

Braiding Geometry and Topology to Study Shapes and Data

by

Christopher Fillmore

December, 2025

A thesis submitted to the
Graduate School
of the
Institute of Science and Technology Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Committee in charge:

Matthew Robinson, Chair
Herbert Edelsbrunner
Uli Wagner
Timothy Browning
Katharine Turner

The thesis of Christopher Fillmore, titled Braiding Geometry and Topology to Study Shapes and Data, is approved by:

Supervisor: Herbert Edelsbrunner, ISTA, Klosterneuburg, Austria

Signature: _____

Co-supervisor: Uli Wagner, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Timothy Browning, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Katharine Turner, Australian National University, Canberra, Australia

Signature: _____

Defense Chair: Matthew Robinson, ISTA, Klosterneuburg, Austria

Signature: _____

Signed page is on file

© by Christopher Fillmore, December, 2025

CC BY 4.0 The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution 4.0 International License. Under this license, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author.

ISTA Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I accept full responsibility for the content and factual accuracy of this work, including the data and their analysis and presentation, and the text and citation of other work.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Christopher Fillmore
December, 2025

Signed page is on file

Abstract

This thesis examines how geometry and topology intersect in the representation, transformation, and analysis of complex shapes. It considers how continuous manifolds relate to their discrete analogues, how topological structures evolve in persistence vineyards, and how tools from topological data analysis can illuminate problems in mathematical physics. Central to this exploration is the question of how structure, both geometric and topological, persists or changes under approximation, sampling, or deformation. The work develops new approaches to skeletal and grid-based representations of surfaces, reveals the full expressive capacity of persistence vineyards, and applies topological methods to the longstanding problem of equilibria in electrostatic fields. These threads braid together into a broader understanding of how topology and geometry inform one another across theory, computation, and application.

Acknowledgements

I would like to express my deepest gratitude to my advisor, Herbert Edelsbrunner, for his invaluable support and guidance throughout my doctoral studies. His remarkable insight and intuition guided me through many of the most challenging problems in this thesis. I am also grateful for his generosity in allowing me to attend conferences around the world and for encouraging me to make my research tangible through the 3D printing of numerous models.

I also thank my co-advisor, Uli Wagner, whose thoughtful questions and advice over the years contributed meaningfully to the development of my research.

I am grateful to my dissertation committee, Katharine Turner and Timothy Browning, for their time, valuable feedback, and constructive comments, which significantly strengthened this work.

I would also like to thank Mathijs Wintraecken, Erin Wolf Chambers, and Elizabeth Stephenson for inviting me to collaborate on work that forms several chapters of this thesis. It has been a genuine pleasure to work with each of them, and their creativity and insight have made these collaborations both productive and enjoyable.

I would also like to thank the Edelsbrunner and Wagner group members for their camaraderie, stimulating discussions, and the many foosball matches that provided welcome breaks and a sense of community during the long days of research.

Special thanks go to my family and friends for their encouragement and for providing much-needed distractions from work, which made the long days of research and writing far more enjoyable.

Finally, I would like to thank my partner, Valerie, for her care, support, and good humour through both the enjoyable and challenging periods of my doctoral studies.

Funding. The research presented in this thesis was funded by the DFG Collaborative Research Center TRR 109, ‘Discretization in Geometry and Dynamics’.

About the Author

Christopher Fillmore completed a BSc (Honours) in Mathematics and Computational Science at the University of Queensland and an MSc in Mathematics at the University of Bonn before joining the Edelsbrunner and Wagner groups at ISTA in 2019. His research focuses on computational topology and geometry, particularly on understanding how complex, high-dimensional structures can be represented, visualised, and their properties computed. Motivated by a fascination with the interplay between abstract mathematics and visual intuition, he has explored methods for turning topological and geometric ideas into digital and tangible forms, including the 3D printing of mathematical models. Christopher has been delighted to attend and present his work at various conferences around the world, including the Symposium on Computational Geometry (SoCG) and the Algebraic and Topological Methods in Computer Science (ATMCS) conference, where he enjoyed engaging with the broader mathematical community.

List of Collaborators and Publications

Author's publications that are part of this thesis:

1. Erin Wolf Chambers, Christopher Fillmore, Elizabeth Stephenson, and Mathijs Wintraecken. Burning or collapsing the medial axis is unstable. La Matematica, August 2025
2. Erin Chambers, Christopher Fillmore, Elizabeth Stephenson, and Mathijs Wintraecken. Braiding vineyards. arXiv preprint arXiv:2504.11203, 2025
3. Herbert Edelsbrunner, Christopher Fillmore, and Gonçalo Oliveira. Counting equilibria of the electrostatic potential. arXiv preprint arXiv:2501.05315, 2025

Section 3.2 is based on (1). Chapter 5 is based on (2). Chapter 6 is based on (3).

Author's publications that are not part of this thesis:

1. Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. Tight bounds for the learning of homotopy à la Niyogi, Smale, and Weinberger for subsets of Euclidean spaces and of Riemannian manifolds. In 40th International Symposium on Computational Geometry (SoCG 2024). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024
2. Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. The ultimate frontier: An optimality construction for homotopy inference (media exposition). In 40th International Symposium on Computational Geometry (SoCG 2024). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024

Table of Contents

Abstract	vii
Acknowledgements	viii
About the Author	ix
List of Collaborators and Publications	x
Table of Contents	xi
List of Figures	xii
List of Tables	xvii
1 Introduction	1
2 Mathematical Background	3
2.1 Topology and Geometry	3
2.2 Manifolds	3
2.3 Cell Complexes	5
2.4 Voronoi Tessellations and Delaunay Mosaics	6
2.5 Morse Theory	7
2.6 Homology	8
2.7 Persistent Homology	9
2.8 List of Notation	13
3 Medial Axes	17
3.1 Introduction	17
3.2 Instability of the Medial Axis	18
3.3 Higher Order Medial Axes	25
3.4 Discussion	32
4 Delaunay Surfaces and Paths	35
4.1 Introduction	35
4.2 Preliminaries	36
4.3 Distortion of Mean Curvature	42
4.4 Distortion of Writhe	46
4.5 Discussion	53
5 Braiding Vineyards	55

5.1	Introduction	55
5.2	Preliminaries	56
5.3	Monodromy in Vineyards	58
5.4	Avoiding Monodromy in \mathbb{R}^2	74
5.5	Discussion	76
5.A	Example	77
5.B	Monodromy touching the diagonal	78
6	Counting Equilibria of the Electrostatic Potential	83
6.1	Introduction	83
6.2	Method	86
6.3	Regular and Semi-Regular configurations	88
6.4	1-parameter Family of Potentials	96
6.5	Distance Functions	99
6.6	Discussion	102
6.A	Proof of Theorem 13	103
Bibliography		111

List of Figures

2.1	The Reidemeister moves. <u>Left to right</u> : Type I, Type II, Type III.	4
2.2	A neighbourhood of a saddle point in a surface together with the tangent plane at the saddle point and the two planes of principal curvature which intersect along the normal vector.	5
2.3	A sample $A \subseteq \mathbb{R}^2$, light blue points with its corresponding Delaunay tessellation, dark blue edges and triangles. The edges of dual Voronoi mosaic domains are shown as black lines in the background.	7
2.4	A torus $(\mathbb{S}^1 \times \mathbb{S}^1)$ embedded in \mathbb{R}^3 . In the figure black dots depict points of \mathbf{H}_0 -persistence, the red dots depict points in \mathbf{H}_1 -persistence, the arrows indicate cycles that live forever (black for \mathbf{H}_0 , red for \mathbf{H}_1 , and green for \mathbf{H}_2).	10
2.5	A torus $(\mathbb{S}^1 \times \mathbb{S}^1)$ embedded in \mathbb{R}^3 . In the figure ordinary, relative and essential persistence points are depicted as dots, squares and triangles respectively. The colour of a persistence point indicates its degree, black for \mathbf{H}_0 -persistence, red for \mathbf{H}_1 -persistence, and green for \mathbf{H}_2 -persistence.	11
3.1	An ellipse with its medial axis, a straight line which is homotopy equivalent to the inside of the ellipse.	17
3.2	Two similar manifolds to the ellipse in Figure 3.1 with drastically different inside medial axes.	18
3.3	The singularities of the medial axis of a generic surface embedded in \mathbb{R}^3 (including the ‘non-singularity’, that is the surface). The double lines indicate open ends.	19

3.4 The generic transitions of singularities (perestroikas) of the medial axis in \mathbb{R}^3 . Adapted from figures by Bogaevsky [Bog02]. We use the same notation as [GKP08], which in turn follows Arnold. As we will see in Figure 3.9 (see also Figure 3.11) it is the transition of type A_1^4 that causes the (inside) medial axis to become non-collapsible (during the deformation of the cube into a thickened Bing's house).	20
3.5 Various pruning methods, from left to right: Object angles [AM96, DZ04], (circum)radius of the set of closest points [CL05] (the λ -medial axis, also used in our computation), and a burning method proposed in [YSC ⁺ 16], with various undesirable features indicated. The value of the object angle, radius of the set of closest points, and burning time (referred to as erosion thickness in [YSC ⁺ 16]) is indicated in colour on top. Reproduced from [YSC ⁺ 16].	21
3.6 The fire front progression on the medial axis (grey) of a curve (black). As the fire front (indicated by the red dot) hits an unburned junction, it stops. If the junction is already burned (with the colour indicating the burn time) the fire continues. As in Figure 3.5 the burning time is indicated in colour.	21
3.7 The deformation retraction of a solid cube (topological ball) to Bing's house. In the final frame we show the smoothed version of a thickened Bing's house used in the computation.	22
3.8 The various parts of Bing's house indicated.	23
3.9 This deformation (top six panels), which cuts a flap open, makes Bing's house collapsible. The topology of this transition (in the neighbourhood of the point where the flap disappears) is precisely that of the type A_1^4 . The transition of type A_1^4 is depicted in the bottom three panels; the strata are coloured to indicate the correspondence with the construction directly above.	24
3.10 The evolution of the inside medial axis (yellow) as the smoothed solid cube is deformed into a thickened version of Bing's house (blue/purple). The smoothed solid cubes in these five frames coincide with the unsmoothed solid cubes in the first five frames of Figure 3.7. The exact moment at which the medial axis becomes non-collapsible lies between frames four and five and is depicted explicitly in Figure 3.11.	26
3.11 The critical transition of the medial axis is highlighted in red. There are points on the medial axis equidistant to the two parts of the room that wrap around the corridor, the corridor itself, and the exterior wall, which can be avoided by a small perturbation. This transition occurs between frames 4 and 5 of Figure 3.7.	26
3.12 An ellipse (in black) and its evolute (in red). The standard (0th) medial axis (or insphere axis) is the gray horizontal line, and the 1st medial axis (or circumsphere axis) is the vertical gray line.	27
3.13 <u>From left to right</u> : The insphere axis of a piecewise linear ellipsoid; the midsphere axis; and the circumsphere axis. All axes are numerically computed and hence consist of polygons whose edges are drawn in black.	28
3.14 The Hasse diagram of the local neighborhood of a vertex, which is contained in four lines and six planes, such that each line is contained in three of the planes, while each plane contains only two of the lines.	30
3.15 The line arrangement inside a plane (the bisector of points A, B). The line equidistant to points A, B, C is in this plane, but also in the planes that are the bisectors of A, C and B, C , respectively. One of the vertices on this line is equidistant to points A, B, C, D	31

3.16 The sphere is centered at an interior point of a facet in the bisector arrangement and passes through the two points for which this facet belongs to the bisector. The sphere is a mid-sphere of the input surface if the two points are saddles relative to this sphere, and there are separable patches on \mathcal{M} connecting A to B inside the sphere as shown.	31
3.17 A close up of the midsphere axis from Figure 3.13. Solid black lines separate regions where one of the pair A, B changes, whereas light dashed lines separate regions where one of the points in the links of A or B enters or exits the contact sphere.	33
3.18 A two dimensional example of an isotopy guided by the burning of the medial axis.	34
4.1 We may approximate the disk bounded by the red circle with the pixels whose centers lie within the disk.	36
4.2 <u>Left</u> : the Delaunay path of the circle on the square grid sample. <u>Right</u> : the Delaunay path of the same circle on a less regular sample.	37
4.3 <u>Top</u> : Three input piecewise linear surfaces, a Schönhardt polyhedron, a torus and a bunny (red). <u>Middle</u> : The Delaunay surfaces generated from an integer lattice. <u>Bottom</u> : The Delaunay surfaces generated from a Poisson–Delaunay mosaic.	38
4.4 <u>From left to right</u> : pinching at a vertex where the surface, \mathcal{M} , intersects a Voronoi 3-cell in two disks; pinching along an edge because \mathcal{M} intersects a Voronoi polygon in two arcs; pinching along a triangle because \mathcal{M} intersects a Voronoi edge in two points.	39
4.5 The intersection of a Voronoi polygon with a closed surface. The shading indicates the inside of the surface. In all three cases, there is pinching along the corresponding Delaunay edge. <u>Left</u> : pinching of the outside. <u>Middle</u> : pinching of the inside. <u>Right</u> : simultaneous pinching of the outside and the inside.	42
4.6 The expected angle of two Delaunay wedges, first with input a wedge of two triangles that share an edge (orange) and secondly an input wedge of two large rectangles representing a wedge of half planes (green), both with dihedral angle $\pi/2$ (blue).	45
4.7 Sign rule for writhe of a knot diagram: a crossing is assigned $+1$ if the under-crossing strand passes the over-crossing strand from right to left, and -1 if it passes the over-crossing strand from left to right. Viewed from the opposite direction, under- and over-passes switch and so do left and right, so the assigned sign stays the same.	47
4.8 The spherical quadrangle defined by two non-consecutive edges of a polygonal knot. The sign of its area depends on the orientation of the two edges along the knot and is defined by the sign rule illustrated in Figure 4.7.	48
4.9 The barycentric subdivision of the Delaunay mosaic of eight points in \mathbb{R}^2 , before moving its vertices to the selected points of the Voronoi cells on the <u>left</u> and after moving them on the <u>right</u>	50
5.1 Here we see a cover \tilde{X} (in this case a double cover) of the base space X , in this case a circle as well as the curve γ and its lift $\tilde{\gamma}$	57

5.2	An illustration of the map $CV_{\mathcal{M}}$. We indicate the persistence diagrams, which form the fibers, only explicitly in a number of places for reasons of visibility. The knot diagram below, defines ouroboros knots. Here, we define the ouroboros as an unknot whose knot diagram is a (finite or segment of a) spiral with the end points connected with a monotone curve (where monotonicity refers to the relation between the angular and radial coordinates) such that the connecting segment of the knot diagram crosses over all the intermediate strands.	59
5.3	A schematic demonstrating the effect of the elder rule on our construction. Note that the gray part of the knot does not contribute to monodromy, as the elder rule dictates that the first birth is paired with the last death which causes interference with the desired braiding. By adding a new strand to the outside of the annulus via a trivial twist, we ensure that there is an elder vine that sits far away from the diagonal and doesn't interfere with our desired braid. To simplify the image, the gray portion is not shown in the vineyard.	61
5.4	The angle estimate for $\angle \beta - x, s - x$	62
5.5	The construction for the angle $\phi = \angle(s', s - x')$	63
5.6	A figure illustrating the statement of Corollary 2, highlighting the intersection of the ouroboros (in yellow) and a family of 3D growing spheres that highlight the n births and n deaths in \mathbf{H}_0	66
5.7	<u>Top:</u> The closed braid. <u>Bottom:</u> The embedding with equivalent height function. The arrow indicates the direction of the height function. The b^i 's indicate the birth times and D^J , where J is a roman numeral (i.e. $J \in \{I, II, III\}$), the death times. We stress that the picture should be interpreted in a 3D way, and in particular D^{II} does not have to be larger than D^J , with $J \neq II$. We stress that the critical point with the highest value of the Morse function (D^{II} in the figure) corresponds to a death only in extended persistence, in non-extended persistence, only a 1-cycle is born there.	67
5.8	For a given link (top left), we construct a closed braid (that is, a braid with ends identified) embedded near an annulus (bottom left). Note that we add an additional trivial loop in each connected braid component traveling around the exterior of the annulus. We then partially twist the annulus (top right) such that the cylindrical vertical portion in gray and the horizontal portion in yellow are perpendicular to each other. In the top right figure, the top edge is shown in red and the bottom edge in dark blue to clarify this twisting. It looks pinched due to the perspective, but is in fact only twisted. The bottom right figure shows the result of twisting the braid together with the annulus, with some massaging to keep the crossings in the horizontal part. Due to the addition of the exterior trivial loops from the bottom left figure, the vineyard consisting of the persistence diagrams of the distance function to a point following the twisted annulus (bottom right) on the 'outside' contains the input braid with some surgery. See Figure 5.13 for several 3D views of the embedded knot, as well as 5.16 and the corresponding detailed example and discussion in Appendix 5.A.	69
5.9	The braid follows a twisted annulus and the observation loop is in its proper place, but without the manipulation of the strands near the antipodal points of the crossing there can be incorrect crossings in the vineyard.	70
5.10	Figure illustrating the notation (the coordinates $(\theta, \omega, \tilde{\psi})$) in step 3b of Theorem 11. The observation loop is the outer loop shown in dark blue, and (the flat part of) the twisted annulus is shown in grey.	70

5.11 <u>left</u> : We see the particular closed braid that we called the ouroboros. <u>Right</u> : We see the θ and ω coordinates in the vineyard, where the h coordinate is in the direction orthogonal to the plane and \tilde{R} is the distance from the observation loop to the origin in the annulus. Here we identify ω with $\tilde{R} - b^j$ and h with $D^{J'} - 2\tilde{R}$, where $D^{J'}$ is the death time of the cycle born at b^j , corresponding to a strand as indicated in the figure. The elder rule induces surgery indicated with a dashed line (as depicted on the right). Note these are simplified sketches.	72
5.12 Manipulating the $\tilde{\psi}$ changes the death times (vertical section of Figure 5.8, not to scale).	73
5.13 <u>Top</u> : a side view of a 3-dimensional embedding of the link depicted in Figure 5.8, where the crossings and strands are perturbed in a particular way to get the desired monodromy (see the description in Figure 5.8). The red and gold portions are the two sections of the link, and the blue curve is an observation loop for the radial transform, where the dark blue portion corresponds to the section of vineyard visualized in the bottom. See also the proof of Theorem 10 and the corresponding Example in Appendix 5.A for the full details of the perturbation. <u>Middle</u> : Front-angled view of the embedded link. <u>Bottom</u> : A sideways view of our computed vineyard, showing the diagrams computed from the radial transform of observation points taken from the fraction of the full period (from 0 to 2π) depicted in dark blue segment of curve above. This segment of the vineyard captures all crossings and exhibits monodromy of period $2\pi \cdot 3$. To transform this braided vineyard into a closed braid, we would identify the “sides”, or slices, at 0 and 2π	80
5.14 Minimal example of monodromy in the plane. <u>Left</u> : Curve (blue), evolute (green), and symmetry set (red). <u>Right</u> : Corresponding extended persistence vineyard exhibiting monodromy in both H_0 and H_1 demonstrating a non-trivial permutation.	81
5.15 The crossings and the specific correspondence between the deaths or critical points and the strands indicated.	81
5.16 Different views of the Example. <u>Top</u> : together with observation loop (blue) and twisted annulus (transparent gray). <u>Bottom</u> : front view of just the manipulated trefoil-circle link.	82
6.1 One dimensional section of the electrostatic potential of a single charge.	84
6.2 One dimensional section of the electrostatic potential of a single charge capped so as to make it Morse.	85
6.3 One dimensional section of the electrostatic potential of three charges along a line.	85
6.4 Two dimensional section of the electrostatic potential of three charges arranged as the vertices of an equilateral triangle. Lighter contours indicate higher values of the potential. Equilibria are indicated by red crosses.	86
6.5 <u>From left to right</u> : the binary functions on the unit 2-sphere for a non-critical point, a minimum, a 1-saddle, a 2-saddle, and a maximum.	88
6.6 <u>From left to right</u> : Platonic Solids following the order in Table 6.2.	89
6.7 <u>From left to right</u> : the binary functions for the centers of the tetrahedron, cube, octahedron, dodecahedron, and isocahedron (these are degenerate equilibria and so are neither 1-saddles nor 2-saddles).	89
6.8 Plot of $f'(x)$ whose zeroes correspond to electrostatic points lying in the line $P_1 \cap P_2$	91
6.9 <u>From left to right, top to bottom</u> : Archimedean Solids following the order in Table 6.4	92

6.10	From left to right, top to bottom: the Catalan solids following the order in Table 6.5	93
6.11	Anti-prisms for $k = 4, 5, 6$	94
6.12	From left to right: the hexagonal, pentagonal, square anti-prisms with the heights chosen to maximize the number of equilibria. The ratios of equilibria over vertices are $37/12 < 31/10 < 25/8$, respectively. Observe how a ring of alternating 1- and 2-saddles gets successively more concentrated around the center.	95
6.13	Cut-away views of three level sets of V_1 (upper row) and three level sets of $V_{1,3}$ (lower row) defined by point sources at the vertices of the octahedron. From left to right: the values are chosen slightly less than, equal to, and slightly greater than the potential at the center of the octahedron. Removing the front of the surface reveals some of the complication at the center. For V_1 there is a degenerate equilibrium (see Figure 6.5 for the local homology), but for $V_{1,3}$ there is a minimum with a single point in the level set at the center.	97
6.14	Equilibria of the electrostatic potential generated by unit point charges at the vertices of the truncated octahedron. In total there are 36 light blue 2-saddles, 18 dark red 1-saddles, and the degenerate equilibrium at the center. For better visualization, we split these equilibria into two groups, with one 1-saddle and four 2-saddles near each of the six squares displayed in the left panel, and one 1-saddle as well as one 2-saddle for each of the twelve edges shared by two hexagons in the right panel.	100
6.15	The equilibria of V_2 generated by unit point charges at the vertices of the truncated octahedron. Compared with V_1 , we note a drastically reduced number of 1-saddles and a minimum at the origin; see Figure 6.14 where we used two copies of the solid to show all equilibria. While $p = 2$ is still not large, the equilibria are already close to the barycenters of the facets and edges of the solid.	101

List of Tables

2.1	Notation used in the thesis.	15
4.1	Summary of known and unknown distortion factors for intrinsic volumes in low dimensional ambient spaces. The well known intrinsic volumes for a surface are listed, however considering instead a curve in the plane (or space), ϕ_1 would represent the distortion factor of length.	36
4.2	Summary of average and expected distortion factors, $\mathcal{D}_{p,d}$, for p -volumes in d dimensional Euclidean space [EN25].	39
6.1	The ranks in local homology at the points whose binary functions are illustrated in Figure 6.5.	88
6.2	The ranks in local homology at the centers of the Platonic solids, with binary functions shown in the lower row of Figure 6.7. Observe that the last column in Table 6.3 is the alternating sum of these ranks.	89

6.3	Numerical results for the electrostatic potential defined by placing the point charges at the vertices of the five Platonic solids. The first entry in each f -vector is the number of point charges. The next two columns give the number of observed 2- and 1-saddles. In each case, the center is a degenerate equilibrium, for which we give the alternating sum of ranks in local homology. In each case, the number vertices exceeds this alternating sum plus the number of 2-saddles minus the number of 1-saddles by 1.	90
6.4	The equilibria of the electrostatic potential of point charges at the vertices of the thirteen Archimedean solids and the Elongated Square Gyrobicupola. The f -vector gives the number of vertices, edges, facets, and the solid itself, in this sequence. Except possibly in the last case, the center is a degenerate equilibrium, and in each case, the number of 1-saddles is non-zero; compare with Table 6.5. . . .	92
6.5	The equilibria of the electrostatic potential of point charges at the vertices of the thirteen Catalan solids and the Pseudo Deltoidal Icositetrahedron. The f -vector gives the number of vertices, edges, facets, and the solid itself, in this sequence. Except possibly in the last case, the center is a degenerate equilibrium, and in each case, the number of 1-saddles vanishes; compare with Table 6.4.	93
6.6	Growth of the various bounds discussed in this chapter. Note the lower bound values are only representative of the growth rate as the construction is built specifically for $n = 8^l$ vertices where l is the number of recursion layers.	104

Introduction

The mathematical study of shape and space has a long and distinguished history, culminating in two major branches of mathematics today, geometry and topology. Geometry, the elder branch of the two, can trace its origins back to the Greeks, specifically Euclid who systematised an axiomatic approach to geometry in his *Elements* [Euc56]. This work established the framework of rigorous deductive reasoning that would dominate all mathematics for centuries. Subsequently, Archimedes would contribute methods for understanding curvature and surface area, which presaged our modern analytic techniques of calculus [Arc17].

The Enlightenment and early modern period introduced the first arguments that could be deemed topological. Euler's solution to the Königsberg bridge problem in 1736 is widely regarded as the initial branching of topology from geometry, abstracting away from metric considerations to focus exclusively on connectivity [Eul36]. His formulation of the polyhedral theorem,

$$V - E + F = 2$$

was the first proof that certain properties of geometric objects are invariant under continuous deformation.

In the nineteenth century, Gauss profoundly advanced geometry by establishing that Gaussian curvature is intrinsic to a surface [Gau27], meaning it is independent of any particular embedding in \mathbb{R}^3 . Riemann would extend these geometric concepts to higher-dimensional manifolds (analogues of surfaces in higher dimensions) and provided the analytic apparatus for what would become Riemannian geometry [Rie68]. Later, Poincaré instituted an algebraic approach to topology through concepts such as the fundamental group and homology, which emphasise the role of invariants in distinguishing topological spaces [Poi95].

The twentieth century saw further refinement and unification. Notably, contributions from Noether who identified the algebraic structures underlying topological invariants [Noe21]; Milnor who related topology to functions on manifolds [Mil69]; Lefschetz who demonstrated how discrete representations of objects maintained topological properties [Lef42, Lef49]; and Whitehead who was able to reduce topological problems to algorithmic ones [Whi49a, Whi49b]. Together these works culminated in our modern algebraic and computational approaches to the study of shape today.

With the rise in computing power in recent decades, this trajectory has converged with data science in the form of Topological Data Analysis (TDA). At its core, TDA seeks to extract

qualitative, structural information from complex data by encoding its “shape” into topological invariants. Persistent homology, persistence vineyards, and related constructions allow the tracking of features across scales [ELZ02, CSEH05], providing a bridge between discrete representations of data and their underlying geometric structures.

This thesis is intended to continue in this tradition. With this in mind our aims are threefold.

Firstly, we attempt to analyse and numerically compare the topological and geometric properties of smooth manifolds to particular discrete approximations representations. The first such attempt is encountered in Chapter 3, where we encounter the classical medial axis and introduce higher order variants. These are skeletal representations of a manifold which capture essential structural information about the manifold, such as homotopy type. We observe that the classical medial axis cannot be simplified in a stable way and we present a novel algorithm for the computation of the higher order variants. Subsequently, in Chapter 4 we investigate the Delaunay surface which can be considered a grid representation of a manifold generated from a suitable sample of an ambient space. We prefer to work with samples drawn from a Poisson point process which result in randomised grid representations that are unbiased to any particular direction. Interestingly, we find that the geometric properties of these representations in \mathbb{R}^3 are distorted on average by fixed amounts dependent only on the dimensions of the manifold and ambient space.

Secondly, we aim to understand and describe the complex topology of persistence vineyards, which have become of greater interest recently due to efficacy of the Persistent Homology Transform (PHT) [TMB14] in shape analysis. This is covered in Chapter 5 where we demonstrate that any knot or link (or their corresponding closed braid) can be obtained as a persistence vineyard, meaning vineyards are as topologically rich as one could hope for.

Our final goal is to apply tools from TDA to a concrete problem in physics. In Chapter 6 we tackle Maxwell’s electrostatic conjecture [Max73] using persistence. Maxwell conjectured that the number of equilibria points of the electric field generated by a set of n point charges in \mathbb{R}^3 , that is, points in space where forces exerted by all point charges are perfectly balanced, is bounded from above by $(n - 1)^2$. We aim to demonstrate how topological methods can generate novel insights into this long outstanding problem, by finding charge arrangements that exceed currently known upper bounds for the number of equilibria and suggest that a quadratic upper bound is unlikely to be achievable.

CHAPTER 2

Mathematical Background

An effort is made in this Chapter to extract common concepts that recur in subsequent chapters, minimizing redundancy and establishing the necessary preliminaries in geometry, topology, and data analysis.

2.1 Topology and Geometry

A central distinction between geometry and topology lies in the class of transformations under which objects are considered equivalent. In geometry, we focus on transformations like: rotations; translations; and reflections. Under these transformations properties such as length and angle are preserved. Topology loosens these restrictions by considering transformations that are: continuous; bijective; and have a continuous inverse, which we call homeomorphisms. Homeomorphisms maintain the overall connectivity and structure of an object but may drastically alter the aforementioned geometric quantities. The classical example of deforming a doughnut into a coffee mug, without cutting or gluing, demonstrates how drastically the geometry of an object can be distorted by a homeomorphism. In the remainder of this section we will explore some basic properties of 1 and 2-manifolds.

2.2 Manifolds

A d -manifold without boundary is a topological space \mathcal{M} , such that for each point $p \in \mathcal{M}$ there exists a neighbourhood of p , $U \subset \mathcal{M}$, such that U is homeomorphic to \mathbb{R}^d for some $d \in \mathbb{Z}$ with $d > 0$. \mathcal{M} is compact if within every covering of \mathcal{M} with open sets there exists a finite covering.

A d -manifold with boundary is a topological space \mathcal{M} , such that for each point $p \in \mathcal{M}$ there exists a neighbourhood of p , $U \subset \mathcal{M}$, such that U is homeomorphic to \mathbb{R}^d or the closed d -halfspace, $\mathbb{R}_{\geq 0}^d$.

We can quantify the difference between manifolds, or any subsets of a metric space with distance function ρ , using the *Hausdorff distance*,

$$\rho_H(\mathcal{M}, \mathcal{N}) := \max \left\{ \sup_{p \in \mathcal{M}} \inf_{q \in \mathcal{N}} \rho(p, q), \sup_{q \in \mathcal{N}} \inf_{p \in \mathcal{M}} \rho(p, q) \right\}.$$

2.2.1 Knots and Links

The set of distinct compact 1-manifolds without boundary up to homeomorphism is quite limited, containing only \mathbb{S}^1 . Specifically in three dimensions, we can add specificity by considering equivalence classes of these embedded circles.

An (oriented) knot is the equivalence class of oriented closed curves embedded in 3-dimensional Euclidean space, $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$, under continuous deformations that never self intersect. To simplify the discussion, we will abuse notation and write γ for the map as well as its image in \mathbb{R}^3 . A link with n components is a disjoint union of n knots, $L = \gamma_1 \cup \dots \cup \gamma_n \subset \mathbb{R}^3$. We say two knots (or links) are equivalent if there is an ambient isotopy $\mathbb{R}^3 \times I \rightarrow \mathbb{R}^3$ such that one knot (or link) is the image of the other, and the map at each time in the interval is also a homeomorphism. Generally, this equivalence is formalized as a series of Reidemeister moves (see Figure 2.1, as any isotopy between knots can be related by a sequence of these three local moves [AB26, Rei27]).

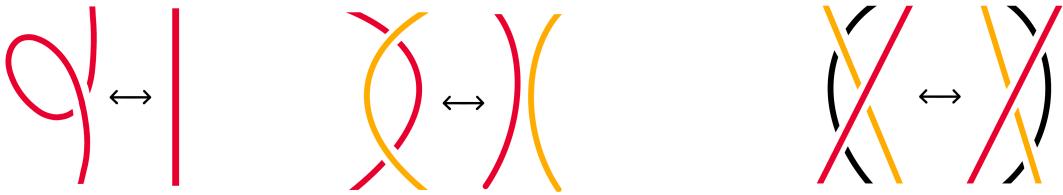


Figure 2.1: The Reidemeister moves. Left to right: Type I, Type II, Type III.

For each $u \in \mathbb{S}^2$, the projection of a knot (or link), γ , in the direction u provides a knot diagram (or link diagram) of γ with crossings determined by the relative heights in direction u . Throughout this thesis we assume diagrams with generic crossings, to avoid self tangencies and infinite numbers of crossings.

2.2.2 Surfaces

The term surface is more commonplace than 2-manifold so it is preferred here. The classification theorem for surfaces without boundary says there exists two infinite families $\mathbb{S}^2, \mathbb{T}, \mathbb{T} \# \mathbb{T} \dots$ and $\mathbb{P}^2, \mathbb{P}^2 \# \mathbb{P}^2, \dots$. Where \mathbb{T} is the torus; \mathbb{P}^2 is the real projective plane, the surface obtained by identifying antipodal points of \mathbb{S}^2 ; and $\#$ represents the connected sum operation which connects two surfaces by removing a disc from each and gluing an end of a cylinder along each of the newly created boundaries. The first family corresponds to the orientable surfaces, these are the surfaces for which one can take a small oriented circle and slide it around any closed loop in the surface and obtain a circle with the same orientation on the circle. For the second family of surfaces there exist loops which reverse this orientation, these are known as the non-orientable surfaces. Due to the classification theorem, the orientable surfaces (which are simply the connected sums of tori) can be embedded in \mathbb{R}^3 . The non-orientable surfaces on the other hand can only be shown to be embedded in \mathbb{R}^4 as a result of the Whitney embedding theorem [Whi44].

Focusing on orientable surfaces in \mathbb{R}^3 , we can consider some of their geometric properties which will be important for later chapters. Consider a smoothly embedded surface $S \subset \mathbb{R}^3$. At each point $p \in S$, there are infinitely many possible normal sections, curves obtained by slicing the surface with planes that contain the surface's normal vector at that point. Each such curve bends with some curvature, the principal curvatures are the maximum and minimum of

those normal curvatures, which we label $\kappa_1(S) \geq \kappa_2(S)$. The sign of the principal curvatures is inherited from the orientation on S . The normal sections corresponding to the principal curvatures together with the tangent plane at p form an orthonormal frame, see Figure 2.2.

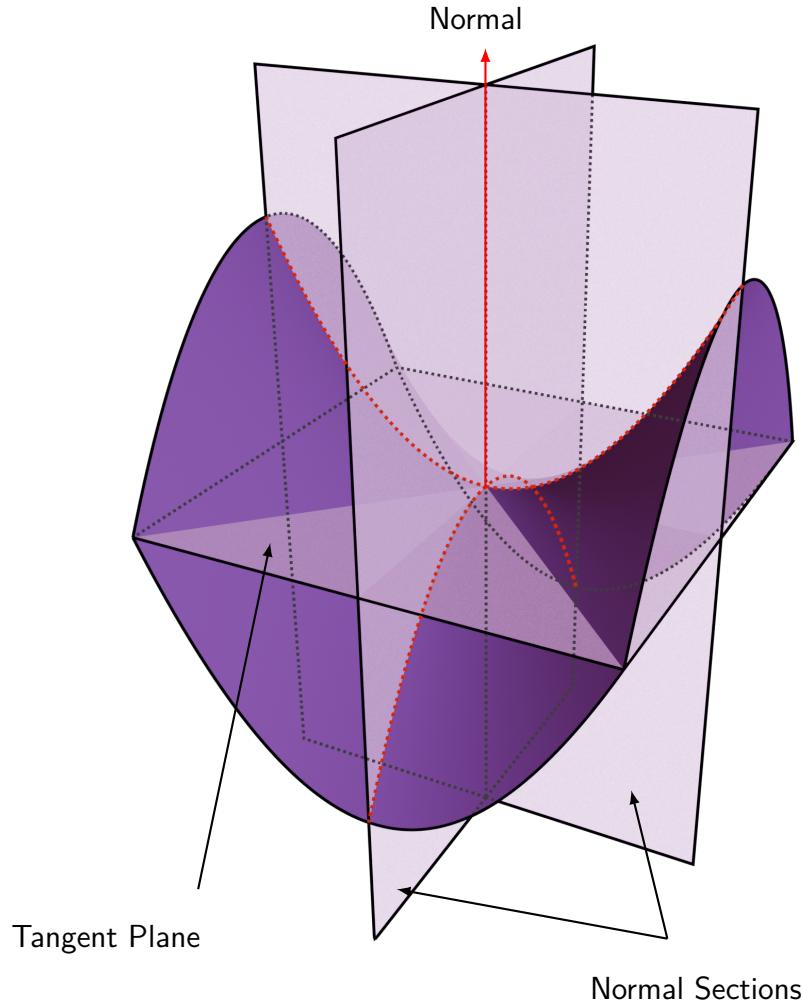


Figure 2.2: A neighbourhood of a saddle point in a surface together with the tangent plane at the saddle point and the two planes of principal curvature which intersect along the normal vector.

Integrating the mean of these curvatures across the entire surface we obtain the *total mean curvature*, $\text{Mean}(S)$, whereas integrating their product we obtain the *Gaussian curvature*, $\text{Gauss}(S)$. Explicitly,

$$\begin{aligned}\text{Mean}(S) &= \frac{1}{2} \int_{p \in S} (\kappa_1(p) + \kappa_2(p)) dp, \\ \text{Gauss}(S) &= \int_{p \in S} \kappa_1(p) \kappa_2(p) dp.\end{aligned}$$

2.3 Cell Complexes

Given a finite set of vertices V , an abstract simplicial complex, K , is a system of subsets of V , $\sigma \in K$ called simplices, which is closed under taking subsets. By which we mean that if $\sigma \in K$ and $\tau \subset \sigma$, then τ must also be an element of K and we call τ a face of σ . If the cardinality

of σ is $p + 1$, $\text{card } \sigma = p + 1$, then we call σ a *p-simplex* or a *p-dimensional simplex* and the $p = 0, 1, 2, 3$ simplices are known as vertices, edges, triangles and tetrahedra respectively. The dimension of the entire simplicial complex is given by the maximum dimension amongst its simplices. Let L be a subset of the simplices of K , again closed under taking subsets, then L is a subcomplex of K . An important collection of subcomplexes are the *p-skeleta* of K , denoted $K^{(p)}$, which include only the p -simplices of K and their lower dimensional faces.

If V is in fact a set of disjoint points in \mathbb{R}^d , in other words $g : V \rightarrow \mathbb{R}^d$ is a injective map, we can realise the simplex as the convex hull of its vertices, $\text{Conv}(\sigma)$, provided its vertices are affinely independent. Additionally, if the intersection of two simplices is either empty or a common face of both, that is, $\text{Conv}(\sigma) \cap \text{Conv}(\tau) = \text{Conv}(\sigma \cap \tau)$, then we obtain a proper realisation of K , whose underlying space is denoted $|K|$.

If we were instead interested in complexes where faces can be arbitrary polyhedra or polytopes we can define polyhedral complexes. Now K is a collection of subsets of $V \subset \mathbb{R}^d$ that: form polytopes; are closed under taking faces; and the intersection of any two faces are empty or another face of K . Notably, polyhedral complexes arise in the context of Voronoi tessellations and cubical complexes which will be important later in this thesis.

2.4 Voronoi Tessellations and Delaunay Mosaics

Let $A \subseteq \mathbb{R}^d$ be locally finite. A is called generic if for every $1 \leq p \leq d$, no $p + 2$ points lie on a common $(p - 1)$ -sphere in \mathbb{R}^d . For each point, $a \in A$, we write $\text{dom}(a, A)$ for the points $x \in \mathbb{R}^d$ that satisfy $\|x - a\| \leq \|x - b\|$ for all $b \in A$. We refer to $\text{dom}(a, A)$ as the Voronoi domain of a in A and note that it is a convex polyhedron with faces of dimension $0 \leq i \leq d$. Any two Voronoi domains have disjoint interiors but they may share common faces of dimension strictly less than d . The interiors of the Voronoi domains and all their faces form a partition of \mathbb{R}^3 into open convex polyhedra. The Voronoi tessellation of A , denoted $\text{Vor}(A)$, is the collection of Voronoi domains [Vor08]. For $0 \leq i \leq d$, we call the i -dimensional faces of the Voronoi domains the i -cells of $\text{Vor}(A)$. While $\text{Vor}(A)$ is not a cell complex, the collection of its cells is. See the cells bordered with black lines in Figure 2.3.

The Delaunay mosaic of A , denoted $\text{Del}(A)$, is dual to the Voronoi tessellation [Del34]. Specifically, for each i -cell of $\text{Vor}(A)$, $\text{Del}(A)$ contains the convex hull of the points generating the Voronoi domains that share the i -cell. By construction, this convex hull has dimension $j = d - i$, and we refer to it as a j -cell of $\text{Del}(A)$. Observe that $\text{Del}(A)$ is a polyhedral cell complex and generically a simplicial complex. We refer to the 1 and 2-cells as (Delaunay) edges and polygons, and use letters e and p to denote individual such edges and polygons. In the familiar setting of \mathbb{R}^3 , the corresponding cells in $\text{Vor}(A)$ are e^* , which is the Voronoi polygon dual to e , and p^* , which is the Voronoi edge dual to p . See the shaded cells bordered with blue lines in Figure 2.3.

The Voronoi tessellation and Delaunay Mosaic will be essential to subsequent chapters. In Section 2.7 we will see that Delaunay mosaics are fundamental to persistent homology which is essential to Chapters 5 and 6. The Alpha complex filtration, which tracks the homotopy type of a union of balls of increasing radii, is a subset the Delaunay mosaic. The medial axis constructions from Chapter 3 can also be interpreted as the continuous limit of the Voronoi tessellations of a sample of points a given manifold (including higher order Voronoi tessellations). While the titular Delaunay surfaces and paths from Chapter 4 obviously owe their name to the Delaunay complex, the Voronoi tessellation is also intrinsic to the definition

due to their mutual duality. Finally, we encounter the Voronoi tessellation and its weighted variant in Chapter 6 as we study the limit of a family of functions that converge to the distance function.

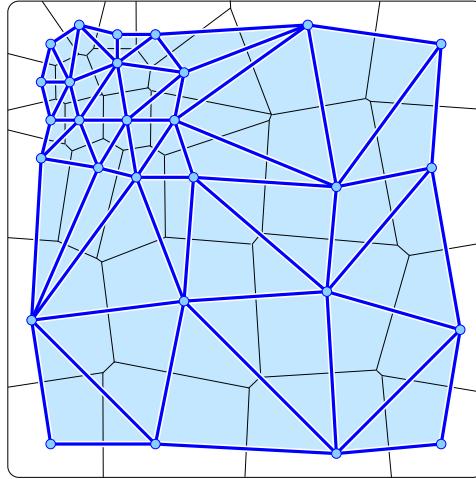


Figure 2.3: A sample $A \subseteq \mathbb{R}^2$, light blue points with its corresponding Delaunay tessellation, dark blue edges and triangles. The edges of dual Voronoi mosaic domains are shown as black lines in the background.

2.5 Morse Theory

Morse theory is a fundamental topic in topology which allows one to discern topological features of a manifold from differentiable functions on that manifold. We follow the work of Milnor [Mil69] for this section. Let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function on a differentiable manifold \mathcal{M} . Write $\nabla f(p)$ and $(H_f)_{i,j}(p) = \frac{\partial^2 f}{\partial x_i \partial x_j}(p)$ for the gradient and Hessian of f at p . We call p and $f(p)$ a critical point and critical value of f respectively, if $\nabla f(p) = 0$. All other values in \mathbb{R} are considered non-critical or regular. A critical point is called non-degenerate if $H_f(p)$ is invertible and degenerate otherwise. Non-degenerate critical points are isolated and functions with only non-degenerate points have only a finite number of critical points. f is generically a Morse function if it satisfies:

1. All critical points are non-degenerate
2. Each critical point has a unique value.

The Hessian of a Morse function determines the nature of its critical points. The number of negative eigenvalues is known as the index of the critical point and for $\dim \mathcal{M} = 3$ indices of 0, 1, 2 correspond to minima, saddles and maxima respectively.

Three fundamental results of Morse theory describe the relation between a Morse function, $f : \mathcal{M} \rightarrow \mathbb{R}$, and the topology of \mathcal{M} [Mil69]. If $[a, b] \subset \mathbb{R}$ is a compact interval containing no critical points of f , then the sublevel sets $\mathcal{M}_a = f^{-1}(-\infty, a]$ and $\mathcal{M}_b = f^{-1}(-\infty, b]$ are diffeomorphic and \mathcal{M}_b deformation retracts onto \mathcal{M}_a . Additionally, if p is a critical point of index k and the interval $[f(p) - \varepsilon, f(p) + \varepsilon]$ contains no other critical points, then $\mathcal{M}_{f(p)+\varepsilon}$ is homotopy equivalent to $\mathcal{M}_{f(p)-\varepsilon}$ with a k -cell attached. Using the two previous results we can deduce that any manifold admitting a Morse function is homotopy equivalent to a CW complex with a k -cell corresponding to each index k -critical point of f .

The number of index k critical points of a Morse function on \mathcal{M} , C_k , is an upper bound on the rank of the k -th homology group, $\beta_k(\mathcal{M}) \leq C_k$ known as the Betti number. The Euler characteristic, which is a topological invariant of \mathcal{M} , can also be computed from the critical points of a Morse function, $\chi(\mathcal{M}) = \sum (-1)^k C_k$. Combining these two facts Milnor was able to prove what are known as the Morse inequalities:

$$C_k - C_{k-1} + \dots + (-1)^k C_0 \geq \beta_k(\mathcal{M}) - \beta_{k-1}(\mathcal{M}) + \dots + (-1)^k \beta_0(\mathcal{M}). \quad (2.1)$$

There is a Morse inequality for each $0 \leq k < d$, and an equation for $k = d$.

2.6 Homology

In this section we present a crash course on simplicial homology with coefficients in $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$ which should provide the reader with a sufficient understanding to proceed with persistent homology in the following sections.

For a simplicial complex K , a p -chain, c , is a formal sum of p -simplices, $\sigma_i \subset K$, denoted $c = \sum a_i \sigma_i$, for $a_i \in \mathbb{Z}_2$. We can add another chain, $c_1 = \sum b_i \sigma_i$, to c in the same way one would add polynomials, $c_1 + c_2 = \sum (a_i + b_i) \sigma_i$, where the addition of the coefficients is performed in \mathbb{Z}_2 . Of course alternative coefficient groups could be selected, however, we restrict ourselves to the simplest example in this brief overview. The set of p -chains in K together with chain addition form the chain group which we denote $(\mathbf{C}_p, +)$, for each $p \in \mathbb{Z}$. The chain group has: identity, the chain with all coefficients equal to zero; the inverse of a chain is the chain itself; and associativity inherited from \mathbb{Z}_2 . Additionally, the chain group inherits commutativity from \mathbb{Z}_2 and is therefore also abelian.

The chain groups are related to each other by a set of boundary maps which intuitively link each simplex to the simplices in its boundary. Writing $\sigma = [v_0, \dots, v_p]$ for the simplex spanned by the vertices v_0, \dots, v_p we define its boundary map as:

$$\partial_p \sigma = \sum_i [v_0, \dots, \hat{v}_i, \dots, v_p],$$

where \hat{v}_i indicates that vertex v_i is not included in the simplex. The boundary map of a chain is then simply the formal sum of the boundaries of its simplices. The boundary map commutes with chain addition, $\partial_p(c + c_1) = \partial_p c + \partial_p c_1$, and is therefore a group homomorphism.

By sequentially mapping between chain groups via boundary maps we can construct what is known as the chain complex of K ,

$$\dots \xrightarrow{\partial_{p+2}} \mathbf{C}_{p+1} \xrightarrow{\partial_{p+1}} \mathbf{C}_p \xrightarrow{\partial_p} \mathbf{C}_{p-1} \xrightarrow{\partial_{p-1}} \dots$$

Each chain group contains two important subgroups, the group of cycles and the group of boundaries. Beginning with the former a p -cycle is a p -chain whose boundary is empty, $\partial_p c = 0$. Thinking of 1-chains the 1-cycles would include: the empty chain; loops; and unions of loops. The set of all p -cycles, denoted \mathbf{Z}_p , is a subgroup of the p -chain group as the p boundary map commutes with addition. We set the boundary of each vertex equal to 0 and let \mathbf{C}_{-1} equal the trivial group. Since all the p -cycles map to 0 under the ∂_p the cycle group is exactly the kernel, $\mathbf{Z}_p = \ker \partial_p$.

A p -boundary on the other hand, is a p -chain which forms the boundary of a $(p+1)$ -chain, $c = \partial_{p+1} c_1$ for $c_1 \in \mathbf{C}_{p+1}$. Again due to the commutativity of chain addition, the set of

p -boundaries, denoted \mathbf{B}_p , forms a subgroup of the p -chain group. By definition the boundary group is exactly the image of the boundary map, $\mathbf{B}_p = \text{im } \partial_{p+1}$.

The fundamental lemma of homology states that applying the boundary map to a chain twice results in the zero chain. Due to associativity we can show this by simply showing it is true for a single simplex. The boundary of a p -simplex, σ , consists of all the $(p-1)$ -faces of σ and each $(p-2)$ -face of σ appears in the exactly two $(p-1)$ -faces. When we compute the formal sum each $(p-1)$ -face will appear twice and cancel out resulting in $\partial_p(\partial_{p-1}\sigma) = 0$. Additionally, this implies that each p -boundary is also a p -cycle and \mathbf{B}_p is a subgroup of \mathbf{Z}_p .

The p -th homology group of K is then the quotient of these subgroups,

$$\mathbf{H}_p = \mathbf{Z}_p / \mathbf{B}_p = \ker \partial_p / \text{im } \partial_{p+1}.$$

Intuitively, it counts the number of p -holes, that is p -spheres that are not filled in by $(p+1)$ -discs. The rank of each group gives the corresponding Betti number of K .

For a subcomplex L of K there is a variant of homology, which will be important to us later, known as the relative homology of the pair (K, L) . Relative homology describes the topology of K after contracting L to a single point. In this setting the chain groups are given by the quotients, $\mathbf{C}_p(K, L) = \mathbf{C}_p(K) / \mathbf{C}_p(L)$ and boundary maps $\partial_p : \mathbf{C}_p(K, L) \rightarrow \mathbf{C}_{p-1}(K, L)$ are induced by those in the regular homology. The relative cycle, boundary and homology groups are then defined in the obvious way,

$$\begin{aligned} \mathbf{Z}_p(K, L) &= \ker(\partial_p : \mathbf{C}_p(K, L) \rightarrow \mathbf{C}_{p-1}(K, L)); \\ \mathbf{B}_p(K, L) &= \text{im } (\partial_{p+1} : \mathbf{C}_{p+1}(K, L) \rightarrow \mathbf{C}_p(K, L)); \\ \mathbf{H}_p(K, L) &= \mathbf{Z}_p(K, L) / \mathbf{B}_p(K, L). \end{aligned}$$

2.7 Persistent Homology

Persistent homology [Fro90, Rob99, ELZ02] is a relatively recent and active area which has its roots in algebraic topology; see [Oud15, DW22] for recent survey books on this active topic.

