

Geometric Studies in Point Cloud Skeletonisation and Their Applications in Robotics

**A thesis submitted in partial fulfilment
of the requirement for the degree of Doctor of Philosophy**

Qingmeng Wen

**Supervised by
Dr. Seyed Amir Tafrishi**

July 2025

**Cardiff University
School of Engineering**

**To People I care
for their patience and support.**

Abstract

Skeletons are commonly known as internal frameworks of humans and animals, but a similar concept is widely used in shape analysis. In this context, a skeleton represents a compact, centred abstraction of an object’s shape, capturing its topological structure. The process of extracting such a representation is known as *skeletonisation*. As a useful tool in shape analysis, skeletonisation has been extensively studied; however, key challenges remain—particularly in reducing computational overhead, establishing robust evaluation metrics, and applying skeletons to surface generalisation.

This thesis addresses these gaps through three major contributions. First, to mitigate the high computational cost of existing skeletonisation methods, we analyse their inefficiencies and propose a novel geometry-based skeletonisation framework. By eliminating redundant computations inherent in prior techniques, our method significantly accelerates the skeletonisation process, outperforming a well-established baseline in terms of speed.

Second, the development and comparison of skeletonisation techniques have been hindered by the absence of comprehensive quantitative evaluation tools. To address this, we systematically investigate the fundamental properties of skeletonisation and introduce a suite of evaluation metrics, along with practical numerical techniques. These metrics are validated for their effectiveness and demonstrated in robotics-related scenarios.

Finally, the thesis explores a novel application of skeletons as priors for modelling

rolling contact kinematics. By extracting skeletons from point clouds, we reconstruct the topological structure of arbitrary surfaces, enabling rolling contact analysis for complex, irregular, and even discontinuous object shapes.

Acknowledgements

I would like to express my deepest gratitude to all those who supported and guided me throughout the completion of this thesis.

First and foremost, I am profoundly grateful to my supervisors, Dr. Seyed Amir Tafrishi, Dr. Ze Ji, and Prof. Yu-Kun Lai, for their invaluable guidance, continuous encouragement, and unwavering support at every stage of this research. Their insightful feedback and constructive criticism were essential to the development of this work.

I would also like to thank the faculty and staff of the School of Engineering at Cardiff University for their kind support and assistance throughout my PhD journey. Special thanks go to Prof. Mikhail Svinin for his helpful discussions and valuable suggestions.

I am sincerely thankful to my colleagues and friends, particularly Baoming Guo, Yating Li, Jun Song, for the stimulating conversations, emotional support, and moments of laughter that made this journey more enjoyable and meaningful.

A very special thanks goes to Shuang Kang, whose presence turned the darkest moments into something endurable, and the brightest ones into something unforgettable. Your calm, support, and steady companionship made all the difference. I am truly thankful.

I would like to acknowledge the financial support provided by the China Scholarship Council, which made this research possible.

Last but certainly not least, I would like to express my heartfelt appreciation to my

family. To my parents, thank you for your unconditional love, endless patience, and steadfast belief in me. Your encouragement has been the foundation of this journey.

This thesis would not have been possible without all of you.

Contents

Abstract	v
Acknowledgements	vii
Contents	ix
List of Publications	xiii
List of Figures	xv
List of Tables	xxi
List of Algorithms	xxiii
List of Acronyms	xxv
1 Introduction	1
1.1 Motivation	2
1.2 Aim and Objectives	4
1.3 Contributions	5

1.4	Outline of the Thesis	7
2	Literature Review	9
2.1	Introduction of Object Skeletons	9
2.1.1	Definition of Skeleton	9
2.1.2	Skeleton Properties	13
2.1.3	Object Skeleton Applications in Robotics Research	16
2.1.4	Object Skeleton Applications in Other Fields	17
2.2	Skeletonisation Methods and Evaluation	19
2.2.1	Boundary-based Methods	19
2.2.2	Region Field Methods	22
2.2.3	Topological Thinning Methods	25
2.2.4	Learning-Based Methods	29
2.2.5	Skeletonisation Evaluation	32
2.3	Rolling Contact Kinematics	35
2.3.1	Application Background	36
2.3.2	Theoretical Research	37
2.4	Summary	38
3	GLSkeleton	41
3.1	Introduction	41
3.2	GLSkeleton Framework	43
3.2.1	Overview of Laplacian-based Skeletonisation	44

3.2.2	Geometric Local Point Reduction Strategy	47
3.2.3	Terminating Condition for Contraction Iterations	50
3.3	Results and Discussion	51
3.3.1	Contraction Speed Comparison	51
3.3.2	Skeletonisation Results	54
3.4	Conclusion	55
4	Desirable Skeletonisation Properties and Skeletonisation Quality Evaluation	59
4.1	Introduction	59
4.2	Characteristics of Stable Contraction-Based Skeletonisation	60
4.3	Skeletonisation Evaluation Metrics	64
4.3.1	Topological Similarity	64
4.3.2	Boundedness	73
4.3.3	Centeredness	78
4.3.4	Smoothness	84
4.4	Results & Discussion	88
4.4.1	Results of the Laplacian-based Skeletonisation Characteristics	88
4.4.2	Quantitative Evaluation Experiments on Skeletonisation Results	93
4.4.3	Desired Properties in Robotic Applications	103
4.5	Conclusion	106

5	Skeleton-Guided Rolling Contact Kinematics	109
5.1	Introduction	109
5.2	Skeletonisation, Slicing and Curve Fitting	111
5.2.1	Point Cloud Skeletonisation	111
5.2.2	Slicing and Curve Fitting	112
5.3	Kinematics Model with Generalised Semi-Convex Surface	116
5.3.1	Modelling by Locally Generalised Semi-Convex Surface . . .	116
5.3.2	Modelling by a Globally Generalised Semi-Convex Surface .	120
5.4	Results and Discussion	122
5.4.1	Simulation Results Using the Locally Generalised Surface Model	123
5.4.2	Simulation Results of the Model with Globally Generalised Surface	125
5.4.3	Comparison	132
5.5	Conclusions	135
6	Conclusion and Future Work	137
6.1	Conclusion	137
6.2	Future Work	138
	Bibliography	141

List of Publications

The work introduced in this thesis is based on the following 5 publications:

Peer-Reviewed Publications:

- **Q. Wen**, S. A. Tafrishi, Z. Ji and Y.-K. Lai, "GLSkeleton: A Geometric Laplacian-Based Skeletonisation Framework for Object Point Clouds", *IEEE Robotics and Automation Letters*, 2024.
With presentation at ICRA@40
- **Q. Wen**, S. A. Tafrishi, Z. Ji and Y.-K. Lai, "What Criteria Define an Ideal Skeletonisation Reference in Object Point Clouds?", *Towards Autonomous Robotic Systems (TAROS)*, 2024.
- **Q. Wen**, Y.-K. Lai, Z. Ji, M. Svinin and S. A. Tafrishi, "Skeleton-Guided Rolling-Contact Kinematics for Arbitrary Point Clouds via Locally Controllable Parameterized Curve Fitting", *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hangzhou, China, 2025*.

Under Review:

- **Q. Wen**, Y.-K. Lai, Z. Ji and S. A. Tafrishi, "Skeletonization Quality Evaluation: Geometric Metrics for Point Cloud Analysis in Robotics".
- **Q. Wen**, Y.-K. Lai, Z. Ji and S. A. Tafrishi, "Differential-Geometry Rolling Contact Kinematics for Point Clouds via Skeleton and Fourier Models".

Other Publications:

- Munguia-Galeano, F., Veeramani, S., Hernandez, J. D., Wen, Q. and Ji, Z. 2023. Affordance-based human-robot interaction with reinforcement learning. IEEE Access 11, pp. 31282-31292. (10.1109/ACCESS.2023.3262450)

List of Figures

2.1	Medial axis of a 2D and 3D object respectively [83, 138].	11
2.2	Curve skeleton generated from the object's mesh and point cloud, respectively.	11
2.3	Voronoi digram (VD) based skeleton [100]. (a) Voronoi diagram; (b) Extracted skeleton; (c) Bounded Voronoi diagram.	21
2.4	Distance field based skeletonisation [100]. (a) Binary shape; (b)Distance field; (c) Extracted skeleton by mean curvature map.	24
2.5	Deletion-based thinning for binary images [30].	26
2.6	Laplacian-based point cloud contraction and skeletonization [34]. (a)-(b) illustrate the progressive Laplacian-based contraction of the input horse point cloud; (c) shows the extracted curve skeleton obtained from the contracted (thinned) point cloud; and (d) presents the final refined skeleton.	28
2.7	Learning-based direct skeletonisation from images [116].	29
2.8	Learning-based skeletonisation from natural images [136].	31
2.9	Example of rolling-contact kinematics application background. (a) An example spherical rolling robot [16]. (b) An example of a rolling-based manipulator [188].	37

3.1	Skeleton-assisted grasp planning. (a) The grasp candidate generated from the object skeleton for a two-finger gripper. It is represented as a straight line corresponding to the gripper opening, with short perpendicular markers at both ends indicating the gripper's orientation. (b) The robot executing a grasp using the two-finger gripper.	42
3.2	Block diagram of the GLSkeleton framework	44
3.3	Each anchor point (in red) and its neighbours (in yellow) are projected onto a local tangent plane (a). The neighbour point ring is inferred from the Delaunay triangulation of the projected points (b), although it is not explicitly visualised.	45
3.4	Illustration of LPR. (a) Original egg model. (b) Result after one contraction iteration; yellow circles indicate improved regions, red points are removed, and blue points remain. (c) Reduction in point number and contraction cost. (d) Evolution of Δe over iterations.	52
3.5	Contraction time comparison between GLSkeleton (blue) and the baseline (red). Top: Contraction times for 216 models with 1024, 4096, and 16384 points. Bottom: Remaining points after contraction.	53
3.6	Curve computation cost from contracted point cloud (a) and overall computation cost (b).	54
3.7	Skeletonisation comparison for selected models. (a)-(f) Results using the baseline; (g)-(l) results using GLSkeleton.	57
3.8	Skeletonisation results for noisy point clouds. Gaussian noise with standard deviations of 5%, 10%, and 15% (relative to half the diagonal of the object's bounding box) is added to each point.	58
4.1	Examples of simplices in various dimensions.	66
4.2	An example of a Vietoris-Rips complex [15].	67

4.3	Persistent homology analysis of simple geometries using H_0 features. Shapes in (a), (b), (c), and (g) contain 200 points, and (h) contains 2000 points. The corresponding barcode diagrams are shown below each shape. The horizontal axis indicates birth and death scales, while the vertical axis simply indexes the bars for visual clarity.	71
4.4	The barcode of persistent homology (H_0 features). The input point clouds are scaled to fit within a cubic bounding box whose diagonal is 1.6 and only the top 5% and bottom 5% of persistence bars are shown for clarity)	73
4.5	Projection of directional vectors onto a spherical surface.	74
4.6	Triangulated projections of points on the unit sphere for (a) a bounded skeletal point and (b) an unbounded skeletal point, based on a horse-shaped point cloud.	77
4.7	Boundedness assessment of contracted skeletal points derived from a horse point cloud. Points exterior to the shape boundary demonstrate lower boundedness metrics, whereas points interior to the boundary exhibit values approaching unity.	77
4.8	Centredness of skeletal points (a) and points sampled from the curve skeleton (b), respectively.	80
4.9	Points (orange) separated by two parallel cutting planes.	81
4.10	Ellipse fitting applied to point sets sampled from simple geometries, demonstrating robust centre estimation.	83
4.11	Smoothness of 2D skeletal points (a) and 2D curve skeleton (b) in different smoothness. The vertices in (b) without valid smoothness value are marked in magenta.	86

4.12	Evolution of normal vector and curvature difference distributions throughout the contraction process. The selected object is a chilli point cloud.	89
4.13	Curvature differences between the input and point cloud after final contraction. The original point cloud (grey) with the contraction results (red) is put at the top right corner of each histogram.	90
4.14	Normal vector differences between the input and contracted point clouds. Insets at the top right of each histogram show the original point cloud (grey) alongside the contraction results (red).	91
4.15	Resulting curve skeletons.	92
4.16	Curvature and normal vector differences between the input and contracted point cloud obtained via the GLSkeleton method. Insets at the top right of each histogram show the original point cloud (grey) alongside the contraction results (yellow).	93
4.17	Skeleton-guided surface reconstruction [178].	104
4.18	Skeleton-based catheter navigation [57].	105
5.1	Illustration of point cloud slicing and curve fitting. (a) shows the extracted section points from the point cloud. (b) and (c) demonstrate the corresponding fitted curves in the transformed 2D coordinate system and in Cartesian coordinates, respectively.	113
5.2	Block diagram illustrating the slicing procedure and Fourier-based curve fitting method.	114
5.3	Parametrisation and frame definitions for rolling contact kinematics between a rolling sphere and a geometrically fitted surface.	115
5.4	Simulation trajectories for the carrot object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.	124

5.5	Simulation trajectories for the banana object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.	126
5.6	Simulation trajectories for the egg object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.	127
5.7	(a) Time responses of states \mathbf{u}_g , \mathbf{u}_s , and ψ for the carrot simulation corresponding to Figure 5.4(b). (b) Variation of radius R_g for the simulations shown in Figures 5.4–5.6.	128
5.8	Simulated contact trajectories of rolling contact between a sphere and the carrot, banana, and egg objects, respectively.	129
5.9	Simulated contact trajectories of rolling contact between a sphere and (A) the banana and (B) the egg objects.	130
5.10	(a) Time responses of the states \mathbf{u}_o , \mathbf{u}_s , and ψ for the example simulation in Fig. 5.8D. (b) Response of the radius R_o for the example simulations in Figs. 5.8A-D and 5.9A-B.	130
5.11	Simulation of rolling in the Indian Tunnel environment [176].	131
5.12	Rolling simulation without slippage of the locally (a) and globally (b) generalised surface based rolling contact kinematics models.	132
5.13	Rolling simulation with slippage of the locally (a) and globally (b) generalised surface based rolling contact kinematics models.	133
5.14	Time responses of the states \mathbf{u}_o , \mathbf{u}_s , and ψ for the example simulation corresponding to Fig. 5.12 ((a) and (b)), and Fig. 5.13 ((c) and (d)) respectively.	135
6.1	Promising potential robotic applications.	139

List of Tables

3.1	Parameter settings for GLSkeleton	52
4.1	Evaluation of skeletonisation results for normal, noise-induced, and sparse point cloud forms, assessed using the proposed metrics. The colour indicates the local scoring along the skeleton.	95
4.2	Quantitative results of skeletonisation evaluation. The input cave point cloud is sectioned and capped from the original data to ensure a closed shape for skeletonisation. Different resolutions of cave and synthetic tree point clouds are acquired by grid-averaged sampling.	100
5.1	Summary of constant parameters used in the simulations.	123

List of Algorithms

3.1	The computation of LPR of GLSkeleton.	48
-----	---	----

List of Acronyms

OCR optical character recognition

1D one dimensional

2D two dimensional

3D three dimensional

MA medial axis

MAT medial axis transform

MAT medial axis transform

LBC Laplacian-based contraction

LPR local point reduction

GLSkeleton geometric Laplacian-based skeletonisation framework for object point clouds

Chapter 1

Introduction

Skeletons are usually associated with vertebrates as real and important components, while they can also be a functional shape representation assisting high-level shape analysis. It is said that the skeletal information plays a critical role when humans distinguish various shapes. For a shape, the skeleton of that shape is the effective and intuitive abstraction of that shape, preserving the topological path of the shape and representing the main feature that makes the shape distinguishable from others [156, 123].

The skeleton of a shape provides information about abstracted topology that makes it easier for shape analysis. For example, it might be utilised for applications including shape manipulation and shape reconstruction [34], image processing [40], animation [156, 40]. Recently, the achievement of skeletonisation algorithm research also demonstrates the potential for robotic applications, such as grasp planning, navigation, SLAM, and rolling contact systems [40, 166, 165, 75, 109]. Despite the promising application potential of shape skeletonisation, several critical issues remain largely unaddressed. For example, most existing point cloud skeletonisation methods are computationally inefficient [34, 155]; however, point clouds are actually the most accessible 3D data in practice [22]. Besides, there are no common-agreed metrics that comprehensively evaluate the skeletonisation results, especially for curve skeletons of 3D shapes [155]. As for application in robotics, previous studies dominantly focus on grasping planning and navigation [165, 166, 109], while the skeleton of shapes also provides the link between geometric properties and rolling kinematics.

In the field of skeletonisation, two primary skeletal representations are commonly used: the medial axis surface and the curve skeleton. Among them, the curve skeleton has emerged as a particularly prominent descriptor of shape, often compared to its counterpart, the medial axis surface [23, 40]. Owing to their simpler topology and reduced computational complexity, curve skeletons are generally preferred in practical applications [123, 165].

Thus, this thesis comprehensively investigates curve skeletonisation methods for shapes represented by point clouds, including the computational efficiency, quantitative evaluation, and application in rolling contact kinematics.

1.1 Motivation

As a powerful tool for shape analysis, skeletonisation has found wide applications across diverse domains, including computer vision, medical imaging, robotics, and agriculture [127, 40, 155]. As an efficient shape representation, skeletons have been applied to tasks such as shape matching, retrieval, decomposition, reconstruction, and segmentation [162, 12, 8, 149, 11, 88, 21]. In the field of medical imaging, skeletons serve as useful tools for vessel analysis and morphological assessment [159, 36]. Additionally, skeleton-based approaches have recently shown promise in agricultural phenotyping [69, 99].

Skeleton representations have shown promise across various robotic domains, including manipulation, path planning, navigation, and SLAM [40, 109, 151, 178, 165, 75]. By capturing the topological and geometric structure of objects, skeletons provide a compact and expressive framework for reasoning about shapes and motion. However, despite their potential, skeletonisation techniques have not yet achieved widespread adoption in practical robotic systems.

Several factors contribute to this gap. In navigation and path planning, for instance, skeletons provide a natural representation of traversable routes [40, 109]; however, the

high computational cost of extracting skeletons limits their use in real-time applications [32]. In manipulation, skeleton-based grasp planning can yield more intuitive and natural grasp candidates [165], yet the lack of unified metrics for quantifying skeleton quality prevents consistent and comparable performance evaluation. Moreover, understanding continuous surface geometry is crucial for rolling contact kinematics [103], but this remains a fundamental challenge because robots typically operate on discrete representations of shapes. In this context, skeleton representations can simplify 3D shape analysis and support geometric reasoning [34]. Specifically, skeletons may serve as a bridge between discrete shape representations and continuous kinematic equations by reducing 3D shape analysis to local 2D cross-sectional shapes, which are intrinsically linked to the underlying skeleton.

Motivated by these challenges, this thesis addresses three critical gaps that currently hinder the broader application of skeleton-based methods in robotics: (1) the high computational cost of skeletonisation, (2) the lack of standardised evaluation metrics, and (3) the limited exploration of skeleton-based approaches in specific robotic tasks.

First, the computational expense of skeletonisation, particularly for 3D shapes, poses a significant challenge for real-world robotic applications. Since most real-world objects are inherently 3D, efficient analysis of 3D shapes is essential in robotics [166, 165, 119]. Point clouds, being one of the most accessible forms of 3D data, are frequently used for shape representation. However, skeleton extraction from point clouds is typically resource-intensive [180]. Therefore, one key goal is to enhance the computational efficiency of skeletonisation methods, especially those targeting robot-friendly applications.

Second, despite extensive research on skeletonisation, there is no universally accepted definition of a skeleton, and thus no standardised method for evaluating skeleton quality [155]. Prior studies have often relied on visual inspection [34, 123, 66] or comparison against manually annotated ground-truth skeletons [101, 49], both of which can be subjective and inconsistent. Although certain performance metrics like recon-

structibility, memory usage, and runtime efficiency have been used [158, 68, 79], a comprehensive and robotics-oriented quantitative evaluation framework for 3D curve skeletonisation is still lacking.

Third, while skeleton-based techniques have been applied in tasks like grasping and navigation, their potential for bridging the gap between discrete shape representations (e.g., point clouds) and continuous kinematic models remains largely unexplored. This is particularly relevant for understanding rolling contact kinematics—a common phenomenon in robotics [150]. Traditional studies often describe rolling motion based on continuous surface models [103, 150]. However, real-world robotic perception often deals with discrete data formats, such as point clouds [22]. Thus, exploring how skeletons can serve as an intermediary between discrete observations and continuous motion models could offer novel insights and practical value for robotics research.

1.2 Aim and Objectives

The object skeletons look promising in application in robotics; however, they face some serious fundamental and practical problems. Thus, this thesis aims to investigate a faster point cloud skeletonisation method, define metrics that numerically evaluate the skeletonisation performance, and bridge the gap between rolling contact kinematics and arbitrary surface geometries with skeleton information. The objectives of this thesis are as follows:

- To conduct a comprehensive and structured literature review on object skeletons, point cloud skeletonisation techniques, and their robotics applications, with a focus on identifying existing limitations and gaps in computational performance and evaluation.
- To design and implement an improved framework for point cloud skeletonisation that enhances computational efficiency while preserving or improving skeleton

quality relative to existing methods.

- To formulate formal, quantitative, and reproducible evaluation metrics that enable comprehensive assessment of skeletonisation quality.
- To develop a modelling approach for rolling-contact kinematics that integrates skeleton information extracted from point cloud data, enabling the approximation of continuous surface interactions from discrete shape representations.
- To develop, document, and release open-source toolboxes that implement the proposed methods and provide accessible resources for the broader research community.

1.3 Contributions

With respect to the objectives, multiple contributions are achieved and summarised in this section.

Addressing the computational efficiency problem of the skeletonisation approaches, work in this thesis makes contributions as follows,

- **Local Geometry-Based Point Reduction (LPR):** Introduced a strategy that accelerates Laplacian-based contraction (LBC) by eliminating redundant points while preserving geometry features, improving robustness against noise and non-uniform point distributions.
- **Novel Global Terminating Condition:** Proposed a new condition, using global features rather than local geometrical properties, for Laplacian-contraction with in-loop reduction, enhancing robustness to local changes.
- **Comprehensive Computational Efficiency Evaluation:** Conducted an evaluation of the geometric method's computational performance using a real-world

dataset, demonstrating its effectiveness in accelerating the skeletonisation process while preserving the quality of the resulting skeletons.

Addressing the skeletonisation evaluation problem, work in this thesis makes the following contribution,

- **Desirable Geometrical Properties:** Investigates the desirable geometrical properties of contraction-based skeletonisation and the quality of the resultant skeleton.
- **Formal and Numerical Evaluation Metrics:** Introduce formal qualitative criteria—such as topological similarity, boundedness, centeredness, and smoothness—and develop numerical and representative metrics to enable comprehensive evaluation of point cloud surfaces and their skeletons.
- **Validation and Application-based Evaluation:** Analyse the performance, robustness, and sensitivity of the proposed metrics, and validate them across diverse point cloud datasets relevant to robotics applications.

As for extending the skeleton application in bridging gaps between the existing continuous rolling contact kinematics theory and discrete object surfaces, the corresponding contributions are:

- **Skeleton-Guided Point Cloud Slicer:** Present an innovative technique for segmenting point cloud data, utilising skeletal structures and curvature-based geometric features to guide the slicing process.
- **Fourier-based Surface Reconstruction:** Design of an effective surface reconstruction framework that transforms point cloud data into smooth and continuous surfaces via Fourier-based curve approximation.

- **Geometry-based Kinematic Model:** Develop two differential geometry-inspired kinematic frameworks tailored for general semi-convex surfaces, formulated within two coordinate systems to ensure adaptability across various object forms. One coordinate system is locally constructed while another one is globally constructed.
- **Simulation and Evaluation:** Assess the performance of the proposed methods through rigorous experiments involving both synthetic and real-world object datasets.

Besides, all toolboxes¹²³ developed as part of this thesis have been released as open-source resources for the research community.

1.4 Outline of the Thesis

This outline gives brief summaries of each chapter in this thesis, which is organised as 6 chapters in total. The overview of each chapter is as follows.

Chapter 2 introduces relevant literature of skeletonisation theory, including multiple skeletonisation approaches, existing skeletonisation evaluation methods, various skeletonisation applications, and the rolling contact theory.

Chapter 3 introduces a geometric Laplacian-based skeletonisation framework for object point clouds, which increases the computational efficiency by removing redundant computation from LBC. The results showed that our methods are remarkably more efficient than the baseline with a trivial sacrifice in skeletonisation quality.

Chapter 4 demonstrates our findings in characteristics of desirable skeletonisation patterns when doing Laplacian-based curve skeletonisation from point cloud data, fol-

¹GLSkeleton Toolbox: <https://github.com/weiqimeng1/GLSkeleton>

²Skeletonisation Evaluation Toolbox: https://github.com/weiqimeng1/PointCloud_Skeletonization_Metrics

³Skeleton-based RCK Toolbox: https://github.com/weiqimeng1/Geometric_RCK

lowed by numerically defined metrics that comprehensively evaluate the skeletonisation qualities from various aspects, with discussion in robotic applications. The results indicate that our proposed metric effectively provides quality scores from different perspectives.

Chapter 5 introduces our work in rolling contact kinematics modelling approach with the information provided by point cloud skeletonisation results, bridging the gaps between rolling contact kinematics and geometrical properties of arbitrary shape. The simulation results of rolling contact motion show that the skeleton information successfully guides the rolling states update, enabling the simulation of rolling trajectories with given motion settings.

Chapter 6 concludes all the contributions of this thesis. Besides, the remained limitations and challenges are summarised, followed by future research directions.

Chapter 2

Literature Review

Over the past decades, researchers have extensively investigated the field of skeletonisation. This chapter presents a literature review from the perspective of both skeletonisation research and its relevant progress in robotics. Specifically, it includes an introduction of object skeletons (Section 2.1), existing methods and evaluation approaches (Section 2.2), and rolling contact kinematics research with its application background (Section 2.3).

2.1 Introduction of Object Skeletons

2.1.1 Definition of Skeleton

Skeleton normally refers to the frame of bones that support the body of vertebrates. However, in this thesis, we refer to “skeleton” as an abstraction of a shape that preserves the topological and geometric properties of the shape. Object skeletons, as a useful shape descriptor, have been extensively researched [155]. However, the definition of skeletons may vary. The common skeleton types are known as medial axis transform (MAT) and curve skeleton. Curve skeletons are gaining wide acceptance for 3D shape representations due to their simplicity, while MATs are still too complex for analysis [123, 34, 165].

Medial Axis Transform

The first formal definition of skeletonisation was proposed by Blum [23], who defined the skeleton as “*the locus of an object shape in \mathbb{R}^n .*” More specifically, this definition corresponds to the concept of the medial axis (MA). The MA is a set of centre points lying within the object, each of which is equidistant to at least two points on the object’s boundary. To enable shape reconstruction, the medial axis transform (MAT) is defined as the set of MA points together with their associated radii of maximal inscribed circles (in 2D) or balls (in 3D).

To provide an intuitive understanding of the MA, Blum introduced the well-known *grassfire transform*. In this analogy, imagine a planar object (like dry grass) whose boundary is set on fire. As the fire fronts propagate inward at a constant speed, they eventually meet and extinguish at specific points within the shape. These meeting points collectively form the MA, representing the loci where fronts originating from different parts of the boundary collide [24]. Assuming a unit propagation speed, the distance from the boundary to each meeting point—equivalently, the radius of the maximal inscribed circle or ball—indicates how far the fire fronts have travelled before meeting. In this sense, the radius in the MAT reflects the local geometric thickness or scale of the shape rather than a temporal quantity.

Generally, as illustrated in Fig. 2.1, the MA of a two dimensional (2D) shape is a collection of one dimensional (1D) manifolds (curve segments), whereas the MA of a three dimensional (3D) shape is a collection of 2D manifolds that may intersect along curves or at singular points. In the context of object skeletons, “1D” indicates that the elemental components of the skeleton are 1D manifolds. According to previous research, the MA can be approximately obtained through three main approaches: by constructing a Voronoi diagram, by iterative thinning or erosion, or by computing a distance transform field [31, 105, 77, 163, 27].

The MAT provides abundant abstract information for shape analysis and is ideal for




Figure 2.1: Medial axis of a 2D and 3D object respectively [83, 138].




Figure 2.2: Curve skeleton generated from the object's mesh and point cloud, respectively.

scenarios with only 2D shapes. However, MAT of 3D shapes is usually composed of multiple intersected sheets with a 2D structure, making it difficult for modelling and manipulation of the represented shape, while 3D models are common in application scenarios. It is worth mentioning that one may refer to the MA and its approximations as *surface skeleton* due to its surface structure [144]. Besides, MAT always retains too many unnecessary details, obstructing the acquisition of useful information [40, 123].

Curve Skeleton

Due to the demand for handling 3D shapes, the curve skeleton is a new powerful tool for 3D shape analysis due to its simpler and more abstracted structure [156, 34]. This thesis mainly focuses on curve skeleton due to its advantages in applications. As shown in Fig. 2.2. Curve skeleton is typically of a simple 1D structure (a collection of 1D manifolds) and can be obtained from common 3D shape input, such as a mesh and point cloud. However, unlike MATs, curve skeletons do not have a strict mathematical definition. Instead, various definitions of curve skeletons are constantly proposed in the literature, stopping from being used in practice[40]. One possible definition proposed by Dey and Sun states that the curve skeleton is a subset of the medial surface, obtained by a non-trivial elimination process [3, 45]. However, as argued by Cornea et al. [40], this definition is rather too strict in practice. Instead, they define the curve skeleton as simply as a 1D representation abstracted from a 3D shape and provide multiple definitions of skeleton properties, depicting expected skeletons. One related shape representation is bone-skeleton, with application in character animation [14, 184] and mesh deformation [139, 174]. The bone skeleton is a structure that is further simpler than the curve skeleton and can be considered as a subset of the curve skeleton.

The various definitions of curve skeleton make it hard to evaluate the desirable skeleton in the skeletonisation field, while quantifying the skeleton quality is useful for judging the effectiveness of the skeletons in application. To help the reader gain an intuitive understanding of curve skeleton, the author would like to give the curve skeleton definition as follows:

A curve skeleton of a shape is a collection of 1D manifolds embedded in \mathbb{R}^3 , composed of one or multiple branches aligned with the object's geometry, where each point approximately lies at the cross-sectional centre of the shape.

2.1.2 Skeleton Properties

In prior research, the definition of a shape's skeleton differs across studies. Without a common definition, skeletons become intangible concepts, especially for the curve skeleton. To address this issue, researchers investigated multiple desirable properties and several of these proposed properties are widely accepted, including homotopy equivalence, thinness, centredness, reconstructibility, and smoothness [155, 40, 144].

Homotopy Equivalence. A primary purpose of a skeleton is to abstract the topology of the original shape; thus, preserving topological structure is essential. Two shapes are considered *homotopy equivalent* if they share the same topological invariants, such as connected components, tunnels, and cavities [82, 129]. For medial axes (MA), this property is well established: Lieutier [82] proved that the MA preserves the topology of the input under mild assumptions. Owing to its rigorous mathematical definition, the topology preservation of the MA is relatively straightforward to analyse.

In contrast, curve skeletons are 1D structures without cavities and lack a universally accepted ground-truth definition, making strict homotopy equivalence difficult to characterise. To address this, Cornea et al. [40] proposed a relaxed notion of topological preservation: a curve skeleton S preserves the topology of an object O if it maintains the same number of connected components and includes at least one loop corresponding to each tunnel or cavity in O . Although this practical criterion is widely adopted, quantitative evaluation remains challenging. Point clouds are unordered and unconnected with non-uniform sampling, whereas curve skeletons are ordered, connected, and one-dimensional, making direct topological comparison difficult. As a result, skeleton quality assessment often relies on visual inspection [143].

Many point cloud skeletonisation methods naturally generate intermediate contracted point sets (often referred to as skeletal points) during the extraction process. Since both the original point cloud and these intermediate point sets share the same representation type, their topological properties can be compared directly, making the

analysis substantially more tractable than attempting to compare a point cloud with a one-dimensional curve skeleton.

Because homotopy equivalence cannot be reliably determined directly from discrete point sets, *persistent homology* is widely adopted as a more feasible alternative for topological analysis. Niyogi et al. [108] demonstrated that the topological invariants of a manifold can be recovered from sufficiently dense point samples. Persistent homology has since been applied to various shape-analysis frameworks, including size functions [167, 56, 168], stability analyses in machine-learning contexts [17], and robotics applications such as caging grasps using nontrivial H_1 features [166].

Thinness: The skeleton is an abstracted representation of a shape. Thus, it should be as thin as it is allowed by the space [40, 155].

Generally, the thinness reflects the degree of abstraction in a skeletal representation, with the ideal skeleton occupying minimal space. While both the MA and curve skeleton are inherently thin, the curve skeleton is superior, as it embodies a more abstracted structure..

Centredness: Each skeleton point of a shape is desired to be located at the centre of the shape. By definition, the MA is defined to be the centre of the shape as a point of MA has equal distance to at least two separate points of the shape boundary. However, the centredness of the curve skeleton is ambiguous. Some researchers suggest that the curve skeleton is defined as a subset of the MA, so each curve skeleton point should satisfy the same distance property as the MA [45, 156]. However, this strict definition may be unnecessary in most applications. For instance, in navigation, a planned route does not need to strictly follow the central line. Moreover, enforcing a strict distance property can make the skeleton highly sensitive to minor perturbations [40]. Furthermore, strictly following the distance property may lead to a sacrifice in smoothness, which may be unexpected in scenarios such as navigation. In fact, the approximate centredness is enough for most known application scenarios. The scenarios may include skeleton as a path for navigation [40], skeleton-assisted analysis in agricultural

or robotic manipulation [147, 165, 75], and agricultural plant phonemics [179, 99].

Reconstructibility: Since a skeleton is one of the shape representation forms, previous researcher argues that a good skeleton should have good reconstructibility [40, 155]. It is the critical property of a skeleton indicating the preserved information of a shape. Thus, researchers in previous work might evaluate their work by demonstrating the capability of their curve skeletonisation method in reconstructing a shape [11, 34]. Good shape reconstructibility may be essential for 3D reconstruction [2], while it is less critical for other applications such as grasping planning (deriving grasping candidates) [166, 165] and navigation (curve skeleton used as a navigation path) [109].

In fact, this property is typically discussed with curve skeletons, while the process of generating MAT generally is invertible.

Smoothness: As a shape representation, the skeleton is not necessarily required to have a smooth structure, since smoothness does not directly affect its capability to encode the shape. However, smoothness becomes a critical criterion in many practical applications [154, 40], particularly in navigation tasks such as camera motion planning or robot path generation. In these scenarios, the skeleton often serves as a reference trajectory for movement, where abrupt changes in direction are undesirable and may lead to unstable or infeasible motion. As argued by Cornea et al. [40], a curve skeleton can be effectively used as a camera translation path only if it is sufficiently smooth. Specifically, the tangent variation along the skeleton should be minimal, ideally ensuring at least C^2 continuity [141, 144]. Such smoothness guarantees that the derived navigation paths are feasible, stable, and physically realisable.

For MA, its smoothness strictly depends on the original shape, potentially leading to abrupt changes in the skeleton path. As for curve skeleton, its structure is more flexible and can be smooth with smoothness constraints while obtaining a curve skeleton [154].

To summarise, an object's skeleton is usually of a lower dimension and simpler than

the object's shape according to the definition. Thus, it is a useful tool in shape analysis, such as shape description, manipulation, matching and retrieval, registration, recognition, and compression [127], and has shown the application potential in various fields, such as image processing, computer vision and graphics, scientific data analysis, and robotics [40, 127, 175]. The typical skeleton forms are MA/MAT and the curve skeleton. The MAT has better shape representation performance, as there is no information loss in transformation between the MAT and the original shape, while the curve skeleton has information loss but of simpler structure when representing a 3D shape. As a powerful tool in real-world applications, the curve skeleton is a more accepted skeletal representation. However, more efforts need to be made on detailing skeleton quality evaluation and improving the skeleton quality, before the curve skeleton becomes reliable to be used in practice.

2.1.3 Object Skeleton Applications in Robotics Research

Since this thesis investigates object skeletons, emphasising their crucial role in robotics, and proposes a novel approach in skeleton-guided rolling contact kinematics, it is important to situate this new perspective within the broader context of skeleton-based applications in robotics, which motivates the following review.

Skeleton-based representations of objects capture essential topological features, which have led to their adoption across a range of application domains, particularly in robotics. Early work by Cornea et al. [40] highlighted the utility of curve skeletons for virtual navigation, noting that their centredness helps to avoid collisions in constrained environments. In robotic navigation, topology-preserving skeletons have been employed to construct free-space bubble graphs that support efficient path planning [109]. More recent research has extended their applicability to challenging scenarios such as autonomous cave exploration, where environments are typically dark, unstructured, and partially observable [151, 185].

Curve skeletons have also been leveraged for dense surface reconstruction in SLAM frameworks [178], as well as for robotic grasping tasks. Their topological properties make them particularly useful for grasping objects with holes or complex geometries, as demonstrated in the works of Pokorny et al. and Stork et al. [119, 147]. Varava et al. [166] further utilised skeletal structures to identify critical features such as forks and necks, enabling caging grasps. Both curve skeletons and MATs have been incorporated into various grasping planning strategies [120, 121, 165]. Moreover, skeletonisation techniques have shown promise in modelling complex rolling contact phenomena [152].

Beyond these applications, skeleton-based representations have demonstrated benefits in agricultural robotics as well. In particular, they have been used to support semantic understanding and manipulation of plants and fruits in unstructured agricultural environments [186, 132, 75].

2.1.4 Object Skeleton Applications in Other Fields

Apart from the robotics research, the object skeleton has been widely utilised in many other areas, which may include computer vision, shape analysis, medical image analysis, animation, and agricultural phenotyping.

As an efficient shape descriptor, skeletons are naturally applied in shape analysis, including shape recognition, matching, decomposition, reconstruction, segmentation, and compression [40, 127, 155]. For example, a skeleton-based shape model is designed by Trinh and Kimia [162], used for object detection on real images. A graph matching algorithm is developed by Bai and Latecki [12], used as part of a shape recognition algorithm. One of the most well-known skeleton applications in shape recognition/matching is the use in OCR, where skeletons might be used in both handwritten characters and characters on contaminated documents [8, 78, 67]. An approach for matching and retrieving shapes is proposed by Sundar et al. [149] based on skel-

eton information. The object shape can be decomposed into several parts with the information of skeleton branches by the methods introduced by Serio et al. [134]. Tree structures might be automatically reconstructed with guidance by the skeleton information extracted from point clouds of the trees [88]. A shape semantic segmentation example might be given by Au et al. [11], who demonstrated the capability of shape segmentation by skeleton along with thickness information. The application in data compression is mentioned as the skeleton is originally defined by Blum [23]. For a more specific example, one may find skeletons useful in image compression [21].

Another important application area of the skeletonisation is medical image analysis. Firstly, centreline skeletons can be applied in vessel image analysis since the skeletal approaches are quite suitable in analysing multi-branch structures. For example, a series of binary skeleton images can be applied to approximate the artery motions [159]. Blood vessel trees can be extracted from the MR images using skeletonisation methods [55]. The extracted centreline skeletons can also be applied to stenosis assessments, such as the work introduced by Nystrom and Smedby [110], where the centreline skeleton is used to enhance the view of stenoses. Besides, finding a reliable path for virtual endoscopy may also benefit from the extracted skeleton information [63]. Secondly, application in medical image analysis can be a numerical assessment of anatomical morphology. For instance, spicules on mammograms can be detected by identifying characterised structures of the extracted skeletons. Bone strength can be analysed by in vivo imaging with skeleton topology analysis [36].

In fact, the applications of the skeleton are not limited to shape and medical analysis. The skeleton may be used in animation where 3D characters' surfaces are linked to the skeleton, utilising skeletal motion data for animation [14]. In agricultural research, the skeleton information may also be useful in plant phenotyping [69, 99]. As for robotics applications, it will be detailed in the next subsection.

Although the author has found skeleton applications in many previous studies, there might be applications beneficial from the object skeleton's power that are not covered

by this thesis. Thus, it is clear that the object skeleton is a quite critical shape representation worth further study.

2.2 Skeletonisation Methods and Evaluation

While the skeleton is a representation of a shape, the process of generating it is known as skeletonisation. Since Blum introduced the concept of the skeleton through the original medial axis (MA) formulation [23], a substantial body of research has focused on developing skeletonisation methods. However, for many practical applications such as navigation, the exact MA remains overly complex, as it retains an excessive level of detail [40, 34]. Consequently, most prior work has concentrated on approximating the MA rather than deriving it exactly [46, 34].

In this section, we first introduce existing skeletonisation methods, categorised into four groups: boundary-based methods, region field methods, topological thinning methods, and, more recently, learning-based methods, as discussed in Subsections 2.2.1 to 2.2.4, respectively. The highly related approach investigated by the thesis is Laplacian-based topological thinning methods, covered in Subsection 2.2.3. Secondly, we comprehensively review the evaluation approaches for skeletonisation methods in Subsection 2.2.5, including a discussion of desirable skeleton properties and current evaluation approaches.

