Topology for Trees and Curves: Theory and Applications

David Beers

Pembroke College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2023

Contents

List of Figures

Chapter 1

Introduction

Problems in data science are often driven by a need to describe the shape of a set. Topological data analysis (TDA) is a field of research primarily concerned with using ideas from topology to extract computable and interpretable descriptors of the geometric structure of data [17, 82, 91]. While methods from TDA have been extremely useful in the analysis of high dimensional data [70], TDA has also seen applications in the study of low dimensional shapes with real-world geometry [3, 94]. In this thesis, we examine several settings wherein we can attain a conceptual or technical advantage by restricting our attention to shapes of low dimension. The discussion presented hereafter can broadly be separated into three topics.

The first of these topics concerns the study of neurons, cells in the nervous system with the structure of rooted trees. Here, we study a descriptor of neuronal data called the topological morphology descriptor (TMD), which can be viewed as a collection of intervals representing the branches of a neuron. The TMD was inspired by *persistent homology*, a popular descriptor in TDA, which takes a topological space equipped with a real valued function as input and returns a multiset of intervals, called a *barcode*. The TMD can be viewed as a conceptual simplification of persistent homology, specialised to rooted trees. We shall use the TMD to study diseased cells, and will investigate the properties of the TMD and its relationship to persistent homology, which has also been used for neuronal data.

The second topic concerns the loss information resulting from applying persistent homology to data. Precisely, we study the space of continuous functions on a topological space whose persistent homology is a given barcode. Describing this space in general is challenging, so, motivated by the previous section, we begin by studying the case where the underlying space is a geometric tree. We then move our attention to manifolds of dimension one, where we can use an alternate methodology to handle the problem in this simpler setting. Finally, we will show that this second approach

generalises to all manifolds, provided that we restrict our space of functions to the well-behaved class of Morse functions.

Lastly, we study a third descriptor called the Euler characteristic transform (ECT), which turns compact subsets in Euclidean space into integer valued functions. Unlike the persistence map, the ECT is an injective transformation, provided we restrict our attention to reasonably well-behaved subsets. We are interested in approximating the ECT of an underlying shape from a noisy sample with increasing accuracy as the sample size increases. Proving that this can be done in general is challenging, and a natural starting point is the case where the underlying shape is one dimensional. We address this case here.

This thesis is organised as follows.

1.1 Chapter 2: Background

In this chapter we introduce concepts that are necessary to understand the remainder of the text. Throughout, we will assume the reader is familiar with algebraic topology, and in particular homology. In Section 2.1.3 (and Section 4.3, for which this earlier section is a prerequisite), we will also assume that the reader has some familiarity with smooth manifolds.

We begin with an introduction to persistent homology, explaining how it assigns barcodes to real valued functions on topological spaces, in Section 2.1. While some of the material presented here is standard, such as the definition of persistence diagrams and the statement of the stability theorem presented in Section 2.1.2, some other content in Section 2.1 is new. For instance, in Section 2.1.1 we explicitly show how persistent homology is related to the local minima of a function on a topological space. In Section 2.1.3 we give a definition for a Morse function on a manifold with boundary (which agrees with the standard definition on a manifold without boundary) and prove that the critical values of Morse functions on manifolds with boundary are related to persistent homology.

In Section 2.2 we begin with some basic definitions regarding trees and a basic lemma regarding the persistent homology of functions on trees. Then, in Section 2.2.1 we define an alternative to persistent homology for trees called the topological morphology descriptor (TMD) first introduced in [51], which will play a major role in Chapter 3. We then discuss persistence images [1], a popular method for mapping barcodes into a vector space, in Section 2.2.2. This type of vectorisation is most typically used in the context of the TMD. In Section 2.2.3 we take a closer look

at persistent homology for functions on discrete trees, in particular discussing how barcodes are computed in the work of Li et al. [59]. This discussion will be put to use in Chapter 3 when we compare the methodologies of [51] and [59].

Section 2.3 is concerned with two similar constructions in TDA, *merge trees* (Section 2.3.1) [68] and what we call a *cellular merge trees* (Section 2.3.2) [30]. Both of these constructions, while slightly different at a technical level, are tree-like spaces designed to record information about the connected components of the sublevel sets of a function on topological space. In Section 2.3.3 we discuss the interleaving distance on merge trees.

Sections 2.4-2.6 are shorter than the previous sections and consist of the necessary background for Chapter 5. In Section 2.4 we recall the definition of a one dimensional CW complex and introduce several related notations that we will use later. In Section 2.5 we introduce the Euler characteristic transform (ECT) [91], a rich descriptor of shape data, and its variants, the smooth Euler characteristic transform (SECT) and the persistent homology transform (PHT). In the final section of this chapter, Section 2.6 we describe Gaussian processes, a statistical model which can be used to smoothly regress noisy data.

1.1.1 Contributions

Section 2.1.1 is based on work coauthored by David Beers and Jacob Leygonie [9, 58], as is the entirety of Section 2.3, the sole lemma in Section 2.2, and the introduction to geometric trees in the same section. The remainder of Section 2.2 is based on work coauthored by David Beers, Despoina Goniotaki, Heather A. Harrington, and Alain Goriely [7, 8]. Sections 2.4-2.6 are based on work coauthored by David Beers and Lewis Marsh [62]. David Beers is the sole author of all other portions of this chapter.

1.2 Chapter 3: Topology for Neuronal Trees

In this chapter we study the TMD from [51], which was introduced in Section 2.2.1. The TMD is an operation takes as input a rooted tree equipped with a real valued function and finds a canonical decomposition of the input tree into branches, which can be viewed as a barcode. The motivation of this construction is that the rooted tree may represent the geometry of a neuron, with the function on the tree describing the distance of various points to the cell body. Under this interpretation, the branch decomposition TMD then relates the distance from the soma that branches initiate

and terminate in a neuron, and is of practical interest. In practice, barcodes obtained by the TMD are often vectorised as persistence images (see Section 2.2.2) [1, 51, 53].

In this work we primarily study the TMD when the function assigned to neurons is the path length to the soma. We show that in this setting, computing the TMD can be recast as a computation in persistent homology. Further, we define *topological morphology functions*, a function which we show can be recovered from the TMD in the path distance. We then draw analogies between topological morphology functions and *Sholl functions* [81] which were shown in the work of Li et al. to be recoverable from certain applications of persistent homology to neuronal data [59]. We further show that both topological morphology functions and Sholl functions can be approximately recovered from relevant persistence images.

Lastly, in Section 3.4.3 we study the stability of persistence images arising from the TMD. There are already results showing that the barcode arising from the TMD is stable against reasonable perturbations of input data. Here, we show that the TMD is stable in the *1-Wasserstein distance* against perturbations to the neuronal data corresponding roughly to:

1. slightly adjusting the position of a neuron's branches; and
2. attaching a new short branch to a neuron.

The significance of our stability results is twofold. First, as a consequence of Adams et al. [1], persistence images are also stable against the same perturbations. Second, we provide Lipschitz bounds for both kinds of stability listed above.

1.2.1 Contributions

Section 3.4.3 of this chapter is based on a paper coauthored with Heather A. Harrington and Alain Goriely [8]. The remainder of this chapter is based on a paper coauthored with Despoina Goniotaki and Diane Hanger, along with Heather A. Harrington and Alain Goriely [7]. David Beers is responsible for all theoretical results and data analysis in this section.

1.3 Chapter 4: Inverse problems for persistent homology

A natural question following from the previous chapter is how descriptive the TMD and persistent homology are. Given a specified tree, one way of reformulating this

question is to ask how many functions result in the same prespecified barcode B. This question has been studied for the TMD in [52]. For persistent homology, this question is a special case of the following more general question: for a topological space X and a space of real valued functions $\mathcal{F}$ on X, what is the subset $\mathrm{PH}^{-1}(B)$ of $\mathcal{F}$ whose persistent homology gives rise to a barcode B?

Returning to trees, we begin in Section 4.1 by studying the case where X is a geometric tree, $\mathcal{F}$ is the space of continuous functions on X, and B is generic. Via an analysis which tracks local minima and other critical sets that we call *saddles*, we are able to relate the homotopy type of $\mathrm{PH}^{-1}(B)$ to specialised configuration spaces which we define in Section 4.1.3. Then in Section 4.1.4 we use this relation to enumerate the number of path components in $\mathrm{PH}^{-1}(B)$ and, for simple barcodes B, compute its homotopy type.

In Section 4.2 we then turn our attention to the case that X is instead the circle, although our methods here would work equally well for the interval. In this short section we show that an alternate approach can be used to study the fiber $\mathrm{PH}^{-1}(B)$. Instead of relating $\mathrm{PH}^{-1}(B)$ to a configuration space on the circle, we study the fiber via the action of automorphisms on it. This lets us compute the homeomorphism type of the path components of $\mathrm{PH}^{-1}(B)$, which in particular allows us to show that the path components of $\mathrm{PH}^{-1}(B)$ are always homotopic to a circle, as long as B is not a barcode arising from a constant function.

We conclude the chapter with Section 4.3 by studying how the intuition of the last section can generalise to higher dimensional manifolds, with or without boundary, provided we work with a nicer space of functions $\mathcal{F}$. Here, we let $\mathcal{F}$ be the space of Morse functions on X which have prescribed constant values on boundary components. In this setting we show that if f is a function with barcode B, then the path component of $\mathrm{PH}^{-1}(B)$ containing f is exactly the orbit of f under the action of diffeomorphisms homotopic to the identity. Since the topology of orbit spaces of Morse functions has already been studied, in particular by Maksymenko [61], this allows us to deduce topological properties of the fiber, especially when X is a surface.

1.3.1 Contributions

This chapter is based on two papers coauthored with Jacob Leygonie [9, 58].

1.4 Chapter 5: The Euler characteristic transform for shapes of dimension one

In Chapter 5 we turn our attention to a transformation called the Euler characteristic transform (ECT). The ECT of a shape $X \subseteq \mathbb{R}^d$ is defined as the function which, for each linear half-space, returns the Euler characteristic of the intersection of X with that half-space. In contrast with persistent homology, which loses information, it has been shown that the ECT is an injective transform for broad classes of inputs, collections of shapes called *constructible sets* [31, 44, 91]. In particular, the ECT is injective on compact simplicial complexes, or more generally compact semialgebraic sets. The ECT is therefore useful for the analysis of shape data as it maps subsets of Euclidean space into a vector space of functions.

How stable is the ECT? Suppose that X is a piece-wise linear curve interpolating n points in $\mathbb{R}^d$, which we assume to be sampled from some underlying curve. If we perturb these points each by a distance of no more than ε, then a result of Skraba and Turner [83] implies that, for many sensible metrics on ECTs, we have a Lipschitz bound on the difference between our previous and perturbed ECT. However, this Lipschitz bound grows in an unbounded fashion when n increases. Since the ECT is used to describe shapes, with larger data (i.e. when n is larger), we need a better estimate of the ECT of a given curve. In this chapter we show that an estimator with this desirable property can be attained by smooth regression instead of interpolation of point cloud data, for shapes of dimension one.

To begin, in Section 5.1 we define a metric on piece-wise C^2 embeddings of the interval (and more generally a finite one-dimensional CW complexes) in Euclidean space, which depends not on any underlying point cloud, but instead on the proximity of embeddings and distortions to arc lengths. In Section 5.2 we use our stability result to show that the theory of Gaussian processes can be used to regress noisy data in such a way that the ECT of the regression probabilistically approaches the ECT of the underlying shape. We conclude with Section 5.3, where we demonstrate our results on generated data.

1.4.1 Contributions

This chapter is based on a paper coauthored with Lewis Marsh [62].

Chapter 2

Background

In this chapter we introduce concepts that that are either standard or folklore in the literature, providing proofs to a few facts which I could not find proven elsewhere. We assume throughout that the reader has some familiarity with algebraic topology. For Section 2.1.3 (and Section 4.3 later on) we will also assume that the reader is familiar with smooth manifolds, although this topic will not be necessary elsewhere. We begin by introducing persistent homology [17, 37, 43, 96], which plays a central role in the next two chapters, and TDA in general.

2.1 Persistent Homology

Consider a sequence of topological spaces $\{X_t\}_{t\in\mathbb{R}}$ nested by inclusion in the following way:

$$X_t \subseteq X_s \qquad \text{whenever} \qquad t \leq s.$$

Persistent homology is a mathematical operation designed to provide an answer to a following question.

Question. *Can we in some way track how the topology of X_t evolves as t varies?*

A good first approach might be to take an algebraic descriptor of the topology of each space X_t, such as the i^{th} homology H_i, pointwise as t varies. Of course, each inclusion $X_t \subseteq X_s$ can be viewed as an inclusion map, and homology is a functor, so moreover we can get a sequence of groups $\{H_i(X_t)\}_{t\in\mathbb{R}}$ connected by commuting homomorphisms

$$H_i(X_t) \to H_i(X_s) \qquad \text{whenever} \qquad t \leq s.$$

When we take homology over a field $\mathbb{F}$ these groups connected by homomorphisms are in fact vector spaces connected by linear maps. So we get an example of the following construction:

Definition 1. *Let R be a totally ordered set. An* R-indexed persistence module *is a sequence of vector spaces $\mathbb{V} = \{\mathbb{V}(t)\}_{t\in R}$ with commuting linear maps*

$$\mathbb{V}(t) \to \mathbb{V}(s) \qquad \textit{whenever} \qquad t \leq s.$$

We say $\mathbb{V}$ is pointwise finite dimensional *or* pfd *if each $\mathbb{V}(t)$ is finite dimensional. We say $\mathbb{V}$ is* tame *if there is a number N such that $\dim \mathbb{V}(t) \leq N$ for all t.*

An especially simple example of an R-indexed persistence module can be constructed in the following way. Let $I \subseteq R$ be an interval, i.e. a subset satisfying that if $x \leq y \leq z$ and $x, z \in I$, then $y \in I$. We define the *interval module* of I, denoted M_I, to be the persistence module whose pointwise entries are given by

$$(M_I)(t) = \begin{cases} \mathbb{F} & t \in I \\ 0 & \text{otherwise,} \end{cases}$$

and whose associated linear maps are the identity where possible, and the zero map everywhere else.

The following result from [29] assures that any reasonably well behaved persistence module $\mathbb{V}$ has a decomposition into interval modules

Theorem 2 (Crawley-Boevey)**.** *Any pfd persistence module is a direct sum of interval modules. Moreover, the multiset of intervals in this decomposition is unique.*

Returning to our sequence $\{H_i(X_t; \mathbb{F})\}_{t\in\mathbb{R}}$, the above theorem says that if this sequence is pfd, it can be described by a multiset of intervals in the real line. We give a name to this structure, calling a multiset of intervals in the real line a *barcode*. We will occasionally refer to the set of barcodes as Bar. We say that a barcode is *finite* if it is a finite multiset of intervals. The interpretation behind a barcode is that each interval records a different homological feature in the sequence $\{X_t\}_{t\in\mathbb{R}}$. The shape of an interval I then determines the range of values of t for which the feature it represents are present. For this reason, the left and right endpoints of an interval are often referred to as its *birth* and *death*, while its length, the difference between the right and left endpoint values, is often called its *persistence*. The barcode obtained from applying H_i to the sequence $\{X_t\}_{t\in\mathbb{R}}$ is called the *i dimensional persistent homology* of $\{X_t\}_{t\in\mathbb{R}}$.

Since Theorem 2 holds whenever the index set is totally ordered, persistent homology is defined for sequences of spaces $\{X_t\}_{t\in R}$ for any index set R. We refer to a sequence of spaces such as $\{X_t\}_{t\in R}$ that is nexted by inclusion as a *filtration*.

Historically, the idea of describing the evolving homology of a sequence of spaces by a barcode can be traced back to a seminal paper of Zomorodian and Carlson [96], which proved that Theorem 2 holds for tame persistence modules indexed by a finite totally ordered set.

