

Reeb Spaces and the Robustness of Preimages

by

Amit Patel

Department of Computer Science
Duke University

Ph.D. Dissertation
2010

Copyright © 2010 by Amit Patel
All rights reserved except the rights granted by the
Creative Commons Attribution-Noncommercial Licence

Abstract

We study how the preimages of a mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$ between manifolds vary under perturbations. First, we consider the preimage of a single point and track the history of its connected component as this point varies in \mathbb{Y} . This information is compactly represented in a generalization of the Reeb graph we call the Reeb space. We study its local and global properties and provide an algorithm for its construction. Using homology, we then consider higher dimensional connectivity of the preimage. We develop a theory quantifying the stability of each homology class under perturbations of the mapping f . This value, called *robustness*, is given to each homology class in the preimage. The robustness of a class is the magnitude of the perturbation necessary to remove it from the preimage. The generality of this theory allows for many applications. We apply this theory to quantify the stability of contours, fixed points, periodic orbits, and more.

Contents

Abstract	iii
List of Tables	vi
List of Figures	vii
Acknowledgements	viii
1 Introduction	1
1.1 Stability	1
1.2 Our Model	2
1.3 History and Contributions	2
2 Background	4
2.1 Smooth Mappings	4
2.1.1 Transversality	4
2.1.2 Oriented Intersection Number	6
2.1.3 Jet Bundles	7
2.1.4 The Whitney C^∞ Topology	8
2.1.5 Singularities	8
2.2 Simplicial Complexes and Homology	10
2.2.1 Simplicial Complexes	10
2.2.2 Homology	11
2.3 Piecewise Linear Mappings	12
2.3.1 PL Mappings	12
2.3.2 Upper and Lower Links	14
2.3.3 Jacobi Sets	14
2.3.4 Degree Mod 2	15
2.4 Persistence	16
2.4.1 Tame functions	16
2.4.2 Stability	17
2.4.3 Computation	18

3	Reeb Spaces	19
3.1	Definition	20
3.2	Local Structure	21
3.3	Global Structure	24
3.4	The Orientable 3-Manifold Case	27
3.5	Remarks	32
4	Robustness of Preimages	34
4.1	Well Groups	34
4.2	Well Modules	37
4.3	Well Diagrams	38
4.4	Stability	39
4.5	Closed Forms	43
	4.5.1 Equidimensional manifolds	44
	4.5.2 Real-valued Functions	45
4.6	Applications	46
	4.6.1 Fixed Points of Mappings	46
	4.6.2 Periodic Orbits	48
	4.6.3 Vector Fields	48
	4.6.4 Simplification of the Contour	49
4.7	Remarks	50
	Bibliography	53
	Biography	56

List of Tables

2.1	Indexing singular vertices	14
4.1	Example: well groups	36

List of Figures

2.1	Transverse intersections	5
2.2	The degree of a mapping	6
2.3	A function to the reals and a mapping to the plane	9
2.4	Preimages of PL mappings	13
2.5	Birth and death	16
2.6	Persistence diagrams	17
3.1	An example of a Reeb space	21
3.2	Subdividing a complex with respect to a preimage	23
3.3	Critical edges	28
3.4	A tube, its core, and its divider	29
3.5	Arcs	31
3.6	Nodes	32
4.1	Preimages and the graph of a mapping	35
4.2	The well module	37
4.3	The left filtration of a well module	38
4.4	Connecting two well modules	42
4.5	Bridges between well modules	43
4.6	Constructing a perturbation	45
4.7	The robustness of a wrinkle	50
4.8	Eroding wrinkles	50
4.9	A curve in the plane	52

Acknowledgements

First, I thank my friends and family for their patience especially in the last two and a half years.

Dmitriy, I enjoyed our conversations in Berlin especially those after a visit to Görlitzer Park. I thank my friends at Görlitzer Park.

Brit and Paul, thanks for keeping me sane while I was writing my thesis in Vienna. Paul, our daily field trips to Billa were very therapeutic.

John, thanks for pointing me to the right books. Valerio, thanks for allowing me to spend the summer in the UC Berkeley library reading.

Thank you Herbert for allowing me the freedom to explore. I am grateful for the opportunity to live in Berlin and Vienna. Berlin changed me forever.

1

Introduction

This dissertation is about mappings between manifolds and their preimages. There are two main contributions: Reeb spaces and well groups.

Given a mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$, the Reeb space is the quotient space identifying two points $x, y \in \mathbb{X}$ if $f(x) = f(y)$ and both points belong to the same component in the preimage $f^{-1}(f(x)) = f^{-1}(f(y))$. The Reeb space captures the connectivity relationship between components in preimages. We study some local and global properties of this topological space.

The second main contribution in this dissertation is the idea of well groups. Using homology, we look beyond connected components in preimages to higher dimensional connectivity, for example tunnels or voids. We are interested in answering the following question. Given a homology class in the preimage of a point in \mathbb{Y} , how much must we perturb our mapping f to remove it from the preimage? The well group makes this question precise and also answers it. The well groups are motivated by the desire for measurably stable features in the preimage.

We start this chapter with a few vague words on stability and the nature of reality. §1.2 assigns math words to the words of the previous section. We conclude this chapter with a few words on the evolution of this work.

1.1 Stability

The ideas expressed here are a combination of the qualitative ideas of René Thom [31, page 1] with the quantitative ideas of persistent homology [13].

A *form* is that which can be named.

A form is *stable* if it is measurable.

Only that which is stable can be named.

A form is stable.

A *scale* is a collection of measurements.

A form can not be named at a scale that does not include its measurement.

Reality is the creation, evolution, and destruction of forms at a given scale.

Science is the study of reality and its purpose is to predict the dynamics of form.

1.2 Our Model

The state space of a system is often modeled as a manifold. Probes assign to each state of the system a point in a another manifold. For example, the intensity of a magnetic field, temperature, and pressure at each point in space gives a mapping from space to \mathbb{R}^3 . Variation is introduced by changes to the parameters of the system or by the measurement device itself. If f is a smooth mapping from a manifold \mathbb{X} to a manifold \mathbb{Y} and \mathbb{A} is a submanifold of \mathbb{Y} , variations to the mapping f introduces dynamics to the inverse $f^{-1}(\mathbb{A})$. The *form* of interest is the homology of the inverse. Let C be the space of all smooth mappings from \mathbb{X} to \mathbb{Y} . Then there is a closed subset K of C , called the *catastrophe set*, such that if f lies outside K , small variations do not effect the homology of the inverse. The qualitative nature of the form does not change under small variations and in this sense, the form is stable. However, we are interested in quantifying stability. That is, how much variation is necessary to remove the form. The well groups *measure* the stability of each form.

1.3 History and Contributions

I started with almost no knowledge of topology. The popularity and simplicity of the Reeb graph made it a good candidate to start with. After reading about time-varying Reeb graphs [10], I started to think about Reeb spaces. The Reeb space is a generalization of the Reeb graph defined for continuous mappings between manifolds. Soon it was clear that the singularity set of the mapping plays an important role in the structure of the Reeb space. This motivated a closer look at the singularities of smooth mapping [18, 32]. A year down the road, I found a wealth of literature on Reeb spaces [2, 22, 25, 29]. Of course, the literature does not name this space the Reeb space. In fact, no one name seems to exist. The literature is restricted to Reeb spaces of smooth mappings. I decided to focus on piecewise linear mappings and their Reeb spaces. I studied both the local and global properties of the Reeb space and gave the first algorithm for its construction.

The next project was to study the contour of a shape. Given a shape in \mathbb{R}^3 and a plane orthogonal to the viewing direction, there is mapping, namely the orthogonal projection, that takes the shape to the plane. The contour is the set of critical values of this mapping. The goal was to find global properties of the contour that remain

stable under perturbations of the mapping. Imagine walking a path in the plane and looking at the preimage of each point along the way. A change in the preimage coincides with the path crossing the contour. In other words, the contour is crossed every time the stability of a component in the preimage goes to zero. This inspired the definition of the well group. However, there are an infinite number of well groups one for each real number. With the help of Dmitriy Morozov, I employed the new idea of zigzag persistence [3] enabling us to connect the sequence of well groups with homomorphisms. This allowed us to relate two well groups and talk about their properties.

We arrived at the well group by studying the contour of a shape, but the well group works for mappings between any pair of manifolds. This leads to more applications. For example, the set of fixed points of a mapping or a vector field can be recast as the preimage of a mapping. Using well groups, we measure the robustness of each fixed point. The robustness is the amount we must perturb the mapping or vector field to remove it. I talk about more applications in §4.

2

Background

We begin with an introduction to smooth mappings, homology, piecewise linear mappings, and persistence. We also use this chapter to establish notation.

2.1 Smooth Mappings

Many of the ideas introduced here can be found in introductory differential topology books. See [18, 19, 21] for more details. For simplicity, we assume all manifolds are without boundary.

2.1.1 Transversality

Let \mathbb{X} and \mathbb{Y} be differentiable manifolds. Two smooth mappings $f, g : \mathbb{X} \rightarrow \mathbb{Y}$ are *homotopic* if there is a smooth mapping $h : \mathbb{X} \times [0, 1] \rightarrow \mathbb{Y}$ such that $h_0(x) = h(x, 0)$ is $f(x)$ and $h_1(x) = h(x, 1)$ is $g(x)$, for every $x \in \mathbb{X}$. We say a property of f is *infinitesimally stable* if for every homotopy h with $h_0 = f$, there is a value $\varepsilon > 0$ such that each mapping h_t with $t < \varepsilon$ has the same property. For example, consider the function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined as $f(x) = x^3$. It has a root, at $x = 0$, and the property that it has at least one root is infinitesimally stable. Furthermore, the derivative of f has two roots both at 0. However, this property is not infinitesimally stable because one can find a homotopy, namely $f_t = x^3 - tx$, that immediately separates the two roots of the derivative. Since no measurement in the real world is perfectly determined, any property that is observable is infinitesimally stable.

Transversality is an infinitesimally stable property of the intersection between two manifolds. Roughly speaking, an intersection is transverse if the two submanifolds intersect in a nonzero angle. For example, consider two curves intersecting in the plane as shown in Figure 2.1. The intersection on the left is non-transverse



FIGURE 2.1: A non-transverse intersection on the left is transverse after a small perturbation of one of the curves.

because the two tangent lines at the intersection point coincide. However, the intersection on the right is transverse because the two tangent lines at each intersection point intersect in a nonzero angle and therefore span the entire plane. There is an arbitrarily small perturbation taking the non-transverse intersection to a transverse intersection.

We now give the formal definition of a transverse intersection. Letting $a = f(x)$, the *derivative*

$$Df(x) : T_x\mathbb{X} \rightarrow T_a\mathbb{Y}$$

of f at x is a linear map from the tangent space at x to the tangent space at a . The mapping f is *transverse* to a submanifold $\mathbb{A} \subseteq \mathbb{Y}$, denoted as $f \pitchfork \mathbb{A}$, if for each $x \in \mathbb{X}$ with $f(x) \in \mathbb{A}$, the image of the derivative of f at x together with the tangent space of \mathbb{A} at $a = f(x)$ spans the tangent space of \mathbb{Y} at a . More formally, $f \pitchfork \mathbb{A}$ if

$$Df(x)(T_x\mathbb{X}) + T_a\mathbb{A} = T_a\mathbb{Y},$$

for all $x \in \mathbb{X}$ with $f(x) = a \in \mathbb{A}$.

At this point, you might wonder how the definition of transverse mappings relates to the intersection of submanifolds. The inverse $f^{-1}(\mathbb{A})$ is homeomorphic to the intersection of two submanifolds in the product space $\mathbb{X} \times \mathbb{Y}$. Write

$$\text{gf } f = \{(x, y) \in \mathbb{X} \times \mathbb{Y} \mid y = f(x)\},$$

for the graph of f . Then $\text{gf } f \cap (\mathbb{X} \times \mathbb{A})$ maps homeomorphically to the preimage $f^{-1}(\mathbb{A})$ by a projection to the first factor. The two tangent spaces at every point in the intersection touch at a nonzero angle iff f is transverse to \mathbb{A} . Amazingly, if $f \pitchfork \mathbb{A}$, the inverse $f^{-1}(\mathbb{A})$ is a manifold of the same codimension as \mathbb{A} .

PREIMAGE THEOREM [19]. If $f : \mathbb{X} \rightarrow \mathbb{Y}$ is transverse to a submanifold $\mathbb{A} \subseteq \mathbb{Y}$, then the preimage $f^{-1}(\mathbb{A})$ is a smoothly embedded submanifold of \mathbb{X} . Moreover, the codimension of $f^{-1}(\mathbb{A})$ in \mathbb{X} equals the codimension of \mathbb{A} in \mathbb{Y} .




FIGURE 2.2: Construct a mapping from the 1-sphere to itself by taking the radial projection of the outside curve to the inside curve. Letting \mathbb{A} be a single point of the inside curve, the intersection number is -1 for the mapping on the left and 0 for the right.

2.1.2 Oriented Intersection Number

Given two ordered bases $\{u_1, \dots, u_m\}$ and $\{v_1, \dots, v_m\}$ on a common real vector space of dimension m , there is a unique linear isomorphism $A : \mathbb{U} \rightarrow \mathbb{V}$ such that $v_i = Au_i$, for each i . The two ordered bases are *equivalently oriented* if the sign of the determinant of A is positive. This partitions the set of all ordered bases into two classes. An *orientation* is an arbitrary assignment of a plus to one class and a minus to the other. The isomorphism A is called *orientation preserving* if the sign of its determinant is positive and *orientation reversing* otherwise. Now let \mathbb{X} be an m -dimensional manifold and assign to each tangent space $T_x\mathbb{X}$ an orientation. There is a local parameterization $\phi : \mathbb{R}^m \rightarrow \mathbb{X}$ for each point $x \in \mathbb{X}$ and its derivative $D\phi(a) : \mathbb{R}^m \rightarrow T_{\phi(a)}\mathbb{X}$ is a linear isomorphism between the two real vector spaces. We say ϕ is an *orientation preserving mapping* if for each $a \in \mathbb{R}^m$, the derivative $D\phi(a)$ is orientation preserving. An *orientation* of \mathbb{X} is an assignment of an orientation to each tangent space $T_x\mathbb{X}$ such that there is an orientation preserving mapping for each $x \in \mathbb{X}$.

Let $f : \mathbb{X} \rightarrow \mathbb{Y}$ be a smooth mapping between two compact orientable manifolds. Now let $\mathbb{A} \subset \mathbb{Y}$ be a closed orientable submanifold such that $\dim \mathbb{X} + \dim \mathbb{A} = \dim \mathbb{Y}$. If $f \bar{\cap} \mathbb{A}$, $f^{-1}(\mathbb{A})$ is a finite collection of points. Assign to each point $x \in f^{-1}(\mathbb{A})$ a plus one if $Df(x)(T_x\mathbb{X}) + T_{f(x)}\mathbb{A}$ is equivalently oriented with $T_{f(x)}\mathbb{Y}$, otherwise assign to x a minus one. The *intersection number*, $I(f, \mathbb{A})$, is the sum of the plus ones and minus ones in the inverse $f^{-1}(\mathbb{A})$. Figure 2.2 shows two mappings from a circle to itself. Choosing \mathbb{A} as a single point on the 1-sphere, the mapping on the left has intersection number -1 and the mapping on the right has intersection number 0 . It is not too hard to believe that the intersection number is homotopy invariant. That is, if g is homotopic to f , then $I(f, \mathbb{A}) = I(g, \mathbb{A})$ [19].

2.1.3 Jet Bundles

The jet bundle allows for a geometric understanding of the derivative. Combined with the idea of transversality, we begin to identify infinitesimally stable properties of smooth mappings.

Let $f, g : \mathbb{X} \rightarrow \mathbb{Y}$ be two smooth mappings taking an m -manifold to an n -manifold such that $f(x) = g(x) = a$. We say f has *first order contact* with g at x if the first derivatives $Df(x)$ and $Dg(x)$ equal. In general, f has *k -th order contact* with g at x if the derivative Df has $(k - 1)$ -st order contact with Dg at every point in $T_x\mathbb{X}$. We write $f \sim_k g$ at x if f has k -th order contact with g . $f \sim_k g$ at x iff the Taylor expansion of the two mappings in a common neighborhood of x are equal up to and including order k . See [18, page 37] for a proof.

Let $J^k(\mathbb{X}, \mathbb{Y})_{x,a}$ be the set of equivalence classes under the equivalence relation “ \sim_k at x .” In other words, identify two mappings f and g with $f(x) = g(x) = a$ if f has k -th order contact with g at x . The *k -th order jet bundle* $J^k(\mathbb{X}, \mathbb{Y})$ is the disjoint union of $J^k(\mathbb{X}, \mathbb{Y})_{x,a}$ over all pairs (x, a) . In symbols,

$$J^k(\mathbb{X}, \mathbb{Y}) = \bigsqcup_{(x,a) \in \mathbb{X} \times \mathbb{Y}} J^k(\mathbb{X}, \mathbb{Y})_{x,a}.$$

For example, for $k = 0$, the jet bundle is in one-to-one correspondence with $\mathbb{X} \times \mathbb{Y}$. Specifically, the point $(x, a) \in \mathbb{X} \times \mathbb{Y}$ corresponds to all mappings $f : \mathbb{X} \rightarrow \mathbb{Y}$ taking x to a . Similarly, for $k = 1$, the jet bundle is in one-to-one correspondence with $\mathbb{X} \times \mathbb{Y} \times \mathbb{R}^{mn}$ because \mathbb{R}^{mn} parameterizes the set of linear mappings from \mathbb{R}^m to \mathbb{R}^n . This can be generalized to $k > 1$, as discussed shortly.

An element σ of $J^k(\mathbb{X}, \mathbb{Y})$ is called a *k -jet*. The k -jet σ belongs to $J^k(\mathbb{X}, \mathbb{Y})_{x,a}$ for some pair of points (x, a) . This gives canonical mappings to \mathbb{X} and \mathbb{Y} called the *source* and *target* mappings $\alpha : J^k(\mathbb{X}, \mathbb{Y}) \rightarrow \mathbb{X}$ and $\beta : J^k(\mathbb{X}, \mathbb{Y}) \rightarrow \mathbb{Y}$ defined by $\alpha(\sigma) = x$ and $\beta(\sigma) = a$. The *spray* of f is the mapping $j_k f : \mathbb{X} \rightarrow J^k(\mathbb{X}, \mathbb{Y})$ that takes each point $x \in \mathbb{X}$ to the k -jet representing the equivalence class of f in $J^k(\mathbb{X}, \mathbb{Y})_{x,f(x)}$. The spray is a section of the jet bundle. For $k = 0$, the image of the spray has codimension n and for $k = 1$, its codimension is $n + mn$.

