiplication rules for testing the inclusion of a point in the union and intersection of two shapes map to the addition and multiplication rules of probability, assuming that inclusion in the two shapes represent two discrete events that are mutually exclusive (i.e., independent). This, in turn, can be used to reason about the probabilistic meaning of convolutions for independent continuum events (e.g., when sweeping a fuzzy shape along a fuzzy surface or curve). See Section 5.1 for an outlook on probabilistic geometric modeling and spatial reasoning.

shapes as long as their lower-dimensional 'dangling' features (e.g., discrete points, curve segments, and surface patches) at the boundary are also captured as the 0−isolevel set of the distance function in (3.3.17):

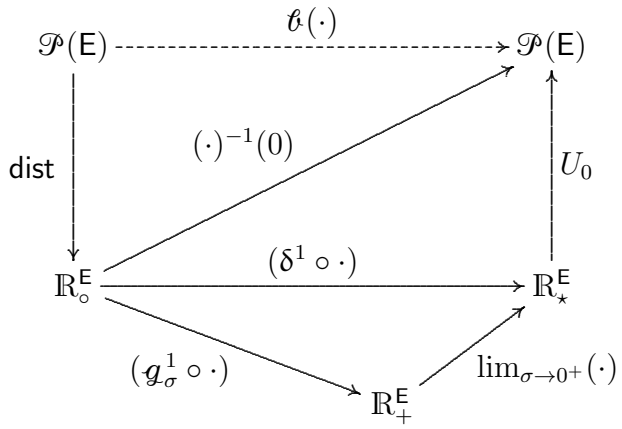

(3.3.25)

The analytic methods of topological regularization and boundary evaluation described in this section for shapes can be extended in principle to motions, by replacing the Lebesgue 3−measure on $\mathsf{E} = \mathbb{R}^3$ with Haar 6−measure on $\mathsf{C} = \mathrm{SE}(3)$. However, defining δ−singularities of various orders over the non-Euclidean nonlinear C−space $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$, despite appearing possible via Lie algebra $\mathfrak{se}(3) \cong \mathfrak{so}(3) \oplus \mathbb{R}^3$, appears far from trivial.

3.4 Morphological Operations

Traditional solid modeling [275, 277, 278] focuses on reasoning about geometric complexities of shapes in the 3−space—with natural generalization to the d−space—while kinematics is typically handled in a more restricted fashion—e.g., with rigid motions restricted to one-parametric trajectories [111,113,161–164] whose geometric complexities do not reach beyond those of piecewise continuous curves in $\mathsf{C} = \mathrm{SE}(3)$.

On the other hand, mathematical morphology, pioneered by Matheron and Serra for image analysis [139, 308] and extended later to group morphology by Roerdink [291, 292] provides a different set of powerful tools to analyze various structures or patterns within images by applying different types of 'filters' or 'structuring elements' defined in terms of Minkowski operators.

Morphological reasoning provides a new set of powerful tools with numerous applications in robotics robotics [88, 174] and solid modeling [215, 218]. Unfortunately, in spite of their versatility, morphological operations have not been widely implemented in solid modeling systems and applications, which can be traced to the computational difficulties that one faces when working in the explicit realm, detailed next.

3.4.1 Explicit Morphology

3.4.1.1 Pointset-Theoretic Euclidean (i.e., Commutative) Morphology:

Classical morphology [139, 308] deals with predominantly translational filters on the 2D plane, in which images and filters are uniformly abstracted as continuous sets in $\mathbb{R}^2$ or their discretized approximations over uniform grids (i.e., 'bitmaps') in $\mathbb{Z}^2$. The natural generalization to the 3D space allows one to study problems that involve shapes in $\mathsf{E} := \mathbb{R}^3$ subjected to translational motions in $\mathsf{T} := \mathrm{T}(3) \cong \mathbb{R}^3$, or their discretized approximations over uniform grids (i.e., 'voxmaps') in $\mathbb{Z}^3$.

A subspace of $\mathscr{P}(\mathsf{E})$ is an 'explicit morphology' on $\mathsf{E} = \mathbb{R}^3$ if it is a (regularized) Boolean algebra, and is also closed under (regularized) Minkowski sum and reflection with respect to the origin.[15] Thus $\mathscr{P}(\mathsf{E})$ and $\mathscr{P}^*(\mathsf{E})$ are both explicit morphologies on

[15]This automatically guarantees closure under Minkowski difference (due to the duality via De Morgan's laws), and derived morphological operations such as dilation/erosion and closing/opening.

the shape space with ordinary and regularized operations, respectively.

Boolean Union and Intersection. As mentioned earlier, the set-theoretic Boolean operations $\cup, \cap : \mathcal{P}(\mathsf{E}) \times \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$ and their regularized counterparts $\cup^*, \cap^* : \mathcal{P}^*(\mathsf{E}) \times \mathcal{P}^*(\mathsf{E}) \to \mathcal{P}^*(\mathsf{E})$ are dual to each via set complement operator $c : \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$ and its regularized counterpart $c^* : \mathcal{P}^*(\mathsf{E}) \to \mathcal{P}^*(\mathsf{E})$, respectively. Composing the diagrams in (2.3.20) and (2.3.22) via De Morgan's laws, one obtains:

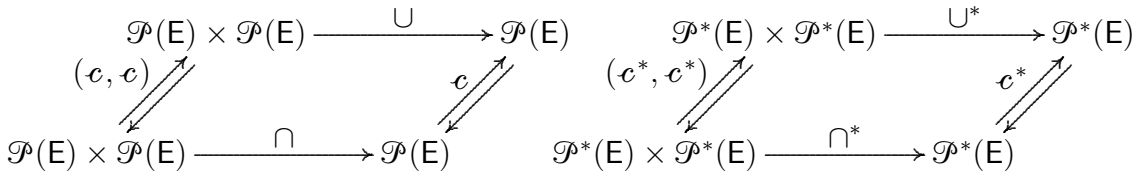

(3.4.1)

The two diagrams above can be connected together via explicit regularization:

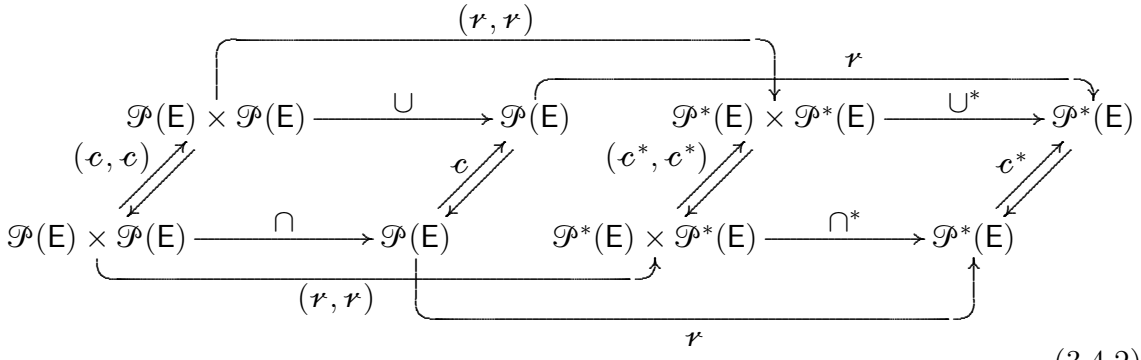

(3.4.2)

Boolean operations were discussed at length in Section 2.2.2.

The classical approach to handling Boolean operations was to use constructive and procedural representation schemes such as constructive solid geometry (CSG) [282] which implement solids as finite sequences of such operations on a set of primitives. The CSG is arguably an implicit representation, as discussed in Section 3.4.2. How-

ever, other representation schemes that are based on embedded manifold cell complexes such as spatial cell decomposition (e.g., volumetric mesh for FEA) and boundary representation (B-rep) require an explicit implementation of Boolean operations, which typically looks something like this:

1. construct a decomposition of the 3−space into a cell complex whose canonical intersection terms (i.e., 'atoms') are obtained from pairwise intersections of the cells of each object's cell complex [314];

2. select the atomic cells that belong to the union or intersection of the given objects by PMC testing a representative point in each cell against each object.

In the case of B-reps, this amounts to computing trimming and merging operations on intersecting boundary elements from each object, whose algorithmic details and computational complexity depends on the type and order of boundary elements.

Note also that such operations are technically ill-defined from a computability perspective [105, 106], as a small perturbation of the boundary can lead to large topological and/or geometric changes in the intersection terms.

Minkowski Sum and Difference. As mentioned in Section 2.2.2, the vector space $\mathsf{E} = \mathbb{R}^3$ together with vector summation and reflection—with respect to an arbitrarily chosen origin $\mathbf{0} \in \mathsf{E}$—forms an Abelian group $\langle \mathsf{E}; +, - \rangle$. The Minkowski operations $\oplus, \ominus : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ can thus be defined as the following dual

combinations of infinite Boolean operations:

$$(S_1 \oplus S_2) := \bigcup_{\mathbf{x}_2 \in S_2} \bigcup_{\mathbf{x}_1 \in S_1} \{\mathbf{x}_1 + \mathbf{x}_2\} = \mathcal{c}(\mathcal{c}(S_1) \ominus S_2), \tag{3.4.3}$$

$$(S_1 \ominus S_2) := \bigcap_{\mathbf{x}_2 \in S_2} \bigcup_{\mathbf{x}_1 \in S_1} \{\mathbf{x}_1 + \mathbf{x}_2\} = \mathcal{c}(\mathcal{c}(S_1) \oplus S_2). \tag{3.4.4}$$

Note that the Minkowski sum is commutative $(S_1 \oplus S_2) = (S_2 \oplus S_1)$ while the Minkowski difference is not $(S_1 \ominus S_2) \neq (S_2 \ominus S_1)$, and the two operators are dual to each other via De Morgan's laws.[16] Note also that containment is preserved under both Minkowski operations in an intuitive sense, i.e., if $S_1 \subseteq S_1'$ and $S_2 \subseteq S_2'$ then $(S_1 \oplus S_2) \subseteq (S_1' \oplus S_2) \subseteq (S_1 \oplus S_2')$ and $(S_1 \ominus S_2') \subseteq (S_1 \ominus S_2) \subseteq (S_1' \ominus S_2)$. Observe also that $(S_1 \ominus S_2) \subseteq S_1 \subseteq (S_1 \oplus S_2)$.

Once the Euclidean metric is introduced (making $\langle \mathsf{E}; \mathrm{d}; +, - \rangle$ a topological group), the regularized Minkowski operations $\oplus^*, \ominus^* : \mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) \to \mathscr{P}^*(\mathsf{E})$ are defined in the same fashion that regularized Boolean operations were defined [343]; namely, as $\odot^* := (\mathcal{r} \circ \odot|_{\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E})})$ in which $\odot \in \{\oplus, \ominus\}$, while many of the basic properties are preserved. It is interesting to note that closed-regular sets are closed under Minkowski sums—which are constructed from unions alone, thus $\oplus^* = \oplus|_{\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E})}$—but are not closed under Minkowski differences—which use both unions and intersections in their definition, meaning that in general $\ominus^* \neq \ominus|_{\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E})}$ and are subject to similar stability issues that were discussed before for computing intersections.

The De Morgan duality on the right-hand sides of (3.4.3) and (3.4.4) can be

[16] Some older texts define the Minkowski difference as $(S_1 \ominus S_2) := (S_1 \oplus (-S_2))$ for obvious notation-driven incentives, instead of $(S_1 \ominus S_2) := \mathcal{c}(\mathcal{c}(S_1) \oplus S_2)$. Inspired by [215, 218], I choose to adopt the latter which is consistent with modern morphology [308] and exhibits a similar type of desirable duality between Minkowski operations that one observes between Boolean operations.

illustrated via the following commutative diagrams:

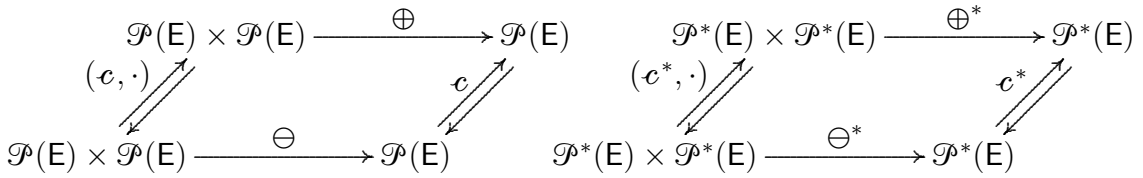

(3.4.5)

Note that unlike (3.4.1) and (3.4.2) for Boolean operations, Minkowski sum and difference of regularized shapes are different from regularization of Minkowski sum and difference, i.e., $\odot \overset{\mathsf{ae}}{\neq} (\odot \circ (\mathit{r}, \mathit{r}))$ therefore $(\mathit{r} \circ \odot) \neq (\odot \circ (\mathit{r}, \mathit{r}))$ for $\odot \in \{\oplus, \ominus\}$. This is because lower-dimensional features (i.e., irregularities) can create full-dimensional features upon these operations, unlike Boolean operations, thus losing them in regularization will alter the result (even beyond equality$-\mathsf{ae}$).

Figures 3.4.1 and 3.4.2 illustrate these operations for simple planar examples. Observe that Minkowski sum and difference operations enable "generalized" *offsetting* for objects of arbitrary shape.

As mentioned in Section 2.2.2, the shape space $\mathsf{E} = \mathbb{R}^3$ is isomorphic to the subgroup $\mathsf{T} := \mathrm{T}(3) \cong \mathbb{R}^3$ of the configuration space $\mathsf{C} = \mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$—or more technically, to the quotient subgroup $\mathrm{T}(3) \cong \mathrm{SE}(3)/\mathrm{SO}(3)$ in which all rotations around the origin are deemed equivalent. The isomorphism is established by a bijective embedding $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$ that maps a point in the shape space to a translation vector in the configuration space with respect to a particular (and arbitrary) choice of an origin $\mathbf{0} \in \mathsf{E}$, thus $\gamma_0(\mathsf{E}) = \mathsf{T}$.

Morphological Dilation and Erosion. Among the fundamental morphological devices—derived directly from γ_0-embeddings of Minkowski sum and difference,

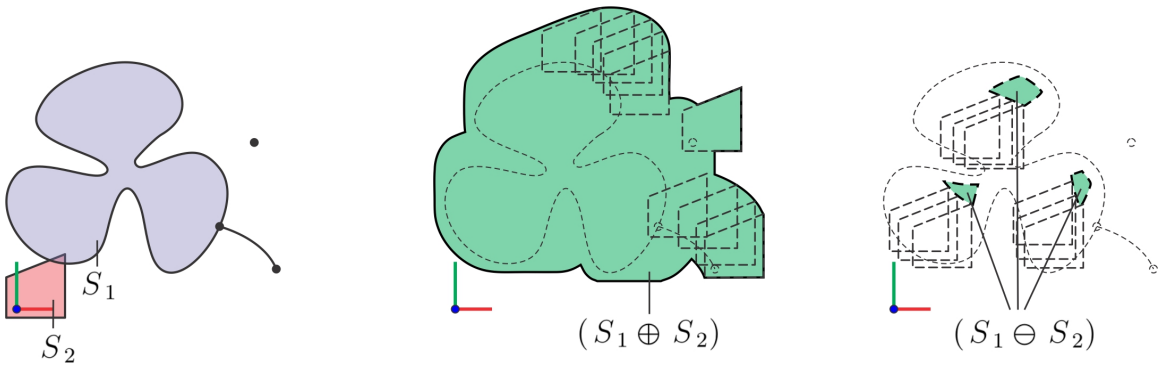

(a) Two Shapes (b) Minkowski Sum (c) Minkowski Difference

FIGURE 3.4.1: The Minkowski sum and difference of two planar shapes. Note that lower-dimensional features can create full-dimensional features regardless of post-regularization.

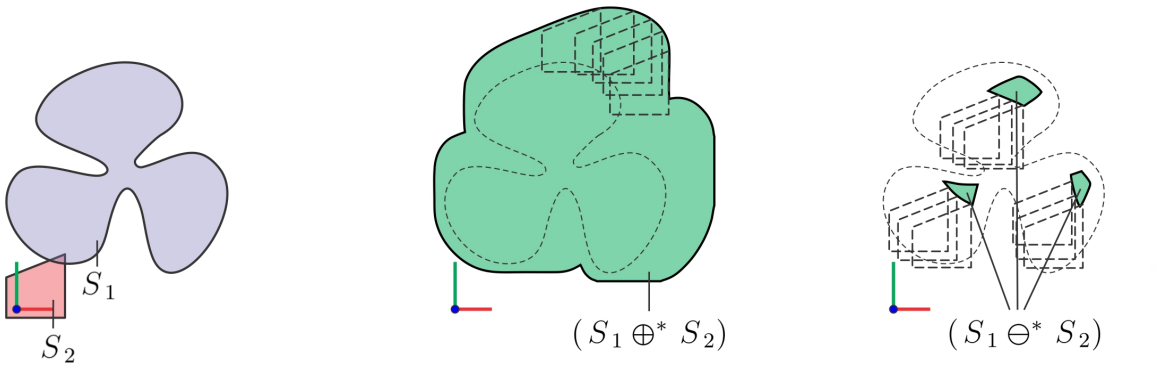

(a) Two Solids (b) Minkowski Sum (c) Minkowski Difference

FIGURE 3.4.2: The Minkowski sum and difference of two planar solids (i.e., full-D shapes).

respectively—are the 'dilation' and 'erosion' operators $\mathsf{dil}, \mathsf{ero} : \mathscr{P}(\mathsf{T}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, which are defined as the following pair of dual functions:

$$\mathsf{dil}\,(T, S) := \bigcup_{\mathbb{t} \in T} \bigcup_{\mathbf{x} \in S} \mathbb{t}^{-1}\mathbf{x} \;= \mathscr{c}\big(\mathsf{ero}(T, \mathscr{c}S)\big) \tag{3.4.6}$$

$$= \big(S \oplus (-\gamma_0^{-1}(T))\big) \;=\; \gamma_0^{-1}\big(T^{-1} \otimes \gamma_0(S)\big), \tag{3.4.7}$$

$$\mathsf{ero}(T, S) := \bigcap_{\mathbb{t} \in T} \bigcup_{\mathbf{x} \in S} \mathbb{t}^{-1}\mathbf{x} \;= \mathscr{c}\big(\mathsf{dil}\,(T, \mathscr{c}S)\big) \tag{3.4.8}$$

$$= \big(S \ominus (-\gamma_0^{-1}(T))\big) \;=\; \gamma_0^{-1}\big(T^{-1} \oslash \gamma_0(S)\big), \tag{3.4.9}$$

where $\mathsf{act}(\mathbb{t}^{-1})(\mathbf{x}) = -\gamma_0^{-1}(\mathbb{t}) + \mathbf{x} = \gamma_0^{-1}(\mathbb{t}^{-1}\gamma_0(\mathbf{x}))$ denotes the inverted group action of $\mathbb{t} \in \mathsf{T}$ on $\mathbf{x} \in \mathsf{E}$, which turns into a vector sum if the former is represented with respect to the origin as $\mathbb{t} \overset{0}{=} \mathbf{t}$ where $\mathbf{t} := \gamma_0^{-1}(\mathbb{t}) \in \mathbb{R}^3$ is called the translation vector. The operators $\otimes, \oslash : \mathscr{P}(\mathsf{T}) \times \mathscr{P}(\mathsf{T}) \to \mathscr{P}(\mathsf{T})$ defined over the translation space are γ_0-isomorphic to $\oplus, \ominus : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ defined over the shape space. One can construct regularized dilation and erosion to achieve closure when working with regular sets. Once again, a commutative diagram can be used to show De Morgan dualities for both ordinary and regularized dilation and erosion operators:

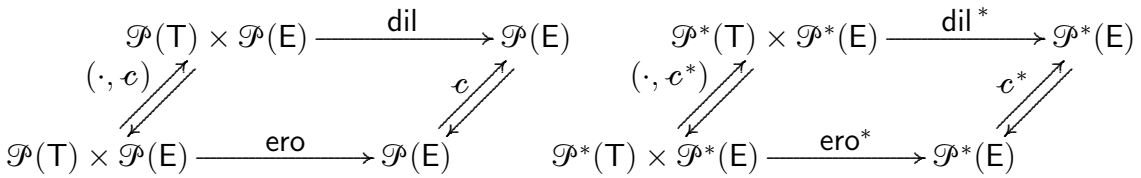

(3.4.10)

Once again, unlike (3.4.1) and (3.4.2) for Boolean operations, morphological dilation and erosion of regularized shapes and motions are different from regularization of dilation and erosion, i.e., $\mathsf{opr} \overset{\mathsf{ae}}{\neq} (\mathsf{opr} \circ (\mathscr{r}, \mathscr{r}))$ therefore $(\mathscr{r} \circ \mathsf{opr}) \neq (\mathsf{opr} \circ (\mathscr{r}, \mathscr{r}))$ for

$\mathsf{opr} \in \{\mathsf{dil}, \mathsf{ero}\}$. This is because lower-dimensional features (i.e., irregularities) can create full-dimensional features upon these operations, unlike Boolean operations, thus losing them in regularization will alter the result (even beyond equality$-\mathsf{ae}$).

The relationships between dilation and erosion and Minkowski operations can be illustrated via the following commutative diagram:

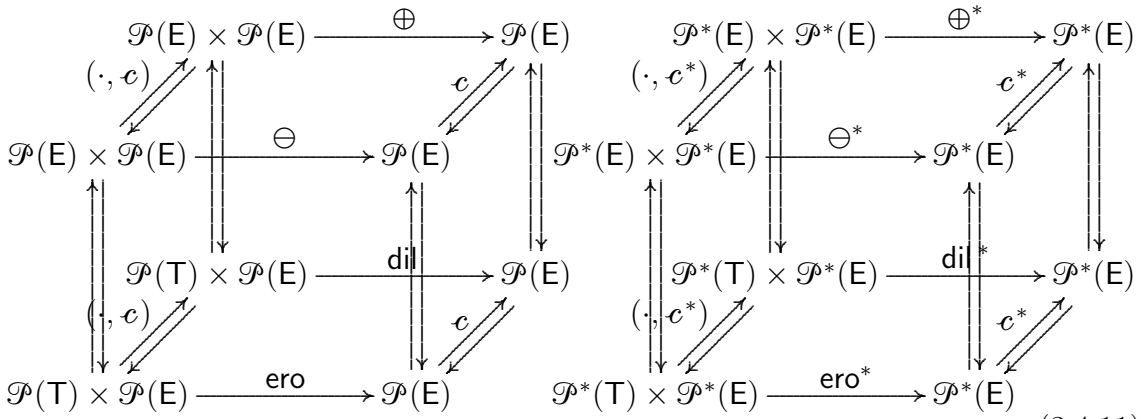

(3.4.11)

in which the labels for the vertical double-arrows should be $(-\gamma_0^{-1}, (\cdot)) \uparrow\downarrow (-\gamma_0, (\cdot))$ and $\gamma_0^{-1} \uparrow\downarrow \gamma_0$, in addition to swapping the order of arguments, omitted from the diagram due to cluttering.

Morphological Closing and Opening. Another pair of important morphological devices—derived from the compositions of dilation and erosion and vice versa, respectively—are the 'closing' and 'opening' operators $\mathsf{cls}, \mathsf{opn} : \mathscr{P}(\mathsf{T}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$,

which are defined as the following pair of dual functions:

$$\mathsf{cls}(T,S) := \mathsf{ero}\big(T, \mathsf{dil}\,(T^{-1},S)\big) = \bigcap_{\mathbb{t}\in T}\bigcup_{\mathbb{t}'\in T}\bigcup_{\mathbf{x}\in S}(\mathbb{t}^{-1}\mathbb{t}')\mathbf{x} = \mathscr{c}\ \mathsf{opn}(T,\mathscr{c}S) \tag{3.4.12}$$

$$= (S \oplus \gamma_0^{-1}(T)) \ominus (-\gamma_0^{-1}(T)) \ = \ \gamma_0^{-1}\big(T^{-1} \oslash (T \otimes \gamma_0(S))\big), \tag{3.4.13}$$

$$\mathsf{opn}(T,S) := \mathsf{dil}\,\big(T, \mathsf{ero}(T^{-1},S)\big) = \bigcup_{\mathbb{t}\in T}\bigcap_{\mathbb{t}'\in T}\bigcup_{\mathbf{x}\in S}(\mathbb{t}^{-1}\mathbb{t}')\mathbf{x} = \mathscr{c}\ \mathsf{cls}(T,\mathscr{c}S) \tag{3.4.14}$$

$$= (S \ominus \gamma_0^{-1}(T)) \oplus (-\gamma_0^{-1}(T)) \ = \ \gamma_0^{-1}\big(T^{-1} \otimes (T \oslash \gamma_0(S))\big), \tag{3.4.15}$$

where $\mathsf{act}(\mathbb{t}^{-1}\mathbb{t}')(\mathbf{x}) = (\gamma_0^{-1}(\mathbb{t}') - \gamma_0^{-1}(\mathbb{t})) + \mathbf{x} = \gamma_0^{-1}((\mathbb{t}^{-1}\mathbb{t}')\gamma_0(\mathbf{x}))$ denotes the displacement of a point $\mathbf{x} \in \mathbb{R}^3$ via translation $(\mathbb{t}^{-1}\mathbb{t}') \overset{0}{=} (\mathbf{t}' - \mathbf{t})$ where $\mathbf{t} := \gamma_0^{-1}(\mathbb{t}) \in \mathbb{R}^3$ and $\mathbf{t}' := \gamma_0^{-1}(\mathbb{t}') \in \mathbb{R}^3$ are translation vectors. Note that the resultant translation vector $(\mathbf{t}' - \mathbf{t}) \in \mathbb{R}^3$ is independent of the choice of origin, which, in turn, makes closing and opening idempotent,[17] unlike dilation and erosion. One can construct regularized closing and opening to achieve closure when working with regular sets.

Once again, a commutative diagram can be used to show De Morgan dualities for both ordinary and regularized closing and opening operators:

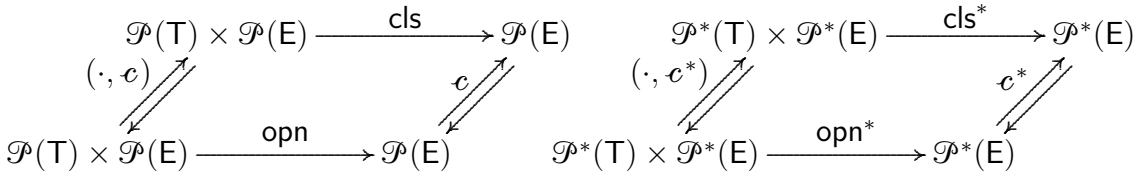

(3.4.16)

Once again, unlike (3.4.1) and (3.4.2) for Boolean operations, morphological closing and opening of regularized shapes and motions are different from regularization of closing and opening, i.e., $\mathsf{opr} \overset{\text{ae}}{\neq} (\mathsf{opr} \circ (\mathscr{r},\mathscr{r}))$ therefore $(\mathscr{r} \circ \mathsf{opr}) \neq (\mathsf{opr} \circ (\mathscr{r},\mathscr{r}))$ for

[17] This means that their subsequent applications using the same filter do not have an affect after the first time, i.e., if $\mathsf{cls}(T,\mathsf{cls}(T,S)) = \mathsf{cls}(T,S)$ and $\mathsf{opn}(T,\mathsf{opn}(T,S)) = \mathsf{opn}(T,S)$.

$\mathsf{opr} \in \{\mathsf{cls}, \mathsf{opn}\}$. This is because lower-dimensional features (i.e., irregularities) can create full-dimensional features upon these operations, unlike Boolean operations, thus losing them in regularization will alter the result (even beyond equality$-\mathsf{ae}$). Note also that because of the two-step nature of closing and opening—each involving two morphological operations: a dilation and an erosion in either order—even when starting from regular shapes and motions, irregularities might appear and disappear between the two steps, which, in turn, may alter the result (even beyond equality$-\mathsf{ae}$).

Minkowski Closing and Opening. Because of the isomorphism between the shape space $\mathsf{E} = \mathbb{R}^3$ and the translational configuration space $\mathsf{T} = \mathrm{T}(3) \cong \mathbb{R}^3$, it is common in classical morphology to ignore the fundamental difference between images (i.e., shapes) and filters (i.e., motions). Thus dilation and erosion are often used interchangeably with Minkowski sum and difference, while closing and opening are defined simply in terms of the compositions of Minkowski operations without bothering with the γ_0-embedding, i.e., denoted by $\oplus, \ominus \colon \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ and defined as the following pair of dual operators:[18]

$$(S_1 \oplus S_2) := (S_1 \oplus S_2) \ominus (-S_2) = \bigcap_{\mathbf{x}_2 \in S_2} \bigcup_{\mathbf{x}_2' \in S_2} \bigcup_{\mathbf{x}_1 \in S_1} \{\mathbf{x}_1 + (\mathbf{x}_2 - \mathbf{x}_2')\}, \tag{3.4.17}$$

$$(S_1 \ominus S_2) := (S_1 \ominus S_2) \oplus (-S_2) = \bigcup_{\mathbf{x}_2 \in S_2} \bigcap_{\mathbf{x}_2' \in S_2} \bigcup_{\mathbf{x}_1 \in S_1} \{\mathbf{x}_1 + (\mathbf{x}_2 - \mathbf{x}_2')\}. \tag{3.4.18}$$

[18] Some older texts define the morphological closing and opening without reflecting the second shape as $(S_1 \oplus S_2) = (S_1 \oplus S_2) \ominus S_2$ and $(S_1 \ominus S_2) = (S_1 \ominus S_2) \oplus S_2$, respectively, which does not strictly satisfy the idempotence property—though it does satisfy a weaker form-idempotence (up to a translation). Confusions like this can be resolved by acknowledging the fundamental difference between shapes and motions [215, 218].

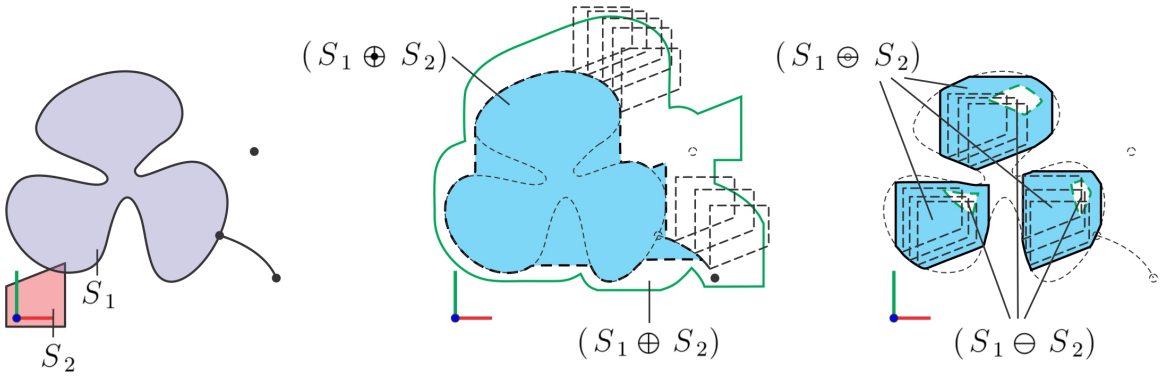

(a) Two Shapes (b) Minkowski Closing (c) Minkowski Opening

FIGURE 3.4.3: The Minkowski closing and opening of two planar shapes. Once again, lower-dimensional features can create full-dimensional features regardless of post-regularization.

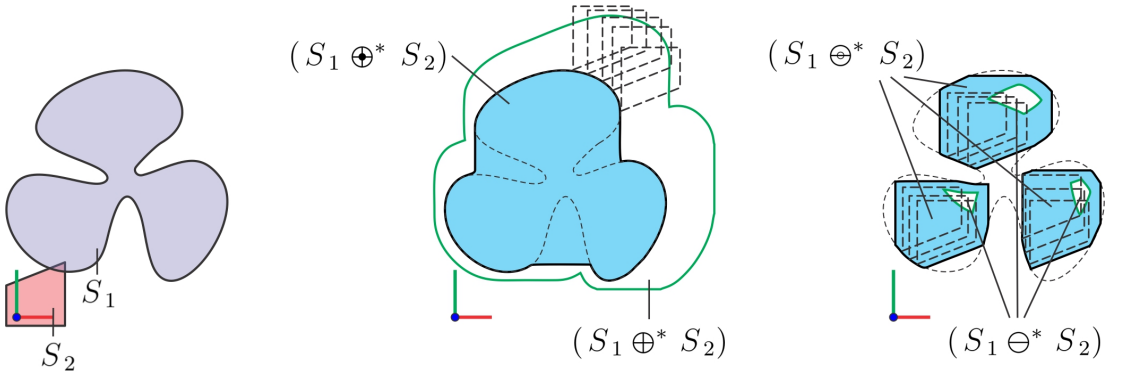

(a) Two Solids (b) Minkowski Closing (c) Minkowski Opening

FIGURE 3.4.4: The Minkowski closing and opening of two planar solids (i.e., full-D shapes).

Note that some of the properties of Minkowski sums and difference are inherited such as containment being analogously preserved, i.e., if $S_1 \subseteq S_1'$ and $S_2 \subseteq S_2'$ then $(S_1 \oplus S_2) \subseteq (S_1' \oplus S_2) \subseteq (S_1 \oplus S_2')$ and $(S_1 \ominus S_2') \subseteq (S_1 \ominus S_2) \subseteq (S_1' \ominus S_2)$. Observe also that $(S_1 \ominus S_2) \subseteq S_1 \subseteq (S_1 \oplus S_2)$.

And of course the Minkowski closing and opening are illustrated as follows:

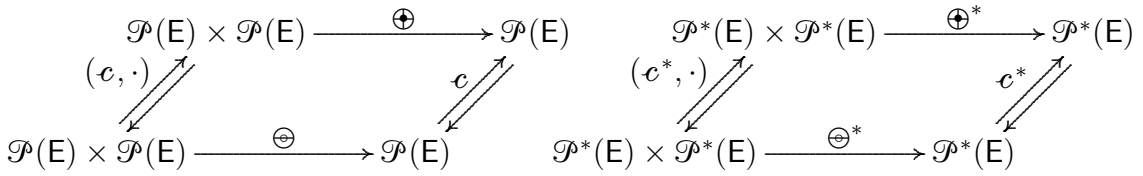

(3.4.19)

Once again, unlike (3.4.1) and (3.4.2) for Boolean operations, Minkowski closing and opening of regularized shapes and motions are different from regularization of closing and opening, i.e., $\odot \overset{\text{ae}}{\neq} (\odot \circ (\mathbf{r}, \mathbf{r}))$ therefore $(\mathbf{r} \circ \odot) \neq (\odot \circ (\mathbf{r}, \mathbf{r}))$ for $\odot \in \{\oplus, \ominus\}$.

Figures 3.4.3 and 3.4.4 illustrate these operations for simple planar examples. Observe that Minkowski closing and opening operations enable "generalized" *blending* for objects of arbitrary shape.

In the following, I present a few special cases by restricting the topology and geometry of the shapes that are important in engineering applications:

Digital Morphology. Given two finite sets of points $S_1^0, S_2^0 \subset \mathsf{E}$ with cardinalities $|S_1^0| = n_1$ and $|S_2^0| = n_2$, respectively, the Minkowski sum $(S_1^0 \oplus S_2^0) \subset \mathsf{E}$ is also a finite set with at most $n_1 n_2$ points, i.e., $|S_1^0 \oplus S_2^0| \leq n_1 n_2$, while the Minkowski difference $(S_1^0 \ominus S_2^0) \subset \mathsf{E}$ is also a finite set with at least $\lceil \frac{n_1}{n_2} \rceil$ points, i.e., $|S_1^0 \ominus S_2^0| \geq \lceil \frac{n_1}{n_2} \rceil$.

If the points are selected from a discrete subgroup of $\mathsf{E} = \mathbb{R}^3$—e.g., if $S_1^0, S_2^0 \subset \mathbb{Z}^3$ thus $(S_1^0 \oplus S_2^0) \subset \mathbb{Z}^3$ and $(S_1^0 \ominus S_2^0) \subset \mathbb{Z}^3$ due to the algebraic closure under vector sums and reflections—which are ϵ_1- and ϵ_2-approximations (in terms of Hausdorff metric) of some continuum shapes $S_1, S_2 \subset \mathsf{E}$, respectively, then it is easily verified that $(S_1^0 \oplus S_2^0)$ and $(S_1^0 \ominus S_2^0)$ are $(\epsilon_1 + \epsilon_2)-$approximations (in terms of Hausdorff metric) of $(S_1 \oplus S_2)$ and $(S_1 \ominus S_2)$, respectively:

Lemma 3.4.1. *The Hausdorff distance*[19] *and Minkowski operations are related as*

$$\mathbb{d}_{\mathrm{H}}(S_1, S_2) < \epsilon \quad \rightleftharpoons \quad S_1 \subseteq (S_2 \oplus B_\epsilon(\mathbf{0})) \text{ and } S_2 \subseteq (S_1 \oplus B_\epsilon(\mathbf{0})), \tag{3.4.20}$$

$$\rightleftharpoons \quad S_1 \supseteq (S_2 \ominus B_\epsilon(\mathbf{0})) \text{ and } S_2 \supseteq (S_1 \ominus B_\epsilon(\mathbf{0})), \tag{3.4.21}$$

for all $S_1, S_2 \subseteq \mathsf{E}$ *and* $\epsilon > 0$*, where* $B_\epsilon(\mathbf{0}) := \{\mathbf{x} \in \mathsf{E} \mid \|\mathbf{x}\|_2 < \epsilon\}$ *is an open* $\epsilon-$*ball.*

The following are immediately resulted if $\mathbb{d}_{\mathrm{H}}(S_1, S_1^0) < \epsilon_1$ and $\mathbb{d}_{\mathrm{H}}(S_2, S_1^0) < \epsilon_2$:

$$(S_1^0 \oplus S_2^0) \subseteq (S_1 \oplus B_{\epsilon_1}) \oplus (S_2 \oplus B_{\epsilon_2}) \tag{3.4.22}$$

$$= (S_1 \oplus S_2) \oplus (B_{\epsilon_1} \oplus B_{\epsilon_2}) = (S_1 \oplus S_2) \oplus B_{\epsilon_1+\epsilon_2}, \tag{3.4.23}$$

$$(S_1 \oplus S_2) \subseteq (S_1^0 \oplus B_{\epsilon_1}) \oplus (S_2^0 \oplus B_{\epsilon_2}) \tag{3.4.24}$$

$$= (S_1^0 \oplus S_2^0) \oplus (B_{\epsilon_1} \oplus B_{\epsilon_2}) = (S_1^0 \oplus S_2^0) \oplus B_{\epsilon_1+\epsilon_2}, \tag{3.4.25}$$

in which the simplified notation $B_\epsilon := B_\epsilon(\mathbf{0})$ is used, noting that the Minkowski sum of two balls is a grown ball of sum of radii, i.e., $(B_{\epsilon_1} \oplus B_{\epsilon_2}) = B_{\epsilon_1+\epsilon_2}$, containment is preserved under Minkowski sums, i.e., $S \subseteq S' \;\Rightarrow\; (S \oplus B_\epsilon) \subseteq (S' \oplus B_\epsilon)$, and that Minkowski sum are associative and commutative. Therefore, one obtains $\mathbb{d}_{\mathrm{H}}\big((S_1 \oplus S_2), (S_1^0 \oplus S_2^0)\big) < (\epsilon_1 + \epsilon_2)$. Similar relationships can be asserted for approximating digitizing the Minkowski difference by noting the dualities of the form $(\boldsymbol{c}(S_1) \ominus S_2) = \boldsymbol{c}(S_1 \oplus S_2) = \boldsymbol{c}(S_2 \oplus S_1) = (\boldsymbol{c}(S_2) \ominus S_1)$, that containment is preserved under Minkowski difference as well, and that Minkowski difference is associative

[19] The symmetric Hausdorff distance between two shapes $\mathbb{d}_{\mathrm{H}} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$—which is a metric over $\mathscr{P}(\mathsf{E})$—is defined as the maximum of the left- and right-Hausdorff distances (not metrics) $\mathbb{d}_{\mathrm{H}}^{\mathrm{L}}, \mathbb{d}_{\mathrm{H}}^{\mathrm{R}} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$, i.e.,

$$\left.\begin{array}{l} \mathbb{d}_{\mathrm{H}}^{\mathrm{L}}(S_1, S_2) := \sup_{\mathbf{x}_1 \in S_1} \inf_{\mathbf{x}_2 \in S_2} \mathbb{d}(\mathbf{x}_1, \mathbf{x}_2) \\ \mathbb{d}_{\mathrm{H}}^{\mathrm{R}}(S_1, S_2) := \sup_{\mathbf{x}_2 \in S_2} \inf_{\mathbf{x}_1 \in S_1} \mathbb{d}(\mathbf{x}_1, \mathbf{x}_2) \end{array}\right\} \text{ and } \mathbb{d}_{\mathrm{H}}(S_1, S_2) := \max\{\mathbb{d}_{\mathrm{H}}^{\mathrm{L}}(S_1, S_2), \mathbb{d}_{\mathrm{H}}^{\mathrm{R}}(S_1, S_2)\}.$$

(though not commutative). Thus one obtains $\mathbb{d}_{\mathrm{H}}\big((S_1 \ominus S_2), (S_1^0 \ominus S_2^0)\big) < (\epsilon_1 + \epsilon_2)$.

Therefore, digital morphology [139, 308] approximately implements continuum morphology, as long as the sampling is equispaced (i.e., over a uniform grid). The above relationships will be extended to incorporate grid-free nonequispaced samples and nonequiradius balls in Chapter 4 to obtain a generalized spherical discretization scheme for analytic methods [36] that tremendously improves over the performance of uniform sampling implementation strategies.

Groupings as Dilations. Given a finite set of points $S^0 = \{\mathbf{x}_i\}_{0 \leq i < n_0} \subset \mathsf{E}$ with $|S^0| = n_0$, and an (closed or open) $\epsilon-$ball $B_r(\mathbf{0})$, the Minkowski sum $S^0 \oplus B_\epsilon(\mathbf{0})$ is a grouping, i.e., a finite union of (possibly overlapping) equiradius balls:

$$\mathfrak{G}_\epsilon^{\uparrow}(S^0) := S^0 \oplus B_\epsilon(\mathbf{0}) = \bigcup_{0 \leq i < n_0} B_\epsilon(\mathbf{x}_i) = \bigcup_{0 \leq i < n_0} \bigcup_{\mathbf{n} \in \mathbb{S}^2} \bigcup_{r \lessdot \epsilon} \{\mathbf{x}_i + r\mathbf{n}\}, \tag{3.4.26}$$

where $B_\epsilon(\mathbf{x}_i) := \big\{\mathbf{x} \in \mathsf{E} \mid \|\mathbf{x} - \mathbf{x}_i\|_2 \lessdot \epsilon\big\}$, denotes an (closed or open) $\epsilon-$ball centered at $\mathbf{x}_i \in \mathsf{E}$ and $\mathbb{S}^2 = \mathscr{b}(B_r(\mathbf{0}))\big|_{r:=1}$ is the unit 2−sphere centered at the origin.[20]

The grouping is said to be an $\epsilon-$approximate 'cover' (or $\epsilon-$cover for short) of a 3D regular shape $S \in \mathscr{P}^*(\mathsf{E})$ if $\mathbb{d}_{\mathrm{H}}\big(S, (S^0 \oplus B_\epsilon)\big) < \epsilon$. Applying Lemma 3.4.1 yields $(S \oplus B_\epsilon) \subseteq (S^0 \oplus B_\epsilon)$ thus $(S \oplus B_\epsilon) \subseteq (S^0 \oplus B_\epsilon)$ due to the fact that containment is preserved under Minkowski difference. But $(S^0 \oplus B_\epsilon) = ((S^0 \oplus B_\epsilon) \ominus B_\epsilon)$ by definition in (3.4.17) and noting that $B_\epsilon = -B_\epsilon$ due to symmetry. The latter is in general heterogeneous in dimensions and contains the original discrete pointset, i.e., $S^0 \subseteq (S^0 \oplus B_\epsilon)$. In particular, the points from the discrete pointset S^0 are classified

[20]The notation $r \lessdot \epsilon$ is used to represent either $r \leq \epsilon$ or $r < \epsilon$ depending on whether closed or open balls are used, respectively.

against the closing $(S^0 \oplus B_\epsilon)$ as follows:

- The internal points in $S_0 \cap \mathit{i}\,(S^0 \oplus B_\epsilon)$ are centers of the balls in the grouping that are 'stably overlapped'—meaning that all corresponding balls are completely covered by the union of other balls in the grouping and slight perturbations of the center positions cannot change the inclusion.

- The boundary points in $S_0 \cap \mathit{b}(S^0 \oplus B_\epsilon)$ are centers of the balls in the grouping that are either 'critically overlapped'—meaning that some balls are completely covered by the union of other balls in the grouping but slight perturbations of the center position can change it—or 'exposed'—meaning that parts of some balls are not intersected by any other ball in the grouping.

The distinction between the latter two conditions can be decided by the dimensionality of the portion of the boundary that includes each ball center, i.e., based on whether the ball center gets eliminated from the closing after closed-regularization. Thus the points from the discrete pointset S^0 are classified against the regularized closing $(S^0 \oplus^* B_\epsilon)$ as follows:

- The internal points in $S_0 \cap \mathit{i}\,(S^0 \oplus^* B_\epsilon) = S_0 \cap \mathit{i}\,(S^0 \oplus B_\epsilon)$ are centers of the balls in the grouping that are 'stably overlapped', as before.

- The boundary points in $S_0 \cap \mathit{b}(S^0 \oplus^* B_\epsilon) \subseteq S_0 \cap \mathit{b}(S^0 \oplus B_\epsilon)$ are centers of the balls in the grouping that are 'critically overlapped'. These points belonged to lower-dimensional portions of the boundary in nonregularized closing that were connected to an interior region.

- The external points in $S_0 \cap \mathit{e}(S^0 \oplus^* B_\epsilon) \supseteq S_0 \cap \mathit{e}(S^0 \oplus B_\epsilon)$ are centers of the

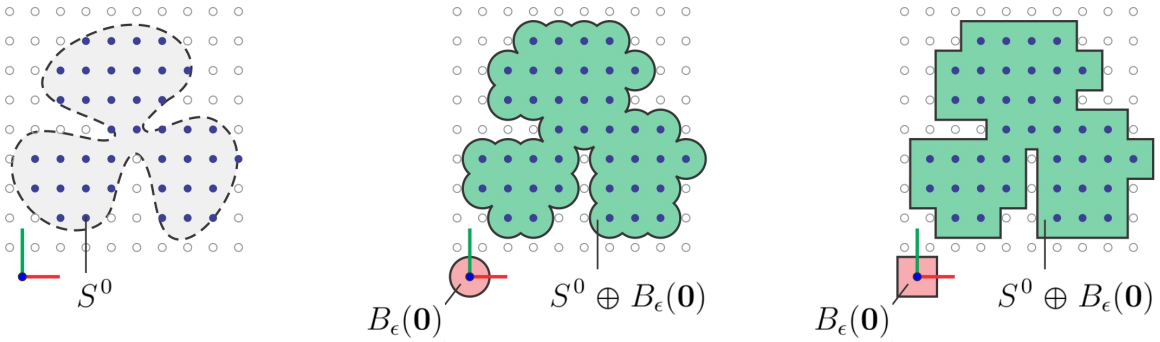

(a) Discrete Pointset (b) Grouping of Balls (c) Grouping of Cubes

FIGURE 3.4.5: A discrete pointset that is an $\epsilon-$sample of a regular shape in (a) is Minkowski summed with an $\epsilon-$ball (e.g., an L^2-ball in (b) or $L^\infty-$ball in (c)) centered at the origin $\mathbf{0} \in \mathbb{R}^2$ to obtain an $\epsilon-$cover of the regular shape.

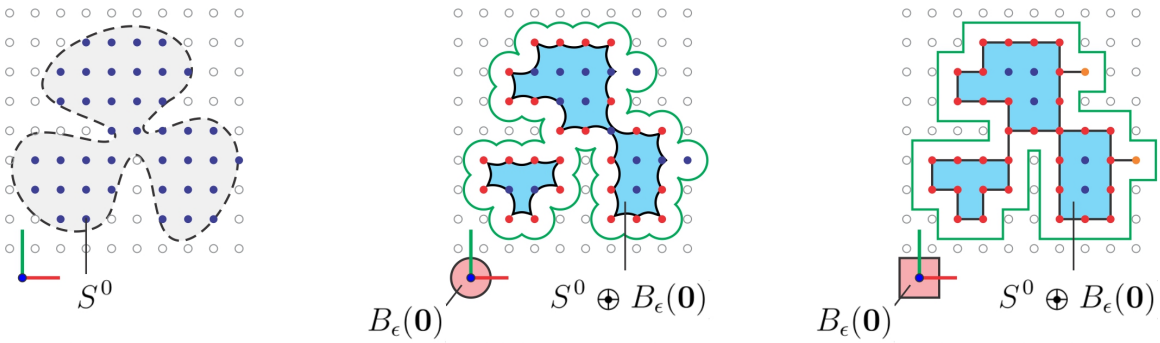

(a) Discrete Pointset (b) Grouping of Balls (c) Grouping of Cubes

FIGURE 3.4.6: The Minkowski closing of an $\epsilon-$cover grouping of $\epsilon-$balls with the $\epsilon-$ball (e.g., an L^2-ball in (b) or $L^\infty-$ball in (c)) returns a superset of the original $\epsilon-$sample in (a). The PMC of the original sample points against this superset (or its regularization) can be used to determine which points are overlapped, critically overlapped, or exposed.

balls in the grouping that are 'exposed', as expected. These points belonged to isolated points of the boundary in nonregularized closing.

An exact explicit solution can be obtained by intersecting the spherical surfaces and identifying individual exposed patches [287], which is expensive. These operations are useful in modeling molecular structures, particularly for characterizing 'core' and 'skin' atoms in for protein docking [25–27, 69].

Figures 3.4.5 and 3.4.6 illustrate how an $\epsilon-$cover can be generated via Minkowski sum of an $\epsilon-$sample with an $\epsilon-$ball (e.g., an L^2-ball or $L^\infty-$ball) and how a point membership classification (PMC) of the sample points against the Minkowski closing can be used to determine overlapped, critically overlapped, and exposed sample points, respectively, for a simple planar example.

Generalized Offsetting. Given an (at most 3D) shape $S \in \mathcal{P}(\mathsf{E})$, which may or may not be regular, and a (closed or open) $r-$ball $B_r := B_r(\mathbf{0})$ centered at the origin, the dilation $(S \oplus B_r)$ and erosion $(S \ominus B_r)$ give positive and negative $r-$*offsets* of the shape, respectively, which bound it as $(S \ominus B_r) \subseteq S \subseteq (S \oplus B_r)$. These envelops 'grow' and 'shrink' the object by $\pm r$ via $\mathcal{O}_r^{\uparrow}, \mathcal{O}_r^{\downarrow} : \mathcal{P}(\mathsf{E}) \times \mathbb{R}_+ \to \mathcal{P}(\mathsf{E})$:

$$\mathcal{O}_r^{\uparrow}(S) := S \oplus B_r(\mathbf{0}) = \bigcup_{\mathbf{n}\in\mathbb{S}^2} \bigcup_{s<r} \bigcup_{\mathbf{x}\in S} \{\mathbf{x} + s\mathbf{n}\} = \complement\mathcal{O}_r^{\downarrow}(\complement S), \tag{3.4.27}$$

$$\mathcal{O}_r^{\downarrow}(S) := S \ominus B_r(\mathbf{0}) = \bigcap_{\mathbf{n}\in\mathbb{S}^2} \bigcap_{s<r} \bigcup_{\mathbf{x}\in S} \{\mathbf{x} + s\mathbf{n}\} = \complement\mathcal{O}_r^{\uparrow}(\complement S), \tag{3.4.28}$$

whose boundaries correspond to the $(\pm r)-$level sets of the signed distance function defined earlier via (3.1.2). The De Morgan duality relationships for offsetting and its

regularized counterparts $\mathcal{O}_r^{\uparrow *}, \mathcal{O}_r^{\downarrow *} : \mathscr{P}^*(\mathsf{E}) \times \mathbb{R}_+ \to \mathscr{P}^*(\mathsf{E})$ are illustrated via:

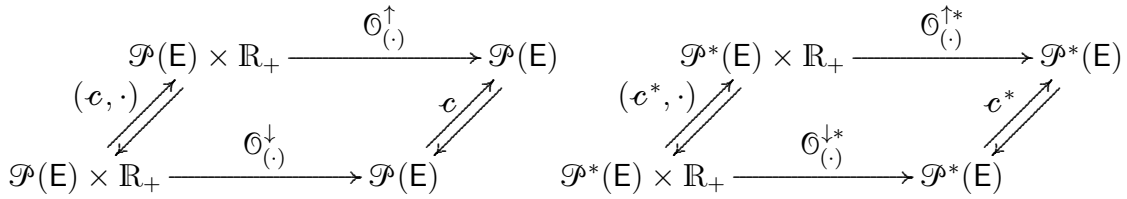

(3.4.29)

As was the case with Minkowski sum and difference in general, $\mathsf{opr} \overset{\mathsf{ae}}{\neq} (\mathsf{opr} \circ \boldsymbol{r})$ therefore $(\boldsymbol{r} \circ \mathsf{opr}) \neq (\mathsf{opr} \circ \boldsymbol{r})$ for $\mathsf{opr} \in \{\mathcal{O}_r^{\uparrow}, \mathcal{O}_r^{\downarrow}\}$ for all $r > 0$. This means that lower-dimensional irregularities will create full-dimensional features upon offsetting, whose elimination before offsetting will change the result. This is not to mention that negative offsetting can also create lower-dimensional features from full-dimensional initial shapes. See Figs. 3.4.7 and 3.4.8 for an illustration.

Offsetting preserves convexity, and offsets of regular sets for a given radius are unique [294], but they are not injective (thus cannot be inverted).

Offsetting (e.g., growing and shrinking) are fundamental geometric modeling operations [257,279,294,295,298,301] and among the basic tools in every solid modeling system's feature set. Other than the obvious applications such as tolerance analysis and clearance testing, cam and follower design, chip design-rule checking (DRC), and safe-margin robot path planning, offsets have been extensively used for CNC cutter path generation in general (reviewed in [101, 141]).

Generalized Blending. Given an (at most) 3D shape $S \in \mathscr{P}(\mathsf{E})$, which may or may not be regular, and a (closed or open) $r-$ball $B_r := B_r(\mathbf{0})$ centered at the origin, the closing $(S \oplus B_r)$ and opening $(S \ominus B_r)$ give positive and negative $r-$*blends* of the shape, respectively, which bound it as $(S \ominus B_r) \subseteq S \subseteq (S \oplus B_r)$. These envelopes

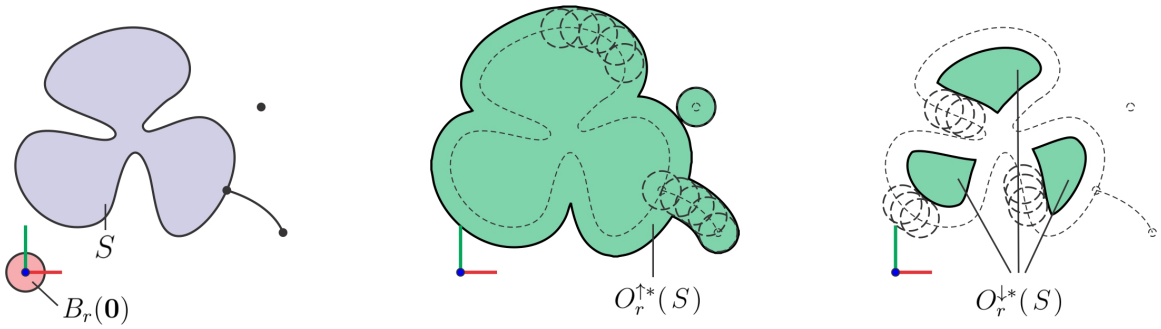

(a) A Shape and a Ball (b) Positive Offset (c) Negative Offset

FIGURE 3.4.7: Minkowski sum and difference of a shape with an $r-$ball give $\pm r-$offsets.

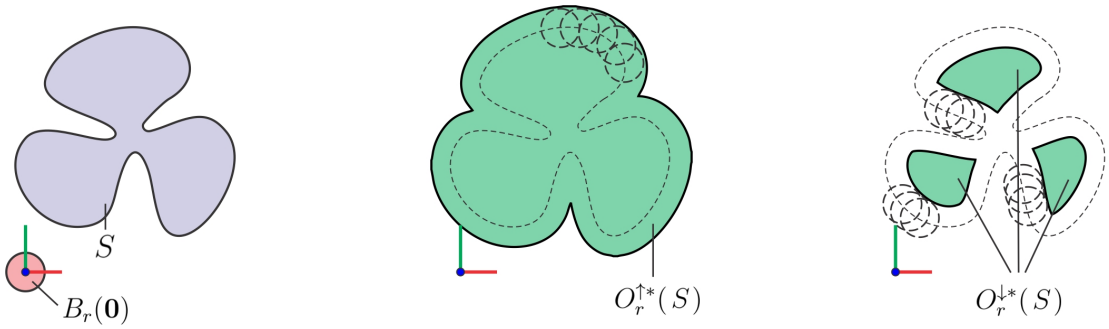

(a) A Solid and a Ball (b) Positive Offset (c) Negative Offset

FIGURE 3.4.8: Minkowski sum and difference of a solid with an $r-$ball give $\pm r-$offsets.

'fillet' and 'round' the corners by $\pm r$ via $\mathscr{R}_r^{\uparrow}, \mathscr{R}_r^{\downarrow} : \mathscr{P}(\mathsf{E}) \times \mathbb{R}_+ \to \mathscr{P}(\mathsf{E})$:

$$\begin{aligned}\mathscr{R}_r^{\uparrow}(S) :=& S \oplus B_r(\mathbf{0}) = (S \oplus B_r(\mathbf{0})) \ominus B_r(\mathbf{0}) = (\mathscr{O}_r^{\downarrow}(\mathscr{O}_r^{\uparrow}(S)) \\ &= \bigcap_{\mathbf{n}\in\mathbb{S}^2} \bigcap_{s\lessdot r} \bigcup_{\mathbf{n}'\in\mathbb{S}^2} \bigcup_{s'\lessdot r} \bigcup_{\mathbf{x}\in S} \{\mathbf{x} + (s\mathbf{n} - s'\mathbf{n}')\} = \mathscr{c}\mathscr{R}_r^{\downarrow}(\mathscr{c}S), \end{aligned} \tag{3.4.30}$$

$$\begin{aligned}\mathscr{R}_r^{\downarrow}(S) :=& S \ominus B_r(\mathbf{0}) = (S \ominus B_r(\mathbf{0})) \oplus B_r(\mathbf{0}) = (\mathscr{O}_r^{\uparrow}(\mathscr{O}_r^{\downarrow}(S)) \\ &= \bigcup_{\mathbf{n}\in\mathbb{S}^2} \bigcup_{s\lessdot r} \bigcap_{\mathbf{n}'\in\mathbb{S}^2} \bigcap_{s'\lessdot r} \bigcup_{\mathbf{x}\in S} \{\mathbf{x} + (s\mathbf{n} - s'\mathbf{n}')\} = \mathscr{c}\mathscr{R}_r^{\uparrow}(\mathscr{c}S), \end{aligned} \tag{3.4.31}$$

which eliminate features that are smaller than r—where local feature size is defined as the minimum distance to the medial axis (MA) defined in (3.5.6) of Section 3.5.1.1. The De Morgan duality relationships for blending and its regularized counterparts $\mathscr{R}_r^{\uparrow *}, \mathscr{R}_r^{\downarrow *} : \mathscr{P}^*(\mathsf{E}) \times \mathbb{R}_+ \to \mathscr{P}^*(\mathsf{E})$ are illustrated via:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{\mathscr{R}_{(\cdot)}^{\uparrow}} & \mathscr{P}(\mathsf{E}) \\ {\scriptstyle (\mathscr{c},\cdot)}\uparrow\downarrow & & \uparrow\downarrow{\scriptstyle \mathscr{c}} \\ \mathscr{P}(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{\mathscr{R}_{(\cdot)}^{\downarrow}} & \mathscr{P}(\mathsf{E}) \end{array} \qquad \begin{array}{ccc} \mathscr{P}^*(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{\mathscr{R}_{(\cdot)}^{\uparrow *}} & \mathscr{P}^*(\mathsf{E}) \\ {\scriptstyle (\mathscr{c}^*,\cdot)}\uparrow\downarrow & & \uparrow\downarrow{\scriptstyle \mathscr{c}^*} \\ \mathscr{P}^*(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{\mathscr{R}_{(\cdot)}^{\downarrow *}} & \mathscr{P}^*(\mathsf{E}) \end{array} \tag{3.4.32}$$

As was the case with Minkowski closing and opening in general, $\mathsf{opr} \overset{\mathsf{ae}}{\neq} (\mathsf{opr} \circ \mathscr{r})$ therefore $(\mathscr{r} \circ \mathsf{opr}) \neq (\mathsf{opr} \circ \mathscr{r})$ for $\mathsf{opr} \in \{\mathscr{R}_r^{\uparrow}, \mathscr{R}_r^{\downarrow}\}$ for all $r > 0$. This means that lower-dimensional irregularities will create full-dimensional features upon offsetting, whose elimination before offsetting will change the result. Once again, irregularities can appear and disappear in the intermediate steps (due to negative offsetting) even when starting from regular shapes. See Figs. 3.4.9 and 3.4.10 for an illustration.

Like offsetting, blending amounts to a loss of information (i.e., is not invertible).

It is rather idempotent, and can be composed from offsetting [294] as

$$\mathscr{R}_r^{\uparrow}(S) := (\mathscr{O}_r^{\downarrow} \circ \mathscr{O}_r^{\uparrow})(S) = (\mathscr{R}_r^{\uparrow} \circ \mathscr{R}_r^{\uparrow})(S) = (\mathscr{R}_r^{\uparrow} \circ \mathscr{R}_r^{\uparrow} \circ \cdots)(S), \tag{3.4.33}$$

$$\mathscr{R}_r^{\downarrow}(S) := (\mathscr{O}_r^{\uparrow} \circ \mathscr{O}_r^{\downarrow})(S) = (\mathscr{R}_r^{\downarrow} \circ \mathscr{R}_r^{\downarrow})(S) = (\mathscr{R}_r^{\downarrow} \circ \mathscr{R}_r^{\downarrow} \circ \cdots)(S). \tag{3.4.34}$$

Blending (e.g., filleting and rounding) are fundamental geometric modeling operations [66,294–297,363] and among the basic tools in every solid modeling system's feature set. Besides the obvious applications in form generation and conceptual design, they are extremely useful in accessibility analysis, design feedback and correction, noise elimination and surface smoothing, and aesthetics.

Although in traditional solid modeling, blending is defined for regular sets (Fig. 3.4.10), as defined in (3.4.30) and (3.4.31), it can be generalized to general shapes of arbitrary homogeneous dimensions (Fig. 3.4.9).

The relationship between offsetting and blending can be captured by a single commutative diagram of the following form:

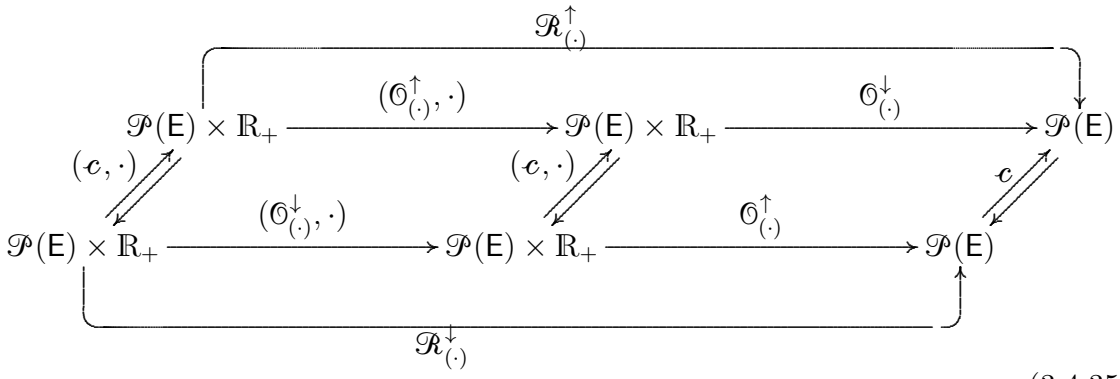

(3.4.35)

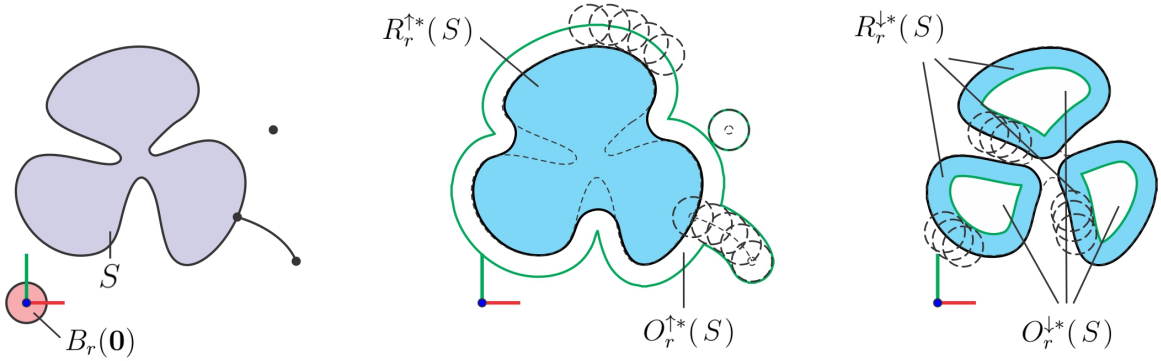

(a) A Shape and a Ball (b) Positive Blend (c) Negative Blend

FIGURE 3.4.9: Minkowski closing and opening of a shape with an $r-$ball give $\pm r-$blends.

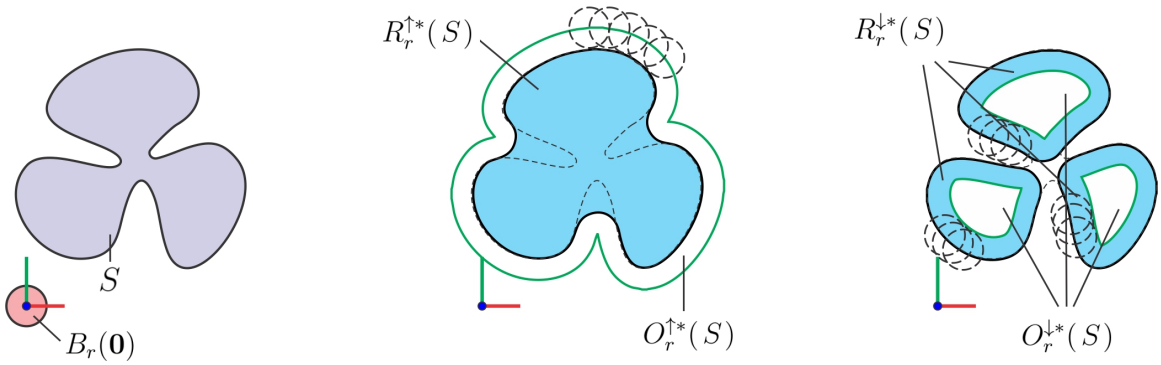

(a) A Solid and a Ball (b) Positive Blend (c) Negative Blend

FIGURE 3.4.10: Minkowski closing and opening of a solid with an $r-$ball give $\pm r-$blends.

and a similar diagram can be constructed for regularized shapes and operations:

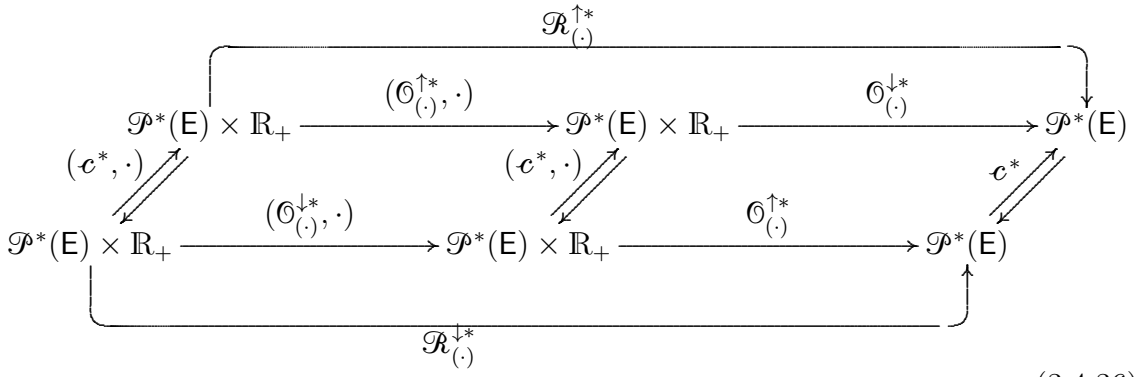

(3.4.36)

Properties of offsets of general curves and surfaces have been studied extensively (reviewed in [257]). Although from an algebraic perspective, offsetting appears to be more fundamental from which blending is derived, the order is typically reversed for computing. An explicit approach to computing offsets typically consists of:

1. offsetting the individual faces of the solid's boundary (e.g., using B-reps) or building blocks (e.g., using CSG);

2. replacing edges at the corners with 'canal' or 'tubular' surfaces—i.e., envelopes of families of balls; and

3. sticking the offset and canal surfaces along the edges and constructing a combinatorial structure;

the last two of which are blending. Often times, intersection curves are approximated by piecewise-linear or circular curves and canal surfaces are approximated by piecewise-cylindrical or toroidal surfaces to improve the numerical performance [294]. Offsets of planar shapes are much easier to compute and are particularly useful for automating contour-parallel pocket milling [70, 144, 167]. Existing explicit meth-

ods typically use Voronoi diagrams for polygonal shapes [142, 186] as well as curved shapes [143], winding numbers [62] or combinations of the two [55]. Although these approached work well for designed geometry with simple surface elements, it produces excessive overhead when dealing with tessellated geometry and noisy data.

Another ongoing research track that has significantly contributed to the developments is in molecular biology, where positive offsetting (i.e., growing) and external blending (i.e., filleting) are the basic tools to generate solvent accessible surface (SAS) and solvent excluded surface (SES) structures defined by Lee and Richards [200, 287] and are fundamental to a myriad of applications in the analysis of protein structure and function [256]. Numerous data structures and algorithms have been developed to compute mesh approximations of protein surfaces among other explicit methods.

Generalized offsetting and blending amount to changing the structuring element (i.e., filter) from a ball to some other virtual 'sculpting' tool (e.g., a cube) in the corresponding morphological definitions. In other words, mathematical morphology provides an 'algebra of manufacturing' that can uniformly model complex subtractive (e.g., 5-axis CNC machining) and additive (e.g., FDM 3D printing) manufacturing processes and tie them back to design.[21]

Explicit computation of general Minkowski sums are very costly, and is usually

[21] At the time when Voelcker and Requicha established the foundations of solid modeling back in 1970s [351, 352], basic CSG with Boolean operations [276, 282] was appealing due to the intuitive mathematical language it provided to think about material removal processes in NC machining.

The basic CSG was later augmented with filleting and blending operations by Rossignac and Requicha [294–297]. Recent proposals by Lysenko et al. [215,218,235] can be interpreted as augmenting the CSG with generalizations of offsetting and blending (i.e., Minkowski operations), and using the more powerful language of mathematical morphology to think about today's complex material removal and deposition processes. This proposal contributes to the mix by augmenting the existing toolset with singularity modeling for heterogeneous dimensions. Noting that CSG is itself an implicit representation at an abstract level [316], one way to think about this school of thought is bringing it back and extending it to formalize geometric modeling into a "functional" style [22]. That being said, B-reps are not going away anytime soon, in my opinion, and every new proposal (including this one) needs to be able to provide tools to interoperate with existing B-rep technology.

limited to 2D polygons [8, 201, 366] or (mostly convex) 3D polyhedra [30, 119, 366]. Minkowski sums preserve convexity, which is a very helpful property that allows using convex hulls to compute the polyhedra B-rep of Minkowski sums from those of constituent solids. They are also distributive over unions, allowing a divide-and-conquer approach to solving nonconvex problems. Therefore, a typical explicit process for computing Minkowski operations looks like this:

1. approximating the constituent parts as unions (e.g., cell decompositions) of convex polyhedra;

2. finding the vertex sets of individual pairs of polyhedra;

3. adding up the position coordinates of pairs of vertices in each pair of polyhedra to obtain the vertex set of their pairwise Minkowski sums;

4. constructing the convex hulls of obtained vertex sets; and

5. computing the unions of the convex hulls to obtain the result.

The Minkowski difference can be computed by duality; however, it explicit computation is subject to the same set of accuracy and stability issues that apply to general intersections—as a direct result of its definition as intersections. In other words, Minkowski differences are not strictly computable [105, 106].

These generalized operations have recently found important applications in manufacture process planning; e.g., to automatically examine designs for tool accessibility and develop CNC machining process plans [232] and to automatically examine designs for 3D printability and correct models via MA-based heuristics [233, 237]. The morphological generalizations eliminates the restrictions on tool geometry.

Explicit morphologies on the shape space help formulating a variety of design and engineering problems. However, their application is limited to translational motions (i.e., filters), while rotations in most practical implementations are handled by auxiliary methods (e.g., sampling). Next, the concept is generalized to explicit morphologies on the motion space, in which the shape space is naturally embedded, to tackle a richer class of shape and motion related problems.

3.4.1.2 Pointset-Theoretic Group (i.e., Noncommutative) Morphology:

An important insight in modern morphology is to acknowledge that images (i.e., shapes) and filters (i.e., motions) are fundamentally distinct types [215, 218] and should be treated as such to avoid some of the confusions that arise in classical morphology (e.g., due to the arbitrariness of the choice of origin).[22] For our purposes, shapes sit in a space $\mathsf{E} = \mathbb{R}^3$ on which the group C acts, while filters reside in the group $\mathsf{C} = \mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$.

A subspace of $\mathscr{P}(\mathsf{C})$ is an 'explicit morphology' on C if it is a (regularized) Boolean algebra, and is also closed under (regularized) Minkowski product and inversion with respect to the identity.[23] Thus $\mathscr{P}(\mathsf{C})$ and $\mathscr{P}^*(\mathsf{C})$ are both explicit morphologies on the motion space with ordinary and regularized operations, respectively.

[22] For example, Minkowski sums of shapes in a vector space defined in (3.4.3) are not translation-invariant and will shift around by choosing different origins for the Cartesian coordinate system. However, the ambiguity is resolved once one adopts the perspective that one of the two sets is a filter that resides in the translation space $\mathsf{T} := \mathrm{T}(3) \cong \mathbb{R}^3$ whose elements $\mathbb{t} \in \mathrm{T}(3)$ are represented by vectors $\mathbf{t} \in \mathbb{R}^3$ with respect to a particular choice of origin as $\mathbb{t} \overset{0}{=} \mathbf{t}$ (though $\mathbb{t} \neq \mathbf{t}$) formalized by the embedding $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$ in the definition of dilations in (3.4.7).

[23] This automatically guarantees closure under Minkowski left- and right-quotients (due to the duality via De Morgan's laws), and other morphological operations such as left- and right- closing and opening. Note that in this case, dilation and erosion—to be defined shortly in (3.4.45) and (3.4.46), respectively—are defined across shapes and motions. Since the former is (isomorphic to) a subgroup of the latter, the morphology is technically closed under these operators as well.

Boolean Union and Intersection. Set operations apply equally to shapes and motions. As before, the set-theoretic Boolean operations $\cup, \cap : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ and their regularized counterparts $\cup^*, \cap^* : \mathscr{P}^*(\mathsf{C}) \times \mathscr{P}^*(\mathsf{C}) \to \mathscr{P}^*(\mathsf{C})$ are dual to each via set complement operator $\boldsymbol{c} : \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ and its regularized counterpart $\boldsymbol{c}^* : \mathscr{P}^*(\mathsf{C}) \to \mathscr{P}^*(\mathsf{C})$, respectively. The De Morgan's laws illustrated in (3.4.1) extend to motions in the obvious way:

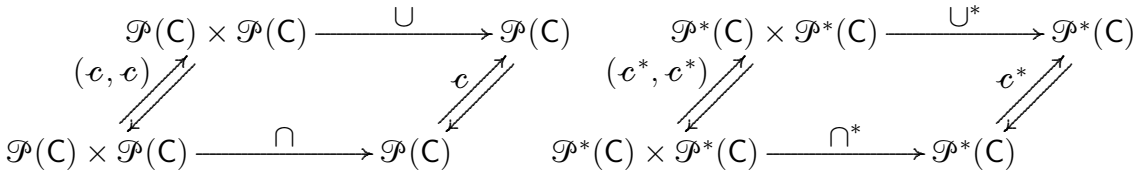

(3.4.37)

The two diagrams above can be connected together via explicit regularization:

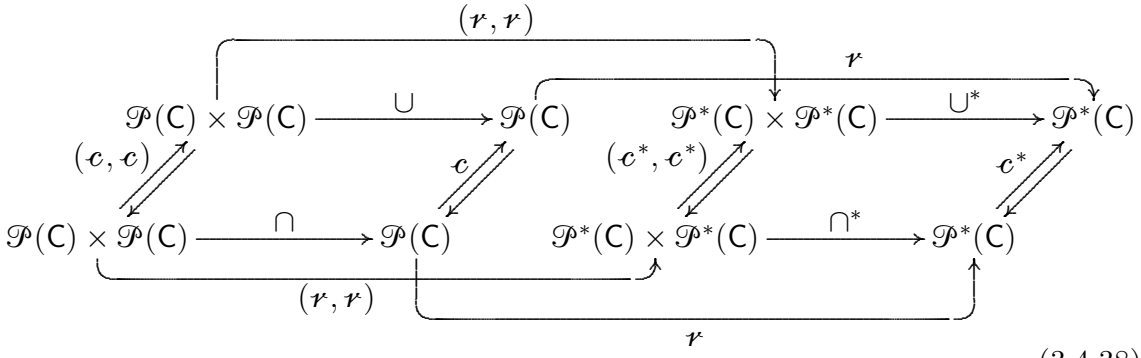

(3.4.38)

The explicit descriptions and combinatorial methods presented for (at most 3D) shapes in Section 3.4.1.1 by decomposition of the space into canonical intersection terms (i.e., 'atoms') can be applied to the $\mathsf{C}-$space constructions as well. However, the explicit approach becomes computationally difficult to implement in the 6D motion space, and eventually intractable in higher-dimensional $\mathsf{C}-$spaces.

Minkowski Product and Quotients. The aforementioned concepts established for vector spaces (and Abelian groups in general) can be extended to noncommutative groups; collectively referred to as group morphology [291, 292]. For $\mathsf{C} := \mathrm{SE}(3)$, the Minkowski sum and difference defined in (3.4.3) and (3.4.3) generalize to Minkowski product and left- and right-quotients $\otimes, \oslash, \obslash : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$:

$$(M_1 \otimes M_2) := \bigcup_{\mathbb{c}_1 \in M_1} \bigcup_{\mathbb{c}_2 \in M_2} \{\mathbb{c}_1 \mathbb{c}_2\} = \boldsymbol{c}(\boldsymbol{c}(M_1) \oslash M_2) = \boldsymbol{c}(M_1 \obslash \boldsymbol{c}(M_2)), \tag{3.4.39}$$

$$(M_1 \oslash M_2) := \bigcap_{\mathbb{c}_1 \in M_1} \bigcup_{\mathbb{c}_2 \in M_2} \{\mathbb{c}_1 \mathbb{c}_2\} = \boldsymbol{c}(\boldsymbol{c}(M_1 \otimes M_2)) = \boldsymbol{c}(\boldsymbol{c}(M_1) \obslash \boldsymbol{c}(M_2)), \tag{3.4.40}$$

$$(M_1 \obslash M_2) := \bigcap_{\mathbb{c}_2 \in M_2} \bigcup_{\mathbb{c}_1 \in M_1} \{\mathbb{c}_1 \mathbb{c}_2\} = \boldsymbol{c}(M_1 \otimes \boldsymbol{c}(M_2)) = \boldsymbol{c}(\boldsymbol{c}(M_1) \oslash \boldsymbol{c}(M_2)). \tag{3.4.41}$$

Note that the Minkowski product is noncommutative $(M_1 \otimes M_2) \neq (M_2 \otimes M_1)$—in spite of the unfortunate symmetric-looking symbol—while the left- and right-quotients are distinct noncommutative operators that are dual to each other and to the product operator through De Morgan's laws. The reflection with respect to the origin in vector spaces is also extended to an inversion with respect to group identity defined as $M^{-1} := \{\mathsf{C}^{-1} \mid \mathsf{C} \in M\}$. Using the semidirect product decomposition, $\mathbb{c}_1 \mathbb{c}_2 \overset{0}{=} (\mathbb{r}_1 \mathbb{r}_2, \mathbf{t}_1 + \mathbb{r}_1 \mathbf{t}_2)$ and $\mathbb{c}^{-1} \overset{0}{=} (\mathbb{r}^{\mathrm{T}}, -\mathbb{r}^{\mathrm{T}} \mathbf{t})$ due to (A.1.1) and (A.1.3), respectively.

The De Morgan duality on the right-hand sides of (3.4.39), (3.4.40), and (3.4.41)

can be illustrated via the following commutative diagrams:

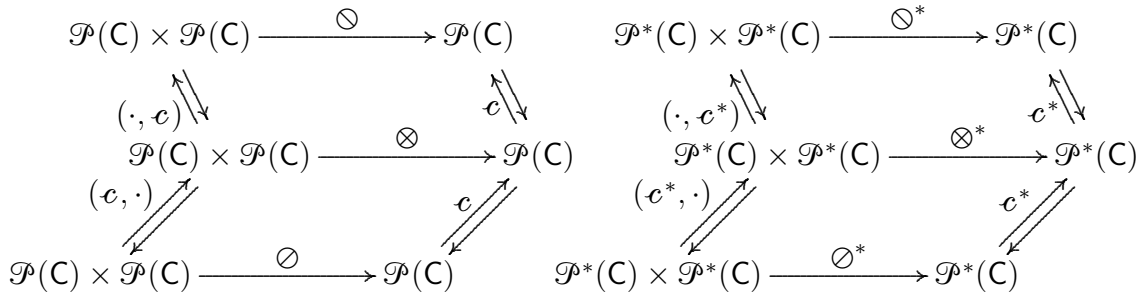

(3.4.42)

For similar reasons to those mentioned after (3.4.5) for Minkowski sum and difference over commutative groups, the above diagrams for ordinary and regularized Minkowski product and quotients over noncommutative groups cannot be connected via pre- and post-regularization as in (3.4.37) and (3.4.38). This is because Minkowski operations do not generally commute with regularization.

Explicit morphologies on the motions space provide unified formulation and solution methods for a variety of forward and inverse engineering problems. The following are a few important special cases for engineering applications:

As elaborated in Section 2.2.2, the shape space $\mathsf{E} = \mathbb{R}^3$ is a 'homogeneous space' of the configuration space $\mathsf{C} = \mathrm{SE}(3)$ since the latter acts on the former by a map $\mathsf{act} : \mathsf{C} \to \mathsf{aut}(\mathsf{E})$ where $\mathsf{aut}(\mathsf{E}) \subset \mathsf{E}^{\mathsf{E}}$ is the set of automorphisms $\tau : \mathsf{E} \to \mathsf{E}$ (i.e., continuous permutations) of E onto itself. In other words, for every $\mathbb{c} \in \mathsf{C}$, its action $\tau := \mathsf{act}(\mathbb{c})$ is a continuous invertible map (i.e., homeomorphism) that satisfies identity preservation (i.e., $\mathsf{act}(\mathbb{c}_{\mathsf{id}}) = \mathsf{id}_{\mathsf{E}}$) and compatibility (i.e., $\mathbb{c}_1(\mathbb{c}_2\mathbf{x}) = (\mathbb{c}_1\mathbb{c}_2)\mathbf{x}$). The action is 'transitive' since for every $\mathbf{x}_1, \mathbf{x}_2 \in \mathsf{E}$ there exists a configuration $\mathbb{c} \in \mathsf{C}$ such that $\mathbf{x}_2 = \mathbb{c}\mathbf{x}_1$, i.e., $\mathbf{x}_1 = \mathbb{c}^{-1}\mathbf{x}_2$, and E is called a 'principal' homogeneous space (i.e., 'torsor') since that $\mathbb{c} \in \mathsf{C}$ is always unique [215, 218].

Embedding, Lifting, and Projection. To convert the morphological operations across shapes and motions to Minkowski operations over groups, the embedding $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$ is no longer sufficient, unlike the case with purely translational motions in (3.4.7) through (3.4.13). After choosing an origin $\mathbf{0} \in \mathsf{E}$, the space of rigid transformations $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$ can be computationally decomposed to

- pure rotations around the origin, i.e., the 'stabilizer' subgroup $\mathrm{SO}(3) = \mathsf{stb}(\mathbf{0})$ defined as the set of transformations that stabilize/fix (i.e., do not move) the (arbitrarily chosen) origin $\mathbf{0} \in \mathsf{E}$; and

- pure translations isomorphic to the quotient subgroup $\mathrm{T}(3) \cong \mathrm{SE}(3)/\mathrm{SO}(3)$ in which all rotations around the origin are deemed equivalent, implying a one-to-one correspondence with $\mathsf{E} = \mathbb{R}^3$.

Accordingly, the 'lifting' $\ell_0 : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{C})$ of the shapes to motions with respect to the arbitrarily chosen origin $\mathbf{0} \in \mathsf{E}$ is defined as

$$\ell_0(S) = \big\{\mathbb{c} \in \mathsf{C} \mid \mathsf{act}(\mathbb{c})(\mathbf{0}) \in \mathsf{E}\big\} \stackrel{0}{=} \big\{(\mathbb{r}, \mathbf{x}) \mid \mathbb{r} \in \mathrm{SO}(3) \text{ and } \mathbf{x} \in \mathsf{E}\big\}, \tag{3.4.43}$$

where $\ell_0(S) \stackrel{0}{=} (\mathrm{SO}(3) \times S)$. For 2D shapes acted on by planar motions, this can be viewed as assigning a circle—i.e., a $2-$sphere isomorphic to SO(2)—to each point of the shape, as illustrated in Fig. 3.4.11. In 3D, this can be pictured by assigning a half $3-$sphere isomorphic to SO(3) to each point, or alternatively, by associating each point with all possible 3D orientations pictured as orthogonal triads of unit 3D vectors. Correspondingly, the 'projection' $\pi_0 : \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{E})$ of the motions to shapes

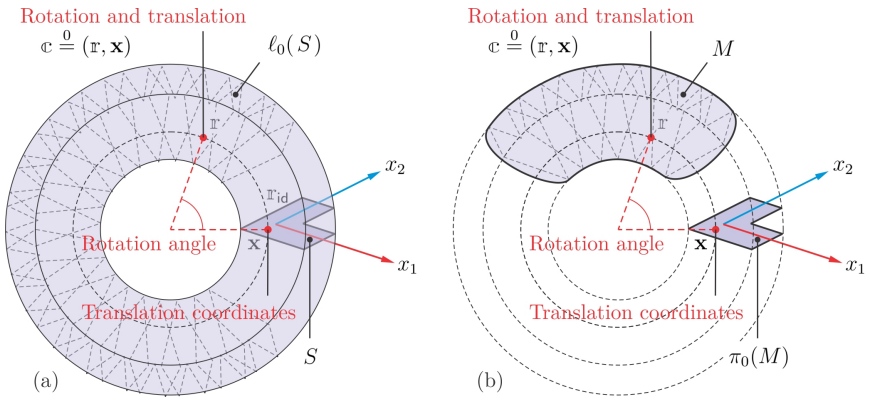

FIGURE 3.4.11: The explicit lifting and projection of a 2D shape in $\mathbb{R}^2$ to and from a 3D motion in $\mathrm{SE}(2) \cong \mathrm{SO}(2) \rtimes \mathrm{T}(2)$ amounts to extruding and collapsing a 1D circle (representing rotations) at each query point or configuration, respectively.

with respect to the origin $\mathbf{0} \in \mathsf{E}$ is defined as

$$\pi_0(M) = \big\{\mathsf{act}(\mathbb{c})(\mathbf{0}) \mid \mathbb{c} \in \mathsf{C}\big\} = \big\{\mathbf{x} \in \mathsf{E} \mid \mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x}) \text{ and } \mathbb{c} \in \mathsf{C}\big\}. \tag{3.4.44}$$

Figure 3.4.11 illustrates the explicit approach to lifting and projection for a simple planar shape, viewed along a direction on the plane (i.e., showing a 1D side-image of it) for ease of illustration. The relationship between a 2D planar shape in $\mathbb{R}^2$ and a 3D planar motion (i.e., 1D rotation and 2D translation) in $\mathrm{SE}(2) \cong \mathrm{SO}(2) \rtimes \mathrm{T}(2)$ with $\mathrm{SO}(2) \cong \mathbb{S}^1$ (i.e., a circle) and $\mathrm{T}(2) \cong \mathbb{R}^2$ (i.e., a plane) is illustrated by revolving the plane to create a toroidal topology. Lifting of a 2D shape to a 3D motion amounts to revolving every point in the shape along the circle, thus revolving the shape into a torus, while the projection of a 3D motion to a 2D shape returns the cross-section of the smallest torus that contains the given motion.

Morphological Dilation and Erosion. The group morphological generalizations of dilation and erosion $\mathsf{dil}\,, \mathsf{ero} : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ are defined as

$$\mathsf{dil}\,(M, S) := \bigcup_{\mathbb{c} \in M} \bigcup_{\mathbf{x} \in S} \mathbb{c}^{-1}\mathbf{x} = \pi_0\big(M^{-1} \otimes \ell_0(S)\big) \neq \gamma_0^{-1}\big(M^{-1} \otimes \gamma_0(S)\big), \tag{3.4.45}$$

$$\mathsf{ero}(M, S) := \bigcap_{\mathbb{c} \in M} \bigcup_{\mathbf{x} \in S} \mathbb{c}^{-1}\mathbf{x} = \pi_0\big(M^{-1} \oslash \ell_0(S)\big) \neq \gamma_0^{-1}\big(M^{-1} \oslash \gamma_0(S)\big), \tag{3.4.46}$$

where the inequalities are given to emphasize the difference with (3.4.7) and (3.4.9). Once again, regularized dilations and erosions can be conceptualized to obtain closure when working with regular sets. The following diagram generalizes (3.4.10):

$$\begin{array}{ccc} \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad\mathsf{dil}\quad} & \mathscr{P}(\mathsf{E}) \\ {\scriptstyle(\cdot,\boldsymbol{c})}\uparrow\downarrow & & \uparrow\downarrow{\scriptstyle\boldsymbol{c}} \\ \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\quad\mathsf{ero}\quad} & \mathscr{P}(\mathsf{E}) \end{array} \qquad \begin{array}{ccc} \mathscr{P}^*(\mathsf{C}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\quad\mathsf{dil}^*\quad} & \mathscr{P}^*(\mathsf{E}) \\ {\scriptstyle(\cdot,\boldsymbol{c}^*)}\uparrow\downarrow & & \uparrow\downarrow{\scriptstyle\boldsymbol{c}^*} \\ \mathscr{P}^*(\mathsf{C}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\quad\mathsf{ero}^*\quad} & \mathscr{P}^*(\mathsf{E}) \end{array} \tag{3.4.47}$$

For similar reasons to those mentioned after (3.4.10) for morphological dilation and erosion for translational motions, the above diagrams for ordinary and regularized dilation and erosion for general motions cannot be connected via pre- and post-regularization as in (3.4.37) and (3.4.38). This is because morphological operations do not generally commute with regularization.

The relationships between dilation and erosion and Minkowski operations can be

generalized from (3.4.11) to the following commutative diagram:

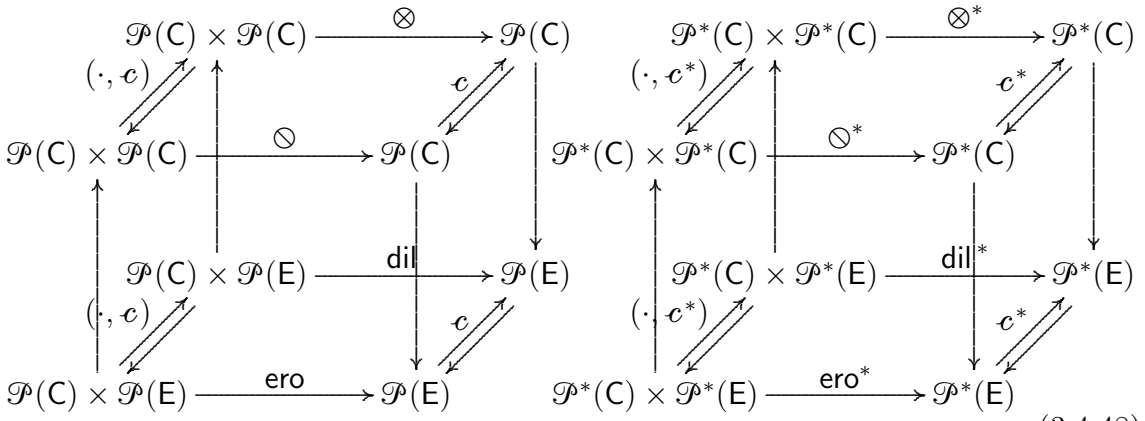

(3.4.48)

in which the labels for the vertical upward arrows should be $((\cdot)^{-1}, \ell_0) \uparrow$ and the labels for the vertical downward arrows should be $\downarrow \pi_0$—unlike the rotation-free special case in (3.4.11)—omitted from the diagram due to cluttering.

Morphological Closing and Opening. The group morphological generalizations of closing and opening $\mathsf{cls}, \mathsf{opn} : \mathcal{P}(\mathsf{C}) \times \mathcal{P}(\mathsf{E}) \to \mathcal{P}(\mathsf{E})$ are defined as

$$\mathsf{cls}(M, S) := \mathsf{ero}\big(M, \mathsf{dil}\,(M^{-1}, S)\big) = \bigcap_{\mathbb{c} \in M} \bigcup_{\mathbb{c}' \in M} \bigcup_{\mathbf{x} \in S} (\mathbb{c}^{-1}\mathbb{c}')\mathbf{x} \tag{3.4.49}$$

$$= \pi_0\big(M^{-1} \oslash (M \otimes \ell_0(S))\big) \neq \gamma_0^{-1}\big(M^{-1} \oslash (M \otimes \gamma_0(S))\big), \tag{3.4.50}$$

$$\mathsf{opn}(M, S) := \mathsf{dil}\,\big(M, \mathsf{ero}(M^{-1}, S)\big) = \bigcup_{\mathbb{c} \in M} \bigcap_{\mathbb{c}' \in M} \bigcup_{\mathbf{x} \in S} (\mathbb{c}^{-1}\mathbb{c}')\mathbf{x} \tag{3.4.51}$$

$$= \pi_0\big(M^{-1} \otimes (M \oslash \ell_0(S))\big) \neq \gamma_0^{-1}\big(M^{-1} \otimes (M \oslash \gamma_0(S))\big), \tag{3.4.52}$$

where the inequalities are given to emphasize the difference with (3.4.15) and (3.4.13). Once again, regularized openings and closings can be conceptualized to obtain closure

when working with regular sets. The following diagram generalizes (3.4.16):

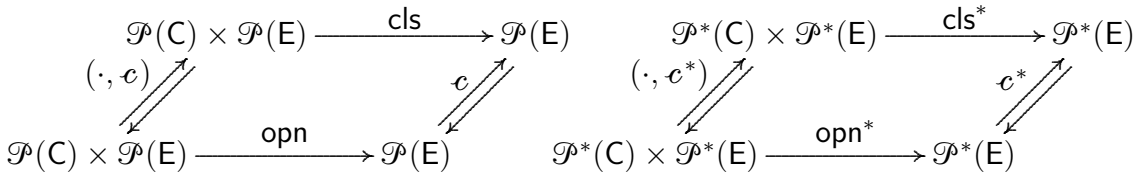

(3.4.53)

For similar reasons to those mentioned after (3.4.16) for morphological closing and opening for translational motions, the above diagrams for ordinary and regularized dilation and closing for opening motions cannot be connected via pre- and post-regularization as in (3.4.37) and (3.4.38). This is because morphological operations do not generally commute with regularization.