The choice of field $\mathbb{F}$ is relevant to the barcode returned by persistent homology[1]. However, zero dimensional persistent homology is not dependent on the choice of $\mathbb{F}$, at least provided a filtration furnishes a finite barcode:

Proposition 3. *Let $\mathbb{F}_1$ and $\mathbb{F}_2$ be fields and $\{X_t\}_{t\in R}$ be a filtration. If one of the persistence modules $\mathbb{V}_1 = \{H_0(X_t;\mathbb{F}_1)\}_{t\in R}$ and $\mathbb{V}_2 = \{H_0(X_t;\mathbb{F}_2)\}_{t\in R}$ is pfd then both are. Moreover if the barcode corresponding to either persistence module is finite, these persistence modules have the same barcode.*

Proof. Suppose without loss of generality that $\mathbb{V}_1$ is pfd. The universal coefficients theorem for homology gives us natural isomorphisms

$$H_0(X;\mathbb{Z}) \otimes \mathbb{F}_i \cong H_0(X;\mathbb{F}_i)$$

for $i = 1, 2$ and arbitrary spaces X. Hence the two persistence modules in the proposition have the same pointwise dimensions. Assuming either persistence module has a finite barcode, we can take a finite filtration $\{X_t\}_{t\in R'}$ for R' consisting of values at and between barcode endpoints. The second statement follows from [71, Corollary 1], which says that the H_0 barcode for R' is independent of the choice of field. □

Remark 1. *We conjecture that the requirement that the barcode be finite in the second statement is superfluous.*

Even for higher order homology, homology will always be assumed to be taken over a field for the remainder of the text, and so we shall often write $H_i(X;\mathbb{F})$ as $H_i(X)$, and will remark upon the choice of field when it is relevant. We now describe a few common examples of how a sequence $\{X_t\}_{t\in\mathbb{R}}$ can arise, the first of which will play a key role in Chapters 3 and 4.

Example. *Let X be a topological space and $f : X \to \mathbb{R}$ be a function. The sublevel set filtration of f is the filtration $\{X_t\}_{t\in\mathbb{R}}$ defined by $X_t = f^{-1}(-\infty, t]$. We say that f is pfd if the persistence module obtained by applying homology in any dimension is*

[1]Most often when computing persistent homology on a computer, one chooses $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$ for the sake of computational speed.

pfd. We define $\mathrm{PH}_i(f)$ *to be the barcode derived from applying* i^{th} *homology to the sublevel set filtration of* f*, should it exist. In situations where we want to emphasise the role that* X *plays, we may also refer to this barcode as* $\mathrm{PH}_i(X, f)$*. For pfd* f*, we then define* $\mathrm{PH}(f)$ *to be the ordered list of barcodes*

$$[\mathrm{PH}_0(f), \mathrm{PH}_1(f), \mathrm{PH}_2(f), \ldots].$$

Again, at times we will refer to $\mathrm{PH}(f)$ *as* $\mathrm{PH}(X, f)$ *when we wish for the role of* X *to be emphasised.*

We can also define the superlevel set filtration *of* f *to be the filtration* $\{X^t\}_{t\in\mathbb{R}}$ *by defining* $X^t = f^{-1}[t, \infty)$*, although this time the inclusions are in the other direction, i.e. we have*

$$X^t \supseteq X^s \qquad \textit{whenever} \qquad t \leq s.$$

Nevertheless taking i^{th} *homology furnishes a persistence module, and we therefore get a barcode via persistent homology provided the persistence module is pfd.*

Sometimes it is useful to consider more complicated totally ordered sets than $\mathbb{R}$. For example, a popular technique in persistent homology is, for a space X with function f, to consider the sequence of pairs

$$\to (X_t, \emptyset) \to \ldots \to (X_s, \emptyset) \to \ldots \to (X, \emptyset) \to \ldots \to (X, X^s) \to \ldots \to (X, X^t) \to .$$

This sequence is indexed by the totally ordered set

$$\mathbb{R} \sqcup \{\infty\} \sqcup \mathbb{R},$$

where the order in the first copy of $\mathbb{R}$ is the standard order, the order in the second copy of $\mathbb{R}$ is the reverse order and the order between sets in the above disjoint union is determined by the order they appear. The barcode obtained by this procedure is called the *extended persistent homology* of f [27]. As a result of our unusual indexing set, this barcode can contain some entries which are not technically intervals in the real line, such as

$$[3, 3) \text{ and } [5, -1).$$

Extended persistence barcodes and other barcodes with similar entries will play a role later in the text. For this reason, we will let a barcode be more generally defined by a multiset of ordered pairs of elements of the extended real numbers

$$\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty, -\infty\},$$

such that each element of every pair is assigned a label either *open* or *closed*. In this more general setting, we will still refer to elements of barcodes as intervals, and notions such as left and right interval endpoints, birth, death, and persistence extend naturally.

Example. *If we have a finite point cloud $S \subseteq \mathbb{R}^d$, we can define spaces S_ε via unions of balls of radius ε around points in S, i.e.*

$$S_\varepsilon := \{x \in \mathbb{R}^d : \text{there exists } s \in S \text{ such that } d(x,s) \leq \varepsilon\}.$$

This defines a filtration of spaces called the Čech filtration *of S. The motivation for studying this filtration is that, if S is drawn from a subset $X \subseteq \mathbb{R}^d$ and is sufficiently dense, then in many cases there should be a range of values ε where the topology of S_ε is identical to that of X up to homotopy. Therefore, the intervals in the barcode of the Čech filtration in the appropriate range should reflect the topology of the underlying set X.*

We remark that the sets S_ε can also be thought of as sublevel sets of the distance function to S:

$$d(x,S) := \min_{s \in S} d(x,s).$$

Example. *Another filtration which is attained from a finite point cloud S is what is called the* Vietoris-Rips *filtration V_ε. Explicitly, V_ε is the simplicial complex consisting of simplices $\sigma = \{p_0, \ldots, p_n\} \subseteq S$ such that if p_i and p_j are both in σ, then $d(p_i, p_j) \leq 2\varepsilon$.*

Again the sets V_ε can be viewed as arising from the sublevel sets of a function, this time on the complete simplicial complex on S. If we define a function f on the simplices of this complex by

$$f(\sigma) = \frac{1}{2} \max_{p_i, p_j \in \sigma} d(p_i, p_j),$$

then the spaces $\{V_\varepsilon\}_{\varepsilon \in \mathbb{R}}$ are exactly the sublevel set filtration of f.

2.1.1 Persistent homology and local minima

In this section we show some relations between the zero dimensional barcode $\mathrm{PH}_0(f)$ and the number of local minima of a function f. The single lemma proven here is largely folklore in the theory of persistent homology, but we provide a proof for completeness. Results presented later in this text hinge upon the following precise definition of a local minimum.

Definition 4. *Given a topological space X and a map $f : X \to \mathbb{R}$, a subset $M \subseteq X$ is a* ***local minimum of*** *f if M is connected, f is constant on M, and any connected M' containing M also contains a point x satisfying $f(x) > f(M)$.*

Note that if f is continuous then its local minima are each closed. Hence, if X is also compact, then the local minima of f are compact.

Lemma 5. *Let X be any topological space, B be a finite barcode, and $f : X \to \mathbb{R}$ with $\mathrm{PH}_0(f) = B$. Then f has finitely many local minima. Moreover, if every connected subset of each sublevel set $f^{-1}(-\infty, t]$ is path connected, then f has exactly as many local minima as intervals in B with closed left endpoints.*

Proof. Let M be a local minimum of f, and $m = f(M)$. Let $\mathbb{V}_0(f)$ be the persistence module obtained by applying H_0 over a field to the sublevel set filtration of f. Assume, seeking contradiction, that there are k local minima M_i with value m, but that there are at most $k-1$ intervals with closed left endpoint m. In this case we have

$$\dim \operatorname{coker} \left(\mathbb{V}_0(f)(m-\varepsilon) \to \mathbb{V}_0(f)(m)\right) < k$$

for sufficiently small ε. Pick $x_i \in M_i$ for each i. The set $\{[x_i]\}$ is linearly independent in $\mathbb{V}_0(f)(m) = H_0(f^{-1}(-\infty, m])$ since the x_i are in different path components of $f^{-1}(-\infty, m]$. Otherwise, we could construct a path from M_i to M_j for some i and j with f values less than or equal to m, contradicting the fact that M_i is a local minimum. Therefore, $\operatorname{span}\{[x_i]\}$ has dimension k and so intersects the image of the map $\mathbb{V}_0(f)(m-\varepsilon) \to \mathbb{V}_0(f)(m)$ for sufficiently small ε. Pick an element $\sum_i c_i[x_i]$ in this image.

We can pull this element back to

$$\sum_j d_j[y_j] \in \mathbb{V}_0(f)(m-\varepsilon),$$

so each y_j must be in the same path component of $f^{-1}(-\infty, m]$ as some x_i. This yields a path γ from $y_j \notin M_i$ to $x_i \in M_i$ on which f takes values less than or equal to m, contradicting that M_i is a local minimum. Therefore there are at least as many intervals with closed left endpoint m as local minima with value M. This proves the first statement of the lemma.

For the second statement, suppose every connected subset of $f^{-1}(-\infty, t]$ is path connected for each t, and that there are k local minima M_i with value m and greater than k intervals with closed left endpoint m. This implies that when ε is small enough,

$$\dim \operatorname{coker} \left(\mathbb{V}_0(f)(m-\varepsilon) \to \mathbb{V}_0(f)(m)\right) > k.$$

Let

$$V := \operatorname{im}\left(\mathbb{V}_0(f)(m-\varepsilon) \to \mathbb{V}_0(f)(m)\right) \qquad \text{for } \varepsilon \text{ suff. small,}$$
$$W := \operatorname{span}\{[x] | x \in M_i \text{ for some } i\} \subseteq \mathbb{V}_0(f)(m).$$

Since each local minimum is connected, each local minimum is path connected. Therefore, $\dim W = k$. Hence, there must be some element of $H_0(f^{-1}(-\infty, m])$ which is not in $V + W$. This implies there is an element $y \in f^{-1}(-\infty, m]$ such that $[y] \in H_0(f^{-1}(-\infty, m])$ that is not in $V + W$. Let $\widetilde{M}$ be the connected component containing y in $f^{-1}(-\infty, m]$. Thus, $\widetilde{M}$ is also path connected. If there is a path in $f^{-1}(-\infty, m]$ from y to an element of $f^{-1}(-\infty, m)$, then y is in V, a contradiction. So f must be constant on $\widetilde{M}$. Any connected $\widetilde{M'}$ strictly containing $\widetilde{M}$ cannot be contained in $f^{-1}(-\infty, m]$, as this would imply it is equal to $\widetilde{M}$ itself or disconnected. Therefore it must intersect $f^{-1}(m, \infty)$. So $\widetilde{M}$ is in fact a local minimum of f, and $y \in W$, a contradiction. This proves the second statement.

□

2.1.2 Persistence diagrams, metrics, and stability

Barcodes can be visualised as multisets of intervals. They can also be visualised as multisets of points in the real plane by sending each interval with left endpoint x and right endpoint y to the point (x, y). By taking a disjoint union of this multiset with each point (x, x) with infinite multiplicity, we get an example of the following construction.

Definition 6. *We refer to the multiset consisting of each point (x, x) for each $x \in \overline{\mathbb{R}}$ with infinite multiplicity as* diag. *A* persistence diagram *is a disjoint union of any multiset of points in $\overline{\mathbb{R}}^2$ with* diag.

We will explain the reason for introducing the multiset diag shortly, but first note that the x and y values of a point in a persistence diagram represent the birth and death of the interval it represents, respectively. Moreover, the height of each point (x, y) above the line $y = x$ is $y - x$, the persistence of the interval represented by (x, y). The diagonal points (x, x) in a persistence diagram are typically either not depicted visually or are represented by a diagonal line. In Figure 2.1, we show a barcode and its resulting persistence diagram.

We still have not addressed why every persistence diagram includes the multiset diag, which at this point may seem like a needless complication of a simple concept.

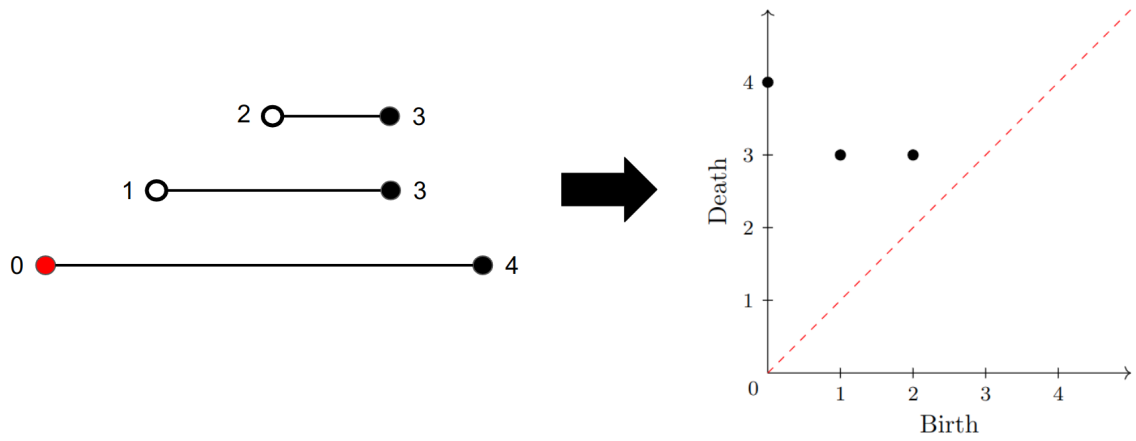

Figure 2.1: A barcode and its resulting persistence diagram.

The reason for these diagonal points is that they allow us to metrise the space of persistence diagrams in ways such that points with small persistence have little impact on distances between diagrams. We define some of the most popular metrics for persistence diagrams below.

Definition 7. *A map ϕ between persistence diagrams D_1 and D_2 is called a matching if it is bijective. The bottleneck distance between D_1 and D_2 is defined to be*

$$d_B(D_1, D_2) := \inf_{\phi} \sup_{x \in D_1} \|x - \phi(x)\|_\infty,$$

where the infimum is taken over all matchings ϕ between D_1 and D_2. The q-Wasserstein distance is defined similarly, as

$$d_q(D_1, D_2) := \inf_{\phi} \left(\sum_{x \in D_1} \|x - \phi(x)\|_\infty^q \right)^{1/q},$$

again with the infimum being over matchings ϕ between D_1 and D_2.

If for a matching ϕ we have

$$\sum_{x \in D_1} \|x - \phi(x)\|_\infty \leq \varepsilon,$$

we will say that ϕ is an ε-matching[2].

For this definition, we take $|\infty - \infty| = 0$ and $|\infty - t| = \infty$ for any $t \in \overline{\mathbb{R}} - \{\infty\}$, and similarly define the absolute value for $-\infty$.

[2]It is more common in the literature for an ε-matching to refer to a matching such that $\sup_{x \in D_1} \|x - \phi(x)\|_\infty \leq \varepsilon$. We use an alternate definition in this work as the 1-Wasserstein distance will be of primary interest when we discuss matchings in Chapter 3.

The bottleneck distance is shown to be an extended metric on the space of persistence diagrams with finitely many off-diagonal points in [37, p. 219]. The q-Wasserstein distances are shown to be extended metrics on the space of persistence diagrams with finitely many off-diagonal points similarly. In general, any time we have a pair of barcodes, we can compare them by computing bottleneck or Wasserstein distances on their associated persistence diagrams. The bottleneck and Wasserstein distances do not induce true extended metrics on barcodes, even finite barcodes, since pairs of barcodes that differ only in their endpoints have the same persistence diagram.

It turns out that the persistent homology of sublevel set filtrations is very stable against the bottleneck distance, a fact that was originally proven in [25]. We state a version of this result that is as general as is both necessary and convenient for this text.

Theorem 8. *Let X be a topological space and $f, g : X \to \mathbb{R}$ be two (potentially discontinuous) functions such that $\mathrm{PH}_i(f)$ and $\mathrm{PH}_i(g)$ exist and are finite, for some i. Then*

$$d_B(\mathrm{PH}_i(f), \mathrm{PH}_i(g)) \leq \|f - g\|_\infty$$

This can be proven, for example, by combining Theorem 4.16 and Theorem 5.1 from [15].