We now give a feeling for the topology on the jet bundle. Again, we refer the reader to [18] for details. Each k -jet is determined by its source, target, and a Taylor expansion up to and including order k , ignoring the constant term. Fix a source $x \in \mathbb{X}$ and a target $a \in \mathbb{Y}$ and let f be a mapping taking x to a . Now decompose f into its n component functions in some neighborhood of x . This gives n Taylor expansions each a polynomial in m variables of degree at most k with the constant term equal to zero. Let \mathbf{A}_m^k be the vector space of polynomials in m variables of degree at most k with the constant term equal to zero. For example, \mathbf{A}_m^1 is isomorphic to \mathbb{R}^m , while \mathbf{A}_m^2 is isomorphic to $\mathbb{R}^{m + \binom{m+1}{2}}$. Let $\mathbf{B}_{m,n}^k = \bigoplus_{i=1}^n \mathbf{A}_m^k$ be n copies of \mathbf{A}_m^k . The reason for defining $\mathbf{B}_{m,n}^k$ is that we have a bijection between $J^k(\mathbb{X}, \mathbb{Y})_{x,a}$ and the vector space $\mathbf{B}_{m,n}^k$. Placing a copy of $\mathbf{B}_{m,n}^k$ above each pair of source and target point

pair, the jet bundle is a manifold of dimension

$$\dim J^k(\mathbb{X}, \mathbb{Y}) = m + n + \dim \mathbf{B}_{m,n}^k.$$

To get a feeling for the dimensionality of this space, we note that the dimension of $\mathbf{B}_{m,n}^k$ is 0 for $k = 0$, mn for $k = 1$, $(m + \binom{m+1}{2})n$ for $k = 2$, and $(m + \binom{m+1}{2} + \binom{m+2}{3})n$ for $k = 3$. The mappings α , β , and $j_k f$ are smooth.

2.1.4 The Whitney C^∞ Topology

The jet bundle induces a topology on the space of all smooth mappings $C^\infty(\mathbb{X}, \mathbb{Y})$. Letting U be an open subset of $J^k(\mathbb{X}, \mathbb{Y})$, define

$$M_U = \{f \in C^\infty(\mathbb{X}, \mathbb{Y}) \mid j_k f(\mathbb{X}) \subset U\}.$$

The collection of sets $\{M_U\}$ for open subsets U of $J^k(\mathbb{X}, \mathbb{Y})$ forms a basis for a topology called the *Whitney C^k topology*. Let \mathcal{W}_k be the Whitney C^k topology on $C^k(\mathbb{X}, \mathbb{Y})$. The *Whitney C^∞ topology* on $C^\infty(\mathbb{X}, \mathbb{Y})$ is the topology generated by the basis $\cup_{k=0}^\infty \mathcal{W}_k$.

Surprisingly, the sprays of most mappings are transverse to a fixed submanifold of the jet bundle. This is essentially the statement of Thom's Transversality Theorem. A few more definitions are required before we can state the theorem. A subset of a topological space \mathbb{F} is *residual* if it is the countable intersection of open dense subsets of \mathbb{F} . We say \mathbb{F} is a *Baire space* if every residual subset is dense [18, page 44]. The space $C^\infty(\mathbb{X}, \mathbb{Y})$ with the Whitney C^∞ topology is a Baire space.

THOM'S TRANSVERSALITY THEOREM [18]. Let \mathbb{X} and \mathbb{Y} be manifolds and \mathbb{W} a submanifold of $J^k(\mathbb{X}, \mathbb{Y})$. The set

$$T_{\mathbb{W}} = \{f \in C^\infty(\mathbb{X}, \mathbb{Y}) \mid j_k f \not\cap \mathbb{W}\}$$

is a residual subset in the Whitney C^∞ topology. If \mathbb{W} is closed, then the set $T_{\mathbb{W}}$ is open.

Thom's Transversality Theorem is useful in arguing about generic properties of smooth mappings.

2.1.5 Singularities

A point $x \in \mathbb{X}$ is a *singular point* of the smooth mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$ if the first derivative $Df(x)$ at x is not surjective. A point $a \in \mathbb{Y}$ is a *regular value* of f if $f \not\cap a$ otherwise, a is a *singular value*. Equivalently, a is a singular value of f if it is the image of a singular point. Define the *singular set* of f , Σf , as the set of singular points of f . Often a singular point is called a *critical point* and a singular value a *critical value*. We will use "singular" and "critical" interchangeably.




FIGURE 2.3: Shown is a height function on a surface (left) and a mapping taking a surface to the plane (right). The singular set of the height function is the set of points where the gradient vanishes. For the mapping to the plane, the singular set is the curve where the surface folds over on itself in the image.

Consider the height function on the surface shown in Figure 2.3 on the left. The singular points are where the gradient of f vanishes. Also shown in Figure 2.3 is a mapping taking a surface to the plane. The singular points are where the surface folds over on itself in the image. One can easily imagine functions and mappings in which the singular set has dimension higher than in the example in Figure 2.3. For example, the singular set of the height function may have an entire line of points where the gradient vanishes. Similarly, the mapping to the plane may take an entire surface patch to a curve in the plane. However, in both cases, there are arbitrarily close mappings with nice singular sets. We use Thom's Transversality Theorem to shed some light on the generic behavior of the singular set. We assume the dimension m of \mathbb{X} is at least the dimension n of \mathbb{Y} , otherwise the singular set is the entire space \mathbb{X} .

Recall that a 1-jet $\sigma \in J^1(\mathbb{X}, \mathbb{Y})$ is an equivalence class of mappings in $C^\infty(\mathbb{X}, \mathbb{Y})$. Letting $x = \alpha(\sigma)$ be the source of σ , any two mappings f and g in σ satisfy $f(x) = g(x)$ and $Df(x) = Dg(x)$. Now take a mapping f in σ . Define the *corank* of σ as the corank of the image of the first derivative $Df(x)$. Recall that the spray $j_1 f : \mathbb{X} \rightarrow J^1(\mathbb{X}, \mathbb{Y})$ takes each point $x \in \mathbb{X}$ to the equivalence class of f in $J^1(\mathbb{X}, \mathbb{Y})_{x, f(x)}$. A point x belongs to the singular set $\Sigma(f)$ iff the corank of $j_1 f(x)$ is greater than zero. Let us introduce the *singular manifold of order r* ,

$$\mathbb{S}_r = \{\sigma \in J^1(\mathbb{X}, \mathbb{Y}) \mid \text{corank } \sigma = r\},$$

as the collection of jets with corank r . As described in [18], the singular manifold \mathbb{S}_r is a submanifold of $J^1(\mathbb{X}, \mathbb{Y})$ with codimension $r(m - n + r)$. The singular set of f is then

$$\Sigma(f) = j_1 f^{-1} \left(\bigcup_{r \geq 1} \mathbb{S}_r \right).$$

Now assume the spray of f is transverse to \mathbb{S}_r . By the Preimage Theorem, the

inverse $j_1 f^{-1}(\mathbb{S}_r)$ is a submanifold of \mathbb{X} with codimension $r(m - n + r)$. Define

$$T_r = \{f \in C^\infty(\mathbb{X}, \mathbb{Y}) \mid j^1 f \bar{\cap} \mathbb{S}_r\}$$

as the set of mappings whose sprays are transverse to \mathbb{S}_r . By Thom's Transversality Theorem, the set T_r is a residual subset of $C^\infty(\mathbb{X}, \mathbb{Y})$ for each r . Recall $C^\infty(\mathbb{X}, \mathbb{Y})$ with the Whitney C^∞ topology is a Baire space. This implies the intersection $T = \bigcap_{r \geq 1} T_r$ is a residual subset and therefore dense. In other words, the set of mappings whose sprays are transverse to \mathbb{S}_r for every r is dense. We say f is *generic* if f belongs to T .

For a sanity check, consider the case $\mathbb{Y} = \mathbb{R}$. If the spray of f is transverse to \mathbb{S}_r , then the inverse $j_1 f^{-1}(\mathbb{S}_r)$ is a submanifold of codimension $r(m - 1 + r)$. The singular manifold \mathbb{S}_r is empty for $r > 1$ and of codimension m for $r = 1$. This means the singular set of a generic function f is a collection of isolated points. A function f is *Morse* if $j_1 f \bar{\cap} \mathbb{S}_1$. In other words, f is Morse iff it is generic.

Now consider the case $\mathbb{Y} = \mathbb{R}^2$. If the spray of f is transverse to \mathbb{S}_r , then the inverse $j_1 f^{-1}(\mathbb{S}_r)$ has codimension greater than m for $r > 1$ and codimension $m - 1$ when $r = 1$. This implies $j_1 f^{-1}(\mathbb{S}_r)$ is empty for $r > 1$. For generic mappings to the plane, the singular set is a smoothly embedded 1-submanifold.

Now consider a generic mapping $f : \mathbb{X} \rightarrow \mathbb{R}^4$. The codimension of the inverse $j_1 f^{-1}(\mathbb{S}_1)$ is $m - 3$ and the codimension of $j_1 f^{-1}(\mathbb{S}_2)$ is $2m - 4$. When $m = 4$, $j_1 f^{-1}(\mathbb{S}_2)$ is empty or of dimension zero. As a consequence, the singular set of a generic mapping need not be a smoothly embedded submanifold.

2.2 Simplicial Complexes and Homology

A simplicial complex is a special type of topological space obtained by gluing together vertices, edges, triangles, and so on. Simplicial homology takes a simplicial complex and makes precise the idea of holes in the space.

2.2.1 Simplicial Complexes

An *i-simplex* σ is the convex hull of $i+1$ affinely independent points in some Euclidean space. Letting u_0, u_1, \dots, u_i be the points, σ is the set of *convex combinations*, that is, points $\sum s_j u_j$ with $\sum s_j = 1$ and $s_j \geq 0$ for all $0 \leq j \leq i$. The *interior* of σ consists of the convex combinations for which all the s_j are strictly positive. The *dimension* of the simplex is $\dim \sigma = i$, which is at most the dimension of the ambient Euclidean space. A *face* of σ is spanned by a non-empty subset of the $i + 1$ points. All faces are *proper* except for σ which is an *improper* face of itself. The *boundary* of the simplex, denoted as $\partial\sigma$, is the union of all its proper faces. If τ and v are two disjoint faces of σ with $\dim \tau + \dim v = \dim \sigma - 1$ then σ is the *join* of the two, $\sigma = \tau * v$, meaning it is the union of line segments connecting points of τ with points of v . Any two of these line segments are either equal, disjoint or meet at a common endpoint.

A *simplicial complex* is a finite set of simplices K such that every face of a simplex in K belongs to K and the intersection of any two simplices in K is either empty or a face of both. K is an m -*complex* if the largest dimension of any of its simplices is m . The *underlying space* of K is the union of the simplices, $|K| = \bigcup_{\sigma \in K} \sigma$, together with the subspace topology inherited from the ambient space. Avoiding any possible confusion, we will sometimes blur the distinction between a complex and its underlying space. A *subcomplex* is a simplicial complex $L \subseteq K$. L is called *full subcomplex* of K if it contains every simplex of K whose vertices lie in L . For every non-negative integer $i \leq m$, the i -*skeleton*, denoted as $K^{(i)}$, is the largest subcomplex of dimension i ; it consists of all simplices of dimension i or less in K . The 0-skeleton is often referred to as the *vertex set*, $\text{Vert } K = K^{(0)}$. The *star* of a simplex σ , denoted as $\text{St } \sigma$, is the set of simplices in K that have σ as a face. We get the *closed star* if we add all faces of simplices in the star. The *link* of σ , denoted as $\text{Lk } \sigma$, consists of all simplices in the closed star that have an empty intersection with σ . Note that the closed star and the link are complexes while the star is generally not a complex. A *subdivision* of K is a simplicial complex with the same underlying space for which every simplex is contained in a simplex in K . Particularly useful is the *barycentric subdivision*, which we denote as $\text{Sd } K$. To describe it, we recall that the *barycenter* of an i -simplex is the average of its $i + 1$ vertices. The barycenters of the simplices in K form the vertex set of $\text{Sd } K$ and a subset of the barycenters spans a simplex iff the corresponding simplices in K form a chain in which every simplex is a proper face of the next in the sequence.

We say K *triangulates* a topological space homeomorphic to its underlying space. If K triangulates an m -manifold then every point of $|K|$ has a neighborhood homeomorphic to \mathbb{R}^m . However, this does not imply that the link of every i -simplex triangulates a sphere of dimension $m - i - 1$. A counterexample to this seemingly plausible property can be found in Edwards [16], see also [30]. We call K a *combinatorial m -manifold* if it satisfies this stronger property, that is, the link of every vertex triangulates the $(m - 1)$ -sphere and is itself a combinatorial $(m - 1)$ -manifold.

2.2.2 Homology

Homology is an algebraic language making precise the idea of holes in a topological space. We define the homology groups of a finite simplicial complex using $\mathbb{Z}/2\mathbb{Z}$ coefficients.

A p -*chain* is a formal sum $a_0\sigma_0 + a_1\sigma_1 + \dots + a_m\sigma_m$ of all p -simplices, with $a_i \in \{0, 1\}$. The sum of two p -chains is another p -chain obtained by adding the corresponding coefficients using $\mathbb{Z}/2\mathbb{Z}$ arithmetic. The p -*th chain group*, denoted as $C_p(K)$, is the abelian group of p -chains. The *boundary* of a p -simplex is the sum of its $(p - 1)$ -dimensional faces and therefore an element of $C_{p-1}(K)$. The p -th boundary homomorphism, $\partial_p : C_p(K) \rightarrow C_{p-1}(K)$, takes each p -chain $c \in C_p(K)$ to the sum of the $(p - 1)$ -faces of the simplices in c . A p -chain c is a *cycle* if its boundary vanishes. The set of cycles, $Z_p(K)$, is the kernel of ∂_p and forms a subgroup of $C_p(K)$. The

boundary of a boundary is always zero. In other words, $\partial_{p-1} \circ \partial_p = 0$. This implies the image of ∂_{p+1} , denoted as $B_p(K)$, is a subgroup of $Z_p(K)$. The p -th *homology group* is defined as the quotient group $H_p(K) = Z_p(K)/B_p(K)$. For $\mathbb{Z}/2\mathbb{Z}$ arithmetic, the homology group $H_p(K)$ is a vector space and its rank, the dimension of the vector space, is the p -th *Betti number* denoted as $\beta_p(K)$.

We now define relative homology. Letting L be a subcomplex of K , the chain group $C_p(L)$ is a subgroup of $C_p(K)$. The quotient group $C_p(K, L) = C_p(K)/C_p(L)$ is the *relative p -th chain group* of the pair (K, L) and its elements are called *relative p -chains*. An element c of $C_p(K, L)$ is typically written $c = d + C_p(L)$ where d is a p -chain containing only those p -simplices in $K - L$. The boundary homomorphism on $C_p(K)$ induces a homomorphism $\partial_p : C_p(K, L) \rightarrow C_{p-1}(K, L)$ between the relative chain groups. A relative p -chain is a *relative p -cycle* if its boundary is an element of $C_p(L)$. A relative p -chain is a *relative p -boundary* if its in the image of ∂_{p+1} . As before, call $Z_p(K, L)$ the kernel of ∂_p and $B_p(K, L)$ the image of ∂_{p+1} . The p -th *relative homology group* is the quotient group $H_p(K, L) = Z_p(K, L)/B_p(K, L)$.

2.3 Piecewise Linear Mappings

A necessary step towards computing properties of general smooth mappings is their discretization. This section introduces piecewise linear mappings and its singularity set, called the Jacobi set.

2.3.1 PL Mappings

Let u_1, u_2, \dots, u_n be the vertices of a simplicial complex K , σ a simplex in K , and x a point of σ . Recall that x is a unique convex combination of the vertices of σ so we can write $x = \sum_{j=1}^n s_j u_j$ with $\sum_{j=1}^n s_j = 1$, $s_j \geq 0$ for all j , and $s_j = 0$ unless u_j is a vertex of σ . The s_j are unique and are called the *barycentric coordinates* of x . We use them to extend a vertex map $\hat{f} : \text{Vert } K \rightarrow \mathbb{R}^n$ by piecewise linear interpolation to a *piecewise linear* or *PL mapping* $f : |K| \rightarrow \mathbb{R}^n$ defined as

$$f(x) = \sum_{j=1}^n s_j \hat{f}(u_j).$$

By construction, the restriction of f to a simplex of K is linear.

We call f a *generic PL mapping* if the images of the vertices have no structural properties that can be removed by arbitrarily small perturbations of the vertex map. In particular, we call $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping only if the restrictions of f to simplices of dimension n or less are injective. That is, the image of every simplex of dimension $i \leq n$ is an i -simplex.

The following theorem is analogous to the Preimage Theorem for smooth mappings.




FIGURE 2.4: We see a vertical edge and the corresponding 1-ball obtained by connecting its midpoint to the respective third vertices of the two triangles in the star. Connecting every point of the 1-ball to the endpoints of the edge gives a decomposition of the closed star. Pieces of the decomposing line segments define a homeomorphism between the 1-ball and a portion of the preimage $f^{-1}(a)$.

PL PREIMAGE THEOREM. Let K be a combinatorial m -manifold, $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping, and $a \in \mathbb{R}^n$ a point not in the image of the $(n - 1)$ -skeleton of K . Then $f^{-1}(a)$ is either empty or a manifold of dimension $m - n$.

PROOF. For $m < n$ the preimage of a is empty and for $m = n$ it is either empty or a finite set of points. In both cases there is nothing left to show. We therefore assume $m > n$ for the remainder of this proof.

Let σ be an i -simplex in K . Since K is a combinatorial m -manifold, the link of σ triangulates a sphere of dimension $m - i - 1$. Letting u be the barycenter of σ , we construct $B_\sigma = u * |\text{Lk } \sigma|$ by drawing a line segment from u to every point in the link. Clearly, B_σ is a PL ball of dimension $m - i$. We further draw a line segment between every point of B_σ and every point of the boundary of σ , as sketched in Figure 2.4. Any two of these line segments are either disjoint or meet at a common endpoint, which is either in B_σ or in $\partial\sigma$. Together, the line segments decompose the closed star of σ .

Next we show that for $i = n$, the portion of $f^{-1}(a)$ inside the closed star of σ is homeomorphic to B_σ . Equivalently, every vertex of the preimage has a neighborhood homeomorphic to \mathbb{R}^{m-n} . This implies that $f^{-1}(a)$ is indeed an $(m - n)$ -manifold. Let σ be a n -simplex in K that contains a point u_σ with $f(u_\sigma) = a$. Because $f^{-1}(a)$ avoids the $(n - 1)$ -skeleton of K , u_σ belongs to the interior of σ . Let $\tau \in \text{St } \sigma$ and let v be its maximal face disjoint from σ . Hence $\sigma * v = \tau$ and $u * v$ is the contribution of τ to B_σ . Letting j be the dimension of τ we have $\dim v = j - n - 1$ and $\dim(u * v) = j - n$. Furthermore, $f^{-1}(a)$ intersects τ in a polytope of dimension $j - n$. The line segments in the decomposition of $B_\sigma * \partial\sigma$ define a piecewise linear homeomorphism from $u * v$ to this polytope; see Figure 2.4. The collection of such then gives a homeomorphism from B_σ to the intersection of $f^{-1}(a)$ with the closed star of σ . \square

2.3.2 Upper and Lower Links

Suppose $h : |K| \rightarrow \mathbb{R}$ is a generic PL function on K . This implies $h(u_i) \neq h(u_j)$ whenever u_i and u_j are the two endpoints of an edge in K . We define the *lower link* of a vertex u_j as the collection of simplices in the link whose vertices all have smaller function value than u_j . Symmetrically, the *upper link* is the collection of simplices in the link whose vertices have larger function value:

$$\begin{aligned} \text{Lk}_- u_j &= \{\sigma \in \text{Lk } u_j \mid a \in \sigma \Rightarrow h(a) < h(u_j)\}; \\ \text{Lk}^+ u_j &= \{\tau \in \text{Lk } u_j \mid a \in \tau \Rightarrow h(a) > h(u_j)\}. \end{aligned}$$

Assuming K is an m -manifold, the link of u_j is a triangulation of the $(m - 1)$ -dimensional sphere, \mathbb{S}^{m-1} . The lower and upper links are full subcomplexes of this triangulation. Note that their union is not necessarily the entire link as there are simplices that have some of their vertices with higher value, and some with lower value.