Whereas the classical homology, discussed in Section 2.6, is useful for studying topological features of data, it lacks the discernment to pick the correct scale to extract features from. Persistent homology remedies this shortcoming by studying nested sequences of sublevel sets called filtrations, thereby extracting features across many scales. Taking a manifold \mathcal{M} and a nice function to filter it with (for example the height function) yields the persistence module comprised of homology groups and linear maps induced by the inclusions between sublevel sets

$$\dots \rightarrow \mathbf{H}(\mathcal{M}_{a_{i-1}}) \rightarrow \mathbf{H}(\mathcal{M}_{a_i}) \rightarrow \dots \rightarrow \mathbf{H}(\mathcal{M}_{a_{j-1}}) \rightarrow \mathbf{H}(\mathcal{M}_{a_j}) \rightarrow \dots \quad (2.2)$$

Composing the maps between consecutive groups, we get a map between any two groups in the module. We say a homology class $\alpha \in \mathbf{H}(\mathcal{M}_{a_i})$ is born at \mathcal{M}_{a_i} if it is not in the image of the map from $\mathbf{H}(\mathcal{M}_{a_{i-1}})$ to $\mathbf{H}(\mathcal{M}_{a_i})$. If α is born at \mathcal{M}_{a_i} , it dies entering \mathcal{M}_{a_j} if the image of the map from $\mathbf{H}(\mathcal{M}_{a_{i-1}})$ to $\mathbf{H}(\mathcal{M}_{a_{j-1}})$ does not contain the image of α , but the image of the map from $\mathbf{H}(\mathcal{M}_{a_{i-1}})$ to $\mathbf{H}(\mathcal{M}_{a_j})$ does. The persistence of α is the difference between the function values at its birth and its death. If the function is a Morse function on a manifold, then precisely one Betti number β_p changes when the threshold passes a critical value. If the index of the corresponding critical point is p , then either a p -dimensional class is born, so β_p increases by one, or a $(p-1)$ -dimensional class dies, so β_{p-1} decreases by one. We use a

persistence diagram $\mathcal{D}_p(f)$ to encode the birth-death information of all of the p -dimensional homology classes of M arising from the sublevel set filtration induced by the filtering function f , where each birth/death pair becomes a point in $\bar{\mathbb{R}}^2 = [-\infty, \infty]^2$. $\mathcal{D}(f)$ is used when referring to all degrees. See Figure 2.4.

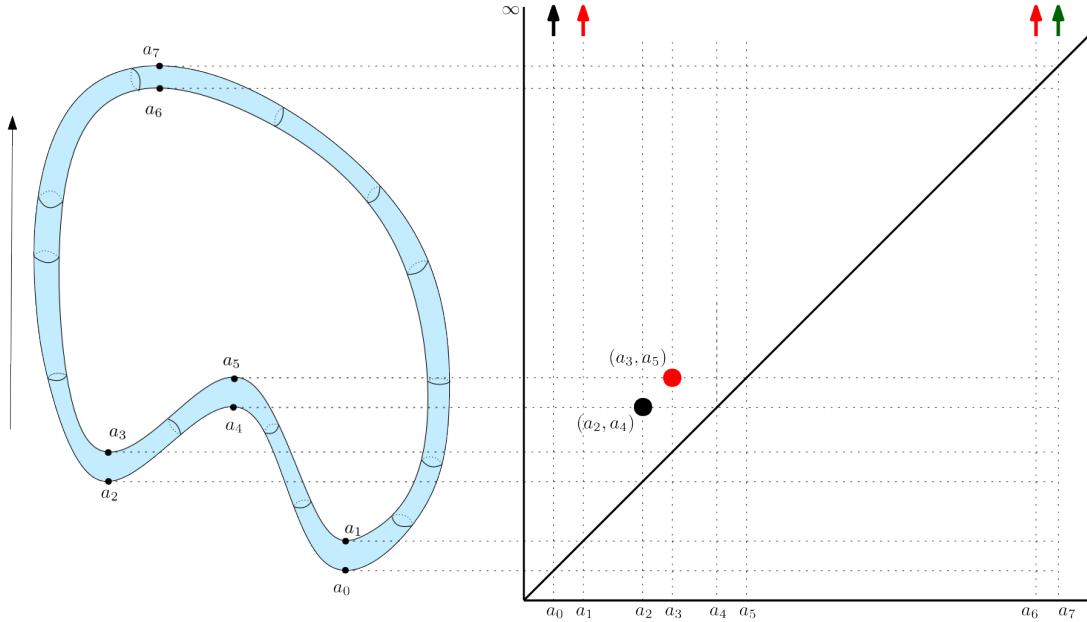


Figure 2.4: A torus ($\mathbb{S}^1 \times \mathbb{S}^1$) embedded in \mathbb{R}^3 . In the figure black dots depict points of \mathbf{H}_0 -persistence, the red dots depict points in \mathbf{H}_1 -persistence, the arrows indicate cycles that live forever (black for \mathbf{H}_0 , red for \mathbf{H}_1 , and green for \mathbf{H}_2).

2.7.1 Extended Persistent Homology

There are some drawbacks to the topological summary given by standard persistence, most notably points at infinity. For example, consider a sublevel set filtration of a shape embedded in \mathbb{R}^2 or \mathbb{R}^3 with nontrivial \mathbf{H}_0 and \mathbf{H}_1 : at some point connected components and loops are born, but never die, as they are present in all sublevel sets after their initial appearance. Even worse, if the input is non-generic and two \mathbf{H}_1 cycles are born at the same height in the sublevel set filtration, then they give rise to identical persistence pairs.

To address this, Agarwal et. al. [AEHW06] established a pairing between all critical points of a height function on a 2-manifold, which Cohen-Steiner et. al. [CSEH08] extended to general manifolds with tame functions, leveraging Poincaré and Lefschetz duality to create a new sequence of homology groups where we begin and end with the trivial group. This guarantees that each homology class that is born will also die at a finite value, replacing all the problematic points paired with ∞ , and guarantees a perfect matching on critical points.

Let $\mathbf{H}_d(\mathcal{M}, \mathcal{M}^a)$ denote the relative homology group of \mathcal{M} with the superlevel set $\mathcal{M}^a = f^{-1}[a, \infty)$. Again assume we have a tame function and set of critical values a_1, \dots, a_k , and note that $\mathbf{H}_p(\mathcal{M}_{a_k}) = \mathbf{H}_p(\mathcal{M}) = \mathbf{H}_p(\mathcal{M}, \mathcal{M}^a)$ for any $a > a_k$. From this, we can create a

new sequence of homology groups

$$\begin{aligned} 0 \rightarrow \mathbf{H}_*(\mathcal{M}_{a_1}) \rightarrow \mathbf{H}_*(\mathcal{M}_{a_2}) \rightarrow \dots \rightarrow \mathbf{H}_*(\mathcal{M}_{a_{k-1}}) \\ \rightarrow \mathbf{H}_*(\mathcal{M}_{a_k}) = \mathbf{H}_*(\mathcal{M}, \emptyset) \rightarrow \mathbf{H}_*(\mathcal{M}, \mathcal{M}^{a_k}) \\ \rightarrow \mathbf{H}_*(\mathcal{M}, \mathcal{M}^{a_{k-1}}) \rightarrow \dots \rightarrow \mathbf{H}_*(\mathcal{M}, \mathcal{M}^{a_1}) = \mathbf{H}_*(\mathcal{M}, \mathcal{M}) = 0, \end{aligned}$$

which we call the extended filtration sequence. We define the sequence $\mathbf{H}_*(\mathcal{M}_{a_1}) \rightarrow \dots \rightarrow \mathbf{H}_*(\mathcal{M}_{a_k})$ as the upwards sequence and the sequence $\mathbf{H}_*(\mathcal{M}, \mathcal{M}^{a_k}) \rightarrow \dots \rightarrow \mathbf{H}_*(\mathcal{M}, \mathcal{M}^{a_1})$ as the downwards sequence.

Note that this sequence fits the structure of a persistent module, so like standard persistence, it also has a unique interval decomposition. The only change is that we interpret the persistence points differently in this setting. Specifically, the points in the persistence diagram can be partitioned into three different groups: 1) the classes which are born and die in the upwards sequence, 2) the classes which are born and die in the downwards sequence, and 3) the classes which are born in the upwards sequence and die in the downwards sequence. Further, we associate the index of birth and death intervals with the value a_i for both $\mathbf{H}_p(\mathcal{M}_{a_i})$ in the upwards sequence and $\mathbf{H}_p(\mathcal{M}, \mathcal{M}^{a_i})$ in the downwards sequence.

This creates three different classes of persistence pairs: those that correspond to a class that is born and dies in the upwards sweep (ordinary persistence points), those that correspond to a class that is born and then dies in the downwards sweep (relative persistence points), and those that correspond to a class that is born in the upwards sweep and dies in the downwards sweep (essential persistence points). We denote the class of ordinary, relative, and essential points as $\text{Ord}_p(f)$, $\text{Rel}_p(f)$, and $\text{Ess}_p(f)$, respectively.

Note that the ordinary diagram is entirely above the diagonal; the relative diagram is entirely below the diagonal; while the essential persistence points can be on either side. See Figure 2.5.

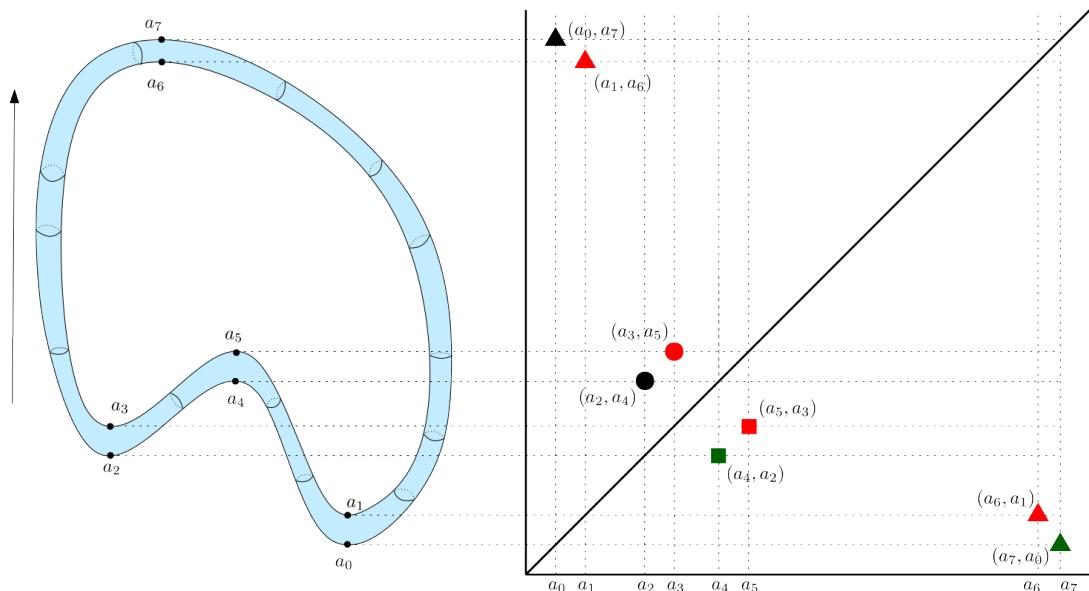


Figure 2.5: A torus ($\mathbb{S}^1 \times \mathbb{S}^1$) embedded in \mathbb{R}^3 . In the figure ordinary, relative and essential persistence points are depicted as dots, squares and triangles respectively. The colour of a persistence point indicates its degree, black for \mathbf{H}_0 -persistence, red for \mathbf{H}_1 -persistence, and green for \mathbf{H}_2 -persistence.

The symmetry of the persistence points across the diagonal observed in Figure 2.5 is independent of the manifold or filtration. In fact, it is a result of the Lefschetz duality between complementary dimensional homology of a sublevel set and relative homology of an n -manifold and the corresponding superlevel set, $\mathbf{H}_p(\mathcal{M}_a) \simeq \mathbf{H}_{n-p}(\mathcal{M}, \mathcal{M}^a)$. The following duality relations arise as a result of Lefschetz's theorem:

$$\begin{aligned}\text{Ord}_p(f) &= \text{Rel}_{n-p}^T(f); \\ \text{Ess}_p(f) &= \text{Ess}_{n-p}^T(f), \\ \text{Rel}_p(f) &= \text{Ord}_{n-p}^T(f),\end{aligned}$$

where the superscript ‘ T ’ represents reflection of the set of points across the diagonal.

2.7.2 Comparing Persistence

For a suitable metric ρ , the space of persistence diagrams is defined as:

$$\text{Dgm}_\rho = \{\mathcal{D} \subset \bar{\mathbb{R}}^2 \mid \rho(\mathcal{D}, \emptyset) < \infty\},$$

where each \mathcal{D} is a multi-set of points in $\bar{\mathbb{R}}^2$, including infinite copies of the points on the diagonal $(a, a) \in \bar{\mathbb{R}}^2$ for $a \in (-\infty, \infty)$. There are many choices of metric for Dgm_ρ , some important examples from the literature are included here. For points in $\bar{\mathbb{R}}^2$ the distance between them is given by $\|x - y\|_\infty = \max\{|x_1 - y_1|, |x_2 - y_2|\}$. For two functions $f, g : M \rightarrow \mathbb{R}$, $q \in \mathbb{R}$ and $p > 0$, the *Wasserstein distance* is defined as,

$$\rho_{W_p}(\mathcal{D}(f), \mathcal{D}(g)) := \inf_{\phi} \left(\sum_{x \in \mathcal{D}(f)} \|x - \phi(x)\|_\infty^{1/p} \right)^p,$$

taken over all bijections $\phi : \mathcal{D}(f) \rightarrow \mathcal{D}(g)$. Letting $p \rightarrow \infty$ we obtain the *Bottleneck distance*,

$$\rho_B(\mathcal{D}(f), \mathcal{D}(g)) := \inf_{\phi} \sup_{x \in \mathcal{D}(f)} \|x - \phi(x)\|_\infty.$$

Finally, we have the *Hausdorff distance*,

$$\rho_H(\mathcal{D}(f), \mathcal{D}(g)) := \max \left\{ \sup_{x \in \mathcal{D}(f)} \inf_{y \in \mathcal{D}(g)} \|x - y\|_\infty, \sup_{y \in \mathcal{D}(g)} \inf_{x \in \mathcal{D}(f)} \|x - y\|_\infty \right\}.$$

In each case ϕ ranges over all possible bijections between $\mathcal{D}(f)$ and $\mathcal{D}(g)$. Bijections include points from the diagonal in order to balance the number of off-diagonal points between diagrams and potentially further minimise the cost of a matching.

In the case of extended persistence diagrams matchings must be further restricted such that each class (ordinary, relative and extended) are matched only to the corresponding class or the diagonal.

Intuitively, the bottleneck distance describes the worst disparity between the best matching of points in persistence diagrams: the worst disparity is the “bottleneck” preventing a smaller distance.

2.7.3 Vineyards

One can also study the evolution of persistence on a manifold, \mathcal{M} , arising from changing functions. As long as the functions are similar enough, the Stability Theorem [CSEH05] of persistent homology asserts that their associated persistence diagrams will also be similar; see also [Tur23] for a discussion of the algebraic details of maps between “nearby” diagrams.

This gives rise to the concept of a vineyard [CSEM06, Tur23]. It formalizes the idea that a feature of g_u is still recognizable in g_v , provided the two functions u and v are not too far apart. Features are points in the diagram, and the association is a matching between the points of g_u and of g_v . Cohen-Steiner et al. [CSEM06] propose the so-called vineyard algorithm that traces the features (points) while continuously deforming g_u into g_v , realized as an update to the reduced matrix arising from g_u . This update can be done in $O(n)$ time, as compared to the worst case $O(n^3)$ time to compute and reduce the new matrix from scratch.

By tracing a point through a continuous family of diagrams we obtain a vine in the vineyard. For our purposes in Chapter 5 it will be advantageous to assume that no vines intersect (which is possible in this setting). However, we note that this is not generally the case. In fact, in some pathological cases intersections may not even be transversal, meaning there is no unique choice on how to extend a vine.

2.8 List of Notation

Mathematical Background

$\mathbb{Z}; \mathbb{R}; \mathbb{R}^d$	the integers; the real numbers; Euclidean d -space
$\mathcal{M} \subset \mathbb{R}^d; \mathbb{S}^d \subset \mathbb{R}^{d+1}$	a manifold; d -sphere
$p \in \mathcal{M}; x \in \mathbb{R}^d$	a point in an ambient space; a point in a manifold; d -sphere
$\ x\ ; \rho: \mathbb{R}^d \rightarrow \mathbb{R}$	Euclidean norm; general distance metric
$D^d; \mathring{D}^d; \mathbb{R}_{\geq 0}^d$	Euclidean disc; open Euclidean disc; closed d -halfspace
$\mathbb{T}; \mathbb{P}^2; S$	torus; real projective plane; a surface
$\sigma, \tau; K, L$	simplices; simplicial complexes
$\mathcal{X} \subset \mathbb{R}^d; \text{Conv}(\mathcal{X}); \mathcal{T}$	set of points in \mathbb{R}^d ; convex hull of \mathcal{X} ; a triangulation
$K^{(p)}; K $	the p -skeleton of K ; the realisation of K
$P \subset \mathbb{R}^d$	a d -polytope
$A \subseteq \mathbb{R}^3, \text{Vor}(A), \text{Del}(A)$	locally finite set, Voronoi tessellation, Delaunay mosaic
$e \in \text{Del}(A); e^* \in \text{Vor}(A)$	a Delaunay edge; a Voronoi polygon dual to e
$f: \mathcal{M} \rightarrow \mathbb{R}; H_f$	a Morse function; Hessian of f
$\mathcal{M}_a; \mathcal{M}^a$	sublevel set of M at a ; superlevel set of M at a
$C_k; \beta_k$	number of k -critical points of f ; k -th Betti number of M
$\mathbf{C}_p; \mathbf{Z}_p; \mathbf{B}_p$	p -th chain group; p -th cycle group; p -th boundary group
$\mathbf{H}_p(M)$	p -th homology group of M with coefficients in $\mathbb{Z}/2\mathbb{Z}$
$\mathcal{D}(f); \mathcal{D}_p(f)$	persistence diagram of f ; persistence diagram of degree p
$\text{Dgm}_p; \text{CV}_{\mathcal{M}}: \mathbb{S}^1 \rightarrow \mathbb{S}^1 \times \text{Dgm}$	space of persistence diagram; closed vineyard map
$\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$	(oriented) knot

Medial Axes

$\mathcal{S}; M[\mathcal{S}]$	a closed subset set of \mathbb{R}^d ; medial axis of \mathcal{S}
$M_\lambda[\mathcal{S}]$	λ -medial axis of \mathcal{S}
$M^0[\mathcal{M}]; M^1[\mathcal{M}]; M^2[\mathcal{M}]$	insphere axis; midsphere axis; circumsphere axis
$R: M^1[\mathcal{M}] \rightarrow \mathbb{R}$	radius function on mid-sphere axis
A_i^j	singularity corresponding to sphere with j type i contact points
$A, B; A_0, A_1, B_0, B_1$	points on sphere; points in link
$u = (ABCD)$	vertex (quadruplet of points)
$e = (A; BCD; E)$	edge on line (triplet of points)
$f = (A; B; CD; E; F)$	facet in plane (pair of points)

Delaunay Surfaces & Paths

$A \subseteq \mathbb{R}^3$	locally finite set
$\text{Del}(\mathcal{M}, A)$	Delaunay mosaic, restricted Delaunay complex
$\text{Del}(\gamma, A); \text{Del}'(\gamma, A); \text{Del}''(\gamma, A)$	Voronoi path; subdivided; moved
$L; \text{Del}(L, A); L'$	line; Delaunay path; skew line
M_ε	Piecewise linear approximation to \mathcal{M}
$e, p; e^*, p^*$	Delaunay edge, polygon; dual Voronoi polygon, edge
$\delta; \eta = \pi - \delta$	dihedral angle; external dihedral angle
$\mu: \mathbb{R}^3 \rightarrow \mathbb{R}^3$	rigid motion
$E \in \text{Graff}_{d,i}$	plane in the affine Grassmannian
$K; K_r$	convex set; r -parallel body
$\phi_i: \mathbb{R}^d \rightarrow \mathbb{R}$	i -th intrinsic volume
$\varphi: \mathbb{R}^d \rightarrow \mathbb{R}$	a geometric measure in \mathbb{R}^d
$\text{Wr}(\gamma); \text{Lk}(\gamma); \text{Tw}(\gamma)$	writhe; linking number; twist
$u \in \mathbb{S}^2; \text{Wr}_u(\gamma)$	direction; writhe of knot diagram in direction u
$T_t = \dot{\gamma}(t)/\ \dot{\gamma}(t)\ $	unit tangent vector; derivative
$T; -T$	tantrix; negative tantrix
$P < \gamma$	inscribed polygon
α_i	turning angle of projected path
$\alpha: \mathbb{G}_{1,3} \times \mathbb{R}^2 \rightarrow [-\pi, \pi]$	turning angle map on Grassmannian of lines
$\chi(P)$	chirality of convex polyhedron
n_+, n_-	numbers of Type-I Reidemeister moves

Braiding Vineyards

$B; B_i : I \rightarrow D^2 \times I$	a braid; a strand of a B
$\rho_{\mathbb{E}}(\cdot, p) _{\mathcal{M}}$	radial distance function

Counting Equilibria

$A_1, A_2, \dots, A_n \in \mathbb{R}^3$	finite number of points
$\zeta_1, \zeta_2, \dots, \zeta_n \in \mathbb{R}$	corresponding charges
$V: \mathbb{R}^3 \rightarrow \mathbb{R}$	electrostatic potential
m_p	number of index- p equilibria
$\theta: \mathbb{R}^3 \rightarrow \mathbb{R}^3$	similarity
$\ell; k; c; n$	no. of layers; no. of vertices; no. of equilibria; no. of point charges
$1 \leq p \leq \infty$	parameter
$V_p: \mathbb{R}^3 \rightarrow \mathbb{R}$	modified potential
$E: \mathbb{R}^3 \rightarrow \mathbb{R}$	(weighted) Euclidean distance function
$w_i; \pi_i(x), \varphi_i(x)$	weight of i -th point; (weighted) Laguerre distance
$\mathbf{i} = \sqrt{-1}$	imaginary unit
$x = (w, z), A_i = (w_i, z_i)$	points in $\mathbb{R}^3 = \mathbb{R}^2 \times \mathbb{R}$
$x = (x_1, x_2, x_3)$	coordinates in \mathbb{R}^3

Table 2.1: Notation used in the thesis.

CHAPTER 3

Medial Axes

3.1 Introduction

The term medial axis was introduced by Blum [Blu67], who suggested its use in the context of shape recognition. However, the concept predates Blum's work, likely the concept was first probed by Erdős [Erd45, Erd46].

The (inside)¹ medial axis is a skeleton that captures many geometric and topological features of an input set. Formally, the medial axis of a closed set \mathcal{S} , denoted $M[\mathcal{S}]$, consists of all points in the ambient space (usually \mathbb{R}^d) that do not have a unique closest point in \mathcal{S} .

In this chapter, the ambient space will be \mathbb{R}^d . See Figure 3.1 for a simple example in the plane. The reach of a (closed) set, $\mathcal{S} \subset \mathbb{R}^d$, was introduced by Federer [Fed59] as the minimum distance from \mathcal{S} to the medial axis, as a means to study properties of the medial axis and its complement. More recently, its closure has been reintroduced as the central set [MW80].

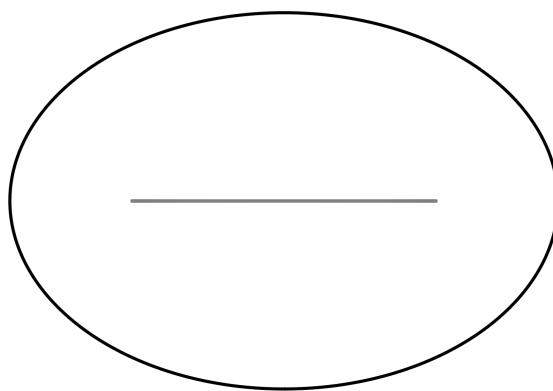


Figure 3.1: An ellipse with its medial axis, a straight line which is homotopy equivalent to the inside of the ellipse.

Lieutier proved in [Lie04] that the medial axis captures the homotopy type of the complement of \mathcal{S} , $M[\mathcal{S}] \simeq (\mathbb{R}^d \setminus \mathcal{S})$. Given the medial axis of a smooth manifold, the original manifold

¹A compact $(d-1)$ -manifold $\mathcal{M} \subset \mathbb{R}^d$ divides \mathbb{R}^d into a compact or inside part and a non-compact or outside part. By the inside medial axis we mean the part of the medial axis that lies on the inside of the compact or interior part.

may be recovered if the distance to the manifold at each point of the medial axis is known [Wol92]. Damon [Dam03, Dam04, Dam07, Dam05] investigated how to recover the differential geometry of the manifold from the geometry of the medial axis and the radius function.

In this chapter we investigate the medial axis, firstly, proving the instability of retraction based approximations; and secondly, generalising the medial axes to higher orders.

3.2 Instability of the Medial Axis

Unfortunately, the medial axis is limited in its use by its (topological) instability under small perturbations [ABE09]. Here small is understood to be small with respect to the Hausdorff distance (recall Equation (2.2)), $\rho_H(\mathcal{S}, \mathcal{S}')$. Compare Figures 3.1 and 3.2 for a standard example of such instabilities.

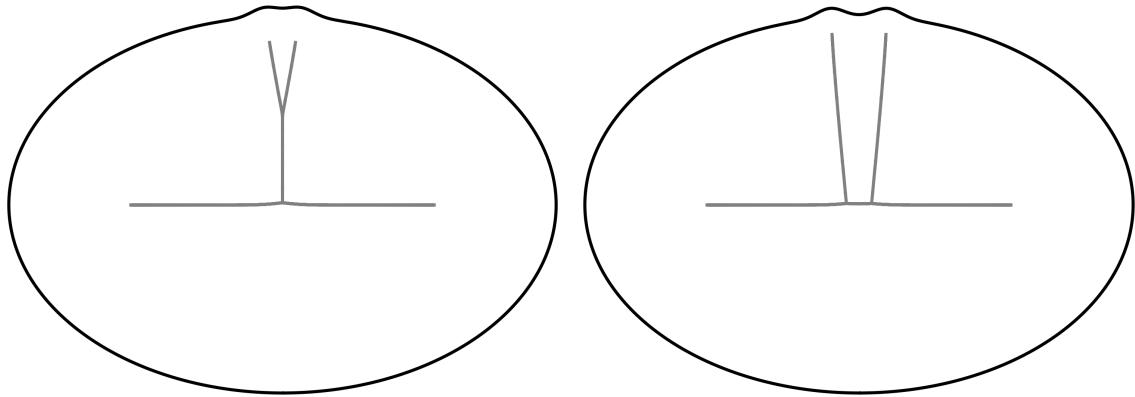


Figure 3.2: Two similar manifolds to the ellipse in Figure 3.1 with drastically different inside medial axes.

If we restrict ourselves to a smaller class of spaces and perturbations, there are stability results available. Chazal and Soufflet [CS04] proved that the medial axis is stable with respect to the Hausdorff distance under ambient diffeomorphisms, under strong conditions. Namely, Chazal and Soufflet assumed that the set, \mathcal{S} , is a C^2 manifold of positive reach and the distortion is a C^2 diffeomorphism of \mathbb{R}^d . Federer [Fed59] proved that the reach is stable under $C^{1,1}$ diffeomorphisms of the ambient space, that is, a diffeomorphism of the ambient space whose derivative and the derivative of the inverse are Lipschitz continuous. Using Federer's result Chazal and Soufflet's work [CS04] was recently improved by Dal Poz Kouřimská, Lieutier, and Wintraecken to the medial axis of any closed set with bounded complement under $C^{1,1}$ diffeomorphisms [DPKLW24].

3.2.1 Singularities, generic stability and transitions

Generically a point on the medial axis in three dimensions has two closest points, in which case the geometry of a neighbourhood (or more precisely the second fundamental form) can be described in terms of the geometry near the two closest points [vM02, Sie99, SSG99, vM03]. This means that the geometry of the medial axis is (generically) locally stable.

The topology of the singularities of the medial axis can also be classified in a generic setting in low dimensions. More precisely, any point in the medial axis of a smoothly embedded generic manifold has a neighbourhood such that the medial axis in this neighbourhood is diffeomorphic

to sets in the classification. The classification of singularities of the medial axis is the same as for the cut locus, which is, roughly speaking, the medial axis for a single point on a manifold, see Figure 3.3.

The A_i^j notation of Figure 3.3 for the types of singularities was introduced by Arnold. The subscript, i , refers to the order of contact the sphere makes with the surface and the superscript, j , indicates the number of this type of contact for this sphere. First order contact, A_1 , implies the surface and sphere share a tangent space; second order contact, A_2 , implies the one of the principal curvatures of the surface matches the curvature of the sphere; third order contact, A_3 , implies the derivative of the curvature of the surface and sphere are also equal. In the case of third order contact this further implies the surface has a critical point in curvature at the contact, since the curvature of the sphere is constant.

This classification was found in a variety of settings [Mat83, Buc78, Yom81, Arn13, Tho72, Loo74, DG14]. In particular, Chapter 13 of [DG14] extensively discusses genericity, and we refer to this chapter for full discussion on the topic. We refer the reader to Van Manen [vM07] for a more in-depth overview of this rich topic and its history, as well as to work by Damon [Dam06]. Unfortunately, the classification of singularities is only explicitly possible in sufficiently low dimensions² and in this chapter we limit ourselves to at most dimension 3.

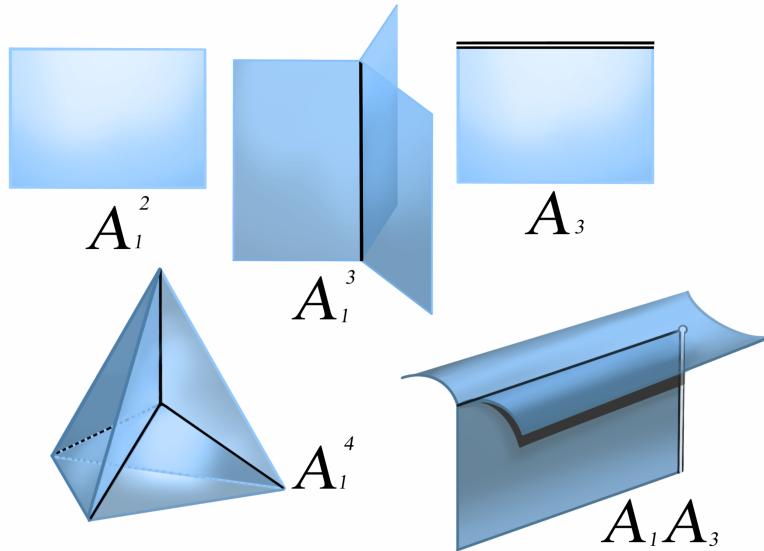


Figure 3.3: The singularities of the medial axis of a generic surface embedded in \mathbb{R}^3 (including the ‘non-singularity’, that is the surface). The double lines indicate open ends.

The generic transitions between singularities of the medial axis, known as perestroikas, in a one parameter family isotopy between (generic) surfaces in \mathbb{R}^3 have been studied by Giblin and Kimia [GKP08], based on earlier work by Bogaevsky [Bog89, Bog02, Bog99, Bog95]. We represent the transitions given by Giblin and Kimia, and Bogaevsky in Figure 3.4. The A_1^4 transition plays a significant role in this section.

²In dimensions 7 or larger, there is no stability of the singularities of the cut locus in the C^∞ sense, only in the C^0 sense, see Buchner [Buc77, Buc78], see also [Wal77, Theorem 4 (i) and (ii), page 737]. This is an essential obstruction to a classification of singularities for the cut locus on manifolds in higher dimensions.

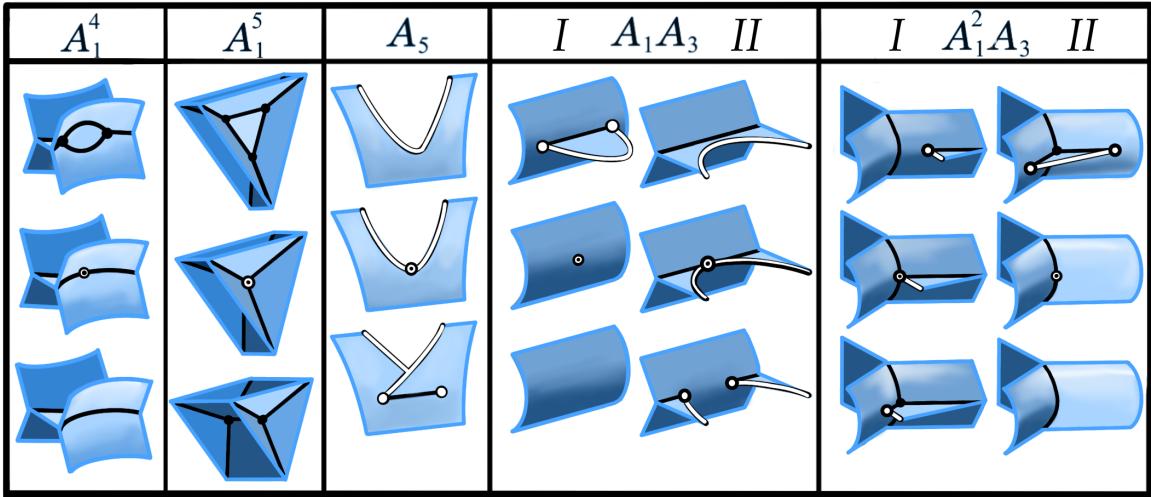


Figure 3.4: The generic transitions of singularities (perestroikas) of the medial axis in \mathbb{R}^3 . Adapted from figures by Bogaevsky [Bog02]. We use the same notation as [GKP08], which in turn follows Arnold. As we will see in Figure 3.9 (see also Figure 3.11) it is the transition of type A_1^4 that causes the (inside) medial axis to become non-collapsible (during the deformation of the cube into a thickened Bing's house).

3.2.2 Prunings

Significant effort has gone into the simplification (pruning) of the medial axis. This was motivated by applications in: graphics (where the medial axis is used as a skeleton, see the surveys [SBdB16, TDS⁺16]); data reduction; shape recognition; and machine learning (see for example [BDC⁺15, CK11, DHL⁺19, HD86, LCLJ11, SKK04, SB98, YSC⁺16]). Various prunings of the medial axis have been proposed in many different settings [ACK01, AM96, BBML⁺18, CL05, DS06, DZ04, GMPW09, LCLJ11, SB98]. See Figure 3.5 for an illustration of some commonly used prunings and their pitfalls: the object angle, which is of historic importance in the community but can disconnect the medial axis; the λ -medial axis, which is used to compute a close approximation of the medial axis but which can truncate “thin” regions undesirably; and the burning method, which we consider in this chapter.

Since the λ -medial axis is essential to our computations we briefly define it now. For $\mathcal{S} \subset \mathbb{R}^d$ and $M[\mathcal{S}]$ its medial axis, at each point $x \in M[\mathcal{S}]$ we have a radius function $r(x) = \inf_{p \in \mathcal{S}} \|x - p\|$, nearest neighbour function $\Pi(x) = \{p \in \mathcal{S} \mid \|x - p\| = r(x)\}$ and separation of contact points given by $\delta(x) = \max_{p,q \in \Pi(x)} \|p - q\|$. Then for a fixed parameter $\lambda > 0$, the λ -medial axis is given by, $M_\lambda[\mathcal{S}] = \{x \in M[\mathcal{S}] \mid \delta(x) \geq \lambda r(x)\}$.

The simplification which we focus on for this work is the burning of the medial axis [LCLJ11, YSC⁺16], which generalizes Blum’s original “grassfire” analogy for the medial axis. The burning of the medial axis removes the extremities of the medial axis by “starting a fire” at the boundary of the medial axis which stops if the fire hits an obstacle, as illustrated in Figure 3.6. Note that this illustration focuses on 2 dimensions, where the definition is simpler [LCLJ11]. The generalization of burning to 3 dimensions is significantly more complex, as the singularities are no longer points. We refer the reader to [YSC⁺16] for the full definition and examples of this generalization, as the full details are not needed for our main result.

Because of the good experimental results, it was conjectured that the burning method of simplification of the medial axis would be stable [LCLJ11], i.e. no discontinuous jumps. The

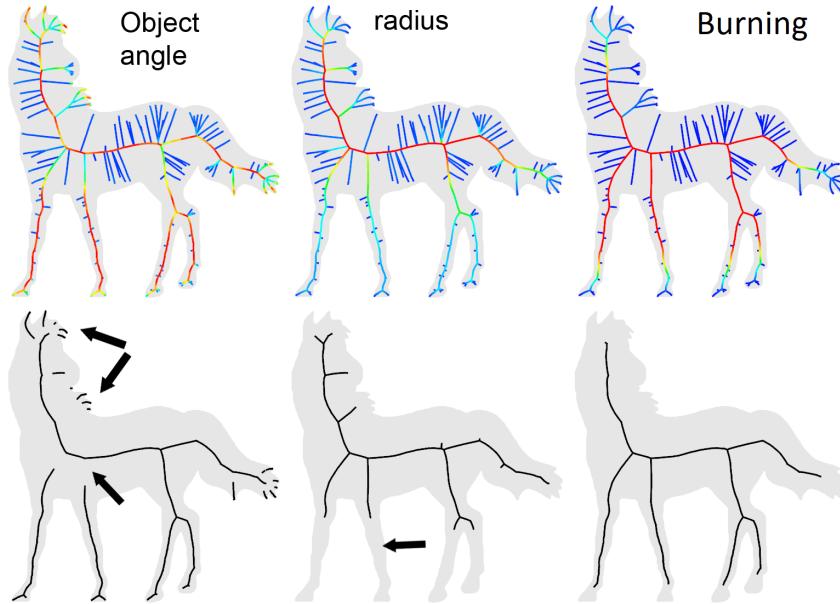


Figure 3.5: Various pruning methods, from left to right: Object angles [AM96, DZ04], (circum)radius of the set of closest points [CL05] (the λ -medial axis, also used in our computation), and a burning method proposed in [YSC⁺16], with various undesirable features indicated. The value of the object angle, radius of the set of closest points, and burning time (referred to as erosion thickness in [YSC⁺16]) is indicated in colour on top. Reproduced from [YSC⁺16].

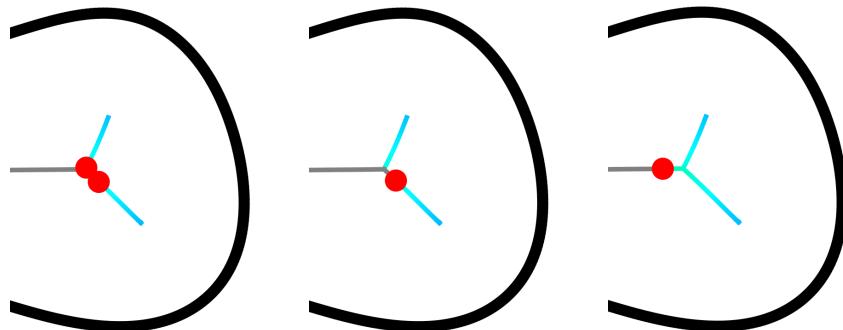


Figure 3.6: The fire front progression on a the medial axis (grey) of a curve (black). As the fire front (indicated by the red dot) hits an unburned junction, it stops. If the junction is already burned (with the colour indicating the burn time) the fire continues. As in Figure 3.5 the burning time is indicated in colour.

remainder of this section is dedicated to showing that this is not the case.

3.2.3 Burning Bing's house

The counterexample is based on the standard deformation retraction from the closed ball to Bing's house with two rooms [Bin64], which is a contractible but not collapsible two dimensional simplicial complex, see Figure 3.7. Bing's house is not collapsible, as there is no boundary from which a collapsing sequence could begin.

Bing constructed his house as a deformation retract from a solid cube, see Figure 3.7. The deformation retraction goes as follows. During the entire construction, with the exception of the final step, the set will remain three dimensional. As mentioned, we start with a solid cube. We then start by digging two corridors. Once we have reached the part of the solid

cube that will become a room, we start hollowing out the room by protruding a rectangular cuboid into the room. This can be continued until we come near the corridor that goes to the opposite room. Of course we cannot interfere with the corridor, instead we engulf the corridor from both sides. Once we are past the corridor, we extend the hollowed space until it almost touches itself near the flap. We finally flatten the three dimensional set to two dimensions, creating Bing's house.

As indicated in Figure 3.7 the deformation retraction of the solid cube to Bing's house can be done in a piecewise linear way. At every point before the final state where the set becomes two dimensional, the boundary is a piecewise linear 2-sphere. Because the boundary of the piecewise linear complex is a 2-sphere and all angles between adjacent faces of the piecewise linear complex are $\pi/2$, the boundary complex can be made $C^{1,1}$ by replacing the edges with parts of cylinders and the vertices with parts of spheres. Such $C^{1,1}$ manifolds can be smoothed, even without decreasing the reach, see [KLW24]. This means in particular that a smooth version of the deformation retraction from the smoothed solid cube to a thickening of Bing's house exists.

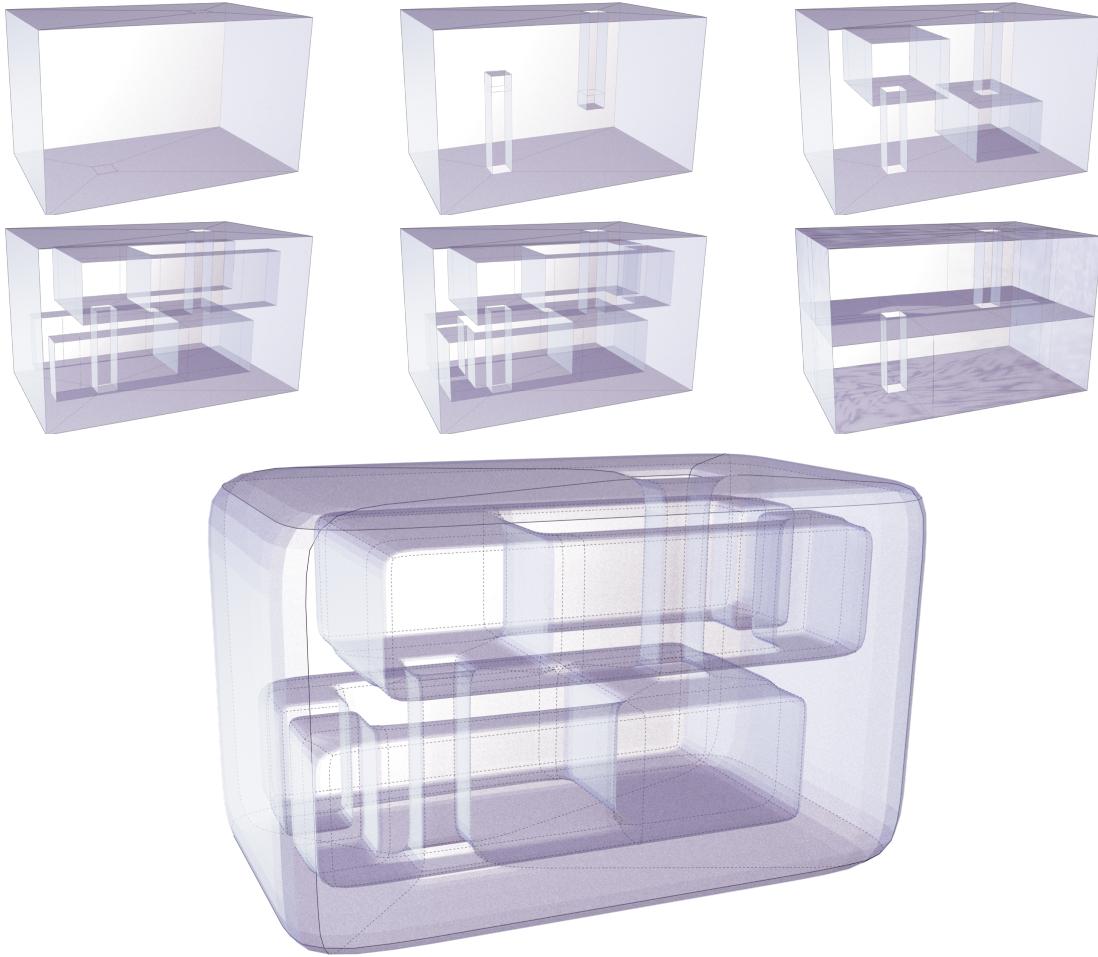


Figure 3.7: The deformation retraction of a solid cube (topological ball) to Bing's house. In the final frame we show the smoothed version of a thickened Bing's house used in the computation.

Before we go into the main statement, we consider a deformation of Bing's house which makes it collapsible. This deformation will be mirrored in the medial axis in our construction and is depicted in Figure 3.9, see Figure 3.8 for the nomenclature. In this construction we cut

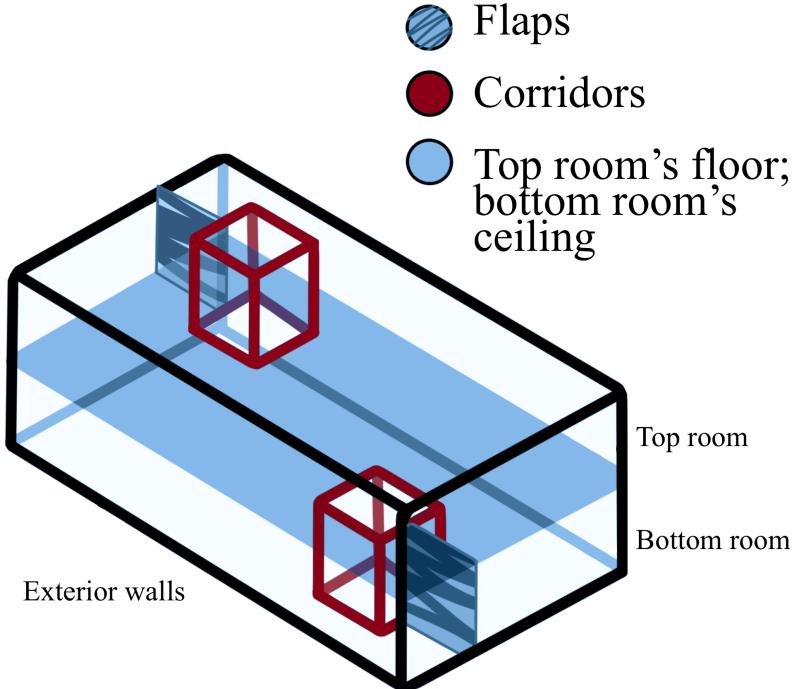


Figure 3.8: The various parts of Bing's house indicated.

a flap open so that the room no longer completely runs around the corridor. This cutting exposes an edge of one of the walls of the corridor and the path that goes from the edge to the bottom room. We can use this edge to collapse along the path into the bottom room, then the bottom room itself, and from this the rest of Bing's house.

We will now describe the counterexample. It follows the construction of Bing's house from a solid cube, but stops just before the set becomes two dimensional. This gives an isotopy of the boundary of the cube, namely \mathbb{S}^2 , such that the medial axis transitions from being collapsible and burnable down to a point, into Bing's house itself.

The precise result is the following:

Theorem 1. *There exists a smooth ambient isotopy $H_t : [0, 1] \times \mathbb{S}^2 \rightarrow \mathbb{R}^3$ such that:*

- *The (inside) medial axis $M[H_0(\mathbb{S}^2)]$ is collapsible/burnable to a single point.*
- *The (inside) medial axis $M[H_1(\mathbb{S}^2)]$ is Bing's house and is therefore non-collapsible/cannot burn (is fireproof).*
- *The burning of $M[H_t(\mathbb{S}^2)]$ is not continuous in t with respect to the Hausdorff distance and Gromov-Hausdorff (as well as the Fréchet distance which is infinite).*
- *The topology of the burned (inside) medial axis changes from a point to Bing's house with two rooms at a single $t_0 \in [0, 1]$.*
- *The isotopy H_t can be chosen to be generic in the sense of singularity theory as developed by Arnold and Thom [Arn13], see in particular [GKP08].*

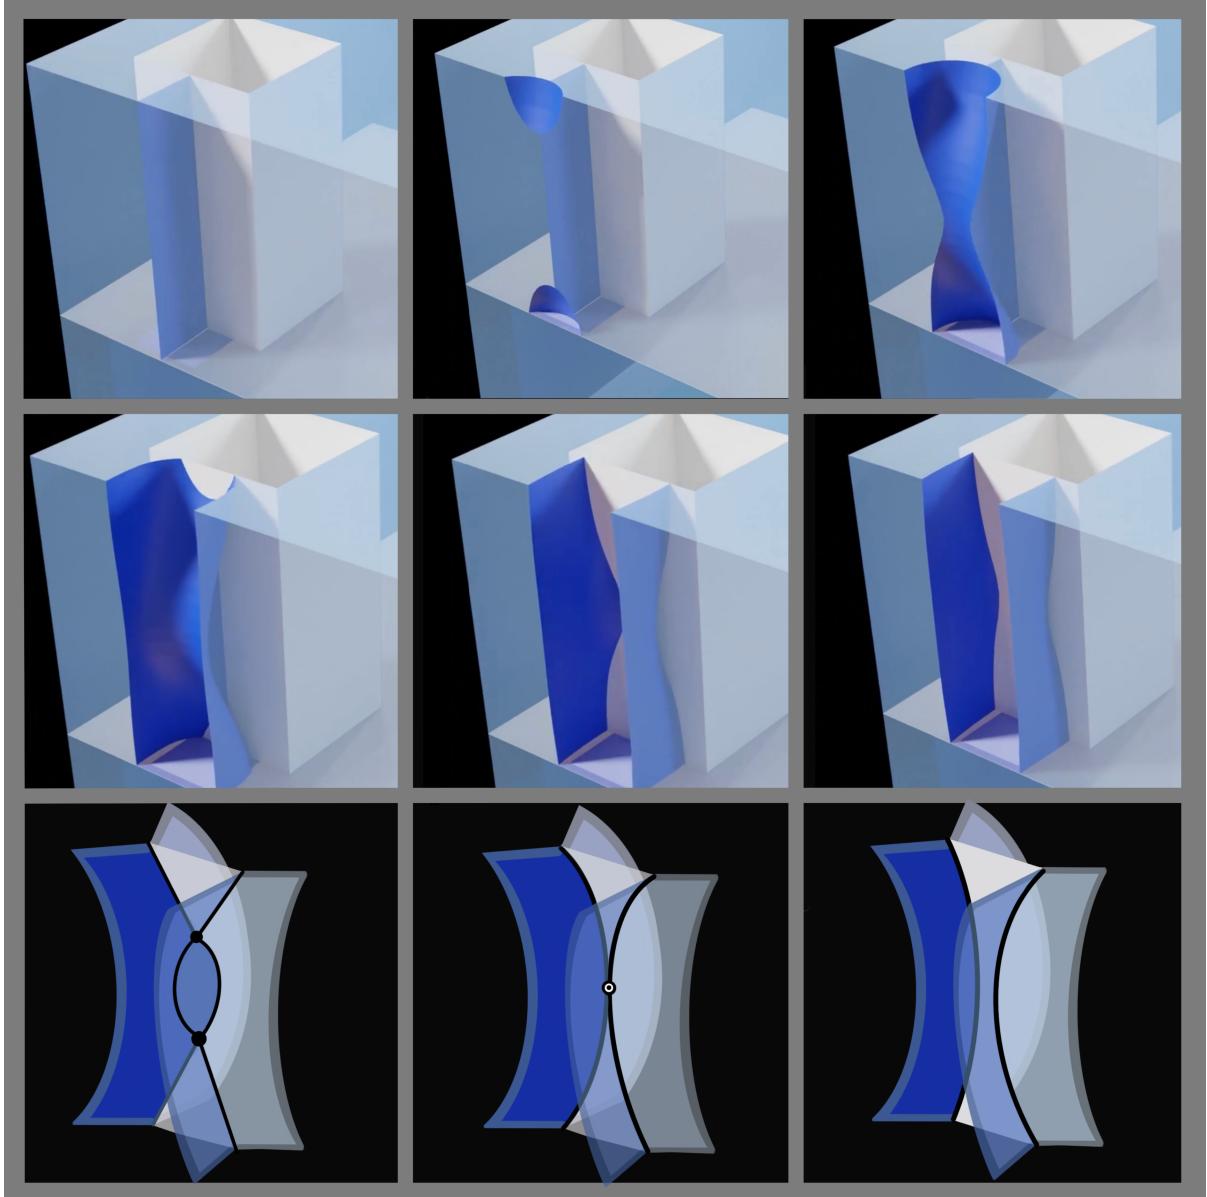


Figure 3.9: This deformation (top six panels), which cuts a flap open, makes Bing's house collapsible. The topology of this transition (in the neighbourhood of the point where the flap disappears) is precisely that of the type A_1^4 . The transition of type A_1^4 is depicted in the bottom three panels; the strata are coloured to indicate the correspondence with the construction directly above.