2.1.3 Morse functions

Decades before persistent homology was developed, a special class of functions on smooth manifolds was identified whose sublevel sets are very topologically well-behaved. These functions are called *Morse functions*.

Definition 9. *Let M be a smooth manifold with or without boundary. A smooth function $f : M \to \mathbb{R}$ is called* Morse *if the following hold:*

1. *All of the critical points of f lie in $M - \partial M$;*
2. *For each critical point p there exists local coordinates $x_1, \ldots, x_n$ such that 0 is identified with p and*

$$f(x_1, \ldots, x_n) = f(p) \pm x_1^2 \ldots \pm x_n^2; \tag{2.1}$$

3. *The restrictions $f_{|\partial M_j}$ to each boundary component ∂M_j are constant with prescribed value a_j.*

The number of times a minus occurs in Equation 2.1 can be shown to be well-defined and is called the index *of the critical point p.*

From this definition it follows that the critical values of f are isolated, and in particular finite if M is compact. Typically, Morse functions are defined on manifolds without boundary. The definition given here agrees with the standard definition for such manifolds. We recite some important facts about manifolds without boundary, which we will use to prove a proposition we will need later, and refer the interested reader to [66, Part I]. Whenever M is understood to be equipped with a specified Morse function f, we let M_t denote its sublevel set $f^{-1}(-\infty, t]$.

Theorem 10. *Let M be without boundary and f be Morse on M. Let $a < b$ and suppose $f^{-1}[a, b]$ is compact. If f has no critical values within the interval $(a, b]$, then M_b deformation retracts onto M_a. Hence the inclusion $M_a \to M_b$ induces isomorphisms on homology. Moreover, if there are no critical values of f in $(a, b]$ except at b, then up to homotopy M_b can be obtained from M_a by attaching a cell to M_a for each critical point with critical value b.*

Given that the sublevel sets of Morse functions are well-behaved, it is perhaps unsurprising that a precursor to persistent homology called the abstract framed Morse complex was studied in the specialised setting of Morse functions [6]. Using the above proposition we shall prove the following result, which we shall need in Chapter 4.

Proposition 11. *Let f be a Morse function on a compact manifold M, potentially with boundary. Then f is pfd. Also the multiset of bounded interval endpoint values of $\mathrm{PH}(f)$ is exactly the disjoint union of the critical values of f and each a_j for which ∂M_j is a local minimum with multiplicity $\sum_i \beta_i(\partial M_j)$. Moreover every interval in $\mathrm{PH}(f)$ has a closed left endpoint and open right endpoint.*

Proof. Since M is compact, it has finitely many critical values and boundary components. By flowing each boundary component along the (possibly negative) gradient we can find a collar neighborhood $\partial M_j \times [0, \varepsilon]$ of ∂M_j on which $f = a_j \pm \pi$, where π denotes projection onto the second coordinate. It follows that we can extend M to be a closed (non compact) manifold $\overline{M}$ without boundary by attaching an open collar $\partial M_j \times [0, \infty)$ to each boundary component ∂M_j. Moreover, on this collar we can extend f to a smooth function $\overline{f}$ such that $\overline{f} = a_j \pm \pi$, where again π denotes projection onto the second coordinate. Depending on the derivative of f on ∂M_j. Thus $\overline{f}$ is Morse on $\overline{M}$ and $\overline{f}^{-1}[a, b]$ is compact for $a < b$.

Fix a real number t and let n be the number of critical points with value t. Let ∂_{up}, ∂_{down}, and ∂_{mid} denote the sets of indices j where a_j is greater than, less than, and equal to t, respectively. We let ∂_- denote the indices j on which ∂M_j is a local minimum. By choosing ε sufficiently small, we can guarantee there are no critical values or boundary values of f in the interval $[t-\varepsilon, t)$. Let ι denote the inclusion $M_{t-\varepsilon} \to M_t$. We have the following commutative diagram of spaces

$$\begin{array}{ccc} M_{t-\varepsilon} & \xrightarrow{\quad\iota\quad} & M_t \\ \downarrow & & \downarrow \\ \overline{M}_{t-\varepsilon} & \longrightarrow & \overline{M}_t \\ \downarrow & & \downarrow \\ M_{t-\varepsilon} \sqcup \bigsqcup_{j\in(\partial_{\text{up}}\cup\partial_{\text{mid}})\cap\partial_-} \partial M_j & \longrightarrow & M_t \sqcup \bigsqcup_{j\in\partial_{\text{up}}\cap\partial_-} \partial M_j. \end{array}$$

The top two horizontal arrows are inclusion maps. The top two vertical arrows are also inclusion maps. The bottom two vertical arrows come from retracting collar neighborhoods indexed by ∂_{down} and, on the left, also ∂_{eq} onto M, which induces retractions on the path components intersecting M onto $M_{t-\varepsilon}$ or M_t. Further we retract the remaining cylinders $\partial M_j \times (-\infty, t]$ onto a space homeomorphic to the boundary component they represent. As a result the composition of the vertical maps is the obvious inclusion map and the bottom two vertical maps are homotopy equivalences. Moreover the bottom horizontal map is an inclusion.

Since $\overline{f}$ is Morse on $\overline{M}$, $\overline{M}_t$ is attained from $\overline{M}_{t-\varepsilon}$ by attaching n different cells to $\overline{M}_{t-\varepsilon}$. From the diagram above we have that the bottom horizontal map is a homotopy equivalence when restricted to path components ∂M_j for $j \in (\partial_{\text{up}} \cup \partial_{\text{mid}}) \cap \partial_-$. On the other hand, a Mayer-Vietoris argument shows that attaching any number of cells to a disjoint union of boundary components $\bigsqcup \partial M_j$ to obtain a new space Y results in the inclusion $\bigsqcup \partial M_j \to Y$ not being a homotopy equivalence. Thus it follows that up to homotopy, M_t is obtained from $M_{t-\varepsilon}$ by attaching n cells to $M_{t-\varepsilon}$ and then taking the disjoint union with ∂M_j for each j in $\partial_{\text{mid}} \cap \partial_-$. Let ι' denote the inclusion map $M_{t-\varepsilon} \to M_t - \bigsqcup_{j\in\partial_{\text{mid}}\cap\partial_-} \partial M_j$. A Mayer-Vietoris argument involving the attached cells shows

$$\dim \ker \iota'_* + \dim \operatorname{coker} \iota'_* = n,$$

where the subscript $*$ indicates that we are taking the induced map on $\bigoplus_k H_k$. Thus,

taking into account the additional inclusions of ∂M_j for j in $\partial_{\text{mid}} \cap \partial_-$, we get

$$\dim \ker \iota_* + \dim \operatorname{coker} \iota_* = n + \sum_{j \in \partial_{\text{mid}} \cap \partial_-} \sum_{i=0}^{\infty} \beta_i(\partial M_j). \tag{2.2}$$

Using this fact to bound the cokernel of inclusion maps in the filtration, we see that the sublevel set filtration of M is pfd. By viewing M as a subset of $\overline{M}$ it can be shown that space $M_{t+\varepsilon}$ retracts onto M_t for ε sufficiently small (depending on t). As a result, in Equation 2.2 the value on the left is the sum of left and right barcode endpoints with value t, and so we have proven the second statement. The last statement also follows from the fact that $M_{t+\varepsilon}$ retracts onto M_t for all t, provided ε is chosen sufficiently small. □

2.2 Trees and barcodes for trees

A tree is a finite connected acyclic graph. We call T a *rooted* tree if T contains a distinguished node r, which we refer to as the *root* of T. From the root we induce an orientation on the edges of T. We orient any edge e incident to the root r away from r. Inductively, if v is incident to an edge we have already oriented, we orient any edges incident to v and v', but not yet oriented, away from v towards v'. If, for a given v and v', in the resulting directed graph there is a directed edge e from v to v', denoted (v, v'), we say that v is the *parent* of v' and v' is a *child* of v. If there is a directed path from v to v' we say v is an *ancestor* of v' and v' is a *descendant* of v. If a vertex v in a tree has three or more incident edges, we say that v is a *branch point*. If a root r of a tree has two or more incident edges then we say that r is a branch point. Similarly, if a vertex v (not equal to r if T is rooted) has exactly one incident edge, we say that v is a *leaf* of T. If the root r in a rooted tree has no incident edges, then r is a leaf. We denote by $N(T)$, $E(T)$, and $L(T)$ the sets of vertices, edges, and leaves of T respectively. We define the *depth* of a vertex $v \in N(T)$ to be the number of nodes, minus one, along the shortest path from r to v. The *depth* of a rooted tree T is the greatest depth of any of its vertices. We say that T is *binary* if each node has at most two children. Suppose that T is rooted such that v' is a child of v, i.e. there is a directed edge $e = (v, v')$. We can associate to v' the induced subtree T' generated by v', its children, its children's children, and so on. We say that the union $T' \cup \{e\}$ is a *child branch* of v. Examples of these definitions are shown in Figure 2.2. The *least common ancestor* of nodes v and v' is the only node in T which has v and

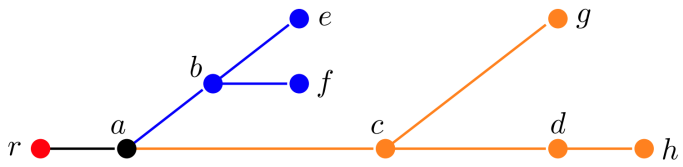

Figure 2.2: A tree with root r. Nodes a, b and c are branch points, while nodes e, f, g, and h are leaves. Nodes r and d are neither. The two child branches of a are colored in blue and orange.

v' as descendants, but no descendant that also has this property. For rooted trees, we denote by $\mathrm{LCA}(v, v')$ the least common ancestor of v and v'.

The *geometric realisation* of a tree T is a topological space given by a copy of the interval $[0, 1]$ for each edge in T with pairs of endpoints quotiented whenever they correspond to the same vertex of T. A *geometric tree* is the geometric realisation of a tree. Between any two points on a tree it is well known that there is a unique non self-intersecting path. We refer to this path as the *shortest path* between a given two points. Indeed any other path connecting a given two points contains the shortest path in its image. When T is geometric, the discussion extends to disjoint closed connected nonempty subsets $A, B \subseteq T$: there is a unique shortest path $\mathrm{ShortPath}(A, B)$ connecting them.

It follows that a subset $S \subseteq T$ of a geometric tree is path-connected if and only if it is connected. Namely, if S is not path connected, then the shortest path between a and b in S is not contained in S. Taking U_1 and U_2 to be the connected components of T minus a point in this shortest path, but not in S, we induce a disjoint open cover of S.

We define the *convex hull* of a collection $\mathcal{C}$ of closed subsets, denoted $\mathrm{Conv}(\mathcal{C})$, as the union of points on shortest paths between elements of sets in C. Clearly, convex hulls are connected. We call these sets convex hulls to suggest a familiarity with convex hulls of sets in Euclidean space. In Euclidean space, the convex hull of a subset $S \subseteq \mathbb{R}^n$, denoted $\mathrm{Conv}(S)$, is the smallest set containing S such that if $x, y \in \mathrm{Conv}(S)$, then the shortest path (i.e. line segment) between x and y is also contained in $\mathrm{Conv}(S)$. One can check that our definition of a convex hull of a geometric tree has this same property for S being the union of elements in $\mathcal{C}$.

In the next two subsections we shall restrict our attention to discrete trees, but before we do this we shall show that under very mild conditions a function on a geometric tree gives rise to a barcode in every dimension via its sublevel set filtration.

Lemma 12. *Let X be a geometric tree and $f : X \to \mathbb{R}$ be a continuous function with finitely many local minima. Then f is pfd.*

Proof. Let $t \in \mathbb{R}$. Let Ω be a path-connected component of $f^{-1}(-\infty, t]$. Consider the following alternative:

1. Either f has constant value t over Ω. Then $M := \Omega$ is a local minimum of f, because any connected strict superset M' will be included in some $f^{-1}(-\infty, t']$ with $t' > t$, but not in $f^{-1}(-\infty, t]$.

2. Or f attains a minimum $t' < t$ over one maximal connected subset $M \subseteq \Omega$. Then M is also a local minimum of f in the whole X.

In both cases, we can find a local minimum of f inside Ω, and since f has finitely many local minima, we deduce that $f^{-1}(-\infty, t]$ has finitely many path-connected components, i.e. $\dim H_0(f^{-1}(-\infty, t]) < +\infty$. This is because a subset of a geometric tree is connected if and only if it is path connected.

Finally, since X is a geometric tree, its subsets are component-wise contractible, hence $\dim H_i(f^{-1}(-\infty, t]) = 0$ for all $i > 0$ and the result follows. □

2.2.1 The topological morphology descriptor

The *topological morphology descriptor* (TMD) is a specialised kind of barcode designed to study cells in the brain and nervous system called *neurons* which are responsible for the transfer of electrical signals. These cells consist of a *soma*, or *cell body*, which contains the nucleus of the cell and is roughly spherical, and various outgrowths called *neurites*. These neurites can be classified into the *axon*, which mostly transmits signals to other neurons, and *dendrites*, which mostly receive signals from other neurons. The soma is typically small relative to the size of the dendrites of a neuron. As such, neurons can be thought of as rooted trees embedded in Euclidean space, where the root represents the location of the soma. Indeed, the popular digital `.swc` representation [86] of neurons takes this approach, viewing a neuron as a point cloud in Euclidean space, with edges which are assumed to linearly interpolate between adjacent vertices.

Let T denote a finite rooted tree with root r and consider a real valued function f on $N(T)$. The topological morphology descriptor (TMD) of T and f, denoted $\mathrm{TMD}(T, f)$, is a barcode obtained via an algorithm developed in [51]. If T is a tree representing a neuron with its root representing the soma and f records a notion of distance of the vertices of T to the root, then $\mathrm{TMD}(T, f)$ is interpreted as a

decomposition of the underlying neuron into its branches, with each interval recording how far a branch's initial and terminal points are from the root.

The process of obtaining a barcode from T and f is as follows:

1. Choose any branch point b in T;
2. Identify one of the children c of b whose child branch maximises f on the vertices that are leaves of T descendant from b;
3. Detach all child branches for children $c' \neq c$ of b. Replace T with the resulting forest;
4. If there are no branch points in the resulting forest, then T is a collection of intervals, and we are done. Otherwise return to step 1.

What will remain at the end of this procedure is a multiset of intervals with right endpoints in bijective correspondence with the leaves of T. We assign numbers to the endpoints of these intervals via f and call the resulting barcode $\mathrm{TMD}(T, f)$. Figure 2.3 shows an example of a barcode computed via the TMD. An efficient algorithm for computing the barcode $\mathrm{TMD}(T, f)$ is given in [51]. Note that in general it may be the case that the number assigned by f to the left endpoint of an interval may be greater than that assigned to the right endpoint of an interval, and so this is a barcode in the more general sense, as defined immediately before Subsection 2.1.1. We then define $\mathrm{TMDg}(T, f)$ to be the persistence diagram obtained from $\mathrm{TMD}(T, f)$, i.e. the disjoint union of a copy $(a, b) \in \mathbb{R}^2$ for every interval in $\mathrm{TMD}(T, f)$ with left endpoint a and right endpoint b with the multiset diag[3]. Notice that $\mathrm{TMDg}(T, f)$ will always contain a point $(f(r), L)$, where L is the maximum of f on the leaves of T.

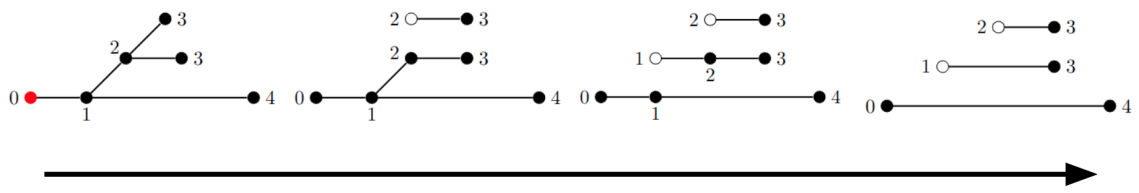

Figure 2.3: Computing the TMD from neuronal data.