2.2.1 Boundary-based Methods

Since the skeleton is closely related to the boundary of a shape, early researchers naturally proposed skeletonisation methods that rely directly on boundary information. A small number of approaches aim to compute the exact MA as originally defined, although most focus on approximating the MA from finite samples of the shape boundary, which tends to offer greater stability [46].

Extracting the Exact Medial Axis: The exact MA, defined as the set of centres of maximal inscribed discs or spheres, is rarely computed directly due to its instability and complexity. Hoffmann [64] introduced a MA generation method for constructive solid geometry (CSG), constructing the skeleton via an increasing distance approach. Culver et al. [43] designed an algorithm and data structure for accurately computing the MA of polyhedra by iteratively identifying neighbouring junctions guided by seam curves. Similarly, Sherbrooke [138] proposed an algorithm for 3D polyhedral solids that classifies different types of MA points, allowing for the construction of a graph of skeletal elements. While these methods may compute the exact MA, it is often argued that exact computation is unnecessary in many applications and that the associated instability limits practical utility [111, 40].

Extracting the Median Line: Rather than computing the exact skeleton, many methods focus on approximating the MA, thereby reducing computational complexity and improving applicability [46, 34]. One common approach is to extract a “median” line from the shape contour as an alternative to the exact skeleton. Bookstein [25] used boundary information, including boundary points and tangents, to discretely extract segments of the symmetry axis as a MA approximation. Shapiro [135] extracted skeletons from sequential boundary data for “ribbon-like” shapes. However, these methods have been criticised for their limited applicability to shapes with multiple branches [31]. For discrete point sets, Tagliasacchi [156] proposed the ROSA method, which extracts a complete one-dimensional *curve skeleton*, defined as the rotational symmetry axis, and is applicable even to incomplete point clouds.

Voronoi Diagram-Based Approaches: As illustrated by Fig. 2.3, the Voronoi diagram (VD) partitions a 2D plane or 3D space into regions based on proximity to a given set of objects, such that every point within a region is closer to one particular object than to any other. Each resulting region is referred to as a Voronoi cell. Since the MA is defined based on proximity to the shape boundary, the Voronoi diagram, which encodes such proximity information, serves as a natural and effective geometric tool




Figure 2.3: Voronoi digram (VD) based skeleton [100]. (a) Voronoi diagram; (b) Extracted skeleton; (c) Bounded Voronoi diagram.

for approximating the MA.

This relationship was first observed by Kirkpatrick [76], who proposed a method for constructing the skeleton of a polygon as a subgraph of the Voronoi diagram. Brandt and Algazi [31] introduced a continuous Voronoi-based MA approximation for binary images, with guaranteed accuracy based on a proposed criterion. Ogniewicz and Ilg [111] proposed a skeletonisation method based on computing the Voronoi diagram of boundary points, yielding a simplification-friendly hierarchical structure referred to as the “Voronoi skeleton”.

While skeletonisation methods were initially researched with a preference for 2D objects, studies have gradually extended to 3D objects. Naf et al. [105] studied 3D Voronoi skeletonisation from shapes represented by large point sets, with applications in the feature extraction and identification of organ shapes. Both 2D and 3D Voronoi diagram-based skeletonisation approaches were investigated by Attali and Montanvert [10] for continuous shapes, introducing an efficient skeleton simplification approach and the formation of a union of balls. For regular 3D shapes such as polyhedra, Voronoi diagrams also provide a solution for generating skeletons as the “proximity skeleton”, introduced by Etzion and Rappoport [52] using octree-assisted space subdivision. In the work by Amenta et al. [3], the “power crust” was introduced as a new

construction of 3D skeletons generated from Voronoi diagrams, which is robust to noise and independent of the boundary sampling quality. Extended by Wang et al. [170], a restricted power diagram is utilised for approximating the MAT. To address the instability of Voronoi diagram-based skeletonisation convergence, a sampling scale- and density-independent Voronoi diagram-based MA approximation method with guaranteed convergence was proposed [46]. More recently, Guo et al. [61] applied a restricted Voronoi diagram for “remeshing” the MA of object point sets.

Overall, boundary-based methods start from the MA definition, working with boundary information to obtain a skeleton representation. Typically, boundary-based approaches, whether approximating or computing the exact skeleton, result in skeletons containing many unnecessary details and therefore require refinement before deriving the final skeleton [138, 10, 46], as illustrated in Fig. 2.3. Among these methods, some derive skeletons from continuous shapes, while others work on discrete boundary shapes, such as the extensively researched Voronoi-based skeletons. Since objects described by discrete samples are very common in practice, Voronoi-based algorithms operating on discrete boundary samples tend to be sensitive to the sampling process [3, 43]. Despite advances in boundary-based approaches, a robust skeletonisation method capable of handling discrete object inputs remains in demand.

2.2.2 Region Field Methods

As the shape skeleton is defined as a structure bounded by the shape’s surface, characterising the interior space of a shape to identify likely skeletal elements can be a viable strategy for skeleton extraction. Accordingly, field-based skeletonisation methods have been developed, particularly for digitised object shapes, with distance field-based approaches being among the most popular.

Distance Transform Approach: As illustrated by Fig. 2.4, the distance transform (also known as the distance map or distance field) is a shape descriptor typically com-

puted from digitised shapes, such as 2D images or 3D voxelised models. As such, it is often classified as a digital skeletonisation approach [127], and in some cases even regarded as a form of skeleton itself [161].

Distance transforms are widely employed for thinning and simplifying shapes represented by binary images, and were initially used to generate 1D skeletons of 2D shapes. Due to the prevalence of early studies and the intertwined development of skeleton and distance transform concepts, Toriwaki and Yokoi [161] provided a systematic analysis that unified the related notions. They defined the distance transform as a transformation of a binary shape where each foreground (1-valued) pixel is assigned the distance to the nearest background (0-valued) pixel, assuming binary shapes composed of rectangular pixels. Importantly, distance transforms are generalisable to arbitrary dimensions, as discussed by Borgefors [26], who later proposed implementations in both 3D and 4D [28, 29].

Toriwaki and Mori [160] further explored Euclidean distance transforms as a means of skeletonising 3D binary images, addressing concerns around topological preservation. Pudney [122] introduced a method using chamfer distances to construct distance fields, which then guided the order of point removal in the skeletonisation process. Similarly, Arcelli et al. [7] employed a weighted distance transform with $\langle 3, 4, 5 \rangle$ weights for robust target point detection and topology-preserving point removal.

Although distance transforms are widely recognised and accepted for skeletonisation, their high computational cost poses a significant challenge. Consequently, substantial research has been devoted to improving computational efficiency. Early methods, such as [4], were criticised for being time-consuming due to their iterative nature. In response, Arcelli and Baja [5] proposed a more efficient method that directly identifies pixels with maximum distance values and exploits shape symmetry for accelerated skeleton extraction. Computational efficiency has also been improved by replacing global distance computation with locally optimised distance transformations [27].

Beyond sequential improvements, researchers have also explored parallel implement-




Figure 2.4: Distance field based skeletonisation [100]. (a) Binary shape; (b) Distance field; (c) Extracted skeleton by mean curvature map.

ations of distance transform computation. For example, Lotufo and Zampieroli [93] interpreted the distance transform as a morphological erosion process, which can be computed in parallel with straightforward implementations. Arcelli et al. [6] proposed a parallel two-subiteration algorithm for generating distance transforms of 3D digital objects.

Other Field Based Approach: Apart from distance transforms, skeletonisation may also be achieved using other types of spatial fields. Ahuja and Chuang [1] presented a method that derives a MA skeleton from a generalised potential field, constructed based on shape boundary segments. Siddiqi et al. [140] utilised the average outward flux within a Hamiltonian system to implement a topological thinning process, resulting in a robust and computationally efficient skeletonisation approach. Cornea et al. [41] applied a repulsive force field to extract curve skeletons from various 3D shape representations, including point sets, polygonal meshes, and volumetric objects. Gradient vector flows have also been employed for skeleton extraction; for example, Hassouna and Farag [62] proposed a robust variational framework for 1D curve skeletonisation of volumetric shapes based on such a flow field. Furthermore, Ma et al. [98] approximated 3D medial axes from point set inputs using a normal field combined with neighbourhood information. Their method was shown to be well-suited for parallel computing and efficient in both runtime and memory usage.

In summary, field-based methods focus on analysing interior spatial features of digitised objects to identify medial structures. However, a major challenge for these methods is their considerable computational cost. Moreover, the digitisation of continuous shapes often results in a loss of geometric detail, which may degrade the accuracy of the resulting skeleton.

2.2.3 Topological Thinning Methods

Generally, the topological thinning methods are an iterative process that continuously shrinks the shape till meeting a terminating condition, in which the resultant shape preserves the shape topology and a thin skeleton structure is ensured at the same time.

Thinning by Deletion: One well-known class of thinning methods operates by iteratively deleting boundary elements of a shape, thereby shrinking the shape body until a desirable skeleton is obtained [77]. These methods typically take as input a digitised representation of an object.

The core principle of deletion-based thinning lies in the proper definition of deletion criteria and termination conditions. Arcelli and Gabriella [5] defined deletable elements based on the contours of the shape. Ronse [124] characterised deletable pixels in a figure through a concept known as k -deletability. Lee and Kashyap [80] proposed a criterion for identifying border points suitable for further deletion, based on 6-connected neighbours. Ronse [125] further introduced a set of minimal test patterns to characterise deletable sets while ensuring the connectivity of the resulting skeleton.

A particularly influential concept in defining deletable elements is the notion of a simple point. As described by Gong and Bertrand [59], a point is deemed simple if its removal from the shape does not alter the topological properties of the shape. This theory was later extended to P-simple points, enabling topology-preserving parallel skeletonisation [18]. Building upon this idea, Palagyi and Kuba [114] proposed a 3D “12-subiteration” thinning algorithm, which Lohou and Bertrand [91] extended




Figure 2.5: Deletion-based thinning for binary images [30].

by incorporating P-simple point deletion to ensure topology preservation throughout the thinning process. Additional work on P-simple point-based thinning can be found in [89, 92]. A similar concept, termed simple cell, has also been applied to cell complexes in order to preserve topological structures during thinning [86, 48].

In addition to defining a deletion strategy, similar to the other approach, the computational complexity is another major concern as heavy computation loads are applied to each iteration of deletion [5, 183, 112]. Thus, work on improving the computational cost is extensively conducted. By tracing contours of the shape, it is argued by Arcelli [4] that the thinning computational time can be reduced and the contour analysis is enabled at the same time, followed by a collaborative work [5] presenting a “fast thinning” algorithm acquiring skeleton after cleaning and characterising the input pictures. In the thinning method proposed by Sadleir and Whelan [126], the computational cost of thinning is reduced via lookup tables, applied to computing the centreline skeleton. Boudaoud et al. [30] developed a hybrid method combining a direction approach and a subfield approach to achieve higher computational efficiency for binary image thinning, as illustrated by Fig. 2.5.

Since deletion of the boundary elements is a repeating and simple process, the computation of deletion-based methods can be easily paralleled [163, 19, 97]. Thus, parallel computing may be utilised to speed up the thinning process for acquiring a skeleton with criteria particularly defined for avoiding unexpected deletion led by parallelisa-

tion [163]. In the work by Bertrand [19], the border points are classified for characterised deletion to ensure precise parallel thinning. Saha et al. [130] presented a method using two versions of input images in paralleled thinning with a discussion of its necessity, demonstrating robustness in synthetic noises. A study by Ma and Sonka [96] discusses the sufficient conditions for topology preservation used in their parallel computing with proof. The approach designed by Nemeth et al. [107] embedded smoothing algorithms in thinning iterations to remove unexpected extremities. In order to precisely delete a specific type of boundary points in parallel thinning, one may separate the thinning iterations into multiple subiterations. For example, Ma [95] developed a 6-subiteration thinning method with a discussion on connectivity preservation conditions. For more subiteration thinning algorithm-related work, please refer to [113, 114, 90, 112, 97, 183]. More recently, Wagner [169], leveraging GPU-based parallel computing in the design of a thinning algorithm approaching a real-time thinning speed of both 2D and 3D binary image skeletonisation.

Thinning by Optimisation: Apart from deletion-based thinning methods, iterative thinning of a shape can be achieved by minimising well-defined energy terms. Existing optimisation-based methods are capable of performing skeletonisation on both mesh models and discrete point sets that represent object boundaries. Unlike other approaches, skeletons generated through optimisation are typically not defined in a precise or canonical form. Nevertheless, their principal function is to approximate the MA and provide a skeletal representation of the shape.

A well-known optimisation-based thinning method is Laplacian-based contraction (LBC) [34], which contracts the input shape by minimising a quadratic energy composed of a point attraction term and a contraction term. For mesh inputs, this method contracts the shape to a thin structure, from which a skeleton can be extracted through connectivity processing, as described by Au et al. [11]. Recognising that noisy and incomplete point clouds are more common in real-world scenarios, Cao et al. [34] extended LBC to handle point cloud inputs. Their algorithm incorporates neighbourhood information




Figure 2.6: Laplacian-based point cloud contraction and skeletonization [34]. (a)-(b) illustrate the progressive Laplacian-based contraction of the input horse point cloud; (c) shows the extracted curve skeleton obtained from the contracted (thinned) point cloud; and (d) presents the final refined skeleton.

into the contraction process to better support curve skeletonisation (see Fig.2.6).

Wang et al. [171] argued that existing curve skeletonisation methods often lack robustness and simplicity. To address this, they proposed a more resilient branch-handling technique for skeleton meshes derived via contraction, followed by a “growing” procedure to construct the curve skeleton. Additionally, Yu et al. [187] improved mesh skeletonisation in terms of homotopy preservation and centredness by applying a pre-processing step that closes open surfaces and introduces a weighting scheme based on one-ring area sequences. The mesh-based skeletonisation method of Au et al. [11] was further extended by Tagliasacchi et al. [154], who proposed the mean curvature skeleton, incorporating a medial energy term to better approximate the MA. More recently, Meyer [101] suggested that Cao’s point cloud-based skeletonisation method could be improved by introducing semantic weighting, resulting in a Laplacian-based semantic skeletonisation approach.

In addition to LBC, other optimisation-based thinning algorithms have also contributed to the shape analysis community. For example, Wang et al. [173] employed an iterative least-squares optimisation technique to generate smooth and robust curve skeletons from volumetric inputs. Jiang et al. [71] combined the Voronoi diagram concept with their optimisation framework, resulting in a graph-contraction-based skeletonisation




Figure 2.7: Learning-based direct skeletonisation from images [116].

algorithm. This was later extended by incorporating surface clustering procedures for contraction regularization [72]. Huang et al. [66] proposed the L_1 -medial skeleton, in which a 1D skeletal structure is iteratively derived by locally applying the L_1 -median, a defined centre of a set of points, to point cloud data.

Deletion-based thinning methods are generally developed for digitised objects and tend to require high-performance computing due to their significant computational cost. Optimisation-based approaches, on the other hand, are more suitable for scattered point cloud data and often require fewer iterations to extract skeletons. Despite their advantages, these methods remain computationally intensive, and few studies have explored parallel or accelerated implementations [127, 155]. Improving computational efficiency remains an important direction for the development of optimisation-based thinning algorithms.

2.2.4 Learning-Based Methods

Recently, with the booming development of machine learning, skeletonisation may be done directly via end-to-end neural networks.

Learning-based methods have been increasingly applied to address traditional skeletonisation problems. Sironi et al. [142] proposed an algorithm for detecting multi-scale centrelines by learning a specialised distance transform. This approach reformulates

centreline detection as a regression problem, thereby overcoming challenges in skeletonising irregular 2D and 3D shapes.

To provide benchmark datasets for geometry understanding in learning-based research, Demir et al. [44] introduced the SkelNetOn 2019 challenge. The challenge includes three skeletonisation tasks—from images, from point clouds, and from parametric shape representations—each with a corresponding dataset, with a focus on 2D space only.

Addressing these challenges, Panichev and Voloshyna [116] adopted a modified end-to-end U-Net architecture (see Fig. 2.7) for direct skeleton extraction from images. Jiang et al. [70] employed a fully convolutional network, specifically, a feature hourglass network, to detect skeletons in both images and point clouds. Tang et al. [157] introduced additional modules, including “smooth distance estimation” and “edge transformation”, to enhance skeleton learning beyond standard end-to-end networks, evaluating performance on both point clouds and images. Nathan and Kansal [106] proposed SkeletonNet, featuring a newly designed decoder aimed at reconnecting broken skeleton structures. Lin et al. [83] redefined skeletonisation as a MAT approximation and used a skeletal mesh as the target representation for point cloud inputs. Recognising the importance of implicit fields in shape analysis, Clemot and Digne[39] incorporated them into a neural network to acquire medial complexes from point clouds. Building on this, Petrov et al. [118] proposed a generative model that uses a set of neural implicit functions with local geometric awareness to extract medial abstractions, a collection of medial points.

Beyond traditional tasks, researchers have also explored skeleton extraction from natural colour images. Tsogkas and Kokkinos [164] created a symmetry-axis skeleton dataset and introduced a learning-based symmetry detection method for natural images. To improve computational efficiency in multi-scale feature extraction, Shen et al. [137] proposed a fully convolutional network utilising scale-associated side outputs for skeleton detection. Shen et al. [136] further advanced this line of work by employing a multi-task learning strategy, showing promising results (see Fig. 2.8). Ke



Figure 2.8: Learning-based skeletonisation from natural images [136].

et al. [74] introduced symmetry-axis fitting using residual units, achieving effective multi-scale symmetry representation.

Learning-based methods have gained considerable momentum in recent years, driven by advances in high-performance computing platforms and machine learning research. Nevertheless, most existing approaches concentrate on 2D shape skeletonisation and still face several significant challenges. A major limitation is the substantial need for labelled skeleton datasets, on which many of these methods depend. In addition, they typically require powerful computing resources for deployment, which can be impractical in certain contexts, such as light-weight mobile robotic platforms. For the learning-based skeletonisation methods that do target 3D shapes, limited generalisation remains a further concern. Moreover, explainability continues to be a critical issue for contemporary learning-based approaches [102].

In summary, numerous skeletonisation methods have been developed for shape analysis and various applications. Early research concentrated on accurately extracting the medial axis, while later work aimed to improve computational efficiency through parallel processing and algorithmic optimisation. The curve skeleton, a more practical representation for 3D shapes, has received particular attention. Despite these advances, computational efficiency remains a major limitation. While learning-based skeleton-

isation methods have advanced quickly, their reliance on large labelled datasets, heavy computational demand, and limited generalisation, especially for 3D shapes, raises concerns about their practical robustness. Traditional methods offer stronger stability and interpretability but also have limitations in complex or noisy settings. A hybrid approach may therefore provide a more balanced solution, improving generalisation, explainability, efficiency, and overall performance. Overall, computational efficiency—particularly for curve skeletons—remains the primary challenge at present.

2.2.5 Skeletonisation Evaluation

Since skeletonisation methods have been extensively studied, a unified quantitative evaluation approach is needed for comparing various methods. In this subsection, we begin by briefly introducing widely accepted skeletonisation properties. Additionally, existing skeletonisation evaluation approaches are discussed.

Skeletonisation Properties

In this thesis, we classify the common expected skeletonisation properties into two groups. The first property group is the resultant skeleton quality or accuracy. This group of properties is preferred to be referred to as skeleton quality, as the criteria for a good skeleton may vary in different application scenarios. Another group is the computational performance of the skeletonisation process. The first group is result-oriented, while another is process-oriented.

Firstly, the main goal of skeletonisation methods is to generate high-quality skeletons from input object shapes. Generally, the properties of the produced skeleton of skeletonisation methods can refer to the skeleton properties introduced in subsection 2.1.2, including homotopy equivalence, thinness, centredness, reconstructibility, and smoothness. Additionally, the author would like to refer to robustness for the skeletonisation process [34, 155, 128]. The desired skeletonisation methods are also expected

to be robust to noise in the input shape. Although one can give a general definition of skeleton properties, variation may exist when discussing specific skeletonisation algorithms. For example, there are multiple types of shape input formats handled by the existing skeletonisation algorithms, such as point clouds, mesh, binary 2D/3D images, and continuous shapes, leading to variance in quantitative comparison. These shape input formats may vary in resolution and precision, leading to variance in judgement of skeleton qualities [144]. In addition, the output of the existing skeletonisation algorithms also varies, obstructing unified skeleton comparisons [143]. Besides, the skeleton quality requirement may vary in different scenarios. For curve skeletons applied to navigation, smoothness and exact centredness are critical while reconstructibility is useless [40]. However, as a shape representation method, reconstructibility and centredness are quite vital for shape analysis, while smoothness might be unnecessary [7].

Secondly, it is also important that the skeletonisation methods are computationally efficient, both in speed and memory usage. Dealing with shape skeletonisation may pose a substantial computational burden to computing platforms, especially when dealing with 3D shapes [191, 68, 155]. Thus, the computational performance is also important when discussing skeletonisation properties. Generally, the skeletonisation methods are desired to be computationally efficient both in computational speed and memory usage [143, 155].

Existing Evaluation Approach

As there are multiple definitions of skeletal representations according to the Literature, thus, existing skeleton evaluation studies may vary in skeletonisation inputs and outputs.

skeletonisation for digitised objects is common in the early stage of skeletonisation research, especially the thinning methods. Thus, early skeletonisation evaluations are conducted for those methods. One of the systematic skeletonisation evaluation studies

in the early stage was conducted by Lee and Suen [79], in which 20 skeletonisation algorithms were evaluated with various types of binary image input in terms of reconstructibility, computation speed, similarity extent to the human-labelled “reference skeleton”, connectivity of the skeleton by visual ranking and parallelism. Arcelli and Baja [5] discussed the computation costs of their fast thinning algorithm via a computation complexity analysis. The noise-based robustness evaluation was discussed in the work by Saha and Majumder [130]. The voxel shape skeletonisation methods were evaluated by Sobiecki et al. [143], where the performance is discussed primarily via visual judgement from perspectives such as centredness and smoothness. The data reduction and the removability of the skeletonisation methods may be demonstrated via quantitatively comparing skeletonisation results of multiple voxel image inputs[7].

For mesh input, curve skeletonisation performance is discussed by Telea and Jalba [158] upon computation time and visual judgement. Visual comparison of curve skeletonisation from mesh inputs might also be found in work by Sobiecki et al. [144] and Jalba et al. [68], while the memory and computation time may be analysed as well.

As for point cloud input, existing methods may use an annotated dataset as ground truth for comparison or defined metrics. Work by Dobbs et al. [49] generates synthetic skeleton datasets as a reference with the assistance of learning-based approaches. Meyer et al. [101] compare their algorithm with the traditional Laplacian-based approach via distance to ground truth skeletons. Comparison by ground truth can also be conducted by comparing the shape reconstructed from skeletal representation and the ground truth reconstruction [84]. Quantitative analysis of skeletonisation performance may be conducted by Bucksch and Lindenberg [33] via a defined distance metric to depict the relationship between skeletons and the corresponding point cloud inputs.

For skeletonisation from natural images, a common approach is dominantly by annotating reference skeletons for comparison [74, 137].

The performance of skeletonisation algorithms may be discussed from application perspectives. For digitised objects, thinning performance evaluation may be conducted

with emphasis on optical character recognition (OCR) application, evaluating thinning algorithms by comparing the character recognition rate of generated skeletons. Focused on the application of OCR, a study by Jaisimha et al. [67] tests thinning algorithms with ribbon images contaminated by noises, with performance analysis conducted by statistical methods. Greenspan et al. [60] evaluate two centre-line extraction algorithms from a Quantitative Coronary Angiography (QCA) application perspective by taking one famous algorithm as reference. Evaluation of centreline extraction of medial images may be evaluated with manually annotated ground truth skeletons [131, 73].

Overall, the existing evaluation mainly relies on visual inspection, especially for modern 3D skeletonisation approaches [155]. Ground truth skeleton might be a solution, while annotating a convincing reference skeleton is still challenging, as the definition may vary in different methods. Although several skeletonisation properties are widely accepted, it is still necessary to propose a systematic numerical skeletonisation evaluation approach, which may address specific application scenarios.

2.3 Rolling Contact Kinematics

Rolling contact kinematics play a central role in this thesis, particularly in the development of the proposed skeleton-guided framework. Understanding continuous surface geometry is essential for formulating rolling contact interactions [103]. However, this remains a fundamental challenge because robotic systems typically operate on discrete shape representations, which limits the direct application of continuous kinematic models. In this context, skeleton representations offer a promising bridge between discrete geometry and continuous surface behaviour: by reducing 3D shape analysis to a family of local 2D cross-sections intrinsically linked to the underlying skeleton [34], they can simplify geometric reasoning and support the integration of continuous kinematic equations.

To provide a comprehensive background for the proposed method, this section is organised into two parts. First, we review the key concepts and application domains in which rolling contact principles have been effectively employed. Then, we introduce the theoretical foundations of rolling contact kinematics, which underpin the models and methods developed in this thesis.

2.3.1 Application Background

The rolling contact motion refers to the relative movement between rigid bodies that remain in continuous contact during motion [103]. This form of motion is intrinsically omnidirectional, offering high flexibility and versatility. Consequently, rolling contact has inspired various robotic designs that mirror principles observed in natural locomotion and have been successfully applied across multiple robotic domains [152]. One of the earliest and most prominent application areas is robotic locomotion. Spherical robots, in particular, exploit their symmetric geometries and rolling-based mobility to enable autonomous inspection and exploration in unstructured or constrained environments [87, 47, 152]. Numerous spherical robot platforms have been developed in this context. For example, ARIES employs cylindrical joints to achieve controllable rolling trajectories (see Fig. 2.9(a)) [16], while Sugiyama and Hirai [148] demonstrated that a deformable soft-bodied robot can crawl and jump, facilitating robust locomotion over rough terrain. Another notable design, Omnicron, integrates three independently controlled omnidirectional wheels within a spherical shell to realise fully omnidirectional movement [38].

Beyond locomotion, rolling contact principles have also influenced robotic manipulation. Notable examples include the rolling contact juggling framework introduced by Woodruff and Lynch [177], and a three-fingered grasper capable of executing smooth in-hand manipulation through controlled rolling (see Fig. 2.9(b)) [189]. Specian et al. [146] further demonstrated the use of edge-rolling mechanisms to transport cylindrical objects. Rolling-based joint designs have been noted for offering several ad-




Figure 2.9: Example of rolling-contact kinematics application background. (a) An example spherical rolling robot [16]. (b) An example of a rolling-based manipulator [188].

vantages, including improved abrasion resistance, simplified control, and greater configuration flexibility [192]. Given the prevalence and geometric convenience of rolling motion, accurate modelling is essential for enabling effective motion planning and control [50, 13, 145, 153].

2.3.2 Theoretical Research

The theoretical foundations of rolling contact kinematics provide a formal description of interactions between rolling bodies, typically governed by their local curvature and relative orientation [104, 42, 20]. Montana’s foundational model [103] introduces multiple coordinate frames at the contact point to represent the relative motion between rigid bodies. This underactuated framework, often characterised by five degrees of freedom, two per body and one for their relative spin, has been widely utilised in continuous-domain motion planning [152, 150, 153]. It serves as a basis for the design and control of rolling systems such as disks, wheels, and omnidirectional robots [152]. However, despite the theoretical maturity of continuous surface contact models, a disconnect remains between these models and the discrete nature of sensory inputs used in practice. Surface geometry captured through perception systems such as point clouds

is inherently discrete, presenting a challenge for integrating continuous rolling contact models into real-world planning and control pipelines.

2.4 Summary

In this chapter, the research background of this thesis was comprehensively introduced, including the notion of object skeletons, skeletonisation methods and their evaluation, and the application background of object skeletons. The insights and research gaps derived from the discussed literature are as follows.

Firstly, the object skeleton is a useful shape representation that efficiently abstracts the topological structures of the shapes. And the applications of the object skeletons vary in many areas, including robotics, demonstrating the great application power of the skeletons. However, the existing definition varies in different studies. Thus, defining properties of a “good skeleton” is not easy, in spite of some commonly accepted characteristics.

Secondly, the skeletonisation algorithm has been extensively researched for decades, from various perspectives. The common research objectives for the skeletonisation algorithms are improving the computational efficiency and the quality of the generated skeletons. In spite of the efforts in previous studies, computational load is a remaining challenge for the skeletonisation methods for some types of defined skeletons, such as curve skeletons derived from point clouds. As for skeleton quality improvement, a good definition of desirable skeletons may be helpful in skeletonisation method comparison and evaluating skeleton quality improvements.

Thirdly, rolling motion is widespread in natural environments, making accurate modelling of rolling contact kinematics particularly valuable for robotics research. Although several models have been proposed, a gap remains between the real-time, discrete shape representations obtained from perception and the continuous models currently

used in analysis. The curve skeleton, as a topological abstraction of shape, may serve as a promising bridge to connect these two domains.

The work conducted in this thesis addresses some of these gaps. First, aiming for more computationally efficient skeletonisation, a new skeletonisation framework is proposed by diagnosing the computational redundancy of the existing Laplacian-based skeletonisation method [34]. Secondly, a numerical skeleton quality evaluation study is conducted from a robotic application perspective, addressing a more solid skeletonisation quality comparison. Thirdly, the object skeleton application in rolling contact kinematics modelling is investigated, aiming to bridge the gap between discrete shape representations and the traditional continuous kinematics model.

GLSkeleton: A Geometric Laplacian-based Skeletonisation Framework for Object Point Clouds

3.1 Introduction

Improving computational efficiency is a key challenge in 3D point cloud skeletonisation due to the high processing demands. This chapter addresses this issue by identifying and reducing computational redundancies in the skeletonisation process.

The Laplacian-based contraction (LBC) method for curve skeletonisation, originally developed for mesh skeletonisation [11], was later extended to point clouds by Cao et al. [34]. This method is robust to incomplete data and well-suited for robotic applications such as grasping planning, as shown in Fig.3.1 [119, 166, 165].

However, the original LBC approach processes all points in each iteration, leading to high computational cost. To improve efficiency, we propose a geometric Laplacian-based skeletonisation framework (GLSkeleton) for object point clouds, a modified framework that selectively retains feature-rich points using a local point reduction (LPR) strategy, reducing computation time by about half when the input cloud point number is 4096 according to the experimental results, and this reduction rate will be greater when the input cloud point number increases. In this framework, a global

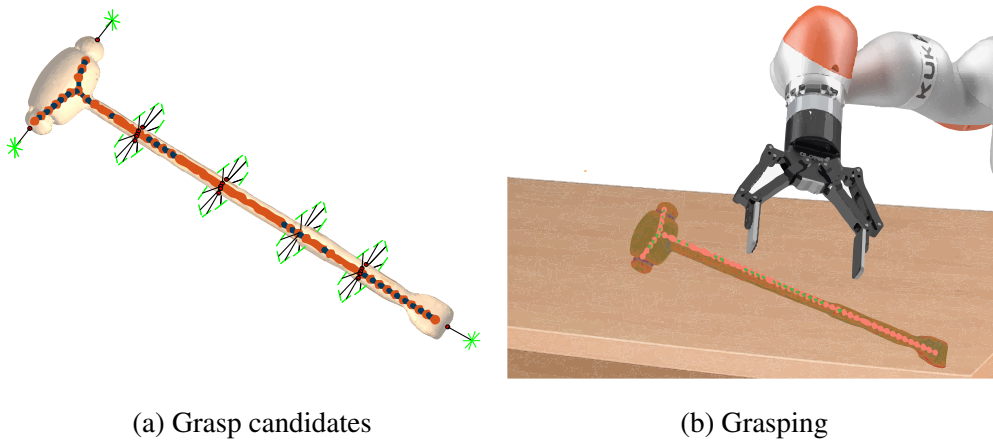


Figure 3.1: Skeleton-assisted grasp planning. (a) The grasp candidate generated from the object skeleton for a two-finger gripper. It is represented as a straight line corresponding to the gripper opening, with short perpendicular markers at both ends indicating the gripper’s orientation. (b) The robot executing a grasp using the two-finger gripper.

feature-based termination criterion is also introduced for more reliable convergence.

In the experiments, GLSkeleton is evaluated against the baseline (Laplacian-based point cloud skeletonisation [34]) using real-world, fully-scanned data from OmniObject3D [181], demonstrating significant improvements in speed with comparable skeleton quality. Besides, a toolbox for the implementation of GLSkeleton is open-sourced for the community¹.

This chapter is organised as follows. The GLSkeleton framework is introduced in Section 3.2, where the proposed local point reduction (LPR) and global feature-based contraction terminating condition are explained. Then the skeletonisation experiments on real-scanned object point clouds with comparison to the baseline methods are presented in Section 3.3. Finally, the work of this chapter is concluded in Section 3.4.

¹<https://github.com/weiqimeng1/GLSkeleton>

3.2 GLSkeleton Framework

The GLSkeleton framework provides an efficient method for extracting the curve skeleton of a shape whose surface is represented by a point cloud. Therefore, before introducing the framework, it is essential to clarify the notions of a point cloud and a curve skeleton. Their formal definitions are provided in Definition 3.1 and Definition 3.2, respectively.

Definition 3.1. A point cloud \mathbf{P} of a shape Ω is defined as a set of points such that $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n]^T$, where each point $\mathbf{p}_i = [p_{ix}, p_{iy}, p_{iz}]^T \in \mathbb{R}^3$.

Definition 3.2. A curve skeleton G_s corresponding to the shape Ω is represented as a graph (V, E) , where $V \subseteq \mathbb{R}^3$ is the set of skeletal vertices and $E \subseteq V \times V$ is the set of edges. Each element $e_i \in E$ or vertex $v_i \in V$ is expected to be bounded by $\partial\Omega$.

With these definitions established, we now outline the GLSkeleton framework. GLSkeleton is a novel geometric skeletonisation framework based on Laplacian contraction, derived from the classical Laplacian smoothing algorithm. The framework consists of four main components, as illustrated in Fig. 3.2. Let \mathbf{P}_o denote the input object point cloud, as defined in Definition 3.1. Prior to the contraction process, the local neighbouring point ring (see Fig. 3.3(b)) of \mathbf{P}_o is computed. The curve skeleton G_s (see Definition 3.2) is subsequently extracted through topological connection and refinement.

It is noteworthy that the computational cost of processing 3D point clouds scales linearly with the number of points. However, many points may become unnecessary or redundant during the iterative contraction. Therefore, reducing the number of points during the contraction process can effectively decrease the computational load of the Laplacian-based skeletonisation. Building upon this observation, we introduce a local point reduction (LPR) procedure within each contraction iteration to remove superfluous points from the local rings, as detailed in Section 3.2.2. Furthermore, given

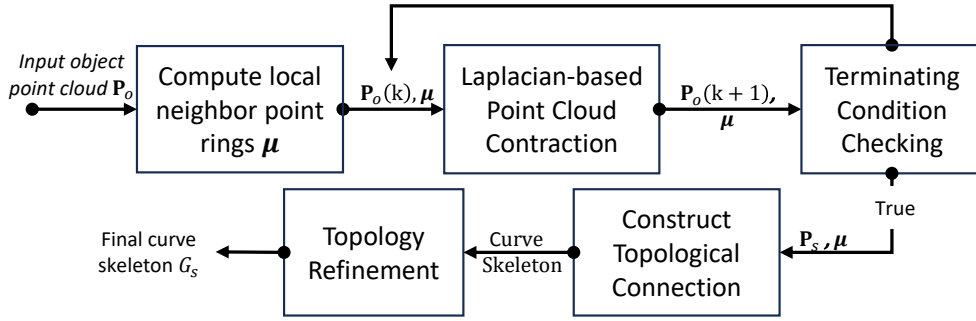


Figure 3.2: Block diagram of the GLSkeleton framework


the progressively reduced point set, a more robust contraction-termination condition is proposed to ensure reliable convergence of the contraction loop.

3.2.1 Overview of Laplacian-based Skeletonisation

Before introducing the details of our proposed framework, we first explain the general concept of Laplacian-based point cloud skeletonisation. Originally developed for mesh data by Au et al. [11], this method was later adapted for point clouds by Cao et al. [34], who introduced a *local point ring* strategy to establish neighbourhood relationships in unstructured point cloud data.

For a point cloud shape P as defined by Definition 3.1. The local ring for an anchor point p_i is constructed by first identifying its k nearest neighbours. These neighbours are projected onto the estimated tangent plane at p_i , and a 2D Delaunay triangulation is then applied to the projections. This process produces a circular (clockwise) ordering of the neighbour indices, denoted as $\mu_i \subseteq 1, 2, \dots, n$, which defines the local point ring. This ordering plays a critical role in preserving local geometric structure during skeletonisation.

The core of the Laplacian-based skeletonisation process is an optimisation framework that balances two competing forces: contraction and attraction. The contraction term



(a) Projected plane

(b) Delaunay triangulation

Figure 3.3: Each anchor point (in red) and its neighbours (in yellow) are projected onto a local tangent plane (a). The neighbour point ring is inferred from the Delaunay triangulation of the projected points (b), although it is not explicitly visualised.

encourages points to move inward toward a medial representation, while the attraction term resists excessive displacement from the original shape.

To this end, cotangent-weighted Laplacian coordinates are computed as an approximation of local curvature flow. For a point cloud \mathbf{P} , the Laplacian coordinates $\boldsymbol{\delta} = [\delta_1, \delta_2, \dots, \delta_n]^T$ are given by:

$$\boldsymbol{\delta} = \mathbf{L} \mathbf{P}, \quad (3.1)$$

where $\mathbf{L} = [\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n]^T$ denotes the Laplacian matrix defined by cotangent weights. Each row $\mathbf{l}_i = [l_{i1}, l_{i2}, \dots, l_{in}]^T$ is given by:

$$L_{ij} = \begin{cases} \omega_{ij} = \cot \alpha_{ij} + \cot \beta_{ij}, & \text{if } j \in \boldsymbol{\mu}_i; \\ -\sum_{k \in \boldsymbol{\mu}_i} \omega_{ik}, & \text{if } i = j; \\ 0, & \text{otherwise,} \end{cases} \quad (3.2)$$

where α_{ij} and β_{ij} are the angles opposite to edge (i, j) in the local triangulation (see Fig. 3.3(b)), and $\boldsymbol{\mu}_i$ denotes the neighbour ring index set for point \mathbf{p}_i .

To moderate the contraction and prevent excessive structural distortion, an attraction term is incorporated by penalising the deviation of each point from its previous posi-

tion. The full contraction process is defined as solving:

$$\begin{bmatrix} \mathbf{W}_L \mathbf{L} \\ \mathbf{W}_H \end{bmatrix} \mathbf{P}_s(k+1) = \begin{bmatrix} \mathbf{0} \\ \mathbf{W}_H \mathbf{P}_s(k) \end{bmatrix}, \quad (3.3)$$

where $\mathbf{P}_s(k)$ and $\mathbf{P}_s(k+1)$ denote the point cloud before and after the k -th contraction iteration, respectively. The diagonal matrices \mathbf{W}_L and \mathbf{W}_H control the influence of contraction and attraction forces.

Although the contraction process can be described by the general formulation in Eq. (3.3), it cannot be solved directly. Instead, it is approached iteratively by minimising the following quadratic energy functional [11]:

$$\|\mathbf{W}_L \mathbf{L} \mathbf{P}_s(k)\|^2 + \sum_i \mathbf{W}_{H,i}^2 \|\mathbf{p}_{s,i}(k+1) - \mathbf{p}_{s,i}(k)\|^2, \quad (3.4)$$

where the first term promotes local geometric contraction, and the second term penalises large displacements to ensure smooth transitions across iterations.

The contraction proceeds iteratively. Initially, the weight matrices \mathbf{W}_L and \mathbf{W}_H are set following the recommendations in [34]. After each iteration, these weights are updated according to the following rule:

$$\begin{bmatrix} \mathbf{W}_L(k+1) \\ \mathbf{W}_H(k+1) \end{bmatrix} = \begin{bmatrix} s_L \mathbf{W}_L(k) \\ \mathbf{W}_H(k) (\mathbf{S}(0)/\mathbf{S}(k)) \end{bmatrix}, \quad (3.5)$$

where s_L is a decay scalar, and the vectors $\mathbf{S}(0)$ and $\mathbf{S}(k)$ represent the average distances from each point to its neighbours in the original and current point clouds, respectively.

Focus on our proposed GLSkeleton framework, illustrated in Fig. 3.2, we incorporate a LPR strategy—described in Section 3.2.2—directly into the contraction loop. This strategy dynamically removes redundant points during optimisation, significantly reducing computational cost.

To improve the reliability of termination, we introduce a global feature-based criterion that robustly determines when the contraction has converged. This termination condition is detailed in Section 3.2.3.

It is important to note that the Laplacian contraction step adjusts only the point positions and does not establish topological connectivity. Therefore, a post-processing stage is required to infer and refine the skeletal graph from the contracted point cloud.