Proof. Bing constructed his house as a deformation retract from a solid cube, see Figure 3.7. The isotopy of the sphere we consider is the boundary of this deformation. However instead of reducing to a two dimensional object we forego the final step so that at any point in the deformation the set remains a topological (solid) ball and its boundary a sphere. The end point of this deformation is a thickened version of Bing's house.

We will only consider the medial axis on the inside of the sphere and not the exterior. The medial axis of a thickened version of Bing's house is Bing's house itself. The deformation is depicted in Figure 3.7. The essential topological change only happens near the end of the deformation when the rooms wrap around their respective corridors, see Figure 3.11. At this moment the bisector between the corridor and the wall disappears and is replaced by the

bisector between the two parts of the room that are wrapping around the corridor and the medial axis becomes non-collapsible.

The transition as depicted in Figure 3.10 is not generic, it is in fact quite symmetrical and at the critical transition, see Figure 3.11, there is an entire edge in the medial axis for which there are 4 nearest points on the boundary of the thickened Bing's house. For each point on the edge, the closest points are a point on the exterior wall, a point on the corridor and two points on either side that lie on the expanding top room.

We can perturb or change the configuration slightly as follows in order to make it generic. Just before the time where the edge(s) in the medial axis for which there are 4 nearest points on the boundary of the thickened Bing's house is created, we can deform a small part of the thickened Bing's house. These edges come in pairs, one for the top room and one for the bottom. For each of the edges the four closest points on the thickening of Bing's house lie on two sides of the growing room, one near the corridor, and one on the outside wall. This gives rise to four curves on the boundary of the thickened Bing's house. These curves are in fact straight lines as the thickening of Bing's house is locally flat in this neighbourhood. In the middle of these curves on the wall and near the corridor, we deform the thickening of Bing's house so that it becomes locally spherical (with a large radius) in a small neighbourhood of the edge. For the curves on the growing rooms we deform in the same manner, that is, we make it locally spherical, but with a much smaller radius.

For all points x near the centre of the configuration, the points $p \in \pi(x)$ in the closest point projection π of x , lie on a part of a sphere. Here the two walls of the expanding top (or bottom) room become closer, while the distance between the corridor and exterior wall remains roughly the same. By focusing on the spherical parts, which locally control the transition, we see this is equivalent to two small spheres moving closer and two spheres remaining at the same distance, which are exactly a transition between singularities of type A_1^4 . Figure 3.9 offers a global view of what happens to the topology.

Figure 3.9 also helps to explain why the medial axis becomes collapsible. By the perestroika, the entire length of the corridor is exposed to the outside. We can therefore start collapsing/burning from the top of the corridor. We then continue to collapse/burn along parts of the corridor and separator that have been exposed by the splitting of the flap. We then can continue into the walls and floor of the bottom room, removing those. This exposes the bottom of the corridor leading to the top room, and we can collapse/burn this too. Finally we can collapse/burn the top room itself to a point. \square

Because Bing's house cannot be pruned or collapsed we immediately have the following corollary:

Corollary 1. *Collapsing or pruning the medial axis of a domain such that it becomes one-dimensional, as proposed in e.g. [BBML⁺18], is not always possible, even if the boundary of the domain is a smooth sphere.*

3.3 Higher Order Medial Axes

The definition of the medial axis as the centres of balls whose boundary touch a manifold in at least two places, otherwise not intersecting the manifold, lends itself to a straight forward generalization. Namely, we could take the centres of balls whose boundary contact a manifold



Figure 3.10: The evolution of the inside medial axis (yellow) as the smoothed solid cube is deformed into a thickened version of Bing's house (blue/purple). The smoothed solid cubes in these five frames coincide with the unsmoothed solid cubes in the first five frames of Figure 3.7. The exact moment at which the medial axis becomes non-collapsible lies between frames four and five and is depicted explicitly in Figure 3.11.

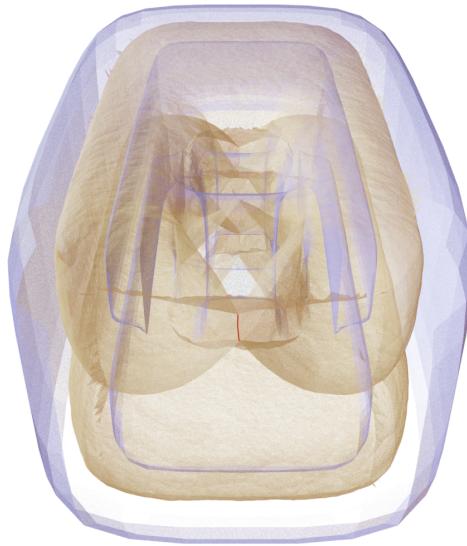


Figure 3.11: The critical transition of the medial axis is highlighted in red. There are points on the medial axis equidistant to the two parts of the room that wrap around the corridor, the corridor itself, and the exterior wall, which can be avoided by a small perturbation. This transition occurs between frames 4 and 5 of Figure 3.7.

in at least two places, while the intersection of the ball with the manifold is the entire manifold. Given the nature of contacts and intersections we could label these axes the insphere and circumsphere axes respectively. The simple case of an ellipse in the plane demonstrates the two possible axes we obtain in two dimensions, see Figure 3.12. When we pair these axes with the radii of the spheres at each point we obtain the insphere and circumsphere axis transforms. This idea was first explored in the master thesis of Elizabeth Stephenson [Ste23a] and the following paper [EST26].

Alternatively to the sphere definition, it is possible to define these medial axes via the Euclidean distance function from a point $x \in \mathbb{R}^d$ restricted to a manifold, $\mathcal{M} \subset \mathbb{R}^d$, $\rho_{\mathbb{E}}(\cdot, x)|_{\mathcal{M}} : \mathbb{R}^d \rightarrow$

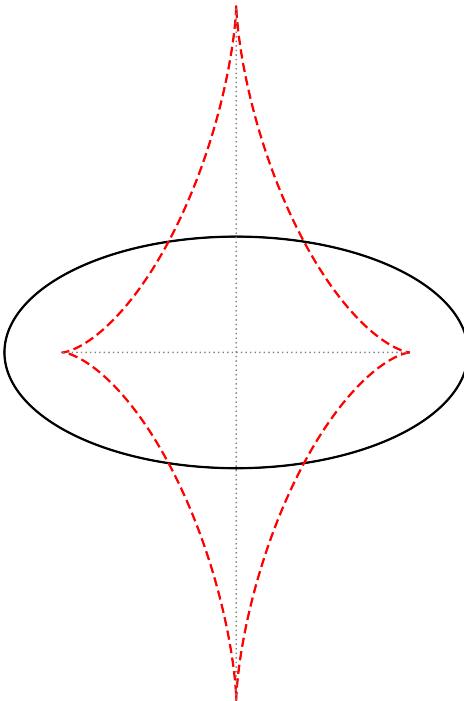


Figure 3.12: An ellipse (in black) and its evolute (in red). The standard (0th) medial axis (or insphere axis) is the gray horizontal line, and the 1st medial axis (or circumsphere axis) is the vertical gray line.

\mathbb{R}^d , where $x \mapsto \|p - x\|$, which we call the restricted distance function or simply the radial distance function. When $d = 2$, the familiar medial axis (insphere axis) is the set points, $x \in \mathbb{R}^2$, for which the first minima of the restricted distance function, with respect to an expanding radius from p , is not unique. The circumsphere axis is then the set of points for which the final maxima of the restricted distance function is not unique. Because the indices of a minima and maxima of functions in \mathbb{R}^2 are 0 and 1 respectively, the insphere and circumsphere axes have also been called the 0-th and 1-st medial axes.

If one relaxes the condition of critical points being the first or last to occur in a radial filtration of \mathcal{M} , and additionally that critical points of the same index must appear at the same radius, the insphere and circumsphere axes are seen to be subsets of the symmetry set of \mathcal{M} . The symmetry set simply being the centres of all spheres which contact \mathcal{M} in at least two points.

In \mathbb{R}^3 index 1 critical points correspond to saddles while maxima have index 2. The order of appearance of saddles with respect to the radial filtration has no clear importance, unlike for minima and maxima. Instead we require additional definitions and constraints. A sphere, S , touches \mathcal{M} at a point $p \in \mathcal{M} \cap S$ if the tangent spaces of \mathcal{M} and S at p are equal. Additionally, S and \mathcal{M} touch generically, when the principal curvatures of \mathcal{M} at p are different from the curvature of S , $1/r$. A pair $(x, r) \in \mathbb{R}^d \times \mathbb{R}$ belongs to the i -th medial axis if:

- (i) The restricted distance function $\rho(\cdot, x)|_{\mathcal{M}}$ has an index i critical value at radius $r \in \mathbb{R}$
- (ii) The rank of the i -th homology of the sublevel set $\rho(\cdot, x)|_{\mathcal{M}}^{-1}([0, r])$ is one
- (iii) The number of points in the level set at r touches at least two distinct points generically.

In \mathbb{R}^3 the 0-th, 1-st and 2-nd medial axes become the insphere, midsphere and circumsphere axes and we label them $M^0[\mathcal{M}], M^1[\mathcal{M}], M^2[\mathcal{M}]$ respectively, see Figure 3.13. It is the midsphere axis we are primarily interested in for this section.

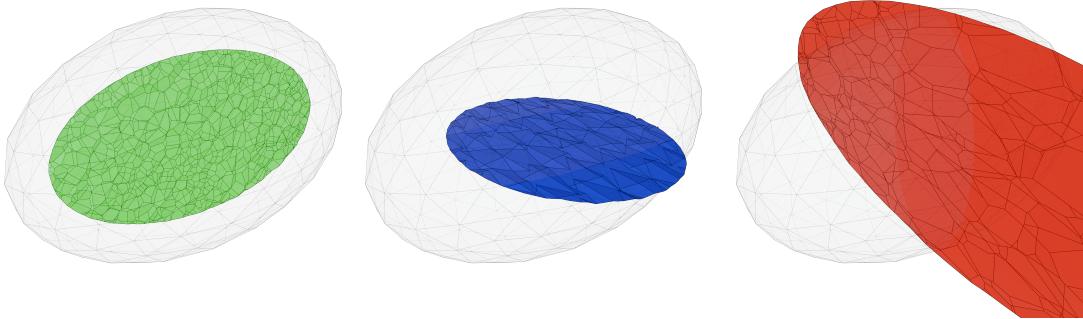


Figure 3.13: From left to right: The insphere axis of a piecewise linear ellipsoid; the midsphere axis; and the circumsphere axis. All axes are numerically computed and hence consist of polygons whose edges are drawn in black.

The focal sets of a manifold, are the sets of spheres whose radii are equal to the radii of principal curvature of points $x \in \mathcal{M}$, such that the derivatives of the spheres and \mathcal{M} at x are equal. The cusps of these focal sets are exactly the critical points of curvature of \mathcal{M} and the limit of branches in the symmetry set (possibly also the various medial axes) as the contact points converge. The midsphere axis terminates in ridges that pass through umbilic points of \mathcal{M} , that is points where the principal curvatures are equal, and hence is bounded partially by each of the focal sets. We call the union of the symmetry set and the focal sets the extended symmetry set.

In \mathbb{R}^2 there is a single focal set, more commonly referred to as the evolute and the symmetry set depicted in Figure 3.12 consists exclusively of medial axis pieces. In \mathbb{R}^3 the critical points of principal curvatures form arcs and loops (referred to as ribs) in the focal sets which correspond to ridge-like features in \mathcal{M} [Por94]. Again in Figure 3.13 the symmetry of the ellipsoid consists exclusively of medial axis pieces. However, in general the symmetry set can be extremely complicated, see [Blu67], potentially obfuscating useful information about a manifold captured by the medial axes.

For the remainder of this section we will explore new and existing algorithms for the computation of the midsphere axis of surfaces in \mathbb{R}^3 .

3.3.1 Vineyard Algorithm

In the original introduction [EST26], the authors presented an algorithm for the computation of the midsphere axis based on so called Faustian interchanges in the vineyard of radial distance function parameterised by \mathbb{R}^3 . A Faustian interchange occurs when two critical points exchange their types, that is, one point that had given rise to a birth in persistence now gives rise to a death and the other which had given rise to a death now gives rise to a birth. If we track the persistence of the radial distance function through a curve that passes through the insphere axis (respectively midsphere or circumsphere axes), the two minima (respectively saddles or maxima) with smallest radii (respectively largest for circumsphere axis) will exchange their types resulting in a Faustian exchange.

The algorithm in [EST26] proceeds by probing for Faustian interchanges in the vineyard of the radial distance function over short intervals in \mathbb{R}^3 . The authors were able to do this efficiently by leveraging the vineyard algorithm introduced in [CSEH05]. When there is a transposition between filtrations of simplices of the same dimension they check if their interchange of types is Faustian. In particular they use a neighbouring vertices in a scaled integer lattice to compute the interval vineyards and include the square that would be the facet between the Voronoi regions of the vertices into the midsphere axis when a Faustian interchange is detected. The lattice is swept through in such a way that the reduced matrices between vertices are recycled, no computations are repeated, and the whole process can be computed in parallel. If n is the number of vertices in the input manifold and g^3 is the number of lattice points, the number of interchanges to be checked is bounded by $O(g^2n^2)$.

Similarly to the λ -medial axis they also employ pruning techniques based on either the distance between interchanging simplices or their persistence. In the former case they implemented Euclidean pruning which, while less accurate than geodesic distance, is simpler to compute. In the latter case the persistence of the features involved with the interchange are thresholded. It was noted in [EST26] that best results were achieved with a combination of both techniques.

3.3.2 Bisector Algorithm

The λ -medial axis, which was discussed and used for computations in Section 3.2, can be seen to demonstrate how the medial axis can be seen as the continuous limit of the Voronoi tessellation. We mean this in the sense that by taking an ε -sample from a manifold, \mathcal{M}_ε and applying a λ -filtering to the vertices of $\text{Vor}(\mathcal{M}_\varepsilon)$ then $\lim_{\varepsilon \rightarrow 0} \rho_H(\text{Vor}_\lambda(\mathcal{M}_\varepsilon), M_\lambda[\mathcal{M}]) \rightarrow 0$. It was this intuition that led us to consider how the midsphere axis of a piecewise linear manifold might exist within the bisector arrangement of its vertices. A schematic of our proposed algorithm is as follows.

Input: a triangulation (piece-wise linear approximation), \mathcal{T} , of a smoothly embedded closed surface, $\mathcal{M} \subset \mathbb{R}^3$. \mathcal{T} consists of vertices, edges, and triangles, with the global connectivity of a 2-manifold (without boundary). We call the vertices of \mathcal{T} points and write \mathcal{X} for the set of points and $n = \text{card } \mathcal{X}$ for their number.

Output: a polyhedral complex, $M_\lambda^1[\mathcal{T}]$, that approximates the midsphere axis of the smooth surface. $M_\lambda^1[\mathcal{T}]$ consists of vertices, edges, and facets (convex polygons), with the global connectivity of a 2-dimensional stratified space. To make it a ‘transform’, we add the piecewise smooth radius function, $R: |M_\lambda^1[\mathcal{T}]| \rightarrow \mathbb{R}$.

Approach: a depth-first search of the bisector arrangement defined by the points in \mathcal{X} , which generalizes the approach of approximating the medial axis using facets of the Voronoi tessellation; see [AC08]. Each point of a bisector is the center of a sphere that passes through the corresponding two points in \mathcal{X} . A crossing between two bisectors is actual if each shared point is the center of the same sphere for both bisectors, and it is spurious otherwise. In other words, the bisectors define an arrangement of spheres; for example a line is the pencil of spheres that all pass through the same three points in \mathcal{X} . Following the description in [Ede86], we locally search this plane arrangement without constructing it, but note that we ignore spurious intersections in this search. The midsphere axis is the subcomplex of visited and then accepted vertices, edges, and facets of this arrangement. For all geometric operations, we assume genericity, which we simulate using simulation of simplicity (SoS) [EM90].

Consider a sphere that passes through a point $A \in \mathcal{X}$, which is a vertex of \mathcal{T} . The link of A consists of the vertices and edges that bound the disk of triangles that share A ; see

Figure 3.16. We say A is a saddle relative to the sphere if the cyclic sequence of the vertices in the link alternate twice between inside and outside the sphere: in-out-in-out and then repeat.

We define the bisector arrangement as the projection of the 2-skeleton of a hyperplane arrangement in \mathbb{R}^4 . The projection introduces extra intersections, which we ignore. While

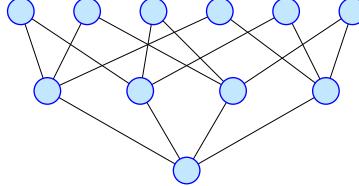


Figure 3.14: The Hasse diagram of the local neighborhood of a vertex, which is contained in four lines and six planes, such that each line is contained in three of the planes, while each plane contains only two of the lines.

this requires that we can distinguish between spurious and actual intersections, we prefer the description in \mathbb{R}^3 as it is easier to imagine. The constituents of the bisector arrangement are:

- vertices, each the center of a sphere that passes through 4 of the n points;
- lines, each consisting of points that are equidistant to 3 of the n points;
- planes, each consisting of points that are equidistant to 2 of the n points;
- edges, each the portion between two consecutive vertices on a line;
- facets, each a connected portion after cutting a plane along its lines.

The local connectivity is important and illustrated in Figure 3.14. In particular, each vertex belongs to six planes and to four lines, and each line belongs to three planes. Note that in an arrangement of planes in general position, a line belongs to only two planes, but in our setting, such a line is a spurious intersection and ignored. If a plane contains a vertex, it also contains exactly two lines that pass through this vertex. So within such a plane, a vertex belongs to only two lines; that is: we have a simple line arrangement, see Figure 3.15.

By definition, any point x in the interior of a facet is equidistant to two unique points A and B . Let A_0, A_1 and B_0, B_1 be pairs of vertices, in the links of A and B respectively, that are inside the sphere and separated from each other along the link by vertices outside the sphere. Hence, Conditions (i) and (ii) of Section 3.3 translate to two necessary conditions for the facet to belong to the midsphere axis:

Saddle Condition A and B are saddles relative to the sphere with center x and radius r that passes through A and B , see Figure 3.16.

Path Condition There are paths on \mathcal{T} inside the sphere that connect A_0 with B_0 and A_1 with B_1 , but there are no such paths connecting A_0 with B_1 and A_1 with B_0 , see the dashed lines in Figure 3.16.

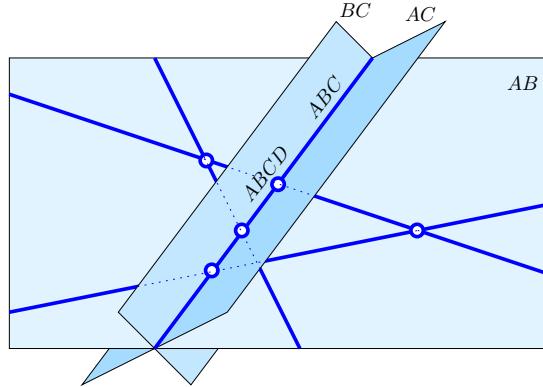


Figure 3.15: The line arrangement inside a plane (the bisector of points A, B). The line equidistant to points A, B, C is in this plane, but also in the planes that are the bisectors of A, C and B, C , respectively. One of the vertices on this line is equidistant to points A, B, C, D .

The first condition implies that the vertices in the link of A alternate twice between inside and outside the sphere. While the second condition ensures that a \mathbf{H}_1 feature is born from this sphere.

We construct the set of facets that satisfy these two conditions, together with all their edges and vertices, as the approximation of the mid-sphere axis. Note that there may be multiple combinations of the neighbours of A and B which lie inside the sphere and satisfy the saddle condition. To get the mid-sphere axis transform, we equip each facet with the function that maps each of its points to the squared radius of the mid-sphere centered at this point, which globally gives a piecewise quadratic function.

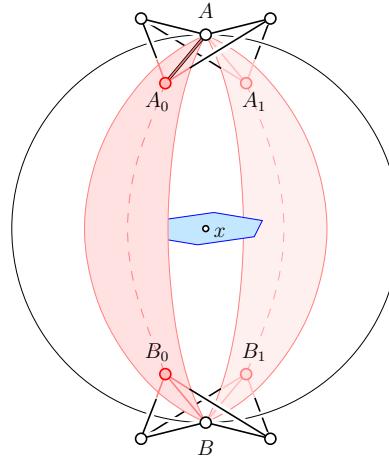


Figure 3.16: The sphere is centered at an interior point of a facet in the bisector arrangement and passes through the two points for which this facet belongs to the bisector. The sphere is a mid-sphere of the input surface if the two points are saddles relative to this sphere, and there are separable patches on \mathcal{M} connecting A to B inside the sphere as shown.

Given a pair of points $A, B \in \mathcal{X}$ and fixed $\lambda \in \mathbb{R}$, $\lambda > 0$, it is possible to decide via a subroutine if there are any facets f satisfying the Saddle and Path Conditions and hence belongs to the mid-sphere axis or not.

Step 1 if $\|A - B\| \geq \lambda$ proceed to Step 2, otherwise return `false`.

Step 2 consider each combination of points in the links of both A and B such that the selected points lying inside of a sphere and unselected points lying outside of the same sphere would give rise to saddles at both A and B .

Step 3 for each such combination take the half spaces bordered by the bisectors of each point in the link of A and B and orient them such that the halfspaces of selected points are oriented towards A (or B) and the halfspaces of unselected points point towards the point itself.

Step 4 if the intersection of these halfspaces along with the bisector plane of A and B themselves is non empty we obtain a facet f for the combination selected in Step 1 and continue to Step 5, otherwise repeat this step with the next combination from Step 2.

Step 5 for each edge e in the boundary of f trace the pencil of spheres defined by e , if any point $C \in \mathcal{X}$ enters or exits the sphere, subdivide f by the edge ABC into subfacets and complete Step 6 with each subfacet in place of f (this represents a branch in the midsphere axis)

Step 6 letting $x \in f$ be the center for a sphere S representing f , take A_0, A_1 and B_0, B_1 as selected points belonging to separate connected components of their respective links inside S . Perform a search in \mathcal{T} inside S starting at A_0 , and another such search starting at A_1 , and return `true` if one search finds B_0 (but not B_1), and the other search finds B_1 (but not B_0).

For each f found to belong to the mid-sphere axis, we may also return the squared radius function, $R^2: |f| \rightarrow \mathbb{R}$, defined by $R^2(y) = \|y - A\|^2$.

Note each bisector AB is subdivided into regions (many of which may be empty) by the different link combinations from Step 2, see Figure 3.17

In order to extend this subroutine into a complete algorithm we require an efficient means to search the space of bisectors and a way to identify a starting pair. To address the former we employ a breadth first search approach. Once we have identified a pair A, B that returns `true` from the subroutine we can continue with the pairs A_i, B and A, B_j , where A_i and B_j are the points in the links of A and B respectively. Meanwhile, in the current instantiation we find a starting pair A_0, B_0 by identifying pairs of ridges, that is, positive (or negative) maximal (minimal) ribs of the minimal (maximal) curvature map, on the surface and selecting one point from each ridge, with minimal pairwise distance. With these additional routines in place we were able to implement a functional, yet far from optimal, algorithm for computing the midsphere axis of a surface.

3.4 Discussion

In Section 3.2 we constructed an isotopy of the sphere such that the inside medial axis transitions from collapsible to non-collapsible.

The frames of the deformation illustrated in Figures 3.7, 3.10, and 3.11 are taken from our Symposium on Computational Geometry Media Exposition submission [CFSW25b] (available at <https://youtu.be/CFmFP6CHVEk>). These animations were made using our own

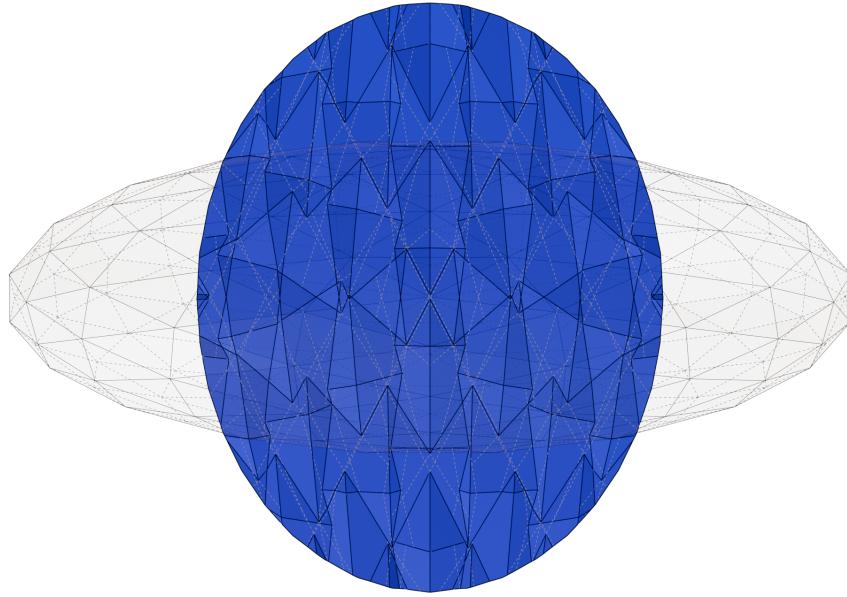


Figure 3.17: A close up of the midsphere axis from Figure 3.13. Solid black lines separate regions where one of the pair A, B changes, whereas light dashed lines separate regions where one of the points in the links of A or B enters or exits the contact sphere.

implementation of the λ -medial axis and the open source-software Blender [Com18]. Here λ is chosen very small to ensure that the λ -medial axis is a good approximation of the medial axis.

The exterior medial axis throughout this isotopy remains collapsible (by which we mean after compactifying the ambient space to \mathbb{S}^3). This automatically leads to some interesting questions.

Does there exist an embedding of \mathbb{S}^2 in \mathbb{S}^3 where it is not possible to burn/partially collapse either connected component of the medial axis of \mathbb{S}^2 at all? We will refer to one of the connected components as the outside and one as the inside, as before, even though this choice is arbitrary (unless \mathbb{S}^3 is explicitly described as the one-point compactification of Euclidean space).

Can we use the burning of either the inside or outside medial axes to guide an isotopy of the sphere? See Figure 3.18 for an example of an isotopy guided by burning.

This leads us to the following stronger conjecture.

Conjecture 1. *Given any embedding of \mathbb{S}^2 in \mathbb{S}^3 , there exists a series of isotopies of \mathbb{S}^2 based on successive burnings of both the inside and outside medial axes, which yields the standard symmetrical embedding of the 2-sphere in \mathbb{S}^3 .*

We point out that some work on the correspondence between the singularities of the inside and outside medial axes has already been done [DG14].

Lieutier and Wintraecken [LW23] recently showed that there is a pruning (the (λ, α) -pruning) of the medial axis that is Hausdorff and Gromov-Hausdorff stable under Hausdorff perturbations of the input set.

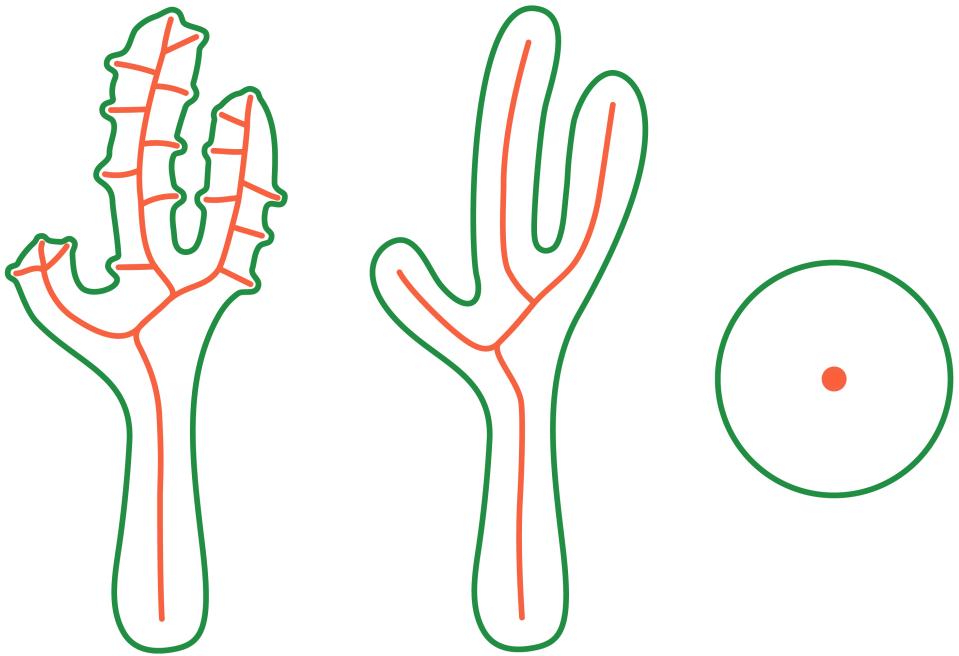


Figure 3.18: A two dimensional example of an isotopy guided by the burning of the medial axis.

Finally, although the algorithm for computing the midsphere was correct and functional, the implementation was not optimised to the extent that triangulations of more than a few thousand vertices were impractical to compute. For example, our implementation must fully recompute the halfspace intersection for each saddle arrangement between a pair of points. In practice most saddle arrangements require flipping the norms of a few halfspaces in the arrangement which could be computed more efficiently. Likewise the Python implementation of the halfspace intersection algorithm only returns the vertices of the intersection without the face relations which is what we really need, necessitating an additional convex hull algorithm. With these optimizations in place the algorithm would already be far more practical for applications.

CHAPTER 4

Delaunay Surfaces and Paths

4.1 Introduction

This chapter contributes credence to the idea that stochastic geometry is an effective tool to measure geometric shapes that are otherwise difficult to measure. As a first example consider a smoothly embedded closed surface, $\mathcal{M} \in \mathbb{R}^3$. Except in rare cases, its area cannot be computed analytically. To compute it numerically, we may sample points on \mathcal{M} and connect them with triangles to get a piecewise linear approximation whose area is easily obtained. Sampling progressively finer while avoiding pitfalls like skinny triangles whose normals do not approach those of the surface, the area of the piecewise linear approximation converges to the area of \mathcal{M} [CMS84].

However, avoiding these pitfalls can be quite difficult in practice and it may not be clear a priori whether the area indeed converges to the area of \mathcal{M} . We refer to [CDS13] for background on mesh generation questions. Alternatively, we may side-step this problem by exploiting an interesting result that the Delaunay surface of \mathcal{M} for a stationary Poisson point process in \mathbb{R}^3 has expected area $3/2$ times the area of \mathcal{M} [EN25]. The Delaunay surface, which we denote, $\text{Del}(\mathcal{M}, A)$, is constructed by collecting the cells in the Delaunay mosaic of a Poisson point process, A , whose dual Voronoi cells have non empty intersection with the surface. The computation is not necessarily faster but it is less demanding and error prone than that of the approximating sequence of piecewise linear surfaces.

The main result of this chapter for Delaunay surfaces, is to extend this result to the integrated mean curvature of \mathcal{M} . Thanks to Hadwiger's theorem [Had57], this result for mean curvature completes the distortion factors for all intrinsic volumes, and hence all rigid motion invariant valuations, in \mathbb{R}^3 . We write $\text{Mean}(\mathcal{M})$ for the (integrated) mean curvature of \mathcal{M} for which we find an expected distortion of $3/2$.

In contrast to the expected area, for which the corresponding equation holds for every intensity [EN25], the expected mean curvature differs from $3/2$ by a small amount, which vanishes as the intensity goes to infinity, see Equation (4.6).

Similar to the area, a distortion factor also holds on average, over all rigid motions of \mathcal{M} , but now only in the limit, so we need an appropriate sequence of point sets. For example, we may formulate the result for $A = r\mathbb{Z}^2$ and take the limit as r goes to 0. This version of the result says that the factor $3/2$ is a property of \mathbb{R}^3 and not of the random point process.

	$d = 1$	2	3
$\phi_0 = \text{Gauss}(\text{Del}(\mathcal{M}, A))$	1	1	1
$\phi_1 = \text{Mean}(\text{Del}(\mathcal{M}, A))$	1	$4/\pi$?
$\phi_2 = \text{Area}(\text{Del}(\mathcal{M}, A))$		1	$3/2$
$\phi_3 = \text{Vol}(\text{Del}(\mathcal{M}, A))$			1

Table 4.1: Summary of known and unknown distortion factors for intrinsic volumes in low dimensional ambient spaces. The well known intrinsic volumes for a surface are listed, however considering instead a curve in the plane (or space), ϕ_1 would represent the distortion factor of length.

The relevant background on Delaunay surfaces; mean curvature and the intrinsic volumes; and Hugger's characterisation theorem are covered in Section 4.2. The main result, Theorem 5, and the necessary lemmas are proved in Section 4.3.

Finally, we also investigate the writhe of Delaunay loops in \mathbb{R}^3 as a first example of non-valuation measures in Section 4.4. Here we work instead with a knot, γ , in \mathbb{R}^3 , $A \subseteq \mathbb{R}^3$ a locally finite set whose convex hull is the entire \mathbb{R}^3 , then $\text{Del}(\gamma, A)$ is the Delaunay loop of γ and A . We are interested in both the average writhe, $\text{Avg}[\text{Wr}(\text{Del}(\gamma, A))]$, over all rigid motions, and the expected writhe, $\text{Exp}[\text{Wr}(\text{Del}(\gamma, A))]$, for the expected writhe over all choices of A . Which leads us to the Theorem 8.

4.2 Preliminaries

4.2.1 Delaunay Surfaces

The Gauss circle problem compares the area of a disk to an approximation of the disk generated by collecting the pixels whose centers lie within the disk as in Figure 4.1 [Har99]. The approximation improves as we decrease the spacing between lattice points. The difference in area for a unit circle with grid spacing α is given by $|\pi - n\alpha^2|$, which vanishes as $\alpha \rightarrow 0$.

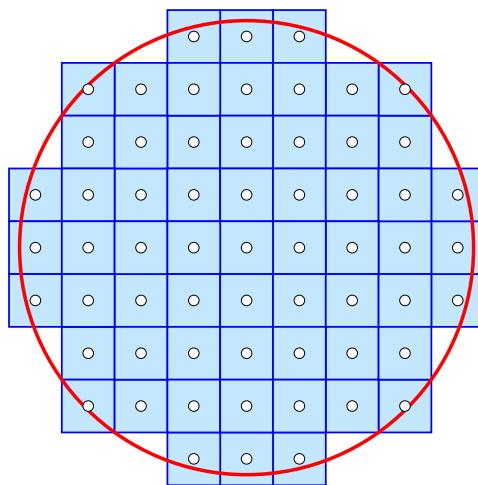


Figure 4.1: We may approximate the disk bounded by the red circle with the pixels whose centers lie within the disk.

In contrast, the perimeter of the union of pixels does not converge to the perimeter of the disk. Note that the boundary consists of four staircases whose combined length converges to

four times the diameter of the unit disk, which as a ratio to the actual perimeter is, $8/2\pi = \frac{4}{\pi}$. On the other hand, replacing the unit disk with the unit square, the approximation converges to the correct perimeter. This difference in behavior is caused by the lack of isotropy of the square grid. This motivates us to view shapes through the lens of an isotropic tiling, and we use randomness to generate it.

Recall the definitions of Voronoi tessellations and Delaunay mosaics from Section 2.4. Given a space $\mathcal{S} \subseteq \mathbb{R}^3$ and a locally finite set $A \subseteq \mathbb{R}^3$, the Voronoi domains of A decompose \mathcal{S} , and we are interested in the dual of this decomposition, which is necessarily a subcomplex of $\text{Del}(A)$. The restricted Delaunay complex of \mathcal{S} and A , denoted $\text{Del}(\mathcal{S}, A)$, consists of all cells in $\text{Del}(A)$ whose dual cells in $\text{Vor}(A)$ have a non-empty intersection with \mathcal{S} . If \mathcal{S} is a closed surface, we write $\mathcal{M} = \mathcal{S}$ and refer to $\text{Del}(\mathcal{M}, A)$ as the (closed) Delaunay surface of \mathcal{M} and A , see Figure 4.3 for several examples.

Unless mentioned otherwise, the set $A \subseteq \mathbb{R}^d$ is a stationary Poisson point process with positive intensity ϱ . It can be characterized by the following two properties: the number of points within a finite collection of pairwise disjoint Borel sets are independent random variables; and the expected number of points within a Borel set is ϱ times its Lebesgue measure, see for example [Kin95]. A Poisson point process with positive intensity in \mathbb{R}^d has with probability 1: infinitely many points; is locally finite; and is in general position. We prefer to generate the Voronoi tessellations with a stationary Poisson point process rather than, for example, an integer lattice, in order to obtain an anisotropic grid.

Given a stationary Poisson point process, $A \subseteq \mathbb{R}^2$, and a line segment, L , Baccelli, Tchoumatchenko and Zuyev consider what they call the Voronoi path of L and A and prove that its expected length is $4/\pi$ times the length of L ; see [BTZ00]. See Figure 4.2 for two example Delaunay paths in the plane. In their work, this path consists of all Delaunay edges whose dual Voronoi edges have a non-empty intersection with L . It is thus a special case of the restricted Delaunay complex as defined in [ES97].



Figure 4.2: Left: the Delaunay path of the circle on the square grid sample. Right: the Delaunay path of the same circle on a less regular sample.

The restricted Delaunay complex of a line segment in midst of a Voronoi tessellation in \mathbb{R}^2 was referred to as a Voronoi path in [dCD18]. The generalisation beyond line segments and two dimensions was subsequently referred to as a Voronoi scape in [EN25]. We feel that the alternative terminology of calling these concepts Delaunay paths, Delaunay scapes, and



Figure 4.3: Top: Three input piecewise linear surfaces, a Schönhardt polyhedron, a torus and a bunny (red). Middle: The Delaunay surfaces generated from an integer lattice. Bottom: The Delaunay surfaces generated from a Poisson–Delaunay mosaic.

in the case of surfaces, Delaunay surfaces is more natural as they consist of cells in the Delaunay mosaic and not the Voronoi tessellation after all. We will continue with our preferred terminology and hope that the reader agrees that this is the less confusing of the choices.

The 2-dimensional result in [BTZ00] was generalised to p -dimensional shapes $\mathcal{S} \subseteq \mathbb{R}^d$ in [EN25], proving that the expected p -dimensional volume of the Delaunay scape of \mathcal{S} is some constant times the p -dimensional volume of \mathcal{S} . Furthermore, [EN25] gives an explicit expression for the constants, $\mathcal{D}_{p,d}$, which are displayed in Table 4.2. In words: the length and area in \mathbb{R}^3 are both distorted by a factor $3/2$. The main result of this chapter states that this is also the case for the mean curvature of a closed surface $\mathcal{M} \in \mathbb{R}^3$, albeit only in the limit.

By definition, a closed surface is a compact 2-manifold (without boundary), so every point $x \in \mathcal{M}$ has a neighbourhood homeomorphic to \mathbb{R}^2 . Since we limit ourselves to closed surfaces in \mathbb{R}^3 , they are also orientable, so there is a unique subset $\mathbb{S} \subseteq \mathbb{R}^3$ for which \mathcal{M} is the boundary. Under some conditions on \mathcal{M} and A , the Delaunay surface has the same properties.

Proposition 1 (Closed Ball Property [ES97]). *We say that \mathcal{M} satisfies the closed ball property*

	$d = 1$	2	3	4	5	6
$p = 1$	1	$4/\pi$	$3/2$	$16/3\pi$	$15/8$	$32/5\pi$
2		1	$3/2$	2	$5/2$	3
3			1	$16/3\pi$	$5/2$	$32/3\pi$
4				1	$15/8$	3
5					1	$32/5\pi$
6						1

Table 4.2: Summary of average and expected distortion factors, $\mathcal{D}_{p,d}$, for p -volumes in d dimensional Euclidean space [EN25].

if the intersection of \mathcal{M} with any i -cell of $\text{Vor}(A)$ is either empty or homeomorphic to an $(i-1)$ -ball. If \mathcal{M} satisfies this condition, then $\text{Del}(\mathcal{M}, A)$ is homeomorphic to \mathcal{M} .

We spend the remainder of this subsection with the consequences of violations of the closed ball property. We distinguish those that occur generically from others that require degenerate position and focus on the former. In many situations, the Voronoi surface is not a 2-manifold. Nonetheless, we are motivated to call it a surface because local reinterpretations of the complex—such as splitting a vertex into two—allow us to see it as a 2-manifold.

Each violation of the closed ball property is an intersection of \mathcal{M} with an i -cell that is not an $(i-1)$ -ball. Figure 4.4 displays violations of the closed ball property for $i = 3$ on the left, $i = 2$ in the middle, and $i = 1$ on the right. Such violations can be avoided by having smaller Voronoi cells or, equivalently, a denser set A . A typical case of pinching along a triangle (right case in Figure 4.4) is a fin, which is a polygon that sticks out of a surface to which it is connected along an edge. Another common violation of the closed ball property is when \mathcal{M} intersects a Voronoi polygon in a topological circle. In this case, $\text{Del}(\mathcal{M}, A)$ contains a spike, which is an edge that is not part of any polygon or 3-cell in $\text{Del}(\mathcal{M}, A)$. Most common are spikes that stick out of a surface, but more complicated scenarios in which the spike connects two surface patches are possible.

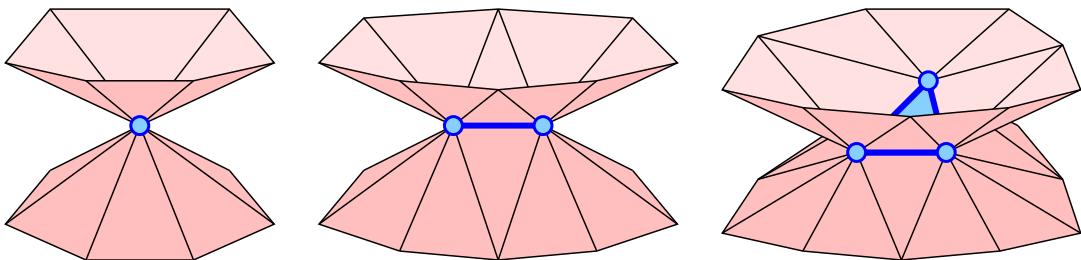


Figure 4.4: From left to right: pinching at a vertex where the surface, \mathcal{M} , intersects a Voronoi 3-cell in two disks; pinching along an edge because \mathcal{M} intersects a Voronoi polygon in two arcs; pinching along a triangle because \mathcal{M} intersects a Voronoi edge in two points.

4.2.2 Intrinsic Volumes

This subsection introduces the intrinsic volumes, in particular the integrated mean curvature, and highlights their importance amongst geometric measures via Hadwiger's characterisation theorem [Had57].

For a convex body, $K \subseteq \mathbb{R}^d$, and a radius, $r \geq 0$, the parallel body is the set of points at distance at most r from K : $K_r = K + rB^d = \{x \in \mathbb{R}^d \mid \min_{y \in K} \|x - y\| \leq r\}$. As shown

by Steiner [Ste13] two centuries ago, the d -dimensional volume of the parallel body follows a polynomial of degree d in r .

Theorem 2 (Steiner Polynomial [Ste13]). *Let K be a compact convex body in \mathbb{R}^d , and write κ_i for the volume of the unit ball in \mathbb{R}^i . Then the volume of the parallel body of K is*

$$\text{Vol}(K_r) = \sum_{i=0}^d r^{d-i} \kappa_{d-i} \phi_i(K), \quad (4.1)$$

in which $\phi_i(K)$ is the i -th intrinsic volume of K .

In \mathbb{R}^3 , the intrinsic volumes are familiar measures: ϕ_0 is the Euler characteristic, ϕ_1 is half the mean curvature, ϕ_2 is twice the surface area, and ϕ_3 is the volume. These measures are called intrinsic because they do not depend on the dimension of the ambient space, by which we mean if $K \subseteq \mathbb{R}^d$ and $\dim K < d$, then $\phi_i(K)$ is the same whether computed in \mathbb{R}^d or the affine hull of K treated as a Euclidean space.

A classic result of Crofton [CS68] expresses the i -th intrinsic volume of a convex body as the average Euler characteristic of its intersection with a $(d-i)$ -flat in \mathbb{R}^d .

Theorem 3 (Crofton's Integral Formula [CS68]). *Let K be a compact convex body in \mathbb{R}^d , and write $\text{Graff}_{d,d-i}$ for the affine Grassmannian of $(d-i)$ -planes in \mathbb{R}^d . Then*

$$\phi_i(K) = \frac{i!(d-i)!}{0!d!} \frac{\kappa_i \kappa_{d-i}}{\kappa_0 \kappa_d} \int_{E \in \text{Graff}_{d,d-i}} \chi(K \cap E) \, dE. \quad (4.2)$$

Here $\text{Graff}_{d,d-i}$ is given a locally finite motion invariant measure, which is normalised such that the measure of flats intersecting the unit ball is κ_i .

A geometric measure, $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$, is called a valuation if it is additive; that is: $\varphi(\emptyset) = 0$ and for all convex sets K and L , we have

$$\varphi(K \cup L) + \varphi(K \cap L) = \varphi(K) + \varphi(L). \quad (4.3)$$

The intrinsic volumes inherit additivity from the Euclidean volumes, so they are valuations. If a valuation is furthermore continuous (under the Hausdorff distance) and invariant under rigid motions, a result of Hadwiger [Had57] says that it is a linear combination of the intrinsic volumes.

Theorem 4 (Hadwiger's Characterisation Theorem [Had57]). *For $K \subseteq \mathbb{R}^d$ a compact convex body, and $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ a continuous valuation that is invariant under rigid motions, we have*

$$\varphi(K) = \sum_{i=0}^d c_i \phi_i(K), \quad (4.4)$$

for some constants c_i independent of K .

Since the intrinsic volumes have these three properties, and these properties are preserved under linear combination, the converse of the implication is also true.

Both Crofton's integral formula and Hadwiger's characterisation theorem extend to smooth finite unions of compact convex sets, see for example, [KRdB97]. The Steiner polynomial and Crofton formula have been extended to compact submanifolds of \mathbb{R}^d , both smooth and piecewise linear, via Weyl's tube formula [Wey39] and normal cycle theory [CMS86], see for example [Mor08]. In these contexts the intrinsic volumes are often labelled i -th Lipschitz-Killing curvatures or i -th mean curvatures of a manifold.

4.2.3 External Dihedral Angles

Consider a piecewise linear closed surface in \mathbb{R}^3 . It is necessarily orientable, so it is well defined which side of the surface is the inside (the space enclosed by the surface), and which is the outside (the space that surrounds the surface and reaches all the way to infinity). For an orientable surface that is not closed, we will talk about the positive and negative sides, defined in such a way that if this surface is a patch of a closed surface, then the inside is on the positive and the outside on the negative side. For a polygon of this surface, the outer unit normal is the unit length vector normal to the plane that contains the polygon and points to the negative side. Let p_1 and p_2 be two polygons that share an edge, e , of the surface. We define the angles at e in terms of a small circle centred at an interior point of e and contained in a plane normal to e . We require that this circle is sufficiently small so that it intersects p_1 and p_2 in a point each.

The dihedral angle between p_1 and p_2 at e is 2π times the fraction of the mentioned small circle on the positive side of the two polygons. Letting δ be this angle, the external dihedral angle at e is $\eta = \pi - \delta$.

By construction, the dihedral angle satisfies $0 \leq \delta \leq 2\pi$. We call e a convex, straight, or reflex edge if $\delta < \pi$, $\delta = \pi$, or $\delta > \pi$, respectively. It follows that the external dihedral angle satisfies $-\pi \leq \eta \leq \pi$: it is positive when e is convex, zero when e is straight, and negative when e is reflex. Indeed, η is the (signed) angle between the outer unit normals of p_1 and p_2 .

For a piecewise linear surface, \mathcal{M} , with 1-skeleton, $M^{(1)}$, the integrated mean curvature has a well known discretisation which depends on the external dihedral angle η ,

$$\text{Mean}(\mathcal{M}) = \sum_{e \in M^{(1)}} \eta_e \|e\|. \quad (4.5)$$

When computing mean curvature of a Delaunay surface it is important to consider the violations to being a 2-manifold that one may encounter. We will see mean curvature can still be computed in these cases and therefore we need not assume the Delaunay surface is indeed a 2-manifold when completing the proof of Theorem 5.

When pinching occurs along a Delaunay edge, e , it either pinches the outside or the inside; see Figure 4.5. In both cases, we compute the external dihedral angle at e as the sum of the external dihedral angles at the edges that are glued to each other to form e . These edges are only imagined, but we can recover the necessary information by inspecting how e^* intersects the surface.

Consider the case in which the intersection consists of two arcs. This corresponds to two glued edges, so e is a face of four polygons. Let $\delta_1 + \delta_2 + \delta_3 + \delta_4 = 2\pi$ be the four dihedral angles between consecutive polygons around e , and assume that δ_1 and δ_2 are angles on the positive side, while δ_3 and δ_4 are angles on the negative side. If e pinches the outside, as in Figure 4.5 on the left, then the external angle at e is computed as $\eta = (\pi - \delta_1) + (\pi - \delta_2) = \delta_3 + \delta_4$, which is necessarily non-negative. If e pinches the inside, as in Figure 4.5 in the middle, then the external dihedral angle at e is computed as $\eta = (\delta_3 - \pi) + (\delta_4 - \pi) = -\delta_1 - \delta_2$, which is necessarily non-positive.

The above case analysis includes the case in which e is an edge of a polygon obtained by pinching, such as the edges of a fin. It can be expanded to the case in which three or more edges are glued. As illustrated in Figure 4.5 on the right, such cases include the possibility that the inside and outside are pinched simultaneously.