Minkowski Closing and Opening. The Minkowski closing and opening defined in (3.4.17) and (3.4.18) are also generalized in a similar fashion to left- and right-closing $\oslash, \osslash \colon \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ and opening $\oslash, \osslash \colon \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ defined respectively as the following different compositions of products and quotients:

$$(M_1 \oslash M_2) := \big(M_1^{-1} \oslash (M_1 \otimes M_2)\big) = \bigcap_{\mathbb{c}_1 \in M_1} \bigcup_{\mathbb{c}_2 \in M_2} \bigcup_{\mathbb{c}_1' \in M_1} \{\mathbb{c}_1'^{-1} \mathbb{c}_1 \mathbb{c}_2\}, \tag{3.4.54}$$

$$(M_1 \osslash M_2) := \big((M_1 \otimes M_2) \oslash M_2^{-1}\big) = \bigcap_{\mathbb{c}_2 \in M_2} \bigcup_{\mathbb{c}_2' \in M_2} \bigcup_{\mathbb{c}_1 \in M_1} \{\mathbb{c}_1 \mathbb{c}_2'^{-1} \mathbb{c}_2\}. \tag{3.4.55}$$

$$(M_1 \oslash M_2) := \big((M_1 \oslash M_2) \otimes M_2^{-1}\big) = \bigcup_{\mathbb{c}_2 \in M_2} \bigcap_{\mathbb{c}_2' \in M_2} \bigcup_{\mathbb{c}_1 \in M_1} \{\mathbb{c}_1 \mathbb{c}_2'^{-1} \mathbb{c}_2\}, \tag{3.4.56}$$

$$(M_1 \osslash M_2) := \big(M_1^{-1} \otimes (M_1 \oslash M_2)\big) = \bigcup_{\mathbb{c}_1 \in M_1} \bigcap_{\mathbb{c}_2 \in M_2} \bigcup_{\mathbb{c}_1' \in M_1} \{\mathbb{c}_1'^{-1} \mathbb{c}_1 \mathbb{c}_2\}, \tag{3.4.57}$$

Once again, the De Morgan dualities can be observed in commutative diagrams:

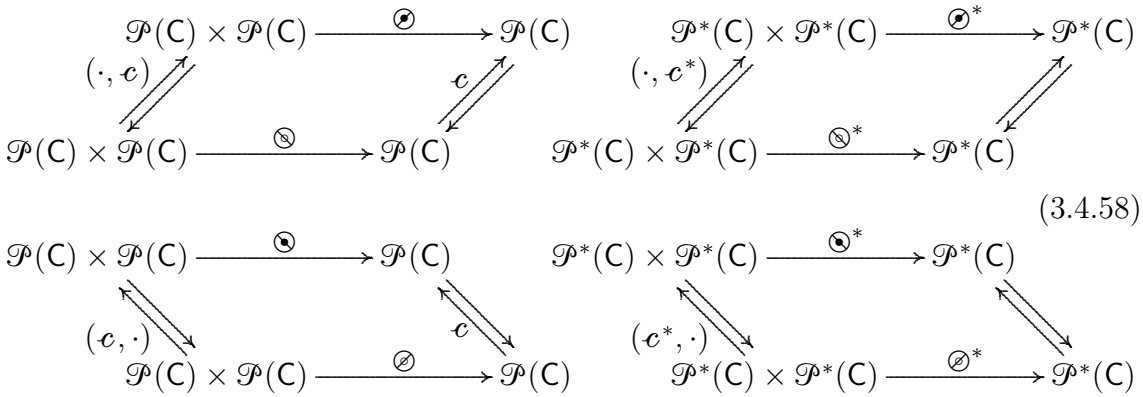

(3.4.58)

(3.4.59)

For similar reasons to those mentioned after (3.4.19) for Minkowski closing and opening over commutative groups, the above diagrams for ordinary and regularized Minkowski closings and openings over noncommutative groups cannot be connected via pre- and post-regularization as in (3.4.37) and (3.4.38). This is because Minkowski operations do not generally commute with regularization.

Homogeneous Product and Quotients. The group-theoretic conceptualization also allows developing Minkowski products and quotients over the shape space that map shapes to motions; namely, the 'homogeneous' Minkowski product and left- and right-quotients $\boxtimes, \boxslash, \boxbslash : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{C})$ defined as

$$(S_1 \boxtimes S_2) := \big\{ \mathbb{c} \in \mathsf{C} \mid (S_1 \cap \mathbb{c} S_2) \neq \emptyset, \text{ i.e., } (\mathbb{c}^{-1} S_1 \cap S_2) \neq \emptyset \big\} \tag{3.4.60}$$

$$= \boldsymbol{c}(\boldsymbol{c}(S_1) \boxslash S_2) = \boldsymbol{c}(S_1 \boxbslash \boldsymbol{c}(S_2)), \tag{3.4.61}$$

$$(S_1 \boxslash S_2) := \big\{ \mathbb{c} \in \mathsf{C} \mid S_1 \supseteq \mathbb{c} S_2, \text{ i.e., } \mathbb{c}^{-1} S_1 \supseteq S_2 \big\} \tag{3.4.62}$$

$$= \boldsymbol{c}(\boldsymbol{c}(S_1) \boxtimes S_2) = \boldsymbol{c}(\boldsymbol{c}(S_1) \boxbslash \boldsymbol{c}(S_2)), \tag{3.4.63}$$

$$(S_1 \boxbslash S_2) := \big\{ \mathbb{c} \in \mathsf{C} \mid S_1 \subseteq \mathbb{c} S_2, \text{ i.e., } \mathbb{c}^{-1} S_1 \subseteq S_2 \big\} \tag{3.4.64}$$

$$= \boldsymbol{c}(S_1 \boxtimes \boldsymbol{c}(S_2)) = \boldsymbol{c}(\boldsymbol{c}(S_1) \boxslash \boldsymbol{c}(S_2)). \tag{3.4.65}$$

The homogeneous product and quotients operations are strongly correlated with Minkowski product and quotients via lifitng/projection maps with respect to a particular origin $\mathbf{0} \in \mathsf{E}$ [215, 218]; namely:

$$(S_1 \boxtimes S_2) = \ell_0(S_1) \otimes (\ell_0(S_2))^{-1} = \ell_0(S_1) \otimes (\ell_0(-S_2)), \tag{3.4.66}$$

$$(S_1 \boxslash S_2) = \ell_0(S_1) \oslash (\ell_0(S_2))^{-1} = \ell_0(S_1) \oslash (\ell_0(-S_2)), \tag{3.4.67}$$

$$(S_1 \boxbslash S_2) = \ell_0(S_1) \obslash (\ell_0(S_2))^{-1} = \ell_0(S_1) \obslash (\ell_0(-S_2)), \tag{3.4.68}$$

remembering that $-S = \{-\mathbf{x} \mid \mathbf{x} \in S\}$ and $M^{-1} = \{\mathbb{c}^{-1} \mid \mathbb{c} \in M\}$ stand for reflection (before lifting, in E) and inversion (after lifting, in C).

The De Morgan duality on the right-hand sides of (3.4.61), (3.4.63), and (3.4.65) can be illustrated via the following commutative diagrams:

$$\begin{array}{ccc}
\mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) & \xrightarrow{\boxbslash} & \mathscr{P}(\mathsf{C}) \\
{\scriptstyle (\cdot, \boldsymbol{c})} \uparrow\downarrow & & \uparrow\downarrow {\scriptstyle \boldsymbol{c}} \\
\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\boxtimes} & \mathscr{P}(\mathsf{C}) \\
{\scriptstyle (\boldsymbol{c}, \cdot)} \uparrow\downarrow & & \uparrow\downarrow {\scriptstyle \boldsymbol{c}} \\
\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) & \xrightarrow{\boxslash} & \mathscr{P}(\mathsf{C})
\end{array}
\qquad\qquad
\begin{array}{ccc}
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\boxbslash^*} & \mathscr{P}^*(\mathsf{C}) \\
{\scriptstyle (\cdot, \boldsymbol{c}^*)} \uparrow\downarrow & & \uparrow\downarrow {\scriptstyle \boldsymbol{c}^*} \\
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\boxtimes^*} & \mathscr{P}^*(\mathsf{C}) \\
{\scriptstyle (\boldsymbol{c}^*, \cdot)} \uparrow\downarrow & & \uparrow\downarrow {\scriptstyle \boldsymbol{c}^*} \\
\mathscr{P}^*(\mathsf{E}) \times \mathscr{P}^*(\mathsf{E}) & \xrightarrow{\boxslash^*} & \mathscr{P}^*(\mathsf{C})
\end{array}
\tag{3.4.69}$$

Once again, the above diagrams for ordinary and regularized homogeneous product

and quotients over noncommutative groups cannot be connected via pre- and post-regularization as in (3.4.37) and (3.4.38), as they too, like other Minkowski and morphological operations, do not generally commute with regularization.

The relationships between homogeneous products and quotients and Minkowski products and quotients are illustrated via:

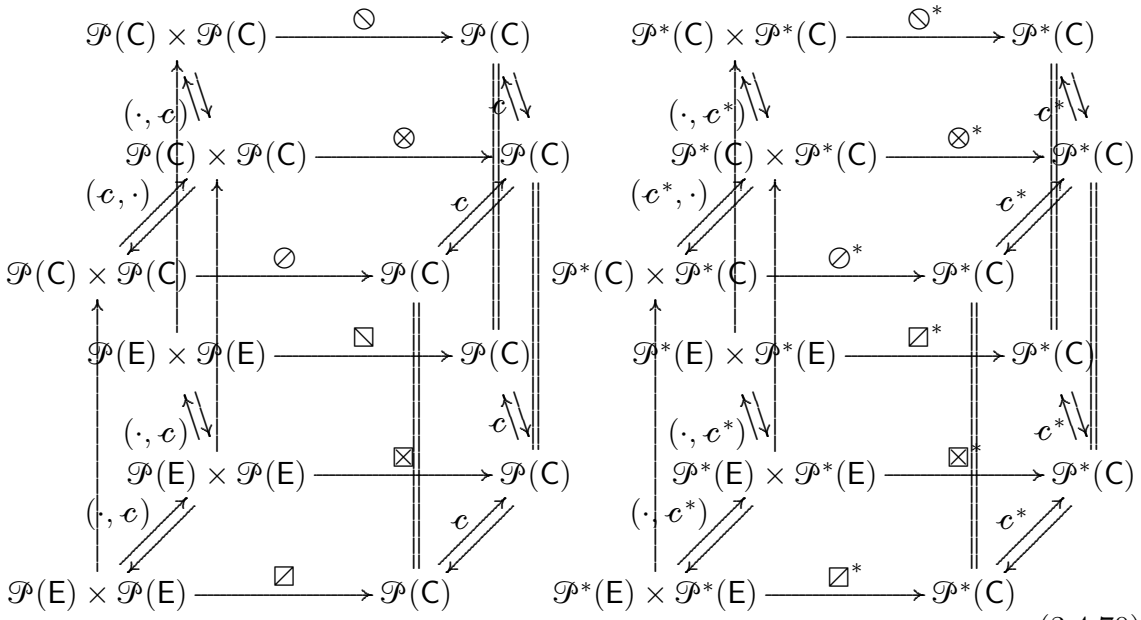

(3.4.70)

in which the labels for the vertical upward arrows should be $(\ell_0, (\ell_0(\cdot))^{-1}) \uparrow$ and the vertical equalities can be consolidated to obtain triangular vertical cycles.

Explicit group morphologies (via regularized and nonregularized) operations provide a uniform formalism for a variety of shape and motion related problems. Below I shall give examples that are important in engineering applications.

$\mathsf{C}-$Obstacles and Free Space. Given two (at most) 3D shapes $S_1, S_2 \in \mathscr{P}(\mathsf{E})$,

- the homogeneous product $(S_1 \boxtimes S_2)$ corresponds to the set of configurations (i.e., relative transformations between the shapes) that leads to a nonempty set-theoretic intersection between them; while
- the homogeneous left- and right-quotients $(S_1 \boxslash S_2)$ and $(S_1 \boxbslash S_2)$ correspond to the sets of configurations that lead to full inclusion of one set in another, which means a nonempty set-theoretic intersection between one shape and the other shape's complement.

Note the emphasis on set-theoretic (not regularized) intersection. Thus the former, defined in (3.4.60) and repeated below, along with its regularized counterpart:

$$(S_1 \boxtimes S_2) = \big\{\mathbb{c} \in \mathsf{C} \mid (S_1 \cap \mathbb{c}S_2) \neq \emptyset, \text{ i.e., } (\mathbb{c}^{-1}S_1 \cap S_2) \neq \emptyset\big\}, \tag{3.4.71}$$

$$(S_1 \boxtimes^* S_2) = \text{r}\big\{\mathbb{c} \in \mathsf{C} \mid (S_1 \cap \mathbb{c}S_2) \neq \emptyset, \text{ i.e., } (\mathbb{c}^{-1}S_1 \cap S_2) \neq \emptyset\big\}, \tag{3.4.72}$$

are both different from the notion of the $\mathsf{C}-$obstacle O_{S_1,S_2} defined in (2.3.51) in Section 2.3.1; in fact $O_{S_1,S_2} \subseteq (S_1 \boxtimes S_2)$, however, $O_{S_1,S_2} =: (S_1 \boxtimes^\ddagger S_2)$, in which the new operator $\boxtimes^\ddagger$: $\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{C})$ is defined as

$$(S_1 \boxtimes^\ddagger S_2) := \big\{\mathbb{c} \in \mathsf{C} \mid (S_1 \cap^* \mathbb{c}S_2) \neq \emptyset, \text{ i.e., } (\mathbb{c}^{-1}S_1 \cap^* S_2) \neq \emptyset\big\}, \tag{3.4.73}$$

which precisely gives the $\mathsf{C}-$obstacle, and can be rewritten as:

$$(S_1 \boxtimes^{\ddagger} S_2) = \big\{\mathbb{c} \in \mathsf{C} \mid (S_1 \cap \mathbb{c}S_2) \overset{\mathsf{ae}}{\neq} \emptyset, \text{ i.e., } (\mathbb{c}^{-1}S_1 \cap S_2) \overset{\mathsf{ae}}{\neq} \emptyset\big\}, \tag{3.4.74}$$

whose set-theoretic (not regularized) complement in the $\mathsf{C}-$space is typically referred to as the collision-free space, or simply 'free space' in robotics [199]:

$$\boldsymbol{c}(S_1 \boxtimes^{\ddagger} S_2) = \big\{\mathbb{c} \in \mathsf{C} \mid \boldsymbol{c}(S_1) \overset{\mathsf{ae}}{\supseteq} \mathbb{c}S_2 \text{ i.e., } \mathbb{c}^{-1}S_1 \overset{\mathsf{ae}}{\subseteq} \boldsymbol{c}(S_2)\big\}, \tag{3.4.75}$$

while their shared boundary $\boldsymbol{b}(S_1 \boxtimes^{\ddagger} S_2)$ is often called the 'contact space'. Accordingly, the related morphological notions homogeneous left- and right-quotient operators defined in (3.4.63) and (3.4.65) can be modified as well to obtain:

$$(S_1 \boxslash^{\ddagger} S_2) := \big\{\mathbb{c} \in \mathsf{C} \mid S_1 \overset{\mathsf{ae}}{\supseteq} \mathbb{c}S_2, \text{ i.e., } \mathbb{c}^{-1}S_1 \overset{\mathsf{ae}}{\supseteq} S_2\big\}, \tag{3.4.76}$$

$$(S_1 \boxbslash^{\ddagger} S_2) := \big\{\mathbb{c} \in \mathsf{C} \mid S_1 \overset{\mathsf{ae}}{\subseteq} \mathbb{c}S_2, \text{ i.e., } \mathbb{c}^{-1}S_1 \overset{\mathsf{ae}}{\subseteq} S_2\big\}. \tag{3.4.77}$$

As expected, the free space is dual to the above operators via $\boldsymbol{c}(S_1 \boxtimes^{\ddagger} S_2) = \boldsymbol{c}(\boldsymbol{c}(S_1) \boxslash^{\ddagger} S_2) = \boldsymbol{c}(S_1 \boxbslash^{\ddagger} \boldsymbol{c}(S_2))$.

It is important to acknowledge the fact that group-theoretic homogeneous product $(S_1 \boxtimes S_2)$, its regularization $(S_1 \boxtimes^{*} S_2)$, and the (potentially irregular) $\mathsf{C}-$obstacle $(S_1 \boxtimes^{\ddagger} S_2)$ are three different sets.[24] They are equal$-\mathsf{ae}$ if the two sets are regular (i.e., $S_1, S_2 \in \mathscr{P}^{*}(\mathsf{E})$) which is *not* the case in general (i.e., $S_1, S_2 \in \mathscr{P}(\mathsf{E})$) in which lower-dimensional irregularities can create full-dimensional features, as illustrated in Fig. 3.4.12. Even in the former case, the differences become significant, for example,

[24] However, they are often taken to be interchangeable notions for the same thing in the literature, when regularization is taken for granted [36, 213, 215].

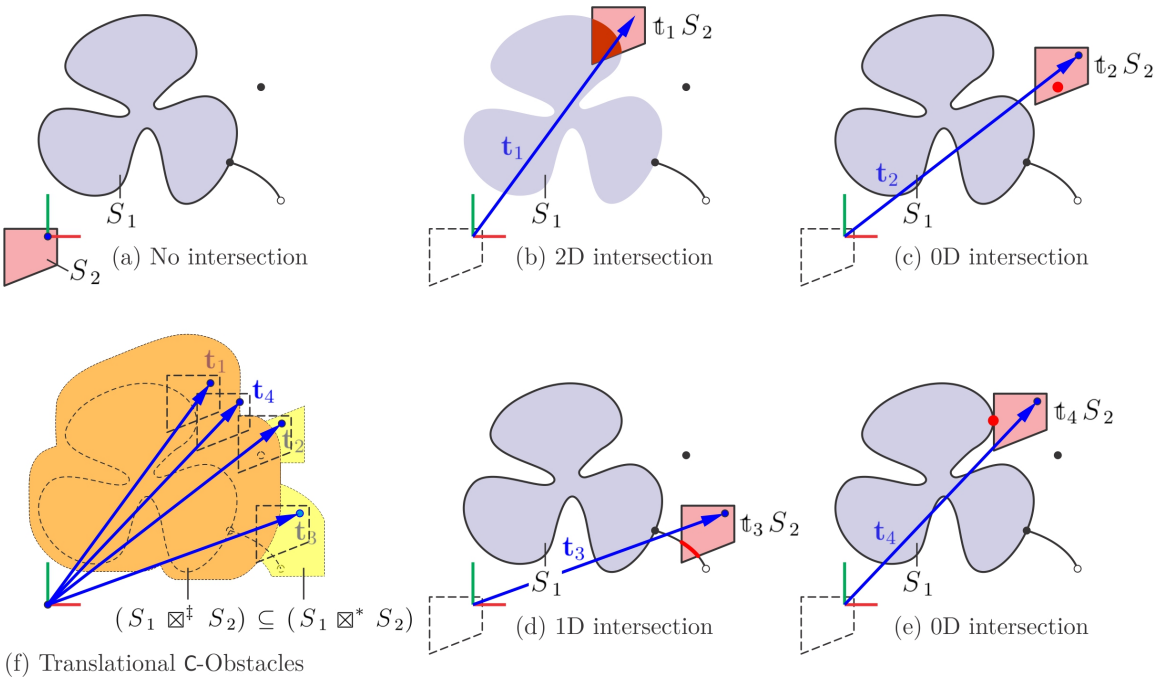

FIGURE 3.4.12: The generalized translational C−obstacles in (f) for a pair of shapes on the 2D plane are the subset of all relative displacements at which there are intersections of nonzero d−measure for d := 0, 1, and 2. Compare to Figs. 2.3.5 and 2.3.15.

when dealing with mechanical assemblies in which the nominal (i.e., not toleranced) mating constraints between parts give rise to degenerate lower-dimensional features in the C−obstacle that deliberately result in lower DOF motion, to which I shall return in Chapter 5.

The modified group morphological operations (denoted by a double-dagger ‡) will be elaborated in Section 3.4.2 where they will be reformulated in detail in measure-theoretic terms such as inner products and convolutions.

Figure 2.3.15 in Section 2.3.1, repeated here in Fig. 3.4.12 with slight modifications, illustrates the difference between the original and modified C−obstacles for a simple pair of planar shapes.

Parametric Sweep and Unsweep. Given an (at most) 3D shape $S \in \mathscr{P}(\mathsf{E})$, which may or may not be regularized (called the 'generator') along with a 1D motion (called the 'trajectory') $M \in \mathscr{P}^1(\mathsf{C})$ which is a piecewise parameterized curve embedded into the $\mathsf{C}-$space as $M := \bigcup_{0 \le i < n_1} \Gamma_i(\Delta^1)$, where $\Gamma_i : \Delta^1 \to \mathsf{C}$ are analytic parameterizations that map the linear interval $\Delta^1 \subset \mathbb{R}$ to the n_1 sub-trajectories $\Gamma_i(\Delta^1) \in \mathscr{P}(\mathsf{C})$ for all $0 \le i < n_1$:

- The dilation under the *inverse* trajectory $M^{-1} \in \mathscr{P}(\mathsf{C})$ corresponds precisely to the classical 'sweep' of the shape S along the given 1D motion M, denoted via $\mathsf{sweep} : \mathscr{P}^1(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is defined as

$$\mathsf{sweep}(M, S) = \mathsf{dil}\,(M^{-1}, S) = \pi_0\big(M \otimes \ell_0(S)\big) \tag{3.4.78}$$

$$= \bigcup_{0 \le i < n_1} \bigcup_{\mathbb{c} \in \Gamma_i(\Delta^1)} \bigcup_{\mathbf{x} \in S} (\mathbb{c}\mathbf{x}). \tag{3.4.79}$$

  Thus the sweep includes all spatial points that will be included in the generator at some time during the motion.

- The erosion under the *forward* trajectory $M^{+1} \in \mathscr{P}(\mathsf{C})$ corresponds precisely to the classical 'unsweep' of the shape S along the given 1D motion M, denoted via $\mathsf{unswp} : \mathscr{P}^1(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is defined as

$$\mathsf{unswp}(M, S) = \mathsf{ero}(M, S) = \pi_0\big(M^{-1} \oslash \ell_0(S)\big) \tag{3.4.80}$$

$$= \bigcap_{0 \le i < n_1} \bigcap_{\mathbb{c} \in \Gamma_i(\Delta^1)} \bigcup_{\mathbf{x} \in S} (\mathbb{c}^{-1}\mathbf{x}). \tag{3.4.81}$$

  Thus the unsweep includes all spatial points that will remain included in the

generator at all times during the motion.[25]

The sweep has been recognized as an important geometric modeling operator—even among the fundamental representation schemes [277,278]—and has been studied extensively, using a myriad of different explicit and implicit methods such as sweep differential equation (SDE) [40], sweep-envelope differential equation (SEDE) [41,42], the Jacobian rank deficiency method [3,4,6,7], point membership classification (PMC) methods [111, 113], and many more (reviewed in [5]) due to their broad applications in geometric modeling and spatial planning.

The unsweep, on the other hand, was introduced as the dual operator [161, 162] with important applications in design [163, 164]. However, the definitions of sweep and unsweep as infinite unions and intersections make their computational evaluation challenging. Most explicit methods attempt to evaluate the boundary of the swept envelope from a knowledge of the boundary of the generator and the trajectory. Such an evaluation typically requires going through the following sequence of steps:[26]

1. characterizing the generator boundary points depending on their contributions to the different subsets of the swept volume—e.g., egress, ingress, and grazing points for SDE [40];

2. approximating the motion, which is an infinite sequence of transformations, by a finite sequence, e.g., by time-stepping along the motion trajectory;

[25]The reason behind the differences in the use of forward and inverse motions is conventional. The dilation and erosion are chosen to be *covariant* with respect to motions, i.e., $\mathsf{dil}\,(\mathbb{c}M, \mathbb{c}S) = \mathsf{dil}\,(M, S)$ and $\mathsf{ero}(\mathbb{c}M, \mathbb{c}S) = \mathsf{ero}(M, S)$ for all $\mathbb{c} \in \mathsf{C}$ which is important form a morphological perspective, while the sweep is chosen to be *contravariant*, i.e., $\mathsf{sweep}^1(\mathbb{c}M, S) = \mathsf{sweep}^1(M, \mathbb{c}S)$, which is a more intuitive choice from an engineering designer's perspective [215, 218].

[26]The author is thankful to Radu Corcodel and Horea Ilies [79] for providing the herein cited references to integration methods. That being said, the responsibility of potential misrepresentation or oversimplification of the literature lies with the author.

3. computing the approximate orbits of a subset of the generator boundary points that may contribute to the swept volume boundary (i.e., the 'candidate set');

4. classifying the candidate set against the swept volume—e.g., as sweep boundary, sweep interior, and fold points [111]; and

5. trimming the fold regions and retaining the nonsingular surface patches that contribute to a manifold swept volume boundary.

Thus, in a sense, most of these methods follow a similar path to that of boundary evaluation (i.e., CSG $\rightarrow$ B-rep representation conversion) [285] except that here one deals with a much larger number of primitives to intersect [314], but the fact that they are all generated from the same generator and a continuous motion are exploited to rule out a large number of points from the candidate sets. One way or another, the solutions are all approximate due to the complexity of the differential equations, the running times are extremely at odds with the performance requirements of real-time applications with limited time budgets, and the behavior of singularities (i.e., fold regions, undercuts, or interferences) are not completely understood.

Figures 3.4.13 and 3.4.14 illustrate 1−parameteric sweep and unsweep along a planar motion. Note that lower-dimensional features may or may not contribute to full-dimensional features in the regularized swept volume.

The notions of one-parametric sweep and unsweep can be naturally extended to d−parameteric motions. This time assuming that the motion is d−dimensional, i.e., $M := \bigcup_{0 \leq i < n_{\mathrm{d}}} \Gamma_i^{\mathrm{d}}(\Delta^{\mathrm{d}})$, where $\Gamma_i^{\mathrm{d}} : \Delta^{\mathrm{d}} \rightarrow \mathsf{C}$ is an analytic parameterization of a subset of the motion, the sweep and unsweep are generalized from (3.4.79) and (3.4.81),

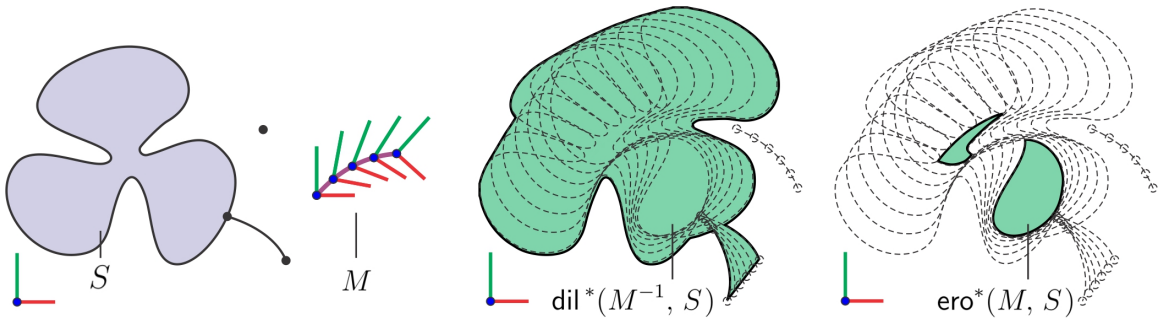

(a) A Shape and a 1-Parameteric Motion (b) Sweep of the Shape (c) Unsweep of the Shape

FIGURE 3.4.13: The sweep and unsweep of a shape along a 1−parametric planar motion.

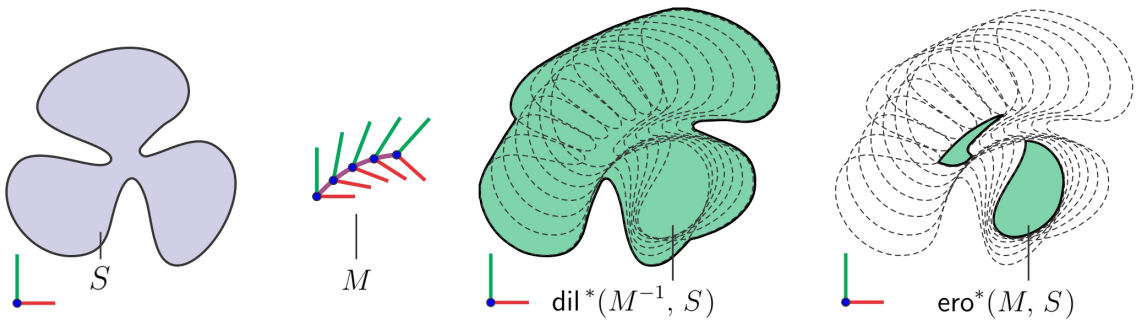

(a) A Solid and a 1-Parameteric Motion (b) Sweep of the Solid (c) Unsweep of the Solid

FIGURE 3.4.14: The sweep and unsweep of a solid along a 1−parametric planar motion.

respectively, in the natural way:

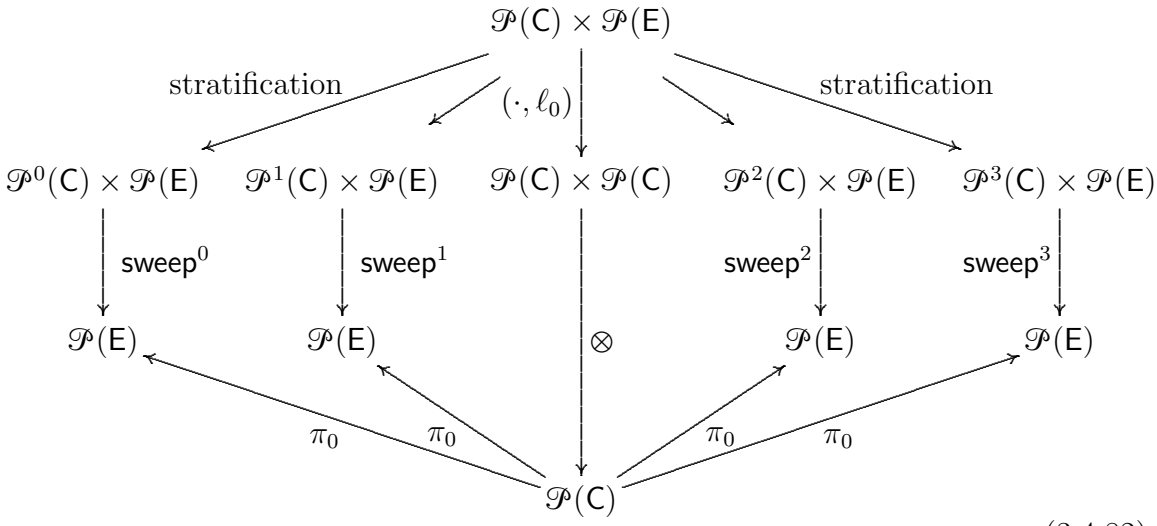

(3.4.82)

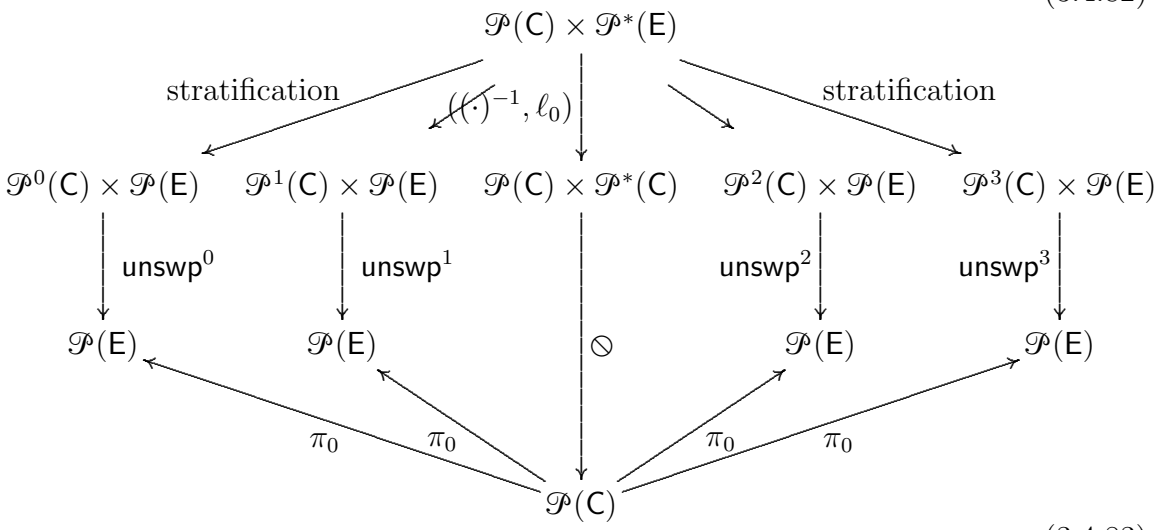

(3.4.83)

the two diagrams being coupled together (not shown) via De Morgan dualities of dilation and erosion in (3.4.48). The following special cases occur:

- For $\mathrm{d} := 0$, the sweep/unsweep of the solid S along a discrete set of configurations $M^0 := \{\mathbb{c}_i\}_{0 \le i < n_0}$ with $|M^0| = n_0$ is simply a finite union/intersection of n_0 moved (i.e., translated and rotated) instances of the same shape, which

generalizes the notion of groupings (e.g., of balls) presented earlier in (3.4.26).

- For $\mathrm{d} := 1$, the motion $M^1 := \bigcup_{0 \leq i < n_1} \Gamma_i(\Delta^1)$ can be viewed as an embedding of a collection of 1D curve segments in the configuration space, where each configuration $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ is obtained by assigning a rotation $\mathbb{r} \in \mathrm{SO}(3)$ to every point $\mathbf{t} \in \mathbb{R}^3$ prescribing translation along the curve.

- For $\mathrm{d} := 2$, the motion $M^2 := \bigcup_{0 \leq i < n_2} \Gamma_i(\Delta^2)$ can be viewed as an embedding of a collection of 2D surface patches in the configuration space, where each configuration $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ is obtained by assigning a rotation $\mathbb{r} \in \mathrm{SO}(3)$ to every point $\mathbf{t} \in \mathbb{R}^3$ prescribing translation over the surface.

These notions can be further generalized by assuming higher-dimensional trajectories (e.g., up to 6D) or by allowing the generator itself to be a lower-dimensional set of the 3−space (e.g., a curve segment or surface patch) or of heterogeneous dimensions—examples of which are ubiquitous in the feature set of most solid modeling systems in the form of the so-called 'extruding', 'revolving', and 'lofting' operations.

Generalized multi-parametric sweep and unsweep operations are rarely studied in a uniform context, and their properties in terms of singularities (e.g., fold regions and undercuts) are not understood [79]. In fact, most existing explicit methods are either limited to simple shapes and lower-dimensional generators and trajectories—which require simple representations and enumeration methods—or use simplifying (e.g., polyhedral or enumerative) approximations.

Configuration Sampling. Given two finite sets of configurations $M_1^0, M_2^0 \subset \mathsf{C}$ with cardinalities $|M_1^0| = n_1$ and $|M_2^0| = n_2$, respectively, the Minkowski product

$(M_1^0 \otimes M_2^0) \subset \mathsf{C}$ is also a finite set with at most $n_1 n_2$ poses, i.e., $|M_1^0 \oplus M_2^0| \leq n_1 n_2$, while the Minkowski quotients $(M_1^0 \oslash M_2^0) \subset \mathsf{C}$ and $(M_1^0 \oslash M_2^0) \subset \mathsf{C}$ are also a finite set with at least $\lceil \frac{n_1}{n_2} \rceil$ points, i.e., $|M_1^0 \oslash M_2^0| \geq \lceil \frac{n_1}{n_2} \rceil$ and $|M_1^0 \oslash M_2^0| \geq \lceil \frac{n_1}{n_2} \rceil$.

Once again, it is desirable to use such discrete samples to approximate continuum motions similar to the case with shapes, with important applications in forward problems such as mechanism workspace computation and inverse problems such as roadmap construction for path planning. However, discretization in the presence of rotations is not as trivial as with translations, due to the following two difficulties:

- There are only finitely many nontrivial discrete rotation subgroups which puts a limit to the degree to which one can uniformly sample the space of rotations in away that guarantees algebraic closure. There are numerous techniques on global sampling [24, 97, 134, 370, 371] and local refining [234, 236] for computations over the rotation group.[27]

- Unlike the Euclidean distance over E, the geodesic distance over SE(3) is not preserved under general rigid transformations due to (A.1.6)—e.g., small rotations can cause large translations depending on the distance to the identity $\mathbb{c}_{\mathsf{id}} \overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{0})$. As a result, even if the discrete motions $M_1^0, M_2^0 \in \mathscr{P}(\mathsf{C})$ are somehow sampled to yield ϵ_1- and ϵ_2-approximations of some continuum motions $M_1, M_2 \in \mathscr{P}(\mathsf{C})$, respectively, then $(M_1^0 \otimes M_2^0)$ is *not* an $(\epsilon_1+\epsilon_2)-$approximation of $(M_1 \otimes M_2)$. Rather, the approximation quality differs from one pose to another depending on the size of the translational components [235, 236].

The following are simpler special cases in which one or both of the motions are

[27] Unlike many other difficulties with explicit methods described so far that are circumvented by switching to implicit methods in Section 3.4.2, there is no way around sampling which underlies every implicit implementation when performing global computations or querying global properties.

restricted to a subgroup of $\mathrm{SE}(3) = \mathrm{SO}(3) \rtimes \mathrm{T}(3)$:

- Given two finite sets of translations $T_1^0, T_2^0 \subset \mathrm{T}(3)$ with $|T_1^0| = n_1$ and $|T_2^0| = n_2$, if they are ϵ_1- and ϵ_2-approximations of some $T_1, T_2 \subset \mathrm{T}(3)$, respectively, then $(T_1^0 \otimes T_2^0)$ is an $(\epsilon_1 + \epsilon_2)-$approximation of $(T_1 \otimes T_2) \subset \mathrm{T}(3)$, since Euclidean distance over $\mathrm{T}(3)$ is preserved under pure translations due to (A.1.6).

- Given two finite sets of rotations $P_1, P_2 \subset \mathrm{SO}(3)$ with $|P_1| = n_1$ and $|P_2| = n_2$, if they are ϵ_1- and ϵ_2-approximations of some $R_1, R_2 \subset \mathrm{SO}(3)$, respectively, then $(R_1^0 \otimes R_2^0)$ is an $(\epsilon_1 + \epsilon_2)-$approximation of $(R_1 \oplus R_2) \subset \mathrm{SO}(3)$, since geodesic distance over $\mathrm{SO}(3)$ is preserved under pure rotations due to (A.1.6).

- Given a finite set of translations $T^0 \subset \mathrm{T}(3)$ with $|T^0| = n_1$ and a finite set of rotations $R^0 \subset \mathrm{SO}(3)$ with $|R^0| = n_2$, if they are ϵ_1- and ϵ_2-approximations of some $T \subset \mathrm{T}(3)$ and $R \subset \mathrm{SO}(3)$, respectively, then $(T^0 \otimes R^0) \subset \mathrm{SE}(3)$ is an $(\epsilon_1 + \epsilon_2)-$approximation of $(T \oplus R) \subset \mathrm{SE}(3)$, while $(R^0 \otimes T^0) \subset \mathrm{SE}(3)$ is an $(\epsilon_1 + r\epsilon_2)-$approximation where $r = \max_{\mathbf{t} \in T^0} \|\mathbf{t}\|_2$. This is because the left-action of a translation on a pair rotations does not alter their geodesic distance in $\mathrm{SO}(3)$ while the left-action of a rotation on a pair of translation does scale the Euclidean distance in $\mathrm{T}(3) \cong \mathbb{R}^3$ depending on distances to the origin.

One can construct regularized versions of all of the Minkowski operations over the configuration space to obtain closure when working with regular motions. Given two 6D regular motions $M_1, M_2 \in \mathscr{P}^*(\mathsf{C})$, the configuration product $(M_1 \otimes^* M_2) \in \mathscr{P}^*(\mathsf{C})$ is the set of all configurations that can be reached by an ordered combination of the operand motions. On the other hand, configuration left- and right-quotients $(M_1 \oslash^* M_2), (M_1 \osslash^* M_2) \in \mathscr{P}^*(\mathsf{C})$ are the maximal sets of all configurations whose left-

and right-combination with M_1 remains inside M_2. Note that if closed-regularization is used, $(M_1 \otimes M_2) = (M_1 \otimes^* M_2)$ but the quotients may result in degeneracies of heterogeneous dimensions.

Unlike the case with Minkowski sums and differences over the shape space, configuration products are rarely studied in the context of geometric modeling. The state-of-the-art in explicit computation relies on sampling [235, 236]; namely:

1. obtaining a dense enough sample of the individual motions;

2. computing the pairwise group products of the sample points; and

3. refining the input samples locally and recursively depending on the criteria for the sampling density at the output.

Such pairwise multiplications are easy to implement in parallel, but can get very costly for general and complex 6D motions even when implemented on high-end GPUs. Adaptive rotation sampling [236], nonuniform translation sampling [36], and other possible alternatives are rarely implemented in practice, even though elegant formal theories are available [24].

The concepts of $\epsilon-$offsetting, and $\epsilon-$blending presented earlier for Euclidean morphology to formulate the well-known solid modeling operators can be extended to configuration morphology with extra care when dealing with 6D $\epsilon-$balls that can scale along some dimensions due to the left-action of rotations on translations, as a result of (A.1.1) and (A.1.6).

3.4.2 Implicit Morphology

3.4.2.1 Measure-Theoretic Euclidean (i.e., Commutative) Morphology:

In this section we will see how the function subspace $\mathbb{R}_+^{\mathsf{E}}$ of constitutes an 'implicit morphology' on $\mathsf{E} = \mathbb{R}^3$ if it is closed, either precisely or up to equality$-\mathsf{ae}$, under summation, multiplication, extended negation, convolution, and reflection with respect to the origin.

In particular, the explicit morphologies $\mathscr{P}(\mathsf{E})$ and $\mathscr{P}^*(\mathsf{E})$ correspond respectively to the precise closure or closure up to equality$-\mathsf{ae}$ under the said analytic operations. In the following, I revisit the morphological operators discussed in Section 3.4.1.1 to present their implicit counterparts.

Boolean Union and Intersection. As demonstrated in Section 2.3.1, the set operations are implemented implicitly in the profile-oblivious and regularity-tolerant fashion via $\mathbb{R}_+-$arithmetic operations $+, \bullet\colon \mathbb{R}_+^{\mathsf{E}} \times \mathbb{R}_+^{\mathsf{E}} \to \mathbb{R}_+^{\mathsf{E}}$ on the real-valued non-negative defining functions:

$$f_{S_1 \cup S_2} \equiv (f_{S_1} + f_{S_2}) \quad \text{i.e.,} \quad (S_1 \cup S_2) = U_0(f_{S_1} + f_{S_2}), \tag{3.4.84}$$

$$f_{S_1 \cap S_2} \equiv (f_{S_1} \bullet f_{S_2}) \quad \text{i.e.,} \quad (S_1 \cap S_2) = U_0(f_{S_1} \bullet f_{S_2}), \tag{3.4.85}$$

which are dual to each other via extended logical negation $\neg : \mathbb{R}^{\mathsf{E}}_+ \to \mathbb{R}^{\mathsf{E}}_+$. Extending the diagrams in (3.4.1) using (2.3.16) one obtains:

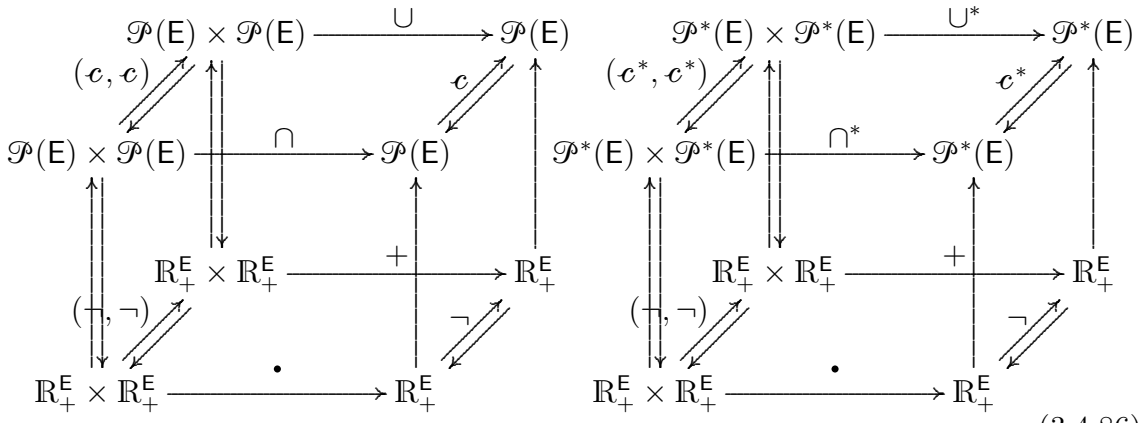

(3.4.86)

in which the vertical double-arrows should be $(U_0, U_0) \uparrow\downarrow \mathsf{desc}$ and $(U_0^*, U_0^*) \uparrow\downarrow \mathsf{desc}$ and vertical upward arrows should be $U_0 \uparrow$ and $U_0^* \uparrow$, for the left and right diagrams, respectively, eliminated due to cluttering. Once again, $U_0^*(\cdot) = \mathscr{r} U_0(\cdot) = U_0(\cdot * \delta^3)$ is the regularized $0-$superlevel set generator.

The regularization can be either taken for granted implicitly by appealing to equivalence via $\overset{*}{\equiv}$, as discussed at length in Section 2.3.1:

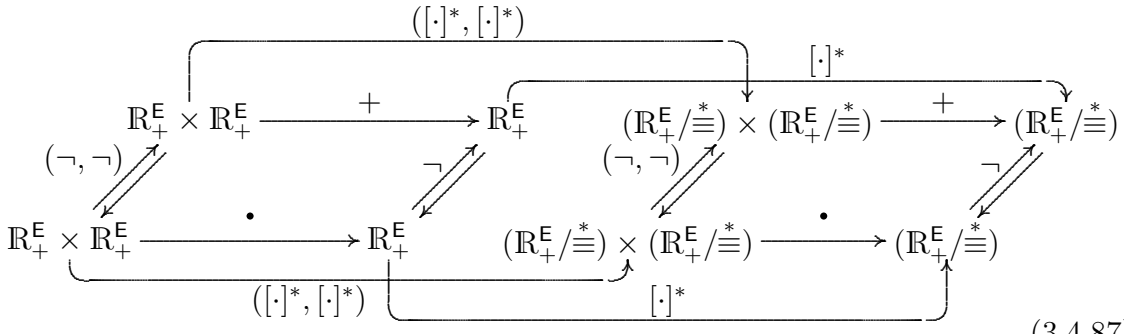

(3.4.87)

or it can achieved by applying the $\delta-$convolution filter as discussed in Section 2.3.3:

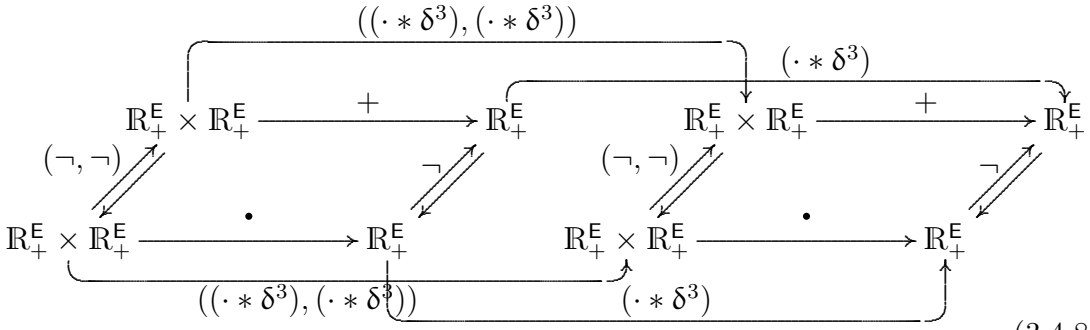

(3.4.88)

As mentioned earlier, constructive solid geometry (CSG) [282]—which was invented as one of the earliest representation schemes for design and manufacturing [277]—follows an implicit approach to modeling regularized Boolean operations. The CSG tree is essentially a network (i.e., a simplified 'Petri net' [255]) in which the allowed operations are regularized Boolean operations and rigid body motions, a finite number of which are applied sequentially starting from primitives (e.g., semialgebraic halfspaces or compact intersections of them) at the leaves of the tree. Thus it gives a mechanism to respond to inclusion and PMC queries about a shape in terms of the same queries on primitives by backtracking the CSG tree. The challenge with CSG is the need for explicit representation of neighborhoods to resolve regularized intersections, which is only solved at the implementation level on a case-by-case basis, as discussed in Section 1.3. The theory of R-functions [302] provides the most elegant formalism to implicit modeling with the level of accuracy that is needed for engineering applications [312, 313, 317], subsuming and extending the various methods that have been extensively applied to computer graphics applications such as min/max-algebras [286], blobby modeling [43], convolution surfaces [44], and a unified functional representation scheme called the F-rep [252]. However, R-functions

are not scalable to complex practical problems with numerous modeling operations. Moreover, most existing implicit methods do not provide morphological tools beyond basic Boolean operations. The rest of this section attempts to bridge this gap.

Infinite Unions as Integrals. To see how the isomorphism of Boolean and $\mathbb{R}_+-$algebras (up to equivalence via $\equiv$) leads naturally to the implementation of Minkowski operations as convolutions, consider a finite set of points $S^0 := \{\mathbf{x}_i\}_{0\leq i<n_0}$ with $|S^0| = n_0$ whose Minkowski sum with a 3D regular set $S \in \mathscr{P}^*(\mathsf{E})$ is of interest. By definition of the Minkowski sum in (3.4.3),

$$(S^0 \oplus S) = \bigcup_{0\leq i<n_0} \bigcup_{\mathbf{x}\in S} (\mathbf{x}_i + \mathbf{x}) = \bigcup_{\mathbf{x}'\in S^0} \bigcup_{\mathbf{x}\in S} (\mathbf{x}' + \mathbf{x}) = \bigcup_{0\leq i<n_0} \mathsf{act}(\gamma_0(\mathbf{x}_i))(S), \quad (3.4.89)$$

where $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$ is the embedding of the shape space $\mathsf{E} = \mathbb{R}^3$ in the translation space $\mathsf{T} = \mathrm{T}(3)$ with respect to the origin $\mathbf{0} \in \mathsf{C}$, thus $\mathbb{t}_i := \gamma_0(\mathbf{x}_i) \in \mathsf{T}$ is a translation corresponding to $\mathbf{x}_i \in \mathsf{E}$. Substituting for the analytic form of the union (i.e., multiplication) from (2.3.19) or (2.3.21) of Section 2.2.2, one obtains

$$f_{S^0\oplus S}(\mathbf{x}) := \sum_{0\leq i<n_0} f_{\mathbb{t}_i S}(\mathbf{x}) = \sum_{0\leq i<n_0} \big(f_S \circ \mathsf{act}(\mathbb{t}_i^{-1})\big)(\mathbf{x}) = \sum_{0\leq i<n_0} f_S(\mathbf{x} - \mathbf{x}_i), \quad (3.4.90)$$

Noting that defining functions are covariant under rigid transformations due to (2.3.17), i.e., $f_{\mathsf{act}(\mathbb{t})(S)} = \big(f_S \circ \mathsf{act}(\mathbb{t}^{-1})\big)$ and that $\mathsf{act}(\mathbb{t}_i^{-1})(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_i)$. The right-most term above is a discrete convolution. Thus a finite union of shifted instances of the 'template' $S \in \mathscr{P}(\mathsf{E})$ amounted to a finite summation of shifted defining functions, all of which can be queried by shifting the argument by an inverse transformation—as a direct result of the duality between explicit and implicit definition and the invert-

ibility of group operations. The discrete convolution (i.e., finite summation) can be converted to a 3D convolution (i.e., Lebesgue 3−integral) by assigning $3^{\text{rd}}-$order $\delta-$singularities at the isolated points as in (2.3.74):

$$f_{S^0 \oplus S}(\mathbf{x}) = \sum_{0 \leq i < n_0} \int_{\mathsf{E}} f_S(\mathbf{x} - \mathbf{x}') \; \delta^3(\mathbf{x}' - \mathbf{x}_i) \; d\mu^3[\mathbf{x}'] \tag{3.4.91}$$

$$= \int_{\mathsf{E}} \sum_{0 \leq i < n_0} f_S(\mathbf{x} - \mathbf{x}') \; \delta^3(\mathbf{x}' - \mathbf{x}_i) \; d\mu^3[\mathbf{x}'], \tag{3.4.92}$$

i.e., $f_{S^0 \oplus S} \equiv (f_S * \tilde{f}_{S^0})$, where $f_{S^0}(\mathbf{x}) = \sum_{0 \leq i < n_0} c_i \delta^3(\mathbf{x} - \mathbf{x}_i)$ as in (2.3.74), in which coefficients $c_i := 1$ have no impact up to equivalence via $\equiv$.

The goal is to extend this result for the finite sums (i.e., 0−integrals) to infinite sums over 1−, 2−, and 3−dimensional motions (i.e., filters) as 1−, 2−, and 3−integrals, respectively, which would then be converted to 3−integrals via proper $\delta-$singularity assignments. Since this is a search for an implicit descriptor, the implicit set-theoretic definitions are more helpful than the explicit ones. In particular, an implicit way to express the explicit union in (3.4.89) is

$$\mathbf{x} \in (S^0 \oplus S) \quad \text{iff} \quad \exists \mathbf{x}_i \in S^0 : (\gamma_0(\mathbf{x}_i))^{-1} \mathbf{x} = (\mathbf{x} - \mathbf{x}_i) \in S. \tag{3.4.93}$$

In other words, $\mathbf{x} \in (S^0 \oplus S)$ iff the set $\gamma_0(\mathbf{x})(-S^0) = \{\mathbf{x} - \mathbf{x}_i\}_{0 \leq i < n_0}$ intersects S in at least one point, i.e., $\gamma_0(\mathbf{x})(-S^0) \cap S = \emptyset$. Testing this amounts to reflecting and moving S to the query point $\mathbf{x} \in \mathsf{E}$ and counting (i.e., 0−measuring) the number of points from S^0 it overlaps. The null-volume lemma (Lemma 2.3.6) and its extension to the null d−overlap lemma (Lemma 2.3.15) allow one to extend this to measuring curve length, surface area, and solid volume of the overlaps.

Minkowski Sum and Difference. The explicit definitions of the Minkowski operations in (3.4.3) and (3.4.3) as (possibly infinite) unions and intersections are dual to the following implicit definitions, whose equivalence to the explicit forms follows from simple logical inference to resolve inclusion:

$$\begin{aligned}
\mathbf{x} \in (S_1 \oplus S_2) \quad &\text{iff} \quad \gamma_0(\mathbf{x})(-S_2) \cap S_1 \neq \emptyset \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_2) \subseteq \complement(S_1), \\
&\text{iff} \quad \gamma_0(\mathbf{x})(-S_1) \cap S_2 \neq \emptyset \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_1) \subseteq \complement(S_2). \\
\mathbf{x} \in (S_1 \ominus S_2) \quad &\text{iff} \quad \gamma_0(\mathbf{x})(-S_2) \subseteq S_1 \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_2) \cap \complement(S_1) = \emptyset, \\
&\text{iff} \quad \gamma_0(\mathbf{x})(-S_1) \subseteq S_2 \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_1) \cap \complement(S_2) = \emptyset. \qquad (3.4.94)
\end{aligned}$$

In other words, the Minkowski sum and difference characterize the two extremes of the interference possibilities between the one set and the reflection of the other set shifted to the query point, the shift being represented by $\mathbb{t} := \gamma_0(\mathbf{x})$. At a first glance, these extremes are beyond the reach of measure-theoretic tools to capture as strictly as it is defined above. However, it is possible to characterize the 'almost everywhere', meaning that the best we can hope for is to capture the 'measurable' Minkowski operations $\oplus^{\ddagger}, \ominus^{\ddagger} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ defined implicitly using the following weaker set-theoretic conditions:

$$\begin{aligned}
\mathbf{x} \in (S_1 \oplus^{\ddagger} S_2) \quad &\text{iff} \quad \gamma_0(\mathbf{x})(-S_2) \cap^{*} S_1 \neq \emptyset \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_2) \overset{\mathsf{ae}}{\subseteq} \complement(S_1), \\
&\text{iff} \quad \gamma_0(\mathbf{x})(-S_1) \cap^{*} S_2 \neq \emptyset \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_1) \overset{\mathsf{ae}}{\subseteq} \complement(S_2). \\
\mathbf{x} \in (S_1 \ominus^{\ddagger} S_2) \quad &\text{iff} \quad \gamma_0(\mathbf{x})(-S_2) \overset{\mathsf{ae}}{\subseteq} S_1 \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_2) \cap^{*} \complement(S_1) = \emptyset, \\
&\text{iff} \quad \gamma_0(\mathbf{x})(-S_1) \overset{\mathsf{ae}}{\subseteq} S_2 \quad \text{i.e.,} \quad \gamma_0(\mathbf{x})(-S_1) \cap^{*} \complement(S_2) = \emptyset. \qquad (3.4.95)
\end{aligned}$$

Lemma 2.3.6 can now be applied to test whether either of the regularized intersection terms above is empty, i.e., whether the nonregularized intersections are empty$-\mathsf{ae}$. For example, the expression $\gamma_0(\mathbf{x})(-S_2) \cap^* S_1 \neq \emptyset$ implies

$$\mathbf{x} \in (S_1 \oplus^\ddagger S_2) \quad \text{iff} \quad \mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1] > 0, \tag{3.4.96}$$

i.e., a query point $\mathbf{x} \in \mathsf{E}$ belongs to the Minkowski sum iff the reflection of the second shape translated via $\mathbb{t} := \gamma_0(\mathbf{x}) \in \mathrm{T}(3)$ is collision-free with the first shape. To compare it with the original definition of the Minkowski sum, note that:

$$\mathbf{x} \in (S_1 \oplus S_2) \quad \text{iff} \quad \mu^0[\gamma_0(\mathbf{x})(-S_2) \cap S_1] > 0. \tag{3.4.97}$$

But the measures $\mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1]$ in (3.4.96) is precisely given by the inner product of indicator functions due to Lemma 2.3.6, i.e.,

$$\mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1] = \Big\langle \mathbf{1}_{S_1}, \big(\mathbf{1}_{S_2} \circ \mathsf{act}(\gamma_0^{-1}(\mathbf{x}))\big) \Big\rangle, \tag{3.4.98}$$

noting the covariance of indicator functions under rigid transformations due to (2.3.6) leading to $\mathsf{indc}\big(\gamma_0(\mathbf{x})(-S_2)\big) = \big(\mathbf{1}_{S_2} \circ \mathsf{act}(\gamma_0^{-1}(\mathbf{x}))\big)$. The expression on the right-hand side of (3.4.98) is precisely the convolution function $(\mathbf{1}_{S_1} * \mathbf{1}_{S_2})(\mathbf{x})$ due to (2.2.7). Therefore, (3.4.96) can be rewritten as

$$\mathbf{x} \in (S_1 \oplus^\ddagger S_2) \quad \text{iff} \quad (\mathbf{1}_{S_1} * \mathbf{1}_{S_2})(\mathbf{x}) \geq 0 \quad \text{iff} \quad (f_{S_1} * f_{S_2})(\mathbf{x}) \geq 0, \tag{3.4.99}$$

in which the latter in terms of the more general nonnegative real-valued defining functions is implied by Lemma 2.3.6. Thus the convolution of indicator or defining

functions is a defining function for the Minkowski sum, i.e., $f_{S_1 \oplus^\ddagger S_2} \equiv (f_{S_1} * \tilde{f}_{S_2})$. The Minkowski difference, on the other hand, can be computed by De Morgan duality:

$$f_{S_1 \oplus^\ddagger S_2} \equiv (f_{S_1} * f_{S_2}) \quad \text{i.e.,} \quad (S_1 \oplus^\ddagger S_2) = U_0(f_{S_1} * f_{S_2}), \tag{3.4.100}$$

$$f_{S_1 \ominus^\ddagger S_2} \equiv (f_{S_1} \diamond f_{S_2}) \quad \text{i.e.,} \quad (S_1 \ominus^\ddagger S_2) = U_0(f_{S_1} \diamond f_{S_2}). \tag{3.4.101}$$

where $*, \diamond : L^1(\mathsf{E}) \times L^1(\mathsf{E}) \to L^1(\mathsf{E})$ are the convolution and its dual operator defined via $(f_1 \diamond f_2) := \neg((\neg f_1) * f_2)$, respectively.

The measurable Minkowski operations defined above are different from the ordinary Minkowski operations (i.e., $\oplus \neq \oplus^\ddagger$ and $\ominus \neq \ominus^\ddagger$). In general, they are not even equal$-\mathsf{ae}$ (i.e., $\oplus \overset{\mathsf{ae}}{\neq} \oplus^\ddagger$ and $\ominus \overset{\mathsf{ae}}{\neq} \ominus^\ddagger$), however, if the pair of shapes under considerations are regular shapes, the two become equal$-\mathsf{ae}$, thus have the same regularization, i.e., $\odot|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})} \overset{\mathsf{ae}}{=} \odot^*|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})} \overset{\mathsf{ae}}{=} \odot^\ddagger|_{\mathscr{P}^*(\mathsf{E})\times\mathscr{P}^*(\mathsf{E})}$ for $\odot \in \{\oplus, \ominus\}$. With that in mind, extending the diagrams in (3.4.5) yields:

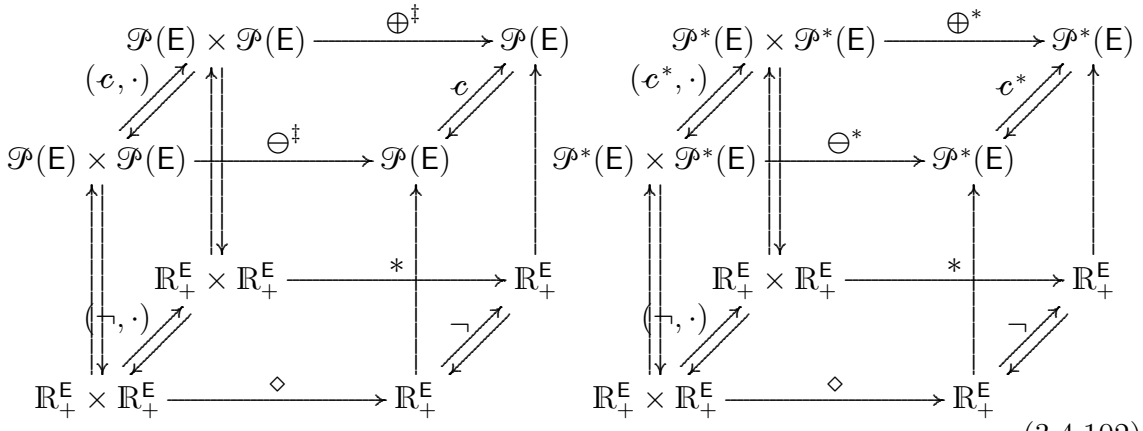

(3.4.102)

in which the vertical double-arrows should be $(U_0, U_0) \uparrow\downarrow \mathsf{desc}$ and $(U_0^*, U_0^*) \uparrow\downarrow \mathsf{desc}$ and vertical upward arrows should be $U_0 \uparrow$ and $U_0^* \uparrow$, for the left and right diagrams, respectively, eliminated due to cluttering. As before, $U_0^*(\cdot) = \mathscr{r}U_0(\cdot) = U_0(\cdot * \delta^3)$ is

the regularized $0-$superlevel set generator.

Once again, note that the measurable Minkowski sum and difference cannot capture lower-dimensional features, e.g., sweep or unsweep of points, curves, and surfaces along points, curves, and surfaces, unless the shape descriptor functions are properly assigned with $\delta-$singularities, making them artificially $3-$measurable, thus observable to the convolution $3-$integral. If all dangling $\mathrm{d}-$elements (e.g., as in Fig. 3.4.1) are properly described by $(3-\mathrm{d})^{\mathrm{th}}-$order $\delta-$singularities, the $(\oplus^{\ddagger}, \ominus^{\ddagger})$ in the above and following diagrams can be replaced with the ordinary $(\oplus, \ominus)$ with some care.

Minkowski Difference Revisited. It is important to recognize that although both of the convolution $(f_1 * f_2)$ and its dual $(f_1 \diamond f_2) = \neg(f_1 * (\neg f_2))$ are computable for compactly supported functions $f_1, f_2 \in L^1(\mathsf{E})$ (e.g., $f_1 := f_{S_1}$ and $f_2 := f_{S_2}$), the intermediate function $\neg(f_1 \diamond f_2) = (f_1 * (\neg f_2))$ is not computable due to the unbounded $0-$superlevel set (i.e., $\|\neg f_{S_2}\|_1 \to \infty$ because $\mu^3[\mathfrak{c} S_2] \to \infty$). Therefore, the De Morgan duality may not be useful in practice to compute the dual of convolution for Minkowski difference operation. The alternative route is to go back to the implicit definitions in (3.4.95) and observe that the expression $\gamma_0(\mathbf{x})(-S_2) \overset{\mathsf{ae}}{\subseteq} S_1$ implies

$$\mathbf{x} \in (S_1 \ominus^{\ddagger} S_2) \quad \text{iff} \quad \mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1] = \mu^3[S_2], \tag{3.4.103}$$

i.e., a query point $\mathbf{x} \in \mathsf{E}$ belongs to the Minkowski difference iff the reflection of the second shape translated via $\mathbb{t} := \gamma_0(\mathbf{x}) \in \mathrm{T}(3)$ is contained$-\mathsf{ae}$ inside the first shape. Noting that in general, the measure of intersection of two sets cannot exceed the measure of either set, i.e., $\mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1] \leq \mu^3[S_2]$, the equality condition in

(3.4.103) can be safely replaced with an inequality, i.e.,

$$\mathbf{x} \in (S_1 \ominus^{\ddagger} S_2) \quad \text{iff} \quad \mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1] \geq \mu^3[S_2], \tag{3.4.104}$$

But $\mu^3[S_2] = \|\mathbf{1}_{S_2}\|_1$ and $\mu^3[\gamma_0(\mathbf{x})(-S_2) \cap S_1]$ is given by the inner product of indicator functions in (3.4.98). The latter is precisely the convolution $(\mathbf{1}_{S_1} * \tilde{\mathbf{1}}_{S_2})(\mathbf{x})$ due to Definition 2.2.8, thus (3.4.103) can be rewritten as

$$\mathbf{x} \in (S_1 \ominus^{\ddagger} S_2) \quad \text{iff} \quad (\mathbf{1}_{S_1} * \mathbf{1}_{S_2})(\mathbf{x}) - \|\mathbf{1}_{S_2}\|_1 \geq \ 0. \tag{3.4.105}$$

To put the expression on the right-hand side in the standard form of $f_{S_1 \ominus^{\ddagger} S_2}(\mathbf{x}) > 0$,

$$\mathbf{x} \in (S_1 \ominus^{\ddagger} S_2) \quad \text{iff} \quad (\mathbf{1}_{S_1} * \mathbf{1}_{S_2})(\mathbf{x}) - \|\mathbf{1}_{S_2}\|_1 > -\epsilon, \tag{3.4.106}$$

for all $\epsilon > 00$. Thus the $(-\epsilon)-$superlevel set of the difference on the right hand side as $\epsilon \to 0^+$ returns the Minkowski difference, and the expression $(\mathbf{1}_{S_1} * \mathbf{1}_{S_2}) - \|\mathbf{1}_{S_2}\|_1 + \epsilon$ is a defining function for the Minkowski difference in the limit $\epsilon \to 0^+$, which can be turned into a nonnegative defining function as:

$$f_{S_1 \ominus^{\ddagger} S_2} \equiv \lim_{\epsilon \to 0^+} \min\big\{0, (\mathbf{1}_{S_1} * \mathbf{1}_{S_2}) - \|\mathbf{1}_{S_2}\|_1 + \epsilon\big\}. \tag{3.4.107}$$

Hence the Minkowski difference can be directly computed from the convolution of indicator functions, without using its dual operation via De Morgan laws:

$$(S_1 \ominus^{\ddagger} S_2) = \lim_{\epsilon \to 0^+} U_{-\epsilon}\big((\mathbf{1}_{S_1} * \mathbf{1}_{S_2}) - \|\mathbf{1}_{S_2}\|_1\big) = \lim_{\epsilon \to 0^+} U_{\|\mathbf{1}_{S_2}\|_1 - \epsilon}(\mathbf{1}_{S_1} * \mathbf{1}_{S_2}), \tag{3.4.108}$$

which can be rewritten as follows, if the indicator functions are retrieved from a different defining function (e.g., obtained from an upstream operation):

$$(S_1 \ominus^{\ddagger} S_2) = U_{\|(\mathsf{sign}\circ\mathsf{desc})(S_2)\|_1^-}\Big((\mathsf{sign}\circ\mathsf{desc})(S_1) * (\mathsf{sign}\circ\mathsf{desc})(S_2)\Big), \tag{3.4.109}$$

in which the short notation $U_{t^\pm}(\cdot) := \lim_{t'\to t^\pm} U_{t'}(\cdot)$ is used. Therefore, both implicit Minkowski sum and difference can be obtained directly from the same convolution function (without appealing to De Morgan duality) using different $t-$level sets generators. The commutative diagram in (3.4.102) is thus replaced with:

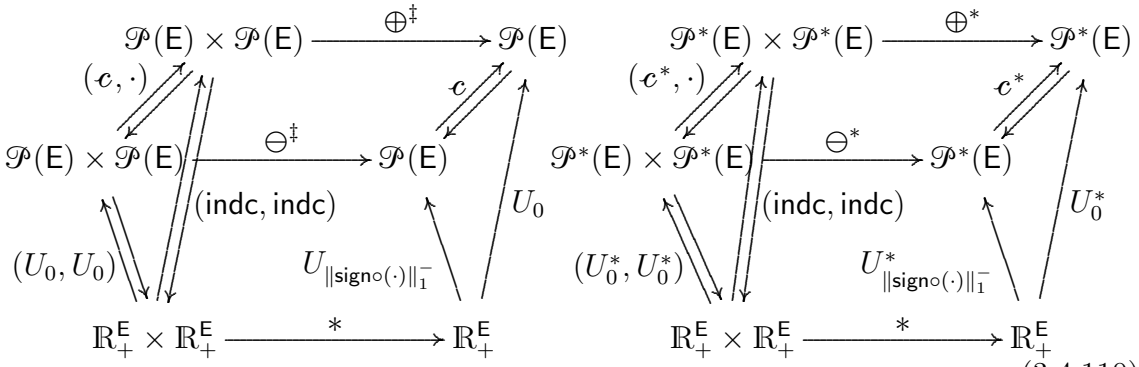

(3.4.110)

To deliver the great promise of the convolution theorem (Theorem 2.2.11), the convolution in (3.4.110) can be computed via Fourier transforms:

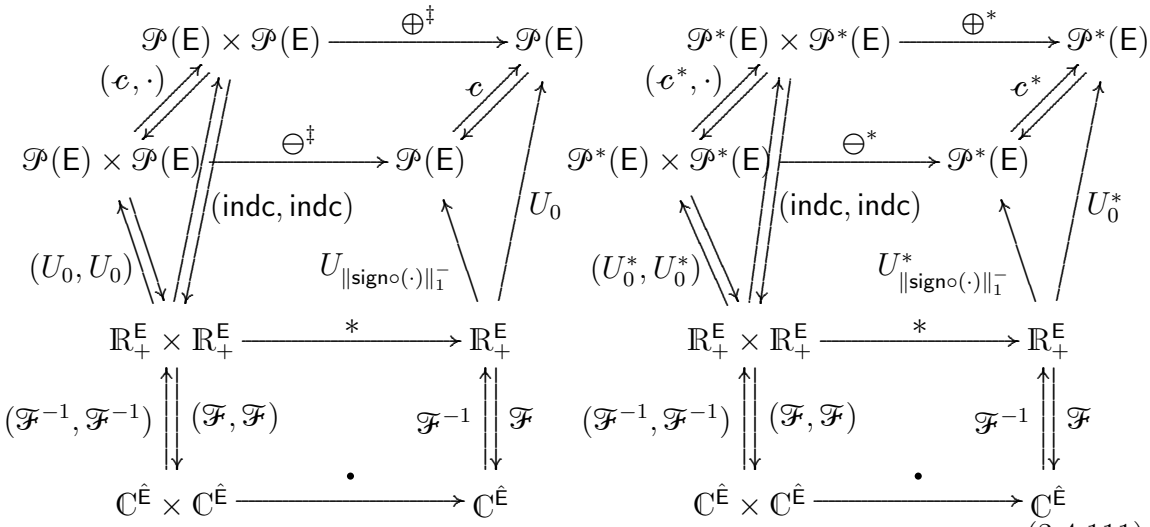

(3.4.111)

Similar to the case with integration in Section 3.2.2.1, Fourier convolution can be approximated by low-pass filtering, e.g., using the indicator function of a compact window $\hat{W}_r \subset \hat{\mathsf{E}}$ around the origin denoted by $\mathbf{1}_{\hat{W}_r} : \hat{\mathsf{E}} \to \{0, 1\}$:

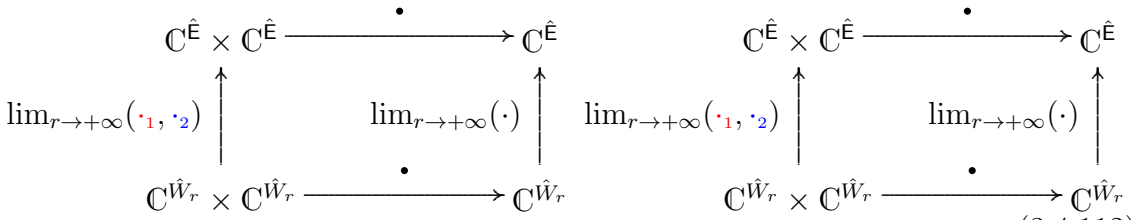

(3.4.112)

Combining (3.4.111) and (3.4.112) into a single diagram illustrates the full picture:

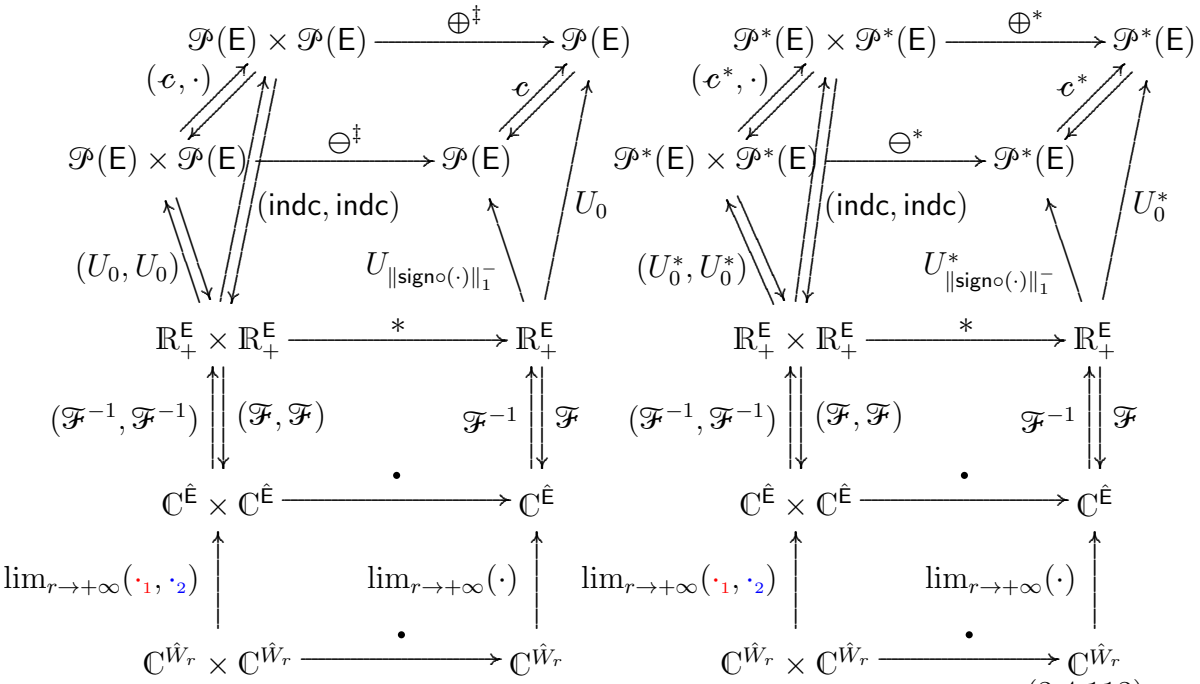

(3.4.113)

Morphological Dilation and Erosion. Having obtained the implicit descriptors of the measurable Minkowski operations, the measurable dilation and erosion $\mathsf{dil}^{\ddagger}, \mathsf{ero}^{\ddagger} : \mathscr{P}(\mathsf{T}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ are implicitly described as

$$f_{\mathsf{dil}^{\ddagger}(T,S)} \equiv f_{S\oplus^{\ddagger}(-\gamma_0^{-1}(T))} \equiv (f_S * \tilde{f}_{\gamma_0^{-1}(T)}) \equiv \big(f_S * (\tilde{f}_T \circ \gamma_0)\big), \tag{3.4.114}$$

$$f_{\mathsf{ero}^{\ddagger}(T,S)} \equiv f_{S\ominus^{\ddagger}(-\gamma_0^{-1}(T))} \equiv (f_S \diamond \tilde{f}_{\gamma_0^{-1}(T)}) \equiv \big(f_S \diamond (\tilde{f}_T \circ \gamma_0)\big), \tag{3.4.115}$$

which can be expressed alternatively in terms of superlevel sets as:

$$\mathsf{dil}^{\ddagger}(T,S) = U_0(f_S * \tilde{f}_{\gamma_0^{-1}(T)}) = U_0\big(f_S * (\tilde{f}_T \circ \gamma_0)\big), \tag{3.4.116}$$

$$\mathsf{ero}^{\ddagger}(T,S) = U_0(f_S \diamond \tilde{f}_{\gamma_0^{-1}(T)}) = U_0\big(f_S \diamond (\tilde{f}_T \circ \gamma_0)\big), \tag{3.4.117}$$

where, once again, the dual of convolution can be expressed as:

$$f_{\mathsf{ero}^{\ddagger}(T,S)} \equiv \big(\mathbf{1}_S \diamond (\tilde{\mathbf{1}}_T \circ \gamma_0)\big) \equiv \min\Big\{0, \big(\mathbf{1}_S * (\tilde{\mathbf{1}}_T \circ \gamma_0)\big) - \|\mathbf{1}_T\|_1\Big\}, \tag{3.4.118}$$

$$\text{i.e.,} \quad \mathsf{ero}^{\ddagger}(T,S) = U_{\|\mathbf{1}_T\|_1}\big(\mathbf{1}_S * (\tilde{\mathbf{1}}_T \circ \gamma_0)\big), \tag{3.4.119}$$

in which, as always, $\mathbf{1}_S = (\mathsf{sign} \circ f_S)$ and $\mathbf{1}_T = (\mathsf{sign} \circ f_T)$ can be used for the generic defining functions that come from upstream operations.

The diagram in (3.4.113) can be combined with (3.4.11) to obtain a convolution-based computation for dilation and erosion:

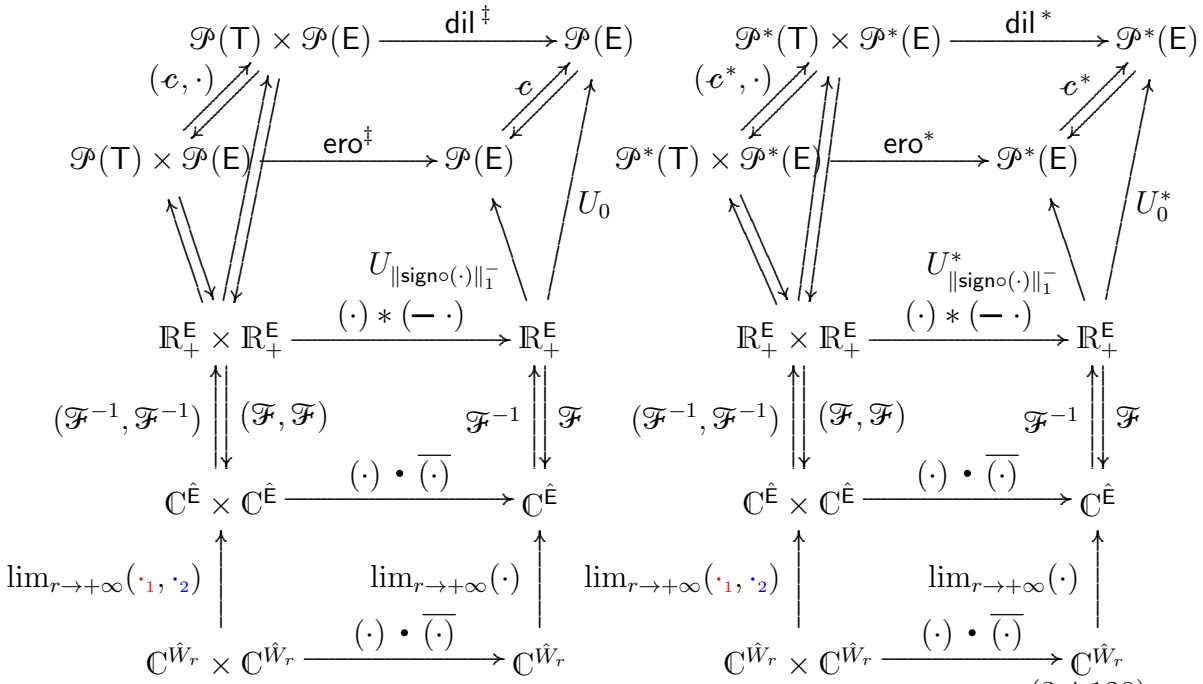

(3.4.120)

in which, unlike (3.4.113), the labels for the left and right tilted double-arrows should be $((\gamma_0 \circ U_0), U_0) \uparrow\downarrow ((\mathsf{indc} \circ \gamma_0^{-1}), \mathsf{indc})$ and $((\gamma_0 \circ U_0^*), U_0^*) \uparrow\downarrow ((\mathsf{indc} \circ \gamma_0^{-1}), \mathsf{indc})$, respectively, in addition to swapping the order of arguments, omitted from the diagram due to cluttering. The Hermitian symmetry of the Fourier transform of the real-valued defining functions was exploited to convert the reflection in the physical domain to a conjugation in the frequency domain $(- \cdot) \overset{\mathscr{F}}{\to} \overline{(\cdot)}$.

Morphological Closing and Opening. Since morphological closing and opening are defined in terms of a sequence of a dilation and an erosion and vice versa in (3.4.13) and (3.4.15), respectively, they can be obtained in the obvious way from a sequence of convolutions—also using the sign function to convert the obtained defining

functions to indicator functions after each convolution. Accordingly, the measurable closing and opening $\mathsf{cls}^{\ddagger}, \mathsf{opn}^{\ddagger} : \mathscr{P}(\mathsf{T}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ are implicitly described as

$$f_{\mathsf{cls}^{\ddagger}(T,S)} \equiv f_{\mathsf{ero}(T,\mathsf{dil}^{\ddagger}(T^{-1},S))} \equiv \big((f_S * \tilde{f}_{\gamma_0^{-1}(T)}) \diamond f_{\gamma_0^{-1}(T)}\big) \tag{3.4.121}$$

$$\equiv \big((f_S * (\tilde{f}_T \circ \gamma_0)) \diamond (f_T \circ \gamma_0)\big), \tag{3.4.122}$$

$$f_{\mathsf{opn}^{\ddagger}(T,S)} \equiv f_{\mathsf{dil}\,(T,\mathsf{ero}^{\ddagger}(T^{-1},S))} \equiv \big((f_S \diamond \tilde{f}_{\gamma_0^{-1}(T)}) * f_{\gamma_0^{-1}(T)}\big) \tag{3.4.123}$$

$$\equiv \big((f_S \diamond (\tilde{f}_T \circ \gamma_0)) * (f_T \circ \gamma_0)\big), \tag{3.4.124}$$

which can be expressed alternatively in terms of superlevel sets as:

$$\mathsf{cls}^{\ddagger}(T,S) = U_0\big((f_S * (\tilde{f}_T \circ \gamma_0)) \diamond (f_T \circ \gamma_0)\big), \tag{3.4.125}$$

$$\mathsf{opn}^{\ddagger}(T,S) = U_0\big((f_S \diamond (\tilde{f}_T \circ \gamma_0)) * (f_T \circ \gamma_0)\big). \tag{3.4.126}$$

Minkowski Closing and Opening. Since Minkowski closing and opening are defined in terms of a sequence of a dilation and an erosion and vice versa in (3.4.17) and (3.4.18), respectively, they can be obtained in the obvious way from a sequence of convolutions—also using the sign function to convert the obtained defining functions to indicator functions after each convolution. Accordingly, the measurable closing and opening $\oplus^{\ddagger}, \ominus^{\ddagger}$: $\mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ are implicitly described as

$$f_{S_1 \oplus^{\ddagger} S_2} \equiv f_{(S_1 \oplus S_2) \ominus (-S_2)} \equiv \big((f_{S_1} * f_{S_2}) \diamond \tilde{f}_{S_2}\big), \tag{3.4.127}$$

$$f_{S_1 \ominus^{\ddagger} S_2} \equiv f_{(S_1 \ominus S_2) \oplus (-S_2)} \equiv \big((f_{S_1} \diamond f_{S_2}) * \tilde{f}_{S_2}\big), \tag{3.4.128}$$

which can be expressed alternatively in terms of superlevel sets as:

$$(S_1 \oplus^{\ddagger} S_2) = U_0\big((f_{S_1} * f_{S_2}) \diamond \tilde{f}_{S_2}\big), \tag{3.4.129}$$

$$(S_1 \ominus^{\ddagger} S_2) = U_0\big((f_{S_1} \diamond f_{S_2}) * \tilde{f}_{S_2}\big). \tag{3.4.130}$$

Digital Morphology. Given two finite sets of points $S_1^0 = \{\mathbf{x}_{i_1}^{(1)}\}_{0\leq i_1<n_1}$ and $S_2^0 = \{\mathbf{x}_{i_2}^{(2)}\}_{0\leq i_2<n_2}$ with cardinalities $|S_1^0| = n_1$ and $|S_2^0| = n_2$, their implicit descriptors are two finite sums of $3^{\text{rd}}-$order $\delta-$functions (i.e., a discrete set of 'impacts'):

$$f^3_{S_1^0}(\mathbf{x}) = \sum_{0\leq i_1<n_1} c_{i_1}^{(1)} \varsigma^3_{\mathbf{x}_{i_1}^{(1)}}(\mathbf{x}) = \sum_{0\leq i_1<n_1} c_{i_1}^{(1)} \delta^3(\mathbf{x} - \mathbf{x}_{i_1}^{(1)}), \tag{3.4.131}$$

$$f^3_{S_2^0}(\mathbf{x}) = \sum_{0\leq i_2<n_2} c_{i_2}^{(2)} \varsigma^3_{\mathbf{x}_{i_2}^{(2)}}(\mathbf{x}) = \sum_{0\leq i_2<n_2} c_{i_2}^{(2)} \delta^3(\mathbf{x} - \mathbf{x}_{i_2}^{(2)}). \tag{3.4.132}$$

Their Minkowski sum's implicit descriptor is obtained by a convolution:

$$f^3_{S_1^0\oplus S_2^0} \equiv (f^3_{S_1} * f^3_{S_2}) = \sum_{0\leq i_1<n_1} \sum_{0\leq i_2<n_2} c_{i_1}^{(1)} c_{i_2}^{(2)} \ (\varsigma^3_{\mathbf{x}_{i_1}^{(1)}} * \varsigma^3_{-\mathbf{x}_{i_2}^{(2)}}), \tag{3.4.133}$$

noting that convolutions distribute over summations due to linearity. But we know that the convolution of the to shifted $\delta-$functions is another shifted $\delta-$function:

$$(\varsigma^3_{\mathbf{x}_{i_1}^{(1)}} * \varsigma^3_{-\mathbf{x}_{i_2}^{(2)}})(\mathbf{x}) = \int_{\mathsf{E}} \delta^3(\mathbf{x}' - \mathbf{x}_{i_1}^{(1)}) \delta^3(\mathbf{x} - (\mathbf{x}' + \mathbf{x}_{i_2}^{(2)}))\ d\mu^3[\mathbf{x}'] \tag{3.4.134}$$

$$= \int_{\mathsf{E}} \delta^3(\mathbf{x}') \delta^3(\mathbf{x} - (\mathbf{x}' + \mathbf{x}_{i_1}^{(1)} + \mathbf{x}_{i_2}^{(2)}))\ d\mu^3[\mathbf{x}'], \tag{3.4.135}$$

obtained from a change of variables $\mathbf{x}' \mapsto (\mathbf{x}' + \mathbf{x}_{i_1}^{(1)})$. Letting $\mathbf{x}_{i_1,i_2} := (\mathbf{x}_{i_1}^{(1)} + \mathbf{x}_{i_2}^{(2)})$ and noting that the first $\delta-$function on the right-hand side sifts the roots of the

second $\delta-$function—i.e., using the sifting lemma (Lemma 2.3.16) with $S := \mathsf{E}$ and $f(\mathbf{x}') := \delta^3(\mathbf{x} - (\mathbf{x}' + \mathbf{x}_{i_1,i_2}))$ thus $f(\mathbf{0}) = \delta^3(\mathbf{x} - \mathbf{x}_{i_1,i_2}) = \varsigma^3_{\mathbf{x}_{i_1,i_2}}(\mathbf{x})$—one obtains $(\varsigma^3_{\mathbf{x}^{(1)}_{i_1}} * \varsigma^3_{-\mathbf{x}^{(2)}_{i_2}}) = \varsigma^3_{\mathbf{x}_{i_1,i_2}}$. Letting $c_{i_1,i_2} := c^{(1)}_{i_1} c^{(2)}_{i_2}$, (3.4.133) can be rewritten as:

$$f^3_{S^0_1 \oplus S^0_2} \equiv (f^3_{S_1} * f^3_{S_2}) = \sum_{0 \le i_1 < n_1} \sum_{0 \le i_2 < n_2} c_{i_1,i_2} \varsigma^3_{\mathbf{x}_{i_1,i_2}} = \sum_{0 \le i < n} c^*_i \varsigma^3_{\mathbf{x}^*_i}, \tag{3.4.136}$$

which is the implicit descriptor of $|S^0_1 \oplus S^0_2| =: n \le n_1 n_2$ discrete points (i.e., another set of 'impacts'), with the possibility of repetitions because one may have $\mathbf{x}_{i_1,i_2} = \mathbf{x}_{i'_1,i'_2}$, i.e., $(\mathbf{x}^{(1)}_{i_1} + \mathbf{x}^{(2)}_{i_2}) = (\mathbf{x}^{(1)}_{i'_1} + \mathbf{x}^{(2)}_{i'_2})$ for some $0 \le i_1, i'_1 < n_1$ and $0 \le i_2, i'_2 < n_2$, in which case the intensities c_{i_1,i_2} and $c_{i'_1,i'_2}$ are added up. Thus $(S^0_1 \oplus S^0_2) = \{\mathbf{x}^*_i\}_{0 \le i < n}$ for which the implicit descriptor has intensities $c^*_i = \sum_{\mathbf{x}^*_i = (\mathbf{x}^{(1)}_{i_1} + \mathbf{x}^{(2)}_{i_2})} c_{i_1,i_2}$ for $0 \le i < n$.

The above reasoning could have been done in the Fourier domain in a much simpler fashion. For the finite summations of $\delta-$functions in (3.4.131) and (3.4.132), the Fourier descriptors are obtained as

$$\hat{f}^3_{S^0_1}(\boldsymbol{\omega}) = \sum_{0 \le i_1 < n_1} c^{(1)}_{i_1} \hat{\varsigma}^3_{\mathbf{x}^{(1)}_{i_1}}(\boldsymbol{\omega}) = \sum_{0 \le i_1 < n_1} c^{(1)}_{i_1} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}^{(1)}_{i_1})}, \tag{3.4.137}$$

$$\hat{f}^3_{S^0_2}(\boldsymbol{\omega}) = \sum_{0 \le i_2 < n_2} c^{(2)}_{i_2} \hat{\varsigma}^3_{\mathbf{x}^{(2)}_{i_2}}(\boldsymbol{\omega}) = \sum_{0 \le i_2 < n_2} c^{(2)}_{i_2} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}^{(2)}_{i_2})}, \tag{3.4.138}$$

noting that the Fourier transform of the shifted $\delta-$function is a sinusoidal basis function $\hat{\varsigma}^3_{\mathbf{x}'}(\boldsymbol{\omega}) = e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}')}$—obtained by applying the sifting lemma (Lemma 2.3.16) to the definition of the Fourier transform in (2.2.11) (Definition 2.2.9). Then the

convolution theorem (Theorem 2.2.11) and the Hermitian symmetry lead to:

$$\hat{f}^3_{S_1^0 \oplus S_2^0} \equiv \mathscr{F}\{f^3_{S_1} * f^3_{S_2}\} = (\hat{f}^3_{S_1} \bullet \hat{f}^3_{S_2}) \tag{3.4.139}$$

$$= \sum_{0 \le i_1 < n_1} \sum_{0 \le i_2 < n_2} c^{(1)}_{i_1} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}^{(1)}_{i_1})} c^{(2)}_{i_2} e^{+2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}^{(2)}_{i_2})} \tag{3.4.140}$$

$$= \sum_{0 \le i_1 < n_1} \sum_{0 \le i_2 < n_2} c^{(1)}_{i_1} c^{(2)}_{i_2} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot (\mathbf{x}^{(1)}_{i_1} + \mathbf{x}^{(2)}_{i_2}))} \tag{3.4.141}$$

$$= \sum_{0 \le i_1 < n_1} \sum_{0 \le i_2 < n_2} c_{i_1,i_2} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}_{i_1,i_2})} = \sum_{0 \le i < n} c^*_i e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}^*_i)}, \tag{3.4.142}$$

which is the Fourier transform of (3.4.136), where, once again, $c_{i_1,i_2} := c^{(1)}_{i_1} c^{(2)}_{i_2}$, $\mathbf{x}_{i_1,i_2} := (\mathbf{x}^{(1)}_{i_1} + \mathbf{x}^{(2)}_{i_2})$ and $c^*_i = \sum_{\mathbf{x}^*_i = (\mathbf{x}^{(1)}_{i_1} + \mathbf{x}^{(2)}_{i_2})} c_{i_1,i_2}$ for $0 \le i < n$.