3.2.2 Geometric Local Point Reduction Strategy

To accelerate the Laplacian computation, we propose a geometric strategy termed LPR. As illustrated in Algorithm 3.1 and Fig. 3.2, after each contraction cycle, the resulting point cloud \mathbf{P}_s undergoes LPR. This process starts by evaluating and removing points based on their point-wise contraction stage function values. Following removal, the disrupted neighbour rings μ are reconstructed via a ring combination mechanism. Meanwhile, the original point cloud \mathbf{P}_o is updated to preserve its one-to-one correspondence with the contracted point cloud.

In LBC for skeletonisation, computational bottlenecks primarily arise from the size of the point cloud. As contraction progresses, many points become redundant due to local densification and increasing structural abstraction. Retaining all points throughout the process is therefore unnecessary. LPR addresses this by removing points deemed redundant after each iteration, thus improving efficiency.

The core idea behind LPR is to relieve local constraints that limit contraction speed by selectively removing points within the ring neighbourhood. From Eqs. (3.4)-(3.5), the contraction weights \mathbf{W}_L and attraction weights \mathbf{W}_H constrain the contraction dynamics. These weights must be updated after each iteration, as the system reaches a local equilibrium where the Laplacian term becomes minimal. In Eq. (3.4), the contraction energy term diminishes while the attraction energy increases during contraction. For convex point sets, this equilibrium satisfies

$$\|\mathbf{W}_L \mathbf{L} \mathbf{P}_s(k)\|^2 = \sum_i \mathbf{W}_{H,i}^2 \|\mathbf{p}_{s,i}(k+1) - \mathbf{p}_{s,i}(k)\|^2 \quad (3.6)$$

To measure progress towards this balance, we introduce a point-wise stage function

Algorithm 3.1 The computation of LPR of GLSkeleton.

```

1: function LPR( $\mathbf{P}_o(k)$ ,  $\mathbf{P}_s(k)$ ,  $\boldsymbol{\mu}(k)$ )
2:   for each point  $\mathbf{p}_{s,i}$  in  $\mathbf{P}_s(k)$  and  $\boldsymbol{\mu}_i$  in  $\boldsymbol{\mu}(k)$  do
3:     Compute stage function value  $F(\mathbf{p}_{s,i})$  by Eq. (3.10)
4:     if the stage function value meets the inequality (3.11) then
5:        $m \leftarrow \underset{j \in \{1, \dots, k_i\}}{\operatorname{argmin}} \|\mathbf{p}_{s,i} - \mathbf{p}_{c,\boldsymbol{\mu}_i(j)}\|$ 
6:       Remove  $\mathbf{p}_{\boldsymbol{\mu}_i(m)}$  and  $\mathbf{p}_{c,\boldsymbol{\mu}_i(m)}$  from point clouds  $\mathbf{P}_o(k)$  and  $\mathbf{P}_s(k)$  re-
       spectively
7:       Compute new  $\boldsymbol{\mu}_i$  by combining rings  $\boldsymbol{\mu}_i$  and  $\boldsymbol{\mu}_m$  with estimated surface
       normals (3.1)
8:       Remove ring  $\boldsymbol{\mu}_m$  from  $\boldsymbol{\mu}(k)$ 
9:       Update all rings in  $\boldsymbol{\mu}(k)$ 
10:    end if
11:  end for
12:   $\{\mathbf{P}_o(k+1), \mathbf{P}_s(k+1), \boldsymbol{\mu}(k+1)\} \leftarrow \{\mathbf{P}_o(k), \mathbf{P}_s(k), \boldsymbol{\mu}(k)\}$ 
13:  return  $\mathbf{P}_o(k+1), \mathbf{P}_s(k+1), \boldsymbol{\mu}(k+1)$ 
14: end function

```

based on the quadratic contraction and attraction energies:

$$\begin{aligned}
V(\mathbf{p}_{s,i}, k+1) &= w_l^2(k+1) \|\boldsymbol{\delta}_i(k+1)\|^2 - w_{h,i}(k+1)^2 \|\Delta \mathbf{p}_{s,i}(k+1)\|^2 \\
&= s_L^{k+2} w_l^2(0) \|\boldsymbol{\delta}_i(k+1)\|^2 - \left(\frac{s_i(0)}{s_i(k+1)} w_{h,0} \right)^2 \|\Delta \mathbf{p}_{s,i}(k+1)\|^2,
\end{aligned} \tag{3.7}$$

where $\Delta \mathbf{p}_{s,i}(k+1) = \mathbf{p}_{s,i}(k+1) - \mathbf{p}_{s,i}(k)$, and $w_{h,i}(k)$ and $w_l(k)$ are point-wise weights derived from $\mathbf{W}_H(k)$ and $\mathbf{W}_L(k)$, respectively.

This function evaluates whether a point has reached the contraction equilibrium in the next iteration. A value of zero implies equilibrium; a positive value indicates over-contraction, while a negative value indicates under-contraction. Let $\psi \in [0, 1]$ denote the average ring size reduction rate in the k -th iteration. Assuming a predicted reduc-

tion $\psi' = \varepsilon\psi$, we write:

$$s_i(k+1) = s_i(k) - \psi' s_i(0), \quad (3.8)$$

and derive:

$$\begin{aligned} \delta_i(k+1) &= \frac{s_i(k+1)}{s_i(0)} \delta_i(k) = \frac{s_i(k) - \psi' s_i(0)}{s_i(0)} \delta_i(k), \\ \Delta \mathbf{p}_{s,i}(k+1) &= [\mathbf{p}_{s,i}(0) - \mathbf{p}_{s,i}(k)] - \frac{s_i(0) - s_i(k+1)}{s_i(0) - s_i(k)} \cdot [\mathbf{p}_{s,i}(0) - \mathbf{p}_{s,i}(k)] \\ &= \frac{\psi' s_i(0)}{s_i(0) - s_i(k)} (\mathbf{p}_{s,i}(k) - \mathbf{p}_{s,i}(0)). \end{aligned} \quad (3.9)$$

Using the above relations and including a curvature preservation term $U(k)$ to protect sharp features (i.e., significant local curvature deviations), the final point evaluation function becomes:

$$\begin{aligned} F(\mathbf{p}_{s,i}(k+1)) &= \omega_1 V_1(k+1) + \omega_2 V_2(k+1) + \omega_3 U(k) \\ V_1(k+1) &= s_L^{k+2} w_l^2(0) \|\eta \delta_i(k)\|^2 \\ V_2(k+1) &= \left(\frac{w_{h,i}(k)}{\eta} \right)^2 \left\| \frac{\psi' s_i(0)}{s_i(0) - s_i(k)} (\mathbf{p}_{s,i}(k) - \mathbf{p}_{s,i}(0)) \right\|^2 \\ U(k) &= \|\delta_i(k) - \frac{1}{n_i} \sum_{j \in \mu_i} \delta_j(k)\|^2 \end{aligned} \quad (3.10)$$

where $\eta = [s_i(k) - \psi' s_i(0)] / s_i(0)$, and $\delta_i(k)$ denotes the Laplacian coordinates, serving as an estimate of $k_{n,i} \cdot \mathbf{n}$. The weights are empirically set to $\omega_1 = 1$, $\omega_2 = -1$, and $\omega_3 = 3$, ensuring stable convergence of contraction (as required by Eq. 3.6) and retention of sharp features.

As summarised in Algorithm 3.1, after each contraction iteration, the point $\mathbf{p}_{s,i}$ and its closest neighbours are removed if:

$$F(\mathbf{p}_{s,i}, k+1) < \phi s_i^2(k) \quad (3.11)$$

where ϕ is a user-defined scaling factor that governs the reduction aggressiveness.

3.2.3 Terminating Condition for Contraction Iterations

In addition to the computational cost, defining an appropriate termination criterion for the contraction iterations becomes more challenging when point reduction is applied. In the method proposed by Cao et al. [34], the contraction process is terminated once the change in ring size—defined as the minimum distance between each point and its neighbouring ring points—falls below a predefined threshold. However, this metric becomes unreliable under point reduction, as ring sizes are directly affected by the removal of points. Therefore, in this work, we introduce a termination strategy based on the proposed explained score (Eq. (3.13)). This score reflects the overall progress of the contraction process and remains robust against decreases in the distance-defined ring size caused by point number reduction.

We employ principal component analysis (PCA) to compute the explained score of the contracted point cloud since it is capable of analysing the data globally. For a point set $\mathbf{X} = \{\mathbf{x}_k\}$, where $\mathbf{x}_k \in \mathbb{R}^3$, let \mathbf{m} denote its centroid. The 3×3 covariance matrix is calculated as

$$\mathbf{C}_{\mathbf{X}} = \sum_k (\mathbf{x}_k - \mathbf{m}) \otimes (\mathbf{x}_k - \mathbf{m}), \quad (3.12)$$

where \otimes denotes the outer product.

Considering the entire contracted model $\mathbf{P}_s = \{\mathbf{p}_{s,i}\}$ as a single point set, we compute the global covariance matrix $\mathbf{C}_{\mathbf{P}_s}$. The explained score \mathbf{e} , representing the percentage of the whole shape information captured on projected planes, is then defined as:

$$\mathbf{e} = \begin{bmatrix} e_1 & e_2 & e_3 \end{bmatrix}^T = \frac{1}{2(\lambda_1 + \lambda_2 + \lambda_3)} \begin{bmatrix} \lambda_1 + \lambda_2 & \lambda_1 + \lambda_3 & \lambda_2 + \lambda_3 \end{bmatrix}^T, \quad (3.13)$$

where $\lambda_1, \lambda_2, \lambda_3$ are the eigenvalues of $\mathbf{C}\mathbf{V}_{\mathbf{P}_s}$, sorted in descending order. In the three projected planes, we assume that the shape information is gradually abstracted into two of them. Building on this, to quantify the overall contraction stage, we compute the mean of the two components with the highest explained score, given by

$$e = \frac{1}{2}(e_1 + e_2) = \frac{2\lambda_1 + \lambda_2 + \lambda_3}{2(\lambda_1 + \lambda_2 + \lambda_3)}, \quad (3.14)$$

where $e \in [0, 1]$. As the contraction progresses, this score generally increases.

To determine the convergence of the contraction process, we evaluate the difference in explained scores between iterations:

$$\Delta e(k) = \frac{|e(k+1) - e(k)|}{1 - (n(k) - n(k+1))/n(0)}, \quad (3.15)$$

where $n(k)$ and $n(k+1)$ are the numbers of points at iterations k and $k+1$, respectively. The denominator accounts for point reduction effects, ensuring the stability of the score change measurement. When $\Delta e(k)$ falls below a predefined threshold $\theta \in [0, 1]$, the contraction loop is terminated. In practice, θ is set to a small value close to zero, under the assumption that the contraction process exhibits no large variance once convergence is reached. The threshold is determined empirically from experimental observations and is fixed across all objects; it is selected as the maximum value that still ensures convergence for most models.

3.3 Results and Discussion

In this section, we evaluate the performance of GLSkeleton using the OmniObject3D dataset [181]. This dataset comprises real-scanned 3D objects from a wide range of categories, characterised by substantial variation in shape and appearance. The point clouds are often noisy and unevenly distributed, which reflects the challenges of real-world data. One limitation of using this dataset is that it is collected offline rather than on a real robotic system. All experiments in this section were conducted in MATLAB on a PC with an Intel(R) Core(TM) i7-3770K CPU.

3.3.1 Contraction Speed Comparison

To assess contraction speed, we compare GLSkeleton with the baseline method proposed by Cao et al. [34]. The primary enhancement introduced by GLSkeleton is the

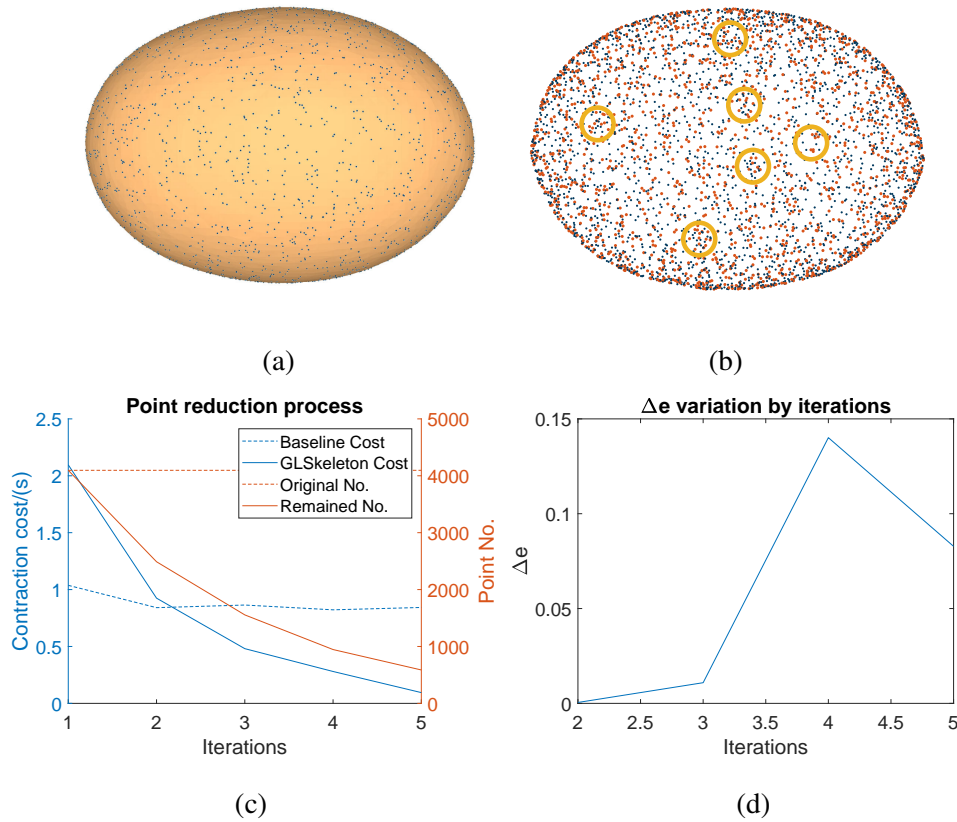


Figure 3.4: Illustration of LPR. (a) Original egg model. (b) Result after one contraction iteration; yellow circles indicate improved regions, red points are removed, and blue points remain. (c) Reduction in point number and contraction cost. (d) Evolution of Δe over iterations.

Table 3.1: Parameter settings for GLSkeleton

Parameter	Value	Description
ε	1.2	ψ'/ψ
ϕ	1×10^5	Parameter in Eq. (3.11)
θ	0.02	Termination threshold of Δe in Eq. (3.15)

in-loop LPR, which accelerates contraction iterations. Values of the constant parameters and thresholds are summarised in Table 3.1.

A total of 216 point cloud models, each sampled at three different resolutions, were selected from the dataset without bias towards object type. Throughout the contraction

process, redundant points that hinder progress were effectively removed, as shown in Fig. 3.4(a). Given that the contraction system in Eq. (3.4) has linear time complexity $O(n)$, the LPR provides a proportional speed-up, verified in Fig. 3.4(c). Furthermore, the proposed explained-score difference Δe increases during the initial iterations and subsequently decreases (Fig. 3.4(d)), reflecting its alignment with the global contraction rate. This behaviour is consistent across the objects we evaluated and is not restricted to the egg object alone. These observations confirm that Δe serves as a reliable indicator for contraction convergence and termination.

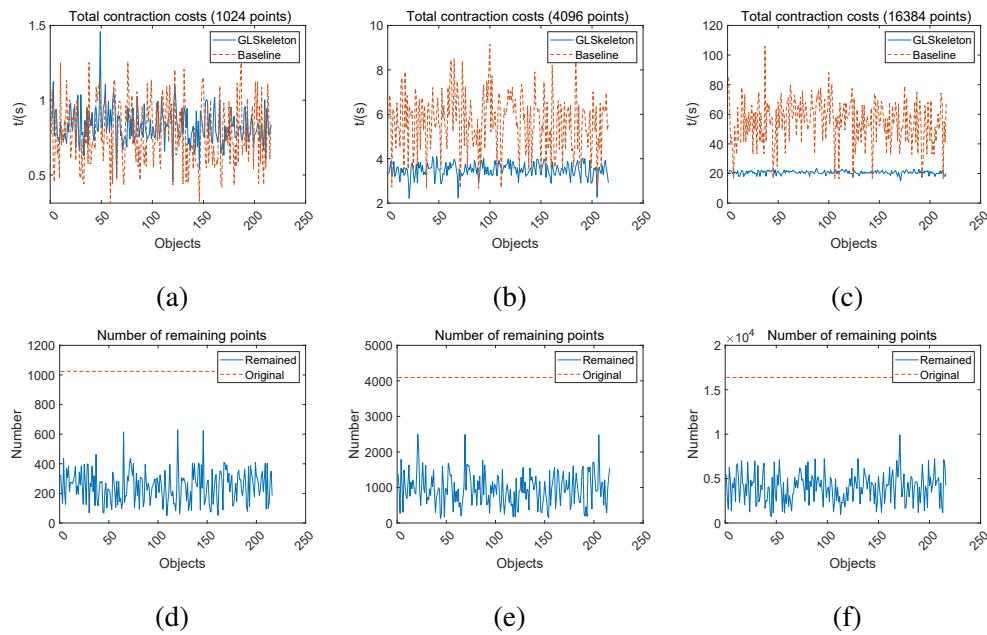


Figure 3.5: Contraction time comparison between GLSkeleton (blue) and the baseline (red). Top: Contraction times for 216 models with 1024, 4096, and 16384 points. Bottom: Remaining points after contraction.

As shown in Fig. 3.5, GLSkeleton performs similarly to the baseline on low-resolution models (1024 points). This is because solving the point cloud reduction also introduces additional computational cost, and this cost is relatively high compared with sparse point cloud contraction. However, as model size increases, GLSkeleton increasingly outperforms the baseline due to significant cloud point number reduction. For models with 4096 and 16384 points, the contraction time is significantly reduced. While

differences are minimal for smaller models (Fig. 3.5(a)), GLSkeleton exhibits lower variance in runtime across different objects (standard deviations of 1.2996 vs. 0.3205), indicating better controllability of contraction speed.

3.3.2 Skeletonisation Results

To assess skeletonisation quality, we selected 77 models (4096 points) from the full set of 216 objects, choosing those that exhibit representative 3D structures. For many of the remaining objects, valid curve skeletons could not be further extracted, either because the Laplacian-based contraction produced unsatisfactory point cloud results or because, for shapes such as spheres or planar surfaces, the expected curve skeleton is inherently ambiguous. Consequently, only the subset of objects with meaningful skeletons was retained for comparison.

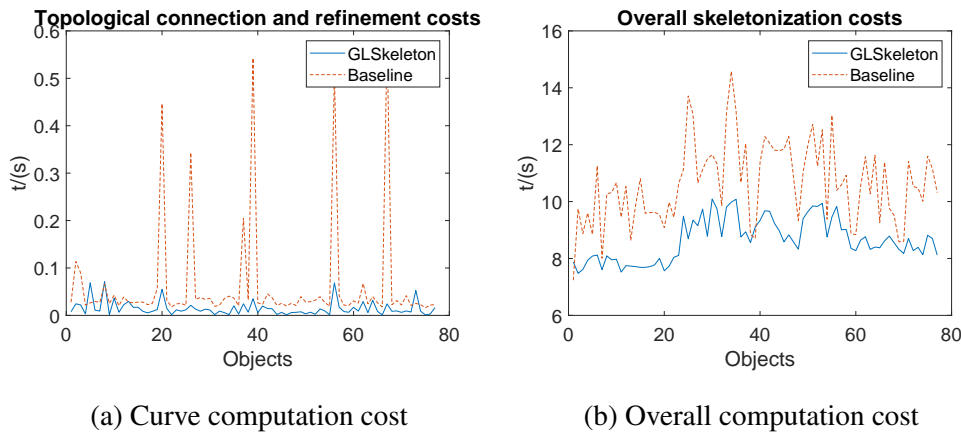


Figure 3.6: Curve computation cost from contracted point cloud (a) and overall computation cost (b).

LPR also accelerates the curve construction stage (Fig. 3.6(a)), as the farthest-point sampling relies on kd-tree range queries with time complexity $O(n^{2/3})$. Therefore, skeletonisation via GLSkeleton is faster than the baseline (Fig. 3.6(b)). It is worth noting that the smaller performance gain here (compared to Fig. 3.5(b)) is due to both methods sharing the same initial neighbour ring computation stage (see Fig. 3.2).

As seen in Fig. 3.7, GLSkeleton successfully preserves topological features in most cases, including models with missing data (e.g., Figs. 3.7(e), (k)). However, removing points during contraction may result in minor geometric distortion or the appearance of non-medial branches, especially in sparse regions with joints or loops (Figs. 3.7(g), (h), (i), (l)). These effects are generally negligible in practical applications. For example, small changes in the skeleton used for generating grasp candidates have little impact unless they alter the local topological direction or the local centeredness.

Fig. 3.8 presents results under varying levels of Gaussian noise. GLSkeleton maintains robustness even with severe noise, occasionally outperforming the baseline, while requiring less computational effort.

The skeletons produced by GLSkeleton provide meaningful vertex structures with clear topological connections, supporting applications such as grasping planning, which may use the topological information to propose grasp candidates [119, 165]. They are also suitable for segmentation and curvature estimation, aiding local geometry analysis. However, real-time use may be hindered by non-medial branches and challenges in segmenting strictly convex shapes like spheres.

3.4 Conclusion

This chapter presents GLSkeleton, a geometric Laplacian-based framework for point cloud skeletonisation that achieves comparable results to existing methods with significantly reduced computational cost. The proposed in-loop local point reduction strategy identifies and removes redundant points that impede contraction, enhancing computational efficiency. Additionally, a robust contraction-convergence indicator, independent of local point density, is introduced to guide termination during iteration.

Extensive evaluation using real-scanned data demonstrates the method’s effectiveness and robustness. However, challenges remain, including the intrinsic limitations of

Laplacian-based methods, such as failing in topology preservation when handling specific structures and the sensitivity to the constant update factor s_L in Eq. (3.5) [34]. Furthermore, the lack of standardised evaluation metrics for skeletonisation limits its practical deployment.

Future work will focus on improving skeleton robustness, publishing a curated dataset, and exploring motion-planning-based skeleton inference instead of purely topological approaches. We will also investigate grasping planning using GLSkeleton, including the relationship between skeleton quality, contact points, and local surface geometry.

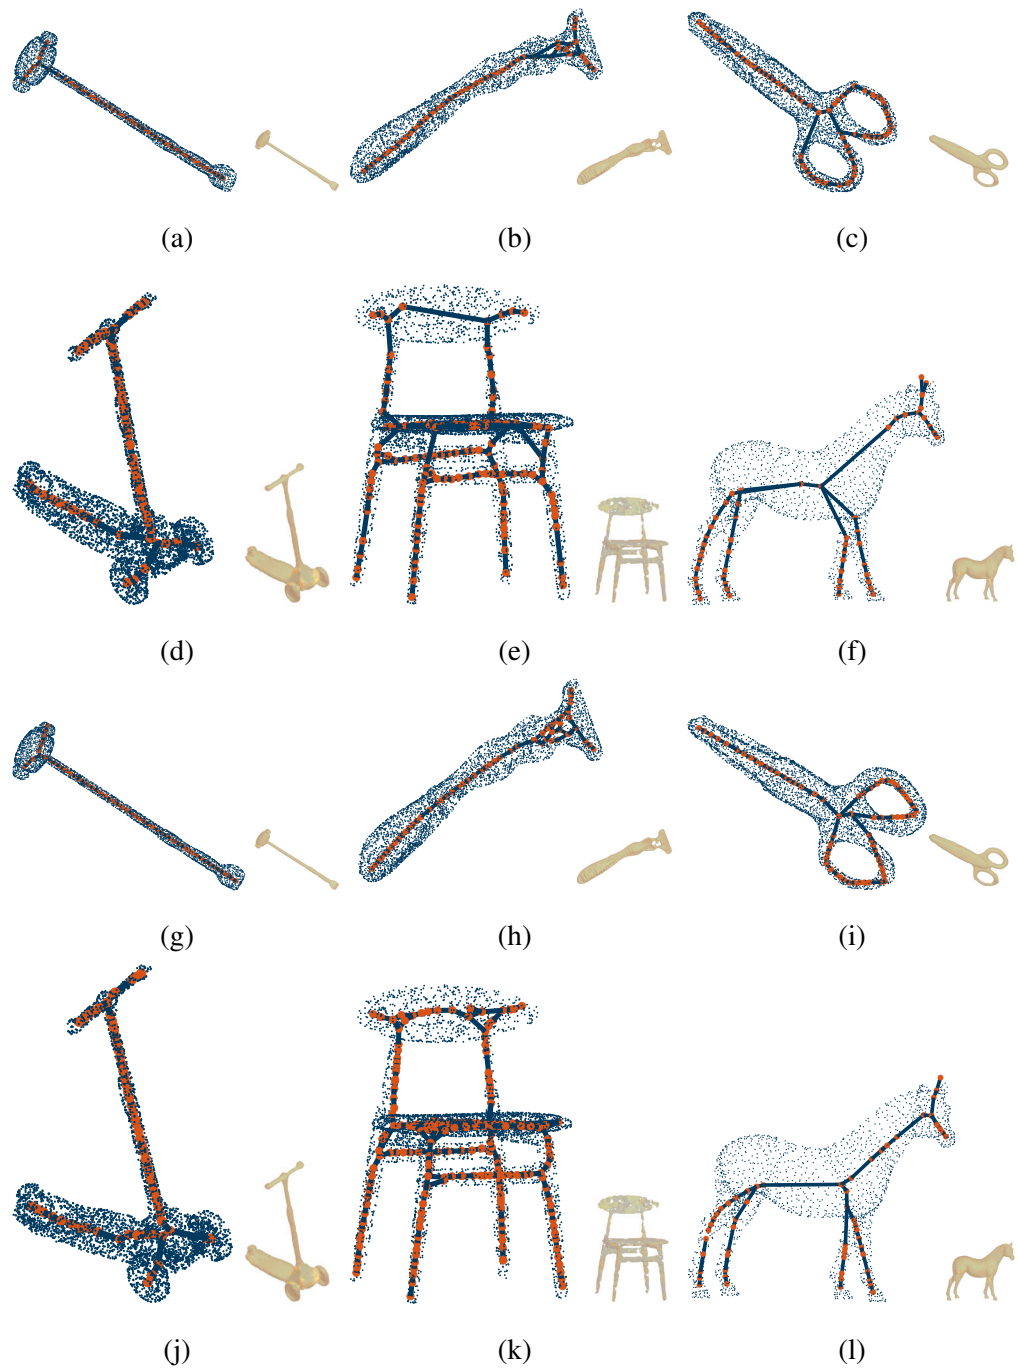


Figure 3.7: Skeletonisation comparison for selected models. (a)-(f) Results using the baseline; (g)-(l) results using GLSkeleton.

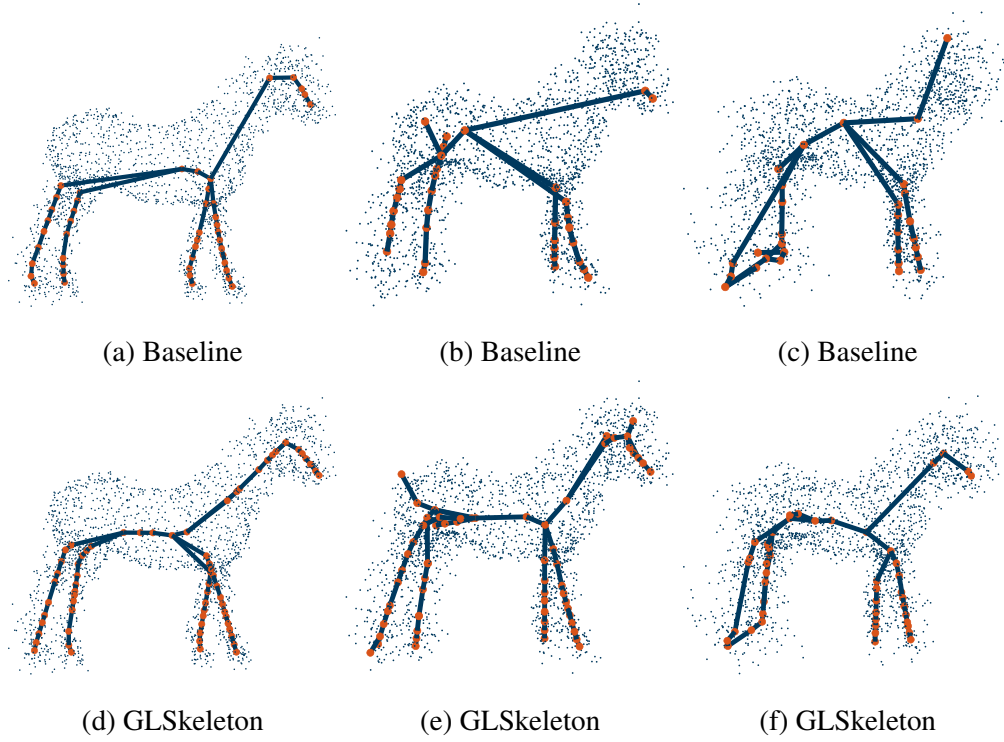


Figure 3.8: Skeletonisation results for noisy point clouds. Gaussian noise with standard deviations of 5%, 10%, and 15% (relative to half the diagonal of the object's bounding box) is added to each point.

Desirable Skeletonisation Properties and Skeletonisation Quality Evaluation

4.1 Introduction

Defining high-quality skeletonisation is a critical problem in evaluating skeletonisation approaches [155]. Despite the importance of quantifying the quality of skeletonisation methods, research on quantifying skeletonisation quality is limited, largely due to the algorithmic complexity of existing approaches and the lack of clearly defined skeleton properties [144]. Much of the existing research still highly relies on visual inspections [34, 143, 66], or ground truth-based evaluation [101, 49]. It is challenging to quantify skeletonisation quality or skeletonisation comparison, due to reasons such as no commonly accepted skeleton definition, difficulty in obtaining “ground truth” skeletons [155].

Addressing this challenge, this chapter firstly investigated the implicit skeletonisation patterns of the desired skeletonisation process. The statistical properties of the contraction-based skeletonisation results are systematically studied with a given criterion for deriving the desired Laplacian-based contraction, followed by the desired skeletonisation.

Apart from the statistical studies, a quantitative skeletonisation quality evaluation method via four critical criteria is investigated, aiming to numerically quantify the quality of the resultant skeletons. The critical skeleton properties, including topology similarity, centredness, boundedness and smoothness, are defined both theoretically and numerically. The effectiveness of the numerically defined evaluation approach is comprehensively discussed via experiments on skeletonisation results of various qualities, with an application background of robotics. The numerical skeletonisation analysis is available for the research community ¹.

In this chapter, the remaining is organised as follows. The study of desirable characteristics of stable contraction-based skeletonisation is introduced in Section 4.2. Then the theoretical and numerical definition of four critical skeletonisation evaluation metrics corresponding to skeletonisation properties are introduced in Section 4.3. All the skeletonisation evaluation experiments and discussions are present in Section 4.4. Finally, the work of this chapter is concluded in Section 4.5.

4.2 Characteristics of Stable Contraction-Based Skeletonisation

This section delves into the significant statistical pattern of geometrical attributes of stable contraction-based skeletonisation, with a focus on addressing the inherent difficulties in assessing skeletonisation outcomes. The contraction-based skeletonisation methodology is exemplified through the Laplacian-based approach [34], which is also introduced in Subsection 3.2.1. The contracted set of points \mathbf{P}_s is herein referred to as the surface skeleton, as it serves as an approximation of the MA; the ultimate skeleton derived from the process is designated as the curve skeleton.

Although the LBC is capable of producing both the surface skeleton and the resultant

¹https://github.com/weiqimeng1/PointCloud_Skeletonization_Metrics

curve skeleton, the task of evaluating the quality of the generated skeleton remains substantially challenging. To elucidate the statistical characteristics of stable contraction, this study investigates the deformation and distortion effects induced by the Laplacian contraction by analysing variations in surface normal vectors and curvature.

To quantify distances between surface normal vectors, the cosine similarity metric is employed due to its superior ability to capture directional relationships compared to alternative measures [182]. For general m -dimensional vectors \mathbf{n}_x and \mathbf{n}_y , the cosine similarity is formally defined as

$$S_C(\mathbf{n}_x, \mathbf{n}_y) = \cos(\theta) = \frac{\mathbf{n}_x \cdot \mathbf{n}_y}{\|\mathbf{n}_x\| \|\mathbf{n}_y\|} = \frac{\sum_{i=1}^m n_{x,i} n_{y,i}}{\sqrt{\sum_{i=1}^m n_{x,i}^2} \sqrt{\sum_{i=1}^m n_{y,i}^2}}, \quad (4.1)$$

where θ denotes the angle subtended between \mathbf{n}_x and \mathbf{n}_y . Given the nonlinear relationship between the cosine value and the angle, the cosine distance is normalised via

$$D_\theta = \frac{\arccos(S_C)}{\pi}, \quad (4.2)$$

such that D_θ represents the normalised angular distance bounded by $0 < D_\theta < 1$. For each point \mathbf{p}_i , the distance between the normal vectors of the original and contracted point clouds is computed as

$$D_{\theta,k,i} = \frac{\arccos(S_C(\mathbf{n}_{k,i}, \mathbf{n}_{0,i}))}{\pi}, \quad (4.3)$$

where $\mathbf{n}_i(k)$ and $\mathbf{n}_i(0)$ correspond to the normal vectors at the original and post- k -th contraction iteration positions, respectively. As presented in Eq. (4.3), it is requisite to obtain the normal vectors prior to distance computation. Analogous to the determination of local neighbour rings, a subset of nearest neighbouring points surrounding each point in the cloud is utilised to estimate the normal vectors. This estimation is realised through the calculation of the covariance matrix \mathbf{C}_n of the k -nearest neighbouring points, as expressed by Eq. (3.12). Upon computing the eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ of \mathbf{C}_n alongside their associated eigenvalues $\lambda_1, \lambda_2, \lambda_3$ arranged such that $\lambda_1 < \lambda_2 < \lambda_3$, the normal vector at the point is identified as \mathbf{v}_1 .

To quantify the differences in curvature between the contracted results and the original point cloud, we herein define the curvature difference metric. Utilising the covariance matrix \mathbf{C} described previously, the curvature at a given point is estimated and normalised as

$$\kappa_n = \frac{\lambda_1}{\sum_{i=1}^3 \lambda_i}. \quad (4.4)$$

Since the curvatures are represented as scalar values, the curvature difference between the input point cloud and its contracted counterpart can be readily computed by

$$\Delta\kappa_{n,k} = \kappa_{n,k} - \kappa_{n,0}, \quad (4.5)$$

where $\kappa_n(k)$ and $\kappa_n(0)$ denote the curvature at the point position after the k -th contraction iteration and at the original point position, respectively. Here k denotes the point cloud contraction iteration number, with $k = 0$ representing the original point cloud.

In this context, we propose a hypothesis concerning the stable convergence behaviour of the contraction process and describe a qualitative evaluation method based on the convergence characteristics of the contracted surface. It is noteworthy that, to the best of our knowledge, there exists no prior research explicitly addressing the desirability or stability of contraction convergence. Consequently, Propositions 4.1–4.2 have been formulated as abstractions of the convergence patterns consistently observed in the contraction results presented in Subsection 4.2, where the characteristics of the Laplacian-based skeletonisation are experimentally examined. These experimental observations provide the empirical foundation upon which the propositions are established.

The first proposition concerns the stability of the contraction process. As illustrated by the iterative contraction process in Fig. 2.6, we empirically observe that, for a stable contraction, every point in the point cloud moves progressively inwards over successive iterations, with the contracted surface remaining well contained within the original shape. Building on these recurrent observations, we derive Proposition 4.1.

Proposition 4.1. Let us consider the original point set of the object, denoted as $\mathbf{P}_o = \{\mathbf{p}_{o,i}\}$, and the contracted point cloud obtained after the k -th contraction iteration,

denoted as $\mathbf{P}_s(k) = \{\mathbf{p}_{s,i}(k)\}$. Here, the index i refers to the i -th point within the point set. The surfaces represented by the point clouds \mathbf{P}_o and $\mathbf{P}_s(k)$ are designated as U_o and $U_{s,k}$, respectively. The stability of the contraction is assessed locally with respect to the boundary region of the subset $\mathbf{X}_o \subset \mathbf{P}_o$. If the iterative contraction results, denoted as $\mathbf{X}_{s,k} \subset \mathbf{P}_s(k)$, satisfy the condition

$$X_{s,k+1} < X_{s,k} \ll X_o,$$

with the contracted surfaces constrained by

$$U_{s,k+1} < U_{s,k} \ll U_o,$$

we define this criterion as *contraction boundedness*, which serves as a measure to evaluate the stability of the contraction process.

Furthermore, since geometric properties are vital to a meaningful shape representation, we examine the stable convergence of the contraction via the analysis of geometric property variances, including the curvature differences $\Delta\kappa_{n,k}$ and the angular differences of the normal vectors $D_{\theta,k,i}$ within local surface regions. Empirically, as shown in the experimental results of Section 4.2, stable contraction convergence patterns exhibit characteristic distributions of these quantities.

Proposition 4.2. For a point cloud contraction process exhibiting stable convergence, the empirical distributions of $D_{\theta,k,i}$ and $\Delta\kappa_{n,k}$ are anticipated to be approximately symmetric and unimodal, resembling a “bell-shaped” profile. Moreover, the means of $D_{\theta,k,i}$ and $\Delta\kappa_{n,k}$ are expected to converge towards values of 0.5 and 0, respectively.

Beyond these shared characteristics, distinct patterns in the distributions may reveal particular geometrical structures inherent to different shape categories or indicate anomalies in shape processing.

Given that the contracted results approximate the surface skeleton, these distributional patterns provide an effective means of evaluating the quality of the resulting surface

skeleton. Additionally, as the Laplacian-based curve skeletonisation relies fundamentally on the stability of contraction results, a stable and convergent contraction pattern is instrumental in producing curve skeletons of superior quality.

4.3 Skeletonisation Evaluation Metrics

Beyond the statistical examination of contraction-based skeletonisation, this section aims to provide a comprehensive theoretical and numerical characterisation of skeletonisation properties that influence the quality of the resultant skeleton. These properties include topological similarity, boundedness, centredness, and smoothness.

4.3.1 Topological Similarity

This subsection addresses a quantitative analysis of the preservation of topology throughout the skeletonisation process. While distances between point cloud datasets may be evaluated using various metrics—such as the Hausdorff or Chamfer distances—disentangling purely shape-related metrics from those based on distance measurements remains a challenging endeavour. Herein, the focus is placed on topological similarity as a fundamental attribute of skeletonisation outcomes. Two shapes are considered topologically similar if they exhibit comparable topological structures. To assess point cloud skeletonisation with respect to topological similarity, it is necessary to develop a method that compares two topological shapes represented by point sets and produces a score quantifying the degree of their topological congruence.

Problem Statement & Preliminaries

An ideal skeleton is expected to preserve the topology of the original shape [40]. To evaluate the extent of topology preservation during skeletonisation, a metric that effectively captures topological shape similarity is required. In this part, we examine the

topological alterations that arise throughout the point cloud skeletonisation process. We assume the existence of an intermediate skeletal point set, point-wise related to the original point cloud, generated as part of the skeletonisation pipeline. As an overview, we approximate topological similarity by analysing the evolution of nontrivial H_0 features, which indicate the emergence or disappearance of connected components during the contraction process.

The problem is formulated with object point clouds as defined by Definition 3.1. Similarly, we define two point cloud shapes relevant to this problem as: \mathbf{P}_o , denoting the original point cloud, and \mathbf{P}_s , representing the skeletal point set, which approximates the number of points in \mathbf{P}_o .

Local geometric features of \mathbf{P} are frequently modelled using simple geometrical constructs known as simplices. Generally, a simplex is the convex hull of its vertices, formalised as follows:

Definition 4.1. A k -simplex σ is the convex hull of its $k + 1$ affinely independent vertices, denoted by

$$\sigma := \text{convh}\{\mathbf{v}_0, \dots, \mathbf{v}_k\},$$

where $\mathbf{v}_0, \dots, \mathbf{v}_k \in \mathbb{R}^3$ are the vertices of σ .

Remark 4.1. The term “affinely independent vertices” signifies that given the vertices $\mathbf{v}_0, \dots, \mathbf{v}_k$, the vectors $\mathbf{v}_1 - \mathbf{v}_0, \dots, \mathbf{v}_k - \mathbf{v}_0$ are linearly independent. As illustrated in Fig. 4.1, simplices in dimensions 0, 1, 2, and 3 correspond respectively to a point, a line segment, a triangle, and a tetrahedron. The convex hull of any subset of the vertices of a simplex constitutes a face of that simplex.