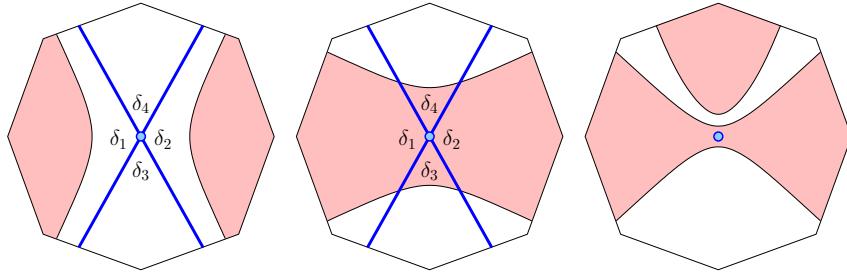


Figure 4.5: The intersection of a Voronoi polygon with a closed surface. The shading indicates the inside of the surface. In all three cases, there is pinching along the corresponding Delaunay edge. Left: pinching of the outside. Middle: pinching of the inside. Right: simultaneous pinching of the outside and the inside.

A spike results from the local strangulation of the closed surface, and similar to the case of pinching, we distinguish between strangulating the inside or the outside. The two cases are again differentiated by inspecting how the surface intersects the corresponding Voronoi polygon, namely in a topological circle that encloses either the inside, in which case the external dihedral angle is 2π , or the outside, in which case the external dihedral angle is -2π .

4.3 Distortion of Mean Curvature

We are now prepared to formally state and prove the main theorem for Delaunay surfaces.

Theorem 5. *Let $A \subseteq \mathbb{R}^3$ be a stationary Poisson point process with intensity $\varrho > 0$, and \mathcal{M} a rectifiable closed surface in \mathbb{R}^3 . Then*

$$\lim_{\varrho \rightarrow \infty} \text{Exp}[\text{Mean}(\text{Del}(\mathcal{M}, A))] = \frac{3}{2} \text{Mean}(\mathcal{M}). \quad (4.6)$$

The proof of the Theorem 5 will be presented in four steps. We prove the average version first and get the probabilistic version from the symmetry of stationary Poisson point processes. Here is an overview.

1. We begin with three geometric lemmas about the average external dihedral angle between the polygons sharing an edge in the Delaunay surface.
2. We leverage the work of Cheeger, Müller and Schrader [CMS84] to obtain a piecewise linear surface that approximates the integrated mean curvature of \mathcal{M} to within some error $\varepsilon > 0$ and exactly in the limit.
3. We recall the work of Edelsbrunner and Nikitenko [EN25], which tells us that the Delaunay path of the 1-skeleton of the piecewise linear surface on average stretches length by a factor of $3/2$.
4. We argue that the length and external dihedral angle are independent so that the distortion of the integrated mean curvature can be obtained by multiplication.

4.3.1 Geometric Lemmas

We begin with the aforementioned geometric lemmas about average external dihedral angles, which are core to our main theorem.

Let $A \subseteq \mathbb{R}^3$ be locally finite and $E \in \text{Graff}_{3,2}$ a plane. Assume for simplicity that E does not pass through any vertex of $\text{Vor}(A)$, in which case $\text{Del}(E, A)$ is 2-dimensional. By convexity, E satisfies the closed ball property, so Proposition 1 implies that $\text{Del}(E, A)$ is homeomorphic to E . Indeed, the orthogonal projection to E is a homeomorphism. After orienting E , we call one half-space the positive and the other the negative side of E . Similarly, $\text{Del}(E, A)$ separates \mathbb{R}^3 itself into positive and a negative sides. Let E' be the same plane but with opposite orientation, so that $\text{Del}(E', A) = \text{Del}(E, A)$ but with positive and negative sides exchanged. Hence, the sum and therefore the average of the two external dihedral angles at any one edge is zero.

For the proof of the main theorem, we need the same conclusion for the more general case of a polygon rather than a plane.

Lemma 1 (Complementary Dihedral Angles). *Let $A \subseteq \mathbb{R}^3$ be locally finite, F a (flat) oriented polygon in \mathbb{R}^3 , and e an edge shared by polygons p and p_0 in $\text{Del}(A)$. Then the average, over all rigid motions, $\mu: \mathbb{R}^3 \rightarrow \mathbb{R}^3$, for which e, p, p_0 belong to $\text{Del}(\mu(F), A)$, of the external dihedral angle on the positive side of p, p_0 vanishes.*

Proof. The edge e and polygons p, p_0 belong to $\text{Del}(F, A)$ iff the dual Voronoi edges, p^* and p_0^* , intersect F in a point each. The orientation of F induces consistent orientations of p and p_0 , so the dihedral angle on the positive side of the two polygons is well defined.

Next we construct a rigid motion for which the dihedral angle between p and p_0 is complementary. Let x and x_0 be the points at which p^* and p_0^* intersect F , and call the line in the plane of F whose points are equidistant to x and x_0 the bisector of the two points. We write μ for the rotation by 180° around the bisector of x and x_0 . By construction, p^* and p_0^* intersect $\mu(F)$, so e, p, p_0 also belong to $\text{Del}(\mu(F), A)$. However, the orientation is now reversed, which implies that the dihedral angle on the positive side of p, p_0 is complementary to what it was before. Since the two dihedral angles are complementary, the sum of the corresponding two external dihedral angles vanishes. \square

The next lemma considers two consistently oriented half-planes, E_1 and E_2 , glued along their common boundary, which is a line, L . We call this a wedge and let α be the dihedral angle on its positive side. We call $\text{Del}(E_1 \cup E_2, A)$ a Delaunay wedge and $\text{Del}(L, A)$ a Delaunay line, noting that the latter is necessarily a subcomplex of the former. Generically, L does not intersect any edge of $\text{Vor}(A)$, in which case L satisfies the closed ball property, by convexity, and $\text{Del}(L, A)$ is homeomorphic to L . In contrast, the Delaunay wedge may or may not be homeomorphic to $E_1 \cup E_2$ because pinching may occur even in the generic case. Nevertheless, the dihedral angle at any edge $e \in \text{Del}(L, A)$ is well defined albeit possibly zero. The lemma asserts that on the average this dihedral angle is α . To formalise this claim, we consider all rotations of the wedge about L . Writing $\mu: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ for such a rotation, we have $\mu(L) = L$ and $\mu(E_1 \cup E_2)$ is still a wedge with dihedral angle α .

Lemma 2 (Average External Dihedral Angles, A). *Let $A \subseteq \mathbb{R}^3$ be locally finite, e an edge in $\text{Del}(A)$, and $E_1 \cup E_2$ a wedge with dihedral angle α such that $L = E_1 \cap E_2$ intersects e^* in a single point in the interior of the Voronoi polygon. Then the average external dihedral angle of the Delaunay wedge at e , over all rotations of the wedge about its line, is $\pi - \alpha$.*

Proof. To compute the average external dihedral angle of the Delaunay wedges at a fixed edge, e , we let p_0, p_1, \dots, p_{k-1} be the cyclic order of Delaunay polygons that share e . Correspondingly,

$p_0^*, p_1^*, \dots, p_{k-1}^*$ are the Voronoi edges in the boundary of e^* . Since L crosses e^* at an interior point, the portion of the boundary on the positive side of the wedge is a connected arc whose endpoints are where E_1 and E_2 cross the boundary of e^* . Suppose E_1 crosses p_0^* in its interior, and E_2 crosses p_j^* in its interior. Writing δ_i for the dihedral angle between p_i and p_{i+1} (taking indices modulo k), the dihedral angle of $\text{Del}(E_1 \cup E_2, A)$ at e is $\delta_0 + \delta_1 + \dots + \delta_{j-1}$.

The average is the integral of this sum divided by 2π . The vertex shared by p_i^* and p_{i+1}^* contributes to the integral while it is on the positive side of the wedge, which it is during an angle α of the full rotation. It follows that the average dihedral angle is $\frac{1}{2\pi} \sum_{i=0}^{k-1} \alpha \delta_i = \alpha$. The average external dihedral angle is therefore $\pi - \alpha$, as claimed. \square

Next, we generalise to a wedge formed by two polygons, F_1 and F_2 , that share an edge, $S = F_1 \cap F_2$. In contrast to the half-planes case, such a wedge does not necessarily intersect the boundary of a Voronoi polygon, e^* , even if S intersects e^* in an interior point. This happens also for small e^* , namely when it intersects S near one of its endpoints. To handle these cases, we formulate the following lemma in terms of a limit process.

Lemma 3 (Average External Dihedral Angles, B). *Let $A \subseteq \mathbb{R}^3$ be a stationary Poisson point process with intensity $\varrho > 0$, e an edge in $\text{Del}(A)$, and $F_1 \cup F_2$ a wedge of polygons with dihedral angle α , such that the line segment $S = F_1 \cap F_2$ intersects e^* at an interior point. Then the limit, for $\varrho \rightarrow \infty$, of the expected external dihedral angle of $\text{Del}(F_1 \cup F_2, A)$ at e , is $\pi - \alpha$.*

Proof. As the intensity of the Poisson process tends to infinity, the sizes of Voronoi polyhedra and polygons tend to zero. Because of this, the contribution of edges of $\text{Del}(S, A)$ whose corresponding Voronoi polygons have boundaries that do not intersect F_1 and F_2 in a point each decreases to zero in the limit. \square

Similar to the main theorem, Lemma 3 can also be formulated for the average instead of the expected external dihedral angle. For this we would need an appropriate sequence of locally finite point sets, such as $r\mathbb{Z}^2$, with r tending to zero.

We provide a computational experiment to give further evidence that a limit statement is indeed necessary for Lemma 3. Figure 4.6 shows the average dihedral angle at the edges of the Delaunay wedge computed for a right-angled wedge of two polygons and, for comparison, a right-angled wedge of two half-planes. Each point is the dihedral angle, averaged (without weights) over the edges in $\text{Del}(S, A)$, in which edges with fewer than two incident Delaunay polygons do not contribute to this average. Note the averaging is taken over a sample of 400 rotations about the wedge axis, causing the average to be slightly off from the true average, which in our setting is $\pi/2$. Beyond averaging over the rotations, we average over several samples of a stationary Poisson point process with progressively larger intensity going from left to right. The difference between the curves for the half-planes and the polygons is proportional to the fraction of Voronoi polygons whose boundaries do not intersect both polygons of the wedge, and this difference is seen to decrease as the intensity of the Poisson point process increases.

4.3.2 Piecewise Linear Approximation and Length Distortion

As stated in Section 4.3, we require a result of Cheeger, Müller and Schrader [CMS84] which defines discrete Lipschitz-Killing curvatures (a generalisation of intrinsic volumes to smooth

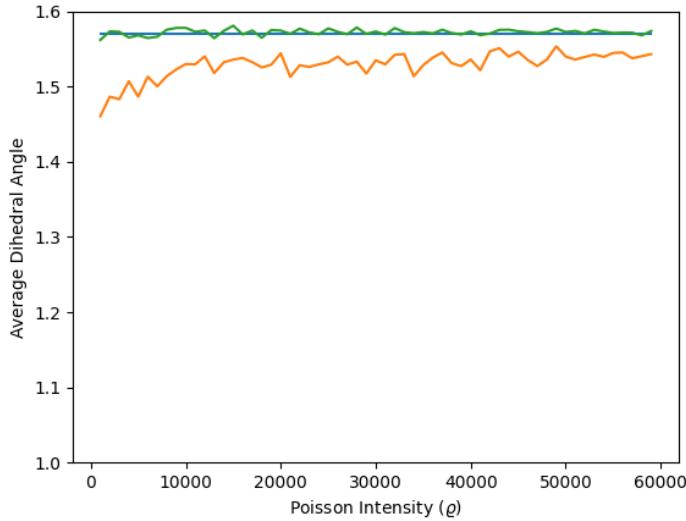


Figure 4.6: The expected angle of two Delaunay wedges, first with input a wedge of two triangles that share an edge (orange) and secondly an input wedge of two large rectangles representing a wedge of half planes (green), both with dihedral angle $\pi/2$ (blue).

closed Riemannian manifolds) for piecewise linear spaces and proves that if a piecewise linear space suitably approximates a smooth space, then the corresponding curvatures are close. We state the result as applicable to our situation.

Theorem 6 (PL Approximation of Surface [CMS84]). *Let $\mathcal{M} \subseteq \mathbb{R}^3$ be a closed surface and let $\{M_\varepsilon\}$ be a family of piecewise linear approximations of \mathcal{M} , for which the supremum of lengths tends to zero and the infimum of angles is bounded away from zero. Let $N_{\mathcal{M}}(x)$ and $N_{M_\varepsilon}(y)$ be the outer unit normals at $x \in \mathcal{M}$ and $y \in M_\varepsilon$, respectively. For each surface in the family, there exists a homeomorphism, $\vartheta_\varepsilon: \mathcal{M} \rightarrow M_\varepsilon$, such that $\|x - \vartheta_\varepsilon(x)\| < \varepsilon$, and for all $x \in \mathcal{M}$ with $\vartheta_\varepsilon(x)$ in the interior of a polygon of M_ε , we have $\|N_{\mathcal{M}}(x) - N_{M_\varepsilon}(\vartheta_\varepsilon(x))\| < \varepsilon$, and $\|\text{Mean}(\mathcal{M}) - \text{Mean}(M_\varepsilon)\| < \varepsilon$.*

Furthermore, we require a result from Edelsbrunner and Nikitenko [EN25], where they compute the distortion factors for volumes of i -dimensional sets embedded in a Euclidean space of dimension $d \geq i$. However, we settle here for lengths of 1-dimensional sets in \mathbb{R}^3 .

Theorem 7 (Edge Expansion [EN25]). *Let $\mathcal{S} \subseteq \mathbb{R}^3$ be a rectifiable 1-dimensional set, and $A \subseteq \mathbb{R}^3$ a locally finite set such that the convex hull of A is \mathbb{R}^3 . Then, the average length of $\text{Del}(\mathcal{S}, A)$, over all rigid motions of \mathcal{S} , is $3/2$ times the length of \mathcal{S} .*

4.3.3 Completing the proof

We have collected all the elements required to complete the proof of the main theorem of this chapter. We begin with the proof for the average, over all rigid motions of \mathcal{M} , and afterwards return to the expectations for a sequence of Poisson point processes.

Beginning with the argument for an average, let $\mathcal{M} \subseteq \mathbb{R}^3$ be a closed surface. Take a piecewise linear approximation, M_ε , which by Theorem 6 sufficiently approximates the integrated mean curvature of \mathcal{M} , and write $M_\varepsilon^{(1)}$ for its 1-skeleton.

Next, compute the restricted Delaunay complexes of M_ε and $M_\varepsilon^{(1)}$, and note that $\text{Del}(M_\varepsilon^{(1)}, A) \subseteq \text{Del}(M_\varepsilon, A)$. Theorem 7 says that the 1-skeleton of $\text{Del}(M_\varepsilon^{(1)}, A)$ is on average $3/2$ times longer than $M_\varepsilon^{(1)}$. The 1-skeleton of $\text{Del}(M_\varepsilon, A)$ is generally even longer as it may contain additional edges.

We also consider the average external dihedral angles at the edge of $\text{Del}(M_\varepsilon, A)$. These split into two cases: angles about the edges in $\text{Del}(M_\varepsilon^{(1)}, A)$, which are, on average and in the limit, equal to the corresponding angles in M_ε as shown in Lemma 3, and edges in $\text{Del}(M_\varepsilon, A) \setminus \text{Del}(M_\varepsilon^{(1)}, A)$, which on average are zero by Lemma 1.

We note the independence of length and angle in the Delaunay surface. Recall that the integrated mean curvature of a PL surface, is the sum, over all edges, of its length times the external dihedral angle at the edge. The average integrated mean curvature, over all rigid motions, is thus a weighted sum of such average products. Each Delaunay edge, e , contributes a number of such average products, namely one for each edge and polygon of M_ε . Since the length of e is constant, each such product is this length times the average external dihedral angle. As argued above, in the limit the latter is 0 in case of a polygon, and the external dihedral angle, in case of an edge of M_ε . We conclude that the average integrated mean curvature is distorted, in the limit, by the same factor as the length, namely $3/2$.

Finally, to argue for the expectation we take a stationary Poisson point process in \mathbb{R}^3 which is invariant under rigid motions. Hence the average integrated mean curvature is also the expected integrated mean curvature. Hence we get $3/2$ as the expected distortion factor in the limit, when the intensity of the process goes to infinity.

4.4 Distortion of Writhe

The writhing number, or simply the writhe, of a space curve is a geometric measure that quantifies the coiling of a cord as it is twisted. Writhe is of interest to us in this section as it is an example of a geometric property that is not an additive valuation and hence not representable by Hadwiger's theorem and outside of our analysis so far in this chapter.

Many natural phenomena exhibit coiling behaviour, including supercoiling of DNA helices, in which chemical and biological properties are characterised by the writhe [KL00]; proteins that coil to fold DNA into compact forms in nuclei [Lev83]; and twisted magnetic fields where the writhe can determine their stability [MR92, PL20, YB03]. Several mathematical models of coiled ribbons have therefore been devised.

Consider a ribbon, (γ, U) , as the annulus formed from an ε offset of an oriented closed curve, $\gamma \subseteq \mathbb{R}^3$, along normals $U_t \subseteq \mathbb{S}^2$ at the points γ_t of the curve. We write $\text{Wr}(\gamma)$ for the writhe of the closed curve, let $\text{Lk}(\gamma)$ be the linking number between the two boundary curves of the ribbon, which is also the integer number of rotations of one of these curves about the other and a topological invariant, and let $\text{Tw}(\gamma)$ be the total number of rotations the ribbon makes around its core axis¹. In this setting, Călugăreanu [Că61], White [Whi69], and Fuller [Ful71] independently proved that the writhe is characterised as the difference between the number of rotations and the rate of rotation of the ribbon about γ :

$$\text{Wr}(\gamma) = \text{Lk}(\gamma) - \text{Tw}(\gamma). \quad (4.7)$$

¹Said otherwise, twist is the angular rate of rotation of γ 's local frame (tangent, normal, binormal) about its tangent, integrated over γ .

Recall the definition of knot diagrams from Section 2.2.1. Given an oriented closed curve $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ each direction $u \in \mathbb{S}^2$ provides a knot diagram of γ with some number of crossings. We can apply a count, ± 1 , to each crossing in the diagram depending on whether a rotation of less than 180 degrees clockwise or anti-clockwise provides the same orientation as the underpass, see Figure 4.7. Calling the sum of these integers the writhe of the knot diagram, denoted $\text{Wr}_u(\gamma)$. If we average the writhe over all possible diagrams we obtain an expression for the writhe independent of the twisting and linking numbers:

$$\text{Wr}(\gamma) = \frac{1}{4\pi} \int_{u \in \mathbb{S}^2} \text{Wr}_u(\gamma) du. \quad (4.8)$$

Observe that reversing the orientation of the knot preserves the signs at the crossings and therefore the total writhe.



Figure 4.7: Sign rule for writhe of a knot diagram: a crossing is assigned $+1$ if the under-crossing strand passes the over-crossing strand from right to left, and -1 if it passes the over-crossing strand from left to right. Viewed from the opposite direction, under- and over-passes switch and so do left and right, so the assigned sign stays the same.

Assume the knot $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ is smooth and the derivative is everywhere non-zero. Write $T_t = \dot{\gamma}(t)/\|\dot{\gamma}(t)\|$ for the unit tangent vector at $\gamma(t)$. These vectors define a smooth but possibly self-intersecting curve on the unit sphere, which we call the tangent indicatrix or simply the tantrix of γ , denoted $T = T(\gamma)$. Together, T and $-T$ partition \mathbb{S}^2 into regions such that the writhe of the knot diagrams along directions in the same region is constant. Let $R \subseteq \mathbb{S}^2$ be such a region, write $\text{Area}(R)$ for its area and $\text{Wr}_R(\gamma)$ for the writhe of the knot diagram along a direction $u \in R$.

Proposition 2 (Tantrix Formula for Writhe). *Let $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ be smooth with everywhere non-zero derivative. Then*

$$\text{Wr}(\gamma) = \frac{1}{4\pi} \sum_R \text{Wr}_R(\gamma) \cdot \text{Area}(R), \quad (4.9)$$

in which the sum is over all regions in the decomposition of \mathbb{S}^2 by T and $-T$.

The writhe of a knot can also be computed from its Gauss map, $g : (\mathbb{S}^1 \times \mathbb{S}^1) \setminus \Delta \rightarrow \mathbb{S}^2$, which sends points $s \neq t$ of the knot to the point $(\gamma(t) - \gamma(s))/\|\gamma(t) - \gamma(s)\|$ on the unit sphere. We note that $(\mathbb{S}^1 \times \mathbb{S}^1) \setminus \Delta$ is the torus with a closed curve removed. Hence, g maps an open surface onto the 2-dimensional sphere. Taking the limit, as we approach the boundary of the open surface, we get the tantrix of γ as the image of the removed curve. By distinguishing whether the map preserves or reverses the orientation of the surface locally at $u \in \mathbb{S}^2$, we get the writhe of the knot diagram in direction u as the signed sum of the points in $g^{-1}(u)$. The writhe of the knot is then the normalised signed area of the map.

Proposition 3 (Double Integral Formula for Writhe [KL00, VALFK79]). *Let $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^3$ be a smooth knot with everywhere non-zero derivative. Then*

$$\text{Wr}(\gamma) = \frac{1}{4\pi} \int_{\mathbb{S}^1} \int_{\mathbb{S}^1} \langle T_t \times T_s, \frac{\gamma(t) - \gamma(s)}{\|\gamma(t) - \gamma(s)\|^3} \rangle dt ds. \quad (4.10)$$

Trace [Tra83] observed that two knot diagrams have the same writhe and Gauss map index iff they are equivalent via only type-II and type-III Reidemeister moves. The condition on the Gauss map can be dropped if we allow equally many type-I Reidemeister moves that increase and decrease the writhe of the knot diagram.

The above formulas have discrete versions for polygonal approximations of a knot. To get an inscribed polygon, $P < \gamma$, we select values $0 \leq t_0 < t_1 < \dots < t_{n-1} < 2\pi$, and let $\gamma_i = \gamma(t_i)$ be the vertices of P . The edges are the line segments between consecutive vertices, with indices read modulo n . To get the tantrix of P , we take the unit vectors $T_i = (\gamma_{i+1} - \gamma_i)/\|\gamma_{i+1} - \gamma_i\|$, which define a spherical polygon in \mathbb{S}^2 , denoted $T = T(P)$. As before, T and $-T$ partition \mathbb{S}^2 into regions within which the writhe of the knot diagrams is constant. We can therefore compute $\text{Wr}(P)$ with the formula given in Proposition 2.

To discretise the double integral formula, let e_i and e_j be two non-consecutive edges of P , and consider the tetrahedron whose vertices are the endpoints a, b of e_i and u, v of e_j ; see Figure 4.8 on the left. Direct the remaining four edges of the tetrahedron from the endpoints of e_j to those of e_i . Furthermore, scale the vectors $a - u, a - v, b - u, b - v$ to unit length to get four points on the unit sphere, and write $A_{i,j}$ for the signed area of the thus defined spherical quadrangle; see Figure 4.8 on the right. Setting $A_{i,i} = 0$ for $0 \leq i < n$, we get

$$\text{Wr}(P) = \frac{1}{4\pi} \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} A_{i,j}, \quad (4.11)$$

as the discrete version of Proposition 3. A fast sweep-line algorithm for polygonal curves, not reliant on the double integral representation, was given in [AEW04].

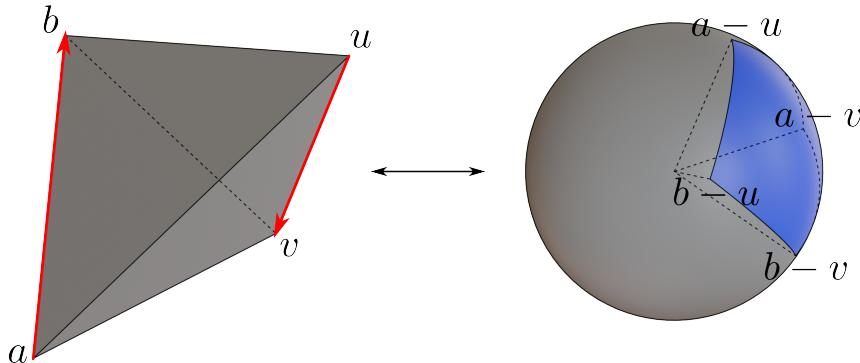


Figure 4.8: The spherical quadrangle defined by two non-consecutive edges of a polygonal knot. The sign of its area depends on the orientation of the two edges along the knot and is defined by the sign rule illustrated in Figure 4.7.

Finally, for many applications, particularly in cell biology where DNA molecules are usually represented as open curves, it would be useful to have an expression for the writhe of an open curve. Some authors attempt to solve this problem by simply connecting the end points of the curve with a straight line segment [Ful71, VM97]. Others have instead connected the endpoints of the tantrix curve with shortest spherical arcs [Sta05]. While [BP06] has defined a method for truly open curves that are anchored at parallel planes or spheres.

4.4.1 Resolving a Knot

We now explain how to turn a knot into a loop consisting of edges in the Delaunay mosaic of a set A . We use the barycentric subdivision of the Delaunay mosaic as a technical tool

to prove the equivalence of the knot and this loop, provided A is sufficiently dense. After explaining how this subdivision can be mapped to a subdivision of the Voronoi tessellation, we formalise what we mean by a sufficiently dense set.

Let $A \subseteq \mathbb{R}^3$ be locally finite and $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$ a knot. We say γ meets the Voronoi tessellation of A generically if γ does not intersect any edge of $\text{Vor}(A)$, and it intersects any 2-cell in at most finitely many points, and at each such point γ crosses the 2-cell transversally.

We may sequentially collect the dual Delaunay edges and vertices, which intersect Voronoi 2 and 3-cells of $\text{Vor}(A)$ respectively, while moving continuously along γ . Hence, every edge is directed, from the previous to the next vertex. As we saw in the case of Delaunay surfaces, it is also possible that an edge is collected multiple times, namely once for each point in which γ crosses the dual Voronoi 2-cell. The resulting multi-set of edges and vertices is an abstract polygonal knot, which is realised in \mathbb{R}^3 by gluing all copies of the same edge or vertex to each other. This realisation is the Delaunay loop, $\text{Del}(\gamma, A)$.

The barycentric subdivision of the Delaunay mosaic of a locally finite set $A \subseteq \mathbb{R}^3$ is the simplicial complex, denoted $\text{Sd Del}(A)$, constructed as follows:

- for each cell, $\sigma \in \text{Del}(A)$, add the average of the vertices of σ , which is a point, $\hat{\sigma}$, in the interior of σ , as a vertex to $\text{Sd Del}(A)$;
- whenever $\sigma_0, \sigma_1, \sigma_2, \sigma_3$ is a sequence of cells such that σ_i is i -dimensional and σ_{i-1} is a face of σ_i , for $1 \leq i \leq 3$, add the tetrahedron with vertices $\hat{\sigma}_0, \hat{\sigma}_1, \hat{\sigma}_2, \hat{\sigma}_3$, as well as its six edges and four triangles to $\text{Sd Del}(A)$.

The simplices of $\text{Sd Del}(A)$ can be grouped to subdivide the cells in $\text{Del}(A)$ or, alternatively, the cells of $\text{Vor}(A)$. The latter grouping is of interest to us, and we make it geometric by moving every vertex $\hat{\sigma}$ to an arbitrary but fixed interior point of σ^* , which we recall is the Voronoi cell dual to σ ; see Figure 4.9. The edges, triangles, and tetrahedra move accordingly.

Importantly, moving the vertices as described preserves the orientation of each tetrahedron. Hence, the tetrahedra fit together without gaps or improper overlap, like they did before the movement. In other words, the movement is an orientation-preserving piecewise linear homeomorphism between the Delaunay mosaic and a portion of the Voronoi tessellation. It can be extended to a piecewise linear homeomorphism from \mathbb{R}^3 to \mathbb{R}^3 by subdividing the space outside the Delaunay mosaic into unbounded polyhedra, and this can be done such that every cell of the Voronoi tessellation is a union of simplices and unbounded polyhedra that share the selected point in the interior of the cell as a vertex, see again Figure 4.9.

Let γ be a knot and $\text{Del}(\gamma, A)$ its Delaunay loop, in which we assume that γ meets $\text{Vor}(A)$ generically. Every edge in $\text{Del}(\gamma, A)$ is subdivided into two edges in the barycentric subdivision of the Delaunay mosaic. Let $\text{Del}'(\gamma, A)$ be the polygonal knot obtained after replacing each edge in $\text{Del}(\gamma, A)$ by the two edges and the vertex they share. It consists of twice as many edges and vertices, but it has the same underlying space in \mathbb{R}^3 . Finally, we move the vertices to the selected interior points of the dual Voronoi cells to get a different geometric realisation of the polygonal knot, which we denote $\text{Del}''(\gamma, A)$.

Lemma 4 (Equivalence of Subdivided Knots). *Let $A \subseteq \mathbb{R}^3$ be locally finite, $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$ a knot, and recall that $\text{Del}'(\gamma, A)$ and $\text{Del}''(\gamma, A)$ are different geometric realisations of the Delaunay loop of γ and A . Then $\text{Del}'(\gamma, A)$ and $\text{Del}''(\gamma, A)$ are equivalent.*

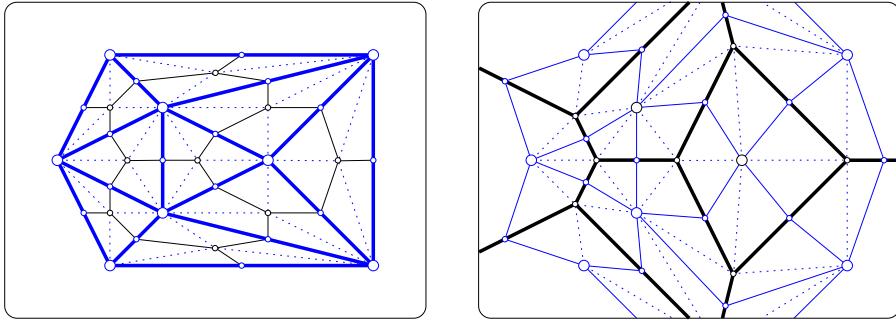


Figure 4.9: The barycentric subdivision of the Delaunay mosaic of eight points in \mathbb{R}^2 , before moving its vertices to the selected points of the Voronoi cells on the left and after moving them on the right.

Proof. Observe that the map from $\text{Del}'(\gamma, A)$ to $\text{Del}''(\gamma, A)$ is the restriction of the orientation-preserving homeomorphism introduced above. The existence of such a homeomorphism is the required condition for the two knots to be equivalent. \square

Let $\eta: [0, 1] \rightarrow B^3$ be an arc embedded in the unit 3-ball such that the endpoints map to different points on the boundary of the 3-ball. We say η can be straightened within B^3 if there is a ambient isotopy $B^3 \rightarrow B^3$ that fixes the boundary and maps η to the straight line segment connecting $\eta(0)$ and $\eta(1)$. We use this notion to specify when we deem a Voronoi tessellation sufficiently fine to consider the Delaunay loop of a knot.

We say $\text{Vor}(A)$ resolves a knot, $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$, if γ meets $\text{Vor}(A)$ generically, its intersection with any 2-cell in $\text{Vor}(A)$ is either empty or a single points, and its intersection with any 3-cell is either empty or a single arc that can be straightened within the 3-cell without crossing.

Lemma 5 (Resolved Knot). *Let $A \subseteq \mathbb{R}^3$ be locally finite and $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$ a knot. If $\text{Vor}(A)$ resolves γ , then $\text{Del}(\gamma, A)$ is a (polygonal) knot, and $\gamma, \text{Del}(\gamma, A)$ are equivalent.*

Proof. First we will prove that γ and $\text{Del}''(\gamma, A)$ are equivalent. As mentioned earlier, $\text{Del}''(\gamma, A)$ and $\text{Del}'(\gamma, A)$ are equivalent, and since $\text{Del}'(\gamma, A)$ and $\text{Del}(\gamma, A)$ have the same underlying space, and are therefore trivially equivalent, this will imply the claim.

To prove the equivalence of γ and $\text{Del}''(\gamma, A)$, consider a 3-cell, ν , generated by a point $a \in A$, that has a non-empty intersection with γ . Since $\text{Vor}(A)$ resolves γ , this intersection is a single arc, $\eta: [0, 1] \rightarrow \nu$, that can be straightened within ν . The homeomorphism that does the straightening does not affect the boundary of ν . We compose this homeomorphism with the straight-line homotopy that maps the line segment with endpoints $\eta(0)$ and $\eta(1)$ on the boundary of ν to the union of the two edges in $\text{Del}''(\gamma, A)$ incident to the point $a \in A$. These two edges share one endpoint, and the remaining two endpoints lie on the boundary of ν , and more specifically in the same 2-cells that contain $\eta(0)$ and $\eta(1)$.

Since the homeomorphism does not affect the boundary of ν , the only effect is the movement of $\eta(0)$ and $\eta(1)$ to selected interior points of their respective 2-cells in the boundary of ν . The homeomorphism within ν can therefore be combined with the homeomorphisms within the neighbouring 3-cells. This implies that there is a homeomorphism that moves γ to $\text{Del}''(\gamma, A)$, which completes the proof of the lemma. This homeomorphism preserves orientation because the homeomorphisms within the cells preserve the boundaries and therefore the orientation. \square

In the remainder of this section we present the proofs for each claim in the following theorem.

Theorem 8. *Let $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$ be a knot, $A \subseteq \mathbb{R}^3$ a locally finite set with $\text{Conv}(A) = \mathbb{R}^3$.*

1. *If A is a lattice, then $\text{Avg}[\text{Wr}(\text{Del}(\gamma, A))] = \text{Wr}(\gamma)$.*
2. *If A is a stationary Poisson point process, then $\text{Exp}[\text{Wr}(\text{Del}(\gamma, A))] = \text{Wr}(\gamma)$.*

However, we will first need to define the chirality of a polyhedron and prove a short lemma on the chirality of centrally symmetric polyhedra.

Specifically, we ask ourselves whether there are convex polyhedra that can arise as a Voronoi cell that have a bias toward left-turns or right-turns, and how big such a bias may be.

4.4.2 Chiral Polyhedra

To formalise this question, we first introduce notation for the turning angle in the projection. Fixing the convex polyhedron, P , we introduce the map $\alpha: \mathbb{G}_{1,3} \times \mathbb{R}^2 \rightarrow (-\pi, \pi]$, in which $\mathbb{G}_{1,3}$ is the linear Grassmannian of lines in \mathbb{R}^3 . Write F_b for the facet with outer normal $b - a$. For a line $L \in \mathbb{G}_{1,3}$ and a point $y \in L^\perp$, the plane orthogonal to L , assume $L + y$ enters P at an interior point of F_b and exits at an interior point of F_c . In this case, we set $\alpha(L + y)$ equal to the turning angle defined by the projection of the ordered triplet b, a, c in the direction parallel to L . Note that reversing the orientation of $L + y$ preserves this angle. We set $\alpha(L + y) = 0$ if $L + y$ does not intersect P or the intersection between $L + y$ and P is not generic.

The chirality of P is the average turning angle in the projection defined by the lines that intersect P :

$$\chi(P) = \frac{1}{2\pi} \int_{L \in \mathbb{G}_{1,3}} \frac{1}{\text{Area}(P_L)} \int_{y \in L^\perp} \alpha(L + y) \, dy \, dL, \quad (4.12)$$

in which we use half the area of the unit sphere as the measure of $\mathbb{G}_{1,3}$, and write P_L for the projection of P onto L^\perp .

For trivial reasons, $-\pi \leq \chi(P) \leq \pi$. It is plausible and also true that $\chi(P) = 0$ if P and its central reflection, $-P$, are congruent. In other words, if there is a rigid motion (not including a reflection across a plane) that maps P to $-P$. These bodies include all centrally symmetric bodies, for which $P = -P$, but also others, such as the regular tetrahedron and, more generally, all regular pyramids.

Lemma 6 (Polyhedra with Congruent Central Reflections have Vanishing Chirality). *Let P be a convex polyhedron with non-empty interior congruent to $-P$ in \mathbb{R}^3 . Then $\chi(P) = 0$.*

Proof. Write $\mu: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ for a rigid motion that maps P to $-P$. Let L be a directed line that enters and exits P at interior points x of F_b and y of F_c . Then $\alpha = \alpha(L)$ is the turning angle of the projection of $b, 0, c$ in the direction of L . Applying the central reflection, we get $-P$ and $-L$ and thus a turning angle of $-\alpha$. Applying μ^{-1} , which is a rigid motion, we turn $-P$ into P and $-L$ into $\mu^{-1}(-L)$, and we preserve the turning angle, $-\alpha$.

Hence, we obtain a pairing of the lines, L paired with $\mu^{-1}(-L)$, such that the turning angles for the two lines cancel. This implies that the chirality of P vanishes, as claimed. \square

Let L be the line of points $p + \lambda u$, with $p \in \mathbb{R}^3$ and $\lambda \in \mathbb{R}$, and let $u \in \mathbb{S}^2$. For a point $x \in \mathbb{R}^3$, we call $\langle x, u \rangle$ the depth of x , noting that it increases from one end of L to the other. We project $\text{Del}(L, A)$ in direction u to get a possibly self-intersecting polygonal curve in the plane. Write e_0, e_1, \dots, e_n for the sequence of edges of $\text{Del}(L, A)$, write v_i for the vertex shared by e_{i-1} and e_i , and let $\alpha_i \in (-\pi, \pi]$ be the turning angle at v_i , for $1 \leq i \leq n$. We have $\alpha_i > 0$ if the ordered triplet v_{i-1}, v_i, v_{i+1} is a left-turn, $\alpha_i < 0$ if it is a right-turn, $\alpha_i = 0$ if the points lie on a straight line in sequence, and $\alpha_i = \pi$ if the points are not in sequence. The total signed curvature of the curve is the sum of the turning angles:

$$\kappa(\text{Del}(L, A)) = \sum_{i=1}^n \alpha_i. \quad (4.13)$$

By assumption of A being finite, e_0 and e_n are parallel, which implies that $\kappa(L)$ is an integer multiple of 2π . We show that this integer is the negative writhe of $\text{Del}(\gamma, A)$ in direction u .

Lemma 7 (Writhe of Path Diagram). *Let $A \subseteq \mathbb{R}^3$ be finite, $u \in \mathbb{S}^2$, and L a line parallel to u that meets $\text{Vor}(A)$ generically. Then the writhe of $\text{Del}(L, A)$ in direction u is minus the total signed curvature of the projection of $\text{Del}(L, A)$ divided by 2π : $\text{Wr}_u(\text{Del}(L, A)) = -1/2\pi\kappa(\text{Del}(L, A))$.*

Proof. Let L' be a line that is neither parallel nor orthogonal to L . The second requirement implies that the depth increases monotonically from one end of L' to the other. Consider the straight-line homotopy from $E_A(L)$ to L' that preserves depth. During this homotopy, the path diagram goes through a sequence of Reidemeister moves. Type-II and Type-III moves affect neither the writhe nor the total signed curvature. Let n_+ be the number of Type-I moves that decrease the writhe, and note that each increases the total signed curvature by 2π . Let n_- be the number of Type-I moves that increase the writhe, and note that each decreases the total signed curvature by 2π .

We have $\text{Wr}_u(L') = 0$ and therefore $\text{Wr}_u(\text{Del}(L, A)) = n_+ - n_-$. Similarly, $\kappa(L') = 0$ and therefore $\kappa(\text{Del}(L, A)) = 2\pi(-n_+ + n_-)$. The claimed relation follows. \square

4.4.3 Proof of Theorem 8

We are now prepared to complete the proofs of the two claims in Theorem 8.

Proof of Claim 1. If A is a lattice, then all 3-cells in $\text{Vor}(A)$ are translates of the 3-cell generated by $0 \in A$, and this 3-cell is centrally symmetric. Lemma 6 implies the total signed curvature of a line through the 3-cell averaged over all rigid motions vanishes, and hence contributes nothing to the total for the entire Delaunay path of a line. Without any total signed curvature Lemma 7 implies there is also no writhe. The latter also holds for knots, which implies Claim 1 of Theorem 8. \square

Proof of Claim 2. Let P be a convex polyhedron with non-empty interior in \mathbb{R}^3 , and let $-P$ be the reflection of P across a plane. Observe that $\chi(-P) = -\chi(P)$. Assuming $A \subseteq \mathbb{R}^3$ is a stationary Poisson point process, every convex polyhedron is as likely as its reflection in the Voronoi tessellation. Fixing a point, $a \in A$, and varying the others, this implies that the expected chirality of the 3-cell generated by a vanishes.

The total signed curvature of the Delaunay path of a line is the sum of turning angles defined by the line and the 3-cells in $\text{Vor}(A)$. The expectation of this sum is the sum of expectations,

which is 0. Applying this result to line segments within an input piecewise linear knot implies there is no additional contribution to writhe from the Delaunay paths of each individual line segment. Therefore, Claim 2 of Theorem 8 holds. \square

4.5 Discussion

We venture that the result for intrinsic volumes generalises to all intrinsic volumes in d -dimensional Euclidean spaces.

Conjecture. *Let $A \subseteq \mathbb{R}^d$ be a stationary Poisson point process with intensity $\varrho > 0$, and $\mathcal{S} \subseteq \mathbb{R}^d$ a rectifiable space. Then there is a constant, $c_{d,i}$, such that the limit of the i -th intrinsic volume of $\text{Del}(\mathcal{S}, A)$, with $\varrho \rightarrow \infty$, is $c_{d,i}$ times the i -th intrinsic volume of \mathcal{S} .*

The current evidence for this conjecture are the proofs for the i -dimensional volume measures in [EN25] and the mean curvature in this paper.

Our original goal for the writhe, which we were not able to prove, was the average non-distortion over rigid motions, which we leave as a conjecture.

Conjecture. *For $\gamma: \mathbb{S}^1 \rightarrow \mathbb{R}^3$ a knot, $A \subseteq \mathbb{R}^3$ a locally finite set with $\text{Conv}(A) = \mathbb{R}^3$,*

$$\text{Avg}[\text{Wr}(\text{Del}(\gamma, A))] = \text{Wr}(\gamma). \quad (4.14)$$

This question originally motivated our investigation of chiral polytopes in Section 4.4.2. Specifically, if the conjecture holds there would not exist chiral polyhedra which tile \mathbb{R}^3 as the Voronoi tessellation of a point set $A \subseteq \mathbb{R}^3$.

Braiding Vineyards

5.1 Introduction

In this chapter we demonstrate a surprising link between persistence vineyards and knot theory. Recall, we introduced persistent vineyards in Section 2.7.3 and used them to study shapes and their higher order medial axes in Section 3.3. Similarly, we introduced the fundamentals of knot theory in Section 2.2.1. Note that there are many computational aspects of knot theory which overlap with the study and characterization of low dimensional manifolds' topology, both of which have a long history of algorithmic development, see Dehn's algorithm [Deh11, Deh12] and the many following algorithmic results in more recent decades on shape and knot recognition in low dimensions [Rub95, Tho94, Has97, BH99, Bur20].

The first goal of this chapter is to study the aforementioned new link between these two important branches of computational topology. Thanks to the work of Alexander [Ale23] we know that every knot or link can be represented as a braid. That is, for every link there is a braid such that if we glue together the ends of the braid, we recover the link. Recall also that for a continuous one parameter family of filtrations, we can "stack" the persistence diagrams of these filtrations; we call the resulting object a vineyard [CSEM06, Tur23, Hic22]. Thanks to the stability of persistence diagrams, the points in the persistence diagram move continuously (even Lipschitz continuously) with the parameter. This means that we can follow a point in the stack of the persistence diagrams and the resulting curve is called a vine. We will prove that for every link there exists an embedded manifold and a family of functions on \mathcal{M} (where each function is induced by the distance to a point in the ambient space, and where in turn each point comes from a curve γ) such that the vineyard of the family of function yields the braid representing the link in the sense of Alexander.

The second goal of this chapter is to show that any type of monodromy can occur in vineyards. This is part of a new research direction in computational topology. In [AGH⁺24], the occurrence of monodromy in the context of the directional persistence transform is studied, more precisely for 0-dimensional persistence modules of objects embedded in \mathbb{R}^2 . Roughly speaking, the directional persistence transform considers the persistence diagram of the height function on a shape for any direction, which in two dimensions gives what we call a closed vineyard (whose precise definition will be given below). It should be noted that monodromy had already been identified much earlier in the context of multidimensional persistence in [CEF13] (see also the Applied Algebraic Topology Research Network video, [Net25]). The authors of [AGH⁺24]

conclude with an open question about demonstrating monodromy in higher dimensions, as well as several interesting and more open ended questions related to better understanding what monodromy captures about the input shape.

The setting of [AGH⁺24] is the directional persistent homology transform, the vineyard generated by the family of height functions along all possible directions. Recently however, this has been placed in a larger context [OOT24], where one studies the distance to all flats of any given dimension. In this chapter we focus on that extreme case of 0-flats, which are simply points. The resulting radial transform is the vineyard generated by the family of distance functions to each point in \mathbb{R}^d , allows for a more geometric understanding of monodromy, and is critical in our construction.

In this chapter we exhibit that any type of monodromy and the braid associated to any link can occur in a vineyard. To make our statement more precise, however, we need to introduce some nomenclature, although full definitions will be deferred until Section 5.2. Intuitively, monodromy is the effect where if one makes a loop in a base space of a covering or fibre bundle, the lifted curve may not end up in the same point as you started out with. We say that the monodromy is of period $2\pi k$ (with $k > 0$) if the lifted curve returns to the starting point after k revolutions in the base space. In our context, the base space is a closed curve or loop $\gamma : [0, 2\pi] \rightarrow \mathbb{R}^d$, it is into this image \mathbb{R}^d that we have embedded a manifold \mathcal{M} (link or some offset of link, which is a modification of the input link). The fibres are the persistence diagrams of the Euclidean distance function restricted to the manifold \mathcal{M} , that is $\rho_t(p)$. The bundle therefore is the vineyard. The lifted curve is a vine $\tilde{\gamma}_{(b_0, d_0)}(t)$ in the vineyard starting at (b_0, d_0) in the persistence diagram of $\rho(x, \gamma(0))_{\mathcal{M}}$ and the periodicity is the smallest $k > 0$, such that for all i , $\tilde{\gamma}_{(b_0, d_0)}(0) = \tilde{\gamma}_{(b_i, d_i)}(2\pi k)$, where we assume that the vine is non-degenerate in the sense that it stays away from the diagonal.

The main results for this chapter are that any type of monodromy and braid can be generated in a vineyard. These are presented in Section 5.3. Additionally, we present interesting results linking monodromy to the extended symmetry set in Section 5.4.

5.2 Preliminaries

5.2.1 Monodromy

Monodromy is an important concept in mathematics that appears in various guises. We refer to the review [Ebe05] (the first part of which is almost a review of reviews) and the other reviews mentioned in that paper for an overview of the various aspects of the theory. In this chapter we will only consider the simplest incarnation, and only in the setting of topological data analysis.

Let \tilde{X} be a covering space of X with covering map $C : \tilde{X} \rightarrow X$, that is for every $x \in X$ there exists an open neighbourhood $x \in U$ and a discrete set J , such that $C^{-1}(U) = \sqcup_{i \in J} V_i$ and $C|_{V_j} : V_j \rightarrow U$ is a homeomorphism for all $j \in J$. We call the inverse images of points $x \in X$ of the map C the fibres. For a curve $\gamma : [0, 2\pi] \rightarrow X$ we write $\tilde{\gamma}$ for (one of) its lift(s), that is a continuous map $\tilde{\gamma} : [0, 2\pi] \rightarrow \tilde{X}$ such that $C \cdot \tilde{\gamma} = \gamma$. If γ is a loop, that is $\gamma(0) = \gamma(2\pi)$, then we say that γ exhibits monodromy (at the starting point $\tilde{\gamma}(0) \in C^{-1}(\gamma(0))$) if we have that the start and end points of its lift $\tilde{\gamma}$ are different, i.e. $\tilde{\gamma}(0) \neq \tilde{\gamma}(2\pi)$. The difference between $\tilde{\gamma}(0)$ and $\tilde{\gamma}(2\pi)$ is also referred to as monodromy (this difference can in certain cases be best represented by a group, see [Ebe05], although we will not need this in our discussion).

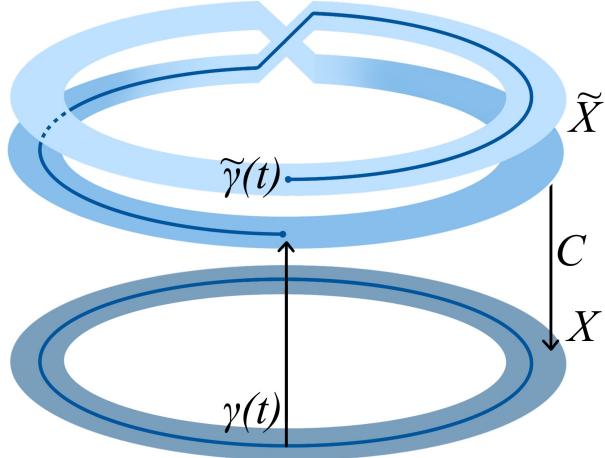


Figure 5.1: Here we see a cover \tilde{X} (in this case a double cover) of the base space X , in this case a circle as well as the curve γ and its lift $\tilde{\gamma}$.

If γ is a loop we can extend it, formally speaking, by concatenating with itself. Here we adopt the convention that if γ and γ' are two curves parametrized by $[0, 2\pi]$, then the concatenation $\gamma \circ \gamma'$ is parametrized by $[0, 4\pi]$. That is, we do not rescale the parametrization interval. We write

$$\gamma^k = \underbrace{\gamma \circ \cdots \circ \gamma}_k$$

and $\tilde{\gamma}^k$ for its lifting. We stress that generally

$$\tilde{\gamma}^k \neq \underbrace{\tilde{\gamma} \circ \cdots \circ \tilde{\gamma}}_k,$$

where $\tilde{\gamma}$ is the lifting of γ . In fact, the right hand side does not even have to be a continuous curve. We say that a loop γ (parametrized by $[0, 2\pi]$) in the base space X exhibits monodromy of order k if k is the smallest positive integer such that the lifted curve $\tilde{\gamma}^k$ satisfies

$$\tilde{\gamma}^k(0) = \tilde{\gamma}^k(2\pi k).$$

If $k = 1$ we say that γ exhibits no or trivial monodromy.

5.2.2 Braids

In this section we formally introduce braids which, as we will see, are closely related to knots and links which were introduced in Section 2.2.1. The reader is advised to recall the definitions from Section 2.2.1 to draw a comparison to braids.

A braid on m strands is the disjoint union of m intervals embedded in a solid cylinder, $D^2 \times I$, where each strand, $B_i: I \rightarrow D^2 \times I$, monotonically increases with respect to I , such that $B_i(0) = (d_i, 0)$ and $B_i(1) = (d_j, 1)$ for each i , where i indexes the strand, so that the set of endpoints of strands is some permutation of the set of origins. Two braids are equivalent if they are related by an ambient isotopy which fixes the endpoints, and at each time during the isotopy the image is also a braid. For each $u \in \mathbb{S}^1$, the projection of a braid in the direction u provides a braid diagram of B with crossings. In this chapter we will assume braid diagrams with generic crossings, and we further use that each braid can be represented in a

piecewise vertical diagram. That is, the strands in the diagram are vertical except in a (small) neighbourhood of a crossing. The fact that braids can always be represented in such a way seems to be folklore; see for example “Reidemeister’s Theorem” and the discussion in [BB05].