[3]In [51], the authors use reversed notation, where an interval with left endpoint a and right endpoint b gets mapped to (b, a). This is likely done to suggest a connection between the TMD and extended persistence of the superlevel set filtration of f.

2.2.2 Persistence images

For data science problems, persistence diagrams are often mapped into vector spaces, which allows them to be analyzed via machine learning [14, 18, 77]. One popular technique of this style is to transform persistence diagrams into matrices called *persistence images* introduced by Adams et al [1]. Persistence images are of particular interest in the setting of the TMD as they are the standard type of vectorisation used in this context [51, 53, 56] in part because they provide interpretable visual summaries of averaged persistence diagrams.

Obtaining a persistence image from a persistence diagram is a process involving three steps. The first step is to transform a persistence diagram D to another multiset of the plane $\widetilde{D}$ by removing points at infinity and mapping each remaining point (x, y) to $(x, y - x)$. Next we fix a positive number σ and define a function f_σ to be the sum of 2D Gaussian functions of standard deviation σ centered over each point in $\widetilde{D}$, and weighted by their distance from the x-axis. The function f_σ is called a *persistence surface*. A persistence image can then be defined by fixing a grid over a relevant range of values and generating a matrix where the $(i, j)^{\text{th}}$ value is the integral of f_σ over the $(i, j)^{\text{th}}$ cell. If the dimensions of each cell in the grid are small relative to σ, a persistence image can be approximated by sampling over a grid of points. Persistence images can be viewed as vectors with a coordinate for each cell in their corresponding grid. Viewing persistence images as vectors, they are stable in their 1, 2, and ∞ norms against the 1-Wasserstein distance of persistence diagrams [1, Theorem 10][4]. We show the pipeline of obtaining the persistence image of the TMD of neuronal data in Figure 2.4.

The first transformation step $(x, y) \mapsto (x, y - x)$ can be removed from the construction of persistence images without disrupting stability guarantees, provided we instead apply a weight of $|y - x|$ to Gaussians centered at (x, y) when constructing persistence surfaces. Indeed, following the proofs of Theorems 9 and 10 from [1], it can be observed that the same reasoning holds. In fact, the bounds we are guaranteed are tighter by a factor of $\sqrt{5}$, thanks to no longer needing to transform coordinates in the proof of their Theorem 9. This can be a slightly convenient alternate version

[4]The observant reader may notice that the Lipschitz bound provided in the cited theorem is dependent on the ∞-norm of our weighting function. This may cause some alarm since we suggest weighting by distance to the x-axis, an unbounded function. However, going through the proof one observes that the bound is only dependent on the ∞-norm of the weight function over points in a fixed dataset, which is always bounded, provided the dataset is finite and each diagram has finitely many nontrivial points. Furthermore, the weighting the authors themselves use in [1, page 8] differs from the distance from the x-axis function by a constant multiple on data.

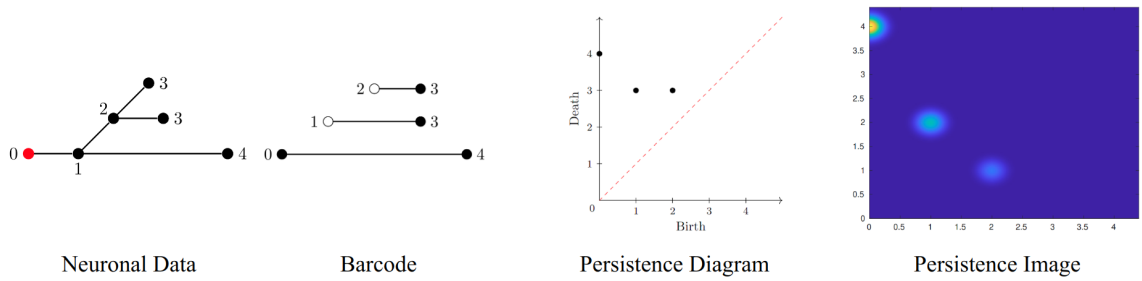

Figure 2.4: The persistence image pipeline applied to neuronal data.

of persistence images in practice, since in certain settings $\mathrm{TMDg}(T, f)$ can be a persistence diagram with points in both regions $y < x$ and $y > x$. Similarly, we can also reflect points in a persistence diagram without the transformation $(x, y) \mapsto (x, y - x)$ about the diagonal $y = x$, or with the transformation about the x-axis while still ensuring persistence images are 1-Wasserstein stable. We mention this since Kanari et al. have already used unweighted analogues of reflected persistence images without the transformation $(x, y) \mapsto (x, y - x)$ in their previous work on the TMD [51, 53]. In summary, any transformation of a tree equipped with a function that is stable for the 1-Wasserstein distance of its TMD is also stable with respect to popular metrics on its associated transformed and untransformed, reflected and unreflected persistence images.

2.2.3 Persistent homology for neuronal trees

TDA has been used in other ways to study the shape of neurons. In the work of Li et al. [59], the authors instead opt to directly apply persistent homology to produce a morphologically relevant barcode describing neuronal geometry.

Suppose we have a discrete tree T representing a neuron and a map $f : N(T) \to \mathbb{R}$. This induces a sublevel set filtration $\{T_t\}_{t\in\mathbb{R}}$ where T_t contains every vertex v of T such that $f(v) \leq t$ and every edge in T that is between vertices in T_t. Analogously we can define a superlevel set filtration $\{T^t\}_{t\in\mathbb{R}}$ of T. We depict an example sublevel set filtration in Figure 2.5. We denote the persistent homology of this sublevel set filtration with the usual notation, $\mathrm{PH}(T, f)$. Since T is a tree, every subset of T consists only of contractible components, and hence the only dimension where $\mathrm{PH}(T, f)$ is nonempty is in dimension zero. It follows from the fact that the sublevel set filtration changes at only finitely many critical values of t, say $t_1, \ldots, t_n$, that every interval in the barcode $\mathrm{PH}(T, f)$ has a closed left endpoint and an open (potentially infinite) right endpoint. In fact, since T is connected, it follows that in $\mathrm{PH}(T, f)$ there is

exactly one infinite right endpoint and that every left endpoint takes a value in the real numbers.

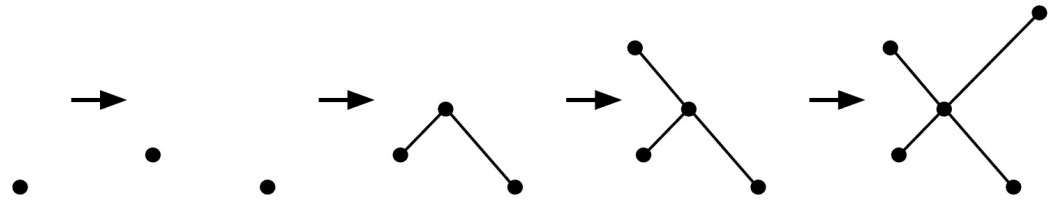

Figure 2.5: The sublevel set filtration of the height function on the nodes of the rightmost tree.

When making persistence images from barcodes, it is convenient to work with barcode consisting only of intervals with real endpoints and we can achieve this, for example, by sending the one instance of ∞ to $\max(f)$ and replacing the open endpoint with a closed endpoint. We call the resulting barcode $\mathrm{EPH}_0(T, f)$. One can check that this new barcode is precisely the zero dimensional extended persistence barcode of T, justifying the notation. We show the barcodes $\mathrm{PH}(T, f)$ and $\mathrm{EPH}_0(T, f)$ of the sublevel set filtration in Figure 2.5 in Figure 2.6.

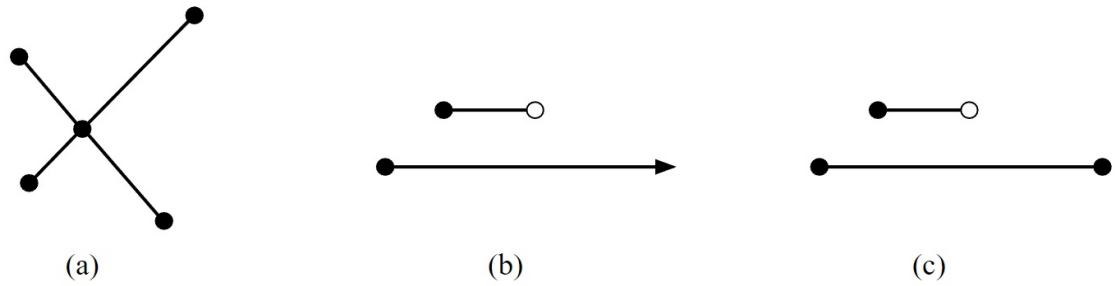

Figure 2.6: (a) The example tree T from Figure 2.5. (b) The barcode $\mathrm{PH}(T, f)$, where f is the height function on the nodes. (c) The barcode $\mathrm{EPH}_0(T, f)$.

In [59], barcodes of bounded intervals for a tree T and a function f are obtained by first computing $\mathrm{PH}(T, f)$ and then transforming the result to the barcode of bounded intervals $\mathrm{EPH}_0(T, f)$ as described above. In this setting $\mathrm{PH}(T, f)$ is computed by the following procedure under the assumption of certain genericity conditions. Since the sublevel set filtration changes finitely many times, it follows that every interval in $\mathrm{PH}(T, f)$ has a closed left endpoint and open right endpoint. It follows from Lemma 5 that there is a bijection between the intervals of $\mathrm{PH}(T, f)$ and the local minima of f. The local minima are the subtrees of T on which f takes exactly one value which is less than its value on all neighboring nodes. If f is generic then all local minima of

f are nodes. The values $f(v)$ for minima are the values of the left endpoints of their associated intervals, since they correspond to t values where their associated connected components first appear. Suppose that at time t_i in the sublevel set filtration of f, two or more connected components $\{C_l\}_{l=1}^N$ of $T_{t_{i-1}}$ merge. Additionally, suppose there exists a unique global minimum of f restricted to each C_l, denoted v_l, and one of the v_l has the lowest f value, without loss of generality $f(v_l) < f(v_1)$ for $l \geq 2$. In this generic case, then the Elder Rule [16, Theorem 4.4] determines that the right endpoint of the interval corresponding to v_l to be $f(t_l)$ for $l \geq 2$, giving rise to the interval $[f(v_l), f(t_i))$. If our assumptions hold whenever connected components merge, we can compute persistent homology by this procedure. Indeed, after applying this process to every merging of connected components in the sublevel set filtration of T, every local minimum whose corresponding right endpoint has yet to be assigned must be the global minimum of its connected component in T. Since T is connected, there must only be one such local minimum. Hence, the value of this local minimum must be paired with an infinite right endpoint, as $H_0(T)$ is a one dimensional vector space.

We can still compute $\mathrm{PH}(T, f)$ in a similar fashion when f is not generic. If $M_1, \ldots, M_k$ are the local minima of f, then we have k intervals in $\mathrm{PH}(T, f)$, each with left endpoint $f(M_k)$. This time inducting on i, suppose again that at time t_i, connected components $\{C_l\}_{l=1}^N$ of $T_{t_{i-1}}$ merge, but by induction suppose exactly one of the local minima in each C_l has not yet been assigned a right endpoint. Indeed there must be exactly one such local minimum for the base case $i = 1$ where every C_i must be a local minimum. For general i, refer to this distinguished local minimum in C_l by M_{n_l}. By reordering indices, we can have that $f(M_{n_1}), \ldots, f(M_{n_N})$ is weakly increasing. The Elder Rule [16, Theorem 4.4] then determines the right endpoint of the interval associated to M_{n_i} to be t_i for $i \geq 2$, leaving the right endpoint of M_{n_1} undecided. There was a choice in our ordering of the M_{n_i} of which local minimum was assigned to M_{n_1} if multiple of the M_{n_i} have $f(M_{n_i}) = f(M_{n_1})$. However, according to the Elder Rule, regardless of the choice we make, we will get the same barcode. Our inductive hypothesis is satisfied, as we have exactly one local minimum, namely M_{n_1}, which has a right endpoint yet to be determined in the connected component of T_{t_i} containing each C_l. As before, only one interval will remain with undecided right endpoint after this procedure, and we label this right endpoint with ∞ to produce the barcode $\mathrm{PH}(T, f)$. Replacing ∞ with $\max(f)$ we can use this algorithm to compute $\mathrm{EPH}_0(T, f)$.

We can easily define the negation operation on barcodes to be the map which swaps and takes the negative value of interval endpoints. For example, we have

$$-\big([1,\infty) \sqcup [2,5) \sqcup [3,4)\big) = (-\infty,-1] \sqcup (-5,-2] \sqcup (-4,-3]$$

Similarly, we can also define a switching map S which switches the endpoints of each interval. For example,

$$S\big([1,\infty) \sqcup [2,5) \sqcup [3,4)\big) = (\infty,1] \sqcup (5,2] \sqcup (4,3].$$

We can also get a barcode via the superlevel set filtration of f, with the caveat that the intervals $[x,y)$ in this barcode satisfy $x \geq y$. Similarly, we can make these intervals finite by replacing the $-\infty$ in the resulting barcode with $\min(f)$, and replacing the open endpoint corresponding to $-\infty$ with a closed endpoint. It is readily verified that applying persistent homology to the superlevel sets of f gives us the barcode $-S(\mathrm{PH}(T,-f))$. Applying persistent homology to the superlevel sets composed with the map $-\infty \mapsto \min(f)$, which changes the infinite endpoint to be closed, is then easily seen to be $-S(\mathrm{EPH}_0(T,-f))$.

2.3 Merge trees and cellular merge trees

A merge tree is a tree-like topological space derived from a real valued function f on topological space X. The idea behind merge trees is that they encode the t values at which connected components appear and merge in the sublevel set filtration of f. Additionally, they also encode which connected components merge when a merging occurs. We will use two distinct definitions of merge trees from the literature, which are both instances of *gauged spaces*:

Definition 13. *A **gauged space** is a topological space X equipped with a continuous map $\pi : X \to \mathbb{R}$.*

*A **morphism** between two gauged spaces (X_1,π_1) and (X_2,π_2) is a continuous map $\phi : X_1 \to X_2$ satisfying $\pi_1 = \pi_2 \circ \phi$. An **isomorphism** of gauged spaces is a morphism that is also a homeomorphism.*

2.3.1 Merge trees

A continuous function $f : X \to \mathbb{R}$ yields a merge tree as defined in [68]:

Definition 14. *For a topological space X with a continuous function f, the associated* ***merge tree*** $\mathrm{MT}(f)$ *is the quotient of the space*

$$\mathrm{epi}(f) := \{(x,t) \in X \times \mathbb{R} : t \geq f(x)\}$$

by the relation $(x,t) \sim (y,t)$ whenever x and y are in the same connected component of $f^{-1}(-\infty, t]$.

Since merge trees inherit a map π_f from the second coordinate projection map on $\mathrm{epi}(f)$, they are naturally viewed as gauged spaces. We illustrate the construction of a merge tree in Figure 2.7.