If all points are assigned with the same weight of unity, i.e., $c^{(1)}_{i_1} = c^{(2)}_{i_2} = 1$ for all $0 \le i_1 < n_1$ and $0 \le i_2 < n_2$, then $f^3_{S_1^0} = \mathbb{1}^3_{S_1^0}$ and $f^3_{S_2^0} = \mathbb{1}^3_{S_2^0}$. Then, $c_{i_1,i_2} = 1$ as well, and the resulting c^*_i counts the multiplicity of the point $\mathbf{x}^*_i$, i.e., the number of times a pair of points are summed to the same point.[28] Note also that the L^1-norm $\left\| f^3_{S_1^0 \oplus S_2^0} \right\|_1 = n = n_1 n_2$ counts the number of points by applying the multiplicities.

The results can be illustrated via the following commutative diagrams:

$$\begin{array}{ccc}
\mathscr{P}^0(\mathsf{E}) \times \mathscr{P}^0(\mathsf{E}) & \xrightarrow{\quad\oplus\quad} & \mathscr{P}^0(\mathsf{E}) \\
(U_0^0, U_0^0) \Big\updownarrow (\mathsf{desc}^0, \mathsf{desc}^0) & & U_0^0 \Big\updownarrow \mathsf{desc}^0 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} \times (\mathbb{R}^0_\star)^{\mathsf{E}} & \xrightarrow{\quad *\quad} & (\mathbb{R}^0_\star)^{\mathsf{E}} \\
\Big\downarrow (\mathscr{F}, \mathscr{F}) & & \mathscr{F}^{-1} \Big\updownarrow \mathscr{F} \\
\mathbb{C}^{\hat{\mathsf{E}}} \times \mathbb{C}^{\hat{\mathsf{E}}} & \xrightarrow{\quad\bullet\quad} & \mathbb{C}^{\hat{\mathsf{E}}}
\end{array} \tag{3.4.143}$$

[28] If the original points are in 'general conditions', overlaps almost surely never happen, thus multiplicities of the resultant points in the Minkowski sum are almost surely always unity, thus $\mathbb{1}^3_{S_1^0 \oplus S_2^0} \equiv f^3_{S_1^0 \oplus S_2^0} := (\mathbb{1}^3_{S_1^0} * \mathbb{1}^3_{S_2^0})$ can almost surely be rewritten as $\mathbb{1}^3_{S_1^0 \oplus S_2^0} = (\mathbb{1}^3_{S_1^0} * \mathbb{1}^3_{S_2^0})$.

Note that low-pass filtering cannot be applied in this case to derive an approximation because the frequency domain representation is oscillatory—i.e., with no decay in the Fourier modes as one looks into higher frequencies.

Let us now consider a different approximation scheme by replacing the $\delta-$functions in (3.4.131) and (3.4.132) with their limit representations:

$$f^3_{S^0_1}(\mathbf{x}) = \lim_{\sigma_1 \to 0^+} \sum_{0 \leq i_1 < n_1} c^{(1)}_{i_1} g^3_{\sigma_1}(\mathbf{x} - \mathbf{x}^{(1)}_{i_1}), \tag{3.4.144}$$

$$f^3_{S^0_2}(\mathbf{x}) = \lim_{\sigma_2 \to 0^+} \sum_{0 \leq i_2 < n_2} c^{(2)}_{i_2} g^3_{\sigma_2}(\mathbf{x} - \mathbf{x}^{(2)}_{i_2}). \tag{3.4.145}$$

It can be shown that in general, the convolution of two nascent $\delta-$functions is also a nascent $\delta-$function (of possibly different type and thickness factor):

$$f^3_{S^0_1 \oplus S^0_2} \equiv (f^3_{S_1} * f^3_{S_2}) = \lim_{\sigma \to 0^+} \sum_{0 \leq i_1 < n_1} \sum_{0 \leq i_2 < n_2} c_{i_1,i_2} g^3_{\sigma'}((\cdot) - \mathbf{x}_{i_1,i_2}), \tag{3.4.146}$$

In fact, if the Gaussian function is used, the convolution also turns out to be a Gaussian with a thickness factor $\sigma' = (\sigma_1^2 + \sigma_2^2)^{\frac{1}{2}}$. Using limit representations to compute implicit Minkowski sums on discrete poinsets can be illustrated via:

$$\begin{array}{ccc}
\mathscr{P}^0(\mathsf{E}) \times \mathscr{P}^0(\mathsf{E}) & \xrightarrow{\quad\oplus\quad} & \mathscr{P}^0(\mathsf{E}) \\
\Big\uparrow\Big\downarrow (U^0_0, U^0_0) \quad\quad (\mathsf{desc}^0, \mathsf{desc}^0) & & U^0_0 \quad \Big\uparrow\Big\downarrow \mathsf{desc}^0 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} \times (\mathbb{R}^0_\star)^{\mathsf{E}} & \xrightarrow{\quad *\quad} & (\mathbb{R}^0_\star)^{\mathsf{E}} \\
\lim_{\sigma \to 0^+ (\cdot_1, \cdot_2)} \Big\uparrow & & \lim_{\sigma \to 0^+ (\cdot)} \Big\uparrow \\
\mathscr{G}^0_\sigma \times \mathscr{G}^0_\sigma & \xrightarrow{\quad *\quad} & \mathscr{G}^0_\sigma
\end{array} \tag{3.4.147}$$

Interestingly, Fourier transforms of Gaussian functions are also Gaussian functions,

leading to a neat mathematical description of the problem in the Fourier domain, which I choose to skip here.

Groupings as Dilations. Given a finite set of points $S^0 = \{\mathbf{x}_i\}_{0\leq i<n_0} \subset \mathsf{E}$ with $|S^0| = n_0$, and an $r-$ball $B_r(\mathbf{0})$, the implicit description of the grouping $S^0 \oplus B_r(\mathbf{0})$ is obtained as the following convolution:

$$f^0_{S^0\oplus B_r(\mathbf{0})} \equiv (f^3_{S_0} * f^0_{B_r(\mathbf{0})}) = \Big(\sum_{0\leq i<n_0} \varsigma^3_{\mathbf{x}_i}\Big) * f^0_{B_r(\mathbf{0})} = \sum_{0\leq i<n_0} (\varsigma^3_{\mathbf{x}_i} * f^0_{B_r(\mathbf{0})}), \tag{3.4.148}$$

noting that convolutions distribute over summations due to linearity. Note that $f^3_{S^0} \in (\mathbb{R}^0_\star)^{\mathsf{E}}$ is the $3^{\text{rd}}-$order $\delta-$singular descriptor of the discrete pointset $S^0 \in \mathscr{P}^0(\mathsf{E})$ while $f^0_{B_r(\mathbf{0})}, f^0_{S^0\oplus B_r(\mathbf{0})} \in (\mathbb{R}^3_\star)^{\mathsf{E}}$ are the nonsingular defining functions of the 3D ball and the grouping of its n_0 copies centered at the discrete pointset. Each individual convolution with a shifted $\delta-$function yields:

$$(\varsigma^3_{\mathbf{x}_i} * f^0_{B_r(\mathbf{0})})(\mathbf{x}) = \int_{\mathsf{E}} \delta^3(\mathbf{x}' - \mathbf{x}_i) f^0_{B_r(\mathbf{0})}(\mathbf{x} - \mathbf{x}')\ d\mu^3[\mathbf{x}'] \tag{3.4.149}$$

$$= \int_{\mathsf{E}} \delta^3(\mathbf{x}') f^0_{B_r(\mathbf{0})}(\mathbf{x} - \mathbf{x}_i - \mathbf{x}')\ d\mu^3[\mathbf{x}'], \tag{3.4.150}$$

obtained from a change of variables $\mathbf{x}' \mapsto (\mathbf{x}' + \mathbf{x}_i)$. Once again, using the sifting lemma (Lemma 2.3.16) with $S := \mathsf{E}$ and $f(\mathbf{x}') := f^0_{B_r(\mathbf{0})}(\mathbf{x} - \mathbf{x}_i - \mathbf{x}')$ which gives $f(\mathbf{0}) = f^0_{B_r(\mathbf{0})}(\mathbf{x} - \mathbf{x}_i) = f^0_{B_r(\mathbf{x}_i)}(\mathbf{x})$, (3.4.148) can be rewritten as

$$f^0_{S^0\oplus B_r(\mathbf{0})} \equiv \sum_{0\leq i<n_0} (\varsigma^3_{\mathbf{x}_i} * f^0_{B_r(\mathbf{0})}) = \sum_{0\leq i<n_0} f^0_{B_r(-\mathbf{x}_i)} \equiv \sum_{0\leq i<n_0} \mathbf{1}^0_{B_r(-\mathbf{x}_i)}, \tag{3.4.151}$$

The same result could be developed in the Fourier domain, where convolution is substituted by pointwise multiplication:

$$\hat{f}^0_{S^0 \oplus B_r(\mathbf{0})} \equiv (\hat{f}^3_{S_0} \bullet \hat{f}^0_{B_r(\mathbf{0})}) = \Big(\sum_{0 \leq i < n_0} \hat{\varsigma}^3_{\mathbf{x}_i} \Big) \bullet \hat{f}^0_{B_r(\mathbf{0})} = \sum_{0 \leq i < n_0} (\hat{\varsigma}^3_{\mathbf{x}_i} \bullet \hat{f}^0_{B_r(\mathbf{0})}), \quad (3.4.152)$$

noting that the Fourier transform of the shifted $\delta-$function is a sinusoidal basis function $\hat{\varsigma}^3_{\mathbf{x}'}(\boldsymbol{\omega}) = e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}')}$—obtained by applying the sifting lemma (Lemma 2.3.16) to the definition of the Fourier transform in (2.2.11) (Definition 2.2.9). It appears that the above equation is a discrete Fourier transform (DFT) using the discrete functional basis defined by the finite pointset.

$$\hat{f}^0_{S^0 \oplus B_r(\mathbf{0})}(\boldsymbol{\omega}) \equiv \sum_{0 \leq i < n_0} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}_i)} \hat{f}^0_{B_r(\mathbf{0})}(\boldsymbol{\omega}) \equiv \sum_{0 \leq i < n_0} e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}_i)} \hat{\mathbf{1}}^0_{B_r(\mathbf{0})}(\boldsymbol{\omega}). \quad (3.4.153)$$

The results can be illustrated via the following commutative diagrams:

$$\begin{array}{ccc}
\mathscr{P}^0(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \oplus B_{(\cdot_2)}(\mathbf{0})} & \mathscr{P}^3(\mathsf{E}) \\
\Big\uparrow\Big\downarrow (U^0_0, \cdot_2) \quad \Big\uparrow\Big\downarrow (\mathsf{desc}^0, \cdot_2) & & \Big\uparrow\Big\downarrow U^3_0 \quad \Big\uparrow\Big\downarrow \mathsf{desc}^3 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) * \mathbf{1}_{B_{(\cdot_2)}(\mathbf{0})}} & (\mathbb{R}^3_\star)^{\mathsf{E}} \\
\Big\downarrow (\mathscr{F}, \cdot_2) & & \mathscr{F}^{-1} \Big\uparrow \mathscr{F} \\
\mathbb{C}^{\hat{\mathsf{E}}} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \bullet \hat{\mathbf{1}}_{B_{(\cdot_2)}(\mathbf{0})}} & \mathbb{C}^{\hat{\mathsf{E}}}
\end{array} \quad (3.4.154)$$

As in (3.4.112), Fourier convolution can be approximated by low-pass filtering:

$$\begin{array}{ccc}
\mathscr{P}^0(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \oplus B_{(\cdot_2)}(\mathbf{0})} & \mathscr{P}^3(\mathsf{E}) \\
\Big\uparrow\Big\downarrow (U_0^0, \cdot_2) \quad (\mathsf{desc}^0, \cdot_2) & & U_0^3 \Big\uparrow\Big\downarrow \mathsf{desc}^3 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) * \mathbf{1}_{B_{(\cdot_2)}(\mathbf{0})}} & (\mathbb{R}^3_\star)^{\mathsf{E}} \\
\Big\downarrow (\mathscr{F}, \cdot_2) & & \mathscr{F}^{-1} \Big\uparrow\Big\downarrow \mathscr{F} \\
\mathbb{C}^{\hat{\mathsf{E}}} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \bullet \hat{\mathbf{1}}_{B_{(\cdot_2)}(\mathbf{0})}} & \mathbb{C}^{\hat{\mathsf{E}}} \\
\Big\uparrow (\lim_{r\to+\infty}(\cdot_1), \cdot_2) & & \lim_{r\to+\infty}(\cdot) \Big\uparrow \\
\mathbb{C}^{\hat{W}_r} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \bullet \hat{\mathbf{1}}_{B_{(\cdot_2)}(\mathbf{0})}} & \mathbb{C}^{\hat{W}_r}
\end{array} \tag{3.4.155}$$

It is possible to approximate the discrete impacts with their limit representations:

$$f^0_{S^0 \oplus B_r(\mathbf{0})} \equiv \lim_{\sigma\to 0^+} \sum_{0\le i<n_0} (g^3_\sigma((\cdot) - \mathbf{x}_i) * f^0_{B_r(\mathbf{0})}), \tag{3.4.156}$$

in which the right-hand side terms are the 'blurred' implicit descriptions of the balls, each having their crisp boundaries replaced with a Gaussian decay. The result is illustrated via the following diagram:

$$\begin{array}{ccc}
\mathscr{P}^0(\mathsf{E}) \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) \oplus B_{(\cdot_2)}(\mathbf{0})} & \mathscr{P}^3(\mathsf{E}) \\
\Big\uparrow\Big\downarrow (U_0^0, \cdot_2) \quad (\mathsf{desc}^0, \cdot_2) & & U_0^3 \Big\uparrow\Big\downarrow \mathsf{desc}^3 \\
(\mathbb{R}^0_\star)^{\mathsf{E}} \times \mathbb{R}_+ & \xrightarrow{(\cdot_1) * \mathbf{1}_{B_{(\cdot_2)}(\mathbf{0})}} & (\mathbb{R}^3_\star)^{\mathsf{E}} \\
\Big\uparrow (\lim_{\sigma\to 0^+(\cdot_1)}, \cdot_2) & & (\lim_{\sigma\to 0^+(\cdot_1)}, \cdot_2) \Big\uparrow \\
\mathscr{G}^0_\sigma & \xrightarrow{(\cdot_1) * \mathbf{1}_{B_{(\cdot_2)}(\mathbf{0})}} & \mathscr{G}^0_\sigma
\end{array} \tag{3.4.157}$$

Recall the PMC of the discrete pointset S^0 against the regularized closing with the balls $(S^0 \oplus^* B_\epsilon)$ into stably overlapped, critically overlapped, and exposed. In the implicit realm, the classification can be performed in a straightforward fashion using the regularization theorem (Lemma 2.3.20). First, the regularized closing is obtained by applying (2.3.159) to (3.4.127):

$$f_{S^0 \oplus^* B_\epsilon} := (\mathbf{1}_{S^0 \oplus B_\epsilon} * \delta^3) = (\mathbf{1}_{(S^0 \oplus B_\epsilon) \ominus B_\epsilon} * \delta^3) \tag{3.4.158}$$

$$= \Big(\big(\mathsf{sign} \circ ((\mathbb{1}^3_{S^0} * \mathbf{1}_{B_\epsilon}) \diamond \mathbf{1}_{B_\epsilon}) \big) * \delta^3 \Big), \tag{3.4.159}$$

noting that $(S^0 \odot B_\epsilon) = (S^0 \odot^\ddagger B_\epsilon)$ for $\odot \in \{\oplus, \ominus, \oplus, \ominus\}$ for the case of groupings, when one set is a finite collection of points and the other set is a 3D ball. Therefore, a point $\mathbf{x}_i \in S^0$ is classified into stably overlapped, critically overlapped, and exposed if the value of the above function $f_{S^0 \oplus^* B_\epsilon}(\mathbf{x}_i)$ is 0, in $(0, 1)$, and 1, respectively. The limit representation provides an approximation, though the precise behavior of the 'blurred' morphological closing is not fully understood. It is conceivable that the classification can be given with some confidence by replacing the $\delta-$function in (3.4.159) with some $g_\sigma-$function with $\sigma > 0$ and checking if the value of the function $f_{S^0 \oplus^* B_\epsilon}(\mathbf{x}_i)$ is in the intervals $[0, \epsilon]$, $(\epsilon, 1-\epsilon)$, and $[1-\epsilon, 1]$, respectively, for some $\epsilon = \ell(\sigma)$. The relationship between the rate of convergence of $\sigma, \epsilon \to 0^+$ is captured by $\ell : \mathbb{R}_+ \to \mathbb{R}_+$, i.e., $\epsilon := \ell(\sigma)$, and may depend on the choice of the $g_\sigma-$kernel.

Generalized Offsetting. The $\pm r-$offsetting operations defined in (3.4.27) and (3.4.28) for regular shapes can be redefined in terms of measurable Minkowski oper-

ation, thus can be expressed implicitly in terms of convolutions:

$$f_{\mathfrak{G}_r^{\uparrow\ddagger}(S)} \equiv (f_S * f_{B_r(\mathbf{0})}) \quad \text{i.e.,} \quad \mathfrak{G}_r^{\uparrow\ddagger}(S) = U_0(f_S * f_{B_r(\mathbf{0})}), \tag{3.4.160}$$

$$f_{\mathfrak{G}_r^{\downarrow\ddagger}(S)} \equiv (f_S \diamond f_{B_r(\mathbf{0})}) \quad \text{i.e.,} \quad \mathfrak{G}_r^{\downarrow\ddagger}(S) = U_0(f_S \diamond f_{B_r(\mathbf{0})}), \tag{3.4.161}$$

noting that $(\oplus^*, \ominus^*) \overset{\mathsf{ae}}{=} (\oplus^\ddagger, \ominus^\ddagger)$ when both shapes are full-dimensional—more precisely, letting $\mathfrak{G}_r^{\uparrow\ddagger}(S) := (S \oplus^\ddagger B_r)$ and $\mathfrak{G}_r^{\downarrow\ddagger}(S) := (S \ominus^\ddagger B_r)$ to look like the original definitions of $\mathfrak{G}_r^{\uparrow*}(S) = (S \oplus^* B_r)$ and $\mathfrak{G}_r^{\downarrow*}(S) = (S \ominus^* B_r)$, respectively, one obtains $\mathfrak{G}_r^{\uparrow\ddagger}(S) = \mathfrak{G}_r^{\uparrow*}(S)$ and $\mathfrak{G}_r^{\downarrow\ddagger}(S) \overset{\mathsf{ae}}{=} \mathfrak{G}_r^{\downarrow*}(S)$. Using the alternative formula in (3.4.108) for the Minkowski difference one obtains

$$f_{\mathfrak{G}_r^{\downarrow\ddagger}(S)} \equiv \min\left\{0, (\mathbf{1}_S * \mathbf{1}_{B_r(\mathbf{0})}) - \frac{4}{3}\pi r^3\right\} \quad \text{i.e.,} \quad \mathfrak{G}_r^{\uparrow\ddagger}(S) = U_{(\frac{4}{3}\pi r^3)^-}(\mathbf{1}_S * \mathbf{1}_{B_r(\mathbf{0})}), \tag{3.4.162}$$

where $\|\mathbf{1}_{B_r}\|_1 = \mu^3[B_r] = \frac{4}{3}\pi r^3$ is the $r-$ball volume, and $U_{t^-}(\cdot) = \lim_{t' \to t^-} U_{t'}(\cdot)$.

Note that the above definitions can be extended from offsetting of regular shapes (Fig. 3.4.8) to those of general shapes of heterogeneous dimensions (Fig. 3.4.7) if the d-dimensional dangling features are properly assigned with $(3-\mathrm{d})^{\mathrm{th}}-$order $\delta-$singular shape descriptors, as prescribed in Section 2.3.2.

Generalized Blending. The $\pm r-$blending operations defined in (3.4.30) and (3.4.31) for regular shapes can be expressed implicitly in terms of convolutions:

$$f_{\mathscr{R}_r^{\uparrow\ddagger}(S)} \equiv f_{(\mathfrak{G}_r^{\downarrow\ddagger} \circ \mathfrak{G}_r^{\uparrow\ddagger})(S)} \equiv \big((f_S * f_{B_r(\mathbf{0})}) \diamond f_{B_r(\mathbf{0})}\big), \tag{3.4.163}$$

$$f_{\mathscr{R}_r^{\downarrow\ddagger}(S)} \equiv f_{(\mathfrak{G}_r^{\uparrow\ddagger} \circ \mathfrak{G}_r^{\downarrow\ddagger})(S)} \equiv \big((f_S \diamond f_{B_r(\mathbf{0})}) * f_{B_r(\mathbf{0})}\big), \tag{3.4.164}$$

which can be expressed alternatively in terms of superlevel sets as:

$$\text{i.e.,}\quad \mathscr{R}_r^{\uparrow\ddagger}(S) = (\mathscr{O}_r^{\downarrow\ddagger} \circ \mathscr{O}_r^{\uparrow\ddagger})(S) = U_0\big((f_S * f_{B_r(\mathbf{0})}) \diamond f_{B_r(\mathbf{0})}\big), \tag{3.4.165}$$

$$\text{i.e.,}\quad \mathscr{R}_r^{\uparrow\ddagger}(S) = (\mathscr{O}_r^{\uparrow\ddagger} \circ \mathscr{O}_r^{\downarrow\ddagger})(S) = U_0\big((f_S \diamond f_{B_r(\mathbf{0})}) * f_{B_r(\mathbf{0})}\big), \tag{3.4.166}$$

noting that $(\oplus^*, \ominus^*) \overset{\mathsf{ae}}{=} (\oplus^\ddagger, \ominus^\ddagger)$ when both shapes are full-dimensional—more precisely, $\mathscr{R}_r^{\uparrow *}(S) := (S \oplus^* B_r) \overset{\mathsf{ae}}{=} (S \oplus^\ddagger B_r)$ and $\mathscr{R}_r^{\downarrow *}(S) := (S \ominus^* B_r) \overset{\mathsf{ae}}{=} (S \ominus^\ddagger B_r)$. Once again, the alternative formula in (3.4.108) for the Minkowski difference, leading to the alternative formula in (3.4.162) for the negative offsetting can be used to uniformly forumate everything in terms of a single convolution operator.

The diagram in (3.4.35) can be augmented as follows to illustrate convolution-based analytic (i.e., functional) offsetting and blending in the implicit realm:

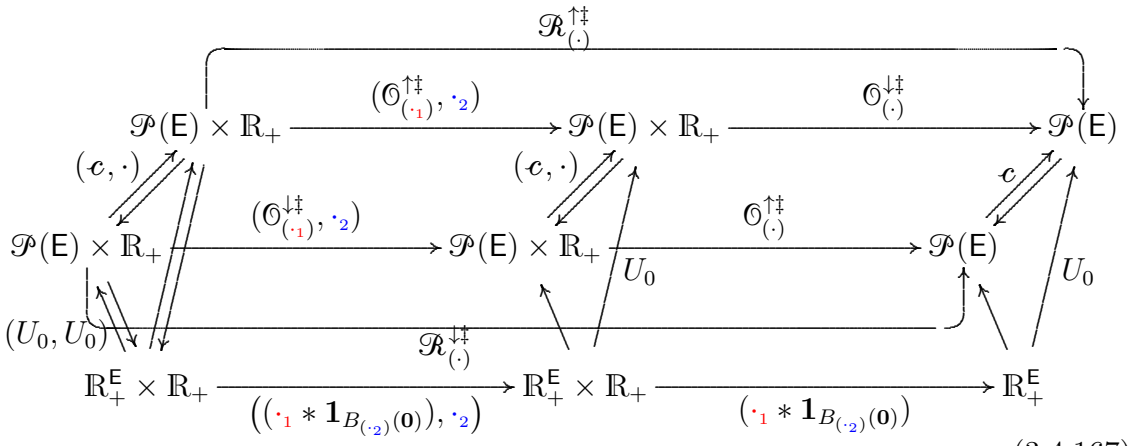

(3.4.167)

in which the vertical arrows are similar to those of (3.4.110). Once again, Fourier transforms can be use to convert the convolutions to pointwise multiplications, which,

in turn, can be approximated by low-pass filtering using a compact window $\hat{W}_r \subset \hat{\mathsf{E}}$:

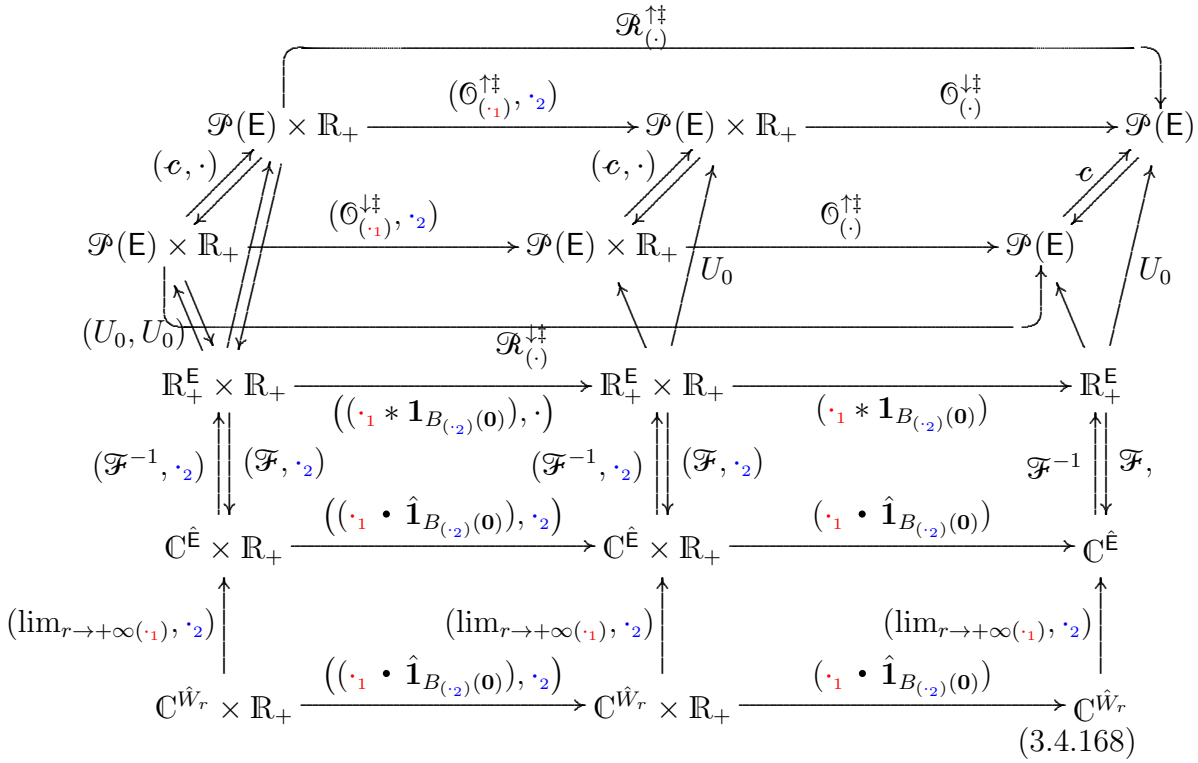

(3.4.168)

To summarize, it must be clear that a variety of morphological operations can be implicitly formulated in a uniform fashion, using convolutions in the physical space, or equivalently, using pointwise multiplications in the frequency domain, if the defining functions of the shapes are viewed as 3D signals. Accordingly, this opens up the possibility of leveraging methods from digital signal processing (DSP) such as low-pass filtering and anti-aliasing.

Next, the analytic methods for Euclidean morphology of shapes are extended to the group morphology of shapes and motions, by extending the convolution to the $\mathsf{C}-$space [215, 218].

3.4.2.2 Measure-Theoretic Group (i.e., Noncommutative) Morphology:

The generalization of the measure-theoretic concepts from Euclidean morphology based on the Lebesgue 3−measure to group morphology based on the Haar 6−measure [215, 218] is briefly presented here. The most notable distinction is that the group convolution operator $* : L^1(\mathsf{C}) \times L^1(\mathsf{C}) \to L^1(\mathsf{C})$ over the C−space, defined in (2.2.9) and (2.2.9) of Section 2.2.2.2 is noncommutative.

In particular, the explicit morphologies $\mathscr{P}(\mathsf{C})$ and $\mathscr{P}^*(\mathsf{C})$ correspond respectively to the precise closure or closure up to equality−ae under the said analytic operations. In the following, I revisit the morphological operators discussed in Section 3.4.1.2 to present their implicit counterparts.

Boolean Union and Intersection. In an identical fashion for shapes and motions. the set operations are implemented implicitly in the profile-oblivious and regularity-tolerant fashion via $\mathbb{R}_+$−arithmetic operations $+, \bullet$: $\mathbb{R}_+^{\mathsf{C}} \times \mathbb{R}_+^{\mathsf{C}} \to \mathbb{R}_+^{\mathsf{C}}$ on the real-valued nonnegative defining functions:

$$f_{M_1 \cup M_2} \equiv (f_{M_1} + f_{M_2}) \quad \text{i.e.,} \quad (M_1 \cup M_2) = U_0(f_{M_1} + f_{M_2}), \tag{3.4.169}$$

$$f_{M_1 \cap M_2} \equiv (f_{M_1} \bullet f_{M_2}) \quad \text{i.e.,} \quad (M_1 \cap M_2) = U_0(f_{M_1} \bullet f_{M_2}), \tag{3.4.170}$$

which are dual to each other via extended logical negation $\neg : \mathbb{R}^{\mathsf{C}}_{+} \to \mathbb{R}^{\mathsf{C}}_{+}$. Extending the diagrams in (3.4.37) using (2.3.16) one obtains:

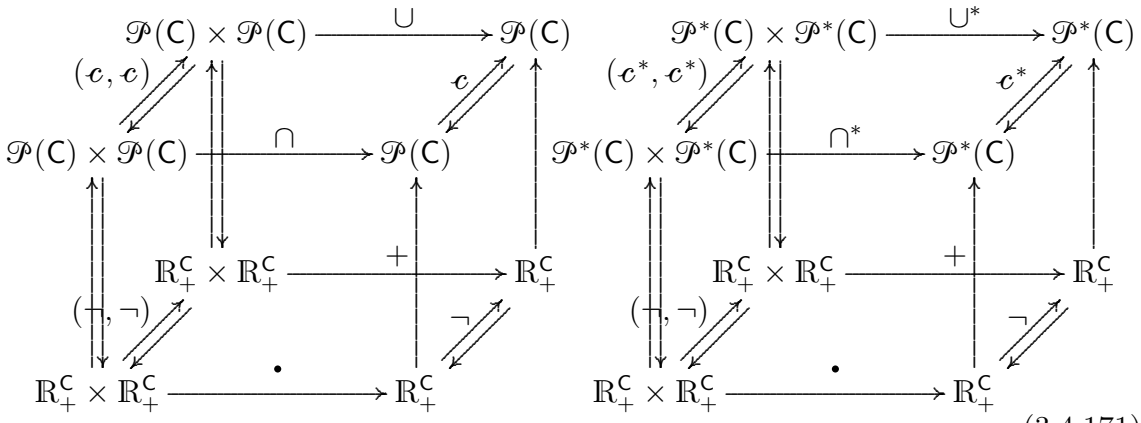

(3.4.171)

in which the vertical double-arrows should be $(U_0, U_0) \uparrow\downarrow \mathsf{desc}$ and $(U_0^*, U_0^*) \uparrow\downarrow \mathsf{desc}$ and vertical upward arrows should be $U_0 \uparrow$ and $U_0^* \uparrow$, for the left and right diagrams, respectively, eliminated due to cluttering. Once again, $U_0^*(\cdot) = \text{r}U_0(\cdot) = U_0(\cdot * \delta^3)$ is the regularized $0-$superlevel set generator.

Similar to the case with shapes, regularization of motions can be taken for granted (modulo $\overset{*}{\equiv}$) as discussed at length in Section 2.3.1:

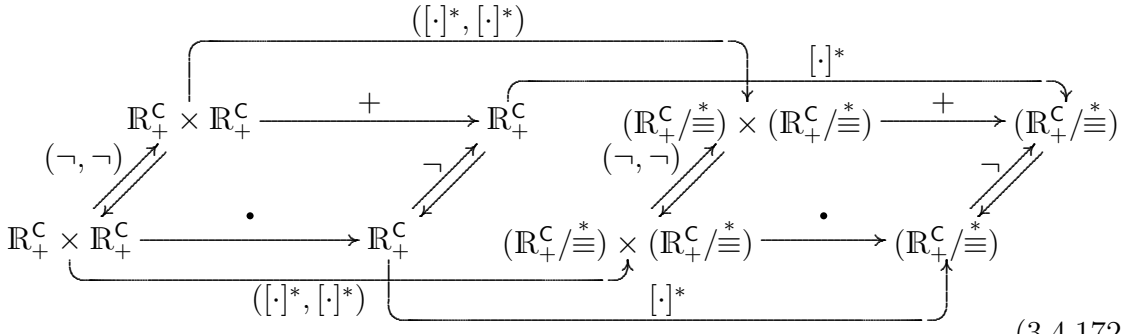

(3.4.172)

or it can achieved by extending the $\delta-$convolution filter (Section 2.3.3) from $\mathsf{E} = \mathbb{R}^3$ based on the Lebesgue 3−measure and 3−integral to $\mathsf{C} = \mathrm{SE}(3)$ based on the Haar

6−measure and 6−integral. 6F regularization is beyond the scope of this thesis.

Minkowski Product and Quotients. To begin with, the implicit formulation of the 'measurable' Minkowski sum and difference in (3.4.100) and (3.4.101) can be extended to the 'measurable' Minkowski product and left- and right-quotients $\otimes^{\ddagger}, \oslash^{\ddagger}, \oslash^{\ddagger} : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ in terms of a noncommutative convolution and its left- and right-dual operators, respectively:

$$f_{M_1 \otimes^{\ddagger} M_2} \equiv (f_{M_1} * f_{M_2}) \quad \text{i.e.,} \quad (M_1 \otimes^{\ddagger} M_2) = U_0(f_{M_1} * f_{M_2}), \tag{3.4.173}$$

$$f_{M_1 \oslash^{\ddagger} M_2} \equiv (f_{M_1} \triangleright f_{M_2}) \quad \text{i.e.,} \quad (M_1 \oslash^{\ddagger} M_2) = U_0(f_{M_1} \triangleright f_{M_2}), \tag{3.4.174}$$

$$f_{M_1 \oslash^{\ddagger} M_2} \equiv (f_{M_1} \triangleleft f_{M_2}) \quad \text{i.e.,} \quad (M_1 \oslash^{\ddagger} M_2) = U_0(f_{M_1} \triangleleft f_{M_2}), \tag{3.4.175}$$

where $*, \triangleright, \triangleleft : L^1(\mathsf{C}) \times L^1(\mathsf{C}) \to L^1(\mathsf{C})$ are the group convolution and its left- and right-dual operators defined as $\neg(f_1 * f_2) = ((\neg f_1) \triangleright f_2) = (f_1 \triangleleft (\neg f_2))$. The above formulae subsume (3.4.100) and (3.4.101) as special cases, when considering the isomorphism between shapes and translational motions via the embedding $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$. Similarly, the alternative computation of the Minkowski different over $\mathsf{E} = \mathbb{R}^3$ in (3.4.107) can be extended to the Minkowski left- and right-quotients over $\mathsf{C} = \mathrm{SE}(3)$ as:

$$f_{M_1 \oslash^{\ddagger} M_2} \equiv \lim_{\epsilon \to 0^+} \min \big\{ 0, (\mathbf{1}_{M_1} * \mathbf{1}_{M_2}) - \|\mathbf{1}_{M_2}\|_1 + \epsilon \big\}. \tag{3.4.176}$$

$$f_{M_1 \oslash^{\ddagger} M_2} \equiv \lim_{\epsilon \to 0^+} \min \big\{ 0, (\mathbf{1}_{M_1} * \mathbf{1}_{M_2}) - \|\mathbf{1}_{M_1}\|_1 + \epsilon \big\}, \tag{3.4.177}$$

based on which the following generalization of (3.4.108) can be obtained:

$$(M_1 \oslash^{\ddagger} M_2) = \lim_{\epsilon \to 0^+} U_{-\epsilon}\big((\mathbf{1}_{M_1} * \mathbf{1}_{M_2}) - \|\mathbf{1}_{M_2}\|_1\big) = \lim_{\epsilon \to 0^+} U_{\|\mathbf{1}_{M_2}\|_1 - \epsilon}(\mathbf{1}_{M_1} * \mathbf{1}_{M_2}),$$
$$(M_1 \ominus^{\ddagger} M_2) = \lim_{\epsilon \to 0^+} U_{-\epsilon}\big((\mathbf{1}_{M_1} * \mathbf{1}_{M_2}) - \|\mathbf{1}_{M_1}\|_1\big) = \lim_{\epsilon \to 0^+} U_{\|\mathbf{1}_{M_1}\|_1 - \epsilon}(\mathbf{1}_{M_1} * \mathbf{1}_{M_2}).$$

A proof can be found in [215, 218]. A similar diagram to that of (3.4.110) can be drawn to extend (3.4.42) to the implicit realm:

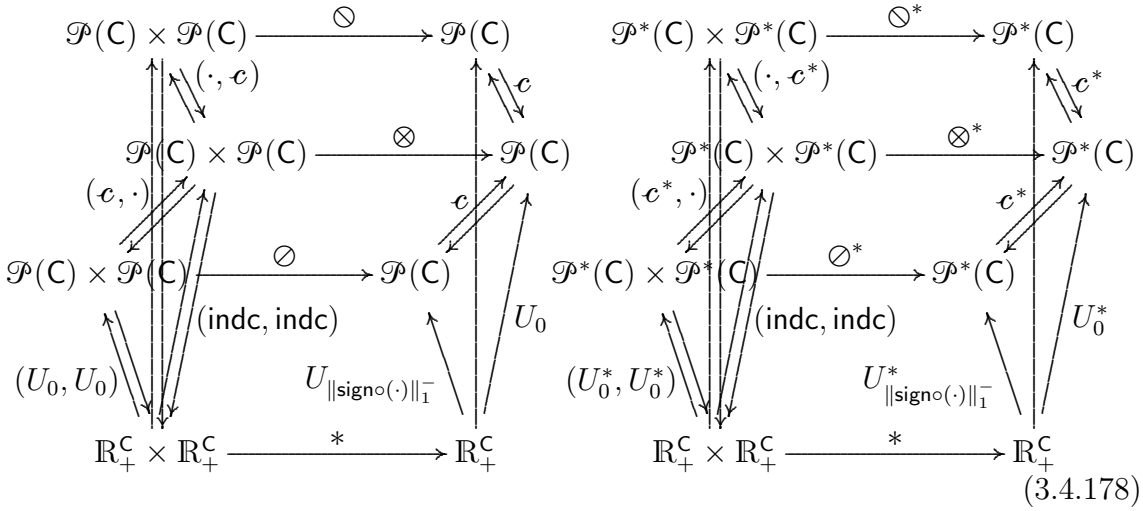

(3.4.178)

However, Fourier transforms are messier for noncommutative groups, leading to infinite-dimensional matrices for the frequency domain representation [65], which are beyond the scope of this thesis.

Embedding, Lifting, and Projection. Before moving on to implicit formulation of the morphological operations that entail interactions of shapes and motions, it is necessary to obtain the implicit formulation of the embedding, lifting, and projection gadgets that I introduced in Section 3.4.1.2. Starting with the bijective

embedding $\gamma_0 : \mathsf{E} \hookrightarrow \mathsf{T}$ extended as $\gamma_0 : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{T})$, it is easy to see that:

$$\mathbb{t} \in \gamma_0(S) \quad \text{iff} \quad \Big[\text{if} \quad \mathbf{x} = \mathsf{act}(\mathbb{t})(\mathbf{0}) \quad \text{then} \quad \mathbf{x} \in S\Big], \tag{3.4.179}$$

for all shapes $S \in \mathscr{P}(\mathsf{E})$. Note that $\mathbf{x} = \gamma_0^{-1}(\mathbb{t}) = \mathsf{act}(\mathbb{t})(\mathbf{0})$ denotes the action of the translational motion on the origin $\mathbf{0} \in \mathsf{E}$. The above expression can be rewritten as

$$\mathbf{x} \in \gamma_0^{-1}(T) \quad \text{iff} \quad \Big[\text{if} \quad \mathbb{t} \overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{x}) \quad \text{then} \quad \mathbb{t} \in T\Big], \tag{3.4.180}$$

for all purely translational motions $T \in \mathscr{P}(\mathsf{T})$. Note that $\mathbb{t} = \gamma_0(\mathbf{x}) \overset{0}{=} (\mathbb{r}_{\mathsf{id}}, \mathbf{x})$ is a translational motion represented with respect to the choice of origin $\mathbf{0} \in \mathsf{E}$. These expressions can be put into equivalence relations between the functional shape and motion descriptors $f_S := \mathsf{desc}(S) \in \mathbb{R}_+^{\mathsf{E}}$ and $f_{\gamma_0(S)} := (\mathsf{desc} \circ \gamma_0)(S) \in \mathbb{R}_+^{\mathsf{T}}$ as:

$$f_S \equiv (f_{\gamma_0(S)} \circ \gamma_0), \quad \text{i.e.,} \quad f_{\gamma_0(S)} \equiv (f_S \circ \gamma_0^{-1}), \tag{3.4.181}$$

where the left- and right-hand side $\equiv$ are over $\mathbb{R}_+^{\mathsf{E}}$ and $\mathbb{R}_+^{\mathsf{T}}$, respectively, which are isomorphic. The relationship can be illustrated via:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) & \underset{\gamma_0(\cdot)}{\overset{\gamma_0^{-1}(\cdot)}{\rightleftarrows}} & \mathscr{P}(\mathsf{T}) \\ U_0 \Big\uparrow \Big\| \mathsf{desc} & & U_0 \Big\uparrow \Big\| \mathsf{desc} \\ (\mathbb{R}_+^{\mathsf{E}}/\equiv) & \underset{(\cdot)\circ\gamma_0^{-1}}{\overset{(\cdot)\circ\gamma_0}{\rightleftarrows}} & (\mathbb{R}_+^{\mathsf{T}}/\equiv) \end{array} \tag{3.4.182}$$

Next, consider the lifting $\ell_0 : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{C})$ in (3.4.43). It is easy to verify that

the definition can be put into implicit form as:

$$\mathbf{x} \in S \quad \text{iff} \quad \Big[\forall \mathbb{r} \in \mathrm{SO}(3) : \ \text{if } \mathbb{c} \overset{0}{:=} (\mathbb{r}, \mathbf{x}) \quad \text{then} \quad \mathbb{c} \in \ell_0(S)\Big]. \tag{3.4.183}$$

Once again, the above relation can be expressed in terms of equivalence relations between $f_S := \mathsf{desc}(S) \in \mathbb{R}^{\mathsf{E}}_+$ and $f_{\ell_0(S)} := (\mathsf{desc} \circ \ell_0)(S) \in \mathbb{R}^{\mathsf{C}}_+$ as:

$$f_S(\mathbf{x}) > 0 \ \rightleftharpoons \ f_{\ell_0(S)}(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})} > 0, \quad \text{i.e.,} \quad f_S \equiv f_{\ell_0(S)}(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)}, \tag{3.4.184}$$

in which $\mathbb{r} \in \mathrm{SO}(3)$ is arbitrary—e.g., if $\mathbb{r} := \mathbb{r}_{\mathsf{id}}$ then $(\mathbb{r}, \cdot) = \gamma_0(\cdot)$ thus the left-hand side can be safely replaced with $f_S \equiv (f_{\ell_0(S)} \circ \gamma_0)$. On the other hand,

$$f_{\ell_0(S)}(\mathbb{c}) > 0 \ \rightleftharpoons \ f_S(\mathsf{act}(\mathbb{c})(\mathbf{0})) > 0, \quad \text{i.e.,} \quad f_{\ell_0(S)} \equiv (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}), \tag{3.4.185}$$

noting that $\mathsf{act}(\cdot)(\mathbf{0})$ projects $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})$ back to $\mathbf{x} \in \mathsf{E}$, i.e., $\mathsf{act}(\cdot)(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})} = \mathbf{x}$. The relationship can be illustrated via:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) & \xrightarrow{\qquad \ell_0(\cdot) \qquad} & \mathscr{P}(\mathsf{C}) \\ \Big\uparrow \Big\| \ \mathsf{desc} & & \Big\uparrow \Big\| \ \mathsf{desc} \\ (\mathbb{R}^{\mathsf{E}}_+/\!\equiv) & \underset{(\cdot_1 \circ \gamma_0)}{\overset{(\cdot_1 \circ \mathsf{act}(\cdot)(\mathbf{0}))}{\rightleftarrows}} & (\mathbb{R}^{\mathsf{C}}_+/\!\equiv) \end{array} \tag{3.4.186}$$

Next, consider the projection $\pi_0 : \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{E})$ in (3.4.44). It is easy to verify

that the definition can be put into implicit form as:

$$\mathbf{x} \in \pi_0(M) \quad \text{iff} \quad \Big[\exists \mathbb{r} \in \mathrm{SO}(3) : \ \text{if } \mathbb{c} \overset{0}{:=} (\mathbb{r}, \mathbf{x}) \quad \text{then} \quad \mathbb{c} \in M\Big]. \tag{3.4.187}$$

Once again, the above relation can be expressed in terms of equivalence relations between $f_M := \mathsf{desc}(M) \in \mathbb{R}^{\mathsf{C}}_+$ and $f_{\pi_0(M)} := (\mathsf{desc} \circ \pi_0)(M) \in \mathbb{R}^{\mathsf{E}}_+$ as:

$$f_M(\mathbb{c}) > 0 \ \Rightarrow \ (f_{\pi_0(M)} \circ \mathsf{act}(\mathbb{c}))(\mathbf{0}) > 0, \quad \text{i.e.,} \quad f_M \preccurlyeq f_{\pi_0(M)}\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)}, \tag{3.4.188}$$

noting that $\mathsf{act}(\mathbb{c})(\mathbf{0}) = \mathbf{x}$ and although $f_M(\mathbb{c}) > 0$ (i.e., $\mathbb{c} \in M$) implies $f_{\pi_0(M)}(\mathbf{x}) > 0$ (i.e., $\mathbf{x} \in \pi_0(M)$), the converse is not necessarily true—i.e., $M \subseteq (\ell_0 \circ \pi_0)(M)$ due to the loss of rotational information in projection, as illustrated in Fig. 3.4.11 repeated here in Fig. 3.4.15 with additional details. For a projected motion to contain a given point $\mathbf{x} \in \mathsf{E}$, it is necessary and sufficient to have a nonempty set of configurations of the form $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})$ with arbitrary $\mathbb{r} \in \mathrm{SO}(3)$. In other words, if $\ell_0(\mathbf{x}) := \ell_0(\{\mathbf{x}\})$, i.e., $\ell_0(\mathbf{x}) \overset{0}{=} \big\{(\mathbb{r}, \mathbf{x}) \mid \mathbb{r} \in \mathrm{SO}(3)\big\}$, then $(M \cap \ell_0(\mathbf{x})) \neq \emptyset$, i.e., $|M \cap \ell_0(\mathbf{x})| > 0$, i.e., $\mu^0[M \cap \ell_0(\mathbf{x})] > 0$. Alternatively, one can characterize the 'measurable' projection by $\mu^3[M \cap \ell_0(\mathbf{x})] > 0$ where the measuring occurs over the rotation space, i.e.,

$$f_{\pi_0^\ddagger(M)}(\mathbf{x}) > 0 \ \Rightarrow \ \int_{M \cap \ell_0(\mathbf{x})} f_M(\mathbb{c}) \ d\mu^3[\mathbb{c}] = \int_{\mathrm{SO}(3)} f_M(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})} \ d\mu^3[\mathbb{r}] > 0. \tag{3.4.189}$$

In other words, the integral function is equivalent via $\equiv$ with the motion descriptor:

$$f_{\pi_0^\ddagger(M)} \equiv \int_{M \cap \ell_0(\cdot)} f_M(\mathbb{c}) \ d\mu^3[\mathbb{c}] \equiv \int_{\mathrm{SO}(3)} f_M(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} \ d\mu^3[\mathbb{r}], \tag{3.4.190}$$

which can be converted to an integral over the entire $\mathsf{C} = \mathrm{SE}(3)$ by using a $3^{\mathrm{rd}}-$order

$\delta-$function over $\mathrm{T}(3) \cong \mathsf{E}$:

$$f_{\pi_0^\ddagger(M)} \equiv \int_{\mathsf{E}} \delta^3(\mathbf{x})\ d\mu^3[\mathbf{x}] \int_{\mathrm{SO}(3)} f_M(\mathbb{r}, \cdot)\ d\mu^3[\mathbb{r}] \tag{3.4.191}$$

$$\equiv \int_{\mathsf{E}} \int_{\mathrm{SO}(3)} \delta^3(\mathbf{x})\ f_M(\mathbb{r}, \cdot)\ d\mu^3[\mathbb{r}]\ d\mu^3[\mathbf{x}], \tag{3.4.192}$$

which is an inner product over the function space on $\mathsf{C} = \mathrm{SE}(3)$:

$$f_{\pi_0^\ddagger(M)} \equiv \int_{\mathrm{SE}(3)} (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0})\ f_M(\mathbb{c})\ d\mu^3[\mathbb{c}] = \big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), f_M \big\rangle, \tag{3.4.193}$$

noting that $\mathsf{act}(\cdot)(\mathbf{0})$ projects $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})$ back to $\mathbf{x} \in \mathsf{E}$, i.e., $\mathsf{act}(\cdot)(\mathbb{c})\big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x})} = \mathbf{x}$. The relationship can be illustrated via:

$$\begin{array}{ccc} \mathscr{P}(\mathsf{E}) & \xleftarrow{\quad \pi_0^\ddagger(\cdot) \quad} & \mathscr{P}(\mathsf{C}) \\ U_0 \big\uparrow \big\downarrow \mathsf{desc} & & U_0 \big\uparrow \big\downarrow \mathsf{desc} \\ (\mathbb{R}_+^{\mathsf{E}}/\equiv) & \underset{(\cdot_1 \circ \gamma_0)}{\overset{\langle (\delta^3 \circ \mathsf{act}(\cdot)), \cdot_1 \rangle}{\rightleftarrows}} & (\mathbb{R}_+^{\mathsf{C}}/\equiv) \end{array} \tag{3.4.194}$$

Figure 3.4.15 illustrates the implicit approach to lifting and projection—i.e., characterizing a query on the lifted motion or projected shape in terms of queries on the original shape and motion, respectively. It shows a simple planar shape, viewed along a direction on the plane (i.e., showing a 1D side-image of it) for ease of illustration. The relationship between a 2D planar shape in $\mathbb{R}^2$ and a 3D planar motion (i.e., 1D rotation and 2D translation) in $\mathrm{SE}(2) \cong \mathrm{SO}(2) \rtimes \mathrm{T}(2)$ with $\mathrm{SO}(2) \cong \mathbb{S}^1$ (i.e., a circle) and $\mathrm{T}(2) \cong \mathbb{R}^2$ (i.e., a plane) is illustrated by revolving the plane to create a toroidal

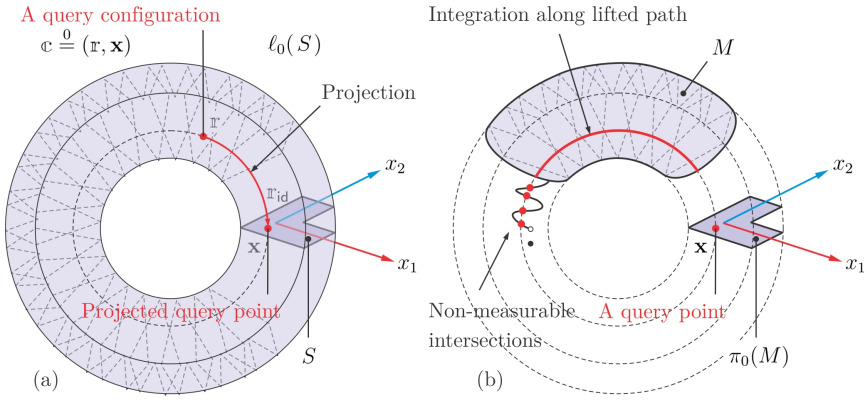

FIGURE 3.4.15: The implicit lifting and projection of a 2D shape in $\mathbb{R}^2$ to and from a 3D motion in $\mathrm{SE}(2) \cong \mathrm{SO}(2) \rtimes \mathrm{T}(2)$ amounts to projecting and integrating along a 1D circle (representing rotations) at each query point or configuration, respectively.

topology. Lifting of a 2D shape to a 3D motion amounts to revolving every point in the shape along the circle, thus revolving the shape into a torus. This means that every membership query on the torus is equal to a membership query on its cross-section, regardless of the rotation angle, as depicted by (3.4.184). On the other hand, projection of a 3D motion to a 2D shape loses the rotational content, which means that every membership query on the projected cross-section can be obtained by an intersection test between the motion and the circle emanating from the query point.

Morphological Dilation and Erosion. The implicit formulation of the measurable morphological dilation and erosion $\mathsf{dil}^{\ddagger}, \mathsf{ero}^{\ddagger} : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ can be formulated in terms of the noncommutative convolution and its right-dual operator. Starting by applying the projection formula in (3.4.193), to the explicit formulation

in (3.4.45) and (3.4.46), one obtains an inner product formulation:

$$f_{\mathsf{dil}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1}\otimes\ell_0(S))} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), f_{M^{-1}\otimes\ell_0(S)} \Big\rangle, \tag{3.4.195}$$

$$f_{\mathsf{ero}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1}\oslash\ell_0(S))} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), f_{M^{-1}\oslash\ell_0(S)} \Big\rangle. \tag{3.4.196}$$

which can be alternatively written in terms of an integral over SO(3) as in (3.4.190):

$$f_{\mathsf{dil}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1}\otimes\ell_0(S))} \equiv \int_{\mathrm{SO}(3)} f_{M^{-1}\otimes\ell_0(S)}(\mathbb{C})\big|_{\mathbb{C}=\underline{0}(\mathbb{r},\cdot)}\, d\mu^3[\mathbb{r}], \tag{3.4.197}$$

$$f_{\mathsf{ero}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1}\oslash\ell_0(S))} \equiv \int_{\mathrm{SO}(3)} f_{M^{-1}\oslash\ell_0(S)}(\mathbb{C})\big|_{\mathbb{C}=\underline{0}(\mathbb{r},\cdot)}\, d\mu^3[\mathbb{r}]. \tag{3.4.198}$$

Next, applying the implicit formulae for the Minkowski operations product and right-quotient operations in (3.4.173) and (3.4.175) yields:

$$f_{\mathsf{dil}^{\ddagger}(M,S)} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), \tilde{f}_M * (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \Big\rangle, \tag{3.4.199}$$

$$f_{\mathsf{ero}^{\ddagger}(M,S)} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), \tilde{f}_M \triangleleft (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \Big\rangle, \tag{3.4.200}$$

which can be expressed alternatively in terms of superlevel sets as:

$$\mathsf{dil}^{\ddagger}(M,S) = U_0\bigg(\int_{\mathrm{SO}(3)} \tilde{f}_M * (f_S \circ \mathsf{act}(\mathbb{C}))(\mathbf{0})\big|_{\mathbb{C}=\underline{0}(\mathbb{r},\cdot)}\, d\mu^3[\mathbb{r}] \bigg), \tag{3.4.201}$$

$$\mathsf{ero}^{\ddagger}(M,S) = U_0\bigg(\int_{\mathrm{SO}(3)} \tilde{f}_M \triangleleft (f_S \circ \mathsf{act}(\mathbb{C}))(\mathbf{0})\big|_{\mathbb{C}=\underline{0}(\mathbb{r},\cdot)}\, d\mu^3[\mathbb{r}] \bigg). \tag{3.4.202}$$

The quotient operator can be alternatively implemented in terms of the convolution operator and $\epsilon-$superlevel sets in the limit, using (3.4.177).

Morphological Closing and Opening. The implicit formulation of the measurable morphological closing and opening $\mathsf{cls}^{\ddagger}, \mathsf{opn}^{\ddagger} : \mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ can be formulated in terms of the noncommutative convolution and its right-dual operator. Once again, by applying the projection formula in (3.4.193), to the explicit formulation in (3.4.50) and (3.4.52), one obtains an inner product formulation:

$$f_{\mathsf{cls}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1} \oslash (M \otimes \ell_0(S)))} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), f_{M^{-1} \oslash (M \otimes \ell_0(S))} \Big\rangle, \tag{3.4.203}$$

$$f_{\mathsf{opn}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1} \otimes (M \oslash \ell_0(S)))} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), f_{M^{-1} \otimes (M \oslash \ell_0(S))} \Big\rangle. \tag{3.4.204}$$

which can be alternatively written in terms of an integral over SO(3) as in (3.4.190):

$$f_{\mathsf{cls}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1} \oslash (M \otimes \ell_0(S)))} \equiv \int_{\mathrm{SO}(3)} f_{M^{-1} \oslash (M \otimes \ell_0(S))}(\mathbb{c}) \big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} \, d\mu^3[\mathbb{r}], \tag{3.4.205}$$

$$f_{\mathsf{opn}^{\ddagger}(M,S)} \equiv f_{\pi_0(M^{-1} \otimes (M \oslash \ell_0(S)))} \equiv \int_{\mathrm{SO}(3)} f_{M^{-1} \otimes (M \oslash \ell_0(S))}(\mathbb{c}) \big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} \, d\mu^3[\mathbb{r}]. \tag{3.4.206}$$

Next, applying the implicit formulae for the Minkowski operations product and right-quotient operations in (3.4.173) and (3.4.175) yields:

$$f_{\mathsf{cls}^{\ddagger}(M,S)} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), \tilde{f}_M \triangleleft \big(f_M * (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \big) \Big\rangle, \tag{3.4.207}$$

$$f_{\mathsf{opn}^{\ddagger}(M,S)} \equiv \Big\langle (\delta^3 \circ \mathsf{act}(\cdot))(\mathbf{0}), \tilde{f}_M * \big(f_M \triangleleft (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \big) \Big\rangle, \tag{3.4.208}$$

which can be expressed alternatively in terms of superlevel sets as:

$$\mathsf{dil}^{\ddagger}(M,S) = U_0 \bigg(\int_{\mathrm{SO}(3)} \tilde{f}_M \triangleleft \big(f_M * (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \big) \big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} \, d\mu^3[\mathbb{r}] \bigg), \tag{3.4.209}$$

$$\mathsf{ero}^{\ddagger}(M,S) = U_0 \bigg(\int_{\mathrm{SO}(3)} \tilde{f}_M * \big(f_M \triangleleft (f_S \circ \mathsf{act}(\cdot))(\mathbf{0}) \big) \big|_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} \, d\mu^3[\mathbb{r}] \bigg). \tag{3.4.210}$$

The quotient operator can be alternatively implemented in terms of the convolution operator and $\epsilon-$superlevel sets in the limit, using (3.4.177).

Minkowski Closing and Opening. The implicit formulation of the measurable Minkowski left- and right-closing $\circledcirc^{\ddagger}, \circledcirc^{\ddagger}$: $\mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ and opening $\oslash^{\ddagger}, \oslash^{\ddagger}$: $\mathscr{P}(\mathsf{C}) \times \mathscr{P}(\mathsf{C}) \to \mathscr{P}(\mathsf{C})$ can be implicitly formulated in terms of the noncommutative convolution and its left- and right-dual operators by applying (3.4.173) through (3.4.175) to their explicit formulation in (3.4.54) through (3.4.57):

$$f_{M_1 \circledcirc^{\ddagger} M_2} \equiv f_{(M_1^{-1} \oslash (M_1 \otimes M_2))} \equiv \big(\tilde{f}_{M_1} \triangleleft (f_{M_1} * f_{M_2})\big), \tag{3.4.211}$$

$$f_{M_1 \circledcirc^{\ddagger} M_2} \equiv f_{((M_1 \otimes M_2) \oslash M_2^{-1})} \equiv \big((f_{M_1} * \tilde{f}_{M_2}) \triangleright f_{M_2}\big), \tag{3.4.212}$$

$$f_{M_1 \oslash^{\ddagger} M_2} \equiv f_{((M_1 \oslash M_2) \otimes M_2^{-1})} \equiv \big((f_{M_1} \triangleright \tilde{f}_{M_2}) * f_{M_2}\big), \tag{3.4.213}$$

$$f_{M_1 \circledcirc^{\ddagger} M_2} \equiv f_{(M_1^{-1} \otimes (M_1 \oslash M_2))} \equiv \big(\tilde{f}_{M_1} * (\tilde{f}_{M_1} \triangleleft f_{M_2})\big), \tag{3.4.214}$$

which can be expressed alternatively in terms of superlevel sets as:

$$(M_1 \circledcirc^{\ddagger} M_2) = U_0\big(\tilde{f}_{M_1} \triangleleft (\tilde{f}_{M_1} * f_{M_2})\big), \tag{3.4.215}$$

$$(M_1 \circledcirc^{\ddagger} M_2) = U_0\big((f_{M_1} * \tilde{f}_{M_2}) \triangleright f_{M_2}\big), \tag{3.4.216}$$

$$(M_1 \oslash^{\ddagger} M_2) = U_0\big((f_{M_1} \triangleright \tilde{f}_{M_2}) * f_{M_2}\big), \tag{3.4.217}$$

$$(M_1 \circledcirc^{\ddagger} M_2) = U_0\big(\tilde{f}_{M_1} * (\tilde{f}_{M_1} \triangleleft f_{M_2})\big), \tag{3.4.218}$$

which can be viewed as extensions of the Minkowski closing and opening over $\mathsf{E} = \mathbb{R}^3$ in (3.4.127) and (3.4.128) to their left- and right-counterparts over $\mathsf{C} = \mathrm{SE}(3)$.

Once again, the left- and right-quotient operators can be alternatively imple-

mented in terms of the convolution operator and $\epsilon-$superlevel sets in the limit, using (3.4.176) and (3.4.177), respectively.

Homogeneous Product and Quotients. As expected, the implicit formulation of the measurable homogeneous product and left- and right-quotients $\boxtimes^{\ddagger}, \boxslash^{\ddagger}, \boxbslash^{\ddagger} : \mathscr{P}(\mathsf{E}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{C})$ can be implicitly formulated in terms of the noncommutative convolution and its left- and right-dual operators by applying (3.4.173) through (3.4.175) and the implicit lifting formula in (3.4.185) to their explicit formulation in (3.4.66) through (3.4.68):

$$f_{S_1 \boxtimes^{\ddagger} S_2} \equiv (f_{\ell_0(S_1)} * \tilde{f}_{\ell_0(S_2)}) \equiv \big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) * (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.219}$$

$$f_{S_1 \boxslash^{\ddagger} S_2} \equiv (f_{\ell_0(S_1)} \triangleright \tilde{f}_{\ell_0(S_2)}) \equiv \big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleright (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.220}$$

$$f_{S_1 \boxbslash^{\ddagger} S_2} \equiv (f_{\ell_0(S_1)} \triangleleft \tilde{f}_{\ell_0(S_2)}) \equiv \big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleleft (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.221}$$

which can be expressed alternatively in terms of superlevel sets as:

$$(S_1 \boxtimes^{\ddagger} S_2) = U_0\big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) * (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.222}$$

$$(S_1 \boxslash^{\ddagger} S_2) = U_0\big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleright (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.223}$$

$$(S_1 \boxbslash^{\ddagger} S_2) = U_0\big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleleft (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big). \tag{3.4.224}$$

Once again, the left- and right-quotient operators can be alternatively implemented in terms of the convolution operator and $\epsilon-$superlevel sets in the limit, using (3.4.176) and (3.4.177), respectively.

$\mathsf{C}-$Obstacles and Free Space. The configuration space obstacles were defined in Section 2.3.1.3 as the set of configurations in $\mathsf{C} = \mathrm{SE}(3)$ at which the two objects $S_1, S_2 \in \mathscr{P}(\mathsf{E})$ do not collide, i.e., have a zero intersection 3−measure. Thus it is precisely given by the measurable homogeneous product operator characterized implicitly in (3.4.219) $O_{S_1,S_2} = (S_1 \boxtimes^{\ddagger} S_2)$, i.e.,

$$f_{O_{S_1,S_2}} \equiv f_{S_1 \boxtimes^{\ddagger} S_2} \equiv \big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) * (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.225}$$

which can be expressed alternatively in terms of superlevel sets as:

$$O_{S_1,S_2} = (S_1 \boxtimes^{\ddagger} S_2) = U_0\big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) * (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big). \tag{3.4.226}$$

Accordingly, its complement (i.e., free space) is characterized as:

$$f_{\mathscr{c}(O_{S_1,S_2})} \equiv \neg f_{S_1 \boxtimes^{\ddagger} S_2} \equiv \neg\big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) * (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.227}$$

which could be alternatively derived using the duality relationships $\mathscr{c}(O_{S_1,S_2}) = \mathscr{c}(S_1 \boxtimes^{\ddagger} S_2) = (\mathscr{c}(S_1) \boxslash^{\ddagger} S_2) = (S_1 \boxbslash^{\ddagger} \mathscr{c}(S_2))$, i.e.,

$$f_{\mathscr{c}(O_{S_1,S_2})} \equiv \big((\neg f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleright (\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big), \tag{3.4.228}$$

$$f_{\mathscr{c}(O_{S_1,S_2})} \equiv \big((f_{S_1} \circ \mathsf{act}(\cdot)(\mathbf{0})) \triangleleft (\neg\tilde{f}_{S_2} \circ \mathsf{act}(\cdot)(\mathbf{0}))\big). \tag{3.4.229}$$

It is helpful to think of the 6D $\mathsf{C}-$obstacle as a collection of slices, each slice being a 3D translational $\mathsf{C}-$obstacle corresponding to the set of translations $\mathbf{t} \in \mathbb{R}^3$ for a fixed rotation $\mathbb{r} \in \mathrm{SO}(3)$ such that $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$, which is enabled by the semiproduct

structure of the configuration space $\mathsf{C} = \mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathrm{T}(3)$:

$$O_{S_1,S_2}|_{\mathbb{r}} = \big\{\mathbf{t} \in \mathbb{R}^3 \mid \mathsf{col}(S_1, (\mathbb{r}, \mathbf{t})S_2) = \mathtt{1}\big\}, \tag{3.4.230}$$

The implicit description of the translational $\mathsf{C}-$obstacle $f_{O_{S_1,S_2}|_{\mathbb{r}}} \equiv f_{O_{S_1,S_2}}(\mathbb{r}, \cdot)$ can be obtained as a noncommutative configuration:

$$f_{O_{S_1,S_2}|_{\mathbb{r}}} \equiv \Big[(f_{S_1} \circ \mathsf{act}(\mathbb{c})(\mathbf{0})) \ * \ (\tilde{f}_{S_2} \circ \mathsf{act}(\mathbb{c})(\mathbf{0}))\Big]_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)}, \tag{3.4.231}$$

It is easy to show that the above noncommutative convolution over the group of 6D rigid configurations can be converted to a commutative convolution over the subgroup of 3D translations for different fixed rotations:

$$f_{O_{S_1,S_2}|_{\mathbb{r}}} \equiv [f_{S_1} \circ \mathsf{act}(\mathbb{c})(\mathbf{0})]_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)} * [\tilde{f}_{S_2} \circ \mathsf{act}(\mathbb{c})(\mathbf{0})]_{\mathbb{c} \overset{0}{=} (\mathbb{r}, \cdot)}, \tag{3.4.232}$$

But $\mathsf{act}(\mathbb{r}, \mathbf{x})(\mathbf{0}) = \mathbf{x}$ implies $(f_S \circ \mathsf{act}(\mathbb{r}, \mathbf{x})(\mathbf{0})) = f_S(\mathbf{x})$, for all $S \in \mathscr{P}(\mathsf{E})$ and $\mathbf{x} \in \mathsf{E}$, i.e., $(f_S \circ \mathsf{act}(\mathbb{r}, \cdot)(\mathbf{0})) \equiv f_S$, therefore:

$$f_{O_{S_1,S_2}|_{\mathbb{r}}} \equiv (f_{S_1} * \tilde{f}_{\mathbb{r}S_2}), \quad \text{i.e.,} \quad O_{S_1,S_2}|_{\mathbb{r}} = U_0(f_{S_1} * \tilde{f}_{\mathbb{r}S_2}), \tag{3.4.233}$$

from which the complete 6D $\mathsf{C}-$obstacle can be implicitly obtained as

$$f_{O_{S_1,S_2}}(\mathbb{r}, \cdot) \equiv (f_{S_1} * \tilde{f}_{\mathbb{r}S_2})(\cdot) \equiv \big(f_{S_1} * (\tilde{f}_{S_2} \circ \mathsf{act}(\mathbb{r}^{-1}))\big)(\cdot), \tag{3.4.234}$$

for all $\mathbb{r} \in \mathrm{SO}(3)$. The explicit description of the 6D $\mathsf{C}-$obstacle can thus be retrieved

by stacking up these 3D slices together for all $\mathbb{r} \in \mathrm{SO}(3)$ which is:

$$O_{S_1,S_2} = \left\{ \mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{x}) \mid \mathbb{r} \in \mathrm{SO}(3) \text{ and } (f_{S_1} * \tilde{f}_{\mathbb{r}S_2})(\mathbf{x}) > 0 \right\}. \tag{3.4.235}$$

The collision detection is concerned with intersections of nonzero 3−merasure, thus it yields the same result for both regular and irregular shapes:

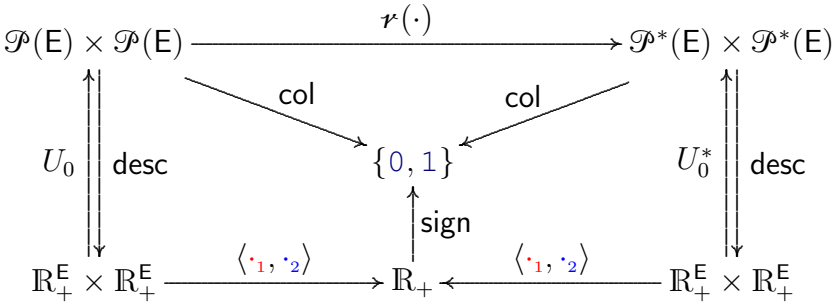

(3.4.236)

in which the inner product can be taken to the Fourier domain, thanks to Parseval's theorem (Theorem 2.2.10), which, in turn, can be approximated by low-pass filtering:

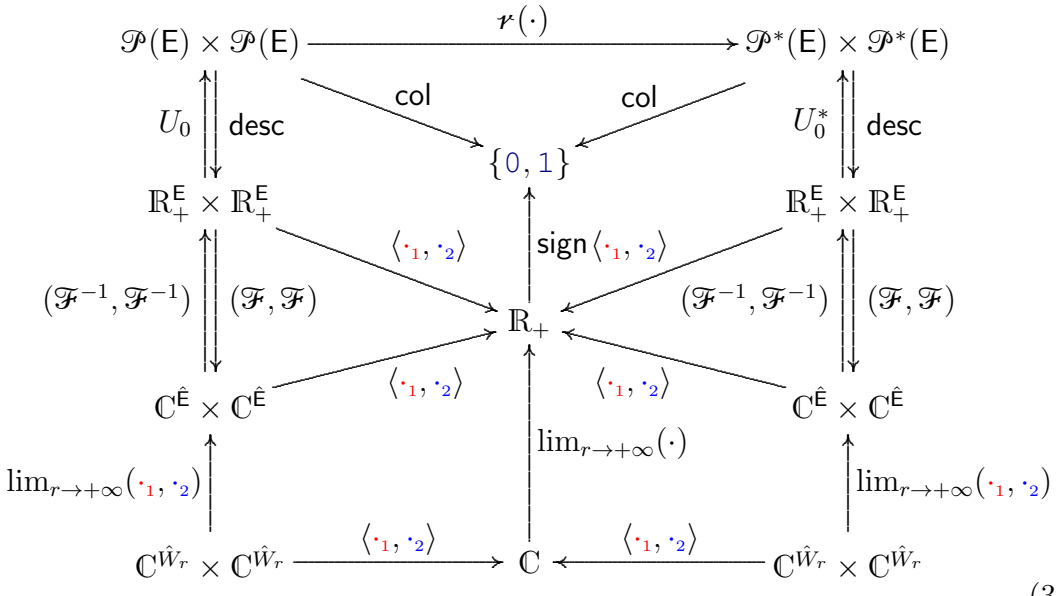

(3.4.237)

Accordingly, the translational $\mathsf{C}-$obstacle for a fixed rotation $\mathbb{r} \in \mathrm{SO}(3)$ formulated implicitly by a commutative convolution over $\mathsf{T} \cong \mathrm{SE}(3)/\mathrm{SO}(3)$ is illustrated via:

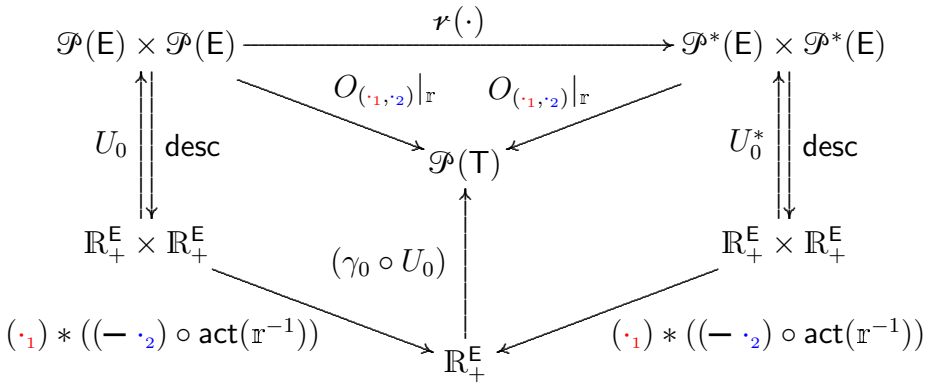

(3.4.238)

Once again, the convolution can be converted to a pointwise multiplication in the Fourier domain, thanks to convolution theorem (Theorem 2.2.11), which, in turn, can be approximated by low-pass filtering:

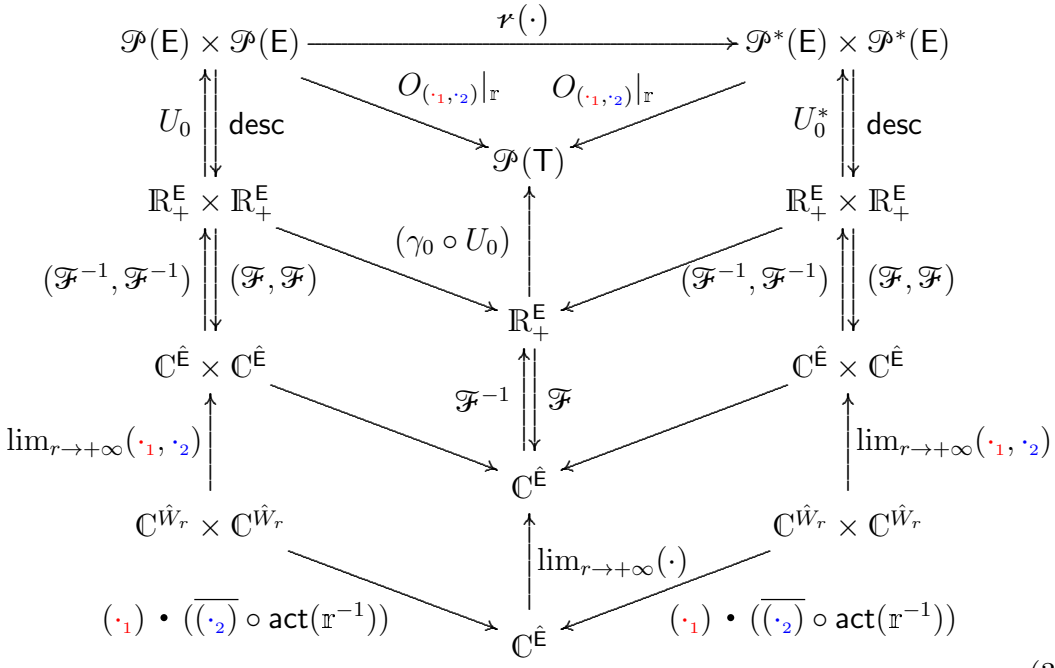

(3.4.239)

in which the labels $(\cdot_1) * ((- \cdot_2) \circ \mathsf{act}(\mathbb{r}^{-1}))$ and $(\cdot_1) \bullet (\overline{(\cdot_2)} \circ \mathsf{act}(\mathbb{r}^{-1}))$ are omitted from the intermediate horizontal arrows to avoid cluttering.

As discussed in depth in Section 2.3.2.4, if the implicit shape descriptors are properly set up to assigned $(3-\mathrm{d})^{\mathrm{th}}-$order $\delta-$singularities to $\mathrm{d}-$dimensional irregularities, the collision detection and $\mathsf{C}-$obstacle notions are generalized to detecting lower-dimensional intersections and characterizing configurations at which they happen, respectively. For example:

- If both shapes are characterized as $3-$manifolds $S_1^3, S_2^3 \in \mathscr{P}^3(\mathsf{E})$, the inner product $\langle \mathbb{1}^0_{S_1^3}, \mathbb{1}^0_{(\mathbb{r},\mathbf{0})S_2^3} \rangle = \langle \mathbb{1}^0_{S_1^3}, (\mathbb{1}^0_{S_2^3} \circ \mathsf{act}(\mathbb{r}^{-1})) \rangle$ measures the intersection solid volume (i.e., $3-$measure: $3 = (3+3)-3$) between them at a relative rotation $\mathbb{r} \in \mathrm{SO}(3)$ for no translation:

$$\mathbb{1}^0_{O^3_{S_1^3,S_2^3}}(\mathbb{r},\mathbf{0}) = \int_{\mathsf{E}} \mathbb{1}^0_{S_1^3}(\mathbf{x}')\ \mathbb{1}^0_{S_2^3}(\mathbf{x}'-\mathbf{0})\ d\mu^3[\mathbf{x}'] = \mu^3[S_1^3 \cap (\mathbb{r},\mathbf{0})S_2^3], \quad (3.4.240)$$

  noting that for transverse intersections, the two open solids are either disjoint or intersect over another open solid (i.e., $3-$manifold: $3 = (3+3)-3$) in general conditions. Thus the convolution $(\mathbb{1}^0_{S_1^3} * \tilde{\mathbb{1}}^0_{\mathbb{r}S_2^3}) = \big(\mathbb{1}^0_{S_1^3} * (\tilde{\mathbb{1}}^0_{S_2^3} \circ \mathsf{act}(\mathbb{r}^{-1}))\big)$ characterizes the nonsingular motion descriptor for the $\mathsf{C}-$obstacle $O^3_{S_1^3,S_2^3}$ at different configurations $\mathbb{c} \in \mathsf{C}$ where $\mathbb{c} \overset{0}{=} (\mathbb{r},\mathbf{x})$:

$$\mathbb{1}^0_{O^3_{S_1^3,S_2^3}}(\mathbb{r},\mathbf{x}) = \int_{\mathsf{E}} \mathbb{1}^0_{S_1^3}(\mathbf{x}')\ \mathbb{1}^0_{S_2^3}(\mathbf{x}'-\mathbf{x})\ d\mu^3[\mathbf{x}'] = \mu^3[S_1^3 \cap (\mathbb{r},\mathbf{x})S_2^3], \quad (3.4.241)$$

  remembering also that $\mathbb{1}^0_{S_1^3} = \mathbf{1}_{S_1^3}$, $\mathbb{1}^0_{S_2^3} = \mathbf{1}_{S_2^3}$ and $\mathbb{1}^0_{O^3_{S_1^3,S_2^3}} = \mathbf{1}_{O^3_{S_1^3,S_2^3}}$.

  The collision detection can be generalized to more interesting interference tests such as contact between surfaces and curves exemplified below.

- If both shapes are characterized as 2−manifolds $S_1^2, S_2^2 \in \mathscr{P}^2(\mathsf{E})$, the inner product $\langle \mathbb{1}^1_{S_1^2}, \mathbb{1}^1_{(\mathbb{r},\mathbf{0})S_2^2} \rangle = \big\langle \mathbb{1}^1_{S_1^2}, (\mathbb{1}^1_{S_2^2} \circ \mathsf{act}(\mathbb{r}^{-1})) \big\rangle$ measures the intersection curve length (i.e., 1−measure: $1 = (2+2)-3$) between them at a relative rotation $\mathbb{r} \in \mathrm{SO}(3)$ for no translation:

$$\mathbb{1}^0_{O^1_{S_1^2,S_2^2}}(\mathbb{r},\mathbf{0}) = \int_{\mathsf{E}} \mathbb{1}^1_{S_1^2}(\mathbf{x}')\ \mathbb{1}^1_{S_2^2}(\mathbf{x}'-\mathbf{0})\ d\mu^3[\mathbf{x}'] = \mu^1[S_1^2 \cap (\mathbb{r},\mathbf{0})S_2^2], \quad (3.4.242)$$

noting that for transverse intersections, the two surfaces intersect over a non-degenerate curve (i.e., 1−manifold: $1 = (2+2)-3$) in general conditions. Thus the convolution $(\mathbb{1}^1_{S_1^2} * \tilde{\mathbb{1}}^1_{\mathbb{r}S_2^2}) = \big(\mathbb{1}^1_{S_1^2} * (\tilde{\mathbb{1}}^1_{S_2^2} \circ \mathsf{act}(\mathbb{r}^{-1}))\big)$ characterizes the nonsingular motion descriptor for the generalized C−obstacle $O^1_{S_1^2,S_2^2}$ at different configurations $\mathbb{c} \in \mathsf{C}$ where $\mathbb{c} \overset{0}{=} (\mathbb{r},\mathbf{x})$:

$$\mathbb{1}^0_{O^1_{S_1^2,S_2^2}}(\mathbb{r},\mathbf{x}) = \int_{\mathsf{E}} \mathbb{1}^1_{S_1^2}(\mathbf{x}')\ \mathbb{1}^1_{S_2^2}(\mathbf{x}'-\mathbf{x})\ d\mu^3[\mathbf{x}'] = \mu^1[S_1^2 \cap (\mathbb{r},\mathbf{x})S_2^2]. \quad (3.4.243)$$

The above equation does not hold at configurations where the intersection is non-transverse. For degenerate conditions such as two surfaces contacting over a nonzero area surface (i.e., 2−manifold: $2 > (2+2)-3$), the above inner product approaches infinity. I conjecture that the above convolution simultaneously characterizes the 1^{st}−order δ−singular motion descriptor for generalized C−obstacle $O^2_{S_1^2,S_2^2}$, i.e.,

$$(\mathbb{1}^1_{O^2_{S_1^2,S_2^2}} + \mathbb{1}^0_{O^1_{S_1^2,S_2^2}}) \overset{?}{\equiv} (\mathbb{1}^1_{S_1^2} * \tilde{\mathbb{1}}^1_{\mathbb{r}S_2^2}) \equiv \big(\mathbb{1}^1_{S_1^2} * (\tilde{\mathbb{1}}^1_{S_2^2} \circ \mathsf{act}(\mathbb{r}^{-1}))\big). \quad (3.4.244)$$

As a result, the L^1−norm of the convolution measures the curve length over the C−space along which such degeneracies happen.

- If both shapes are characterized as $1-$manifolds $S_1^1, S_2^1 \in \mathscr{P}^1(\mathsf{E})$, the inner product $\langle \mathbb{1}^2_{S_1^1}, \mathbb{1}^2_{(\mathbb{r},\mathbf{0})S_2^1} \rangle = \left\langle \mathbb{1}^2_{S_1^1}, (\mathbb{1}^2_{S_2^1} \circ \mathsf{act}(\mathbb{r}^{-1})) \right\rangle$ counts the number of intersection points (i.e., $0-$measure: $0 \leq (1+1)-3$) between them at a relative rotation $\mathbb{r} \in \mathrm{SO}(3)$ for no translation:

$$\mathbb{1}^0_{O^0_{S_1^1,S_2^1}}(\mathbb{r},\mathbf{0}) = \int_{\mathsf{E}} \mathbb{1}^2_{S_1^1}(\mathbf{x}')\ \mathbb{1}^2_{S_2^1}(\mathbf{x}'-\mathbf{0})\ d\mu^3[\mathbf{x}'] = \mu^0[S_1^1 \cap (\mathbb{r},\mathbf{0})S_2^1], \quad (3.4.245)$$

noting that for transverse intersections, the two curves intersect over a finite set of points (i.e., $0-$manifold: $0 \leq (1+1)-3$) in general conditions. Thus the convolution $(\mathbb{1}^2_{S_1^1} * \tilde{\mathbb{1}}^2_{\mathbb{r}S_2^1}) = \left(\mathbb{1}^2_{S_1^1} * (\tilde{\mathbb{1}}^2_{S_2^1} \circ \mathsf{act}(\mathbb{r}^{-1}))\right)$ characterizes the non-singular motion descriptor for the generalized $\mathsf{C}-$obstacle $O^0_{S_1^1,S_2^1}$ at different configurations $\mathbb{c} \in \mathsf{C}$ where $\mathbb{c} \overset{0}{=} (\mathbb{r},\mathbf{x})$:

$$\mathbb{1}^0_{O^0_{S_1^1,S_2^1}}(\mathbb{r},\mathbf{x}) = \int_{\mathsf{E}} \mathbb{1}^2_{S_1^1}(\mathbf{x}')\ \mathbb{1}^2_{S_2^1}(\mathbf{x}'-\mathbf{x})\ d\mu^3[\mathbf{x}'] = \mu^0[S_1^1 \cap (\mathbb{r},\mathbf{x})S_2^1]. \quad (3.4.246)$$

The above equation does not hold at configurations where the intersection is non-transverse. For degenerate conditions such as two curves contacting over a nonzero curve length (i.e., $1-$manifold: $1 > (1+1)-3$), the above inner product approaches infinity. I conjecture that the above convolution simultaneously characterizes the $1^{\text{st}}-$order $\delta-$singular motion descriptor for generalized $\mathsf{C}-$obstacle $O^1_{S_1^1,S_2^1}$, i.e.,

$$(\mathbb{1}^1_{O^1_{S_1^1,S_2^1}} + \mathbb{1}^0_{O^0_{S_1^1,S_2^1}}) \overset{?}{\equiv} (\mathbb{1}^2_{S_1^1} * \tilde{\mathbb{1}}^2_{\mathbb{r}S_2^1}) \equiv \left(\mathbb{1}^2_{S_1^1} * (\tilde{\mathbb{1}}^2_{S_2^1} \circ \mathsf{act}(\mathbb{r}^{-1}))\right). \quad (3.4.247)$$

As a result, the L^1-norm of the convolution counts the number of configurations over the $\mathsf{C}-$space at which such degeneracies happen.

A general pattern is emerging, which requires further research. In general, I conjecture that given two shapes of different dimensions $S_1^{\mathrm{d}_1} \in \mathscr{P}^{\mathrm{d}_1}(\mathsf{E})$ and $S_2^{\mathrm{d}_2} \in \mathscr{P}^{\mathrm{d}_2}(\mathsf{E})$ for arbitrary $0 \leq \mathrm{d}_1, \mathrm{d}_2 < 4$, the different segments of the $\mathsf{C}-$obstacle characterizing intersections of various dimensions can be implicitly characterized by the convolution of $\delta-$singular shape descriptors. In each case, transverse intersections are $\mathrm{d}-$dimensional for $\mathrm{d} = (\mathrm{d}_1 + \mathrm{d}_2) - 3$. However, the conjecture implies that non-transverse intersections that are $\mathrm{d}'-$dimensional for $\mathrm{d}' > (\mathrm{d}_1 + \mathrm{d}_2) - 3$ lead to infinite $\mathrm{d}-$measures, which, in turn, can be conceptualized as $(\mathrm{d}'-\mathrm{d})^{\mathrm{th}}-$order $\delta-$singularities. Although it strikes me as an intuitive proposition, rigorous proof appears nontrivial.

1$-$**Parametric Sweep and Unsweep.** The one-parametric sweep and unsweep of a given shape $S \in \mathscr{P}(\mathsf{E})$ along a given 1D motion trajectory $M \in \mathscr{P}^1(\mathsf{C})$ defined in (3.4.79) and (3.4.81), respectively, can be characterized using convolutions and $\delta-$singularities. They can be characterized implicitly by measuring the inverse trajectory's intersection with the generator shape:

$$\mathbf{x} \in \mathsf{sweep}(M, S) \quad \text{iff} \quad (M^{-1}\mathbf{x}) \cap S \neq \emptyset, \quad \text{i.e.,} \quad (M^{-1}\mathbf{x}) \not\subseteq \mathscr{c}(S), \tag{3.4.248}$$

$$\mathbf{x} \in \mathsf{unswp}(M, S) \quad \text{iff} \quad (M^{+1}\mathbf{x}) \subseteq S, \quad \text{i.e.,} \quad (M^{+1}\mathbf{x}) \cap \mathscr{c}(S) = \emptyset, \tag{3.4.249}$$

In measure-theoretic terms, the original operators can be defined in terms of the sign of the $0-$measure of the above intersections:

$$\mathbf{x} \in \mathsf{sweep}(M, S) \quad \text{iff} \quad \mu^0[(M^{-1}\mathbf{x}) \cap S] > 0, \quad \text{i.e.,} \quad \mu^0[(M^{-1}\mathbf{x}) \cap \mathscr{c}(S)] < \mu^0[M^{-1}],$$

$$\mathbf{x} \in \mathsf{unswp}(M, S) \quad \text{iff} \quad \mu^0[(M^{+1}\mathbf{x}) \cap S] = \mu^0[M^{+1}], \quad \text{i.e.,} \quad \mu^0[(M^{+1}\mathbf{x}) \cap \mathscr{c}(S)] = 0.$$

Since the motion trajectory is 1−measurable, we can define the modified operators $\mathsf{sweep}^{\ddagger}, \mathsf{unswp}^{\ddagger} : \mathscr{P}^1(\mathsf{C}) \times \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ in terms of the sign of the (more manageable) 1−measure of the above intersections:

$$\begin{aligned}
&\mathbf{x} \in \mathsf{sweep}^{\ddagger}(M,S) \quad \text{iff} \quad \mu^1[(M^{-1}\mathbf{x}) \cap S] > 0, \quad \text{i.e.,} \quad \mu^1[(M^{-1}\mathbf{x}) \cap \mathscr{C}(S)] < \mu^0[M^{-1}],\\
&\mathbf{x} \in \mathsf{unswp}^{\ddagger}(M,S) \quad \text{iff} \quad \mu^1[(M^{+1}\mathbf{x}) \cap S] = \mu^1[M^{+1}], \quad \text{i.e.,} \quad \mu^1[M^{+1}\mathbf{x}) \cap \mathscr{C}(S)] = 0.
\end{aligned}$$

The modified sweep and unsweep can thus be characterized in terms of convolutions of δ−singular descriptors of motion trajectories with nonsingular descriptors of solids. To see how it works, note that the measure of intersection between the inverse trajectory and the generator can be obtained as:

$$\mu^1[(M^{\pm 1}\mathbf{x}) \cap S] = \int_{(M^{\pm 1}\mathbf{x}) \cap S} d\mu^1[\mathbf{x}'] = \int_{M^{\pm 1}\mathbf{x}} \mathbf{1}_S(\mathbf{x}')\; d\mu^1[\mathbf{x}'], \tag{3.4.250}$$

which can be rewritten in terms of an inner product if the Euclidean curve $(M^{\pm 1}\mathbf{x}) = \left\{ \mathbb{c}^{\pm 1}\mathbf{x} \mid \mathbb{c} \in M \right\} \in \mathscr{P}^1(\mathsf{E})$ is characterized by a 2^{nd}−order δ−singular shape descriptor:

$$\mu^1[(M^{\pm 1}\mathbf{x}) \cap S] = \int_{\mathsf{E}} \mathbb{1}^2_{M^{\pm 1}\mathbf{x}}(\mathbf{x}')\; \mathbf{1}_S(\mathbf{x}')\; d\mu^3[\mathbf{x}'] = \langle \mathbb{1}^2_{M^{\pm 1}\mathbf{x}}, \mathbf{1}_S \rangle. \tag{3.4.251}$$

Therefore, the one-parametric sweep can be characterized via:

$$f_{\mathsf{sweep}^{\ddagger}(M,S)} \equiv \langle \mathbb{1}^2_{M^{-1}(\cdot)}, \mathbf{1}_S \rangle, \quad \text{i.e.,} \quad \mathsf{sweep}^{\ddagger}(M,S) = U_0\big(\langle \mathbb{1}^2_{M^{-1}(\cdot)}, \mathbf{1}_S \rangle\big). \tag{3.4.252}$$

Noting that $\mathsf{sweep}^*(M,S) \overset{\mathsf{ae}}{=} \mathsf{sweep}^\ddagger(M,S)$ when the generator is a solid (i.e., has no lower-dimensional irregularities) $S \in \mathscr{P}(\mathsf{E})$:

$$f_{\mathsf{sweep}^*(M,S)} \overset{*}{\equiv} \langle \mathbb{1}^2_{M^{-1}(\cdot)}, \mathbf{1}_S \rangle, \quad \text{i.e.,} \quad \mathsf{sweep}^*(M,S) = U_0^*\big(\langle \mathbb{1}^2_{M^{-1}(\cdot)}, \mathbf{1}_S \rangle\big). \tag{3.4.253}$$

On the other hand, the one-parametric unsweep can be characterized via:

$$f_{\mathsf{unswp}^\ddagger(M,S)} \equiv \lim_{\epsilon \to 0^+} \min\big\{0, \langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle - \|\mathbb{1}^2_{M(\cdot)}\|_1 + \epsilon\big\}, \tag{3.4.254}$$

which can alternatively be stated in terms of superlevel sets as in the limit:

$$\mathsf{unswp}^\ddagger(M,S) = \lim_{\epsilon \to 0^+} U_{-\epsilon}\langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle - \|\mathbb{1}^2_{M(\cdot)}\|_1 = \lim_{\epsilon \to 0^+} U_{\|\mathbb{1}^2_{M(\cdot)}\|_1 - \epsilon}\langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle. \tag{3.4.255}$$

Noting that $\mathsf{unswp}^*(M,S) \overset{\mathsf{ae}}{=} \mathsf{unswp}^\ddagger(M,S)$ when the generator is a solid (i.e., has no lower-dimensional irregularities) $S \in \mathscr{P}(\mathsf{E})$:

$$f_{\mathsf{unswp}^*(M,S)} \overset{*}{\equiv} \lim_{\epsilon \to 0^+} \min\big\{0, \langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle - \|\mathbb{1}^2_{M(\cdot)}\|_1 + \epsilon\big\}, \tag{3.4.256}$$

which can alternatively be stated in terms of superlevel sets as in the limit:

$$\mathsf{unswp}^*(M,S) = \lim_{\epsilon \to 0^+} U^*_{-\epsilon}\langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle - \|\mathbb{1}^2_{M(\cdot)}\|_1 = \lim_{\epsilon \to 0^+} U_{\|\mathbb{1}^2_{M(\cdot)}\|_1 - \epsilon}\langle \mathbb{1}^2_{M(\cdot)}, \mathbf{1}_S \rangle. \tag{3.4.257}$$

In other words, measurable sweep is characterized by the inverse trajectory intersecting the generator over a curve of nonzero $1-$measure, while measurable unsweep is characterized by the forward trajectory remaining inside$-\mathsf{ae}$ of the generator.