For more intricate structures, simplicial complexes are defined as unions of simplices. We adopt this framework to describe sets of point clouds (including skeletal point sets) as follows:

Definition 4.2. A *simplicial complex* K is a finite collection of simplices such that for




Figure 4.1: Examples of simplices in various dimensions.

any $\sigma, \sigma' \in K$, their intersection $\sigma \cap \sigma'$ is either empty or a face common to both σ and σ' .

An *abstract simplicial complex* \hat{K} is a collection of sets formed from the geometric elements of simplicial complex K .

Remark 4.2. In an abstract simplicial complex, every non-empty subset of a set contained in \hat{K} is also an element of \hat{K} . Effectively, this abstraction disregards certain geometric attributes of the simplicial complex, preserving only the combinatorial relationships among the elements.

The Vietoris-Rips complex provides a natural method for constructing an abstract simplicial complex from a finite metric space and can be used to extract topological features through complex filtration [9]. The given finite metric space serves as a guide for combinations of geometrical elements.

Definition 4.3. A finite metric space $(\mathbf{X}, d_{\mathbf{X}})$ of the discrete space of the point set \mathbf{X} is a metric space such that the distance between a pair of points \mathbf{x}_i and \mathbf{x}_j ($\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}$) is given by $d_{\mathbf{X}}(\mathbf{x}_i, \mathbf{x}_j)$.

Definition 4.4. Given a finite metric space $(\mathbf{X}, d_{\mathbf{X}})$ and a fixed radius ϵ , the Vietoris-Rips complex $\text{VR}_{\epsilon, \mathbf{P}}(\mathbf{X}, d_{\mathbf{X}})$ of point cloud shape \mathbf{P} is an abstract simplicial complex where the vertices are the points in \mathbf{P} , and each k -simplex $\sigma = \text{convh}\{\mathbf{v}_0, \dots, \mathbf{v}_k\} \in \text{VR}_{\epsilon, \mathbf{P}}(\mathbf{X}, d_{\mathbf{X}})$ satisfies

$$d_{\mathbf{X}}(\mathbf{v}_i, \mathbf{v}_j) < \epsilon, \quad \text{for all } 0 \leq i < j \leq k.$$

Homology in a topological space is characterised by its homology groups, defined using the boundary homomorphism [15]. It serves as a topological invariant, allowing the comparison of point cloud shapes based on their topological properties.




Figure 4.2: An example of a Vietoris-Rips complex [15].

Definition 4.5. A boundary operator (or homomorphism) ∂_d is a homomorphism connecting the chain complexes in dimension d and $d - 1$, expressed as

$$\partial_d : C_d(K) \rightarrow C_{d-1}(K),$$

where for $i < 0$, $C_i \equiv 0$.

Definition 4.6. Let the d -th boundaries and cycles of \mathbf{P} constructed in point space be denoted as

$$B_d(\mathbf{P}) := \text{im}(\partial_{d+1}), \quad Z_d(\mathbf{P}) := \ker(\partial_d),$$

respectively. The d -th homology group of the topological space relative to the point cloud \mathbf{P} is defined by

$$H_d(K) := \frac{\ker \partial_d}{\text{im } \partial_{d+1}} = \frac{Z_d(\mathbf{P})}{B_d(\mathbf{P})}.$$

Remark 4.3. The d -th homology group $H_d(K)$ of a point cloud shape represents the number of d -dimensional holes in the corresponding simplicial complex, thereby describing its topological structure.

Definitions 4.4–4.6 provide the mathematical foundation for the homology groups of point cloud shapes. As a topological invariant, the homology group enables the abstraction and comparison of topological properties between shapes. For point clouds, this is practically achieved using persistent homology.

Methodology

Persistent homology constitutes a robust framework for analysing topological variations in point clouds. Conventional skeletonisation techniques characterise a “good” skeleton as one that preserves topology by remaining homotopic to the original shape [40, 143], although formal quantification of this property is non-trivial. Motivated by Edelsbrunner et al. [51], we apply persistent homology to extract topological features from both the skeleton and the original point cloud, enabling quantification of their dissimilarity via topological distances.

Let ϵ denote the radius parameter defining ϵ -neighbourhoods used to construct complexes. Persistent homology tracks the emergence and disappearance of topological features as ϵ increases. At $\epsilon = 0$, each point is isolated, representing an H_0 feature. As ϵ increases, points connect to form higher-dimensional simplices, generating loops (H_1) and cavities (H_2). These features appear and vanish at different scales: minor features vanish rapidly, while significant ones endure longer. We model complex growth using the Vietoris-Rips complex (Definition 4.4). The persistence of features is encoded by their birth and death times and can be visualised via persistence barcodes [35]. As depicted in Fig. 4.4, the start and end points of each bar represent the birth and death of an H_0 feature, respectively.

Within barcode space, bottleneck and Wasserstein distances serve as standard metrics for quantifying dissimilarity between two barcodes. For skeletonisation, the two barcodes are assumed to originate from the original point cloud and its skeletal representation. Given two persistence intervals $[a_1, b_1)$ and $[a_2, b_2)$, the ∞ -distance between them is defined as:

$$d_\infty([a_1, b_1), [a_2, b_2)) = \max(|a_1 - a_2|, |b_1 - b_2|).$$

Since preservation typically concerns only major topological features exhibiting high persistence, it is essential to filter out minor local features prior to comparison. Let

\mathbf{P}_o and \mathbf{P}_s denote the original point cloud and the skeletal point set, respectively. The maximal nearest-neighbour distance for any point $\mathbf{p}_i \in \mathbf{P}_o$ is defined as ϵ^* :

$$\epsilon^* = \sup_{\mathbf{p}_o, i \in \mathbf{P}_o} \|\mathbf{p}_i - \psi(\mathbf{p}_i)\|, \quad (4.6)$$

where ψ maps each point \mathbf{p}_i to its nearest neighbour in \mathbf{P}_o . We assume that minor local features, such as small connected components, disappear once each point is connected to at least one neighbour. For example, a point forming the smallest connected component is born at $\epsilon = 0$ and dies when linked to its nearest neighbour. Hence, bars corresponding to minor features in persistence barcodes are eliminated by discarding those with persistence below ϵ^* .

Let B_o and B_s denote the filtered barcodes of \mathbf{P}_o and \mathbf{P}_s , respectively. The point sets \mathbf{P}_o and \mathbf{P}_s are normalised to fit within a standard cubic bounding box to facilitate comparison. The *bottleneck distance* between the filtered barcodes is defined by:

$$d_B(B_o, B_s) = \inf_{\phi} \sup_{Z \in B_o} d_{\infty}(Z, \phi(Z)), \quad (4.7)$$

where ϕ ranges over all bijections between B_o and B_s . The normalised *p-Wasserstein distance* is given by:

$$d_{W_p}(B_o, B_s) = \frac{1}{n_b} \left(\inf_{\phi} \sum_{Z \in B_o} d_{\infty}(Z, \phi(Z))^p \right)^{\frac{1}{p}}, \quad (4.8)$$

where n_b is the number of selected bars.

Proposition 4.3. Consider two normalised point clouds \mathbf{P}_o and \mathbf{P}_s enclosed within a bounding box of diagonal length ϵ_{\max} . Their topological similarity is characterised by the distance between their most persistent homology features, computed via the Vietoris-Rips complex filtration (Definition 4.4). Persistence patterns are represented in barcode space and filtered using the threshold ϵ^* defined in Equation (4.6). Let B_o and B_s be the filtered persistence barcodes corresponding to \mathbf{P}_o and \mathbf{P}_s , respectively, with bars representing the birth and death of H_0 features. The dissimilarity $d_{o,s}$ between B_o and B_s is quantified by distance metrics such as the bottleneck distance (Equation (4.7)) or the Wasserstein distance (Equation (4.8)) [35].

Let N_o be the number of points in \mathbf{P}_o . The topological similarity between \mathbf{P}_o and \mathbf{P}_s is then expressed as:

$$\begin{cases} \text{High similarity,} & \text{if } d_{o,s} < d^* \text{ with } 0 < d^* \leq \epsilon_{\max}, \\ \text{Low similarity,} & \text{if } d_{o,s} \geq d^*. \end{cases}$$

As $d^* \rightarrow 0$ when $N_o \rightarrow \infty$, this criterion yields an increasingly accurate measure of topological similarity.

It has been established by Niyogi et al. [108] that the homology of a shape can be reliably recovered from sufficiently dense discrete samples. As a topological invariant, the homology group serves as a fundamental tool for shape analysis [167, 168, 56].

Proof. Consider a discrete point set \mathbf{P}_o composed of N_o points sampled uniformly from the surface U_o of a shape o . Define ϵ^* as the largest nearest-neighbour distance among all points $\mathbf{p}_i \in \mathbf{P}_o$, as specified in Eq. 4.6. Assuming that \mathbf{P}_o represents valid samples of U_o , the surface U_o may be approximated by connecting each point \mathbf{p}_i to its neighbours within radius ϵ , forming a reconstructed surface U_r .

As the number of points increases without bound, $N_o \rightarrow \infty$, the maximum neighbour distance shrinks, $\epsilon^* \rightarrow 0$, and the reconstructed surface converges, $U_r \rightarrow U_o$. This convergence implies that the homology features computed from the point samples approach those of the original shape. Additionally, as $\epsilon^* \rightarrow 0$, the persistence associated with minor H_0 features—commonly noise—rapidly decreases. Since ϵ^* measures local proximity, it predominantly affects small-scale connectivity and not the global topology. Thus, constructing the ϵ -Vietoris-Rips complex for sufficiently small ϵ effectively filters out these insignificant local features while preserving key global topological characteristics. This phenomenon is demonstrated in Figs. 4.3(b), (e), (h), and (j). \square

Remark 4.4. The examination of simple geometric shapes reveals a direct correlation between topological alterations and persistence barcodes of homology. As depicted

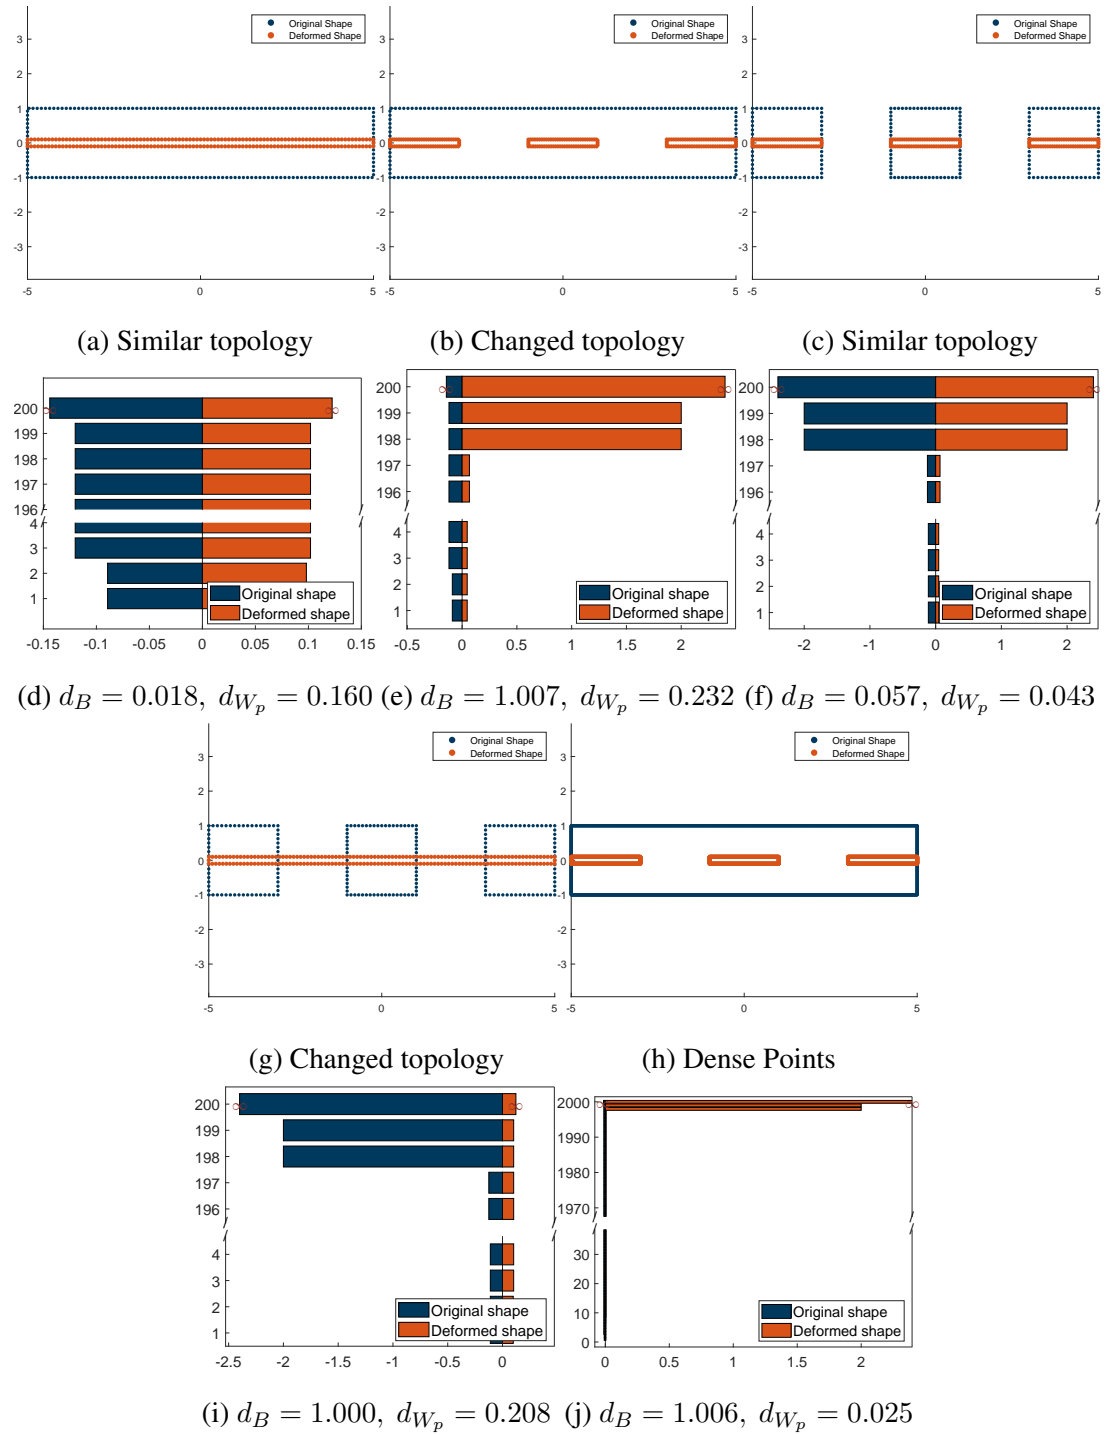


Figure 4.3: Persistent homology analysis of simple geometries using H_0 features. Shapes in (a), (b), (c), and (g) contain 200 points, and (h) contains 2000 points. The corresponding barcode diagrams are shown below each shape. The horizontal axis indicates birth and death scales, while the vertical axis simply indexes the bars for visual clarity.

in Fig. 4.3, topological changes affecting connectivity (e.g. Figs. 4.3(b), (g)) produce notable variations in the lengths of the longest barcode intervals. In contrast, topology-preserving deformations (e.g. Figs. 4.3(a), (c)) maintain consistent patterns in these dominant bars. Quantitative measures such as the Wasserstein and bottleneck distances effectively capture these effects. In particular, the bottleneck distance exhibits enhanced robustness against non-topological variations, confirming its reliability as a metric for topological dissimilarity.

Remark 4.5. The proposed topological similarity metric is applicable solely to point-wise deformations of point cloud shapes. Furthermore, it is sensitive to variations in point density, as the dominant bars are filtered according to the largest nearest-neighbour distance in the original point cloud.

Fig. 4.4 presents the persistence barcodes of homological features for both the original and skeletal point sets of the hammer and biscuit shapes. In Fig. 4.4(a), inward displacement of the skeletal points reduces neighbour distances, causing connected components to merge earlier and thus shortening the persistence of H_0 features. This reflects quicker local connections forming while global topological features remain intact. The displacement is primarily radial, reducing distances more significantly in the radial than the axial (medial) direction. Effective skeletonisation preserves these medial-axis topological features.

In the hammer shape (Fig. 4.4(a)), longer-persisting homology features signal limited topological alteration compared to the biscuit shape (Fig. 4.4(b)), where the topology is considerably disrupted, indicated by irregular spikes. The extent of topological preservation through skeletonisation is quantitatively captured by the distances between barcode patterns: the hammer shape contraction yields smaller bottleneck and Wasserstein distances than the biscuit shape, as reported in the figure captions.

In summary, persistent homology enables quantification of topological similarity by comparing the persistence intervals of homology features between original and skeletal point sets, using metrics such as bottleneck and Wasserstein distances.

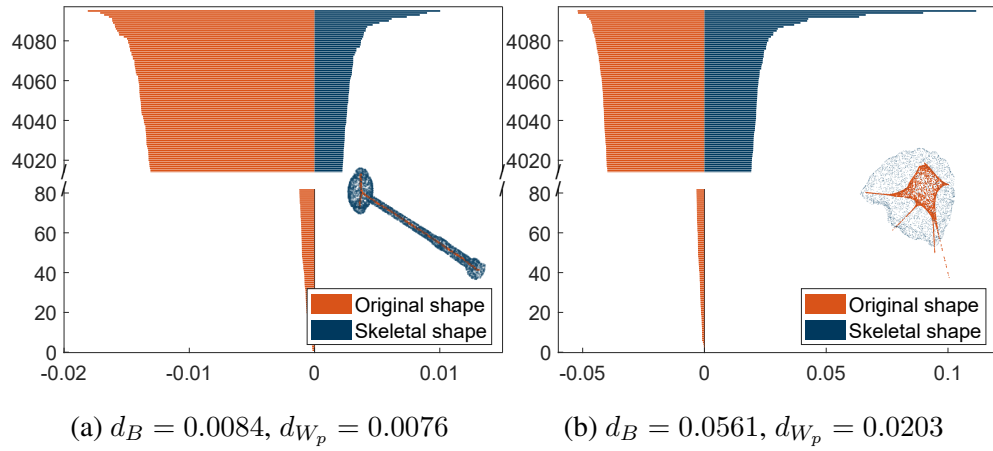


Figure 4.4: The barcode of persistent homology (H_0 features). The input point clouds are scaled to fit within a cubic bounding box whose diagonal is 1.6 and only the top 5% and bottom 5% of persistence bars are shown for clarity)

4.3.2 Boundedness

The skeleton of a shape represented by a point cloud is expected to lie within the shape, or equivalently, to be bounded by the original shape surface [144, 40], before considerations of other properties, such as centredness, can be made. However, the principal challenge is that the surface defined by point cloud data is implicit and not explicitly given. This subsection addresses the boundedness of skeletal points or vertices. Given any point in the point cloud space, our method assigns a score reflecting the bounding status of that point relative to the shape's surface. Aggregating these results enables an assessment of the overall boundedness of the derived skeletal shape.

Problem statement

As exemplified by the skeletonisation result at the bottom-right of Fig. 4.4(b), skeletal points or vertices extending beyond the shape's surface boundary are generally undesirable and represent incorrect convergence or transformation. *Boundedness* denotes the property whereby all skeletal elements lie within the original shape. For point cloud shapes, this is challenging because the surface is ambiguously represented by discrete

points.

Here, boundedness is defined at the point level. Let $\mathbf{P}_s \subset \mathbb{R}^{m \times 3}$ denote the skeletal point set, where each skeletal point $\mathbf{p}_{s,i} \in \mathbb{R}^3$ corresponds to an original point cloud $\mathbf{P}_o \subset \mathbb{R}^{n \times 3}$. An example skeletal point set is the contracted point cloud as shown in Fig. 3.2. The shape represented by \mathbf{P}_o is denoted as Ω . Each skeletal point $\mathbf{p}_{s,i} \in \mathbf{P}_s$ is expected to lie within the shape boundary $\partial\Omega$.

A curve skeleton G_s is as the definition given by Definition 3.1. Each element $e_i \in E$ or vertex $\mathbf{v}_i \in V$ is expected to be bounded by $\partial\Omega$.

Methodology

Let \mathbf{P} be an n -point cloud within the discrete space $\mathbf{X} \subseteq \mathbb{R}^{m \times 3}$, and $\mathbf{x}_i \in \mathbf{X}$ be an arbitrary point. We define the direction vector from \mathbf{x}_i to \mathbf{p}_i as

$$\mathbf{d}(\mathbf{x}_i, \mathbf{p}_i) := \frac{\mathbf{p}_i - \mathbf{x}_i}{\|\mathbf{p}_i - \mathbf{x}_i\|}, \quad (4.9)$$

where $\mathbf{p}_i \in \mathbf{P}$ is any point in \mathbf{P} . Interpreting these normalised vectors from all $\mathbf{p}_i \in \mathbf{P}$ relative to \mathbf{x}_i as coordinates yields points projected onto a unit sphere. If \mathbf{x}_i is fully enclosed by the shape defined by \mathbf{P} , this projection produces a nearly complete spherical surface, as illustrated in Fig. 4.5. Conversely, if \mathbf{x}_i is only partially enclosed,




Figure 4.5: Projection of directional vectors onto a spherical surface.

regions of the sphere remain empty, manifesting as holes on the sphere's surface. This observation motivates the following definitions for boundedness metrics:

Definition 4.7. Assuming the shape represented by \mathbf{P} is convex, the boundedness of \mathbf{x}_i relative to the point cloud \mathbf{P} can be evaluated by the proportion of the sphere's surface area occupied by the projected points. Let $\mathbf{p}_{\mathbf{x}_i, \mathbf{P}} = \mathbf{d}(\mathbf{x}_i, \mathbf{p}_i)$ (from Eq. 4.9) denote the projection of \mathbf{p}_i with respect to \mathbf{x}_i , and $\mathcal{P}_{\mathbf{x}_i, \mathbf{P}}$ be the corresponding set of points on the sphere. The boundedness metric of \mathbf{x}_i with respect to \mathbf{P} is defined by

$$\beta_{\mathbf{x}_i \circ \mathbf{P}} := \frac{S_{\mathbf{x}_i, \mathbf{P}}}{4\pi r^2}, \quad (4.10)$$

where $r = 1$ for the unit sphere and $S_{\mathbf{x}_i, \mathbf{P}}$ denotes the total area covered by points $\mathbf{p}_{\mathbf{x}_i, \mathbf{P}} \in \mathcal{P}_{\mathbf{x}_i, \mathbf{P}}$.

Definition 4.8. Given a curve skeleton $G_s = (V, E)$, an edge $e_i = (\mathbf{v}_j, \mathbf{v}_k) \in E$ can be parametrised by

$$e_i(t) = t\mathbf{v}_j + (1 - t)\mathbf{v}_k, \quad t \in [0, 1].$$

A point \mathbf{p}_g on the skeleton is given by $\mathbf{p}_g = e_i(t)$, for some $t \in [0, 1]$ and edge $e_i \in E$.

Definition 4.9. Let $\mathbf{P}_o, \mathbf{P}_s \subseteq \mathbf{X}$ be the original and skeletal point sets in the discrete space \mathbf{X} . Denote by N_s the total number of skeletal points, and by $N_{s,b}$ the number of skeletal points bounded by \mathbf{P}_o (see Proposition 4.4). The boundedness of \mathbf{P}_s relative to \mathbf{P}_o is

$$\mathcal{B}_{\mathbf{P}_s \circ \mathbf{P}_o} := \frac{N_{s,b}}{N_s}. \quad (4.11)$$

For a curve skeleton G_s , sample N_{sp} points evenly according to Definition 4.8, with $N_{sp,b}$ bounded points. Then the curve skeleton boundedness is defined as

$$\mathcal{B}_{G_s \circ \mathbf{P}_o} := \frac{N_{sp,b}}{N_{sp}}. \quad (4.12)$$

Exact computation of the area covered by points on the sphere's surface is difficult. To simplify, we project the sphere points onto a 2D plane to perform Delaunay triangulation efficiently.

For an arbitrary point on the unit sphere with coordinates $[x_s, y_s, z_s]^T$, its planar coordinates $[x_p, y_p]$ under the sinusoidal projection [133] are given by

$$\begin{cases} x_p = \sqrt{x_s^2 + y_s^2} \cdot \tan^{-1} \left(\frac{y_s}{x_s} \right), \\ y_p = r \cdot \tan^{-1} \left(\frac{z_s}{\sqrt{x_s^2 + y_s^2}} \right), \end{cases}$$

where $r = 1$ is the sphere radius.

After projecting onto the plane, Delaunay triangulation is applied. The triangulation is then mapped back to the sphere to approximate the covered area $S_{\mathbf{x}_i, \mathbf{P}}$ in Eq. (4.10) by summing the areas of all triangles, as illustrated in Fig. 4.6.

Proposition 4.4. Suppose the shape $\Omega \subset \mathbb{R}^3$ with boundary $\partial\Omega$ is convex and represented by the point cloud \mathbf{P}_o containing N_o points. Let S_Δ denote the approximate covered area computed by summing the areas of all triangulated regions (Fig. 4.6). Then the approximated boundedness

$$\hat{\beta}_{\mathbf{x}' \circ \mathbf{P}_o} = \frac{S_\Delta}{4\pi r^2}$$

serves as an estimate of the true boundedness $\beta_{\mathbf{x}' \circ \mathbf{P}_o}$ (Eq. (4.10)). A point $\mathbf{x}' \in \mathbf{X}$ is considered bounded by $\partial\Omega$ if $\hat{\beta}_{\mathbf{x}' \circ \mathbf{P}_o} \geq \beta^*$, where the threshold $0 < \beta^* \leq 1$ depends on the point density of \mathbf{P}_o . As $N_o \rightarrow \infty$, $\beta^* \rightarrow 1$ and the approximation becomes exact.

Proof. Referring to Fig. 4.6, let a_i denote the area of each triangulated patch on the sphere generated by points projected from \mathbf{P}_o relative to \mathbf{x}' . Define $a^* = \sup_i a_i$ as the largest triangle area. As $N_o \rightarrow \infty$, the triangulation becomes finer and $a^* \rightarrow 0$. Consequently, the approximate area $\hat{S}_{\mathbf{x}', \mathbf{P}_o} = \sum_i a_i$ converges to the true area $S_{\mathbf{x}', \mathbf{P}_o}$.

If \mathbf{x}' is fully enclosed by $\partial\Omega$, then $S_{\mathbf{x}', \mathbf{P}_o} = 4\pi r^2$, the total area of the unit sphere. Hence, $\hat{\beta}_{\mathbf{x}' \circ \mathbf{P}_o} \rightarrow 1$ as $N_o \rightarrow \infty$, completing the proof. \square

In practical applications, as illustrated in Fig. 4.7, the point density of point clouds is typically lower than the theoretical ideal. Consequently, the boundedness threshold β^*




Figure 4.6: Triangulated projections of points on the unit sphere for (a) a bounded skeletal point and (b) an unbounded skeletal point, based on a horse-shaped point cloud.

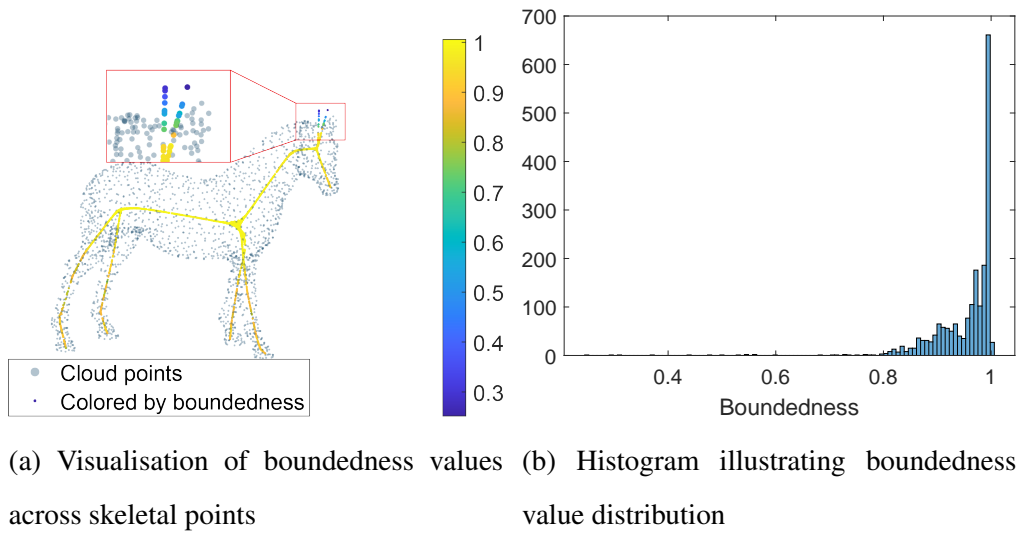


Figure 4.7: Boundedness assessment of contracted skeletal points derived from a horse point cloud. Points exterior to the shape boundary demonstrate lower boundedness metrics, whereas points interior to the boundary exhibit values approaching unity.

is generally observed to be less than one. Moreover, skeletal points situated beyond the shape boundary consistently present markedly reduced boundedness values, while those contained within the boundary maintain boundedness values close to unity.

4.3.3 Centeredness

Centeredness is widely regarded as one of the most fundamental properties in point cloud skeletonisation, serving as a key attribute across numerous applications. However, its quantitative assessment remains challenging. This difficulty arises not only because the object's surface represented by a point cloud is inherently uncertain, but also because skeletal structures have been defined in markedly different ways throughout the literature. To address these challenges, this subsection introduces a metric that evaluates the centredness of skeletonisation outcomes by integrating medial-axis-based principles from both surface and curve skeleton frameworks.

Problem Statement & Preliminaries

Multiple interpretations of skeletonisation have resulted in a diverse set of definitions for what constitutes an ideal skeleton. Although there is broad conceptual agreement that centredness is a desirable and often essential property, the field lacks a unified quantitative formulation of centredness [40, 143]. As noted in recent surveys, no consensus mathematical definition is currently established [155]. Existing works typically express centredness only implicitly, most commonly through the expectation that a well-formed skeleton should approximate the medial axis (MA) of the underlying shape [115, 58, 143, 158]. This medial-axis perspective, therefore, provides the foundational basis from which centredness is understood.

Building on the conceptual descriptions found in prior work [40, 143], a rigorous centredness definition for both the skeletal point set \mathbf{P}_s and the curve skeleton G_s of a shape Ω can be stated as follows. The skeletal point set $\mathbf{P}_s = \{\mathbf{p}_{s,i} \mid \mathbf{p}_{s,i} \in \mathbb{R}^3\}$ corresponding to shape Ω should have every skeletal point $\mathbf{p}_{s,i}$ located on the MA \mathcal{M}_Ω defined in Eq. 4.13. Similarly, the curve skeleton G_s , represented as a graph (V, E) , must satisfy that all vertices $\mathbf{v}_i \in V$ and every point $e_i(t) \in E$ along its edges lie on the MA \mathcal{M}_Ω of the shape, where t parametrises the edges.

It is important to note that exact centredness is not always obligatory; some dimensions may be well-centred while others are not, and the level of centredness required varies depending on the application. For instance, approximate centredness is often sufficient for virtual navigation tasks, as the MA can contain overly detailed and unnecessary features [40]. Furthermore, the preceding definition presupposes knowledge of the shape's surface, which is typically unavailable for point cloud data. Therefore, before proposing our metric, we reiterate the formal definitions of the MA for known surfaces and subsequently extend these to a numerical approximation suitable for point cloud shapes.

Definition 4.10. For a shape $\Omega \subset \mathbb{R}^3$ with boundary $\partial\Omega$, the distance transform $DT_{\partial\Omega} : \mathbb{R}^3 \rightarrow \mathbb{R}_{>0}$ is defined as

$$DT_{\partial\Omega}(\mathbf{p}_x \in \Omega) = \min_{\mathbf{p}_y \in \partial\Omega} \|\mathbf{p}_x - \mathbf{p}_y\|.$$

Definition 4.11. The MA \mathcal{M}_Ω (or medial surface S_Ω) of shape Ω is the locus within the shape satisfying

$$\begin{aligned} \mathcal{M}_\Omega = \{ \mathbf{p}_x \in \Omega \mid \exists \mathbf{m}_1, \mathbf{m}_2 \in \partial\Omega, \mathbf{m}_1 \neq \mathbf{m}_2, \\ \text{such that } \|\mathbf{p}_x - \mathbf{m}_1\| = \|\mathbf{p}_x - \mathbf{m}_2\| \}. \end{aligned} \quad (4.13)$$

Methodology

Given that the surface of a real-scanned point cloud is typically uncertain and subject to noise, directly quantifying centredness as per Definition 4.11 is infeasible. However, under the assumption that opposite points within a cross-section of the shape cluster together due to the nature of skeletonisation, it is possible to estimate the proximity of the skeleton to the medial surface.

For skeletal point sets, centredness can be assessed by assuming each skeletal point $\mathbf{p}_{s,i} \in \mathbf{P}_s$ corresponds to a matching point $\mathbf{p}_{o,i}$ in the original point cloud \mathbf{P}_o . If a direct one-to-one correspondence is not available, it may be approximated through

reverse projection or by associating the skeleton points with neighbourhoods in the object point cloud.

Definition 4.12. Let $\mathbf{p}_{s,i} \in \mathbf{P}_s$ be a skeletal point. Its k nearest neighbours are found via the K-Nearest Neighbours (KNN) algorithm and denoted by $\{\mathbf{p}_{s,i \odot j}\}$, where $j = 1, \dots, k$. Define φ as the mapping from a skeletal point $\mathbf{p}_{s,i}$ to its corresponding point $\mathbf{p}_{o,i} \in \mathbf{P}_o$. The centredness $c(\mathbf{p}_{s,i})$ of $\mathbf{p}_{s,i}$, quantifying how centrally located it is within the shape, is expressed as

$$c(\mathbf{p}_{s,i}) := 1 - \frac{\left\| \sum_{j=1}^k \varphi(\mathbf{p}_{s,i \odot j}) - \sum_{j=1}^k \mathbf{p}_{s,i \odot j} \right\|}{\sum_{j=1}^k \left\| \varphi(\mathbf{p}_{s,i \odot j}) - \frac{1}{k} \sum_{j=1}^k \mathbf{p}_{s,i \odot j} \right\|}. \quad (4.14)$$

Remark 4.6. As demonstrated in Fig. 4.8, the centredness of skeletal points computed via Eq. (4.14) for the Laplacian-based skeletonisation method [34] reveals that points shifting towards the shape's centre exhibit higher centredness scores. Nevertheless, this approach inadequately handles points corresponding to skeletal joints and endpoints. Moreover, the centredness metric becomes invalid for points that lie outside the shape's boundary, such as the skeletal points representing the horse's ears shown in Fig. 4.8(a).

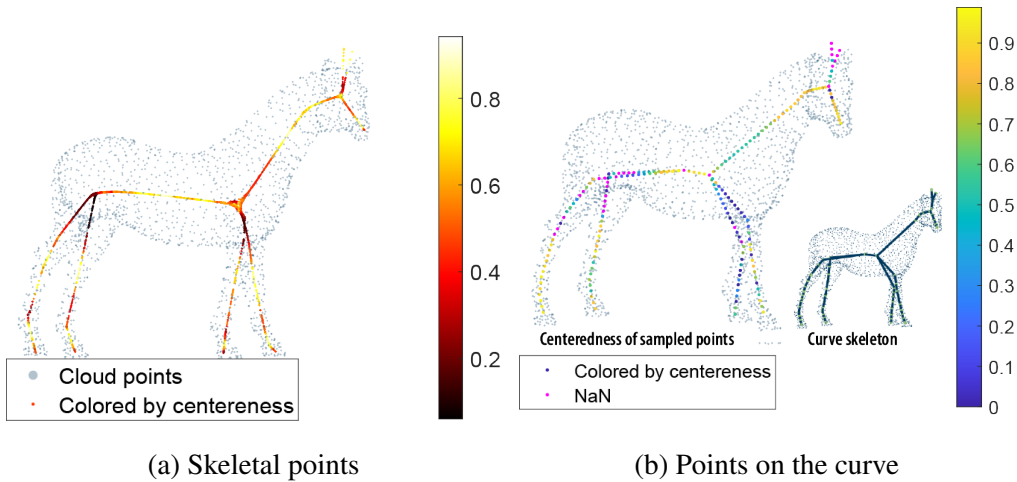


Figure 4.8: Centredness of skeletal points (a) and points sampled from the curve skeleton (b), respectively.

Considering the distinct physical and geometric properties of curve skeletons—where the curve is represented by a continuous sequence of connected vertices rather than discrete points—the centredness of a point on a curve skeleton $G_s = (V, E)$ is defined relative to its associated original cloud points, which are those points from the original point cloud in the vicinity of the curve point. As illustrated in Fig. 4.9, the original cloud points can be partitioned by two parallel cutting planes M_1 and M_2 , which share the same normal vector aligned with the curve's local direction.

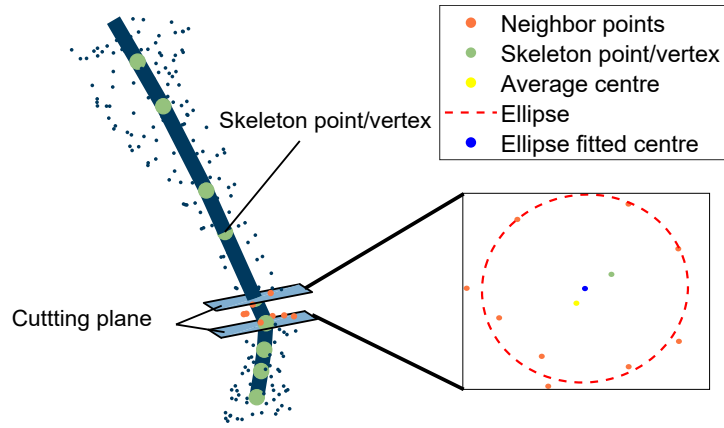


Figure 4.9: Points (orange) separated by two parallel cutting planes.

Definition 4.13. Consider two adjacent edges $e_1 = (\mathbf{v}_0, \mathbf{v}_1)$ and $e_2 = (\mathbf{v}_1, \mathbf{v}_2)$ in the graph G_s , where $e_1, e_2 \in E$ and $\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2 \in V$.

Let $\mathbf{p}_{g,1} = e_1(t_1)$ be a point on edge e_1 with parameter $t_1 \in [0, 1]$ (Definition 4.8). The curve direction vector at $\mathbf{p}_{g,1}$, denoted \mathbf{u}_1 , is defined as:

1. For $t_1 \in (0, 1)$:

$$\mathbf{u}_1 = \frac{\mathbf{v}_1 - \mathbf{v}_0}{\|\mathbf{v}_1 - \mathbf{v}_0\|}.$$

2. For $t_1 = 1$: \mathbf{u}_1 is approximated by the tangent vector at vertex \mathbf{v}_1 , computed as the tangent to the circle passing through $\mathbf{v}_0, \mathbf{v}_1$, and \mathbf{v}_2 .

The associated point set $\mathbf{Q}_1 \subset \mathbf{P}_o$ corresponding to $\mathbf{p}_{g,1}$ comprises all points enclosed between two parallel cutting planes, oriented by \mathbf{u}_1 and separated by an interval ϵ_p , with the planes positioned relative to $\mathbf{p}_{g,1}$.

Given the curve direction at the point and the plane separation ϵ_p , these cutting planes segment the original point cloud into associated point sets as illustrated in Fig. 4.9. The interval ϵ_p is set as $\epsilon_p = \alpha \inf \|\mathbf{v}_i - \mathbf{v}_j\|$, where $\alpha \in (0, 1)$, dependent on the minimum distance between adjacent vertices in G_s .

After extracting the associated points \mathbf{Q}_i for a point $\mathbf{p}_{g,i} \in G_s$, the centre of these points is needed to estimate the centredness of $\mathbf{p}_{g,i}$.

Definition 4.14. Let $\mathbf{Q}_i \subset \mathbf{P}_o \subset \mathbb{R}^{n \times 3}$ be the associated points (Definition 4.13) of $\mathbf{p}_{g,i}$, a sampled point on curve skeleton G_s (Definition 4.8). Let \mathbf{g}_i and \mathbf{h}_i be two orthonormal vectors perpendicular to the curve direction \mathbf{u}_i at $\mathbf{p}_{g,i}$. The projected 2D points are

$$\hat{\mathbf{Q}}_i = \mathbf{Q}_i \cdot [\mathbf{g}_i, \mathbf{h}_i],$$

and similarly,

$$\hat{\mathbf{p}}_{g,i} = \mathbf{p}_{g,i} \cdot [\mathbf{g}_i, \mathbf{h}_i].$$

Let $\hat{\mathbf{q}}_c$ denote the centre of the ellipse fitted to the projected points $\hat{\mathbf{Q}}_i$. The centredness $\mathbf{c}(\mathbf{p}_{g,i})$ of point $\mathbf{p}_{g,i}$ is then defined as

$$\mathbf{c}(\mathbf{p}_{g,i}) := 1 - \frac{\|\hat{\mathbf{p}}_{g,i} - \hat{\mathbf{q}}_c\|}{0.5 \cdot (l_a + l_b)}, \quad (4.15)$$

where l_a and l_b are the lengths of the semi-major and semi-minor axes of the fitted ellipse, respectively. If $\mathbf{c}(\mathbf{p}_{g,i}) < 0$, it is reassigned to zero to ensure non-negativity.