We measure the way the lower link is connected to the rest of the complex using reduced homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients. Following the usual convention, we write $\tilde{\beta}_i$ for the rank of the dimension i reduced homology group. Denoting the ranks of the non-reduced homology groups by β_i , we have $\tilde{\beta}_i = \beta_i$ unless $i < 1$. Furthermore $\tilde{\beta}_0 = \beta_0 - 1$ and $\tilde{\beta}_{-1} = 0$ unless the lower link is empty in which case we have $\tilde{\beta}_0 = \beta_0 = 0$ and $\tilde{\beta}_{-1} = 1$. All $\tilde{\beta}_i$ are non-negative integers. We call u_j a *regular vertex* of h if all reduced Betti numbers of its lower link vanish and a *singular* or *critical vertex*, otherwise. It is a *simple* singular vertex if $\sum \tilde{\beta}_i = 1$. Simple singular points are conveniently classified by the *index* that exceeds the dimension of the non-zero reduced homology group by one; see Table 2.1. For $m = 3$, it is common to refer to simple singular vertices of index 0, 1, 2, 3 as *minima*, *1-saddles*, *2-saddles*, *maxima*.

type	index	$\tilde{\beta}_{-1}$	$\tilde{\beta}_0$	$\tilde{\beta}_1$	$\tilde{\beta}_2$
regular		0	0	0	0
minimum	0	1	0	0	0
1-saddle	1	0	1	0	0
2-saddle	2	0	0	1	0
maximum	3	0	0	0	1

Table 2.1: A simple singular vertex of index i is characterized by $\tilde{\beta}_{i-1} = 1$ and $\tilde{\beta}_j = 0$ for all $j \neq i - 1$.

2.3.3 Jacobi Sets

We now return to a multivariate generic PL mapping $f : |K| \rightarrow \mathbb{R}^n$. Following [8], we consider all linear combinations of the components of f . Let \vec{u} be a unit vector

in \mathbb{S}^{n-1} and consider the PL function $h_{\vec{u}} : |K| \rightarrow \mathbb{R}$ defined by

$$h_{\vec{u}}(x) = \langle f(x), \vec{u} \rangle,$$

the height of the image of the point x in the direction \vec{u} . Assuming $h_{\vec{u}}$ is constant on a simplex τ in K , we can define its *lower link* the same way as for a vertex, namely as the collection of simplices in the link whose vertices have function value less than the points of τ . The *upper link* of τ is similarly defined. Assuming the upper and lower links exhaust all vertices of $\text{Lk } \tau$, we use the reduced homology of the lower link to decide whether τ is *regular* or *singular* for $h_{\vec{u}}$, and if it is singular whether or not it is *simple*.

If τ is an $(n-1)$ -simplex, then there are exactly two unit vectors for which the height functions they define are constant on τ , namely the unit normals \vec{u} and $-\vec{u}$ of the image of τ in \mathbb{R}^n . The lower link of τ under one height function is its upper link under the other. Recall the link of τ is an $(m-n)$ -sphere. Alexander duality says that if a submanifold of the sphere has nontrivial reduced Betti numbers, then so does its complement. This implies τ is singular for $h_{\vec{u}}$ iff it is singular for $h_{-\vec{u}}$. In other words, τ has only one chance to be singular. Finally, we define the *Jacobi set* of f as the collection of singular $(n-1)$ -simplices together with their faces. These simplices form a subcomplex of K which we denote as J_f .

2.3.4 Degree Mod 2

Consider a mapping between compact manifolds of the same dimension. Two examples are the mappings described in Figure 2.2. Roughly speaking, the mod 2 degree of a mapping is zero if the mapping covers the range space an even number of times and one if it covers the image an odd number of times. The advantage of restricting ourselves to zero and one as opposed to the integers is that the manifolds involved need not be orientable. We use homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients to define the degree mod 2 of a mapping.

Let K and L be triangulations of two compact manifolds without boundary. As before, we assume the two manifolds have the same dimension, m . Using $\mathbb{Z}/2\mathbb{Z}$ coefficients, the top dimensional homology groups $H_m(K)$ and $H_m(L)$ are both isomorphic to $\mathbb{Z}/2\mathbb{Z}$. A *simplicial mapping* $f : K \rightarrow L$ taking simplices in K to simplices in L induces a homomorphism $f_m : H_m(K) \rightarrow H_m(L)$ between the two homology groups. The *degree mod 2* of f , $\deg f$, is the image of f_m which is 0 or 1.

For our purposes, the degree of a mapping between to pairs of spaces will be more useful. Let B^c be the complement of an open ball in \mathbb{R}^m and $f : |K| \rightarrow \mathbb{R}^m$ a continuous PL mapping. Now let C be a component of $f^{-1}(B)$. We think of C as an m -chain and assume its boundary is empty or if not empty, maps to B^c via f . This means the group $H_m(C, \partial C)$ is isomorphic to $\mathbb{Z}/2\mathbb{Z}$. The group $H_m(\mathbb{R}^m, B^c)$ is also isomorphic to $\mathbb{Z}/2\mathbb{Z}$. We define the degree mod 2 of the mapping f restricted to C as the image of the homomorphism $f'_m : H_m(C, \partial C) \rightarrow H_m(\mathbb{R}^m, B^c)$. An important and later useful property of the degree mod 2 is that it is homotopy invariant. That




FIGURE 2.5: A class α is born at \mathbb{X}_i and dies entering \mathbb{X}_j .

is, if one is to continuously deform a mapping $f : (C, \partial C) \rightarrow (\mathbb{R}^n, B^c)$ making sure that ∂C maps to B^c at all times, then the degree mod 2 does not change. Note that if the inverse $f^{-1}(0)$ does not intersect the component C , then the degree mod 2 of f restricted to C is zero.

2.4 Persistence

Persistent homology [5, 9, 13, 33] is an algebraic language that tracks the history of homology classes in a filtration of a topological space. This history is compactly represented in a diagram called the persistence diagram. An important property of persistence is that the persistence diagram is stable.

To simplify notation, we suppress dimensions and refer to $H(\mathbb{X})$ as the direct sum $\bigoplus H_p(\mathbb{X})$ of all the homology groups.

2.4.1 Tame functions

Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be a continuous function on a compact manifold \mathbb{X} . Denote by \mathbb{X}_r the sublevel set $f^{-1}(-\infty, r]$. For $r \leq s$, the inclusion of \mathbb{X}_r into \mathbb{X}_s induces a homomorphism $f_{r,s} : H(\mathbb{X}_r) \rightarrow H(\mathbb{X}_s)$ between the two homology groups. A value r is a *homological critical value* of f if for every sufficiently small $\delta > 0$, the homomorphism $f_{r-\delta, r+\delta}$ is not an isomorphism. The function f is *tame* if the homology group of each sublevel set has finite rank and there are only finitely many homological critical values.

Choose $a_0 < a_1 < \dots < a_\ell$ regular values straddling the ℓ homological critical values, $r_1 < \dots < r_\ell$, of f . Letting $\mathbb{X}_i = f^{-1}(-\infty, a_i]$, we have the following sequence of homology groups connected by maps induced by inclusion:

$$0 = H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \dots \rightarrow H(\mathbb{X}_\ell) = H(\mathbb{X}). \quad (2.1)$$

We say a homology class α is *born* at \mathbb{X}_i if it does not belong to the image of $f_{i-1,i}$. The class α *dies entering* \mathbb{X}_j if the image of $f_{i-1,j-1}$ does not contain $f_{i,j-1}(\alpha)$ but the image of $f_{i-1,j}$ does contain $f_{i,j}(\alpha)$. Note that if a class α is born at \mathbb{X}_i , then every class in the coset $[\alpha] = \alpha + \text{im } f_{i-1,i}$ is born at \mathbb{X}_i . Similarly, if the class α dies entering \mathbb{X}_j , the entire coset $[\alpha]$ dies entering \mathbb{X}_j . See Figure 2.5.

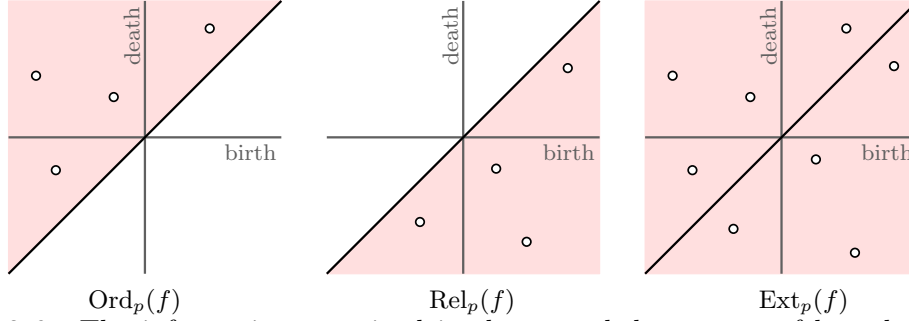


FIGURE 2.6: The information contained in the extended sequence of homology groups (2.2) is represented in the above three diagrams. The points in each diagram are restricted to the shaded region. For reasons of stability, each diagram contains the diagonal.

The sequence in (2.1) starts with zero but ends in a possibly nonzero group. This allows for the possibility of a class that never dies. We call these *essential* classes as they represent the actual homology of the space. To define the persistence of an essential class, we follow [5] and extend (2.1) using relative homology groups. Define the *superlevel set* \mathbb{X}^i as $f^{-1}[a_{\ell-i}, \infty]$. For $i < j$, the inclusion $\mathbb{X}^i \hookrightarrow \mathbb{X}^j$ induces a homomorphism $H(\mathbb{X}, \mathbb{X}^i) \rightarrow H(\mathbb{X}, \mathbb{X}^j)$. We have the following extended sequence of homology groups:

$$0 = H(\mathbb{X}_0) \rightarrow H(\mathbb{X}_1) \rightarrow \cdots \rightarrow H(\mathbb{X}_\ell) = H(\mathbb{X}) = H(\mathbb{X}, \mathbb{X}^0) \rightarrow H(\mathbb{X}, \mathbb{X}^1) \rightarrow \cdots \rightarrow H(\mathbb{X}, \mathbb{X}^\ell) = 0. \quad (2.2)$$

Applying the definition of birth and death to the extended sequence, all classes eventually die.

2.4.2 Stability

The information contained within (2.2) can be compactly represented by the *persistence diagrams*, $\text{Dgm}_p(f)$, of f , one for each dimension p . Each diagram is a multiset of points in the plane containing one point (r_i, r_j) for each coset of classes born at \mathbb{X}_i or $(\mathbb{X}, \mathbb{X}^{\ell-i-1})$, and dying entering \mathbb{X}_j or $(\mathbb{X}, \mathbb{X}^{\ell-j-1})$. For reasons of stability, we also add the points on the diagonal to $\text{Dgm}_p(f)$. There are three types of points in the diagram corresponding to the three different locations of its births and deaths. The *ordinary subdiagram*, $\text{Ord}_p(f)$, is the submultiset of points representing the cosets born and dying in the first half of (2.2). The *relative subdiagram*, $\text{Rel}_p(f)$, is the submultiset of points representing cosets born and dying in the second half of (2.2). Finally, the *extended subdiagram*, $\text{Ext}_p(f)$, is the submultiset of points representing cosets born in the first half and dying in the second half of the extended filtration. The points in $\text{Ord}_p(f)$ lie above the diagonal while the points in $\text{Rel}_p(f)$ lie below the diagonal. The points in $\text{Ext}_p(f)$ may lie on either side of the diagonal. See Figure 2.6.

Define the distance between two tame functions $f, g : \mathbb{X} \rightarrow \mathbb{R}$, as

$$\|f - g\|_\infty = \sup_{x \in \mathbb{X}} |f(x) - g(x)|.$$

The L_∞ -distance between two points $a = (a_1, a_2)$ and $b = (b_1, b_2)$ in the plane is $\|a - b\|_\infty = \max\{|b_1 - a_1|, |b_2 - a_2|\}$. Now let $\eta : \text{Dgm}_p(f) \rightarrow \text{Dgm}_p(g)$ be a bijection between the two persistence diagrams for dimension p . We take the supremum L_∞ -distance between the matched points and define the *bottleneck distance* as the infimum over the supremums,

$$W_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) = \inf_{\eta} \sup_a \|a - \eta(a)\|_\infty.$$

The bottleneck distance between the two diagrams is at most the distance between the two functions.

STABILITY OF PERSISTENCE DIAGRAMS [4, 5]. Let $f, g : \mathbb{X} \rightarrow \mathbb{R}$ be tame functions on a triangulable space \mathbb{X} . Then for each dimension p ,

$$W_\infty(\text{Dgm}_p(f), \text{Dgm}_p(g)) \leq \|f - g\|_\infty.$$

2.4.3 Computation

The history of births and deaths is efficiently computable. Assume we are given a filtration

$$\emptyset = K_0 \subset K_1 \subset \cdots \subset K_\ell = K$$

of a simplicial complex K such that the difference between two consecutive complexes is a single simplex. Applying the homology functor, we get a sequence of homology groups connected by homomorphisms induced by inclusion. The filtration is neatly encoded in a single boundary matrix. Here the i -th row and column represents a single simplex, namely $K_i - K_{i-1}$. We place a 1 in entry (i, j) of the boundary matrix if the simplex i belongs to the boundary of simplex j otherwise we place a 0. A reduction of this matrix similar to the Smith normal form gives all the birth and death information [9, 6, 13]. The running time of the reduction algorithm is on the order of the number of simplices in the complex K cubed. There are examples requiring the worst case cubic time [27].

3

Reeb Spaces

Given a function $f : \mathbb{X} \rightarrow \mathbb{R}$, its Reeb graph represents the components of the level sets as points and expresses their relationship by forming a 1-dimensional space. The situation for mappings to \mathbb{R}^n , for $n > 1$, is significantly more complicated and the topic of this chapter.

There is some prior work on the extension of Reeb graphs to Reeb spaces. The existing work is limited to bivariate, generic, smooth mappings.

- Burlet and de Rham study smooth, bivariate mappings on orientable 3-manifolds [2]. Under the assumption that every point of the singular set is definite (appears as a minimum for some linear combination of the two components), they establish relationships between the topology of the 3-manifold and that of the Reeb space. Porto and Furuya extend this work to orientable d -manifolds for $d \geq 3$ [29].
- Motivated by the study of immersions of 3-manifolds in \mathbb{R}^4 , Levine and coauthors give a complete local classification of points in the Reeb space of bivariate, generic, smooth mappings on orientable as well as non-orientable 3-manifolds [24, 25]. Furuya extends this work to orientable 4-manifolds [17] and Kobayashi and Saeki extend it further to d -manifolds for $d \geq 3$ [22].

In the piecewise linear literature, we find only one paper that goes beyond Reeb graphs [10]. It gives a dynamic algorithm for maintaining the Reeb graph in time for piecewise linear functions.

In this chapter, we consider generic, piecewise linear mappings from a combinatorial m -manifold to \mathbb{R}^n . Following the work on generic, smooth mappings, we characterize points of the Reeb space, proving that their neighborhoods are cones over Reeb spaces of one lower dimension. Complementing the local analysis, we show that Reeb spaces have triangulations and coarsest stratifications. Their existence is established constructively. In the case of the triangulation, this leads to a

polynomial-time algorithm while the construction of the coarsest stratification contains an undecidable subproblem and leads to algorithms only for $n \leq 4$.

3.1 Definition

Intuitively, the Reeb space of f parameterizes the set of components of preimages of points $c \in \mathbb{R}^n$. By the PL Preimage Lemma, all but a measure zero subset of these components are manifolds of dimension $m - n$. As we vary c without crossing the image of the $(n - 1)$ -skeleton, these manifolds vary without a change to their homeomorphism type. Since c has n degrees of freedom, this variation has locally the structure of an n -manifold. Only when c belongs to the image of the $(n - 1)$ -skeleton can we have violations of the manifold property and get shapes that appear as transitions between manifolds of possibly different global connectivity. In summary, we imagine the Reeb space as a collection of n -manifolds glued to each other in possibly complicated ways. The remainder of this chapter shows that this is indeed the right intuition. We do this by first formally introducing the Reeb space and then studying its local and global topological properties.

Call two points x and y in $|K|$ *equivalent*, denoted by $x \sim y$, if $f(x) = f(y)$ and x and y belong to the same component of the preimage $f^{-1}(f(x)) = f^{-1}(f(y))$. The *Reeb space* is the quotient space obtained by identifying equivalent points, $\mathbb{W}_f = |K| / \sim$, together with the quotient topology inherited from $|K|$. We already have a map from $|K|$ to \mathbb{R}^n , namely f , and one from $|K|$ to \mathbb{W}_f called the *quotient map*, ρ_f . The *Stein factorization* adds another map π_f from \mathbb{W}_f to \mathbb{R}^n such that the triangle commutes:

$$\begin{array}{ccc} |K| & \xrightarrow{f} & \mathbb{R}^n \\ \rho_f \searrow & & \nearrow \pi_f \\ & \mathbb{W}_f & \end{array}$$

It is not difficult to prove that the Reeb space is Hausdorff, that is, any two different points in \mathbb{W}_f have disjoint neighborhoods.

An example. We illustrate the definition with a mapping from a 3-manifold to the plane. It is convenient to describe a smooth mapping. The extra details that appear in the case of a PL approximation are not difficult. Consider $f : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ defined by its two component functions $f_1(x_1, x_2, x_3) = x_2^3 - x_1x_2 + x_3^2$ and $f_2(x_1, x_2, x_3) = x_1$. The preimage of a point $c = (s, t)$ is the intersection of two level surfaces, $f_1^{-1}(s) \cap f_2^{-1}(t)$. Setting the two components to s and t we get $x_1 = t$ and $x_3^2 = \gamma(x_2)$ where $\gamma(x_2) = s - x_2^3 + tx_2$. For $t > 0$, γ has a minimum and a maximum, for $t = 0$ it has a single degenerate critical point, and for $t < 0$ it has no critical points. We are interested in the number of roots and in particular, the values of x_2 where $\gamma(x_2) \geq 0$ because only for these values do we get a solution to $x_3^2 = \gamma(x_2)$. The odd degree




FIGURE 3.1: Above the fold, the function γ has three roots and thus two intervals where γ is positive. Below the fold, γ has one root and only one interval where γ is positive. The Reeb space of this function is a fin glued to a disk along one of the fold curves.

polynomial γ has either 1 or 3 roots except where γ and its derivative have a common zero, which occurs along the fold $27s^2 = 4t^3$. As illustrated in Figure 3.1, this curve decomposes the (s, t) plane into two regions, γ has three roots above, two roots on, and one root below the fold. Accordingly, $f^{-1}(c)$ has two components above and one component below the fold. It consists of a curve and an isolated point for c on the left branch and of two touching curves for c on the right branch of the fold. It should be clear how these cases transition between each other as we vary c in the plane. The fold is the image of the singularity set under the mapping f .