By noting that sweep and unsweep are just special cases of dilation and erosion,

as per their definitions in (3.4.79) and (3.4.81), the above expressions can be recast as noncommutative convolutions over $\mathsf{C} = \mathrm{SE}(3)$ if the motion trajectory itself is described by a higher-order $\delta-$singularity over the $\mathsf{C}-$space, whose detailed analysis is beyond the scope of this thesis. Here I focus on the more useful attempt to formulate sweep and unsweep implicitly via commutative convolutions over $\mathsf{E} = \mathbb{R}^3$, which, in turn, can be computed efficiently using Fourier transforms.

The next question is, how can we obtain the $2^{\mathrm{nd}}-$order $\delta-$singular descriptor of the curves $(M^{\pm 1}\mathbf{x}) \in \mathscr{P}^1(\mathsf{E})$ that are traced by the forward and inverse motion of an arbitrary query point $\mathbf{x} \in \mathsf{E}$? This is a case example in which an explicit parametrization for a the curve is known, meaning that we can directly use the results of Section 2.3.2.2.

For a one-parametric sweep, one often assumes that the motion is parameterized as $M = \bigcup_{0 \leq i < n_1} \Gamma_i(\Delta^1)$ over an interval $\Delta^1 \subseteq \mathbb{R}$ via $\Gamma_i|_{\mathbb{R}} : \mathbb{R} \to \mathsf{C}$. In that case, the forward and inverse trajectory $(M^{\pm 1}\mathbf{x}) \in \mathscr{P}^1(\mathsf{E})$ for a given query point $\mathbf{x} \in \mathsf{E}$ is also parameterized as:

$$M^{\pm 1}\mathbf{x} = \bigcup_{0 \leq i < n_1} \bigcup_{t_1 \in \Delta^1} \mathsf{act}(\Gamma_i(t_1, 0, 0)^{\pm 1})(\mathbf{x}) = \bigcup_{0 \leq i < n_1} \Gamma^{\pm}_{\mathbf{x},i}(\Delta^1), \tag{3.4.258}$$

where $\Gamma^{\pm}_{\mathbf{x},i} : \mathsf{E} \to \mathsf{E}$ are defined as $\Gamma^{\pm}_{\mathbf{x},i} := \mathsf{act}(\Gamma_i(\cdot)^{\pm 1})(\mathbf{x})$ for all $0 \leq i < n_1$. Substituting for $S^1 := (M^{\pm 1}\mathbf{x})$ and $\Gamma^{\pm}_{\mathbf{x},i}$ for a fixed query point $\mathbf{x} \in \mathsf{E}$ into (2.3.89) of Section 2.3.2.2, one obtains:

$$\mathbb{1}^2_{M^{\pm 1}\mathbf{x}}(\mathbf{x}') := |\mathfrak{J}(t_1, 0, 0)|^{-1} \; \delta^2(t_2, t_3) \; \big|_{\mathbf{x}' = \mathsf{act}(\Gamma_i(t_1, t_2, t_3)^{\pm 1})(\mathbf{x})}. \tag{3.4.259}$$

Note that the above formulation runs into difficulties as soon as the curve has non-

manifold self-intersections; more precisely, if there are more than one parameter point $\mathbf{t} = (t_1, t_2, t_3) \in \mathbb{R}^3$ that is mapped to the same $\mathbf{x}' = \mathsf{act}(\Gamma_i(\mathbf{t})^{\pm 1})(\mathbf{x})$ via the same $\Gamma^{\pm}_{\mathbf{x},i}-$map for some $0 \leq i < n_0$. In other words, if there are non-manifold points on the Euclidean curve $(M^{\pm 1}\mathbf{x}) \in \mathscr{P}^1(\mathsf{E})$ that do not belong to the (potentially non-manifold) connection points between a pair of segments in the $\mathsf{C}-$space curve $M^{\pm 1} \in \mathscr{P}^1(\mathsf{C})$, one needs to re-parameterize the motion with larger number of segments $n_1' > n_1$ to separate the self-intersections into different submanifolds.

The forward and inverse trajectory measurement method explained above can be approximated by using limit representations (e.g., selected from Table 2.3.1) for $\sigma > 0$ to substitute the $2^{\text{nd}}-$order $\delta-$function in (3.4.259). For example, if the normalized indicator function of a (closed or open) $\sigma-$square or $\sigma-$disk is used, the forward and inverse trajectories are replaced with thickened approximations. In this case, the results can be viewed as *probabilistic* sweep and unsweep, in which the uncertainty of the trajectories are modeled by a $\sigma-$tolerance zones with uniform probability distribution over the square or disk cross-section. If a Gaussian $\mathscr{g}_\sigma-$kernel is used, on the other hand, the results can be interpreted as the probability distribution of the sweep and unsweep using a normal distribution for the uncertainty of the trajectories.

The generalization of one-parametric sweep and unsweep, described via (3.4.82) and (3.4.83) in the explicit realm can be carried to the implicit realm as well. For different dimensionality of the generator shape and motion trajectory, different orders of $\delta-$signularities are required. For example, if we assume that the generator is a solid (i.e., regular shape) $S \in \mathscr{P}^*(\mathsf{E})$ and the motion is a $\mathrm{d}-$parameteric trajectory $M = \bigcup_{0 \leq i < n_{\mathrm{d}}} \Gamma_i(\Delta^{\mathrm{d}})$ that is the embedding of a finite collection of $\mathrm{d}-$cells in the $\mathsf{C}-$space for $\mathrm{d} = 0, 1, 2, \cdots$:

- For d := 0, one has $M = \{\mathbb{c}_i\}_{0\leq i<n_0} \in \mathscr{P}^0(\mathsf{C})$, thus the forward and inverse trajectories reduce to finite sets of points $M^{\pm 1}\mathbf{x} = \{\mathbb{c}_i^{\pm 1}\mathbf{x}\}_{0\leq i<n_0} \in \mathscr{P}^0(\mathsf{E})$. Thus $0-$measruing (i.e., counting) the trajectories' intersections with the generator can be computed by $\langle \mathbb{1}^3_{M^{\pm 1}\mathbf{x}}, \mathbf{1}_S\rangle$ at every query point $\mathbf{x} \in \mathsf{E}$, whose $0-$ and $\|\mathbb{1}^3_{M^{\pm 1}}\|_1^--$superlevel sets are the discrete sweep and unsweep. The $3^{\text{rd}}-$order $\delta-$singular trajectory descriptors are obtained using (2.3.73):

$$\mathbb{1}^3_{M^{\pm 1}\mathbf{x}}(\mathbf{x}') = \sum_{0\leq i<n_0} \varsigma^3_{\mathbb{c}_i^{\pm 1}\mathbf{x}}(\mathbf{x}') = \sum_{0\leq i<n_0} \delta^3(\mathbf{x}' - \mathbb{c}_i^{\pm 1}\mathbf{x}). \tag{3.4.260}$$

- For d := 1, the forward and inverse trajectories are a finite union of curve segments, whose intersections $1-$measure with the generator can be computed by $\langle \mathbb{1}^2_{M^{\pm 1}\mathbf{x}}, \mathbf{1}_S\rangle$ at every query point $\mathbf{x} \in \mathsf{E}$, whose $0-$ and $\|\mathbb{1}^2_{M^{\pm 1}}\|_1^--$superlevel sets are the $1-$parameteric sweep and unsweep. The $2^{\text{nd}}-$order $\delta-$singular trajectory descriptors are obtained using (2.3.89).

- For d := 2, the forward and inverse trajectories are a finite union of surface patches, whose intersections $2-$measure with the generator can be computed by $\langle \mathbb{1}^1_{M^{\pm 1}\mathbf{x}}, \mathbf{1}_S\rangle$ at every query point $\mathbf{x} \in \mathsf{E}$, whose $0-$ and $\|\mathbb{1}^1_{M^{\pm 1}}\|_1^--$superlevel sets are the $2-$parameteric sweep and unsweep. The $1^{\text{st}}-$order $\delta-$singular trajectory descriptors are obtained using (2.3.100).

The development can be generalied to lower-dimensional generators—e.g., to implicitly characterize solid modeling operations such as 'extruding', 'revolving', 'lofting', and others supported via explicit methods in most CAD systems. All it takes is to use the proper order of $\delta-$singular shape descriptors for the generator as well. In each scenario, approximations can be obtained by using limit representations of the

$\delta-$functions (e.g., Gaussian or mollifier) to model shape and motion uncertainty, and conceptualize fuzzy or probabilistic sweep and unsweep operations.

Configuration Sampling. To replicate the implicit digital morphology discussed in Section 3.4.2.1, given two finite sets of configurations $M_1^0 = \{\mathbb{c}_{i_1}^{(1)}\}_{0\leq i_1<n_1}$ and $M_2^0 = \{\mathbb{c}_{i_2}^{(2)}\}_{0\leq i_2<n_2}$ with cardinalities $|M_1^0| = n_1$ and $|M_2^0| = n_2$, their implicit descriptors can be expressed in terms of $6^{\text{th}}-$order $\delta-$singularities residing in the $\mathsf{C}-$space, formulated in a similar fashion to (3.4.131) and (3.4.132)—in which subtraction is to be replaced with left- or right-multiplication with inverse of the discrete set of 'impacts' at $\mathbb{c}_{i_1}^{(1)}, \mathbb{c}_{i_2}^{(2)} \in \mathsf{C}$ and the $\delta-$functions are defined in terms of left- and right-Haar $6-$measures and $6-$integrals. Accordingly, the Minkowski product $(M_1^0 \otimes M_2^0)$ and quotients $(M_1^0 \oslash M_2^0)$ and $(M_1^0 \oslash M_2^0)$ can be expressed in terms of convolutions across $\delta-$functions. The convolutions can be transformed to the Fourier domain (in spite of the complications), and the impact functions can be approximated using limit representations. The details are beyond the scope of this thesis.

3.5 Skeletal Transforms

The skeletonization of a shape is of theoretical and practical significance as it captures its connectivity [17]. The different types of shape skeletons and their simplified and stablized modifications in relation with Voronoi diagrams [20] of boundary discretizations has found important applications in a number of areas that deal with shape analysis. In this section, I shall study the medial axis (MA), briefly overview its stability and approximation, and discuss its implicitization using on $\delta-$calculus.

3.5.1 Explicit Skeletonization

3.5.1.1 Pointset-Theoretic Perspective for Evaluating Shape Skeletons:

Definition 3.5.1. The 'core' of a general set $S \in \mathscr{P}(\mathsf{E})$ is defined as the collection of all maximal balls (partially ordered via containment) that are inscribed in the set [67]:

$$\mathscr{C}(S) := \big\{ B_r(\mathbf{x}) \mid B_r(\mathbf{x}) \subseteq B_{r'}(\mathbf{x}') \subseteq S \;\Rightarrow\; (\mathbf{x}, r) = (\mathbf{x}', r') \big\}, \tag{3.5.1}$$

where $B_r(\mathbf{x}) = \big\{ \mathbf{x}' \in \mathsf{E} \mid \|\mathbf{x}' - \mathbf{x}\|_2 < r \big\}$ is an open 3−ball of radius $r > 0$ centered at $\mathbf{x} \in \mathsf{E}$. If the set is open, it can be retrieved from the core as $S = \bigcup_{B \in \mathscr{C}(S)} B$.

Definition 3.5.2. (Skeleton) The 'skeleton' of a general set $S \in \mathscr{P}(\mathsf{E})$ is defined as the set of centers of maximal balls that are inscribed in it [17], i.e.,

$$\mathscr{S}(S) := \big\{ \mathbf{x} \in S \mid \exists r > 0 : B_r(\mathbf{x}) \in \mathscr{C}(S) \big\}. \tag{3.5.2}$$

Thus the mapping $\mathscr{S} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is referred to as 'skeletonization'.

An alternative (and more popular) explicit definition of a skeleton-like object, is the medial axis (MA), which is related to but different from the above definition of shape skeleton, and is defined next:[29]

Definition 3.5.3. (Nearest Neighbors) The set of 'nearest neighbors' (NN) of $\mathbf{x} \in \mathsf{E}$ in a given set $S' \in \mathscr{P}(\mathsf{E})$ is defined as

$$C(S', \mathbf{x}) := \big\{ \mathbf{x}' \in S' \mid \|\mathbf{x} - \mathbf{x}'\|_2 = \mathrm{d}(S', \mathbf{x}) \big\}, \tag{3.5.3}$$

[29] In some literature [67] the object defined in (3.5.2) is used to define MA and skeleton as interchangeable notions, while in most other references (including but not limited to [17, 59, 207]) the distinction between the two is emphasized.

where $\mathbb{d}(S', \mathbf{x}) = \inf_{\mathbf{x}' \in S'} \|\mathbf{x} - \mathbf{x}'\|_2$, thus $C(S', \mathbf{x}) \neq \emptyset$ iff the infimum is also a *minimum*, and all points $\mathbf{x}' \in S'$ that are a minimum distant away from the query point are its NN, all of which lie on a $2-$sphere centered at the query point:

$$C(S', \mathbf{x}) = \mathscr{b}(B_r(\mathbf{x})) \cap S'\big|_{r:=\mathbb{d}(S', \mathbf{x})} \subseteq \mathscr{b}(B_r(\mathbf{x})), \tag{3.5.4}$$

where $r := \mathbb{d}(S', \mathbf{x})$ is often referred to as the 'radius function' [59, 207].

Definition 3.5.4. (Medial Axis) The 'medial axis' (MA) of a general set $S \in \mathscr{P}(\mathsf{E})$ is defined as the set of all points inside the set that have strictly more than one (i.e., at least two) NNs in its complement $\mathscr{c}(S)$ [17]:

$$\mathscr{M}(S) := \big\{\mathbf{x} \in S \mid |C(\mathscr{c}(S), \mathbf{x})| \geq 2\big\}, \tag{3.5.5}$$

Thus the mapping $\mathscr{M} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is (another form of) 'skeletonization'.

Note that in (3.5.3), it is easy to verify that $\mathbb{d}(S', \mathbf{x}) = 0$ iff $\mathbf{x} \in \mathscr{k}(S')$, in which case the following two scenarios can occur:

- $\mathbf{x} \in S'$ in which case $C(S', \mathbf{x}) = \{\mathbf{x}\}$ and $|C(S', \mathbf{x})| = 1$; or
- $\mathbf{x} \in (\mathscr{k}(S') - S') = (\mathscr{b}(S') - S')$ in which case $C(S', \mathbf{x}) = \emptyset$ and $|C(S', \mathbf{x})| = 0$.

Therefore, if $S' := \mathscr{c}(S)$ as in (3.5.5) above, for a given $\mathbf{x} \notin S$ (i.e., $\mathbf{x} \in S'$) one always has $|C(S', \mathbf{x})| < 2$, violating the defining condition for the MA. Thus $\mathbf{x} \in S$ in (3.5.5) can be safely replaced with $\mathbf{x} \in \mathsf{E}$:

$$\mathscr{M}(S) := \big\{\mathbf{x} \in \mathsf{E} \mid |C(\mathscr{c}(S), \mathbf{x})| \geq 2\big\}, \tag{3.5.6}$$

Although I used general sets in the above definitions, it is customary to define and reason about the skeleton and MA only for the open sets in the literature [17, 59, 67, 207, 325, 365], which simplifies the above discussion since $S' := \mathscr{c}(S)$ is a closed set, for which $\mathbb{d}(S', \mathbf{x}) = \inf_{\mathbf{x}' \in S'} \|\mathbf{x} - \mathbf{x}'\|_2 = \min_{\mathbf{x}' \in S'} \|\mathbf{x} - \mathbf{x}'\|_2$ meaning that there is always at least one NN (i.e., $C(S', \mathbf{x}) \neq \emptyset$) even when the query point is on the boundary, i.e., when $\mathbf{x} \in \mathscr{b}(S) = \mathscr{b}(S') \subseteq S'$.

Lemma 3.5.5. *For a closed set $S' = \mathscr{k}(S')$ with manifold boundary $\mathscr{b}(S')$, the NN-set is either singular (i.e., containing the query point alone) or belong to the boundary:*

$$C(S', \mathbf{x}) = C(\mathscr{k}(S'), \mathbf{x}) = \begin{cases} C(\mathscr{b}(S'), \mathbf{x}) & \text{if } \mathbf{x} \notin S', \\ \{\mathbf{x}\} & \text{if } \mathbf{x} \in S'. \end{cases} \tag{3.5.7}$$

For the case $\mathbf{x} \notin S'$ one has $\mathbf{x} \notin \mathscr{b}(S'))$ and (3.5.4) can be altered as follows:

$$C(S', \mathbf{x}) = \mathscr{b}(B_r(\mathbf{x})) \cap \mathscr{b}(S')\big|_{r := \mathbb{d}(\mathscr{b}(S'), \mathbf{x})} \subseteq \mathscr{b}(B_r(\mathbf{x})), \tag{3.5.8}$$

with a nonzero radius function, i.e., $0 < r = \mathbb{d}(S', \mathbf{x}) = \mathbb{d}(\mathscr{b}(S'), \mathbf{x})$. For the other case with a zero radius function, i.e., $0 = r = \mathbb{d}(S', \mathbf{x}) \neq \mathbb{d}(\mathscr{b}(S'), \mathbf{x})$, the spherical superset $\mathscr{b}(B_r(\mathbf{x}))$ and $C(S', \mathbf{x})$ itself degenerate to a single point $\mathbf{x} \in S'$.

Proof. For the case when $\mathbf{x} \in \mathscr{c}(S)$, the proof is trivial. For the case when $\mathbf{x} \in S$, the proof is by contradiction using the Jordan-Brouwer separation theorem [10] and noting that the NN-sets are subsets of concentric $2-$spheres. □

As a result, the definition in (3.5.6) can be altered safely as:

$$\mathscr{M}(S) := \big\{\mathbf{x} \in S \mid |C(\mathscr{b}(S), \mathbf{x})| \geq 2\big\}, \tag{3.5.9}$$

noting that the condition $\mathbf{x} \in S$ is no longer redundant, because if $\mathbf{x} \in \mathscr{c}(S)$ then $C(\mathscr{c}(S), \mathbf{x}) \neq C(\mathscr{b}(S), \mathbf{x})$ and even though $|C(\mathscr{c}(S), \mathbf{x})| = 1$, it might be the case that $|C(\mathscr{b}(S), \mathbf{x})| > 1$. In fact, it is easy to show that removing the condition $\mathbf{x} \in S$ from (3.5.9) would turn it into $\mathscr{M}(S) \cup \mathscr{M}(\mathscr{e}(S))$.

The notions of skeleton and MA are intimately related but different; in fact, it can be shown that the MA is a subset of the skeleton [17], which, in turn, is a subset of the closure of MA (often called the 'cut locus') [365]: $\mathscr{M}(S) \subseteq \mathscr{S}(S) \subseteq \mathscr{k}(\mathscr{M}(S))$.

Definition 3.5.6. (Medial Axis Transform) The medial axis transform (MAT) is defined as an embedding of the MA in the 4D halfspace $\mathsf{E} \times \mathbb{R}_+$ as [67, 325]

$$\mathscr{T}(S) := \big\{(\mathbf{x}, \mathrm{d}(\mathscr{b}(S), \mathbf{x})) \mid \mathbf{x} \in \mathscr{M}(S)\big\}, \tag{3.5.10}$$

Thus the MAT is a lower-dimensional set (i.e., of $0-$, $1-$ or $2-$dimensional pieces) embedded in a higher-dimensional (i.e., 4D) space, which uniquely describes the original open set as a dilation:

$$S = \bigcup_{(\mathbf{x},r)\in\mathscr{T}(S)} B_r(\mathbf{x}) = \mathsf{dil}\left(\mathscr{T}(S), B_1(\mathbf{0})\right) = \pi_0\big(\mathscr{T}(S) \otimes \ell_0(B_1(\mathbf{0}))\big), \tag{3.5.11}$$

where the later expression should be interpreted as a morphological dilation of the MAT—viewed as a collection of configurations $\mathbb{c} \overset{0}{=} (\mathbf{x}, r)$ embedded into a group of translations and uniform scaling $\mathsf{G} \cong \mathrm{T}(3) \ltimes \mathbb{R}_+$ acting on the $3-$space—with a ball of unit radius centered at the origin $B_1(\mathbf{0})$. The dilation, in turn, is expressed in terms of a Minkowski product using the lifting/projection scheme similar to the one presented in Section 3.4.1.2—except that here with the nonrigid scaling, the 3D balls are lifted to 4D halfcones, whose different slices along the 4^{th} dimension gives balls of

varying sizes. The dilation can be conceptualized as sweeping an open 3D ball of unit radius, as it is translated along the MA and resized according to the radius function. Thus the MAT is a more compact description of the shape (including topology and geometry), a discretized approximation of which using a countable sample of points on the MAT and its analytic interpretation will be central to Chapter 4.[30]

The concept of MA was first introduced by Blum [45] as a tool for image analysis, and has been used extensively ever since in computer vision, surface reconstruction, and solid modeling. He conceptualized it in terms of 'wave fronts' and 'corners' and showed how it captures the essential features of form qualitatively. Lieutier [207] showed that the MA of every bounded open set is homotopy equivalent to the set itself—weaker forms of which were proved earlier by other researchers using some regularity assumptions [67,325,365]—meaning that it captures the connectivity characteristics in the most compact structure. Moreover, the distance to the MA (i.e., the radius function) characterizes the 'local feature size' [59] with numerous recent design and manufacturing applications such as tool path planning for subtractive manufacturing [116,336] and design correction for additive manufacturing [233,237].

Since solid modeling typically deals with (closed or open) regular semianalytic sets $S \in \mathscr{P}^*(\mathsf{E})$, in that context one usually speaks of the MAs of a given solid's interior $(\mathscr{M} \circ \mathscr{i})(S)$ and/or exterior $(\mathscr{M} \circ \mathscr{e})(S)$, both of which are open and semianalytic:

$$\mathscr{M}(\mathscr{i}(S)) = \big\{\mathbf{x} \in \mathscr{i}(S) \mid |C(\mathscr{b}(S), \mathbf{x})| \geq 2\big\}, \tag{3.5.12}$$

$$\mathscr{M}(\mathscr{e}(S)) = \big\{\mathbf{x} \in \mathscr{e}(S) \mid |C(\mathscr{b}(S), \mathbf{x})| \geq 2\big\}, \tag{3.5.13}$$

[30] Skeletonization is often considered a fundamental operation in mathematical morphology [139, 308]. One can alternatively think of this dilation as the 3D slice of a 4D Minkowski sum of the MAT with a canonical 4D half-cone that results from lifting a 4D ball [34, 36]. Skeletonization is often considered a fundamental operation in mathematical morphology [139, 308].

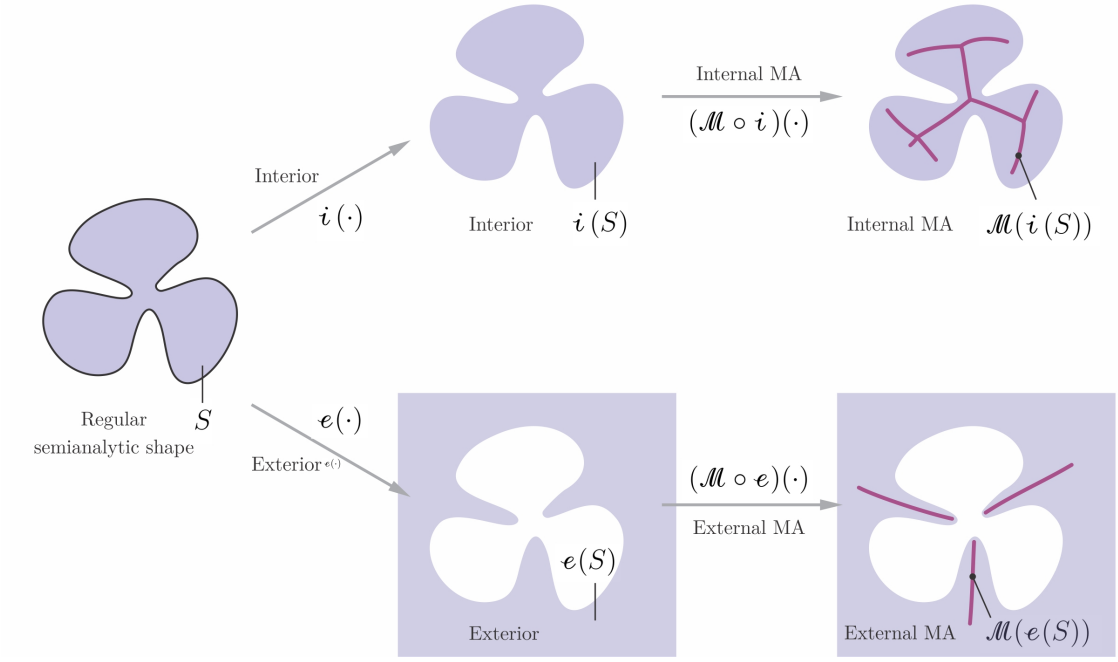

FIGURE 3.5.1: The MA of the interior and exterior of a given shape is the set of points that are equidistant from at least two points (i.e., NNs) on the boundary of the shape.

and their MATs $(\mathscr{T} \circ \mathscr{i})(S)$ and/or $(\mathscr{T} \circ \mathscr{e})(S)$, which are also open and semianalytic:

$$\mathscr{T}(\mathscr{i}(S)) = \big\{(\mathbf{x}, \mathrm{d}(\mathscr{b}(S), \mathbf{x})) \mid \mathbf{x} \in \mathscr{M}(\mathscr{i}(S))\big\}, \tag{3.5.14}$$

$$\mathscr{T}(\mathscr{e}(S)) = \big\{(\mathbf{x}, \mathrm{d}(\mathscr{b}(S), \mathbf{x})) \mid \mathbf{x} \in \mathscr{M}(\mathscr{e}(S))\big\}, \tag{3.5.15}$$

noting that the interior, the exterior, and the shape itself share the same boundary, i.e., $\mathscr{b}(S) = (\mathscr{b} \circ \mathscr{e})(S) = (\mathscr{b} \circ \mathscr{i})(S)$. Figure 3.5.1 illustrates the MA of interior and exterior for a simple planar shape.

Chazal and Soufflet [59] showed that MAs of semianalytic sets are (at most 2D) subanalytic sets[31]—thus admit finite analytic stratifications to vertices, curved edges,

[31] In fact, they showed that this is almost the best one can hope for, since even simple $C^\infty-$manifold surfaces can produce infinitely many MA branches, thus restricting to $C^\omega-$manifolds is the least conservative choice. Earlier, Choi et al. showed a similar result for 2D shapes [67].

and curved triangles—and in particular, MAs of semialgebraic sets are (at most 2D) semialgebraic sets which can be explicitly computed by solving a finite set of algebraic equations. Therefore, from a computational perspective, the explicit evaluation of the MA(T) is a question of mapping a finite set-theoretic representation (FSR) of the shape to a FSR of its MA(T)—most commonly shape B-rep or spherical grouping to MA(T) strata, due to the links between boundaries, inscribed balls, and the MA(T).

Although exact computation of the MA(T) is possible in principle for general semialgebraic sets, the most advanced and effective implementations are limited to planar shapes [268, 269], polygons in 2D [53, 54, 171], polyhedra in 3D [85–87], and finite groupings of balls [13] or their combinations, which compute the FSR of the MA from Voronoi or Apollonious diagrams of the finite number of primitives in the approximate FSR of the shape [20, 21].[32] For most practical cases with arbitrary shapes, the MA is computed approximately through

1. approximating the shape with another 'surrogate' shape from a class of constructions for which the MA can be computed exactly—e.g., most commonly approximating its boundary with a finite sample of points that is dense enough to capture the smallest features to some adequate resolution;

2. computing the MA of the approximate surrogate—e.g., the MA of the complement of the discrete pointset sampled on the boundary (i.e., the 'punctured' 3−space [17]), which is the same as the sample's Voronoi graph, i.e., the union of the 2−cells; namely, faces, edges, and vertices of the Voronoi diagram; and

[32]Decades of research in computational geometry have resulted in optimal sequential and parallel algorithms for Voronoi diagrams, power diagrams, Apollonious diagrams, Johnson-Mehl diagrams, etc. Among the most effective methods in practice are randomized parallel algorithms (e.g., via 'polling') by Reif and Sen [267, 272–274] and Rajasekaran and Ramaswami [264–266].

3. pruning the MA of the surrogate to a particular subset according to a set of rules to obtain a polyhedral approximation of the original shape's MA.

The notion of approximation used in the last step varies across different implementations of the approximation method in the first step.[33] Once again, there are notoriously long-recognized—despite only recently well-understood—computability difficulties with approximating the MA due to its local instability issues, e.g., with respect to C^0- and C^1-perturbations, even though it is stable with respect to C^2- and higher-order perturbations [59]. Informally, although small modifications of a shape (e.g., due to noise/errors) do not globally affect the entire medial axis, it can create fluctuating branches (e.g, in 2D) and fluctuating spikes (e.g., in 3D) added to or removed from the otherwise stable structure—formalized by the fact that MA is 'semi-continuous' with respect to the Hausdorff metric [17]. Therefore, MA is not strictly computable [105, 106] with approximate methods, because there is no guarantee (in general) that the surrogate MA approximates the original shape's MA up to some reasonable Hausdorff distance error. The good news is that when using punctured spaces, under certain technical conditions on the boundary (e.g., C^1-continuity) and/or sampling (e.g., surface incidence)—which are much simpler in 2D [49] than in 3D [12, 93–95]—the MA can be approximated for the Hausdorff distance from a subset of the Voronoi graph of the sample. However, in most general cases the MA needs to be 'pruned' by shortening the peripheral branches, trying to capture the stable part. Typically, this is done by successful removal of borderline-stable points using a variety of heuristic stopping conditions, either based on the deviation of the

[33] It either means that the surrogate shape is the image of the original shape under a small C^n-perturbation for some $n \geq 0$, or that the Hausdorff distance between the (complements of) the two is small [59].

shape reconstructed from the pruned MA and the original shape [50, 68, 96, 98, 310], or based on some local estimate of stability of the MA [18, 19, 121, 122, 247].

Overall, in spite of the ubiquity and versatility of MA(T) as powerful aids for shape analysis, their explicit computation for objects of arbitrary shapes is an ill-posed problem from a finite-precision computational perspective, and has contributed to their unpopularity within some academic circles.

The inverse problem of retrieving a solid from the interior or exterior MAT corresponds to computing the sweep of a ball along the MA as its size scales with the radius function (i.e., the 4^{th} coordinate of the MAT), or alternatively, the 3D slice through the sweep of a 4D canonical halfcone along the MAT in the lifted 4D space—which forms the basis for a finite spherical decomposition of a shape [34, 36] for implementation purposes in Chapter 4.

3.5.2 Implicit Skeletonization

3.5.2.1 Measure-Theoretic Perspective for Evaluating Shape Skeletons:

The definition of the MA given in (3.5.5) can be expressed implicitly by *counting* the number of nearest neighbors (NN) on the complement of a given general set $S \in \mathscr{P}(\mathsf{E})$ at a given query point $\mathbf{x} \in \mathsf{E}$:

$$\mathbf{x} \in \mathscr{M}(S) \quad \text{iff} \quad \mathbf{x} \in S \quad \text{and} \quad |C(\mathscr{c}(S), \mathbf{x})| \geq 2, \tag{3.5.16}$$

for which I argued in Section 3.5.1 that the condition $\mathbf{x} \in S$ is redundant, noting that $\mathbf{x} \notin S$ implies $|C(\mathscr{c}(S), \mathbf{x})| \leq 1$. For an open set—i.e., one with closed complement $\mathscr{c}(S) = (\mathscr{k} \circ \mathscr{c})(S) = \mathscr{h}(S)$—the two of which are assumed to be separated by a man-

ifold boundary $\partial(S) = \overline{\mathcal{K}}(S) \cap \overline{\mathcal{K}}(S)$, Lemma 3.5.5 allows for an alternative definition in terms of the boundary NN count instead of the complement NN count:

$$\mathbf{x} \in \mathcal{M}(S) \quad \text{iff} \quad \mathbf{x} \in S \quad \text{and} \quad |C(\partial(S), \mathbf{x})| \geq 2, \tag{3.5.17}$$

The NN-set defined in Definition 3.5.3, repeated below for $\mathbf{x} \in S$:

$$C(\partial(S), \mathbf{x}) = \partial(B_r(\mathbf{x})) \cap \partial(S)\big|_{r:=\mathrm{d}(\partial(S),\mathbf{x})}, \tag{3.5.18}$$

is the subset of a 2−sphere $\partial(B_r(\mathbf{x}))$ of radius $r = \mathrm{d}(\partial(S), \mathbf{x})$ that is overlapped with the boundary $\partial(S)$—which shrinks to a single point $\mathbf{x} = \partial(B_0(\mathbf{x}))$ when $\mathbf{x} \in \mathcal{C}(S)$. The MAT is subsequently characterized implicitly as

$$(\mathbf{x}, r) \in \mathcal{T}(S) \quad \text{iff} \quad \mathbf{x} \in \mathcal{M}(S) \quad \text{and} \quad r = \mathrm{d}(\partial(S), \mathbf{x}). \tag{3.5.19}$$

Given a 3D (closed or open) regular set $S \in \mathcal{P}^*(\mathsf{E})$, it partitions the space into three sets, one of which is a closed set; namely, the 2D boundary, studied in Section 3.3.2.2. The other two are open sets; namely, the 3D interior and the 3D exterior (i.e., interior of complement) whose MAs carry a great deal of information regarding the topological and geometric properties of the shape of three sets:

$$\mathbf{x} \in \mathcal{M}(\mathcal{I}(S)) \quad \text{iff} \quad \mathbf{x} \in \mathcal{I}(S) \quad \text{and} \quad |C(\partial(S), \mathbf{x})| \geq 2, \tag{3.5.20}$$

$$\mathbf{x} \in \mathcal{M}(\mathcal{E}(S)) \quad \text{iff} \quad \mathbf{x} \in \mathcal{E}(S) \quad \text{and} \quad |C(\partial(S), \mathbf{x})| \geq 2, \tag{3.5.21}$$

whose corresponding MATs can be characterized implicitly as

$$(\mathbf{x}, r) \in \mathscr{T}(\mathscr{i}(S)) \quad \text{iff} \quad \mathbf{x} \in \mathscr{M}[\mathscr{i}(S)] \quad \text{and} \quad r = \mathbb{d}(\mathscr{b}(S), \mathbf{x}), \tag{3.5.22}$$

$$(\mathbf{x}, r) \in \mathscr{T}(\mathscr{e}(S)) \quad \text{iff} \quad \mathbf{x} \in \mathscr{M}[\mathscr{e}(S)] \quad \text{and} \quad r = \mathbb{d}(\mathscr{b}(S), \mathbf{x}), \tag{3.5.23}$$

noting that the interior, the exterior, and the shape itself share the same boundary, i.e., $\mathscr{b}(S) = (\mathscr{b} \circ \mathscr{e})(S) = (\mathscr{b} \circ \mathscr{i})(S)$.

On the other hand, the minimum distance to the boundary $\mathbb{d}(\mathscr{b}(S), \mathbf{x})$ can be viewed as the absolute value of the signed distance query defined at the beginning of the chapter in (3.1.2). Thus (3.5.18) can be rewritten as

$$C(\mathscr{b}(S), \mathbf{x}) = \mathscr{b}(B_r(\mathbf{x})) \cap \mathscr{b}(S)\big|_{r:=|\mathrm{d}_S(\mathbf{x})|}. \tag{3.5.24}$$

In general, the NN-set $C := C(\mathscr{b}(S), \mathbf{x})$ can be a single point (i.e., $|C| = 1$), a finite or countable number of point ($1 < |C| \leq |\mathbb{N}|$), or an uncountable continuum set (i.e., $|C| = |\mathbb{R}|$) comprising geodesic curves and/or spherical surface patches on $\mathscr{b}(B_r(\mathbf{x}))$. Thus an implicit definition of the MA amounts to *counting*, i.e., 0−measuring the NN-set. In other words, answering the query $(\mathbf{x} \overset{?}{\in} \mathscr{M}(S))$ amounts to checking if $|C| = \mu^0[C] \overset{?}{\geq} 2$, i.e., if $|C| = \mu^0[C] \overset{?}{>} 1$ meaning that $(\mu^0[C] - 1)$ clearly gives a nonnegative real-valued defining function for the MA. The challenge is in quantifying $\mu^0[C]$ in terms of the simple (e.g., inclusion and distance) queries on the shape.

It turns out to be significantly simpler to start with higher-dimensional measures such as $\mu^1[C]$ or $\mu^2[C]$, thus the following definitions become useful:

Definition 3.5.7. (d−Medial Axis) The 'd−medial axis' (d−MA) is the subset of

the MA for which the NN-set has a nonzero d−measure, i.e.,

$$\mathbf{x} \in \mathscr{M}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathbf{x} \in \mathscr{M}(S) \quad \text{and} \quad \mu^{\mathrm{d}}[C(\mathscr{c}(S), \mathbf{x})] > 0, \tag{3.5.25}$$

Thus the mapping $\mathscr{M}^{\mathrm{d}} : \mathscr{P}(\mathsf{E}) \to \mathscr{P}(\mathsf{E})$ is referred to as 'd−skeletonization'.

Once again, for open sets with manifold boundaries, Lemma 3.5.5 allows for an alternative definition in terms of the boundary NN d−measure instead of the complement NN d−measure:

$$\mathbf{x} \in \mathscr{M}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathbf{x} \in \mathscr{M}(S) \quad \text{and} \quad \mu^{\mathrm{d}}[C(\mathscr{b}(S), \mathbf{x})] > 0, \tag{3.5.26}$$

Note that $\emptyset = \mathscr{M}^3(S) \subseteq \mathscr{M}^2(S) \subseteq \mathscr{M}^1(S) \subseteq \mathscr{M}^0(S) = \mathscr{M}(S)$. Intuitively,[34]

- The 2−MA only includes vertices on the MA whose NN-sets are spherical patches on the boundary;[35]
- The 1−MA contains 2−MA but also contains 1D curve segments on the MA whose NN-sets are geodesic arcs on the boundary;
- The 0−MA contains 1−MA but also contains 2D surface patches on the MA whose NN-sets are isolated points on the boundary;

Definition 3.5.8. (d−Medial Axis Transform) The 'd−medial axis transform' (d−MAT)

[34]Note that the MA topology can be extremely difficult to reason formally about when it comes to 3D shapes bounded partially with canal (i.e., tubular) surfaces. See [59] for a simple yet nonintuitive example of the union of a sphere and a cylinder.

[35]Such point are rare in general shapes with arbitrary surfaces, though are ubiquitous in mechanical parts with spherical fillets (where two cylindrical fillets meet at a corner), and molecular models (where collections of atoms are modeled as groupings of balls).

is defined as an embedding of the d−MA in the 4D halfspace $\mathsf{E} \times \mathbb{R}_+$, i.e.,

$$(\mathbf{x}, r) \in \mathcal{T}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathbf{x} \in \mathcal{M}^{\mathrm{d}}(S) \quad \text{and} \quad r = |\mathrm{d}_S(\mathbf{x})|, \tag{3.5.27}$$

As long as we have a gadget to measure d−overlaps between the shape's boundary $\mathcal{b}(S)$ and the 2−sphere $\mathcal{b}(B_r(\mathbf{x}))$ for $r = \mathrm{d}(\mathcal{b}(S), \mathbf{x})$, we can characterize the d−MA(T)—with some limitations. In particular, applying Definition 2.3.14 of Section 2.3.2.4 for the d−overlap predicate to the intersection of surfaces in (3.5.24) and substituting it into (3.5.25) yields the following for $\mathrm{d} > 0$:

$$\mathbf{x} \in \mathcal{M}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathbf{x} \in S \text{ and } \mathsf{col}^{\mathrm{d}}\big(\mathcal{b}(S), \mathcal{b}(B_r(\mathbf{x}))|_{r:=|\mathrm{d}_S(\mathbf{x})|}\big) = 1, \tag{3.5.28}$$

noting that for $\mathrm{d} > 0$, having a nonzero d−measure (i.e., $\mu^{\mathrm{d}}[C] > 0$) automatically implies having an infinite 0−measure (i.e., cardinality $|C| = \mu^0[C] > 0$), which means that the condition $\mathbf{x} \in \mathcal{M}(S) = \mathcal{M}^0(S)$ on the right-hand side of (3.5.25) can be reduced to the inclusion $\mathbf{x} \in S$ in (3.5.17), which, in turn, can be incorporated by demanding the signed distance function to have a nonpositive value:

$$\mathbf{x} \in \mathcal{M}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathcal{b}(S), \mathcal{b}(B_r(\mathbf{x}))|_{r:=-\mathrm{d}_S(\mathbf{x})}\big) = 1 \text{ and } r > 0. \tag{3.5.29}$$

Accordingly, one obtains for the d−MAT within the 4D halfspace $\mathsf{E} \times \mathbb{R}_+$:

$$(\mathbf{x}, r) \in \mathcal{T}^{\mathrm{d}}(S) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathcal{b}(S), \mathcal{b}(B_r(\mathbf{x}))\big) = 1 \text{ and } r = -\mathrm{d}_S(\mathbf{x}), \tag{3.5.30}$$

in which $r > 0$ is automatically implied from $(\mathbf{x}, r) \in (\mathsf{E} \times \mathbb{R}_+)$.

For a regular solid $S \in \mathcal{P}^*(\mathsf{E})$, the implicit definitions of the internal and external

d−MAs are expressed in terms of the shared boundary as

$$\mathbf{x} \in \mathscr{M}^{\mathrm{d}}(\mathit{i}\,(S)) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathscr{b}(S), \mathscr{b}(B_r(\mathbf{x}))|_{r:=-\mathrm{d}_S(\mathbf{x})}\big) = 1 \text{ and } r > 0, \tag{3.5.31}$$

$$\mathbf{x} \in \mathscr{M}^{\mathrm{d}}(\mathit{e}(S)) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathscr{b}(S), \mathscr{b}(B_r(\mathbf{x}))|_{r:=+\mathrm{d}_S(\mathbf{x})}\big) = 1 \text{ and } r > 0, \tag{3.5.32}$$

whose corresponding d−MATs can be characterized implicitly as

$$(\mathbf{x}, r) \in \mathscr{T}\,(\mathit{i}\,(S)) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathscr{b}(S), \mathscr{b}(B_r(\mathbf{x}))\big) = 1 \text{ and } r = -\mathrm{d}_S(\mathbf{x}), \tag{3.5.33}$$

$$(\mathbf{x}, r) \in \mathscr{T}(\mathit{e}(S)) \quad \text{iff} \quad \mathsf{col}^{\mathrm{d}}\big(\mathscr{b}(S), \mathscr{b}(B_r(\mathbf{x}))\big) = 1 \text{ and } r = +\mathrm{d}_S(\mathbf{x}). \tag{3.5.34}$$

The interesting observation is that every maximal ball $B_r(\mathbf{x})$ can be viewed as an instance of a ball of unit radius centered at the origin $B_1(\mathbf{0})$, after applying a translation and uniform scaling $\mathbb{c} \overset{0}{=} (\mathbf{x}, r)$, where the tuple $(\mathbf{x}, r) \in (\mathsf{E} \times \mathbb{R}_+)$ in the 4D halfspace is viewed as a representation of the configurations of the ball $\mathbb{c} \in \mathsf{G}$ where the Lie group $\mathsf{G} := \mathrm{T}(3) \ltimes \mathbb{R}_+$ acts on the 3−space. Then (3.5.30) implies that the d−MAT can be obtained as the (representation of) a proper subset of the following G−space d−obstacle

$$O^{\mathrm{d}}_{\mathscr{b}(S),\mathbb{S}^2} = \big\{\mathbb{c} \in \mathsf{G} \mid \mathsf{col}^{\mathrm{d}}\big(\mathscr{b}(S), \mathbb{c}\mathbb{S}^2\big) \text{ i.e., } \mu^{\mathrm{d}}[\mathscr{b}(S) \cap (\mathbb{c}\mathbb{S}^2)] > 0\big\}, \tag{3.5.35}$$

in which the unit 2−sphere $\mathbb{S}^2 = \big\{\mathbf{x}' \in \mathsf{E} \mid \|\mathbf{x}'\|_2 = 1\big\} = \mathscr{b}(B_1(\mathbf{0}))$ is translated and scaled as $\mathbb{c}\mathbb{S}^2 = \mathscr{b}(B_r(\mathbf{x}))$ for $\mathbb{c} \overset{0}{=} (\mathbf{x}, r)$ onto the boundary of the maximal ball. The d−MATs are proper subsets of $O^{\mathrm{d}}_{\mathscr{b}(S),\mathbb{S}^2}$ for which the additional conditions of maximality of the instantiated 2−sphere (i.e., $r = |\mathrm{d}_S(\mathbf{x})|$) in (3.5.33) and (3.5.34) hold. If we think of the 4D graph of the distance function as a lifting of the 3D shape

into the 4^{th} dimension as[36]

$$\ell_0(S) = \big\{\mathbb{c} \in \mathsf{G} \mid \mathbb{c} \overset{0}{=} (\mathbf{x}, r) \text{ and } r = \mathbb{d}_S(\mathbf{x})\big\} = \big\{\gamma_0(\mathbf{x}, \mathbb{d}_S(\mathbf{x})) \mid \mathbf{x} \in \mathsf{E}\big\}, \tag{3.5.36}$$

where $\gamma_0 : (\mathsf{E} \times \mathbb{R}_+) \hookrightarrow \mathsf{G}$ is the obvious embedding of the 4D halfspace into the configuration space G with respect to a particular choice of origin $\mathbf{0} \in \mathsf{E}$, then

$$\mathscr{T}^{\mathrm{d}}(\dot{\imath}\,(S)) = \gamma_0^{-1}\big(O^{\mathrm{d}}_{\mathring{\mathbb{b}}(S),\mathbb{S}^2} \cap \ell_0(S)\big) \cap \dot{\imath}\,(S), \tag{3.5.37}$$

$$\mathscr{T}^{\mathrm{d}}(\dot{\boldsymbol{e}}(S)) = \gamma_0^{-1}\big(O^{\mathrm{d}}_{\mathring{\mathbb{b}}(S),\mathbb{S}^2} \cap \ell_0(S)\big) \cap \dot{\boldsymbol{e}}(S), \tag{3.5.38}$$

whose union gives the (representation of) the entire $\mathsf{G}-$space $\mathrm{d}-$obstacle restricted to the 4D graph of the distance function:

$$O^{\mathrm{d}}_{\mathring{\mathbb{b}}(S),\mathbb{S}^2} \cap \ell_0(S) \overset{0}{=} \mathscr{T}^{\mathrm{d}}(\dot{\imath}\,(S)) \cup \mathscr{T}^{\mathrm{d}}(\dot{\boldsymbol{e}}(S)). \tag{3.5.39}$$

Let us first consider the simplest case with $\mathrm{d} := 1$ or 2, for which an implicit characterization of the $\mathrm{d}-$MA(T)s requires checking if the intersection of the shape's boundary and the maximal sphere has a nonzero $1-$ or $2-$measure (i.e., contact over a continuum set), respectively. We already know how to obtain the $1^{\text{st}}-$order $\delta-$singluar surface descriptor of the shape's boundary $\mathbb{1}^1_{\mathring{\mathbb{b}}(S)} : \mathsf{E} \to \mathbb{R}^1_\star$ from Section 3.3.2.2, i.e., by $\delta-$sifting the signed distance function's $0-$isolevel set:

$$\mathbb{1}^1_{\mathring{\mathbb{b}}(S)}(\mathbf{x}') = (\delta^1 \circ \mathbb{d}_S)(\mathbf{x}') = \delta^1\big(\mathbb{d}_S(\mathbf{x}')\big). \tag{3.5.40}$$

[36] Think of the trivariate function $\mathbb{d}_S : \mathbb{R}^3 \to \mathbb{R}$ being plotted as a 3D hypersurface in a 4D graph whose ordinate is along the 4^{th} dimension

The same can be applied to obtain the surface descriptor of the maximal ball, hereafter denoted via the simplified notation $B := B_r(\mathbf{x})$, in which the query point $\mathbf{x} \in \mathsf{E}$ and $r := \mathbb{d}_S(\mathbf{x})$ are viewed as fixed; thus $\mathbb{1}^1_{\mathscr{b}(B)} : \mathsf{E} \to \mathbb{R}^1_\star$ is obtained as:

$$\mathbb{1}^1_{\mathscr{b}(B)}(\mathbf{x}') = (\delta^1 \circ \mathbb{d}_B)(\mathbf{x}') = \delta^1\big(\mathbb{d}_B(\mathbf{x}')\big), \tag{3.5.41}$$

in which the distance function $\mathbb{d}_B(\mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2 - r$ for the 3−ball $B = B_r(\mathbf{x})$ is a radial basis function (RBF) corresponding to a fixed $(\mathbf{x}, r) \in (\mathbb{R}^3 \times \mathbb{R}_+)$. The following definition will be helpful in the upcoming discussions:

Definition 3.5.9. (ζ-descriptor) Let $\mathsf{zmap} : \mathscr{P}^*(\mathsf{E}) \to (\mathbb{R} \times \mathbb{R}_+)^{\mathsf{E}}$ map a given shape $S \in \mathscr{P}^*(\mathsf{E})$ to its so-called '$\zeta-$descriptor' defined as:

$$\zeta_S := \mathsf{zmap}(S) \in (\mathbb{R} \times \mathbb{R}_+)^{\mathsf{E}}, \quad \text{where} \quad \zeta_S(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2 - |\mathbb{d}_S(\mathbf{x})|. \tag{3.5.42}$$

The bivariate $\zeta-$descriptor can be viewed as a family of RBFs, i.e., for a fixed query point $\mathbf{x} \in \mathsf{E}$, one has $\zeta_S(\mathbf{x}, \mathbf{x}') = \mathbb{d}_B(\mathbf{x}')$ which is an RBF centered at the query point, whose 0−isolevel is the 2−sphere $\mathscr{b}(B) = \mathscr{b}(B_r(\mathbf{x}))$ of radius $r := \mathbb{d}_S(\mathbf{x})$. Accordingly, (3.5.41) can be rewritten as $\mathbb{1}^1_{\mathscr{b}(B)} = (\delta^1 \circ \zeta_S(\mathbf{x}, \cdot))$.

Using the null-overlap lemma (Lemma 2.3.15) with $\mathrm{d} := 1$, a 1−overlap predicate for the boundary $\mathscr{b}(S)$ and the 2−sphere $\mathscr{b}(B) = \mathscr{b}(B_r(\mathbf{x}))$ can be obtained as

$$\langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B)} \rangle > 0 \;\rightleftharpoons\; \mu^1[C(\mathscr{b}(S), \mathbf{x})] = \mu^1[\mathscr{b}(S) \cap \mathscr{b}(B)] > 0, \tag{3.5.43}$$

meaning that the surfaces have contact over a set of nonzero (finite or infinite) curve length (i.e., $\mu^1[C] > 0$) characterized as a 1−collision (i.e., $\mathsf{col}^1\big(\mathscr{b}(S), \mathscr{b}(B_r(\mathbf{x}))\big) = 1$)

iff the inner product of their $1^{\text{st}}-$order $\delta-$singular surface descriptors is nonzero (i.e., $\langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B)} \rangle > 0$), which resolves the first condition in (3.5.31) through (3.5.34) for the case of $\mathrm{d} = 1$, while the maximality condition (i.e., $r = \pm \mathrm{d}_S(\mathbf{x})$) is enforced explicitly. Thus the inner product $\langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B)} \rangle$ is a defining function for the $1-$MA of the set's interior and exterior, which are to be separated by enforcing containment in the original set using its indicator function. It can be rewritten in terms of the boundary descriptor and $\zeta-$descriptor as

$$\langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B)} \rangle = \big\langle (\delta^1 \circ \mathrm{d}_S), (\delta^1 \circ \mathrm{d}_B) \big\rangle = \big\langle (\delta^1 \circ \mathrm{d}_S), (\delta^1 \circ \zeta_S(\mathbf{x}, \cdot)) \big\rangle, \tag{3.5.44}$$

which can be also interpreted as a convolution over the Lie group $\mathsf{G} \cong \mathrm{T}(3) \rtimes \mathbb{R}_+$.

To informally describe its measure-theoretic properties, note that the extended-real-valued 4D function $g(\mathbf{x}, r) := \langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B_r(\mathbf{x}))} \rangle$ over the $4-$space $(\mathsf{E} \times \mathbb{R}_+)$ or its 3D projection $f(\mathbf{x}) := g(\mathbf{x}, |\mathrm{d}_S(\mathbf{x})|) = \langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B_r(\mathbf{x}))} \rangle|_{r=|\mathrm{d}_S(\mathbf{x})|}$ into the $3-$space E reduce to a volumetric $3-$integral over the contact set $C(\mathscr{b}(S), \mathbf{x}) = \mathscr{b}(S) \cap \mathscr{b}(B_r(\mathbf{x}))$ where the two $1^{\text{st}}-$order $\delta-$singularities magnify each other to $1 + 1 = 2^{\text{nd}}-$order, whose integration over the contact set C yields

- $f(\mathbf{x}) \to +\infty$ if $\mu^1[C] \to +\infty$, i.e., if C contains nonzero-area surface patches;
- $0 < f(\mathbf{x}) < +\infty$ if $0 < \mu^1[C] < +\infty$, i.e., if C contains at most nonzero-length curve segments; and
- $f(\mathbf{x}) = 0$ if $\mu^1[C] = 0$, i.e., if C is a countable discrete pointset;

Thus it appears that not only are f and g real-valued defining functions of $1-$MA(T)s, but they are also $\delta-$singular extended-real-valued defining functions of $2-$MA(T)s,

of the interior and exterior, respectively. For the 1− and 2−MAs one has

$$f_{(\mathscr{M}^1\circ i)(S)} = f^0_{\mathscr{M}^1(i(S))} \equiv f^1_{\mathscr{M}^2(i(S))} \equiv \quad \mathbf{1}_S \bullet \langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B_r(\cdot))} \rangle \big|_{r=|\mathrm{d}_S(\cdot)|}, \tag{3.5.45}$$

$$f_{(\mathscr{M}^1\circ e)(S)} = f^0_{\mathscr{M}^1(e(S))} \equiv f^1_{\mathscr{M}^2(e(S))} \equiv \neg\mathbf{1}_S \bullet \langle \mathbb{1}^1_{\mathscr{b}(S)}, \mathbb{1}^1_{\mathscr{b}(B_r(\cdot))} \rangle \big|_{r=|\mathrm{d}_S(\cdot)|}. \tag{3.5.46}$$

The first expression can be illustrated with the following commutative diagram:

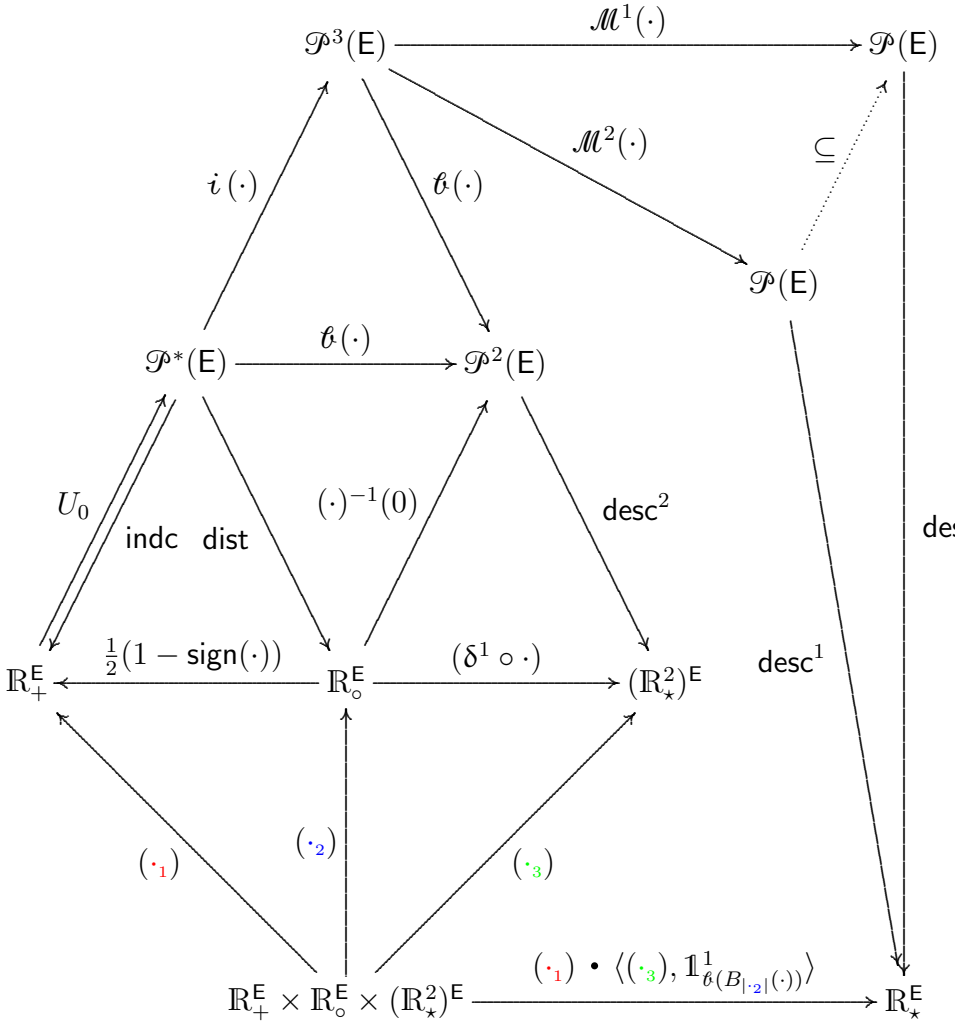

(3.5.47)

which uses (3.3.23) of Section 3.3.2.2 for boundary evaluation. A similar diagram can

be obtained by replacing $\dot{\imath}(\cdot)$ with $\dot{e}(\cdot)$ at the top, and $(\cdot_1)$ with $\neg(\cdot_1)$ on the bottom horizontal arrow, partially illustrated below:

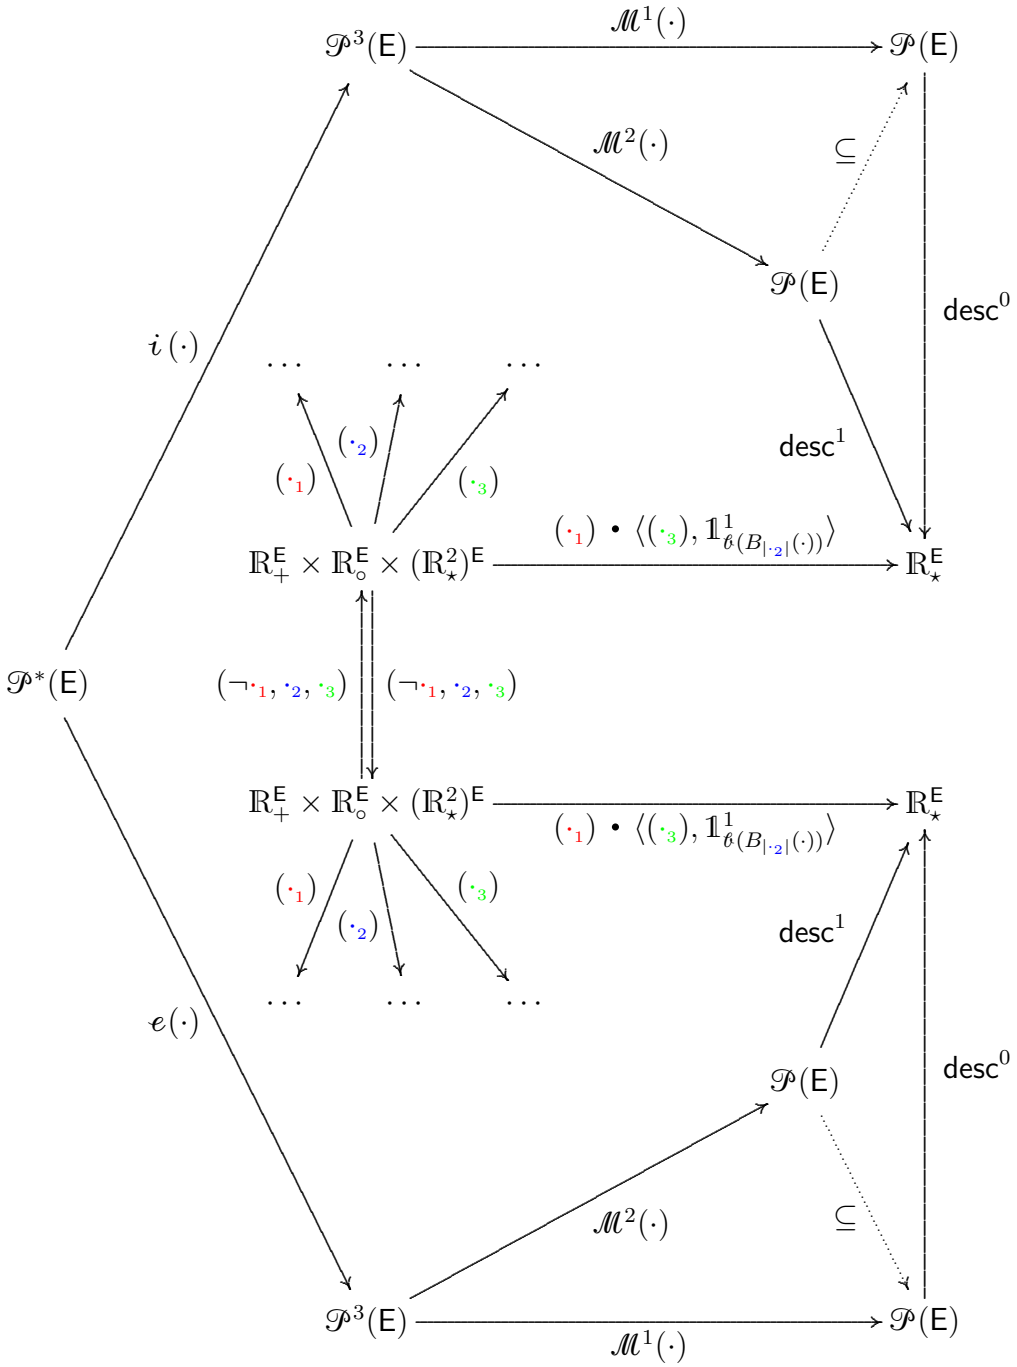

(3.5.48)

In obtaining (3.5.45) and (3.5.46), it should be noted that the inner product is zero on the boundary (since $r = 0$ when $\mathbf{x} \in \partial(S)$), thus $\mathbf{1}_S \equiv \mathbf{1}_{i(S)}$ and $\neg\mathbf{1}_S = \mathbf{1}_{e(S)} \equiv \mathbf{1}_{e(S)}$ are used instead of $\mathbf{1}_{i(S)}$ and $\mathbf{1}_{e(S)}$, respectively, to separate the internal and external 1−MAs, without worrying about the differences on the boundary.

Unfortunately, the above functions diminish over the most part of the MA(T) characterized by 0−MA(T) with more than one (but finite or countable) NN-set, i.e., when $\mu^0[C] = |C| \geq 2$ but $\mu^1[C] = \mu^2[C] = 0$ thus zero inner product in (3.5.44). Characterizing 0−MA(T) implicitly turns out to be much harder, requiring an appeal to more advanced δ−sifting properties.

I conjecture that the generalization of the δ−sifting property (Corollary 2.3.22) allows one to use the δ−function to count (i.e., 0−measure) the roots of a given C^1−continuous real-valued function $f : \Omega \to \mathbb{R}$ over a piecewise 2−manifold (i.e., surface) $\Omega \in \mathscr{P}^2(\mathsf{E})$ embedded in the 3−space by integrating the composition $(\delta^1 \circ f) : \Omega \to \mathbb{R}^1_\star$ over the surface. In particular, I conjecture that under certain conditions, (2.3.153) of Section 2.3.3.2 can be extended to:

$$\int_\Omega (\delta^1 \circ f)(\mathbf{x}')\ d\mu^2[\mathbf{x}'] \stackrel{?}{=} \int_{Z\cap\Omega} \frac{1}{\|\nabla_\Omega f(\mathbf{x}')\|_2}\ d\mu^1[\mathbf{x}'], \tag{3.5.49}$$

where the restricted root set $(Z \cap \Omega) = (f^{-1}(0) \cap \Omega)$ is assumed to be a curve (i.e., the 0−isolevel curve on Ω) and $\nabla_\Omega f(\mathbf{x}')$ is the gradient over the surface, assumed to be well-defined and nonzero for all $\mathbf{x}' \in (Z \cap \Omega)$ along the curve. In other words, if $\Omega := \bigcup_{0\leq i<n} \Gamma(\Delta^2)$ is a finite union of embedded 2−strata:

$$\sum_{0\leq i<n}\iint_{\Gamma_i(\Delta^2)}(\delta^1\circ f\circ\Gamma)(t_1,t_2)|\mathfrak{J}(t_1,t_2)|\;dt_1dt_2 \tag{3.5.50}$$

$$\stackrel{?}{=}\sum_{0\leq i<n}\int_{Z\cap\Gamma_i(\Delta^2)}|\mathfrak{J}(t_1,t_2)|\left(\frac{(dt_1)^2+(dt_2)^2}{(\frac{\partial f}{\partial t_1})^2+(\frac{\partial f}{\partial t_2})^2}\right)^{\frac{1}{2}}. \tag{3.5.51}$$

Now, if we use the function $f(\mathbf{x}') := \zeta_S(\mathbf{x},\mathbf{x}') = \|\mathbf{x}-\mathbf{x}'\|_2 - r$ for a fixed $r>0$, the root set is obtain by solving $f(\mathbf{x}')=0$, i.e., $Z = f^{-1}(0) = \partial(B_r(\mathbf{x}))$ which is an $r-$sphere centered at the query point $\mathbf{x}\in\mathsf{E}$. Then $(Z\cap\Omega)$ for $\Omega := \partial(S)$ is a subset of the $r-$sphere that is obtained by intersecting it with the boundary of the solid $S\in\mathscr{P}^*(\mathsf{E})$. Therefore, the integral in (3.5.49) measures the curve length of this intersection. Now if $r > \mathrm{d}_S(\mathbf{x})$ for a fixed query point $\mathbf{x}\in\mathsf{E}$, i.e., $r := (r_0+\epsilon)$ with $r_0 := \mathrm{d}_S(\mathbf{x})$ for some $\epsilon>0$, the $r-$sphere is guaranteed to intersect the boundary over a curve of nonzero length. As $\epsilon\to 0^+$, one has $r\to r_0 = |\mathrm{d}_S(\mathbf{x})|$ and the set $(Z\cap\Omega) = (Z\cap\partial(B_r(\mathbf{x})))$ approaches the NN-set $C(\partial(S),\mathbf{x}) = (Z\cap\partial(B_{r_0}(\mathbf{x})))$.

Now, assuming that the NN-set does not contain 'sharp' or 'dull' corners of the piecewise $2-$manifold and that the transverse intersection curves in $(Z\cap\Omega)$ shrink to tangent points in the most general conditions, the $1-$measurable $0-$isolevel curve approaches a $0-$measurable (i.e., countable) number of points at which the gradient approaches zero (due to tangency). I conjecture that under these conditions, the integrals in (3.5.49) or (3.5.51) give a $\delta-$singular descriptor of the MA, whose coefficient counts the NN-set. More research is required to make these claims precise and prove them for implicit computation of the MA(T) for objects of arbitrary shape.

Chapter 4

Discretization Schemes

4.1 Representing Functions

In Chapter 3, I showed that implicit descriptions and analytic methods offer tremendous benefits over explicit descriptions and combinatorial methods in uniformly formulating and efficiently computing fundamental geometric modeling and spatial reasoning operations. These operations, in turn, are central to solving important problems in robot path planning [210], mechanism workspace design [235], virtual reality (graphics/haptics) [35,37], protein docking [25], packaging and nesting [64], and more, some of which will be exemplified in Chapter 5.

On several occasions I emphasized that implicit models are to be viewed as the *ability to respond* to a set of basic queries about shapes and motions, viewing them as 'black boxes' with clear input/output semantics. One question that remains, is how are we going to *represent* functions for implicit models and how will the operations defined to operate on functions be *implemented* on a computer?

Let us take a look at the explicit paradigm illustrated in (1.4.1), repeated below:

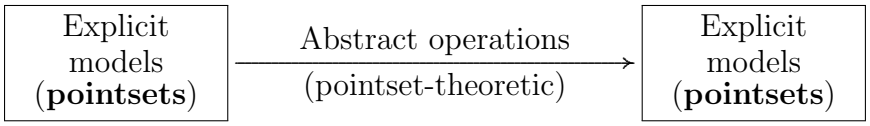

(4.1.1)

Numerous representation schemes (e.g. CSG, B-rep, mesh, voxel grid, etc.) have been developed to store the information about explicit descriptions of continuum shapes and motions into finite symbol structures. These schemes come with representation-specific algorithms that leverage their combinatorial structure to implement geometric operations. Thus models and operations in (4.1.1) are represented via:

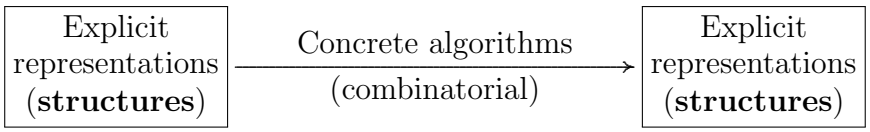

(4.1.2)

The syntactic and semantic *validity* of such representations—when interpreted as mathematical models—and the *correctness* of the algorithms—when interpreted as abstract operations—according to the rules of a representation scheme requires:

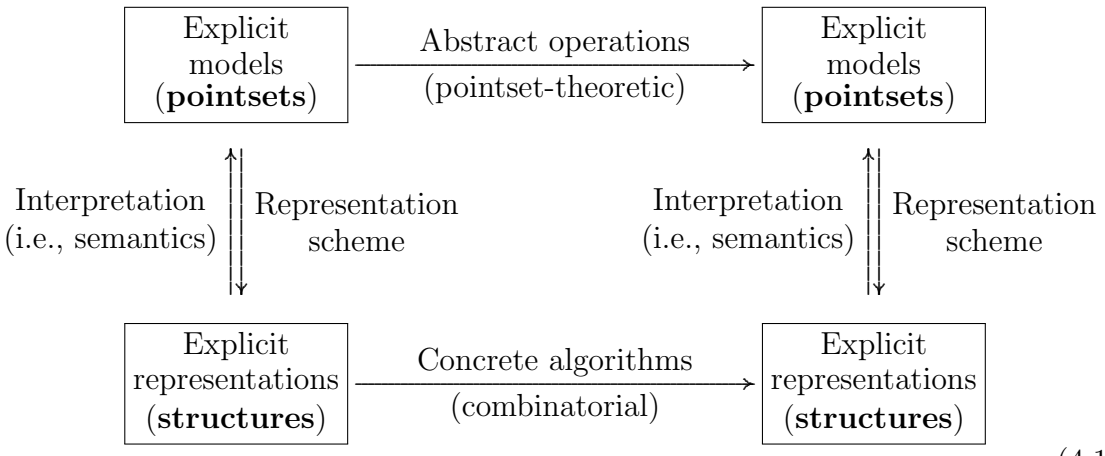

(4.1.3)

Similarly, consider the implicit paradigm illustrated in (1.4.1), repeated below:

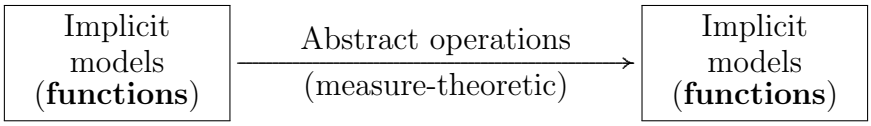

(4.1.4)

for which we would like to develop representations and algorithms. Examples of implicit representations are 'sampling' schemes in physical or frequency domain using finite-element (FE) (e.g., voxel- or mesh-based discretization), radial-basis functions, Fourier basis functions, etc. Examples of analytic algorithms that operate on these samples are finite sequences of operations that implement discretized integration (e.g., Runge Kutta methods), Discrete Fourier transforms (DFT), etc.

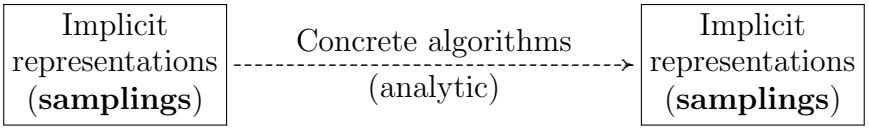

(4.1.5)

Once again, the validity and correctness amounts to demonstrating that the following diagram commutes for a given representation and algorithmic scheme:

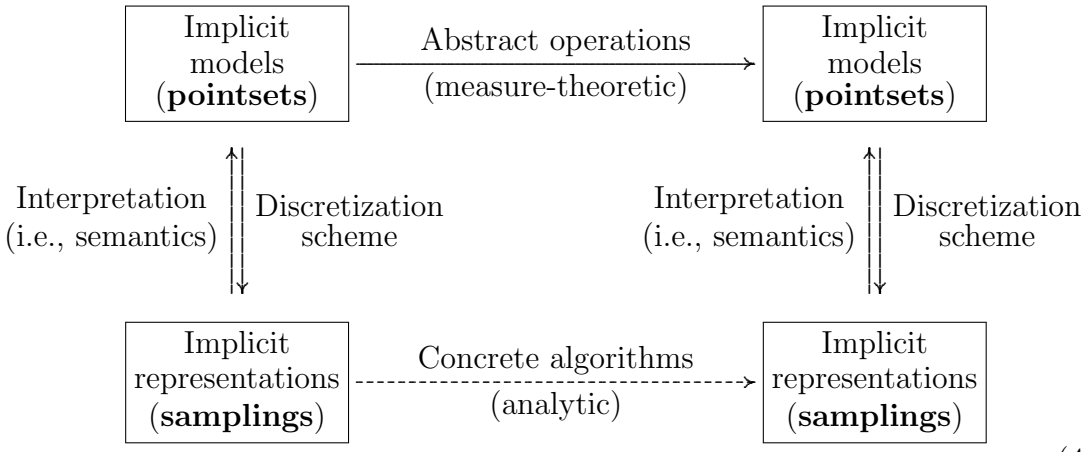

(4.1.6)

Developing proper representations that support intuitive, effective, and efficient (in

addition to, obviously, *correct*) algorithms for the dashed horizontal arrow in the above diagram is the subject of interest in this chapter. The two commutative diagrams in (4.1.5) and (4.1.8) can be viewed as explicit and implicit 'representation theorems', whose composition via (1.4.3) yields:

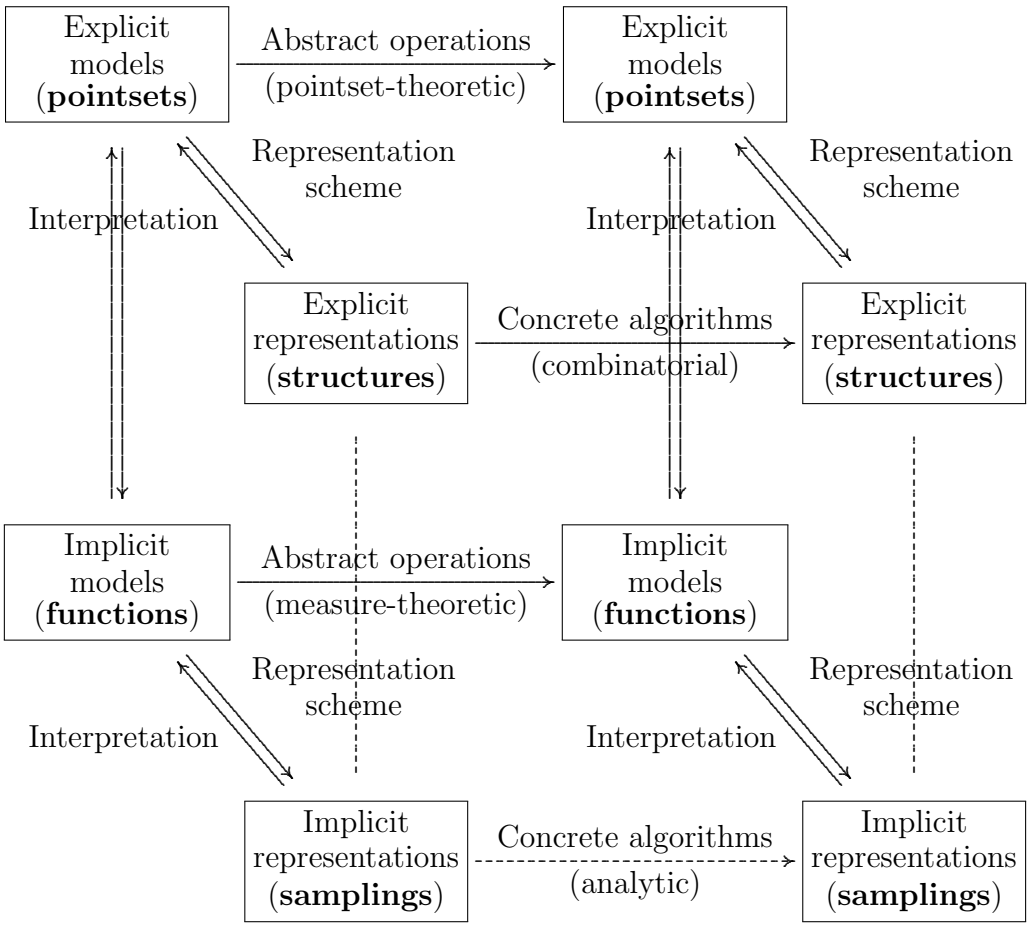

(4.1.7)

in which the solid vertical arrow ↑↓ on the back stand for reconstruction (i.e., fitting) and characterization (i.e., querying), respectively, as in (1.4.3). The dashed vertical lines on the front represent the relationship between the explicit and implicit discretizations, which are not of interest. What is important here is that the concrete combinatorial algorithms that operate on explicit enumerative geometric

representations can be replaced with concrete analytic algorithms that operate on finite samplings of implicit functional descriptions:[1]

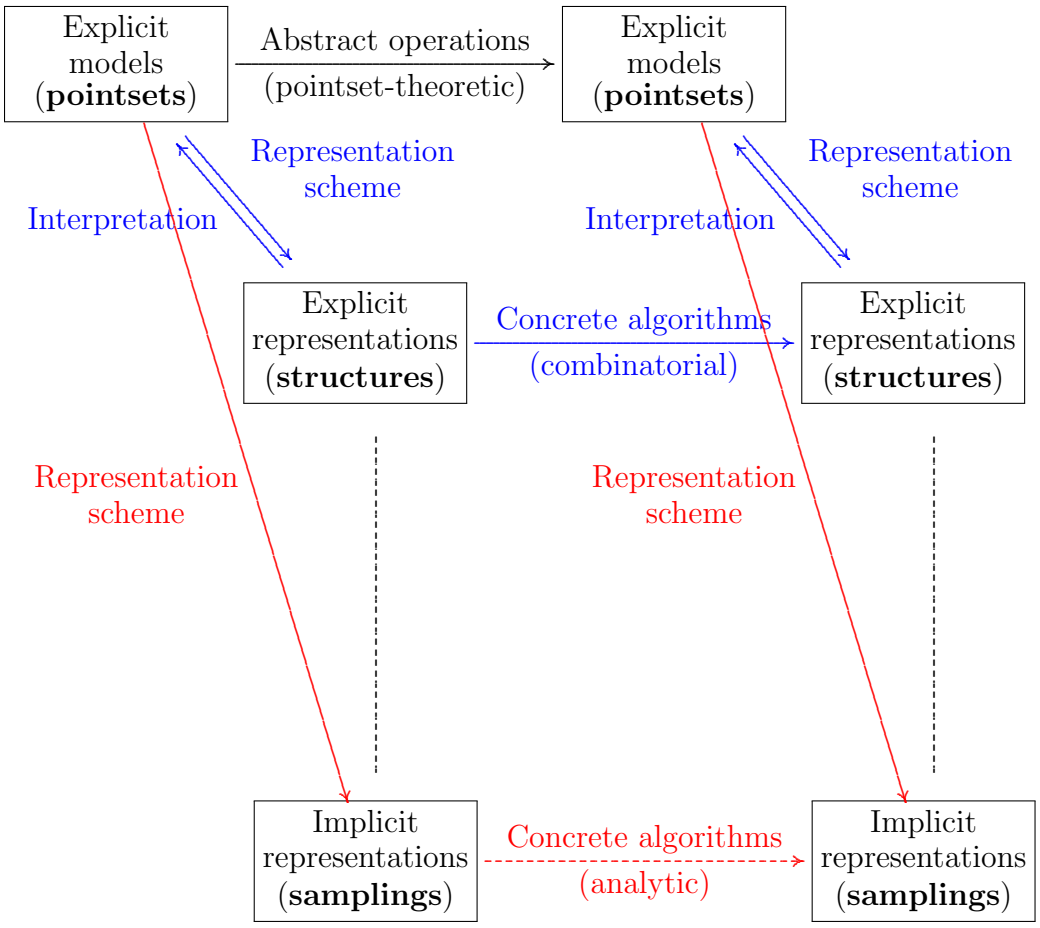

(4.1.8)

The goal is to take the implicit path (colored red) to re-implement the explicit path (colored blue) more effectively and efficiently. For example, the classical implementation of Minkowski operations—and derived morphological operations discussed at length in Section 3.4.1—by applying trimming and merging operations on explicit B-reps of polyhedra by their vertices, edges, and faces can be replaced by an alter-

[1]Note that although it is possible to represent functions via similar explicit enumerative representations—e.g., looking at $f : \mathsf{E} \to \mathbb{R}_+$ as a 4D surface and using NURBS or mesh to represent it—it goes against the implicit, black box view presented earlier.

native implementation using convolution integrals, discretized into convolution sums, and computed using DFT. The DFT can in turn be implemented efficiently using the well-known fast Fourier transforms (FFT), if the sampling is over a uniform grid.

To the best of my knowledge, the existing convolution-based analytic methods in solid and configuration modeling rely on uniform sampling over equispaced grids—with the exception of [236] that uses adaptive sampling for inverse problems, and [25] that is restricted to molecular models as unions of balls of fixed radius. In what follows, I present a nonuniform discretization scheme[2] that relies on progressive spherical approximations with balls of different sizes. I show that the spherical discretization scheme offers an algebraic structure that is closed under Minkowski sum and product (but not difference and quotients) operations, and at the same time offers more appealing properties than uniform grid- or octree-based discretizations. As the continuous geometry (of both shapes and motions) is abstracted away by the balls, the computational implementation solely relies on convolution algebra over discrete sets specified completely by ball center coordinates and radii, allowing the use of the efficient nonequispaced FFT (NFFT) algorithm [258] on the massively parallel graphics processing units (GPU) architectures [193].

Unlike combinatorial methods whose complexities typically depend on the syntactic size of the representation (e.g., polygon count) fixed upfront, this method allows for a choice of complexity in real-time based on the affordable resources, by proceeding deep enough down the sphere-tree—which can be constructed using any algorithm of choice, such as [48, 156, 248, 362] or my own presented in Section 4.6. On the other

[2]What I mean by a 'discretization scheme' is not a particular decomposition algorithm or approximation method, but a generic formalism for reconciling such a nonuniform discretization (in contrast to the extensively used uniform sampling) to analytic modeling, using Minkowski sums and convolutions. I do present one new algorithm in Section 4.6; nevertheless, other methods [48, 156, 248, 362] are also applicable under the same scheme.

hand, unlike grid-based analytic methods whose arithmetic complexities scale with object size and grid resolution, this method enables filling large regions with large balls and efficiently allocating more primitives to capture features of smaller size with higher fidelity. Finally, by working in the Fourier domain, aside from converting convolutions to simple pointwise multiplications, preserving inner products and linear maps (e.g., rotations and reflections), and converting translations to multipliers, the method allows for 'graceful' degradation of the accuracy by truncating the frequency domain representations, enabling another trade-off mechanism between running time and numerical precision.

After providing the basic motivation in Section 4.2, I present the formalism for discretization of an arbitrary shape as a countable union of balls, its interpretation as a Minkowski operation, and its analytic description as a convolution in Section 4.3. In Section 4.4 I formulate correlation predicates in terms of Minkowski operation across multiple shapes relatively positioned and oriented in arbitrary poses together with their convolution forms, and use the results from the previous section to carry the discretization into the configuration space. In Section 4.5 I demonstrate how efficient CPU and GPU implementations of this method outperform the state-of-the-art. In Section 4.6, a particular sampling algorithm is presented and compared with a popular existing algorithm for ball fitting in terms of storage complexity, though the general approach is applicable regardless of the choice of sampling algorithm.

4.2 Discretization Techniques

One of the advantages of the analytic framework is that the simple algebraic structure presented in Section 3.1 provides a natural route to the choice of representation and sampling schemes. Recall that the basic algebraic operations—from which all other operations were derived—were limited to

- *arithmetics* of real-valued functions, including extended-real-valued defining functions (subsuming indicator functions) and distance functions;
- *convolutions* including convolutions among defining functions themselves and between defining functions and $\delta-$function (i.e., $\delta-$regularization); and
- *compositions* with the $\delta-$function (i.e., $\delta-$sifting).

Arithmetic operations are *local* computations that impose little burden on the representation scheme. For example, An inclusion query on the Boolean constructions such as $\mathbf{x} \stackrel{?}{\in} (S_1 \cup S_2)$ or $\mathbf{x} \stackrel{?}{\in} (S_1 \cap S_2)$ which amounts to computing the values of $(\mathbf{1}_{S_1} + \mathbf{1}_{S_2})(\mathbf{x})$ or $(\mathbf{1}_{S_1} \cdot \mathbf{1}_{S_2})(\mathbf{x})$, respectively, are $O(1)-$time operations; namely, query $(\mathbf{x} \stackrel{?}{\in} S_1)$ and $(\mathbf{x} \stackrel{?}{\in} S_2)$ and instantly compute the arithmetic sum/product of the outputs without a need for sampling any point(s) other than $\mathbf{x} \in \mathsf{E}$ itself.

Convolutions, on the other hand, are *global* computations that require integrating over the entire space—and so do inner products and L^1-norms which are instantiations of convolutions. For example, an inclusion query on the Minkowski constructions such as $\mathbf{x} \stackrel{?}{\in} (S_1 \oplus S_2)$ or $\mathbf{x} \stackrel{?}{\in} (S_1 \ominus S_2)$ which amounts to computing the values of $(\mathbf{1}_{S_1} * \mathbf{1}_{S_2})(\mathbf{x})$ or $(\mathbf{1}_{S_1} \diamond \mathbf{1}_{S_2})(\mathbf{x})$, respectively, are $O(m)-$time operations, where m is the sample size of the discrete enumeration of the convolution integrals, which could be sampled uniformly (e.g., Riemann sum approximation over a grid) or carefully

selected nonuniformly (e.g., adaptive octree sampling or Gauss points). Thus it takes $O(m)$ computations for a single query on a convolution at $\mathbf{x} \in \mathsf{E}$.

Now, if the convolution is to be used in downstream operations, it has to be sampled at many such query points—ideally at m query points if it is to be used in a second convolution—or any other integration-based operation such as inner products or L^1-norms—which takes $O(m^2)$.

4.2.1 FFT-Based Convolution

Harmonic analysis provides remedies to the complexity problems in several different ways, and provides access to digital signal processing (DSP) techniques that are the reasons to popularize analytic methods in the first place:

First, the convolution theorem (Theorem 2.2.11) allows one to convert the convolution (i.e., a *global* computation) in the physical domain to a pointwise multiplication (i.e., a *local* computation) in the frequency domain. Although the forward and inverse Fourier transforms are integrals themselves (thus global computations), they can be performed in $O(m \log m)$ time using the radix-2 fast Fourier transform [76], in contrast to the $O(m^2)-$time cascade method, *only if uniform sampling is used*, in both physical and frequency domains. Thus convolutions can be computed cumulatively (i.e., at one shot) for m query points after two $O(m \log m)-$time forward FFTs, an $O(m)-$time pointwise multiplication, and another $O(m \log m)-$time inverse FFT, yielding a total asymptotic time complexity of $O(m \log m)$ instead of $O(m^2)$.[3]

If a single query on the convolution (i.e., one inner product) is enough—e.g.,

[3] As I show in Section 4.5, the $O(m)-$time pointwise multiplication can take longer than the $O(m \log m)-$time FFTs, simply because FFTs have been optimized in standard CPU- and GPU-based implementations due to their wide applications in various fields, and the $\log m$ factor is almost as small as $O(1)$ for moderate resolutions.

in real-time collision detection, querying the $\mathsf{C}-$obstracle for one configuration per frame—Parseval's theorem (Theorem 2.2.10) implies that the complexity of computing inner products in both physical and configuration spaces is asymptotically the same (namely, $O(m)$). Although it is simpler to compute inner products over real-valued and compactly supported shape descriptors in the physical domain in contrast to their complex-valued and band-unlimited transforms in the frequency domain, there are other advantages to working in the Fourier space, described below.

First of all, forward FFTs can be computed once for every shape descriptor, and reused many times as the shape goes through translational and/or rotational motions, thanks to the straightforward relations in (2.2.16) through (2.2.20) of Section **??**. Moreover, if multiple *consecutive* convolutions, inner products, and L^1-norms are to be computed in the physical domain—which correspond to multiplications, inner products, and L^1-norms in the frequency domain, respectively—one can stay in the Fourier space without applying inverse FFTs at intermediate states, of course until other operations are called, such as arithmetics for Boolean operations in the physical space, which are harder to compute over Fourier representations.[4]

The computational advantage of FFT-based convolutions have made them popular in many related applications such as path planning [88, 174] and protein docking [60, 61]. However, the uniform sampling requirement forces a non-optimal strategy for capturing the shape details, wasting numerous sample points in the interior of the object while most of the geometric details are near the boundary.[5]

[4]For example, implicit Boolean intersection is a pointwise multiplication in the physical domain (i.e., a local computation) but converts to a convolution in the frequency domain (i.e., a global computation).

[5]For sufficiently 'globular' objects, it is expected to have only $O(m^{\frac{2}{3}})$ sample points next to the boundary and as many as $O(m)$ sample points in the interior, which do not provide any additional information if the objects are assumed to be structurally homogeneous.

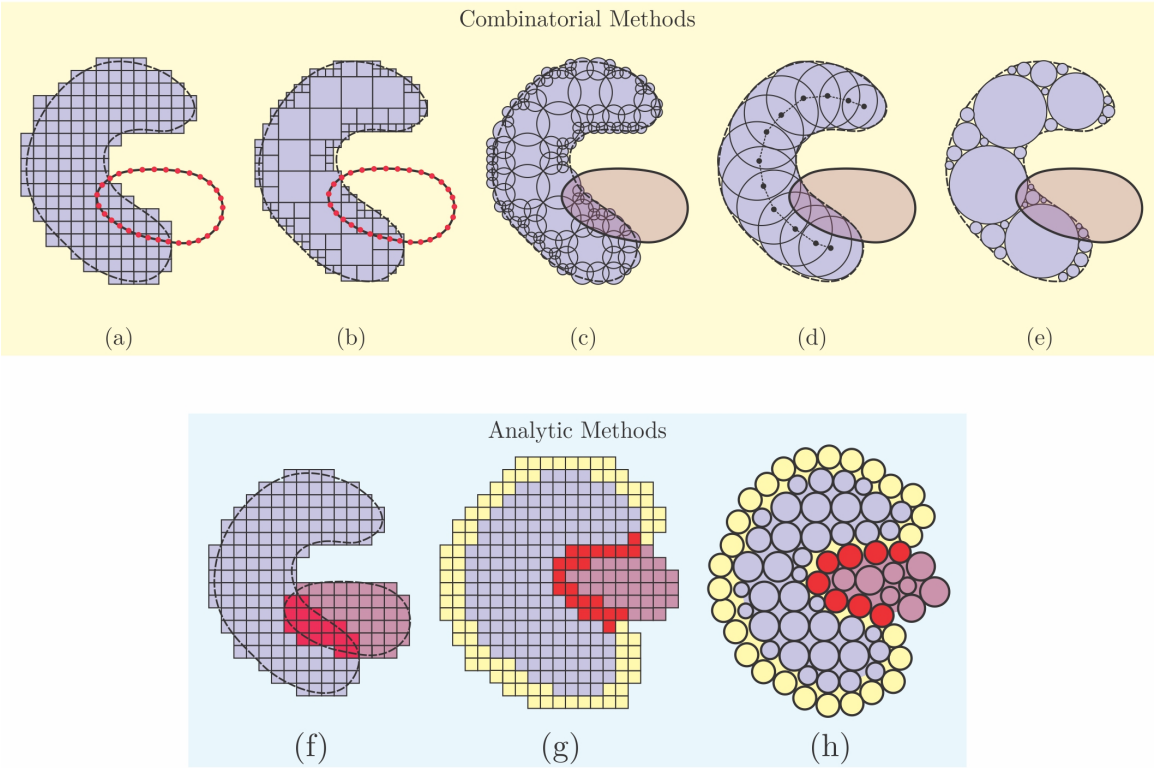

FIGURE 4.2.1: Research in combinatorial methods has shown that CD tests over uniform grids (i.e., 'voxmaps') [304] (a) and octrees (special case of OBB-trees) [133] (b) can be made more efficient if voxels are replaced with spherical primitives, e.g., built around octrees [248] (c), sampled over the MA [48, 156] (d) or packed inside using distance fields [362] (e). However, the more nascent analytic methods are still mostly reliant on uniform grids for, e.g., CD testing for solids by integrating the intersection [174] (f), and SC scoring for proteins by integrating skin overlaps [61] (g). The latter has been outperformed by grid-free correlations of atoms grouped with equal radii [25] (h). I show that constructions in (c–e) with arbitrary radii can also be interpreted analytically as a convolution and solved by nonuniform FFTs after a geometric lifting.