A closed braid or braided link is the image of a braid under the map from the solid cylinder to the solid torus, $D^2 \times I \rightarrow D^2 \times S^1 \subset \mathbb{R}^2 \times \mathbb{C} \simeq \mathbb{R}^4$, which sends $(x, y) \mapsto (x, e^{iy})$, where e^{iy} gives the standard embedding from $\mathbb{R}/2\pi\mathbb{Z}$ into $\mathbb{C} \simeq \mathbb{R}^2$. The braid index is the minimum number of strands required to form a closed braid equivalent to a given link.

Under the standard embedding \mathbb{T} of the solid torus in 3-dimensional Euclidean space, where the embedding is rotationally symmetric around the z -axis and where each strand is oriented positively, a braid with n strands, once closed, can be considered a link with m components, where $m \leq n$. Note if we concatenate a braid enough times such that the permutation of its endpoints is trivial, then we have $m = n$. Importantly, the orientation of each resulting component is aligned to a positive orientation on the core circle of the solid torus at all points.

We note that previous work has connected braids and braid groups with monodromy [CA11, CS97, Sal23, Sal24], further motivating the connection which we explore in this chapter.

The following result of Alexander will be essential to our result.

Theorem 9 (Alexander 1923 [Ale23]). *Every knot or link is equivalent to a closed braid.*

Of course this correspondence is not bijective as each link may be equivalent to many closed braids. An algorithmic alternative proof of this result was later given in [Vog90]. The complexity of the algorithm depends on the number of Seifert circles, which are defined as follows. Given an oriented link diagram, by eliminating each crossing and connecting each incoming strand with its adjacent outgoing strand we obtain a diagram of oriented circles known as Seifert circles. By construction the circles do not intersect.

The number of elementary operations of the algorithm in [Vog90] to obtain a braid diagram from a given link diagram with c crossings and s Seifert circles is at most $(s-1)(s-2)/2$ and the number of crossings in the resulting braid is at most $c + (s-1)(s-2)$.

The braid index of a link is the smallest number of strands needed for a closed braid representation of the link. The braid index is equal to the minimal number of Seifert circles over all diagram of the braid [Yam87].

5.3 Monodromy in Vineyards

In the context of persistence of some topological space with the induced distance function from a point, we need to introduce a number of conventions for monodromy to be well defined. In fact we will introduce monodromy both associated to an entire vineyard as well as to a single vine in the vineyard. We will then present the main results of this chapter, demonstrating monodromy and braiding in detail through several examples.

5.3.1 Monodromy in vineyards: a geometric viewpoint

Let \mathcal{M} be a manifold embedded in \mathbb{R}^d and $\gamma : [0, 2\pi] \rightarrow \mathbb{R}^d$ be a parametrization of a loop γ . Let $\rho_t(p) : \mathcal{M} \rightarrow \mathbb{R}$ be the function $p \mapsto \|p - \gamma(t)\|$, note that we exclusively refer to points in \mathcal{M} as p . We will now identify the ends of the interval $[0, 2\pi]$ – that is, we pass to \mathbb{S}^1 . For each t the function $\rho_t(p)$ induces a filtration on \mathcal{M} , by the sub-level sets of the

function. Therefore we can consider (for each t) the degree l persistence diagram of this filtration $\mathcal{D}_l(\rho_t(p))$. The map

$$\begin{aligned} \text{CV}_{\mathcal{M}} : \mathbb{S}^1 &\rightarrow \mathbb{S}^1 \times \text{Dgm} \\ t &\mapsto (t, \mathcal{D}_l(\rho_t(p))), \end{aligned}$$

where Dgm is the space of persistence diagrams, is (trivially) a covering space of \mathbb{S}^1 . We refer to $\text{CV}_{\mathcal{M}}$ as the closed vineyard map. As proved in [CSEH05] the points in the persistence diagram are (Lipschitz) continuous with respect to t . The map $\text{CV}_{\mathcal{M}}$ is illustrated in Figure 5.2.

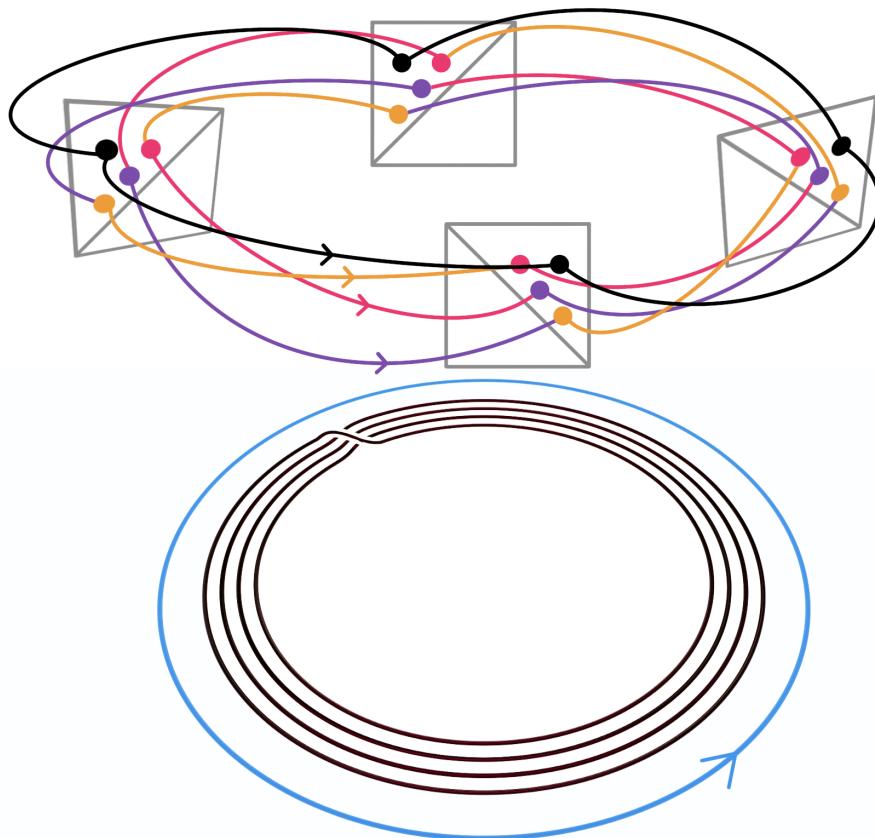


Figure 5.2: An illustration of the map $\text{CV}_{\mathcal{M}}$. We indicate the persistence diagrams, which form the fibers, only explicitly in a number of places for reasons of visibility. The knot diagram below, defines ouroboros knots. Here, we define the ouroboros as an unknot whose knot diagram is a (finite or segment of a) spiral with the end points connected with a monotone curve (where monotonicity refers to the relation between the angular and radial coordinates) such that the connecting segment of the knot diagram crosses over all the intermediate strands.

Remark 1 (Radial transform). *We follow the convention from [OOT24] and define the radial transform as the map that associates to any point $p \in \mathbb{R}^d$ the diagram $\mathcal{D}_l(\rho_t(p)) \in \text{Dgm}$, which is, slightly more general than the map $\text{CV}_{\mathcal{M}}$ introduced above.*

We note that generically [Tur23, page 3], there are no points of multiplicity greater than one in a persistence diagram of $\rho_t(p)$, and hence the connected components in the image of $\text{CV}_{\mathcal{M}}$, that is the vines, are non-intersecting curves. This justifies our earlier assumption from

Section 2.7.3, that our vineyards are generic in this sense, which we will use throughout this chapter.

This leads us to the definition of monodromy in the simplest setting (where we stay clear of the diagonal), which nonetheless is all that is needed for our constructions. Intuitively, we want to codify monodromy in terms of an integer, which encodes how many of copies of the vineyard must be glued together in order to have points return to their origins. More formally:

Assume that for all points $\gamma(t)$, there are no points of multiplicity higher than 1 in $\mathcal{D}_l(\rho_t(p))$, which means that all the vines are non-intersecting curves. Assume moreover that all of the vines are disjoint from the diagonal. We define the following:

Individual vine monodromy: Given a point in the persistence diagram $\mathcal{D}_l(\rho_\tau(p))$ for some $\tau \in \mathbb{S}^1$, write $\mathcal{V}(t)$ for the lift of $\gamma(t)$, which continuously assigns a point in $\mathcal{D}_l(\rho_t(p))$ for every t and yields the given point for $t = \tau$. We will now assume without loss of generality that $\tau = 0$. We say that the vine $\mathcal{V}(t)$ exhibits monodromy if the start and end point of the lifted curve do not coincide. Similar to $\mathcal{V}(t)$, write $\tilde{\mathcal{V}}^k(t)$ for the analogous lift of $\gamma^k(t)$. Then $\mathcal{V}(t)$ exhibits monodromy of order k if k is the smallest integer strictly larger than 0, such that $\tilde{\mathcal{V}}^k(0) = \tilde{\mathcal{V}}^k(2\pi k)$. In other words, the order of the monodromy is the number of points above any t in the connected component in the image of CV_M that contains the given point in $\mathcal{D}_l(\rho_\tau(p))$.

Vineyard monodromy: Following an individual vine (with increasing time t) for a given point on $\mathcal{D}_l(\rho_0(p))$ yields a point in $\mathcal{D}_l(\rho_{2\pi}(p)) = \mathcal{D}_l(\rho_0(p))$. In other words, the vines or vineyard induce a map P_V from $\mathcal{D}_l(\rho_0(p))$ to itself, which permutes the points in the persistence diagram. We say that the vineyard has monodromy of order k if k is the order of the permutation, that is, the smallest integer $k > 0$ such that applying this permutation k times yields the identity permutation.

See Figure 5.3 for a simple sketch of a knot exhibiting monodromy.

We note that monodromy can be defined when the vines touch the diagonal, which corresponds more closely to the notion introduced in [AGH⁺24]. We refer the interested reader to Appendix 5.B, where we formalize this definition, but we do not need it for our results. In addition, there may be self-intersections on the diagonal of the completed vineyard. If the vineyards are non-generic, the situation is significantly more complex as vines can intersect [Tur23]. As this non-generic case is not needed, we will not consider it any further for this chapter.

We conclude this section by stating the two main theorems for this chapter.

Theorem 10 (Vineyard Monodromy). *The radial transform in \mathbb{R}^d can exhibit monodromy for persistence up to degree $(d - 2)$ homology and for extended persistence up to degree $(d - 1)$ homology. Moreover the order of the monodromy can be k for any integer $k \geq 2$.*

Theorem 11 (Vineyard Links). *Given a link L , and integers $d \geq 3$, $l \geq 0$, with $l < d - 2$, there exists a submanifold $M \subset \mathbb{R}^d$ and a closed curve $\gamma \subset \mathbb{R}^d$ such that identifying the ends of the degree l vineyard of $\rho_t(p)$ yields a link which contains the given link as a subset. That is, the vineyard is topologically equivalent to the link we were given after removing some spurious connected components, i.e. the output is a link $L' = \tilde{L} \sqcup \mathcal{Y}$, where \sqcup denotes the disjoint union of different connected components and \tilde{L} is ambient isotopic to L .*

We remark that our construction hinges not only on M but also on a careful choice of γ . Indeed, fixing M , there can exist closed curves which yield no links in the vineyard.

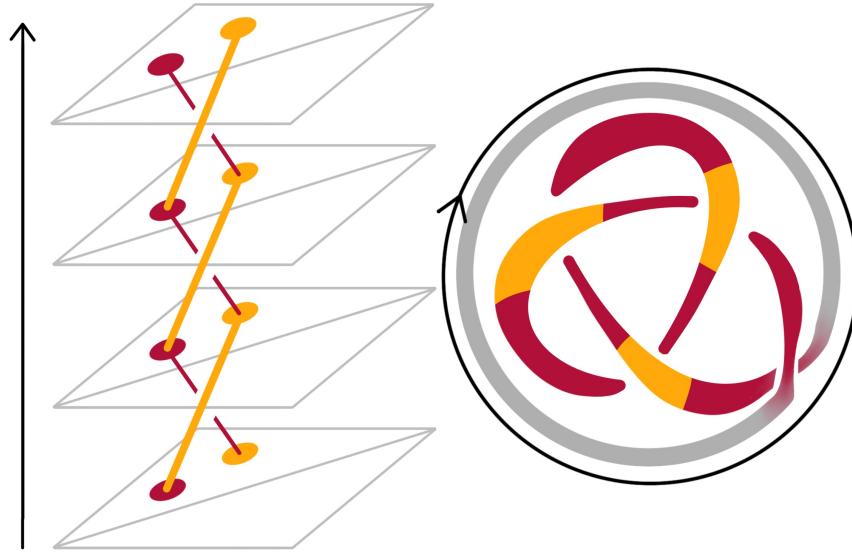


Figure 5.3: A schematic demonstrating the effect of the elder rule on our construction. Note that the gray part of the knot does not contribute to monodromy, as the elder rule dictates that the first birth is paired with the last death which causes interference with the desired braiding. By adding a new strand to the outside of the annulus via a trivial twist, we ensure that there is an elder vine that sits far away from the diagonal and doesn't interfere with our desired braid. To simplify the image, the gray portion is not shown in the vineyard.

5.3.2 Braids and Morse theory

We say that a closed braid B is (ϵ, R) -embedded in \mathbb{R}^3 if the closed braid is contained in the ϵ -thickening of the circle $C_h(0, R)$ of radius R contained in the horizontal plane. We will assume throughout this chapter that $\epsilon \leq R$.

For a point $b \in B$ let P be the plane spanned by p and the z -axis such that the normal to P , n_P is contained in the horizontal plane. Further, let t_p be the tangent vector to B at p . Then the maximal braid angle is given by, $\theta_B = \max_{p \in B} \angle T_p B, n_P$.

Intuitively, θ_B measures the maximal deviation of the tangent direction of the braid from the angular direction around the z -axis, and quantifies how closely the braid follows a circular trajectory.

We can consider a small $(l + 1)$ -dimensional α -offset \mathcal{M} of a closed braid B . We interpret ‘small’ here as follows: The offset should be small enough such that no self-intersection, nor intersections with the z -axis occur. In particular, it suffices for the offset to be small compared to the reach [Fed59] of the closed braid. Formally, the offset is defined as follows: For $l = 0$, we do not add an offset. For $l \geq 1$, we consider the embedding $\mathbb{T} \times 0$ in $\mathbb{R}^3 \times \mathbb{R}^{l-1}$, where \mathbb{R}^3 is the space that contains the standard embedding of the torus, in the sense of a closed braid. We take \mathcal{M} to be the offset of the braid $B \times 0$ in $\mathbb{R}^3 \times \mathbb{R}^{l-1}$, that is the boundary of $(\mathbb{T} \times 0) \oplus B(0, \alpha)$, where \oplus denotes the Minkowski sum. The resulting manifold \mathcal{M} can then be embedded in \mathbb{R}^d where the \mathbb{R}^3 that contained \mathbb{T} corresponds to the first three coordinates. The maximal braid angle in this context is defined as the maximum over the braid angles: Let $p \in \mathcal{M}$ and $T_p \mathcal{M}$ be its tangent space, then the braid angle is the angle between the normal n of the hyperplane $P \times \mathbb{R}^{d-3}$, where P is the plane defined by the z -direction and the point p .

Lemma 8 (Clustering of critical points). *Let B be an (ϵ, R) -embedded closed braid and suppose that \mathcal{M} is its $(l+1)$ -dimensional α -offset, with braid angle θ_B . If $x \in \mathbb{R}^d$ satisfies $\rho(x, C_h(0, R)) \leq \eta$, $C_h(0, R)$ is parametrized by $s(\theta)$, and the closest point x' of x on $C_h(0, R)$ satisfies $x' = s(0)$, then the function $p \mapsto \|p - x\|$ has no critical points at the closest point $\beta(\theta)$ on B to p as long as,*

$$\frac{\theta}{2} + \theta_B + \arcsin \frac{\epsilon}{2R \sin(\frac{\theta}{2}) - \eta} + \arcsin \frac{\eta}{2R \sin(\frac{\theta}{2})} < \frac{\pi}{2}. \quad (5.1)$$

This implies in particular that there is no topological change as long as (5.1) is satisfied.

See Figure 5.6 for a illustration of this configuration.

The proof of this lemma depends on one of the Morse theorems referenced in Section 2.5, we state it explicitly as a theorem now.

Theorem 12 ([Mil69]). *Suppose f is a smooth real-valued function on \mathcal{M} , $a < b$, $f^{-1}[a, b]$ is compact, and there are no critical values between a and b . Then \mathcal{M}^a is diffeomorphic to \mathcal{M}^b , and \mathcal{M}^b deformation retracts onto \mathcal{M}^a .*

Proof of Lemma 8. We write $\beta(\theta)$ for a parametrization of B according to the angle of the circle $C_h(0, R)$. We parameterize the circle as $s(\theta) = (R \sin \theta, R \cos \theta, 0, \dots, 0)$. Using this notation we note that the gradient of the function $\mathcal{M} \rightarrow \mathbb{R}$ given by $\|q - x\|$ is zero at $q = p$ if for its closest point on B , that is $\beta(\theta)$, we have that $\langle \beta', \beta - x \rangle = 0$, where β' denotes the derivative of β with respect to θ . This is equivalent to

$$\angle \beta', \beta - x = \pi/2.$$

By assumption we have that

$$\angle \beta', s' \leq \theta_B.$$

Moreover, because $|\beta(\theta) - s(\theta)| \leq \epsilon$, we also have

$$\sin(\angle \beta - x, s - x) \leq \frac{\epsilon}{|s - x|},$$

as can be seen from Figure 5.4.

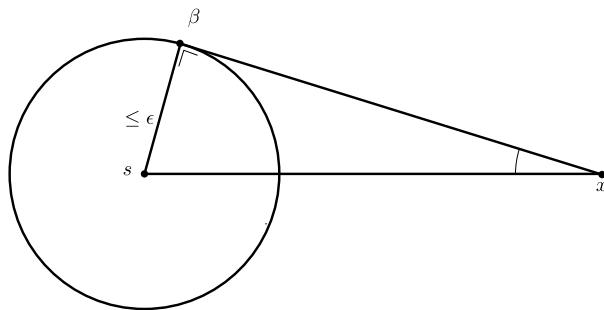


Figure 5.4: The angle estimate for $\angle \beta - x, s - x$.

Using the triangle inequality of angles (or points on the sphere) we find that

$$|\angle(\beta', \beta - x) - \angle(s', s - x)| \leq \theta_B + \arcsin \frac{\epsilon}{|s - x|}. \quad (5.2)$$

Let us now write x' for the closest point projection of x on $C_h(0, R)$ and $\phi = \angle(s', s - x')$. By reparameterization we can assume that $x' = s(0)$. With this assumption we see by the construction in Figure 5.5 that $\phi = \angle(s', s - p') = \frac{\theta}{2}$. Moreover we have that $|x' - s| = 2R \sin(\frac{\theta}{2})$. By the same argument as given in Figure 5.4 we have that

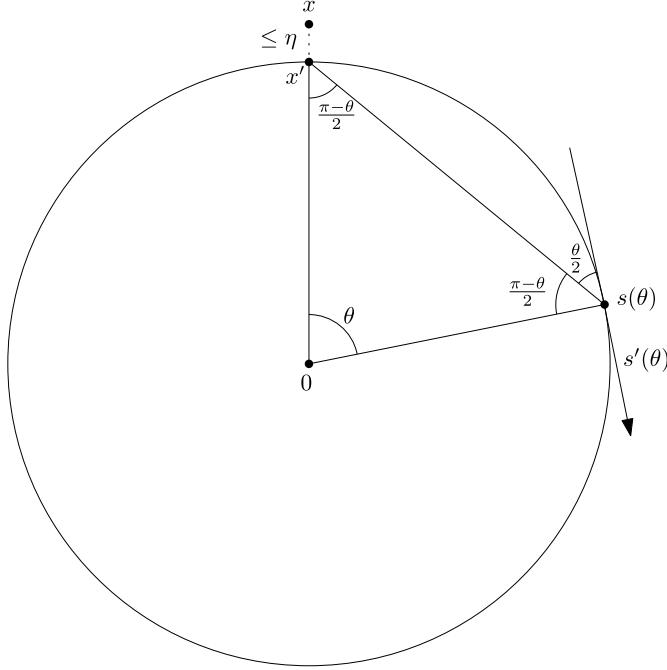


Figure 5.5: The construction for the angle $\phi = \angle(s', s - x')$.

$$\sin \angle(s - x', s - x) \leq \frac{\eta}{|s - x'|},$$

so that together with (5.2) we find that

$$\begin{aligned} \left| \angle(\beta', \beta - x) - \frac{\theta}{2} \right| &\leq \theta_B + \arcsin \frac{\epsilon}{|s - x|} + \arcsin \frac{\eta}{|s - x'|} \\ &\leq \theta_B + \arcsin \frac{\epsilon}{|s - x'| - \eta} + \arcsin \frac{\eta}{|s - x'|} \\ &\quad \text{(by the triangle inequality)} \\ &\leq \theta_B + \arcsin \frac{\epsilon}{2R \sin(\frac{\theta}{2}) - \eta} + \arcsin \frac{\eta}{2R \sin(\frac{\theta}{2})} \end{aligned}$$

This means that if

$$\frac{\theta}{2} + \theta_B + \arcsin \frac{\epsilon}{2R \sin(\frac{\theta}{2}) - \eta} + \arcsin \frac{\eta}{2R \sin(\frac{\theta}{2})} < \frac{\pi}{2}$$

then $\angle(\beta', \beta - x) \neq \pi/2$ and hence there are no critical points. The fact that there is no topological change follows from Theorem 12. \square

Remark 2. The important conclusion from the bound of Lemma 8 is that by choosing $\epsilon < \eta \simeq \theta_B$ small compared to $\min\{R, \pi\}$, the bound (5.1) is satisfied as long as $\theta_B \ll \theta < \pi - 4\theta_B$.

We say that an oriented closed braid B that is (ϵ, R) -embedded is δ -circular if its parametrization according to arc length β satisfies $\left| \frac{\beta(t)}{R^2} + \ddot{\beta}(t) \right| < \delta$, where we use Newton's notation for the derivative.

Lemma 9 (Morse indices of clustered critical points). *Let B be an oriented (ϵ, R) -embedded, δ -circular closed braid, with braid angle θ_B . Suppose $x \in \mathbb{R}^d$ satisfies $\rho(x, C_h(0, R)) \leq \eta$, $C_h(0, R)$ is parametrized by $s(\theta)$, and the closest point x' of x on $C_h(0, R)$ satisfies $x' = s(0)$.*

If an oriented (ϵ, R) -embedded, δ -circular closed braid with n strands, with braid angle θ_B ,

$\epsilon \ll R$, and

$$\frac{6\epsilon}{R} + \delta(R + \eta) \leq \frac{1}{R^2}(R - \epsilon)(R - \eta),$$

and $\epsilon < \eta \simeq \theta_B$ are small compared to $\min\{R, \pi\}$, then the function $B \rightarrow \mathbb{R}$ given by $b \rightarrow \|b - x\|$ has $2n$ critical points, n maxima and n minima. Let $l \geq 1$. If \mathcal{M} is a $(l + 1)$ -dimensional α offset \mathcal{M} of the same type of braid satisfying the same conditions then the function $\mathcal{M} \rightarrow \mathbb{R}$ given by $p \mapsto \|p, x\|$ has $4n$ critical points, n maxima and n minima and n saddle points of Morse index l and n saddle points of Morse index 1.

Proof. Because the conditions of Lemma 8 are satisfied we know that there are no Morse critical points unless $\theta \simeq 0$ or $\theta \simeq \pi$ so we focus on establishing that there is only one critical point per strand at $\theta \simeq 0$ or $\theta \simeq \pi$ for B and or two per strand at $\theta \simeq 0$ or $\theta \simeq \pi$ in the case of \mathcal{M} .

The proof for the braid case is the difficult step, and we will see below that the statement for \mathcal{M} follows immediately. The idea of the proof is the following: Because the closed braid is δ -circular its parametrization β satisfies $\left| \frac{\beta(t)}{R^2} + \ddot{\beta}(t) \right| < \delta$. This means that β is forced to turn inward towards the centre of $C_h(0, R)$, following the circle $C_h(0, R)$, where 0 is the origin of Euclidean space. This in turn implies that β cannot satisfy $\frac{d}{dt}|\beta - x|^2(t) = 0$ for two times t that are relatively close. In other words the curve cannot be tangent to some sphere (not necessarily of the same radius) centred at x .

The way that we establish that $\frac{d}{dt}|\beta - x|^2(t)$ cannot be zero for two nearby values is by establishing that if $\frac{d}{dt}|\beta - x|^2(t) = 0$ then

$$\frac{d}{dt} \left(\frac{d}{dt} |\beta - x|^2 \right) (t) = \frac{d^2}{dt^2} |\beta - x|^2(t)$$

is large. If we write $\frac{\beta(t)}{R^2} + \ddot{\beta}(t) = \Delta(t)$ and suppress t from the notation, we see that

$$\begin{aligned} \frac{d^2}{dt^2} |\beta - x|^2(t) &= 2\langle \ddot{\beta}, \beta - x \rangle + 2\langle \dot{\beta}, \dot{\beta} \rangle \\ &= 2\langle \ddot{\beta}, \beta - x \rangle + 2 && (\text{because } |\dot{\beta}| = 1) \\ &= 2 \left\langle -\frac{\beta}{R^2} + \Delta, \beta - x \right\rangle + 2 \\ &= -2 \left\langle \frac{\beta}{R^2}, \beta \right\rangle + 2 \langle \Delta, \beta - x \rangle + 2 + 2 \left\langle \frac{\beta}{R^2}, x \right\rangle \end{aligned} \tag{5.3}$$

We can now examine the first three terms in (5.3):

- Because B is (ϵ, R) -embedded $R - \epsilon \leq |\beta| \leq R + \epsilon$, so that

$$-2 \left(1 + \frac{\epsilon}{R} \right)^2 \leq -2 \left\langle \frac{\beta}{R^2}, \beta \right\rangle \leq -2 \left(1 - \frac{\epsilon}{R} \right)^2,$$

which, if $\epsilon \ll R$, simplifies to

$$-2 - \frac{6\epsilon}{R} \leq -2 \left\langle \frac{\beta}{R^2}, \beta \right\rangle \leq -2 + \frac{6\epsilon}{R}.$$

- Because $|\Delta(t)| = |\frac{\beta(t)}{R^2} + \ddot{\beta}(t)| \leq \delta$, Cauchy-Schwarz yields that $2|\langle \Delta, \beta - x \rangle| \leq \delta|\beta - x| \leq \delta(R + \eta)$.

This implies that these three terms are close to zero, i.e.

$$\frac{d^2}{dt^2} |\beta - x|^2(t) \simeq 2 \left\langle \frac{\beta}{R^2}, x \right\rangle.$$

Because $\left| 2 \left\langle \frac{\beta}{R^2}, x \right\rangle \right|$ is lower bounded by $\frac{1}{R^2}(R - \epsilon)(R - \eta)$ assuming that the angle between β and x is no more than 45 degrees (or more than 135), the first part of the result now follows if, the angle between β and x is no more than 45 degrees (or more than 135), $\epsilon \ll R$,

$$\frac{6\epsilon}{R} + \delta(R + \eta) \leq \frac{1}{R^2}(R - \epsilon)(R - \omega_{\max}).$$

For the second part, note that there is a one-to-one correspondence between pairs of critical points of the distance function to a fixed point on an offset and the critical points of the distance function to the same fixed point on the curve β itself. Here minima of β correspond to a pair of a minimum and a critical point of index l on \mathcal{M} , while maxima correspond to a pair of a maximum and a critical point of index 1. \square

Given these $4n$ critical points, we can now consider the persistence diagram that results. Although we work with extended persistence throughout this chapter in order to avoid points at infinity and establish a perfect pairing of critical points, we note that in fact for the purposes of establishing monodromy in our knot offset, it suffices to restrict our attention to the behavior of the ordinary points as well as one single extended point, all above the diagonal, as those that contribute to the construction of knot or link as mentioned in Theorem 11 are born and die ‘on the way up’.

Corollary 2. *Under the same assumptions as in Lemma 9, the maxima and minima of the functions $B \rightarrow \mathbb{R}$ given by $b \mapsto \|b - x\|$ and $M \rightarrow \mathbb{R}$ given by $p \mapsto \|p - x\|$, as well as the saddle points of the latter function can be divided into two groups, one group occurring at low values and corresponding to \mathbf{H}_0 births of cycles in the persistence diagram and one group occurring at high values and corresponding to \mathbf{H}_0 deaths in the persistence diagram. Here, under the assumption¹ of Remark 2, low means $\lesssim R\theta_B$, where the \lesssim hides a constant, and high means $\geq 2R - 6R\theta_B$. For \mathbf{H}_l the situation is identical except for one change, namely that a single birth occurs at a high value. In ‘ordinary’ persistence theory this cycle lives forever, while in extended persistence it dies at the global minimum, and in fact lies below the diagonal.*

¹ More generally, Equation (5.1) of Lemma 8 divides \mathcal{M} (B respectively) into three regions, the part close to x where Morse critical points can occur, a large region where no Morse critical points can be found, and finally the part furthest from x where again one may find Morse critical points. The high and low in this corollary should be interpreted as such.

Proof. The only thing in this corollary that requires an extra argument on top of Lemma 9, is the correspondence of the critical points with the births and deaths respectively: Because, by Remark 2 there are no Morse critical points unless $\theta \simeq 0$ or $\theta \simeq \pi$, we know that $B(x, r) \cap B$ (respectively $B(x, r) \cap \mathcal{M}$) with $r \simeq R$ consists of n topological line segments (topological cylinder segments $\mathbb{S}^{l+1} \times [0, 1]$). While at for $r > 2R + \eta + \epsilon$, the set $B(p, r) \cap B$ (respectively $B(p, r) \cap \mathcal{M}$) consists of a topological circle or knot (its offset respectively). See Figure 5.6. The only way this can be achieved with the number of critical points we found in Lemma 9 is if the births and deaths occur as described in the statement of the corollary, by a simple counting argument or the pigeonhole principle. We distinguish the two cases, namely B and \mathcal{M} : For B we have n minima and n maxima, and all n minima are needed to create the n topological line segments (which we know exist if $r \simeq R$). Similarly, all n maxima are needed to create the handle attachments that recover the circle (which we know exist if $r > 2R + \eta + \epsilon$). For \mathcal{M} , we have n maxima and n minima, as well as n saddle points of index l and n of index 1. We need n minima and n saddle points of index l to create the n topological cylinder segments $\mathbb{S}^{l+1} \times [0, 1]$ (which we know exist if $r \simeq R$). To form $\mathbb{S}^{l+1} \times \mathbb{S}^1$ we need connect these segments to each other for which we need all the critical points of index 1, so that we end up with $\mathbb{S}^{l+1} \times \mathbb{S}^1$ with n punctures. We need all the maxima to fill all the punctures. \square

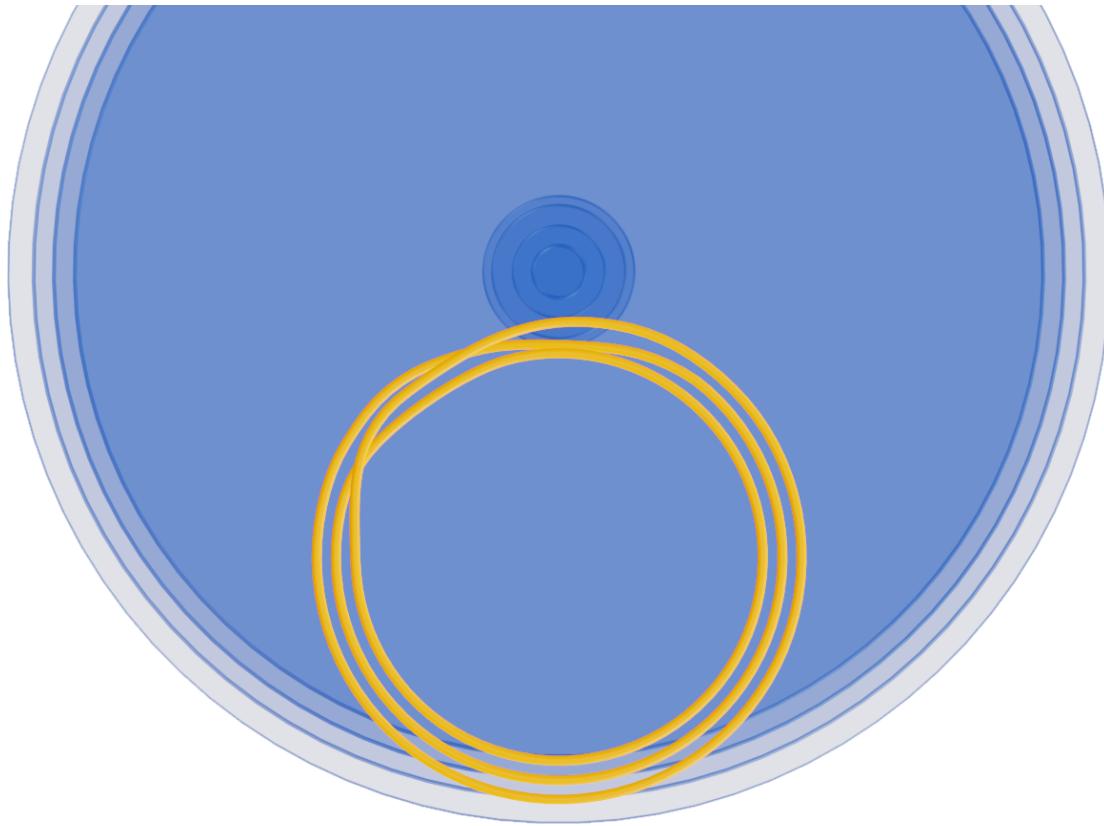


Figure 5.6: A figure illustrating the statement of Corollary 2, highlighting the intersection of the ouroboros (in yellow) and a family of 3D growing spheres that highlight the n births and n deaths in \mathbf{H}_0 .

Let f_1 and f_2 be two functions on the same manifold \mathcal{M} . We say that these two functions are handle-equivalent if the handle decomposition is the same and the times of insertion of these handles are also identical. We have the following observation.

Corollary 3. Let B be (ϵ, R) -embedded closed braid and suppose that \mathcal{M} is its $(l+1)$ -dimensional α -offset, with braid angle θ_B . We have that the B is a circle and the manifold \mathcal{M} is diffeomorphic to the torus $\mathbb{S}^{l+1} \times \mathbb{S}^1$. Under the same conditions as Lemma 9 (and Corollary 2), we have that, the maps $B \rightarrow \mathbb{R}$ given by $b \mapsto \|b - x\|$ and $M \rightarrow \mathbb{R}$ given by $p \mapsto \|p - x\|$ are handle equivalent to the height function of the embedding depicted in Figure 5.7.

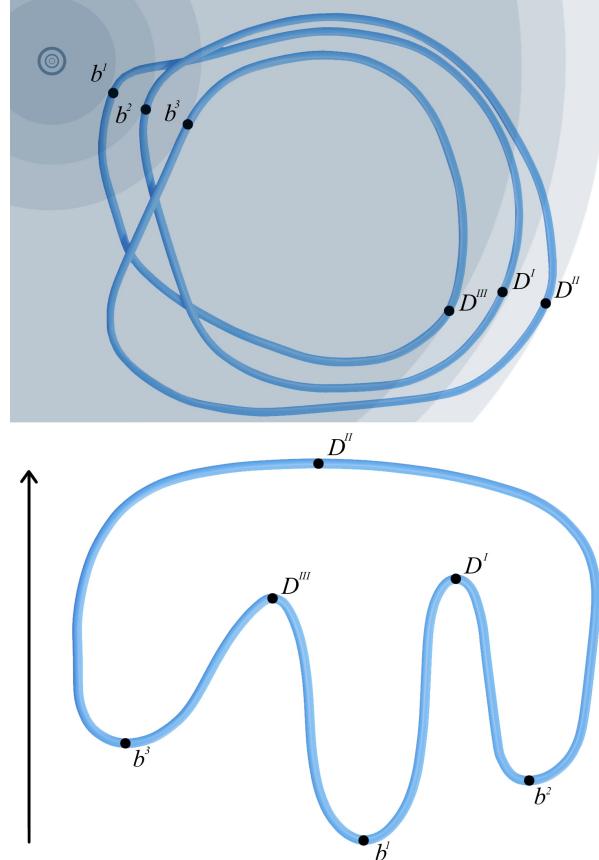


Figure 5.7: Top: The closed braid. Bottom: The embedding with equivalent height function. The arrow indicates the direction of the height function. The b^i 's indicate the birth times and D^J , where J is a roman numeral (i.e. $J \in \{I, II, III\}$), the death times. We stress that the picture should be interpreted in a 3D way, and in particular D^{II} does not have to be larger than D^J , with $J \neq II$. We stress that the critical point with the highest value of the Morse function (D^{II} in the figure) corresponds to a death only in extended persistence, in non-extended persistence, only a 1-cycle is born there.

Remark 3. We stress that given a closed braid B it is not difficult to adjust the embedding such that it is (ϵ, R) -embedded and δ -circular with ϵ and δ as small and R as large as you like. Because if B is a closed braid, its parametrization $\beta(t)$ can be written as $\beta(\theta) = n(\theta) + \omega(\sin(\theta), \cos(\theta), 0, \dots, 0)$ for some $\omega > 0$, with n normal to $(\cos(\theta), -\sin(\theta), 0, \dots, 0)$. By redefining $\beta(\theta) = \tilde{\epsilon}n(\theta) + R(\sin(\theta), \cos(\theta), 0, \dots, 0)$ for sufficiently small $\tilde{\epsilon}$ the δ -circular (ϵ, R) -embedding can be achieved. By the same argument the observation loop (to be defined in the proof of Theorem 11) can be made arbitrarily close in terms of η to a circle with radius R , so that it has reach as close to R as one likes, by [Fed59, Theorem 4.19], and hence has a nice tubular neighbourhood of that size.

5.3.3 Vineyard Braiding

We are now prepared to prove the first main result of this chapter, namely Theorem 11.

Proof of Theorem 11. The proof of this theorem is constructive. Thanks to Alexander's theorem, Theorem 9, a given link can be represented as a closed braid. To every connected component of the closed braid we add an extra loop (using a Reidemeister move of type I, so as not to disturb the equivalence class of the closed braid) as in Figure 5.8. This Reidemeister move may introduce extra crossings, but at most $\mathcal{O}(s \cdot K)$, where s denotes the number of strands in the original closed braid and K the number of connected components of the loop. We now write C for the number of crossings of the resulting closed braid. We write n for the resulting number of strands, for which we see that $n = s + K$. We define/construct a particular embedding of the closed braid such that the original braid appears in the vineyard. Essentially, the closed braid diagram is parametrized by ω and θ , while the embedding of the closed braid in \mathbb{R}^3 is parametrized by ω , θ , and h , as in Figure 5.11. Roughly speaking the parameter ω corresponds to the birth time, θ is the parameter of the observation loop γ (to be defined in Step 3) and h corresponds to the death time, again see Figure 5.11.

Our construction proceeds stepwise:

- **STEP 1** We start with an embedding which is close to its annular braid diagram, by which we mean that the embedding of the closed braid lies in a neighbourhood of an annulus in the plane and the braid is planar with the exception of small neighbourhoods of the points of crossing. We write θ, r for the coordinates of the annulus, which are polar coordinates in the plane restricted to the annulus.
- **STEP 2** We then modify (if necessary) the braid such that the crossings are equally parsed on one side of the annulus, that is, if θ is the angle that parameterizes the annulus, see Figure 5.10, then the crossings are contained in the interval $[0 + \frac{\pi}{8(C+1)}, \pi - \frac{\pi}{8(C+1)}]$, where C is the number of crossings. By equally parsed we mean that there is only one crossing in each of the intervals $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ where $\tilde{j} \in \{8, 16, \dots, 8C\}$.
- **STEP 3a** We now modify the embedding of the braid in an angular interval $[\pi - \frac{\pi}{4(C+1)}, 2\pi] \cup [0, \frac{\pi}{2(C+1)}]$. This interval should be interpreted in a periodic manner. We do so by twisting the annulus (and by extension the almost annular braid) 90 degrees in the direction orthogonal to the plane into which the annulus was originally embedded, see Figure 5.8. We do so in such a way that the twisted annulus and by extension part of the braid in the angular interval $[\pi + \frac{\pi}{8(C+1)}, 2\pi - \frac{\pi}{8(C+1)}]$ is now close to a cylinder. We do this in a way that preserves cylindrical coordinates, that is, if θ was the planar angular coordinate of a point on the annulus, after twisting θ is the cylindrical coordinate of the corresponding point. We denote the resulting twisted annulus by \mathcal{A}_T .
- **STEP 3b** We define an observation loop γ to be the curve that follows that twisted annulus on the outside at a constant distance (less than η , with η as in Lemmas 8 and 9), see e.g. Figure 5.9. We also define coordinates $\tilde{\theta}, \omega, \tilde{\psi}$ with respect to this observation loop in a tubular neighbourhood of size $\mathcal{O}(\eta)$ of the observation loop. This is possible thanks to Remark 3. The coordinate ω of a point y in the tubular neighbourhood is the distance to γ . The coordinate $\tilde{\psi}$ is the angle between a point y in the tubular neighbourhood of the curve, its closest point on curve $\pi_\gamma(y)$, and the closest point of this point on the annulus $\pi_{\mathcal{A}_T}(\pi_\gamma(y))$. See Figure 5.10.

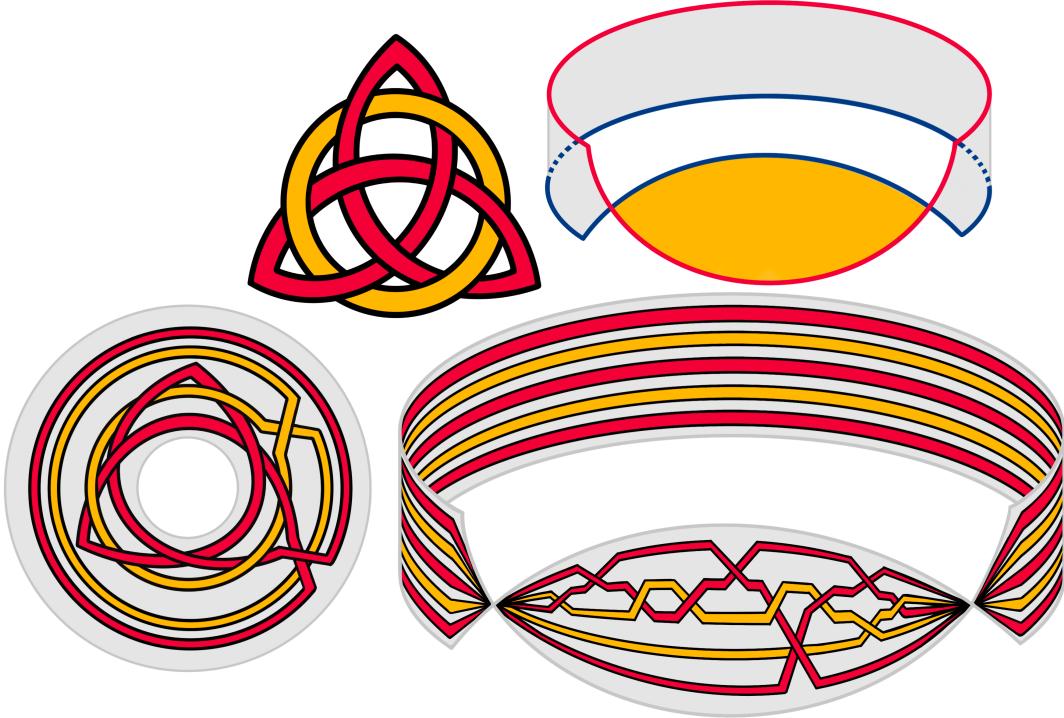


Figure 5.8: For a given link (top left), we construct a closed braid (that is, a braid with ends identified) embedded near an annulus (bottom left). Note that we add an additional trivial loop in each connected braid component traveling around the exterior of the annulus. We then partially twist the annulus (top right) such that the cylindrical vertical portion in gray and the horizontal portion in yellow are perpendicular to each other. In the top right figure, the top edge is shown in red and the bottom edge in dark blue to clarify this twisting. It looks pinched due to the perspective, but is in fact only twisted. The bottom right figure shows the result of twisting the braid together with the annulus, with some massaging to keep the crossings in the horizontal part. Due to the addition of the exterior trivial loops from the bottom left figure, the vineyard consisting of the persistence diagrams of the distance function to a point following the twisted annulus (bottom right) on the ‘outside’ contains the input braid with some surgery. See Figure 5.13 for several 3D views of the embedded knot, as well as 5.16 and the corresponding detailed example and discussion in Appendix 5.A.

Intermezzo: The persistence diagram for a point on the observation curve

Before we continue with the construction we now discuss the persistence diagram for a given point. In the following step we will further modify the closed braid in the angular interval $[\pi - \frac{\pi}{4(C+1)}, 2\pi] \cup [0, \frac{\pi}{2(C+1)}]$, but for now we consider the braid fixed. We also assume for now that the link has only one connected component, that is, it is a knot. Thanks to Lemmas 8 and Corollaries 2-3 we know that the braid and observation curve are chosen in such a way that all the births occur first and then all the deaths occur, in both distance order as well as in order along the braid. Let us in particular consider the equivalent embedding of Corollary 3, where we let b^1, b^2, \dots, b^n be the births as they occur in order following along the braid B , as in Figure 5.7; note that we are slightly abusing notation here, as we are identifying the births with the Morse critical points. Next, let $b_0^1, b_0^2, \dots, b_0^n$ respectively $b_l^1, b_l^2, \dots, b_l^n$, where the subscript indicates if the birth occurs in 0 or l -homology, be the corresponding ordered births of \mathcal{M} . Note that we always have n births unless $l = 1$, when instead there are $n + 1$ births. Again we emphasize that the births are ordered they occur following B starting

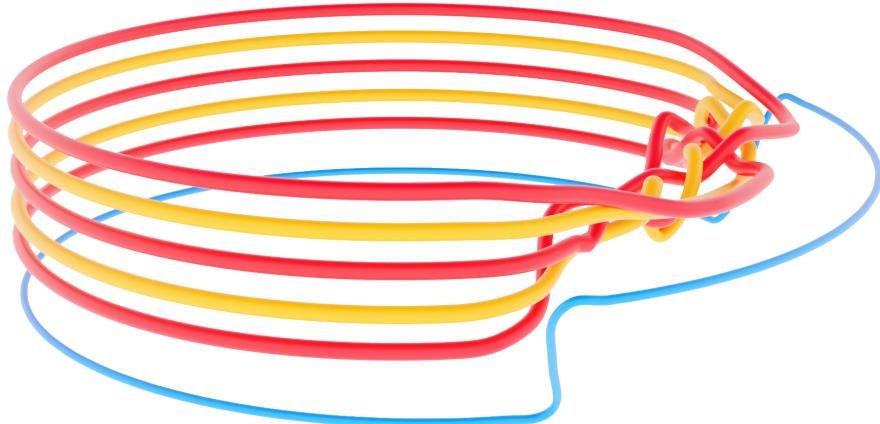


Figure 5.9: The braid follows a twisted annulus and the observation loop is in its proper place, but without the manipulation of the strands near the antipodal points of the crossing there can be incorrect crossings in the vineyard.

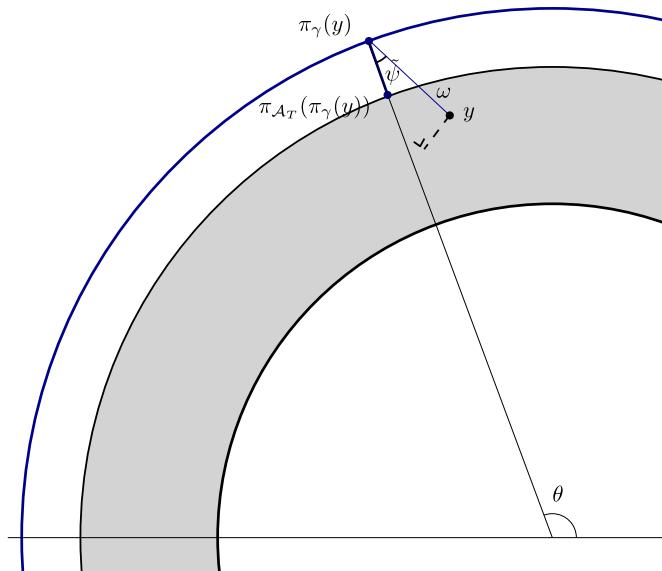


Figure 5.10: Figure illustrating the notation (the coordinates $(\theta, \omega, \tilde{\psi})$) in step 3b of Theorem 11. The observation loop is the outer loop shown in dark blue, and (the flat part of) the twisted annulus is shown in grey.

with the first birth, not in order of birth time according to the distance filtration; See Figure 5.7. Similarly, let D^I, D^{II}, \dots, D^N be the deaths as they occur in order if we consider B , and $D_0^I, D_0^{II}, \dots, D_0^N$ respectively $D_l^I, D_l^{II}, \dots, D_l^N$ be the deaths if we consider them on the offset of B , \mathcal{M} . Again we assume that the deaths are ordered as they occur following B , not in order of deaths time.

We will assume without loss of generality that b^1 is the lowest birth value. Because of the elder, rule the death of the cycle created at b^1 dies (in extended persistence) at the maximal death value, that is $\max_J D^J$, for B . Similarly b_0^1, b_l^1 respectively die at $\max_J D_0^J, \max_J D_l^J$ respectively for \mathcal{M} .

The other persistence points are less straightforward, as they follow the mergers of the sublevel sets of the equivalent embedding of Corollary 3; see Figure 5.7. However, the most important case for us will be the following: Consider first any ordinary \mathbf{H}_0 points, so we have b^1, b^j, b^k

with $1 \neq j < k \neq 1$. Assume that b^1 is the earliest birth. Suppose that the death times are all larger than all birth times and are ordered as follows

$$\begin{aligned} \max\{D^I, D^{II}, \dots, D^{J-2}, D^{K+1}, D^{K+2}, \dots, D^N\} &< \min\{D^{J-1}, D^K\} \\ \max\{D^{J-1}, D^K\} &< \min\{D^J, \dots, D^{K-1}\}, \end{aligned} \quad (5.4)$$

then the connected component born at time b^j merges with the connected component born at time b^1 at time D^{J-1} and the connected component born at time b^k merges with the connected component born at time b^1 at time D^K for B . In other words, under these conditions there are points (b^j, D^{J-1}) and (b^k, D^K) in the persistence diagram.