Remark 4.7. Following Fitzgibbon et al. [54], ellipse fitting offers a robust estimate of the centre $\hat{\mathbf{q}}_c$ of the projected points $\hat{\mathbf{Q}}_i$. As shown in Fig. 4.9, ellipse-fitting centres align more closely with human visual assessment compared to barycentres computed as simple averages of all points. This robustness is further confirmed by experiments on simple geometries illustrated in Fig. 4.10, which depict the geometric centres of 2D point sets from various shapes.

As visualised in Fig. 4.8(b), points are uniformly sampled along the curve skeleton, with centredness values colour-coded. Not all curve points possess valid centredness

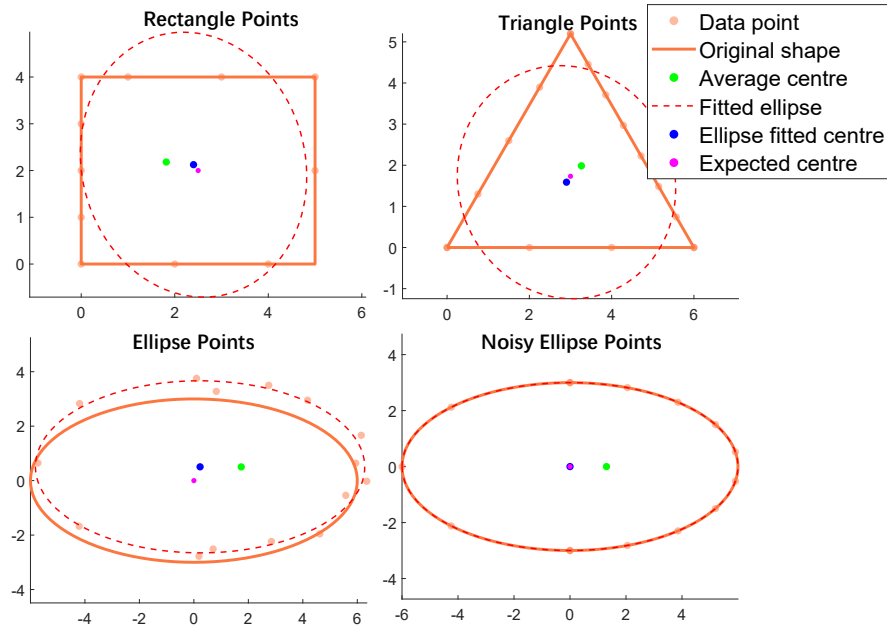


Figure 4.10: Ellipse fitting applied to point sets sampled from simple geometries, demonstrating robust centre estimation.

values; those without are marked in magenta. This set includes points coinciding with joint vertices—defined as vertices connected to more than two edges—and points with fewer than three associated points.

Definition 4.15. The overall quantified centredness of a set of skeletal points or a curve skeleton is defined as the proportion of points that are sufficiently centred. A skeletal point $\mathbf{p}_{s,i} \in \mathbf{P}_s$ is deemed sufficiently centred if its centredness value $c(\mathbf{p}_{s,i})$, as given in Eq. (4.14), satisfies $c(\mathbf{p}_{s,i}) > c^*$. Similarly, a sampled point \mathbf{q}_e on the curve skeleton G_s is sufficiently centred if its centredness value $c(\mathbf{q}_e)$, as defined in Eq. (4.15), satisfies $c(\mathbf{q}_e) > c^*$.

The threshold values c^* and c^* depend on the application-specific centredness requirements, as discussed in Section 4.4.3. Let $N_{s,c}$ denote the number of sufficiently centred points in \mathbf{P}_s or sufficiently centred sampled points in G_s , and N_s the total number of considered points. The overall centredness is then defined as

$$\mathcal{C}_s := \frac{N_{s,c}}{N_s}. \quad (4.16)$$

Remark 4.8. The centredness definition for curve skeleton points may also be applicable to a skeletal point set if the skeleton forms a thin, line-like structure. In such cases, the principal direction at each skeletal point can be estimated via Principal Component Analysis (PCA) on its neighbouring points, as detailed in Proposition 4.5.

4.3.4 Smoothness

The smoothness of a curve skeleton is characterised by the variation in the tangent direction along the curve [40], a feature which plays a crucial role in applications such as navigation and motion planning. As highlighted by Sobiecki et al. [143], both the skeletal point set manifold and the curve skeleton should ideally possess at least C^2 continuity, ensuring desirable curvature continuity.

Building upon the conceptual description of curve-skeleton smoothness introduced by Cornea et al. [40], we convert their qualitative notion into a quantitatively defined smoothness metric that enables explicit numerical assessment. This quantitative formulation, which is not provided in [40], offers a concrete means of evaluating curve-skeleton smoothness. This metric is particularly relevant for path-planning applications, which will be further discussed in Section 4.4.3.

Problem Statement

Given that smoothness significantly influences navigation and motion planning, where directional changes are critical, the focus is on assessing variations in tangent directions along the skeleton's curved branches. We assume that skeletonisation, whether producing a curve skeleton $G_s = (V, E)$ or a skeletal point set \mathbf{P}_s , yields a geometrically slender, line-like structure that approximates the medial axes of the object.

Methodology

The smoothness metric is developed by analysing the tangent vectors along the skeleton. The approach involves three key steps:

1. Estimation of tangent vectors at skeletal points or vertices.
2. Quantification of directional changes between successive tangent vectors.
3. Assignment of a smoothness score based on these directional variations.

Estimation of Tangent Vectors For a curve skeleton $G_s = (V, E)$, directional changes occur exclusively at vertices. The tangent vectors at points on the skeleton are estimated according to Definition 4.13.

In the case of skeletal point sets forming curve-like structures, the tangent vector at any point can be estimated by applying Principal Component Analysis (PCA) to its local neighbourhood, as depicted in Fig. 4.11(a). This procedure is formalised in the following proposition.

Proposition 4.5. Consider a skeletal point set $\mathbf{P}_s \subset \mathbb{R}^{n \times 3}$. For a given point $\mathbf{p}_{s,i} = [x_{s,i}, y_{s,i}, z_{s,i}]^\top \in \mathbb{R}^3$, denote its k nearest neighbours as $\{\mathbf{p}_{s,i \odot j}\}$, with $\mathbf{p}_{s,i}, \mathbf{p}_{s,i \odot j} \in \mathbf{P}_s$, for $j = 1, \dots, k$. The covariance matrix of these neighbours is computed by [75]:

$$\mathbf{C}_i = \frac{1}{k} \sum_{j=1}^k (\mathbf{p}_{s,i \odot j} - \mathbf{p}_{s,i})(\mathbf{p}_{s,i \odot j} - \mathbf{p}_{s,i})^\top.$$

The tangent vector at $\mathbf{p}_{s,i}$, denoted $\mathbf{t}_{\mathbf{p}_{s,i}}$, is identified as the eigenvector corresponding to the largest eigenvalue of \mathbf{C}_i .

Proof. The local neighbourhood of k points approximates a short linear segment. The covariance matrix \mathbf{C}_i captures this local structure, and the eigenvector with the greatest eigenvalue aligns with the direction of maximal variance, which corresponds to the tangent direction when the points approximate a line. \square

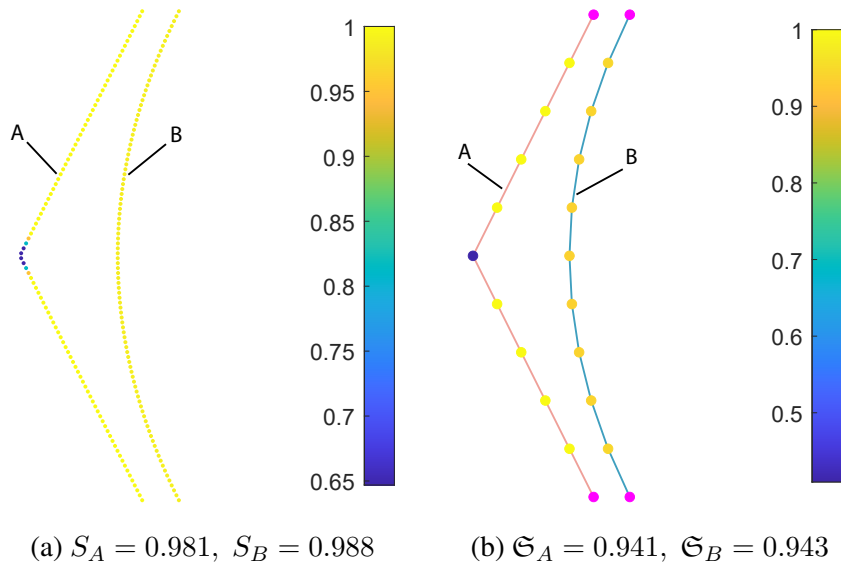


Figure 4.11: Smoothness of 2D skeletal points (a) and 2D curve skeleton (b) in different smoothness. The vertices in (b) without valid smoothness value are marked in magenta.

Measurement of Tangent Direction Variation To quantify the difference between two tangent vectors \mathbf{t}_1 and \mathbf{t}_2 , a normalised cosine distance metric is utilised (similar usage can be found in computation of surface normal difference explained by Eqs. (4.1)-(4.3)), following [182]:

$$D_n(\mathbf{t}_1, \mathbf{t}_2) = \frac{1}{\pi} \arccos \left(\frac{\mathbf{t}_1 \cdot \mathbf{t}_2}{\|\mathbf{t}_1\| \|\mathbf{t}_2\|} \right). \quad (4.17)$$

Here, D_n provides a normalised measure of angular disparity between the two tangent vectors, with values ranging from 0 (identical directions) to 1 (opposite directions).

Definition 4.16. For a point $\mathbf{p}_{s,i} \in \mathbf{P}_s$ with m neighbouring points $\{\mathbf{p}_{s,i \odot j}\}$, the local smoothness is defined as

$$s(\mathbf{p}_{s,i}) := \min_{j=1, \dots, m} \left| 1 - 2 \cdot D_n(\mathbf{t}_{\mathbf{p}_{s,i}}, \mathbf{t}_{\mathbf{p}_{s,i \odot j}}) \right|, \quad (4.18)$$

where $s(\mathbf{p}_{s,i}) \in [0, 1]$ and tangent vectors are computed as per Proposition 4.5.

For a vertex $\mathbf{v}_i \in V$ in the curve skeleton $G_s = (V, E)$, let $N_e(\mathbf{v}_i)$ denote the number of edges connected to \mathbf{v}_i . If $N_e(\mathbf{v}_i) \neq 2$, \mathbf{v}_i is an endpoint, junction, or isolated vertex,

and smoothness is assigned the maximum value 1. Otherwise, if $N_e(\mathbf{v}_i) = 2$, the smoothness at \mathbf{v}_i is

$$\mathfrak{s}(\mathbf{v}_i) := |1 - 2 \cdot D_n(\mathbf{v}_{i-1} - \mathbf{v}_i, \mathbf{v}_{i+1} - \mathbf{v}_i)|, \quad (4.19)$$

with $\mathfrak{s}(\mathbf{v}_i) \in [0, 1]$.

The overall smoothness metrics for the skeletal point set and the curve skeleton are defined as follows:

Definition 4.17. Let $\mathbf{P}_s = \{\mathbf{p}_{s,i}\}_{i=1}^{N_s}$ be a skeletal point set. Its overall smoothness is

$$S(\mathbf{P}_s) = \frac{1}{N_s} \sum_{i=1}^{N_s} s(\mathbf{p}_{s,i}). \quad (4.20)$$

For the curve skeleton $G_s = (V, E)$, define the set of vertices with degree two as $\{\hat{v}_i\}_{i=1}^{\hat{N}_v}$. The overall smoothness is the weighted average

$$\mathfrak{S}(G_s) = 1 - \frac{1}{W} \sum_{i=1}^{\hat{N}_v} w_i \cdot (1 - \mathfrak{s}(\hat{v}_i)), \quad (4.21)$$

where w_i is the total length of the two half-edges incident on \hat{v}_i , and $W = \sum_{e_i \in E} \text{length}(e_i)$ denotes the total length of all edges.

Remark 4.9. The overall smoothness of the skeletal point set is computed as the arithmetic mean of local smoothness values at all points. For the curve skeleton, smoothness is aggregated as an edge-length weighted mean of local smoothness at vertices with degree two.

Fig. 4.11 demonstrates that the local smoothness measures, as defined in Eqs. (4.18) and (4.19), effectively capture abrupt directional changes. Points or vertices exhibiting sharp bends correspond to higher smoothness values, whereas gently curving regions produce lower values.

Although the global smoothness metrics (Definition 4.17) provide a holistic measure of skeletal quality, the local smoothness characteristics are often more critical. In robotic

navigation [109] or grasping planning [165], local directional smoothness directly influences trajectory feasibility and manipulator alignment. Poor local smoothness may lead to infeasible trajectories or singularities, thereby compromising the continuity and reliability of motion planners and control algorithms.

4.4 Results & Discussion

This section covers two parts of the results with their corresponding discussion. Firstly, the statistically stable contraction findings are experimented and discussed with contraction patterns corresponding to the skeletonisation quality of different levels (Subsection 4.4.1). The second part presents a thorough evaluation and discussion of skeletonisation outcomes based on the metrics proposed in Subsection 4.4.2, followed by a discourse on the desirable properties for robotic applications in Subsection 4.4.3.

4.4.1 Results of the Laplacian-based Skeletonisation Characteristics

Here, we examine patterns in normal vector and curvature differences, alongside results from various object skeletonisation methods. The point cloud models used originate from the OmniObject3D dataset [181], which consists of real-world scanned data. Consequently, the point clouds exhibit inherent noise and imperfections. The default skeletonisation method applied is the baseline Laplacian-based approach [34]. Additionally, to validate our observations, we tested GLSkeleton (see Chapter 3), a skeletonisation technique introduced in our prior work. Multiple object point clouds from the dataset were evaluated, revealing consistent patterns; therefore, only topologically distinct objects are selected for presentation as representative examples.

To investigate geometric property variations across contraction iterations, histograms of curvature and normal vector differences were computed between each iterative con-

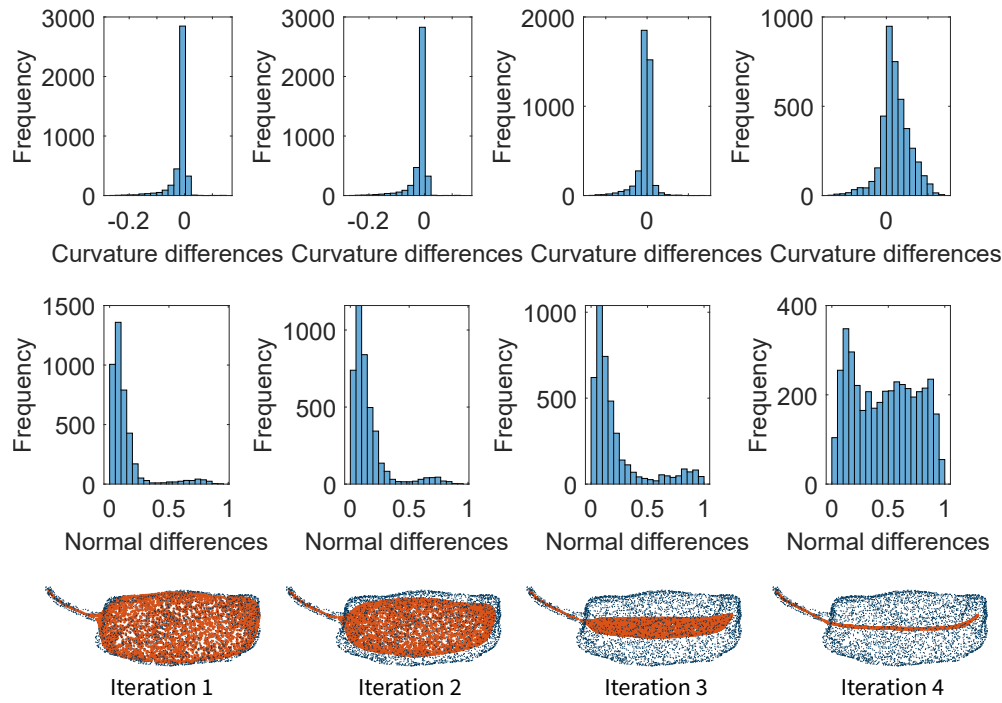


Figure 4.12: Evolution of normal vector and curvature difference distributions throughout the contraction process. The selected object is a chilli point cloud.

traction output and the original point cloud. As shown in Fig. 4.12, the evolving histogram patterns indicate the geometric convergence behaviour during contraction.

As the object’s shape deforms inward, both curvature and normal vector difference distributions evolve accordingly. Overall, these different distributions progressively become more symmetric over iterations, signalling stable convergence. Specifically, the curvature differences (first row of Fig. 4.12) remain minimal throughout, retaining a bell-shaped distribution whose normality increases steadily. This suggests that curvature converges smoothly and reliably during the entire contraction process.

In contrast, the normal vector differences (second row of Fig. 4.12) exhibit a distinct progression. Initially, the distribution is left-skewed, indicating instability in the early contraction stages. However, as contraction and attraction weights (cf. Equation 3.3) are adaptively updated, the distribution gradually becomes symmetric and unimodal, reflecting increasing stability.

These statistical changes correlate with the geometric evolution illustrated in the third row of Fig. 4.12: the contracted surface initially deviates from the original surface and extends beyond its boundaries, but is progressively pushed inward, culminating in the formation of the surface skeleton.

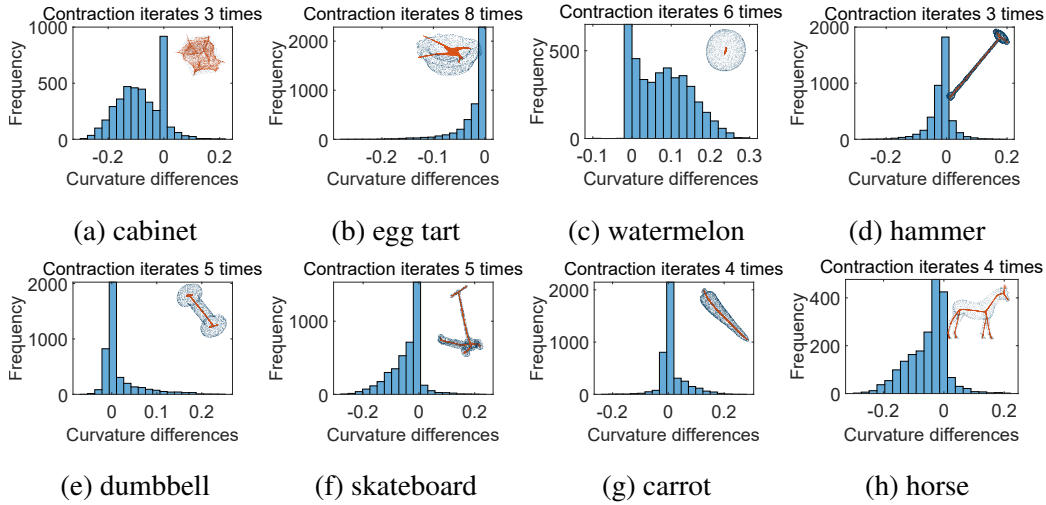


Figure 4.13: Curvature differences between the input and point cloud after final contraction. The original point cloud (grey) with the contraction results (red) is put at the top right corner of each histogram.

The evaluation of final contraction results is conducted through a comparative analysis of curvature and normal vector differences between the contracted and original point clouds, as depicted in Figs. 4.13–4.14. Notably, contraction outcomes for certain objects—such as the cabinet, egg tart, and watermelon—exhibit topological deviations from the original shapes, reflecting lower quality. Conversely, other objects display well-contracted forms with distributions more closely aligned to the stability criteria outlined in Section 4.2.

Regarding curvature differences, the distributions for the cabinet, egg tart, and watermelon show either skewed or sharply peaked profiles, with average curvature differences significantly diverging from zero. In contrast, curvature distributions for the remaining objects tend towards a bell-shaped curve centred around zero, indicating superior preservation of the original geometry.

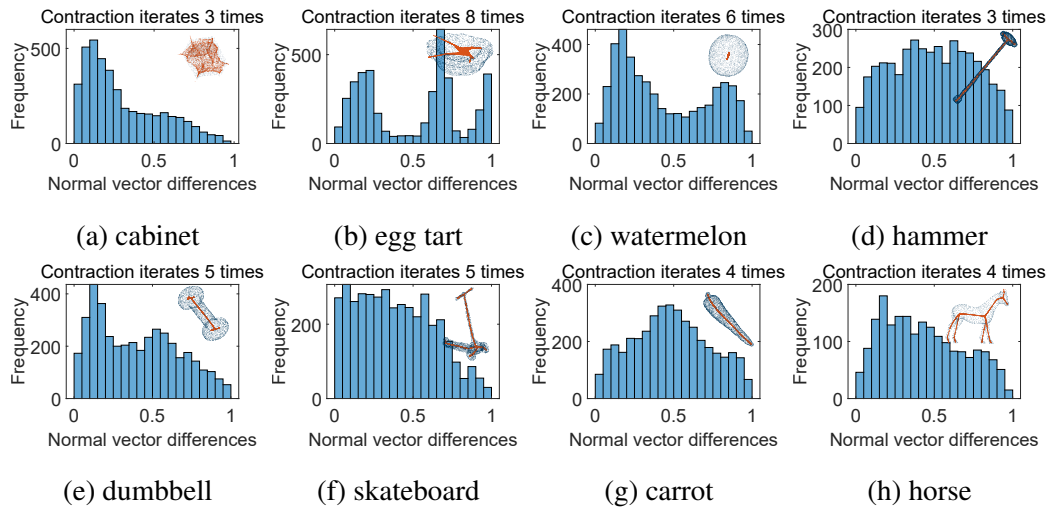


Figure 4.14: Normal vector differences between the input and contracted point clouds. Insets at the top right of each histogram show the original point cloud (grey) alongside the contraction results (red).

For normal vector differences, objects exhibiting poorer contraction similarly display less symmetric histogram patterns. The contraction process for the cabinet, for instance, terminated prematurely, resulting in a left-skewed distribution. Likewise, the left-skewed peak observed in the dumbbell’s histogram corresponds to incomplete contraction on the left side of the point cloud. Here, only frequency peaks that decline by at least 2% on both sides are considered significant, accounting for potential errors in surface normal estimation and data noise. The multiple spikes apparent in the egg tart’s distribution reflect uneven contracted surfaces, manifesting as pronounced peaks. In contrast, the skateboard, carrot, and horse exhibit histograms that are more symmetric and centred closer to 0.5 on average, consistent with their superior contraction outcomes.

It is important to note that differences in normal vectors on the skeleton surfaces lead to variations in the distributions, often producing multiple peaks within histograms. These features may provide useful insights into the topology of the contracting surface, indicating segments that have contracted improperly.

Given that the curve skeletons are extracted from the contraction results via the Laplacian-

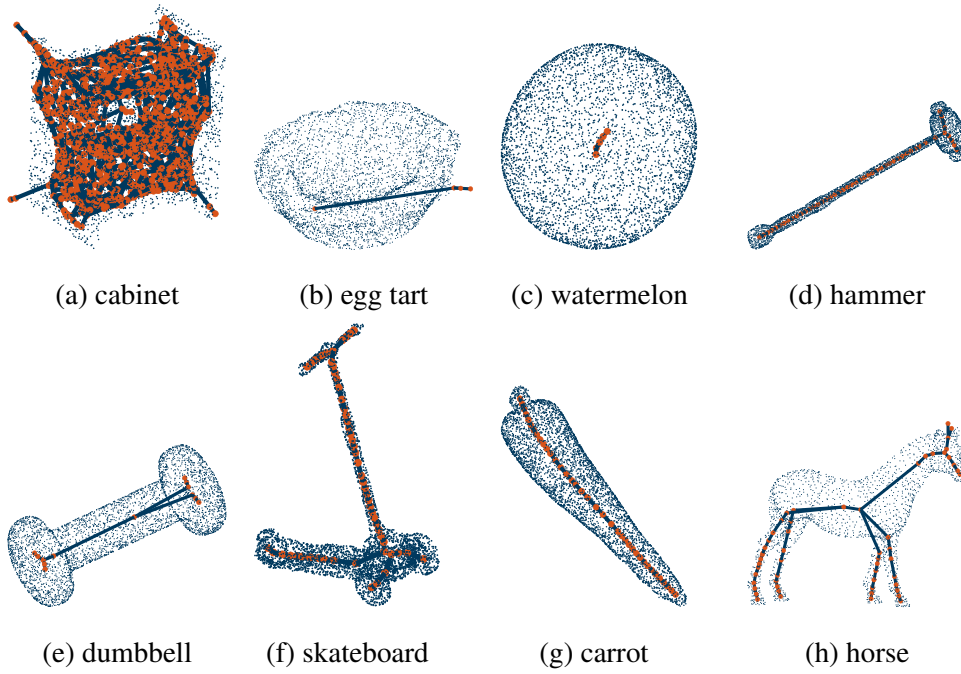


Figure 4.15: Resulting curve skeletons.

based skeletonisation method, their quality heavily depends on the contraction outcomes. Consequently, the geometric change patterns observed during contraction effectively indicate the performance of the resultant curve skeletons. As shown in Fig. 4.15, curve skeletons generated for objects such as the hammer, skateboard, carrot, and horse exhibit superior quality, consistent with their contraction results. Conversely, the skeletons for the cabinet, egg tart, and watermelon lack meaningful topological detail, reflecting the shortcomings in their respective contraction processes. Moreover, the dumbbell's incomplete contraction is evident in its skeleton as a bifurcation occurring at an unintended location.

Furthermore, we extended this analysis to the GLSkeleton approach [175], which progressively reduces points during point cloud shrinking. Since points are iteratively removed by this algorithm and differences are computed point-wise, only those points consistently retained across iterations were used for contraction pattern analysis. Fig. 4.16 demonstrates similar contraction patterns with GLSkeleton, indicating that the surface skeletons produced by this method exhibit comparable performance,

thus corroborating our propositions.

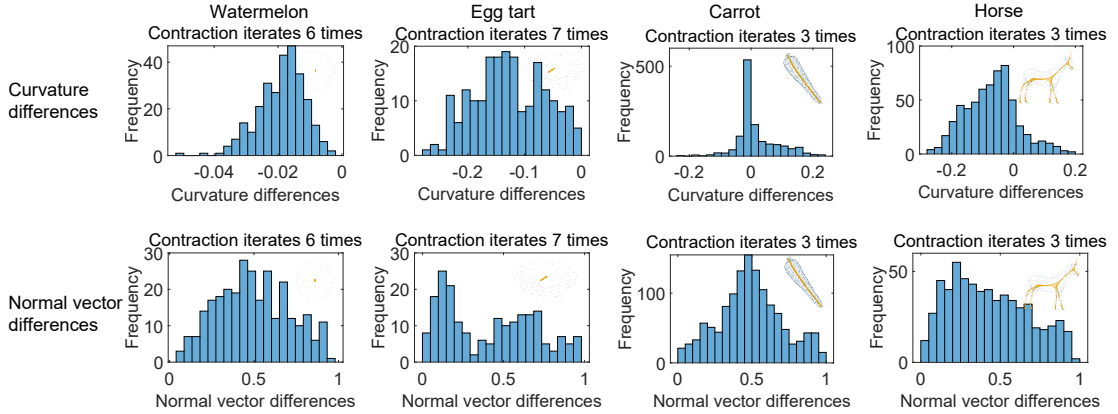


Figure 4.16: Curvature and normal vector differences between the input and contracted point cloud obtained via the GLSkeleton method. Insets at the top right of each histogram show the original point cloud (grey) alongside the contraction results (yellow).

In summary, the stable convergence of contraction and Laplacian-based curve skeletonisation is characterised by the distribution and evolution of curvature and surface normal vector differences throughout the contraction iterations. Moreover, with the adaptive control of contraction and attraction weights in the Laplacian-based method, these distribution patterns can be regulated and improved during the contraction process.

4.4.2 Quantitative Evaluation Experiments on Skeletonisation Results

The initial evaluation focused on skeletonisation results produced via the LBC method [34]. Importantly, the proposed evaluation metric is designed to be versatile and applicable to a range of skeletonisation techniques by importing the skeletal surface—whether represented as a point set or curve—into the open-access toolbox provided.

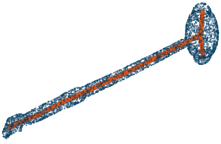
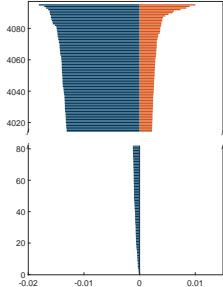
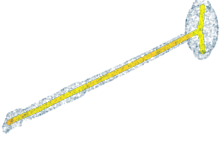
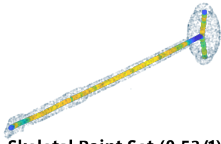
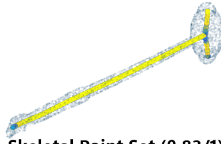
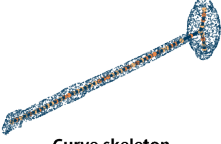
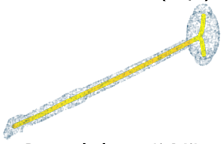
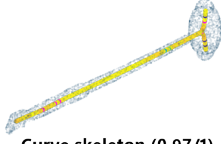

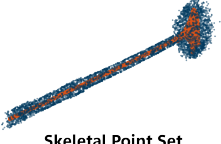
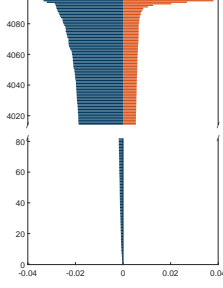
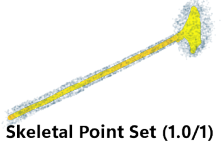
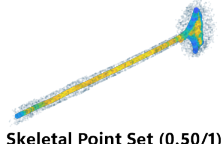
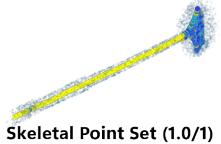
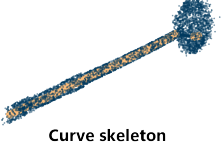
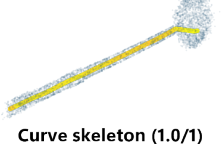
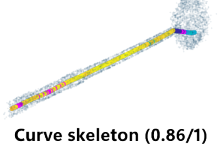
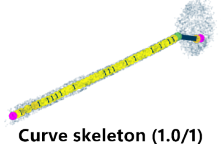
To broaden the generality and applicability of our metrics, controlled degradations

were applied to the input point clouds, including the addition of noise and increased sparsity. These controlled variations aim to test the sensitivity of the evaluation method and its capability to discriminate variations in skeletonisation quality. It is anticipated that such degradations will lead to diminished skeletonisation performance.

All experiments were conducted using MATLAB on a system equipped with an Intel i5-13500H CPU and 16 GB of RAM. The threshold parameters c^* , \mathfrak{c}^* (as defined in Definition 4.15), and β^* (from Proposition 4.4) were uniformly set to 0.75. Furthermore, the spatial extent of each point cloud was normalised such that the diagonal length of the bounding box, denoted ϵ_{\max} in Proposition 4.3, was fixed at 1.6.

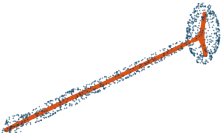
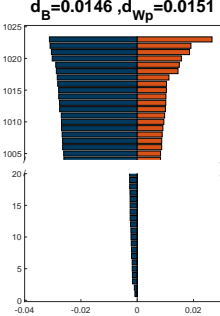
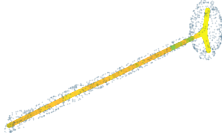
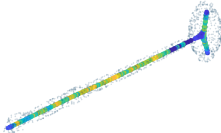
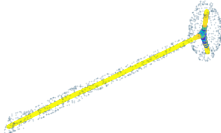
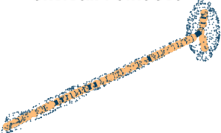
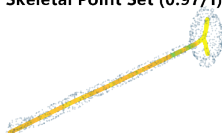
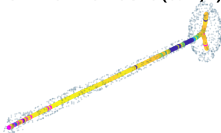
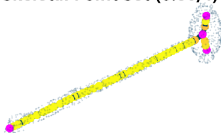
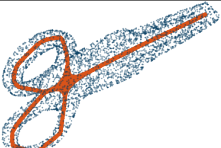
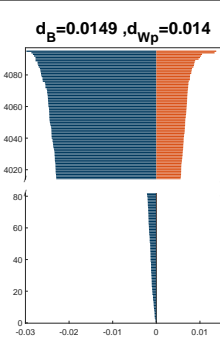

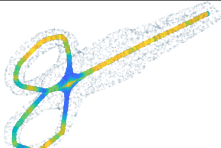
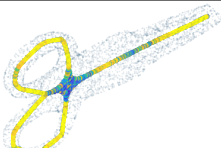

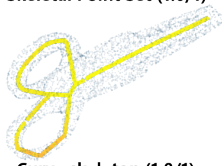
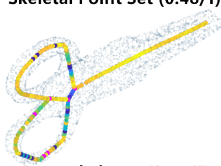
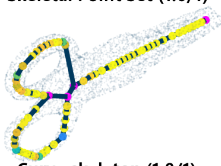
It should be noted that certain metrics, such as centredness and smoothness, may be unavailable for some point clouds under the applied computational methods. Values for these incomputable cases are considered invalid and are highlighted in magenta in Table 4.1.

Table 4.1: Evaluation of skeletonisation results for normal, noise-induced, and sparse point cloud forms, assessed using the proposed metrics. The colour indicates the local scoring along the skeleton.

Categories	Skeletonisation Results	Topological Similarity	Boundedness	Centeredness	Smoothness
4096 pts		$d_B=0.00841, d_{WP}=0.00762$ 			
					
5% noise		$d_B=0.0142, d_{WP}=0.0115$ 			
					

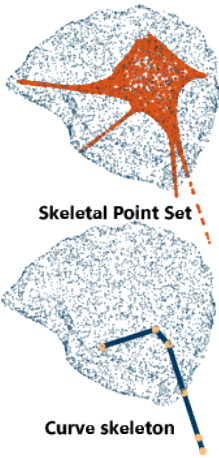
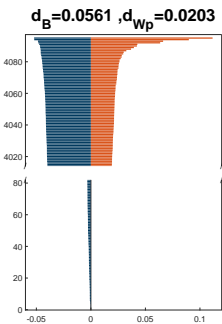
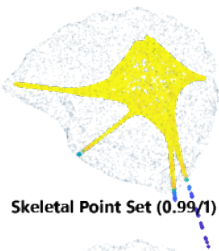
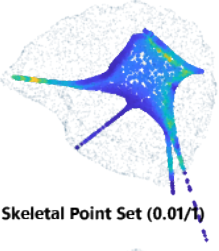
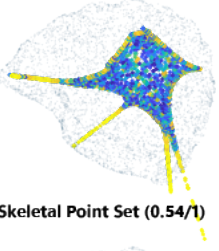
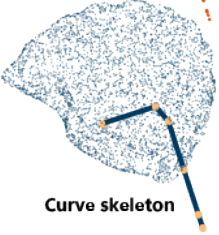
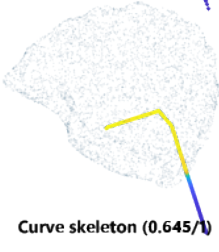
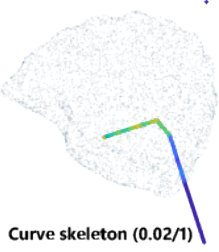
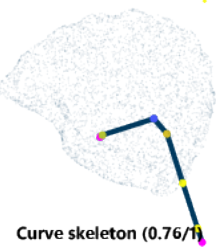
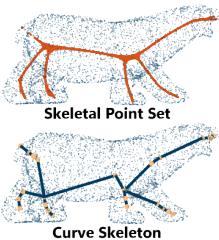
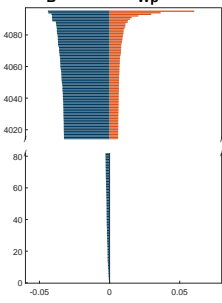
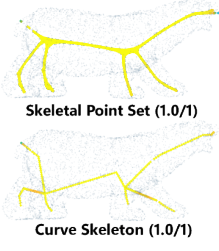
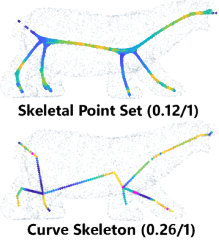
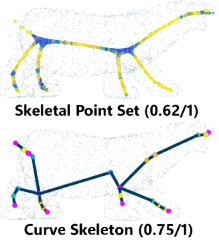

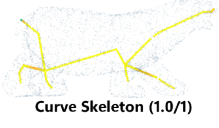


Continued on next page

Table 4.1 – Continued from previous page

Categories	Skeletonisation Results	Topological Similarity	Boundedness	Centeredness	Smoothness
1024 pts		 $d_B=0.0146, d_{Wp}=0.0151$			
					
4096 pts		 $d_B=0.0149, d_{Wp}=0.014$			
					

Continued on next page

Table 4.1 – Continued from previous page

Categories	Skeletonisation Results	Topological Similarity	Boundedness	Centeredness	Smoothness
4096 pts					
					
4096 pts					
					

Original PointCloud

Skeletal Point Set

0

0.2

0.4

0.6

0.8

1

Invalid

The study was conducted on a variety of inanimate objects and animal shapes using point cloud data sourced from multiple datasets [181, 49, 176]. Selected examples, encompassing both well-formed skeletons and those exhibiting structural deficiencies, are presented in Table 4.1 and Table 4.2. These cases are analysed to highlight the differences, illustrating how the proposed metrics facilitate their identification.

Topology preservation of the skeletal shape is quantified via a topological distance score, computed by comparing the barcodes of the persistent homology features of the original shape and the skeleton. The procedure for generating these barcodes and computing their distances is described in Proposition 4.3. As shown in Table 4.1, inferior topological alignment corresponds to higher topological distance scores, as measured by the bottleneck distance (4.7) and Wasserstein distance (4.8). Under the current normalisation, these distance values are capped by the diagonal length of the shape’s bounding box, set at 1.6. Lower distance values indicate greater topological similarity. The bottleneck distance is sensitive to significant topological changes, while the Wasserstein distance captures average topological variations.

For skeletal point sets derived from inputs of 4096 points, topology preservation is considered satisfactory if both distance measures fall below 0.02 ($d^* < 0.02$ as defined in Proposition 4.3). For instance, the skeletal point set resulting from the hammer input with 4096 points (first row of Table 4.1) demonstrates good topological preservation. Introducing 5% Gaussian noise to the hammer point cloud of 4096 points or reducing the input density to 1024 points approximately doubles both topological distances between the skeletal output and the input (rows 1-3 of Table 4.1). In the skeletal point set of the bear toy (last row), larger structural deviations under identical input conditions yield more than twice the topological distances compared to the hammer example.

In the case of the biscuit shape (fifth row), significant shape alterations along edge regions produce a notably large bottleneck distance (0.0561), whereas the Wasserstein distance remains comparatively lower (0.0203), reflecting that the majority of

the skeletal planar structure remains aligned with the input shape. Since topological similarity is based on H_0 features (connected components), the spatial neighbourhood relations influence the scores, which can cause the difference in topological distances between sparse and dense inputs (third versus first row) to be larger than might be anticipated.

Table 4.2 further corroborates that sparse and noisy inputs deteriorate topology preservation, resulting in increased topological distance scores.

Table 4.2: Quantitative results of skeletonisation evaluation. The input cave point cloud is sectioned and capped from the original data to ensure a closed shape for skeletonisation. Different resolutions of cave and synthetic tree point clouds are acquired by grid-averaged sampling.