4.2.2 Spherical Decomposition

Taking collision detection (CD) as a representative example, at one end of the research spectrum are combinatorial techniques that use surface meshes or higher-order algebraic parametrizations to resolve collisions or identify matching features. Examples are polyhedral CD methods based on Voronoi-clipping/marching [108, 109, 227] and oriented bounding box (OBB) trees [72, 133, 157], or spatial enumeration-based techniques such as the Voxmap PointShell (VPS)™ [222–224, 304, 355]. VPS™ works by a pairwise test between a shell of vertices for the moving object against a map of voxels that discretizes the stationary obstacle, and is popular in physically-based modeling in virtual environments [185]. Others have identified more efficient techniques for time-critical CD by using hierarchical bounding spheres sampled on octrees [248] or on the medial axes (MA) [48, 154–156], and interior sphere packing guided by distance fields [360–362] (Fig. 4.2.1 (a–e)). These 'sphere-tree' based methods have been shown to outperform voxmap or OBB-tree based techniques in real-time applications [361, 362]—particularly because primitive collision predicates are simplified to center-distance tests as a result of the orientational symmetry of balls—and are considered state-of-the-art in practice. For comprehensive surveys on CD methods, I refer the reader to [169, 188, 208, 339].

Taking shape complementarity (SC) analysis as another representative example, on the other end of the research spectrum are analytic methods that have been more popular in robotics [210]. Unlike the combinatorial approach that searches for a collision certificate point (or lack thereof) in the intersection of the objects, the analytic approach treats the collision predicate as a Boolean combination of inequalities over some configuration space of the objects [214]. The obstacle avoidance in path

planning is, for example, treated as an optimization problem subjected to holonomic collision constraints formulated analytically as a convolution of the robot and its workspace [88, 174]. Most convolution-based methods have so far focused on generating uniformly sampled configuration bitmaps for all spatial positions and orientations simultaneously, which can be cumulatively computed with asymptotically optimal FFTs [76]. However, a complete description of the configuration obstacle is overkill for real-time CD. A recent work [214], also reliant on uniform grid-based sampling, formally reframes the approach for time-critical CD (Fig. 4.2.1 (f)), but has not yet been compared with sphere-tree methods, nor applied to real-time applications.

In an independent line of research, numerous analytic methods for molecular surface analysis and SC-based protein docking have been developed, whose outcomes are platforms that use grid-based occupancy enumeration and leverage classical FFTs [76] such as ZDock [61], or more recent grid-free techniques that rely on nonuniform FFTs [258], such as F^2Dock [25] (Fig. 4.2.1 (g, h)). The latter exploits the spherical shape of the atomic building blocks and implicitly represents the proteins as summations of radial kernels centered around atoms, assigning different weights to core and skin atoms. The SC score is obtained by cross-correlating these functions from different proteins, which turns into a convolution discretized over the center points. It has been shown that grid-free methods outperform uniform grid-based methods [25] by taking advantage of the spherical geometry. For comprehensive surveys on advances in protein docking, see [47, 288, 338, 346].

The Best of Both Worlds. Although objects of arbitrary shape, unlike molecules, cannot be represented *exactly* as finite unions of balls, the sphere-tree methods for time-critical CD were shown to be more successful in progressively approximating the shape, when compared to uniform grid- or octree-based voxelization, with a faster convergence and a better use of computational resources [156]. Motivated by this observation, I present a generic framework for representing arbitrary shapes with finite (or countably infinite, in the limit) radial kernels, formulated as a convolution of a discrete pointset and the primitive kernel in a higher-dimensional space. The latter is described as a *geometric lifting* trick in Section 4.3, and is deemed necessary due to the inevitable size difference between primitive balls, unlike the simpler case for the proteins. We show that this approach offers 'the best of both worlds' by combining the computational efficiency of the sphere-tree techniques for time-critical applications (i.e., with a single configuration query) with that of the analytic methods for cumulative configuration space constructions (i.e., requiring a complete map for all spatial relations), unified under a single paradigm with analytic formalism that applies to a multitude of applications.

4.3 Radial Basis Functions

It is rarely the case in practical applications with arbitrarily complex shapes to have the defining function in closed form. However, one can always decompose the solid into a finite grouping of simpler primitives, as an immediate consequence of its finite describability postulate given in Section 2.2.1, and apply a finite sequence of set-theoretic union operations to combine the primitive defining functions—to which I

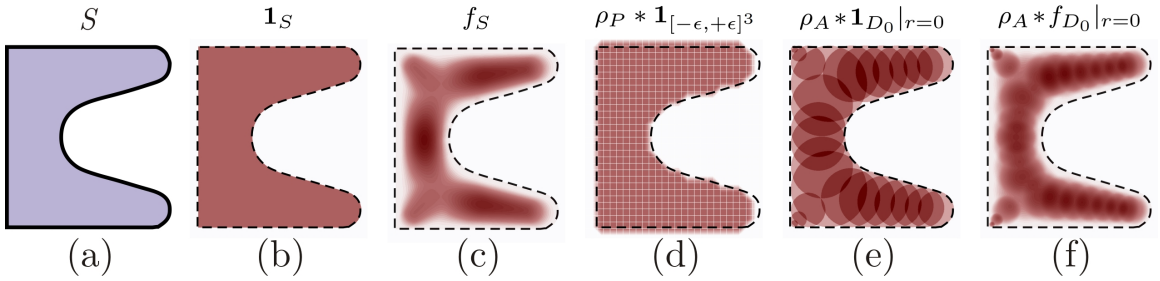

FIGURE 4.3.1: An r-set (a), its indicator function (b) and bump function (c), discretized using grid-based bitmap sampling (d), and grid-free spherical sampling corresponding to $\alpha \to \infty$ (e) and $\alpha := 2$ (f) in (4.3.9).

refer as 'discretization'. However, obtaining exact discretizations (e.g., curvilinear cell decompositions) from the popular CSG or B-rep schemes [278] is not trivial. An alternative is to use approximate discretizations (e.g., spatial enumerations via uniform grids or octrees) which converge to the r-set in the limit. Next, I introduce a more general discretization scheme that subsumes these enumeration methods with non-intersecting cubic primitives (i.e., voxelization) as special cases, and enables other types of (possibly intersecting) primitives.

4.3.1 Spherical Expansions

Consider the case when an r-set $S \in \mathscr{P}^*(\mathsf{E})$ can be decomposed as a finite grouping $S = \bigcup_{0 \le i < n} B_i$ where $B_i \in \mathscr{P}^*(\mathsf{E})$ for all $0 \le i < n$, hence $f_S(\mathbf{x}) = \sum_{0 \le i < n} f_{B_i}(\mathbf{x})$ for all $\mathbf{x} \in \mathsf{E}$. The following two cases, called 'equiradius' and 'inequiradius' decompositions are of prime significance for the purposes of this chapter:

Equiradius Decomposition. First, let B_i be relocated instances of the same base shape $B_0 \in \mathscr{P}^*(\mathsf{E})$ instantiated by the different translations $\mathbf{x} \mapsto (\mathbf{x} + \mathbf{x}_i)$:

$$B_i = B_0 \oplus \{\mathbf{x}_i\} = \big\{(\mathbf{x} + \mathbf{x}_i) \mid \mathbf{x} \in B_0\big\}, \quad 0 \leq i < n. \tag{4.3.1}$$

Each such translation corresponds to a 3D point $\mathbf{x}_i \in \mathbb{R}^3$, hence the discrete pointset $P := \{\mathbf{x}_i\}_{0 \leq i < n} \subset \mathbb{R}^3$ (of cardinality $|P| = n$) and the base primitive B_0 contain all the information to reconstruct the solid. We can use the notation $S = S_n(P)$, as if $S_n : \mathscr{P}^0(\mathbb{R}^3) \to \mathscr{P}(\mathbb{R}^3)$ is a mapping from the discrete space of n points to the shape space, illustrated in Fig. 4.3.2. A crucial observation is that this mapping can be viewed as a Minkowski sum $S_n(P) = (P \oplus B_0)$. If the original shape cannot be exactly decomposed into a finite union of transformed B_0 instances, then $S \approx S_n(P)$ such that $S = \lim_{n\to\infty} S_n(P) = \lim_{|P|\to\infty}(P \oplus B_0)$ in the Hausdorff topology:

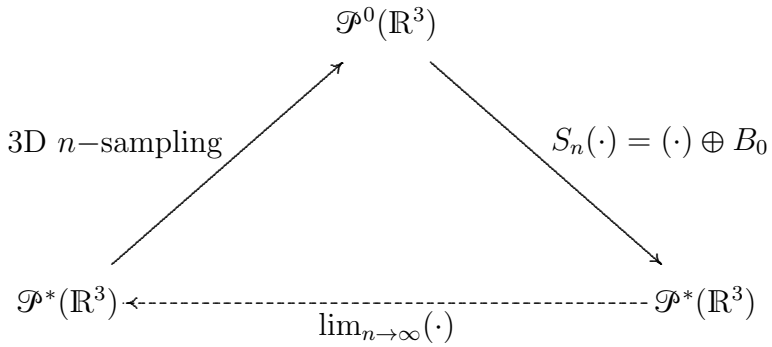

(4.3.2)

in which the particular sampling algorithm/strategy $\mathscr{P}(\mathbb{R}^3) \to \mathscr{P}^0(\mathbb{R}^3)$ is not of interest for now—see Section 4.6 for a particular example.

To reconcile with the analytic model, let $B_0 = U_0^*(f_{B_0})$ hence a defining function

for each primitive instance can be obtained as

$$f_{B_i}(\mathbf{x}) := c_i f_{\mathbb{t}_i B_0}(\mathbf{x}) = c_i(f_{B_0} \circ \mathsf{act}(\mathbb{t}_i^{-1}))(\mathbf{x}) = c_i f_{B_0}(\mathbf{x} - \mathbf{x}_i), \tag{4.3.3}$$

as a result of the covariance of defining functions under translations. The weight coefficients $c_i > 0$ are arbitrarily assigned (for now). Thus the Boolean union (i.e., summation, in the implicit domain) of the balls takes the form

$$f_{S_n(P)}(\mathbf{x}) = \sum_{0 \le i < n} f_{B_i}(\mathbf{x}) = \sum_{0 \le i < n} c_i f_{B_0}(\mathbf{x} - \mathbf{x}_i), \tag{4.3.4}$$

which can be viewed as a discrete convolution. To make it compatible with continuous convolutions, let us rewrite (4.3.4) as an integral

$$f_{S_n(P)}(\mathbf{x}) = \sum_{0 \le i < n} \int_{\mathsf{E}} \delta^3(\mathbf{x}' - \mathbf{x}_i) \big[c_i f_{B_0}(\mathbf{x} - \mathbf{x}')\big] d\mu^3[\mathbf{x}'],$$

In accordance with the concepts presented in Section 2.3.2, if we assume a $3^{\text{rd}}-$order $\delta-$singular density function of the form

$$\mathbb{1}_P^3(\mathbf{x}) := \sum_{0 \le i < n} c_i \varsigma_{\mathbf{x}_i}(\mathbf{x}) = \sum_{0 \le i < n} c_i \delta^3(\mathbf{x} - \mathbf{x}_i), \tag{4.3.5}$$

where $\varsigma_{\mathbf{x}_i}(\mathbf{x}) = \delta^3(\mathbf{x} - \mathbf{x}_i)$, to represent the discrete pointset P as a collection of spatial impulses of intensities c_i at each $\mathbf{x}_i \in P$—i.e., each point carrying a 'lumped' artificial volume of c_i—the Minkowski sum can be analytically expressed as a convolution:

$$f_{S_n(P)} = (\mathbb{1}_P^3 * f_{B_0}), \quad \text{i.e.,} \quad (P \oplus B_0) = U_0^*(\mathbb{1}_P^3 * f_{B_0}). \tag{4.3.6}$$

Once again, if we have an approximate discretization scheme $S \approx S_n(P)$ such that $S = \lim_{n\to\infty} S_n(P) = \lim_{|P|\to\infty}(P \oplus B_0)$:

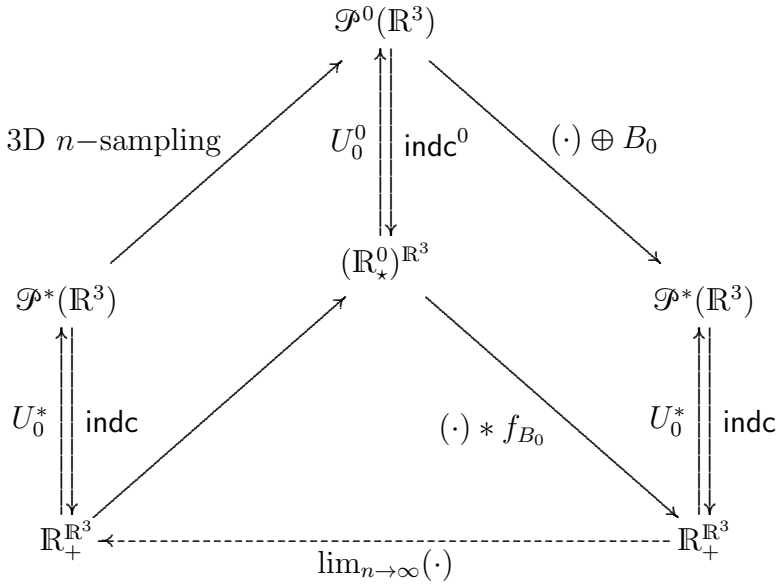

(4.3.7)

A particularly favorable choice for the primitive shape (for reasons to be explained in Section 4.3) is a closed 3−ball, i.e.,

$$B_i := B_r(\mathbf{x}_i) = \big\{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x} - \mathbf{x}_i\|_2 \leq r\big\}, \quad 0 \leq i < n, \tag{4.3.8}$$

are balls of constant radius $r > 0$ centered at $\mathbf{x}_i \in P$.

Let $f_{B_0}(\mathbf{x}) := \psi_\alpha(\frac{1}{r}\|\mathbf{x}\|_2)$ where the function $\psi_\alpha : \mathbb{R} \to [0, 1]$ is a generic cut-off kernel (also referred to as the 'mollifier' or the 'bump') with the closed form

$$\psi_\alpha(x) = \begin{cases} e^{(1-|x|^{-\alpha})^{-1}} & \text{if } |x| < 1, \\ 0 & \text{otherwise,} \end{cases} \tag{4.3.9}$$

which can be thought of as a smoothed extension of the discontinuous cut-off function

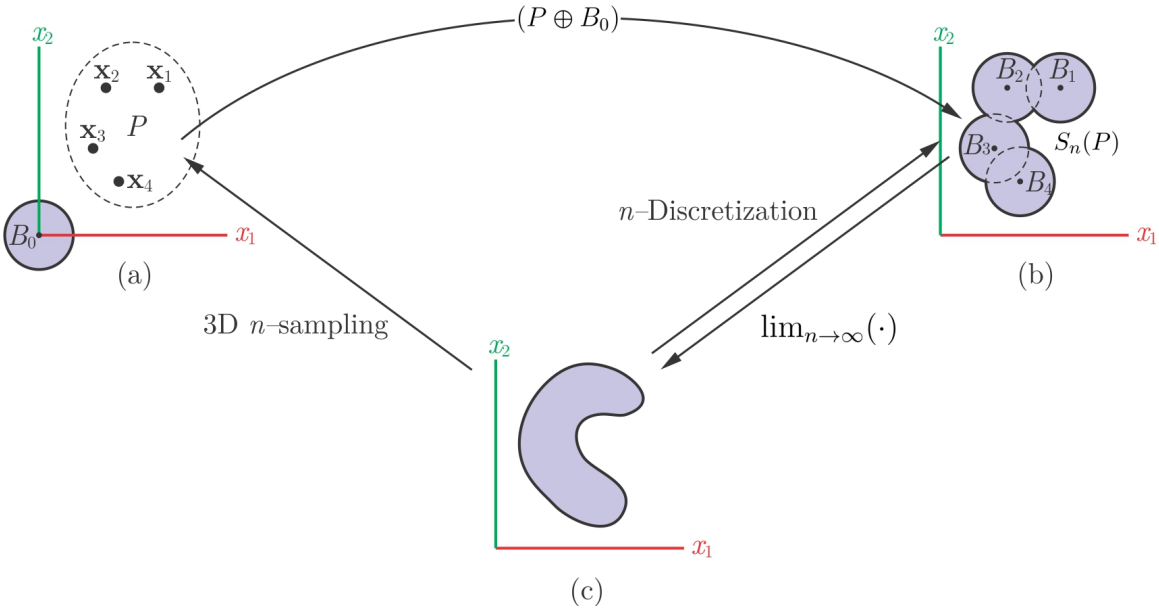

FIGURE 4.3.2: An equiradius sampling of a 3D r-set can be viewed as a 3D Minkowski sum $S_n(P) = (P \oplus B_0) = U_0^*(\mathbb{1}_P^3 * f_{B_0})$. Here, the lifting and slicing are illustrated for 2D r-sets.

$\mathbf{1}_{(-1,+1)}(x) = \lim_{\alpha\to\infty} \psi_\alpha(x)$, and the resulting $f_{B_0}(\mathbf{x}) = \psi_\alpha(\frac{1}{r}\|\mathbf{x}\|_2)$ is a mollified extension of the binary indicator function $\mathbf{1}_{i(B_0)}(\mathbf{x}) = \lim_{\alpha\to\infty} f_{B_0}(\mathbf{x})$.

The spatial enumeration schemes over uniform grids—ranging from bitmap encoding for path planning [174] to rasterized density functions for protein docking [61], illustrated in Fig. 4.2.1 (f, g)—can be viewed as special cases of this scheme with cubic primitives (i.e., $L^\infty-$ instead of L^2-balls) with an additional disjointness condition that is unnecessary for our purposes. Grid-free molecular modeling based on Gaussian densities for protein surface reconstruction [103] and protein docking [25] (Fig. 4.2.1 (h)) are more closely related to the scheme proposed here, as they use spherical primitives.[6]

Generalizing this grid-free discretization to solid objects of arbitrary geometric

[6]Except that these methods use Gaussian kernels rather than compactly supported cut-off kernels such as the one in (4.3.9), and use the $1-$ (instead of the $0-$) isosurface to define the molecular surface, as a matter of convenience.

complexity would enable more efficient use of the computational resources by adaptively approximating the shape, filling large interior regions with fewer primitives, and allocating resources to capture the details of surface features.

A simple solution is to use a recursive decomposition (e.g., an octree) and take the leaf cells (or balls enclosed by their bounding spheres) as primitives (Fig. 4.2.1 (c)), collected into groups of constant radii according to their level in the tree, for (4.3.6) to apply. Hubbard [154–156] showed that octree-based spherical approximation is non-optimal in terms of convergence in the Hausdorff distance, and compares poorly with sampling the centers of primitive balls over the MA (Fig. 4.2.1 (d)). However, the latter requires a generalization of (4.3.6) that supports different sizes for the balls, which I address next.

Nonequiradius Decomposition. The generalization is enabled by a simple *geometric lifting* trick. This time, let B_i be translated and scaled instances of the base shape $B_0 \in \mathscr{P}^*(\mathsf{E})$, i.e., instantiated by the affine transformation $\mathbf{x} \mapsto \mathbf{x}_i + (r_i\mathbf{x})$:

$$B_i = (r_i B_0) \oplus \{\mathbf{x}_i\} = \big\{((r_i\mathbf{x}) + \mathbf{x}_i) \mid \mathbf{x} \in B_0\big\}, \quad 0 \le i < n. \tag{4.3.10}$$

Each such transformation corresponds to a 4D point $\mathbf{a}_i := (\mathbf{x}_i, r_i) \in (\mathsf{E} \times \mathbb{R}_+)$, hence the discrete pointset $A := \{\mathbf{a}_i\}_{0 \le i < n} \subset (\mathsf{E} \times \mathbb{R}_+)$ (of cardinality $|A| = n$) contains all the information to reconstruct the solid, and $S = S_n(A)$ can be viewed as a Minkowski product [291] of the form $S_n(A) = A \otimes \gamma_0(B_0)$, defined over the group of the aforementioned instance transformations $\mathsf{G} \cong \mathbb{R}^4$, where $\gamma_r : \mathbb{R}^3 \hookrightarrow \mathbb{R}^4$, $\gamma_r(\mathbf{x}) = (\mathbf{x}, r)$, and $\gamma_0(B_0) = B_0 \times \{0\}$ is a trivial embedding in G. Roerdink [291] formulated the notion of Minkowski products over general groups (e.g., motion, affine, and projective groups),

a noncummutative convolution interpretation of which was discovered by Lysenko et al. [218] and is related to (4.3.16). This scheme is extensible to allow for rotated, reflected or nonuniformly scaled instances using generalized Minkowski operations [291] along with an analytic expression using noncommutative convolutions [218]. However, to exploit the radial symmetry and rotational invariance properties as well as the extensively studied medial axis (MA) formulated based on the L^2-norm, hereon I restrict the primitive shapes to spherical balls, whose instantiation is completely described by translation and uniform scaling.

A more helpful way of looking at this formulation is to think of each primitive $B_i \subset \mathbb{R}^3$ as a cross-section (i.e., a 3D 'slice' orthogonal to the $r-$axis at $r = 0$) through a hypothetical hypercone $C_i \subset \mathbb{R}^4$ whose apex is located at $(\mathbf{x}_i, r_i) \in A$, illustrated in Fig 4.3.3 (c, d):

$$C_i := C_{r_i}(\mathbf{x}_i) = \{(\mathbf{x}, r) \in \mathbb{R}^4 \mid \|\mathbf{x} - \mathbf{x}_i\|_2 \leq |r - r_i|\}, \quad 0 \leq i < n. \tag{4.3.11}$$

To ensure compactness of the 4D objects, let us replace the unbounded cones C_i with *trimmed* half-cones $D_i \subset \mathbb{R}^4$:

$$D_i := D_{r_i}(\mathbf{x}_i) = C_{r_i}(\mathbf{x}_i) \cap (\mathbb{R}^3 \times [r_i - L, r_i]), \quad 0 \leq i < n, \tag{4.3.12}$$

where $L > \max_{0 \leq i < n} r_i$ to guarantee that all displaced half-cones will intersect the $r = 0$ hyperplane. If we define $K_n := \bigcup_{0 \leq i < n} D_i$, the 3D solid $S_n(A)$ becomes a slice of the 4D solid $K_n(A)$ at $r = 0$, where $K_n : \mathcal{P}^0(\mathbb{R}^4) \to \mathcal{P}(\mathbb{R}^4)$ is the discretization mapping illustrated in Fig. 4.3.3 (a, b). The key observation is that unlike the scaled primitives B_i that have different sizes, their cones C_i and D_i are all translated instances of the

same base shapes C_0 and D_0, respectively, whose apexes are at the origin. Their union can thus be viewed as a Minkowski sum $K_n(A) = (A \oplus D_0)$. The 3D solid is then obtained as a 3D slice $S = S_n(A) = K_n(A)|_{r=0} := \gamma_0^{-1}(K_n(A) \cap \gamma_0(\mathbb{R}^3))$,[7] illustrated in Fig. 4.3.3 (c, d). Once again, if the original shape cannot be exactly decomposed into a finite union of transformed B_0 instances, then $S \approx S_n(A)$ such that $S = \lim_{n\to\infty} S_n(A) = \lim_{|A|\to\infty}(A \oplus D_0)|_{r=0}$ in the Hausdorff topology:

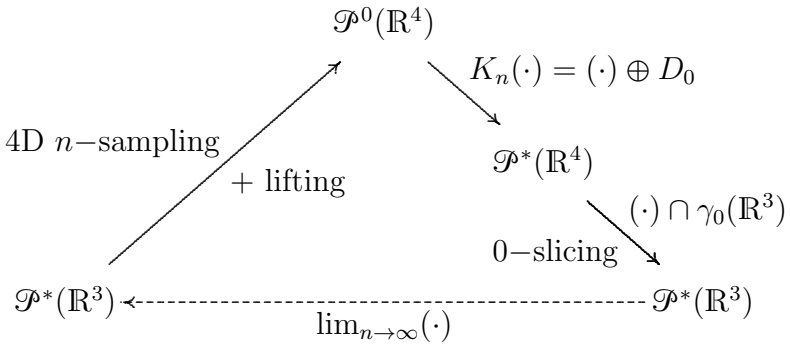

(4.3.13)

in which the particular sampling algorithm/strategy $\mathscr{P}(\mathbb{R}^3) \to \mathscr{P}^0(\mathbb{R}^4)$ is not of interest for now—see Section 4.6 for a particular example.

To obtain an analytic expression for the 4D geometry similar to that of (4.3.6), let $D_0 = U_0^*(f_{D_0})$, hence a defining function for each compact cone is obtained as

$$f_{D_i}(\mathbf{a}) := c_i f_{\mathbb{t}_i D_0}(\mathbf{x}) = c_i(f_{D_0} \circ \mathsf{act}(\mathbb{t}_i^{-1}))(\mathbf{a}) = c_i f_{B_0}(\mathbf{a} - \mathbf{a}_i), \tag{4.3.14}$$

where $\mathbf{a} = (\mathbf{x}, r) \in \mathbb{R}^4$ represents a point in the 4D space where the lifted geometry

[7] For a set $K \subset \mathbb{R}^4$, I use the simplified notation $K|_{r=r_0}$ for its $r = r_0$ slice projected to $\mathbb{R}^3$, i.e., $K|_{r=r_0} := \gamma_{r_0}^{-1}(K \cap \gamma_{r_0}(\mathbb{R}^3))$, where $\gamma_r : \mathbb{R}^3 \hookrightarrow \mathbb{R}^4$ is defined as $\gamma_r(\mathbf{x}) = (\mathbf{x}, r)$ hence $\gamma_r^{-1}(\mathbf{x}, r) = \mathbf{x}$.

resides. If we form an impulsive density function similar to (4.3.5),

$$\mathbb{1}_A^4(\mathbf{a}) := \sum_{0\leq i<n} c_i \varsigma_{\mathbf{a}_i}(\mathbf{a}) = \sum_{0\leq i<n} c_i \delta^4(\mathbf{a}-\mathbf{a}_i), \tag{4.3.15}$$

where $\delta^4(\mathbf{a}) = \delta^3(\mathbf{x})\delta(r)$, the convolution in (4.3.6) generalizes to the 4−space as

$$f_{K_n(A)} = (\mathbb{1}_A^4 * f_{D_0}), \text{ i.e., } (A \oplus D_0) = U_0^*(\mathbb{1}_A^4 * f_{D_0}), \tag{4.3.16}$$

whose domain restriction to the $r = 0$ hyperplane gives a bump function for the 3D solid as $f_{S_n(A)} = f_{K_n(A)}|_{r=0}$,[8] therefore $S_n(A) = U_0^*(f_{K_n(A)}|_{r=0})$. As before, we can let $f_{C_0}(\mathbf{x}, r) := \psi_\alpha(\frac{1}{r}\|\mathbf{x}\|_{\mathrm{p}})$ whose unbounded support is smoothly trimmed along the 4^{th} dimension to $r \in (-L, 0)$ as $f_{D_0}(\mathbf{x}, r) := \psi_\alpha(\frac{1}{r}\|\mathbf{x}\|_{\mathrm{p}})\psi_\alpha(1 - 2\frac{r}{L})$.

Once again, if we have an approximate discretization scheme $S \approx S_n(A)$ such that $S = \lim_{n\to\infty} S_n(A) = \lim_{|A|\to\infty}(A \oplus D_0)|_{r=0}$:

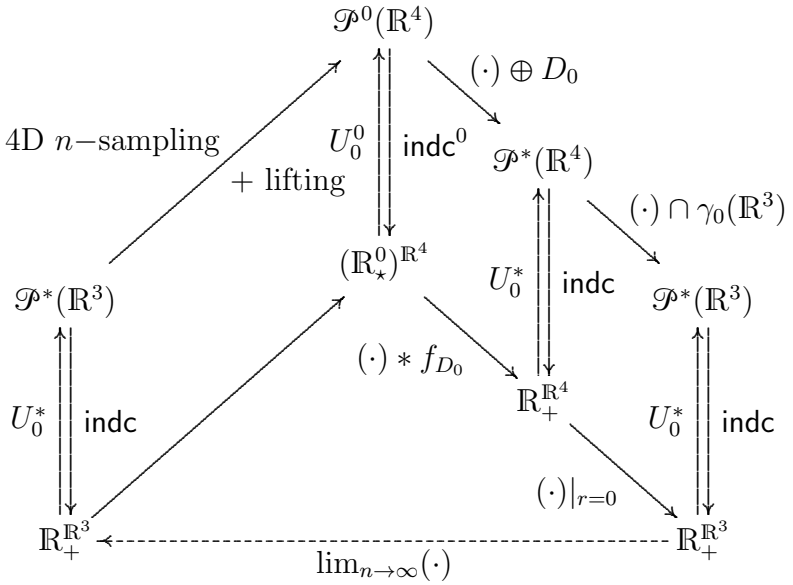

(4.3.17)

[8] For a function $f_K : \mathbb{R}^4 \to \mathbb{R}$, I use the notation $f_K|_{r=r_0} : \mathbb{R}^3 \to \mathbb{R}$ to denote the restriction (and trivial projection) of its domain to the $r = r_0$ hyperplane, i.e., $f_K|_{r=r_0}(\mathbf{x}) = f_K(\mathbf{x}, r_0)$.

Sampling Strategies. Clearly, a solid $S \in \mathcal{P}^*(\mathsf{E})$ of arbitrary shape cannot in general be exactly constructed as a finite union of balls, i.e., a 3D slice of $K_n(A) = (A \oplus D_0)|_{r=0}$. However, a similar construction is possible by replacing the finite set of ball centers $P \subset \mathbb{R}^3$ and its radius-lift $A \subset \mathbb{R}^4$, with the MA/MAT of the shape's interior, denoted by $(\mathscr{M} \circ \imath)(S)$ and $(\mathscr{T} \circ \imath)(S)$, respectively. In fact, a shape can be reconstructed as a Minkowski product of its MAT, residing in the group of translations and uniform scaling $\mathsf{G} \cong \mathbb{R}^4$, with a unit ball, i.e., a 3D slice of $K = ((\mathscr{T} \circ \imath)(S)) \oplus^* D_0$.[9] This construction can be thought of as sweeping a resizeable ball along the MA (with prespecified scaling for the balls along the MA trajectory), or equivalently, sweeping a rigid cone along the MAT in 4D followed by a 3D slicing. Unfortunately, the convolution formulation is not as simple in this case, because MA and MAT are not homogeneous, but are in general made of $2-$, $1-$ or $0-$dimensional subanalytic components for 3D solids [59]. Although it is possible to generalize the density function $\mathbb{1}_A^4$ defined in (4.3.15) to $\delta-$singularities functions of various orders over different strata of $(\mathscr{T} \circ \imath)(S)$ depending on their dimensionalities in accordance with Convention 2.3.11, it is desirable to use finite (or at most countable) constructs in this stage for the sake of computer representation.

Here I take a simple approach, by assuming sequences of finite samples $A \subset \mathbb{R}^4$ of different sizes $|A| = 1, 2, \cdots$ that progressively approximate the shape. The set $S_n(A)$ is called an $\epsilon-$approximation of S if $\mathbb{d}_{\mathrm{H}}(S_n(A), S) \leq \epsilon$, where $\mathbb{d}_{\mathrm{H}}$ denotes the Hausdorff L^2-metric. It is important to emphasize that the formulation does not impose any theoretical restriction on the sampling algorithm, as long as it guarantees that as $n \to \infty$, $S_n(A)$ converges to S (i.e., $\lim_{|A|\to\infty} \mathbb{d}_{\mathrm{H}}(S_n(A), S) = 0$), and the

[9] The MA/MAT of an r-set are not necessarily closed, and neither are their Minkowski sums with a closed ball/cone, which is why the regularized Minkowski sum/product (denoted by $\oplus^*/\otimes^*$) need to be used.

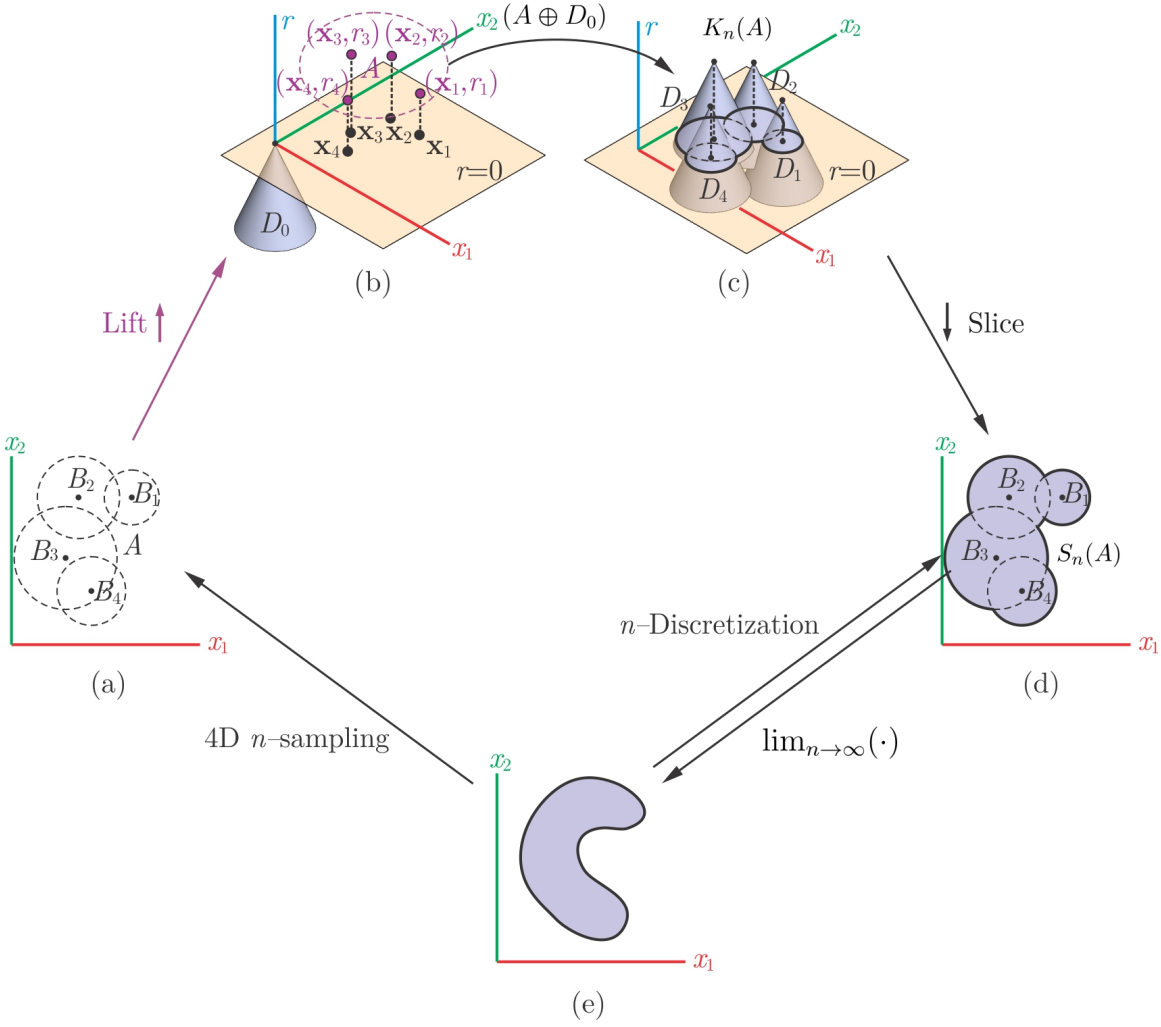

FIGURE 4.3.3: A nonequiradius sampling of a 3D r-set is a 3D slice through a 4D Minkowski sum $K_n(A) = (A \oplus D_0) = U_0^*(\mathbb{1}_A^4 * f_{D_0})$ of lifted primitives. Here, the lifting and slicing are illustrated for 2D r-sets.

convolution in (4.3.16) holds in the limit for the *countably* infinite set of knots A. A variety of methods that have been in use in the CD literature [48, 156, 248, 362] (Figs. 4.2.1 (c–e)) can be used, two of which are briefly reviewed here due to their theoretical significance and computational relationship with the algorithm that I present in great detail in Section 4.6.

Hubbard [154–156] proposed an algorithm that populates the maximal balls over the MA (Fig. 4.2.1 (d)), obtained from pruning the Voronoi diagram of a dense sampling over the boundary, and follows a principle of conservative coverage to create a bounding sphere-tree. In terms of our formulation, this is equivalent to selecting $A \subset \mathcal{T}(S)$ and has been shown in [156] to converge to the shape faster than octree-based sampling (Fig. 4.2.1 (c)). However, MA and MAT are unstable with respect to C^0- and C^1-perturbations of the boundary [59], making their computations extremely difficult in the presence of noise/errors. Weller and Zachmann [360–362] proposed the inner-sphere tree (IST) method that precomputes the distance function over a uniform grid and uses a greedy algorithm to pack the interior of the solid, giving priority to the largest ball that fits at each step. This approach has been proven effective for real-time applications [362], but leaves out void spaces in the interior of the set that are undesirable for analytic modeling, and is non-optimal for thin objects.

I use a similar greedy algorithm in Section 4.6 that is guided by the SDF field, which creates spherical samples that are similar to the outcomes of MA-based algorithms [48, 156], without the need to compute and prune the MA, and is capable of producing better approximations than distance-based sphere packing [362] with fewer number of balls. The presented algorithm guarantees bounds on the Hausdorff distance from the original shape that are proportional to the SDF grid resolution. This property is used in Section 4.5 as a basis for time complexity comparisons be-

tween operations on uniform samples versus spherical samples generated with the same input grid resolution. To prevent distraction from this article's main focus on Minkowski discretizations and their Fourier reconciliations, the details of our spherical decomposition algorithm along with its topological properties and approximation error bounds are postponed to Section 4.6.

Why Spherical Primitives? The advantage of using primitives with spherical symmetry becomes evident in the light of the isometric property of $\mathsf{C} = \mathrm{SE}(3)$. A 3D ball $B_0 := B_r(\mathbf{0})$ is invariant under 3D rotations, i.e., $\mathbb{r}B_0 = B_0$ for all $\mathbb{r} \in \mathrm{SO}(3)$ hence $(\mathbb{r}, \mathbf{t})B_0 = B_r(\mathbf{t})$ for all $\mathbb{c} \in \mathrm{SE}(3)$ represented via $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ The same invariance property can be asserted for 4D cones C_0 and D_0 whose axes stay parallel to the $r-$axis after 3D translations and rotations. Accordingly, the transformation of $S_n(A)$ and $K_n(A)$ amounts only to a relocation of the center and apex positions. For the equiradius case, the Minkowski sum in (4.3.6) for the transformed solid is given by

$$\mathbb{c}S_n(P) = (\mathbb{r}, \mathbf{t})S_n(P) \ = S_n((\mathbb{r}, \mathbf{t})P) = ((\mathbb{r}, \mathbf{t})P) \oplus B_0, \tag{4.3.18}$$

whose analytic expression is given by the real-valued defining function

$$f_{\mathbb{c}S_n(P)} = f_{(\mathbb{r},\mathbf{t})S_n(P)} = \mathbb{1}^3_{(\mathbb{r},\mathbf{t})P} * f_{B_0} = \left[\mathbb{1}^3_P \circ \mathsf{act}((\mathbb{r}, \mathbf{t})^{-1})\right] * f_{B_0}, \tag{4.3.19}$$

where the lumped density (i.e., finitely many $3^{\mathrm{rd}}-$order $\delta-$singularities)

$$\mathbb{1}^3_{(\mathbb{r},\mathbf{t})P}(\mathbf{x}) = \sum_{0 \le i < n} c_i \delta^3(\mathbf{x} - (\mathbb{r}, \mathbf{t})\mathbf{x}_i) = \sum_{0 \le i < n} c_i \delta^3(\mathbf{x} - (\mathbb{r}\mathbf{x}_i) - \mathbf{t}), \tag{4.3.20}$$

is an implicit representation of the transformed set of 3D knots

$$\mathbb{c}P = (\mathbb{r}, \mathbf{t})P = \big\{(\mathbb{r}, \mathbf{t})\mathbf{x} = (\mathbb{r}\mathbf{x}) + \mathbf{t} \mid \mathbf{x} \in P\big\}. \tag{4.3.21}$$

In a similar fashion, for nonequiradius case the Minkowski sum in (4.3.16) for the transformed lifted geometry is given by

$$\mathbb{c}K_n(A) = (\mathbb{r}, \mathbf{t})K_n(A) = K_n((\mathbb{r}, \mathbf{t})A) = ((\mathbb{r}, \mathbf{t})A) \oplus D_0, \tag{4.3.22}$$

whose analytic expression is given by the the real-valued defining function

$$f_{\mathbb{c}K_n(A)} = f_{(\mathbb{r},\mathbf{t})K_n(A)} = \mathbb{1}^4_{(\mathbb{r},\mathbf{t})A} * f_{D_0} = \big[\mathbb{1}^4_A \circ \mathsf{act}((\mathbb{r}, \mathbf{t})^{-1})\big] * f_{D_0}, \tag{4.3.23}$$

where the lumped density (i.e., finitely many $4^{\text{rd}}-$order $\delta-$singularities)

$$\mathbb{1}^4_{(\mathbb{r},\mathbf{t})A}(\mathbf{a}) = \sum_{0\le i<n} c_i \delta^4(\mathbf{a} - (\mathbb{r}, \mathbf{t})\mathbf{a}_i) = \sum_{0\le i<n} c_i \delta^4(\mathbf{a} - (\mathbb{r}\mathbf{a}_i) - \mathbf{t}), \tag{4.3.24}$$

is an implicit representation of the transformed set of 4D knots

$$\mathbb{c}A = (\mathbb{r}, \mathbf{t})A = \big\{(\mathbb{r}, \mathbf{t})\mathbf{a} = (\mathbb{r}\mathbf{a}) + \mathbf{t} \mid \mathbf{a} \in A\big\}, \tag{4.3.25}$$

using the trivial extension $(\mathbb{r}, \mathbf{t})\mathbf{a} = ((\mathbb{r}, \mathbf{t})\mathbf{x}, r)$ for $\mathbf{a} = (\mathbf{x}, r) \in \mathbb{R}^4$.

The strength of the discretization schemes that transform according to (4.3.18) through (4.3.24) lies in the rotation invariance of the primitive sets B_0 or D_0 and their radial kernels f_{B_0} or f_{D_0}, which appear in the same form in all equations before and after motion. I show in Section 4.4 that the same form is conserved across

Minkowski sums and related cross-correlations between pairs of discretized objects. The practical implication is that the primitives do not need to take an explicit role in the computations, and numerical algorithms deal only with the discrete sets of points P or A or their density functions $\mathbb{1}_P^3$ or $\mathbb{1}_A^4$, respectively.

4.3.2 Fourier Expansions

For most applications, such as the ones discussed in Chapter 5, the current analytic methods rely on discretizing the defining functions over a dense sample of points in the interior [61,174,214], which turns the continuous Fourier transform (CFT) defined in Definition 2.2.9 into a discrete Fourier transform (DFT). If the sampling is over a uniform grid, the DFT can be implemented very efficiently using the well-known fast Fourier transform (FFT) algorithm, first discovered by Cooley and Tukey [76]. The rotations are handled by an interpolation over the frequency grid and translations are embedded into convolutions.[10] However, when the objects are discretized with spherical primitives whose centers form a nonuniform set of points, the CFT interpolates a nonequispaced DFT (NDFT), to which the classical FFT algorithms do not apply. Potts et al. [258] developed a nonequispaced FFT (NFFT) algorithm for efficient implementation of NDFT sums in asymptotically similar running times with the classical FFT. The NFFT offers a formal unification of the so-called 'gridding' algorithms that were in use for years in medical imaging [165].

If an r-set $S \in \mathcal{P}^*(\mathsf{E})$ moves via $\mathbb{c} \in \mathsf{C} = \mathrm{SE}(3)$ represented as $\mathbb{c} \stackrel{0}{=} (\mathbb{r}, \mathbf{t})$ its

[10] The complete affine transformation can be conceptualized as a 6D noncommutative convolution [218] but the corresponding Fourier analysis would require different basis functions, and its cumulative computation over 6D grids is computationally impractical.

transformed defining function changes in the frequency domain as

$$\hat{f}_{\mathbb{c}S} = \hat{f}_{(\mathbb{r},\mathbf{t})S} = (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{f}_{\mathbb{r}S}) = \hat{\varsigma}_{\mathbf{t}} \bullet (\hat{f}_S \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})), \tag{4.3.26}$$

where $\varsigma_{\mathbf{t}}(\mathbf{x}) := \delta^3(\mathbf{x}-\mathbf{t})$ denotes a shifted $\delta-$function that transfers to the sinusoidal function $\hat{\varsigma}_{\mathbf{t}}(\boldsymbol{\omega}) = e^{-2\pi \mathbf{i}(\boldsymbol{\omega}\cdot\mathbf{t})}$. For an equiradius discretization $S_n(P) = (P \oplus B_0)$ the Fourier expansion of $f_{S_n(P)} = (\mathbb{1}^3_P * f_{B_0})$ is a simple product $\hat{f}_{S_n(P)} = (\hat{\mathbb{1}}_P \bullet \hat{f}_{B_0})$. Thus the diagram in (4.3.7) can be expanded to the Fourier domain as:

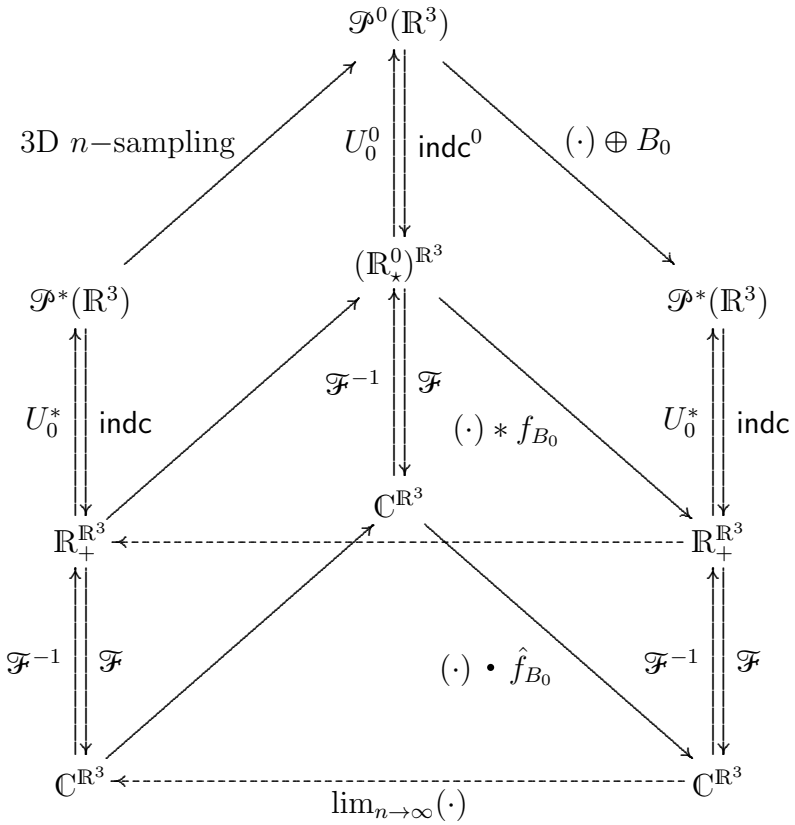

(4.3.27)

in which the left and right $\mathscr{F}^{-1} \uparrow\downarrow \mathscr{F}$ represent 3D CFTs for the continuum descriptors of the shape and its ball union approximation, respectively, while the middle $\mathscr{F}^{-1} \uparrow\downarrow \mathscr{F}$ represents 3D NDFT for the discretized function over the ball centers (i.e., knots).

For nonequiradius discretization $S_n(A) = K_n(A)|_{r=0}$ where $K_n(A) = (A \oplus D_0)$, the Fourier expansion of $f_{K_n(A)} = (\mathbb{1}_A^4 * f_{D_0})$ is a simple product $\hat{f}_{K_n(A)} = (\hat{\mathbb{1}}_A \bullet \hat{f}_{D_0})$. Thus the diagram in (4.3.17) can be expanded to the Fourier domain as:

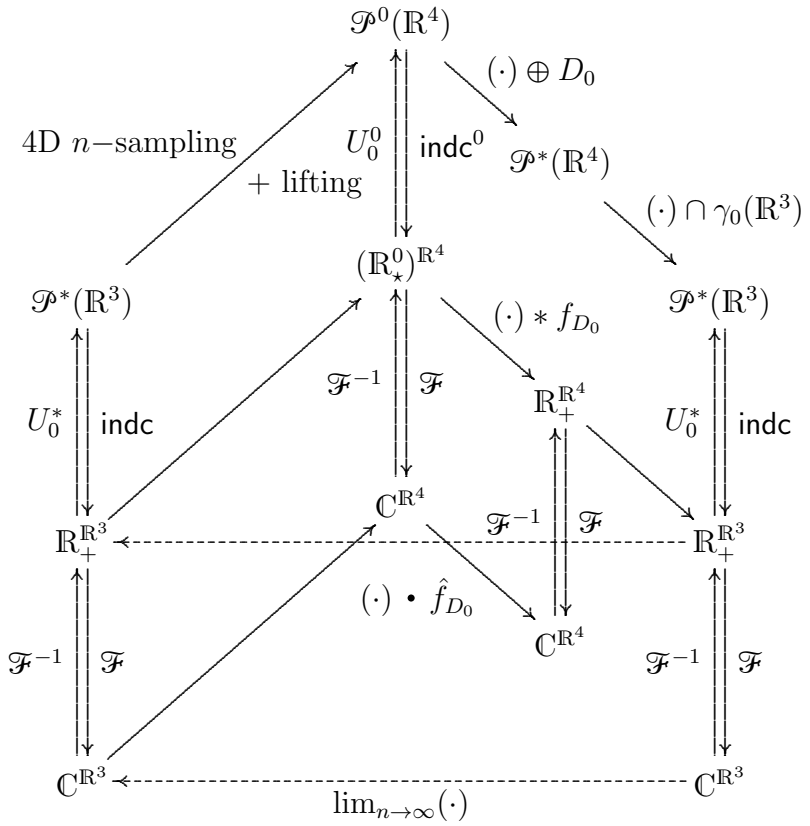

(4.3.28)

in which the vertical arrows $U_0^* \uparrow\downarrow \mathsf{indc}$ between $\mathscr{P}(\mathbb{R}^4)$ and $\mathbb{R}_+^{\mathbb{R}^4}$ for the finite collection of cones are not shown to avoid cluttering. Once again, the left and right $\mathscr{F}^{-1} \uparrow\downarrow \mathscr{F}$ represent 3D CFTs for the continuum descriptors of the shape and its ball union approximation, respectively, while the middle $\mathscr{F}^{-1} \uparrow\downarrow \mathscr{F}$ represents 4D NDFT for the discretized function over the cone apexes (i.e., knots), which approximates the 4D CFT illustrated via middle-right $\mathscr{F}^{-1} \uparrow\downarrow \mathscr{F}$. In this case, the relationship between the 3D and 4D CFTs is nontrivial.

To see how rigid motions come into play for equiradius discretizations, applying the CFT to (4.3.19) gives

$$\hat{f}_{\mathbb{C}S_n(P)} = \hat{f}_{(\mathbb{r},\mathbf{t})S_n(P)} = (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}_{\mathbb{r}P} \bullet \hat{f}_{B_0}) = \hat{\varsigma}_{\mathbf{t}} \bullet \left(\hat{\mathbb{1}}_P \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})\right) \bullet \hat{f}_{B_0}, \tag{4.3.29}$$

where the density function given in (4.3.5) is transferred to the following form:

$$\hat{\mathbb{1}}_P(\boldsymbol{\omega}) = \sum_{0 \le i < n} c_i \hat{\varsigma}_{\mathbf{x}_i}(\boldsymbol{\omega}) = \sum_{0 \le i < n} c_i e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{x}_i)}. \tag{4.3.30}$$

The evaluation of (4.3.30) from a nonuniform set of 3D knots to a uniform 3D frequency grid amounts to a one-sided 3D NDFT computation.

To see how rigid motions come into play for nonequiradius discretizations, applying the CFT to (4.3.23) gives

$$\hat{f}_{\mathbb{C}K_n(A)} = \hat{f}_{(\mathbb{r},\mathbf{t})K_n(A)} = (\hat{\varsigma}_{\mathbf{a}} \bullet \hat{\mathbb{1}}_{\mathbb{r}A} \bullet \hat{f}_{D_0}) = \hat{\varsigma}_{\mathbf{a}} \bullet \left(\hat{\mathbb{1}}_A \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})\right) \bullet \hat{f}_{D_0}, \tag{4.3.31}$$

where $\mathbf{a} = (\mathbf{t}, r) \in \mathbb{R}^4$ represents a lifted translation, and the density function given in (4.3.15) is transferred to the following form:

$$\hat{\mathbb{1}}_A(\boldsymbol{\upsilon}) = \sum_{0 \le i < n} c_i \hat{\varsigma}_{\mathbf{a}_i}(\boldsymbol{\upsilon}) = \sum_{0 \le i < n} c_i e^{-2\pi \mathbf{i}(\boldsymbol{\upsilon} \cdot \mathbf{a}_i)}, \tag{4.3.32}$$

in which $\hat{\varsigma}_{\mathbf{a}}(\boldsymbol{\upsilon}) = \hat{\varsigma}_{\mathbf{t}}(\boldsymbol{\omega}) e^{-2\pi \mathbf{i}(\eta r)}$ for the lifted physical domain $\mathbf{a}_i = (\mathbf{x}_i, r_i) \in A$ is obtained in a similar fashion to (4.3.30) using NDFTs, except that the frequency domain is also lifted to 4D as $\boldsymbol{\upsilon} = (\boldsymbol{\omega}, \eta) \in \mathbb{R}^4$.

4.4 Radial Basis Correlations

Having defined a spherical sampling in terms of a Minkowski sum of discrete knots and balls/cones alongside their analytic formulation in Section 4.3, I now investigate how they embed into correlations between pairs of objects. I show that spherical discretization structure is preserved and carried into configuration pointsets, whose Fourier formulation is given towards the end of this section.

4.4.1 Spherical Correlations

Given two r-sets $S_1, S_2 \in \mathscr{P}^*(\mathsf{E})$, I define their rigid body 'correlation function' as $\mathfrak{G}_{S_1,S_2} : \mathsf{C} \times \mathsf{C} \to \mathbb{R}$ (or $\mathbb{C} := \mathbb{R} + \mathbf{i}\mathbb{R} \cong \mathbb{R}^2$ in general) where

$$\mathfrak{G}_{S_1,S_2}(\mathbb{c}_1, \mathbb{c}_2) = \langle f_{\mathbb{c}_1 S_1}, f_{\mathbb{c}_2 S_2} \rangle = \big\langle (f_{S_1} \circ \mathsf{act}(\mathbb{c}_1^{-1})), (f_{S_2} \circ \mathsf{act}(\mathbb{c}_2^{-1})) \big\rangle, \tag{4.4.1}$$

accumulates the pointwise multiplication of the overlapped shape descriptor functions of S_1 and S_2 moved using $\mathbb{c}_1 \overset{0}{=} (\mathbb{r}_1, \mathbf{t}_1)$ and $\mathbb{c}_2 \overset{0}{=} (\mathbb{r}_2, \mathbf{t}_1)$, respectively. The function in (4.4.1) can formulate, for example, a holonomic collision constraint, a shape similarity or complementarity metric, an implicit morphological operator, or any other analytic combination of shape descriptors that can be described as a convolution. For instance, comparing (4.4.1) with (2.3.47) of Lemma 2.3.6 of Section 2.3.1 shows that if $S_1 = U_0^*(f_{S_1})$ and $S_2 = U_0^*(f_{S_2})$, then $\mathfrak{G}_{S_1,S_2}(\mathbb{c}_1, \mathbb{c}_2)$ defines a collision predicate for the moved solids, i.e., $((\mathbb{c}_1 S_1) \cap^* (\mathbb{c}_2 S_2)) \neq \emptyset$ iff $\mathfrak{G}_{S_1,S_2}(\mathbb{c}_1, \mathbb{c}_2) > 0$.

It is easy to show that the correlation function only depends on the *relative* configuration $\mathbb{c} := (\mathbb{c}_1^{-1}\mathbb{c}_2)$, i.e., the instantaneous pose of S_2 observed from a coordinate frame attached to S_1. Letting $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t}) = (\mathbb{r}_1, \mathbf{t}_1)^{-1}(\mathbb{r}_2, \mathbf{t}_2) = (\mathbb{r}_1^{\mathrm{T}}\mathbb{r}_2, \mathbb{r}_1^{\mathrm{T}}(\mathbf{t}_2 - \mathbf{t}_1))$, the

alternative formulation $\mathfrak{G}_{S_1,S_2} : \mathsf{C} \to \mathbb{R}$ becomes:

$$\mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t}) = \big\langle f_{S_1}, (f_{\mathbb{r}S_2} \circ \mathsf{act}(\mathbb{t}^{-1})) \big\rangle = \big\langle (f_{\mathbb{r}^{\mathrm{T}}S_1} \circ \mathsf{act}(\mathbb{t})), f_{S_2} \big\rangle, \tag{4.4.2}$$

also noting that $f_{\mathbb{c}S} = (f_{\mathbb{r}S} \circ \mathsf{act}(\mathbb{t}^{-1}))$—where $\mathbb{t} \in \mathrm{T}(3)$ is the translation corresponding to $\mathbf{t} \in \mathbb{R}^3$ with respect to a chosen origin $\mathbf{0} \in \mathsf{E}$—which, in turn, can be further decoupled into $f_{\mathbb{r}S} = (f_S \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}}))$. The inner product in (4.4.2) can be viewed as a 6D noncommutative convolution over $\mathsf{C} = \mathrm{SE}(3)$ [218]. For numerical tractability, I decompose the motion into rotational and translational parts, and view the latter as a 3D commutative convolution:

$$\mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t}) = (f_{S_1} * \tilde{f}_{\mathbb{r}S_2})(\mathbf{t}) = (f_{S_1} * (\tilde{f}_{S_2} \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}}))(\mathbf{t}), \tag{4.4.3}$$

where $\tilde{f}_S = f_{-S}$ and $-S = \{-\mathbf{x} \mid \mathbf{x} \in S\}$ denotes a reflection with respect to the origin $\mathbf{0} \in \mathsf{E}$. This defines a *relative* collision predicate, i.e., $(S_1 \cap^* (\mathbb{r}, \mathbf{t})S_2) \neq \emptyset$ iff $\mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t}) > 0$, and brings us back to the important concept of $\mathsf{C}-$obstacle:

$$O_{S_1,S_2} = \big\{ \mathbb{c} \stackrel{0}{=} (\mathbb{r}, \mathbf{t}) \in \mathsf{C} \mid \mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t}) > 0 \big\}, \tag{4.4.4}$$

whose set-theoretic complement $\complement(O_{S_1,S_2})$ is the 'free space', as mentioned earlier in Section 3.4.1.2. For a fixed rotation $\mathbb{r} \in \mathrm{SO}(3)$ the translational $\mathsf{C}-$obstacle is

$$O_{S_1,S_2}|_{\mathbb{r}} := \big\{ \mathbf{t} \in \mathbb{R}^3 \cong \mathsf{E} \mid \mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t}) > 0 \big\}, \tag{4.4.5}$$

is a 3D slice through the 6D obstacle, obtained by offsetting S_1 with the rotated and

reflected $-\mathbb{r}S_2$ which, in turn, is given by a Minkowski sum:

$$O_{S_1,S_2}|_{\mathbb{r}} = S_1 \oplus^{\ddagger} (-\mathbb{r}S_2) = U_0(f_{S_1} * \tilde{f}_{\mathbb{r}S_2}). \tag{4.4.6}$$

As a result of the definition in (4.4.4), $\mathfrak{G}_{S_1,S_2}(\mathbb{r}, \mathbf{t})$ serves as a defining function of the $\mathsf{C}-$obstacle, which, in turn, is retrieved as the $0-$superlevel set $O_{S_1,S_2} = U_0(\mathfrak{G}_{S_1,S_2})$.

Note that depending on the choice of the shape descriptors $f_{S_1}, f_{S_2} \in \mathbb{R}_+^{\mathsf{E}}$, their correlation function have different meanings. For example:

- Using nonsingular interior shape descriptors $f_{S_1} := \mathbb{1}^0_{i(S_1)}$ and $f_{S_2} := \mathbb{1}^0_{i(S_2)}$, the correlation function $\mathfrak{G}_{S_1,S_2} = (f_{S_1} * \tilde{f}_{S_2})$ measures the collision volume, thus is a motion descriptor (i.e., defining function) of the $\mathsf{C}-$obstacle as depicted above.

- Using singular boundary shape descriptors $f_{S_1} := \mathbb{1}^1_{\acute{b}(S_1)}$ and $f_{S_2} := \mathbb{1}^1_{\acute{b}(S_2)}$, the correlation function $\mathfrak{G}_{S_1,S_2} = (f_{S_1} * \tilde{f}_{S_2})$ measures the contact area, thus is a motion descriptor (i.e., defining function) of the contact space.

- Using shape descriptors of medial axis (MA) of interiors and exteriors, of appropriate orders of $\delta-$singularities at various branches of different dimensionalities, the correlation function $\mathfrak{G}_{S_1,S_2} = (f_{S_1} * \tilde{f}_{S_2})$ measures the similarity and/or complementarity of the shapes in terms of internal and external skeletal overlaps.

Next I investigate the discretization of the $\mathsf{C}-$obstacle (and the Minkowski sums in general) when spherical sampling is used for the constituent parts. The ideas described for the $\mathsf{C}-$obstacle can be readily extended to other interpretations of the correlation function exemplified above.

Equiradius Correlations. First, let $S_1 = S_{n_1}(P_1)$ and $S_2 = S_{n_2}(P_2)$ be composed of instances of balls denoted by $B_1 := B_{r_1}(\mathbf{0})$ and $B_2 := B_{r_2}(\mathbf{0})$, respectively. Substituting from (4.3.18) in (4.4.6), and noting the invariance of the balls with respect to reflection and rotation (i.e., $-\mathbb{r}B_2 = B_2$) and the commutativity of Minkowski sums, we obtain

$$O_{S_1,S_2}|_{\mathbb{r}} = (P_1 \oplus B_1) \oplus ((-\mathbb{r}P_2) \oplus B_2) \tag{4.4.7}$$

$$= (P_1 \oplus (-\mathbb{r}P_2)) \oplus (B_1 \oplus B_2) \tag{4.4.8}$$

$$= (P_O|_{\mathbb{r}} \oplus B_O). \tag{4.4.9}$$

The first term $P_O|_{\mathbb{r}} := P_1 \oplus (-\mathbb{r}P_2)$ is a finite set of $n_1 n_2$ points in the 3D translation space obtained from pairwise summations of ball centers in P_1 and $-\mathbb{r}P_2$. It represents the discrete translational obstacle $O_{P_1,P_2}|_{\mathbb{r}}$, a 3D slice through $P_O := O_{P_1,P_2} \subset \mathrm{SE}(3)$ defining a collection of $n_1 n_2$ curves. The second term $B_O := (B_1 \oplus B_2) = B_{r_O}(\mathbf{0})$ is a ball of radius $r_O = (r_1 + r_2)$ in the translational space, representing the 'primitive obstacle' $O_{B_1,B_2}|_{\mathbb{r}}$, which is a 3D slice of $O_{B_1,B_2} \subset \mathrm{SE}(3)$. Therefore, the total obstacle itself is a finite union of balls, i.e., discretized with the same scheme as the original objects: $O_{S_1,S_2}|_{\mathbb{r}} = S_{n_1 n_2}(P_O|_{\mathbb{r}})$.

Combining the two diagrams of the form (4.3.2) for the two constituent shapes $S_1, S_2 \in \mathscr{P}^*(\mathbb{R}^3)$ and their spherical approximations $S_{n_1}(P_1), S_{n_1}(P_1) \in \mathscr{P}^*(\mathbb{R}^3)$ for $P_1, P_2 \in \mathscr{P}^0(\mathbb{R}^3)$ via (4.4.7) through (4.4.9) yields:

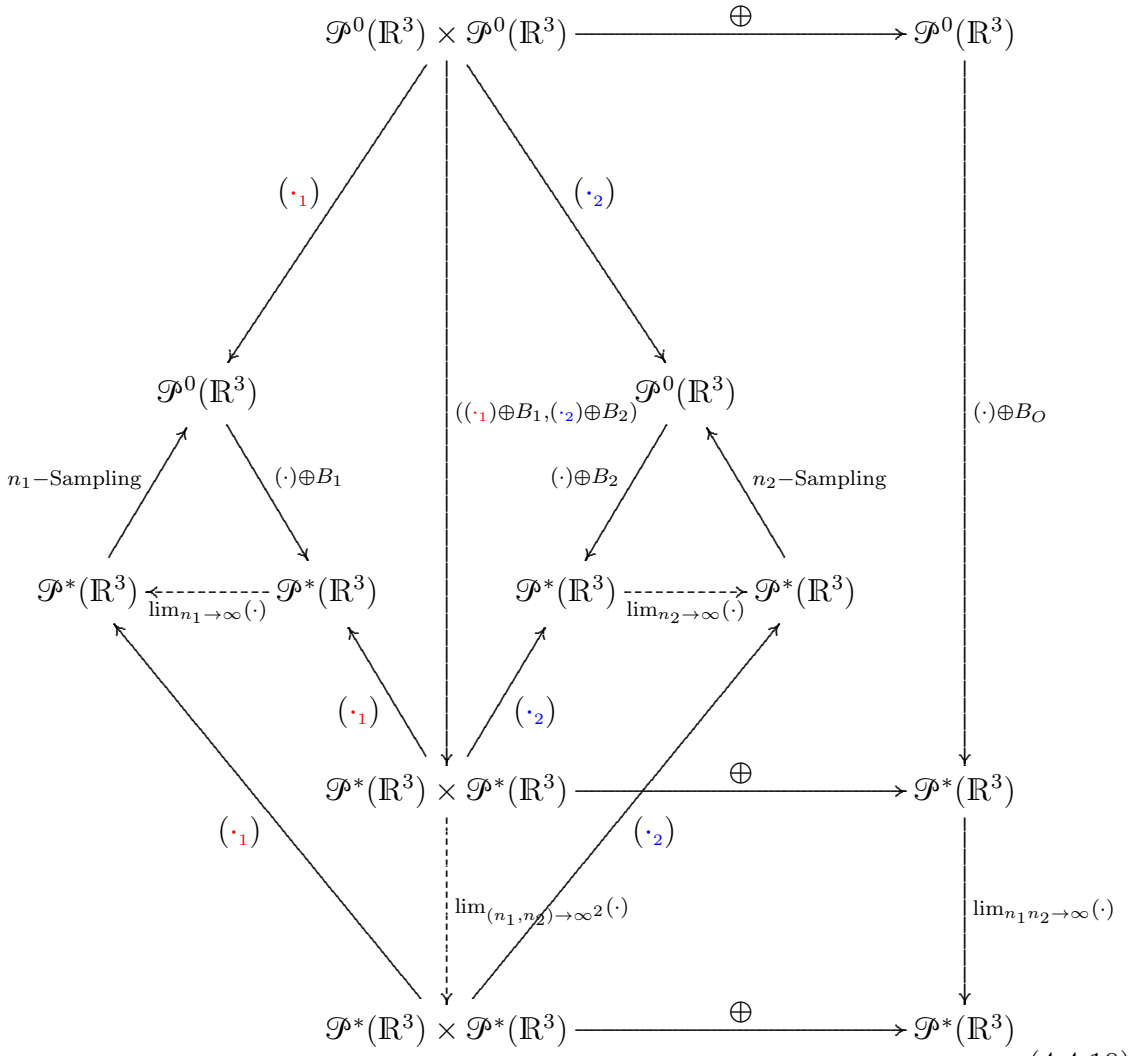

(4.4.10)

This diagram means that computing explicit Minkowski operations—and various derived morphological operations—can be performed in a finite number of steps (i.e., takes $O(n_1 n_2)-$time) if the two shapes are approximated via equiradius collections of n_1 and n_2 balls, respectively. As I discussed in Section 3.4.1.1, if the two approximations are ϵ_1- and ϵ_2-covers of the shapes—i.e., the Hausdorff distances are bounded as $\mathrm{d_H}(S_1, S_{n_1}(P_1)) \leq \epsilon_1$ and $\mathrm{d_H}(S_2, S_{n_2}(P_2)) \leq \epsilon_2$—then the implementation yields an

$(\epsilon_1+\epsilon_2)-$cover of the Minkowski sum, i.e., $\mathbb{d}_{\mathrm{H}}(S_1 \oplus S_2, S_{n_1}(P_1) \oplus S_{n_2}(P_2)) \leq (\epsilon_1+\epsilon_2)$. The rotational invariance of the balls is the key when generalizations of Minkowski sums to group morphological operations are of interest, e.g., $\mathsf{C}-$obstacle computation in (4.4.9). An illustration is given in Fig. 4.4.1 (a–c).

The analytic formulation in (4.4.3) develops in parallel as:

$$\mathfrak{G}_{S_1,S_2}|_R = (\mathbb{1}^3_{P_1} * f_{B_1}) * (\mathbb{1}^3_{-\mathbb{r}P_2} * f_{B_2}) \tag{4.4.11}$$

$$= (\mathbb{1}^3_{P_1} * \mathbb{1}^3_{-\mathbb{r}P_2}) * (f_{B_1} * f_{B_2}), \tag{4.4.12}$$

$$= (\mathbb{1}^3_{P_{O|_{\mathbb{r}}}} * f_{B_O}), \tag{4.4.13}$$

where $\mathbb{1}^3_{P_{O|_{\mathbb{r}}}} := (\mathbb{1}^3_{P_1} * \mathbb{1}^3_{-\mathbb{r}P_2})$ is the impulsive density function of the discrete pointset $P_O|_{\mathbb{r}}$ made of $n_1 n_2$ impulses corresponding to pairwise convolution of $\delta-$terms from the constituents,[11] while $f_{B_O} := (f_{B_1} * f_{B_2})$ is a convolution of two radial bumps of radii r_1 and r_2, respectively, leading to another radial bump of radius $r_O = (r_1 + r_2)$ that defines the obstacle ball B_O. Note that if we choose $f_{B_1} := \psi_\alpha(\frac{1}{r_1}\|\cdot\|_2)$ and $f_{B_2} := \psi_\alpha(\frac{1}{r_2}\|\cdot\|_2)$ using the form in (4.3.9), their convolution does not take the same form, i.e., $f_{B_O} \neq \psi_\alpha(\frac{1}{r_O}\|\cdot\|_2)$. Nevertheless, $f_{B_O}(\cdot) \equiv \psi_\alpha(\frac{1}{r_O}\|\cdot\|_2)$, thus the latter can be safely replaced for the last term in (4.4.13) without changing the obstacle $O_{S_1,S_2}|_{\mathbb{r}} = U_0^*(\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}})$, noting that the choice of the defining function as a radial basis summation was arbitrary in the first place, as long as $B_O = U_0^*(f_{B_O})$.

The diagram in (4.4.10) can be implicitized, by combining (4.3.7) for the two constituent shape descriptors $f_{S_1}, f_{S_2} \in \mathbb{R}_+^{\mathbb{R}^3}$ and their spherical approximations $f_{S_{n_1}(P_1)}, f_{S_{n_1}(P_1)} \in \mathbb{R}_+^{\mathbb{R}^3}$ with $\mathbb{1}^3_{P_1}, \mathbb{1}^3_{P_2} \in (\mathbb{R}_\star^0)^{\mathbb{R}^3}$ via (4.4.11) through (4.4.13) yields:

[11] Note that if we let $\mathbb{1}^3_1(\mathbf{x}) := \delta^3(\mathbf{x}-\mathbf{x}_1)$, $\mathbb{1}^3_2(\mathbf{x}) := \delta^3(\mathbf{x}-\mathbf{x}_2)$, and $\tilde{\mathbb{1}}(\mathbf{x}) := \mathbb{1}(-\mathbf{x})$, then $(\mathbb{1}^3_1 * \tilde{\mathbb{1}}_2)(\mathbf{t}) = \delta^3(\mathbf{t}-(\mathbf{x}_1+\mathbf{x}_2))$.

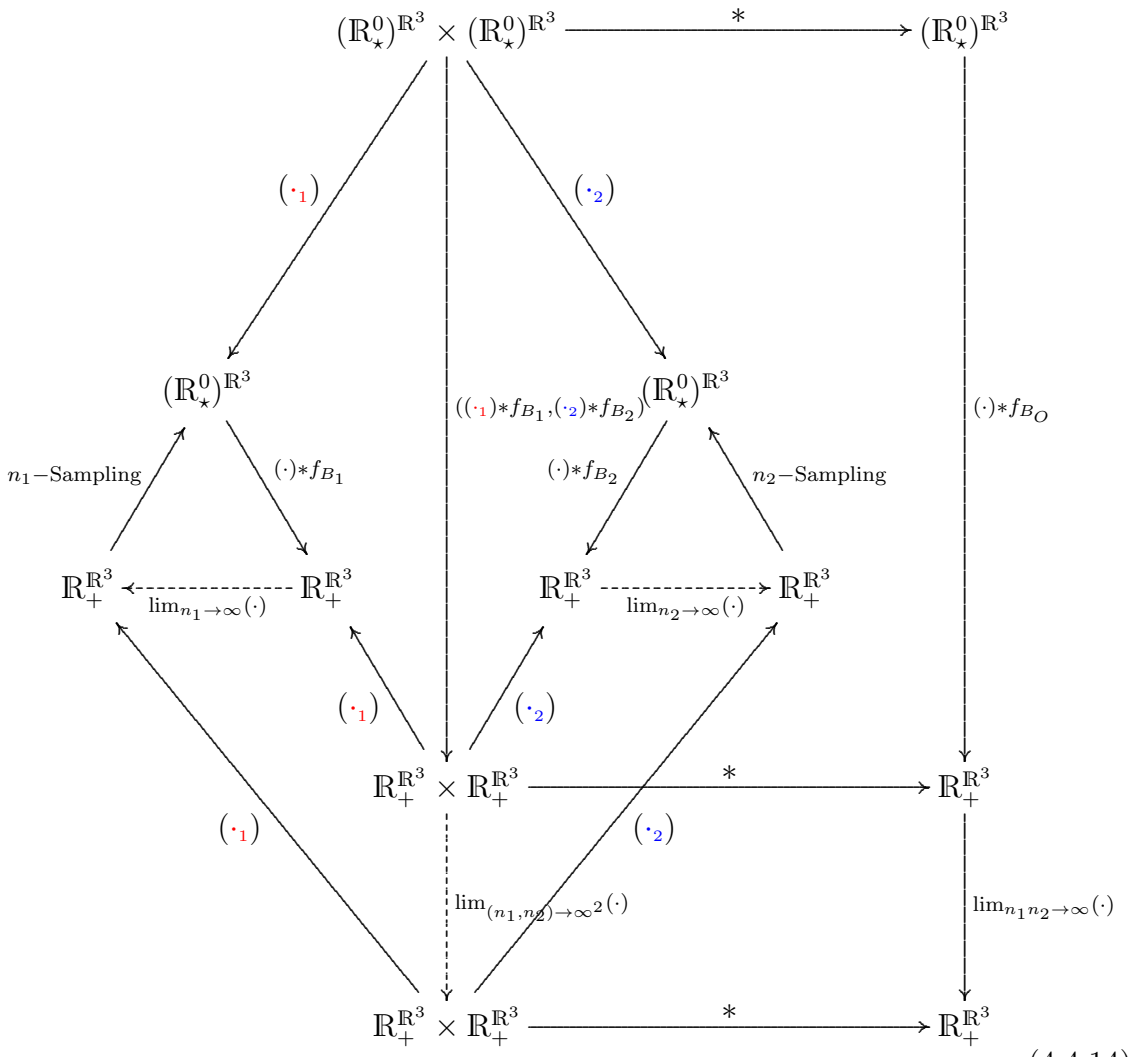

(4.4.14)

The explicit and implicit Minkowski sum implementations in (4.4.10) and (4.4.14), respectively, can be connected term-by-term via $U_0 \uparrow\downarrow \mathsf{desc}$, as expected. The benefit with implicit implementation is that application of 3D CFTs to the continuum convolution and its approximation (the bottom two horizontal arrows) can be implemented by application of 3D NDFTs to the discrete convolution (the top horizontal arrow), converting them into pointwise multiplications (Theorem 2.2.11).

Nonequiradius Correlations The generalization to implementing correlations for $S_1 = S_{n_1}(A_1)$ and $S_2 = S_{n_2}(A_2)$ made of primitive balls of different sizes is not straightforward. This is because the commutativity of the discretization Minkowski sum $S_n(P) = (P \oplus B_0)$ that led from (4.4.7) to (4.4.8) does not extend to the Minkowski product $S_n(A) = (A \otimes \gamma_0(B_0))$. In terms of the lifted geometry, this manifests as the observation that the 3D Minkowski sum of the cross-sections is not equal to the cross-section of the 4D Minkowski sum; or in other words, a collision between $K_1 = K_{n_1}(A_1)$ and $K_2 = K_{n_2}(A_2)$ does not necessarily imply a collision between the cross-sections $S_1 = K_{n_1}(A_1)|_{r=0}$ and $S_2 = K_{n_2}(A_2)|_{r=0}$.

At the primitive level, this is because the 4D half-cones, despite being invariant under 3D rotations and reflections, are *not* invariant under 4D reflections, hence $D_0 \neq -D_0$ and the sum $(D_0 \oplus (-D_0))$ (that appears in $K_1 \oplus (-\mathbb{r}K_2)$) does *not* give a half-cone in the $\mathsf{C}-$space. Fortunately, this can be solved by a pre-reflection with respect to the $r = 0$ hyperplane of one of the two lifted shapes. If we let

$$\breve{K} = \{(\mathbf{x}, -r) \mid (\mathbf{x}, r) \in K\}, \quad (K, \breve{K} \subset \mathbb{R}^4) \tag{4.4.15}$$

denote the $r-$mirror image of the 4D set K, the 3D set can be retrieved from both of them as $S = K|_{r=0} = \breve{K}|_{r=0}$. Then the nonequiradius discretization scheme in (4.3.22) gives $\breve{K}_n(\breve{A}) = (\breve{A} \oplus \breve{D}_0)$ and $-\breve{K}_n(\breve{A}) = ((-\breve{A}) \oplus D_0)$, noting that $D_0 = -\breve{D}_0$ and the sum $(D_0 \oplus (-\breve{D}_0))$ (that appears in $K_1 \oplus (-\mathbb{r}\breve{K}_2)$) is a half-cone of double the size in the $\mathsf{C}-$space. Furthermore, it is easy to prove that in this case, for two collections of half-cones of *opposite* directions that intersect the $r = 0$ hyperplane, a collision between the 4D solids does in fact imply a collision between the 3D slices. This is made precise by the following important lemma:

Lemma 4.4.1. (Cone Collision Lemma) *The collision predicates for the nonequiradius balls in 3D and the lifted half-cones in 4D, upon inversion of half-cones of one of the two shapes along the $r-$direction, are equivalent, i.e., for a given $\mathbb{C} \overset{0}{=} (\mathbb{r}, \mathbf{t})$,*

$$S_{1,2} := S_{n_{1,2}}(A_{1,2}) = \left[K_{n_{1,2}}(A_{1,2})\right]_{r=0} : \quad S_1 \cap^* (\mathbb{c}S_2) \neq \emptyset \;\rightleftharpoons\; K_1 \cap^* (\mathbb{c}\breve{K}_2) \neq \emptyset, \tag{4.4.16}$$

Proof. The proof is straightforward, by noting that for every pair of 4D half-cones $D_1 := D_{r_1}(\mathbf{x}_1)$ and $D_2 := D_{r_2}(\mathbf{x}_2)$ corresponding to $(\mathbf{x}_1, r_1) \in A_1$ and $(\mathbf{x}_2, r_2) \in A_2$, respectively, they intersect after flipping one of them upside down (i.e., $(D_1 \cap^* \breve{D}_2) \neq \emptyset$) if and only if their 3D slices intersect (i.e., $(D_1|_{r=0} \cap^* \breve{D}_2|_{r=0}) \neq \emptyset$). The assertion is very easy to picture for 3D cones whose slices are 2D disks. □

As a direct corollary, we can define a 4D translational $\mathsf{C}-$obstacle that is discretized with the same scheme as

$$O_{K_1,\breve{K}_2}|_{\mathbb{r}} = (A_1 \oplus D_0) \oplus ((-\mathbb{r}\breve{A}_2) \oplus D_0) \tag{4.4.17}$$

$$= (A_1 \oplus (-\mathbb{r}\breve{A}_2)) \oplus (D_0 \oplus D_0) \tag{4.4.18}$$

$$= (A_O|_{\mathbb{r}} \oplus D_O), \tag{4.4.19}$$

and the 3D obstacle is a slice $O_{S_1,S_2}|_{\mathbb{r}} = \left[O_{K_1,\breve{K}_2}|_{\mathbb{r}}\right]_{r=0}$. The first term $A_O|_{\mathbb{r}} := A_1 \oplus (-\mathbb{r}\breve{A}_2)$ is a finite set of $n_1 n_2$ points in the 4D translation space obtained from pairwise summations of cone apexes in A_1 and $-\mathbb{r}\breve{A}_2$, which is the same as $O^*_{A_1,\breve{A}_2}|_{\mathbb{r}}$. In this case, the primitive obstacle $D_O := (D_0 \oplus D_0)$ is a larger half-cone with a height of $2L$, which is equal to $O^*_{D_0,D_0}|_{\mathbb{r}}$. Therefore, the $\mathsf{C}-$obstacle is discretized with the same scheme as $O^*_{S_1,S_2}|_{\mathbb{r}} = S_{n_1 n_2}(A_O|_{\mathbb{r}}) = \left[K_{n_1 n_2}(A_O|_{\mathbb{r}})\right]_{r=0}$.

Combining the two diagrams of the form (4.3.13) for the two constituent shapes $S_1, S_2 \in \mathscr{P}^*(\mathbb{R}^3)$ and their spherical approximations $S_{n_1}(A_1), S_{n_1}(A_1) \in \mathscr{P}^*(\mathbb{R}^3)$ for $A_1, A_2 \in \mathscr{P}^0(\mathbb{R}^4)$ via (4.4.17) through (4.4.19) yields:

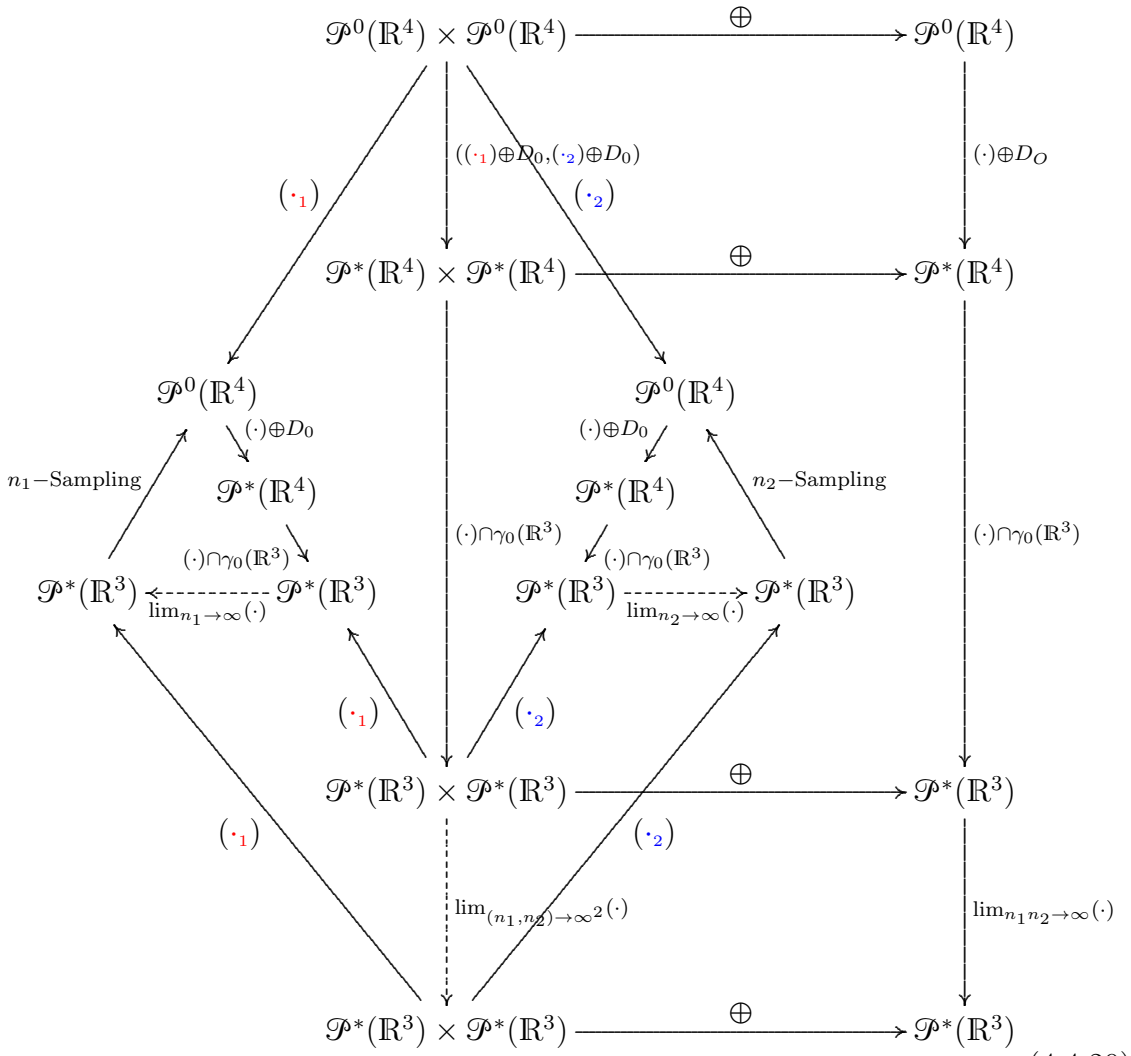

(4.4.20)

This diagram means that computing explicit Minkowski operations—and various derived morphological operations—can be performed in a finite number of steps (i.e., takes $O(n_1 n_2)-$time) if the two shapes are approximated via nonequiradius collec-

tions of n_1 and n_2 balls, respectively. Once again, error bound guarantees can be obtained for the implementation in terms of the Hausdorff distance in a straightforward fashion. The rotational invariance of the balls is the key when generalizations of Minkowski sums to group morphological operations are of interest, e.g., $\mathsf{C}-$obstacle computation in (4.4.19). An illustration is given in Fig. 4.4.1 (d–f).

The analytic formulation in (4.4.3) develops in parallel as

$$\mathfrak{G}_{K_1,\check{K}_2}|_{\mathbb{r}} = (\mathbb{1}^4_{A_1} * f_{D_0}) * (\mathbb{1}^4_{-\mathbb{r}\check{A}_2} * f_{D_0}) \tag{4.4.21}$$

$$= (\mathbb{1}^4_{A_1} * \mathbb{1}^4_{-\mathbb{r}\check{A}_2}) * (f_{D_0} * f_{D_0}) \tag{4.4.22}$$

$$= \mathbb{1}^4_{A_{O|_{\mathbb{r}}}} * f_{D_O}, \tag{4.4.23}$$

whose restriction to $r = 0$ gives $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}} = \left[\mathfrak{G}_{K_1,\check{K}_2}|_{\mathbb{r}}\right]_{r=0}$. Similar to the equiradius case, $\mathbb{1}^4_{A_{O|_{\mathbb{r}}}} := (\mathbb{1}^4_{A_1} * \mathbb{1}^4_{-\mathbb{r}\check{A}_2})$ is the impulsive density function of the discrete pointset $A_O|_{\mathbb{r}}$ made of $n_1 n_2$ impulses corresponding to cross-correlation of pairs of shifted Dirac $\delta-$functions, while $f_{D_O} := (f_{D_0} * f_{D_0})$ is an auto-correlation, which can be arbitrarily modified from the original convolved form to $f_{D_O}(\mathbf{t}, r) := \psi_\alpha(\frac{1}{r}\|\mathbf{t}\|_2)\psi_\alpha(1-\frac{r}{L})$ without changing the obstacle $O_{S_1,S_2}|_{\mathbb{r}} = U_0^*(\mathfrak{G}_{K_1,\check{K}_2}|_{r=0})$.

The diagram in (4.4.20) can be implicitized, by combining (4.3.17) for the two constituent shape descriptors $f_{S_1}, f_{S_2} \in \mathbb{R}_+^{\mathbb{R}^3}$ and their spherical approximations $f_{S_{n_1}(A_1)}, f_{S_{n_1}(A_1)} \in \mathbb{R}_+^{\mathbb{R}^3}$ with $\mathbb{1}^4_{A_1}, \mathbb{1}^4_{A_2} \in (\mathbb{R}_\star^0)^{\mathbb{R}^4}$ via (4.4.21) through (4.4.23) yields:

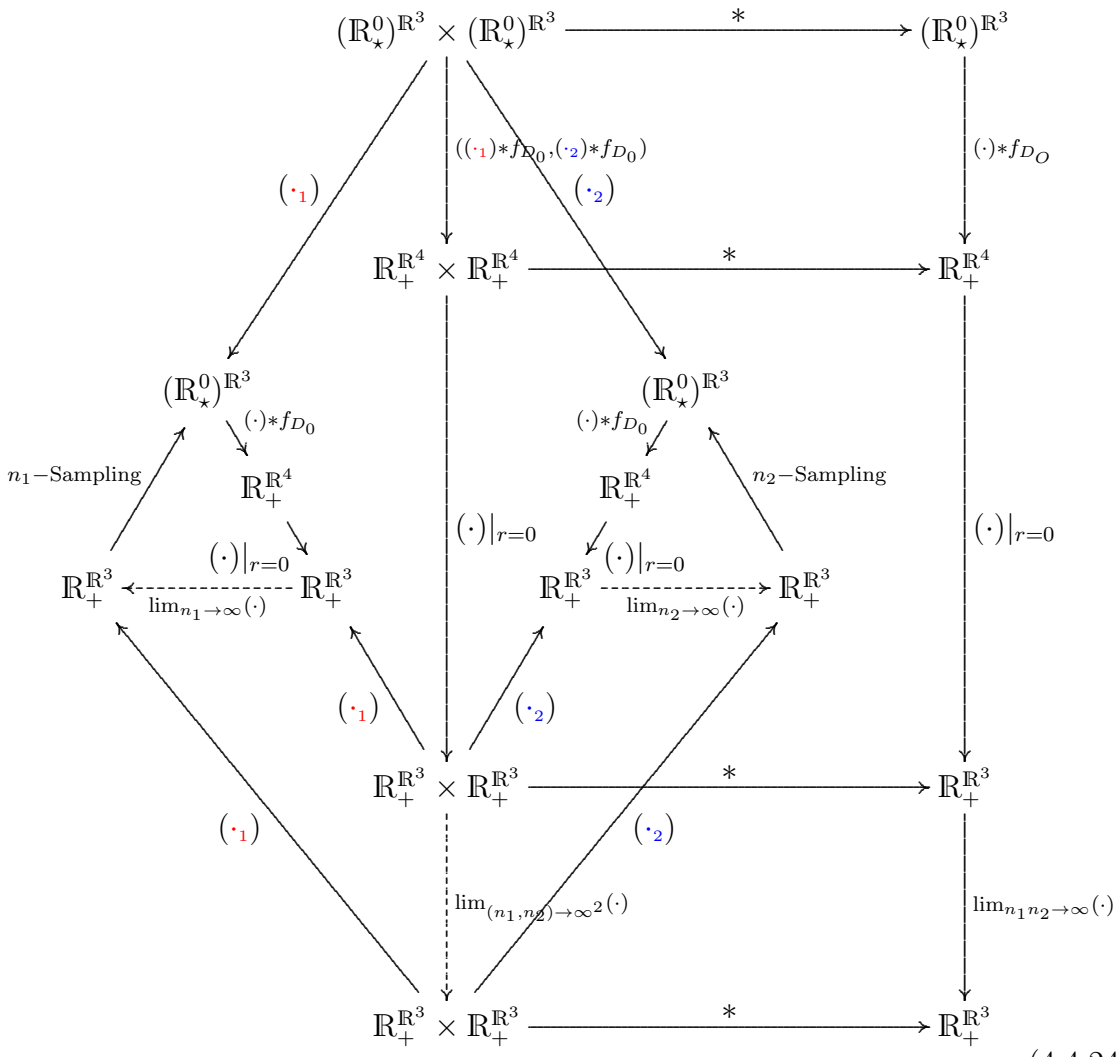

(4.4.24)

The explicit and implicit Minkowski sum implementations in (4.4.20) and (4.4.24), respectively, can be connected term-by-term via $U_0 \uparrow\downarrow \mathsf{desc}$, as expected. The benefit with implicit implementation is that application of 3D CFTs to the continuum convolution and its approximation (the bottom two horizontal arrows), viewed as a 0−slice of 4D CFTs applied to the continuum convolution in the lifted space, can be implemented by application of 4D NDFTs to the discrete convolution (the top

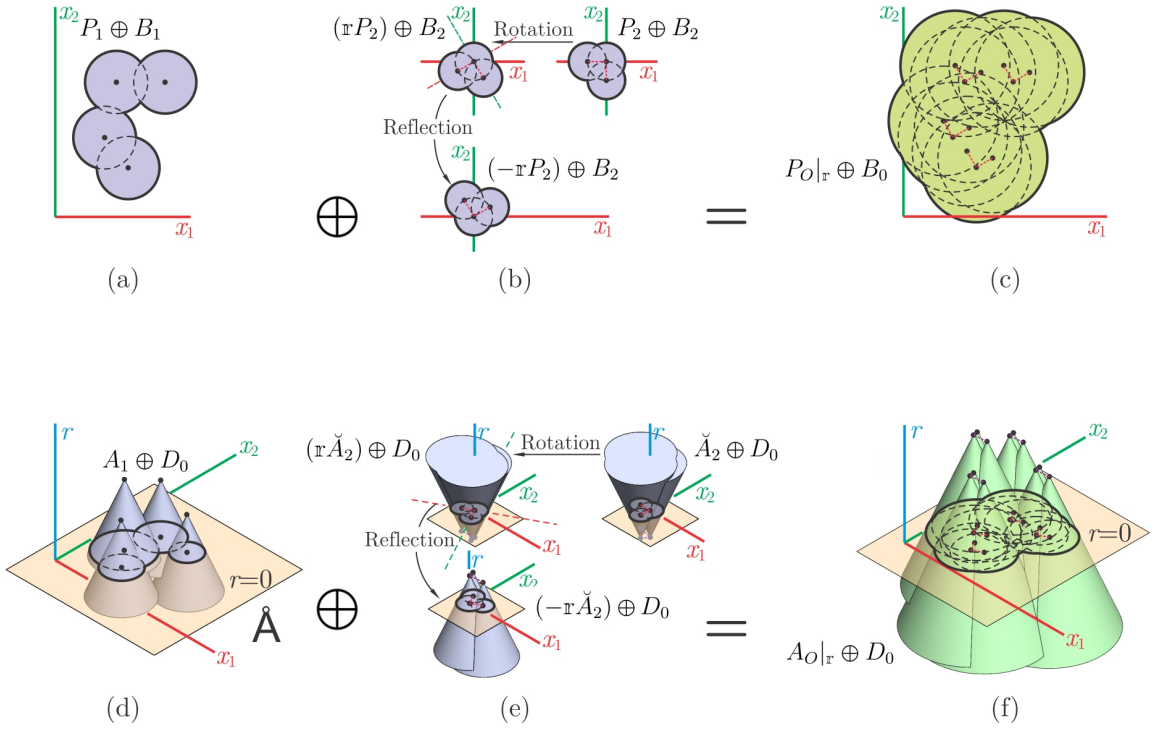

FIGURE 4.4.1: The C−obstacle is obtained as a Minkowski sum. The discretization scheme is closed under Minkowski sums for both equiradius and nonequiradius samples. For the latter, the summands must have the same cone orientation along the $r-$axis, requiring a pre-reflection. The lifting and slicing are illustrated for 2D r-sets.

horizontal arrow), converting them into pointwise multiplications (Theorem 2.2.11).