We now consider l -homology persistent homology of the manifold \mathcal{M} . In general, as in Figure 2.5, the births in l -homology follow the births in 0-homology closely. By this we mean the following: As before we write B for the braid and \mathcal{M} for its $(l+1)$ -offset, and we compare the same Morse function (distance to a point, or height if we consider the equivalent embedding) on these two spaces (B and \mathcal{M} , respectively). The birth of 0-cycles (b^1, b^2, b^3 in Figure 5.7) on the braid B are close both geometrically and with regard to the value of the Morse function (the distance to a point or the height function, where in the latter case we consider the equivalent embedding) to the critical points that give rise to the birth of 0-cycles and l -cycles on \mathcal{M} . There is also a simple correspondence with one exception between deaths of 0-cycles for the braid and deaths of 0- and l -cycles on \mathcal{M} , meaning that the critical points that correspond to deaths for B are close to a pair of critical points (one maximum and one saddle) that correspond to deaths in 0- and l -homology. The exception is the last death of a 0-cycle in extended persistence on B ; here we instead have a Morse critical point which corresponds to the birth of a 1-cycle for the braid and which lives forever in the non-extended persistence, but corresponds to a death of a 0-cycle in extended persistence. There are again two corresponding nearby critical points on \mathcal{M} for this final critical point on B , however in this case the saddle corresponds to the birth of a 1-cycle (which corresponds to \mathbb{S}^1 in $\mathbb{S}^1 \times \mathbb{S}^l$). This means that there is an extra point in the l -persistence diagram if $l = 1$. Most importantly this point can be distinguished by the fact that its birth time is much higher than all other points in the persistence diagram. However, because the births and deaths of B and \mathcal{M} are so intimately linked (except for the final death), we have the following: Consider b_l^1, b_l^j, b_l^k and $1 \neq j < k \neq 1$. Assume that b^1 is the earliest birth. If $l = 1$ further assume that $b_{l=1}^m$ is the final birth and $j \leq m \leq k-1$. If now moreover the death times are ordered as follows

$$\begin{aligned} \max\{D_l^I, D_l^{II}, \dots, D_l^{J-2}, D_l^{K+1}, D_l^{K+2}, \dots, D_l^N\} &< \min\{D_l^{J-1}, D_l^K\} \\ \max\{D_l^{J-1}, D_l^K\} &< \min\{D_l^J, \dots, D_l^{K-1}\}, \end{aligned} \quad (5.5)$$

then the l -cycle born at time b_l^j merges with the l -cycle born at time b_l^1 at time D_l^{J-1} and the l -cycle born at time b_l^k merges with the l -cycle born at time b_l^1 at time D_l^K for \mathcal{M} . In other words under these conditions there are points (b_l^j, D_l^{J-1}) and (b_l^k, D_l^K) in the persistence diagram.

- **STEP 4** As discussed in the intermezzo, the first birth is always coupled to the last death in the persistence diagram. We write $b_k^{j,c}(t)$ and $D_k^{j,c}(t)$ for the births and deaths respectively of $\rho(\cdot, \gamma(t))|_{\mathcal{M}}$ per connected component c , where we stick to the convention that $b_k^{1,c}(t)$ is the first birth (in k -homology) for each connected component c . We use similar notation for B . Here, we note that although this distance value is continuous, the Morse critical point where this minimum is attained is not continuous. A similar

effect was called a Faustian interchange in [Ste23b]. With this notation we can conclude that our first observation implies that for each c there is a point $(b_k^{1,c}(t), \max_J D_k^{J,c}(t))$ in the vineyard at level t and the vine consisting of these points if closed (by identifying the vineyard at times 0 and 2π) will yield a circle for each c . To put it differently this will lead to a surgery as depicted in Figure 5.11. Finally, we also note that if $l = 1$ there is an additional 1-cycle, that is born much later than all the other 1-cycles.

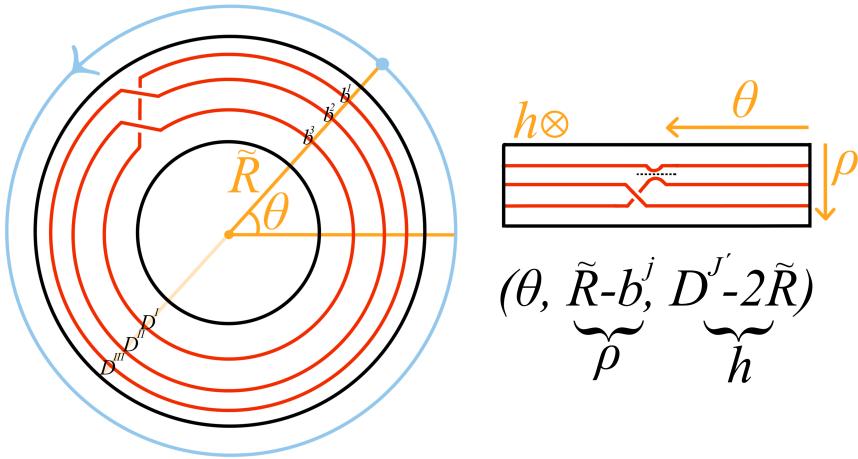


Figure 5.11: left: We see the particular closed braid that we called the ouroboros. Right: We see the θ and ω coordinates in the vineyard, where the h coordinate is in the direction orthogonal to the plane and \tilde{R} is the distance from the observation loop to the origin in the annulus. Here we identify ω with $\tilde{R} - b^j$ and h with $D^{J'} - 2\tilde{R}$, where $D^{J'}$ is the death time of the cycle born at b^j , corresponding to a strand as indicated in the figure. The elder rule induces surgery indicated with a dashed line (as depicted on the right). Note these are simplified sketches.

We further note that because vines are continuous (and even Lipschitz) thanks to [CSEH05], the only thing which we need to worry about is the crossings, because if all the birth values are distinct the birth values give precisely the ω coordinates in Figure 5.11, death times corresponding to the coordinate h do not matter.

We recall that there is only one crossing per (angular) interval $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ where $\tilde{j} \in \{8, 16, \dots, 8C\}$ and no crossings in $[\pi + \frac{\pi}{8(C+1)}, 2\pi - \frac{\pi}{8(C+1)}]$. Because the death times only matter for the crossing we can change the death times between crossings without creating topological problems. For a crossing in the interval $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ we dictate the death times by the geometry of the strands in the angular interval $[\pi \frac{\tilde{j}-1}{8(C+1)} + \pi, \pi \frac{\tilde{j}+2}{8(C+1)} + \pi]$. We do so by changing the $\tilde{\psi}$ coordinates of the strands in the interval, so that the death times for the interval $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ change, but the birth times in the interval $[\pi + \frac{\pi}{8(C+1)}, 2\pi - \frac{\pi}{8(C+1)}]$ remain the same, see Figure 5.12. This in particular ensures that we do not introduce (extra) crossings in the vineyard that are not present in the closed braid we start with.

Between these intervals we interpolate between the different geometries of strands (again by changing the $\tilde{\psi}$ coordinates), which we can do because as mentioned the death times do not influence the topology of the closure of the braid appearing in the persistence diagram.



Figure 5.12: Manipulating the $\tilde{\psi}$ changes the death times (vertical section of Figure 5.8, not to scale).

Now let us consider a crossing of two strands in the interval $[\pi \frac{j}{8(C+1)}, \pi \frac{j+1}{8(C+1)}]$ that are both not born first. As in the intermezzo, we denote with a little bit of abuse of notation by $b^{j,c}(t)$ and $b^{k,\tilde{c}}(t)$ (for B , or $b_0^{j,c}(t)$, $b_0^{k,\tilde{c}}(t)$, $b_l^{j,\tilde{c}}(t)$ and $b_l^{k,c}(t)$ respectively in the case of \mathcal{M}) both the birth times for these strands as well as the Morse critical points, where c, \tilde{c} denote different connected components. We further stress that as in the intermezzo the j and k indices in $b^{j,c}(t)$ and $b^{k,\tilde{c}}(t)$ refer to order along the closed braid not the order of insertion. In the case where $l = 1$ we further assume that neither j nor k corresponds to the birth with the very high birth value, that is the point where in non-extended persistence the second 1-cycle that lives forever is born.

We distinguish two different cases, one where $c = \tilde{c}$ and one where $c \neq \tilde{c}$. We start with the latter: We will focus on the (somewhat simpler) B case, as the \mathcal{M} case is virtually identical. Because the birth times correspond to Morse critical points on different connected components, changing the order does not influence the pairing between Morse critical points and birth and death in the persistence diagrams. This means that as long as the death times are distinct in $[\pi \frac{j}{8(C+1)}, \pi \frac{j+1}{8(C+1)}]$ and the cycle born at $b^{j,c}(t)$ dies before the one born at $b^{k,\tilde{c}}(t)$ for one t in this interval, then the cycle born at $b^{j,c}(t)$ dies before the one born at $b^{k,\tilde{c}}(t)$ for all t in this interval and at continuous and distinct death times. Let us write $D^{J',c}(t)$ and $D^{K',\tilde{c}}(t)$ for the death times. This means that the persistence diagram (at level t in the vineyard) contains the points $(b^{j,c}(t), D^{J',c}(t))$ and $(b^{k,\tilde{c}}(t), D^{K',\tilde{c}}(t))$, each of which locally describe a vine. This in turn implies that if (locally) $D^{J',c}(t) > D^{K',\tilde{c}}(t)$ then the vine $(b^{j,c}(t), D^{J',c}(t))$ in the vineyard crosses under the vine $(b^{k,\tilde{c}}(t), D^{K',\tilde{c}}(t))$ and the reversed order of death corresponds to an under crossing. This is achieved (as mentioned) by manipulating the $\tilde{\psi}$ coordinate. An easy way to achieve this if the component c passes under \tilde{c} is to push all strands of c in and all strands of \tilde{c} out and the reverse for an over pass. If one would like to repeat this discussion for \mathcal{M} one only need to add a lower order 0 or l respectively.

We note that in the case where $l = 1$, the birth with the very high birth value, that is the point where in non-extended persistence the second 1-cycle that lives forever is born and was denoted by m in the intermezzo, is separate (because of the high birth value) and therefore gives a disconnected circle in the closure of the vineyard.

Now we consider the case that $c = \tilde{c}$. To simplify notation, we will drop the index c from the notation altogether. We again focus on the B case, the case \mathcal{M} is almost identical. If the death times satisfy (5.4) for each t in $[\pi \frac{j}{8(C+1)}, \pi \frac{j+1}{8(C+1)}]$, then there are points $(b^j(t), D^{J-1}(t))$ and $(b^k(t), D^K(t))$ in the persistence diagram. We can ensure that (5.4) holds by changing the $\tilde{\psi}$ coordinates as before, see Figure 5.12. Because the assumption

(5.4) does not constrain the relative order of $b^j(t)$ and $b^k(t)$ nor the relative order of $D^{J-1}(t)$ and $D^K(t)$. This means that we can change the order of birth time (which occurs thanks to the crossing in the closed braid B) during the course of the interval $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ and we can fix the relative order of $D^{J-1}(t)$ and $D^K(t)$ as needed (which we can again do by manipulating $\tilde{\psi}$, see Figure 5.12). If $D^{J-1}(t) > D^K(t)$ in $[\pi \frac{\tilde{j}}{8(C+1)}, \pi \frac{\tilde{j}+1}{8(C+1)}]$ then the b_j vine crosses over the b_k vine, while for the reversed order the b_j vine crosses under the b_k vine. We refer to the appendix for an extensive example of this procedure in the case of the braid depicted in Figure 5.8. This means that regardless of whether we have an over or under crossing, we locally push the strands with the critical points $D^I, D^{II}, \dots, D^{J-2}, D^{K+1}, D^{K+2}, \dots, D^N$ in (towards the centre of the annulus) by a lot, and the strands with critical point D^J, \dots, D^{K-1} out (from the centre of the annulus) by a lot. This leaves the strands with D^{J-1} and D^K in the middle, and one may push the strand with D^{J-1} a little out and the D^K a little in if we want the b^j vine segment to pass over the b^k part of the vine, with the reverse pushing for the under crossing. This in particular shows that the braid B can be faithfully reconstructed in the persistence diagram, although we also introduce some one extra loop per connected component of the link if $l \neq 1$ and two if $l = 1$.

The only thing left to remark is that the conditions of Lemma 8 and Corollary 2 can always be satisfied thanks to Remark 3. \square

5.3.4 Vineyard Monodromy

In this section we give the second main result of this chapter, namely that every order of monodromy can be found in a vineyard and in any dimension l of \mathbf{H}_l , with $l \leq d-2$. This result follows almost immediately from Theorem 11, and answers one of the main open questions of Arya et al. [AGH⁺24], albeit in a slightly different context.

Proof of Theorem 10. The result follows by applying Theorem 11 to the (l -offset) of the ouroborus knot, that is the closure of the braid with $k+1$ strands and k over crossings, as depicted in Figure 5.11. \square

5.4 Avoiding Monodromy in \mathbb{R}^2

Recall that the authors of [AGH⁺24] were able to demonstrate that so-called star shaped objects cannot create monodromy in the persistent homology transform. In this section we aim to show sufficient conditions on the observation loop in \mathbb{R}^2 so as to avoid the generation of monodromy in the vineyard of the restricted distance function. Our conditions are based entirely on the geometry of the objects as described by their extended symmetry sets.

5.4.1 Persistence and the Extended Symmetry set

Consider a simple closed curve $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ and a point $p = \gamma(t)$. The distance function $\rho_x : \gamma \rightarrow \mathbb{R}$ has critical points where the radius and tangent of γ are perpendicular, that is $\langle \gamma'(t), \gamma(t) - x \rangle = 0$. Let $e \in \mathbb{R}^2$ be a point on the evolute of γ , that is, e is centre of the osculating circle of γ at p and a degenerate critical point of ρ_x . Additionally, let $e_{\text{in}}, e_{\text{out}}$ be two points in a small neighbourhood of e which lie on opposite sides of the evolute. At e_{in}

the degenerate critical points splits into two disjoint critical points, while conversely at e_{out} the degenerate critical point resolves into a single critical point. Hence crossing the evolute results in the birth or death of a persistence point depending on the direction of traversal. The converse statement also holds as in order to generate a new persistence point you must pass through a degenerate radius function which corresponds exactly to points of the evolute. Indeed, the evolute partitions the plane into regions where the combinatorics of the persistence diagram of the radial distance function are fixed.

Additionally, we can further subdivide the plane using the symmetry set (recall from Section 3.3 the union of the symmetry set and the evolute is known as the extended symmetry set). Each point in the symmetry set is generated by a circle that touches the curve in at least two points and leaves of the symmetry set terminate in A_3 contacts (extrema of curvature in γ and cusps of the evolute). The nature of these contacts, whether the contact point on the curve has greater or lesser curvature compared to the contact sphere, can be used to partition the manifold like pieces of the symmetry set by their impact on the radial distance persistence diagram. When both contact points have less curvature than the curve it implies that a radial distance persistence diagram from this point will produce two points with equal death coordinate. Said differently, crossing such a piece of the symmetry set will exchange the ordering of the births of two pieces in the persistence diagram. Similarly, if both contact points have lesser curvature the ordering of death coordinates will be exchanged, and if the contact points have different types the ordering of a birth and death will be exchanged.

From this we can conclude the following.

Lemma 10 (Preventing monodromy). *Under the conditions described above if the observation loop does not enclose a singularity of the extended symmetry set, it cannot produce monodromy.*

Proof. We first assume the observation loop is contained completely within a cell of the extended symmetry set. Further assume for the sake of contradiction that monodromy exists for such an observation loop. For monodromy to exist after one full rotation of the observation loop a permutation σ is applied to the points in the persistence diagram. Consider the difference between birth values of distinct points in the persistence diagram $h_{i,j} := b_i(t) - b_j(t)$, where t is the observation loop parameter, which is a continuous function. Since we do not cross the symmetry set birth values never coincide so the relative order of the births remains the same and $h_{i,j} \neq 0$ for all i, j . However, if σ is not the identity permutation then there are persistence points such that $b_i(2\pi) = b_{\sigma(i)}(0)$ and $b_j(2\pi) = b_{\sigma(j)}(0)$, with $i \neq j$, such that the order of b_i and b_j at $t = 0$ and $t = 2\pi$ differs. This implies that $h_{i,j}(0)$ and $h_{i,j}(2\pi)$ have opposite signs and therefore by the Intermediate Value Theorem that there must exist $t \in (0, 2\pi)$ with $h_{i,j}(t) = 0$. This contradicts the uniqueness of birth values within a cell of the extended symmetry set and with an identical argument for the death values implies that no nontrivial permutation and hence no monodromy can occur.

Now we allow the observation loop to pass through pieces of the extended symmetry set but restrict it from enclosing singularities. We examine the combinatorial properties of the intersection of the loop and the extended symmetry set. A loop that does not enclose a singularity of the extended symmetry set intersects each piece of the symmetry an even number of times, and the order of intersecting faces creates a well formed expression. In so doing the exchanges in orderings of births and deaths (and creation and deletion of points when crossing the evolute) are cancelled out starting from the centre of the palindrome such that the only possible permutation is the identity.

□

It remains to prove which type or combination of types of singularities in the extended symmetry set which when enclosed by an observation loop result in monodromy. However, we were able to find the curve depicted in Figure 5.14, which exhibits monodromy without a particularly complicated singularity structure. This example should be a relatively simple case to investigate in future work in this direction.

5.5 Discussion

We conclude by reflecting that the construction in this chapter implies that vineyards are as topologically complex as one may hope or fear, opening up a number of interesting directions for future work. We note that as stated in [OOT24], the radial transform presents some significant advantages over the more well-studied directional transform, and we feel it warrants greater study. In particular, it is quite critical in our construction, and allowed for better understanding of monodromy and braiding in higher dimensions.

Theorem 11 implies that comparing closed vineyards with a measure that completely reflects the topology is likely to be a difficult problem, computationally speaking. By this we mean the following: Suppose that we want to compare periodic evolving phenomena using persistence, and thus are faced with closed vineyards. Now suppose that we want a distance between two closed vineyards that respects the topology. As these vineyards can contain arbitrary knots and links embedded in them, solving this problem would necessarily mean recognizing those structures within. While the exact complexity of knot recognition is formally still open, in practice it has proven to be quite difficult, as it sits at the intersection of NP [INT23, Has97] and co-NP [Lac21] and is connected to several known hardness results [KT21]. In particular, it is known to be NP-hard to test unknottedness if the number of simplifying Reidemeister moves allowed while unknotting is specified as a part of the input [dMRST21].

That said, the richness of structure present in vineyards may still allow for subtle distinctions and invariants to be calculated. In fact, our results in some sense show that vineyards are more than the sum of parts, as the choice of base loop in our example is in fact critical for the resulting vineyard to have any braiding, and the resulting vineyard can exhibit quite complex behavior. However, these subtleties come at a high computational cost, and finding a compromise between topological fidelity and computational complexity will be an important challenge for future work.

We conclude with a few specific open questions:

- Can we describe monodromy simply by the ordering and type of the pieces of the symmetry set that are crossed or by the type of the singularities enclosed by an observation loop?
- Are there good combinations of link invariants and distances, like the Wasserstein distance, that still capture a lot of the topology, provide a practical similarity measure, and are relatively easy to compute? If so, could these invariants prove useful in topological data analysis, given the rising use of the persistent homology transform in real world applications?
- Are there geometric conditions (that are easy to verify) on a pair of closed vineyards that imply that the closed vineyards are isomorphic?

- Could tools from topological data analysis, and in particular statistical work on vineyards [MTB⁺15], provide new practical insights into knot and link recognition?
- On the application side, it would be interesting to ask: Do even simple non-trivial knots or links (like the trefoil knot, Hopf link, or Borromean rings) exist in persistence diagrams of real data? There is some indication that knots are present in proteins [Let12], otherwise we would suggest looking in persistence diagrams of biological systems containing circular DNA [NCMRA24].

5.A Example

As an example of pushing outward and inward on the strands by manipulating the $\tilde{\psi}$ coordinates in Step 4 of the proof of Theorem 11 we discuss the example shown in Figure 5.8 in detail. We number the crossings c_1, \dots, c_{13} as indicated in Figure 5.15. The manipulation proceeds as below. We emphasize that all of the death values associated to points near a crossing need to be generic, that is no two identical death values. We will not repeat this for every crossing.

- At c_1 yellow crosses under red, hence we push all red strands out and all yellow ones in. We will refer to this as yellow before red.
- At c_2 the strand on which b^2 lies crosses under the strand with b^3 , in accordance with (5.4) we push D^{III} in and D^{II} out so that D^I lies in between (blue in the figure), what you do with the yellow strands doesn't matter (as long as it is generic). The way that we push is also indicated in blue in the figure.
- At c_3 red crosses under yellow, hence we push the yellow strands out and the red ones in, i.e. red before yellow.
- At c_4 red crosses over yellow, hence push the red strands out and the yellow ones in, i.e. yellow before red.
- Similarly to the crossing at c_2 (but at c_5 , the strand with b^2 crosses under the strand b^3), at c_5 we push D^I in and push D^{II} out so that D^{III} (green in the figure) lies in the middle (there is no condition on the yellow strands except genericity).
- At c_6 red goes before yellow.
- The crossings c_7 , c_8 , and c_9 almost coincide in the figure, however, the red crosses in all cases over yellow, so that yellow needs to go before red², because the first birth in red exchanges strand the coupling is automatic and leads to a disconnected component (i.e. surgery is performed).
- At c_{10} yellow goes before red.
- At c_{11} the order of the strands really doesn't matter as long as all death values are distinct (generic), because this is another first birth interchange (on yellow this time).

²Strictly speaking it is not necessary that yellow needs to go before red as long as you choose consistently for both crossings because at the red crossing the first birth changes from one strand to another, which leads to surgery so that the outer strand disconnects, which means that after a Reidemeister II move you are fine. However to fit with the text in the proof it is best to have yellow before red.

- The crossing at c_{12} is again similar to the crossings at c_5 with the strand containing b^2 crossing under b^3 . We stress that because we label the critical points in their order along the braid, and the first birth has exchanged strands, the strand that contains b^2 is not the same strand that contained b^2 at c_5 (where by the same we mean identification via shortest paths on the link). The strands that are associated to D^I , D^{II} , and D^{III} change as a consequence as well, see Figure 5.15 (purple). This having been said, we follow the same procedure as at c_5 : At c_{12} push D^{III} out and D^{II} in, so that D^I ends up in the middle (there is no condition on the yellow strands except genericity), indicated in purple in the figure.
- At c_{13} red goes before yellow.

5.B Monodromy touching the diagonal

If we drop the assumption that the vines do not touch the diagonal, but still assume that none of the persistence diagrams contain points with higher multiplicity, it is possible to formulate a notion of monodromy that is more equivalent to that contained in prior work [AGH⁺24]. To make sense of the definition given below and to provide an alternative formulation to the one in the body of this chapter, we need the following observation: Let $\tilde{\mathcal{V}}_{\max}(t)$ be a vine defined on its maximal domain (t_{\min}, t_{\max}) , where $t_{\min}, t_{\max} \in \mathbb{R} \cup \{\pm\infty\}$, that contains the point $\mathcal{V}(0)$.

More formally $\tilde{\mathcal{V}}_{\max}(t)$ is the restriction to the maximal open interval such that the lift of the curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}/2\pi\mathbb{Z} : t \mapsto t \bmod 2\pi$, contains no limit points on the diagonal for this open interval.

If we are given a vine $\mathcal{V}(t)$ defined on $[0, 2\pi]$ we call the vine $\tilde{\mathcal{V}}_{\max}(t)$ defined on its maximal domain that coincides with $\mathcal{V}(t)$ on the interval $[0, 2\pi]$, the maximal extension of $\mathcal{V}(t)$. If the domain of a maximal extension is finite, then we say that the vine has a finite maximal extension.

If the number of points in each persistence diagram is finite (as we assume), then either both t_{\min} and t_{\max} are (plus/minus) infinite or neither of them are. This is clear because if one of them is infinite we must return to the same point in the persistence diagram after some time $2\pi k$ (since the number of points in the persistence diagram is finite), in which case $\tilde{\mathcal{V}}(t)$ is periodic and is defined on $(-\infty, \infty)$.

Using this definition we may reformulate our definition of the order of the monodromy from Section 5.3:

Remark 4. *Assume that there are no points of higher multiplicity in $\mathcal{D}_l(\rho_t(p))$, for all $\gamma(t)$ and by extension the vines are non-intersecting curves. Assume moreover that all of the vines are disjoint from the diagonal. The vine $\mathcal{V}(t)$ exhibits monodromy of order k iff k is the smallest positive integer such that $\tilde{\mathcal{V}}_{\max}(0) = \tilde{\mathcal{V}}_{\max}(2\pi k)$, where $\tilde{\mathcal{V}}_{\max}(t)$ is the maximal extension of $\mathcal{V}(t)$.*

In this setting, that is where we allow vines to contain points in the diagonal (as limit points), and the goal again is to glue together in a way that we have the exhibited periodicity of monodromy as an integer. More formally:

Assume that there are no points of multiplicity higher than 1 in $\mathcal{D}_l(\rho_t(p))$, for all $\gamma(t)$ and by extension the vines are non-intersecting curves. We will follow the same notation as above.

Individual vine monodromy: Let $\mathcal{V}(t)$ be the vine and $\tilde{\mathcal{V}}_{\max}(t)$ is its maximal extension. If its maximal domain is $(-\infty, \infty)$, then the vine does not touch the diagonal and the definition from Section 5.3 and the reformulation in Remark 4 both apply. Let us now assume that t_{\min}, t_{\max} are finite. By reparametrizing we can assume (without loss of generality) that $t_{\min} = 0$. The order of monodromy is now defined as the smallest positive integer k , such that $t_{\max} < 2\pi k$. As before, we will say that the monodromy is trivial if $k = 1$.

We will also define the completion along the diagonal of a vine \mathcal{V} . Let $\tilde{\mathcal{V}}_{\max} : (t_{\min}, t_{\max}) \rightarrow \mathcal{D}_l(\rho_t(p))$ be the maximal extension of \mathcal{V} , and write

$$v_{\min} = \lim_{t \searrow t_{\min}} \tilde{\mathcal{V}}_{\max}(t)$$

$$v_{\max} = \lim_{t \nearrow t_{\max}} \tilde{\mathcal{V}}_{\max}(t)$$

for the two limit points of the vine on the diagonal. Let k be the order of monodromy of \mathcal{V} we define the completing diagonal vine \mathcal{V}_D as

$$\mathcal{V}_D : (t_{\max}, t_{\min} + 2\pi k) \rightarrow \mathcal{D}_l(\rho_t(p))$$

$$t \mapsto \left(t, v_{\max} \left(1 - \frac{t - t_{\max}}{t_{\min} - t_{\max} + 2\pi k} \right) + v_{\min} \frac{t - t_{\max}}{t_{\min} - t_{\max} + 2\pi k} \right).$$

The start and end points of the concatenation $\mathcal{V}_D \circ \tilde{\mathcal{V}}_{\max}$ coincide and therefore by identifying t_{\min} and $t_{\min} + 2\pi k$ we can consider this to be a map on $\mathbb{R}/2\pi k\mathbb{Z}$. The set $\mathbb{R}/2\pi k\mathbb{Z}$ can be viewed as k cover of the circle $\mathbb{R}/2\pi\mathbb{Z}$, where we identify $\mathbb{R}/2\pi\mathbb{Z}$ with the loop γ . Composing with this cover map gives a map $\mathcal{V}_C : \gamma(t) \mapsto \mathcal{D}_l(\rho_t(p))$, which exhibits monodromy in the way we defined above. We call \mathcal{V}_C the completion of the vine \mathcal{V} .

Vineyard monodromy: We now call the vines whose maximal extension do not intersect the diagonal (or equivalently those whose maximal domain is $(-\infty, \infty)$) non-rooted vines, while we call those that intersect the diagonal (or equivalently those whose maximal domain has finite length) rooted vines. We write k_1, \dots, k_n for the orders of monodromy of non-rooted vines and write l_1, \dots, l_m for the orders of monodromy of the rooted vines. Let $l_{\max} = \max\{l_1, \dots, l_m\}$, then we define the order of monodromy of the vineyard as the smallest k that is a common multiple of k_1, \dots, k_n and is larger than l_{\max} . We also call k the common order of monodromy of all the vines.

We can complete the vineyard in the same way as before, but with k the common order of monodromy of all the vines, instead of an individual vine. We call the result the completed vineyard.

We stress that there may be self-intersections on the diagonal of the completed vineyard. We note again here that if the vineyards are non-generic, the situation is significantly more complex as vines can intersect [Tur23]. We have not considered the non-generic case in this work, however it may be of interest for future study.

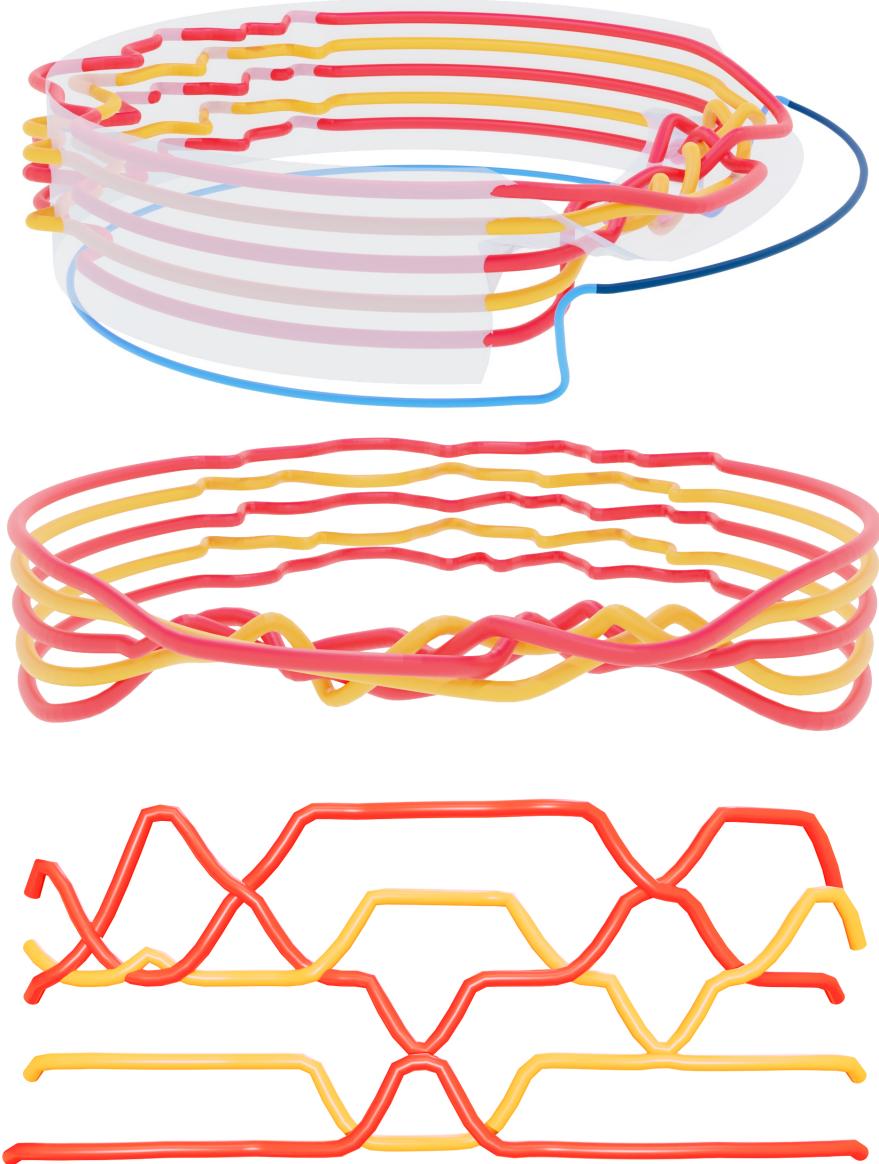


Figure 5.13: Top: a side view of a 3-dimensional embedding of the link depicted in Figure 5.8, where the crossings and strands are perturbed in a particular way to get the desired monodromy (see the description in Figure 5.8). The red and gold portions are the two sections of the link, and the blue curve is an observation loop for the radial transform, where the dark blue portion corresponds to the section of vineyard visualized in the bottom. See also the proof of Theorem 10 and the corresponding Example in Appendix 5.A for the full details of the perturbation. Middle: Front-angled view of the embedded link. Bottom: A sideways view of our computed vineyard, showing the diagrams computed from the radial transform of observation points taken from the fraction of the full period (from 0 to 2π) depicted in dark blue segment of curve above. This segment of the vineyard captures all crossings and exhibits monodromy of period $2\pi \cdot 3$. To transform this braided vineyard into a closed braid, we would identify the “sides”, or slices, at 0 and 2π .

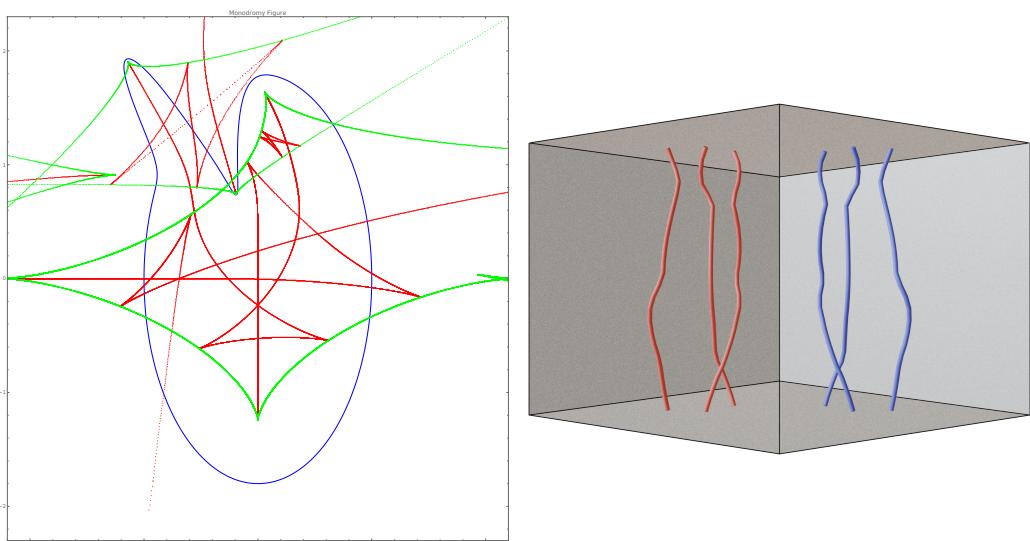


Figure 5.14: Minimal example of monodromy in the plane. Left: Curve (blue), evolute (green), and symmetry set (red). Right: Corresponding extended persistence vineyard exhibiting monodromy in both H_0 and H_1 demonstrating a non-trivial permutation.

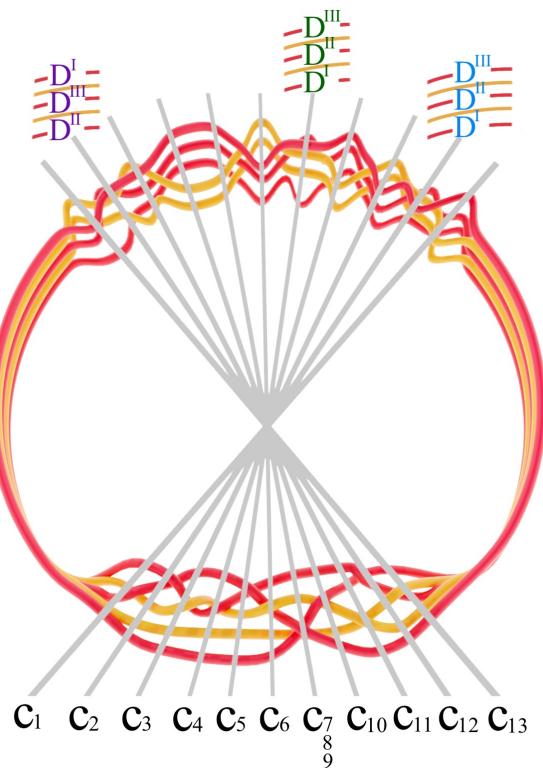


Figure 5.15: The crossings and the specific correspondence between the deaths or critical points and the strands indicated.

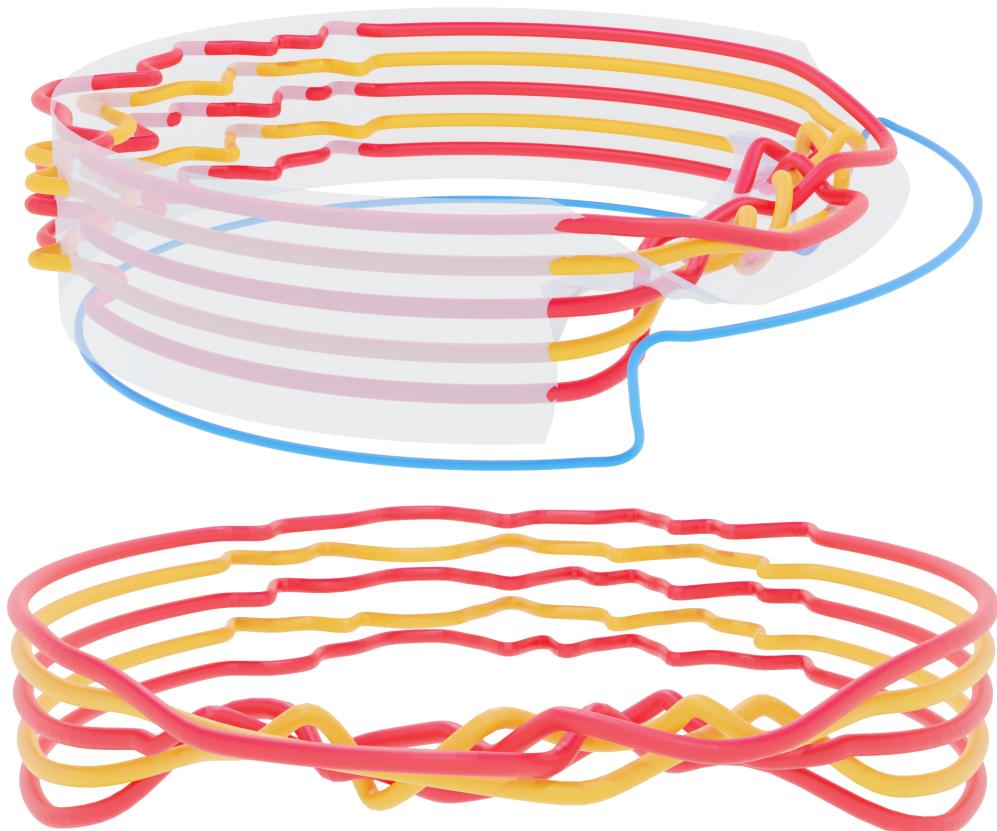


Figure 5.16: Different views of the Example. Top: together with observation loop (blue) and twisted annulus (transparent gray). Bottom: front view of just the manipulated trefoil-circle link.

CHAPTER 6

Counting Equilibria of the Electrostatic Potential

6.1 Introduction

In 1873, James C. Maxwell [Max73] conjectured that the electric field generated by n point charges in generic position has at most $(n - 1)^2$ isolated zeroes.

The first (non-optimal) upper bound was only obtained in 2007 by Gabrielov, Novikov and Shapiro, who also posed two additional interesting conjectures [GNS07]. This upper bound was improved upon in [EFO25]. Prior to this work however, little was known about arrangements of charges that could approach even Maxwell's conjectured upper bounds.

In this chapter, inspired by [GNS07] and based upon our article [EFO25], we explore examples and construct configurations of charges achieving the highest ratios of the number of electric field zeroes to point charges found to this day. These configurations fall drastically short of the best known upper bounds and even that of Maxwell's conjecture lending credence to the possibility that a quadratic upper bound is not tight.

Additionally, we will explore the computational and explicit methods for finding these equilibria.

Finally, we construct a counterexample to Conjecture 1.8 in [GNS07] that the number of equilibria cannot exceed those of the distance function defined by the unit point charges.

6.1.1 The Electrostatic Potential

The electrostatic potential generated by n point charges located at $A_1, A_2, \dots, A_n \in \mathbb{R}^3$ of magnitudes $\zeta_1, \zeta_2, \dots, \zeta_n \in \mathbb{R}$ is the function $V: \mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\} \rightarrow \mathbb{R}$ defined by

$$V(x) = \sum_{i=1}^n \frac{\zeta_i}{\|x - A_i\|}, \quad (6.1)$$

in which $\|x - A_i\|$ is the Euclidean distance between the two points. A one dimensional slice of the electrostatic potential for a single unit charge is depicted in Figure 6.1. The electric field is the gradient of this potential.

In this chapter, we only consider the case in which all the charges ζ_i are unit charges. The zeros of the electric field are the equilibria or the critical points of the potential. They are also

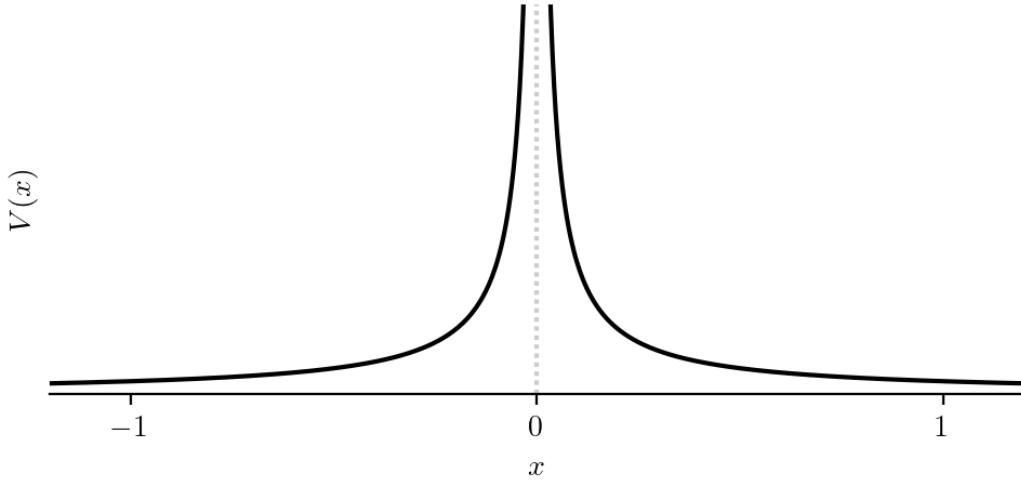


Figure 6.1: One dimensional section of the electrostatic potential of a single charge.

sometimes referred to as electrostatic points as they are the points at which the electric field vanishes, thus corresponding to equilibrium positions for test charges.

We recall briefly some essentials of Morse theory necessary for our analysis (see Section 2.5 for a more complete summary). An equilibrium of a smooth function is non-degenerate if the Hessian at the point is invertible. A useful result of Morse–Cairns in [MC69] yields that for the generic location of a single point charge, all equilibria of V are non-degenerate. In this case, the Hessian has 0, 1, 2, or 3 negative eigenvalues, corresponding to the number of directions of principal curvature in which the function is decreasing. In Morse theory, this number is called the index; see e.g. [Mil69]. Equilibria of index 0, 1, 2, 3 are referred to as minima, 1-saddles, 2-saddles, and maxima respectively.

Importantly, the electrostatic potential in \mathbb{R}^3 is harmonic; that is: its Laplacian vanishes at every point. To explain, it is not difficult to check that the function $x \mapsto 1/\|x\|$ is harmonic on $\mathbb{R}^3 \setminus \{0\}$. Sums, translates, and multiples of harmonic functions are again harmonic, which implies that V is harmonic. It follows that V enjoys the mean value property: $V(x)$ is the average of the values of V on any sphere centered at x that neither passes through nor encloses any of the A_i . Hence, if nonzero, V has neither minima nor maxima in $\mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\}$.

Although the electrostatic potential itself is not Morse, since $\mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\}$ is not compact, the equilibria are in correspondence with those of a Morse function. Namely, we can consider the function $f : \mathbb{R}^3 \cup \{\infty\} \rightarrow \mathbb{R}$ which agrees with V except in small neighbourhoods of each charge where we cap the electrostatic potential with maxima and extend V by continuity at infinity by setting $f(\infty) = 0$ which is a minima (see Figure 6.2). If the neighbourhoods are chosen small enough this has no impact on the existing 1 and 2-saddles of V . Additionally, any degenerate equilibria of V will disappear under a slight perturbation of the charges, again without influencing the existing equilibria.

An early application of Morse theory, carried out by Morse himself [MC69], gives a lower bound on the number of equilibria of V .

Let m_p be the number of index- p equilibria of M . After subtracting the one minimum at infinity and the n maxima at A_1 to A_n , we get $m_0 + m_1 + m_2 + m_3 - (n + 1)$ as the number of equilibria of V . By the Euler–Poincaré relation, we have $m_0 - m_1 + m_2 - m_3 = 0$ because the Euler characteristic of $\mathbb{S}^3 \simeq \mathbb{R}^3 \cup \{\infty\}$ vanishes. By construction, $m_3 = n$ and $m_0 = 1$,

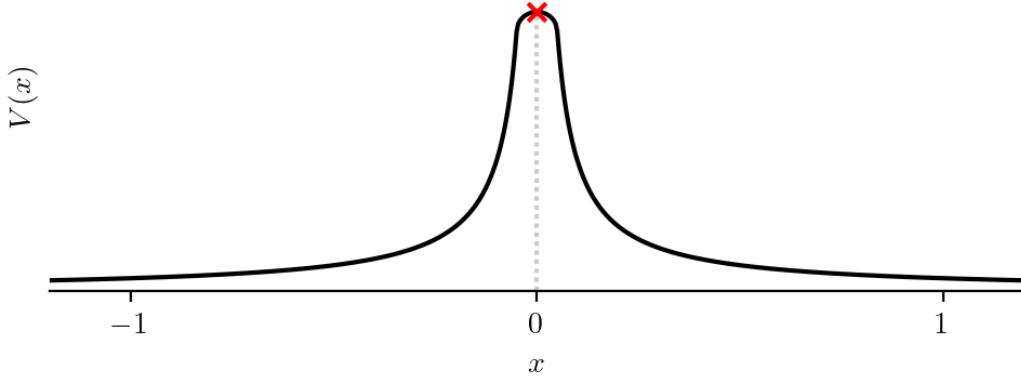


Figure 6.2: One dimensional section of the electrostatic potential of a single charge capped so as to make it Morse.

which thus implies $m_2 - m_1 = n - 1$. Hence, the number of equilibria of V is

$$m_1 + m_2 = 2m_1 + n - 1 \geq n - 1. \quad (6.2)$$

This lower bound is tight since placing n point charges along a straight line in \mathbb{R}^3 defines a potential with only $n - 1$ equilibria, all of which are 2-saddles, see Figure 6.3.

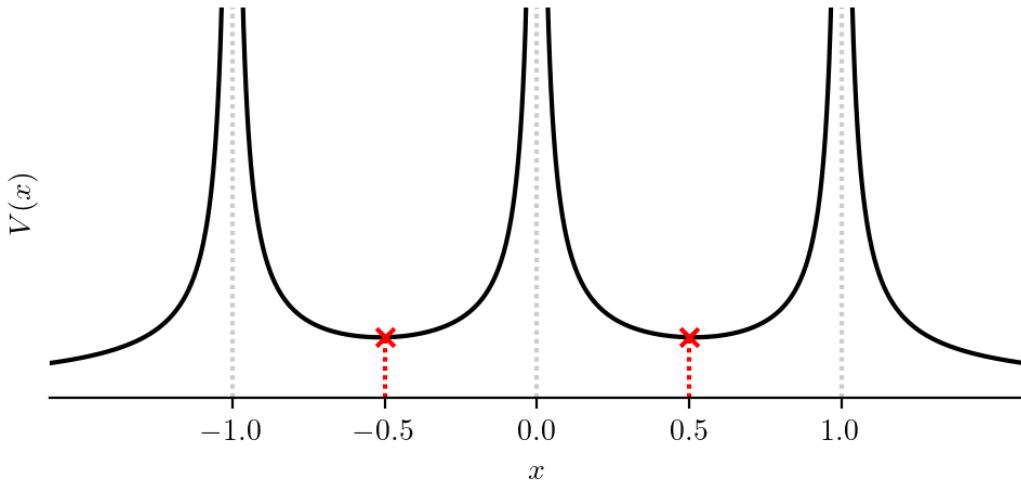


Figure 6.3: One dimensional section of the electrostatic potential of three charges along a line.

There are however configurations with more than $n - 1$ equilibria, and as mentioned in the introduction to this chapter, the quadratic upper bound $(n - 1)^2$ was conjectured by Maxwell [Max73]. This bound is realised for $n = 3$, for example when placing charges at the vertices of an equilateral triangle, see Figure 6.4.

However, configurations that achieve this bound are not known for $n > 3$, and numerical experiments suggest that for randomly placed charges the ratio of the number of equilibria over point charges is bounded and rather low.

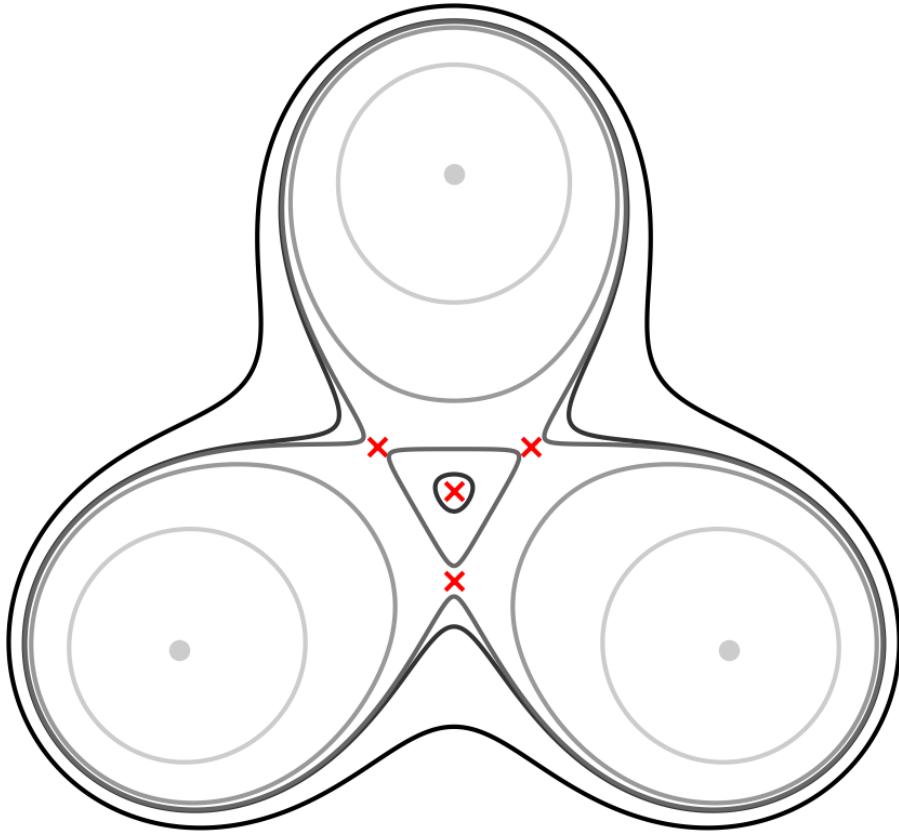


Figure 6.4: Two dimensional section of the electrostatic potential of three charges arranged as the vertices of an equilateral triangle. Lighter contours indicate higher values of the potential. Equilibria are indicated by red crosses.

The major goal in this chapter, which is carried out in Section 6.3, is therefore to study special configurations of point charges, such as those obtained by placing them at the vertices of certain solids.

6.2 Method

Finding critical points of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is a simple problem familiar to high school mathematics students. The electrostatic potential, in principle, is no more difficult. We need to find the points $x \in \mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\}$ where the electric field vanishes. However, one will quickly realise the complexity of solving this problem explicitly when considering more than a handful of charges. We are therefore forced to experiment with computational techniques.