Shapes	$d_B(4.7)/d_{W_p}(4.8) \downarrow$	$\mathcal{B}_{\mathbf{P}_s \circ \mathbf{P}_o}(4.11)/\mathcal{B}_{G_s \circ \mathbf{P}_o}(4.12) \uparrow$	$\mathcal{C}_s(\mathbf{P}_s)/\mathcal{C}_s(G_s)(4.16) \uparrow$	$S(\mathbf{P}_s)(4.20)/\mathfrak{S}_{G_s}(4.21) \uparrow$
Dumbbell (4096 pts)[181]	0.0218/0.0179	1.0/1.0	0.146/0.803	0.592/0.766
Dumbbell (5% noise)	0.0347/0.0239	0.994/1.0	0.060/0.295	0.559/0.914
Dumbbell (1024 pts)	0.0418/0.0379	0.967/0.939	0.021/0.364	0.601/0.974
Steamed Bun (4096 pts)	0.0309/0.0290	0.993/0.910	0.001/0.000	0.464/0.953
Steamed Bun (5% noise)	0.0439/0.0349	0.981/0.804	0.0/0.0	0.368/1.0
Steamed Bun (1024 pts)	0.0987/0.0340	0.979/0.983	0.008/0.023	0.347/0.854
Toy Plant (4096 pts)	0.0230/0.0214	0.993/0.965	0.277/0.358	0.583/0.755
Toy Plant (5% noise)	0.0350/0.0246	0.998/1.0	0.076/0.177	0.542/0.898
Toy Plant (1024 pts)	0.0535/0.0444	0.970/0.891	0.044/0.288	0.654/0.962
Skateboard (4096 pts)	0.0157/0.0144	1.0/1.0	0.223/0.636	0.734/0.918
Skateboard (5% noise)	0.0473/0.0173	0.999/0.995	0.185/0.509	0.687/0.953
Skateboard (1024 pts)	0.0247/0.0241	0.831/0.588	0.051/0.140	0.787/0.957
Continued on next page				

Table 4.2 – continued from previous page

Shapes	$d_B(4.7)/d_{W_p}(4.8) \downarrow$	$\mathcal{B}_{\mathbf{P}_s \circ \mathbf{P}_o}(4.11)/\mathcal{B}_{G_s \circ \mathbf{P}_o}(4.12) \uparrow$	$\mathcal{C}_s(\mathbf{P}_s)/\mathcal{C}_s(G_s)(4.16) \uparrow$	$S(\mathbf{P}_s)(4.20)/\mathfrak{S}_{G_s}(4.21) \uparrow$
Banana (4096 pts)	0.0218/0.0154	1.0/1.0	0.788/0.969	0.704/0.828
Banana (5% noise)	0.1033/0.0188	0.992/0.855	0.157/0.136	0.557/0.934
Banana (1024 pts)	0.0321/0.0322	0.982/0.895	0.165/0.147	0.559/0.844
Knife (4096 pts)	0.0170/0.0112	0.986/0.980	0.255/0.629	0.705/0.880
Knife (5% noise)	0.0209/0.0125	0.968/0.980	0.248/0.311	0.680/0.846
Knife (1024 pts)	0.0212/0.0219	0.851/0.813	0.124/0.337	0.716/0.940
Synthetic tree (8936 pts) [49]	0.0046/0.0048	0.830/0.721	0.284/0.102	0.904/0.922
Synthetic tree (5% noise)	0.0272/0.0158	0.996/1.0	0.233/0.163	0.628/0.867
Synthetic tree (1730 pts)	0.0077/0.0060	0.653/0.503	0.124/0.045	0.921/0.946
Cave (9769 pts) [176]	0.0297/0.0136	1.0/1.0	0.029/0.147	0.567/0.937
Cave (5% noise)	0.0745/0.0183	0.986/0.905	0.010/0.031	0.415/0.934
Cave (1730 pts)	0.261/0.0287	0.993/0.920	0.035/0.390	0.660/0.890

As discussed in the preceding section, local performance constitutes the most significant aspect of skeletonisation evaluation. Table 4.1 further visualises the quality of local skeletal components through colour intensity mapping based on multiple metrics. For instance, in the hammer and scissor models, all skeletal components remain well bounded, whereas in the biscuit and animal bear models, certain skeletal components extend beyond the surface boundary, as indicated by the darker colouration. Regarding sensitivity, boundedness effectively responds to distortions or deformations in the skeleton caused by sparse point clouds, thereby reflecting such perturbations with clarity.

In terms of centredness, the hammerhead exhibits decreased centredness as the input point cloud becomes sparser, demonstrated by cooler colour tones. Moreover, when noise is introduced into the input point cloud, contraction is impaired, resulting in some points within the resultant skeletal point set being less centred relative to the noisy input. The proposed metric also captures variations in centredness across distinct regions, exemplified by the scissor model. The visualisation confirms that Laplacian-based skeletonisation is highly sensitive to point cloud density, with sparse data exerting a more detrimental effect than noise, corroborating findings from prior studies. Importantly, these metrics consistently reflect such changes, thereby validating their sensitivity and utility for quality assessment and convergence stability analysis.

One limitation of the centredness metric lies in the differential sensitivity observed between the centredness evaluations of the skeletal point set and the curve skeleton, attributable to the approximation methods employed for these distinct shape representations. Centredness for skeletal points is computed based on all neighbouring points, whereas for the curve skeleton, it is derived solely from radial neighbours. As demonstrated in Table 4.1—notably in the third row—the centredness values of the skeletal point set exhibit greater variation among components and tend to be lower than those of the curve skeleton, especially near the skeleton termini. Two types of points yield invalid centredness values: firstly, curve points coincident with curve joint vertices, for which centredness is inherently ambiguous; secondly, points corresponding to skeletal

components whose input shape points are extremely sparse, resulting in ambiguous surface representation and uncertain locus determination.

Similarly, the smoothness metric effectively captures local rates of curvature change. Skeletal components characterised by abrupt directional changes—such as the legs of the animal bear (last row)—display lower smoothness values, whereas smoother regions, including the body, manifest higher values. As shown in Table 4.2 and Table 4.1, sparse and noisy inputs generally reduce boundedness and centredness scores, albeit with minor exceptions. In contrast, smoothness depends more strongly on the intrinsic skeleton structure than on input data quality, exhibiting weaker correlations with input variations, as evidenced in both tables. The smoothness metric is undefined at curve skeleton endpoints and joint vertices, where such calculations lack meaningful interpretation.

Regarding the overall evaluation, skeletonisation results exhibiting consistently strong performance across all four metrics are deemed good skeletons, though specific applications may impose further detailed requirements. While our approach effectively distinguishes skeletonisation quality in terms of shape representation, limitations persist, including sensitivity to point cloud density for topological similarity and boundedness, as well as the inability to compute local centredness in regions of extreme input sparsity, which leads to ambiguous surface representations and corresponding skeletal loci.

4.4.3 Desired Properties in Robotic Applications

Although a well-bounded skeleton is generally a fundamental requirement, the importance of other skeletonisation metrics differs depending on the specific robotic application. Topological similarity is especially vital for robotic manipulation tasks that rely on skeleton information [166, 165]. As demonstrated in Table 4.1, skeletons extracted from objects such as the hammer and scissors maintain the original topological




Figure 4.17: Skeleton-guided surface reconstruction [178].

characteristics and yield low topological distance scores based on persistent homology analyses. These results indicate that, under full object scans, the extracted skeletons preserve topology well and are therefore promising for manipulation. Since real-world robotic perception often yields partial scans, the robustness of skeletonisation under occlusions warrants further investigation. It is worth noting that the Laplacian-based skeletonisation method tends to preserve topology better for objects with predominantly cylindrical shapes.

Centeredness is another critical metric, particularly in grasping planning, where local grasp candidate generation depends on accurately centred skeletal structures [165]. For instance, in the hammer model (refer to the first row of Table 4.1), the handle and parts of the hammerhead demonstrate higher centredness values, correlating with regions of increased grasping reliability. Conversely, skeletal endpoints and joint areas exhibit reduced centredness, which may complicate grasp analysis. Thus, skeletal components with higher centredness are generally more favourable for grasp generation.

In agricultural robotics, semantic information encompassing both semantic labelling and topological structure is of paramount importance [186]. Our evaluation of a synthetic tree model (see Table 4.2) reveals that noise substantially undermines topology preservation, an effect exacerbated by the delicate, thin structure of tree branches. This sensitivity aligns with practical challenges encountered in agricultural sensing environments.

For surface reconstruction applications, prior studies such as Wu et al. [178] emphas-




Figure 4.18: Skeleton-based catheter navigation [57].

ise the role of skeleton topology in guiding the exploration of the underlying point cloud (see Fig. 4.17). In robotic navigation, however, centredness plays a comparatively minor role, while smoothness and topological similarity become more relevant [40, 57], as shown by Fig. 4.18. As shown in Table 4.2, the resolution of the input point cloud has a marked effect on topology preservation. Although higher resolutions lead to increased computational cost, excessively sparse point sets substantially reduce skeleton quality. These observations highlight the importance of the proposed metrics for evaluating the robustness of skeleton extraction across varying data conditions.

Smoothness plays a crucial role in navigation, as the number of vertices along a curve skeleton directly influences the path's continuity and smoothness. This often necessitates additional smoothing operations to enable effective path planning. For real-time applications, including both manipulation and navigation, computational efficiency is also essential. Our experiments show that evaluating skeleton quality for point clouds containing 4096 points requires approximately 34 seconds on average, with topological similarity and boundedness computations accounting for about 76% and 23% of the total processing time, respectively. Although the current computational speed is relatively slow for real-time application, we believe that the proposed evaluation method remains valuable for offline comparisons of skeletonisation techniques. Moreover, the computation time may be further reduced in the future through parallelisation.

In summary, firstly, the investigated stable contraction analysis provides a statistical

criterion for a stable contraction-based skeletonisation. Secondly, the proposed evaluation framework provides a comprehensive assessment of skeletal point sets and curve skeletons from multiple perspectives. However, sensitivity to point cloud density impacts the boundedness and topological similarity metrics, and limitations persist for centredness evaluation in scenarios of extreme sparsity. Furthermore, the relative significance of each metric varies by application, suggesting that skeletonisation methods may perform differently depending on the task. Overall, this work lays the foundation for further research into quantifying skeleton quality to improve robotic performance across domains such as grasping and navigation.

4.5 Conclusion

This chapter presents two main contributions:

Firstly, we investigate the stability of skeletonisation using contraction-based methods, particularly the Laplacian-based approach, by statistically analysing changes in curvature and surface normals during stable contraction phases. The proposed metrics, focusing on surface normal and curvature variations, serve as indicators of contraction stability. Although the observed distribution patterns remain theoretical at this stage, the subsequent quantitative evaluation in this chapter establishes a preliminary basis for defining criteria for high-quality skeletons.

Secondly, we introduce a novel, formal geometric evaluation metric for skeletonisation assessment, drawing inspiration from existing metrics but extending them into a comprehensive numerical framework. Unlike conventional subjective visual assessments, this method quantifies skeleton performance across multiple geometric facets. Evaluations conducted on real scanned point clouds under varying resolutions and noise conditions demonstrate that the proposed metric effectively captures aspects including topology preservation, boundedness, centredness, and smoothness. Additionally, we examine skeletonisation performance across diverse robotic application scenarios and

discuss the relative importance of different metrics within these contexts.

To the best of our knowledge, this work is the first to provide a quantitative framework for skeleton quality evaluation supported by application-driven analysis. Nevertheless, several limitations remain. The metrics exhibit sensitivity to point cloud resolution, centredness estimation becomes unreliable for extremely sparse data, and the current computational speed is insufficient for real-time scenarios. Moreover, some metrics rely on specific modelling assumptions, which may limit their applicability across all skeletonisation methods. These limitations indicate that further refinement is required to broaden the robustness and generality of the framework.

Future work will focus on improving computational efficiency for robotic applications such as grasp planning, enhancing robustness to variations in point cloud density, and extending the evaluation framework to accommodate a wider range of skeleton definitions. An additional direction is the development of more rigorous validation methodologies for the proposed metrics.

Chapter 5

Comprehensive Analysis of Skeleton-guided Rolling Contact Kinematics from Arbitrary Point Clouds

5.1 Introduction

The skeleton is a powerful tool in shape analysis and has consequently found wide applicability across geometry processing, robotics, and computer vision. Its ability to capture the intrinsic structure of a shape in a compact and topologically meaningful form makes it particularly valuable when working with irregular or discrete data. This capability is especially relevant for rolling contact kinematics, an area in which existing formulations rely heavily on continuous surface representations. Classical rolling-contact models [103] assume smooth, differentiable surfaces, whereas real-world objects are often available only as noisy, irregular point clouds. By abstracting a point cloud into a low-dimensional structural representation, the skeleton provides a bridge between discrete sensing data and the geometric requirements of rolling-contact formulation. Moreover, the structural abstraction offered by the skeleton facilitates the surface parametrisation of irregular rigid bodies, a task that is otherwise challenging

when working directly with raw point clouds. This motivates our objective: to develop a rolling-contact kinematics model that incorporates skeleton information extracted from point cloud data, enabling continuous surface interactions to be approximated from discrete representations.

Rolling contact is a fundamental motion primitive in robotics, with applications ranging from spherical mobile robots to dexterous in-hand manipulation [94, 152]. Montana’s model [103] remains the most influential framework for describing rolling interactions, yet its reliance on continuous surfaces and precise contact localisation limits its applicability in point-cloud-based perception. This disconnect highlights the need for rolling-contact formulations capable of operating on irregular, non-differentiable rigid bodies reconstructed from discrete data.

To address this gap and challenge, this chapter combines advances in point cloud skeletonisation [34, 175], differential geometry, and Fourier-based surface reconstruction to develop a geometry-driven kinematic framework that is robust to sparse and incomplete data. With the surface reconstruction, we provide two approaches, local and global reparametrisation, to solve surface coordinate mapping. Rolling contact kinematics models based on both formulations are derived and simulated using physical object shapes and scanned environments for reliable validation. Our method bridges discrete perception and continuous control, enabling effective rolling contact modelling even without explicit contact information.

In this chapter, we first address the challenge of modelling discrete point cloud surfaces through curve-skeleton-guided surface sectioning and reconstruction using Fourier series-based curve fitting, as detailed in Section 5.2. We then present both local and global generalised surface formulations, along with the corresponding rolling contact kinematics models in Section 5.3. The proposed method is further validated through simulations of a rolling sphere interacting with objects and environments represented by real-world point clouds, discussed in Section 5.4. Finally, we summarise our findings and contributions in Section 5.5.

5.2 Skeletonisation, Slicing and Curve Fitting

Point cloud data is widely utilised as a practical approach to capture the real-time geometry of physical objects, enabling robotic systems to perceive their surroundings for tasks such as manipulation and navigation. Typically, 3D sensing devices, including RGB-D cameras, acquire only surface information, representing an object's shape as a discrete collection of points, known as a point cloud. Although point clouds are inherently discrete, modelling continuous rolling contact motions requires a continuous representation. To address this, we employ the curve skeleton of the object as an intermediate abstraction that connects the discrete point cloud data to continuous rolling contact kinematics. This curve skeleton acts as a *skeleton frame* within robotic applications, offering a structured foundation for motion modelling.

This section begins by introducing the fundamental concepts of skeletonisation from point cloud data, followed by a description of our curve-fitting approach based on the skeleton, and concludes with an overview of the complete computational procedure.

5.2.1 Point Cloud Skeletonisation

A shape's skeleton provides an abstract topological summary that encapsulates its fundamental structure [127]. In rolling contact systems, it is essential to define a differentiable manifold to facilitate dynamic analysis. Curve skeletonisation fulfils this role by representing the shape's intrinsic topological pathways within 3D space, thus enabling sectional examination of point cloud models [165]. Typically, a curve skeleton consists of a connected set of medial points that capture the shape's topology, as defined by Definition 3.2

On the basis of the curve skeleton definition and assuming the skeleton of convex shapes can be accurately extracted from their point cloud representations, we utilise a Laplacian-based skeletonisation method to derive curve skeletons from object point clouds (please refer to [34, 175] and Chapter 3). In this framework, the curve skeleton

serves two primary functions. Firstly, it provides a reference structure that facilitates the tracking of rolling motion along the skeleton's trajectory. Secondly, it supports sectional analysis by enabling slicing orthogonal to the local tangent vector of the skeleton curve. Each resulting slice produces a set of points approximating a closed curve on the object surface, which is subsequently approximated using a Fourier series to yield a continuous mathematical description of the cross-sectional geometry.

5.2.2 Slicing and Curve Fitting

Although existing methods allow straightforward extraction of curve skeletons, the skeleton obtained is typically a discrete set of vertices connected by straight line segments, resulting in a non-differentiable representation, as shown in Fig. 5.1(a). To facilitate continuous and differentiable modelling necessary for subsequent parametrisation, we apply a curve fitting procedure. This process is illustrated in Fig. 5.2, which will be described in detail subsequently.

Under the assumption that the curve skeleton either contains no branching points or can be segmented into separate branches, each branch is modelled independently. We employ a Fourier series fitting approach, adaptively selecting the order of the series to minimise the mean squared error (MSE) between the fitted curve and the original discrete skeleton points. The fitted curve for a branch is represented as

$$\theta_g(v_o) = (x, y, z) = (p_x(t), p_y(t), p_z(t)),$$

where the spatial coordinates (x, y, z) are parameterised by the variable t .

The fitting process uses the original skeleton vertices as input data. For locally generalised surfaces (see Subsection 5.3.1), the parameter t corresponds to the arc length scaled by a factor α , i.e., $t = \alpha v_o$. For globally generalised surfaces (see Subsection 5.3.2), $t = v_g = v_o$, and $\theta_g(v_o) = f_g(v_o)$ defines the skeletal mapping function. To reduce computational redundancy and improve fitting quality, the parameter t is chosen to coincide with one of the Cartesian coordinates of the point cloud vertices.

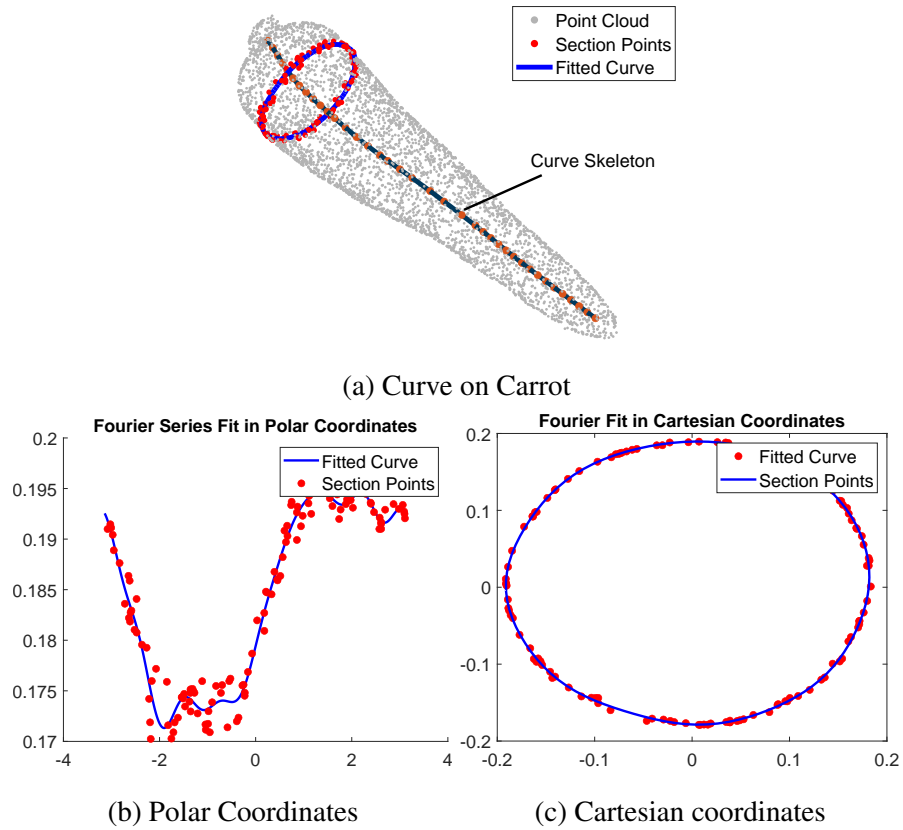


Figure 5.1: Illustration of point cloud slicing and curve fitting. (a) shows the extracted section points from the point cloud. (b) and (c) demonstrate the corresponding fitted curves in the transformed 2D coordinate system and in Cartesian coordinates, respectively.

This choice streamlines the slicing operation and supports the parametrisation of the shape by associating each slice curve with a position along the skeleton.

Fig. 5.1(a) exemplifies this approach by sectioning a carrot-shaped point cloud along the fitted skeleton at fixed intervals. Each section is formed by slicing the point cloud orthogonally to the local tangent of the skeleton curve.

Following curve fitting, the point cloud branch can be sliced continuously along the skeleton curve, enabling the extraction of sectional point sets. These points, highlighted in red in Fig. 5.1(a), approximate the surface contour at each slice. Provided the section thickness is sufficiently small, the local surface can be approximated as a

two-dimensional closed curve. The points are projected onto a plane perpendicular to the skeleton tangent at the slice location.

Since directly modelling closed curves from discrete points is challenging, the projected points are subsequently fitted using a Fourier series representation, yielding a smooth and continuous description of the sectional shape, as shown in Figs. 5.1(b) and 5.1(c).

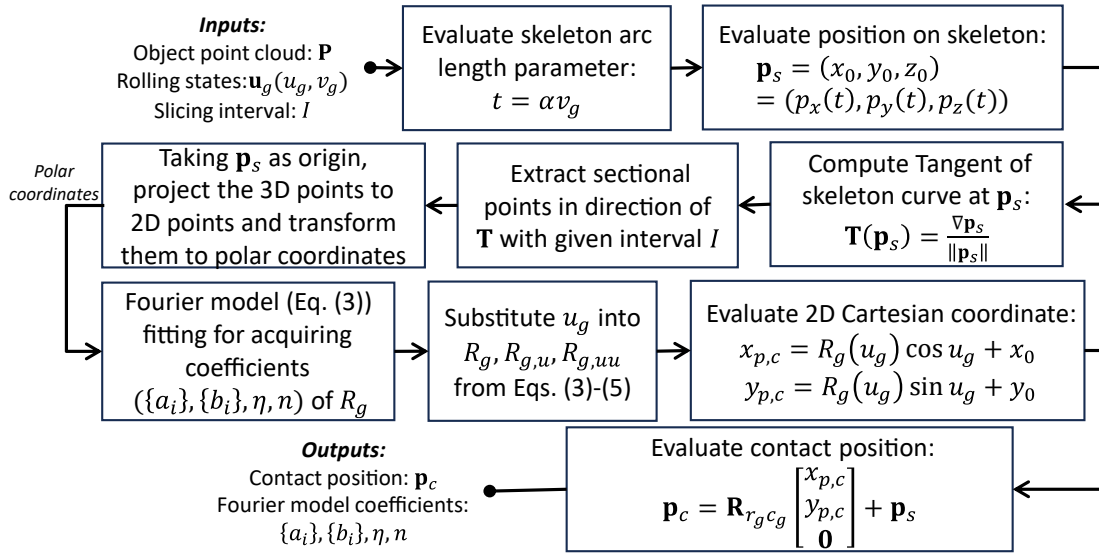


Figure 5.2: Block diagram illustrating the slicing procedure and Fourier-based curve fitting method.

We proceed to approximate the closed curve formed by the two-dimensional sectional points extracted from the point cloud. Following the approach of [117], a Fourier series representation is employed relative to the skeleton reference frame C_g , as depicted in Fig. 5.2 and illustrated in Fig. 5.3, where the section points are expressed in polar coordinates. The radius of the section, denoted R_o , at the slice angle indexed by v_o is modelled as

$$R_o = a_0 + \sum_{i=1}^n (a_i \sin(i\eta u_o) + b_i \cos(i\eta u_o)), \quad (5.1)$$

where u_o is the angular coordinate in the polar system relative to the centre of the section. The Fourier coefficients $\{a_0, a_i, b_i\}$ are obtained through curve fitting, with

n representing the number of harmonics included and η the fundamental frequency. The order n is selected adaptively to optimise fitting accuracy by minimising the loss function defined as

$$L(R_o) = w_1 E_{\text{RMS}} + w_2 \|\mathbf{p}_s - \mathbf{p}_e\|, \quad (5.2)$$

where E_{RMS} denotes the root mean square error between the fitted model $R_o(u_o)$ and the data points, and $\mathbf{p}_s, \mathbf{p}_e$ are the start and end points of the closed contour in polar coordinates. The terms correspond to the fitting error and the closure constraint, weighted by factors w_1 and w_2 , respectively. The fitting process is considered to have converged when $L(R_o(u_o))$ falls below a threshold ϵ . In practice, the polar coordinate data are duplicated to reinforce periodicity and enhance fitting precision.




Figure 5.3: Parametrisation and frame definitions for rolling contact kinematics between a rolling sphere and a geometrically fitted surface.

This Fourier fitting of the projected section points successfully captures intricate curvature features of the surface represented by the point cloud, whilst mitigating noise to produce a smooth, continuous curve.

5.3 Kinematics Model with Generalised Semi-Convex Surface

5.3.1 Modelling by Locally Generalised Semi-Convex Surface

This subsection presents a continuous local two-dimensional manifold coordinate system designed to facilitate the derivation of rolling contact kinematics from a curve-fitted point cloud representation. To simplify the analysis, the rolling object is modelled as a sphere interacting with a locally defined semi-convex 2D manifold. This simplification reduces the complexity of the rolling-contact analysis by enabling local surface interactions to be treated using well-defined curvature and normal properties. In practice, it assumes that the shape does not contain highly folded, self-intersecting, or tightly curved regions that would violate local semi-convexity. While this assumption supports a tractable formulation, it also limits the generalisation of the model to more complex or strongly non-convex geometries.

Consider a spin-rolling sphere Σ_S moving over a geometrically fitted surface Σ_O , as illustrated in Fig. 5.3. The object is associated with two fixed frames: a global frame C_{r_o} and a local frame $C_{l_o}(t)$, as well as two moving frames: a skeleton frame C_g moving along the skeleton curve and a contact frame C_{c_o} moving along the contact path. Notably, C_{c_o} coincides with $C_{l_o}(t)$ at time t . Similarly, the sphere has corresponding frames $C_{r_s}, C_{l_s}, C_{c_s}$, where C_{l_s} and C_{c_s} are Gauss frames as defined by Montana et al. [103], whereas C_{l_o} and C_{c_o} are not.

The rolling contact is characterised by the sphere's angular velocity $\dot{\mathbf{u}}_s$ and the geometric surface's angular velocity $\dot{\mathbf{u}}_o$, along with a relative spin angle ψ between their respective contact coordinate frames. The contact points on both surfaces, denoted Σ_{CS} and Σ_{CO} , coincide precisely. By treating the sphere radius as a generalised coordinate

for a symmetric surface, the parametric mappings are defined as:

$$\begin{aligned} f_o : U_O \rightarrow \mathbb{R}^3, \quad \mathbf{u}_o(u_o, v_o) &\mapsto [R_o(u_o) \cos u_o, R_o(u_o) \sin u_o, \alpha v_o]^T, \\ f_s : U_S \rightarrow \mathbb{R}^3, \quad \mathbf{u}_s(u_s, v_s) &\mapsto [R_s \cos u_s \cos v_s, -R_s \cos u_s \sin v_s, R_s \sin u_s]^T, \end{aligned} \quad (5.3)$$

where the domains are $U_G = \{(u_o, v_o) \mid -\pi < u_o < \pi, 0 < v_o < \pi\}$ and $U_S = \{(u_s, v_s) \mid -\frac{\pi}{2} < u_s < \frac{\pi}{2}, -\pi < v_s < \pi\}$, and $\alpha = L_s/\pi$ with L_s representing the total length of the skeleton. The function R_o is the fitted radius profile from Eq. (5.1), and v_o indexes the local slice for f_o .

We adopt the *Gauss frame* formalism [103] to describe rolling kinematics. The Gauss frame basis vectors at parameter \mathbf{u} are given by

$$\mathbf{x}(\mathbf{u}) = \frac{f_u(\mathbf{u})}{\|f_u(\mathbf{u})\|}, \quad \mathbf{y}(\mathbf{u}) = \frac{f_v(\mathbf{u})}{\|f_v(\mathbf{u})\|}, \quad \mathbf{z}(\mathbf{u}) = \mathbf{x}(\mathbf{u}) \times \mathbf{y}(\mathbf{u}), \quad (5.4)$$

where f_u and f_v denote partial derivatives of f with respect to u and v . Due to the irregularity of the object surface Σ_G , no explicit outward normal function g exists; hence, \mathbf{z}_o is computed via the cross product $\mathbf{x}_o \times \mathbf{y}_o$.

Explicitly, the orthonormal Gauss frame on Σ_G is

$$\begin{aligned} \mathbf{x}_o(\mathbf{u}_o) &= \frac{1}{\sqrt{R_o^2 + R_{g,u}^2}} \begin{bmatrix} R_{o,u} \cos u_o - R_o \sin u_o \\ R_{o,u} \sin u_o + R_o \cos u_o \\ 0 \end{bmatrix}, \\ \mathbf{y}_o(\mathbf{u}_o) &= \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \\ \mathbf{z}_o(\mathbf{u}_o) &= \frac{1}{\sqrt{R_o^2 + R_{o,u}^2}} \begin{bmatrix} R_{o,u} \sin u_o + R_o \cos u_o \\ R_o \sin u_o - R_{o,u} \cos u_o \\ 0 \end{bmatrix}. \end{aligned} \quad (5.5)$$

Introducing the metric tensor $\mathbf{M}_o = \text{diag}(\|f_{o,u}\|, \|f_{o,v}\|)$, curvature matrix \mathbf{K}_o , and

torsion vector \mathbf{T}_o as

$$\mathbf{M}_o = \text{diag} \left(\sqrt{R_o^2 + R_{o,u}^2}, \alpha \right), \quad (5.6)$$

$$\mathbf{K}_o = \begin{bmatrix} k_{nu}^g & \tau_{nu}^g \\ \tau_{nv}^g & k_{nv}^g \end{bmatrix}, \quad \mathbf{T}_o = \begin{bmatrix} k_{ou}^g & k_{ov}^g \end{bmatrix}, \quad (5.7)$$

where the nonzero curvature term is

$$k_{nu}^g = \frac{R_o^2 - R_o R_{o,uu} + 2R_{o,u}^2}{(R_o^2 + R_{o,u}^2)^{3/2}},$$

and all torsion and other curvature terms are zero: $\tau_{nu}^g = \tau_{nv}^g = k_{nv}^g = k_{ou}^g = k_{ov}^g = 0$.

Analogously, for the sphere surface Σ_S , the metric, curvature and torsion matrices are

$$\mathbf{M}_s = \begin{bmatrix} R_s & 0 \\ 0 & R_s \cos u_s \end{bmatrix}, \quad \mathbf{K}_s = \begin{bmatrix} \frac{1}{R_s} & 0 \\ 0 & \frac{1}{R_s} \end{bmatrix}, \quad \mathbf{T}_s = \begin{bmatrix} 0, & -\frac{\tan u_s}{R_s} \end{bmatrix}. \quad (5.8)$$

Following [103], the rolling contact kinematics equations are expressed as

$$\begin{aligned} \dot{\mathbf{u}}_o &= \mathbf{M}_o^{-1}(\mathbf{K}_o + \tilde{\mathbf{K}}_s)^{-1} \left(\begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} - \tilde{\mathbf{K}}_s \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\mathbf{u}}_s &= \mathbf{M}_s^{-1} \mathbf{R}_\psi (\mathbf{K}_o + \tilde{\mathbf{K}}_s)^{-1} \left(\begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} + \mathbf{K}_o \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\psi} &= \omega_z + \mathbf{T}_o \mathbf{M}_o \dot{\mathbf{u}}_o + \mathbf{T}_s \mathbf{M}_s \dot{\mathbf{u}}_s, \quad v_z = 0, \end{aligned} \quad (5.9)$$

where $\boldsymbol{\omega}_{\text{rel}} = [\omega_x, \omega_y, \omega_z]^T$ and $\mathbf{V}_{\text{rel}} = [v_x, v_y, v_z]^T$ denote the relative angular and linear velocities in the orthonormal Gauss frame at the contact point.

Substituting Eqs. (5.7) and (5.8) into Eq. (5.9), the final kinematic equations for a sphere rolling on a rigid object surface are derived as

$$\begin{aligned} \dot{\mathbf{u}}_o &= \begin{bmatrix} \frac{R_s S_1}{S_1^{3/2} + R_s S_2} & 0 \\ 0 & \frac{R_s}{\alpha} \end{bmatrix} \left(\begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} - \begin{bmatrix} \frac{1}{R_s} & 0 \\ 0 & \frac{1}{R_s} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\mathbf{u}}_s &= \begin{bmatrix} \frac{\cos \psi S_1^{3/2}}{S_1^{3/2} + R_s S_2} & -\sin \psi \\ \frac{-\sin \psi S_1^{3/2}}{\cos u_s (S_1^{3/2} + R_s S_2)} & -\frac{\cos \psi}{\cos u_s} \end{bmatrix} \left(\begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} + \begin{bmatrix} S_1^{-3/2} S_2 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\psi} &= \omega_z + \begin{bmatrix} 0 & 0 \end{bmatrix} \dot{\mathbf{u}}_o + \begin{bmatrix} 0 & -\sin u_s \end{bmatrix} \dot{\mathbf{u}}_s, \end{aligned} \quad (5.10)$$

where

$$S_1 = R_o^2 + R_{o,u}^2, \quad S_2 = R_o^2 - R_o R_{o,uu} + 2R_{o,u}^2.$$

For the spin-rolling motion, define the state vector as $\mathbf{x} = [u_o, v_o, u_s, v_s, \psi]^T$. Under the no-slip condition, where the relative linear velocity satisfies $\mathbf{V}_{\text{rel}} = 0$, the kinematic equations can be expressed in the form

$$\dot{\mathbf{x}} = \mathbf{o}_1 \omega_x + \mathbf{o}_2 \omega_y + \mathbf{g}_3 \omega_z, \quad (5.11)$$

with the vector fields given by

$$\begin{aligned} \mathbf{g}_1 &= \left[0, \frac{R_s}{\alpha}, -\sin \psi, -\frac{\cos \psi}{\cos u_s}, \tan u_s \cos \psi \right]^T, \\ \mathbf{g}_2 &= \frac{1}{R_s S_2 + S_1^{3/2}} \left[-R_s S_1, 0, -S_1^{3/2} \cos \psi, \frac{S_1^{3/2} \sin \psi}{\cos u_s}, \frac{S_1^{3/2} \sin \psi}{\cos u_s} \right]^T, \\ \mathbf{g}_3 &= \left[0, 0, 0, 0, 1 \right]^T, \end{aligned}$$

where singularities arise only if the denominator satisfies $R_s S_2 + S_1^{3/2} = 0$.

We introduce the distribution

$$Q(\mathbf{x}) = \{\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3, [\mathbf{g}_1, \mathbf{g}_3], [\mathbf{g}_2, \mathbf{g}_3]\},$$

where $[\cdot, \cdot]$ denotes the Lie bracket of vector fields. According to the controllability criterion in [81], the system is locally controllable if the dimension of the distribution Q satisfies $\dim(Q) = 5$.

By evaluating the determinant of the matrix formed by these vector fields, the rank condition is explicitly given by

$$\det(Q) = -\frac{R_s S_1^2}{\alpha \cos u_s \left(S_1^{3/2} + R_s S_2 \right)^2}. \quad (5.12)$$

Since $\det(Q) \neq 0$ holds in general, the distribution Q is full rank, implying that the system is controllable under the prescribed spin-rolling kinematics with no slip.

5.3.2 Modelling by a Globally Generalised Semi-Convex Surface

This section introduces a continuous two-dimensional manifold coordinate system that facilitates the determination of rolling contact kinematics via a curve-fitted point cloud model. To this end, a locally semi-convex two-dimensional manifold is constructed, with the rolling object assumed to be a sphere.

The frame definition and geometric definitions are similar to the locally generalised approach shown in Fig. 5.3.

Unlike the locally defined approach, the globally generalised surface considers the parametrisation without simplifying the problem by considering the surface globally with the coordinate mappings defined as:

$$\begin{aligned}
 f_g : U_g &\rightarrow \mathbb{R}^3, & v_g &\mapsto [p_x(v_g), p_y(v_g), p_z(v_g)]^T, \\
 f_o : U_o &\rightarrow \mathbb{R}^3, & \mathbf{u}_o(u_o, v_o) &\mapsto f_g(v_o) + \mathbf{R}_{rog} \cdot \begin{bmatrix} 0, & R_o(u_o, v_o) \cos u_o, & R_o(u_o, v_o) \sin u_o \end{bmatrix}^T, \\
 f_s : U_s &\rightarrow \mathbb{R}^3, & \mathbf{u}_s(u_s, v_s) &\mapsto \begin{bmatrix} R_s \cos u_s \cos v_s, & -R_s \cos u_s \sin v_s, & R_s \sin u_s \end{bmatrix}^T,
 \end{aligned} \tag{5.13}$$

where f_g maps the parameter v_g to points on the skeleton curve. The object surface mapping f_o is obtained by substituting $v_o = v_g$ into $f_g(v_g)$ and applying the fitted curve model of the sectional object points, as illustrated in Fig. 5.2 and Fig. 5.1. The sphere mapping f_s follows Montana's definition.

Since the sphere's coordinate system in this global formulation closely resembles the local coordinate system, the metric tensor \mathbf{M}_s , curvature form \mathbf{K}_s , and torsion form remain as defined in Eqs. (5.8).

The coordinate system (f_s, U_s) is orthogonal, allowing direct computation of the corresponding form matrices. However, the object surface parametrised by the skeleton curve is generally non-orthogonal, rendering direct calculation of the form matrices infeasible. To overcome this, differential-geometric quantities are approximated via

frame projections as follows:

$$\begin{aligned}\tilde{\mathbf{M}}_o &= \begin{bmatrix} \mathbf{x}_o(\mathbf{u}_o(t)) & \mathbf{y}_o(\mathbf{u}_o(t)) \end{bmatrix}^T \cdot \begin{bmatrix} f_{o,u}(\mathbf{u}_o) & f_{o,v}(\mathbf{u}_o) \end{bmatrix}, \\ \tilde{\mathbf{K}}_o &= \begin{bmatrix} \mathbf{x}_o(\mathbf{u}_o(t)) & \mathbf{y}_o(\mathbf{u}_o(t)) \end{bmatrix}^T \cdot \begin{bmatrix} \mathbf{z}_{o,u}(\mathbf{u}_o) & \mathbf{z}_{o,v}(\mathbf{u}_o) \end{bmatrix}, \\ \tilde{\mathbf{T}}_o &= \mathbf{y}_o(\mathbf{u}_o(t))^T \cdot \begin{bmatrix} \mathbf{z}_{o,u}(\mathbf{u}_o) & \mathbf{z}_{o,v}(\mathbf{u}_o) \end{bmatrix},\end{aligned}\quad (5.14)$$

where $\{\mathbf{x}_o, \mathbf{y}_o, \mathbf{z}_o\}$ denote the axes of the moving contact frame C_{c_o} defined on the object surface. Here, $\mathbf{z}_o = \mathbf{n}_o$ is the unit normal vector at the contact point, computed via principal component analysis on the k -nearest neighbours and the contact point itself within the object point cloud [65]. The vector \mathbf{x}_o is the tangent vector at the contact point on the cross-sectional contour. Although \mathbf{x}_o could be derived from the fitted Fourier curve in Eq. (5.1), the accuracy of such fitting is not guaranteed. Instead, \mathbf{x}_o is obtained as the cross product of the unit tangent vector of the skeleton curve, \mathbf{x}_g , and the normal vector \mathbf{z}_o , utilising the orthogonality condition $\mathbf{x}_g \perp \mathbf{x}_o$, ensured by the slicing strategy depicted in Fig. 5.2. The vector \mathbf{y}_o is then defined as $\mathbf{y}_o = \mathbf{z}_o \times \mathbf{x}_o$.

With reference to the gauss frame given by Eq. (5.4), the coordinate axes of the skeleton frame C_g and the object contact frame C_{c_o} are thus given by

$$\begin{aligned}\mathbf{x}_g(v_g) &= \frac{\dot{\mathbf{p}}_g}{\|\dot{\mathbf{p}}_g\|}, \quad \mathbf{y}_g(v_g) = \frac{\mathbf{x}_g \times \mathbf{m}_g}{\|\mathbf{x}_g \times \mathbf{m}_g\|}, \quad \mathbf{z}_g(v_g) = \frac{\mathbf{x}_g \times \mathbf{y}_g}{\|\mathbf{x}_g \times \mathbf{y}_g\|}, \\ \text{where } \mathbf{m}_g &\neq \mathbf{x}_g, \quad \|\mathbf{m}_g\| \neq 0,\end{aligned}\quad (5.15)$$

$$\mathbf{x}_o(\mathbf{u}_o) = \frac{\mathbf{x}_g \times \mathbf{n}_o}{\|\mathbf{x}_g \times \mathbf{n}_o\|}, \quad \mathbf{y}_o(\mathbf{u}_o) = \frac{\mathbf{n}_o \times \mathbf{x}_o}{\|\mathbf{n}_o \times \mathbf{x}_o\|}, \quad \mathbf{z}_o(\mathbf{u}_o) = \mathbf{n}_o. \quad (5.16)$$

Here, \mathbf{m}_g is an arbitrary vector parallel to the cross-sectional plane. If \mathbf{m}_g is chosen as the unit normal vector of the skeleton curve, then the frame C_g corresponds to the classical *Frenet-Serret frame*.