Although one only needs the $r = 0$ slice to retrieve the C−obstacle, the other slices carry useful information. In fact, any $r \neq 0$ slice corresponding to $r \in (-L, +L)$ gives the obstacle for a pair of *offset* 3D solids defined as r_1-slice of K_1 (i.e., shrinking S_1's primitives by r_1) and r_2-slice of $\breve{K}_2$ (i.e., expanding S_2's primitives by r_2) giving a total offset of $-r = -(r_1 + r_2)$. These 'offset obstacles' can be used, for example, to incorporate tolerances for machine tooling, guarantee safety margins for path planning, or construct skin layers for protein shape complementarity modeling.

4.4.2 Fourier Correlations

To take advantage of the convolution theorem (Theorem 2.2.11) for a significantly faster computation of correlation functions, I present an analysis of the gap function in the Fourier domain. For two r-sets $S_1, S_2 \in \mathscr{P}^*(\mathsf{E})$, applying CFT to (4.4.3) yields the Fourier correlation function as

$$\hat{\mathfrak{G}}_{S_1,S_2}|_{\mathbb{r}} = (\hat{f}_{S_1} \bullet \bar{\hat{f}}_{\mathbb{r}S_2}) = \hat{f}_{S_1} \bullet \left(\bar{\hat{f}}_{S_2} \circ \mathsf{act}(\mathbb{r}^{\mathrm{T}})\right), \tag{4.4.25}$$

noting that $\hat{f}_{-\mathbb{r}S} = \bar{\hat{f}}_{\mathbb{r}S}$ for real defining functions, i.e., CFT converts reflection (in both physical and frequency domains) to conjugation in the frequency domain, a property known as Hermitian symmetry. $\hat{\mathfrak{G}}_{S_1,S_2}|_{\mathbb{r}} = \mathscr{F}\{\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}\}$ is the CFT of $\mathfrak{G}_{S_1,S_2}(\mathbb{r},\mathbf{t})$ only with respect to translation at a fixed $\mathbb{r} \in \mathrm{SO}(3)$.

For the equiradius discretizations $S_{n_1}(P_1) = (P_1 \oplus B_1)$ and $S_{n_2}(P_2) = (P_2 \oplus B_2)$ with their Fourier representations $\hat{f}_{S_{n_1}(P_1)} = (\hat{\mathbb{1}}^3_{P_1} \bullet \hat{f}_{B_1})$ and $\hat{f}_{S_{n_2}(P_2)} = (\hat{\mathbb{1}}^3_{P_2} \bullet \hat{f}_{B_2})$, respectively, their Fourier correlation is obtained by substituting (4.3.29) in (4.4.25), or alternatively, by directly applying the CFT to (4.4.13) as

$$\hat{\mathfrak{G}}_{S_1,S_2}|_{\mathbb{r}} = \left(\hat{\mathbb{1}}^3_{P_1} \bullet \hat{f}_{B_1}\right) \bullet \left(\bar{\hat{\mathbb{1}}}^3_{\mathbb{r}P_2} \bullet \bar{\hat{f}}_{B_2}\right) \tag{4.4.26}$$

$$= \left(\hat{\mathbb{1}}^3_{P_1} \bullet \bar{\hat{\mathbb{1}}}^3_{\mathbb{r}P_2}\right) \bullet \left(\hat{f}_{B_1} \bullet \hat{f}_{B_2}\right) = (\hat{\mathbb{1}}^3_{P_O|_{\mathbb{r}}} \bullet \hat{f}_{B_O}), \tag{4.4.27}$$

noting that $f_{B_{1,2}} = f_{-B_{1,2}}$ hence $\hat{f}_{B_{1,2}} = \bar{\hat{f}}_{B_{1,2}}$ (i.e., are both real-valued), and so is $\hat{f}_{B_O} = (\hat{f}_{B_1} \bullet \hat{f}_{B_2})$. As expected, $\hat{\mathbb{1}}^3_{P_O|_{\mathbb{r}}} = (\hat{\mathbb{1}}^3_{P_1} \bullet \bar{\hat{\mathbb{1}}}^3_{\mathbb{r}P_2})$ is computed from a pointwise multiplication of the NDFTs over the knots P_1 and P_2.

Analogously, for the nonequiradius discretizations $S_{n_1}(A_1) = K_{n_1}(A_1)|_{r=0}$ and $S_{n_2}(A_2) = \breve{K}_{n_2}(\breve{A}_2)|_{r=0}$, in which $K_{n_1}(A_1) = (A_1 \oplus D_0)$ and $\breve{K}_{n_2}(\breve{A}_2) = (\breve{A}_2 \oplus \breve{D}_0)$

with their Fourier representations $\hat{f}_{K_{n_1}(A_1)} = (\hat{\mathbb{1}}^4_{A_1} \bullet \hat{f}_{D_0})$ and $\hat{f}_{\breve{K}_{n_2}(\breve{A}_2)} = (\hat{\mathbb{1}}^4_{\breve{A}_2} \bullet \hat{f}_{\breve{D}_0})$, respectively, their Fourier correlation is obtained by substituting (4.3.31) in (4.4.25), or alternatively, by directly applying CFT to (4.4.23) as

$$\hat{\mathfrak{G}}_{K_1,\breve{K}_2|_{\mathbb{r}}} = \big(\hat{\mathbb{1}}^4_{A_1} \bullet \hat{f}_{D_0}\big) \bullet \big(\bar{\hat{\mathbb{1}}}^4_{\mathbb{r}\breve{A}_2} \bullet \bar{\hat{f}}_{\breve{D}_0}\big) \tag{4.4.28}$$

$$= \big(\hat{\mathbb{1}}^4_{A_1} \bullet \bar{\hat{\mathbb{1}}}^4_{\mathbb{r}\breve{A}_2}\big) \bullet \big(\hat{f}_{D_0} \bullet \hat{f}_{D_0}\big) = (\hat{\mathbb{1}}^4_{A_O|_{\mathbb{r}}} \bullet \hat{f}_{D_O}), \tag{4.4.29}$$

where $\hat{f}_{D_O} = (\hat{f}_{D_0} \bullet \hat{f}_{D_0}) = (\bar{\hat{f}}_{\breve{D}_0} \bullet \bar{\hat{f}}_{\breve{D}_0})$, noting that $\hat{f}_{D_0} = \bar{\hat{f}}_{\breve{D}_0}$ as a result of the reflective duality $f_{D_0} = f_{-\breve{D}_0}$. In a similar fashion, $\hat{\mathbb{1}}^4_{A_O|_{\mathbb{r}}} = (\hat{\mathbb{1}}^4_{A_1} \bullet \bar{\hat{\mathbb{1}}}^4_{\mathbb{r}\breve{A}_2})$ is computed from a pointwise multiplication of the NDFTs over the knots A_1 and $\breve{A}_2$.

A critical observation is that the computational implementation relies only on the discrete knots P and A (or $\breve{A}$), expressed in the Fourier domain by the NDFTs $\hat{\mathbb{1}}^3_P$ in (4.3.30) and $\hat{\mathbb{1}}^4_A$ (or $\hat{\mathbb{1}}^4_{\breve{A}}$) in (4.3.32), respectively. The continuous geometry is completely embodied by the primitives that are implicit in $\hat{f}_{B_O} = (\hat{f}_{B_1} \bullet \hat{f}_{B_2})$ or $\hat{f}_{D_O} = (\hat{f}_{D_0} \bullet \hat{f}_{D_0}) = (\bar{\hat{f}}_{\breve{D}_0} \bullet \bar{\hat{f}}_{\breve{D}_0})$. However, despite appearing in equations, they do *not* explicitly participate into the numerical algorithms, and that reflects the true power of this particular discretization scheme.

4.4.3 Possible Extensions

Roerdink [291] generalized the concept of Minkowski sums/differences to Minkowski products/quotients over general groups, whose noncommutative convolutional formulation was presented by Lysenko et al. [218]. Nelaturi and Shapiro [235] applied the concept to SE(3) (in this context referred to as 'configuration products/quotients') and showed its applicability to direct and inverse $\mathsf{C}-$space problems ranging from

computing general sweeps to solving for maximal shapes and motions subject to containment constraints. The method embeds the solids in SE(3), and uses a uniform sampling over translations and rotations followed by pairwise matrix multiplications across the two samples to compute the $\mathsf{C}-$products and quotients. The presented non-unform sampling and spherical discretization scheme can readily be applied to more efficiently sample the translation space for different rotational sections through the 6D domain, as will be demonstrated in Section 4.5.1.

An interesting extension of the method would be to formulate 6D spherical sampling of the subsets of the Riemannian manifold SE(3) based on geodesic distances, whose potential benefits are unclear at this stage. One possible application is in machine tool path planning in the presence of tolerances [31], where the embedded workpiece complements the configuration product of the motion trajectory and tool profile. The tolerances can be introduced into either set by Minkowski operations between the 'nominal' geometry and primitive tolerance sets, e.g., Euclidean (for translational tolerances) and geodesic (for rotational tolerances) disks, cylinders, balls, or tori, all of which can be more efficiently discretized via spherical primitives than uniform samples. In a similar fashion to the constructions in Section 4.4.1, the tool's swept volume is then given by a 3D projection of the 6D Minkowski product of the embedded tool profile and its motion, each of which are described by a Minkowski product of sample points on their nominal sets with the primitive tolerance sets. Rearranging the terms (similar to (4.4.7) through (4.4.9)) abstracts the tolerances away into a 6D configuration space tolerance set, and allows working with lower-dimensional nominal sets only. Unfortunately, the corresponding Fourier analysis in this case becomes quite tedious, whose potential benefits are unclear at this stage. See the example in [31] for an elaborate discussion.

A more important research question remains regarding the extension of spherical sampling to dual operations, i.e., Minkowski differences or quotients, which would open up the opportunity to extend the benefits of this approach to inverse problems in configuration modeling.

4.5 Numerical Results

In this section I demonstrate how spherical sampling outperforms uniform sampling for Minkowski computations, that are central to the range of applications discussed in Chapter 5, and validate the additional performance improvement by using FFT algorithms. The method is implemented as a C++ API that reads triangular meshes, generates spherical decompositions (using Algorithm 1 in Section 4.6), converts the geometry to an analytic representation in the physical and/or frequency domains, and computes the correlations in either domain. I report on both CPU- and GPU-parallel computing, implemented using Boost C++ [305] and CUDA-C libraries, respectively.

The numerical experiments were conducted on a desktop computer with Intel® Xeon® E5-2687W CPU (32 cores, 3.10 GHz clock-rate, 64GB host memory) and NVIDIA® Tesla® K20c GPU (2,496 CUDA cores, 5GB device memory).

4.5.1 Combinatorial Advantage

Figure 4.5.1 illustrates the C−obstacle construction for a pair of objects discretized via nonequiradius spherical sampling developed in Section 4.3. Algorithm 1 is repeated with different grid sizes m, and compare the arithmetic complexity of the result with that of the uniform sampling of the same grid dimensions used in [235]. As described

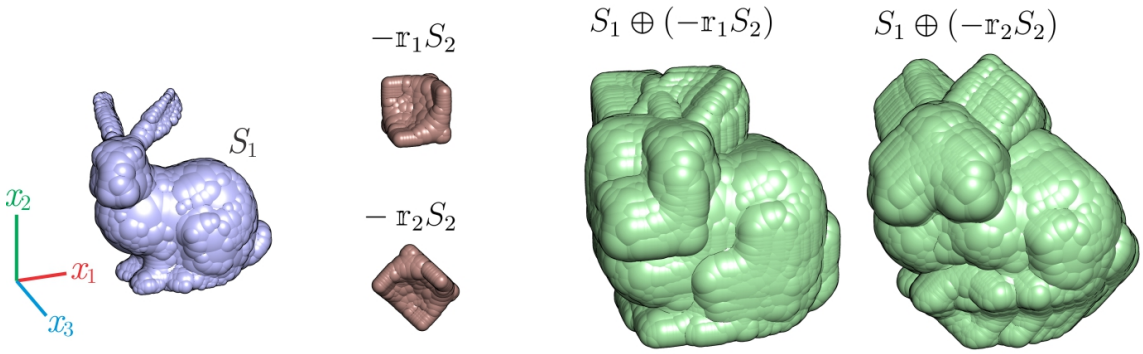

(a) $n_1 = 1{,}024$ (b) $n_2 = 340$ (c, d) $n = n_1 n_2 = 348{,}160$

FIGURE 4.5.1: The Minkowski sum of two spherical samples (a, b) for two sampled rotations, which are sections through the 6D Minkowski product, obtained from pairwise Minkowski sum of primitives (c, d).

TABLE 4.5.1: Comparison of the sample size between grid-based uniform and gird-free spherical samplings. The ratio scales rapidly with size.

	Uniform Sampling			Spherical Sampling			Ratio
m	n'_1	n'_2	$n' = n'_1 n'_2$	n_1	n_2	$n = n_1 n_2$	n'/n
2^{12}	666	44	29,304	49	26	1,274	23.0
2^{15}	5,921	689	4.08×10^6	159	83	13,197	309.1
2^{18}	49,981	3,867	1.93×10^8	1,024	340	3.48×10^5	764.5
2^{21}	409,058	36,874	1.5×10^{10}	3,686	1,081	3.98×10^6	3,785.5

in 4.6, this guarantees that the Hausdorff metric-based approximation error of the spherical discretization is upperbounded by that of the uniform sample, which is $\epsilon = \sqrt{3}(L/m^{\frac{1}{3}})$. Therefore, the "initial" grid size m will be used as a measure of resolution for both methods for the purpose of comparison.

As reported in Table 4.5.1, our method offers a clear advantage, decreasing the complexity by several orders of magnitude. Figure 4.5.2 (a, b) plots the memory requirement and running times, respectively, for Minkowski sum computation (viewed as a translational cross-section of the configuration product, by pairwise summations in the nonuniform 4D sample space, compared to pairwise summations in the uniform 3D sample space used in [235]. The speed-ups of our method scale significantly with

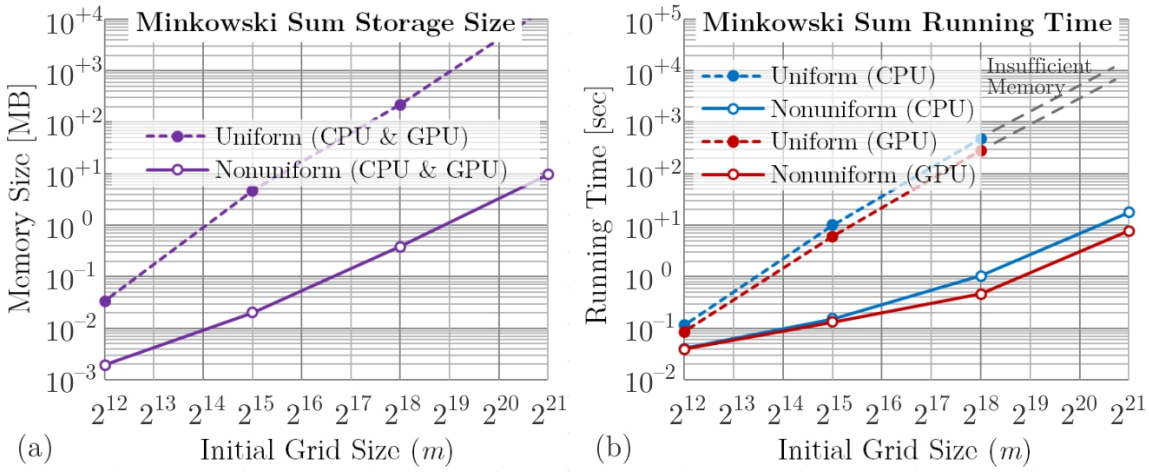

FIGURE 4.5.2: Nonuniform spherical sampling significantly outperforms uniform sampling, by efficient use of memory and time resources.

resolution, and reaches the range 400–600× (on both CPU and GPU)[12] for a grid size of $m := 2^{18} = 262{,}144$, decreasing the CPU/GPU running time from 28/48 seconds to 104/47 milliseconds. For larger initial grid sizes such as $m := 2^{21} = 2{,}097{,}152$, the memory cannot accommodate the Minkowski sum of uniform samples, while our method succeeds and carries out the sum in less than a second.

4.5.2 Analytic Advantage

We next consider the computational performance of the analytic method, using both uniform and nonuniform sampling. Given the spherical decomposition of the two solids, their bump functions as a sum of radial kernels can be rasterized on uniform grids in the physical domain. The gap function representation of the Minkowski sum can then be computed by two forward FFTs, a pointwise multiplication over the frequency grids, and an inverse FFT to retrieve the result, whose running times

[12] In each individual scenario, the GPU runs are only slightly faster than CPU runs (1–3×) due to extensive global memory references, but can be improved in future versions by memory optimization.

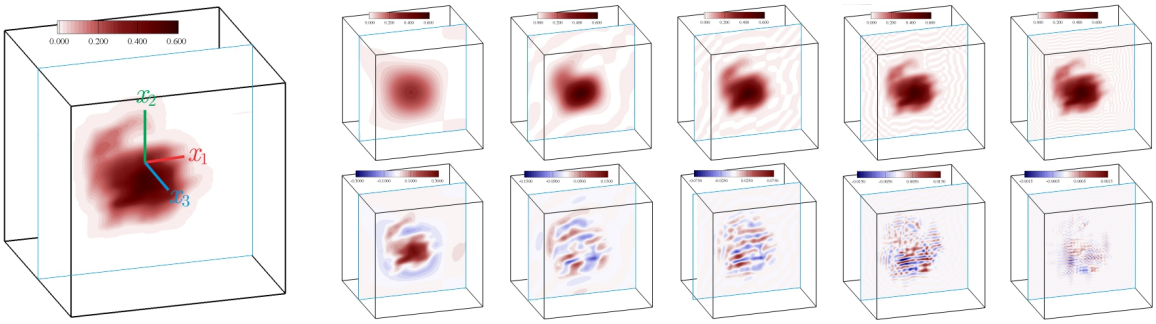

(a) $m'{=}m{=}128^3$ (b) $m'{=}4^3$ (c) $m'{=}8^3$ (d) $m'{=}16^3$ (e) $m'{=}32^3$ (f) $m'{=}64^3$

FIGURE 4.5.3: A section through the gap function representation (a) of the Minkowski sum in Fig. 4.5.1; its approximations (b–f) with truncated Fourier expansions (top), and the residual error (bottom). Uniform grid-based FFT implementation of the convolution outperforms the pairwise primitive multiplication method by two orders of magnitude.

are separately plotted in Fig. 4.5.3 (g). The GPU implementation in this case offers significant speed-ups of 400–800× over its CPU counterpart.[13] Comparing the results with Fig. 4.5.2 (b) shows an improvement of 30–80× over the pairwise computations in the physical domain. For a grid size of $m := 2^{21} = 2{,}097{,}152$, accurate computation of the convolution takes less than 80 milliseconds on the GPU.

It appears that the pointwise multiplication step is the bottleneck in the FFT-based convolutions. However, by performing this step (and the following inverse FFT) over a small subset of size $m' \ll m$ of the frequency grid in the neighborhood of the dominant modes, one could decide on the amount of computation time to spend in a trade-off with accuracy depicted in Fig. 4.5.3 (a–f). It is clear that small gap function errors do not necessarily imply small geometric discrepancies of the $0-$superlevel set in terms of Haudorff metric. However, it was shown by Lysenko [214] that it is also possible to impose upperbounds on the Hausdorff distance-based error as a function of the number of retained frequencies.

[13]The FFT is implemented using FFTW [124] on the CPU and using cuFFT(W) on the GPU. With the exception of FFTW, all other CPU and GPU routines were written in parallel.

As depicted in Sections 4.3 and 4.4, the uniform 3D grid-based FFT can be replaced with a nonuniform 4D grid-free NFFT. As the difference between the number of sample points in each method grows according to Table 4.5.1, even a cascade 4D NDFT over the spherical discretization can be faster, with the additional flexibility it offers in choosing the frequency domain grid size on-the-fly independently of the physical domain sample size. The NDFT does not require the additional step of bump function rasterization over the uniform grid, which is basically a cascade computation of the convolution of the knots and the conical kernel in (4.3.16). It rather incorporates that step as a pointwise multiplication with the kernel's frequency domain representation in (4.3.31) which can be precomputed to full precision. The comparison is shown in Fig. 4.5.4 for different number of modes m' over the 4D frequency grid, which demonstrates an advantage to NDFT for $m > 2^{15} = 32{,}768$ and $m' < 2^{16} = 65{,}536$ on the CPU, and for $m > 2^{18} = 262{,}144$ and $m' < 2^{12} = 4{,}096$ on the GPU. This can be further improved using the optimal NFFT [258]. Unfortunately, NFFTs have been implemented on the GPU [193] for 1, 2, and 3D, while at present the 4D NFFT is available only on the CPU [176].

Lastly, I test the performance for time-critical computation of the correlation predicate for a single configuration, via truncated frequency grid integration. Figure 4.5.5 shows sequential integration time on the CPU for different choices of the number of retained modes $m' \leq m$. An almost linear speed-up of m/m' (as expected from the theory) is achieved, and the collision predicate is computed in less than a millisecond for $m' < 2^{12} = 4{,}096$. This enables fast physically-based modeling and multibody dynamics simulations in real-time applications that require a refresh rate of 1 kHz for graphics and haptics feedback [362]. An important research question concerns the development of a "hybrid" method that further improves the performance by using

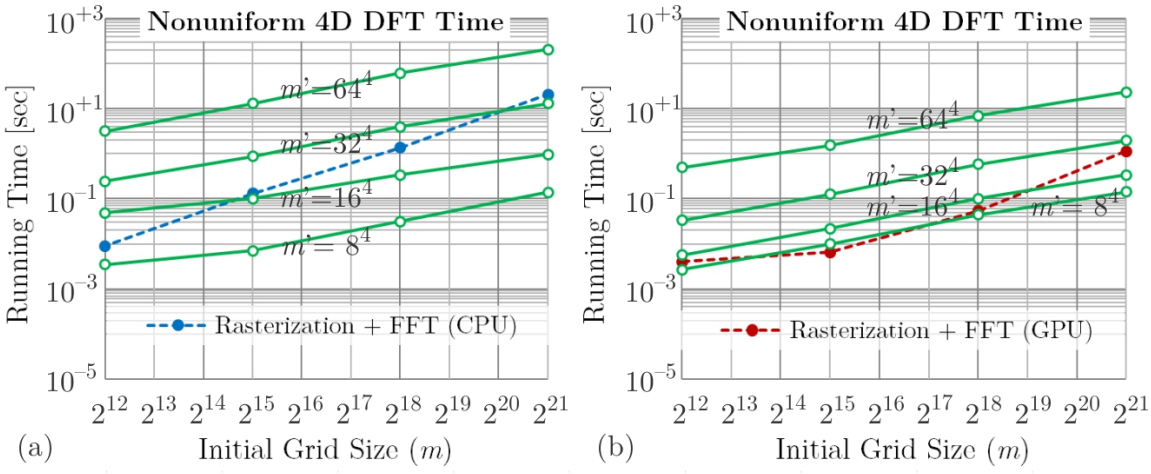

FIGURE 4.5.4: A comparison of 1) bump function rasterization (i.e., cascade convolution of knots with conical kernels) + 3D FFT; and 2) 4D NDFT + pointwise multiplication with kernel's frequency domain representation, for different numbers of needed frequencies $m' \leq m$.

this method alongside a sphere-tree traversal used in [48,156,248,362], and limits the integration over the NDFT of fewer primitives at the tree leaves.

4.6 A Sampling Algorithm

Although the main focus of this article is on how to work with given spherical decompositions in the analytic realm regardless of the method used for their generation [48, 156, 248, 362], a 'good' spherical sampling algorithm is essential to gain practical advantage over uniform sampling. Here I present the details of our new algorithm used to generate the decomposition in Table 4.5.1 and Figs. 4.5.1 and 4.5.2 of Section 4.5, along with its topological properties and geometric error bounds (Propositions 4.6.2 to 4.6.4). The algorithm is compared, in terms of these qualitative properties as well as arithmetic complexity, to the sphere-packing algorithm by Weller and Zachmann [362] (hereon abbreviated as W&Z), which is considered state-

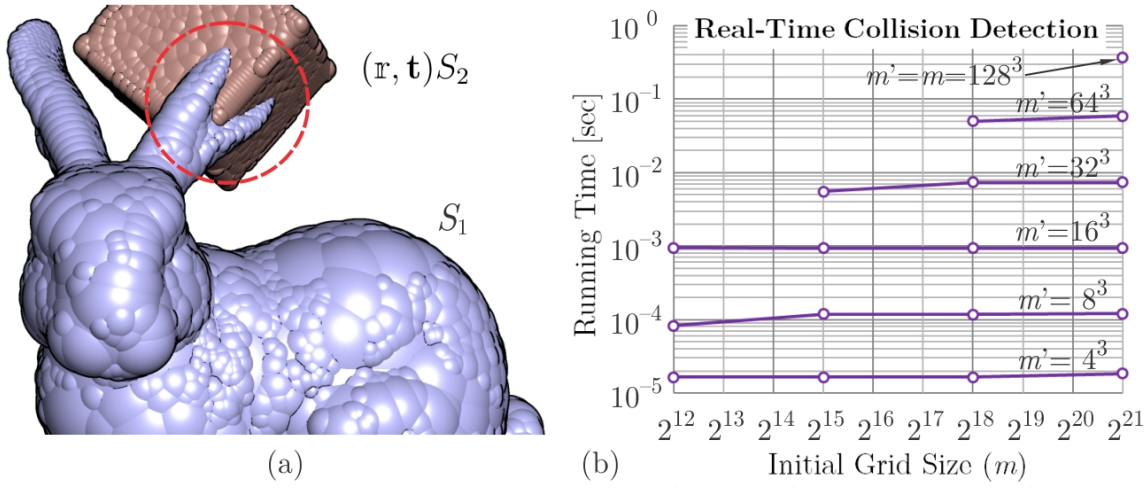

FIGURE 4.5.5: Time-critical collision detection for a single configuration query by integrating over the $m' \leq m$ dominant modes in the frequency domain. The sub-millisecond region is considered 'real-time'.

of-the-art for collision detection and proximity queries, and is successfully applied to real-time physically-based modeling and virtual reality (graphics/haptics) [362].

To approximate an arbitrary r-set $S \in \mathcal{P}^*(\mathsf{E})$ with a finite union of n (possibly overlapping) nonequiradius spherical balls $S_n(A) = \bigcup_{0 \leq i < n} B_{r_i}(\mathbf{x}_i)$, where the balls are encoded by $A = \{(\mathbf{x}_i, r_i)\}_{0 \leq i < n} \subset \mathbb{R}^4$, I use a greedy algorithm similar to W&Z [362]. One of the main differences is that W&Z uses the distance field as the greedy criterion, while our algorithm employs the SDF to create superior decompositions similar to the outcomes of MA-based methods [48,156] without the need to explicitly compute the numerically unstable MA/MAT. The SDF is a real-valued function that can be thought of as a 'dissipated' continuous extension of the indicator function of the MA (Fig. 4.6.1 (b)), whose ridges constitute the interior regions with extensive approximate nearest neighbors on the boundary $\partial(S)$, and is computed directly from the distance field [33].

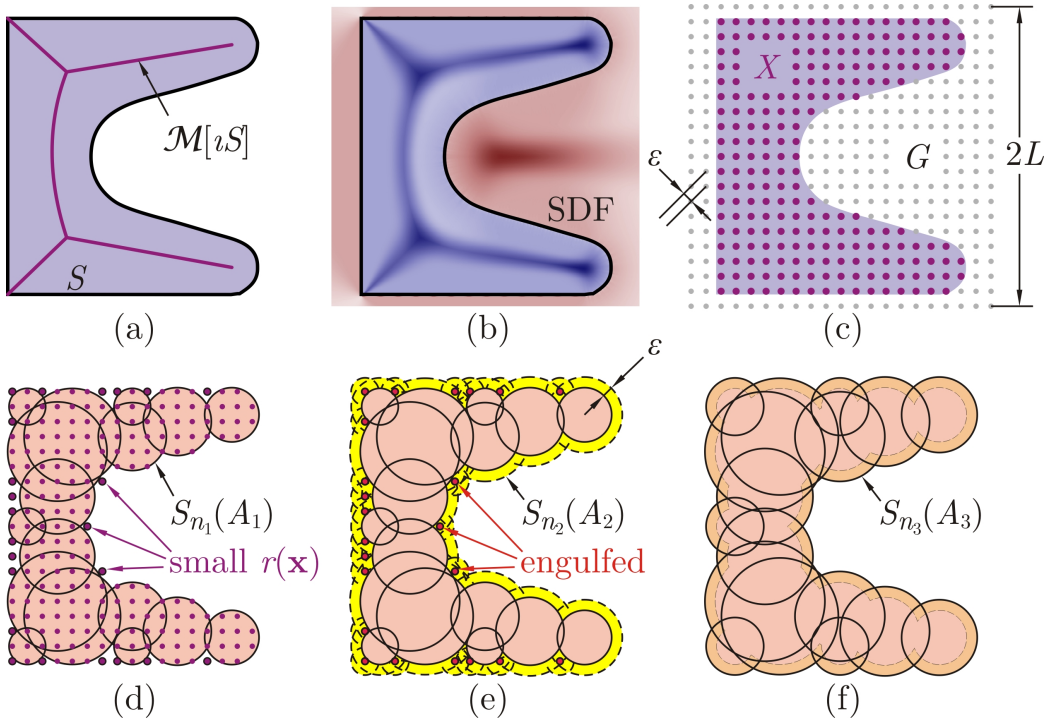

FIGURE 4.6.1: The various geometric constructs in Algorithm 1. It precomputes SDF on a grid G (a–c), then generates balls centered at grid nodes $X = G \cap S$ in descending SDF order to obtain A_1 (d), expands the balls by ϵ to obtain $A_2 = A_1 + (\mathbf{0}, \epsilon)$ (e), and cleans up the engulfed balls to obtain $A_3 \subset A_2$ (f).

Algorithm 1: The SDF-based spherical sampling algorithm B&I : $(S, G; L, \mu) \to (A_1, A_2, A_3)$.

Input: $S \in \mathcal{P}^*(\mathsf{E})$ (represented with any scheme that supports distance and inclusion queries), $G \subset \mathbb{R}^3$, $L > 0$, and $0 \leq \mu \leq 1$;
Output: $A_1, A_2, A_3 \subset \mathbb{R}^4$;

Step 0: Precomputations:
Constant $\epsilon \leftarrow \sqrt{3}(L/|G|^{\frac{1}{3}})$;
Using the inclusion test, separate $X = G \cap S$;
Compute the distance function over all X;
Compute and sort† the SDF over all X;
Step 1: Decomposition:
Initialize $Y \leftarrow \emptyset$; $A_1 \leftarrow \emptyset$;
while $|Y| < |X|$ (implies $Y \subset X$ but $Y \neq X$) **do**
  Select† $\mathbf{x} \in X$ with the maximal SDF;
  Modify $A_1 \leftarrow (A_1 \cup \{(\mathbf{x}, r(\mathbf{x}))\})$; //sampled
  for $\mathbf{x}' \in X$ **do**
    if $\mathbf{x}' \in B(\mathbf{x}, r(\mathbf{x}))$ **then**
      Modify* $Y \leftarrow (Y \cup \{\mathbf{x}'\})$; //covered
    if $\|\mathbf{x} - \mathbf{x}'\|_2 - |r(\mathbf{x}) - r(\mathbf{x}')| \leq \mu r(\mathbf{x}')$ **then**
      Modify† SDF$[\mathbf{x}'']$ $\leftarrow -\infty$; //popped
Step 2: Expansion:
Initialize $A_2 \leftarrow \emptyset$;
for $(\mathbf{x}_1, r_1) \in A_1$ **do**
  $A_2 \leftarrow (A_2 \cup \{(\mathbf{x}_1, r_1 + \epsilon)\})$; //expanded
Step 3: Reduction:
Initialize $A_3 \leftarrow A_2$;
for $(\mathbf{x}_1, r_1) \in A_1$ **do**
  for $(\mathbf{x}_2, r_2) \in A_2$ **do**
    if $\|\mathbf{x}_2 - \mathbf{x}_1\|_2 \leq (r_2 - r_1)$ **then**
      $A_3 \leftarrow (A_3 - \{(\mathbf{x}_1, r_1 + \epsilon)\})$; //popped
      break for;

//Implemented via a priority queue (†) and a binary array (*).

Step 0: Precomputations First, the algorithm precomputes the distance function and SDF over a uniform grid of m nodes sampled to cover the r-set's bounding volume. Without loss of generality, I assume a bounding box of edge length $2L > 0$ and a grid resolution of $2L/m^{\frac{1}{3}}$. Let $G = (2L/m^{\frac{1}{3}})\mathbb{Z}^3 \cap [-L, +L)^3$ with $|G| = m$ be the collection of grid nodes, and $X = (G \cap S)$ be the subset of the nodes that encodes a bitmap approximation of the r-set. The sample X gives an $\epsilon-$covering of S where $\epsilon = \sqrt{3}(L/m^{\frac{1}{3}})$ is the diagonal half-length of a grid cell, since for every point $\mathbf{x}' \in S$ there exists a node $\mathbf{x} \in X$ such that $\|\mathbf{x} - \mathbf{x}'\|_2 \leq \epsilon$ [235], which implies

Lemma 4.6.1. ($\epsilon-$Sampling Lemma) *The Hausdorff distance is upperbonded as* $\mathrm{d_H}(X, S) \leq \epsilon$, *for an* $\epsilon-$*sample* $X = (G \cap S)$ *where* G *is the set of nodes on a uniform grid* $G = (2L/m^{\frac{1}{3}})\mathbb{Z}^3 \cap [-L, +L)^3$ *and* $\epsilon = \sqrt{3}(L/m^{\frac{1}{3}})$.

Proof. By definition $X \subset S \subset (S \oplus B_\epsilon(\mathbf{0}))$. Additionally, $\|\mathbf{x} - \mathbf{x}'\|_2 \leq \epsilon$ for all $\mathbf{x} \in X$ and $\mathbf{x}' \in S$, hence $S \subset (X \oplus B_\epsilon(\mathbf{0}))$. Applying Lemma 3.4.1 of Section 3.4.1.1 yields the result. The equality holds (i.e., $\mathrm{d_H}(X, S) = \epsilon$) unless none of the grid cell centers in between the nodes belongs to S. □

I will show that the same error bound holds for the spherical decomposition, which forms a basis of comparison in Section 4.5 between the uniform sample and the spherical sample, in terms of their deviation from the original solid measured by ϵ (or equivalently, the initial grid size m).

Step 1: Decomposition The algorithm sorts X with respect to the SDF values and arranges the nodes into a priority queue, followed by n rounds of selection. At each round, the grid node $\mathbf{x} \in X$ with maximal SDF is selected, and a ball $B(\mathbf{x}, r(\mathbf{x}))$

is generated, i.e., the 4−tuple $(\mathbf{x}, r(\mathbf{x}))$ is added to the (initially empty) output set A_1, where $r(\mathbf{x}) = \mathrm{d}(\mathbf{x}, \mathcal{B}(S)) = \inf_{\mathbf{x}' \in \mathcal{B}(S)} \|\mathbf{x} - \mathbf{x}'\|_2$ is the distance function. A subset of the grid points inside this new ball are identified with the following condition, and eliminated from the SDF-sorted queue for the next rounds:

$$Q(\mathbf{x}; \mu) = \{\mathbf{x}' \in X \mid \|\mathbf{x} - \mathbf{x}'\|_2 - |r(\mathbf{x}) - r(\mathbf{x}')| \leq \mu r(\mathbf{x}')\}\,. \tag{4.6.1}$$

The left-hand side of the inequality gives a measure of how $B_{r(\mathbf{x}')}(\mathbf{x}')$ protrudes outside $B_{r(\mathbf{x})}(\mathbf{x})$, which is normalized by the smaller radius $r(\mathbf{x}')$ and compared to the 'protrusion factor' $0 \leq \mu \leq 1$. On the one end, $Q(\mathbf{x}; 0)$ includes only the grid nodes $\mathbf{x}' \in (X \cap B_{r(\mathbf{x})}(\mathbf{x}))$ with $B_{r(\mathbf{x}')}(\mathbf{x}') \subset B_{r(\mathbf{x})}(\mathbf{x})$, i.e., nodes that would contribute redundant balls already covered. On the other end, $Q(\mathbf{x}; 1)$ includes the entire $(X \cap B_{r(\mathbf{x})}(\mathbf{x}))$ and eliminates all covered points from the future rounds. The former gives a more conservative coverage and creates better approximations to the original shape in terms of topology (e.g., fewer void spaces) and geometry (e.g., similar surface curvature) in practice. However, it generates a larger number of balls than the latter—especially over the thinner features of the shape—thus takes longer to finish. More importantly, a too small choice (e.g., $\mu < 0.1$) might defeat the advantage of spherical sampling with $|A_1| = n_1$ over uniform sampling with $|X| = O(m)$ reliant on the assumption $n_1 \ll m$. Our experiments suggest a choice of $\mu = 0.25$–0.30 to be a good trade-off—see Table 4.6.1 for experimental results to support this.

At each round of the algorithm, all $\mathbf{x}' \in B_{r(\mathbf{x})}(\mathbf{x})$ are marked by Boolean flags that indicate their containment in at least one ball. The algorithm repeats this process until all points in the uniform sample X are covered by the balls, i.e., until $Y = X$ where $Y = (G \cap S_{n_1}(A_1))$ is the set of grid nodes covered so far.

TABLE 4.6.1: Arithmetic complexities of spherical samples generated by our SDF-based algorithm (B&I) using different protrusion factors, compared to those of W&Z [362].

Initial Grid		B&I, $\mu = 0.25$		B&I, $\mu = 0.50$		B&I, $\mu = 1.00$		W&Z
m	ϵ	$n_{1,2}$	n_3	$n_{1,2}$	n_3	$n_{1,2}$	n_3	n
2^{12}	0.0541	318	36	307	53	275	67	334
2^{15}	0.0271	1,552	161	1,445	251	1,295	324	1,024
2^{18}	0.0135	6,842	1,024	6,374	1,048	5,436	1,311	15,592
2^{21}	0.0068	26,217	3,686	23,891	4,422	20,674	5,160	89,030

Proposition 4.6.2. $\mathbb{d}_{\mathrm{H}}(S_{n_1}(A_1), S) \leq \epsilon$.

Proof. At the end of step 1, we have $X \subset S_{n_1}(A_1)$ hence for every point $\mathbf{x}' \in S$, $\mathbb{d}(\mathbf{x}', S_{n_1}(A_1)) \leq \mathbb{d}(\mathbf{x}', X)$, and since $\mathbb{d}(\mathbf{x}', X) \leq \epsilon$, $S \subset (S_{n_1}(A_1) \oplus B_\epsilon(\mathbf{0}))$. On the other hand, choosing the minimum distance to the boundary as ball radii guarantees $S_{n_1}(A_1) \subset S \subset (S \oplus B_\epsilon(\mathbf{0}))$. Applying Lemma 3.4.1 completes the proof. □

It is worthwhile noting that the above proof only makes use of $G \subset S_{n_1}(A_1) \subset S$, hence Proposition 4.6.2 is valid for the W&Z algorithm [362] as well, providing a basis of comparison between the two (see Table 4.6.1).

An inevitable difficulty with this approach, regardless of the greedy criterion (e.g., distance-based [362] or SDF-based) are topological discrepancies between the solid and its spherical approximation due to the possible void spaces (of feature sizes no larger than 2ϵ) left in between covered grid nodes. Although letting $m \gg 1$ and $\mu \ll 1$ can alleviate the problem, it does not necessarily eliminate it especially near the boundary. In addition to the trapped internal cavities, the algorithm generates numerous small disconnected balls near the boundary where $r(\mathbf{x}) < \epsilon$ as it approaches the final rounds. The next two steps are meant to solve these problems.

Step 2: Expansion The previous step generates a decomposition that is strictly contained in the original solid, which is unnecessary for our purposes. In the next step, the algorithm expands all the balls in the previous step by increasing their radii (i.e., offsetting them) with the constant $\epsilon = \sqrt{3}(L/m^{\frac{1}{3}})$, which gives:

$$A_2 = \big\{\mathbf{a}_2 = (\mathbf{x}_1, r_1 + \epsilon) \mid \mathbf{a}_1 = (\mathbf{x}_1, r_1) \in A_1\big\}, \tag{4.6.2}$$

i.e., $S_{n_2}(A_2) = (S_{n_1}(A_1) \oplus B_\epsilon(\mathbf{0}))$. It is easy to show that $S_{n_2}(A_2)$ is in fact as good of an approximation as $S_{n_1}(A_1)$ in terms of Hausdorff error measure.

Proposition 4.6.3. $\mathbb{d}_{\mathrm{H}}(S_{n_2}(A_2), S) \leq \epsilon$.

Proof. On the one hand we have $S_{n_1}(A_1) \subset S \subset (S \oplus B_\epsilon(\mathbf{0}))$. On the other hand we have $S \subset (S_{n_1}(A_1) \oplus B_\epsilon(\mathbf{0})) \subset (S_{n_2}(A_2) \oplus B_\epsilon(\mathbf{0}))$. Once again, applying Lemma 3.4.1 completes the proof. □

Moreover, the $\epsilon-$expansion "repairs" the inter-cellular cavities and disconnected balls, noting that the distance between every pair of balls corresponding to A_1 is upperbounded by 2ϵ. Therefore, although the algorithm cannot guarantee topological equivalence (i.e., homeomorphism) between S and any decomposition, simply because it only relies on the incomplete shape representation provided by the discrete set X—which fails to capture topological information of the features that are smaller than the sampling resolution—$S_{n_2}(A_2)$ does guarantee a weaker equivalence, which is homeomorphism with a *voxelization* of the shape, implied by the uniform sample X, while $S_{n_1}(A_1)$ does not.

Step 3: Reduction Although $S_{n_2}(A_2)$ exhibits relatively more desirable topological properties than $S_{n_1}(A_1)$, they both have the same complexity $n_1 = n_2$. The former satisfies 'conservative coverage' ($S_{n_2}(A_2) \supset S$) while the latter fulfills 'strict containment' ($S_{n_1}(A_1) \subset S$). The main purpose of the expansion, however, was to eliminate the small balls at the grid nodes near the surface with $r(\mathbf{x}) < \epsilon$ that could not be covered by any other ball in A_1. These artifacts can comprise a significant fraction of all balls, adding to the numerical complexity with little contribution to the volume, while they adversely affect the quality of discretization by introducing high local curvatures. The question is whether I can obtain a third discretization in between the two, i.e., a sample A_3 such that

$$S_{n_1}(A_1) \subset S_{n_3}(A_3) \subset S_{n_2}(A_2), \tag{4.6.3}$$

which retains the desired topological properties, eliminates the undesired small balls, guarantees the same error bounds, and is smaller in size ($|A_3| = n_3 < n_{1,2}$).

The next step of the algorithm reduces A_2 into A_3 by eliminating the sample points whose corresponding original (i.e., not expanded) balls encoded by the tuples $(\mathbf{x}_1, r_1) \in A_1$ were contained in some other expanded ball corresponding to the tuples $(\mathbf{x}'_1, r'_1 + \epsilon) \in A_2$, $(\mathbf{x}_1, r_1) \neq (\mathbf{x}'_1, r'_1)$:

$$\begin{aligned} A_3 = \big\{ \mathbf{a}_2 = (\mathbf{x}_1, r_1 + \epsilon) \mid \mathbf{a}_1 = (\mathbf{x}_1, r_1) \in A_1, \\ B_{r_1}(\mathbf{x}_1) \not\subset S_{n_2-1}(A_2 - \{\mathbf{a}_2\}) \big\}, \end{aligned} \tag{4.6.4}$$

A significant number of small balls in A_1 near the surface are thus 'engulfed' by the expanded neighbor balls in A_2—e.g., constituting $\%75 - \%85$ of all balls for the

Stanford Bunny in Fig. 4.6.2 with $\mu = 0.25$–1.00, which grows by decreasing μ. Although the new decomposition neither contains nor is contained in the original solid, the approximation error bound still holds.

Proposition 4.6.4. $\mathbb{d}_{\mathrm{H}}(S_{n_3}(A_3), S) \leq \epsilon$.

Proof. By definitions in (4.6.2) and (4.6.4), $A_3 \subset A_2$ from which it becomes lear that $S_{n_3}(A_3) \subset S_{n_2}(A_2) \subset (S \oplus B_\epsilon(\mathbf{0}))$. On the other hand, every ball encoded by $(\mathbf{x}_1, r_1) \in A_1$ is contained in at least one ball corresponding to $(\mathbf{x}_2, r_2 + \epsilon) \in A_2$. If this is realized only by $(\mathbf{x}_2, r_2) := (\mathbf{x}_1, r_1)$, noting that always $B_{r_1}(\mathbf{x}_1) \subset B_{r_1+\epsilon}(\mathbf{x}_1)$, then its $\epsilon-$expansion is included in A_3 by definition in (4.6.4). Otherwise, there exists at least another ball represented by $(\mathbf{x}_2, r_2) \neq (\mathbf{x}_1, r_1)$ whose $\epsilon-$expansion is included in A_3, i.e., $(\mathbf{x}_2, r_2 + \epsilon) \in A_3$, and $B_{r_1}(\mathbf{x}_1) \subset B_{r_2+\epsilon}(\mathbf{x}_2)$. In either case, $B_{r_1}(\mathbf{x}_1)$ is included in $S_{n_3}(A_3)$, hence $S_{n_1}(A_1) \subset S_{n_3}(A_3)$. Therefore, $S \subset (S_{n_1}(A_1) \oplus B_\epsilon(\mathbf{0}))$ implies $S \subset (S_{n_3}(A_3) \oplus B_\epsilon(\mathbf{0}))$. Applying Lemma 3.4.1 completes the proof. □

The 3 aforementioned steps—namely, decomposition, expansion, and reductions—are summarized in Algorithm 1. Figure 4.6.1 illustrates the different geometric objects involved in the algorithm, including the output decompositions of each step. Once the decomposition is generated, a variety of methods can be used to create sphere-tree hierarchies—e.g., the simple bottom-up merging algorithm in [156] that groups and merges a pre-specified number of balls that are near each other in A_3 (i.e., the children) into their smallest enclosing ball (i.e., the parent).

Table 4.6.1 compares the numerical complexity of each decomposition $|A_1| = |A_2|$, and $|A_3|$ with that of W&Z [362] for the Stanford Bunny in Fig. 4.5.1 (a) using different grid resolutions and protrusion factors. It is observed that unlike the case for $n_1 = n_2$, n_3 decreases with more conservative choices of μ, making the smaller

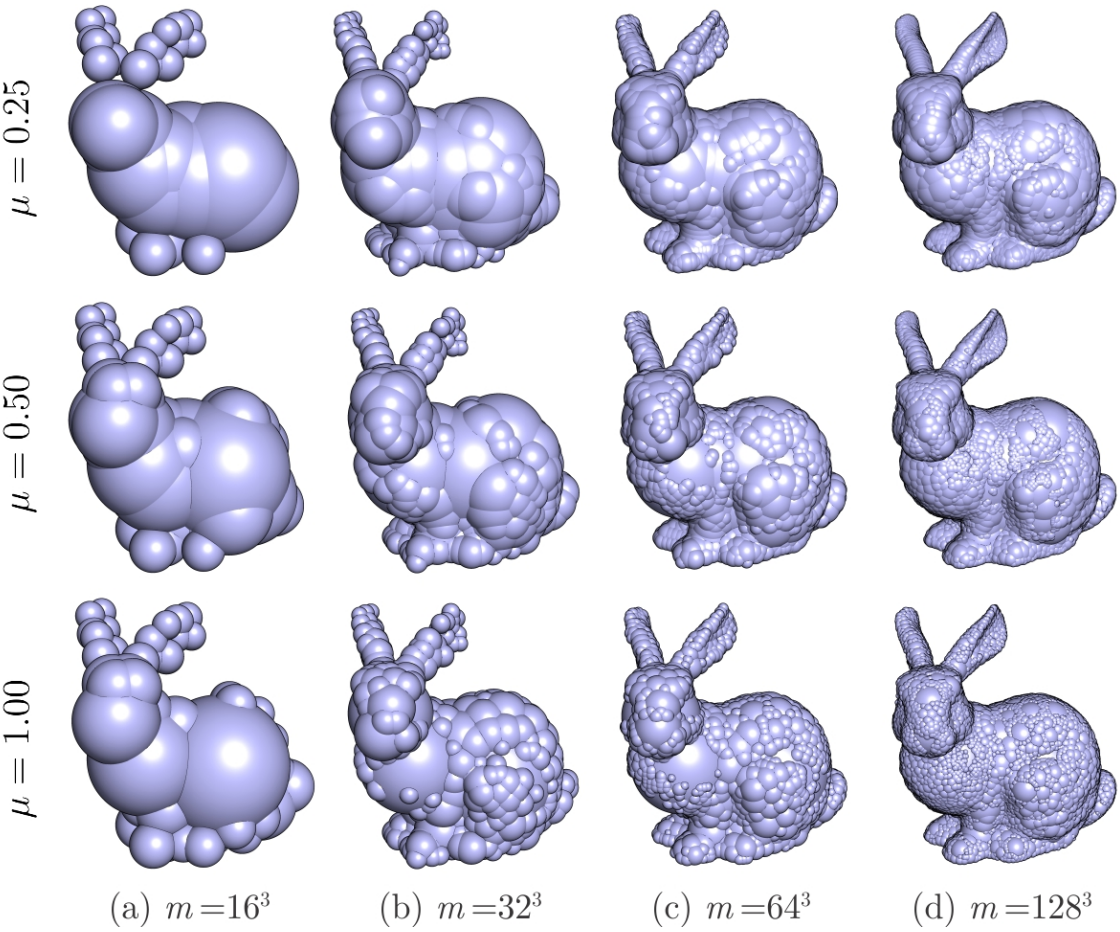

FIGURE 4.6.2: The SDF-based spherical decomposition $S_{n_3}(A_3)$ for different initial grid sizes and protrusion factors in Table 4.6.1.

value of $\mu = 0.25$ even more favorable. The same value was used to generate the results in Table 4.5.1 and Figs. 4.5.1 and 4.5.2 of Section 4.5. After this reduction, the output of our algorithm contains fewer balls than W&Z [362] by an order of magnitude, allowing faster termination—in a few minutes for the slowest case. Figure 4.6.2 illustrates $S_{n_3}(A_3)$ corresponding to the parameters in Table 4.6.1.

Chapter 5

Engineering Applications

5.1 Fuzzy Geometric Modeling

This chapter presents some important engineering applications to demonstrate how the analytic (i.e., functional) paradigm developed in the previous chapters allows one to develop powerful software tools that can transform computer-aided design (CAD) in addressing a wide class of shape and motion related problems.

To recap the development, in Chapter 2 I described a paradigm shift from explicit to implicit (i.e., functional) abstractions. I presented functional shape and motion descriptors ranging from ordinary (i.e., nonsingular) real-valued defining functions, whose $0-$superlevel sets correspond to the geometric pointsets, to $\delta-$singular functional descriptors of various orders for embedded lower-dimensional strata. In chapter 3 I used these novel abstractions to re-conceptualize geometric operations in terms of functional combining forms rooted in measure theory. These two undertakings were illustrated by the blue and red arrows below, respectively:

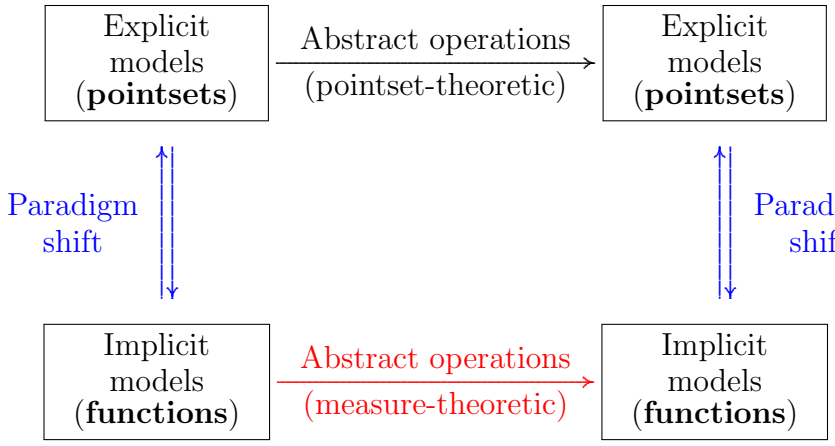

(5.1.1)

In Chapter 4, the computational implementation of the above formalism was presented for ordinary functions that can be sampled uniformly or non-uniformly:

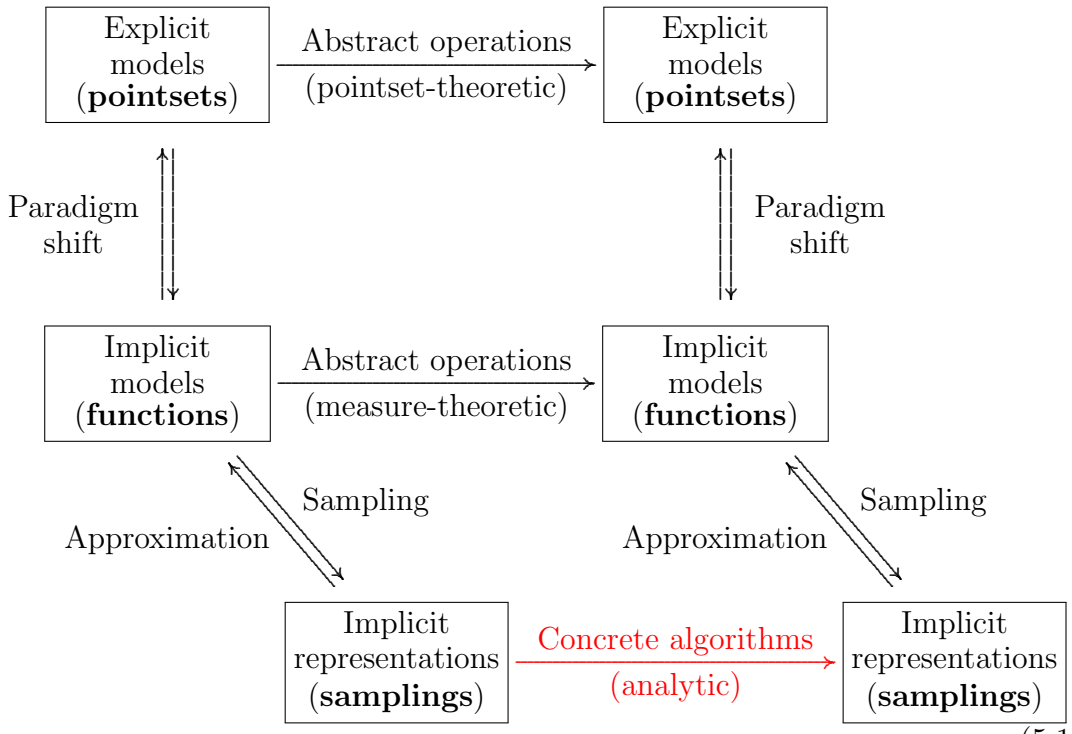

(5.1.2)

in which the interpretation of the samples (e.g., spherical decomposition [36]) returns an approximation of the sampled abstract objects. However, another approximation scheme is necessary to implement the conceptual functional descriptors that use the $\delta-$function, which is naturally given by limit representations (i.e., 'nascent' $\delta-$functions) presented in Chapter 2:

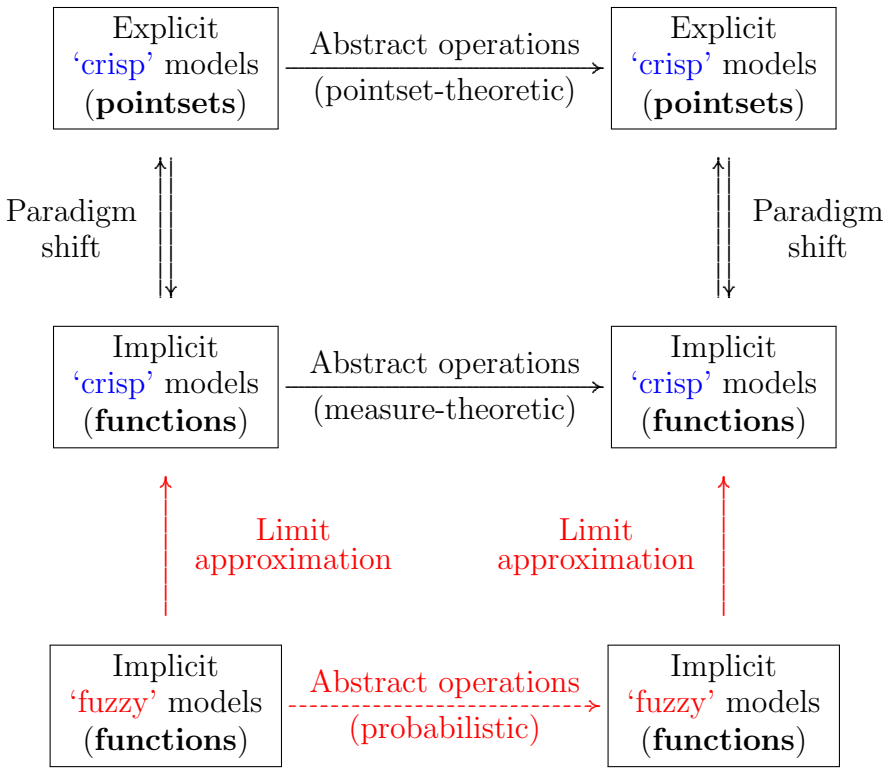

(5.1.3)

There are at least two fundamental advantages with this approach. First, it reveals the theoretical limitations in terms of computability and robustness [104–106] of geometric operations that involve lower-dimensional features. For example:

- The contact measure (e.g., surface area or length) between solids that intersect over a lower-dimensional region is non-computable [217], since small perturbations of the boundary geometry (e.g., due to errors/noise) can lead to substantial

changes in the outcome, as I discussed in multiple earlier occasions.

- The similarity and complementarity of solids, measured in terms of overlaps of internal and external skeletons or medial axes (MA) of the objects is non-computable for a similar reason [38].

- The latter is not to mention that the MA itself is non-computable [17], also because it is formulated in terms of degenerate lower-dimensional intersections of a maximal ball with the boundary.

Importantly, when lower-dimensional features are modeled implicitly as the limit representations of nascent $\delta-$functions, such limitations are clarified and acknowledged with rigorous formalism. Thus by substituting them with 'dissipated' density functions such as pulse function, Gaussian function, bump function (i.e., 'mollifier'), and others (e.g., selected from Table 2.3.1), one can decide on the trade-off between fidelity of the computation to the abstraction in the limit and its robustness against perturbations. In a sense, this approach can be viewed as a 'dissipation' of the lower-dimensional geometry, to an adjustable extent, to the adjacent 3D volume.

The second advantage of this approach is that it provides a basis for *probabilistic* geometric modeling. If the dissipation of the lower-dimensional features is viewed as the model of uncertainty—e.g., (proportional to) the probability distribution function (PDF) of a point's membership in the feature—then one may generalize 'crisp' geometric modeling to 'fuzzy' functional modeling (borrowing terminology from [219]) if one can reason about the corresponding analytic operations in terms of the PDF of the outcome of geometric operations. The key is to prove that these operations, illustrated via the lower horizontal arrow in (5.1.3) do in fact correspond to meaningful probabilistic interpretations. For example:

- The sweep/unsweep of the crisp 3D solid along a crisp 1D or 2D trajectory, modeled respectively via $2^{\text{nd}}-$ and $1^{\text{st}}-$order $\delta-$singularities, can be extended to those of a crisp or fuzzy 3D solid along a fuzzy trajectory, using convolutions of density functions based on the nascent $\delta-$functions.

- The $\mathsf{C}-$space obstacle and free space computation for crisp shapes and obstacles can be extended to those of uncertain shapes and obstacles, using convolutions of density function based on the nascent $\delta-$functions.

- To generalize the latter, measures of collision, contact, similarity, complementarity, and symmetry between two shapes for different configurations can be extended to take uncertainty into account, using cross-correlation of density functions based on the nascent $\delta-$functions.

The important realization is that the fundamental operations presented in Chapter 2 to express explicit Boolean operations in terms of implicit arithmetic operations—namely, unions as pointwise summations and intersections as pointwise multiplications of defining functions—extend naturally to the probabilistic realm (replacing defining functions with PDFs), when the membership probability/grade in the two sets conceptualized in terms of these PDFs are viewed as *independent* events. These so-called summation and multiplication laws of probability naturally extend to convolutions, and much of what I presented in the previous chapters remains valid from the viewpoint of probabilistic interpretations.

Once the special $\delta-$singular functional descriptors are replaced with ordinary density functions, the implementation methods based on uniform or non-uniform sampling discussed in Chapter 4 can be applied for numerical computation:

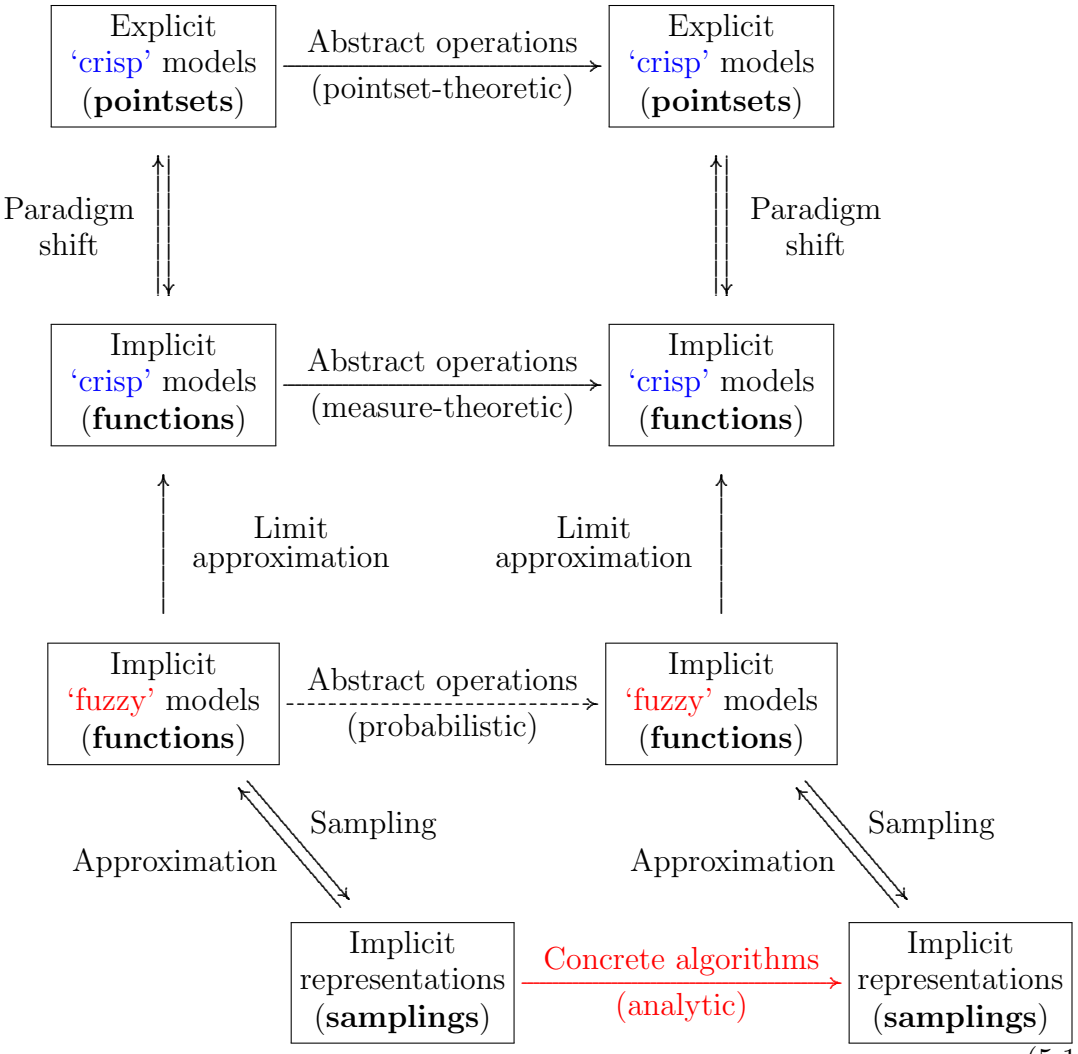

(5.1.4)

After presenting some of the most useful density functions in Section 5.1.1, a few engineering applications that can benefit from their functional combinations are discussed in the rest of the chapter; namely:

- In Section 5.2 I show the applications of cross-correlations of density functions in formulating contact and collision measures for real-time applications.

- In Section 5.3 I discuss the extension of the correlation paradigm to detecting similarities and complementarities between objects of arbitrary shape in different relative configurations, useful in mechanical assembly and protein folding.

- In Section 5.4 I overview an important application of the correlation paradigm to virtual reality systems with graphics- and haptics-enabled assembly guidance.

- In Section 5.5 I show how morphological shape synthesis (e.g., Minkowski operations, sweeping and unsweeping, etc.) can be extended by customizing the density functions of the participating objects using local concentrations, which, in turn, allow one to deform the boundaries of resultant superlevel sets.

- In Section 5.6 I show how manufacturability analysis and feedback for design correction can be addressed, with example from subtractive and additive manufacturing (SM/AM).

- I conclude with remarks on how this thesis can provide the foundations for query-based geometric algorithm interoperability proposed in [147, 148].

5.1.1 Density Functions

In this section I will review some of the useful density functions, produced via nascent $\delta-$functions with $\sigma > 0$, and their interpretations in the limit as $\sigma \to 0^+$.

5.1.1.1 Membership Density Functions

The notions of membership density function (MDF)—or membership grade functions as per terminology in [219]—is a natural generalization of the point membership classi-

fication (PMC). The 'dissipated' interior/exterior of an arbitrary shape can be characterized implicitly in a variety of ways. Using the indicator function (i.e. membership), one can define the internal/external membership density functions (I-MDF/E-MDF) $\wp^{\mathsf{I\text{-}MDF}}_{\sigma,S}, \wp^{\mathsf{E\text{-}MDF}}_{\sigma,S} : \mathsf{E} \to \mathbb{R}$ as:

$$\wp^{\mathsf{I\text{-}MDF}}_{\sigma,S}(\mathbf{x}) := (\mathbb{1}^0_{i(S)} * \mathscr{g}^3_\sigma)(\mathbf{x}) = \int_{\mathsf{E}} \mathbb{1}^0_{i(S)}(\mathbf{x}') \; \mathscr{g}^3_\sigma(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}'], \tag{5.1.5}$$

$$\wp^{\mathsf{E\text{-}MDF}}_{\sigma,S}(\mathbf{x}) := (\mathbb{1}^0_{e(S)} * \mathscr{g}^3_\sigma)(\mathbf{x}) = \int_{\mathsf{E}} \mathbb{1}^0_{e(S)}(\mathbf{x}') \; \mathscr{g}^3_\sigma(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}']. \tag{5.1.6}$$

Noting that $\mathbb{1}^0_{i(S)} \overset{*}{\equiv} \mathbf{1}_S$ and $\mathbb{1}^0_{e(S)} \overset{*}{\equiv} \mathbf{1}_{c(S)} \equiv \neg\mathbf{1}_S$, one obtains the simpler formulae:

$$\wp^{\mathsf{I\text{-}MDF}}_{\sigma,S}(\mathbf{x}) = (\mathbf{1}_S * \mathscr{g}^3_\sigma)(\mathbf{x}) = \int_{\mathsf{E}} \mathbf{1}_S(\mathbf{x}') \; \mathscr{g}^3_\sigma(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}'], \tag{5.1.7}$$

$$\wp^{\mathsf{E\text{-}MDF}}_{\sigma,S}(\mathbf{x}) = (\mathbf{1}_S \diamond \mathscr{g}^3_\sigma)(\mathbf{x}) = \int_{\mathsf{E}} \neg\mathbf{1}_S(\mathbf{x}') \; \mathscr{g}^3_\sigma(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}'], \tag{5.1.8}$$

where the set membership $\mathbf{1}_S = (\mathbf{x} \overset{?}{\in} S)$ is a primitive query, which is reasonable to expect from every geometric modeling system to deliver with universally accepted semantics. The convolution with the $\mathscr{g}_\sigma-$kernel can represent offsetting by a thin tolerance zone $\propto O(\sigma)$, Gaussian blurring, bump mollification, etc. whose limit as $\sigma \to 0^+$ is the $\delta-$regularization of the interior/exterior as depicted in Section 3.3.2.1, i.e., the I-MDF/E-MDF are the nascent implicit forms of the closure of interior/exterior:

$$\mathbf{1}_{(\mathbb{k}\circ i)(S)}(\mathbf{x}) = \lim_{\sigma\to 0^+} \wp^{\mathsf{I\text{-}MDF}}_{\sigma,S} = (\mathbf{1}_S * \delta^3)(\mathbf{x}) = \int_{\mathsf{E}} \mathbf{1}_S(\mathbf{x}') \; \delta^3(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}'], \tag{5.1.9}$$

$$\mathbf{1}_{(\mathbb{k}\circ e)(S)}(\mathbf{x}) = \lim_{\sigma\to 0^+} \wp^{\mathsf{E\text{-}MDF}}_{\sigma,S} = (\mathbf{1}_S \diamond \delta^3)(\mathbf{x}) = \int_{\mathsf{E}} \neg\mathbf{1}_S(\mathbf{x}') \; \delta^3(\mathbf{x} - \mathbf{x}') \; d\mu^3[\mathbf{x}']. \tag{5.1.10}$$

One can use a custom distribution for the I-MDF/E-MDF by substituting the indicator functions with an arbitrary nonnegative real-valued defining function, which is the same as multiplying a custom strictly positive weight function $\rho : \mathsf{E} \to (0, +\infty)$:

$$\mathscr{f}_{\sigma,S}^{\mathsf{I\text{-}MDF}} := \big(\rho \bullet (\mathbf{1}_S * \mathscr{g}_\sigma^3)\big) = (f_S * \mathscr{g}_\sigma^3) \equiv (\mathbf{1}_S * \mathscr{g}_\sigma^3), \tag{5.1.11}$$

$$\mathscr{f}_{\sigma,S}^{\mathsf{E\text{-}MDF}} := \big(\rho \bullet (\mathbf{1}_S \diamond \mathscr{g}_\sigma^3)\big) = (f_S \diamond \mathscr{g}_\sigma^3) \equiv (\mathbf{1}_S * \mathscr{g}_\sigma^3). \tag{5.1.12}$$

Alternatively, one could define the I-MDF/E-MDF in terms of the distance function:

$$\mathscr{f}_{\sigma,S}^{\mathsf{I\text{-}MDF}} := \frac{\rho(\cdot)}{O(\sigma)} \bullet \Big(\mathscr{g}_\sigma^1 \circ \min\big\{0, \mathbb{d}_S(\cdot)\big\}\Big) \equiv \frac{\rho(\cdot)}{O(\sigma)}\Big(\mathscr{g}_\sigma^1 \circ \min\big\{0, \mathbb{d}_S(\cdot)\big\}\Big), \tag{5.1.13}$$

$$\mathscr{f}_{\sigma,S}^{\mathsf{E\text{-}MDF}} := \frac{\rho(\cdot)}{O(\sigma)} \bullet \Big(\mathscr{g}_\sigma^1 \circ \max\big\{0, \mathbb{d}_S(\cdot)\big\}\Big) \equiv \frac{\rho(\cdot)}{O(\sigma)}\Big(\mathscr{g}_\sigma^1 \circ \max\big\{0, \mathbb{d}_S(\cdot)\big\}\Big), \tag{5.1.14}$$

where the denominators $\mathscr{g}_\sigma^1(0) = O(\frac{1}{\sigma})$ are meant to normalize the growth of the nominators to infinity for $\mathbb{d}_S(\mathbf{x}) < 0$ (interior) and $\mathbb{d}_S(\mathbf{x}) > 0$ (exterior). In this case, the signed distance function $\mathbb{d}_S(\mathbf{x}) = \big[1 - 2(\mathbf{1}_S(\mathbf{x}))\big] \inf_{\mathbf{x}' \in S} \|\mathbf{x} - \mathbf{x}'\|_2$ is taken as a primitive query whose support is expected from most geometric modelers, in compliance with fairly universal semantics.[1]

5.1.1.2 Boundary Density Functions

The boundary density function (BDF) $\mathscr{f}_{\sigma,S}^{\mathsf{BDF}} : \mathsf{E} \to \mathbb{R}$ can be defined as:

$$\mathscr{f}_{\sigma,S}^{\mathsf{BDF}} := \big(\rho \bullet (\mathscr{g}_\sigma^1 \circ \mathbb{d}_S)\big) = \big(\rho \bullet (\mathscr{g}_\sigma^1 \circ |\mathbb{d}_S|)\big) = \big(\rho \bullet (\mathscr{g}_\sigma^1 \circ \mathbb{d}(S, \cdot))\big), \tag{5.1.15}$$

[1] The sign convention is not universal; some systems may assign internal/external query points with $-/+$ signs (this thesis' convention) or vice versa, in which case the min/max in the above equations need to be swapped.

whose limit as $\sigma \to 0^+$ yields the $\delta-$singular boundary descriptor presented in Section 3.3.2.2. Once again, for finite $\sigma > 0$, the above function can be viewed as an $O(\sigma)-$offset (i.e., thickening), Gaussian dissipation, etc. of the lower-dimensional boundary into the adjacent 3D volume. An alternative form for the BDF can be obtained using the I-MDF and E-MDF presented earlier:

$$\mathscr{f}_{\sigma,S}^{\mathsf{BDF}} := \frac{\rho(\cdot)}{O(\sigma)} \bullet \mathscr{f}_{\sigma,S}^{\mathsf{I\text{-}MDF}} \bullet \mathscr{f}_{\sigma,S}^{\mathsf{E\text{-}MDF}} = \frac{\rho(\cdot)}{O(\sigma)} \bullet (\mathbf{1}_S * \mathscr{g}_\sigma^3) \bullet (\mathbf{1}_S \diamond \mathscr{g}_\sigma^3), \tag{5.1.16}$$

which can be expanded into the following integral form:

$$\mathscr{f}_{\sigma,S}^{\mathsf{BDF}}(\mathbf{x}) = \frac{\rho(\cdot)}{O(\sigma)} \int_{S \times c(S)} \mathscr{g}_\sigma^3(\mathbf{x} - \mathbf{x}') \; \mathscr{g}_\sigma^3(\mathbf{x} - \mathbf{x}'') \; d\mu^3[\mathbf{x}'] d\mu^3[\mathbf{x}'']. \tag{5.1.17}$$

If the multiplication of the IDF and EDF is interpreted as a Boolean intersection, the above equation can be viewed as defining the thickened boundary as the difference of the $\pm O(\sigma)-$offsets $\mathfrak{G}_{O(\sigma)}^{\uparrow}(S) = (S \oplus B_{O(\sigma)})$ and $\mathfrak{G}_{O(\sigma)}^{\downarrow}(S) = (S \ominus B_{O(\sigma)})$, which, in turn, is the intersection of the offsets of interior and exterior. The denominator of $O(\sigma)$ is included to convert the function in the limit into a $1^{\text{st}}-$order $\delta-$singularity.[2]

Yet another alternative formulation of the BDF can be obtained by directly offsetting the boundary $\mathfrak{G}_{O(\sigma)}^{\uparrow}(\mathscr{b}(S))$, which can be implicitly characterized by a convolution of the $\delta-$singular boundary descriptor with the nascent $\delta-$function:

$$\mathscr{f}_{\sigma,S}^{\mathsf{BDF}} := (f_{\mathscr{b}(S)}^1 * \mathscr{g}_\sigma^3) = \big(\rho \bullet (\mathbb{1}_{\mathscr{b}(S)}^1 * \mathscr{g}_\sigma^3)\big) = \big(\rho \bullet ((\delta^1 \circ \mathrm{d}_S) * \mathscr{g}_\sigma^3)\big). \tag{5.1.18}$$

The presented formulations of the BDF are not equivalent, but may be used inter-

[2]Note that this conceptualization is useful for regular 3D solids only, whose boundaries are piecewise $2-$manifolds that are properly describable via $1^{\text{st}}-$order $\delta-$singularities.

changeability in applications where 2D surfaces bounding 3D solids are to be approximated with density functions.

5.1.1.3 Skeletal Density Functions

The skeletal density function (SDF) $\mathscr{f}^{\mathsf{SDF}}_{\sigma,S} : \mathsf{E} \to \mathbb{R}$ can be defined as:

$$\mathscr{f}^{\mathsf{SDF}}_{\sigma,S}(\mathbf{x}) := \int_{\mathscr{b}(S)} \rho(\mathbf{x}) \; \mathscr{g}^3_\sigma\big(\|\mathbf{x}-\mathbf{x}'\|_2 - |\mathbb{d}_S(\mathbf{x})|\big) \; d\mu^2[\mathbf{x}'], \tag{5.1.19}$$

The surface 2−integral can be converted into a volume 3−integral by using the 1^{st}−order δ−singular boundary descriptor (i.e., an inner product) as discussed in Section 3.5.2.1. The internal/external skeletal density functions (I-SDF/E-SDF) $\mathscr{f}^{\mathsf{I\text{-}SDF}}_{\sigma,S}, \mathscr{f}^{\mathsf{E\text{-}SDF}}_{\sigma,S} : \mathsf{E} \to \mathbb{R}$ can be separated as:

$$\mathscr{f}^{\mathsf{I\text{-}SDF}}_{\sigma,S} := (\mathbb{1}^0_{i\,(S)} \bullet \mathscr{f}^{\mathsf{SDF}}_{\sigma,S}) = \int_{\mathscr{b}(S)} f_S \bullet \mathscr{g}^3_\sigma\big(\|\cdot-\mathbf{x}'\|_2 - |\mathbb{d}_S(\cdot)|\big) \; d\mu^2[\mathbf{x}'], \tag{5.1.20}$$

$$\mathscr{f}^{\mathsf{E\text{-}SDF}}_{\sigma,S} := (\mathbb{1}^0_{e(S)} \bullet \mathscr{f}^{\mathsf{SDF}}_{\sigma,S}) = \int_{\mathscr{b}(S)} \neg f_S \bullet \mathscr{g}^3_\sigma\big(\|\cdot-\mathbf{x}'\|_2 - |\mathbb{d}_S(\cdot)|\big) \; d\mu^2[\mathbf{x}'], \tag{5.1.21}$$

noting that $f_S := (\rho \bullet \mathbf{1}_S)$ and that $\mathbb{1}^0_{i\,(S)} \overset{*}{\equiv} \mathbf{1}_S$ and $\mathbb{1}^0_{e(S)} \overset{*}{\equiv} \mathbf{1}_{e(S)} \equiv \neg\mathbf{1}_S$. Alternatively, the SDFs can be constructed by $O(\sigma)$−offsetting the MAs of interior/exterior, i.e., convolution of the $\mathscr{g}_\sigma$−kernel with a δ−singular shape descriptors of the MAs (e.g., obtained via composition of a δ−function with the distance to MAs). But that is less desirable since distance to MA is difficult to compute—due to the instability issues of the MA [17] discussed at several occasions—and unlike distance to boundary, cannot be expected from a typical geometric modeler to support as a basic query.

5.1.1.4 Complex Density Functions

It is sometimes useful to combine the real-valued internal/external MDF or SDF into a single complex-valued function $\ell^{\mathsf{MDF}}_{\sigma,S}, \ell^{\mathsf{SDF}}_{\sigma,S} : \mathsf{E} \to \mathbb{C}$, from which the internal/external density functions can be retrieved as real- and imaginary-parts (or vice versa). The benefit of this trick is in formulating score functions in terms of cross-correlations. For example, Bajaj et al. [25, 69] used a similar trick to characterize protein molecules with Gaussian densities at each atom's center—which can be viewed as the MDF of the discrete pointset—using real coefficients for the skin atoms and imaginary coefficients for the core atoms. When the so-obtained shape descriptors were cross-correlated to obtain a score function for their shape complementarity, the skin-skin overlaps resulted in a positive-real award (+real $\times$ +real = +real) while core-core overlaps resulted in negative-real penalty (+imaginary $\times$ +imaginary = $-$real). I used a similar trick in formulating shape complementarity score functions for objects of arbitrary shapes in terms of SDFs [33, 38] and applied it to haptic assembly [32, 33, 35, 37], as will be briefly overviewed in Section 5.4.

In the following sections, I demonstrate a few engineering applications of real- and complex-valued density functions and their correlations.

5.2 Collision and Contact Predicates

Analytic collision detection (CD) can be traced to the work by Comba [73] on convex sets. Kavraki [174] discovered the interpretation of the translational $\mathsf{C}-$obstacle as a convolution of the objects—the robot and its workspace in the context of path planning [210]—along with the application of the FFT. Both objects $S_1, S_2 \in \mathscr{P}^*(\mathsf{E})$ are

represented by binary indicators $\mathbf{1}_{S_1}, \mathbf{1}_{S_2} : \mathbb{R}^3 \to \{0, 1\}$, discretized as bitmaps, and the integer map of the translational $\mathsf{C}-$obstacle obtained as $\mathfrak{G}_{S_1,S_2}(\mathbf{t}) = (\mathbf{1}_{S_1} * \tilde{\mathbf{1}}_{S_2})(\mathbf{t})$ simply counts the number of grid cells that overlap at a relative translation $\mathbf{t} \in \mathbb{R}^3$. The algorithm performs two forward FFTs to obtain $\hat{\mathbf{1}}_{S_1}$ and $\hat{\mathbf{1}}_{S_2}$, a pairwise multiplication to obtain $\hat{\mathfrak{G}}_{S_1,S_2} = \hat{\mathbf{1}}_{S_1} \bar{\hat{\mathbf{1}}}_{S_2}$, and an inverse FFT to retrieve the obstacle map in $O(m \log m)$ time, where m is the grid size. Although the algorithm is asymptotically optimal to obtain a complete description of the obstacle for all possible translations in a given discretized domain, it is rarely useful for time-critical CD (e.g., in real-time simulations and physically-based modeling) where a *single* configuration is queried.

Lysenko [214] recently generalized the approach by using bump functions to facilitate differentiation, and proposed techniques to enable time-critical CD for a single-configuration query via truncated Fourier expansions, along with an analytic groundwork for early-hit/miss tests. Noting that the inner product structure is preserved by the CFT according to Parseval's theorem (Theorem 2.2.10), the collision predicate for a single relative configuration $(\mathbb{r}, \mathbf{t}) \in \mathrm{SE}(3)$ can be obtained as

$$\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) = \Big\langle f_{S_1}, (f_{\mathbb{r}S_2} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{f}_{S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{f}_{\mathbb{r}S_2}) \Big\rangle, \tag{5.2.1}$$

noting that $f_{(\mathbb{r},\mathbf{t})S} = (f_{\mathbb{r}S} \circ \mathsf{act}(-\mathbb{t})) = ((f_S \circ \mathbb{r}^{\mathrm{T}}) \circ \mathsf{act}(-\mathbb{t}))$ which transforms to $\hat{f}_{(\mathbb{r},\mathbf{t})S} = (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{f}_{\mathbb{r}S}) = (\hat{\varsigma}_{\mathbf{t}} \bullet (\hat{f}_S \circ \mathbb{r}^{\mathrm{T}}))$. As mentioned earlier, $\mathsf{act}(-\mathbb{t})(\mathbf{x}) = (\mathbf{x} - \mathbf{t})$ is the shift function whose Fourier operator $\hat{\varsigma}_{\mathbf{t}}(\boldsymbol{\omega}) = e^{-2\pi \mathbf{i}(\boldsymbol{\omega} \cdot \mathbf{t})}$ is the CFT of the shifted $\delta-$function $\varsigma_{\mathbf{t}}(\mathbf{x}) = (\delta^3 \circ \mathsf{act}(-\mathbb{t}))(\mathbf{x}) = \delta^3(\mathbf{x} - \mathbf{t})$.

If a grid-based discretization is used, the rotation can be incorporated by a trilinear interpolation in either domain. Although a brute-force computation of the physical domain inner product over a grid of size m takes $O(m)$—without a simple way of re-

ducing the complexity once the grid resolution is fixed upfront—the frequency domain integral can be computed in $O(m')$ over a truncated grid of much smaller size $m' \ll m$ specified on-the-fly. This provides a mechanism for trading off accuracy with time, by a spiral traversal of the frequency grid starting from the dominant modes until the available time is over. On the other hand, the numerous combinatorial CD methods developed over the years (reviewed in [169]) exploit a variety of data structures to avoid brute-force testing in the physical domain, the likes of which are not available in the frequency domain. The sphere-tree methods [48, 156, 248, 362] are among the most efficient, which enable another trade-off mechanism by descending down the tree until the time allocated to CD is consumed. Our framework enables exploiting the existing combinatorial techniques alongside the recent analytic methods in both domains. The details pertaining to the following are beyond the scope of this article and will be presented elsewhere

1. *early-hit test* by limiting the integration of (5.2.1) to an intersection with a ball, which is a simple multiplication in the physical domain;

2. *early-miss test* by offsetting (i.e., Minkowski sum) with a single ball, which becomes a simple multiplication in the frequency domain; and

3. *differentiation* of (5.2.1) for contact force/torque computation, using pairwise spherical primitive interactions.

Figure 5.2.1 illustrates the analytic CD between two planar shapes, and a plot of the collision $2-$measure (i.e., area) between them at different translational configurations, whose $0-$superlevel set is the translational $\mathsf{C}-$obstacle. The figure also shows low-pass filtering, which reduces the computational complexity of the FFT-based convolution from $O(m \log m)$ to $O(m' \log m')$ with $m' \ll m$ retained dominant

modes in the Fourier domain, at the expense of a small compromise of accuracy—e.g., obtaining small nonzero collision measure outside the $\mathsf{C}-$obstacle due to the rippling effect of truncating Fourier expansions.

The expression in (5.2.1) can be extended to compute overlaps between features of heterogeneous dimensions. Without changing the outcome, we can rewrite (5.2.1) in a way that it cross-correlates the nonsingular shape descriptors of the 3D interiors $\mathit{i}\,(S_1), \mathit{i}\,(S_2) \in \mathscr{P}^3(\mathsf{E})$ to measure the volume of internal interference:

$$\mathfrak{G}^{\textsf{CD-II}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \Big\langle \mathbb{1}^0_{\mathit{i}\,(S_1)}, (\mathbb{1}^0_{\mathit{i}\,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathbb{1}}^0_{\mathit{i}\,(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}^0_{\mathit{i}\,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.2}$$

An approximation of the correlation function is given by substituting the crisp memberships $\mathbb{1}^0_{\mathit{i}\,(S_1)}, \mathbb{1}^0_{\mathit{i}\,(S_2)} \in (\mathbb{R}^3_\star)^{\mathsf{E}}$ with their fuzzy approximations $\mathscr{f}^{\textsf{I-MDF}}_{\sigma,S_1}, \mathscr{f}^{\textsf{I-MDF}}_{\sigma,S_2} \in \mathbb{R}^{\mathsf{E}}$:

$$\mathfrak{G}^{\textsf{CD-II}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \mathscr{f}^{\textsf{I-MDF}}_{\sigma,S_1}, (\mathscr{f}^{\textsf{I-MDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathscr{f}}^{\textsf{I-MDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathscr{f}}^{\textsf{I-MDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.3}$$

if $0 < \sigma \ll 1$, where the exact value can be obtained by letting $\sigma \to 0^+$.

Alternatively, the correlation function could be rewritten in a way that it cross-correlates the $1^{\text{st}}-$order $\delta-$singular shape descriptor of one of the two 2D boundaries $\mathit{b}\,(S_1), \mathit{b}\,(S_2) \in \mathscr{P}^2(\mathsf{E})$ with one of the 3D interiors $\mathit{i}\,(S_1), \mathit{i}\,(S_2) \in \mathscr{P}^3(\mathsf{E})$ to measure the surface area of the former overlapped by the latter:

$$\mathfrak{G}^{\textsf{CD-BI}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \Big\langle \mathbb{1}^1_{\mathit{b}\,(S_1)}, (\mathbb{1}^0_{\mathit{i}\,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathbb{1}}^1_{\mathit{b}\,(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}^0_{\mathit{i}\,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.4}$$

$$\mathfrak{G}^{\textsf{CD-IB}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \Big\langle \mathbb{1}^0_{\mathit{i}\,(S_1)}, (\mathbb{1}^1_{\mathit{b}\,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathbb{1}}^0_{\mathit{i}\,(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}^1_{\mathit{b}\,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.5}$$

Once again, an approximation of the correlation function can be given by substituting

the crisp functional descriptors $\mathbb{1}^1_{\mathscr{b}(S_1)} \in (\mathbb{R}^2_\star)^{\mathsf{E}}$ and $\mathbb{1}^0_{i(S_2)} \in (\mathbb{R}^3_\star)^{\mathsf{E}}$ with their fuzzy approximations $\mathscr{f}^{\mathsf{BDF}}_{\sigma,S_1}, \mathscr{f}^{\mathsf{I\text{-}MDF}}_{\sigma,S_2} \in \mathbb{R}^{\mathsf{E}}$:

$$\mathfrak{G}^{\mathsf{CD\text{-}BI}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \mathscr{f}^{\ \mathsf{BDF}}_{\sigma,S_1}, (\mathscr{f}^{\mathsf{I\text{-}MDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathscr{f}}^{\ \mathsf{BDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathscr{f}}^{\mathsf{I\text{-}MDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.6}$$

$$\mathfrak{G}^{\mathsf{CD\text{-}IB}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \mathscr{f}^{\mathsf{I\text{-}MDF}}_{\sigma,S_1}, (\mathscr{f}^{\ \mathsf{BDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathscr{f}}^{\mathsf{I\text{-}MDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathscr{f}}^{\ \mathsf{BDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.7}$$

if $0 < \sigma \ll 1$, where the exact value can be obtained by letting $\sigma \to 0^+$.

Last but not least, the correlation function could be rewritten in a way that it cross-correlates the $1^{\text{st}}-$order $\delta-$singular shape descriptors of the 2D boundaries $\mathscr{b}(S_1), \mathscr{b}(S_2) \in \mathscr{P}^2(\mathsf{E})$ to measure the surface area of contact:

$$\mathfrak{G}^{\mathsf{CD\text{-}BB}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \Big\langle \mathbb{1}^1_{\mathscr{b}(S_1)}, (\mathbb{1}^1_{\mathscr{b}(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathbb{1}}^1_{\mathscr{b}(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}^1_{\mathscr{b}(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.8}$$

Once again, an approximation of the correlation function can be given by substituting the crisp functional descriptors $\mathbb{1}^1_{\mathscr{b}(S_1)}, \mathbb{1}^1_{\mathscr{b}(S_2)} \in (\mathbb{R}^2_\star)^{\mathsf{E}}$ with their fuzzy approximations $\mathscr{f}^{\mathsf{BDF}}_{\sigma,S_1}, \mathscr{f}^{\mathsf{BDF}}_{\sigma,S_2} \in \mathbb{R}^{\mathsf{E}}$:

$$\mathfrak{G}^{\mathsf{CD\text{-}BB}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \mathscr{f}^{\ \mathsf{BDF}}_{\sigma,S_1}, (\mathscr{f}^{\ \mathsf{BDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\mathscr{f}}^{\ \mathsf{BDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathscr{f}}^{\ \mathsf{BDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.2.9}$$

if $0 < \sigma \ll 1$, where the exact value can be obtained by letting $\sigma \to 0^+$.

Figures 5.2.2 and 5.2.3 illustrate the above two cases for a simple pair of planar shapes, along with the effect of low-pass filtering in the accuracy of the computation for a fixed value of $\sigma > 0$.

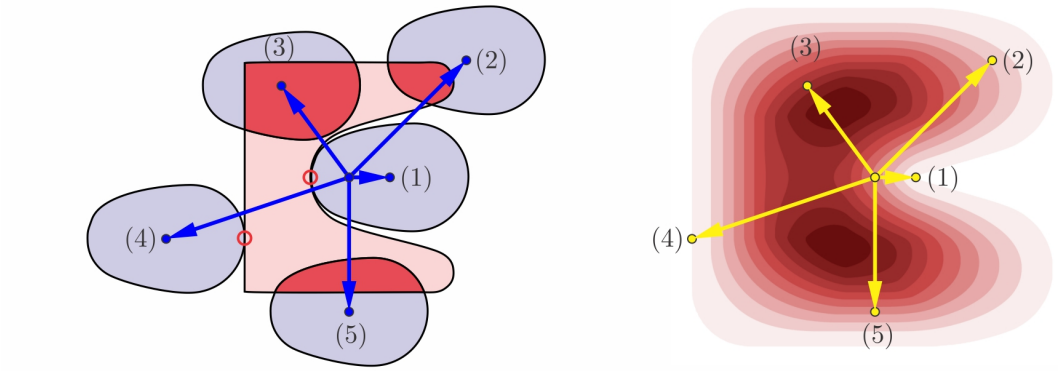

(a) Different configurations (b) Collision measure map

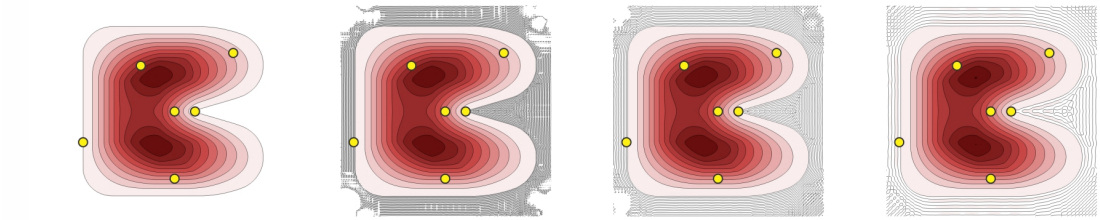

(c) $m' = 2^9$ (d) $m' = 2^8$ (e) $m' = 2^7$ (f) $m' = 2^6$

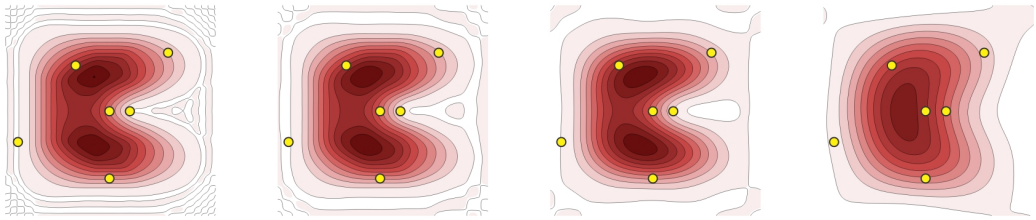

(g) $m' = 2^5$ (h) $m' = 2^4$ (i) $m' = 2^3$ (j) $m' = 2^2$

FIGURE 5.2.1: Measuring interior-interior overlap between two planar parts at different translational configurations (1–6) in (a). The FFT-based convolution of indicator functions returns the map in (b), which is a plot of the intersection 2−measures. Low-pass filtered approximations for different numbers of retained dominant modes are plotted in (c–j).