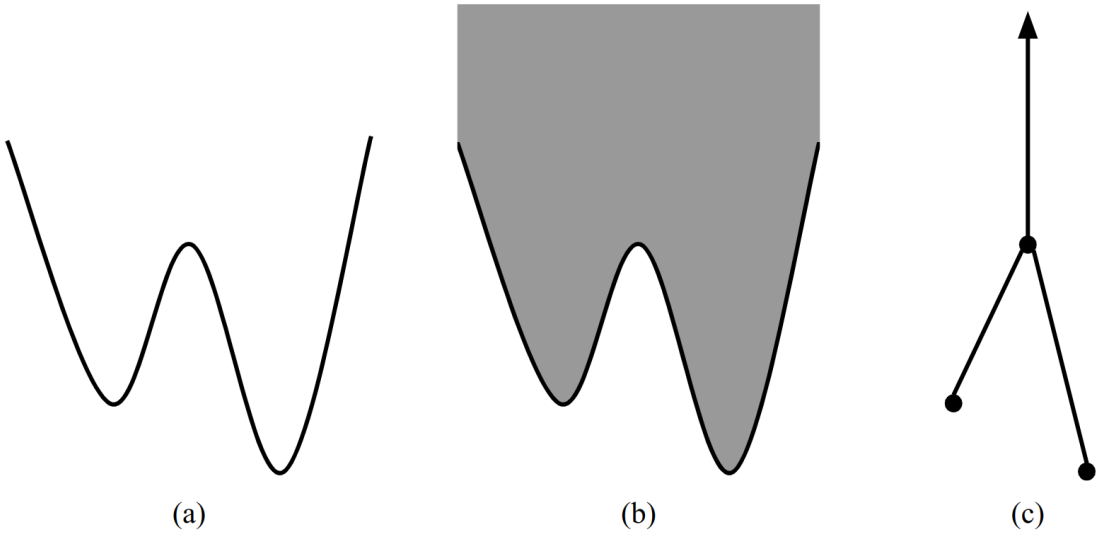

Figure 2.7: The construction of a merge tree from a function. (a) The graph of a function f on an interval. (b) Shaded is the content of $\mathrm{epi}(f)$ strictly above the graph of f. (c) By sending connected components of horizontal slices to points we obtain $\mathrm{MT}(f)$. It happens that this merge tree has a cellular structure, although in general this may not be the case.

2.3.2 Cellular merge trees

The other definition of merge trees that we will use appears for example in [30]. We provide an analogous definition here, along with the notion of a labeling from [42].

Definition 15. *A* ***cellular merge tree*** $(\mathcal{T}, \pi)$ *is the quotient space of a geometric rooted tree (T, r) and a half open interval:*

$$T \sqcup [0,1)/(r \sim 0),$$

equipped with a real-valued map π satisfying

- *π is strictly decreasing along edges in T oriented from the root r.*
- *π is strictly increasing to infinity along the half open interval $[0, 1)$.*

If we fix a cellular structure on $\mathcal{T}$, the nodes of a cellular merge tree $\mathcal{T}$ are endowed with a partial order $\preceq$ where $v \preceq v'$ whenever v is a descendent of v'.

Given a cellular merge tree $(\mathcal{T}, f)$, the map from $x \in f^{-1}(t)$ to the connected component of $f^{-1}(-\infty, t]$ containing x is a bijection from $f^{-1}(t)$ to $\pi_0\big(f^{-1}(-\infty, t]\big)$. Further, if there is a path from x to y in $\mathcal{T}$ along which f is increasing, then for $t \leq f(y)$, the connected component containing x in $f^{-1}(-\infty, t]$ is contained in the connected component containing y in $f^{-1}(-\infty, f(y)]$. From this it follows that $\mathcal{T} = \mathrm{MT}(f)$ and $f = \pi_f$. Therefore, through the continuous injection $(\mathcal{T}, f) \mapsto (\mathrm{MT}(f), \pi_f)$, cellular merge trees form a subspace of regular merge trees.

For indexing convenience, we will often work with *labelled* cellular merge trees, for which an arbitrary ordering of the leaves $l_1, \cdots, l_n$ and of the nodes $v_1, \cdots, v_m$ has been fixed. Although by default we assume that these orderings do not allow repetitions of the leaves and of the nodes, we will sometimes explicitly allow repetitions to make use of more general definitions and results from [42]. For instance, the following definition allows labels with repetitions:

Definition 16. *The **induced matrix** of a labelled cellular merge tree $(\mathcal{T}, \pi)$ is given by*

$$\mathcal{M}(\mathcal{T})_{ij} := \pi(\mathrm{LCA}(l_i, l_j)).$$

To simplify notations, we will often write $\mathrm{MT}(f)$ given a function f to designate the gauged space $(\mathrm{MT}(f), \pi_f)$, and similarly $\mathcal{T}$ to designate a cellular merge tree $(\mathcal{T}, \pi)$.

2.3.3 The interleaving distance for merge trees

In this section, we define the interleaving distance for merge trees, a construction first studied in [68]. Aside from the map π_f, merge trees also come equipped with *ε-shift* maps, for $\varepsilon \geq 0$:

$$i_f^\varepsilon : (x, t) \in \mathrm{MT}(f) \longmapsto (x, t + \varepsilon) \in \mathrm{MT}(f).$$

These maps define an action of the semigroup $\mathbb{R}_{\geq 0}$ on $\mathrm{MT}(f)$. We will often omit subscripts, writing i_f^ε as i^ε.

Definition 17. *Let f and g be two continuous functions on a topological space X. An ε-**interleaving** between $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ is a pair of continuous functions $\alpha^\varepsilon : \mathrm{MT}(f) \to \mathrm{MT}(g)$, $\beta^\varepsilon : \mathrm{MT}(g) \to \mathrm{MT}(f)$ satisfying the following equations*

$$\beta^\varepsilon \circ \alpha^\varepsilon = i^{2\varepsilon} \qquad \pi_g(\alpha^\varepsilon(x)) = \pi_f(x) + \varepsilon$$
$$\alpha^\varepsilon \circ \beta^\varepsilon = i^{2\varepsilon} \qquad \pi_f(\beta^\varepsilon(y)) = \pi_g(y) + \varepsilon.$$

*The **interleaving distance** $d_I(\mathrm{MT}(f), \mathrm{MT}(g))$ is defined as the infimum of values ε such that $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are ε-interleaved.*

We illustrate an interleaving between two merge trees in Figure 2.8. The interleaving distance is an extended pseudo metric on merge trees, i.e. it satisfies all the properties of a metric except that it sometimes takes the value ∞, which is not in $\mathbb{R}$, and may potentially take the value 0 between non-isomorphic merge trees [68, Lemma 1]. Since cellular merge trees can be viewed as merge trees, the interleaving distance also defines an extended pseudo metric between cellular merge trees. In Chapter 4 we will show that the interleaving distance is actually a true metric on isomorphism classes of cellular merge trees.

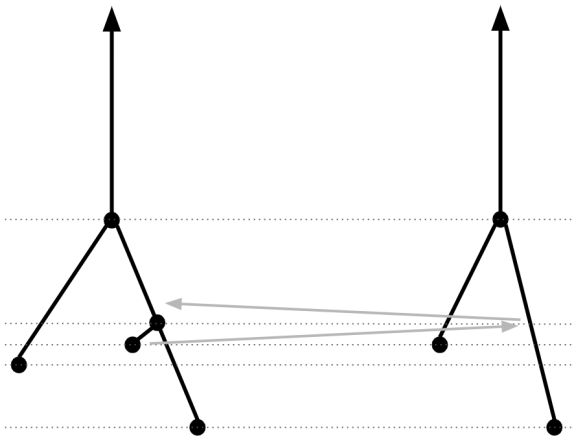

Figure 2.8: An interleaving between two merge trees.

2.4 One-Dimensional CW Complexes

The remaining content of this chapter is only necessary for the material in Chapter 5. We begin by introducing finite one-dimensional CW complexes, a class of shape slightly more general than geometric trees. While we assume the reader is already familiar with this class of topological space, we introduce some notations here that we will use later on.

A topological space Z is called a one-dimensional CW complex if it is of the form

$$Z = \left(Z_0 \sqcup \bigsqcup_{\lambda \in \Lambda} [0,1] \right) / \phi, \tag{2.3}$$

where Z_0 is a set with the discrete topology and ϕ is some map from the endpoints of the intervals in $\bigsqcup_{\lambda\in\Lambda}[0,1]$ to Z_0. We refer to the map sending the λ^{th} copy of $[0,1]$ into Z by Φ_λ (note that Φ_λ must be injective everywhere except possibly the endpoints of the interval). The space Z is said to be a *finite* one-dimensional CW complex if it can be written as in Equation (2.3) with Z_0 and Λ finite. We refer to points in Z that are in the image of $Z_0 \to Z$ as 0-cells and subsets of Z that are the image of a map Φ_λ as 1-cells. For convenience, we denote the set of 1-cells of Z by Z_1. It may be possible for a space Z to be written as in Equation (2.3) in many different ways. For instance, the circle admits the structure of a one-dimensional CW complex with n 0-cells and n 1-cells for any n. Sometimes, we need to fix a cellular decomposition of a shape. When fixing a choice of Z_0 and $\{\Phi_\lambda\}_{\lambda\in\Lambda}$, we refer to $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ as a *CW structure* on Z.

We are primarily interested in shapes with this structure that are subsets of $\mathbb{R}^d$. To this end, we say that $f : Z \to X \subseteq \mathbb{R}^d$ is a C^r map under Z^* if f is C^r on its restriction to each copy of $[0,1]$ in Equation (2.3). We denote the set of such maps f by $\mathcal{F}^r(Z^*, d)$. We denote the subset of $\mathcal{F}^r(Z^*, d)$ of maps that are also embeddings by $\mathcal{E}^r(Z^*, d)$ and the set of images of these embeddings by $\mathcal{G}^r(Z^*, d)$.

Recall that the curvature of differentiable curve is the magnitude of the derivative of its arc length reparametrisation. For $r \geq 2$, we say that $f \in \mathcal{F}^r(Z^*, d)$ has curvature bounded by M if the curvature of the map $[0,1] \to Z \to X$ is bounded by M for every copy of $[0,1]$ in Equation (2.3). By compactness of the unit interval, it follows that every $f \in \mathcal{G}^r(Z^*, d)$ has curvature bounded by some constant M whenever Z is a finite one-dimensional CW complex and $r \geq 2$. We say $X \in \mathcal{G}^r(Z^*, d)$ has curvature bounded by M under Z^* if the curvature of any map $h \in \mathcal{E}^r(Z^*, d)$ with image X has curvature bounded by M. It is straightforward to show that if X has curvature bounded by M under Z^*, then every map $h \in \mathcal{E}^r(Z^*, d)$ with image X has curvature bounded by M.

2.5 The Euler Characteristic Transform

The Euler characteristic of a topological space X with finitely generated homology is defined as the following alternating sum, which is homotopy invariant

$$\chi(X) := \sum_{k=0}^{\infty} (-1)^k \text{ rank } H_i(X; \mathbb{Z}).$$

Here, the rank of a finitely generated abelian group is the number of $\mathbb{Z}$ summands in its canonical decomposition. If X is homotopy equivalent to a finite one-dimensional CW complex containing c_k cells of dimension k for $k = 0, 1$, then the following alternate formula for $\chi(X)$ also holds:

$$\chi(X) = c_0 - c_1.$$

In particular, if every path component of X is contractible, then $\chi(X)$ is the number of path components of X.

For a subset $X \subseteq \mathbb{R}^d$, we define the *Euler characteristic transform* (ECT) of X to be the following map:

$$\begin{aligned} \text{ECT}_X : \mathbb{S}^{d-1} \times \mathbb{R} &\longrightarrow \mathbb{Z} \cup \{\infty\} \\ (v, t) &\longmapsto \chi(\{x \in X : \langle x, v \rangle \leq t\}). \end{aligned}$$

In words, the Euler characteristic transform of a shape X encodes the Euler characteristic of the intersection of X with every closed half-space with affine boundary. When $\text{ECT}_X(v, t)$ is not well defined, we set $\text{ECT}_X(v, t) = \infty$. In this context, we set $\infty + \infty = \infty$, $\infty - \infty = \infty$, and $\infty + n = \infty$ for any integer n.

The Euler characteristic transform has been of interest in the context of studying geometric data since, perhaps surprisingly, the map $X \mapsto \text{ECT}_X$ is injective when X is a simplicial complex. Hence, researchers are able to convert geometric data into functional data, which can be more easily studied on a computer and by methods of classical statistics and machine learning. Injectivity is not difficult to show in dimension one. In dimensions 2 and 3 it was first shown in [91, Theorem 3.1, Corollary 3.2]. This result was generalised using Euler calculus to any dimension independently in [31, Theorem 3.5] and [44, Theorem 5]. In fact, both of these papers show that the ECT is injective on families of what are called *constructible* sets. In particular, the ECT is injective on compact semi-algebraic sets.

Often one restricts to constructible families of subsets of $\mathbb{R}^d$ when studying the ECT, however, for our purposes these assumptions are unnecessary and further mention of constructible sets will be limited to remarks which are inessential to the overall text.

If the subset $X \subset \mathbb{R}^d$ is bounded, as is the case for any C^r-embedding of a finite one-dimensional CW complex, we can restrict the ECT to being a function on $\mathbb{S}^{d-1} \times [-a, a]$, where $a \in \mathbb{R}$ is greater than the bound of X (in the Euclidean norm). For a fixed direction $v \in \mathbb{S}^{d-1}$, we then define the Euler characteristic curve (ECC) as $\mathrm{ECC}_{X,v}(t) = \mathrm{ECT}_X(v, t)$ for each $t \in [-a, a]$. The *smooth Euler characteristic transform* (SECT), introduced in [91], is then defined as

$$\begin{aligned} \mathrm{SECT}_X : \mathbb{S}^{d-1} \times [-a, a] &\longrightarrow \mathbb{R} \\ (v, t) &\longmapsto \int_{-a}^{t} \mathrm{ECT}_X(v, x) - \overline{\mathrm{ECC}_{X,v}} \, \mathrm{d}x, \end{aligned}$$

where $\overline{\mathrm{ECC}_{X,v}}$ is the mean of $\mathrm{ECC}_{X,v}$ over $[-a, a]$. The SECT contains the same information on X as the ECT. However, the SECT lies in the Hilbert space $\mathcal{L}_2(\mathbb{S}^{d-1} \times [-a, a])$ and is, therefore, more amenable to the application of further statistical methods [94].

We endow the ECT (viewed as a function on $\mathbb{S}^{d-1} \times [-a, a]$) with the norm

$$\|\mathrm{ECT}_X\| := \sup_{v \in \mathbb{S}^{d-1}} \int_{-a}^{a} |\mathrm{ECT}_X(v, t)| \, \mathrm{d}t. \tag{2.4}$$

The norm is defined and considered analogously for the SECT.

It is also useful for us to define Euler characteristic transforms of functions f from topological spaces into $\mathbb{R}^d$. We define

$$\begin{aligned} \mathrm{ECT}_f : \mathbb{S}^{d-1} \times \mathbb{R} &\longrightarrow \mathbb{Z} \\ (v, t) &\longmapsto \chi\big(f^{-1}\{x \in \mathbb{R}^d : \langle x, v \rangle \leq t\}\big). \end{aligned}$$

Note that if f is a homeomorphism it is immediate that $\mathrm{ECT}_f = \mathrm{ECT}_{\mathrm{im}\, f}$. We prescribe norms to Euler characteristic transforms of functions as before: restricting to functions bounded by a, we let

$$\|\mathrm{ECT}_f\| := \sup_{v \in \mathbb{S}^{d-1}} \int_{-a}^{a} |\mathrm{ECT}_f(v, t)| \, \mathrm{d}t.$$

We will at times compare the ECT to a related shape transform called the *persistent homology transform* (PHT), which unsurprisingly is defined via persistent homology. If we define

$$h_v : x \in \mathbb{R}^d \mapsto \langle x, v \rangle \in \mathbb{R},$$

then we define the persistent homology transform of $X \subseteq \mathbb{R}^d$ by

$$\begin{aligned} \mathrm{PHT}_X : \mathbb{S}^{d-1} &\longrightarrow \mathrm{Bar} \\ v &\longmapsto \mathrm{PH}(h_v|_X), \end{aligned}$$

should h_v be pfd for each v. Similarly we can define the extended persistent homology transform via extended persistence.

Since the ECT of a shape can be recovered from its PHT, any injectivity result that holds for the ECT must also hold for the PHT, when the PHT is defined.