With the alternative form matrices given by (5.14) and the sphere form matrices in

(5.8), the rolling kinematics between the two objects can be expressed as

$$\begin{aligned}\dot{\mathbf{u}}_o &= \left(\mathbf{R}_\psi \mathbf{K}_s \mathbf{R}_\psi \tilde{\mathbf{M}}_o + \tilde{\mathbf{K}}_o \right)^{-1} \left(\begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} - \mathbf{R}_\psi \mathbf{K}_s \mathbf{R}_\psi \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\mathbf{u}}_s &= \mathbf{M}_s^{-1} \left(\tilde{\mathbf{M}}_o^{-1} \mathbf{R}_\psi + \tilde{\mathbf{K}}_o^{-1} \mathbf{R}_\psi \mathbf{K}_s \right)^{-1} \cdot \left(\tilde{\mathbf{K}}_o^{-1} \begin{bmatrix} -\omega_y \\ \omega_x \end{bmatrix} + \tilde{\mathbf{M}}_o^{-1} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \right), \\ \dot{\psi}_i &= \omega_z + \tilde{\mathbf{T}}_o \dot{\mathbf{u}}_o + \mathbf{T}_s \mathbf{M}_s \dot{\mathbf{u}}_s, \quad v_z = 0,\end{aligned}\tag{5.17}$$

where $\boldsymbol{\omega}_{l_s l_o} = [\omega_x, \omega_y, \omega_z]^T$ and $\mathbf{V}_{l_s l_o} = [v_x, v_y, v_z]^T$ denote the relative angular and linear velocities between the rolling bodies, with C_{l_o} expressed relative to C_{l_s} in frame C_{l_o} . The rotation matrix \mathbf{R}_ψ is given by

$$\mathbf{R}_\psi = \begin{bmatrix} \cos \psi & -\sin \psi \\ -\sin \psi & -\cos \psi \end{bmatrix}.$$

An alternative form of Eq. (5.17) can be derived in terms of the sphere angular and linear velocities by substituting $\boldsymbol{\omega}_{l_o l_s} = -\mathbf{R}_\psi \boldsymbol{\omega}_{l_s l_o}$ and $\mathbf{V}_{l_o l_s} = -\mathbf{R}_\psi \mathbf{V}_{l_s l_o}$ into the equation.

The contact point position on the object surface and the position of the corresponding skeleton point relative to the fixed frame C_{r_o} are expressed as

$$\mathbf{p}_{r_{og}} = [p_x(v_g), p_y(v_g), p_z(v_g)]^T, \quad \mathbf{p}_{r_{oc_o}} = \mathbf{R}_{r_{og}} \cdot \mathbf{p}_{gc_o},\tag{5.18}$$

where

$$\mathbf{p}_{gc_o} = \left[0, \quad R_o(u_o) \cos u_o, \quad R_o(u_o) \sin u_o \right]^T, \quad \mathbf{R}_{r_{og}} = [\mathbf{x}_g, \mathbf{y}_g, \mathbf{z}_g].\tag{5.19}$$

5.4 Results and Discussion

To assess the performance of the two proposed rolling contact kinematics models, simulations were carried out independently for each method, as demonstrated by Subsection 5.4.1 and 5.4.2 respectively. Besides, we made comparison simulations of them in Subsection 5.4.3, followed by a discussion on applications.

All simulations were implemented in MATLAB, utilising the ODE45 solver on a 13th Generation Intel(R) Core(TM) i5-13500H processor running at 2.60 GHz. The simulation outcomes, which include the rolling trajectories on both the spherical and object surfaces, are depicted in Figures 5.4, 5.5, 5.6, 5.8, 5.9, and 5.11. The objects examined in the study comprised a carrot, a banana, and an egg. The rolling environment was modelled after the Indian Tunnel dataset [176].

Input parameters for the simulations consisted of the relative angular velocity $\boldsymbol{\omega}_{l_s l_o} = [\omega_x, \omega_y, \omega_z]^T$ and relative linear velocity $\mathbf{V}_{l_s l_o} = [v_x, v_y, v_z]^T$ between the contacting bodies. The state variables computed during the simulations were $\{\mathbf{u}_o, \mathbf{u}_s, \psi\}$, as specified in Equation (5.17). The point cloud slicing interval I , introduced in Figure 5.2, was set to 0.02 times the diagonal length of the object's point cloud bounding box. Additional simulation parameters are summarised in Table 5.1. The point cloud models used for the objects were obtained from real-world scanned data [181].

Parameter	Value
Initial states $\mathbf{u}_o(0), \mathbf{v}_o(0), \psi(0)$	$[0.1, \pi/2, 0.1, 0.1, \pi/2]^T$
Relative error tolerance (solver)	1×10^{-3}
Absolute error tolerance (solver)	1×10^{-5}
Fourier fitting error (Eq. 5.2)	1×10^{-3}
Fourier fitting order range	$[3, 8]$
Vector \mathbf{m}_g (Eq. 5.15)	$[0, 0, 1]^T$
Neighbourhood size k for estimating \mathbf{n}_o (Eq. 5.16)	20

Table 5.1: Summary of constant parameters used in the simulations.

5.4.1 Simulation Results Using the Locally Generalised Surface Model

This section presents and analyses simulation results obtained using the rolling contact kinematics model based on the locally generalised surface formulation introduced in

Subsection 5.3.1. The simulations consider three objects rolling against a sphere, both with and without slippage.

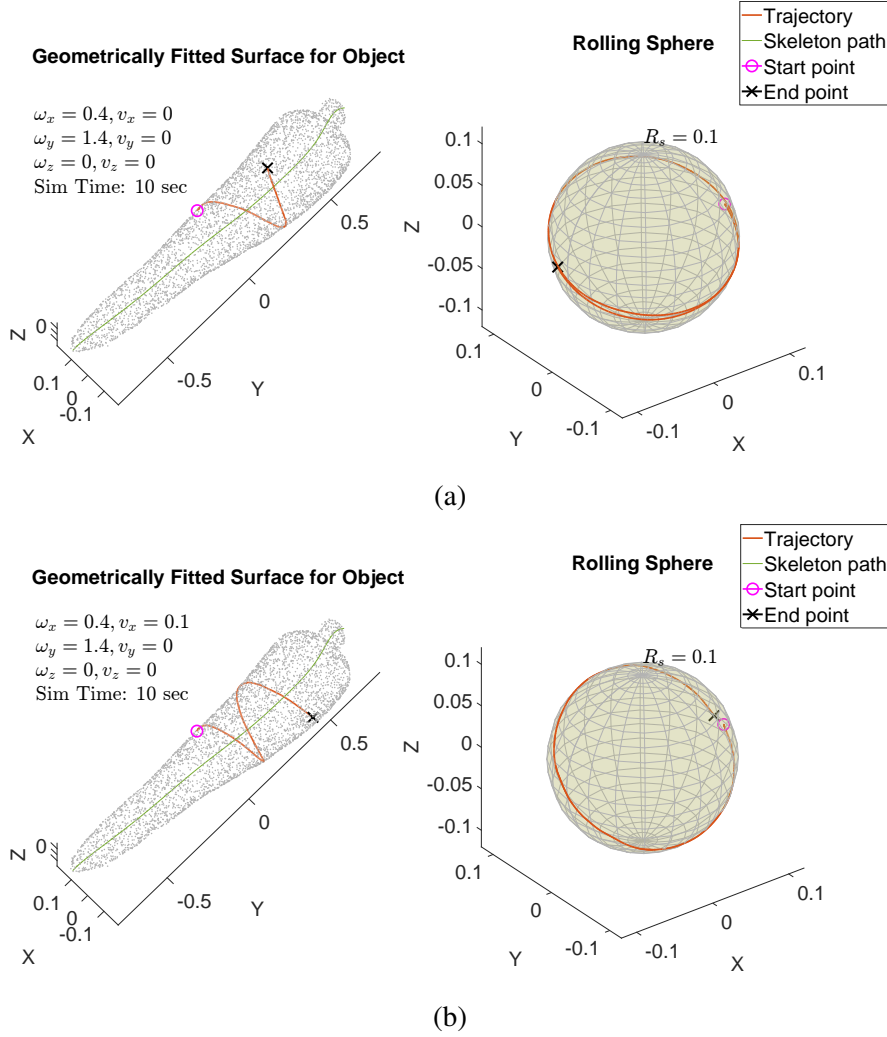


Figure 5.4: Simulation trajectories for the carrot object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.

The first case concerns the carrot, with the sphere initialised near $v_g = \frac{\pi}{2}$, corresponding to the central region of the object. The results indicate that the sphere trajectories are smooth and physically consistent under both pure rolling and rolling with sliding conditions, as illustrated in Figure 5.4. The state variables evolve smoothly, demonstrating stable behaviour without any irregularities or singularities, as shown in Figure 5.10(a). Introducing a small sliding velocity component produces longer trajec-

ories in the sliding direction compared to the pure rolling case, validating the model's ability to capture slip effects accurately.

Next, the banana object presents a more complex geometry characterised by a significantly curved skeleton curve, as shown in Figure 5.5. The simulation explored the effects of altering the sliding direction, demonstrating consistent changes in rolling motion direction as expected, with extended trajectories along the skeleton curve under sliding, confirming the model's robustness.

Similarly, Figure 5.6 displays results for the egg object, which has a markedly different geometric structure. The simulations confirm that sliding in orthogonal directions induces longer displacements both along and perpendicular to the skeleton curve, aligning well with physical intuition.

The variation of R_g , shown in Fig. 5.7(b), reflects the sphere's motion along the object surfaces and validates the effectiveness of the Fourier series-based closed curve fitting. In the banana case (B-NS), R_g remains nearly constant as the sphere moves along a short path with a stable radius. In contrast, the carrot case (C-NS) shows an increasing R_g as the sphere moves toward a region with a larger radius. For the egg, R_g decreases steadily as the contact point shifts from the middle toward the edge. Fluctuations in R_g , particularly in the egg with slippage (E-S) and the carrot simulations (C-NS, C-S), are due to the sensitivity of the fitted model to data quality, including noise and point cloud density. In comparison, the banana's dense and consistent point cloud results in reduced fluctuations. This highlights the need for further study on the effects of noise and sparsity.

5.4.2 Simulation Results of the Model with Globally Generalised Surface

This subsection presents and analyses the proposed method for modelling globally generalised object point clouds, exemplified by a rolling sphere interacting with multiple

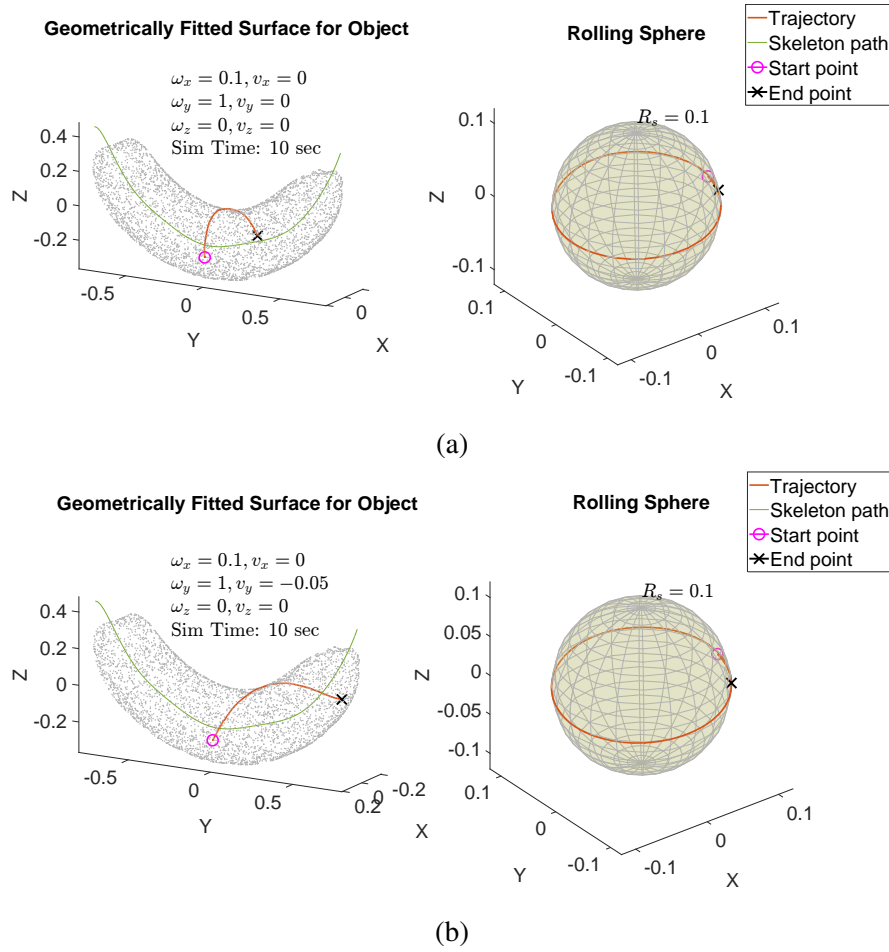


Figure 5.5: Simulation trajectories for the banana object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.

surfaces. The globally generalised surface formulation, detailed in Subsection 5.3.2, is employed herein. Furthermore, the performance and capabilities of the resultant geometric kinematic model are evaluated. To intuitively demonstrate the effectiveness of our approach, comprehensive simulations were conducted on one of the three selected objects, namely the carrot, under various angular and linear velocity settings. Additionally, point clouds of two other objects – a banana and a real-scanned cave – were utilised to provide more general simulations and a broader discussion.

Simulation results in Fig. 5.8A show that applying relative angular velocity solely about the y -axis produces a trajectory that follows the fitted curve. Conversely, when

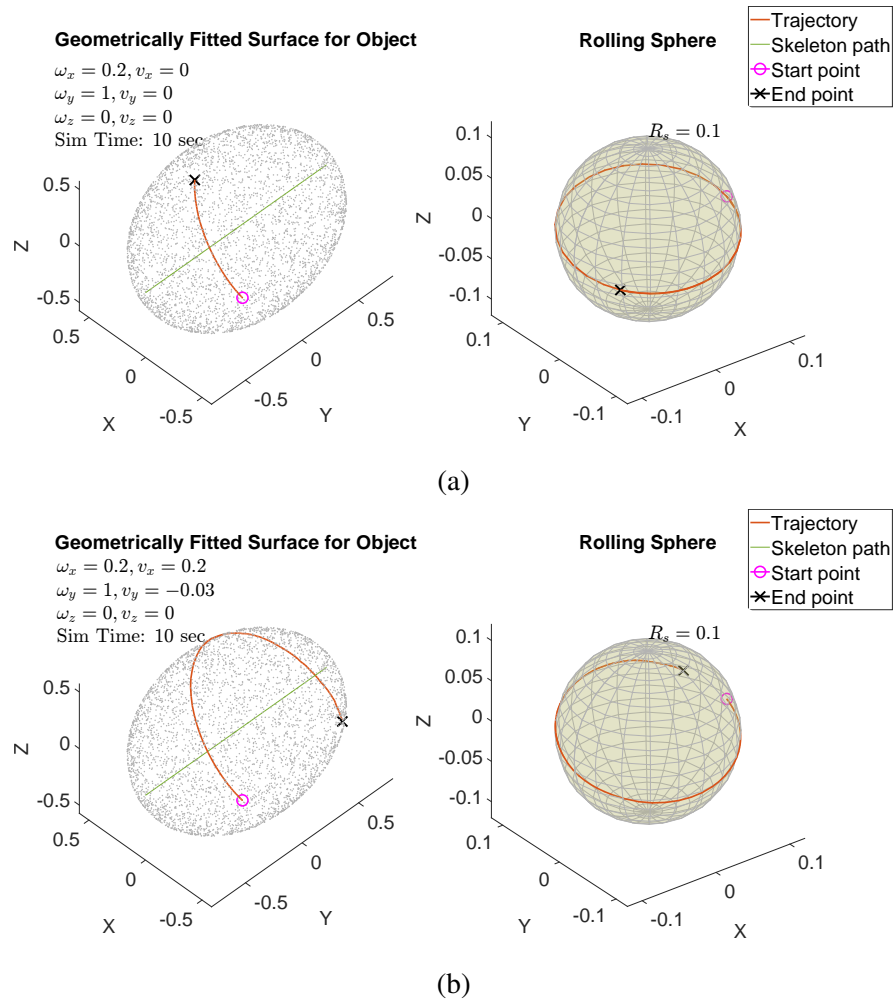


Figure 5.6: Simulation trajectories for the egg object and rolling sphere, under (a) pure rolling and (b) rolling with slippage conditions.

the angular velocity is applied only about the x -axis, the trajectory remains perpendicular to the cross-sectional contours and aligns with the skeleton path, as illustrated in Fig. 5.8B. The choice of the vector \mathbf{m}_g in Eq. (5.15) influences the surface parameterisation and thus alters the computed path. When angular velocities are applied about both the x - and y -axes simultaneously, the contact trajectory becomes helical on the approximately cylindrical carrot surface, as shown in Fig. 5.8C and D. These results confirm that the kinematics follow the control inputs accurately and demonstrate the approach's capability to bridge discrete point-cloud data with a continuous rolling-contact model.




Figure 5.7: (a) Time responses of states u_g , u_s , and ψ for the carrot simulation corresponding to Figure 5.4(b). (b) Variation of radius R_g for the simulations shown in Figures 5.4–5.6.

Slippage, represented by relative linear velocity, affects the movement of the contact point, resulting in longer or shorter paths along the direction of slippage on both the sphere and the object surfaces, as depicted in Fig. 5.8C and D. According to differential geometry, the trajectory on the sphere is expected to vary with different slippage rates, depending on local surface curvature variations. However, the simulated trajectories on the sphere in Fig. 5.8C-D do not exhibit the anticipated differences. This is attributed to errors in the surface normal vector estimation method. As shown in the simulation state response corresponding to Fig. 5.8D (see Fig. 5.10(a)), the rates of change of u_o and v_s are faster than those of v_o and u_s , which aligns with the initial spin angle ψ and the specified angular and linear velocity inputs.

Simulations were also conducted on the point cloud data of the banana and egg objects, illustrated in Fig. 5.9A and B. Since the skeleton path of the banana exhibits greater curvature than that of the carrot and egg, there are more pronounced surface curvature variations along the y -axis of C_{c_o} . Consequently, the path on the sphere no longer follows a fixed circle, unlike the carrot and egg cases, where trajectories remain qualitatively similar.

Fig. 5.10(b) illustrates the response of R_o across all simulations. The R_o corresponding

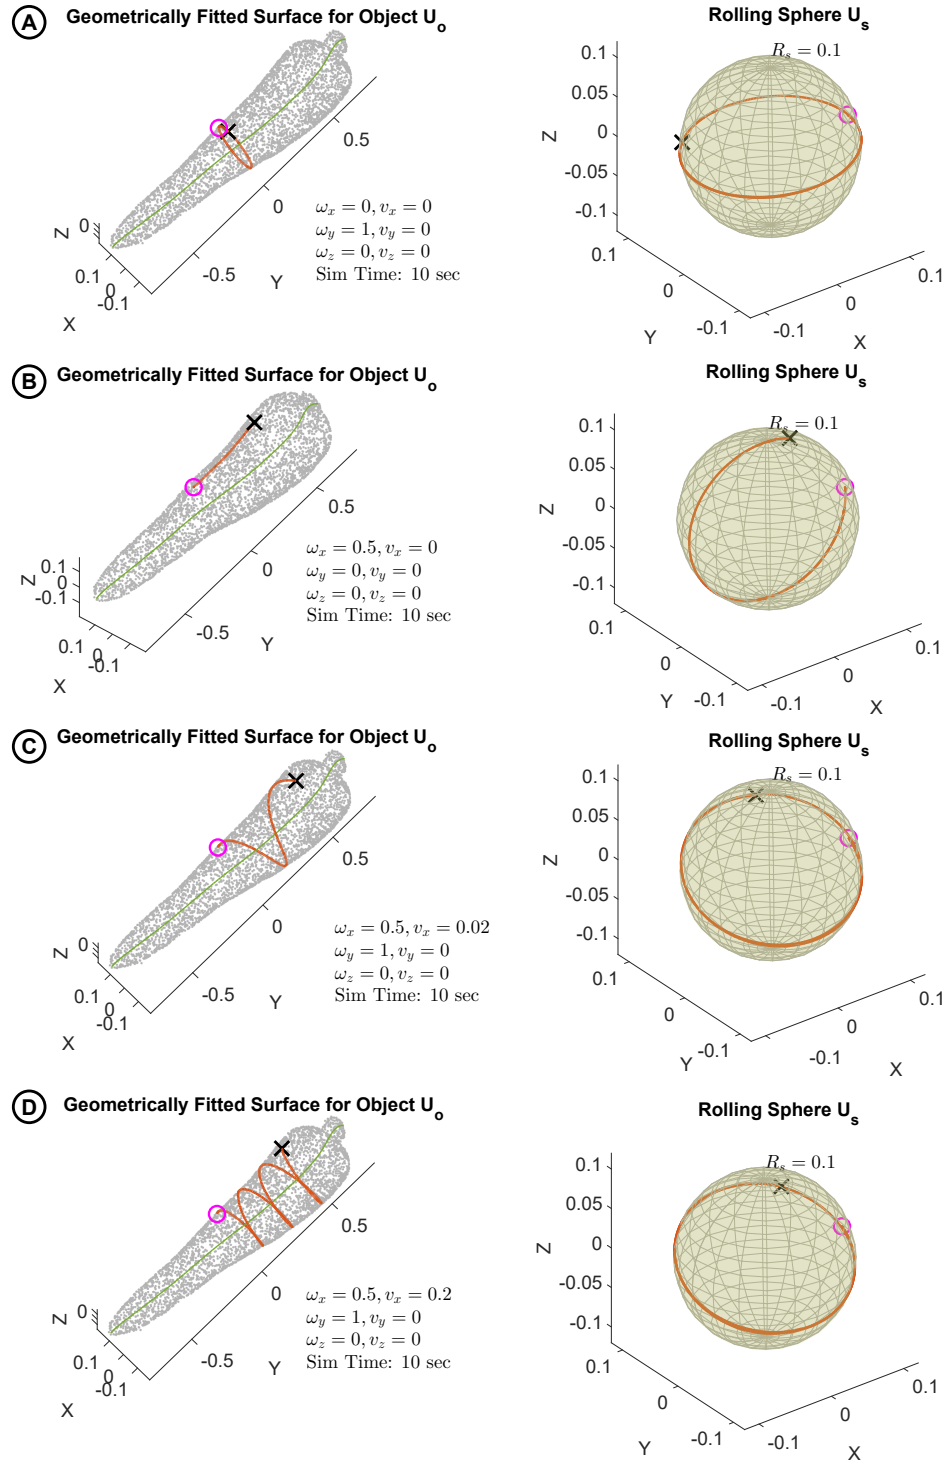


Figure 5.8: Simulated contact trajectories of rolling contact between a sphere and the carrot, banana, and egg objects, respectively.

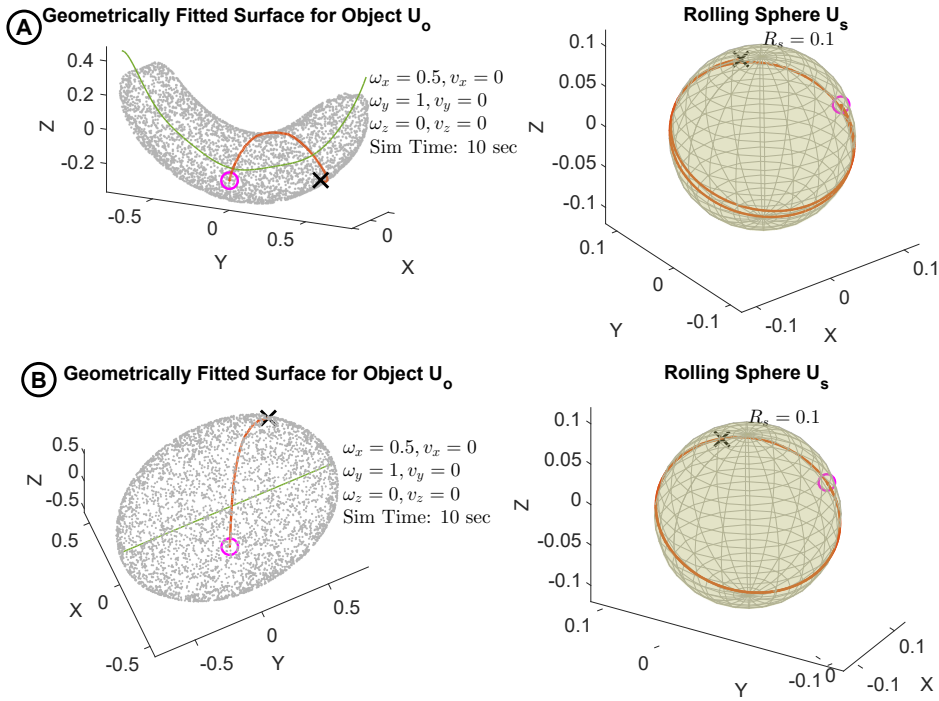


Figure 5.9: Simulated contact trajectories of rolling contact between a sphere and (A) the banana and (B) the egg objects.




Figure 5.10: (a) Time responses of the states u_o , u_s , and ψ for the example simulation in Fig. 5.8D. (b) Response of the radius R_o for the example simulations in Figs. 5.8A-D and 5.9A-B.

to Fig. 5.8A remains within a limited range throughout, whereas other R_o responses either increase or decrease with minor fluctuations. This behaviour indicates that the

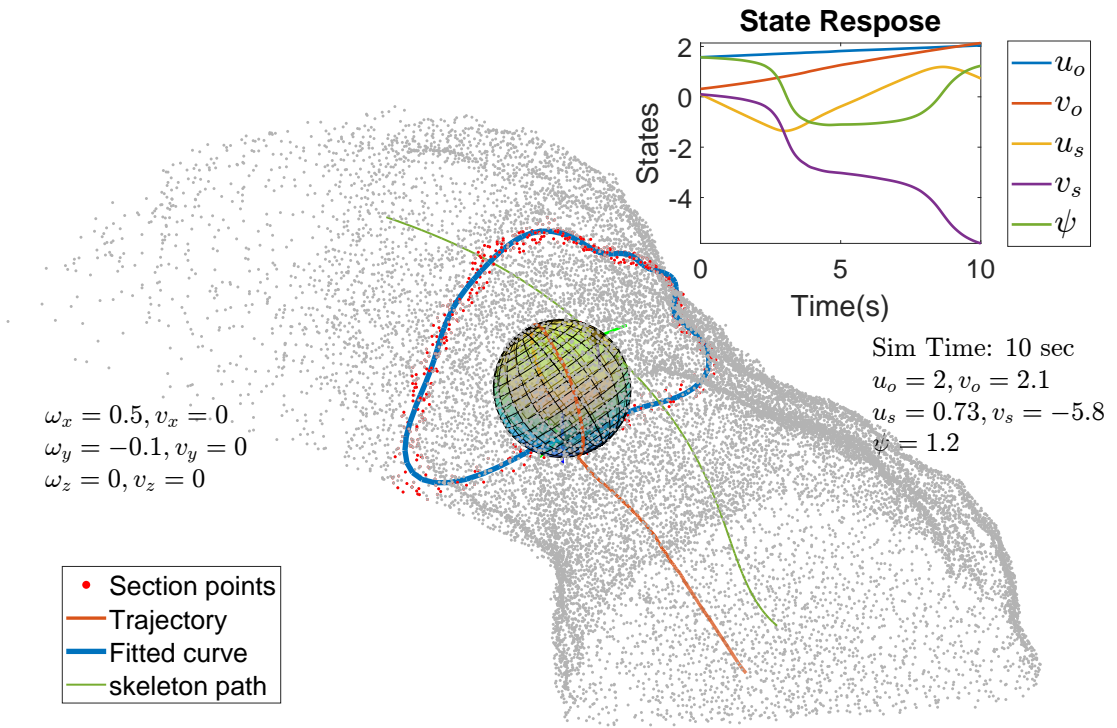


Figure 5.11: Simulation of rolling in the Indian Tunnel environment [176].

geometric radius changes successfully reflect curvature or shape variations, since the condition $\omega_x = 0$ ensures v_o remains invariant. Minor fluctuations primarily arise from errors in surface normal estimation and the inherent imprecision and non-uniform distribution of the point-cloud data.

To demonstrate real-world feasibility, we simulated a sphere rolling within a natural cave environment (Fig. 5.11). The initial contact states (u_o, v_o) were set to $(\frac{\pi}{2}, \frac{\pi}{10})$. The Fourier model (Eq. (5.1)) accurately reconstructs the cross-sectional contour. Trajectory fluctuations caused by the uneven cave floor are faithfully captured by our kinematic model, while minor deviations stem from point-cloud resolution and surface-normal estimation errors.

5.4.3 Comparison

To compare the locally and globally generalised surface-based rolling contact kinematics models, we conducted simulations using both approaches under identical experimental settings and object inputs, as illustrated in Fig. 5.12, Fig. 5.13, and Fig. 5.14.

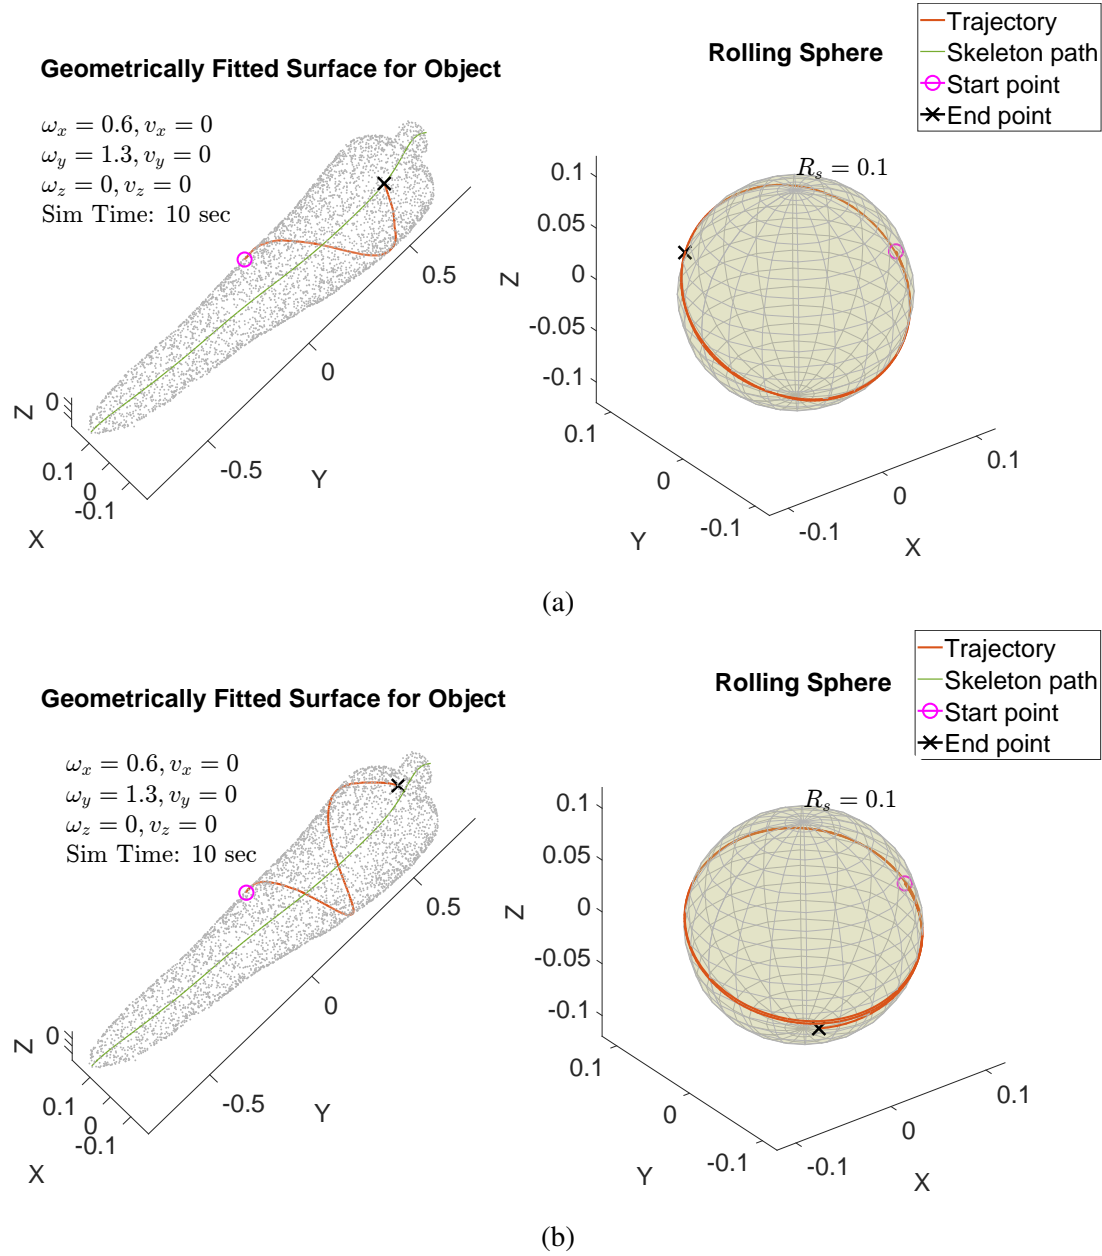


Figure 5.12: Rolling simulation without slippage of the locally (a) and globally (b) generalised surface based rolling contact kinematics models.



All simulations were performed on the same computing platform (Intel(R) Core(TM) i5-13500H processor, 2.60 GHz). The locally generalised surface-based method demonstrated a significant computational advantage, achieving approximately 10 times faster execution compared to the globally generalised surface-based approach. For instance, the computation times for Fig. 5.13(a) and (b) were 7.527 seconds and 76.497 seconds,

respectively.

Trajectory results shown in Fig. 5.12 and Fig. 5.13 reveal noticeable differences between the two methods, particularly in simulations involving slippage. This can be attributed to the underlying assumptions of the models: the locally generalised surface-based approach treats each local section as a thin cylindrical segment with approximately constant curvature along the skeleton path, whereas the globally generalised surface model accounts for full geometric detail, resulting in higher computational complexity. Since slippage dynamics are highly sensitive to local geometric variations, the differences between the two approaches become more pronounced under slipping conditions. This is further reflected in the time responses of the state variables shown in Fig. 5.14, where the responses in Fig. 5.14(a) and (b), corresponding to Fig. 5.12, show similar patterns, while the responses in Fig. 5.14(c) and (d), corresponding to Fig. 5.13, diverge significantly.

Apart from their differences, the two approaches exhibit similar patterns in the simulation results. Both effectively capture the rolling contact motion in alignment with the rolling velocities. Moreover, the accuracy and resolution of the simulations are comparable, which serves as a form of mutual validation.

It is noteworthy that the controllability of locally generalised surface-based methods can be verified, whereas for globally generalised surface-based methods, the complexity makes such verification impractical. Additionally, we highlight that while the globally generalised surface-based methods are slower, they tend to be relatively more accurate compared to the locally generalised surface-based methods.

Results in this section illustrate the core behaviour of the proposed method, indicating its potential relevance to applications such as dexterous manipulation and mobile robot navigation. Although many existing robotic hands employ multiple fingers or complex contact surfaces beyond a single sphere [189, 190, 37], our model provides a general and extendable framework. It can serve as a foundation for modelling more complex contact geometries and multi-contact kinematics [152], which we identify as an im-

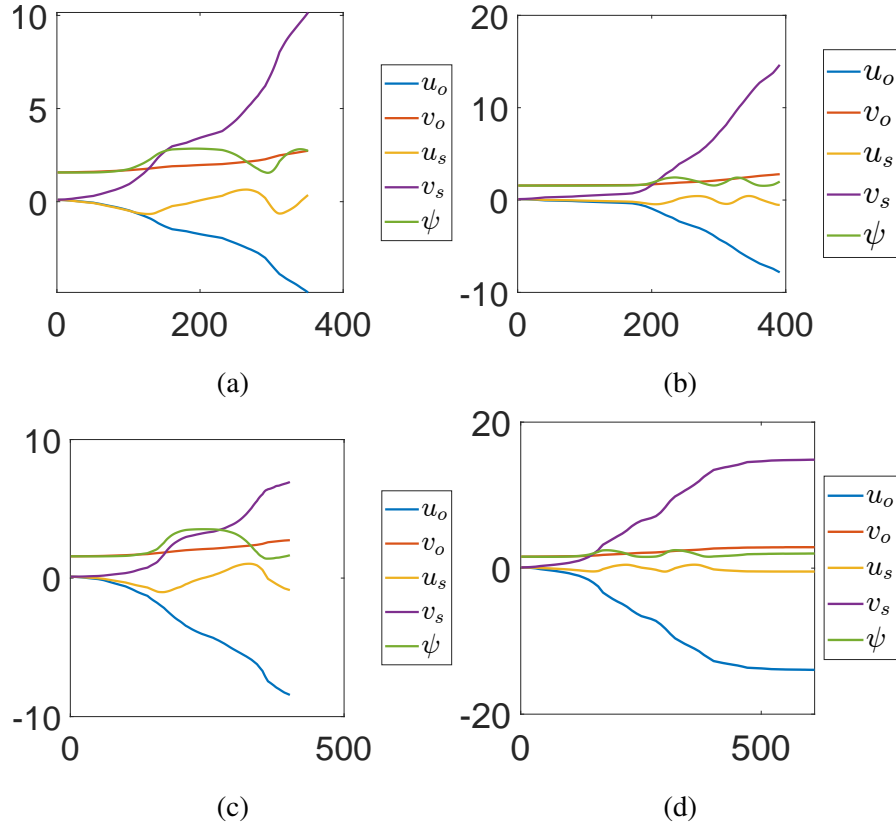


Figure 5.14: Time responses of the states u_o , u_s , and ψ for the example simulation corresponding to Fig. 5.12 ((a) and (b)), and Fig. 5.13 ((c) and (d)) respectively.

portant direction for future work. The residual errors observed in our evaluation are primarily due to the resolution limitations of the input point clouds.

5.5 Conclusions

In this study, we proposed a novel method for geometrically fitting arbitrary point cloud data by employing a skeleton-guided reference frame, combined with a semi-convex spin-rolling-sliding kinematic model grounded in differential geometry. The surface used to construct the kinematic model was generalised in two complementary ways: globally and locally. We presented skeleton-based kinematic formulations for both generalisation approaches and discussed their respective characteristics.

Our method was validated through simulations involving representative objects—including a carrot, banana, and egg—using experimental data from real-scanned point clouds. Additionally, a simulation of a sphere rolling through a real-scanned cave environment further demonstrated the model’s applicability to irregular and naturalistic geometries. The results showed that both pure rolling and rolling with slippage yield smooth and physically plausible contact trajectories, supported by consistent state responses within the proposed skeleton-guided framework.

Importantly, our findings highlight the critical role of point cloud density in determining accuracy and robustness. Denser point clouds enable more reliable slicing and lead to more precise kinematic modelling, especially in scenarios involving slippage, where trajectory sensitivity to local curvature becomes more pronounced.

Looking ahead, we plan to develop more robust slicing techniques that incorporate detailed curvature analysis to better handle complex and multi-branch skeleton structures, such as those found in trees or tunnels. This will broaden the scope of applications for our approach. Additionally, we aim to implement and experimentally evaluate motion planning strategies, with a particular focus on assessing the proposed kinematic model’s performance in rolling fingertip manipulation tasks.

Conclusion and Future Work

6.1 Conclusion

This thesis addresses three key challenges that limit the practical use of skeletonisation techniques in robotics: computational efficiency, skeleton quality evaluation, and the integration of skeletons into rolling-contact kinematics.

First, regarding the computational cost of curve skeletonisation, a Laplacian-based method was adopted as the baseline. Through a series of algorithmic optimisations, redundant computations were eliminated, leading to a more efficient skeletonisation framework with negligible degradation in skeleton quality. Although these improvements significantly reduce computational cost, further acceleration remains possible through parallel computing or learning-based strategies.

Second, to address the challenge of skeleton evaluation, this thesis investigates the statistical behaviour of a stable Laplacian contraction-based skeletonisation method. Building on this analysis, a comprehensive quantitative evaluation framework was proposed, encompassing four key skeleton quality properties defined both theoretically and numerically. To the best of our knowledge, this represents the first attempt to formulate numerical definitions for these evaluation criteria, albeit with certain assumptions such as treating the skeletal point set as an intermediate representation for topological similarity assessment. The relevance of the framework was further demonstrated through discussions of robotics applications, highlighting how different tasks

impose varying levels of quality requirements. For example, using a skeleton as a navigation path requires smoothness, whereas grasp planning does not.