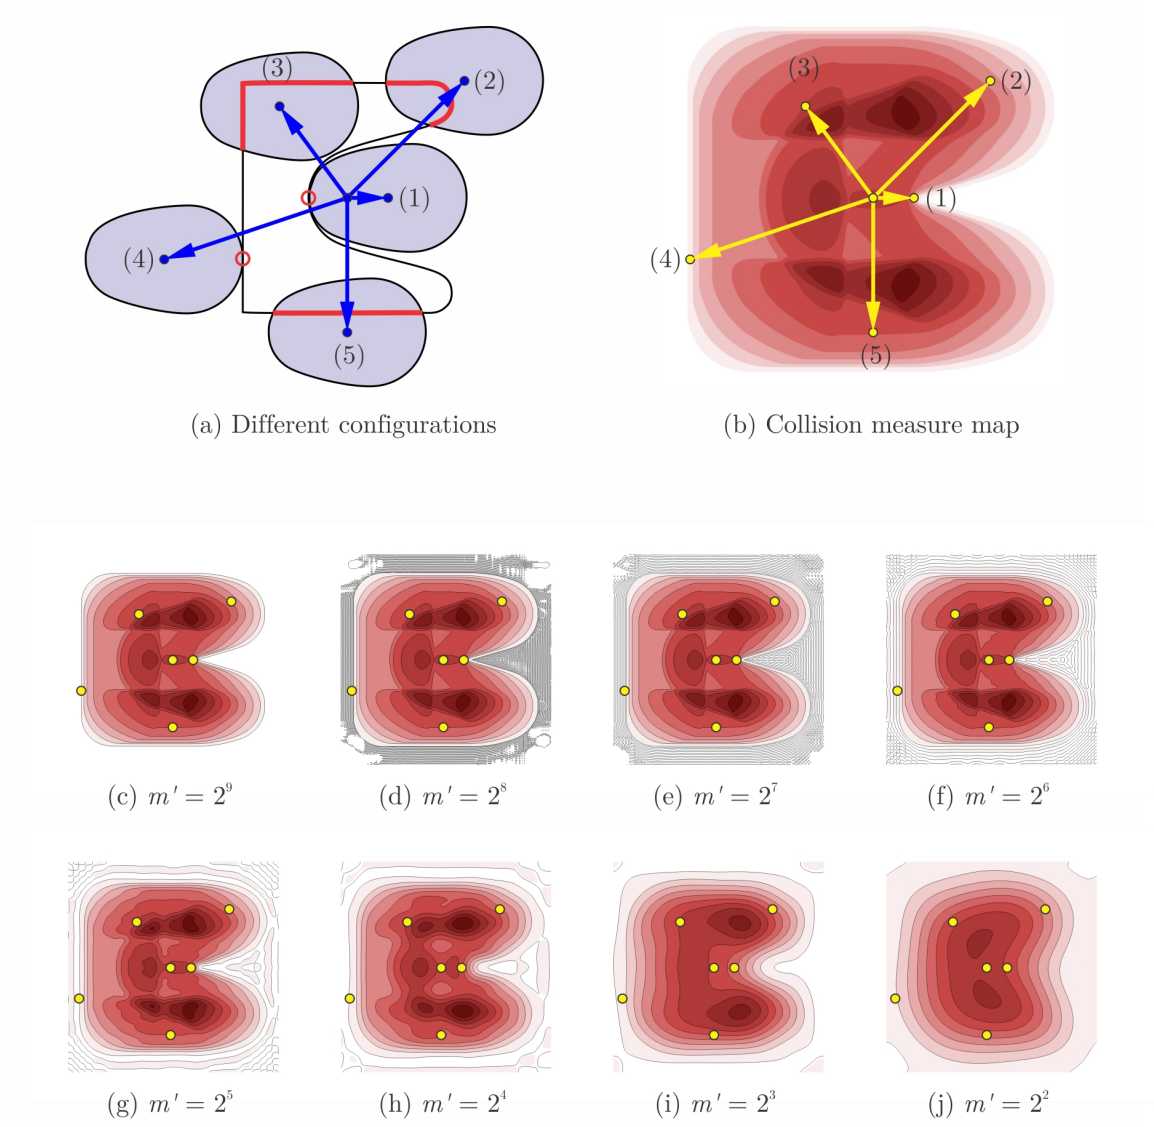

(a) Different configurations (b) Collision measure map

(c) $m' = 2^9$ (d) $m' = 2^8$ (e) $m' = 2^7$ (f) $m' = 2^6$

(g) $m' = 2^5$ (h) $m' = 2^4$ (i) $m' = 2^3$ (j) $m' = 2^2$

FIGURE 5.2.2: Measuring the boundary-interior overlap between two planar parts at different translational configurations (1–6) in (a). The FFT-based convolution between fuzzy shape descriptors for finite $\sigma > 0$ returns the map in (b), which is a plot of the approximation of the intersection 1−measures. Low-pass filtered approximations for different numbers of retained dominant modes are plotted in (c–j).

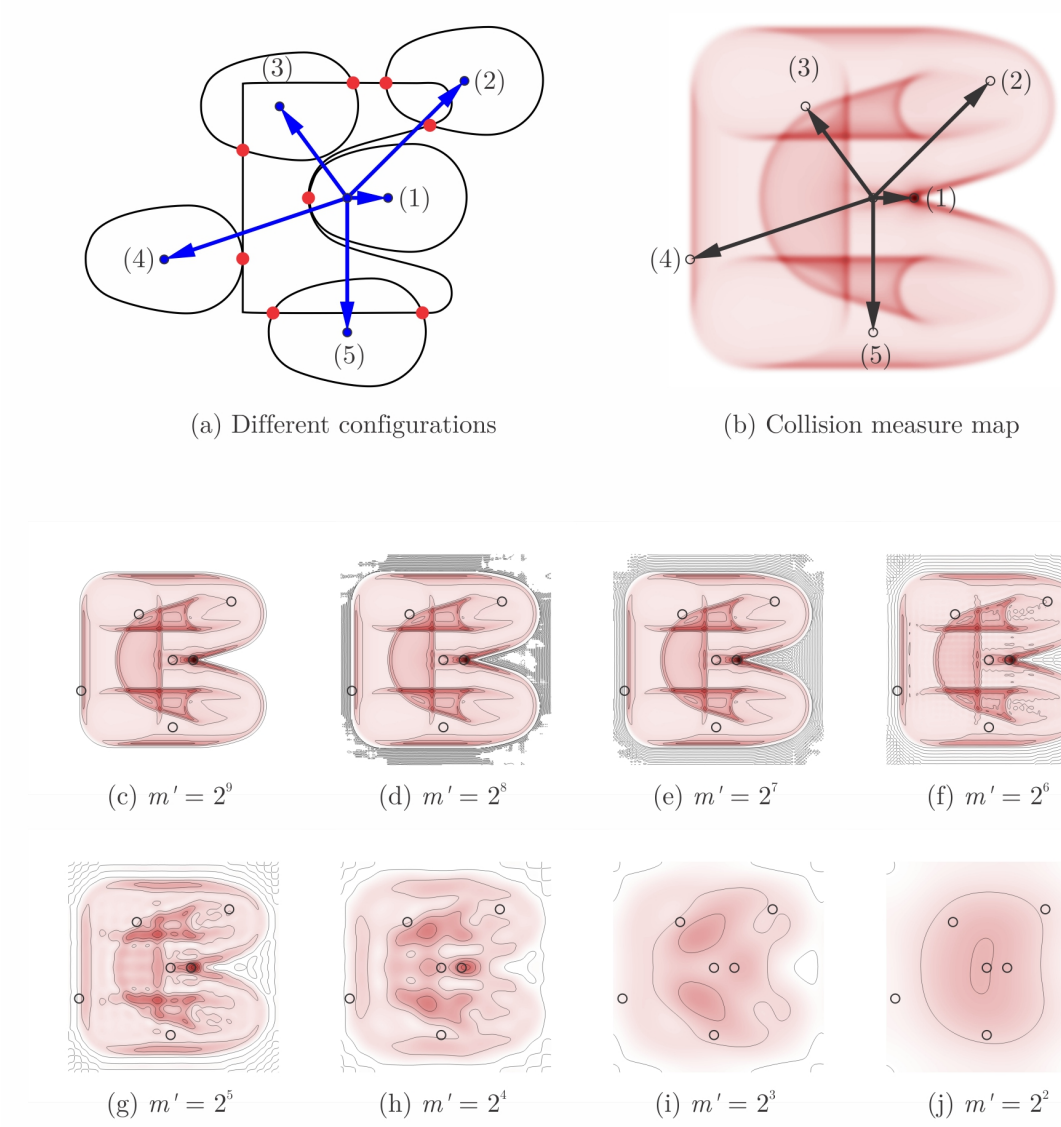

(a) Different configurations (b) Collision measure map

(c) $m' = 2^9$ (d) $m' = 2^8$ (e) $m' = 2^7$ (f) $m' = 2^6$

(g) $m' = 2^5$ (h) $m' = 2^4$ (i) $m' = 2^3$ (j) $m' = 2^2$

FIGURE 5.2.3: Measuring the boundary-boundary overlap between two planar parts at different translational configurations (1–6) in (a). The FFT-based convolution between fuzzy shape descriptors for finite $\sigma > 0$ returns the map in (b), which is a plot of the approximation of the intersection 0−measures. Low-pass filtered approximations for different numbers of retained dominant modes are plotted in (c–j).

5.3 Comparative Shape Analysis

Obtaining measures of shape similarity (SS) or shape complementarity (SC)—which are arguably dual to each other—over the configuration space of relative motions is important in many applications. For example, in automatic (e.g., robot-assisted) solution of jigsaw puzzles [52, 132, 135, 146, 238], whose practical applications extend to automatic and semi-automatic assembly of mechanical parts, bin packing, protein docking, and material cost optimization for manufacturing—ranging from apparel/leather cutting and sheet metal layout optimization (in 2D) to build volume optimization for additive manufacturing (in 3D).

Figure 5.3.1 illustrates the correlation function in (5.2.1) for a pair of planar shapes. Using indicator functions $f_{S_1} := \mathbf{1}_{S_1}$ and $f_{S_2} := \mathbf{1}_{S_2}$, the cross-correlation $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t})$ for a fixed rotation $\mathbb{r} \in \mathrm{SO}(2)$ implicitly characterizes the $\mathsf{C}-$obstacle by the collision $2-$measure (i.e., area) at different displacements $\mathbf{t} \in \mathbb{R}^2$. Consequently, $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) = 0$ characterizes the collection of feasible configurations $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ corresponding to zero intersection measure $\mu^2[S_1 \cap (\mathbb{r}, \mathbf{t})S_2] = 0$, including unassembled (no-contact) and assembled (proper point/curve/surface contact), while $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) > 0$ implicitly defines the C-obstacle. In other words, the correlation function penalizes collision as in positions A and B (i.e., $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) > 0$), but does not differentiate point contact in C and separation in D from proper fit/contact in E (i.e., $\mathfrak{G}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) = 0$). On the other hand, if more sophisticated shaped descriptors $f_{S_1} \neq \mathbf{1}_{S_1}$ and $f_{S_2} \neq \mathbf{1}_{S_2}$ are used, which capture the skeletal density of the two shapes, the correlation function can distinguish between the 'best fit' configuration in E, in which there is maximal skeletal overlap, and other non-colliding contact configurations. This is motivated by the observation that SS and SC can be characterized by the extent of overlap

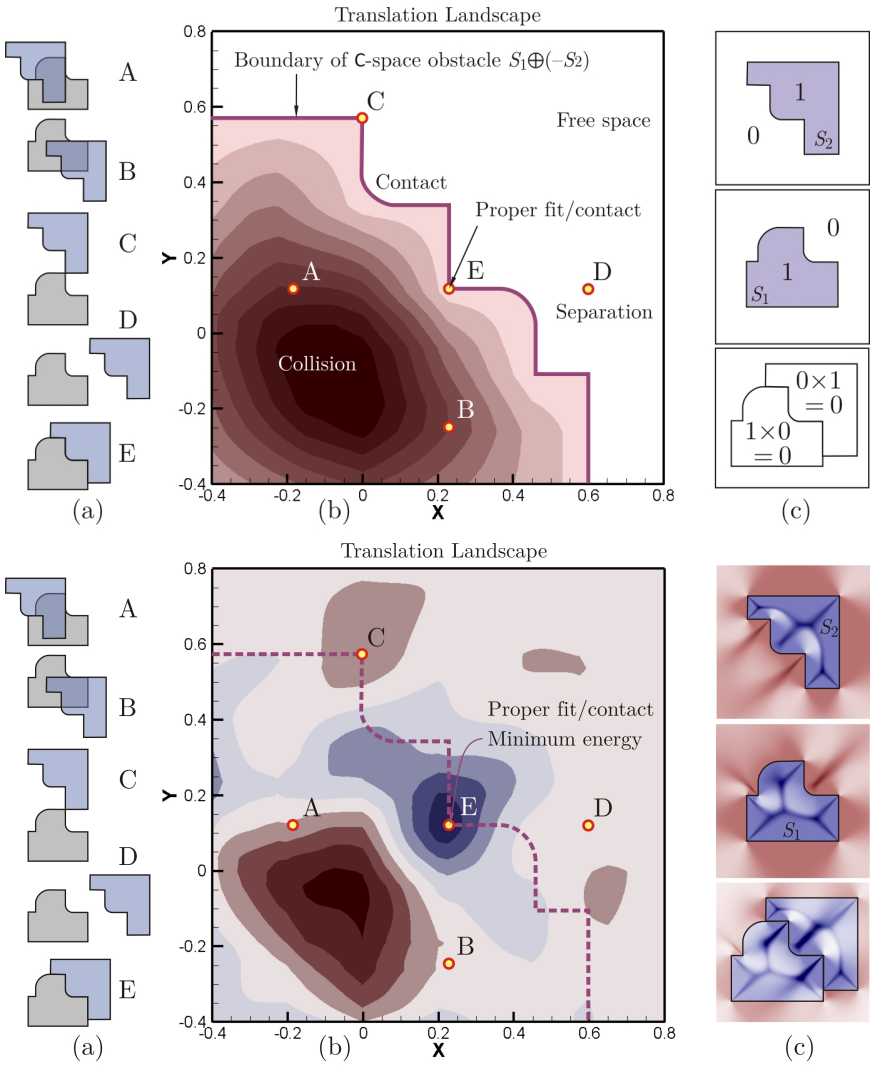

FIGURE 5.3.1: Different configurations (a) are evaluated using correlations of different functions; namely, indicator functions (top) and modified SDFs (bottom) [32, 35, 37] (b) whose overlaps in the best fit configuration are illustrated in (c).

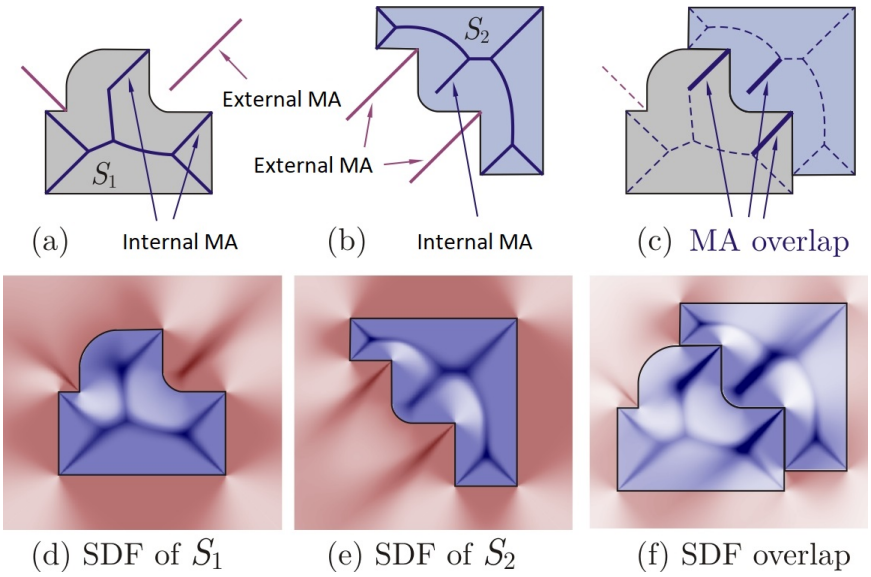

FIGURE 5.3.2: The extent of overlap between skeletal branches of the two shapes in (a, b) at a given configuration (c) gives a measure of complementarity between boundary features. These skeletal branches can be captured implicitly by complex density functions in (d, e) whose inner product in (f) yields the score function.

between MA branches—internal-internal and external-external MA overlaps for SS scoring, and internal-external and external-internal MA overlaps for SC scoring:

$$\mathfrak{G}^{\mathsf{SS\text{-}II}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \left\langle \mathbb{1}_{(\mathscr{M}\circ i)(S_1)}, (\mathbb{1}_{(\mathscr{M}\circ i)(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \right\rangle = \left\langle \hat{\mathbb{1}}_{(\mathscr{M}\circ i)(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}_{(\mathscr{M}\circ i)(\mathbb{r}S_2)}) \right\rangle,$$

$$\mathfrak{G}^{\mathsf{SS\text{-}EE}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \left\langle \mathbb{1}_{(\mathscr{M}\circ e)(S_1)}, (\mathbb{1}_{(\mathscr{M}\circ e)(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \right\rangle = \left\langle \hat{\mathbb{1}}_{(\mathscr{M}\circ e)(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}_{(\mathscr{M}\circ e)(\mathbb{r}S_2)}) \right\rangle,$$

$$\mathfrak{G}^{\mathsf{SC\text{-}IE}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \left\langle \mathbb{1}_{(\mathscr{M}\circ i)(S_1)}, (\mathbb{1}_{(\mathscr{M}\circ e)(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \right\rangle = \left\langle \hat{\mathbb{1}}_{(\mathscr{M}\circ i)(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}_{(\mathscr{M}\circ e)(\mathbb{r}S_2)}) \right\rangle,$$

$$\mathfrak{G}^{\mathsf{SC\text{-}EI}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) := \left\langle \mathbb{1}_{(\mathscr{M}\circ e)(S_1)}, (\mathbb{1}_{(\mathscr{M}\circ i)(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \right\rangle = \left\langle \hat{\mathbb{1}}_{(\mathscr{M}\circ e)(S_1)}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\mathbb{1}}_{(\mathscr{M}\circ i)(\mathbb{r}S_2)}) \right\rangle,$$

whose approximation can be given by substituting the crisp functional descriptors of

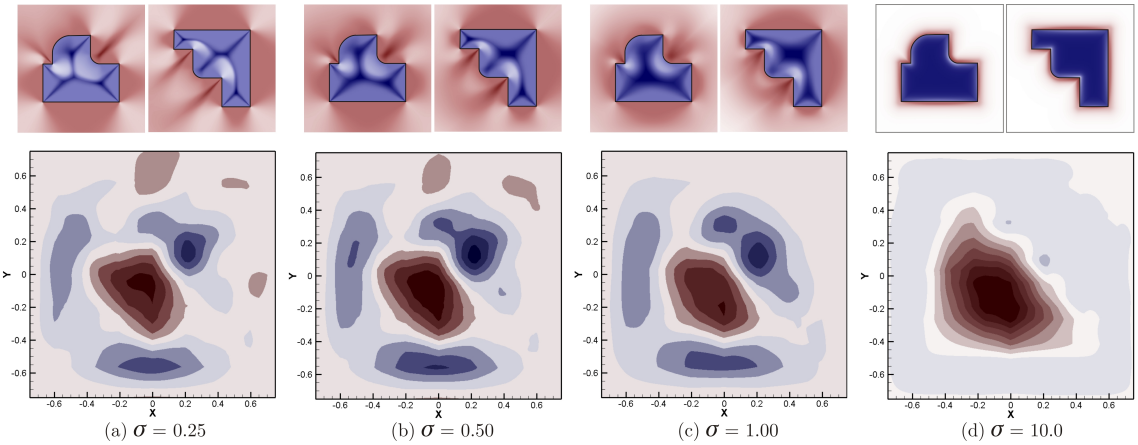

(a) $\sigma = 0.25$ (b) $\sigma = 0.50$ (c) $\sigma = 1.00$ (d) $\sigma = 10.0$

FIGURE 5.3.3: The extent of geometric details captured by the skeletal density distribution is adjustable by the dissipation factor σ.

the internal and external MAs with their fuzzy approximations:

$$\mathfrak{G}^{\mathsf{SS\text{-}II}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,S_1}, (\hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.3.1}$$

$$\mathfrak{G}^{\mathsf{SS\text{-}EE}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,S_1}, (\hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.3.2}$$

$$\mathfrak{G}^{\mathsf{SC\text{-}IE}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,S_1}, (\hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle, \tag{5.3.3}$$

$$\mathfrak{G}^{\mathsf{SC\text{-}EI}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \Big\langle \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,S_1}, (\hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \Big\rangle = \Big\langle \hat{\ell}^{\mathsf{E\text{-}SDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\ell}^{\mathsf{I\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)}) \Big\rangle. \tag{5.3.4}$$

The idea is illustrated in 5.3.2 for the SC scoring of a pair of simple 2D objects. Figures 5.3.3 and 5.3.4 illustrate the effects of changing the dissipation factor $\sigma > 0$ and low-pass filtering on computing the artificial "energy field" obtained from the SC score, whose application to haptic assembly is discussed at length in [32, 35, 37].

For the particular application of analytic comparative shape analysis—i.e., measuring SS and SC scores for different rigid configurations—the real-valued SDF shape descriptors in (5.1.20) and (5.1.21) are modified and compounded into a single complex-

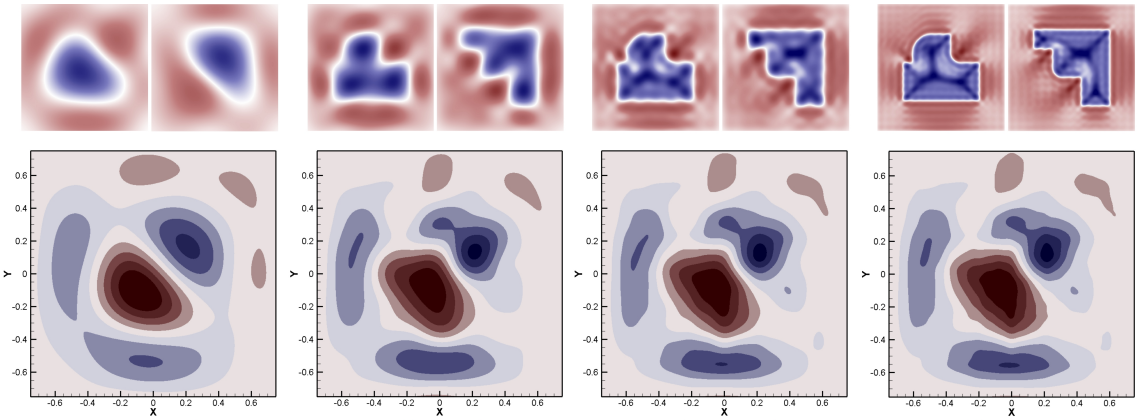

(a) $m' = 4^2$ (b) $m' = 8^2$ (c) $m' = 16^2$ (d) $m' = 32^2$

FIGURE 5.3.4: Frequency domain representation allows for a systematic means of successive approximation of the energy field.

valued SDF $\mathscr{f}^{\mathsf{C\text{-}SDF}}_{\sigma,S} : (\mathsf{E} - \mathscr{b}(S)) \to \mathbb{C}$ defined in terms of the following surface integral:

$$\mathscr{f}^{\mathsf{C\text{-}SDF}}_{\sigma,S}(\mathbf{x}) := \int_{\mathscr{b}(S)} \frac{(\lambda \circ \mathsf{sign} \circ \mathrm{d}_S)(\mathbf{x})}{4\pi \|\mathbf{x} - \mathbf{x}'\|_2^2} \mathscr{g}^3_\sigma\big(\frac{\|\mathbf{x} - \mathbf{x}'\|_2}{|\mathrm{d}_S(\mathbf{x})|} - 1\big)\, d\mu^2_\perp[\mathbf{x}'], \tag{5.3.5}$$

where $d\mu^2_\perp[\mathbf{x}'] = \cos(\theta(\mathbf{x}'))\ d\mu^2[\mathbf{x}']$ in which $\cos(\theta(\mathbf{x}')) = \mathbf{n}(\mathbf{x}') \cdot \frac{(\mathbf{x}' - \mathbf{x})}{\|\mathbf{x}' - \mathbf{x}\|_2}$ projects the area element $d\mu^2[\mathbf{x}']$ with whose unit normal vector is $\mathbf{n}(\mathbf{x}') \in \mathbb{S}^2$ along $(\mathbf{x}' - \mathbf{x})$ for a query point $\mathbf{x} \in (\mathsf{E} - \mathscr{b}(S))$ and boundary point $\mathbf{x}' \in \mathscr{b}(S)$. The inverse-square decay $(4\pi\|\mathbf{x} - \mathbf{x}'\|_2^2)^{-1}$ makes sure the density function fades away as the query point moves away from the boundary features.[3] The complex coefficients given by the function $\lambda : \{-1, +1\} \to \mathbb{C}$ which depends on the sign of the distance function—i.e., $\lambda_{\mathsf{I}} := \lambda(-1)$ and $\lambda_{\mathsf{E}} := \lambda(+1)$ for the interior $(\mathbf{x} \in \mathscr{i}(S))$ and exterior $(\mathbf{x} \in \mathscr{e}(S))$ with $\mathsf{sign}(\mathrm{d}_S(\mathbf{x})) = -1$ and $+1$, respectively—adjust the intensity, sign, and type (i.e., real versus imaginary) of the SDF function in such a way that the max/min of

[3]Note that $(4\pi\|\mathbf{x} - \mathbf{x}'\|_2^2)^{-1}\ d\mu^2_\perp[\mathbf{x}']$ is the signed infinitesimal spatial angle with which the query point observes the boundary element.

the correlation function corresponds to best SS/SC configurations, respectively:

$$\mathfrak{G}^{\mathsf{SS/SC}}_{S_1,S_2}|_{\mathbb{r}}(\mathbf{t}) \approx \left\langle \ell^{\mathsf{C\text{-}SDF}}_{\sigma,S_1}, (\ell^{\mathsf{C\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)} \circ \mathsf{act}(-\mathbb{t})) \right\rangle = \left\langle \hat{\ell}^{\mathsf{C\text{-}SDF}}_{\sigma,S_1}, (\hat{\varsigma}_{\mathbf{t}} \bullet \hat{\ell}^{\mathsf{C\text{-}SDF}}_{\sigma,(\mathbb{r}S_2)}) \right\rangle. \tag{5.3.6}$$

For more technical details on the above formulation, see [32] (Chapter 3).

Figure 5.3.5 illustrates how the correlation function in (5.3.6) scores the SS/SC for a pair of planar jigsaw puzzle pieces. The plots in panels (c, d) are the score function versus relative displacement for a pair of different fixed relative rotations.

SC measures have been found useful in prediction of native binding configurations for protein molecules, and is critical in early-stage lead compound generation for rational drug design. The numerous FFT-based correlation techniques developed over the years (reviewed in [288]) use analytic correlations similar to the one given in (5.3.6), with the exception of their choice of the shape descriptor functions. For example, popular 'double-skin layer' (DSL) approach [25] integrates the skin-skin intersections to obtain a SC score function and subtracts core-core collisions as penalty, which is described as a convolution of a finite summation of Gaussian descriptors for individual atoms. Chen and Weng [61] described successful heuristics for weight assignment rasterized on a uniform grid along with the use of FFT. Bajaj et al. [25] proposed a faster grid-free method along with the use of NFFT, which has been highly influential in the development of our ideas.

Figure 5.3.6 illustrates the SC score variations for a pair of real proteins—namely, the nuclear transport factor 2 (NTF2) and the GDP-bound form of the Ras-family GTPase Ran [334] (PDB code 1A2K)—over the translational $\mathsf{C}-$space using the DSL shape descriptors, which leads to a nonzero score over a narrow subspace, making it difficult to discover by gradient-descent optimization algorithms. Figure 5.3.7 illus-

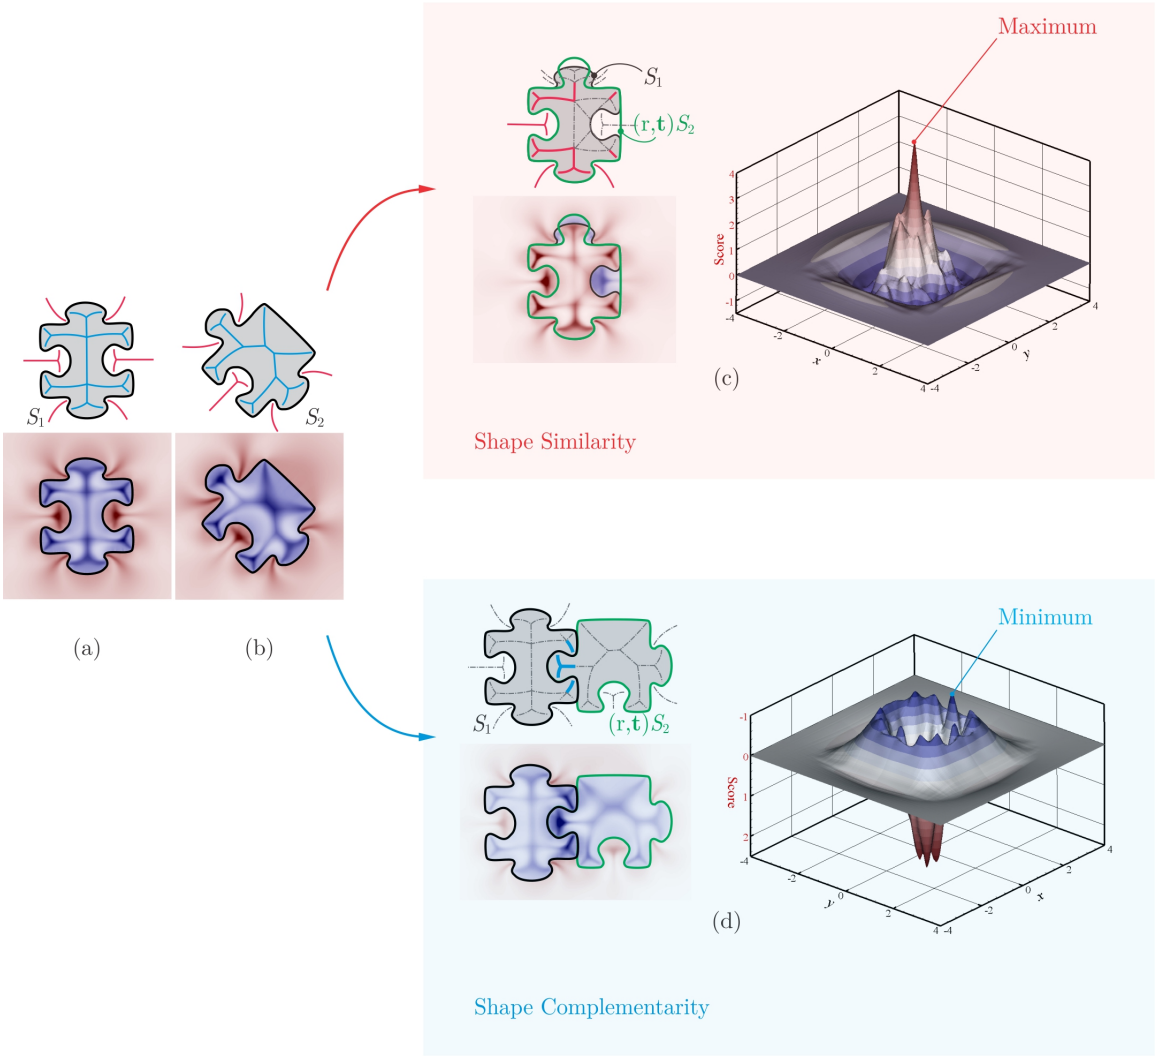

FIGURE 5.3.5: The max/min of the score function in (5.3.6) corresponds to the best shape similarity (SS) and shape complementarity (SC) configurations, respectively, upon the proper calibration of the $\lambda-$coefficients in complex SDF formula in (5.3.5).

trates the SC score variations for the same pair of proteins, using the SDF shape descriptors. The latter provides a more distributed score function, larger support over the $\mathsf{C}-$space, and stronger clues for gradient-descent optimization—essentially guiding the binding process along the skeletal branches—without changing the maximum score configuration much from the native pose.

The SDF-based correlations can also be used to guide the assembly of mechanical parts, as illustrated in Fig. 5.3.8. Although geometric interfaces in mechanical assemblies appear to be simpler than those of protein complexes, an additional form of complication is introduced when the nominal size and geometric features of the individual parts are designed to restrict the motion to 1 or 2 degrees of freedom (DOF)—forming lower kinematic pairs operating within Lie subgroups of SE(3) [242]. This leads to lower-dimensional subsets of the collision-free space over which the score function is of interest. An example application in which one deals with assembling mechanical components in a virtual environment is briefly introduced in the next section, and discussed in more detail in [32].

5.4 Virtual (i.e., Digital) Prototyping

A premier example of virtual prototyping is graphics- and haptics-enabled semi-automatic virtual assembly of mechanical parts, detailed in [32, 35, 37] and briefly overviewed below.

Consider the parts in Fig. 5.4.1 (slightly more detailed version of Fig. 5.3.8). Unlike protein molecules, they are made of simple planar, cylindrical, spherical, and toroidal surface patches; nonetheless, an automatic identification of the correspon-

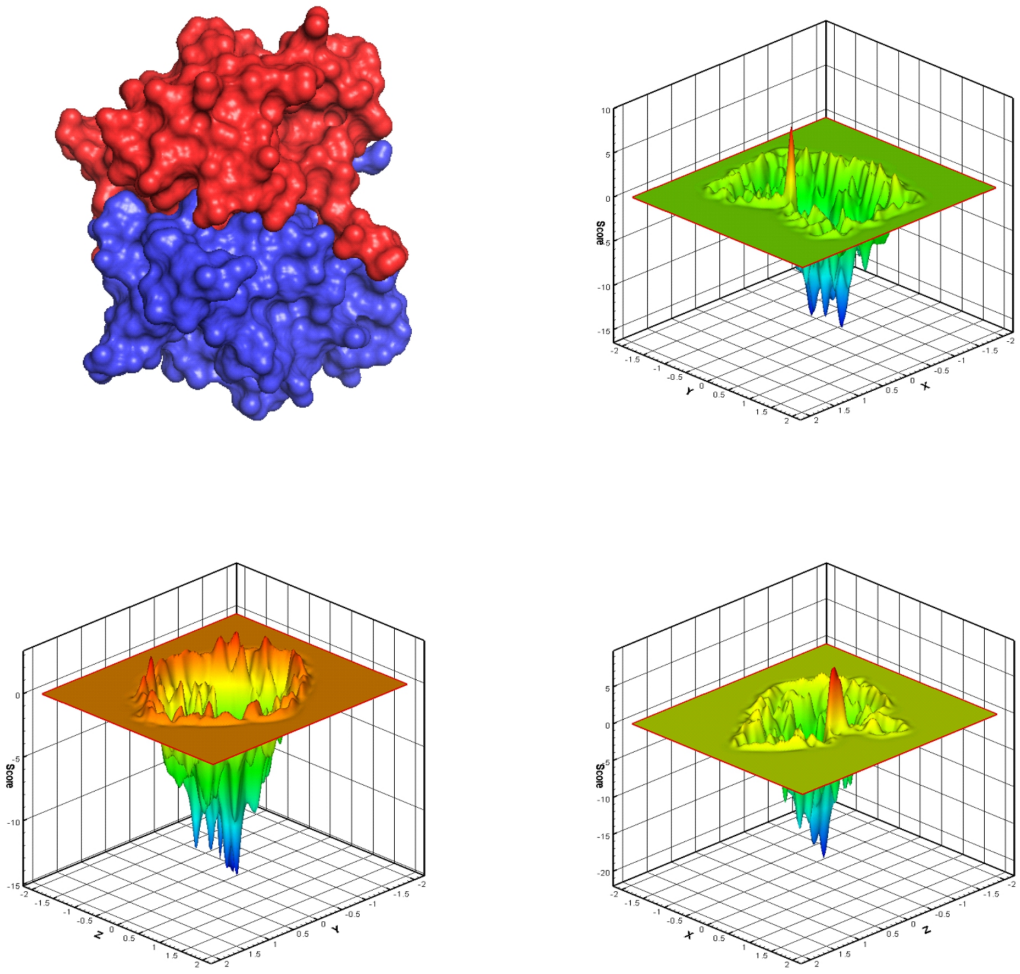

FIGURE 5.3.6: DSL-based shape complementarity (SC) score function for a protein complex (PDB code 1A2K) with respect to relative translations—plotted versus displacements along 2 Cartesian axis at-a-time, for zero translation along the 3rd axis and fixed rotation.

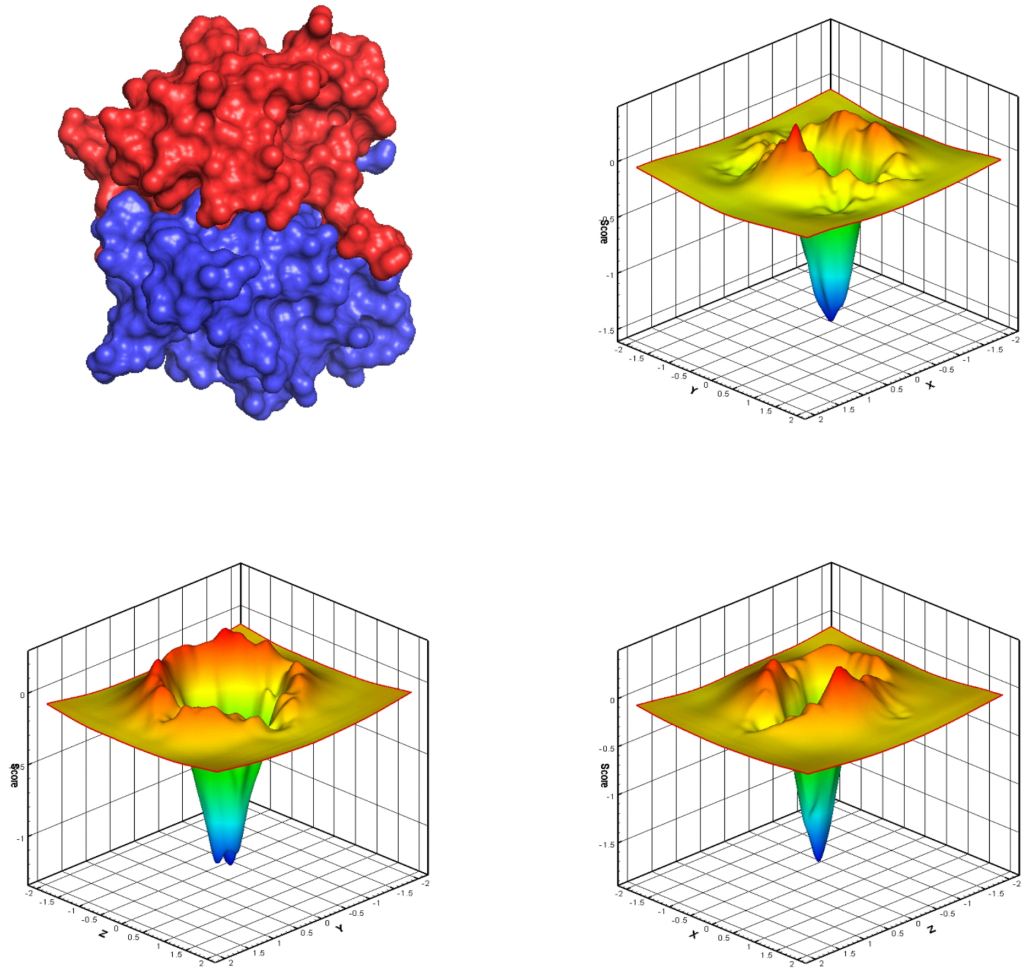

FIGURE 5.3.7: SDF-based shape complementarity (SC) score function for a protein complex (PDB code 1A2K) with respect to relative translations—plotted versus displacements along 2 Cartesian axis at-a-time, for zero translation along the 3rd axis and fixed rotation.

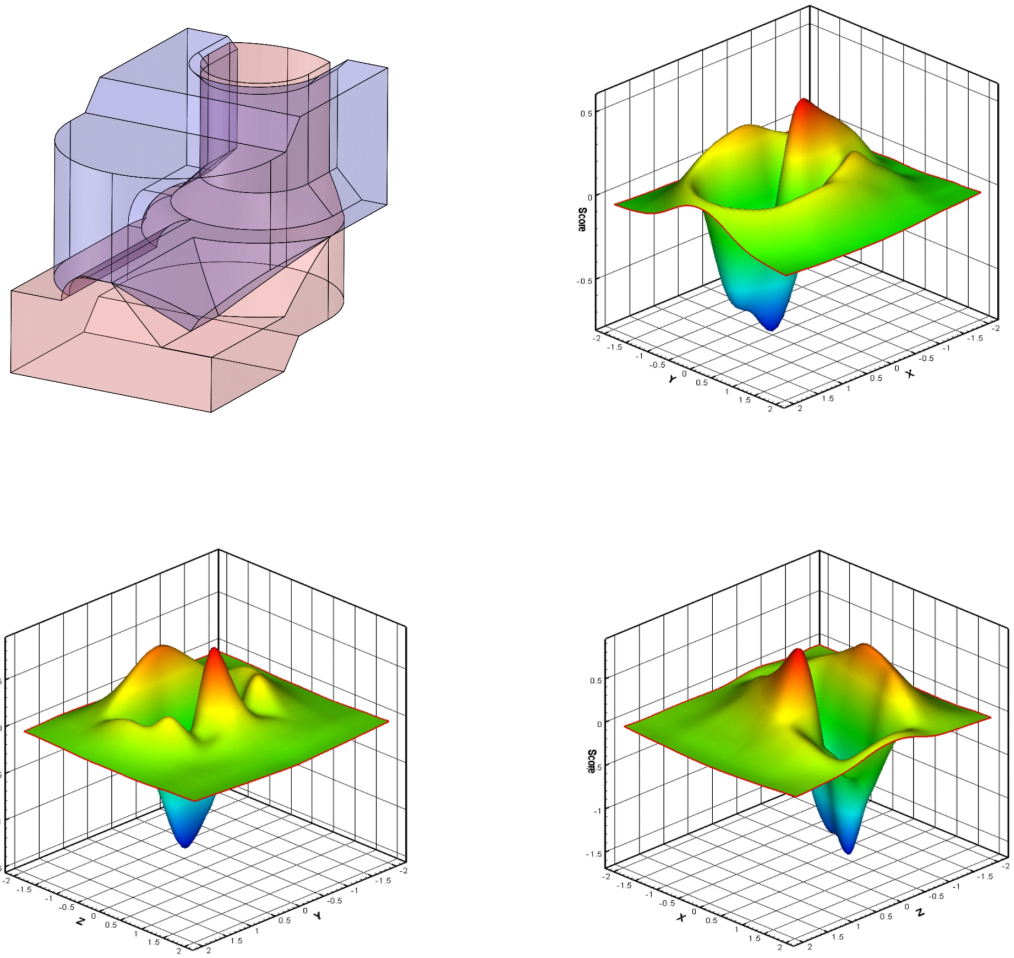

FIGURE 5.3.8: SDF-based shape complementarity (SC) score function for a pair of 3D toy mechanical parts with respect to relative translations—plotted versus displacements along 2 Cartesian axis at-a-time, for zero translation along the 3rd axis and fixed rotation.

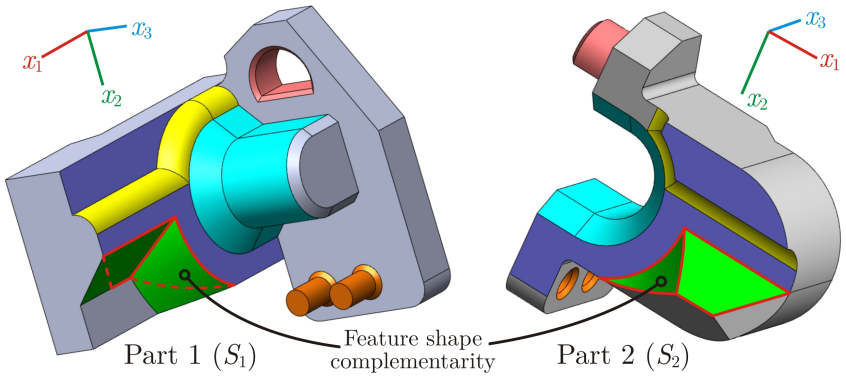

FIGURE 5.4.1: A non-trivial, zero-clearance assembly pair.

dence between the mating features (depicted with different colors) is not trivial from an algorithmic perspective—e.g., recognition and matching of the partially complementary features connected by a curve. Furthermore, there are 3 pairs of pegs and holes with zero clearance, making this example sufficiently challenging.

Figures 5.4.2 shows the effects of successive low-pass filtering for different numbers of retained dominant modes $m' \leq m$ on the SDF (only imaginary part plotted on top rows). It also plots the score function on the bottom row, over a 2D section corresponding to a biaxial relative translation along the x_1x_2-plane through the 6D convolution. As more frequency domain data is kept, the geometric details—e.g., pertaining to the small pairs of cylindrical pegs and holes depicted by their axis lines in panel (a)—start to emerge in the SDF shortly after $m'/m \approx 0.2\%$ in panel (d). An important observation is that the maximum score (i.e., minimum energy) configuration (denoted by B) does not change much even with very few number of frequencies in panels (b) and (c). However, the slopes and curvatures of the energy profile characterizing the forces/torques and the stiffness of combined physical and geometric constraints do change significantly. For example, for the uniaxial motion from A to B, filtering with $m'/m < 1\%$ results in a relaxed collision response and

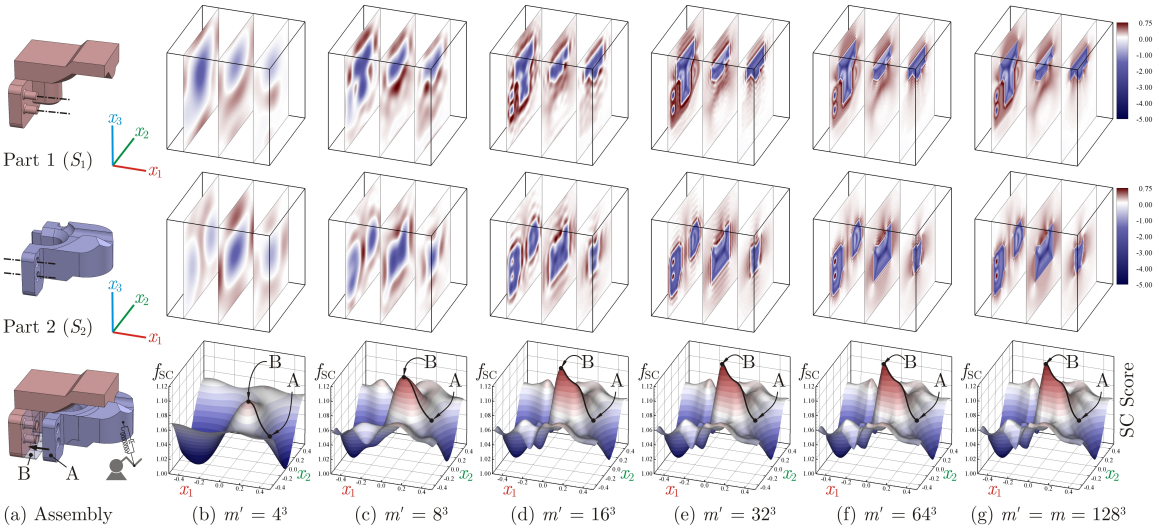

(a) Assembly (b) $m' = 4^3$ (c) $m' = 8^3$ (d) $m' = 16^3$ (e) $m' = 32^3$ (f) $m' = 64^3$ (g) $m' = m = 128^3$

FIGURE 5.4.2: The effect of FFT filtering on part SDFs (top) and score variations versus biaxial relative translation (bottom).

geometric guidance along BA with 'soft' snapping at B, as a result of faded geometric details. By increasing the precision with $m'/m \geq 1\%$, a brisker response is imposed by larger transverse slopes along the AB trajectory with 'hard' snapping at B due to sharper curvature. However, the changes are insignificant after $m'/m > 2\%$ with $\sigma = 0.5$, which enables speed-ups of two orders of magnitude by disposing of 98% of the frequency data. For more details, see [32] (Chapter 5).

5.5 Morphological Synthesis

Among the most important shape morphing problems that have been an integral part of solid modelers for a long time are offsetting (i.e., *growing* and *shrinking*), blending (i.e., *filleting* and *rounding*), sweeping (i.e., *extruding*, *revolving*, *lofting*, and general 2D and 3D feature sweeping), and similar operations along with their so-called 'dual' operations. These constructions are unified and generalized using

only a handful of algebraic operators defined in the setup of mathematical and group morphology and Dirac $\delta-$calculus, along with computational strategies such as low-pass filtering, anti-aliasing, limit representation via nascent $\delta-$functions, and other techniques from digital single processing (DSP). These suite of mathematical tools form the basis for 'spectral' geometric modeling, and opens up opportunities for developing software with unprecedented design and synthesis capabilities. Among the countless applications, I present a few morphological examples to demonstrate how a purely implicit geometric modeler can be developed from the presented analytic toolset (with a focus on $\delta-$calculus).

Figures 5.5.1 and 5.5.2 illustrate the simple Minkowski sum and difference, viewed as the generalizations of $\pm-$offsetting operation, and the morphological closing and opening, viewed as the generalizations of $\pm-$blending, of a 2D bunny with an oval shape. All operations are implemented implicitly in terms of convolutions of indicator functions. As long as one can respond to the membership query for each object, one can compute the membership query for the results of these operations directly from convolution of the constituent queries, either within a single CAD kernel or across two different CAD kernels. Furthermore, these operations can be altered by changing the shape descriptors from simple indicator functions (i.e., membership classification queries) to customized shape descriptors (e.g., membership grade queries) as illustrated by Fig. 5.5.3.

Figure 5.5.4 illustrates the sweep of the same 2D oval shape along a 1D trajectory, followed by an unsweep to obtain a closing of the trajectory in a similar fashion to the previous example with the bunny.

Figure 5.5.5 illustrates the sweep of the same 2D oval shape along a 1D trajectory, with and without rotation. A careful investigation of the implicit sweeps leads one to

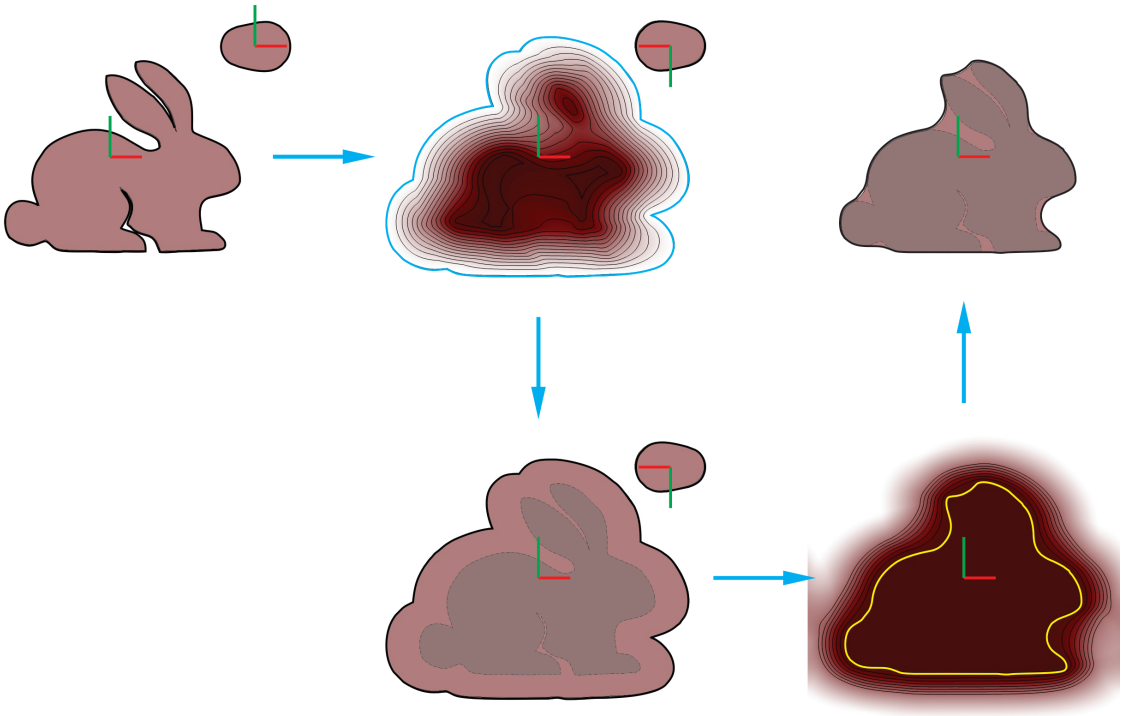

FIGURE 5.5.1: The Minkowski sum of a 2D bunny with an oval shape is the min-superlevel set of the convolution of their indicator functions. The Minkowski difference of the resulting shape's indicator function, in turn, with the same oval shape is the morphological closing, obtained as the max-superlevel set of a second convolution.

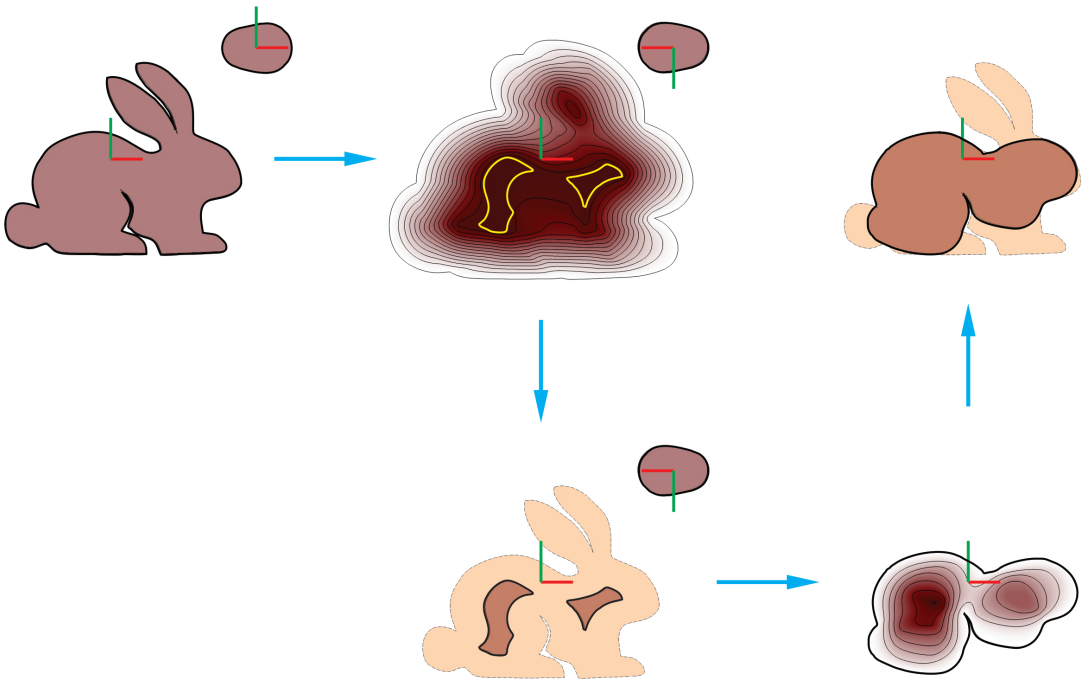

FIGURE 5.5.2: The Minkowski different of a 2D bunny with an oval shape is the max-superlevel set of the convolution of their indicator functions. The Minkowski sum of the resulting shape's indicator function, in turn, with the same oval shape is the morphological opening, obtained as the min-superlevel set of a second convolution.

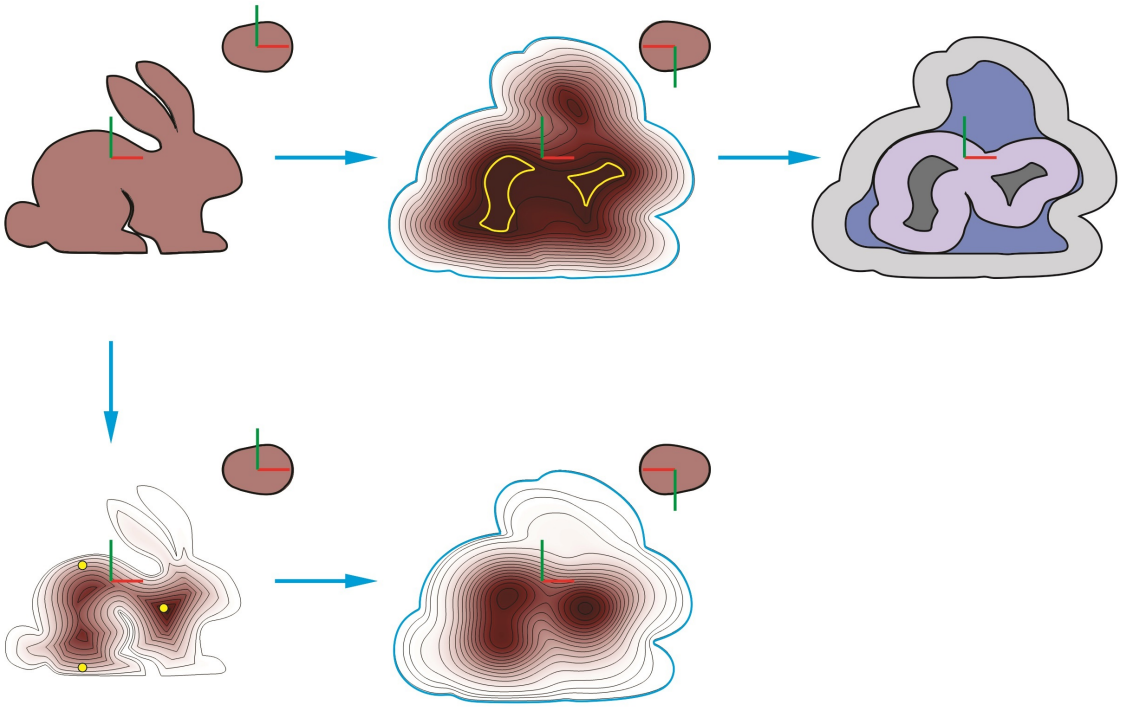

FIGURE 5.5.3: The effect of altering the shape descriptor by adding pressure from the indicator functions (i.e., membership classification queries) to to customized shape descriptors (e.g., membership grade queries) on the level sets of the convolution function and the resulting morphological constructs.

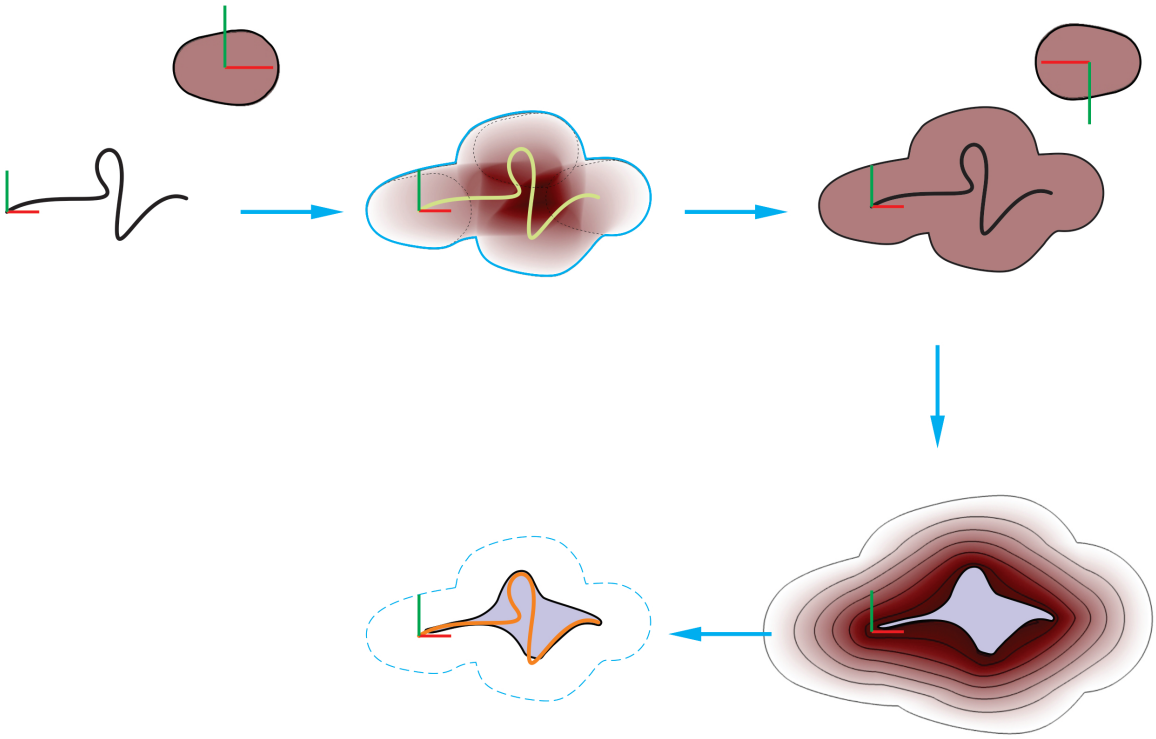

FIGURE 5.5.4: Sweep and unsweep operations performed back-to-back to obtain the closing of a 1D trajectory with a 2D oval shape.

the realization that fold regions and their boundaries can be distinguished by abrupt changes in the convolution function. Figure 5.5.5 also shows the approximation of the 1D curve and $0-$superlevel set reconstruction via a 2D Gaussian distribution centered along the curve followed by an $\epsilon-$superlevel set reconstruction, in order to convert the line integral into a 2D convolution.

Let us next consider a few results in 3D. Figure 5.5.6 shows implicit Minkowski operations on the Stanford bunny and a smaller shape, using convolutions of indicator functions $\mathbf{1}_{S_1}$ and $\mathbf{1}_{S_2}$ (illustrated via slices). To illustrate the degree of practical

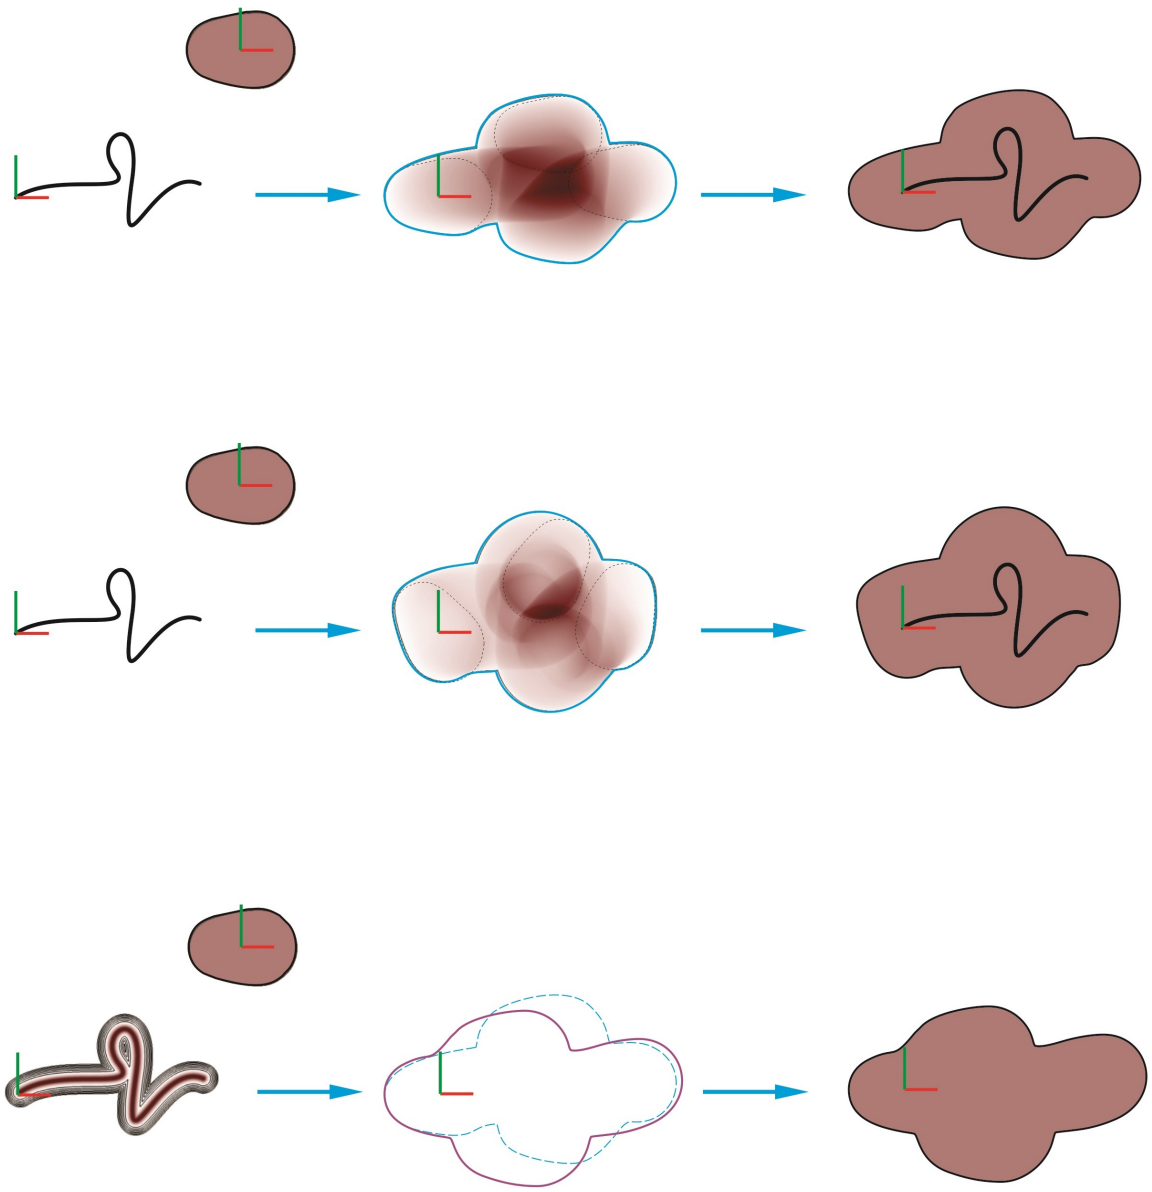

FIGURE 5.5.5: The sweep of a 2D oval shape along a 1D motion trajectory, without rotation (top) and with rotation (middle). The dissipation of the translational trajectory (top) into a Gaussian distribution yields an approximate sweep (bottom).

flexibility one gets from implicit modeling, consider now altering the bunny's implicit description from $\mathbf{1}_{S_1}$ to $f_{S_1} := \mathbf{1}_{S_1} \bullet (1 + \rho_0(\delta^3 \circ \mathsf{act}^{-1}(\mathbf{x}_0)))$, i.e., placing a $\delta-$singular point at $\mathbf{x}_0 \in S_1$. The effect of this 'pressure' point on the convolution function is depicted by Fig. 5.5.7 for a nascent $\delta-$function (namely, a mollifier). The $0-$level set is fixed in its place because I intentionally chose $f_{S_1} \equiv \mathbf{1}_{S_1}$, which could be chosen differently to achieve a different effect (e.g., fixing another level set). More elaborate synthesis activities are enabled by countless combinations of pressure curves and surfaces that can be added/removed, whose effects are quickly rendered by exploiting the convolution's properties (e.g., linearity), adding a level of flexibility that is unmatched by traditional solid modelers based on explicit descriptions. More importantly, these deformations are more than arbitrary aesthetic 'sculpting' with qualitative visual appeal, but are also useful for engineering synthesis as their precise semantics are quantifiable in terms of intersection measures.

To showcase the effectiveness of the strong morphology we developed to deal with heterogeneous dimensions, consider the same bunny, this time using an arbitrary subset cut out of its boundary along which to sweep the other shape. A $\delta-$distribution is obtained easily from the distance function as $(\delta^1 \circ \mathbb{d}_{S_1}) = \lim_{\sigma \to 0^+}(\mathcal{g}^1_\sigma \circ \mathbb{d}_{S_1})$, whose convolution with $\mathbf{1}_{S_2}$ yields a distribution for the sweep. Once again, if the bunny's surface descriptor is interpreted as a probability (e.g., Gaussian or 'normal') distribution, and the other shape's indicator function is normalized to give 'uniform' distribution, the convolution can be interpreted as the probability distribution for the sweep, whose different superlevel sets correspond to different confidence thresholds.[4]

Lastly, Fig. 5.5.9 shows a sweep operation in which both constituents are lower-

[4]This is only true if the inclusion in the two sets is viewed as independent events, to justify the multiplication of probabilities in the convolution formulation.

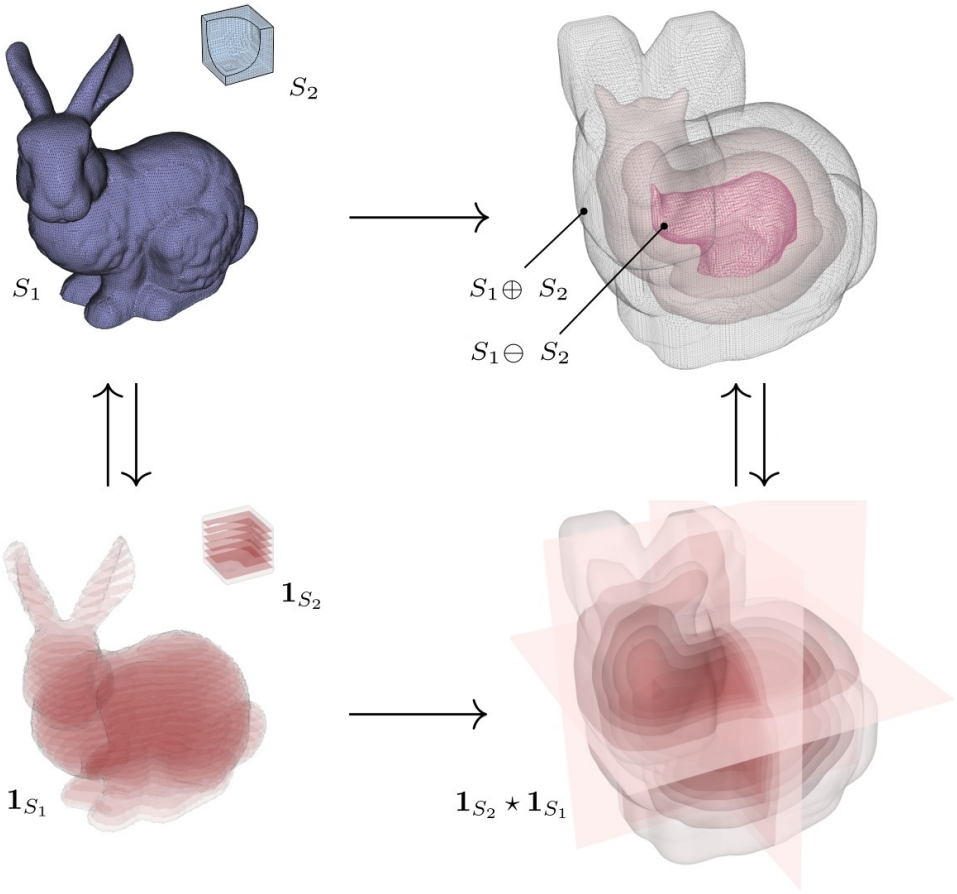

FIGURE 5.5.6: Full-dimensional implicit morphological synthesis. The min/max superlevel sets correspond to Minkowski operations.

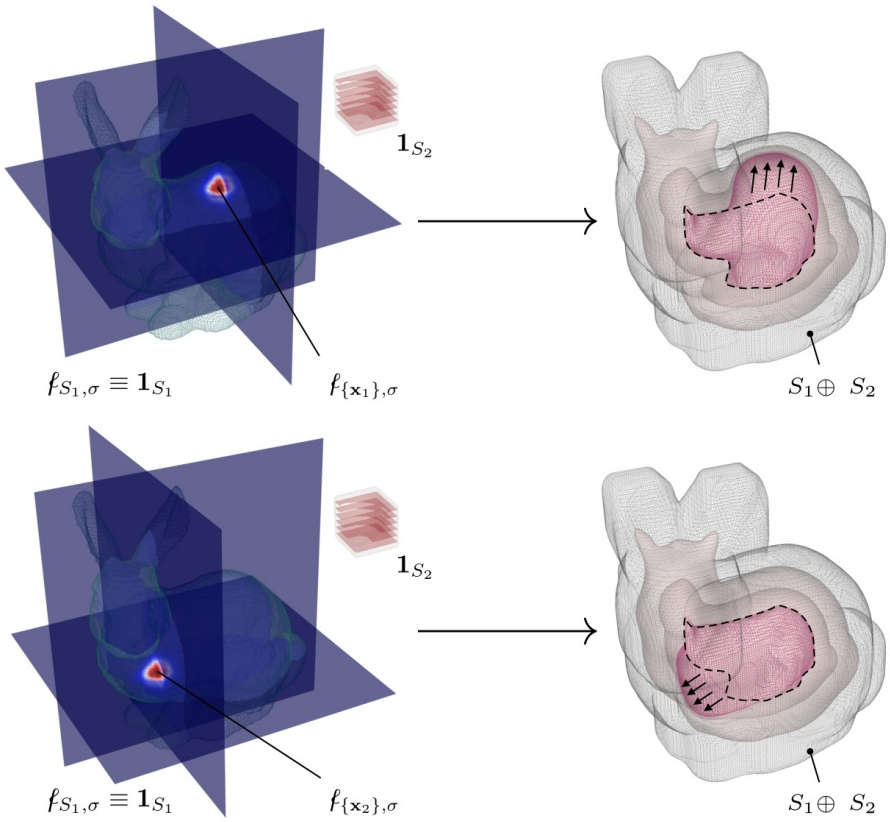

FIGURE 5.5.7: Singularities (e.g., pressure points) pull the level-sets of the convolution function towards the singularity.

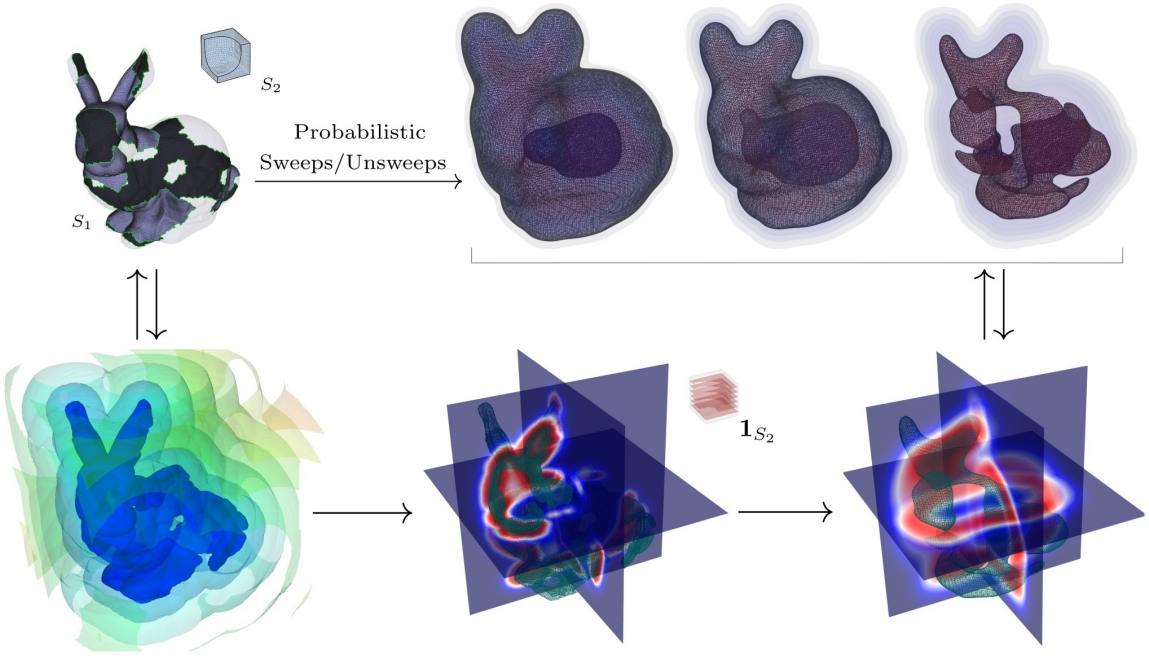

FIGURE 5.5.8: Sweeping a 3D shape along the 2D (partial) surface of the bunny, described as a Gaussian distribution (i.e., normal probability).

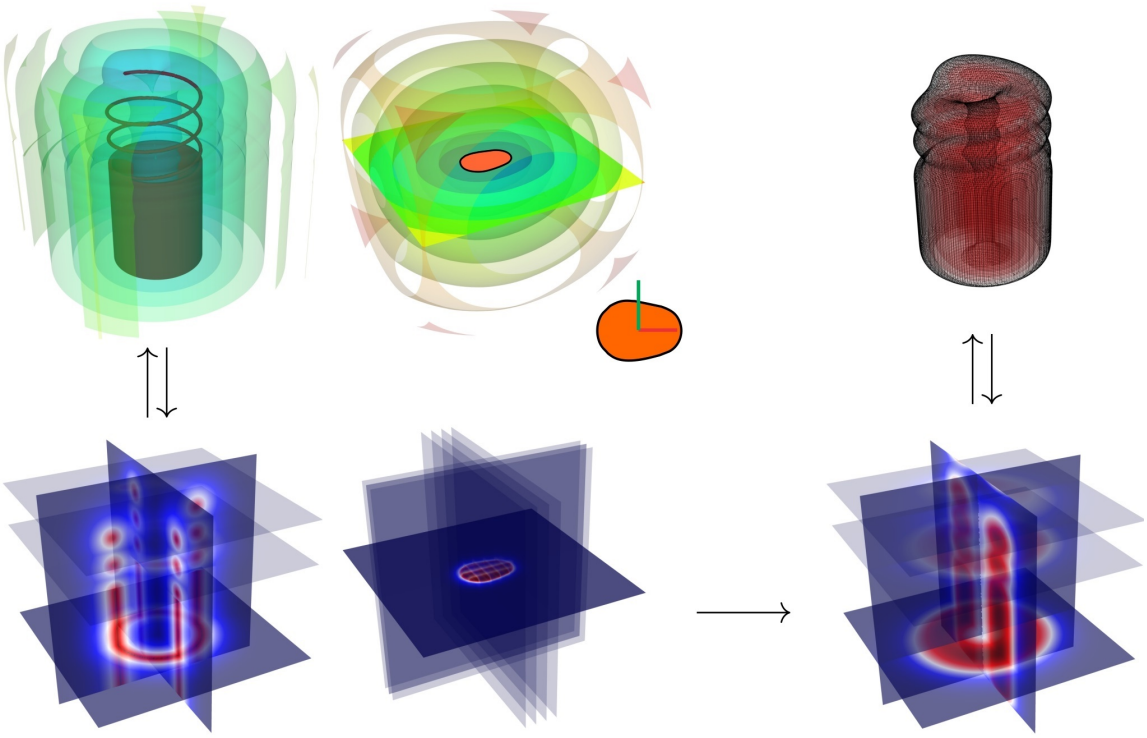

FIGURE 5.5.9: Sweeping a 2D shape along a 2D cylindrical surface and a 1D helical curve, described as a mollifier (i.e., compact tolerance zone).

dimensional elements: one is a cylindrical shell given by an implicit algebraic equation $x_1^2 + x_2^2 = 0.2^2$ for $-0.5 \leq x_3 < 0$ and the other is a helix given by a parametric equation as $x_1 = 0.2\cos(50t)$, $x_2 = 0.2\sin(50t)$, and $x_3 = t$ for $0 \leq t < +0.7$, which are characterized by $1^{\text{st}}-$ and $2^{\text{nd}}-$order $\delta-$singular shape descriptors, respectively. The other shape is the same planar oval given as a bitmap. Regardless of the representations, functional descriptors are computed one way or another based on distance functions or parametric forms, and the convolution is obtained for a finite diffusion. This time, instead of using a normal distribution, a compactly supported bump function is used to ensure bounded level sets for the convolution, leading to a tolerance zone picture of uncertainty.

5.6 Manufacturing Planning

Another important application of analytic methods is in developing powerful tools and robust systems for design feedback and modification [233, 237], manufacture process planning [125, 232], and tolerancing [253, 279]. Here I illustrate some of the benefits in both subtractive and additive manufacturing (SM/AM).

Figure 5.6.1 shows a simple 2D workpiece whose $2-$axis milling with a cutter, modeled by its circular cross-section, is of interest. The part has features that are smaller than the cutter, hence an accessibility analysis needs to be performed to obtain the maximum volume that can be removed from a block of material to reduce it as close as possible to the final part. Depending on whether the geometric representation is of the part itself or its negative image (e.g., in pocket milling), the Minkowski sum and difference represent the collection of displacements of the cutter center for

maximal volume removal, and the morphological closing and opening represent the maximum removed volume, respectively. As before, these operations can be computed implicitly via convolutions, from which the explicit volume representations can be recovered as superlevel sets. More importantly, it does not matter if the workpiece geometry and cutter information are represented in different modelers, as long as they both support membership queries needed to compute the convolutions via sampling, FFT, etc. Figure 5.6.3 depicts the same accessibility analysis performed for milling the exterior or interior of the same shape using different cutter sizes.

It is sometimes important to characterize the contact area between the cutter and the workpiece, which affects the surface quality, tolerancing, etc. As depicted by Fig. 5.6.2, the extent of instantaneous contact between the cutter and the final shape of the workpiece can be obtained by a cross-correlation of their boundary descriptors, approximated via Gaussian boundary density functions. The ability to compute the correlations rapidly on the GPU via FFT allows dynamic real-time monitoring as the part's geometry changes during the machining process, to obtain a lot more information than that of an offline analysis against the final profile.

Figure 5.6.4 illustrates the geometry of a single layer of material deposition for a 3D printed part, using a printer with a simple circular nozzle cross-section. In this case, in addition to shape analysis for obtaining maximum printable subset at each layer, we are also interested in shape synthesis; namely, to make near-minimal changes to the shape to remove small features (i.e., thicken them) just enough for the printer to be able to print them. This can be performed by sweeping the nozzle cross-section along the shape skeleton, which, in turn, can be computed as a convolution of the skeletal density function with the nozzle descriptor. The union of the sweep shape with the original shape (i.e., functional summation) of the layer results in the

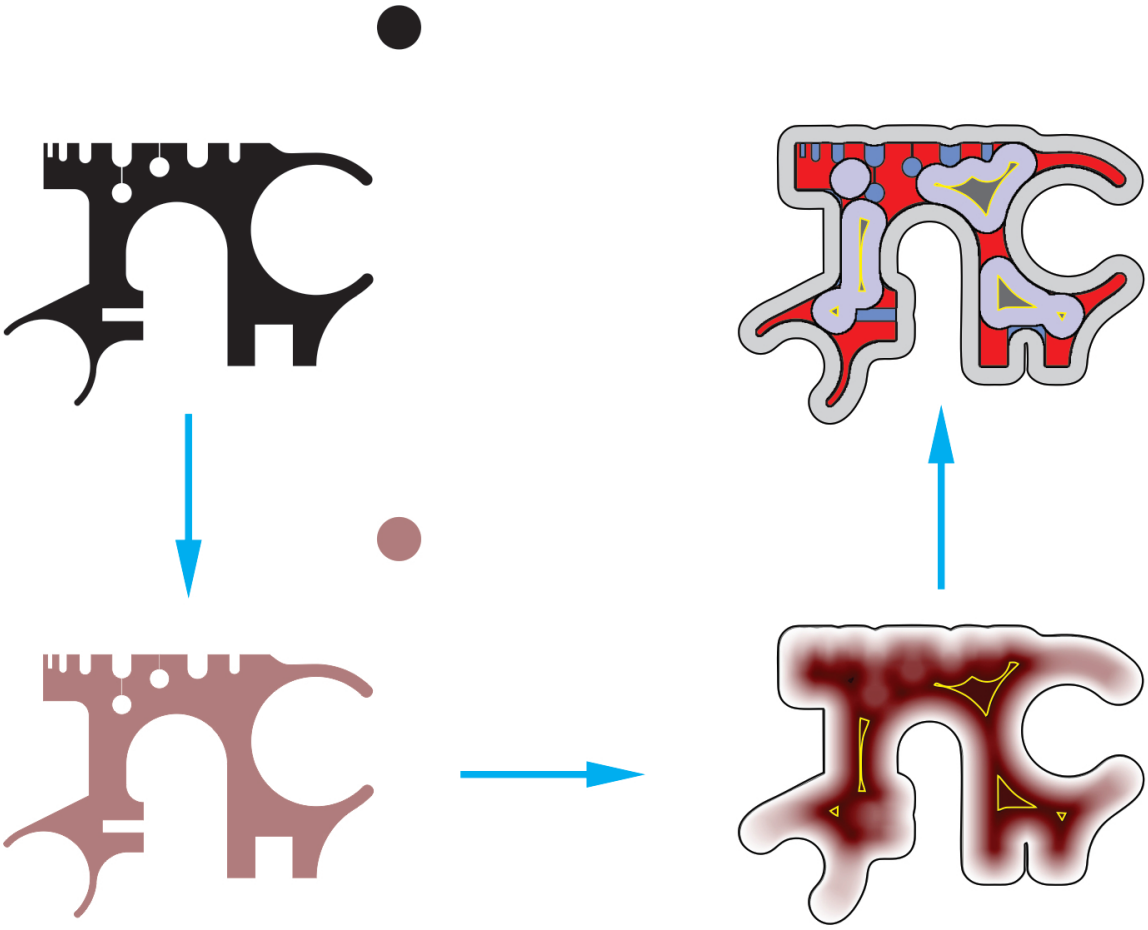

FIGURE 5.6.1: The cutter displacement space and maximum removable volume for 2−axis milling can be formulated in terms of Minkowski operations in the explicit realm, which are superlevel sets of convolutions in the implicit realm.

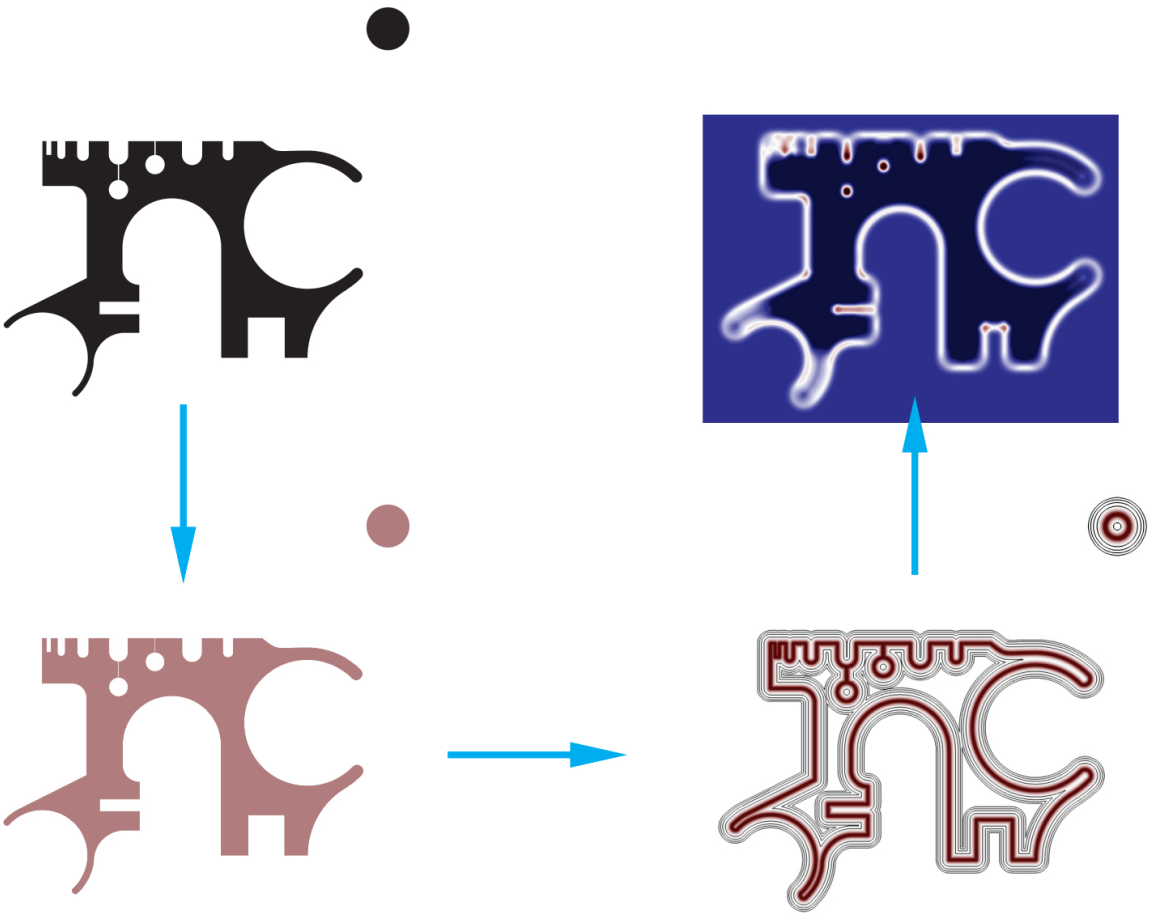

FIGURE 5.6.2: The extent of contact between the workpiece (in its final form) and the cutter can be characterized by cross-correlating their Gaussian boundary density functions, which, in turn, can be interpreted as the probabilistic description of the uncertainties in manufacturing (e.g., tolerancing cutter profile and movement).

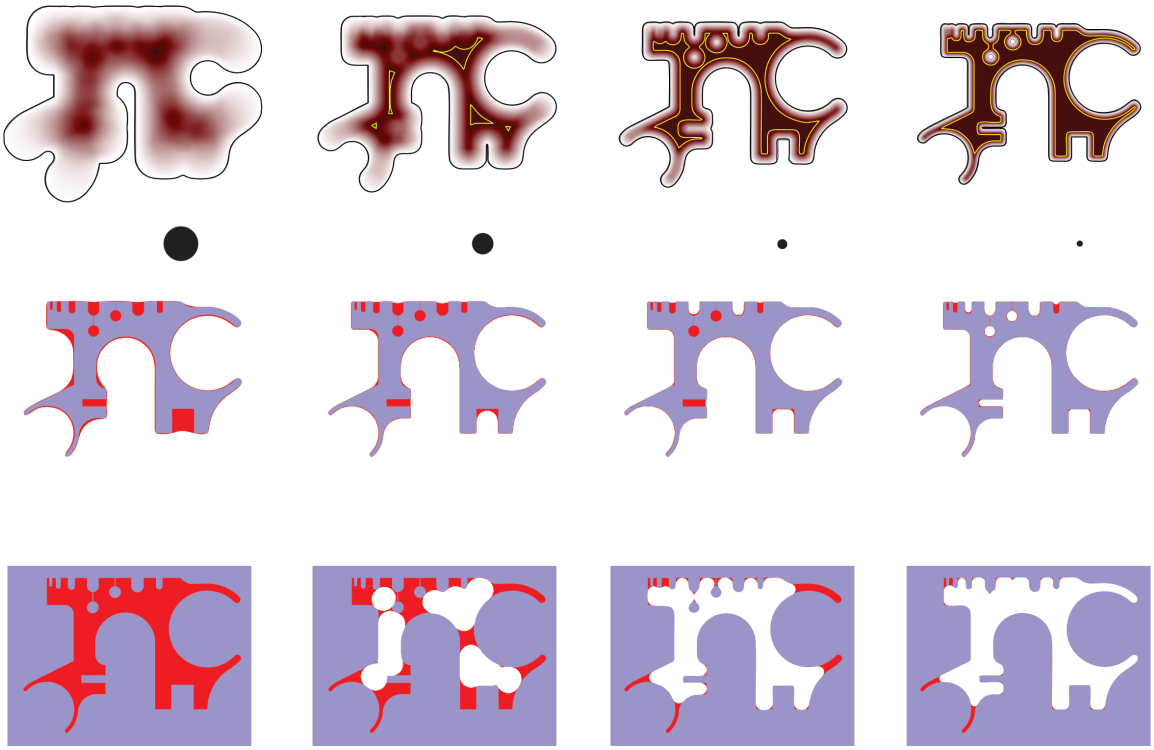

FIGURE 5.6.3: The morphological closing and opening of a given shape with different cutter cross-sections reveal the maximum removable volume from the interior and exterior, respectively. The top row shows the convolution profile and its min/max-level sets corresponding to the Minkowski sum and difference boundaries, whose second convolution with the cutter cross-section gives the closing (middle) and opening (bottom).

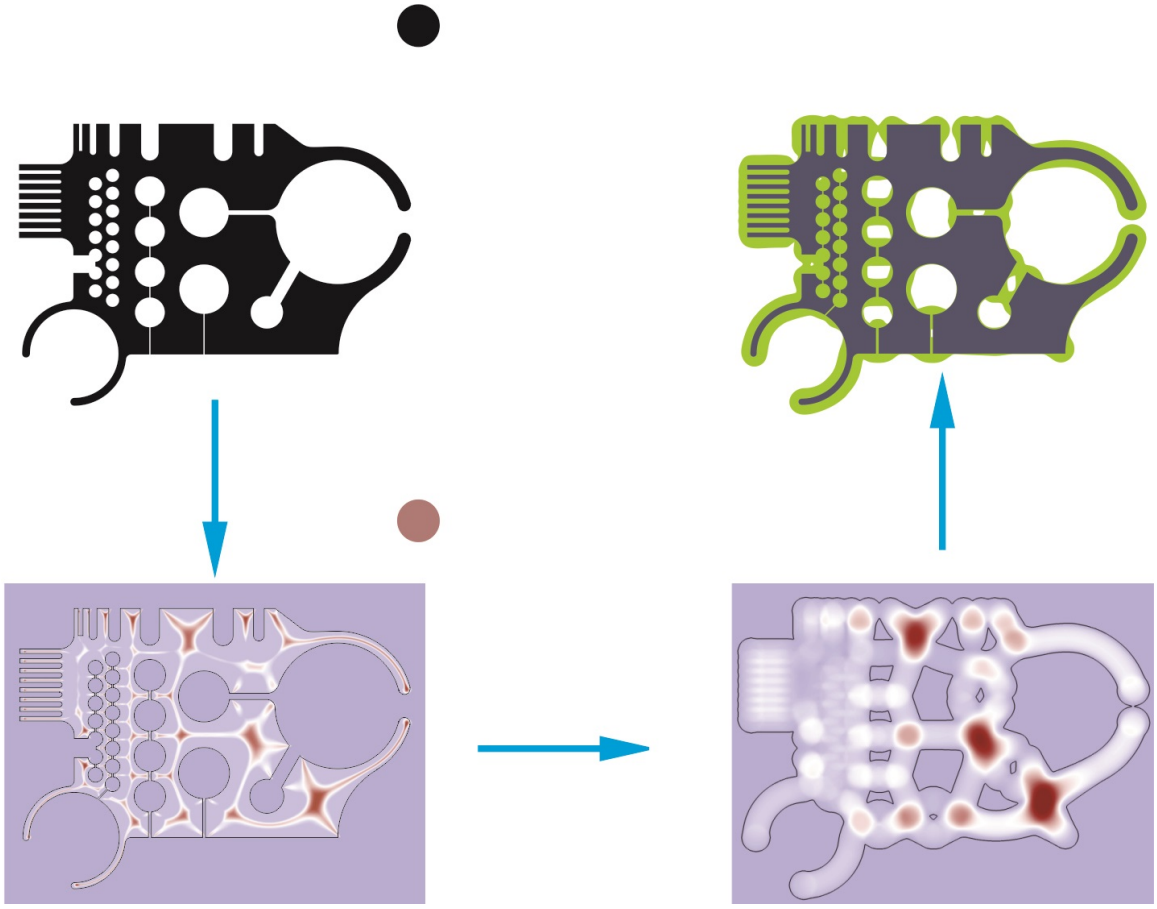

FIGURE 5.6.4: A 2D layer of a 3D printed workpiece can be corrected for its sharp corners and small non-printable features by a convolution of its skeletal density function with the printer head's indicator function.

desired synthesis. This is a simple example of how to use analytic modeling to tie shape synthesis to manufacturability. The same analysis and synthesis is repeated for multiple nozzle sizes in Fig. 5.6.5.