2.6 Kernel Methods and Gaussian Processes

Gaussian processes (GPs) are a model for random functions. Before defining GPs, we introduce the notion of a *kernel*:

Definition 18. *Let X be as set. A kernel on X is a symmetric function $k : X \times X \to \mathbb{R}$ such that for all $x_1, ..., x_n \in X$ and all $a_1, ..., a_n \in \mathbb{R}$ we get*

$$\sum_{i,j=1}^{n} a_i a_j k\left(x_i, x_j\right) \geq 0. \tag{2.5}$$

Given finite subsets $X' = \{x'_1, ..., x'_m\}$ and $X^* = \{x^*_1, ..., x^*_n\}$ of X, we denote by $K(X', X^*)$ the $m \times n$ matrix with i, j-entry $K(X', X^*)_{ij} = k\left(x'_i, x^*_j\right)$, which is called the Gram matrix *of k at X' and X^*.*

To each kernel k, we can associate its reproducing kernel Hilbert space *(RKHS): define a vector space of functions on X by $\mathcal{H}_0 = \operatorname{span}_{\mathbb{R}}\{k(\,\cdot\,, x) \,|\, x \in X\}$ with inner-product induced by*

$$\langle k(\,\cdot\,, x), k(\,\cdot\,, y)\rangle := k(x, y). \tag{2.6}$$

Then the RKHS of k is defined to be $\mathcal{H}_k = \overline{\mathcal{H}_0}$, the completion of $\mathcal{H}_0$.

Note that Equation (2.5) is equivalent to each Gram matrix of the form $K(X', X')$ being positive-definite.

Definition 19. *Let X be a set, $\mu : X \to \mathbb{R}$ a function and $k : X \times X \to \mathbb{R}$ be a kernel. The Gaussian process (GP) on X with mean function μ and kernel k is defined to be the random function $f : X \to \mathbb{R}$ such that for each finite set $X' = \{x'_1, ..., x'_n\} \subseteq X$ we get*

$$\begin{bmatrix} f(x'_1) \\ \vdots \\ f(x'_m) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu(x'_1) \\ \vdots \\ \mu(x'_n) \end{bmatrix}, K\left(X', X'\right) \right). \tag{2.7}$$

The theory of GPs can be used to estimate a deterministic function $f : X \to \mathbb{R}$ given noisy observations of f at points $\{x_1, ..., x_n\} \subseteq X$. Most commonly this is done by a *Gaussian process regression* (GPR), a nonparametric Bayesian method, which

models f as a random function. When performing a GPR with a given kernel k, one typically constructs a *prior distribution* by assuming

$$\begin{bmatrix} f(x'_1) \\ \vdots \\ f(x'_m) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, K\left(X', X'\right)\right) \tag{2.8}$$

for any finite subset $X' \subseteq X$ [75]. Assume we make n observations of the form $y_i = f(x_i^*) + \zeta_i$, where $\zeta_i \sim \mathcal{N}(0, \sigma^2)$ i.i.d, for $i = 1, ..., n$ and $\sigma > 0$. Importantly, ζ_i does not depend on f in any way. Then, by our prior assumption and by the introduction of the shorthand $\mathbf{f}(X') := (f(x_1), ..., f(x_m))^T$ and $\boldsymbol{\zeta} := (\zeta_1, ..., \zeta_m)^T$ we get that

$$\begin{bmatrix} \mathbf{f}(X^*) + \boldsymbol{\zeta} \\ \mathbf{f}(X') \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K(X^*, X^*) + \sigma^2 I_n & K(X^*, X') \\ K(X', X^*) & K(X', X') \end{bmatrix}\right).$$

Thus, by conditioning the above multivariate normal distribution of $\mathbf{f}$ on the observations $\mathbf{y} := (y_1, ..., y_n)^T$, we get [75]

$$\begin{aligned} \mathbf{f}(X')|\mathbf{f}(X^*) + \boldsymbol{\zeta} = \mathbf{y} \sim & \mathcal{N}(K(X', X^*)(K(X^*, X^*) + \sigma^2 I_n)^{-1}\mathbf{y}, \\ & K(X', X') - K(X', X^*)(K(X^*, X^*) + \sigma^2 I_n)^{-1} K(X^*, X')). \end{aligned} \tag{2.9}$$

The above distribution, for any finite $X' \subseteq X$, is the *posterior distribution* of f given X'. In the context of Bayesian modelling, we first summarise our knowledge in the values of f by the prior distribution: unless we gain further information, we assume f to be mean 0 with covariance K (Equation (2.8)). For any (noisy) observation of f we make, we update our belief in the values of f by conditioning our prior distribution on our observations. The posterior density at inputs X' can then be interpreted as how strongly we believe an output value to be the true output of f at X', given our observations and modelling assumptions.

If $m = 1$ (i.e., $X' = \{t\}$ for some $t \in X$), we denote the mean of the above conditional normal distribution by $\hat{f}_n(t)$ and its variance by $v_n(t)$. We henceforth call $\hat{f}_n(t)$ the *Gaussian smoothing of* f (on the set X^* of size n). When needed, we explicitly denote the dependence of $\hat{f}_n(t)$ on X^* and f by writing $\hat{f}_n(t, X^*, f)$. From Equation (2.9), it follows that $\hat{f}_n$ always lies in $\mathcal{H}_k$, the RKHS of k.

Under certain assumptions, one can show that $\hat{f}_n \to f$ in mean. In Chapter 5, we will use results by Koepernik and Pfaff [55], which give strong probabilistic convergence results in the case of X being a compact metric space.

We note that computing $\hat{f}_n$ requires the inversion of an $n \times n$-matrix and thus has a runtime of $\mathcal{O}(n^3)$. By using the ECT on $\hat{f}_n$ we thus lose some of the ECTs runtime

advantage (compared to the PHT and extended PHT). However, both versions of the PHT require $\mathcal{O}(n_s^3)$ computations *per direction* [37], where $n_s \geq n$ is the number of simplices in the triangulation of a shape, which still gives a combined Gaussian process and ECT pipeline an edge in terms of runtime. More importantly, it is common practice to approximate the (inverse) Gram matrix by a low-rank matrix approximation method, such as the Nystroem method [95] or random Fourier features [73]. Such methods run in $\mathcal{O}(l^3 + l^2 n)$, where $l \ll n$ is the approximate rank of the Gram matrix and thus are significantly faster than $\mathcal{O}(n^3)$.

Chapter 3

Topology for Neuronal Trees

Neurons are essential for cognitive function, and their activity is dependent on their morphology [60]. For example, when two simultaneous signals reach a neuron's dendrites, the mutual Euclidean distance between received signals affects the combined signal that reaches the soma, the cell body. In particular, nearby signals have a weaker total effect than those that are far apart [74]. Similarly, the length of the dendrites is known to affect the signal received by the soma, with signals that have traveled a long distance along the dendrites being weaker and more spread out than those that travel a short distance to the soma [40]. To further confound the study of neurons, it is known that neuronal morphology- even within the same animal [45]- can be highly heterogeneous, suggesting that different neurons have morphologies suited for different functions. These observations have naturally led to attempts to classify large sets of neurons via their geometry. As we have discussed, neurons are often represented by trees, collections of vertices in two or three dimensional Euclidean space connected by linear edges, with one such point denoting the location of the soma. Comparison of large classes of neurons can then be achieved by extracting numerical morphological descriptors, such as the number of bifurcations and branching angles, from these reconstructions [5, 34, 56, 80]. Another approach consists in computing density maps that capture the distribution of neurites in a neuron [50, 85, 89]. One drawback of this approach is that if we have two classes of neurons with different density maps, it is difficult to accurately measure the difference between two density maps (i.e., minimise the integral of their difference over all possible rotations and translations).

A classical example morphological descriptor is the *Sholl function* of a neuron [81], which assigns for every positive number R the number $s(R)$ of times a spherical shell of radius R intersects a given neuron. The Sholl function can be interpreted visually

when plotted as a function of R. By taking averages of Sholl functions, researchers are able to visualise the average structure of large classes of neurons.

We have already discussed how two independently developed TDA methods, persistent homology and the TMD, can be used to study neuronal morphology [51, 59]. Li et al. [59], who use extended persistence as described in Section 2.2.3, show that the Sholl function of a neuron can be computed from the union of two particular barcodes derived from the function which records the Euclidean distance of points on a neuron to the soma. Kanari et al. [51], develop and use the TMD instead of persistent homology in their work, converting barcodes into unweighted persistence images for the purpose of data analysis.

Here, we build on the work of both groups. We compute TMDs, but we use the path distance, or *intrinsic distance* to the soma rather than the radial (Euclidean) distance from the soma as done in [51]. The intrinsic distance to the soma has already been suggested as a relevant choice of function [51, 59] and briefly demonstrated on an example in [51, SI]; however, to our knowledge the use of this function has not been thoroughly investigated for neuronal morphology.

We show that different regions in the persistence images of TMDs generated with the path distance to the soma function can be interpreted to understand the morphology of large groups of neurons. Specifically, we discuss how weighted persistence images record information about neuronal branches, including their length within a population of neurons as well as their distance from the soma. Based on this framework we have the following contributions.

- We prove that for a class of functions containing the path distance to the soma the barcodes of [51] and [59] are equivalent (Theorem 20).
- We define a new class of morphology descriptors, the *topological morphology functions*, and show that they can be recovered from barcodes derived using path length. Moreover, we show that topological morphology functions can be approximated from persistence surfaces via integration (Theorem 21).
- In Section 3.2 we apply these techniques to neurons from control mice and mice that model tauopathies, a form of neurodegenerative disorders in which tau protein forms deposits in the brain. This topological framework finds that neurons in mice with a postsymptomatic tauopathy are on average shorter than controls; furthermore, these neurons exhibit less branching far from the soma than controls

- We prove that the weighted persistence image of a TMD is stable for the L_1, L_2, and L_∞ metrics under perturbations of the function f, or the addition or removal a short branch (Theorem 25).

3.1 The TMD for path distance

3.1.1 Interpretation

We restrict our attention to barcodes given by $\mathrm{TMD}(T, d)$, where for a node v in $N(T)$, $d(v)$ is the intrinsic distance from r to v in T as a subset of Euclidean space. In other words, the function d returns the length of the path along the neuron from the point given by v to the soma. Since the TMD algorithm decomposes T into a collection of branches, the sum of the lengths of all the intervals gives the total branch length of the neuron. Further, the only closed interval in $\mathrm{TMD}(T, d)$ will be $[0, L]$, where L is the greatest value of d, since d attains all of its local maxima on leaves and $d(r) = 0$. All other intervals $(x, y]$ represent a branch in the barcode decomposition of T that initiates at an intrinsic distance x from the root and terminates at an intrinsic distance of y from the root. Therefore, the persistence of each interval in $\mathrm{TMD}(T, d)$ is exactly the length of the branch that the interval represents, which must be positive. Hence the area under the persistence surface of $\mathrm{TMD}(T, d)$ is exactly the cumulative branch length of T.

A persistence surface of $\mathrm{TMD}(T, d)$ can be analyzed and interpreted back to its associated neuron in other ways as well. By identifying regions in the xy-plane where persistence surface $F(x, y)$ is large, we can read off the types of branches in $\mathrm{TMD}(T, f)$ that contribute to the total branch length. Indeed, regions where F is large with large y values correspond to longer branches, while regions where F is large with small y values correspond to the contribution of shorter branches. Similarly, regions where F is large with large x values correspond to branches that initiate close to the root. By contrast, regions where F is large with small x values correspond to branches that initiate far from the root, in the sense that the path from the root to the branch point at which they initiate is long. Note that the persistence image will always have a Gaussian contribution of weight L centered at $(0, L)$, since $\mathrm{TMD}(T, d)$ must contain the interval $[0, L]$. Since persistence images are discrete approximations of persistence surfaces, they inherit similar interpretations. These concepts are illustrated with artificial neurons, each with the same total branch length, in Figure 3.1. In Figure 3.2 we show how persistence images can summarise the traits of real morphologically distinct

neurons. From average persistence images we can identify the types of branches that tend to appear in an entire class of neurons.

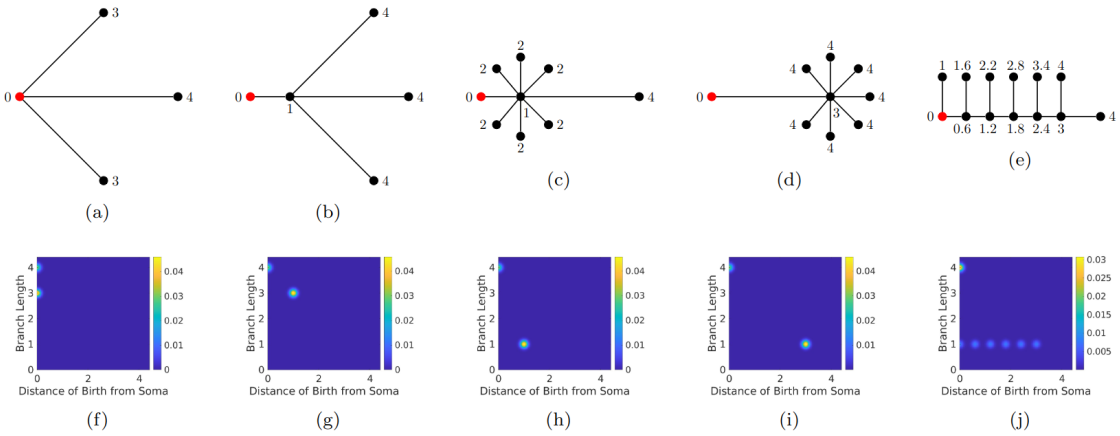

Figure 3.1: Persistence images of toy neuronal trees. Top row: Five toy neuronal trees with roots indicated by a red vertex. Bottom row: Persistence images corresponding to the $\mathrm{TMD}(T, d)$ of the toy neuronal trees in the top row. All persistence images of these trees have a feature centered at $(0, 4)$, indicating that the furthest path length from the root (soma) of the toy neurons is 4 units. The remaining pixels with positive values correspond to the remaining branches of each toy neuron. (a) There are two additional branches of total length 6 that originate at the soma, which correspond to a feature centered at the coordinates (0,3) in (f). The relative pixel intensities reflect that these branches correspond to a total length of 6 while the longest branch only corresponds to a total length of 4. (b) Two shorter branches are a distance of 1 unit from the soma, which corresponds to a birth value of 1 and branch length 3 in (g). (c) Six branches of length 1 all initiate at a distance of 1 from the soma, which corresponds to birth and persistence value of 1. In (d) and (i), six branches of length 1 that initiate at distance 3 from the root causes an increase in pixel intensity near the coordinates (3,1). (e) and (j), Six branches of length 1 initiate at different distances from the root which correspond to positive pixels along the line with a y value of one, however, since they initiate at different distances from the root, the pixel intensity is dispersed across birth values.