We can reinterpret Figure 3.1 as a picture of the Reeb space of f . Indeed, it is the image of \mathbb{W}_f under the map π_f in the Stein factorization. The region above the fold is covered twice and the region below is covered once. Correspondingly, the Reeb space consists of two sheets, one covering the entire plane and the other covering the region above the fold. The latter connects to the former along the right branch of the fold where the two components of $f^{-1}(c)$ come together to merge into a single component. The left branch of the fold is the image of a boundary piece of the second sheet and has nothing to do with the first sheet. In summary, the Reeb space consist of the plane with another two-dimensional sheet attached to it, like a fin sticking out of a fish as in Figure 3.1, right.

3.2 Local Structure

In this section, we prove that every point of the Reeb space has a neighborhood homeomorphic to a cone over a Reeb space of dimension one lower. In stating and proving this result, we follow the work on generic, smooth mappings in [22].

Tubes, cores, and cones. As usual, we let K be a combinatorial m -manifold and $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping. Let p be a point of \mathbb{W}_f and $c = \pi_f(x)$ its image in \mathbb{R}^n . Let B be a closed ball centered at c that is sufficiently small so it intersects the image of a simplex iff this image includes c . Considering the preimages of c and of B , we are interested in the component C of $f^{-1}(c)$ whose image under the quotient map is p and in the component T of $f^{-1}(B)$ that contains C . We call T a *tube* and C its *core*. Define the boundary $\text{Bd } T$ of the tube T to be the intersection of T with $f^{-1}(\text{Bd } B)$, and let $r : \text{Bd } T \rightarrow \mathbb{R}^k$ be the restriction of f . The corresponding restriction, ρ_r , of ρ_f maps $\text{Bd } T$ to $\mathbb{W}_r \subset \mathbb{W}_f$. Since K is finite, $|K|$ is compact. This implies that $\text{Bd } T$ is compact and so is \mathbb{W}_r . The (*closed*) *cone* over \mathbb{W}_r is the space

$$\text{cone}(\mathbb{W}_r) = (\mathbb{W}_r \times [0, 1]) / (\mathbb{W}_r \times 1)$$

and its *cone point* is $(\mathbb{W}_r \times 1) / (\mathbb{W}_r \times 1)$. We are now ready to state the first structural result of this paper.

CONE NEIGHBORHOOD THEOREM. Let K be a combinatorial m -manifold, $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping, and \mathbb{W}_f the Reeb space of f . Then each point $x \in \mathbb{W}_f$ has a homeomorphism from $\text{cone}(\mathbb{W}_r)$ to a closed neighborhood that maps the cone point to x .

To prove this theorem, we use that the cone over \mathbb{W}_r is compact and that \mathbb{W}_f is Hausdorff. Every continuous injection from a compact to a Hausdorff space is an embedding [28, page 167]. This means that the compact space and its image are homeomorphic. It thus suffices to construct a continuous injection $\eta : \text{cone}(\mathbb{W}_r) \rightarrow \mathbb{W}_f$ that maps the cone point to x . The next paragraph does exactly that.

Constructing an embedding. We begin by constructing the barycentric subdivision of K , slightly modified by placing the new vertices not always at the barycenters of the simplices. The connecting simplices are the same as in the standard definition. Specifically, if C intersects the interior of a simplex σ in K then we choose a point $u_\sigma \in C \cap \text{int } \sigma$ as the vertex in $\text{Sd } K$ that represents σ . If C does not intersect the interior of σ , then we choose a point $u_\sigma = \text{int } \sigma - T$. This point exists because B is sufficiently small. By construction, there is a subcomplex L of $\text{Sd } K$ whose underlying space is the core, $|L| = C$; as illustrated in Figure 3.2. It is not difficult to prove that L is a full subcomplex. Extending the concept of a star, we write $\text{St } L$ for the set of simplices in $\text{Sd } K$ that have a face in L . The tube is covered in its entirety by the interiors of the simplices in $\text{St } L$.

We first use the barycentric subdivision to establish a continuous map from $\text{Bd } T \times [0, 1]$ to T whose restriction to $\text{Bd } T \times [0, 1)$ is a homeomorphism onto $T - C$. Let τ be a simplex in $\text{St } L$ but not in L and let σ be the maximal face of τ that belongs to L . We observe that σ is unique because L is full. Let v be the maximal face of τ that is disjoint from σ and note that $\tau = \sigma * v$, the join of its two faces. Writing all simplices τ as joins, we get a decomposition of the closed star of L into line segments.




FIGURE 3.2: A piece of the barycentric subdivision of K . The white dots mark the vertices of K and the black and shaded dots mark the new vertices of $\text{Sd } K$. The shaded path is a piece of the core which is subdivided by the subcomplex L of $\text{Sd } K$. The corridor along the path is a piece of the tube which is contained in $\text{St } L$.

As usual, any two line segments in this decomposition are either disjoint or meet at a common endpoint. Each point of the core is an endpoint of a collection of line segments. In contrast, a point x of the boundary of the tube belongs to exactly one line segment. Letting y be the endpoint of this line segment in the core we let $\lambda_x : [0, 1] \rightarrow T$ be the straight line mapping $\lambda_x(t) = (1 - t)x + ty$. Combining the maps λ_x over all $x \in \text{Bd } T$ gives the map $\lambda : \text{Bd } T \times [0, 1] \rightarrow T$. As anticipated, the restriction of λ to $\text{Bd } T \times [0, 1) \rightarrow T - C$ is a homeomorphism and λ itself is continuous. Finally, define $\gamma : \text{Bd } T \times [0, 1] \rightarrow \mathbb{W}_f$, by setting $\gamma = \rho_f \circ \lambda$. The new map γ takes $\text{Bd } T \times 1$ to the point p . The preimage of every other point q in the image of γ is of the form $U \times t$, where U is the preimage of a point in \mathbb{W}_r and t is in $[0, 1)$.

Next we map $\text{Bd } T \times [0, 1]$ to the cone over \mathbb{W}_r . Recall that $r : \text{Bd } T \rightarrow \mathbb{R}^n$ is the restriction of f to the boundary of the tube and $\rho_r : \text{Bd } T \rightarrow \mathbb{W}_r$ is the corresponding restriction of ρ_f . We extend ρ_r to a map from $\text{Bd } T \times [0, 1]$ to $\mathbb{W}_r \times [0, 1]$ by taking the product with the identity on the unit interval. Composing this product map with the quotient map $\mathbb{W}_r \times [0, 1] \rightarrow \text{cone}(\mathbb{W}_r)$, we get $\rho : \text{Bd } T \times [0, 1] \rightarrow \text{cone}(\mathbb{W}_r)$ mapping $\text{Bd } T \times 1$ to the cone point. The preimage of every other point q in $\text{cone}(\mathbb{W}_r)$ is of the form $U \times t$, where U is the preimage of a point in \mathbb{W}_r and t is in $[0, 1)$, as before. This finally induces a unique map, η , from the cone to the Reeb space that makes

the triangle commute:

$$\begin{array}{ccc}
 & \text{Bd } T \times [0, 1] & \\
 \swarrow \rho & & \searrow \gamma \\
 \text{cone}(\mathbb{W}_r) & \xrightarrow{\eta} & \mathbb{W}_f.
 \end{array}$$

To finish the proof of the Cone Neighborhood Theorem we just need to realize that η has the two required properties. It is continuous because both γ and ρ are continuous and it is injective because the preimages of points in the images of ρ and of γ are the same sets $U \times t$.

3.3 Global Structure

The Reeb space has a canonical stratification into manifolds. We give a construction in two steps, first triangulating the Reeb space and second by grouping simplices to form the strata. The triangulation of the Reeb space implies the Cone Neighborhood Theorem. However, the construction of the triangulation is very different in flavor from the method used in the previous section to prove the Cone Neighborhood Theorem.

Refining arrangement. As before, let K be a combinatorial m -manifold and $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping. We also assume $n < m$, otherwise K is itself a triangulation of the Reeb space. To prepare the construction of a triangulation, we refine K by decomposing its simplices into prisms aligned with the preimages of f . Specifically, we take the images of the $(n - 1)$ -simplices of K in \mathbb{R}^n , dissect the space with their affine hulls, and decompose the simplices using the preimage of the dissection. By assumption of genericity, the image of every $(n - 1)$ -simplex $\sigma \in K$ is a $(n - 1)$ -simplex and its affine hull is a $(n - 1)$ -dimensional plane in \mathbb{R}^n . The collection of such planes dissects \mathbb{R}^n into closed chambers, each a convex polyhedron of dimension n . We call this the *arrangement* defined by the planes [7]. To refine K , we take each simplex τ and decompose it into sets of points that map into a common chamber or a common intersection of chambers. For an n -simplex τ , these sets are n -dimensional convex polytopes, the same as the chambers. For an $(n + 1)$ -simplex τ , these sets are $(n + 1)$ -dimensional prisms each uniquely determined by its top and bottom faces of dimension n . It is allowed that the top and bottom faces touch each other along a common face, generating a partially degenerate prism in between. We show that it is not necessary to study decompositions of simplices of dimension beyond $n + 1$.

SKELETON LEMMA. The Reeb space of $f : |K| \rightarrow \mathbb{R}^n$ is homeomorphic to the Reeb space of the restriction of f to the $(n + 1)$ -skeleton of K .

PROOF. Let $e : |K^{(n+1)}| \rightarrow \mathbb{R}^n$ be the restriction of f to the $(n+1)$ -skeleton and recall that points x and y are equivalent if they map to the same image, $e(x) = e(y) = c \in \mathbb{R}^n$, and belong to the same component of the preimage, $e^{-1}(c)$. By assumption of genericity, this preimage is a complex whose maximal elements are edges, each a line intersecting a $(n+1)$ -simplex. In contrast, $f^{-1}(c)$ is a complex whose maximal elements are $(m-n)$ -dimensional convex polytopes, each the intersection of a $(m-n)$ -dimensional plane with a m -simplex. Since $e^{-1}(c)$ is the 1-skeleton of $f^{-1}(c)$, there is a bijection between the components of the two preimages. Hence there is a bijection between \mathbb{W}_e and \mathbb{W}_f . Finally, we observe that the quotient topologies are equivalent implying that the bijection is a homeomorphism between the two Reeb spaces. \square

Triangulation. We use the decompositions of the skeleta of K to construct a triangulation of the Reeb space. Let Q be the collection of preimages of chambers decomposing the n -skeleton of K and call two of these polytopes *incident* if they share a common $(n-1)$ -dimensional face. Let P be the collection of prisms decomposing the $(n+1)$ -skeleton and recall that each $\varphi \in P$ has two n -dimensional faces in Q , its *top face* φ_t and its *bottom face* φ_b . The algorithm partitions Q into blocks, starting with the partition into singletons, $\mathcal{Q} = \{\{\psi\} \mid \psi \in Q\}$. We write Q_ψ for the block that contains ψ .

```

for each prism  $\varphi \in P$  do
  if  $Q_{\varphi_t} \neq Q_{\varphi_b}$  then
    merge the two blocks into one
  endif
endfor.

```

When we merge two blocks, we remove both from \mathcal{Q} and add their union as a new block to \mathcal{Q} . By construction, all polytopes in a block are preimages of the same chamber in the arrangement. We say two blocks Q_ψ and $Q_{\psi'}$ are *incident* if ψ and ψ' are preimages of different but incident chambers in the arrangement and there are at least two preimages, one in each block, that are incident.

A complex representing the Reeb space of f is readily obtained from the partition into blocks. Specifically, for each block Q_ψ in \mathcal{Q} , we take a copy of the chamber $f(\psi)$ and we glue these copies along shared $(n-1)$ -faces to reflect the incidence relation among the blocks. We further decompose each polytope into simplices and thus finally get a simplicial complex we denote as W_f . In summary, we have an algorithm that triangulates the Reeb space of a generic PL mapping f from a combinatorial m -manifold to \mathbb{R}^n . Assuming m is a constant, the size of the triangulation and the running time of the algorithm are both polynomial in the size of the combinatorial manifold.

Stratification. In general, the complex W_f will be significantly finer than necessary to represent the Reeb space. In a first step towards coarsening the representation,

we group simplices to form manifolds. The result will be a *stratification* of \mathbb{W}_f , that is, a filtration

$$\emptyset = W^{-1} \subseteq W^0 \subseteq \dots \subseteq W^n = W_f$$

such that each W^j is a subcomplex of W_f and $S^j = W^j - W^{j-1}$ is either empty or a j -manifold. We call S^j the *j -stratum* of the stratification and each of its components a *j -dimensional piece*. In addition to being a j -manifold, we require that all points of a piece are topologically equivalent. By this we mean that any two points x and y of a piece have closed neighborhoods $N(x)$ and $N(y)$ in $|W_f|$ and a homeomorphism from one to the other that maps x to y and whose restriction to the piece is again a homeomorphism. By the Cone Neighborhood Theorem the closed neighborhoods are cones over $(n - 1)$ -dimensional Reeb spaces. The requirement of topological equivalence can therefore be reformulated in terms of these spaces. Consider the 2-dimensional Reeb space described in the example of §3.1. Its 2-stratum consists of two sheets (mapping to the plane and to the fin), its 1-stratum consists of two curves (mapping to the two branches of the fold), and its 0-stratum is one point (mapping to the origin).

We construct the stratification in the order of decreasing dimension. At the top dimension, we initialize S^n to the set of n -simplices, each a piece by itself. Then we add simplices of lower dimension effectively merging and enlarging the pieces. For this, we use a boolean subroutine DOESBLEND that decides whether or not a simplex fits into a piece or between pieces of the current stratum. We will prove shortly that each iteration starts with a complex W^j of dimension at most j . Following the same pattern as before, we can therefore construct the j -stratum of W_f as the top dimension stratum of W^j .

```

Set  $W^n = W_f$ ;
for  $j = n$  downto 0 do
  initialize  $W^{j-1}$  to the  $(j - 1)$ -skeleton of  $W^j$ 
  and  $S^j$  to  $W^j - W^{j-1}$ ;
  for  $i = j - 1$  downto 0 do
    for each  $i$ -simplex  $\zeta \in W^{j-1}$  do
      if DOESBLEND( $\zeta, S^j$ ) then
        add  $\zeta$  to  $S^j$  and remove it from  $W^{j-1}$ 
      endif
    endfor
  endfor
endfor.

```

Note that $S^j = W^j - W^{j-1}$ is maintained throughout the algorithm. We still need to establish that the algorithm constructs what we promise but this depends primarily on the boolean subroutine that decides upon which simplices to add to a stratum.

Recognition. According to the definition of a stratification, we need to satisfy two conditions when we add an i -simplex ζ to the current set S^j , the first guaranteeing

that we have a j -manifold and the second that points in the same piece have homeomorphic neighborhoods. We formalize both conditions by considering the second barycentric subdivision and comparing links of vertices in this subdivision. Recall that the first barycentric subdivision contains a vertex $\hat{\xi}$ for each simplex $\xi \in W_f$. We refer to it as a *first generation* vertex of $\text{Sd}^2 W_f = \text{Sd Sd } W_f$, noting that all its neighbors are second generation vertices in $\text{Sd}^2 W_f$. The link of $\hat{\xi}$ is a model for the boundary of the closed neighborhood of any point in the interior of ξ . Let $\text{Sd}^2 S^j$ be the subset of simplices in $\text{Sd}^2 W_f$ whose interiors are contained in $|S^j|$. We accept ζ as a new simplex in the j -stratum if the following two conditions are satisfied:

1. The link of $\hat{\zeta}$ in $\text{Sd}^2 S^j$ is a $(j - 1)$ -sphere.
2. There is a homeomorphism that maps the link of $\hat{\zeta}$ to the link of $\hat{\xi}$ in $\text{Sd } W_f$, where ξ is already in S^j and belongs to the star of ζ . We also require that the restriction of this homeomorphism to $\text{Sd}^2 S^j$ is a homeomorphism between the two links. By Condition 1, both links are $(j - 1)$ -spheres.

It is clear that this implementation of the boolean subroutine `DOESBLEND` maintains S^j as a j -manifold. For the top dimension, $j = n$, this implies that whenever S^n contains a simplex then it also contains the simplices in its star. Symmetrically, whenever $W^{n-1} = W^n - S^n$ contains a simplex it also contains its faces. In other words, W^{n-1} is a complex. We can now use induction over the dimension and prove that W^j is a complex for all j . Similarly, whenever S^j contains a simplex then it also contains its star within W^j . Hence, if S^j is non-empty, then it is a j -manifold, and this is true for every j . Finally, we notice that the result of the algorithm does not depend on the order in which the simplices are processed. Indeed, the test of the i -simplex ζ does not depend on whether or not any other i -simplices belong to the j -stratum. We thus have a constructive proof of a global property of the Reeb space.

STRATIFICATION THEOREM. Let K be a combinatorial m -manifold, $n \geq 1$, and $f : |K| \rightarrow \mathbb{R}^n$ a generic PL mapping. Then the Reeb space \mathbb{W}_f of f is a stratified space and the W^j as constructed by the algorithm form its coarsest stratification.

We note that the constructive proof is really an algorithm only for $n < 5$. Otherwise, the boolean subroutine attempts to recognize when two triangulated spaces of dimension $n - 1 \geq 4$ are homeomorphic. This problem is undecidable as proved by Markov [26].

3.4 The Orientable 3-Manifold Case

We take a closer look at the Reeb spaces for PL mappings from an orientable 3-manifold to the plane. In particular, we want an analysis of local cones. To keep the number of local cones small, we assume simple generic mappings. Here we call




FIGURE 3.3: From left to right: a regular edge and three simple critical edges. Each edge is shown with a cross-section of its star and its lower link in bold.

a generic PL mapping $f : |K| \rightarrow \mathbb{R}^n$ is a *simple* if every $(n - 1)$ -simplex in J_f is a simple singular simplex.

Genericity and simplicity. Let K be a compact combinatorial orientable 3-manifold without boundary and $f : |K| \rightarrow \mathbb{R}^2$ a PL mapping. We assume that f is generic and simple. Specifically, we require that

- I'. the intersection of a level set of f with $|K^{(1)}|$ is empty, one point, or two points each in the interior of an edge;
- II'. the Jacobi set of f is a 1-manifold, that is, each edge of J_f is a simple critical edge and each vertex of J_f is the endpoint of exactly two edges in J_f .