Finally, to bridge the gap between continuous rolling-contact kinematics and the discrete perception typical in robotics, this work explores the use of skeletons to guide rolling-contact kinematic modelling. This contribution enables continuous kinematic theory to be applied to discrete shape representations such as point clouds, providing a valuable link between perception and action within robotic systems.

6.2 Future Work


Despite the progress made, several open questions and promising research directions remain. In particular, integrating skeletonisation techniques into broader robotic perception-action pipelines offers substantial opportunities for further development.

A primary challenge is real-time performance. For deployment in real-world robotic systems, the computational cost of skeletonisation must be further reduced. Potential solutions include leveraging parallel computing architectures or adopting learning-based approaches, where traditional methods may support data generation or labelling.


Another outstanding challenge concerns skeleton evaluation. While this thesis proposes general quantitative criteria, the relevance of individual metrics can vary significantly across robotic tasks. Future work should therefore aim to develop task-specific evaluation frameworks tailored to domains such as manipulation, locomotion, or soft robotics.

Additionally, the skeleton-guided rolling-contact kinematics model introduced in this thesis has thus far been validated through theory and simulation only. Physical implementation on real robotic platforms—where rolling interactions may be more complex—will be essential for assessing its robustness and practical applicability.

A further challenge arises from imperfections in real-world perception, which complicate reliable skeletonisation. Issues such as occlusion, non-uniform point distribution, and measurement noise can severely degrade skeleton quality. Although the Laplacian-based method exhibits robustness to mild imperfections, more severe forms of degradation remain problematic. A promising direction is the combination of traditional skeletonisation algorithms with learning-based techniques for point cloud completion, denoising, and upsampling [53, 172, 85] prior to skeleton extraction.



(a) Rolling grasper [190]



(b) Spherical rolling robots [145]

Figure 6.1: Promising potential robotic applications.

Looking ahead, skeleton-based methods hold promise for a wide range of robotic applications. In manipulation, skeletons may assist in generating grasp candidates, including grasp poses and positions [165, 119]. They may also support future directions such as object pose or shape estimation through the analysis of the skeleton path and its orientation. In the context of grasp planning, a full scan of the object may not be strictly necessary; however, only limited missing data can be tolerated, as long as the overall topology is preserved. Moreover, skeleton-based grasp planning may be less efficient and less robust than learning-based methods, although it benefits from improved explainability and the lack of a need for labelled training data. A promising future direction is to combine skeleton-based and learning-based approaches to leverage the strengths of both. In motion planning and navigation, skeletal representations offer

valuable structural cues for generating smooth, collision-free trajectories, particularly within elongated or tunnel-like environments such as caves. In such settings, local sectional scans of the surroundings are sufficient to derive feasible paths. However, because the computation must be performed online, achieving real-time performance is essential. Skeleton-based representations may also enhance object reconstruction using visual SLAM data [178]. These directions underscore the broad potential of skeletonisation as a foundational tool for future robotic systems.

In summary, this thesis lays the groundwork for a more robust, efficient, and application-oriented approach to skeletonisation in robotics, while identifying several promising avenues for continued research.

Bibliography

- [1] Narendra Ahuja and Jen-Hui Chuang. Shape representation using a generalized potential field model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):169–176, 1997.
- [2] Nina Amenta and Marshall Bern. Surface reconstruction by voronoi filtering. In *Proceedings of the fourteenth annual symposium on Computational geometry*, pages 39–48, 1998.
- [3] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, pages 249–266, 2001.
- [4] Carlo Arcelli. Pattern thinning by contour tracing. *Computer Graphics and Image Processing*, 17(2):130–144, 1981.
- [5] Carlo Arcelli and Gabriella Sanniti Di Baja. A width-independent fast thinning algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (4):463–474, 1985.
- [6] Carlo Arcelli, Gabriella Sanniti di Baja, and Luca Serino. A parallel algorithm to skeletonize the distance transform of 3d objects. *Image and Vision Computing*, 27(6):666–672, 2009.
- [7] Carlo Arcelli, Gabriella Sanniti Di Baja, and Luca Serino. Distance-driven skeletonization in voxel images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):709–720, 2010.
- [8] Nafiz Arica and Fatos T Yarman-Vural. One-dimensional representation of two-dimensional information for hmm based handwriting recognition. *Pattern recognition letters*, 21(6-7):583–592, 2000.

- [9] Dominique Attali, André Lieutier, and David Salinas. Vietoris-rips complexes also provide topologically correct reconstructions of sampled shapes. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 491–500, 2011.
- [10] Dominique Attali and Annick Montanvert. Computing and simplifying 2d and 3d continuous skeletons. *Computer vision and image understanding*, 67(3):261–273, 1997.
- [11] Oscar Kin-Chung Au, Chiew-Lan Tai, Hung-Kuo Chu, Daniel Cohen-Or, and Tong-Yee Lee. Skeleton extraction by mesh contraction. *ACM Trans. Graph.*, 27(3):1–10, 2008.
- [12] Xiang Bai and Longin Jan Latecki. Path similarity skeleton graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 30(7):1282–1292, 2008.
- [13] Yang Bai, Mikhail Svinin, and Motoji Yamamoto. Motion planning for a pendulum-driven rolling robot tracing spherical contact curves. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4053–4058. IEEE, 2015.
- [14] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM Trans. Graph.*, 26(3):72–es, 2007.
- [15] William J Beks and Nikolaos Papanikolopoulos. 3d point cloud segmentation using topological persistence. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5046–5051. IEEE, 2016.
- [16] Bruno Belzile and David St-Onge. Aries: Cylindrical pendulum actuated explorer sphere. *IEEE/ASME Transactions on Mechatronics*, 27(4):2142–2150, 2022.
- [17] Mattia G Bergomi, Patrizio Frosini, Daniela Giorgi, and Nicola Quercioli. Towards a topological–geometrical theory of group equivariant non-expansive operators for data analysis and machine learning. *Nat. Mach. Intell.*, 1(9):423–433, 2019.
- [18] Gilles Bertrand. P-simple points: A solution for parallel thinning. In *5th Conf. on Discrete Geometry*, number 1, pages 233–242, 1995.

- [19] Gilles Bertrand. A parallel thinning algorithm for medial surfaces. *Pattern recognition letters*, 16(9):979–986, 1995.
- [20] Antonio Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Trans. Robot. Autom.*, 16(6):652–662, 2000.
- [21] Albert Bijaoui, Yves Bobichon, and Li Huang. Digital image compression in astronomy morphology or wavelets. *Vistas in astronomy*, 40(4):587–594, 1996.
- [22] Teng Bin, Hanming Yan, Ning Wang, Milutin N Nikolić, Jianming Yao, and Tianwei Zhang. A survey on the visual perception of humanoid robot. *Biometric Intelligence and Robotics*, 5(1):100197, 2025.
- [23] H. Blum. *A Transformation for Extracting New Descriptors of Shape*. M.I.T. Press, 1967.
- [24] Harry Blum and Roger N Nagel. Shape description using weighted symmetric axis features. *Pattern recognition*, 10(3):167–180, 1978.
- [25] Fred L Bookstein. The line-skeleton. *Computer Graphics and Image Processing*, 11(2):123–137, 1979.
- [26] Gunilla Borgefors. Distance transformations in arbitrary dimensions. *Computer vision, graphics, and image processing*, 27(3):321–345, 1984.
- [27] Gunilla Borgefors. Distance transformations in digital images. *Computer vision, graphics, and image processing*, 34(3):344–371, 1986.
- [28] Gunilla Borgefors. On digital distance transforms in three dimensions. *Computer vision and image understanding*, 64(3):368–376, 1996.
- [29] Gunilla Borgefors. Weighted digital distance transforms in four dimensions. *Discrete Applied Mathematics*, 125(1):161–176, 2003.
- [30] Lynda Ben Boudaoud, Abderrahmane Sider, and Abdelkamel Tari. A new thinning algorithm for binary images. In *2015 3rd international conference on control, engineering & information technology (CEIT)*, pages 1–6. IEEE, 2015.
- [31] Jonathan W Brandt and V Ralph Algazi. Continuous skeleton computation by voronoi diagram. *CVGIP: Image understanding*, 55(3):329–338, 1992.

- [32] James Bruce and Manuela Veloso. Real-time randomized path planning for robot navigation. In *IEEE/RSJ international conference on intelligent robots and systems*, volume 3, pages 2383–2388. IEEE, 2002.
- [33] Alexander Bucksch and Roderik Lindenbergh. Campino—a skeletonization method for point cloud processing. *ISPRS journal of photogrammetry and remote sensing*, 63(1):115–127, 2008.
- [34] Junjie Cao, Andrea Tagliasacchi, Matt Olson, Hao Zhang, and Zhinxun Su. Point Cloud Skeletons via Laplacian Based Contraction. In *2010 Shape Modeling International Conference*, pages 187–197, June 2010.
- [35] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas Guibas. Persistence barcodes for shapes. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 124–135, 2004.
- [36] Gregory Chang, S Kubilay Pakin, Mark E Schweitzer, Punam K Saha, and Ravinder R Regatte. Adaptations in trabecular bone microarchitecture in olympic athletes determined by 7t mri. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 27(5):1089–1095, 2008.
- [37] Jayden Chapman, Gal Gorjup, Anany Dwivedi, Saori Matsunaga, Toshisada Mariyama, Bruce MacDonald, and Minas Liarokapis. A locally-adaptive, parallel-jaw gripper with clamping and rolling capable, soft fingertips for fine manipulation of flexible flat cables. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6941–6947. IEEE, 2021.
- [38] Wei-Hsi Chen, Ching-Pei Chen, Wei-Shun Yu, Chang-Hao Lin, and Pei-Chun Lin. Design and implementation of an omnidirectional spherical robot omnicon. In *2012 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 719–724. IEEE, 2012.
- [39] Mattéo Clémot and Julie Digne. Neural skeleton: Implicit neural representation away from the surface. *Computers & Graphics*, 114:368–378, 2023.
- [40] Nicu D. Cornea, Deborah Silver, and Patrick Min. Curve-Skeleton Properties, Applications, and Algorithms. *IEEE Trans. Vis. Comput. Graph.*, 13(3):530–548, May 2007.

- [41] Nicu D Cornea, Deborah Silver, Xiaosong Yuan, and Raman Balasubramanian. Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer*, 21:945–955, 2005.
- [42] Lei Cui and Jian S Dai. A darboux-frame-based formulation of spin-rolling motion of rigid objects with point contact. *IEEE Trans. Robot.*, 26(2):383–388, 2010.
- [43] Tim Culver, John Keyser, and Dinesh Manocha. Accurate computation of the medial axis of a polyhedron. In *Proceedings of the fifth ACM symposium on Solid modeling and applications*, pages 179–190, 1999.
- [44] Ilke Demir, Camilla Hahn, Kathryn Leonard, Geraldine Morin, Dana Rahbani, Athina Panotopoulou, Amelie Fondevilla, Elena Balashova, Bastien Durix, and Adam Kortylewski. Skelneton 2019: Dataset and challenge on deep learning for geometric shape understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- [45] Tamal K Dey and Jian Sun. Defining and computing curve-skeletons with medial geodesic function. In *Symposium on geometry processing*, volume 6, pages 143–152, 2006.
- [46] Tamal K Dey and Wulue Zhao. Approximating the medial axis from the voronoi diagram with a convergence guarantee. *Algorithmica*, 38:179–200, 2004.
- [47] Aminata Diouf, Bruno Belzile, Maarouf Saad, and David St-Onge. Spherical rolling robots—design, modeling, and control: A systematic literature review. *Robot. Auton. Syst.*, page 104657, 2024.
- [48] Paweł Dłotko and Ruben Specogna. Topology preserving thinning of cell complexes. *IEEE Transactions on Image Processing*, 23(10):4486–4495, 2014.
- [49] Harry Dobbs, Oliver Batchelor, Casey Peat, James Atlas, and Richard Green. Quantifying robustness: 3d tree point cloud skeletonization with smart-tree in noisy domains. *Pattern Analysis and Applications*, 27(1):28, 2024.
- [50] Zoe Doulgeri and John Fasoulas. Stable grasping control under gravity by dual robotic fingers with soft rolling contacts. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 1681–1686. IEEE, 2002.

- [51] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & computational geometry*, 28:511–533, 2002.
- [52] Michal Etzion and Ari Rappoport. Computing voronoi skeletons of a 3-d polyhedron by space subdivision. *Computational Geometry*, 21(3):87–120, 2002.
- [53] Ben Fei, Weidong Yang, Wen-Ming Chen, Zhijun Li, Yikang Li, Tao Ma, Xing Hu, and Lipeng Ma. Comprehensive review of deep learning-based 3d point cloud completion processing and analysis. *IEEE Transactions on Intelligent Transportation Systems*, 23(12):22862–22883, 2022.
- [54] Andrew Fitzgibbon, Maurizio Pilu, and Robert B Fisher. Direct least square fitting of ellipses. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(5):476–480, 1999.
- [55] Yonatan Fridman, Stephen M Pizer, Stephen Aylward, and Elizabeth Bullitt. Extracting branching tubular object geometry via cores. *Medical Image Analysis*, 8(3):169–176, 2004.
- [56] Patrizio Frosini. Measuring shapes by size functions. In *Intelligent Robots and Computer Vision X: Algorithms and Techniques*, volume 1607, pages 122–133. SPIE, 1992.
- [57] Yili Fu, Hao Liu, Shuguo Wang, Wei Deng, Xianling Li, and Zhaoguang Liang. Skeleton-based active catheter navigation. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 5(2):125–135, 2009.
- [58] Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. The scale axis transform. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 106–115, 2009.
- [59] Weixin Gong and Gilles Bertrand. A simple parallel 3d thinning algorithm. In *[1990] Proceedings. 10th International Conference on Pattern Recognition*, volume 1, pages 188–190. IEEE, 1990.
- [60] Hayit Greenspan, Moshe Laifenfeld, Shmuel Einav, and Ofer Barnea. Evaluation of center-line extraction algorithms in quantitative coronary angiography. *IEEE Transactions on Medical Imaging*, 20(9):928–941, 2002.
- [61] Minghao Guo, Bohan Wang, and Wojciech Matusik. Medial skeletal diagram: A generalized medial axis approach for compact 3d shape representation. *ACM Transactions on Graphics (TOG)*, 43(6):1–23, 2024.

- [62] Mohamed Sabry Hassouna and Aly A Farag. Variational curve skeletons using gradient vector flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2257–2274, 2008.
- [63] Taosong He, Lichan Hong, Dongqing Chen, and Zhengrong Liang. Reliable path for virtual endoscopy: Ensuring complete examination of human organs. *IEEE transactions on visualization and computer graphics*, 7(4):333–342, 2001.
- [64] Christoph M Hoffmann. How to construct the skeleton of csg objects. 1990.
- [65] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *Proceedings of the 19th annual conference on computer graphics and interactive techniques*, pages 71–78, 1992.
- [66] Hui Huang, Shihao Wu, Daniel Cohen-Or, Minglun Gong, Hao Zhang, Guiqing Li, and Baoquan Chen. L1-medial skeleton of point cloud. *ACM Trans. Graph.*, 32(4), jul 2013.
- [67] Jaisimha, Haralick, and Dori. Quantitative performance evaluation of thinning algorithms under noisy conditions. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 678–683. IEEE, 1994.
- [68] Andrei C Jalba, Jacek Kustra, and Alexandru C Telea. Surface and curve skeletonization of large 3d models on the gpu. *IEEE transactions on pattern analysis and machine intelligence*, 35(6):1495–1508, 2012.
- [69] Katherine Margaret Frances James and Grzegorz Cielniak. Unsupervised clustering with geometric shape priors for improved occlusion handling in plant stem phenotyping. In *Annual Conference Towards Autonomous Robotic Systems*, pages 355–366. Springer, 2024.
- [70] Nan Jiang, Yifei Zhang, Dezha Luo, Chang Liu, Yu Zhou, and Zhenjun Han. Feature hourglass network for skeleton detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [71] Wei Jiang, Kai Xu, Zhi-Quan Cheng, Ralph R Martin, and Gang Dang. Curve skeleton extraction by graph contraction. In *Computational Visual Media: First*

- International Conference, CVM 2012, Beijing, China, November 8-10, 2012. Proceedings*, pages 178–185. Springer, 2012.
- [72] Wei Jiang, Kai Xu, Zhi-Quan Cheng, Ralph R Martin, and Gang Dang. Curve skeleton extraction by coupled graph contraction and surface clustering. *Graphical Models*, 75(3):137–148, 2013.
- [73] Dakai Jin and Punam K Saha. A new fuzzy skeletonization algorithm and its applications to medical imaging. In *International Conference on Image Analysis and Processing*, pages 662–671. Springer, 2013.
- [74] Wei Ke, Jie Chen, Jianbin Jiao, Guoying Zhao, and Qixiang Ye. Srn: Side-output residual network for object symmetry detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1068–1076, 2017.
- [75] Chung Hee Kim and George Kantor. Occlusion reasoning for skeleton extraction of self-occluded tree canopies. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9580–9586. IEEE, 2023.
- [76] David G Kirkpatrick. Efficient computation of continuous skeletons. In *20th Annual Symposium on Foundations of Computer Science (sfcs 1979)*, pages 18–27. IEEE, 1979.
- [77] Louisa Lam, Seong-Whan Lee, and Ching Y Suen. Thinning methodologies-a comprehensive survey. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 14(09):869–885, 1992.
- [78] Louisa Lam and Ching Y. Suen. An evaluation of parallel thinning algorithms for character recognition. *IEEE Transactions on pattern analysis and machine intelligence*, 17(9):914–919, 2002.
- [79] Seong-Whan Lee, Louisa Lam, and Ching Y Suen. A systematic evaluation of skeletonization algorithms. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(05):1203–1225, 1993.
- [80] Ta-Chih Lee, Rangasami L Kashyap, and Chong-Nam Chu. Building skeleton models via 3-d medial surface axis thinning algorithms. *CVGIP: graphical models and image processing*, 56(6):462–478, 1994.

- [81] Zexiang Li and John Canny. Motion of two rigid bodies with rolling constraint. *IEEE Trans. Robot. Autom.*, 6(1):62–72, 1990.
- [82] A Lieutier. Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. In *SM 03*. 2003.
- [83] Cheng Lin, Changjian Li, Yuan Liu, Nenglun Chen, Yi-King Choi, and Wenping Wang. Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4277–4286, June 2021.
- [84] Cheng Lin, Changjian Li, Yuan Liu, Nenglun Chen, Yi-King Choi, and Wenping Wang. Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4277–4286, 2021.
- [85] Hao Liu, Hui Yuan, Raouf Hamzaoui, Wei Gao, and Shuai Li. Pu-refiner: A geometry refiner with adversarial learning for point cloud upsampling. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2270–2274. IEEE, 2022.
- [86] Lu Liu, Erin W Chambers, David Letscher, and Tao Ju. A simple and robust thinning algorithm on cell complexes. In *Computer Graphics Forum*, volume 29, pages 2253–2260. Wiley Online Library, 2010.
- [87] Zhiwei Liu and Seyed Amir Tafrishi. Monorollbot: 3-dof spherical robot with underactuated single compliant actuator design. In *2025 IEEE 8th International Conference on Soft Robotics (RoboSoft)*, pages 1–6. IEEE, 2025.
- [88] Yotam Livny, Feilong Yan, Matt Olson, Baoquan Chen, Hao Zhang, and Jihad El-Sana. Automatic reconstruction of tree skeletal structures from point clouds. In *ACM SIGGRAPH Asia 2010 papers*, pages 1–8. 2010.
- [89] Christophe Lohou. Detection of the non-topology preservation of ma’s 3d surface-thinning algorithm, by the use of p-simple points. *Pattern Recognition Letters*, 29(6):822–827, 2008.
- [90] Christophe Lohou and Gilles Bertrand. A new 3d 12-subiteration thinning algorithm based on p-simple points. *Electronic Notes in Theoretical Computer Science*, 46:33–52, 2001.

- [91] Christophe Lohou and Gilles Bertrand. Two symmetrical thinning algorithms for 3d binary images, based on p-simple points. *Pattern Recognition*, 40(8):2301–2314, 2007.
- [92] Christophe Lohou and Julien Dehos. Automatic correction of ma and sonka’s thinning algorithm using p-simple points. *IEEE transactions on pattern analysis and machine intelligence*, 32(6):1148–1152, 2010.
- [93] Roberto A Lotufo and Francisco A Zampirolli. Fast multidimensional parallel euclidean distance transform based on mathematical morphology. In *Proceedings XIV Brazilian Symposium on Computer Graphics and Image Processing*, pages 100–105. IEEE, 2001.
- [94] Kevin M Lynch and Todd D Murphey. Control of nonprehensile manipulation. In *Control problems in robotics*, pages 39–57. Springer, 2003.
- [95] C Min Ma. Connectivity preservation of 3d 6-subiteration thinning algorithms. *Graphical models and image processing*, 58(4):382–386, 1996.
- [96] C Min Ma and Milan Sonka. A fully parallel 3d thinning algorithm and its applications. *Computer vision and image understanding*, 64(3):420–433, 1996.
- [97] Cherng-Min Ma and Shu-Yen Wan. Parallel thinning algorithms on 3d (18, 6) binary images. *Computer Vision and Image Understanding*, 80(3):364–378, 2000.
- [98] Jaehwan Ma, Sang Won Bae, and Sunghee Choi. 3d medial axis point approximation using nearest neighbors and the normal field. *The Visual Computer*, 28:7–19, 2012.
- [99] Zhihong Ma, Ruiming Du, Jiayang Xie, Dawei Sun, Hui Fang, Lixi Jiang, and Haiyan Cen. Phenotyping of silique morphology in oilseed rape using skeletonization with hierarchical segmentation. *Plant Phenomics*, 5:0027, 2023.
- [100] Djamel Merad. Skeletonization of two-dimesional regions using hybrid method. 2007.
- [101] Lukas Meyer, Andreas Gilson, Oliver Scholz, and Marc Stamminger. Cherry-picker: Semantic skeletonization and topological reconstruction of cherry trees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6243–6252, 2023.

- [102] Dang Minh, H Xiang Wang, Y Fen Li, and Tan N Nguyen. Explainable artificial intelligence: a comprehensive review. *Artificial Intelligence Review*, 55(5):3503–3568, 2022.
- [103] David J Montana. The kinematics of contact and grasp. *Int. J. Robot. Res.*, 7(3):17–32, 1988.
- [104] David J Montana. Contact stability for two-fingered grasps. *IEEE Trans. Robot. Autom.*, 8(4):421–430, 1992.
- [105] Martin Näf, Gábor Székely, Ron Kikinis, Martha Elizabeth Shenton, and Olaf Kübler. 3d voronoi skeletons and their usage for the characterization and recognition of 3d organ shape. *Computer Vision and Image Understanding*, 66(2):147–161, 1997.
- [106] Sabari Nathan and Priya Kansal. Skeletonnet: Shape pixel to skeleton pixel. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [107] Gábor Németh, Péter Kardos, and Kálmán Palágyi. Thinning combined with iteration-by-iteration smoothing for 3d binary images. *Graphical Models*, 73(6):335–345, 2011.
- [108] Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39:419–441, 2008.
- [109] Thibault Noël, Antoine Lehuger, Eric Marchand, and François Chaumette. Skeleton disk-graph roadmap: A sparse deterministic roadmap for safe 2d navigation and exploration. *IEEE Robot. Autom. Lett.*, 9(1):555–562, 2023.
- [110] Ingela Nyström and Örjan Smedby. Skeletonization of volumetric vascular images—distance information utilized for visualization. *Journal of Combinatorial Optimization*, 5(1):27–41, 2001.
- [111] Robert L Ogniewicz and Markus Ilg. Voronoi skeletons: theory and applications. In *CVPR*, volume 92, pages 63–69, 1992.
- [112] Kálmán Palágyi. A 3-subiteration 3d thinning algorithm for extracting medial surfaces. *Pattern Recognition Letters*, 23(6):663–675, 2002.

- [113] Kálmán Palágyi and Attila Kuba. Directional 3d thinning using 8 subiterations. In *Discrete Geometry for Computer Imagery: 8th International Conference, DGCI'99 Marne-la-Vallée, France, March 17–19, 1999 Proceedings* 8, pages 325–336. Springer, 1999.
- [114] Kálmán Palágyi and Attila Kuba. A parallel 3d 12-subiteration thinning algorithm. *Graphical Models and Image Processing*, 61(4):199–221, 1999.
- [115] Kálmán Palágyi, Juerg Tschirren, Eric A Hoffman, and Milan Sonka. Quantitative analysis of pulmonary airway tree structures. *Comput. Biol. Med.*, 36(9):974–996, 2006.
- [116] Oleg Panichev and Alona Voloshyna. U-net based convolutional neural network for skeleton extraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.
- [117] Eric Persoon and King-Sun Fu. Shape discrimination using fourier descriptors. *IEEE Trans. Syst. Man Cybern.*, 7(3):170–179, 1977.
- [118] Dmitry Petrov, Pradyumn Goyal, Vikas Thamizharasan, Vladimir Kim, Matheus Gadelha, Melinos Averkiou, Siddhartha Chaudhuri, and Evangelos Kalogerakis. Gem3d: Generative medial abstractions for 3d shape synthesis. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- [119] Florian T Pokorny, Johannes A Stork, and Danica Kragic. Grasping objects with holes: A topological approach. In *2013 IEEE international conference on robotics and automation*, pages 1100–1107. IEEE, 2013.
- [120] Markus Przybylski, Tamim Asfour, and Rüdiger Dillmann. Planning grasps for robotic hands using a novel object representation based on the medial axis transform. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1781–1788, 2011.
- [121] Markus Przybylski, Mirko Wächter, Tamim Asfour, and Rüdiger Dillmann. A skeleton-based approach to grasp known objects with a humanoid robot. In *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*, pages 376–383. IEEE, 2012.

- [122] Chris Pudney. Distance-ordered homotopic thinning: a skeletonization algorithm for 3d digital images. *Computer vision and image understanding*, 72(3):404–413, 1998.
- [123] Hongxing Qin, Jia Han, Ning Li, Hui Huang, and Baoquan Chen. Mass-driven topology-aware curve skeleton extraction from incomplete point clouds. *IEEE Trans. Vis. Comput. Graph.*, 26(9):2805–2817, 2020.
- [124] Christian Ronse. A topological characterization of thinning. *Theoretical Computer Science*, 43:31–41, 1986.
- [125] Christian Ronse. Minimal test patterns for connectivity preservation in parallel thinning algorithms for binary digital images. *Discrete Applied Mathematics*, 21(1):67–79, 1988.
- [126] Robert JT Sadleir and Paul F Whelan. Fast colon centreline calculation using optimised 3d topological thinning. *Computerized Medical Imaging and Graphics*, 29(4):251–258, 2005.
- [127] Punam K Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja. A survey on skeletonization algorithms and their applications. *Pattern Recognit. Lett.*, 76:3–12, 2016.
- [128] Punam K Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja. Skeletonization and its applications—a review. *Skeletonization*, pages 3–42, 2017.
- [129] Punam K Saha and Bidyut Baran Chaudhuri. 3d digital topology under binary transformation with applications. *Computer vision and image understanding*, 63(3):418–429, 1996.
- [130] Punam K Saha, Bidyut Baran Chaudhuri, and D Dutta Majumder. A new shape preserving parallel thinning algorithm for 3d digital images. *Pattern recognition*, 30(12):1939–1955, 1997.
- [131] Michiel Schaap, Coert T Metz, Theo van Walsum, Alina G van der Giessen, Annick C Weustink, Nico R Mollet, Christian Bauer, Hrvoje Bogunović, Carlos Castro, Xiang Deng, et al. Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical image analysis*, 13(5):701–714, 2009.

- [132] Eric Schneider, Sushanth Jayanth, Abhisesh Silwal, and George Kantor. 3d skeletonization of complex grapevines for robotic pruning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3278–3283. IEEE, 2023.
- [133] Jeong Chang Seong, Karen A Mulcahy, and E Lynn Usery. The sinusoidal projection: A new importance in relation to global image data. *Prof. Geogr.*, 54(2):218–225, 2002.
- [134] Luca Serino, Carlo Arcelli, and Gabriella Sanniti di Baja. From skeleton branches to object parts. *Computer Vision and Image Understanding*, 129:42–51, 2014.
- [135] Bruce Shapiro, Jim Pisa, and Jack Sklansky. Skeleton generation from x, y boundary sequences. *Computer Graphics and Image Processing*, 15(2):136–153, 1981.
- [136] Wei Shen, Kai Zhao, Yuan Jiang, Yan Wang, Xiang Bai, and Alan Yuille. Deep-skeleton: Learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images. *IEEE Transactions on Image Processing*, 26(11):5298–5311, 2017.
- [137] Wei Shen, Kai Zhao, Yuan Jiang, Yan Wang, Zhijiang Zhang, and Xiang Bai. Object skeleton extraction in natural images by fusing scale-associated deep side outputs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 222–230, 2016.
- [138] Evan C Sherbrooke, Nicholas M Patrikalakis, and Erik Brisson. An algorithm for the medial axis transform of 3d polyhedral solids. *IEEE transactions on visualization and computer graphics*, 2(1):44–61, 1996.
- [139] Xiaohan Shi, Kun Zhou, Yiying Tong, Mathieu Desbrun, Hujun Bao, and Bain-ing Guo. Mesh puppetry: cascading optimization of mesh deformation with inverse kinematics. In *ACM SIGGRAPH 2007 papers*, pages 81–es. 2007.
- [140] Kaleem Siddiqi, Sylvain Bouix, Allen Tannenbaum, and Steven W Zucker. Hamilton-jacobi skeletons. *International Journal of Computer Vision*, 48:215–231, 2002.

- [141] Kaleem Siddiqi and Stephen Pizer. *Medial representations: mathematics, algorithms and applications*, volume 37. Springer Science & Business Media, 2008.
- [142] Amos Sironi, Vincent Lepetit, and Pascal Fua. Multiscale centerline detection by learning a scale-space distance transform. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2697–2704, 2014.
- [143] André Sobiecki, Andrei Jalba, and Alexandru Telea. Comparison of curve and surface skeletonization methods for voxel shapes. *Pattern Recognit. Lett.*, 47:147–156, 2014.
- [144] André Sobiecki, Haluk C Yasan, Andrei C Jalba, and Alexandru C Telea. Qualitative comparison of contraction-based curve skeletonization methods. In *Mathematical Morphology and Its Applications to Signal and Image Processing: 11th International Symposium, ISMM 2013, Uppsala, Sweden, May 27-29, 2013. Proceedings 11*, pages 425–439. Springer, 2013.
- [145] Gim Song Soh, Shaohui Foong, Kristin Wood, et al. De-coupled dynamics control of a spherical rolling robot for waypoint navigation. In *2017 IEEE international conference on cybernetics and intelligent systems (CIS) and IEEE conference on robotics, automation and mechatronics (RAM)*, pages 562–567. IEEE, 2017.
- [146] Andrew Specian, Caio Mucchiani, Mark Yim, and Jungwon Seo. Robotic edge-rolling manipulation: A grasp planning approach. *IEEE Robotics and Automation Letters*, 3(4):3137–3144, 2018.
- [147] Johannes A Stork, Florian T Pokorny, and Danica Kragic. A topology-based object representation for clasping, latching and hooking. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 138–145. IEEE, 2013.
- [148] Yuuta Sugiyama and Shinichi Hirai. Crawling and jumping by a deformable robot. *The International journal of robotics research*, 25(5-6):603–620, 2006.
- [149] Hari Sundar, Deborah Silver, Nikhil Gagvani, and Sven Dickinson. Skeleton based shape matching and retrieval. In *2003 Shape Modeling International.*, pages 130–139. IEEE, 2003.

- [150] Mikhail Svinin and Shigeyuki Hosoe. Motion planning algorithms for a rolling sphere with limited contact area. *IEEE Trans. Robot.*, 24(3):612–625, 2008.
- [151] Wennie Tabib, Kshitij Goel, John Yao, Curtis Boirum, and Nathan Michael. Autonomous cave surveying with an aerial robot. *IEEE Transactions on Robotics*, 38(2):1016–1032, 2021.
- [152] Seyed Amir Tafrishi, Mikhail Svinin, and Kenji Tahara. A survey on path planning problem of rolling contacts: Approaches, applications and future challenges. *arXiv preprint arXiv:2501.04442*, 2025.
- [153] Seyed Amir Tafrishi, Mikhail Svinin, Motoji Yamamoto, and Yasuhisa Hirata. A geometric motion planning for a spin-rolling sphere on a plane. *Applied Mathematical Modelling*, 121:542–561, 2023.
- [154] Andrea Tagliasacchi, Ibraheem Alhashim, Matt Olson, and Hao Zhang. Mean curvature skeletons. In *Computer Graphics Forum*, volume 31, pages 1735–1744. Wiley Online Library, 2012.
- [155] Andrea Tagliasacchi, Thomas Delame, Michela Spagnuolo, Nina Amenta, and Alexandru Telea. 3d skeletons: A state-of-the-art report. In *Computer Graphics Forum*, volume 35, pages 573–597. Wiley Online Library, 2016.
- [156] Andrea Tagliasacchi, Hao Zhang, and Daniel Cohen-Or. Curve skeleton extraction from incomplete point cloud. *ACM Trans. Graph.*, 28(3):1–9, July 2009.
- [157] Xiaojun Tang, Rui Zheng, and Yinghao Wang. Distance and edge transform for skeleton extraction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2136–2141, 2021.
- [158] Alexandru C Telea and Andrei C Jalba. Computing curve skeletons from medial surfaces of 3d shapes. In *Theory and Practice of Computer Graphics*. Eurographics UK, 2012.
- [159] Brian CS Tom, Serafim N Efstratiadis, and Aggelos K Katsaggelos. Motion estimation of skeletonized angiographic images using elastic registration. *IEEE Transactions on medical imaging*, 13(3):450–460, 1994.
- [160] Jun-ichiro Toriwaki and Kensaku Mori. Distance transformation and skeletonization of 3d pictures and their applications to medical images. In *Digital and Image Geometry: Advanced Lectures*, pages 412–429. Springer, 2002.

- [161] Jun-ichiro Toriwaki and Shigeki Yokoi. Distance transformations and skeletons of digitized pictures with applications. In *Progress in pattern recognition*, pages 187–264. Elsevier, 1981.
- [162] Nhon H Trinh and Benjamin B Kimia. Skeleton search: Category-specific object recognition and segmentation using a skeletal shape model. *International Journal of Computer Vision*, 94(2):215–240, 2011.
- [163] Yea-Fu Tsao and King Sun Fu. A parallel thinning algorithm for 3-d pictures. *Computer graphics and image processing*, 17(4):315–331, 1981.
- [164] Stavros Tsogkas and Iasonas Kokkinos. Learning-based symmetry detection in natural images. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VII 12*, pages 41–54. Springer, 2012.
- [165] Nikolaus Vahrenkamp, Eduard Koch, Mirko Waechter, and Tamim Asfour. Planning high-quality grasps using mean curvature object skeletons. *IEEE Robot. Autom. Lett.*, 3(2):911–918, 2018.
- [166] Anastasiia Varava, Danica Kragic, and Florian T Pokorny. Caging grasps of rigid and partially deformable 3-d objects with double fork and neck features. *IEEE Trans. Robot.*, 32(6):1479–1497, 2016.
- [167] Alessandro Verri and Claudio Uras. Metric-topological approach to shape representation and recognition. *Image Vis. Comput.*, 14(3):189–207, 1996.
- [168] Alessandro Verri, Claudio Uras, Patrizio Frosini, and Massimo Ferri. On the use of size functions for shape analysis. *Biol. Cybern.*, 70(2):99–107, 1993.
- [169] Martin G Wagner. Real-time thinning algorithms for 2d and 3d images using gpu processors. *Journal of real-time image processing*, 17(5):1255–1266, 2020.
- [170] Ningna Wang, Bin Wang, Wenping Wang, and Xiaohu Guo. Computing medial axis transform with feature preservation via restricted power diagram. *ACM Transactions on Graphics (TOG)*, 41(6):1–18, 2022.
- [171] Sen Wang, Jianhuang Wu, Mingqiang Wei, and Xin Ma. Robust curve skeleton extraction for vascular structures. *Graphical Models*, 74(4):109–120, 2012.

- [172] Xingtao Wang, Wenxue Cui, Ruiqin Xiong, Xiaopeng Fan, and Debin Zhao. Fcnet: Learning noise-free features for point cloud denoising. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(11):6288–6301, 2023.
- [173] Yu-Shuen Wang and Tong-Yee Lee. Curve-skeleton extraction using iterative least squares optimization. *IEEE Transactions on Visualization and Computer Graphics*, 14(4):926–936, 2008.
- [174] Ofir Weber, Olga Sorkine, Yaron Lipman, and Craig Gotsman. Context-aware skeletal shape deformation. In *Computer Graphics Forum*, volume 26, pages 265–274. Wiley Online Library, 2007.
- [175] Qingmeng Wen, Seyed Amir Tafrishi, Ze Ji, and Yu-Kun Lai. Glskeleton: A geometric laplacian-based skeletonisation framework for object point clouds. *IEEE Robot. Autom. Lett.*, pages 1–7, 2024.
- [176] Uland Wong and William Whittaker. Planetary pits and caves 3d analog dataset, 2019. Accessed: 2025-02-06.
- [177] James Zachary Woodruff and Kevin M. Lynch. Robotic contact juggling. *IEEE Trans. Robot.*, 39(3):1964–1981, 2023.
- [178] Lan Wu, Raphael Falque, Victor Perez-Puchalt, Liyang Liu, Nico Pietroni, and Teresa Vidal-Calleja. Skeleton-based conditionally independent gaussian process implicit surfaces for fusion in sparse to dense 3d reconstruction. *IEEE Robot. Autom. Lett.*, 5(2):1532–1539, 2020.
- [179] Sheng Wu, Weiliang Wen, Boxiang Xiao, Xinyu Guo, Jianjun Du, Chuanyu Wang, and Yongjian Wang. An accurate skeleton extraction approach from 3d point clouds of maize plants. *Frontiers in plant science*, 10:248, 2019.
- [180] Sidong Wu, Qingqing Yang, Liuquan Ren, Jiajia Liu, Jianying Yuan, and Zhenbao Luo. 3d human curve skeleton extraction based on solid-state lidar. *IEEE Access*, 2024.
- [181] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 803–814, June 2023.

- [182] Peipei Xia, Li Zhang, and Fanzhang Li. Learning similarity with cosine similarity ensemble. *Information sciences*, 307:39–52, 2015.
- [183] Wenjie Xie, Robert P Thompson, and Renato Perucchio. A topology-preserving parallel 3d thinning algorithm for extracting the curve skeleton. *Pattern Recognition*, 36(7):1529–1544, 2003.
- [184] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. Rignet: Neural rigging for articulated characters. *arXiv preprint arXiv:2005.00559*, 2020.
- [185] Qiao Yang, Zhizhong Kang, Teng Hu, Zhen Cao, Chenming Ye, Dongming Liu, Haoxiang Hu, and Shuai Shao. Cross-section extraction and model reconstruction of lava tube based on 11-medial skeleton. *International Journal of Applied Earth Observation and Geoinformation*, 132:104062, 2024.
- [186] Alexander You, Cindy Grimm, Abhisesh Silwal, and Joseph R Davidson. Semantics-guided skeletonization of upright fruiting offshoot trees for robotic pruning. *Comput. Electron. Agric.*, 192:106622, 2022.
- [187] Liyang Yu, Qi Han, and Xiamu Niu. An improved contraction-based method for mesh skeleton extraction. *Multimedia tools and applications*, 73:1709–1722, 2014.
- [188] Shenli Yuan, Lin Shao, Connor L Yako, Alex Gruebele, and J Kenneth Salisbury. Design and control of roller grasper v2 for in-hand manipulation. In *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pages 9151–9158, 2020.
- [189] Shenli Yuan, Lin Shao, Connor L Yako, Alex Gruebele, and J Kenneth Salisbury. Design and control of roller grasper v2 for in-hand manipulation. In *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, pages 9151–9158, 2020.
- [190] Shenli Yuan, Shaoxiong Wang, Radhen Patel, Megha Tippur, Connor L Yako, Mark R Cutkosky, Edward Adelson, and Kenneth Salisbury. Tactile-reactive roller grasper. *IEEE Trans. Robot.*, 2025.
- [191] Yong Zhou and Arthur W Toga. Efficient skeletonization of volumetric objects. *IEEE Transactions on visualization and computer graphics*, 5(3):196–209, 1999.

- [192] Lijun Zong, Guanqi Liang, and Tin Lun Lam. Kinematics modeling and control of spherical rolling contact joint and manipulator. *IEEE Trans. Robot.*, 39(1):738–754, 2022.