3.1.2 The TMD vs. persistent homology

As previously mentioned, Li et al. [59] instead assign the barcode $\mathrm{EPH}_0(T, f)$ to a tree T equipped with a function f. We show that when f is the path distance to the root function d, the methods of Li et al. and Kanari et al. are highly related:

Theorem 20. *Let T be a finite rooted tree with root r, and f be a function which is strictly increasing along the directed edges induced by r. We have that*

$$\mathrm{TMD}(T, f) = -\mathrm{EPH}_0(T, -f).$$

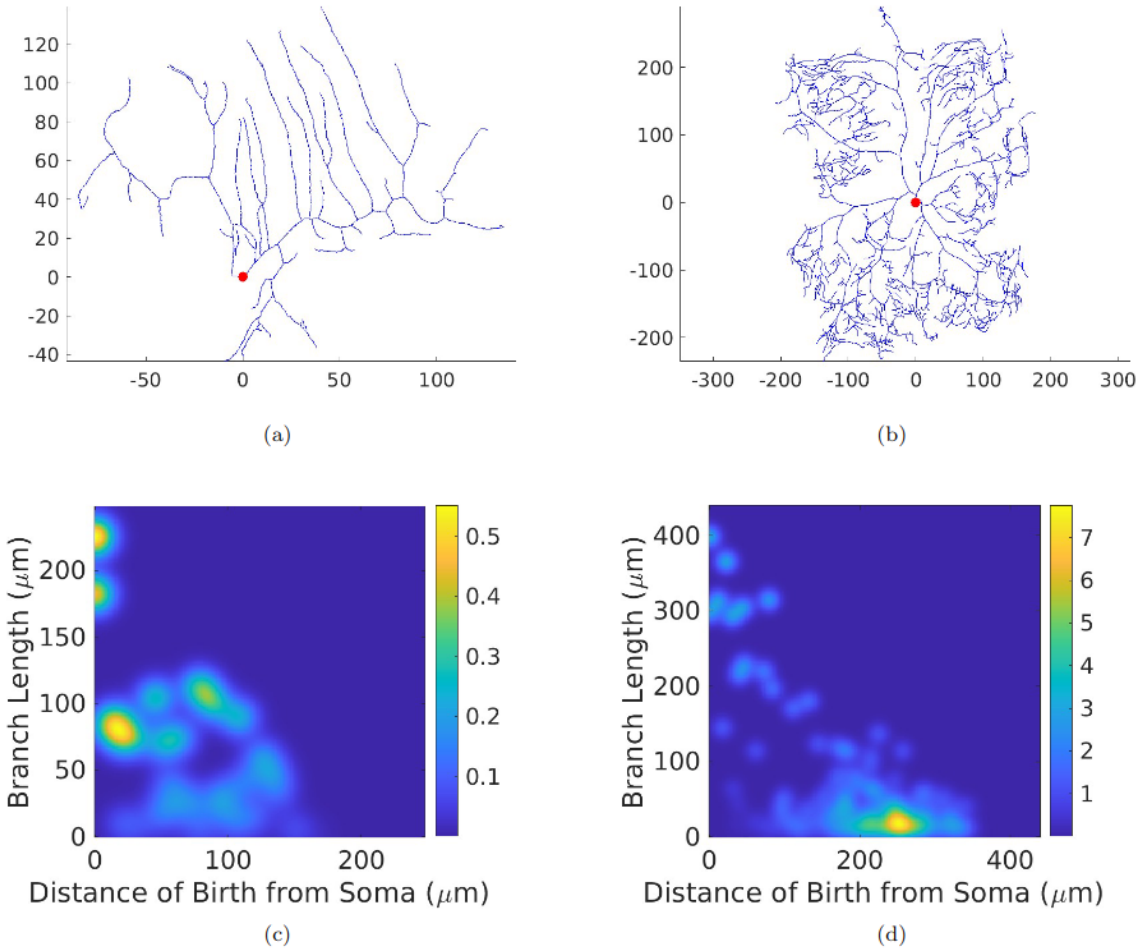

Figure 3.2: Persistence images of *Drosophila* neurons. (a) A class I neuron characterised by few main branches and many long secondary branches growing perpendicularly to the main branch. (b) An area-covering class IV sensory neuron with many small branches far from the root. In both (a) and (b) the soma is indicated by a red node. Below each in (c) and (d) are the corresponding persistence images. There are several bright regions substantially above the x-axis in (c) resulting from the many long secondary branches in (a). The brightest region of (d) is far to the right and close to the x-axis since many of the branches of (b) are far from the root and short. The digital reconstructions of these neurons from [87] are freely available at NeuroMorpho.org [4]. These neurons are named 02-16-09-Class1-B40X and 02-16-09-ClassIV on the site respectively.

In particular, d is a function satisfying the requirements of this theorem, and so the calculation of $\mathrm{TMD}(T, d)$ can be recast as a calculation using the methods of Li et al. [59]. We defer the proof of Theorem 20 to Section 3.4.1. When f does not satisfy the hypothesis of the theorem, there is no guarantee that these two methods are related. For example, when f is radial distance, the existence of branches that grow towards the soma may cause the equality of Theorem 20 to not hold.

3.1.3 Topological morphology functions

In [59], the authors show that a neuron's Sholl function can be reconstructed from a barcode obtained by twice applying EPH_0. By reasoning analogous to theirs, we will show that a similar function can be recovered from the barcode $\mathrm{TMD}(T, d)$. With the function d we lose information about the radial distance of branches from the soma necessary for the construction of Sholl functions. However, we can construct another family of neuron-derived functions which similarly describe branching morphology. Explicitly, we define the *topological morphology function* p of a neuron to be the function which associates to each positive number t the number of points on the neuron which have an intrinsic distance of t to the soma. This is an example of a Sholl descriptor, a kind of morphological descriptor defined and studied in [54]. The value $p(t)$ is the number of intervals in $\mathrm{TMD}(T, d)$ containing t, since this is the number of branches in the TMD decomposition of T containing points a distance of t from the soma. We call a number t *generic* if it is not an endpoint of an interval in $\mathrm{TMD}(T, d)$. For generic t, the number $p(t)$ is easily obtained from both persistence diagrams and transformed persistence diagrams, where the y value of a point denotes an interval's persistence instead of death, as shown in Figure 3.3. In a persistence diagram, $p(t)$ is the number of points above and to the left of (t, t) for generic t. In a transformed persistence diagram, $p(t)$ is the number of points in the region R_t, given by the equations

$$x < t, \quad y > t - x, \tag{3.1}$$

for generic t.

The following theorem shows that the topological morphology function p of a neuron represented by T can be approximated from persistence surfaces of $\mathrm{TMD}(T, d)$ via integration.

Theorem 21. *Let p be the topological morphology function associated to a neuron represented by a tree T. Let $d : N(T) \to [0, \infty)$ be the intrinsic distance to the soma.*

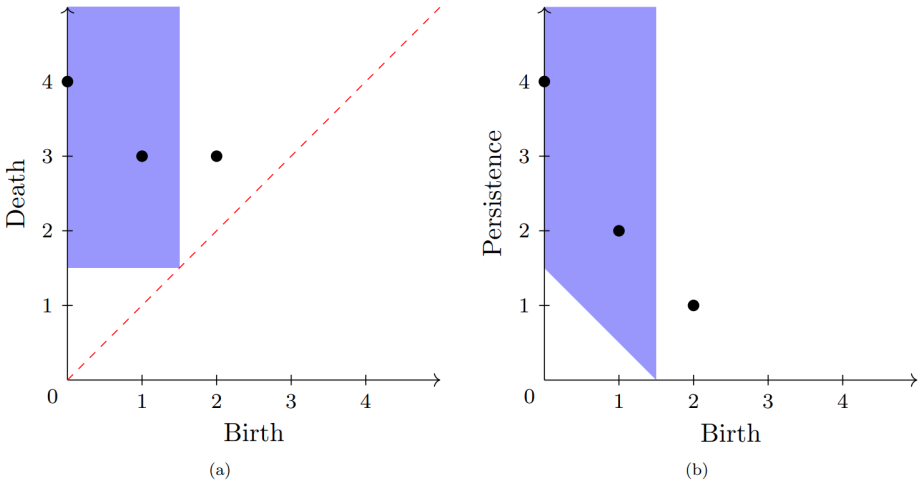

Figure 3.3: Recovering $p(1.5)$ from a standard and transformed persistence diagram. On the left, $p(1.5)$ is recovered from the persistence diagram by counting the number of points above and to the left of the point (1.5, 1.5). This region is shaded blue. On the right, $p(1.5)$ is recovered from the corresponding transformed persistence diagram by counting the number of points in $R_{1.5}$, shaded blue. In both cases, the number of points in the blue region is 2, so $p(1.5) = 2$.

Let $F_\sigma(x, y)$ be the persistence surface corresponding to $\mathrm{TMD}(T, d)$ *constructed with Gaussian functions of standard deviation σ. For any generic positive number t,*

$$p(t) = \lim_{\sigma \to 0} \int_{R_t} \frac{F_\sigma(x, y)}{y} \, dx \, dy. \tag{3.2}$$

Further the integral in the above expression converges and is an infinitely differentiable function of t for all $\sigma > 0$.

This theorem shows that the persistence surface of $\mathrm{TMD}(T, d)$ encodes a smooth approximation of the topological morphology function, which improves in accuracy as the standard deviation σ used to generate the persistence surface approaches zero. We show on an example in Figure 3.4 that approximation of the integral in Theorem 21 by numerical integration over a persistence image can also recover an approximation of the function p. Note that if every branch of a neuron grows in a straight line exactly radially outward from the soma, then the topological morphology function of the neuron is equal to the neuron's Sholl function s. In situations where this is a reasonable approximation, we expect that p and s are similar. Under such an assumption regarding the direction of neurite outgrowths, this theorem shows that we can use a persistence image to approximate the Sholl function s smoothly.

We defer the proof of Theorem 21 to Section 3.4.2, where we also prove an analogous result for the barcodes of Li et al. [59].

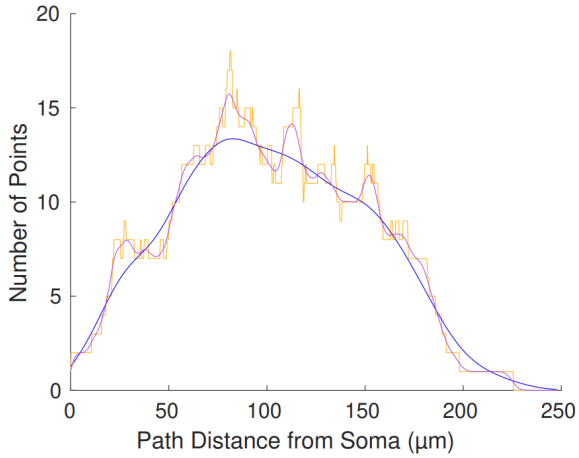

Figure 3.4: The topological morphology function (orange) for the neuron in Figure 3.2 (a) along with the approximate topological morphology function obtained by numerical integration of the persistence image for $\sigma = 10\mu m$ (blue) and $\sigma = 2\mu m$ (purple).

3.2 Methods

3.2.1 Data

We study Tau35 mice, a transgenic mouse line which has been used as a model to reproduce biological and cognitive features of human tauopathies [12]. In particular, we analyze topologically a set of 316 digitally reconstructed mouse neurons (as `.DAT` files), 170 wild type (WT) and 146 Tau35 mouse variant, to study the effect of tauopathy on morphology. These reconstructed neurons are further divided into groups by age (e.g., presymptomatic 4 month neurons or postsymptomatic 10-12 month neurons) and type (e.g., cortical or hippocampal neurons). All neurons in the dataset were grown in live mice and measured via a staining technique. We refer the reader to Section A of the appendix for the details of the data acquisition. A portion of this data is recorded via its projection onto the xy plane, while remainder of the data is recorded as a shape in three dimensional space.

In this data set, all mice are genetically identical, aside from variability between WT and tau35 mice in genes specifically altered to produce the tau35 strain. For each of the 8 groups determined by age, brain region, and strain of mice, neuronal data was determined from between 3 and 7 mice.

3.2.2 Preprocessing

We convert the data from `.DAT` to `.swc` format using NLMorphologyConverter. In `.swc` format, neurons are represented by a point cloud of points on the cell, with adjacent points on neurites indicated. Each `.swc` representation also features a collection of points representing the boundary of the soma. Since our methods do not take into account the morphology of the soma, we further preprocess the data by contracting these boundary points to a single point at their center of mass. By visual inspection, we find that the converted `.swc` converted files often attach neurites to the soma at inaccurate locations. To correct for this, we further preprocess the data by removing edges between each neurite and the soma, and adding a new connection between the point representing the soma and each neurite, at the point on the neurite that is closest to the contracted soma point.

3.2.3 Computations

Since the number of mice from which neurons are derived in each group in small, we assume that average neuronal morphology varies little between genetically identical mice. Since a portion of the data is projected into two dimensions, we use projected path length to the soma as a proxy for path length to the soma. We define the projected distance between any two adjacent points in a `.swc` representation of a neuron to be their mutual distance once projected in the xy plane. The projected path length between a point z and the soma is then defined to be the sum of projected distances between pairs of adjacent points along the shortest path between z and the point representing the soma. Letting $\widetilde{d}$ denote projected path length to the soma, we compute persistence images of $\mathrm{TMD}(T, \widetilde{d})$ associated to each neuron in the data set and take in-group averages. Each persistence image is generated from the persistence surfaces[1] of Gaussian functions with a standard deviation of $10\mu m$. We hypothesise that the difference between group averages of WT and Tau35 neurons of the same age and type is significant. To test this hypothesis, we compute the L^1 distance between average persistence images of the Tau35 and WT neurons in question. We then compare this L^1 distance to the L^1 distances for averages of 10000 randomised regroupings of the data.

[1]For computations we generated persistence images with 231×231 entries. The entries of these persistence images evenly cover the region $[-0.15L_{\max}, L_{\max}]^2 \subseteq \mathbb{R}^2$, where $L_{\max}$ is 1.1 times the maximum in-group projected branch length, across Tau35 and WT neurons. Note that the persistence images we use for computation sample values outside the first quadrant to assure that relevant information of points on the associated persistence diagram near the x or y axes is captured.

Table 3.1: The counts of neurons of different types in the dataset. The rightmost column denotes the percentage of random regroupings which had an L^1 distance of group averages less than the L^1 distances of the WT and Tau35 averages.

Age(Months)	Region	WT Count	Tau35 Count	Percent
10-12	Cortex	64	43	99.06
10-12	Hippocampus	38	49	99.68
4	Cortex	38	33	98.54
4	Hippocampus	30	21	94.87

3.3 Results for diseased vs wild type neurons

To test the hypothesis that WT neurons and tauopathy neurons are on average morphologically distinct, we average the persistence images of each group of neurons and perform a randomisation test as explained in the Methods. We show the results of our randomisation tests in Table 3.1, where the rightmost column is the percent of times the randomised L^1 distance was less than or equal to the same distance between the average Tau35 and WT neurons. We set a significance level of $\alpha = 0.05 = 5\%$. For each of the four randomisation tests we choose to accept the between-average difference as significant if less than $\frac{5}{4}$ percent of the regroupings had a greater between-average distance. We divide by 4 to account for the multiple comparisons problem, since we have performed 4 tests. Under this criterion, the between-average difference of the persistence images of Tau35 and WT neurons was significant for the 10-12 month groups, but not significant for the 4 month groups. In Figure 3.5, we show the average persistence images of the Tau35 and WT neurons of each 10-12 month group with their in-group differences.

We analyzed in-group averages of persistence images to identify features which distinguish WT neurons from Tau35 neurons on average at 10-12 months. For both hippocampal and cortical neurons in this age range (Figure 3.5 (c) and (f) respectively), we observe that the average Tau35 persistence images have a greater intensity towards the origin, while WT persistence images have greater pixel intensity for larger x and y values. From these observations we deduce that branches tend to not only be longer in WT neurons, but also often initiate further along the neuron than in Tau35 neurons.