The simplest computational approach to this problem would be to compute the squared norm of the electric field, $\|\nabla V\|^2$, in a grid of sample points around the charges and then count the number of connected components as you remove grid elements above a decreasing threshold. However, our experimentation with solids revealed that some equilibria are highly degenerate,

meaning there are many equilibria in a small volume and are extremely shallow, by which we mean that the threshold required to separate nearby equilibria can be smaller than what is practical.

6.2.1 Cubical Persistence

To remedy the degeneracy and depth of the equilibria a more topological approach was taken to the problem. As outlined in Section 2.5, the critical points of a Morse function, $f : \mathcal{M} \rightarrow \mathbb{R}$, occur precisely at the points where topological changes in the level sets occur. Additionally, from Section 2.7 the number of critical points is twice the number of persistence pairs in the filtration by f and critical edges and critical faces correspond to 1 and 2-saddles respectively.

We followed the efficient algorithm for computing persistence of cubical data outlined in [WCV12] and implemented in [Ču20]. When interrogating topological changes depicted, such as in Figure 6.13, we used the marching cubes algorithm [LC87] as implemented in Python's sci-kit image package [VdWSNI⁺14].

While persistence is still vulnerable to the degeneracy of the electrostatic potential, the topological nature of the technique provides us good evidence when further scrutiny of an identified equilibria is required.

6.2.2 Local Homology

Once a potential equilibria is identified from the persistence of V we use the local homology at a point to determine whether it is in fact an equilibrium or not, and if so, the level of degeneracy it may have.

For $x \in \mathbb{R}^3 \setminus \{A_0, A_1, \dots, A_n\}$, let $\varepsilon > 0$ be smaller than the distance between x and any of the n point charges, and write $B(x, \varepsilon)$ and $S(x, \varepsilon)$ for the closed ball and sphere with center x and radius ε .

The local homology of V at x is the limit, for ε going to zero, of the relative homology of the pair $(B^-(x, \varepsilon), S^-(x, \varepsilon))$, in which $B^-(x, \varepsilon) \subseteq B(x, \varepsilon)$ and $S^-(x, \varepsilon) \subseteq S(x, \varepsilon)$ are the subsets of points y with $V(y) \leq V(x)$.

To visually represent these groups, we define a binary function on the unit sphere, which maps $u \in \mathbb{S}^2$ to black if $V(x) < V(x + \varepsilon u)$, and to white if $V(x) \geq V(x + \varepsilon u)$ for every sufficiently small $\varepsilon > 0$.

In words: the potential increases if we leave x in a black direction; and it decreases or stays the same if we leave x in a white direction. See Figure 6.5 for the binary functions of a non-critical point and the four types of non-degenerate equilibria and Table 6.1 for the corresponding ranks in the local homology. Writing $W \subseteq \mathbb{S}^2$ for the white points, the local homology of x is the homology of the cone over W relative to W .

Observe that every geometric realization of the solid yields the same local homology. Indeed, a similarity, which is a map $\theta : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that is the composition of: a scaling; a rotation; a translation; and possibly a reflection, preserves the types of equilibria. To prove this claim, we write $x' = \theta(x)$ and $V' : \mathbb{R}^3 \setminus \{A'_1, A'_2, \dots, A'_n\} \rightarrow \mathbb{R}$ for the potential defined by the point A'_i with charges $\zeta'_i = \zeta_i$.

Lemma 11. *Let V, V' be the potentials defined by a finite set of point charges, before and after applying a similarity, respectively. Then there is a bijection between the equilibria of V and V' such that corresponding equilibria have the same type.*



Figure 6.5: From left to right: the binary functions on the unit 2-sphere for a non-critical point, a minimum, a 1-saddle, a 2-saddle, and a maximum.

	$p = 0$	1	2	3
non-critical	0	0	0	0
minimum	1	0	0	0
1-saddle	0	1	0	0
2-saddle	0	0	1	0
maximum	0	0	0	1

Table 6.1: The ranks in local homology at the points whose binary functions are illustrated in Figure 6.5.

Proof. A rotation, translation, and reflection does not change the value at a point, except that it transfers this value to the image of that point. This is not true for a scaling. Letting $s > 0$ be the scaling factor of the similarity, we have

$$V'(x') = \sum_i \frac{\zeta'_i}{\|x' - A'_i\|} = \sum_i \frac{\zeta_i}{s\|x - A_i\|} = \frac{1}{s}V(x), \quad (6.3)$$

at every $x \in \mathbb{R}^3$. Since multiplication with a positive constant preserves the local homology at a point, this establishes the claim for the bijection that maps an equilibrium to its image under the similarity. \square

As we will see in the following section the ranks of local homology of degenerate equilibria of regular and semi-regular solids can be quite large so the local homology is quite useful in order to quantify degeneracy.

6.3 Regular and Semi-Regular configurations

We consider the vertices of the Platonic solids and semi-regular polyhedra as potentially interesting point charge configurations. We cover the families of these solids individually in the follow subsections.

The numbers of equilibria given in Tables 6.3, 6.4 and 6.5 were obtained by first computing cubical persistence and local homology, further investigating graphically with level sets around critical values and finally, where possible solving explicitly the location of equilibria. We discovered that all equilibria are non-degenerate, except at the center of the solid, which unfolds into two or more such equilibria.

6.3.1 Platonic Solids

A Platonic solid is a convex polytope such that

- all facets are the same regular polygon;
- all vertices have the same local shape.

There are 5 Platonic solids, see Figure 6.6.



Figure 6.6: From left to right: Platonic Solids following the order in Table 6.2.



Figure 6.7: From left to right: the binary functions for the centers of the tetrahedron, cube, octahedron, dodecahedron, and icosahedron (these are degenerate equilibria and so are neither 1-saddles nor 2-saddles).

The rank vectors of each platonic solid is listed in Table 6.7. As an example, the tetrahedron has a rank vector of $(0, 0, 3, 0)$, which is consistent with three 2-saddles co-located at its center. It is also consistent with one minimum, four 1-saddles, and six 2-saddles co-located at the center, where the four 1-saddles cancel with the minimum and three of the 2-saddles. The latter choice is suggested by the face structure of the tetrahedron, which has one tetrahedron, four triangles, and six edges.

	$p = 0$	1	2	3
tetrahedron	0	0	3	0
cube	0	5	0	0
octahedron	0	0	5	0
dodecahedron	0	11	0	0
icosahedron	0	0	11	0

Table 6.2: The ranks in local homology at the centers of the Platonic solids, with binary functions shown in the lower row of Figure 6.7. Observe that the last column in Table 6.3 is the alternating sum of these ranks.

solid	f -vector	2-saddles	1-saddles	center
Tetrahedron	(4,6,4,1)	0	0	3
Cube	(8,12,6,1)	12	0	-5
Octahedron	(6,12,8,1)	0	0	5
Dodecahedron	(20,30,12,1)	30	0	-11
Icosahedron	(12,30,20,1)	0	0	11

Table 6.3: Numerical results for the electrostatic potential defined by placing the point charges at the vertices of the five Platonic solids. The first entry in each f -vector is the number of point charges. The next two columns give the number of observed 2- and 1-saddles. In each case, the center is a degenerate equilibrium, for which we give the alternating sum of ranks in local homology. In each case, the number vertices exceeds this alternating sum plus the number of 2-saddles minus the number of 1-saddles by 1.

Through lengthy but straightforward computations and arguments, we were then able to rigorously show the existence of the claimed equilibria. As an example, we present the full proof for the case of the cube, but the argument generalizes to the other Platonic solid with more electrostatic points beside the center, e.g. the dodecahedron.

Proposition 4. *The electrostatic potential defined by the (eight) unit point charges located at the vertices of a cube in \mathbb{R}^3 has twelve equilibria beside the one at the cube's center. These occur along the segments connecting the cube's center to the midpoints of its edges.*

Proof. Assume without loss of generality, that the vertices of the cube are located at the points $(\pm 1, \pm 1, \pm 1)$. The cube is invariant by a discrete group of symmetries, G , which contains, in particular, the reflections on the planes $P_1 = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_2 = 0\}$ and $P_2 = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_1 + x_3 = 0\}$. Denote these reflections $\sigma_1, \sigma_2: \mathbb{R}^3 \rightarrow \mathbb{R}^3$, respectively. Letting $v_1 = (0, 1, 0)$ and $v_2 = (1, 0, 1)$ be vectors orthogonal to these planes, the reflections change their signs: $\sigma_1(v_1) = -v_1$ and $\sigma_2(v_2) = -v_2$. Furthermore, the electrostatic potential V generated by unit charges at the vertices of this cube is invariant by these symmetries, i.e. $V \circ \sigma_1 = V \circ \sigma_2 = V$. Hence, we find from the symmetry that, along P_i ,

$$\langle \nabla V, v_i \rangle = \langle \sigma_i(\nabla V), \sigma_i(v_i) \rangle = -\langle \nabla V, v_i \rangle, \quad (6.4)$$

and therefore $\langle \nabla V, v_i \rangle = 0$, for $i = 1, 2$. In particular, along

$$P_1 \cap P_2 = \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_2 = 0, x_1 + x_3 = 0\} \quad (6.5)$$

$$= \{(x, 0, -x) \in \mathbb{R}^3 \mid x \in \mathbb{R}\} \quad (6.6)$$

we have $\langle \nabla V, v_1 \rangle = 0$ and $\langle \nabla V, v_2 \rangle = 0$, which we write more explicitly as

$$\frac{\partial V}{\partial x_2}(x, 0, -x) = \frac{\partial V}{\partial x_1}(x, 0, -x) + \frac{\partial V}{\partial x_3}(x, 0, -x) = 0. \quad (6.7)$$

We now prove the existence of an equilibrium of V by finding one in $P_1 \cap P_2$. Due to Equation (6.7), we need only prove that there is an $x \in \mathbb{R}$ such that the function $f(x) = V(x, 0, -x)$ has a critical point, i.e. we must find x such that $f'(x) = 0$; that is:

$$\frac{\partial V}{\partial x_1}(x, 0, -x) - \frac{\partial V}{\partial x_3}(x, 0, -x) = 0. \quad (6.8)$$

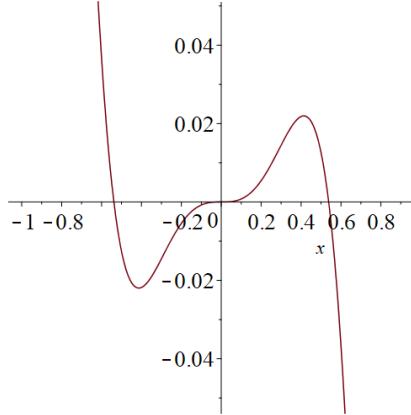


Figure 6.8: Plot of $f'(x)$ whose zeroes correspond to electrostatic points lying in the line $P_1 \cap P_2$.

The function $f'(x)$ is odd and varies continuously with x . We can explicitly compute it and therefore evaluate it at any point; see Figure 6.8. Evaluating the explicit formula for $f'(x)$ we find that

$$f'(-1) = \frac{8}{27} + \frac{8\sqrt{5}}{25} > 0, \quad (6.9)$$

$$f'(-1/2) = -\frac{4}{63} \left[\sqrt{7} \left(\sqrt{3} - \frac{27\sqrt{11}}{121} \right) - \frac{18}{7} \right] < 0, \quad (6.10)$$

$$f'(1/2) = \frac{4}{63} \left[\sqrt{7} \left(\sqrt{3} - \frac{27\sqrt{11}}{121} \right) - \frac{18}{7} \right] > 0, \quad (6.11)$$

$$f'(1) = \frac{8}{27} + \frac{8\sqrt{5}}{25} > 0. \quad (6.12)$$

We therefore conclude from the intermediate value theorem that there are two extra equilibria of V (zeroes of $f'(x)$) located at points $(x_*, 0, -x_*)$ and $(-x_*, 0, x_*)$ for $x_* \in (1/2, 1)$. Notice, in particular, that such points lie along the straight lines passing through the cube's center and the midpoints in its edges.

These two points are in the same orbit of the group of symmetries of the cube and this orbit has cardinality 12 thus resulting in such a number of extra equilibria. \square

Recall that the unit point charges at the vertices of the equilateral triangle in \mathbb{R}^3 define four equilibria, which implies that the ratio of equilibria over the number of vertices is $4/3$. Based on the numbers given in Table 6.3, we deduce that the corresponding ratios for the Platonic solids are $1/4$, $13/8$, $1/6$, $31/20$, and $1/12$, respectively, with the highest ratio for the cube.

6.3.2 Archimedean and Catalan Solids

An Archimedean solid is a convex polytope such that

- all facets are regular polygons;
- all vertices have the same local shape.

Conventionally, there are 13 Archimedean solids, but depending on the interpretation of the second condition, there may be a 14-th, see Figure 6.9.

In particular, we may or may not require vertex-transitivity (for any two vertices, there is a symmetry that maps one vertex into the other), which holds for the 13 Archimedean solids but

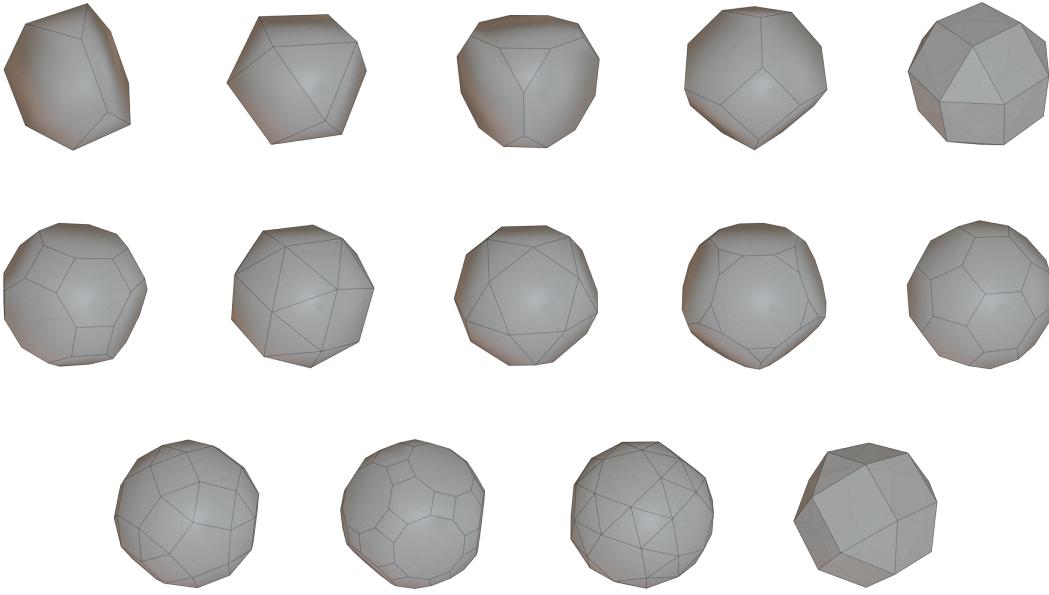


Figure 6.9: From left to right, top to bottom: Archimedean Solids following the order in Table 6.4

not for the Elongated Square Gyrobicupola. The latter is similar to the Rhombicuboctahedron but has one of the square cupola rotated by 45° relative to the opposite cupola. Comparing rows 5 and 14 in Table 6.4, we see that the Elongated Square Gyrobicupola generates fewer equilibria than the Rhombicuboctahedron.

solid	f -vector	2-saddles	1-saddles	center
Truncated Tetrahedron	(12,18,8,1)	18	4	-3
Cuboctahedron	(12,24,14,1)	24	8	-5
Truncated Cube	(24,36,14,1)	36	8	-5
Truncated Octahedron	(24,36,14,1)	36	18	5
Rhombicuboctahedron	(24,48,26,1)	36	8	-5
Truncated Cuboctahedron	(48,72,26,1)	72	20	-5
Snub Cube	(24,60,38,1)	36	8	-5
Icosidodecahedron	(30,60,32,1)	60	20	-11
Truncated Dodecahedron	(60,90,32,1)	90	20	-11
Truncated Icosahedron	(60,90,32,1)	90	12	-19
Rhombicosidodecahedron	(60,120,62,1)	90	20	-11
Truncated Icosidodecahedron	(120,180,62,1)	180	50	-11
Snub Dodecahedron	(60,150,92,1)	90	20	-11
Elongated Square Gyrobicupola	(24,48,26,1)	32	8	-1

Table 6.4: The equilibria of the electrostatic potential of point charges at the vertices of the thirteen Archimedean solids and the Elongated Square Gyrobicupola. The f -vector gives the number of vertices, edges, facets, and the solid itself, in this sequence. Except possibly in the last case, the center is a degenerate equilibrium, and in each case, the number of 1-saddles is non-zero; compare with Table 6.5.

The Catalan solids are dual to the Archimedean solids, so the vertices have multiple possible local shapes, and the facets are congruent to each other. There are again 13 classical examples, and the Pseudo Deltoidal Icositetrahedron as a 14-th solid that is not facet-transitive.

Comparing rows 5 and 14 in Table 6.5, we see that the Pseudo Deltoidal Icositetrahedron generates again fewer equilibria than the similar Deltoidal Icositetrahedron.

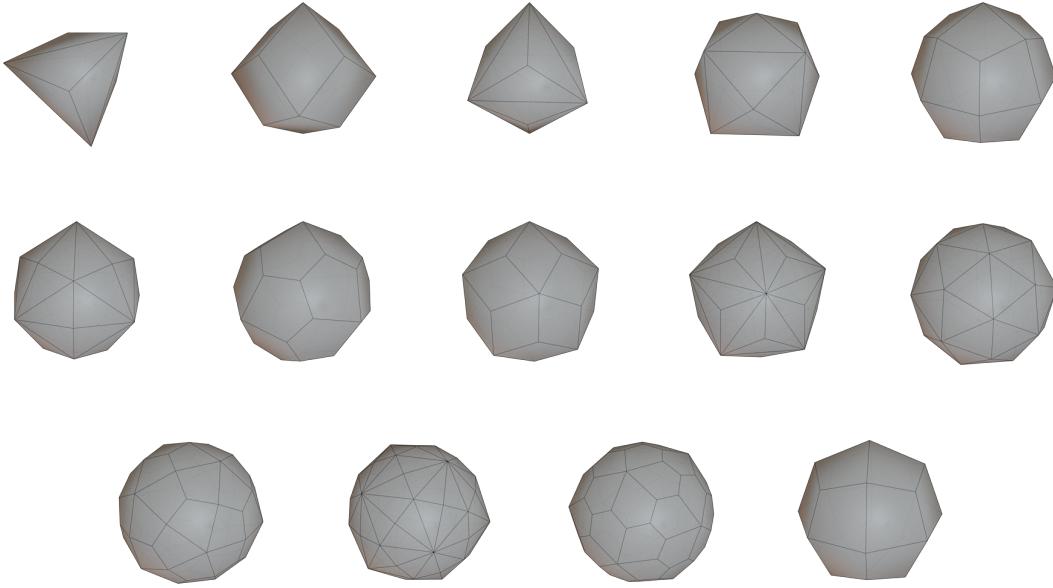


Figure 6.10: From left to right, top to bottom: the Catalan solids following the order in Table 6.5

solid	f -vector	2-saddles	1-saddles	center
Triakis Tetrahedron	(8,18,12,1)	10	0	-3
Rhombic Dodecahedron	(14,24,12,1)	18	0	-5
Triakis Octahedron	(14,36,24,1)	18	0	-5
Tetrakis Hexahedron	(14,36,24,1)	20	0	-7
Deltoidal Icositetrahedron	(26,48,24,1)	30	0	-5
Disdyakis Dodecahedron	(26,72,48,1)	30	0	-5
Pentagonal Icositetrahedron	(38,60,24,1)	48	0	-11
Rhombic Triacontahedron	(32,60,30,1)	42	0	-11
Triakis Icosahedron	(32,90,60,1)	42	0	-11
Pentakis Dodecahedron	(32,90,60,1)	50	0	-19
Deltoidal Hexecontahedron	(62,120,60,1)	72	0	-11
Disdyakis Triacontahedron	(62,180,120,1)	72	0	-11
Pentagonal Hexecontahedron	(92,150,60,1)	102	0	-11
Pseudo Deltoidal Icositetrahedron	(26,48,24,1)	26	0	-1

Table 6.5: The equilibria of the electrostatic potential of point charges at the vertices of the thirteen Catalan solids and the Pseudo Deltoidal Icositetrahedron. The f -vector gives the number of vertices, edges, facets, and the solid itself, in this sequence. Except possibly in the last case, the center is a degenerate equilibrium, and in each case, the number of 1-saddles vanishes; compare with Table 6.4.

We note that all Archimedean solids define electrostatic potentials with a positive number of (non-degenerate) 1-saddles, while all Catalan solids generate no such 1-saddles; compare Tables 6.4 and 6.5. Curiously, also the Platonic solids generate no 1-saddles; see Table 6.3. At this time, there is no explanation for this curious observation.

Finally, recall that the amongst the Platonic solids the Cube gave the highest ratio of equilibria to charges. Amongst the Archimedean and Catalan solids we obtain an even higher ratio of

$33/12$ for the cuboctahedron; see the second row of Table 6.4. Indeed, this is the highest ratio we observe for the regular and semi-regular polytopes surveyed thus far.

6.3.3 Anti-prisms

Beyond the families of solids seen thus far, we wish to consider the prisms and anti-prisms, which are parametrized by an integer, k , and a positive real number. A k -sided prism is the Cartesian product of a regular k -gon and an interval. It has two regular k -gons and k rectangles as facets. We call the aspect ratio of the rectangles the relative height of the prism. Lemma 11 applies only if two prisms agree on k and on the relative height, i.e. there exists a similarity. Indeed, for the same k , the number of equilibria generated by unit point charges at the vertices of the prism depends on the relative height. We shed light on this dependence for the more interesting second family. A k -sided anti-prism has two regular k -gons in parallel planes as facets, one rotated relative to the other so that there are $2k$ isosceles triangles that connect the k -gons and complete the list of boundary facets; see Figure 6.11 for examples. The relative height of the anti-prism is the distance between the two k -gons over the length of their edges. Note that for $k = 3$ and the relative height chosen so that the isosceles triangles are equilateral, the 3-sided or triangular anti-prism is the (regular) octahedron.



Figure 6.11: Anti-prisms for $k = 4, 5, 6$.

Our experiments show that for $k \geq 4$, the relative height of the k -sided anti-prism can be chosen so that the electrostatic potential for unit point charges placed at the vertices has one 2-saddle for each edge, one 1-saddle for each isosceles triangle, and one additional 1-saddle at the center; see Figure 6.12 for illustrations and Appendix 6.A for analytic proofs of some of these equilibria. For $k \geq 4$, the ratio of equilibria over vertices is therefore

$$\frac{1}{2k}(4k + 2k + 1) = 3 + \frac{1}{2k}, \quad (6.13)$$

For $k = 6$, this is the same ratio we get for the cuboctahedron after perturbing the point charges so that the degenerate equilibrium at the center unfolds into individual non-degenerate equilibria; see Section 6.3.4 for details. To maximize this ratio, we minimize the number of sides, which suggests that the 4-sided or square anti-prism is our best choice. In contrast, the equilibria for the 3-sided or triangular anti-prism are either too close to distinguish visually or are absorbed by a degenerate equilibrium at the center.

In addition to the results on the maximum number of equilibria for anti-prisms, we discover an interesting universal phenomenon: the existence of a transition that occurs when the relative height of the prism or anti-prism is $\sqrt{2}$. This transition is universal in the sense that it occurs for both prisms and anti-prisms formed with any regular k -gon and always at relative height $\sqrt{2}$.

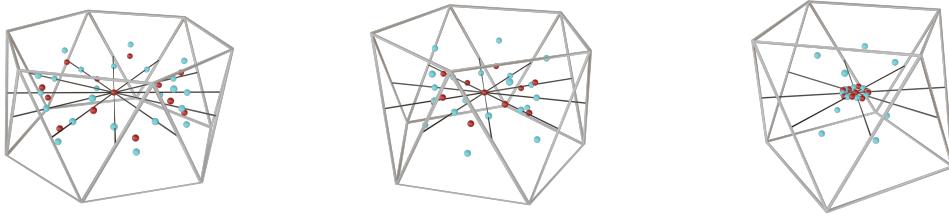


Figure 6.12: From left to right: the hexagonal, pentagonal, square anti-prisms with the heights chosen to maximize the number of equilibria. The ratios of equilibria over vertices are $37/12 < 31/10 < 25/8$, respectively. Observe how a ring of alternating 1- and 2-saddles gets successively more concentrated around the center.

Theorem 13. *Consider the electrostatic potential generated by unit point charges at the vertices of any prism or anti-prism.*

1. *Its center is a non-degenerate equilibrium iff the relative height is not $\sqrt{2}$, in which case the center is a 1-saddle for relative heights less than $\sqrt{2}$ and a 2-saddle for relative heights larger than $\sqrt{2}$.*
2. *Moreover, if the relative height is larger than $\sqrt{2}$, there are two additional 1-saddles along the axis of the prism or anti-prism.*
3. *Furthermore assuming a square or triangular anti-prism and a relative height smaller than $\sqrt{2}$, there is an equilibrium on every line segment connecting the center of the anti-prism to the midpoint of a lateral edge.*

The full proof is provided in Appendix 6.A.

6.3.4 Increasing the Ratio

We can further increase the ratio in two steps. The first of these arises from the realization that the equilibrium at the center of a Platonic solid is necessarily degenerate and we can perturb the location of the charges so that it unfolds into several non-degenerate ones. Indeed, by Theorem 6.3 in [MC69], the point charges can be perturbed such that all equilibria are non-degenerate and arbitrarily close to the degenerate equilibria that give rise to them.

Setting $N = m_1 - m_2 + n - 1$, we note that this is the alternating sum of local homology ranks at the center. It is also a lower bound for the number of 1-saddles and 2-saddles the degenerate equilibrium unfolds into.

For example, we get three 2-saddles near the center of the tetrahedron, five 1-saddles near the center of the cube, and so on. Perturbing the vertices of the five Platonic solid thus increases the ratios to $3/4$, $17/8$, $5/6$, $41/20$, and $11/12$, respectively.

Perturbing the vertices of the cuboctahedron increases the ratio to $37/12$. Compare this with the ratios $37/12 < 31/10 < 25/8$ for the hexagonal, pentagonal, square anti-prisms displayed in Figure 6.12, which are obtained without any perturbation.

In the second step, we iteratively substitute the vertices of a much smaller copy of the solid for each vertex of the original solid. Iterating $\ell - 1$ times, we get a configuration with ℓ layers.

Letting n be the number of vertices of the solid and m the number of equilibria defined by these n charges, the number of point charges in the ℓ layers is $n_\ell = n^\ell$, and the number of equilibria is

$$m \cdot (n^{\ell-1} + n^{\ell-2} + \dots + 1) = m \cdot \frac{n^\ell - 1}{n - 1} = \frac{m}{n - 1} \cdot (n_\ell - 1). \quad (6.14)$$

In particular, the ratio of the number of equilibria by electric charges in the ℓ -th iteration is given by $\frac{m}{n-1} \cdot \frac{n_\ell-1}{n_\ell}$, which for $\ell = 1$ coincides with the initial ratio of $\frac{m}{n}$, and as $\ell \rightarrow +\infty$ converges to $\frac{m}{n-1}$.

In all examples we have computed, the iterated square anti-prism achieves the largest fraction $25/7$ in the limit, when $\ell \rightarrow \infty$. This is the highest ratio ever observed by a wide margin. In particular, we find that for any $\varepsilon > 0$, there is an iteration ℓ_* such that all further iterations $\ell > \ell_*$ have a ratio greater than $25/7 - \varepsilon$.

This establishes the iterated square anti-prism as the greatest lower bound for the maximum ratio of equilibria to charges for the electrostatic potential and naturally leads to the question of whether one can find an example of a configuration which achieves a greater ratio.

6.4 1-parameter Family of Potentials

In the same way as in [GNS07], we generalize the set-up by introducing a real parameter, $p > 0$, which modifies the effect of the distance to the point charges on the potential function:

$$V_p(x) = \sum_{i=1}^n \left(\frac{\zeta_i}{\|x - A_i\|} \right)^p. \quad (6.15)$$

Comparing Equation (6.15) with Equation (6.1), we see that $V = V_1$. The main purpose of this 1-parameter family of functions is to interpolate between the electrostatic potential and the (weighted) Euclidean distance function.

The latter is the limit of one over the p -th root of V_p , as p goes to infinity:

$$E(x) = \lim_{p \rightarrow \infty} \frac{1}{\sqrt[p]{V_p(x)}} = \lim_{p \rightarrow \infty} \left(\sum_{i=1}^n \frac{\zeta_i^p}{\|x - A_i\|^p} \right)^{-1/p} = \min_{1 \leq i \leq n} \frac{\|x - A_i\|}{\zeta_i}. \quad (6.16)$$

6.4.1 No Maxima

For $p = 1$, V_p is harmonic and therefore has neither minima nor maxima. A weaker property still holds for more general values of p . Specifically, if $p \geq 1$ and all charges are positive, then V_p does not have any maxima:

Proposition 5. *Let $V_p: \mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\} \rightarrow \mathbb{R}$ as defined in Equation (6.15), and assume $\zeta_i > 0$ for every $1 \leq i \leq n$. Then for every $p \geq 1$, V_p has no maximum.*

Proof of Proposition 5. The case $p = 1$ follows immediately from the fact that $V_1 = V$ is harmonic. Hence, it suffices to consider the case $p > 1$. For this, we show that the Laplacian

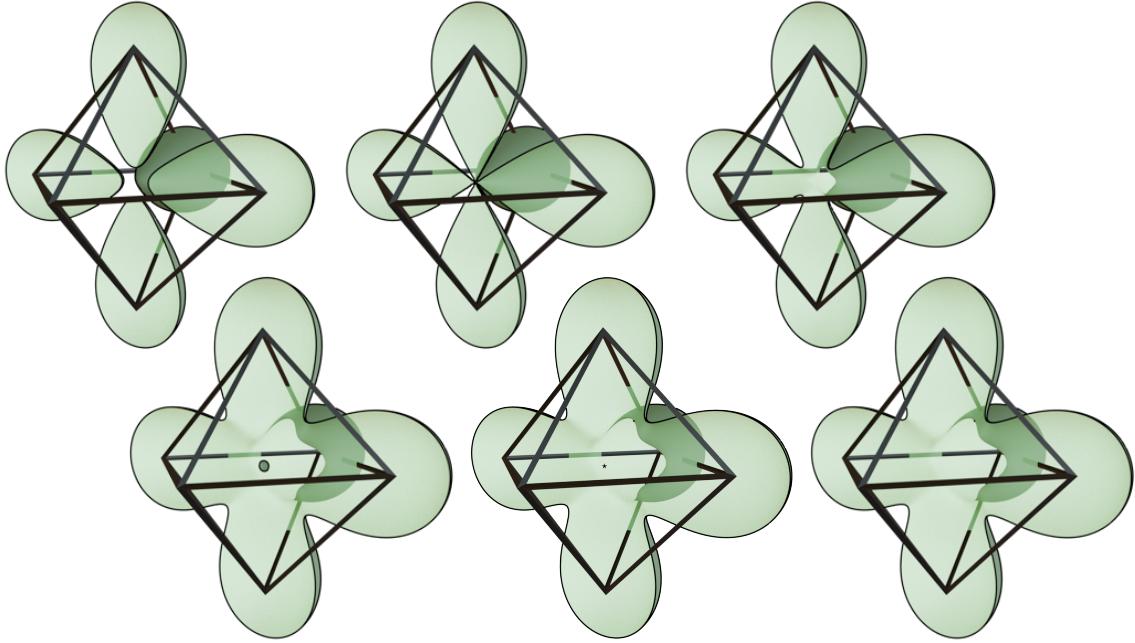


Figure 6.13: Cut-away views of three level sets of V_1 (upper row) and three level sets of $V_{1.3}$ (lower row) defined by point sources at the vertices of the octahedron. From left to right: the values are chosen slightly less than, equal to, and slightly greater than the potential at the center of the octahedron. Removing the front of the surface reveals some of the complication at the center. For V_1 there is a degenerate equilibrium (see Figure 6.5 for the local homology), but for $V_{1.3}$ there is a minimum with a single point in the level set at the center.

satisfies $\Delta V_p > 0$. Since it is the trace of the Hessian and thus the sum of the diagonal eigenvalues, there must be at least one positive eigenvalue, which is incompatible with the existence of a maximum.

To carry out the computation of ΔV_p , we write $r_i(x) = \|x - A_i\|$ and thus $V_p(x) = \sum_{i=1}^n \zeta_i^p / r_i(x)^p$. Let x_1, x_2, x_3 be the coordinates in \mathbb{R}^3 and $A_i = (a_{i1}, a_{i2}, a_{i3})$. Then, using $\frac{\partial r_i}{\partial x_j}(x) = \frac{x_j - a_{ij}}{r_i}$, we compute

$$\frac{\partial V_p}{\partial x_j} = -p \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+1}} \frac{\partial r_i}{\partial x_j} = -p \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+2}} (x_j - a_{ij}); \quad (6.17)$$

$$\frac{\partial^2 V_p}{\partial x_j^2} = -p \sum_{i=1}^n \zeta_i^p \left(\frac{1}{r_i^{p+2}} - (p+2) \frac{x_j - a_{ij}}{r_i^{p+3}} \frac{\partial r_i}{\partial x_j} \right) \quad (6.18)$$

$$= -p \sum_{i=1}^n \zeta_i^p \left(\frac{1}{r_i^{p+2}} - (p+2) \frac{(x_j - a_{ij})^2}{r_i^{p+4}} \right) \quad (6.19)$$

$$= -p \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+2}} \left(1 - (p+2) \frac{(x_j - a_{ij})^2}{r_i^2} \right). \quad (6.20)$$

Finally, summing over the three coordinate directions, we find that

$$\Delta V_p = \sum_{j=1}^3 \frac{\partial^2 V_p}{\partial x_j^2} = -p \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+2}} \left(3 - (p+2) \sum_{j=1}^3 \frac{(x_j - a_{ij})^2}{r_i^2} \right) \quad (6.21)$$

$$= -p \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+2}} (3 - (p+2)) = p(p-1) \sum_{i=1}^n \frac{\zeta_i^p}{r_i^{p+2}}, \quad (6.22)$$

which is positive if all ζ_i are positive. \square

If one trades the assumption that $p \geq 1$ for $p \leq 1$ we can actually conclude, from the same proof, that there are no maxima. On the other hand, if one continues to assume that $p \geq 1$, but that instead that all ζ_i^p are negative, then the same proof shows that there are no minima when $p \geq 1$.

These conclusions are particular cases of the following more general statement which can be inferred from the same proof. Suppose that $p > 0$, then:

- if $(p - 1)\zeta_i^p > 0$ for all $i \in \{1, 2, \dots, n\}$, then there are no local maxima;
- if $p = 1$, then there are neither local minima nor maxima;
- if $(p - 1)\zeta_i^p < 0$ for all $i \in \{1, 2, \dots, n\}$, then there are no local minima.

Figure 6.13 shows a few level sets of V_1 and $V_{1.3}$. The tiny sphere around the center of the octahedron in the lower left level set suggests that the center is a genuine minimum of $V_{1.3}$, which cannot occur if $p = 1$ since this would contradict the harmonicity of V_1 . In fact, we can prove the following result showing that for $p > 1$, the origin is indeed a local minimum of the potential.

Proposition 6. *Let $\{A_1, A_2, \dots, A_6\}$ be the vertices of an octahedron and for $p \geq 1$*

$$V_p(x) = \sum_{i=1}^6 \frac{1}{\|x - A_i\|^p}.$$

Then, the center of the octahedron is a local minimum of V_p for all $p > 1$.

Proof of Proposition 6. With no loss of generality, we can choose the vertices of the octahedron at the points $(\pm 1, 0, 0)$, $(0, \pm 1, 0)$, $(0, 0, \pm 1)$, in which case the center of the octahedron coincides with the origin, $(0, 0, 0)$. Then, using the computations carried out in the proof of Proposition 5, we find that

$$\frac{\partial^2 V_p}{\partial x_i \partial x_j}(0, 0, 0) = 0, \quad (6.23)$$

$$\frac{\partial^2 V_p}{\partial x_i^2}(0, 0, 0) = 2p(p - 1), \quad (6.24)$$

for $i \neq j$ and $i \in \{1, 2, 3\}$, respectively. Hence, the Hessian of V_p at the origin is positive definite whenever $p > 1$ and the result follows. \square

Finally, we point out that in this and other examples, the origin becomes a non-degenerate equilibrium for $p > 1$. Thus, in terms of non-degeneracy, increasing the power, p , can have a similar effect as perturbing the point charges; see [MC69, Thm 6.3].

6.4.2 Counterexample

For positive and equal charges, the function $E: \mathbb{R}^3 \rightarrow \mathbb{R}$ defined in Equation (6.16) is commonly referred to as the Euclidean distance function of the A_i and studied through the Voronoi tessellation, which assigns to each A_i the region of points that are at least as close to A_i as to any of the other points. If the charges are not necessarily the same (but still

positive), then $E: \mathbb{R}^3 \rightarrow \mathbb{R}$ is a weighted version of the Euclidean distance function, namely with multiplicative weights $1/\zeta_i$. The corresponding tessellation of \mathbb{R}^3 is the multiplicatively weighted Voronoi tessellation [AE84, OBS⁺00]. Both types of tessellations are described in more detail in Section 6.5.

The hope expressed as Conjecture 1.8 in [GNS07] is that $V: \mathbb{R}^3 \rightarrow \mathbb{R}$ cannot have more equilibria than $E: \mathbb{R}^3 \rightarrow \mathbb{R}$, in which the latter are defined using the limit process in Equation (6.16). We restate the conjecture now, letting $\#_j$ denote this limiting number of equilibria of V_p for $p \gg 1$ and $j \in \{0, 1, 2, 3\}$.

Conjecture 2 (Conjecture 1.8 (a) in [GNS07]). *For n unit point charges in generic position, and all $p \geq 1$, the number of index j equilibria of V_p is at most $\#_j$.*

However, we were able to show this conjecture is false using one of the Archimedean solids from Section 6.3.2, which we demonstrate now.

Let A_1 to A_{24} be the unit point charges at the vertices of the truncated octahedron, the Voronoi domain in the body centered cubic lattice, and consider the resulting 1-parameter family of potentials,

$$V_p(x) = \sum_{i=1}^n \frac{1}{\|x - A_i\|^p}. \quad (6.25)$$

Figure 6.14 shows truncated octahedron together with the equilibria of the electrostatic potential defined by placing unit point charges at its vertices. This polytope has 14 faces (6 squares and 8 hexagons) and 36 edges (24 shared by a square and a hexagon and 12 shared by two hexagons). Correspondingly, E has 14 1-saddles and 36 2-saddles, which we compare to the 18 1-saddles and 36 2-saddles of V . Hence, the Euclidean distance function has fewer 1-saddles than the electrostatic potential of the same unit point charges, and it has equally many 2-saddles. In total, V has more equilibria than the limiting Euclidean distance function, which contradicts Conjecture 1.8 in [GNS07].

Making use of the obvious symmetries, we may associate equilibria of the electrostatic potential with the facets and edges of the polytope. For example, the line that passes through the centers of two opposite squares is the intersection of four planes of symmetry, and it passes through three equilibria: the center and a 1-saddle near the square on either side.

A more interesting example is the line that passes through the midpoints of two opposite edges shared by two hexagons each, which is the intersection of two planes of symmetry, see Figure 6.14. We observe that such a line passes through five equilibria: the center and two saddles on each side of the center. It is the only example we have so far, in which an edge of the polytope seems to be associated with more than one equilibrium. Furthermore, equilibria located along lines obtained as the common fixed locus of at least two reflections can be deduced to exist using the method implemented in Proposition 4 for the cube.

6.5 Distance Functions

Conjectures 1.8 and 1.9 in [GNS07] motivate us to take a closer look at the Voronoi tessellations defined by the Euclidean distance function and its weighted version. The former relates to the case in which all points have equal charges, while the latter models the situation in which all charges are positive but not necessarily equal to each other.

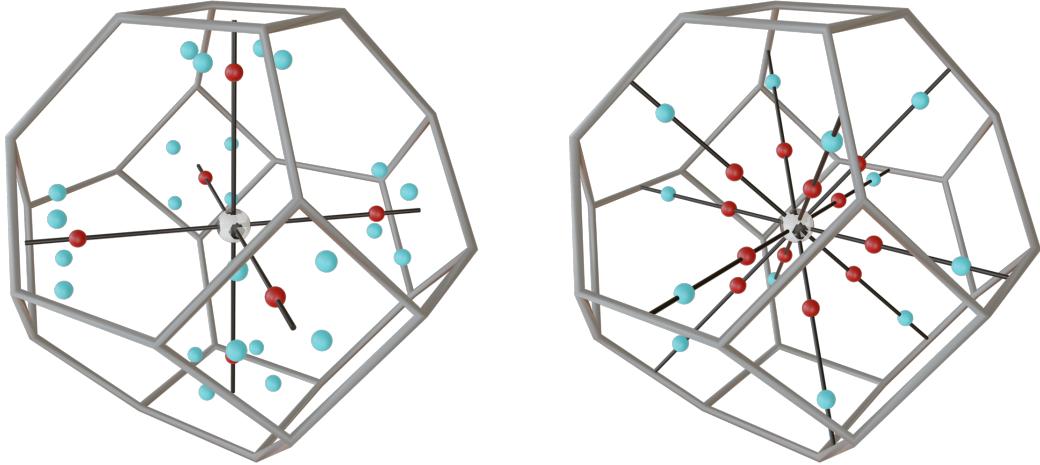


Figure 6.14: Equilibria of the electrostatic potential generated by unit point charges at the vertices of the truncated octahedron. In total there are 36 light blue 2-saddles, 18 dark red 1-saddles, and the degenerate equilibrium at the center. For better visualization, we split these equilibria into two groups, with one 1-saddle and four 2-saddles near each of the six squares displayed in the left panel, and one 1-saddle as well as one 2-saddle for each of the twelve edges shared by two hexagons in the right panel.

6.5.1 Euclidean Distance

If the number of equilibria were monotonically non-decreasing with growing p , then the number of such points of V could be bounded from above by the number of such points of E . To study the latter we recall the definition of Voronoi tessellations and Delaunay mosaics from Section 2.4.

Following [GNS07], we call a Voronoi cell effective if its interior has a non-empty intersection with the interior of the corresponding Delaunay cell. Corresponding cells have complementary dimensions and lie in orthogonal affine subspaces, which implies that the intersection is either empty or a point. Effective cells are interesting because they are in bijection with the equilibria of E . Indeed, with increasing p , the equilibria tend toward the intersections of corresponding Voronoi and Delaunay cells. Figure 6.15 illustrates that already for $p = 2$ the equilibria are barely distinguishable from these intersection points. The Upper bound Theorem for convex polytopes [McM71] implies that the Voronoi tessellation of n points in \mathbb{R}^3 has at most $\mathcal{O}(n^2)$ cells, and thus at most that many effective cells. Recently, [EP24, Theorem 3.1] proved that this bound is asymptotically tight:

Theorem 14. *For every $k \geq 2$, there exist $2k + 2$ points in \mathbb{R}^3 such that the Voronoi tessellation has $(k + 1)^2 - 1$ effective 1-dimensional and k^2 effective 2-dimensional cells.*

Does this lower bound translate to any meaningful lower bound for the electrostatic potential defined by the same number of points with equal charges? Perhaps not, since the construction in [EP24] is delicate and we have not been able to obtain a quadratic number of equilibria of V for similarly placed points.

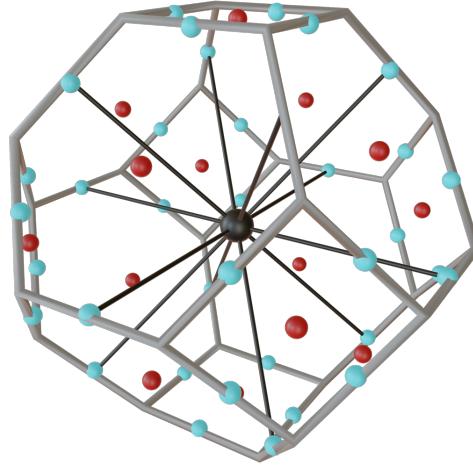


Figure 6.15: The equilibria of V_2 generated by unit point charges at the vertices of the truncated octahedron. Compared with V_1 , we note a drastically reduced number of 1-saddles and a minimum at the origin; see Figure 6.14 where we used two copies of the solid to show all equilibria. While $p = 2$ is still not large, the equilibria are already close to the barycenters of the facets and edges of the solid.

6.5.2 Weighted Euclidean Distance

If we allow points with different charges, then E is the multiplicatively weighted Euclidean distance function; see Equation (6.16). We can still define Voronoi domains, tessellation, and cells, but there are significant differences:

- a domain is no longer the common intersection of closed half-spaces but of $n - 1$ closed balls and closed complements of balls, and thus not necessarily connected;
- already in \mathbb{R}^2 , the tessellation may consist of $\Theta(n^2)$ vertices, (circular) edges, and (connected) regions [AE84];
- there is no established dual structure, like the Delaunay mosaic in the unweighted case.

In the plane, the maximum numbers of cells in the Voronoi tessellation differ even asymptotically between the unweighted and the weighted cases: $\Theta(n)$ versus $\Theta(n^2)$. In \mathbb{R}^3 on the other hand, the maximum number of cells is $\Theta(n^2)$ in both cases. There is at most one equilibrium of E per Voronoi cell, so we have at most $\Theta(n^2)$ such points. However, the difference in the maximum number of cells in \mathbb{R}^2 begs the question whether the maximum number of equilibria differ in the same way. We think not.

Conjecture 3. *The number of equilibria of the weighted Euclidean distance function defined by n points with positive real weights in \mathbb{R}^2 is at most some constant times n .*

The number of equilibria in the plane is relevant because they can be turned into higher index equilibria in three dimensions.

6.5.3 Slices

Consider the Voronoi tessellation that represents the weighted or unweighted Euclidean distance function defined by n weighted or unweighted points in \mathbb{R}^3 . We call the restriction to a plane

or a line a slice. These restrictions look a lot like the lower-dimensional tessellations, but they are more general and can be defined by adding a weight to the respective squared distance to the i -th point:

$$\pi_i(x) = \|x - A_i\|^2 + w_i; \quad (6.26)$$

$$\varphi_i(x) = \left(\frac{\|x - A_i\|}{\zeta_i} \right)^2 + w_i, \quad (6.27)$$

in which w_i is the squared distance of A_i from the plane or line. The weighted squared distance in Equation (6.26) is known as the power or Laguerre distance, and while it is a special case of Equation (6.27), it received substantially more attention in the mathematical literature. In both cases, the 1-dimensional tessellation has at most $O(n)$ cells and thus at most that many equilibria.¹ This is also true for the 2-dimensional tessellation defined by the weighted squared distance in Equation (6.26), but see Conjecture 3 for the more general weighted squared distance in Equation (6.27). By reducing to the same denominator and appealing to the fundamental theorem of algebra, one obtains that, for even $p > 0$ and V_p defined by n unit point charges, its restriction to a straight line in \mathbb{R}^3 has at most $p(n-1) + 2n - 1$ equilibria. While we have no proof, we venture that this bound can be strengthened to at most $O(n)$ equilibria for any $p > 0$, and extended to 2-dimensional slices. To focus, we formulate a more restricted version:

Conjecture 4. *Let $V: \mathbb{R}^3 \setminus \{A_1, A_2, \dots, A_n\} \rightarrow \mathbb{R}$ be the electrostatic potential as in Equation (6.1). Then the restriction of V to any straight line or flat plane in \mathbb{R}^3 has at most $O(n)$ equilibria.*

6.6 Discussion

Inspired by the article of Gabrielov, Novikov and Shapiro [GNS07], we studied the connection between the electrostatic potential and the (weighted) Euclidean distance function defined by a finite collection of point charges. We focus on the 3-dimensional case and on positive charges, which may or may not all be the same. We have two main results: the discovery of a best known lower bound for the maximum ratio of equilibria to charges; and a counterexample to Gabrielov, Novikov and Shapiro's conjecture bounding the number of equilibria of the electrostatic potential by those of the Euclidean distance function.

A more nuanced formulation of the conjecture inspired by a result about the heat-flow in [CE11], could however still hold:

Is the L_1 -norm of the persistence diagram of V_p always smaller than that of V_q for any $1 \leq p \leq q$? In other words, does the combined "strength" of the equilibria monotonically decrease with shrinking parameter, p ?

Such decay in strength does not contradict the increase in number of equilibria, which can indeed happen, as shown by our counterexample. The main question of determining the

¹Indeed, for the weighted squared distance in Equation (6.26), each point generates at most one interval along the line. This is not true for the weighted squared distance in Equation (6.27), but it is not difficult to see that the point with minimum charge generates a single convex cell in three dimensions and therefore at most one interval along the line. The linear bound now follows by induction.

maximum number of equilibria of an electrostatic potential defined by n point charges in \mathbb{R}^3 remains open.

We also tested placing a charge at the centre of the various solids discussed in this chapter. In some instances this did improve the number of equilibria but not enough to exceed the best results already discussed.

Further experiments were carried out to count the number of equilibria of arrangements with randomly placed charges and charges of different signs. Several variations of random arrangements were trialed: a uniform random distribution of n unit charges; a uniform random distribution of n unit positive and n unit negative charges; a uniform random distribution of n uniform $[0, 1]$ random charges; and a uniform distribution of n charges such that $\sum \zeta_i = 0$. The results of this experiment were not particularly insightful in terms of Maxwell's conjecture as the average number of equilibria in the best case (the positive random charges) grew at a near constant rate with n . In particular, a pair of positive charges will result in an equilibria at the midpoint between the two while a positive and negative charge will not, so one would expect fewer equilibria with arrangements of both signs. Proving the expected number of equilibria under these circumstances would however be an interesting result in itself, but beyond the scope of this thesis.

Additionally, in the course of this investigation Oliveira was able to improve upon the upper bound on the number of equilibria Gabrielov, Novikov and Shapiro obtained in [GNS07] using a similar technique involving Bézout's theorem which we outline now.

Theorem 15. *Let $n \in \mathbb{N}$, $A_2, \dots, A_n \in \mathbb{R}^3$, and $\zeta_1, \dots, \zeta_n \in \mathbb{R} \setminus \{0\}$. Then, for the generic position of $A_1 \in \mathbb{R}^3$, the number of isolated critical points of V is at most*

$$2^n \times (3n - 2)^3.$$

The proof begins by computing the system of equations which define a critical point, and clearing out their denominators. The resulting equations can be made into polynomials by a trick that formally introduces extra variables to encode the quantities given by square roots. Next he shows that the non-degenerate critical points of V correspond to non-isolated zeroes of the polynomial system. Finally, he employs Bézout's theorem which provides an upper bound on the number of isolated solutions to these polynomial equations by the degrees of the relevant polynomial equations.

Finally, we summarise the results discussed and obtained in this chapter in Table 6.6. The disparity between our best obtained lower bounds and the best known upper bound is still quite drastic.

6.A Proof of Theorem 13

In this appendix, we provide details about the equilibria of the electrostatic potential for unit point charges placed at the vertices of regular prisms and anti-prisms. After introducing the building blocks of the analysis, we focus on the square and triangular anti-prisms. Together these complete the proof of Theorem 13.