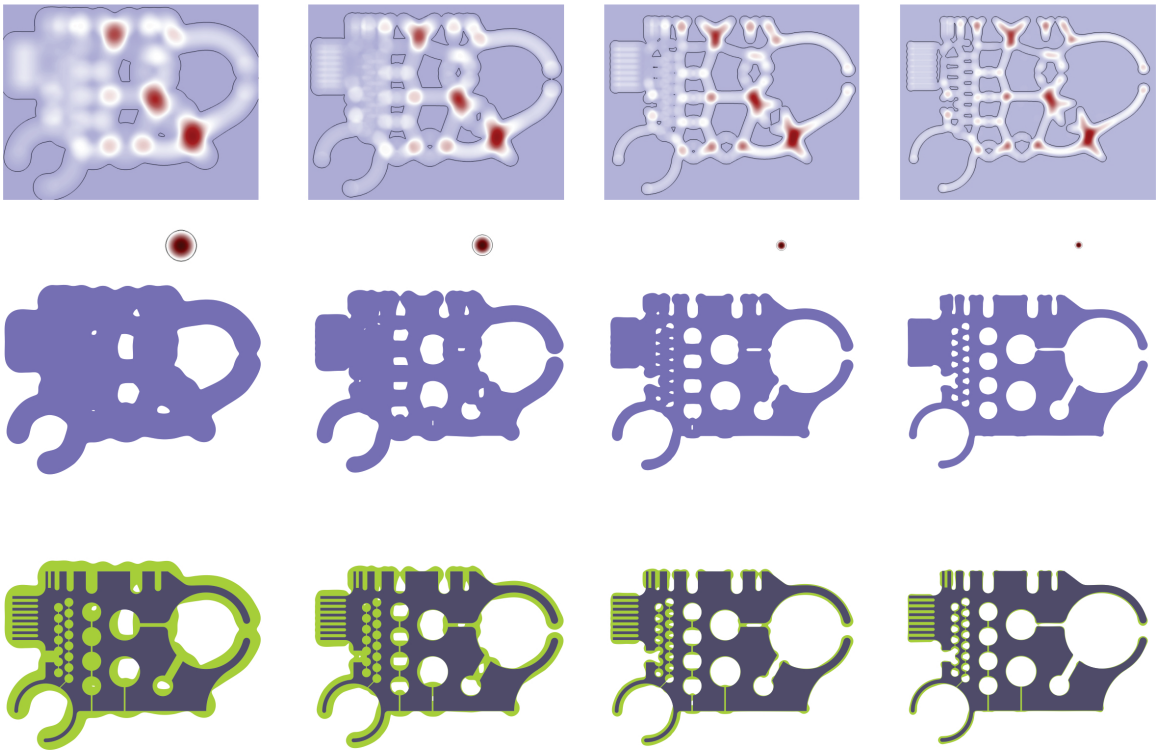

FIGURE 5.6.5: The convolution of the skeletal density function with the printer nozzle's shape descriptor—this time characterized by a Gaussian disk of various radii—yields the extra material needed to correct the non-printable features.

5.7 Query-Based Interoperability

Perhaps the most important reason to take the analytic (i.e., functional) approach to geometric computing seriously, is how it provides a basis for query-based interoperability [147, 148]. The existing geometric modeling implementations are a collection of representations, authored with respect to a particular representation scheme and operated on by representation-specific algorithms:

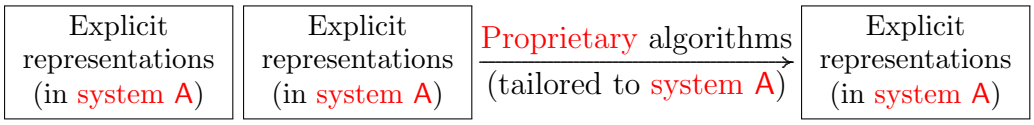

For example, a Boolean operation designed for a particular CSG scheme obviously cannot be used on a B-rep, or a sweep computing algorithm on a triangular mesh B-rep cannot possibly work on a NURBS B-rep, and so on. For commercial systems, the algorithms are often proprietary, as depicted above.

The same can be realized for an implementation based on implicit modeling:

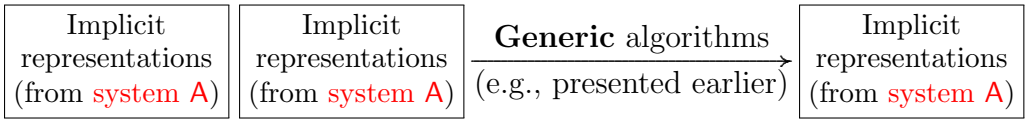

In both of the above scenarios, the inner working of the algorithm and its correctness relies on the semantics of the representation scheme—obviously making it ineffective on a different data structure with a different set of rules and semantics.

Now, an implementation based on implicit models may be a completely new system, developed from scratch (e.g., new software platform), or it can be an encapsulation of an existing implementation (perhaps explicit), wrapped into a 'black box' with which one communicates only by functional input/output semantics. For example, if

the implicit modeling is based on membership or distance functions, which I showed to be sufficient for virtually all important computations that one cares about (at least as much as discussed in this thesis), then any system (implicit or explicit) can be abstracted in terms of its functional interface, as long as it guarantees support for those two queries, providing some well-specified guarantees for its correctness (e.g., in terms of error tolerances). The internal storage and implementation, both for data structures and algorithms, are irrelevant, as long as the interface guarantees are delivered by the system. Such an encapsulation and data hiding is extremely valuable from a commercial and business strategic perspective, as it protects proprietary information. The encapsulation process can be illustrated via

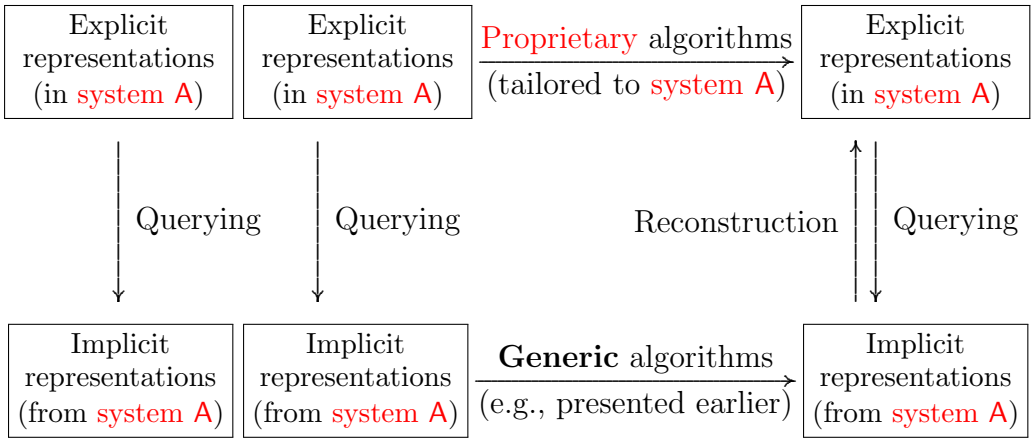

The generic algorithms are combinations of the basic functional combining forms discussed throughout this thesis, i.e., arithmetics of functions, convolution and its relatives (e.g., cross-correlations, inner products, norms, and more generally, integrals) computed via DFTs (or FFTs) or NDFTs (or NFFTs), and composition with various orders of $\delta-$function computed using its limit approximations.

Importantly, now we can do something that was not possible with the explicit implementations; namely, not only we can apply these generic algorithms to implicit

encapsulations of multiple objects from the same system, but also we can apply them to multiple objects *across* different systems:

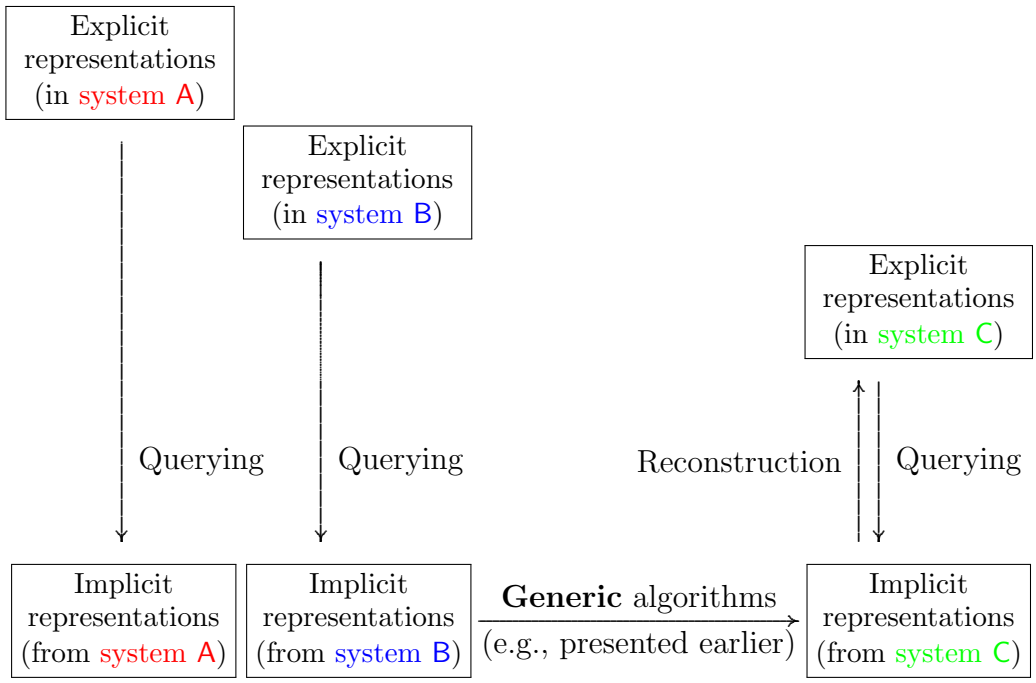

For example, now we can compute the (queries on) Boolean operations across a CSG system and a B-rep system, provided by different vendors, as long as both support membership queries with consistent semantics—namely, by querying their indicator or nonnegative defining functions at a given point whose coordinates are specified with respect to an agreed-upon world coordinate system, adding (for union) or multiplying (or intersection) the two, regularizing by $\delta-$convolution (if necessary), and returning the result as the membership query response against the hypothetical new object. Similarly, one can compute the Minkowski sum of a workpiece specified in one system and a cutting tool specified in a different system by querying both for a dense enough sample of points, computing their convolution function via FFT (over a uniform grid) or NFFT (using spherical decomposition of Chapter 4), and returning the result as the

membership query response against the hypothetical Minkowski sum (e.g., for CNC machining as in Section 5.6). Of course if boundary evaluation or skeletonization are required, the systems need to support signed distance functions as well (e.g., for 3D printing as in Section 5.6). This provides the basis for interoperable design and manufacturing planning.

To summarize, the proposed paradigm provides a basis for standardization of interfaces between computational design and manufacturing systems in terms of basic queries—e.g., membership and distance, and if needed, various density functions and shape descriptors presented in Section 5.1.1 for various applications in subsequent sections. These queries can be combined, based on which the system can be composed together via generic functional combining forms—as in functional programming [22]—that are as simple as basic arithmetics, convolution algebras, and compositions with the various orders of the $\delta-$function. For ultimate performance and parallel implementation using the modern high-performance computing (HPC) systems (e.g., multi-core CPUs and many-core GPUs), standard techniques from digital signal processing (DSP) can be employed as discussed numerous times throughout this thesis.

Current limitations on interoperability and compositionality are among the major bottlenecks whose effective solution requires identification of the proper semantics for the interfaces between computational tools, human operators, and physical products and systems. Similar to the transition from imperative (i.e., von Neumann's) style to declarative (e.g., object-oriented and functional) style in the realm of computer programming and software engineering, a similar transition seems to be in order for geometric modeling and CAD/CAM tools. The query-based approach appears to provide a sound basis for this transition.

Conclusions

> "We must face the fact that we are on the brink of times when man may be able to magnify his intellectual and inventive capability, just as in the nineteenth century he used machines to magnify his physical capacity. Again, as then, our innocence is lost. And again, of course, the innocence, once lost, cannot be regained. The loss demands attention, not denial."
>
> Christopher Alexander, 1964 [9]

Or at least, times when the existing models, representations, and software systems for computational design and manufacturing are not sufficient to support the growing engineering requirements and democratizing manufacturing capabilities. This thesis was an effort to advocate for a paradigm shift in the way we think about abstraction, implementation, and reasoning for computational design and manufacturing to open up new opportunities for research, development, and commercialization.

An alternative paradigm is emerging to perform geometric and spatial computations on implicit descriptions as spatial distribution fields (i.e., signals) instead of homogeneous pointsets. It provides access to a set of measure-theoretic power tools with substantial theoretical and computational advantages over the traditional discrete geometric counterparts. It is particularly attractive due to the formalism it provides to support query-based interoperability as well as its generic and wide

applicability, computational efficiency, and robustness in the presence of errors/noise.

I showed that various shape and motion related problems can be uniformly conceptualized in terms of (a systems ability to respond to) an handful of basic functions (i.e., queries) such as membership and distance functions—whose semantics are universal—and a small collection of functional combinators limited to arithmetics, convolution algebra, and Dirac $\delta-$calculus. In other words, virtually all fundamental geometric modeling and spatial reasoning operations that are relevant to design and manufacturing can be expressed as an algebra of shape descriptors, which can be approximated or generalized into various density functions, providing a basis for moving from crisp solid modeling to a more versatile fuzzy and probabilistic reasoning. I demonstrated the effectiveness, robustness, and flexibility of the proposed methodology on examples ranging from morphological synthesis and manufacture planning to haptic assembly and protein docking.

Most of the basic knowledge for such a formulation is not new—at least as old as basic measure theory, distribution theory, harmonic analysis, and other tools leveraged by digital signal processing (DSP). However, their application to geometric modeling, computing, and reasoning has been limited, especially due to the seeming inability to handle geometric feature of heterogeneous dimensions. In this regard, this thesis provides a practical approach to removing one of the most crippling drawbacks of this analytic paradigm, which was its inability to compute on lower-dimensional elements that are invisible to the measure-theoretic lens. In particular, the convolution algebra in its traditional form is a week implicit substitute for the morphological operators in the presence of lower-dimensional elements. These elements include curves and surfaces that can serve as trajectories for sweep/unsweep and other morphological operations. I presented a $\delta-$calculus to extend the convolution algebra into

a strong morphology by assigning "artificial" volume singularities of proper orders to lower-dimensional elements and showed how to construct such descriptors from both explicit (e.g., parametric) and implicit (e.g., distance-based) descriptions. I also argued how this formalism provides a natural mechanism to turn traditionally non-computable problems into computable ones by 'dissipating' the $\delta-$singularities into the surrounding volume, providing a robust computational framework with clear trade-offs against fidelity to geometric details.

This research opens up promising theoretical and computational directions for future studies on analytic methods that are emerging in solid and physical modeling. Moving forward, I find the following directions worthy of further investigation:

- Category theory appears to be a useful tool to explore, especially with regards to formalizing the transformations and diagrams presented throughout this thesis.

- As multi-scale material structures—which I did not touch in this document—start appearing in implicit models, their description may require extending functions and fields (i.e., continuum forms) into co-chains (i.e., discrete forms). Form an implementation perspective, this corresponds to extending point queries (e.g., PMC or distance at a single point) with measure queries on various spatial elements (e.g., points, line segments, triangles, and tetrahedra) as inputs.

- The re-conceptualization of geometric modeling as high-dimensional signal processing, transitioning from explicit to implicit modeling, from exhaustive to lazy evaluation, and from imperative to declarative conceptualization appear to be inevitable in the face of increasing complexity of design.

Appendix A

Rigid Motion Group

This appendix reviews some basic concepts related to the 'special Euclidean group' $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathbb{R}^3$, which is the group of rigid motions with the customized notation that I have adopted in this thesis. The main purpose is to provide a quick guide to the essentials, without getting into level of rigor and detail that a formal text on the subject would. The readers is referred to [138, 242] for more details.

A.1 Group Operations

The (noncommutative) group operation over the 'special Euclidean group' $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathbb{R}^3$ denoted by $(\mathbb{c}_1, \mathbb{c}_2) \mapsto \mathbb{c}_1\mathbb{c}_2$ is represented, after fixing an origin, in terms of the (noncommutative) group operation over the 'special orthogonal group' $\mathrm{SO}(3)$ denoted by $(\mathbb{r}_1, \mathbb{r}_2) \mapsto \mathbb{r}_1\mathbb{r}_2$, the group action of $\mathrm{SO}(3)$ on $\mathrm{T}(3) \cong \mathbb{R}^3$ denoted by $(\mathbb{r}_1, \mathbf{t}_2) \mapsto \mathbb{r}_1\mathbf{t}_2$, and the (commutative) group operation over $\mathrm{T}(3) \cong \mathbb{R}^3$ which is a

vector sum $(\mathbf{t}_1, \mathbb{r}_1\mathbf{t}_2) \mapsto \mathbf{t}_1 + (\mathbb{r}_1\mathbf{t}_2)$ in the mentioned order:

$$\mathbb{c}_1\mathbb{c}_2 \overset{0}{=} (\mathbb{r}_1, \mathbf{t}_1)(\mathbb{r}_2, \mathbf{t}_2) := (\mathbb{r}_1\mathbb{r}_2, \mathbf{t}_1 + (\mathbb{r}_1\mathbf{t}_2)), \tag{A.1.1}$$

noting that the translation part is affected by one of the rotations and both translations. All three of them can be computed using matrix multiplications if represented by transformation matrices and column vectors; namely, $[\mathbb{c}_1\mathbb{c}_2]_{4\times4} = [\mathbb{c}_1]_{4\times4}[\mathbb{c}_2]_{4\times4}$, $[\mathbb{r}_1\mathbb{r}_2]_{3\times3} = [\mathbb{r}_1]_{3\times3}[\mathbb{r}_2]_{3\times3}$, and $[\mathbb{r}_1\mathbf{t}_2]_{3\times1} = [\mathbb{r}_1]_{3\times3}[\mathbf{t}_2]_{3\times1}$, respectively. As mentioned before, I choose to use the latter two, tupled together as $\big([\mathbb{r}_1\mathbb{r}_2]_{3\times3}, [\mathbf{t}_1]_{3\times1} + [\mathbb{r}_1\mathbf{t}_2]_{3\times1}\big)$ instead of $[\mathbb{c}_1\mathbb{c}_2]_{4\times4}$ to represent $\mathbb{c}_1\mathbb{c}_2$ as it lends itself better to the lifting/projection and slicing techniques detailed in Section 3.4. However, both representations are different arrangement of the same real numbers from an implementation perspective with the same computational work required to compute group operations.

The group inversion over $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathbb{R}^3$ denoted by $\mathbb{c} \mapsto \mathbb{c}^{-1}$ is represented, after fixing an origin, in terms of the group inversion over $\mathrm{SO}(3)$ denoted by $\mathbb{r} \mapsto \mathbb{r}^{-1}$ and the group action of $\mathrm{SO}(3)$ on $\mathrm{T}(3) \cong \mathbb{R}^3$ denoted by $(\mathbb{r}^{-1}, -\mathbf{t}) \mapsto \mathbb{r}^{-1}(-\mathbf{t})$ as

$$\mathbb{c}^{-1} \overset{0}{=} (\mathbb{r}, \mathbf{t})^{-1} := (\mathbb{r}^{-1}, \mathbb{r}^{-1}(-\mathbf{t})) = (\mathbb{r}^{-1}, -\mathbb{r}^{-1}\mathbf{t}). \tag{A.1.2}$$

Since the latter is a linear map acting on the vector space $\mathbb{R}^3$ (i.e., rank-2 tensor) which is special orthogonal (i.e., preserving lengths and signed angles), group inversion can be replaced with tensor transposition:

$$\mathbb{c}^{-1} \overset{0}{=} (\mathbb{r}, \mathbf{t})^{-1} := (\mathbb{r}^{\mathrm{T}}, \mathbb{r}^{\mathrm{T}}(-\mathbf{t})) = (\mathbb{r}^{\mathrm{T}}, -\mathbb{r}^{\mathrm{T}}\mathbf{t}), \tag{A.1.3}$$

noting that the translation part is affected by inverse rotation and reflection. All three

of them can be computed using matrix multiplications if represented by transformation matrices and column vectors; namely, $[\mathbb{c}^{-1}]_{4\times4} = [\mathbb{c}]_{4\times4}^{-1}$, $[\mathbb{r}^{-1}]_{3\times3} = [\mathbb{r}]_{3\times3}^{-1}$, i.e., $[\mathbb{r}^{\mathrm{T}}]_{3\times3} = [\mathbb{r}]_{3\times3}^{\mathrm{T}}$, and $[\mathbb{r}^{-1}(-\mathbf{t})]_{3\times1} = -[\mathbb{r}]_{3\times3}^{-1}[\mathbf{t}]_{3\times1}$, i.e., $[\mathbb{r}^{\mathrm{T}}(-\mathbf{t})]_{3\times1} = -[\mathbb{r}]_{3\times3}^{\mathrm{T}}[\mathbf{t}]_{3\times1}$, respectively. Note that there is a difference between tensor inversion and transposition (defined more abstractly) with matrix inversion and transposition (defined in terms of elements), respectively, though the latter are representations of the former for rank−2 tensors [195].

The group action of $\mathrm{SE}(3) \cong \mathrm{SO}(3) \rtimes \mathbb{R}^3$ on $\mathbb{R}^3$ denoted by $(\mathbb{c}, \mathbf{x}) \mapsto \mathbb{c}\mathbf{x}$ is represented, after fixing an origin, in terms of the group action of $\mathrm{SO}(3)$ on $\mathbb{R}^3$ denoted by $(\mathbb{r}, \mathbf{x}) \mapsto \mathbb{r}\mathbf{x}$ and the group action of $\mathrm{T}(3) \cong \mathbb{R}^3$ on $\mathbb{R}^3$ which is a vector sum $(\mathbb{r}, \mathbf{x}) \mapsto (\mathbb{r} + \mathbf{x})$ as

$$\mathbb{c}\mathbf{x} \overset{0}{=} (\mathbb{r}, \mathbf{t})\mathbf{x} := (\mathbb{r}\mathbf{x}) + \mathbf{t} \quad \rightleftharpoons \quad \mathbb{c}^{-1}\mathbf{x} \overset{0}{=} (\mathbb{r}, \mathbf{t})^{-1}\mathbf{x} = \mathbb{r}^{\mathrm{T}}(\mathbf{x} - \mathbf{t}), \tag{A.1.4}$$

noting that the (noncommutative) sequence of applying rotations and translations is inverted when the motion itself is inverted.

Like with every Riemannian manifold, a geodesic distance function over the $\mathrm{SE}(3)$ denoted by $\mathbb{d}(\mathbb{c}_1, \mathbb{c}_2)$ is computable (after fixing an origin) in terms of the distance function over $\mathrm{SO}(3)$ denoted by $\mathbb{d}(\mathbb{r}_1, \mathbb{r}_2)$ and the distance function over $\mathrm{T}(3) \cong \mathbb{R}^3$ denoted by $\mathbb{d}(\mathbf{t}_1, \mathbf{t}_2)$ for $\mathbb{c}_1 \overset{0}{=} (\mathbb{r}_1, \mathbf{t}_1)$ and $\mathbb{c}_2 \overset{0}{=} (\mathbb{r}_2, \mathbf{t}_2)$ via generalized Pythagorean formula (i.e., cosine rule):

$$\mathbb{d}^2(\mathbb{c}_1, \mathbb{c}_2) = \left\| \log(\mathbb{c}_1^{-1}\mathbb{c}_2) \right\|_2^2 = \left\| \log(\mathbb{r}_1^{\mathrm{T}}\mathbb{r}_2) + (\mathbf{t}_2 - \mathbf{t}_1) \right\|_2^2 \tag{A.1.5}$$

$$= \left\| \log(\mathbb{r}_1^{\mathrm{T}}\mathbb{r}_2) \right\|_2^2 + \left\| \mathbf{t}_2 - \mathbf{t}_1 \right\|_2^2 + 2\log(\mathbb{r}_1^{\mathrm{T}}\mathbb{r}_2) \bullet (\mathbf{t}_2 - \mathbf{t}_1), \tag{A.1.6}$$

where $+, \bullet : \mathfrak{so}(3) \times \mathbb{R}^3 \to \mathbb{R}^3$ should be understood as vector sum and dot product between the dual vector of a skew-symmetric matrix in $\log(\mathbb{r}) \in \mathfrak{so}(3)$ (i.e., an axis-angle representation of rotation $\mathbb{r} \in \mathrm{SO}(3)$) and a translation vector $\mathbf{t} \in \mathbb{R}^3$, noting also that $\log(\mathbb{c}) = \log(\mathbb{r}) + \mathbf{t}$ for $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$.

A.2 Representations

Representations of rigid motions of the 3−space are typically given by 3×3 matrices (for rotations) tupled with a 3×1 column vector (for translations), 4×4 matrices (for rotations and translations), dual quaternions, etc.

In linear algebra literature, SO(3) itself is sometimes viewed as being identical (up to isomorphism) with the collection of 3×3 real matrices $R \in \mathbb{R}^{3\times3}$ that are orthogonal—i.e., $R^{\mathrm{T}} = R^{-1}$ meaning that $\det(R) = \pm1$—and reflection-free (i.e., special)—i.e., $\det(R) > 0$ therefore $\det(R) = +1$. Then the Lie algebra $\mathfrak{so}(3)$ becomes the collection of 3×3 real matrices $\Omega \in \mathbb{R}^{3\times3}$ that are skew-symmetric—i.e., $\Omega = -\Omega^{\mathrm{T}}$ meaning that its 3 diagonal elements are zero, and of the other 6 elements, 3 are redundant. The dual vector of Ω obtained from the nonzero and nonredundant elements provides an axis-angle representation of the rotation $R := e^{\Omega}$. Rotations can also be represented by unit quaternions, whose collection is isomorphic to the unit 3−sphere embedded in $\mathbb{R}^4$. In fact, quaternions are in 1:1 homeomorphism with the elements of SU(2) defined as the collection of 2×2 complex matrices $U \in \mathbb{C}^{2\times2}$ that are unitary—i.e., $R^{\mathrm{T}} = \bar{R}^{-1}$ meaning that $|\det(U)| = 1$—which, in turn, are in 2:1 homeomorphism with the elements of SO(3). Other options are popular in dynamics and aerospace such as Euler angles and Tait-Bryan angles. Each represen-

tation has its own benefits and drawbacks, thus I prefer an abstract group-theoretic conceptualization without appealing to a particular representation—hence the notation $\mathbb{r} \in \mathrm{SO}(3)$ for an abstract rotation and $R = [\mathbb{r}]_{3\times3}$ for its matrix representation, instead of using rotations and rotation matrices interchangeably.

Note that a similar discussion applies to translations, which can be represented via Cartesian, polar-cylindrical, polar-spherical, and other coordinate systems. However, unlike rotations and rotation matrices, translations and translation vectors are used almost interchangeably, i.e., $\mathbb{t} \in \mathrm{T}(3)$ represented by $[\mathbf{t}]_{3\times3} \in \mathbb{R}^3$ are thought to be the same entities. I find this convention a good balance between rigor and simplification, as long as Cartesian coordinate systems are used.

Accordingly, SE(3) is sometimes viewed as being identical (up to isomorphism) with the collection of 4×4 real matrices $T \in \mathbb{R}^{4\times4}$ that are called 'homogeneous transformation matrices':

$$\mathrm{SE}(3) \cong \left\{ \begin{pmatrix} R & \mathbf{t} \\ \mathbf{0}^{\mathrm{T}} & 1 \end{pmatrix} \;\middle|\; R \in \mathrm{SO}(3) \text{ and } \mathbf{t} \in \mathbb{R}^3 \right\}. \tag{A.2.1}$$

Alternatively, dual quaternions or screws can be used to represent rigid body transformations, each with their own benefits and drawbacks. Once again, I prefer an abstract group-theoretic conceptualization without appealing to a particular representation—hence the notation $\mathbb{c} \in \mathrm{SE}(3)$ for an abstract transformation, $T = [\mathbb{c}]_{4\times4}$ for its matrix representation, and $\mathbb{c} \overset{0}{=} (\mathbb{r}, \mathbf{t})$ for its rotation-translation representation, instead of using transformations and transformation matrices interchangeably. The tuple representation $(\mathbb{r}, \mathbf{t}) \in \mathrm{SO}(3) \times \mathbb{R}^3$ is special in the sense that it allows exploiting the lifting/projection and slicing techniques detailed in Section 3.4.

Bibliography

[1] ISO 10303. industrial automation systems and integration – product data representation and exchange. Technical report, International Organization for Standardization (ISO), 1994.

[2] ISO-STEP Part 42. product data representation and exchange. Technical report, International Organization for Standardization (ISO), 1998.

[3] K. Abdel-Malek, J. Yang, and D. Blackmore. Closed-form swept volume of implicit surfaces. In *Proceedings of the 2000 ASME Design Engineering Technical Conferences (DETC'2000)*, volume 2, 2000.

[4] K. Abdel-Malek, J. Yang, and D. Blackmore. On swept volume formulations: Implicit surfaces. *Computer-Aided Design*, 33(1):113–121, 2001.

[5] K. Abdel-Malek, J. Yang, D. Blackmore, and K. Joy. Swept volumes: Foundation, perspectives, and applications. *International Journal of Shape Modeling*, 12(01):87–127, 2006.

[6] K. Abdel-Malek and H. J. Yeh. Geometric representation of the swept volume using Jacobian rank-deficiency conditions. *Computer-Aided Design*, 29(6):457–468, 1997.

[7] K. Abdel-Malek and H. J. Yeh. On the determination of starting points for parametric surface intersections. *Computer-Aided Design*, 29(1):21–35, 1997.

[8] P. K. Agarwal, E. Flato, and D. Halperin. Polygon decomposition for efficient construction of Minkowski sums. *Computational Geometry*, 21(1–2):39–61, 2002. Special Issue for the 16th European Workshop on Computational Geometry (EuroCG'2000).

[9] C. Alexander. *Notes on the Synthesis of Form*, volume 5. Harvard University Press, 1964.

[10] J. W. Alexander. A proof and extension of the Jordan-Brouwer separation theorem. *Transactions of the American Mathematical Society*, 23(4):333–349, 1922.

[11] J. W. Alexander. An example of a simply connected surface bounding a region which is not simply connected. *Proceedings of the National Academy of Sciences*, 10(1):8–10, 1924.

[12] N. Amenta, S. Choi, and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry*, 19(2):127–153, 2001.

[13] N. Amenta and R. V. Kolluri. The medial axis of a union of balls. *Computational Geometry*, 20(1):25–37, 2001.

[14] F. Amirouche. *Fundamentals of Multibody Dynamics: Theory and Applications*. Springer Science & Business Media, 2007.

[15] T. M. Apostol. Mathematical analysis. 1974.

[16] F. Arbab. Set models and Boolean operations for solids and assemblies. *IEEE Computer Graphics and Applications*, 10(6):76–86, 1990.

[17] D. Attali, J. D. Boissonnat, and H. Edelsbrunner. Stability and computation of medial axes: A state-of-the-art report. *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, pages 109–125, 2009.

[18] D. Attali and J. O. Lachaud. Combinatorial curves and surfaces delaunay conforming isosurface, skeletonextraction and noise removal. *Computational Geometry*, 19(2):175–189, 2001.

[19] Dominique Attali and Annick Montanvert. Modeling noise for a better simplification of skeletons. In *Proceedings of the 1996 International Conference on Image Processing*, volume 3, pages 13–16. IEEE, 1996.

[20] F. Aurenhammer. Voronoi diagrams–a survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)*, 23(3):345–405, 1991.

[21] F. Aurenhammer and H. Imai. Geometric relations among Voronoi diagrams. *Geometriae Dedicata*, 27(1):65–75, 1988.

[22] J. Backus. Can programming be liberated from the von Neumann style? a functional style and its algebra of programs. *Communications of the ACM*, 21(8):613–641, 1978. 1978 Turing Award Lecture.

[23] C. L. Bajaj, B. Bauer, R. Bettadapura, and A. Vollrath. Nonuniform Fourier transforms for rigid-body and multidimensional rotational correlations. *SIAM Journal on Scientific Computing*, 35(4):B821–B845, 2013.

[24] C. L. Bajaj, A. Bhowmick, E. Chattopadhyay, and D. Zuckerman. On low discrepancy samplings in product spaces of motion groups. *arXiv preprint arXiv:1411.7753*, 2014.

[25] C. L. Bajaj, R. Chowdhury, and V. Siddahanavalli. F2Dock: Fast Fourier protein-protein docking. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 8(1):45–58, 2011.

[26] C. L. Bajaj, R. A. Chowdhury, and V. Siddavanahalli. F3Dock: A fast, flexible and Fourier based approach to protein-protein docking. Technical report, University of Texas, Austin, 2008.

[27] C. L. Bajaj and V. Siddavanahalli. F2Dock: A fast and Fourier based error-bounded approach to protein-protein docking. Technical report, University of Texas, Austin, 2006.

[28] C. L. Bajaj, G. Xu, and Q. Zhang. A fast variational method for the construction of resolution adaptive c^2-smooth molecular surfaces. *Computer Methods in Applied Mechanics and Engineering*, 198(21):1684–1690, 2009.

[29] G. Barequet, C. A. Duncan, and S. Kumar. RSVP: A geometric toolkit for controlled repair of solid models. *IEEE Transactions on Visualization and Computer Graphics*, 4(2):162–177, 1998.

[30] H. Barki, F. Denis, and F. Dupont. Contributing vertices-based Minkowski sum computation of convex polyhedra. *Computer-Aided Design*, 41(7):525–538, 2009.

[31] M. Behandish. Lie groups, Lie algebras, and correlation gradients: A tutorial. Technical report, University of Connecticut, 2015.

[32] M. Behandish. *Geometric Energies for Haptic Assembly*. Master's thesis, 2016.

[33] M. Behandish and H. T. Ilieş. Shape complementarity analysis for objects of arbitrary shape. Technical Report CDL-TR-14-01, University of Connecticut, 2014.

[34] M. Behandish and H. T. Ilieş. Analytic methods for geometric modeling via spherical decomposition. In *2015 SIAM/ACM Symposium on Solid and Physical Modeling (GD/SPM'2015)*, volume 70, 2015.

[35] M. Behandish and H. T. Ilieş. Peg-in-hole revisited: A generic force model for haptic assembly. *Journal of Computing and Information Science in Engineering (JCISE)*, 15(4):041004:1–11, 2015.

[36] M. Behandish and H. T. Ilieş. Analytic methods for geometric modeling via spherical decomposition. *Computer-Aided Design*, 70:100–115, 2016. 2015 SIAM/ACM Symposium on Solid and Physical Modeling (GD/SPM'2015).

[37] M. Behandish and H. T. Ilieş. Haptic assembly using skeletal densities and Fourier transforms. *Journal of Computing and Information Science in Engineering (JCISE)*, 16(2):021002:1–11, 2016.

[38] M. Behandish and H. T. Ilieş. Comparative shape analysis using skeletal density functions. Manuscript in preparation, University of Connecticut, 2017.

[39] F. Bernardini. Integration of polynomials over $n-$dimensional polyhedra. *Computer-Aided Design*, 23(1):51–58, 1991.

[40] D. Blackmore and M. C. Leu. A differential equation approach to swept volumes. In *Proceedings of Rensselaer's Second International Conference on Computer Integrated Manufacturing*, pages 143–149, 1990.

[41] D. Blackmore, M. C. Leu, and L. P. Wang. The sweep-envelope differential equation algorithm and its application to NC machining verification. *Computer-Aided Design*, 29(9):629–637, 1997.

[42] D. Blackmore, R. Samulyak, and M. C. Leu. Trimming swept volumes. *Computer-Aided Design*, 31(3):215–223, 1999.

[43] J. F. Blinn. A generalization of algebraic surface drawing. *ACM Transactions on Graphics (TOG)*, 1(3):235–256, 1982.

[44] J. Bloomenthal and K. Shoemake. Convolution surfaces. *SIGGRAPH Computer Graphics*, 25(4):251–256, 1991. Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'1991).

[45] H. Blum. A transformation for extracting new descriptors of shape. *Models for the Perception of Speech and Visual Form*, 19(5):362–380, 1967.

[46] J. H. Bohn. Removing zero-volume parts from CAD models for layered manufacturing. *IEEE Computer Graphics and Applications*, 15(6):27–34, 1995.

[47] A. M. J. J. Bonvin. Flexible protein–protein docking. *Current Opinion in Structural Biology*, 16(2):194–200, 2006.

[48] G. Bradshaw and C. O'Sullivan. Adaptive medial-axis approximation for sphere-tree construction. *ACM Transactions on Graphics (TOG)*, 23(1):1–26, 2004.

[49] J. W. Brandt. Convergence and continuity criteria for discrete approximations of the continuous planar skeleton. *CVGIP: Image Understanding*, 59(1):116—124, 1994.

[50] J. W. Brandt and V. R. Algazi. Continuous skeleton computation by Voronoi diagram. *CVGIP: Image Understanding*, 55(3):329–338, 1992.

[51] A. M. Bronstein, M. M. Bronstein, L. J. Guibas, and M. Ovsjanikov. Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics*, 30(1):1:1–1:20, 2011.

[52] B. G. Burdea and H. J. Wolfson. Solving jigsaw puzzles by a robot. *IEEE Transactions on Robotics and Automation*, 5(6):752–764, 1989.

[53] C. Burnikel. *Exact Computation of Voronoi diagrams and line Segment Intersections*. Ph.d. dissertation, Universität des Saarlandes Saarbrücken, 1996.

[54] C. Burnikel, K. Mehlhorn, and S. Schirra. How to compute the Voronoi diagram of line segments: Theoretical and experimental results. In *Proceedings of the 2ND Annual European Symposium on Algorithms (ESA'1994)*, pages 227–239. Springer-Verlag, 1994.

[55] G. Burton. A hybrid approach to polygon offsetting using winding numbers and partial computation of the Voronoi diagram. In *The 2014 ASMEInternational Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE'2014)*, pages V02BT03A014–V02BT03A014. American Society of Mechanical Engineers (ASME), 2014.

[56] S. R. Buss. *3D Computer Graphics: A Mathematical Introduction with OpenGL*. Cambridge University Press, 1st edition edition, 2003.

[57] J. E. C. Candas, V. Siddavanahalli, and C. L. Bajaj. Nonequispaced Fourier transforms for protein-protein docking. Technical report, 2005.

[58] C.o Cattani and A. Paoluzzi. Boundary integration over linear polyhedra. *Computer-Aided Design*, 22(2):130–135, 1990.

[59] F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems*, 10(2):149–170, 2004.

[60] R. Chen, L. Li, and Z. Weng. ZDOCK: An initial-stage protein-docking algorithm. *Proteins: Structure, Function, and Bioinformatics*, 52(1):80–87, 2003.

[61] R. Chen and Z. Weng. A novel shape complementarity scoring function for protein-protein docking. *Proteins: Structure, Function, and Bioinformatics*, 51(3):397–408, 2003.

[62] X. Chen and S. McMains. Polygon offsetting by computing winding numbers. In *The 2005 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE'2005)*, pages 565–575. American Society of Mechanical Engineers (ASME), 2005.

[63] J. Cherfils and J. Janin. Protein docking algorithms: Simulating molecular recognition. *Current Opinion in Structural Biology*, 3(2):265–269, 1993.

[64] N. Chernov, Y. Stoyan, and T. Romanova. Mathematical model and efficient algorithms for object packing problem. *Computational Geometry*, 43(5):535 – 553, 2010.

[65] G. S. Chirikjian and A. B. Kyatkin. *Engineering Applications of Noncommutative Harmonic Analysis: with Emphasis on Rotation and Motion Groups*. CRC Press, 1st edition, 2010.

[66] B. K. Choi and S. Y. Ju. Constant-radius blending in surface modelling. *Computer Aided Design*, 21(4):213–220, 1989.

[67] H. In Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific Journal of Mathematics*, 181(1):57–88, 1997.

[68] S. W. Choi and H. P. Seidel. Linear one-sided stability of MAT for weakly injective 3D domain. *Computer-Aided Design*, 36(2):95–109, 2004.

[69] R. Chowdhury, M. Rasheed, D. Keidel, M. Moussalem, A. Olson, M. Sanner, and C. L. Bajaj. Protein-protein docking with F2Dock 2.0 and GB-rerank. *PloS One*, 8(3):e51307, 2013.

[70] H. S. Choy and K. W. Chan. A corner-looping based tool path for pocket milling. *Computer-Aided Design*, 35(2):155–166, 2003.

[71] A. Church. *The Calculi of Lambda-Conversion*. Princeton University Press, 6th edition edition, 1985.

[72] J. D. Cohen, M. C. Lin, D. Manocha, and M. Ponamgi. I-COLLIDE: An interactive and exact collision detection system for large-scale environments. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics (I3D'1995)*, New York, NY, USA, 1995.

[73] P G. Comba. A procedure for detecting intersections of three-dimensional objects. *The ACM Journal*, 15(3):354–366, 1968.

[74] M L. Connolly. Shape complementarity at the hemoglobin $\alpha 1 \beta 1$ subunit interface. *Biopolymers*, 25(7):1229–1247, 1986.

[75] W. R. Cook. On understanding data abstraction, revisited. *ACM SIGPLAN Notices*, 44(10):557–572, 2009.

[76] Tukey J. Cooley, J. W. An algorithm for the machine calculation of complex Fourier series. *Mathematics of CSomputation*, 19(90):297–301, 1965.

[77] R. Cools. An encyclopaedia of cubature formulas. *Journal of Complexity*, 19(3):445–453, 2003. Oberwolfach Special Issue.

[78] R. Cools and P. Rabinowitz. Monomial cubature rules since "stroud": A compilation. *Journal of Computational and Applied Mathematics*, 48(3):309–326, 1993.

[79] R. Corcodel and H. T. Ilieş. *Solid Sweeping with Multi-Parametric Motions*. Ph.d. proposal, 2016.

[80] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 3rd edition edition, 2009.

[81] S. F. Cortizo. On Dirac's delta calculus. In *arXiv Preprint Funct-An/9510004*. 1995.

[82] G. A. Crocker and W. F. Reinke. Boundary evaluation of non-convex primitives to produce parametric trimmed surfaces. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'1987)*, pages 129–136. ACM, 1987.

[83] G. A. Crocker and W. F. Reinke. Boundary evaluation of non-convex primitives to produce parametric trimmed surfaces. *SIGGRAPH Computer Graphics*, 21(4):129–136, 1987.

[84] G. W. Cui, Y. X. Yao, and Y. Wang. Assembly process simulation of toleranced products. *International Journal of Internet Manufacturing and Services*, 2(3–4):248–264, 2010.

[85] T. Culver. *Computing the Medial Axis of a Polyhedron Reliably and Efficiently*. Ph.d. dissertation, 2000.

[86] T. Culver, J. C. Keyser, and D. Manocha. Accurate computation of the medial axis of a polyhedron. In *Proceedings of the 5th ACM Symposium on Solid Modeling and Applications (SMA'1999)*, pages 179–190. ACM, 1999.

[87] T. Culver, J. C. Keyser, and D. Manocha. Exact computation of the medial axis of a polyhedron. *Computer Aided Geomtric Design*, 21(1):65–98, 2004.

[88] B. Curto, V. Moreno, and F. J. Blanco. A general method for C-space evaluation and its application to articulated robots. *IEEE Transactions on Robotics and Automation*, 18(1):24–31, 2002.

[89] O. de Berg, M. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer, 3rd edition edition, 2000.

[90] P. J. Denning. Ubiquity symposium: 'what is computation?'. *Ubiquity*, 2010, 2010.

[91] S. L. Devadoss and J. O'Rourke. *Discrete and Computational Geometry*. Princeton University Press, 2011.

[92] T. K. Dey, K. Li, C. Luo, P. Ranjan, I. Safa, and Y. Wang. Persistent heat signature for pose-oblivious matching of incomplete models. In *Computer Graphics Forum*, volume 29, pages 1545–1554. Wiley Online Library, 2010.

[93] T. K. Dey and W. Zhao. Approximate medial axis as a Voronoi subcomplex. In *Proceedings of the 7th ACM Symposium on Solid Modeling and Applications (SMA'2002)*, pages 356–366. ACM, 2002.

[94] T. K. Dey and W. Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38(1):179–200, 2003.

[95] T. K. Dey and W. Zhao. Approximate medial axis as a Voronoi subcomplex. *Computer-Aided Design*, 36(2):195–202, 2004.

[96] G. S. di Baja and E. Thiel. A multiresolution shape description algorithm. In D. Chetverikov and W. G. Kropatsch, editors, *Proceedings of the 5th International Conference on Computer Analysis of Images and Patterns (CAIP'1993)*, pages 208–215. Springer Berlin Heidelberg, 1993.

[97] P. Diaconis and M. Shahshahani. The subgroup algorithm for generating uniform random variables. *Probability in the Engineering and Informational Sciences*, 1(01):15–32, 1987.

[98] A. R. Dill, M. D. Levine, and P. B. Noble. Multiple resolution skeletons. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI-9)*, PAMI-9(4):495–504, 1987.

[99] P. A. M. Dirac. *The Principles of Quantum Mechanics*. Number 27 in The International Series of Monographs on Physics. Oxford University Press, 4th edition, 1981.

[100] D. Dobkin, L. Guibas, J. Hershberger, and J. Snoeyink. An efficient algorithm for finding the CSG representation of a simple polygon. *Algorithmica*, 10(1):1–23, 1993.

[101] D. Dragomatz and S. Mann. A classified bibliography of literature on NC milling path generation. *Computer-Aided Design*, 29(3):239–247, 1997.

[102] D. Duhovny, R. Nussinov, and H. J. Wolfson. Efficient unbound docking of rigid molecules. In *Algorithms in Bioinformatics*, pages 185–200. Springer Berlin Heidelberg, 2002.

[103] B. S. Duncan and A. J. Olson. Approximation and characterization of molecular surfaces. *Biopolymers*, 33(2):219–229, 1993.

[104] A. Edalat, A. A. Khanban, and A. Lieutier. Computability in computational geometry. In *Conference on Computability in Europe*, pages 117–127. Springer, 2005.

[105] A. Edalat and A. Lieutier. Foundation of a computable solid modeling. In *Proceedings of the 5th ACM Symposium on Solid Modeling and Applications (SMA'1999)*, pages 278–284. Association for Computing Machinery (ACM), 1999.

[106] A. Edalat and A. Lieutier. Foundation of a computable solid modelling. *Theoretical Computer Science*, 284(2):319–345, 2002.

[107] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. American Mathematical Socoety, 2010.

[108] S. A. Ehmann and M. C. Lin. Accelerated proximity queries between convex polyhedra by multi-level Voronoi marching. In *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2000)*, volume 3, 2000.

[109] S. A. Ehmann and M. C. Lin. Accurate and fast proximity queries between polyhedra using convex surface decomposition. *Computer Graphics Forum*, 20(3):500–511, 2001.

[110] M. Eisenstein and E. Katchalski-Katzir. On proteins, grids, correlations, and docking. *Comptes Rendus Biologies*, 327(5):409–420, 2004.

[111] H. Erdim and H. T. Ilieş. Detecting and quantifying envelope singularities in the plane. *Computer-Aided Design*, 39(10):829–840, 2007.

[112] H. Erdim and H. T. Ilieş. A comparison of sampling strategies for computing general sweeps. *Computer-Aided Design*, 42(8):657–669, 2010.

[113] H. Erdim and Horea T. Ilieş. Classifying points for sweeping solids. *Computer-Aided Design*, 40(9):987–998, 2008.

[114] T. J. A. Ewing, S. Makino, A. G. Skillman, and I. D. Kuntz. DOCK 4.0: Search strategies for automated molecular docking of flexible molecule databases. *Journal of Computer-Aided Molecular Design*, 15(5):411–428, 2001.

[115] G. E. Farin, J. Hoschek, and M. S. Kim. *Handbook of Computer Aided Geometric Design*. North Holland, 1st edition edition, 2002.

[116] J. C. E. Ferreira and D. M. Ochoa. A method for generating trochoidal tool paths for $2\frac{1}{2}$D pocket milling process planning with multiple tools. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, page 0954405413487897, 2013.

[117] D. Fischer, S. L. Lin, H. L. Wolfson, and R. Nussinov. A geometry-based suite of molecular docking processes. *Journal of Molecular Biology*, 248(2):459–477, 1995.

[118] D. Fischer, R. Norel, R. Nussinov, and H. J. Wolfson. 3-D docking of protein molecules. In *Combinatorial Pattern Matching*, volume 684, pages 20–34. Springer Berlin Heidelberg, 1993.

[119] E. Fogel and D. Halperin. Exact and efficient construction of Minkowski sums of convex polyhedra with applications. *Computer-Aided Design*, 39(11):929–940, 2007.

[120] S. Fortune. Polyhedral modelling with exact arithmetic. In *Proceedings of the Third ACM Symposium on Solid Modeling and Applications (SMA'1995)*, pages 225–234. ACM, 1995.

[121] M. Foskey, M. C. Lin, and D. Manocha. Efficient computation of a simplified medial axis. *Journal of Computing and Information Science in Engineering*, 3(4):274–284, 2003.

[122] M. Foskey, M. C. Lin, and D. Manocha. Efficient computation of a simplified medial axis. In *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications (SMA'2003)*, pages 96–107. ACM, 2003.

[123] M. Foskey, D. Manocha, T. Culver, J. C. Keyser, and S. Krishnan. Reliable geometric computations with algebraic primitives and predicates. In *Uncertainty in Geometric Computations*, pages 91–106. Springer, 2002.

[124] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE, Special issue on Program Generation, Optimization, and Platform Adaptation*, 93(2):216–231, 2005.

[125] W. Fu, S. Nelaturi, A. Rangarajan, and T. Kurtoglu. Tolerance analysis for validating manufacturing process plans. In *The 2014 ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE'2014)*, pages V01AT02A055–V01AT02A055. American Society of Mechanical Engineers (ASME), 2014.

[126] H. A. Gabb, R. M. Jackson, M. J. E. Sternberg, et al. Modelling protein docking using shape complementarity, electrostatics and biochemical information. *Journal of Molecular Biology*, 272(1):106–120, 1997.

[127] R. R. Gabdoulline and R. C. Wade. Analytically defined surfaces to analyze molecular interaction properties. *Journal of Molecular Graphics*, 14(6):341–353, 1996.

[128] D. Geiger, T. L. Liu, and R. V. Kohn. Representation and self-similarity of shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(1):86–99, 2003.

[129] H. Ghaderi. *The Rare Decay of the Neutral Pion into a Dielectron.* Master's thesis, Uppsala University, 2013.

[130] H. Ghaderi. Triangle loop in scalar decay and cutting rules. Technical report, Uppsala University, 2013.

[131] W. B. Goh. Strategies for shape matching using skeletons. *Computer Vision and Image Understanding*, 110(3):326–345, 2008. Similarity Matching in Computer Vision and Multimedia.

[132] D. Goldberg, C. Malon, and M. Bern. A global approach to automatic solution of jigsaw puzzles. In *Proceedings of the 18th Annual Symposium on Computational Geometry*, pages 82–87. ACM, 2002.

[133] S. Gottschalk, M. C. Lin, and D. Manocha. OBBTree: A hierarchical structure for rapid interference detection. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, pages 171–180, New York, NY, USA, 1996.

[134] M. Gräf and D. Potts. Sampling sets and quadrature formulae on the rotation group. *Numerical Functional Analysis and Optimization*, 30(7-8):665–688, 2009.

[135] A. Grim, T. OConnor, P. J. Olver, C. Shakiban, R. Slechta, and R. Thompson. Automatic reassembly of three-dimensional jigsaw puzzles. *International Journal of Image and Graphics*, 16(02):1650009, 2016.

[136] W. Guoping, S. Jiaguang, and H. Xuanji. The sweep-envelope differential equation algorithm for general deformed swept volumes. *Computer Aided Geometric Design*, 17(5):399–418, 2000.

[137] E. L. Gursoz, Y. Choi, and F. B. Prinz. Vertex-based representation of non-manifold boundaries. *Geometric Modeling for Product Engineering*, 23(1):107–130, 1990.

[138] B. C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, volume 222. Springer, 2015.

[139] R. M. Haralick, S. R. Sternberg, and X. Zhuang. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(4):532–550, 1987.

[140] R. S. Hartenberg and J. Denavit. *Kinematic Synthesis of Linkages*. McGraw-Hill, 1964.

[141] A. Hatna, R. J. Grieve, and P. Broomhead. Automatic CNC milling of pockets: Geometric and technological issues. *Computer Integrated Manufacturing Systems*, 11(4):309–330, 1998.

[142] M. Held. Computational geometry and computer-aided design and manufacturing voronoi diagrams and offset curves of curvilinear polygons. *Computer-Aided Design*, 30(4):287–300, 1998.

[143] M. Held. Voronoi diagrams and offset curves of curvilinear polygons. *Computer-Aided Design*, 30(4):287–300, 1998.

[144] M. Held, G. Lukács, and L. Andor. NC machining and cutter-path generation pocket machining based on contour-parallel tool paths generated by means of proximity maps. *Computer-Aided Design*, 26(3):189–203, 1994.

[145] J. C. Himmelstein, E. Ferre, and J. P. Laumond. Swept volume approximation of polygon soups. *IEEE Transactions on Automation Science and Engineering*, 7(1):177–183, 2010.

[146] D. J. Hoff and P. J. Olver. Automatic solution of jigsaw puzzles. *Journal of Mathematical Imaging and Vision*, 49(1):234–250, 2014.

[147] C. M. Hoffmann, V. Shapiro, and V. Srinivasan. Geometric interoperability via queries. *Computer-Aided Design*, 46:148–159, 2014.

[148] C. M. Hoffmann, Vadim V. Shapiro, and V. Srinivasan. Geometric interoperability for resilient manufacturing. Technical Report 11-015, 2011.

[149] L. Hong, X. C. Zhang, J. A. Hartsuck, and J. Tang. Crystal structure of an *in vivo* HIV-1 protease mutant in complex with saquinavir: Insights into the mechanisms of drug resistance. *Protein Science*, 9(10):1898–1904, 2000.

[150] L. Hörmander. *The Analysis of Linear Partial Differential Operators I: Distribution Theory and Fourier Analysis*, volume 274. Springer Science & Business Media, 2003.

[151] E. Horowitz, S. Sahni, and S. Rajasekaran. *Computer Algorithms*. Computer Science Press, 2nd edition edition, 2007.

[152] T. Hou and H. Qin. Robust dense registration of partial nonrigid shapes. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1268–1280, 2012.

[153] K. Hrbacek and T. Jech. *Introduction to Set Theory*. Chapman & Hall/CRC Pure and Applied Mathematics. CRC Press, 1999.

[154] P. M. Hubbard. Interactive collision detection. In *Proceedings of the 1993 IEEE Symposium on Research Frontiers in Virtual Reality*, pages 24–31, 1993.

[155] P. M. Hubbard. Collision detection for interactive graphics applications. *IEEE Transactions on Visualization and Computer Graphics*, 1(3):218–230, 1995.

[156] P. M. Hubbard. Approximating polyhedra with spheres for time-critical collision detection. *ACM Transac*, 15(3):179–210, 1996.

[157] T. C. Hudson, M. C. Lin, J. Cohen, S. Gottschalk, and D. Manocha. V-COLLIDE: Accelerated collision detection for VRML. In *Proceedings of the 2nd Symposium on Virtual Reality Modeling Language*, New York, NY, USA, 1997.

[158] J. F. Hughes, A. van Dam, M. McGuire, J. D. Foley, S. K. Feiner, and K. Akeley. *Computer Graphics: Principles and Practice*. Addison-Wesley Professional, 3rd edition edition, 2013.

[159] H. T. Ilieş. *On Shaping Moving Mehcanical Parts*. Ph.d. dissertation, 2000.

[160] H. T. Ilieş and V. Shapiro. *An Approach to Systematic Part Design*, volume 2, pages 383–392. 1996.

[161] H. T. Ilieş and V. Shapiro. UNSWEEP: Formulation and computational properties. In *Proceedings of the 4th ACM Symposium on Solid Modeling and Applications (SMA'1997)*, pages 155–167. Association for Computing Machinery (ACM), 1997.

[162] H. T. Ilieş and V. Shapiro. The dual of sweep. *Computer-Aided Design*, 31(3):185–201, 1999.

[163] H. T. Ilieş and V. Shapiro. On shaping with motion. *Journal of Mechanical Design*, 122(4):567–574, 2000.

[164] H. T. Ilieş and V. Shapiro. A class of forms from function: The case of parts moving in contact. *Research in Engineering Design*, 13(3):157–166, 2002.

[165] J. I. Jackson, C. H. Meyer, D. G. Nishimura, and A. Macovski. Selection of a convolution function for Fourier inversion using gridding. *IEEE Transactions on Medical Imaging*, 10(3):473–478, 1991.

[166] J. Janin and C. Chothia. The structure of protein-protein recognition sites. *Journal of Biological Chemistry*, 265, 1990.

[167] J. Jeong and K. Kim. Generation of tool paths for machining free-form pockets with islands using distance maps. *The International Journal of Advanced Manufacturing Technology*, 15(5):311–316, 1999.

[168] F. Jiang and S. H. Kim. "soft docking": Matching of molecular surface cubes. *Journal of Molecular Biology*, 219(1):79–102, 1991.

[169] P. Jimenez, F. Thomas, and Torras. 3D collision detection: A survey. *Computers and Graphics*, 25(2):269–285, 2001.

[170] L. Joskowicz and E. Sacks. Computer-aided mechanical design using configuration spaces. *Computing in Science and Engineering*, 1(6):14–21, 1999.

[171] M. I. Karavelas. A robust and efficient implementation for the segment Voronoi diagram. In *International Symposium on Voronoi Diagrams in Science and Engineering*, pages 51–62. Citeseer, 2004.

[172] E. Katchalski-Katzir, I. Shariv, M. Eisenstein, A. A Friesem, C. Aflalo, and I. A. Vakser. Molecular surface recognition: Determination of geometric fit between proteins and their ligands by correlation techniques. *Proceedings of the National Academy of Sciences*, 89(6):2195–2199, 1992.

[173] Y. Katznelson. *An Introduction to Harmonic Analysis*. Cambridge University Press, 3rd edition, 2004.

[174] L. E. Kavraki. Computation of configuration-space obstacles using the fast Fourier transform. *IEEE Transactions on Robotics and Automation*, 11(3):408–413, 1995.

[175] A. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer Science & Business Media, 1995.

[176] J. Keiner, S. Kunis, and D. Potts. Using NFFT 3—a software library for various nonequispaced fast Fourier transforms. *ACM Transaction Mathematical Software (TMS)*, 36(4):19:1–19:30, 2009.

[177] J. C. Kendrew, G. B., H. M. Dintzis, R. G. Parrish, H. Wyckoff, and D. C. Phillips. A three-dimensional model of the Myoglobin molecule obtained by X-ray analysis. *Nature*, 181(4610):662–666, 1958.

[178] J. C. Keyser. *Exact Boundary Evaluation for Curved Solids*. Ph.d. dissertation, University of North Carolina at Chapel Hill, 2000.

[179] J. C. Keyser, T. Culver, M. Foskey, S. Krishnan, and D. Manocha. ESOLID–a system for exact boundary evaluation. *Computer-Aided Design*, 36(2):175–193, 2004.

[180] J. C. Keyser, T. Culver, D. Manocha, and S. Krishnan. MAPC: A library for efficient and exact manipulation of algebraic points and curves. In *Proceedings of the 15th annual Symposium on Computational Geometry (SCM'1999)*, pages 360–369. ACM, 1999.

[181] J. C. Keyser, T. Culver, D. Manocha, and S. Krishnan. Efficient and exact manipulation of algebraic points and curves. *Computer-Aided Design*, 32(11):6490–662, 2000.

[182] J. C. Keyser, S. Krishnan, and D. Manocha. Efficient and accurate B-rep generation of low degree sculptured solids using exact arithmetic. In *Proceedings of the Fourth ACM Symposium on Solid Modeling and Applications (SMA'1997)*, pages 42–55. ACM, 1997.

[183] J. C. Keyser, S. Krishnan, and D. Manocha. Efficient and accurate B-rep generation of low degree sculptured solids using exact arithmetic: I–representations. *Computer Aided Geometric Design*, 16(9):841–859, 1999.

[184] J. C. Keyser, S. Krishnan, and D. Manocha. Efficient and accurate B-rep generation of low degree sculptured solids using exact arithmetic: II–computation. *Computer Aided Geometric Design*, 16(9):861–882, 1999.

[185] C. E. Kim and J. M. Vance. Collision detection and part interaction modeling to facilitate immersive virtual assembly methods. *Journal of Computing and Information Science in Engineering*, 4(2):83–90, 2004.

[186] D. S. Kim. Polygon offsetting using a Voronoi diagram and two stacks. *Computer-Aided Design*, 30(14):1069–1076, 1998.

[187] Y. J. Kim, G. Varadhan, M. C. Lin, and D. Manocha. Fast swept volume approximation of complex polyhedral models. *Computer-Aided Design*, 36(11):1013–1027, 2004. Solid Modeling Theory and Applications.

[188] S. Kockara, T. Halic, K. Iqbal, C. Bayrak, and R. Rowe. Collision detection: A survey. In *Proceedings of the 2007 IEEE International Conference on Systems, Man and Cybernetics (ISIC'2007)*, pages 4046–4051, 2007.

[189] D.E. Koshland. The key–lock theory and the induced fit theory. *Angewandte Chemie International (English Edition)*, 33(23–24):2375–2378, 1995.

[190] J. A. Kovacs, P. Chacon, Y. Cong, E. Metwally, and W. Wriggers. Fast rotational matching of rigid bodies by fast Fouriertransform acceleration of five degrees of freedom. *Acta Crystallographica Section D: Biological Crystallography*, 59(8):1371–1376, 2003.

[191] D. Kozakov, R. Brenke, S. R. Comeau, and S. Vajda. PIPER: An FFT-based protein docking program with pairwise potentials. *Proteins: Structure, Function, and Bioinformatics*, 65(2):392–406, 2006.

[192] S. Krishnan, D. Manocha, M. Gopi, T. Culver, and J. C. Keyser. BOOLE: A boundary evaluation system for boolean combinations of sculptured solids. *International Journal of Computational Geometry & Applications*, 11(01):105–144, 2001.

[193] S. Kunis and S. Kunis. The nonequispaced FFT on graphics processing units. *Proceedings in Applied Mathematics and Mechanics (PAMM)*, 12(1):7–10, 2012.

[194] I. D. Kuntz, J. M. Blaney, S. J. Oatley, and T. E. Langridge, R.and Ferrin. A geometric approach to macromolecule-ligand interactions. *Journal of Molecular Biology*, 161(2):269–288, 1982.

[195] W. M. Lai, D. H. Rubin, D. Rubin, and E. Krempl. *Introduction to Continuum Mechanics*. Butterworth-Heinemann, 2009.

[196] Y. Lamdan, J. T. Schwartz, and H. J. Wolfson. Object recognition by affine invariant matching. In *Proceedings of the Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'1988)*, pages 335–344, 1988.

[197] Y. Lamdan, J. T. Schwartz, and H. J. Wolfson. Affine invariant model-based object recognition. *IEEE Transactions on Robotics and Automation*, 6(5):578–589, 1990.

[198] Y. Lamdan and H. J. Wolfson. Geometric hashing: A general and efficient model-based recognition scheme. In *2nd IEEE International Conference on Computer Vision*, pages 238–249, 1988.

[199] J. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.

[200] B. Lee and F. M. Richards. The interpretation of protein structures: Estimation of static accessibility. *Journal of Molecular Biology*, 55(3):379–400, 1971.

[201] I. K. Lee, M. S. Kim, and G. Elber. Polynomial/rational approximation of Minkowski sum boundary curves. *Graphical Models and Image Processing*, 60(2):136–165, 1998.

[202] Y. T. Lee and A. A. G. Requicha. Algorithms for computing the volume and other integral properties of solids. i. known methods and open issues. *Communications of the ACM*, 25(9):635–641, 1982.

[203] Y. T. Lee and A. A. G. Requicha. Algorithms for computing the volume and other integral properties of solids. ii. a family of algorithms based on representation conversion and cellular approximation. *Communications of the ACM*, 25(9):642–650, 1982.

[204] H. P. Lenhof. New contact measures for the protein docking problem. In *Proceedings of the 1st Annual International Conference on Computational Molecular Biology*, pages 182–191, 1997.

[205] J. M. Lien. Covering Minkowski sum boundary using points with applications. *Computer Aided Geometric Design*, 25(8):652–666, 2008. Selected papers from the 15th Pacific Conference on Computer Graphics and Applications.

[206] S. L. Lien and J. T. Kajiya. A symbolic method for calculating the integral properties of arbitrary nonconvex polyhedra. *IEEE Computer Graphics and Applications*, 4(10):35–42, 1984.

[207] A. Lieutier. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.

[208] M. Lin and S. Gottschalk. Collision detection between geometric models: A survey. In *Proceedings of the 1998 IMA Conference on Mathematics of Surfaces*, volume 1, pages 37–56, 1998.

[209] S. Lojasiewicz. Triangulation of semi-analytic sets. *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze*, 18(4):449–474, 1964.

[210] T. Lozano-Perez. Spatial planning: A configuration space approach. *IEEE Transactions on Computers*, C-32(2):108–120, 1983.

[211] B. Luft, V. Shapiro, and I. Tsukanov. Geometrically adaptive numerical integration. In *Proceedings of the 2008 ACM Symposium on Solid and Physical Modeling (SPM'2008)*, pages 147–157. ACM, 2008.

[212] E. D. Lutz. *Numerical Methods for Hypersingular and Near-Singular Boundary Integrals in Fracture Mechanics*. Ph.d. dissertation, 1991.

[213] M. Lysenko. Fourier collision detection. Technical report, University of Wisconsin-Madison, 2011.

[214] M. Lysenko. Fourier collision detection. *International Journal of Robotics Research*, 32(4):483–503, 2013.

[215] M. Lysenko, S. Nelaturi, and V. Shapiro. Group morphology with convolution algebras. In *Proceedings of the 2010 ACM Symposium on Solid and Physical Modeling (SPM'2010)*, pages 11–22, New York, NY, USA, 2010.

[216] M. Lysenko and V. Shapiro. Effective contact measures. In *2015 SIAM/ACM Symposium on Solid and Physical Modeling (GD/SPM'2015)*, volume 70, pages 134–143, 2015.

[217] M. Lysenko and V. Shapiro. Effective contact measures. *Computer-Aided Design*, 70:134–143, 2016. 2015 SIAM/ACM Symposium on Solid and Physical Modeling (GD/SPM'2015).

[218] M. Lysenko, V. Shapiro, and Nelaturi. Non-commutative morphology: Shapes, filters, and convolutions. *Computer Aided Geometric Design*, 28(8):497–522, 2011.

[219] R. Martin. Modelling inexact shapes with fuzzy sets. In *Proceedings of the Computer Science and Graphics Conference (CSG'1994)*, pages 1–26, 1994.

[220] N. L. Max and E. D. Getzoff. Spherical harmonic molecular surfaces. *Computer Graphics and Applications*, 8(4):42–50, 1988.

[221] B. Mazur. On embeddings of spheres. *Bulletin of the American Mathematical Society*, 65(2):59–65, 1959.

[222] W. A. McNeely, K. D. Puterbaugh, and J. J. Troy. Advances in voxel-based 6-DOF haptic rendering. In *Proceedings of the ACM SIGGRAPH'2005 Courses*, 2005.

[223] W. A. McNeely, K. D. Puterbaugh, and J. J. Troy. Six degree-of-freedom haptic rendering using voxel sampling. In *Proceedings of the ACM SIGGRAPH'2005 Courses*, New York, NY, USA, 2005.

[224] W. A. McNeely, K. D. Puterbaugh, and J. J. Troy. Voxel-based 6-DOF haptic rendering improvements. *Haptics-e*, 3(7), 2006.

[225] M. D. Miller, S. K. Kearsley, D. J. Underwood, and R. P. Sheridan. FLOG:a system to select 'quasi-flexible' ligands complementary to a receptor of known three-dimensional structure. *Journal of Computer-Aided Molecular Design*, 8(2):153–174, 1994.

[226] B. Mirtich. Fast and accurate computation of polyhedral mass properties. *journal of Graphics Tools*, 1(2):31–50, 1996.

[227] B. Mirtich. V-Clip: Fast and robust polyhedral collision detection. *ACM Transactions on Graphics (TOG)*, 17(3):177–208, 1998.

[228] K. Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice-Hall Englewood Cliffs, 1994.

[229] J. R. Munkres. *Topology*. Prentice Hall, 2nd edition edition, 2000.

[230] T. M. Murali and T. A. Funkhouser. Consistent solid and boundary representations from arbitrary polygonal data. In *Proceedings of the 1997 Symposium on Interactive 3D Graphics (I3D'1997)*, pages 155–ff. ACM, 1997.

[231] S. Nelaturi. *Configuration Modeling*. Ph.D. dissertation, University of Wisconsin-Madison, 2012.

[232] S. Nelaturi, G. Burton, C. Fritz, and T. Kurtoglu. Automatic spatial planning for machining operations. In *2015 IEEE International Conference on Automation Science and Engineering (CASE'2015)*, pages 677–682, 2015.

[233] S. Nelaturi, W. Kim, and T. Kurtoglu. Manufacturability feedback and model correction for additive manufacturing. *Journal of Manufacturing Science and Engineering*, 137(2):021015, 2015.

[234] S. Nelaturi, M. Lysenko, and V. Shapiro. Rapid mapping and exploration of configuration space. *Journal of Computing and Information Science in Engineering (JCISE)*, 12(2):021007, 2012.

[235] S. Nelaturi and V. Shapiro. Configuration products and quotients in geometric modeling. *Computer-Aided Design*, 43(7):781–794, 2011. The 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling.

[236] S. Nelaturi and V. Shapiro. Solving inverse configuration space problems by adaptive sampling. *Computer-Aided Design*, 45(2):373–382, 2013.

[237] S. Nelaturi and V. Shapiro. Representation and analysis of additively manufactured parts. *Computer-Aided Design*, 67-68:13–23, 2015.

[238] T. R. Nielsen, P. Drewsen, and K. Hansen. Solving jigsaw puzzles using image features. *Pattern Recognition Letters*, 29(14):1924–1933, 2008.

[239] R. Norel, D. Fischer, H. J. Wolfson, and R. Nussinov. Molecular surface recognition by a computer vision-based technique. *Protein engineering*, 7(1):39–46, 1994.

[240] R. Norel, S. L. Lin, H. J. Wolfson, and R. Nussinov. Shape complementarity at protein–protein interfaces. *Biopolymers*, 34(7):933–940, 1994.

[241] R. Norel, S. L. Lin, H. J. Wolfson, R. Nussinov, et al. Molecular surface complementarity at protein-protein interfaces: the critical role played by surface normals at well placed, sparse, points in docking. *Journal of Molecular Biology*, 252(2):263–273, 1995.

[242] M. O'Connor, V. Srinivasan, and A. Jones. Connected Lie and symmetry subgroups of the rigid motions: Foundations and classification. Research Report 20512, IBM, 1996.

[243] Staff of the Production Automation Project. An introduction to PADL. Production Automation Project, Technical Memo. No. 22 (TM-22), University of Rochester, 1974.

[244] Defense Sciences Office. Complex adaptive system composition and design environment (CASCADE). Broad Agency Announcement (BAA) Solicitation No. DARPA-BAA-16-11, Defense Advanced Research Projects Agency (DARPA), 2015.

[245] Defense Sciences Office. Design for advanced materials and manufacturing. Request for Information (RFI) Solicitation No. DARPA-SN-16-03, Defense Advanced Research Projects Agency (DARPA), 2015.

[246] Defense Sciences Office. TRAnsformative DESign (TRADES). Broad Agency Announcement (BAA) Solicitation No. DARPA-BAA-16-39, Defense Advanced Research Projects Agency (DARPA), 2016.

[247] R. L. Ogniewicz. A multiscale MAT from Voronoi diagrams: The skeleton-space and its application to shape description and decomposition. *Aspects of Visual Form Processing*, pages 430–439, 1994.

[248] C. O'Sullivan and J. Dingliana. Real-time collision detection and response using sphere-trees. In *Spring Conference on Computer Graphics*, pages 83–92, 1999.

[249] K. Ouchi and J. C. Keyser. Handling degeneracies in exact boundary evaluation. In *Proceedings of the 9th ACM Symposium on Solid Modeling and Applications (SMA'2004)*, pages 321–326. Eurographics Association, 2004.

[250] M. Ovsjanikov, Q. Mérigot, F. Mémoli, and L. Guibas. One point isometric matching with the heat kernel. 29(5):1555–1564, 2010.

[251] R. S. Palmer and Vadim Shapiro. Chain models of physical behavior for engineering analysis and design. *Research in Engineering Design*, 5(3):161–184.

[252] A. Pasko, V. Adzhiev, A. Sourin, and V. Savchenko. Function representation in geometric modeling: concepts, implementation and applications. *The Visual Computer*, 11(8):429–446, 1995.

[253] T. M. K. Pasupathy, E. P. Morse, and R. G. Wilhelm. A survey of mathematical methods for the construction of geometric tolerance zones. *Journal of Computing and Information Science in Engineering*, 3:64–75, 2003.

[254] M. Peternell, H. Pottmann, T. Steiner, and T. Zhao. Swept volumes. *Computer-Aided Design and Applications*, 2(5):599–608, 2005.

[255] C. A. Petri. *Kommunikation mit Automaten*. Ph.d. dissertation, 1962.

[256] G. A. Petsko and D. Ringe. *Protein Structure and Function*. New Science Press, 2004.

[257] B. Pham. Offset curves and surfaces: A brief survey. *Computer-Aided Design*, 24(4):223–229, 1992.

[258] D. Potts, G. Steidl, and M. Tasche. Fast Fourier transforms for nonequispaced data: A tutorial. In *Modern Sampling Theory*, pages 247–270. Springer-Verlag, 2001.

[259] F. P. Preparata and M. Shamos. *Computational Geometry: An Introduction*. Monographs in Computer Science. Springer Science & Business Media, 1st edition edition, 2012.

[260] J. Qi and V. Shapiro. Epsilon-regular sets and intervals. In *2005 International Conference on Shape Modeling and Applications*, pages 308–317. Elsevier, 2005.

[261] J. Qi and V. Shapiro. $\epsilon-$topological formulation of tolerant solid modeling. *Computer-Aided Design*, 38(4):367–377, 2006. ACM Symposium on Solid and Physical Modeling 2005 (SPM'2005).

[262] J. Qi and V. Shapiro. Geometric interoperability with epsilon solidity. *Journal of Computing and Information Science in Engineering (JCISE)*, 6(3):213–220, 2006.

[263] B. Qian, S. Raman, R. Das, P. Bradley, A. J. McCoy, R. J. Read, and D. Baker. High-resolution structure prediction and the crystallographic phase problem. *Nature*, 450(7167):259–264, 2007.

[264] S. Rajasekaran and S. Ramaswami. Optimal parallel randomized algorithms for the Voronoi diagram of line segments in the plane and related problems. In *Proceedings of the 10th Annual Symposium on Computational Geometry (SCG'1994)*, pages 57–66. ACM, 1994.

[265] S. Rajasekaran and S. Ramaswami. Optimal mesh algorithms for the Voronoi diagram of line segments and motion planning in the plane. *Journal of Parallel and Distributed Computing*, 26(1):99–115, 1995.

[266] S. Rajasekaran and S. Ramaswami. Optimal parallel randomized algorithms for the Voronoi diagram of line segments in the plane. *Algorithmica*, 33:436–460, 2002. 10.1007/s00453-001-0115-6.

[267] S. Rajasekaran and S. Sen. *Synthesis of Parallel Algorithms*, chapter Random Sampling Techniques and Parallel Algorithms Design, pages 411–451. Morgan-Kaufmann Publishers, San Mateo, California, 1993.

[268] R. Ramamurthy and R. T. Farouki. Computational methods in computer graphics Voronoi diagram and medial axis algorithm for planar domains with curved boundaries I. theoretical foundations. *Journal of Computational and Applied Mathematics*, 102(1):119–141, 1999.

[269] R. Ramamurthy and R. T. Farouki. Voronoi diagram and medial axis algorithm for planar domains with curved boundaries II: Detailed algorithm description. *Journal of Computational and Applied Mathematics*, 102(2):253 – 277, 1999.

[270] A. Rappoport. An architecture for universal CAD data exchange. In *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications (SMA'2003)*, pages 266–269. ACM, 2003.

[271] D. Raviv, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Volumetric heat kernel signatures. In *Proceedings of the ACM Workshop on 3D Object Retrieval (3DOR'2010)*, pages 39–44. Association for Computing Machinery (ACM), 2010.

[272] J. H. Reif and S. Sen. Polling: a new randomized sampling technique for computational geometry. In *Proceedings of the 21st annual ACM symposium on Theory of Computing (STOC'1989)*, pages 394–404. ACM, 1989.

[273] J. H. Reif and S. Sen. Optimal parallel randomized algorithms for three-dimensional convex hulls and related problems. *SIAM Journal on Computing*, 21(3):466–485, 1992.

[274] J. H. Reif and S. Sen. Optimal randomized parallel algorithms for computational geometry. *Algorithmica*, 7:91–117, 1992.

[275] A. A. G. Requicha. Mathematical models of rigid solid objects. Production Automation Project, Technical Memo. No. 28 (TM-28), University of Rochester, 1977.

[276] A. A. G. Requicha. Mathematical foundations of constructive solid geometry: General topology of closed regular sets. Production Automation Project, Technical Memo. No. 27 (TM-27), University of Rochester, 1978.

[277] A. A. G. Requicha. Representations for rigid solids: Theory, methods, and systems. *Journal of ACM Computational Surveys*, 12(4):437–464, 1980.

[278] A. A. G. Requicha. Representations of rigid solid objects. Production Automation Project, Technical Memo. No. 29 (TM-29), University of Rochester, 1980.

[279] A. A G. Requicha. Toward a theory of geometric tolerancing. *The International Journal of Robotics Research*, 2(4):45–60, 1983.

[280] A. A. G. Requicha and J. R. Rossignac. Solid modeling and beyond. *IEEE Computer Society*, 12(05):31–44, 1992.

[281] A. A. G. Requicha, N. M. Samuel, and H. B. Voelcker. Part and assembly description languages–II. Production Automation Project, Technical Memo. No. 20a (TM-20a), University of Rochester, 1974.

[282] A. A. G. Requicha and H. B. Voelcker. Constructive solid geometry. Production Automation Project, Technical Memo. No. 25 (TM-25), University of Rochester, 1977.

[283] A. A. G. Requicha and H. B. Voelcker. *Advances in Information Systems Science: Volume 8*, chapter An Introduction to Geometric Modeling and Its Applications in Mechanical Design and Production, pages 293–328. Springer, 1981.

[284] A. A. G. Requicha and H. B. Voelcker. Solid modeling: A historical summary and contemporary assessment. *IEEE Computer Graphics and Applications*, 2(2):9–24, 1982.

[285] A. A. G. Requicha and H. B. Voelcker. Boolean operations in solid modeling: Boundary evaluation and merging algorithms. *Proceedings of the IEEE*, 73(1):30–44, 1985.

[286] A. Ricci. A constructive geometry for computer graphics. *The Computer Journal*, 16(2):157–160, 1973.

[287] T. J. Richmond. Solvent accessible surface area and excluded volume in proteins: Analytical equations for overlapping spheres and implications for the hydrophobic effect. *Journal of Molecular Biology*, 178(1):63–89, 1984.

[288] D. Ritchie. Recent progress and future directions in protein-protein docking. *Current Protein and Peptide Science*, 9(1):1–15, 2008.

[289] D. W. Ritchie and G. J. L. Kemp. Protein docking using spherical polar Fourier correlations. *Proteins: Structure, Function, and Bioinformatics*, 39(2):178–194, 2000.

[290] D. W. Ritchie, D. Kozakov, and S. Vajda. Accelerating and focusing protein-protein docking correlations using multi-dimensional rotational FFT generating functions. *Bioinformatics*, 24(17):1865–1873, 2008.

[291] J. B. T. M. Roerdink. Group morphology. *Pattern Recognition*, 33(6):877–895, 2000.

[292] J. B. T. M. Roerdink. Adaptivity and group invariance in mathematical morphology. In *Proceedings of the 16th IEEE International Conference on Image Processing (ICIP)*, pages 2253–2256, 2009.

[293] N. Rosenauer and A. H. Willis. *Kinematics of Mechanisms*. Dover Publications, 1953.

[294] J. Rossignac. *Blending and Offsetting Solid Models*. Ph.d. dissertation, 1984.

[295] J. Rossignac. Blending and offsetting solid models. Production Automation Project, Technical Memo. No. 54 (TM-54), University of Rochester, 1984.

[296] J. Rossignac and A. A. G. Requicha. Constant-radius blending in solid modelling. *ASME Computers in Mechanical Engineering (CIME)*, 3:65–73, 1984.

[297] J. Rossignac and A. A. G. Requicha. Constant-radius blending in solid modelling. Production Automation Project, Technical Memo. No. 42 (TM-42), University of Rochester, 1984.

[298] J. Rossignac and A. A. G. Requicha. Offsetting operations in solid modelling. Production Automation Project, Technical Memo. No. 53 (TM-53), University of Rochester, 1985.

[299] J. R. Rossignac. Specification, representation, and construction of non-manifold geometric structures. In *ACM SIGGRAPH Courses*, volume 29, chapter Course 29: Representation of Geometry for Computer Graphics. Association for Computing Machinery (ACM), 1996.

[300] J. R. Rossignac and M. A. O'Connor. *SGC: A Dimension-Independent Model for Pointsets with Internal Structures and Incomplete Boundaries*. IBM Thomas J. Watson Research Division, 1989.

[301] J. R. Rossignac and A. A. G. Requicha. Offsetting operations in solid modelling. *Computer-Aided Geometric Design*, 3(2):129–148, 1986.

[302] V. L. Rvachev. Theory of R-functions and some applications. *Naukova Dumka*, 552, 1982.

[303] E. Sacks and L. Joskowicz. *The Configuration Space Method for Kinematic Design of Mechanisms*. MIT Press, 2010.

[304] M. Sagardia, T. Stouraitis, and J. L. e Silva. A new fast and robust collision detection and force computation algorithm applied to the physics engine bullet: Method, integration, and evaluation. In *Proceedings of the 2014 Conference and Exhibition of the European Association of Virtual and Augmented Reality (EuroVR'2014)*, 2014.

[305] B. Schling. *The Boost C++ Libraries*. XML Press, 2011.

[306] L. Schwartz. *Théorie des Distributions*, volume 1, 2 of *Publications de Institut de Mathématique, Université de Strasbourg*. 1966.

[307] T. B. Sebastian, P. N. Klein, and B. B. Kimia. Recognition of shapes by editing their shock graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):550–571, 2004.

[308] J. Serra. *Image Analysis and Mathematical Morphology*. Academic Press, Inc., 1983.

[309] A. A. Shabana. *Dynamics of Multibody Systems*. Cambridge University Press, 4th edition edition, 2013.

[310] D. Shaked and A. M. Bruckstein. Pruning medial axes. *Computer Vision and Image Understanding*, 69(2):156–169, 1998.

[311] V. Shapiro. *Representations of Semi-Algebraic Sets in Finite Algebras Generated by Space Decompositions*. Ph.d. dissertation, 1991.

[312] V. Shapiro. Theory of R-functions and applications: A primer. Technical Report CPA88-3, Cornell University, 1991.

[313] V. Shapiro. Real functions for representation of rigid solids. *Computer Aided Geometric Design*, 11(2):153–175, 1994.

[314] V. Shapiro. Maintenance of geometric representations through space decompositions. *International Journal of Computational Geometry & Applications*, 7(04):383–418, 1997.

[315] V. Shapiro. Well-formed set representations of solids. *International Journal of Computational Geometry & Applications*, 9(02):125–150, 1999.

[316] V. Shapiro. *Handbook of Computer Aided Geometric Design*, chapter Solid Modeling, pages 473–518. North Holland, 1st edition edition, 2001.

[317] V. Shapiro. Semi-analytic geometry with R-functions. *ACTA numerica*, 16:239, 2007.

[318] V. Shapiro. Symmetry: Reality's riddle. In *TEDGlobal 2009*, 2009.

[319] V. Shapiro. Invited presentation (2015 Bézier award recipient): Engineering through abstraction. In *The 2015 SIAM Conference on Geometric and Physical Modeling (GD/SPM'2015)*, 2015.

[320] V. Shapiro, I. Tsukanov, and A. Grishin. Geometric issues in computer aided design/computer aided engineering integration. *Journal of Computing and Information Science in Engineering (JCISE)*, 11(2):021005, 2011.

[321] V. Shapiro and H. Voelcker. On the role of geometry in mechanical design. *Research in Engineering Design*, 1(1):69–73, 1989.

[322] V. Shapiro and D. L. Vossler. Construction and optimization of CSG representations. *Computer Aided Design*, 23(1):4–20, 1991.

[323] V. Shapiro and D. L. Vossler. Efficient CSG representations of two-dimensional solids. *Journal of Mechanical Design*, 113(3):292–305, 1991.

[324] V. Shapiro and D. L. Vossler. Separation for boundary to CSG conversion. *ACM Transactions on Graphics (TOG)*, 12(1):35–55, 1993.

[325] E. C. Sherbrooke, N. M. Patrikalakis, and F. E. Wolter. Differential and topological properties of medial axis transforms. *Graphical Models and Image Processing*, 58(6):574–592, 1996.

[326] B. K. Shoichet and I. D. Kuntz. Protein docking and complementarity. *Journal of Molecular Biology*, 221(1):327–346, 1991.

[327] B. K. Shoichet, I. D. Kuntz, and D. L. Bodian. Molecular docking using shape descriptors. *Journal of Computational Chemistry*, 13(3):380–397, 1992.

[328] K. Siddiqi, A. Shokoufandeh, S. J. Dickenson, and S. W. Zucker. Shock graphs and shape matching. In *Proceedings of the 6th IEEE conference on Computer Vision*, pages 222–229. Institute of Electrical and Electronics Engineers (IEEE), 1998.

[329] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, 1999.

[330] R. M. Solovay. A model of set-theory in which every set of reals is Lebesgue measurable. *Annals of Mathematics*, pages 1–56, 1970.

[331] S. Spitz and A. Rappoport. Integrated feature-based and geometric CAD data exchange. In *Proceedings of the 9th ACM Symposium on Solid Modeling and Applications (SMA'2004)*, pages 183–190. Eurographics Association, 2004.

[332] D. I. Spivak. *Category Theory for the Sciences*. MIT Press, 2014.

[333] H. Stachowiak. *Allgemeine Modelltheorie*. Springer-Verlag, 1973.

[334] M. Stewart, H. M. Kent, and A. J. McCoy. Structural basis for molecular recognition between nuclear transport factor 2 (NTF2) and the GDP-bound form of the Ras-family GTPase Ran. *Journal of Molecular Biology*, 277(3):635–646, 1998.

[335] A. H. Stroud. *Approximate Calculation of Multiple Integrals*. Prentice-Hall, 1971.

[336] C. Sun, Y. H. Wang, and N. D. Huang. A new plunge milling tool path generation method for radial depth control using medial axis transform. *The International Journal of Advanced Manufacturing Technology*, 76(9-12):1575–1582, 2015.

[337] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *Proceedings of the 2003 IEEE International Conference on Shape Modeling*, pages 130–139. Institute of Electrical and Electronics Engineers (IEEE), 2003.

[338] R. D. Taylor, P. J. Jewsbury, and J. W. Essex. A review of protein-small molecule docking methods. *Journal ofComputer-Aided Molecular Design*, 16(3):151–166, 2002.

[339] M. Teschner, S. Kimmerle, B. Heidelberger, G. Zachmann, L. Raghupathi, A. Fuhrmann, M. P. Cani, F. Faure, N. Magnenat-Thalmann, W. Strasser, and P. Volino. Collision detection for deformable objects. *Computer Graphics Forum*, 24(1):61–81, 2005.

[340] V. Thiagarajan and V. Shapiro. Adaptively weighted numerical integration over arbitrary domains. *Computers & Mathematics with Applications*, 67(9):1682–1702, 2014.

[341] V. Thiagarajan and V. Shapiro. Open issues in immersed boundary methods—a review. Technical report (unpublished), University of Wisconsin-Madison, 2014.

[342] R. B. Tilove. Set membership classification: A unified approach to geometric intersection problems. *IEEE Transactions on Computers*, 100(10):874–883, 1980.

[343] R. B. Tilove and A. A. G. Requicha. Closure of Boolean operations on geometric entities. *Computer-Aided Design*, 12(5):219–220, 1980.

[344] E. Tonti. *On the Formal Structure of Physical Theories*. Istituto di matematica del Politecnico di Milano, 1975.

[345] E. Tonti. *The Mathematical Structure of Classical and Relativistic Physics*. Springer, 2013.

[346] N. Tuncbag, G. Kar, O. Keskin, A. Gursoy, and R. Nussinov. A survey of available tools and web servers for analysis of protein–protein interactions and interfaces. *Briefings in Bioinformatics*, 10(3):217–232, 2009.

[347] A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. In *Proceedings of the London Mathematical Society*, volume 42 of *2*, 1936.

[348] J. H. Vandenbrande. Invited presentation: When solid modeling stopped being solid: Custom materials and additive processes. In *The 2015 SIAM Conference on Geometric and Physical Modeling (GD/SPM'2015)*, 2015.

[349] G. Varadhan and D. Manocha. Accurate Minkowski sum approximation of polyhedral models. *Graphical Models*, 68(4):343–355, 2006.

[350] H. Voelcker. *Design and Analysis of Integrated Manufacturing Systems*. National Academy Press, 1988.

[351] H. B. Voelcker and A. A. G. Requicha. Geometric modeling of mechanical parts and processes. *Computer*, 10(12):48–57, 1977.

[352] H. B. Voelcker and A. A. G. Requicha. Geometric modeling of mechanical parts and processes. Production Automation Project, Technical Memo. No. 23 (TM-23), University of Rochester, 1977.

[353] H. B. Voelcker and A. A. G. Requicha. Research in solid modeling at the university of rochester: 1972–1987. In L. Piegl, editor, *Fundamental Developments of Computer-Aided Geometric Modeling*, pages 203–254. Academic Press Ltd., London, UK, 1993.

[354] A. von Dziegielewski, M. Hemmer, and E. Schomer. High quality conservative surface mesh generation for swept volumes. In *Proceedings of the 2012 IEEE International Conference on Robotics and Automation (ICRA'2012)*, pages 764–769, 2012.

[355] M. Wan and W. A. McNeely. Quasi-static approximation for 6 degrees-of-freedom haptic rendering. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, pages 257–262, Washington, DC, USA, 2003.

[356] G. G Wang. Definition and review of virtual prototyping. *Journal of Computing and Information Science in Engineering*, 2(3):232–236, 2002.

[357] H. Wang. Grid-search molecular accessible surface algorithm for solving the protein docking problem. *Journal of Computational Chemistry*, 12(6):746–750, 1991.

[358] K. Weihe and T. Willhalm. Reconstructing the topology of a CAD model—a discrete approach. *Algorithmica*, 26(1):126–147, 2000.

[359] K. J. Weiler. *Topological Structures for Geometric Modeling*. Ph.d. dissertation, Rensselaer Polytechnic Institute, 1986.

[360] R. Weller. Inner sphere trees. In *New Geometric Data Structures for Collision Detection and Haptics*, pages 113–144. Springer International Publishing, 2013.

[361] R. Weller and G. Zachmann. Inner sphere trees for proximity and penetration queries. In *Proceedings of the 2009 Robotics: Science and Systems Conference (RSS'2009)*, volume 2, 2009.

[362] R. Weller and G. Zachmann. Inner sphere trees and their application to collision detection. In G. Brunnett, S. Coquillart, and G. Welch, editors, *Virtual Realities*, pages 181–201. Springer Vienna, 2011.

[363] B. Whited and J. Rossignac. Relative blending. *Computer-Aided Design*, 41(6):456–462, 2009.

[364] H. Whitney. *Local Properties of Analytic Varieties*, pages 497–536. Birkhäuser Boston, 1992.

[365] F. E. Wolter. Cut locus and medial axis in global shape interrogation and representation. *Computer Aided Geometric Design*, pages 92–2, 1992.

[366] Y. Wu, J. J. Shah, and J. K. Davidson. Improvements to algorithms for computing the Minkowski sum of 3-polytopes. *Computer-Aided Design*, 35(13):1181–1192, 2003.

[367] X. X. Sheng and I. R. Meier. Generating topological structures for surface models. *IEEE Computer Graphics and Applications*, 15(6):35–41, 1995.

[368] Chee-Keng Y. Towards exact geometric computation. *Computational Geometry*, 7(1):3–23, 1997.

[369] Y. Yamaguchi, H. Nakamura, and F. Kimura. Probabilistic solid modeling: A new approach for handling uncertain shapes. In *Selected and Expanded Papers from the IFIP TC5/WG5.2 Working Conference on Geometric Modeling for Product Realization*, pages 95–108. Association for Computing Machinery (ACM), 1992.

[370] A. Yershova, S. Jain, S. M. Lavalle, and J. C. Mitchell. Generating uniform incremental grids on $so(3)$ using the Hopf fibration. *The International Journal of Robotics Research*, 2009.

[371] A. Yershova and S. M. LaValle. Deterministic sampling methods for spheres and $so(3)$. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA'2004)*, volume 4, pages 3974–3980, 2004.

[372] X. Zhang, Y. J. Kim, and D. Manocha. Reliable sweeps. In *Proceedings of the 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling (SPM'2009)*, pages 373–378. Association for Computing Machinery (ACM), 2009.