Recall that to define the lower link of an edge, we use the function $h_{\vec{u}} : |K| \rightarrow \mathbb{R}$ mapping a point x to the height of $f(x)$ in the direction $\vec{u} \in \mathbb{S}^1$ normal to the edge. A critical edge is simple iff all reduced Betti numbers of this lower link vanish, except for one, which is equal to 1. There are three possibilities: $\tilde{\beta}_{-1} = 1$ (the cross-section of the edge is a minimum), $\tilde{\beta}_0 = 1$ (a saddle), and $\tilde{\beta}_1 = 1$ (a maximum); see Figure 3.3. Condition II' implies that J_f contains no duplicate edges and no duplicate vertices, where by the latter we mean that each endpoint of an edge in J_f belongs to exactly one other edge in J_f ; see also [8]. It is important to note that Condition II' is not generic. In other words, there may not be a PL mapping arbitrarily close to the given PL mapping satisfying Condition II'. Condition I' is generic.

Walks and sheets. To enumerate the topologically different types of cones that may arise, we let $p \in \mathbb{W}_f$ be a point of the Reeb space and B a small closed disk with center $c = \pi_f(p)$ in the plane, as in §3.2. Furthermore, the core, C , is the component of $f^{-1}(c)$ whose image under ρ_f is p , and the tube, T , is the component of $f^{-1}(B)$ that contains C . Recall that $r : \text{Bd } T \rightarrow \mathbb{R}^2$ is the restriction of f that maps the boundary of the tube to the boundary of the disk and that p has a closed neighborhood in \mathbb{W}_f that is homeomorphic to the cone over $\mathbb{W}_r = \rho_f(\text{Bd } T)$. It thus suffices to understand the structure of \mathbb{W}_r . We get insight into this structure by walking around the circle




FIGURE 3.4: Left: the tube, its core, and the annulus that divides the tube into two. Right: the closed neighborhood of p in the Reeb space and its image in the plane.

bounding B in a counter-clockwise order using $\alpha : [0, 1] \rightarrow \mathbb{R}^2$ with $\text{im } \alpha = \text{Bd } B$. The walk begins and ends at the point $a = \alpha(0) = \alpha(1)$. Letting b be the antipodal point, we also walk along the straight diameter using $\beta : [0, 1] \rightarrow \mathbb{R}^2$ with $a = \beta(0)$, $c = \beta(\frac{1}{2})$, and $b = \beta(1)$. Note that each point of the two walks is the image of a curve in the tube. The two one-parameter families sweep out the boundary of the tube and another surface we refer to as the *divider*, $D = T \cap f^{-1} \circ \beta[0, 1]$. To describe the two sweeps, we define

$$\begin{aligned} \text{Bd } T_s &= T \cap f^{-1} \circ \alpha[0, s]; \\ D_s &= T \cap f^{-1} \circ \beta[0, s], \end{aligned}$$

for each $0 \leq s \leq 1$. The simplest of all possible cases is illustrated in Figure 3.4. There, none of the points in the disk B is critical. Hence, the preimage of every point of $\alpha[0, 1]$ is a closed curve, and the same is true for the preimage of every point of $\beta[0, 1]$. It follows that $\text{Bd } T_s$ is an annulus, for every $0 < s < 1$, that closes up to form a torus when s reaches 1. Similarly, D_s is an annulus, for every $0 < s < 1$, and it remains one until the end. The divider, D , is therefore an annulus bounded by the two closed curves shared with $\text{Bd } T$. Since D also contains the core, C , the only possible configuration is the one depicted in Figure 3.4. In this particular case, the point p belongs to a *sheet*, that is, a piece of the 2-stratum of the Reeb space. The points in the neighborhood correspond to closed curves forming a fibration of the tube.

Choosing the walks. For the more complicated cases, it will be convenient to choose the two walks such that the point a on the circle has a connected preimage and the preimages of a and of b avoid the 1-skeleton of the barycentric subdivision of K . We use the curves sweeping out the divider to prove that such points a and b exist.

ENDPOINT LEMMA. There exist antipodal points a and b of $\text{Bd } B$ such that $r^{-1}(a)$ is connected and $r^{-1}(a)$ and $r^{-1}(b)$ both have empty intersection with $(\text{Sd } K)^{(1)}$.

PROOF. Each pair of antipodal points corresponds to a direction $\vec{u} \in \mathbb{S}^1$ such that \vec{u} is a positive multiple of $b - a$. Let $\beta_{\vec{u}} : [0, 1] \rightarrow \mathbb{R}^2$ be the corresponding diagonal walk and $D(\vec{u})$ the corresponding divider. Note that the dividers all share the core but are otherwise disjoint.

Fixing a direction \vec{u} and a point x in the core, we consider how the curve $T \cap f^{-1} \circ \beta_{\vec{u}}(s)$, which sweeps out $D(\vec{u})$ as s goes from 0 to 1, intersects a sufficiently small neighborhood $N(x)$ of x in $|K|$. If $x \notin |J_f|$ then the curve looks locally like a line that sweeps over x , passing it at $s = \frac{1}{2}$. Hence $N(x)$ intersects the curve in a connected piece, if at all. If $x \in |J_f|$ then a curve approaches x , pinching off to a single point or recombines leaving x in two different directions. In the former case, we see a closed curve shrinking to a point or the other way round. In the later case, locally we see the usual saddle picture of two pieces that look like the two branches of a hyperbola passing through its pair of asymptotic lines. The two pieces are globally connected along a component of the curve before meeting at x but are not connected after meeting at x , or the other way round. There is an open semi-circle of directions \vec{u} such that $N(x)$ intersects a single component of the curve. This semi-circle is determined by the image of the edge or edges in J_f that contains the point x . By Condition I', there are at most two points in the core that belong to $|J_f|$ and by Condition II' at most two edges in J_f are adjacent to a vertex in the Jacobi set. The corresponding two semi-circles are defined by the images of two different edges in J_f . It follows that the two line segments intersect at c and the corresponding semi-circles intersect in an arc of non-zero length. Picking a on this arc implies that r^{-1} is connected. To satisfy the second requirement of avoiding the 1-skeleton of $\text{Sd } K$, we just need to choose a outside a measure zero subset of the arc. \square

Arcs. Beyond sheets, the next more complicated case is when the boundary of the tube meets the Jacobi set in two points, x' and x'' , in the interior of a single edge or in the interior of two different edges. The core intersects $|J_f|$ in a single point, x . The points x' and x'' belong to different edges when x is a vertex of J_f and in this case, we assume both edges are definite or both indefinite. The case in which the two edges have different types will be discussed later. The assumption allows for two cases and in both the point p belongs to an *arc*, that is, a piece of the 1-stratum. The point x is the sole interior critical point of $\beta^{-1} \circ f : D \rightarrow [0, 1]$ and the points x' and x'' the sole interior critical points of $\alpha^{-1} \circ f : \text{Bd } T \rightarrow [0, 1]$.

CASE A.1. The edge of J_f that contains x' and x'' is definite. The tube is a ball obtained by thickening the point x . The divider, D , depends on the choice of the diameter since the preimage of a may be empty or one curve. Walking along the circle, we start with a single curve that shrinks to the point x' and then reappears from the point x'' and returns to its original position. The curve then sweeps out a sphere from one pole to the other. The Reeb space is locally a half-plane, like at a point on the left branch of the fold in Figure 3.5.




FIGURE 3.5: The local cones for cases A.1 and A.2. The cone point belongs to an arc of the Reeb space.

CASE A.2. The edge of J_f that contains x' and x'' is indefinite. The tube is a solid double torus obtained by thickening the figure-8 curve that crosses itself at x . The divider, D , depends on the choice of the diameter since the preimage of b may consist of one or two curves. Walking around the circle, we start with a single curve that splits into two at x' that later merges at x'' to form again a single curve. The Reeb space is locally a book with three pages, like at a point on the right branch of the fold in Figure 3.5.

Nodes. Next, we consider the case when the core meets J_f at a vertex, x , one of the incident edges of J_f is definite and the other incident edge is indefinite. By Condition II', the boundary of the tube meets the Jacobi set in two points x' and x'' . The point x is the sole interior critical point of $\beta^{-1} \circ f : D \rightarrow [0, 1]$ and the points x' and x'' are the sole interior critical points of $\alpha^{-1} \circ f : \text{Bd } T \rightarrow [0, 1]$.

CASE N.1. Assuming x' belongs to the indefinite and x'' to the definite edge, the walk around the circle starts with a single curve that splits into two at x' of which one shrinks to a point at x'' . The divider, D , depends on the choice of the diameter. Specifically, the preimage of the endpoint b may consist of one or of two curves. In the former case, we have one curve that persists along the entire diameter. In the latter case, we start with one curve and get another expanding around x . In either case, the core is a single curve and the Reeb space is locally a disk with a fin sticking out, like at the origin in Figure 3.1; see also Figure 3.6, left. The point p is a *node* of the Reeb space, that is, a piece of the 0-stratum.

In the most complicated case, the core meets the Jacobi set in two points, x and y . Each of the two points lies in the interior of an indefinite edge, else the core would be disconnected. The boundary of the tube meets the Jacobi set in four points, x' , x'' , y' , and y'' . Here, x and y are the sole interior critical points of $\beta^{-1} \circ f : D \rightarrow [0, 1]$ and x' , x'' , y' , y'' are the sole interior critical points of $\alpha^{-1} \circ f : \text{Bd } T \rightarrow [0, 1]$. There are two cases and in both the point p is a node.

CASE N.2. Walking along the diameter, we start with a single curve that gets pinched at x and at y with the net effect that it remains a single curve. Know-

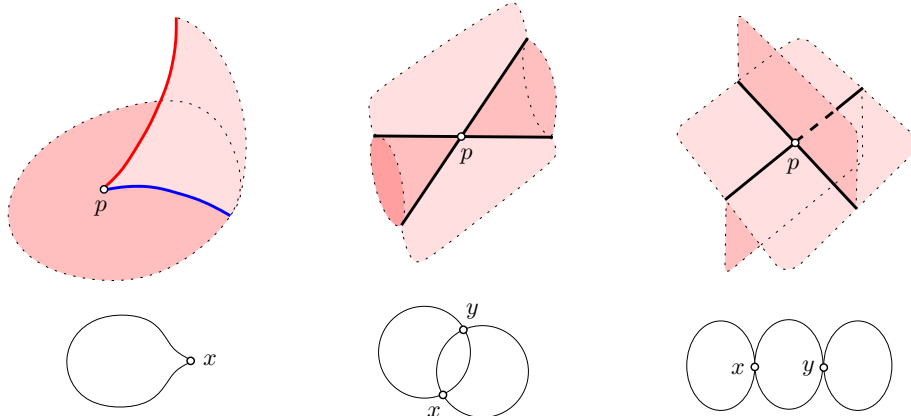


FIGURE 3.6: From left to right: the local cone and the core for the Cases N.1, N.2, N.3 where the point p is a node of the Reeb space.

ing that the preimage of b is a single curve determines the boundary of the tube. Using α to walk the circle, we start with a single curve that splits into two curves at x' and then merges into a single curve at y' . As we continue, the single curve splits into two curves at x'' and once again it merges into a single curve at y'' . The core consists of two circles that meet at two points, x and y , and the tube is a solid triple torus; see Figure 3.6, middle.

CASE N.3. Walking along the diameter, we start with a single curve that splits into three curves at x and y . Knowing that the preimage of b consists of three curves again determines the boundary of the tube. As we walk along the circle, the single curve splits into two at x' , one of the two splits into two at y' , giving a total of three curves. As we continue, two of the three curves merge at x'' and the remaining two merge at y'' . The core is a double figure-8 and the tube a solid triple torus; see Figure 3.6, right.

3.5 Remarks

We described an algorithm to construct the Reeb space of a piecewise linear mapping. Let $f : |K| \rightarrow \mathbb{R}^n$ be a PL mapping. The algorithm involves extending the image of each $(n - 1)$ -simplex to an $(n - 1)$ -dimensional affine subspace. We then use the arrangement of these affine subspaces to divide each $(n + 1)$ -simplex into chambers. The complexity of this arrangement is on the order of the number of $(n - 1)$ -simplices to the power n . Is there a better way to compute the Reeb space? Recently, Harvey et. al. invented a fast randomized algorithm to compute the Reeb graph of a PL function [20]. This algorithm has a flavor similar to that of our algorithm described in §3.3. Is it possible to extend the randomized algorithm to higher dimensional Reeb spaces? Will a fast algorithm for Reeb spaces encourage their use in data analysis?

The last section, §3.4, gives a local classification of the Reeb space for PL map-

pings taking an orientable 3-manifold to the plane. However, these mapping must satisfy a non-generic condition, namely Condition II'. What if Condition II' is not satisfied? That is, what if the lower link of each edge in the Jacobi set is not simple. For example, Figure 3.5 shows the two possible neighborhoods of a point in a 1-stratum assuming Condition II'. The 1-stratum has one sheet attached to it or three. Without Condition II', it is possible to construct a PL mapping with an arbitrary number sheets attached to the 1-stratum. It would be nice to have a local classification of the Reeb space for generic PL mappings.

The definition of the Reeb space works for continuous mappings between arbitrary manifolds and even arbitrary topological spaces. Is it worth studying the Reeb space of mappings where the range space is something more complicated than Euclidean space?

Robustness of Preimages

The Reeb space captures 0-dimensional connectivity of preimages ignoring higher-dimensional connectivity such as tunnels or voids. Homology makes the idea of higher-dimensional connectivity precise. In this chapter, we look not only at the higher dimensional connectivity of the preimage but ask about its stability. We develop a method to measure the stability or robustness of each homology class in the preimage. Roughly speaking, the robustness of a homology class is the amount of perturbation necessary to remove it from the preimage.

4.1 Well Groups

Let \mathbb{X} be an m -manifold, \mathbb{Y} a Riemannian n -manifold, and \mathbb{A} a k -dimensional submanifold of \mathbb{Y} . A continuous mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$ is *admissible* if $f^{-1}(\mathbb{A})$ has a finite rank homology groups. Writing $\|a - b\|_{\mathbb{Y}}$ for the distance between the points $a, b \in \mathbb{Y}$ define $f_{\mathbb{A}} : \mathbb{X} \rightarrow \mathbb{R}$ as $f_{\mathbb{A}}(x) = \inf_{a \in \mathbb{A}} \|f(x) - a\|_{\mathbb{Y}}$. The *distance function* $f_{\mathbb{A}}$ maps each point in \mathbb{X} to the distance between \mathbb{A} and its image under f . The *level set* of $f_{\mathbb{A}}$ at a value r is the preimage of that value, $f_{\mathbb{A}}^{-1}(r)$. The *sublevel set* for the same value, r , is the preimage of $[0, r]$. Writing \mathbb{A}^r for the set of points in \mathbb{Y} at distance r or less from \mathbb{A} , we have $f_{\mathbb{A}}^{-1}[0, r] = f^{-1}(\mathbb{A}^r)$.

A mapping $h : \mathbb{X} \rightarrow \mathbb{Y}$ homotopic to f is a ρ -*perturbation* of f if $\|h - f\|_{\infty} \leq \rho$, where the norm of the difference is the supremum over all $x \in \mathbb{X}$ of the distance between $h(x)$ and $f(x)$ in \mathbb{Y} . The preimage of \mathbb{A} under a ρ -perturbation is contained in the preimage of \mathbb{A}^{ρ} under f . Writing this in terms of distance functions, we have $h_{\mathbb{A}}^{-1}(0) \subseteq f_{\mathbb{A}}^{-1}[0, \rho]$. This inclusion induces a homomorphism between the corresponding homology groups,

$$j_h : H(h_{\mathbb{A}}^{-1}(0)) \rightarrow F(\rho),$$




FIGURE 4.1: The preimage of a , consisting of four points on the horizontal axis representing \mathbb{X} , is homeomorphic to the intersection of the curve with the horizontal line passing through a . The well diagram consists of four values, one for each point in the preimage.

where we simplify notation by writing $F(\rho)$ for $H(f_{\mathbb{A}}^{-1}[0, \rho])$. The image of this map, denoted as $\text{im } j_h$, is a subgroup of $F(\rho)$. The intersection of subgroups is again a subgroup.

DEFINITION. The *well group* of $f_{\mathbb{A}}^{-1}[0, r]$ is the largest subgroup $U(r) \subseteq F(r)$ such that the image of $U(r)$ in $F(\rho)$ is contained in $\bigcap_{h: \mathbb{X} \rightarrow \mathbb{Y}} \text{im } j_h$, where h ranges over all ρ -perturbations of f and $\rho = r + \delta$ for a sufficiently small $\delta > 0$.

The reason for using ρ - instead of r -perturbations is technical and will become clear later. The requirement that the perturbations be homotopic to f is not used in the proofs and can therefore be dropped. However, removing the requirement changes the well groups and therefore the meaning of our results. Similarly, we may obtain additional variants of our results by modifying the definition of a ρ -perturbation in other ways. For example, we may restrict the set of perturbations to a subset of the space of all mappings.

Example. To illustrate the definitions, let us consider the example in Figure 4.1. In this example, f is a mapping taking the reals into the reals and $\mathbb{A} = \{a\}$ is a point. The preimage of $\mathbb{A} = \{a\}$ is a set of four points separated by three critical points of f . From left to right, the values of f at these critical points are $a + r_1$, $a - r_2$, $a + r_3$. Correspondingly, the distance function, $f_a : \mathbb{X} \rightarrow \mathbb{R}$, has three critical values, namely $r_1 > r_2 > r_3$. Table 4.1 shows the ranks of $F(r)$ and $U(r)$ for values of r in the four intervals delimited by the critical values. Starting with $r = 0$, we have four points, each forming a component represented by a class in the homology group and in the well group of the sublevel set of f_a . Therefore, both groups are the same and have rank four, see the first column in Table 4.1. Growing r turns the points into intervals but leaves the groups the same until r reaches r_3 , the smallest of the three critical values. At this time, the two right intervals merge into one, so the rank of the homology group drops to three. We can find an $(r_3 + \delta)$ -perturbation whose level set at a consists of the left two points of $f^{-1}(a)$ but the right two points have

	$[0, r_3)$	$[r_3, r_2)$	$[r_2, r_1)$	$[r_1, \infty)$
$F(r)$	4	3	2	1
$U(r)$	4	2	2	0

Table 4.1: The ranks of the homology and well groups defined for the mapping f and the submanifold $\mathbb{A} = \{a\}$ in Figure 4.1.

disappeared. Indeed, the level set of every $(r_3 + \delta)$ -perturbation has a non-empty intersection with the first two but can have empty intersection with the merged interval on the right. Hence, the left two intervals have a representation in the well group, the merged interval does not, and the rank of the well group is two; see the second column in Table 4.1. The next change happens when r reaches r_2 . At this time, the middle interval merges with the merged interval on the right. The rank of the homology group drops to two, while the rank of the well group remains unchanged at two; see the third column in Table 4.1. Finally, when r reaches r_1 , the remaining two intervals merge into one, so the rank of the homology group drops to one. We can find an $(r_1 + \delta)$ -perturbation whose level set at a is empty, so the rank of the well group drops to zero; see the last column in Table 4.1.