6.A.1 Regular polygons

Let V_0 be the electric potential generated by unit point charges located at the vertices of a regular N -gon centered at the origin, $0 \in \mathbb{R}^3$, and placed on the horizontal plane spanned by

n	New lower bound	Conjectured upper bound	New upper bound	Former Upper bound
	$3.52(n-1)$	$(n-1)^2$	$2^n(3n-2)^3$	$2^{2n^2}(3n)^{2n}$
4	10.56	9	16 000	1.84676×10^{18}
5	14.08	16	70 304	6.49251×10^{26}
6	17.6	25	262 144	5.46298×10^{36}
7	21.12	36	877 952	1.02804×10^{48}
8	24.64	49	2 725 888	4.12306×10^{60}
9	28.16	64	8 000 000	3.39944×10^{74}
10	31.68	81	22 478 848	5.60305×10^{89}
11	35.2	100	61 011 968	1.80537×10^{106}
12	38.72	121	160 989 184	1.1166×10^{124}

Table 6.6: Growth of the various bounds discussed in this chapter. Note the lower bound values are only representative of the growth rate as the construction is built specifically for $n = 8^l$ vertices where l is the number of recursion layers.

the first two coordinate vectors. Denote by R the radius of the circle that passes through its vertices, and by $\alpha = \frac{2\pi}{N}$ the smallest angle that is a rotation symmetry of the N -gon. We start by proving that the partial derivatives are symmetric with respect to the horizontal plane, and they all vanish at the origin.

Lemma 12. *Let $V_0: \mathbb{R}^3 \rightarrow \mathbb{R}$ be the electrostatic potential of unit point charges at the vertices of a regular N -gon, as described, and write $x = (x_1, x_2, x_3)$ for a point with its Cartesian coordinates in \mathbb{R}^3 . Then*

$$\frac{\partial V_0}{\partial x_3}(x_1, x_2, x_3) = -\frac{\partial V_0}{\partial x_3}(x_1, x_2, -x_3) \quad \text{and} \quad \frac{\partial V_0}{\partial x_1}(0, 0, x_3) = \frac{\partial V_0}{\partial x_2}(0, 0, x_3) = 0. \quad (6.28)$$

Proof. The configuration is invariant by reflection through the $x_3 = 0$ plane, which contains the N -gon. Hence, $V_0(x_1, x_2, x_3) = V_0(x_1, x_2, -x_3)$ and differentiating with respect to x_3 gives the first relation. To see the second relation, note that the invariance of the N -gon under rotations of the horizontal plane by an angle α implies

$$V_0(x_1, x_2, x_3) = V_0(x_1 \cos \alpha + x_2 \sin \alpha, -x_1 \sin \alpha + x_2 \cos \alpha, x_3), \quad (6.29)$$

for all $x = (x_1, x_2, x_3) \in \mathbb{R}^3$. Differentiating with respect to x_1 and x_2 , respectively, and evaluating at $x_1 = x_2 = 0$ yields

$$\frac{\partial V_0}{\partial x_1}(0, 0, x_3) = \cos \alpha \frac{\partial V_0}{\partial x_1}(0, 0, x_3) - \sin \alpha \frac{\partial V_0}{\partial x_2}(0, 0, x_3); \quad (6.30)$$

$$\frac{\partial V_0}{\partial x_2}(0, 0, x_3) = \sin \alpha \frac{\partial V_0}{\partial x_1}(0, 0, x_3) + \cos \alpha \frac{\partial V_0}{\partial x_2}(0, 0, x_3). \quad (6.31)$$

This gives a system of linear equations for $\frac{\partial V_0}{\partial x_1}(0, 0, x_3)$ and $\frac{\partial V_0}{\partial x_2}(0, 0, x_3)$ whose only solution is when both vanish, as claimed. \square

It will also be useful to have the second partial derivative at the origin in vertical direction.

Lemma 13. *For $V_0: \mathbb{R}^3 \rightarrow \mathbb{R}$ as before and all $x_3 \in \mathbb{R}$, we have*

$$\frac{\partial^2 V_0}{(\partial x_3)^2}(0, 0, x_3) = -\frac{N}{(R^2 + x_3^2)^{5/2}} (R^2 - 2x_3^2). \quad (6.32)$$

Proof. Denote the vertices of the N -gon by $A_i = (a_{i,1}, a_{i,2}, 0)$, for $i = 1, 2, \dots, N$, and write r_i for the distance of A_i from a point $x = (x_1, x_2, x_3)$. Then the first and second partial derivatives along the vertical direction are

$$\frac{\partial V_0}{\partial x_3} = -\sum_{i=1}^N \frac{1}{r_i^2} \frac{\partial r_i}{\partial x_3} = -\sum_{i=1}^N \frac{x_3}{r_i^3}; \quad (6.33)$$

$$\frac{\partial^2 V_0}{(\partial x_3)^2} = -\sum_{i=1}^N \left(\frac{1}{r_i^3} - 3 \frac{x_3}{r_i^4} \frac{\partial r_i}{\partial x_3} \right) = -\sum_{i=1}^N \left(\frac{1}{r_i^3} - 3 \frac{x_3^2}{r_i^5} \right) = -\sum_{i=1}^N \frac{1}{r_i^5} (r_i^2 - 3x_3^2). \quad (6.34)$$

Evaluating the second partial derivative at $(0, 0, x_3)$, we have $r_i^2 = R^2 + x_3^2$ independent of i , which implies Equation (6.32). \square

6.A.2 Rotated prisms

We get a prism by connecting a regular N -gon with a copy of itself translated in an orthogonal direction. If we first rotate the copy by an angle β —around its center and within its plane—we call the convex hull of the two N -gons a β -rotated prism. For $\beta = \frac{\alpha}{2}$, this is an anti-prism. Its axis is the line that passes through the centers of the two N -gons, its radius is the distance of the vertices from this line, and its height is the distance between the planes that contain the two N -gons. We now consider the electrostatic potential obtained by placing unit point charges at the vertices of such a β -rotated prism. Supposing the N -gons lie in horizontal planes at distance $\frac{h}{2}$ above and below the origin, the electrostatic potential is

$$\begin{aligned} V(x_1, x_2, x_3) &= V_0(x_1, x_2, x_3 + \frac{h}{2}) \\ &+ V_0(x_1 \cos \beta + x_2 \sin \beta, -x_1 \sin \beta + x_2 \cos \beta, x_3 - \frac{h}{2}). \end{aligned} \quad (6.35)$$

Not surprisingly, the type of the equilibrium at the origin depends on the height, but with the exception of a particular height, it does not depend on the angle of the rotation.

Theorem 16. *Let $\beta \in \mathbb{R}/2\pi\mathbb{Z}$, $R > 0$, and $h > 0$. Consider the electrostatic potential generated by unit point charges at the vertices of a β -rotated prism with radius R and height h . Then, its center is a non-degenerate equilibrium iff $h \neq \sqrt{2}R$, in which case the center is a 1-saddle if $h < \sqrt{2}R$ and a 2-saddle if $h > \sqrt{2}R$. Moreover, in the latter case there are two additional 1-saddles along the axis of the β -rotated prism.*

Proof. To use the form of the potential in Equation (6.35), we assume that the two N -gons of the β -rotated prism lie in horizontal planes and its center is the origin in \mathbb{R}^3 . We start by confirming that the origin is always an equilibrium. Using Lemma 12, we obtain

$$\frac{\partial V}{\partial x_i}(0, 0, 0) = \frac{\partial V_0}{\partial x_i}(0, 0, \frac{h}{2}) + \frac{\partial V_0}{\partial x_i}(0, 0, -\frac{h}{2}) = 0, \quad (6.36)$$

which for $i = 1, 2$ is true because both partial derivatives vanish, and for $i = 3$ follows from the symmetry across the horizontal coordinate plane. Next we investigate the Hessian at the origin. By differentiating the identity $\frac{\partial V_0}{\partial x_3}(x_1, x_2, x_3) = -\frac{\partial V_0}{\partial x_3}(x_1, x_2, -x_3)$ with respect to x_1 , x_2 , and x_3 , we obtain

$$\frac{\partial^2 V_0}{\partial x_i \partial x_3}(x_1, x_2, x_3) = -\frac{\partial^2 V_0}{\partial x_i \partial x_3}(x_1, x_2, -x_3); \quad (6.37)$$

$$\frac{\partial^2 V_0}{(\partial x_3)^2}(x_1, x_2, x_3) = \frac{\partial^2 V_0}{(\partial x_3)^2}(x_1, x_2, -x_3), \quad (6.38)$$

for $i = 1, 2$. Hence, using the first of these identities, we have

$$\frac{\partial^2 V}{\partial x_i \partial x_3}(0, 0, 0) = \frac{\partial^2 V_0}{\partial x_i \partial x_3}(0, 0, \frac{h}{2}) + \frac{\partial^2 V_0}{\partial x_i \partial x_3}(0, 0, -\frac{h}{2}) = 0. \quad (6.39)$$

This shows that $\lambda_3 = \frac{\partial^2 V}{(\partial x_3)^2}(0, 0, 0)$ is an eigenvalue of the Hessian at the origin. Using the second identity above, we see that this eigenvalue satisfies

$$\lambda_3 = \frac{\partial^2 V}{(\partial x_3)^2}(0, 0, 0) = \frac{\partial^2 V_0}{\partial x_i \partial x_3}(0, 0, \frac{h}{2}) + \frac{\partial^2 V_0}{\partial x_i \partial x_3}(0, 0, -\frac{h}{2}) = 2 \frac{\partial^2 V_0}{(\partial x_3)^2}(0, 0, \frac{h}{2}). \quad (6.40)$$

The rotational symmetry by $\beta \neq 0$ then implies that the remaining two eigenvalues must coincide: $\lambda_1 = \lambda_2 = \lambda$. Furthermore, since V is harmonic, we have $\lambda = -\frac{\partial^2 V_0}{(\partial x_3)^2}(0, 0, \frac{h}{2})$.

It remains to compute this quantity and show it vanishes precisely when $h = \sqrt{2}R$ as claimed in the statement. This can be done by making use of Lemma 13 to obtain

$$\frac{\partial^2 V}{(\partial x_3)^2}(0, 0, 0) = -\frac{2N}{(R^2 + x_3^2)^{5/2}} \left(R^2 - \frac{h^2}{2} \right), \quad (6.41)$$

which vanishes iff $h = \sqrt{2}R$. Furthermore, the eigenvalues of the Hessian at the origin are

$$\lambda_3 = -\frac{2N}{(R^2 + x_3^2)^{5/2}} \left(R^2 - \frac{h^2}{2} \right); \quad (6.42)$$

$$\lambda_1 = \lambda_2 = \frac{N}{(R^2 + x_3^2)^{5/2}} \left(R^2 - \frac{h^2}{2} \right). \quad (6.43)$$

We therefore conclude that for $h < \sqrt{2}R$ we have $\lambda_1 = \lambda_2 > 0$ and $\lambda_3 < 0$, so the origin is a 1-saddle. On the other hand, if $h > \sqrt{2}R$ we have $\lambda_1 = \lambda_2 < 0$ and $\lambda_3 > 0$, and so the origin is a 2-saddle. However, in this case there are at least two additional equilibria along the axis of the β -rotated prism, which by construction is the x_3 -axis. To prove their existences, we compute the derivatives of V along this axis:

$$\frac{\partial V}{\partial x_i}(0, 0, x_3) = \frac{\partial V_0}{\partial x_i}(0, 0, x_3 + \frac{h}{2}) + \frac{\partial V_0}{\partial x_i}(0, 0, x_3 - \frac{h}{2}) = 0 + 0 = 0, \quad (6.44)$$

for $i = 1, 2$. Hence, we need only look for zeroes of $\frac{\partial V}{\partial x_3}(0, 0, x_3)$, which we compute as

$$\frac{\partial V}{\partial x_3}(0, 0, x_3) = \frac{\partial V_0}{\partial x_3}(0, 0, x_3 + \frac{h}{2}) + \frac{\partial V_0}{\partial x_3}(0, 0, x_3 - \frac{h}{2}). \quad (6.45)$$

To compute the two terms on the right-hand side, we use Lemma 13 and find that $\frac{\partial V_0}{\partial x_3} = -\sum_{i=1}^N \frac{x_3}{r_i^3}$, where $r_i(x_1, x_2, x_3)^2 = (x_1 + a_{i1})^2 + (x_2 - a_{i2})^2 + x_3^2$. From this it follows that

$$\frac{\partial V}{\partial x_3}(0, 0, x_3) = -\sum_{i=1}^N \frac{x_3 + \frac{h}{2}}{r_i^3(0, 0, x_3 + \frac{h}{2})} - \sum_{i=1}^N \frac{x_3 - \frac{h}{2}}{r_i^3(0, 0, x_3 - \frac{h}{2})} \quad (6.46)$$

$$= -N \frac{x_3 + \frac{h}{2}}{\left(R^2 + (x_3 + \frac{h}{2})^2 \right)^{3/2}} - N \frac{x_3 - \frac{h}{2}}{\left(R^2 + (x_3 - \frac{h}{2})^2 \right)^{3/2}}. \quad (6.47)$$

Evaluating this at the top and bottom of the β -rotated prism, we obtain

$$\frac{\partial V}{\partial x_3}(0, 0, \frac{h}{2}) = -\frac{h}{(R^2 + h^2)^{3/2}} < 0 \quad \text{and} \quad \frac{\partial V}{\partial x_3}(0, 0, -\frac{h}{2}) = \frac{h}{(R^2 + h^2)^{3/2}} > 0. \quad (6.48)$$

On the other hand, further differentiating with respect to x_3 , as in the proof of Theorem 16, we find that

$$\frac{\partial^2 V}{\partial x_3^2}(0, 0, 0) = -\frac{2N}{(R^2 + x_3^2)^{5/2}} \left(R^2 - \frac{h^2}{2} \right), \quad (6.49)$$

which is positive when $h > \sqrt{2}R$. Hence, there exists $\varepsilon > 0$, such that $\frac{\partial V}{\partial x_3}(0, 0, t)$ is positive for $t \in (0, \varepsilon)$ and negative for $t \in (-\varepsilon, 0)$. Combining this with the previously computed values of $\frac{\partial V}{\partial x_3}$ at the centers of the two N -gons, the intermediate value theorem implies that $\frac{\partial V}{\partial x_3}$ has two additional zeroes along the x_3 -axis.

We now characterize these equilibria by showing that they are 1-saddles. At both, $V(0, 0, x_3)$ increases on the left and decreases on the right, so they are local maxima of the potential restricted to the x_3 -axis. Given that the potential is harmonic, at least one of the remaining eigenvalues of the Hessian must be positive, and by rotational symmetry so must the other. As the Hessian has a single negative eigenvalue, both equilibria are 1-saddles. \square

6.A.3 Existence results for equilibria in anti-prisms

We consider anti-prisms whose top and bottom faces are N -gons and denote their minimum angle of symmetry by $\alpha = \frac{2\pi}{N}$. Let $R > 0$ be the radius of these N -gons and $h > 0$ the height of the anti-prism, then its relative height is defined by $\frac{h}{R}$.

Identify $\mathbb{R}^3 \cong \mathbb{C} \times \mathbb{R}$, then for $\ell \in \{1, \dots, N\}$ we denote the vertices of the top N -gon as $A_\ell = (Re^{i(\ell\alpha+\frac{\alpha}{4})}, \frac{h}{2})$ and those at the bottom N -gon as $B_\ell = (Re^{-i(\ell\alpha+\frac{\alpha}{4})}, -\frac{h}{2})$. Then, the electrostatic potential generated by placing unit point charges at the vertices of this anti-prism is

$$V(x) = V_u(x) + V_d(x),$$

where

$$V_u(x) = \sum_{\ell=1}^N \frac{1}{|x - A_\ell|}, \quad V_d(x) = \sum_{\ell=1}^N \frac{1}{|x - B_\ell|}.$$

We now state and prove the following result.

Theorem 17. *Let V be the electric potential generated by placing unit point charges at the vertices of an anti-prism as above and suppose its relative height satisfies $\frac{h}{R} < \sqrt{2}$. Then, there are electrostatic points located along the line segments connecting the center of the anti-prism to the midpoints of its lateral edges.*

Proof. We consider the case of the line segment connecting the center to the midpoint of the edge connecting A_N to B_N . The remaining cases can be handled similarly.

The midpoint between A_N and B_N is $\frac{A_N + B_N}{2} = (R \cos(\frac{\alpha}{4}), 0, 0)$ and we can parameterise the segment that connects the origin to it using $t \in [0, 1] \mapsto (tR \cos(\frac{\alpha}{4}), 0, 0)$. The derivatives of $V(x)$ in the normal directions $n_1 = (0, 1, 0)$ and $n_2 = (0, 0, 1)$ are given by

$$\begin{aligned} \langle \nabla V, n_1 \rangle &= \frac{\partial V_u}{\partial x_2} + \frac{\partial V_d}{\partial x_2} \\ \langle \nabla V, n_2 \rangle &= \frac{\partial V_u}{\partial x_3} + \frac{\partial V_d}{\partial x_3}. \end{aligned}$$

Now, we compute each of the terms in the right hand side in turn. We start with

$$\begin{aligned}\frac{\partial V_u}{\partial x_2} &= -\sum_{\ell=1}^N \frac{(x_2 - R \sin(\ell\alpha + \frac{\alpha}{4}))}{|x - A_\ell|^3} \\ \frac{\partial V_d}{\partial x_2} &= -\sum_{\ell=1}^N \frac{(x_2 + R \sin(\ell\alpha + \frac{\alpha}{4}))}{|x - B_\ell|^3},\end{aligned}$$

on the other hand, along the line connecting 0 to $\frac{A_N + B_N}{2}$, we have $x_2 = 0 = x_3$ and

$$\begin{aligned}|x - A_\ell| &= |(x_1 - R \cos(\ell\alpha + \frac{\alpha}{4}), R \sin(\ell\alpha + \frac{\alpha}{4}), 0)| \\ &= |(x_1 - R \cos(-\ell\alpha - \frac{\alpha}{4}), R \sin(-\ell\alpha - \frac{\alpha}{4}), 0)| \\ &= |x - B_\ell|.\end{aligned}$$

which shows that along such a line

$$\frac{\partial V_u}{\partial x_2} = -\frac{\partial V_d}{\partial x_2}.$$

On the other hand,

$$\begin{aligned}\frac{\partial V_u}{\partial x_3} &= -\sum_{\ell=1}^N \frac{(x_3 - \frac{h}{2})}{|x - A_\ell|^3} \\ \frac{\partial V_d}{\partial x_3} &= -\sum_{\ell=1}^N \frac{(x_3 + \frac{h}{2})}{|x - B_\ell|^3},\end{aligned}$$

and, again, along the line $x = 2 = 0 = x_3$ we have

$$\frac{\partial V_u}{\partial x_3} = -\frac{\partial V_d}{\partial x_3}.$$

Hence,

$$\langle \nabla V, n_1 \rangle = 0 = \langle \nabla V, n_2 \rangle.$$

Now, we finally consider the directional derivative in the direction of the tangent vector to the line under consideration $v = (1, 0, 0)$

$$\begin{aligned}\langle \nabla V, v \rangle &= \frac{\partial V}{\partial x_1} \\ &= \frac{\partial V_u}{\partial x_1} + \frac{\partial V_d}{\partial x_1} \\ &= -\sum_{\ell=1}^N \left(\frac{(x_1 - R \cos(\ell\alpha + \frac{\alpha}{4}))}{|x - A_\ell|^3} + \frac{(x_1 - R \cos(-\ell\alpha - \frac{\alpha}{4}))}{|x - B_\ell|^3} \right).\end{aligned}$$

Furthermore, along the line $x_2 = 0 = x_3$, we have

$$\begin{aligned}\langle \nabla V, v \rangle &= -\sum_{\ell=1}^N \left(\frac{(x_1 - R \cos(\ell\alpha + \frac{\alpha}{4}))}{|x - A_\ell|^3} + \frac{(x_1 - R \cos(-\ell\alpha - \frac{\alpha}{4}))}{|x - B_\ell|^3} \right) \\ &= -2 \sum_{\ell=1}^N \frac{(x_1 - R \cos(\ell\alpha + \frac{\alpha}{4}))}{|x - A_\ell|^3}\end{aligned}$$

and inserting $x_1 = tR \cos(\frac{\alpha}{4})$ gives

$$\langle \nabla V, v \rangle = -\frac{2}{R^2} \sum_{\ell=1}^N \frac{t \cos(\frac{\alpha}{4}) - \cos(\ell\alpha + \frac{\alpha}{4})}{\left((t \cos(\frac{\alpha}{4}) - \cos(\ell\alpha + \frac{\alpha}{4}))^2 + (\sin(\ell\alpha + \frac{\alpha}{4}))^2 + \frac{h^2}{4}\right)^{\frac{3}{2}}}.$$

Now, notice that $\cos(\frac{\alpha}{4}) \geq \cos(\ell\alpha + \frac{\alpha}{4})$ with strict inequality for $\ell \neq N$, which shows that for $t = 1$ we have $\langle \nabla V, v \rangle < 0$. On the other hand, for evaluating this quantity near $t = 0$, we compute its Taylor series at $t = 0$

$$\begin{aligned} \langle \nabla V, v \rangle &= \frac{16}{R^2} \frac{1}{((4R^2 + h^2)^{\frac{3}{2}})} \sum_{\ell=1}^N \cos(\ell\alpha + \frac{\alpha}{4}) \\ &\quad + \frac{16t}{R^2} \frac{\cos(\frac{\alpha}{4})}{(4R^2 + h^2)^{\frac{5}{2}}} \sum_{\ell=1}^N \left(12R^2 \cos(\ell\alpha + \frac{\alpha}{4})^2 - 4R^2 - h^2\right) + O(t^2). \end{aligned}$$

Now, using the facts that

$$\sum_{\ell=1}^N \cos(\ell\alpha + \frac{\alpha}{4}) = 0, \quad \sum_{\ell=1}^N \cos(\ell\alpha + \frac{\alpha}{4})^2 = \frac{N}{2},$$

we find that

$$\begin{aligned} \langle \nabla V, v \rangle &= \frac{16t}{R^2} \frac{\cos(\frac{\alpha}{4})}{(4R^2 + h^2)^{\frac{5}{2}}} (6NR^2 - 4NR^2 - Nh^2) + O(t^2) \\ &= \frac{16Nt}{R^2} \frac{\cos(\frac{\alpha}{4})}{(4R^2 + h^2)^{\frac{5}{2}}} (2R^2 - h^2) + O(t^2). \end{aligned}$$

Hence, if $h < \sqrt{2}R$ we have that along this line $\langle \nabla V, v \rangle$ is positive for small positive $t > 0$. Combining this with the previously proven fact that it is negative for $t = 1$ we find, from the intermediate value theorem, that $\langle \nabla V, v \rangle$ has a zero at some $t \in (0, 1)$. This finishes the proof that ∇V has a zero along the segment connecting 0 to $\frac{A_N + B_N}{2}$. \square

For $h = \sqrt{2}R$, the complexity of the degenerate equilibrium at the origin is encapsulated by the Taylor expansion of the potential at the origin in homogeneous harmonic polynomials. As we shall see, not only the degree-2 terms vanish², but also the degree-3 terms do.

Lemma 14. *Let $V: \mathbb{R}^3 \rightarrow \mathbb{R}$ be the electrostatic potential generated by unit point charges placed at the vertices of the square anti-prism. For $h = \sqrt{2}R$, its Taylor series at the origin is given by*

$$V(x_1, x_2, x_3) = \frac{8\sqrt{6}}{3} \left(1 - \frac{7}{108}(x_1^4 + x_2^4) - \frac{14}{81}x_3^4 - \frac{7}{54}x_1^2x_2^2 + \frac{14}{27}x_3^2(x_1^2 + x_2^2)\right) + \dots, \quad (6.50)$$

with terms of order exceeding 4 not shown.

Again, we compute the Taylor expansion at the origin as a way to understand the degenerate critical point there for $h = \sqrt{2}R$.

Lemma 15. *Let $V: \mathbb{R}^3 \rightarrow \mathbb{R}$ be the electrostatic potential generated by unit point charges placed at the vertices of the triangular anti-prism as specified at the beginning of this subsection. For $h = \sqrt{2}R$, its Taylor series at the origin is given by*

$$\sqrt{3} - \frac{7\sqrt{3}}{216} \left(\frac{x_1^4}{8} + \frac{x_2^4}{8} + \frac{x_3^2}{3} + \frac{x_1^2x_2^2}{4} - x_1^2x_3^2 - x_2^2x_3^2 + \frac{5\sqrt{2}}{2}x_2x_3(x_2^2 - x_1^2)\right) + \dots, \quad (6.51)$$

with terms of order exceeding 4 not shown.

²We already knew this from the vanishing of the Hessian in the proof of Theorem 16.

Bibliography

[AB26] J. W. Alexander and G. B. Briggs. On types of knotted curves. *The Annals of Mathematics*, 28(1/4):562, 1926.

[ABE09] Dominique Attali, Jean-Daniel Boissonnat, and Herbert Edelsbrunner. *Stability and Computation of Medial Axes - a State-of-the-Art Report*, pages 109–125. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[AC08] Nina Amenta and Sunghee Choi. *Voronoi Methods for 3D Medial Axis Approximation*, pages 223–239. Springer Netherlands, Dordrecht, 2008.

[ACK01] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19(2-3):127–153, 2001.

[ADPKF⁺24a] Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. Tight bounds for the learning of homotopy à la Niyogi, Smale, and Weinberger for subsets of Euclidean spaces and of Riemannian manifolds. In *40th International Symposium on Computational Geometry (SoCG 2024)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

[ADPKF⁺24b] Dominique Attali, Hana Dal Poz Kouřimská, Christopher Fillmore, Ishika Ghosh, André Lieutier, Elizabeth Stephenson, and Mathijs Wintraecken. The ultimate frontier: An optimality construction for homotopy inference (media exposition). In *40th International Symposium on Computational Geometry (SoCG 2024)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024.

[AE84] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition*, 17(2):251–257, January 1984.

[AEHW06] Pankaj K. Agarwal, Herbert Edelsbrunner, John Harer, and Yusu Wang. Extreme elevation on a 2-manifold. *Discrete & Computational Geometry*, 36(4):553–572, September 2006.

[AEW04] Pankaj K. Agarwal, Herbert Edelsbrunner, and Yusu Wang. Computing the writhing number of a polygonal knot. *Discrete Comput. Geom.*, 32(1):37–53, 2004.

[AGH⁺24] Shreya Arya, Barbara Giunti, Abigail Hickok, Lida Kanari, Sarah McGuire, and Katharine Turner. Decomposing the persistent homology transform of star-shaped objects. *arXiv preprint arXiv:2408.14995*, 2024.

[Ale23] J. W. Alexander. A lemma on systems of knotted curves. *Proceedings of the National Academy of Sciences*, 9(3):93–95, March 1923.

[AM96] Dominique Attali and Annick Montanvert. Modeling noise for a better simplification of skeletons. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 3, pages 13–16. IEEE, 1996.

[Arc17] Archimedes. *The Works of Archimedes: Translation and Commentary*. Cambridge University Press, 2017.

[Arn13] Vladimir Arnold. *Singularities of caustics and wave fronts*, volume 62 of *Mathematics and its Applications*. Springer Science & Business Media, Berlin, Heidelberg, 2013.

[BB05] Joan S. Birman and Tara E. Brendle. *Braids*, page 19–103. Elsevier, 2005.

[BBML⁺18] Thibault Blanc-Beyne, Géraldine Morin, Kathryn Leonard, Stefanie Hahmann, and Axel Carlier. A salience measure for 3D shape decomposition and sub-parts classification. *Graphical Models*, 99:22–30, 2018.

[BDC⁺15] Gulce Bal, Julia Diebold, Erin Wolf Chambers, Ellen Gasparovic, Ruizhen Hu, Kathryn Leonard, Matineh Shaker, and Carola Wenk. Skeleton-based recognition of shapes in images via longest path matching. In *Research in Shape Modeling*, pages 81–99, Cham, 2015. Springer International Publishing.

[BH99] Joan S Birman and Michael D Hirsch. A new algorithm for recognizing the unknot. *Geometry & Topology*, 2(1):175–220, 1999.

[Bin64] RH Bing. Some aspects of the topology of 3-manifolds related to the Poincaré conjecture. *Lectures on modern mathematics*, 2:93–128, 1964.

[Blu67] Harry Blum. A Transformation for Extracting New Descriptors of Shape. In *Models for the Perception of Speech and Visual Form*, pages 362–380. MIT Press, Cambridge, 1967.

[Bog89] Ilya A Bogaevsky. Metamorphoses of singularities of minimum functions, and bifurcations of shock waves of the burgers equation with vanishing viscosity. *Algebra i Analiz*, 4:1–16, 1989. English translation in St.Petersburg (Leningrad) Mathematical Journal, 1 (1990), no.4, 807–823.

[Bog95] Ilya A Bogaevsky. Perestroikas of fronts in evolutionary families. *Trudy Matematicheskogo Instituta imeni V.A. Steklova*, 209:65–83, 1995. English translation in *Proceedings of the Steklov Institute of Mathematics* 209 (1995), 57–72.

[Bog99] Ilya A Bogaevsky. Singularities of viscosity solutions of Hamilton-Jacobi equation (singularity theory and differential equations). *Suri kaiseki kenkyujo kokyuroku*, 1111:138–143, 1999.

[Bog02] Ilya A Bogaevsky. Perestroikas of shocks and singularities of minimum functions. *Physica D: Nonlinear Phenomena*, 173(1-2):1–28, 2002.

[BP06] Mitchell A Berger and Chris Prior. The writhe of open and closed curves. *Journal of Physics A: Mathematical and General*, 39(26):8321, jun 2006.

[BTZ00] F. Bacelli, K. Tchoumatchenko, and S.A. Zuyev. Markov paths on the Poisson-Delaunay graph with applications to routing in mobile networks. *Adv. Appl. Probab.*, 32, 2000.

[Buc77] Michael A. Buchner. Stability of the cut locus in dimensions less than or equal to 6. *Inventiones Mathematicae*, 43(3):199–231, October 1977.

[Buc78] Michael A. Buchner. The structure of the cut locus in dimension less than or equal to six. *Compositio Mathematica*, 37(1):103–119, 1978.

[Bur20] Benjamin A. Burton. The next 350 million knots. In *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[CA11] José Ignacio Cogolludo-Agustín. Braid monodromy of algebraic curves. *Annales mathématiques Blaise Pascal*, 18(1):141–209, 2011.

[CDS13] S.-W. Cheng, T.R. Dey, and J. Shewchuk. *Delaunay Mesh Generation*. Chapman and Hall / CRC computer and information science series. CRC Press, 2013.

[CE11] C. Chen and H. Edelsbrunner. Diffusion runs low on persistence fast. In *2011 International Conference on Computer Vision*, page 423–430. IEEE, November 2011.

[CEF13] Andrea Cerri, Marc Ethier, and Patrizio Frosini. A study of monodromy in the computation of multidimensional persistence. In *Discrete Geometry for Computer Imagery*, pages 192–202, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[CFSW25a] Erin Chambers, Christopher Fillmore, Elizabeth Stephenson, and Mathijs Wintraecken. Braiding vineyards. *arXiv preprint arXiv:2504.11203*, 2025.

[CFSW25b] Erin Wolf Chambers, Christopher Fillmore, Elizabeth Stephenson, and Mathijs Wintraecken. Burning or collapsing the medial axis is unstable. *La Matematica*, August 2025.

[CK11] Ming-Ching Chang and Benjamin B. Kimia. Measuring 3D shape similarity by graph-based matching of the medial scaffolds. *Computer Vision and Image Understanding*, 115(5):707–720, May 2011.

[CL05] Frédéric Chazal and André Lieutier. The “ λ -medial axis”. *Graphical Models*, 67(4):304 – 331, 2005.

[CMS84] J. Cheeger, W. Müller, and R. Schrader. On the curvature of piecewise flat spaces. *Comm. Math. Phys.*, 92(3):405–454, 1984.

[CMS86] J. Cheeger, W. Müller, and R. Schrader. Kinematic and tube formulas for piecewise linear spaces. *Indiana University Mathematics Journal*, 35(4):737–754, 1986.

[Com18] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.

[CS68] M.W. Crofton and J.J. Sylvester. On the theory of local probability. *Trans. Royal Soc.*, 158:181–199, 1868.

[CS97] Daniel C. Cohen and Alexander I. Suciu. The braid monodromy of plane algebraic curves and hyperplane arrangements. *Commentarii Mathematici Helvetici*, 72(2):285–315, June 1997.

[CS04] F. Chazal and R. Soufflet. Stability and finiteness properties of medial axis and skeleton. *Journal of Dynamical and Control Systems*, 10(2):149–170, 2004.

[CSEH05] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. In *Proceedings of the twenty-first annual symposium on Computational geometry*, pages 263–271, 2005.

[CSEH08] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, April 2008.

[CSEM06] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the twenty-second annual symposium on Computational geometry*, SoCG06, page 119–126. ACM, June 2006.

[Că61] G. Călugăreanu. Sur les classes d’isotopie des noeuds tridimensionnels et leurs invariants. *Czechoslovak Mathematical Journal*, 11(4):588–625, 1961.

[Dam03] James Damon. Smoothness and geometry of boundaries associated to skeletal structures I: Sufficient conditions for smoothness. In *Annales de l’Institut Fourier*, volume 53, pages 1941–1985, 2003.

[Dam04] James Damon. Smoothness and geometry of boundaries associated to skeletal structures, II: Geometry in the Blum case. *Compositio Mathematica*, 140(6):1657–1674, 2004.

[Dam05] James Damon. Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision*, 63(1):45–64, 2005.

[Dam06] James Damon. The global medial structure of regions in \mathbb{R}^3 . *Geometry & Topology*, 10(4):2385–2429, 2006.

[Dam07] James Damon. Global geometry of regions and boundaries via skeletal and medial integrals. *Communications in Analysis and Geometry*, 15(2):307–358, 2007.

[dCD18] P.M. de Castro and O. Devillers. Expected length of the Voronoi path in a high dimensional Poisson-Delaunay triangulation. *Discret. Comput. Geom.*, 60(1):200–219, 2018.

[Deh11] M. Dehn. Über unendliche diskontinuierliche gruppen. *Mathematische Annalen*, 71(1):116–144, mar 1911.

[Deh12] M. Dehn. Transformation der kurven auf zweiseitigen flächen. *Mathematische Annalen*, 72(3):413–421, sep 1912.

[Del34] B. Delaunay. Sur la sphère vide. *Bulletin de l'Académie des Sciences de l'URSS. Classe des sciences mathématiques et na*, 1934(6):793–800, 1934.

[DG14] James Damon and Ellen Gasparovic. Medial/skeletal linking structures for multi-region configurations. *Memoirs of the American Mathematical Society*, 250, 02 2014.

[DHL⁺19] 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 *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1143–1151, 2019.

[dMRST21] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. The unbearable hardness of unknotting. *Advances in Mathematics*, 381:107648, April 2021.

[DPKLW24] Hana Dal Poz Kouřimská, André Lieutier, and Mathijs Wintraecken. The medial axis of any closed bounded set is Lipschitz stable with respect to the Hausdorff distance under ambient diffeomorphisms. In *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 69:1–69:18, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[DS06] Tamal K. Dey and Jian Sun. Defining and computing curve-skeletons with medial geodesic function. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*, SGP '06, page 143–152, Goslar, DEU, 2006. Eurographics Association.

[DW22] Tamal Krishna Dey and Yusu Wang. *Computational Topology for Data Analysis*. Cambridge University Press, February 2022.

[DZ04] Tamal K Dey and Wulue Zhao. Approximating the medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, 38(1):179–200, 2004.

[Ebe05] Wolfgang Ebeling. Monodromy. *arXiv preprint arXiv:0507171*, 2005.

[Ede86] H. Edelsbrunner. Edge-skeletons in arrangements with applications. *Algorithmica*, 1(1–4):93–109, November 1986.

[EFO25] Herbert Edelsbrunner, Christopher Fillmore, and Gonçalo Oliveira. Counting equilibria of the electrostatic potential. *arXiv preprint arXiv:2501.05315*, 2025.

[ELZ02] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, November 2002.

[EM90] Herbert Edelsbrunner and Ernst Peter Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, January 1990.

[EN25] H. Edelsbrunner and A. Nikitenko. Average and expected distortion of Voronoi paths and scapes. *Discret. Comput. Geom.*, 73:490–499, 2025.

[EP24] H. Edelsbrunner and J. Pach. Maximum Betti numbers of Čech complexes. In *40th International Symposium on Computational Geometry, SoCG 2024, June 11–14, 2024, Athens, Greece*, volume 293 of *LIPICS*, pages 53:1–53:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[Erd45] Paul Erdős. Some remarks on the measurability of certain sets. *Bulletin of the American Mathematical Society*, 51(10):728–731, 1945.

[Erd46] Paul Erdős. On the Hausdorff dimension of some sets in Euclidean space. *Bulletin of the American Mathematical Society*, 52(2):107 – 109, 1946.

[ES97] H. Edelsbrunner and N.R. Shah. Triangulating topological spaces. *Int. J. Comput. Geom. Appl.*, 7(4):365–378, 1997.

[EST26] Herbert Edelsbrunner, Elizabeth Stephenson, and Martin Hafskjold Thoresen. The mid-sphere cousin of the medial axis transform. In *Discrete Geometry and Mathematical Morphology*, pages 133–147, Cham, 2026. Springer Nature Switzerland.

[Euc56] Euclid. *The thirteen books of the elements, vol. 1*. Dover Books on Mathematics. Dover Publications, Mineola, NY, 2 edition, June 1956.

[Eul36] Leonhard Euler. Solutio problematis ad geometriam situs pertinentis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 8:128–140, 1736.

[Fed59] H. Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93:418–491, 1959.

[Fro90] Patrizio Frosini. A distance for similarity classes of submanifolds of a Euclidean space. *Bulletin of the Australian Mathematical Society*, 42(3):407–415, December 1990.

[Ful71] F. Brock Fuller. The writhing number of a space curve. *Proceedings of the National Academy of Sciences*, 68(4):815–819, 1971.

[Gau27] Carl Friedrich Gauss. *Disquisitiones generales circa superficies curvas*. Commentarii societatis regiae scientiarum Gottingensis recentiores, Göttingen, 1827.

[GKP08] Peter J Giblin, Benjamin B Kimia, and Anthony J Pollitt. Transitions of the 3D medial axis under a one-parameter family of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):900–918, 2008.

[GMPW09] Joachim Giesen, Balint Miklos, Mark Pauly, and Camille Wormser. The scale axis transform. In *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry*, page 106–115, New York, NY, USA, 2009. Association for Computing Machinery.

[GNS07] A. Gabrielov, D. Novikov, and B. Shapiro. Mystery of point charges. *Proceedings of the London Mathematical Society*, 95(2):443–472, September 2007.

[Had57] H. Hadwiger. Vorlesungen über Inhalt, Oberfläche und Isoperimetrie. Die Grundlehren der mathematischen Wissenschaften ; 93. Springer, Berlin, 1957.

[Har99] Godfrey H Hardy. Ramanujan. AMS Chelsea Publishing. American Mathematical Society, Providence, RI, November 1999.

[Has97] Joel Hass. Algorithms for recognizing knots and 3-manifolds. arXiv preprint math/9712269, 1997.

[HD86] Seng-Beng Ho and Charles R Dyer. Shape smoothing using medial axis properties. IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-8(4):512–520, 1986.

[Hic22] Abigail Hickok. Persistence diagram bundles: A multidimensional generalization of vineyards. arXiv preprint arXiv:2210.05124, 2022.

[INT23] Kazuhiro Ichihara, Yuya Nishimura, and Seiichi Tani. The computational complexity of classical knot recognition. Journal of Knot Theory and Its Ramifications, 32(11), October 2023.

[Kin95] J. F. C. Kingman. Poisson processes. J. F. C. Kingman. Clarendon Press, 1995.

[KL00] Konstantin Klenin and Jörg Langowski. Computation of writhe in modeling of supercoiled DNA. Biopolymers, 54(5):307–317, 2000.

[KLW24] Hana Dal Poz Kouřimská, André Lieutier, and Mathijs Wintraecken. A free lunch: manifolds of positive reach can be smoothed without decreasing the reach. arXiv preprint arXiv:2412.04899, 2024.

[KRdB97] D.A. Klain, G.C. Rota, and L.A.R. di Brozolo. Introduction to Geometric Probability. Lezioni Lincee. Cambridge University Press, 1997.

[KT21] Dale Koenig and Anastasiia Tsvietkova. NP-hard problems naturally arising in knot theory. Transactions of the American Mathematical Society, Series B, 8(15):420–441, May 2021.

[Lac21] Marc Lackenby. The efficient certification of knottedness and Thurston norm. Advances in Mathematics, 387:107796, August 2021.

[LC87] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery.

[LCLJ11] Lu Liu, Erin W. Chambers, David Letscher, and Tao Ju. Extended grassfire transform on medial axes of 2D shapes. Computer-Aided Design, 43(11):1496 – 1505, 2011. Solid and Physical Modeling 2011.

[Lef42] Solomon Lefschetz. Algebraic Topology. American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, July 1942.

[Lef49] Solomon Lefschetz. Introduction to Topology. Princeton University Press, December 1949.

[Let12] David Letscher. On persistent homotopy, knotted complexes and the Alexander module. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12, page 428–441. ACM, January 2012.

[Lev83] Michael Levitt. Protein folding by restrained energy minimization and molecular dynamics. Journal of Molecular Biology, 170(3):723–764, 1983.

[Lie04] André Lieutier. Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. Computer-Aided Design, 36(11):1029 – 1046, 2004. Solid Modeling Theory and Applications.

[Loo74] E. J. N. Looijenga. Structural Stability of Smooth Families of C^∞ functions. PhD thesis, University of Amsterdam, 1974.

[LW23] André Lieutier and Mathijs Wintraecken. Hausdorff and Gromov-Hausdorff stable subsets of the medial axis. In Proceedings of the 55th Annual ACM Symposium on Theory of Computing, pages 1768–1776, 2023.

[Mat83] John N Mather. Distance from a submanifold in Euclidean-space. In Proceedings of symposia in pure mathematics, volume 40, pages 199–216. American Mathematical Society, 1983.

[Max73] J.C. Maxwell. A Treatise on Electricity and Magnetism, Vol. 1. Clarendon Press, 1873.

[MC69] M. Morse and S. S. Cairns. Critical Point Theory in Global Analysis and Differential Topology: An introduction. Pure and Applied Mathematics. Academic Press, 1969.

[McM71] P. McMullen. On the upper-bound conjecture for convex polytopes. Journal of Combinatorial Theory, Series B, 10(3):187–200, 1971.

[Mil69] J. Milnor. Morse Theory. Princeton University Press, 1969.

[Mor08] Jean-Marie Morvan. Generalized Curvatures. Springer Publishing Company, Incorporated, 1 edition, 2008.

[MR92] Keith Moffatt and Renzo Ricca. Helicity and the Calugareanu invariant. Proceedings of The Royal Society A: Mathematical, Physical and Engineering Sciences, 439:411–429, 11 1992.

[MTB⁺15] Elizabeth Munch, Katharine Turner, Paul Bendich, Sayan Mukherjee, Jonathan Mattingly, and John Harer. Probabilistic Fréchet means for time varying persistence diagrams. Electronic Journal of Statistics, 9(1), January 2015.

[MW80] David Milman and Zeev Waksman. On topological properties of the central set of a bounded domain in \mathbb{R}^m . Journal of Geometry, 15(1):1–7, 1980.

[NCMRA24] Philip Neill, Natalie Crist, Ryan McGorty, and Rae Robertson-Anderson. Enzymatic cleaving of entangled DNA rings drives scale-dependent rheological trajectories. Soft Matter, 20(12):2750–2766, 2024.

[Net25] Applied Algebraic Topology Network. Sara scaramuccia (04/23/25): Monodromy in bi-parameter persistence modules, 2025.

[Noe21] Emmy Noether. Idealtheorie in ringbereichen. *Mathematische Annalen*, 83(1–2):24–66, March 1921.

[OBS⁺00] A. Okabe, B. Boots, K. Sugihara, S.N. Chiu, and D.G. Kendall. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Second Edition. Wiley Series in Probability and Mathematical Statistics. Wiley, 2000.

[OOT24] Adam Onus, Nina Otter, and Renata Turkes. Shoving tubes through shapes gives a sufficient and efficient shape statistic. [arXiv preprint arXiv:2412.18452](https://arxiv.org/abs/2412.18452), 2024.

[Oud15] Steve Oudot. *Persistence Theory: From Quiver Representations to Data Analysis*. American Mathematical Society, December 2015.

[PL20] Eric Priest and D. Longcope. The creation of twist by reconnection of flux tubes. *Solar Physics*, 295, 03 2020.

[Poi95] Henri Poincaré. Analysis situs. *Journal de l’École polytechnique*, 1:1–121, 1895.

[Por94] I R Porteous. *Geometric differentiation*. Cambridge University Press, Cambridge, England, June 1994.

[Rei27] Kurt Reidemeister. Elementare begründung der knotentheorie. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 5(1):24–32, December 1927.

[Rie68] B. Riemann. über die hypothesen, welche der geometrie zu grunde liegen. (mitgetheilt durch r. dedekind). *Abhandlungen der Königlichen Gesellschaft der Wissenschaften in Göttingen*, 13:133–152, 1868.

[Rob99] Vanessa Robins. Towards computing homology from approximations. *Topology Proceedings*, 24, 01 1999.

[Rub95] Joachim H. Rubinstein. An algorithm to recognize the 3-sphere. In S. D. Chatterji, editor, *Proceedings of the International Congress of Mathematicians*, pages 601–611, Basel, 1995. Birkhäuser Basel.

[Sal23] Nick Salter. Stratified braid groups: monodromy. [arXiv preprint arXiv:2304.04627](https://arxiv.org/abs/2304.04627), 2023.

[Sal24] Nick Salter. Monodromy of stratified braid groups, ii. *Research in the Mathematical Sciences*, 11(4), October 2024.

[SB98] Doron Shaked and Alfred M. Bruckstein. Pruning medial axes. *Computer Vision and Image Understanding*, 69(2):156 – 169, 1998.

[SBdB16] Punam K Saha, Gunilla Borgefors, and Gabriella Sanniti di Baja. A survey on skeletonization algorithms and their applications. *Pattern recognition letters*, 76:3–12, 2016.

[Sie99] Dirk Siersma. Properties of conflict sets in the plane. *Banach Center Publications*, 50:267–276, 1999.

[SKK04] T.B. Sebastian, P.N. Klein, and B.B. Kimia. Recognition of shapes by editing their shock graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):550–571, 2004.

[SSG99] Jorge Sotomayor, Dirk Siersma, and Ronaldo Garcia. Curvatures of conflict surfaces in Euclidean 3-space. *Banach Center Publications*, 50(1):277–285, 1999.

[Sta05] E. L. Starostin. *On the writhing number of a non-closed curve*, pages 525–545. World Scientific Publishing, 2005.

[Ste13] J. Steiner. Über parallele flächen. In *Jacob Steiner's Gesammelte Werke*, pages 171–176. Cambridge University Press, 2013.

[Ste23a] Elizabeth R Stephenson. *Generalizing medial axes with homology switches*. PhD thesis, Institute of Science and Technology Austria, 2023.

[Ste23b] Elizabeth R Stephenson. Generalizing medial axes with homology switches. Master's thesis, Institute of Science and Technology Austria, 2023.

[TDS⁺16] 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.

[Tho72] R. Thom. Sur le cut-locus d'une variété plongée. *Journal of Differential Geometry*, 6(4):577–586, 1972.

[Tho94] Abigail Thompson. Thin position and the recognition problem for \mathbb{S}^3 . *Mathematical Research Letters*, 1(5):613–630, 1994.

[TMB14] Katharine Turner, Sayan Mukherjee, and Doug M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 12 2014.

[Tra83] Bruce Trace. On the Reidemeister moves of a classical knot. *Proceedings of the American Mathematical Society*, 89(4):722–724, 1983.

[Tur23] Katharine Turner. Representing vineyard modules. *arXiv preprint arXiv:2307.06020*, 2023.

[VALFK79] Alexander Vologodskii, Vadim Anshelevich, AV Lukashin, and Maxim Frank-Kamenetskii. Statistical mechanics of supercoils and the torsional stiffness of the DNA double helix. *Nature*, 280:294–8, 08 1979.

[VdWSNI⁺14] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image processing in Python. *PeerJ*, 2:e453, 2014.

[VM97] A.V. Vologodskii and J.F. Marko. Extension of torsionally stressed DNA by external force. *Biophysical Journal*, 73(1):123–132, 1997.

[vM02] Martijn van Manen. The geometry of conflict sets. PhD thesis, Universiteit Utrecht, Faculteit der Wiskunde en Informatica, 2002.

[vM03] Martijn van Manen. Curvature and torsion formulas for conflict sets. Banach Center Publications, 62:209–222, 2003.

[vM07] M. van Manen. Singularities in Geometry and Topology, chapter Maxwell strata and Caustics, pages 787–824. WORLD SCIENTIFIC, 2007.

[Vog90] Pierre Vogel. Representation of links by braids: A new algorithm. Commentarii Mathematici Helvetici, 65(1):104–113, 1990.

[Vor08] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Journal für die reine und angewandte Mathematik, 133:97–178, 1908.

[Wal77] Charles T. C. Wall. Geometric properties of generic differentiable manifolds. In Geometry and Topology, pages 707–774, Berlin, Heidelberg, 1977. Springer Berlin Heidelberg.

[WCV12] Hubert Wagner, Chao Chen, and Erald Vuçini. Efficient Computation of Persistent Homology for Cubical Data, pages 91–106. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[Wey39] Hermann Weyl. On the volume of tubes. American Journal of Mathematics, 61(2):461–472, 1939.

[Whi44] Hassler Whitney. The self-intersections of a smooth n -manifold in $2n$ -space. The Annals of Mathematics, 45(2):220, April 1944.

[Whi49a] J. H. C. Whitehead. Combinatorial homotopy. i. Bulletin of the American Mathematical Society, 55(3):213–245, 1949.

[Whi49b] J. H. C. Whitehead. Combinatorial homotopy. ii. Bulletin of the American Mathematical Society, 55(5):453–496, 1949.

[Whi69] James H. White. Self-linking and the Gauss integral in higher dimensions. American Journal of Mathematics, 91(3):693–728, 1969.

[Wol92] Franz-Erich Wolter. Cut locus and medial axis in global shape interrogation and representation. In MIT Design Laboratory Memorandum 92-2 and MIT Sea Grant Report, 1992.

[Yam87] Shuji Yamada. The minimal number of Seifert circles equals the braid index of a link. Inventiones Mathematicae, 89(2):347–356, Jun 1987.

[YB03] Yousef, T. A. and Brandenburg, A. Relaxation of writhe and twist of a bi-helical magnetic field. A&A, 407(1):7–12, 2003.

[Yom81] Yosef Yomdin. On the local structure of a generic central set. Compositio Mathematica, 43(2):225–238, 1981.

[YSC⁺16] Yajie Yan, Kyle Sykes, Erin Chambers, David Letscher, and Tao Ju. Erosion thickness on medial axes of 3D shapes. ACM Transactions on Graphics, 35:38:1–38:12, July 2016.

[Ču20] Matija Čufar. Ripserer.jl: flexible and efficient persistent homology computation in Julia. *Journal of Open Source Software*, 5(54):2614, 2020.