Terminal critical values. Recall that we assume the mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$ is admissible. The initial homology group, $F(0) = H(f_{\mathbb{A}}^{-1}(0))$, has therefore finite rank, and because $U(0) \subseteq F(0)$, the initial well group has finite rank. Imagine we grow the sublevel set by gradually increasing r from zero to infinity. Since the admissibility of f does not imply the tameness of the distance function, this leaves open the possibility that $f_{\mathbb{A}}$ has infinitely many homological critical values. We call a radius, r , a *terminal critical value* of $f_{\mathbb{A}}$ if for every sufficiently small $\delta > 0$ the homomorphism from $F(r - \delta)$ to $F(r + \delta)$ applied to $U(r - \delta)$ does not give $U(r + \delta)$. In contrast to the homological critical values, there can only be a finite number of terminal critical values. To see this, we note that the set of images whose common intersection is the well group cannot decrease and the rank of the well group can therefore not increase. To state this relationship between well groups more formally, we write $f(r, s) : F(r) \rightarrow F(s)$ for the homomorphism induced by inclusion.

SHRINKING WELLNESS LEMMA. For each choice of radii $0 \leq r \leq s$, the image of the well group at r contains the well group at s , that is, $U(s) \subseteq f(r, s)(U(r))$.

It follows the only way the well group can change is by lowering its rank. Since we start with a finite rank well group at $r = 0$, there can only be finitely many terminal critical values, which we denote as $u_1 < u_2 < \dots < u_l$. To this sequence, we add $u_0 = 0$ on the left and $u_{l+1} = \infty$ on the right. It is convenient to index the homology groups and the well groups accordingly, writing $F_i = F(u_i)$ and $U_i = U(u_i)$ for all i . To these sequence, we add $F_{-1} = U_{-1} = 0$ on the left and $F_{l+2} = U_{l+2} = 0$ on the right. Furthermore, we write $f_{i,j} : F_i \rightarrow F_j$ for all feasible choices of $i \leq j$.

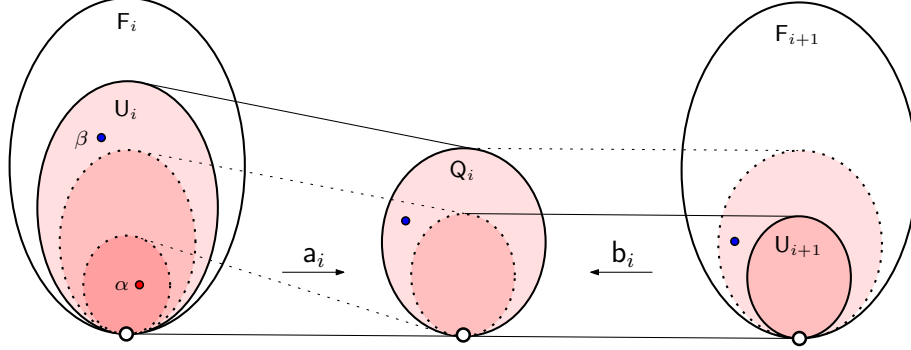


FIGURE 4.2: Connecting two consecutive well groups to the quotient group introduced between them. The class α dies a conventional death and the class β dies an unconventional death.

4.2 Well Modules

In contrast to the homology groups, the well groups of the sublevel sets do not form a filtration. Instead, they form a zigzag module. By definition of terminal critical values, the rank of U_i exceeds the rank of U_{i+1} . The rank of the image, $f_{i,i+1}(U_i)$, is somewhere between these two ranks. We call a difference between U_i and its image a *conventional death*, in which a class maps to zero, and a difference between the image and U_{i+1} an *unconventional death*, in which the image of a class lies outside the next well group. We capture both cases by inserting a new group between the contiguous well groups; see Figure 4.2. To this end, we consider the restriction of $f_{i,i+1}$ to U_i and in particular its kernel, $K_i = U_i \cap \ker f_{i,i+1}$, which we refer to as the *vanishing subgroup* of U_i . Using this subgroup, we construct $Q_i = U_i/K_i$. The forward map, $a_i : U_i \rightarrow Q_i$, is defined by mapping a class ξ to $\xi + K_i$. It is clearly surjective. The backward map, $b_i : U_{i+1} \rightarrow Q_i$, is defined by mapping a class η to $\xi + K_i$, where ξ belongs to $f_{i,i+1}^{-1}(\eta)$. This map is clearly injective. Instead of a filtration in which all maps go from left to right, we get a sequence in which the maps alternate between going forward and backward. As indicated below, every other group in the sequence is a subgroup of the corresponding homology group,

$$\begin{array}{ccccccc}
 Q_{i-1} & \xleftarrow{b_{i-1}} & U_i & \xrightarrow{a_i} & Q_i & \xleftarrow{b_i} & U_{i+1} & \xrightarrow{a_{i+1}} & Q_{i+1} \\
 & & \downarrow & & & & \downarrow & & \\
 \rightarrow & & F_i & \rightarrow & & & F_{i+1} & \rightarrow &
 \end{array}$$

We call this sequence the *well module* of f and \mathbb{A} denoted as \mathbf{U} . We remark that \mathbf{U} is a special case of a zigzag module as introduced in [3]. It is special because all forward maps are surjective and all backward maps are injective. Equivalently, there are no births other than at U_0 .




FIGURE 4.3: The left filtration decomposes U_0 into the preimages of the well groups and the preimages of their vanishing subgroups.

4.3 Well Diagrams

Well groups shrink at a finite number of radii. The history of the well groups can be captured visually in something we call the well diagram.

Left filtration. Perhaps surprisingly, the evolution of the homology classes can still be fully described by pairing births with deaths, just like for a filtration. To shed light on this construction, we follow [3] and turn a zigzag module into a filtration. In our case, all births happen at U_0 , so this transformation is easier than for general zigzag modules. Write $u_{0,i} : U_0 \rightarrow F_i$ for the restriction of $f_{0,i}$ to the initial well group. By the Shrinking Wellness Lemma, the image of this map contains the i -th well group, that is, $U_i \subseteq u_{0,i}(U_0)$. We consider the preimages of the well groups in U_0 together with the preimages of their vanishing subgroups, $A_i = u_{0,i}^{-1}(K_i)$ and $B_i = u_{0,i}^{-1}(U_i)$; see Figure 4.3. We note that $A_i/A_{i-1} \simeq \ker a_i$ and $B_i/B_{i+1} \simeq \text{cok } b_i$. In words, the first quotient represents the homology classes that die a conventional death and the second quotient represents the homology classes that die an unconventional death. As illustrated in Figure 4.3, the preimages form a nested sequence of subgroups of U_0 . Together with the inclusion maps, this gives the *left filtration* of the zigzag module,

$$0 \rightarrow A_0 \rightarrow \dots \rightarrow A_{l+1} = B_{l+1} \rightarrow \dots \rightarrow B_0 = U_0.$$

We can recover the well groups with $U_i \simeq B_i/A_{i-1}$. Recall that $U_{l+2} = 0$, which implies $K_{l+1} = U_{l+1}$. It follows that the middle two groups in the left filtration, A_{l+1} and B_{l+1} , are indeed equal.

Compatible bases. A useful property of the left filtration is the existence of compatible bases of all its groups. By this we mean a basis of U_0 that contains a basis for each A_i and each B_i . Specifically, we rewrite U_0 as a direct sum of kernels of forward maps and cokernels of backward maps:

$$U_0 \simeq \ker a_0 \oplus \dots \oplus \ker a_{l+1} \oplus \text{cok } b_l \oplus \dots \oplus \text{cok } b_0.$$

Reading this decomposition from left to right, we encounter the A_i and the B_i in the sequence they occur in the left filtration. Choosing a basis for each kernel and each

cokernel, we thus get compatible bases for all groups in the left filtration. We call this the *left filtration basis* of U_0 . It is unique up to choosing bases for the kernels and cokernels.

Consider now a homology class α in U_0 and its representation as a sum of basis vectors. We write $\alpha(\mathbf{a}_i)$ for the projection of α to the kernel of the i -th forward map, which is obtained by removing all vectors that do not belong to the basis of $\ker \mathbf{a}_i$. Similarly, we write $\alpha(\mathbf{b}_i)$ for the projection of α to the preimage of $\text{cok } \mathbf{b}_i$. Letting j be the minimum index such that $\alpha(\mathbf{a}_i) = \alpha(\mathbf{b}_i) = 0$ for all $i \geq j$, we say that α *falls ill* at u_{j+1} .

Well diagrams. Constructing the birth-death pairs that describe the well module is now easy. All classes are born at U_0 . However, to distinguish the changes in the well group from those in the homology group, we say instead that all the classes *get well* at U_0 . They fall ill later, and once they fall ill, they do not get well any more. The drop in rank from U_{i-1} to U_i is $\mu_i = \text{rank}(\ker \mathbf{a}_{i-1}) + \text{rank}(\text{cok } \mathbf{b}_{i-1})$. We thus have μ_i copies of the point $(0, u_i)$ in the diagram. There is no information in the first coordinates, which are all zero. We thus define the *well diagram* as the multiset of values u_i with multiplicities μ_i , denoting it as $\text{Dgm}(U)$. For technical reasons that will become obvious in the next section, we add infinitely many copies of 0 to this diagram. Hence, each value in $\text{Dgm}(U)$ is either 0, a positive real number, or ∞ , and the diagram itself is a multiset of values on the extended line, $\bar{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$. It has infinitely many points at 0 and a finite number of non-zero points. Figure 4.1 shows the well diagram for the example considered in §4.1.

As suggested by the heading of this chapter, we think of each point in the diagram as a measure for how resistant a homology class of $f^{-1}(\mathbb{A})$ is against perturbations of the mapping. At each well group U_i , an entire set of homology classes falls ill, and we call u_i the *robustness* of each class α in this set, denoting it as $\rho(\alpha) = u_i$.

4.4 Stability

We are interested in relating the difference between mappings to the difference between their well diagrams. After quantifying these differences, we connect parallel well modules to form new modules, and we finally prove that the well diagram is stable.

Distance between functions. Let \mathbb{X} be an m -manifold, \mathbb{Y} a Riemannian n -manifold, and $\mathbb{A} \subseteq \mathbb{Y}$ a k -manifold. Let $f, g : \mathbb{X} \rightarrow \mathbb{Y}$ be two admissible mappings and assume they are homotopic. Recall that the distance between f and g is quantified by taking the largest distance between corresponding images in \mathbb{Y} , that is,

$$\|f - g\|_\infty = \sup_{x \in \mathbb{X}} \|f(x) - g(x)\|_{\mathbb{Y}}.$$

Using \mathbb{A} , we get two functions, $f_{\mathbb{A}}, g_{\mathbb{A}} : \mathbb{X} \rightarrow \mathbb{R}$. Similar to the mappings, the distance

between them is the largest difference between corresponding values, that is,

$$\|f_{\mathbb{A}} - g_{\mathbb{A}}\|_{\infty} = \sup_{x \in \mathbb{X}} |f_{\mathbb{A}}(x) - g_{\mathbb{A}}(x)|.$$

The two distances are related. Specifically, the distance between the functions cannot exceed the distance between the mappings.

DISTANCE LEMMA. Let $f_{\mathbb{A}}, g_{\mathbb{A}} : \mathbb{X} \rightarrow \mathbb{R}$ be the functions defined by the mappings $f, g : \mathbb{X} \rightarrow \mathbb{Y}$ and the submanifold $\mathbb{A} \subseteq \mathbb{Y}$. Then $\|f_{\mathbb{A}} - g_{\mathbb{A}}\|_{\infty} \leq \|f - g\|_{\infty}$.

PROOF. We prove a stronger result, namely that the claimed inequality holds everywhere, that is,

$$|f_{\mathbb{A}}(x) - g_{\mathbb{A}}(x)| \leq \|f(x) - g(x)\|_{\mathbb{Y}} \quad (4.1)$$

at every point $x \in \mathbb{X}$. We may simplify this inequality by assuming that $f_{\mathbb{A}}(x) - g_{\mathbb{A}}(x)$ is non-negative. Suppose there exists a point $a \in \mathbb{A}$ for which $g_{\mathbb{A}}(x) = \|a - g(x)\|_{\mathbb{Y}}$. Being a metric, the distance in \mathbb{Y} obeys the triangle inequality, and in particular

$$\|a - g(x)\|_{\mathbb{Y}} + \|g(x) - f(x)\|_{\mathbb{Y}} \geq \|a - f(x)\|_{\mathbb{Y}}.$$

The right hand side is an upper bound on $f_{\mathbb{A}}(x)$ which implies (4.1). Since we did not assume that \mathbb{A} is compact, there might not be a point at which $g(x)$ attains its distance to \mathbb{A} . But for every $\delta > 0$, there is a point $a \in \mathbb{A}$ such that $g_{\mathbb{A}}(x) + \delta \geq \|a - g(x)\|_{\mathbb{Y}}$. Plugging this into the triangle inequality above gives $f_{\mathbb{A}}(x) - g_{\mathbb{A}}(x) - \delta \leq \|f(x) - g(x)\|_{\mathbb{Y}}$. Letting δ go to zero, we get (4.1). \square

Distance between diagrams. Let $\mathbf{G}(r)$ be the homology group and $\mathbf{V}(r) \subseteq \mathbf{G}(r)$ the well group of $g_{\mathbb{A}}^{-1}[0, r]$. As for f , we insert quotients between contiguous well groups and connect them with forward and backward maps to form a well module, denoted as \mathbf{V} . The corresponding well diagram, $\text{Dgm}(\mathbf{V})$, is again a multiset of points in $\bar{\mathbb{R}}$, consisting of infinitely many copies of 0 and finitely many non-zero points. Recall that the bottleneck distance between the diagrams of f and g is the length of the longest edge in the minimizing matching. Because our diagrams are one-dimensional, the bottleneck distance is easy to compute. To describe the algorithm, we order the positive points in both diagrams, getting

$$\begin{aligned} 0 &\leq u_1 \leq u_2 \leq \dots \leq u_M; \\ 0 &\leq v_1 \leq v_2 \leq \dots \leq v_M, \end{aligned}$$

where we add zeros to make sure we have two sequences of the same length. The *inversion-free matching* pairs u_i with v_i for all i . We prove that this matching gives the bottleneck distance.

MATCHING LEMMA. Assuming the above notation, the bottleneck distance between $\text{Dgm}(\mathbf{U})$ and $\text{Dgm}(\mathbf{V})$ is equal to $\max_{1 \leq i \leq M} |u_i - v_i|$.

PROOF. For a given matching, we consider the vector of absolute differences, which we sort largest first. Comparing two such vectors lexicographically, we now prove that the inversion-free matching gives the minimum vector. This implies the claimed inequality,

$$W_\infty(\text{Dgm}(\mathbf{U}), \text{Dgm}(\mathbf{V})) = \max_{1 \leq i \leq M} |u_i - v_i|,$$

To prove minimality, we consider a matching that has at least one inversion, that is, pairs (u_i, v_t) and (u_j, v_s) with $i < j$ and $s < t$. If $u_i = u_j$ or $v_s = v_t$ then switching to the pairs (u_i, v_s) and (u_j, v_t) preserves the sorted vector of absolute differences. Otherwise, the new vector is lexicographically smaller than the old vector. Indeed, the minimum of the four points is u_i or v_s and the maximum is u_j or v_t . If the minimum and the maximum are from opposite diagrams then they delimit the largest of the four absolute differences, and this largest difference belongs to the old vector. Otherwise, both absolute differences shrink when we switch the pairs. Repeatedly removing inversions as described eventually leads to the inversion-free matching, which shows that it minimizes the vector and its largest entry is the bottleneck distance. \square

Bridges. The main tool in the proof of stability is the concept of a short bridge between parallel filtrations. The length of these bridges relates to the distance between the functions defining the filtrations. Let $\varepsilon = \|f - g\|_\infty$. By the Distance Lemma, we have $\|f_{\mathbb{A}} - g_{\mathbb{A}}\|_\infty \leq \varepsilon$, which implies that the sublevel set of $g_{\mathbb{A}}$ for radius r is contained in the sublevel set of $f_{\mathbb{A}}$ for radius $r + \varepsilon$. Hence, there is a homomorphism $\mathcal{B}_r : \mathbf{G}(r) \rightarrow \mathbf{F}(r + \varepsilon)$, which we call the *bridge* from \mathbf{G} to \mathbf{F} at radius r . We use the bridge to connect the initial segment of \mathbf{G} to the terminal segment of \mathbf{F} . The endpoints of the bridge satisfy the property expressed in the Shrinking Wellness Lemma.

BRIDGE LEMMA. Let $\mathcal{B}_r : \mathbf{G}(r) \rightarrow \mathbf{F}(r + \varepsilon)$ be the bridge at r , where $\varepsilon = \|f - g\|_\infty$. Then $\mathbf{U}(r + \varepsilon) \subseteq \mathcal{B}_r(\mathbf{V}(r))$.

PROOF. Let α be a homology group in $\mathbf{U}(r + \varepsilon)$. By definition of well group, there is a sufficiently small $\delta > 0$ such that α belongs to the image of $\mathbf{H}(h^{-1}(\mathbb{A}))$ in $\mathbf{F}(r + \varepsilon)$ for every $(r + \varepsilon + \delta)$ -perturbation h of f . This includes all $(r + \delta)$ -perturbations of g . It follows that the preimage of α in $\mathbf{G}(r)$ belongs to the well group, that is, $\mathcal{B}_r^{-1}(\alpha) \in \mathbf{V}(r)$. \square

Everything we said about bridges is of course symmetric in \mathbf{F} and \mathbf{G} . In other words, $f_{\mathbb{A}}^{-1}[0, r] \subseteq g_{\mathbb{A}}^{-1}[0, r + \varepsilon]$ and there is a bridge from $\mathbf{F}(r)$ to $\mathbf{G}(r + \varepsilon)$ for every $r \geq 0$.




FIGURE 4.4: The zigzag module obtained by connecting an initial segment of V to a terminal segment of U .

New modules. We use the Bridge Lemma to construct new zigzag modules from the well modules of f and g . Specifically, we use \mathcal{B}_r to connect the initial segment of V , from $V(0)$ to $V(r)$, to the terminal segment of U , from $U(r + \varepsilon)$ to $U(\infty)$. To complete the module, we insert $Q(r) = V(r)/(\mathcal{B}_r \cap \ker \mathcal{B}_r)$ between $V(r)$ and $U(r + \varepsilon)$. The forward map, from $V(r)$ to $Q(r)$, is surjective, and the backward map, from $U(r + \varepsilon)$ to $Q(r)$ is injective; see Figure 4.4. The new zigzag module is thus of the same type as the well modules implying it has a left filtration basis that gives rise to a family of compatible bases for the groups in the left filtration.

A particular construction starts with the filtrations $F(0) \rightarrow \dots \rightarrow F(\infty)$ and $G(0) \rightarrow \dots \rightarrow G(\infty)$ and adds $\mathcal{B}_0 : G(0) \rightarrow F(\varepsilon)$. Following the bridge from G to F at 0, we get a new filtration and a new zigzag module, denoting the latter as W ; see Figure 4.5. The decomposition of $W(0) = V(0)$ by the left filtration of W is similar to the decomposition of $U(0)$ by the left filtration of U ; see Figure 4.3. Letting i be the index such that $u_i \leq \varepsilon < u_{i+1}$, we have $F(\varepsilon) = F_i$ and $U(\varepsilon) = U_i$. The classes in A_{i-1} and in U_0/B_i die before we reach $F(\varepsilon)$. The remaining classes form $U(\varepsilon) \simeq B_i/A_{i-1}$. Correspondingly, there are homology classes in $W(0)$ that die before we reach $F(\varepsilon)$, namely the ones in the kernel of the forward map, from $W(0)$ to $Q(0)$, and in the preimage of the cokernel of the backward map, from $U(\varepsilon)$ to $Q(0)$. The remaining classes form $W(\varepsilon) \simeq \mathcal{B}_0^{-1}(U(\varepsilon))/(\mathcal{B}_0 \cap \ker \mathcal{B}_0)$. The two quotient groups, $U(\varepsilon)$ and $W(\varepsilon)$, are decomposed in parallel so that choosing a basis for $U(\varepsilon)$ gives one for $W(\varepsilon)$. This will be useful shortly.

Main result. We are now ready to state and prove the stability of the well diagram.

STABILITY THEOREM FOR WELL DIAGRAMS. Let U, V be the well modules of the functions $f_{\mathbb{A}}, g_{\mathbb{A}}$ defined by the admissible, homotopic mappings $f, g : \mathbb{X} \rightarrow \mathbb{Y}$, where \mathbb{X}, \mathbb{Y} , and $\mathbb{A} \subseteq \mathbb{Y}$ are manifolds of finite dimension and \mathbb{Y} is Riemannian. Then $W_{\infty}(\text{Dgm}(U), \text{Dgm}(V)) \leq \|f - g\|_{\infty}$.

PROOF. We construct a bijection from $\text{Dgm}(U)$ to $\text{Dgm}(V)$ such that the L_{∞} -distance between matched points is at most $\varepsilon = \|f - g\|_{\infty}$. Specifically, we match each point




FIGURE 4.5: The four curves represent four filtrations as well as four the zigzag modules. The middle two are constructed from the outer two by adding bridges connecting the dots.

$u \leq \varepsilon$ in $\text{Dgm}(\mathbf{U})$ with a copy of 0 in $\text{Dgm}(\mathbf{V})$, and we use the parallel bases of $\mathbf{U}(\varepsilon)$ and $\mathbf{W}(\varepsilon)$ for the rest, where \mathbf{W} is the zigzag module obtained by adding the bridge from \mathbf{G} to \mathbf{F} at radius 0, as described above.

Let α belong to the left filtration basis of $\mathbf{U}(0)$ such that its image belongs to the basis of $\mathbf{U}(\varepsilon)$. Let r be the value at which α falls ill and note that $r > \varepsilon$. Let β belong to the left filtration basis of $\mathbf{V}(0) = \mathbf{W}(0)$ such that the images of α and β in $\mathbf{W}(\varepsilon) = \mathbf{U}(\varepsilon)$ coincide. We now construct yet another zigzag module, by adding a first bridge from $\mathbf{G}(r - \varepsilon - \delta)$ to $\mathbf{F}(r - \delta)$ and a second bridge from $\mathbf{F}(r + \delta)$ back to $\mathbf{G}(r + \varepsilon + \delta)$, where $\delta > 0$ is sufficiently small such that there are no deaths in the interval $[r - \delta, r + \delta]$, except possibly at r . We denote the resulting module by \mathbf{X} ; see Figure 4.5. We note that all maps between groups are induced by inclusion so that the diagram formed by the filtrations and the bridges between them commutes.

By construction, the image of β in $\mathbf{F}(r - \delta)$ is non-zero and belongs to $\mathbf{U}(r - \delta)$. In contrast, the image of β in $\mathbf{F}(r + \delta)$ is either zero or lies outside $\mathbf{U}(r + \delta)$. Applying the Bridge Lemma going backward along the first bridge, we note that the image of $\beta \in \mathbf{W}(0) = \mathbf{X}(0)$ in $\mathbf{G}(r - \varepsilon - \delta)$ is non-zero and belongs to $\mathbf{V}(r - \varepsilon - \delta)$. Applying the Bridge Lemma going forward along the second bridge, we note that the image of β in $\mathbf{G}(r + \varepsilon + \delta)$ is either zero or lies outside $\mathbf{V}(r + \varepsilon + \delta)$. Since we can choose $\delta > 0$ as small as we like, this implies that β falls ill somewhere in the interval $[r - \varepsilon, r + \varepsilon]$. In the matching, this radius is paired with r , the radius at which α falls ill in \mathbf{U} . The absolute difference between the two radii is at most ε , as required. \square

4.5 Closed Forms

The definition of the well group involves intersecting an infinite number of images. We take a closer look at two classes of mappings and provide a closed form expression for their well groups.

4.5.1 Equidimensional manifolds

Let K be a triangulation of a compact m -manifold without boundary and $f : |K| \rightarrow \mathbb{R}^m$ a generic PL mapping. Assume no vertex of K maps to 0 and set $\mathbb{A} = \{0\}$. Now subdivide K so that the preimage, $f^{-1}(0)$, is a subcomplex of the subdivision, $\text{Sd } K$. Define the PL mapping $f_{\mathbb{A}} : |\text{Sd } K| \rightarrow \mathbb{R}$ by setting each vertex v to $\|f(v)\|_2$. An arbitrarily small perturbation to f ensures that no two vertices of $\text{Sd } K - f^{-1}(0)$ map to the same value under $f_{\mathbb{A}}$.

By assumption, each point $x \in f^{-1}(\mathbb{A})$ lies in the interior of an m -simplex $\sigma \in K$. As the threshold value r increases, the inverse $f_{\mathbb{A}}^{-1}[0, r]$ grows from a collection of isolated points into a collection of m -dimensional components. Let C be a connected component of $f_{\mathbb{A}}^{-1}[0, r]$. The *degree mod 2* of f restricted to C is the number of points in the intersection $C \cap f^{-1}(0) \bmod 2$. We denote this number as $\deg f|_C$. We call the component C *well* if $\deg f|_C$ is non-zero otherwise it is *ill*. See §2.3.4 for more details about the degree mod 2 of a mapping.

CLOSED FORM 1. The rank of the well group $U(r)$ is the number of well components in $f_{\mathbb{A}}^{-1}[0, r]$.

PROOF. There are two parts to the proof. We pick a component C_r in $f_{\mathbb{A}}^{-1}[0, r]$. If C_r is well, then we show that the preimage $h^{-1}(0)$ of every r -perturbation h of f intersects C_r . If C_r is ill, then we construct an r -perturbation h of f such that $h^{-1}(0) \cap C_r = \emptyset$.

Let C_r be a component of $f_{\mathbb{A}}^{-1}[0, r]$. Assuming C_r is well, the component $C_{r+\rho}$ in $\mathbb{X}_{r+\rho}$ containing C_r is also well, for any $0 < \rho \leq \delta$ where δ is a sufficiently small value greater than zero. Now let $h : |K| \rightarrow \mathbb{R}^m$ be an $(r + \rho)$ -perturbation of f , where $0 < \rho < \delta$. Letting $g_t = (1 - t)f + th$ be the straight line homotopy taking f to h , the image of $\text{Bd } C_{r+\delta}$ under g_t never touches \mathbb{A} simply because the image of $\text{Bd } C_{r+\delta}$ under f is too far from \mathbb{A} . Therefore, the sum of the intersection numbers of each point in $g_t^{-1}(\mathbb{A}) \cap C_{r+\delta}$ remains constant. In other words, the degree of $C_{r+\delta}$ restricted to h is the same as the degree of f restricted to $C_{r+\delta}$. We now know $h^{-1}(0)$ intersects $C_{r+\delta}$, for every $(r + \rho)$ -perturbation of f . This implies that the class in $H_0(\mathbb{X}_{r+\delta})$ representing the component $C_{r+\delta}$ belongs to the well group.

We now assume C_r is ill and construct a perturbation h of f such that $h^{-1}(0) \cap C_r = \emptyset$. As the radius increases from zero to r , at some point two well components contained in C_r merge at a critical vertex, say v . The resulting component C_0 is ill. The perturbation h needs to move v beyond \mathbb{A} which it can do without changing the f values of vertices outside C_0 . See Figure 4.6. The number of points in C_r mapping to \mathbb{A} is reduced by two. As the radius increases further, the same thing may happen again. That is, two well components merge to form an ill component C_1 . As before, construct a perturbation further reducing the number of points in C_r mapping to \mathbb{A} . The two perturbations differ from each other in two disjoint vertex subsets of C_r . The continuation of this process results in the desired perturbation h . \square




FIGURE 4.6: The perturbation moves the vertex v beyond \mathbb{A} .

The well diagram is easily computable. A simple union-find data structure is sufficient to keep track of the components in the sublevel set. Create a singleton set for each point in $f^{-1}(\mathbb{A})$ and mark each set well. As r increases, one of two things happen to the components of the sublevel set $f_{\mathbb{A}}^{-1}[0, r]$: two components merge or a new component is born. In the case a new component is born, create a new set and mark it ill. In the case two components merge, perform a find on the two sets representing the two components and mark the union of the two sets well if exactly one of them is well otherwise mark it ill. For each component that goes ill, place a point on the well diagram at the radius it fell ill.

4.5.2 Real-valued Functions

Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be a real-valued function on a compact manifold \mathbb{X} without boundary and $\mathbb{A} = \{0\}$. The well diagram of f and \mathbb{A} encodes the robustness the level set $f^{-1}(\mathbb{A})$. We show that the well group is determined by the image of exactly two perturbations. This is stated precisely in Closed Form 2, which is a simple consequence of the Mayer-Vietoris Lemma.

MAYER-VIETORIS LEMMA. Suppose we write a topological space \mathbb{Y} as $\mathbb{Y} = \mathbb{C} \cup \mathbb{D}$ with $\mathbb{E} = \mathbb{C} \cap \mathbb{D}$. If a class $\alpha \in \mathbf{H}(\mathbb{Y})$ is supported by both \mathbb{C} and \mathbb{D} , then it is also supported by \mathbb{E} .

PROOF. Let $\mathbf{c} : \mathbf{H}(\mathbb{C}) \rightarrow \mathbf{H}(\mathbb{Y})$, $\mathbf{d} : \mathbf{H}(\mathbb{D}) \rightarrow \mathbf{H}(\mathbb{Y})$, $\mathbf{e}_{\mathbb{C}} : \mathbf{H}(\mathbb{E}) \rightarrow \mathbf{H}(\mathbb{C})$, and $\mathbf{e}_{\mathbb{D}} : \mathbf{H}(\mathbb{E}) \rightarrow \mathbf{H}(\mathbb{D})$ be the homomorphisms induced by inclusion. Consider now the following portion of the Mayer-Vietoris long exact sequence:

$$\mathbf{H}(\mathbb{E}) \xrightarrow{(\mathbf{e}_{\mathbb{C}}, \mathbf{e}_{\mathbb{D}})} \mathbf{H}(\mathbb{C}) \oplus \mathbf{H}(\mathbb{D}) \xrightarrow{\mathbf{c} - \mathbf{d}} \mathbf{H}(\mathbb{Y}).$$

Choose $\alpha_{\mathbb{C}} \in \mathbf{H}(\mathbb{C})$ and $\alpha_{\mathbb{D}} \in \mathbf{H}(\mathbb{D})$ such that $\mathbf{c}(\alpha_{\mathbb{C}}) = \alpha$ and $\mathbf{d}(\alpha_{\mathbb{D}}) = \alpha$. The homomorphism $\mathbf{c} - \mathbf{d}$ takes the pair $(\alpha_{\mathbb{C}}, \alpha_{\mathbb{D}})$ to zero implying the pair belongs to the image of $(\mathbf{e}_{\mathbb{C}}, \mathbf{e}_{\mathbb{D}})$. This of course implies that α is supported by \mathbb{E} . \square

Recall, \mathbb{X}_r is the sublevel set $f_{\mathbb{A}}^{-1}[0, r]$. If $h : \mathbb{X} \rightarrow \mathbb{R}$ is an r -perturbation of f , then the inclusion of $h^{-1}(0)$ into \mathbb{X}_r induces a homomorphism $j_h : \mathbf{H}(h^{-1}(0)) \rightarrow \mathbf{H}(\mathbb{X}_r)$. The well group $\mathbf{U}(r) \subseteq \mathbf{H}(\mathbb{X}_r)$ is

$$\mathbf{U}(r) = \bigcap_{\|f-h\|_{\infty} \leq r} \text{im } j_h.$$

This definition of the well group is slightly different from the original definition given in §4.5.1 as it puts the death of a class right before a terminal critical value while in this definition, it happens *at* the value. The well diagrams are the same under both definitions. Now define $\mathbb{B} = f^{-1}(-r)$ and $\mathbb{T} = f^{-1}(r)$ and let $\mathbf{b} : \mathbf{H}(\mathbb{B}) \rightarrow \mathbf{H}(\mathbb{X}_r)$ and $\mathbf{t} : \mathbf{H}(\mathbb{T}) \rightarrow \mathbf{H}(\mathbb{X}_r)$ be the homomorphisms induced by inclusion.

CLOSED FORM 2. $\mathbf{U}(r) = \text{im } \mathbf{b} \cap \text{im } \mathbf{t}$, for every $r \geq 0$.

PROOF. To show $\mathbf{U}(r) \subseteq \text{im } \mathbf{b} \cap \text{im } \mathbf{t}$, consider a class $\alpha \in \mathbf{U}(r)$. Define $h_0 = f + r$ and $h_1 = f - r$ and note that they are r -perturbations of f with $h_0^{-1}(0) = \mathbb{B}$ and $h_1^{-1}(0) = \mathbb{T}$. By definition of well groups, α is supported by every r -perturbation of f and therefore by h_0 and h_1 . It follows $\alpha \in \text{im } \mathbf{b} \cap \text{im } \mathbf{t}$.

To show $\text{im } \mathbf{b} \cap \text{im } \mathbf{t} \subseteq \mathbf{U}(r)$, we consider a class $\alpha \in \text{im } \mathbf{b} \cap \text{im } \mathbf{t}$ and show α is supported by $h^{-1}(0)$ for each r -perturbation h of f . We define $\mathbb{C} = \{x \in \mathbb{X}_r \mid h(x) \leq 0\}$ and $\mathbb{D} = \{x \in \mathbb{X}_r \mid h(x) \geq 0\}$. Note that $\mathbb{C} \cup \mathbb{D} = \mathbb{X}_r$ while $\mathbb{C} \cap \mathbb{D} = h^{-1}(0)$. Furthermore, the inequality $\|f - h\|_{\infty} \leq r$ implies that $\mathbb{B} \subseteq \mathbb{C}$ and $\mathbb{T} \subseteq \mathbb{D}$. By the Mayer-Vietoris Lemma, α is supported by $h^{-1}(0)$ as required. \square

4.6 Applications

We now explore a few applications of the well groups.

4.6.1 Fixed Points of Mappings

A *fixed point* of a continuous mapping from a topological space to itself is a point that is its own image. Assuming this space is the m -dimensional Euclidean space and b is the mapping, we introduce a mapping $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$ defined by $f(x) = x - b(x)$. A fixed point of b is a root of f , that is, $f(x) = 0$. Writing $\mathbb{X} = \mathbb{Y} = \mathbb{R}^m$ and $\mathbb{A} = \{0\}$, the origin of \mathbb{R}^m , we get the setting studied in this chapter. Each fixed point x of b corresponds to a class in the 0-dimensional homology group of $f^{-1}(0)$. Using well groups, we assign a non-negative robustness measure, $\varrho(x)$, to x . It gives the magnitude of perturbation necessary to remove this fixed point. This does not mean that a perturbation of smaller magnitude has a fixed point at precisely the same location but rather that it has one or more fixed points in lieu of x . If the fixed point x has the maximum robustness of any fixed point of b , then every ρ -perturbation of f , for $\rho < \varrho(x)$, has at least one fixed point. This implication suffices to give a new

proof of a classic topological result on fixed points. Let \mathbb{B}^m be the closed unit ball in \mathbb{R}^m .

BROUWER'S FIXED POINT THEOREM. Every continuous mapping $b : \mathbb{B}^m \rightarrow \mathbb{B}^m$ has a fixed point.

PROOF. Extend b to a mapping from \mathbb{R}^m to \mathbb{R}^m by defining $b(x)$ equal to its value at $x/\|x\|_2$ whenever $x \notin \mathbb{B}^m$. Define $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$ as $f(x) = x - b(x)$ and let $g : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be the identity map, that is $g(x) = x$. We may assume that f is admissible, else the homology group of $f^{-1}(0)$ has infinite rank and f has infinitely many roots. The other mapping, g , is clearly admissible, with a single root at $x = 0$. The distance between the two mappings is

$$\begin{aligned} \|f - g\|_\infty &= \sup_{x \in \mathbb{R}^m} \|f(x) - g(x)\|_2 \\ &= \sup_{x \in \mathbb{R}^m} \|b(x)\|_2, \end{aligned}$$

which is at most 1. The well diagram of the identity map, g , consists of a single, non-zero point at plus infinity. The Stability Theorem for Well Diagrams implies that the well diagram of f also has a point at plus infinity. But this implies that f has a root and, equivalently, that b has a fixed point. \square

The above reduction of fixed points to a transversality setting uses the difference between two points, an operation not available if the mapping $b : \mathbb{M} \rightarrow \mathbb{M}$ is defined on a general Riemannian manifold. In this case, we can use the correspondence between the fixed points of b and the intersection points between the graph of b and the diagonal in $\mathbb{M} \times \mathbb{M}$. Let $\mathbb{X} = \mathbb{M}$, let $\mathbb{Y} = \mathbb{M} \times \mathbb{M}$, and let $\mathbb{A} = \{(x, x) \mid x \in \mathbb{M}\}$ be the diagonal. We think of \mathbb{Y} as a fiber bundle over the base space \mathbb{M} with fiber \mathbb{M} . The mapping b defines a section $f : \mathbb{X} \rightarrow \mathbb{Y}$ of the fiber bundle. A *section* is a continuous assignment a point x in the base space, \mathbb{M} , to a point in the fiber above x . Define the section $f : \mathbb{X} \rightarrow \mathbb{Y}$ by taking each point $x \in \mathbb{X}$ to $f(x) = (x, b(x))$. The set of fixed points of b is homeomorphic to the inverse $f^{-1}(\mathbb{A})$. We are now closer to the well group setting. The only thing that is missing is a metric on \mathbb{Y} . Define the distance between two points $x = (x', x'')$ and $y = (y', y'')$ in $\mathbb{Y} = \mathbb{M} \times \mathbb{M}$ equal to

$$\|x - y\|_{\mathbb{Y}} = \begin{cases} \infty & \text{if } x' \neq y'; \\ \|x'' - y''\|_{\mathbb{M}} & \text{if } x' = y'. \end{cases}$$

The distance between two sections $f, g : \mathbb{X} \rightarrow \mathbb{Y}$ is then $\sup_{x \in \mathbb{X}} \|f(x) - g(x)\|_{\mathbb{Y}}$. Note that there is nothing special about this metric. For example, one may decide to stretch the metric over a subset of the base space. This will result in different robust values.

4.6.2 Periodic Orbits

We now consider iterations of self mappings. Letting \mathbb{M} be a Riemannian manifold and $f : \mathbb{M} \rightarrow \mathbb{M}$ a mapping, we write $f^j : \mathbb{M} \rightarrow \mathbb{M}$ for the j -fold composition of f with itself. A sequence

$$F_j(x) = (x, f(x), f^2(x), \dots, f^{j-1}(x))$$

is an *order- j periodic orbit* of f if $f^j(x) = f \circ f^{j-1}(x) = x$. It is straightforward to see the following relationship between f and its j -fold composite.

ORBIT LEMMA. A point $x \in \mathbb{M}$ is a fixed point of f^j iff $F_j(x)$ is an order- j periodic orbit of f .

We can therefore use well groups to measure the robustness of x , that is, to determine how much f^j needs to be perturbed to remove the fixed point. However, it would be more interesting to measure how much f needs to be perturbed to remove the periodic orbit. This is different because not every mapping can be written as the j -fold composite of another mapping. Adapting the framework accordingly is not difficult. Substituting perturbations h of f for those of f^j , we intersect the images of the homomorphisms induced by h^j . Call the resulting values the robustness of the periodic orbits of order j .

4.6.3 Vector Fields

A vector field is the continuous assignment of a point to each fiber of the tangent bundle. More precisely, a vector field is a continuous *section* $v : \mathbb{M} \rightarrow \text{T}\mathbb{M}$ of the tangent bundle. A point $x \in \mathbb{M}$ is a *fixed point* of the vector field if $v(x) = 0$. Writing $\mathbb{X} = \mathbb{M}$, $\mathbb{Y} = \text{T}\mathbb{M}$, and \mathbb{A} equal to the zero section of the tangent bundle, the intersection of the image of v with \mathbb{A} is homeomorphic to the set of fixed points of v . Define the distance between a point $a \in \text{T}_x\mathbb{M}$, denoted as (x, a) , and a point $b \in \text{T}_y\mathbb{M}$, denoted as (y, b) , equal to

$$\|(x, a) - (y, b)\|_{\mathbb{Y}} = \begin{cases} \infty & \text{if } x \neq y; \\ \|a - b\|_2 & \text{if } x = y. \end{cases}$$

A vector field $w : \mathbb{M} \rightarrow \text{T}\mathbb{M}$ is an r -perturbation of v if $\sup_{x \in \mathbb{M}} \|w(x) - v(x)\|_{\mathbb{Y}} \leq r$. We are now in the setting of well groups. Each fixed point of v corresponds to a 0-dimensional homology class of $v^{-1}(\mathbb{A})$ and the robustness of the homology class gives the magnitude of the perturbation necessary to remove it. Once again, this metric is nothing special. One may just as easily use another metric on the tangent bundle.

Now let us consider k vector fields on \mathbb{M} where k is at most the dimension, m , of \mathbb{M} . We describe the k vector fields as a single section $s : \mathbb{M} \rightarrow \oplus_k \text{T}\mathbb{M}$ of the k -times direct sum of $\text{T}\mathbb{M}$. The points of interest are the points $x \in \mathbb{M}$ where the k

vectors fail to span a k -dimensional subspace of the tangent space $T_x\mathbb{M}$. Denote by $\Sigma(s) \subseteq \mathbb{M}$ the collection of these points. Letting $\mathbb{S} \subset \oplus_k T\mathbb{M}$ be the set of k -vectors failing to span a k -dimensional subspace, $\Sigma(s) = s^{-1}(\mathbb{S})$. Once $\oplus_k T\mathbb{M}$ is equipped with a metric, we can use well groups to assign to each homology class of $\Sigma(s)$ a robustness. This may have applications in correlating multiple vector fields on the same domain [11].

4.6.4 Simplification of the Contour

Imagine an object in \mathbb{R}^3 . We get cues of its shape from the curves on its boundary where the surface normal is orthogonal to the viewing direction [23]. The view is an orthogonal projection to the plane and its *contour*, is the set of critical values. Now imagine the object is wrinkly, for example a prune. The contour curve defining the outline of the prune is more stable than the contour curve defined by a single wrinkle. This is because we can remove the outline of the prune by taking the entire prune to a single point in the plane but a smaller perturbation smoothes out the wrinkle. The amount of perturbation necessary to remove a single wrinkle depends on the thickness of the wrinkle. We use well groups to simplify the contour of f by removing contour lines defined by thin wrinkles.

Assume $f : \mathbb{M} \rightarrow \mathbb{R}^2$ is a generic smooth mapping from an orientable compact 2-manifold without boundary to the plane. In particular, this means that the inverse of any point in the plane is a finite number of points. Denote by $\Upsilon(f)$ the contour of f or in other words, the set of critical values of f . Choose a point $x \in \mathbb{M}$ and set $\mathbb{A} = f(x)$. The well groups assign to each homology class of $f^{-1}(\mathbb{A})$ a robustness. Let $\varrho(x)$ be the robustness of the component in $f^{-1}(\mathbb{A})$ containing x . Recall the robustness of x is the radius where the well component containing x merges with another well component at a singular point of f . The robustness of x is then the distance between the image of x and the image of this singular point. Let $\varrho : \mathbb{M} \rightarrow \mathbb{R}$ be the mapping assigning to each point $x \in \mathbb{M}$ its robustness. The singular set, $\Sigma(f)$, of f is the zero set of ϱ . The mapping ϱ is usually not continuous. See Figure 4.7 for an example.

A thin wrinkle is a component of $\mathbb{M} - \Sigma(f)$ with small robustness. By removing points with small robustness, we remove the thin wrinkles and the contour curves defined by them. The super level set $\mathbb{M}^r = \varrho^{-1}[r, \infty)$ is \mathbb{M} eroded from its singularity set by an amount r . The image, $f(\varrho^{-1}(r))$, of the level set $\varrho^{-1}(r)$ is a *simplification* of the contour $\Upsilon(f)$. See Figure 4.8. One drawback of this definition is that we erode away from every point in the singular set and therefore erode into the more robust wrinkles. It would be nice to see a method that erodes only the thin wrinkles leaving the more robust wrinkles intact.




FIGURE 4.7: Shown is a wrinkle, its cross section, and a graph of ρ along the cross section. The robustness of each singular point is zero. Recall, the robustness of each point x , y , and z is the radius where the well component containing each point goes ill. In this case, the three well components merge at the same radius producing a well component. This phenomenon explains the discontinuity at these three points.




FIGURE 4.8: By removing points with low robustness, the surface erodes away from the singular set. Here we see the cross section of the wrinkle eroding away from the singular set.

4.7 Remarks

We took the idea of persistence to the very general framework of well groups. Now that we have the framework in place, many questions arise. We discuss a few now.

Transversality and well groups. Let $f : \mathbb{X} \rightarrow \mathbb{Y}$ be a smooth mapping between two manifolds and $\mathbb{A} \subset \mathbb{Y}$ a submanifold. The graph of f , $\text{gf } f$, is a submanifold of the product space $\mathbb{X} \times \mathbb{Y}$ and the inverse $f^{-1}(\mathbb{A})$ is homeomorphic to the intersection $\text{gf } f \cap (\mathbb{X} \times \mathbb{A})$. Recall, the mapping f is transverse to \mathbb{A} if at every point in the intersection of the two subspaces, the two tangent spaces meet in a non-zero angle. Transverse intersections are, by definition, stable. That is, given $f \pitchfork \mathbb{A}$, there is an open neighborhood of f in $C^\infty(\mathbb{X}, \mathbb{Y})$ such that every mapping in this neighborhood is transverse to \mathbb{A} . Recall, the basis of the topology on $C^\infty(\mathbb{X}, \mathbb{Y})$ are sets of mappings with close derivatives. Transverse intersections are stable under infinitesimal perturbations.

Well groups, on the other hand, offer a global analysis of the intersection. Here we assume a continuous mapping $f : \mathbb{X} \rightarrow \mathbb{Y}$, $\mathbb{A} \subseteq \mathbb{Y}$ a subspace, and \mathbb{Y} has a metric. The three spaces \mathbb{X} , \mathbb{Y} , and \mathbb{A} need not be manifolds. The only requirement is that the inverse, $f^{-1}(\mathbb{A})$, has finite rank homology. The well module assigns to each

homology class in the preimage of \mathbb{A} a robustness. The robustness of a homology class is a real value that is the minimum distance to a perturbation of f that removes that homology class from the intersection. The robustness information is captured in the well diagram. By definition of well groups, the well diagrams are stable. In other words, we can approximate the well diagram of f with the well diagram of another mapping g close to f . In fact, the closer the mapping g is to f the better the approximation. Here, the closeness of two mappings is not determined by their derivatives but by the distance in \mathbb{Y} between the two images of each point in \mathbb{X} .

Computing well groups. The well group is the intersection of an infinite number of images. As we have seen, for some cases, the well groups are efficiently computable. For the rest, the answer is not so clear.

Well groups are computable in the case of PL mappings $f : |K| \rightarrow \mathbb{R}^n$ and $\mathbb{A} = \{0\}$ a single point. If q is the number of vertices in K , then the space C of all PL mappings is isomorphic to \mathbb{R}^{qn} . Now let $h : |K| \rightarrow \mathbb{R}^n$ be a perturbation of f and $H : |K| \times [0, 1] \rightarrow \mathbb{R}^n$ the straight line homotopy connecting f to h . If, in the homotopy, the image of any $(n-1)$ -simplex never crosses the origin, then the images of the two homomorphisms $j_f : \mathbf{H}(f^{-1}(0)) \rightarrow \mathbf{H}(f_{\mathbb{A}}^{-1}[0, r])$ and $j_h : \mathbf{H}(h^{-1}(0)) \rightarrow \mathbf{H}(f_{\mathbb{A}}^{-1}[0, r])$ equal. Therefore, we can use the $(n-1)$ -simplices to partition C into a finite number of cells and consider just a single perturbation per cell. Finding these cells is costly because it involves computing the arrangement of hyper-surfaces in \mathbb{R}^{qn} . Needless to say, a better algorithm would be nice.

Contours. Let $f : \mathbb{X} \rightarrow \mathbb{R}$ be a Morse function on a surface \mathbb{X} . The set of critical values, $\Upsilon(f)$, is a collection of points. As the mapping f varies, the set $\Upsilon(f)$ varies with points disappearing and new ones appearing. Ordinary persistence tells us that these critical values appear and disappear in pairs. The robustness of each critical value is its persistence. That is, if a critical value v_i is paired with v_j , then the persistence of each is $|v_i - v_j|$. If v_i is paired with infinity, then its persistence is infinity. The stability of persistence diagrams tells us that the persistence of each critical value is stable.

Now we come back to contours. Let $f : \mathbb{X} \rightarrow \mathbb{R}^2$ be a smooth mapping from a surface \mathbb{X} to the plane and $\Upsilon(f)$ its critical values. $\Upsilon(f)$ is a collection of curves in \mathbb{R}^2 and it varies as the mapping f varies. A small perturbation of f may introduce many curves, but in the sense we discussed in §4.6.3, these curves are not very stable. Can we assign to each homology class in $\mathbf{H}(\Upsilon(f))$ a robustness? It is clear that some classes in $\mathbf{H}(\Upsilon(f))$ are more stable than others. That is, some classes require a larger perturbation to remove than others. How can we use well groups to assign robustness to each class in the homology of $\Upsilon(f)$?

Untangling curves. Imagine a closed curve in the plane. The curve is a mapping $f : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ from the circle to the plane. We assume the mapping f is continuous but not necessarily an embedding. In other words, the curve is allowed to cross




FIGURE 4.9: There are three pairs of crossing points. Untangling the loop on the right requires a smaller perturbation than the one on the left. In this sense, the loop on the right is more stable than the loop on the left.

itself. Some of these crossings are more stable than others. See Figure 4.9. Using well groups, we can assign to each crossing a robustness. Define $F : \mathbb{S}^1 \times \mathbb{S}^1 \rightarrow \mathbb{R}^2 \times \mathbb{R}^2$ as $F(x, y) = (f(x), f(y))$. Letting $\mathbb{A} = \{(a, a) \in \mathbb{R}^2 \times \mathbb{R}^2 \mid a \in \mathbb{R}^2\}$ be the diagonal, the inverse $F^{-1}(\mathbb{A})$ is the set of pairs (x, y) such that $f(x) = f(y)$. In other words, $F^{-1}(\mathbb{A})$ are the *crossing points* of the mapping f plus the diagonal pairs (x, x) . Equipping $\mathbb{R}^2 \times \mathbb{R}^2$ with a suitable metric, the well groups assign to each pair of crossing points a robustness. However, the well groups are computed using perturbations of the mapping F not f . A perturbation of F , say H , may not be realizable by a curve. That is, it is not always possible to write $H(x, y) = (h(x), h(y))$ for some continuous mapping $h : \mathbb{S}^1 \rightarrow \mathbb{R}^2$. It is therefore unclear what the robustness of each pair of crossing points really means. In particular, what is the relationship between the robustness of a crossing and the amount we must perturb the mapping to remove that crossing?

This idea can be generalized to study the self-intersection of a surface in \mathbb{R}^3 . Given a mapping $f : \mathbb{X} \rightarrow \mathbb{R}^3$ from a surface to \mathbb{R}^3 , we can setup a mapping similar to F allowing us to assign to each homology class in the self-intersection a robustness. Again, we run into the same problem. How does one interpret the robustness of each homology class in the self-intersection?

Bibliography

- [1] P. Bendich, H. Edelsbrunner, D. Morozov, and A. Patel. Robustness of level sets and interval preimages. Manuscript, Computer Science, Duke University.
- [2] O. Burchard and G. de Rham. Sur certaines applications génériques d'une variété close à 3 dimensions dans le plan. *L'Enseign. Math.*, 20:275–292, 1974.
- [3] G. Carlsson and V. de Silva. *Zigzag persistence*. Dept. Mathematics, Stanford Univ., Stanford, California, 2008. Manuscript.
- [4] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37:103–120, 2007.
- [5] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Extending persistence using Poincaré and Lefschetz duality. *Found. Comput. Math.*, 9(1):79–103, 2009.
- [6] D. Cohen-Steiner, H. Edelsbrunner, and D. Morozov. Vines and vineyards by updating persistence in linear time. In *SCG '06*, pages 119–126, New York, NY, USA, 2006. ACM.
- [7] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [8] H. Edelsbrunner and J. Harer. Jacobi sets of multiple Morse functions. In *Foundations of Computational Mathematics, Minneapolis 2002*, pages 35–57. Cambridge University Press, 2004.
- [9] H. Edelsbrunner and J. Harer. *Persistent homology – a Survey*, volume 453. American Mathematical Society, 2008.
- [10] H. Edelsbrunner, J. Harer, A. Mascarenhas, V. Pascucci, and J. Snoeyink. Time-varying Reeb graphs for continuous space-time data. *Comput. Geom. Theory Appl.*, 41:149–166, 2008.
- [11] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci. Local and global comparison of continuous functions. In *VIS '04: Proceedings of the conference on Visualization '04*, pages 275–280, Washington, DC, USA, 2004. IEEE Computer Society.

- [12] H. Edelsbrunner, J. Harer, and A. Patel. Reeb spaces of piecewise linear functions. In *SCG '08*, pages 242–250, New York, NY, USA, 2008. ACM.
- [13] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28:511–533, 2002.
- [14] H. Edelsbrunner, D. Morozov, and A. Patel. Quantifying transversality by measuring the robustness of intersections. Manuscript, Computer Science, Duke University.
- [15] H. Edelsbrunner, D. Morozov, and A. Patel. The stability of the apparent contour of an orientable 2-manifold. In V. Pascucci and J. Tierny, editors, *Topological Methods in Data Analysis and Visualization: Theory, Algorithms, and Applications*. Springer-Verlag, 2010.
- [16] R. D. Edwards. Approximating certain cell-like maps by homeomorphisms. *Notices Amer. Math. Soc.*, 24, 1977.
- [17] Y. K. S. Furuya. Sobre aplicações genéricas $M^4 \rightarrow \mathbb{R}^2$. Technical report, Tese, Departamento de Matemática, ICMSC-USP, São Carlos, Brazil, 1986.
- [18] V. Guillemin and M. Golubitsky. *Stable Mappings and their Singularities*. Springer-Verlag, New York Inc., 1973.
- [19] V. Guillemin and A. Pollack. *Differential Topology*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1974.
- [20] W. Harvey, Y. Wang, and R. Wenger. A randomized $O(m \log m)$ time algorithm for computing Reeb graph of arbitrary simplicial complexes. In *SCG '10*, New York, NY, USA, 2010. ACM.
- [21] M. Hirsch. *Differential Topology*. Springer-Verlag, New York Inc., 1976.
- [22] M. Kobayashi and O. Saeki. Simplifying stable mappings into the plane from a global viewpoint. *Trans. Amer. Math. Soc.*, 348:2607–2636, 1996.
- [23] J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.
- [24] L. Kushner, H. Levine, and J. P. Porto. Mapping three-manifolds into the plane I. *Boletín de las sociedad matemática Mexicana*, 29:11–33, 1984.
- [25] H. Levine. *Classifying Immersions into \mathbb{R}^4 over Stable Maps of 3-Manifolds in \mathbb{R}^2* . Lecture Notes in Mathematics 1157. Springer-Verlag, Berlin, Germany, 1985.
- [26] A. A. Markov. Insolubility of the problem of homeomorphy. *Proc. Int. Congr. Math.*, pages 14–21, 1958.

- [27] D. Morozov. Persistence algorithm takes cubic time in worst case. In *BioGeometry News*. ACM, Dept. Comput. Sci., Duke Univ., Durham, North Carolina, 2005.
- [28] J. R. Munkres. *Topology*. Prentice Hall, Englewood Cliffs, New Jersey, second edition, 2000.
- [29] J. P. Porto and Y. K. S. Furuya. On special generic maps from a closed manifold into the plane. *Topology Appl.*, 35:41–52, 1990.
- [30] D. P. Sullivan. On the hauptvermutung for manifolds. *Bull. Amer. Math. Soc.*, 73:598–600, 1967.
- [31] R. Thom. *Structural Stability and Morphogenesis*. Westview Press, 1988.
- [32] H. Whitney. On singularities of mappings of euclidean spaces. I. mappings of the plane into the plane hassler whitney. *Annals of Math.*, 62(3):374–410, 1955.
- [33] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33:249–274, 2005.

Biography

Amit Kanubhai Patel was born in Chicago, Illinois on May 24, 1979. He received a B.S. and an M.S. in Computer Science from the University of Illinois at Urbana-Champaign in 2001 and 2003. While at Duke University, he coauthored [1, 12, 14, 15].