In Figure 3.6 we plot the associated Sholl and topological morphology functions of the Tau35 and WT 10-12 month cortical and hippocampal neurons. We observe from both the Sholl and persistence functions that there are less branches far from

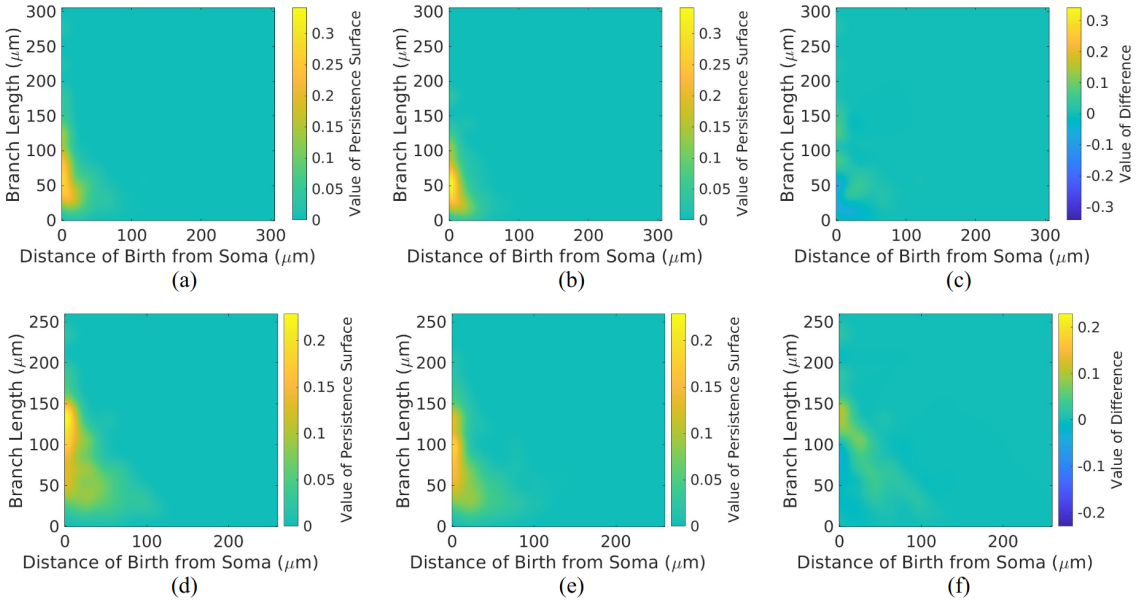

Figure 3.5: In-group average persistence images and the difference between Tau and WT groups for 10-12 month cortical (first row) and hippocampal (second row) neurons. (First column) average WT persistence images, (second column) average Tau35 persistence images, (third column) difference between WT and Tau35 average persistence images.

the root in Tau35 neurons than WT neurons. The observation that the topological morphology and Sholl functions are increasing at greater distances from the soma for WT neurons suggests that branching tends to occur at greater distances from the root in WT neurons. The analysis of the persistence images, which capture the joint distribution between branch initiation and branch termination confirms this conclusion.

A limitation of this study is that, as mentioned in the Methods section, we assume individual differences in the average neuronal morphology of genetically identical mice are small, since the data in consideration has few mice per group. While this could be accounted for by averaging persistence images per mouse, and then averaging across mice, this second approach has its own disadvantage. Since the number of neurons derived from each mouse is heterogeneous in the data, such an apprach may cause few neurons to drive the results.

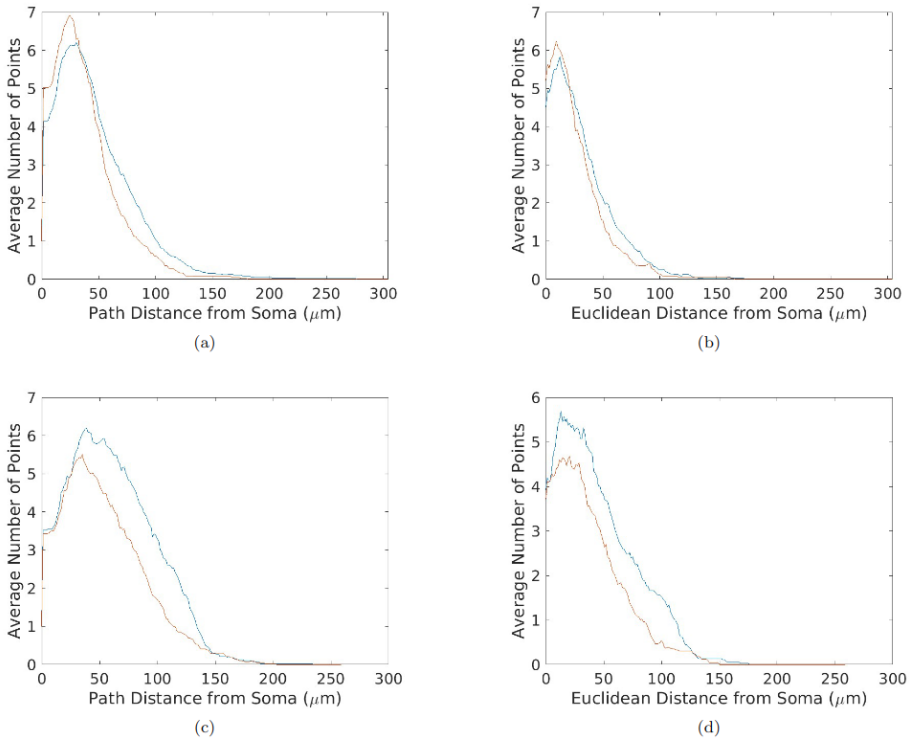

Figure 3.6: Average (left column) topological morphology and (right column) Sholl functions for (top row) 10-12 month cortical and (bottom row) 10-12 month hippocampal neurons of both WT (blue) and Tau35 (orange) classes.

3.4 Technical results

3.4.1 The TMD vs persistent homology revisited

In Section 3.1.2, we claimed that when a tree T is equipped with a function f which increases along paths away from the root, the TMD algorithm can be recast as an extended persistence computation, a fact that has been remarked and alluded to in prior related work, for example [52] and [59]. In this section we prove this fact. Before we do this, we recall that $\mathrm{TMD}(T, f)$ can be computed as follows

1. Choose any branch point b in T;
2. Identify one of the children c of b whose child branch maximises f on the vertices that are leaves of T descendant from b;
3. Detach all child branches for children $c' \neq c$ of b. Replace T with the resulting forest;
4. If there are no branch points in the resulting forest, then T is a collection of intervals with endpoints labelled via f, and we are done. Otherwise return to step 1.

For convenience we restate our theorem below.

Theorem 20. *Let T be a finite rooted tree with root r, and f be a function which is strictly increasing along the directed edges induced by r. We have that*

$$\mathrm{TMD}(T, f) = -\mathrm{EPH}_0(T, -f).$$

Proof. Let l_0 be one of the leaves with the largest f value, and let $f(l_0) = L$. From the fact that f is increasing with respect to the directed edges of T, it follows that the leaves of T are the only local minima of $-f$, and so the intervals of $\mathrm{PH}(T, -f)$ bijectively correspond to leaves of T. Similarly, the TMD algorithm also associates an interval to each leaf of T. Connected components merge in the sublevel set filtration of T and $-f$ only when a branch point b is included, and the components that are merged correspond to the child branches of b. We show the theorem holds by computing $\mathrm{TMD}(T, f)$ and $\mathrm{PH}(T, -f)$ concurrently. For each iteration of the TMD algorithm, choose b, one of the branch points with the greatest value of f, in order to ensure that every detached child branch is actually an interval, and thus contains a single leaf. Then, choose one of the child branches which attains the largest value of f on its leaf l. The Elder Rule [16, Theorem 4.4] then dictates that each leaf

$l' \neq l$ of a child branch of b may be associated to the interval $[-f(l'), -f(b))$ in $\mathrm{PH}(T, -f)$. Meanwhile the TMD algorithm associates to the child branch containing $l' \neq l$ the interval $(f(b), f(l')]$. The TMD algorithm then dictates that we detach the child branch containing each $l' \neq l$. This does not change the left endpoint given by the Elder Rule for an interval corresponding to any other leaf since removing the branch from b to l' does not change the minimal value of $-f$ at any of the remaining child branches of branch points of T, nor the value of $-f$ the remaining branch points of T themselves. After completing this procedure for every branch point of T, what remains to be computed is the infinite interval $[-L, \infty)$ in $\mathrm{PH}(T, -f)$ and the finite interval $[f(r), L]$ in $\mathrm{TMD}(T, f)$. From the monotonicity assumption, $f(r) = \max(-f)$, and so it is immediate that $\mathrm{TMD}(T, f) = -\mathrm{EPH}_0(T, -f)$. □

An immediate consequence of this result is that the methods of [51] and [59] can be used record identical information for functions f satisfying the requirements of the theorem. In particular, the function d, the path distance to the root, satisfies the requirements of this theorem, and so $\mathrm{TMD}(T, d) = -\mathrm{EPH}_0(T, -d)$.

3.4.2 Topological morphology functions from persistence surfaces

For a neuron represented by a tree T, recall that the associated topological morphology function $p(t)$ returns the number of points on the neuron with an intrinsic distance of t from the root.

Let us also recall how the barcode $\mathrm{TMD}(T, d)$ for the function d on T gives rise to a persistence surface. We can transform the persistence diagram of $\mathrm{TMD}(T, f)$ by the map $(x, y) \mapsto (x, y - x)$. If d is the intrinsic distance from the root function, we have observed that for t not an endpoint of an interval of $\mathrm{TMD}(T, d)$, $p(t)$ is the number of points in transformed persistence diagram of $\mathrm{TMD}(T, d)$ in the region given by $x < t$ and $y > t - x$, which we call R_t. Adding together two dimensional Gaussian functions with standard deviation σ centered at each point on the transformed persistence diagram of $\mathrm{TMD}(T, d)$, each weighted by the y value of their center, we produce the persistence surface $F_\sigma(x, y)$. Theorem 21, which remains to be proven, shows that an approximation of p can be constructed from F_σ.

We will formally restate this theorem shortly, but first we will need a definition and a lemma.

Definition 22. *Consider the family of boxes $B_\delta(\mu)$ of width and height δ centered around $\mu = (\mu_x, \mu_y)$. Let $g_\sigma(x, y; \mu)$ be a family of functions for positive σ satisfying the following properties:*

1. *The function $g_\sigma(x, y; \mu)$ is positive, and is bounded for fixed σ.*
2. *$\int_{\mathbb{R}^2} g_\sigma(x, y; \mu)\, dx\, dy = 1$.*
3. *$g_\sigma(x, y; \mu) \to 0$ uniformly as $\sigma \to 0$ on $B_\delta^C(\mu)$, the complement of $B_\delta(\mu)$.*
4. *$\int_{B_\delta(\mu)} g_\sigma(x, y; \mu)\, dx\, dy \to 1$ as $\sigma \to 0$.*
5. *Every partial derivative of $g_\sigma(x, y; \mu)$ with respect to x is continuous, and for fixed σ and μ is bounded in absolute value both by some constant and by $M/(x^2 + y^2)$ for some other constant M.*

Then we say that $g_\sigma(x, y; \mu) \in \mathcal{F}$.

Lemma 23. *If $g_\sigma(x, y; \mu) \in \mathcal{F}$, then*

$$\lim_{\sigma \to 0} \int_{R_t} \frac{g_\sigma(x, y; \mu)}{y}\, dx\, dy \tag{3.3}$$

is equal to zero if μ is in the exterior of R_t and is equal to μ_y^{-1} if μ is in the interior of R_t. Further, for each positive σ,

$$\int_{R_t} \frac{g_\sigma(x, y; \mu)}{y}\, dx\, dy$$

is infinitely differentiable as a function of t.

Proof. We only use property 5 to show the differentiability condition holds at the end. From properties 2 and 4 we immediately have that $\int_{B_\delta^C(\mu)} g_\sigma(x, y; \mu)\, dx\, dy \to 0$ as $\sigma \to 0$.

We first show that the integral in the theorem converges regardless of t. For positive h let Δ_h denote the triangular subset of R_t of points (x, y) additionally satisfying that $y < h$. Let M be a bound for $g_\sigma(x, y; \mu)$ for a given σ. We have

$$\int_{\Delta_h} \frac{g_\sigma(x, y; \mu)}{y}\, dx\, dy = \int_0^h \int_{t-y}^t \frac{g_\sigma(x, y; \mu)}{y}\, dx\, dy \leq \int_0^h \int_{t-y}^t \frac{M}{y}\, dx\, dy = Mh,$$
$$\int_{R_t - \Delta_h} \frac{g_\sigma(x, y; \mu)}{y}\, dx\, dy \leq \int_{R_t - \Delta_h} \frac{g_\sigma(x, y; \mu)}{h}\, dx\, dy \leq \frac{1}{h}.$$

To establish the limiting value in the theorem, first consider $g_\sigma(x, y; \mu)$ fixing μ in the exterior of R_t. For such μ there exists a δ sufficiently small that $B_\delta(\mu)$ is

disjoint from R_t. Let $S(\sigma)$ be the supremum of $g_\sigma(x,y;\mu)$ on $B_\delta^C(\mu)$ as a function of σ. Property 3 implies that this approaches 0 as σ does. We have

$$\int_{\Delta_h} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy = \int_0^h \int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_0^h \int_{t-y}^t \frac{S(\sigma)}{y}\,dx\,dy = hS(\sigma),$$
$$\int_{R_t-\Delta_h} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_{R_t-\Delta_h} \frac{g_\sigma(x,y;\mu)}{h}\,dx\,dy \le \frac{1}{h}.$$

Letting $h = S(\sigma)^{-1/2}$ we observe

$$0 \le \int_{R_t} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le 2S(\sigma)^{1/2},$$

which tends to zero as σ does.

If on the other hand μ lies in R_t, again we can choose δ small enough that $B_\delta(\mu)$ is a subset of R_t. Define $S(\sigma)$ as before. Once again we obtain bounds

$$\int_{\Delta_h - B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_{\Delta_h} \frac{S(\sigma)}{y}\,dx\,dy = \int_0^h \int_{t-y}^t \frac{S(\sigma)}{y}\,dx\,dy = hS(\sigma),$$
$$\int_{R_t-\Delta_h\cup B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_{R_t-\Delta_h} \frac{g_\sigma(x,y;\mu)}{h}\,dx\,dy \le \frac{1}{h}.$$

Combining these two results, and letting $h = S(\sigma)^{-1/2}$ we observe

$$0 \le \int_{R_t - B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le 2S(\sigma)^{1/2},$$

which approaches zero as σ approaches zero. Meanwhile, we also have

$$\int_{B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_{B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{\mu_y - \delta}\,dx\,dy \le \frac{1}{\mu_y - \delta}$$
$$\int_{B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \ge \int_{B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{\mu_y + \delta}\,dx\,dy \to \frac{1}{\mu_y + \delta},$$

with the limit in the last line taken as $\sigma \to 0$. Since δ can be taken arbitrarily small, this implies

$$\int_{B_\delta(\mu)} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \to \frac{1}{\mu_y}.$$

Hence,

$$\int_{R_t} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \to \frac{1}{\mu_y}.$$

In summary, we have shown that the above integral approaches μ_y^{-1} when μ is interior to R_t and approaches zero when μ is exterior to R_t.

All that remains to be shown is the differentiability statement of the lemma, i.e. we must show that the integral

$$I(t;\mu) := \int_{R_t} \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy = \int_0^\infty \int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy$$

is an infinitely differentiable function of t.

To begin, notice that by continuity of g we can extend the integrand to $y = 0$ via

$$\lim_{y\to 0} \int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{y}\,dx = g_\sigma(t,y;\mu).$$

Further, we have the estimate for positive C:

$$\int_C^\infty \int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{y}\,dx\,dy \le \int_C^\infty \int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{C}\,dx\,dy \le \int_{\mathbb{R}^2} \frac{g_\sigma(x,y;\mu)}{C}\,dx\,dy \le \frac{1}{C},$$

showing uniform convergence of the same integral with C taken to be 0, since g is positive valued. We also have that

$$\begin{aligned}\frac{\partial^n}{\partial t^n}\int_{t-y}^t \frac{g_\sigma(x,y;\mu)}{y}\,dx &= \frac{\partial^{n-1}}{\partial t^{n-1}}\frac{g_\sigma(t,y;\mu) - g_\sigma(t-y,y;\mu)}{y} \\ &= \frac{1}{y}\Big[\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t,y;\mu) - \frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t-y,y;\mu)\Big].\end{aligned}$$

This expression is bounded by any constant bounding of $|\frac{\partial^n}{\partial t^n}g_\sigma(t,y;\mu)|$. Hence, applying $\int_0^\infty dy$ to this expression, the lower limit converges uniformly. Meanwhile, letting M be a bound for the integral of $|\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t,y;\mu)|$ over linear domains in t and y, we observe that the upper limit converges uniformly as well since

$$\begin{aligned}&\Big|\int_C^\infty \frac{1}{y}\Big[\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t,y;\mu) - \frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t-y,y;\mu)\Big]\,dy\Big| \\ &\le \int_C^\infty \frac{1}{y}\Big[\Big|\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t,y;\mu)\Big| + \Big|\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t-y,y;\mu)\Big|\Big]\,dy \\ &\le \int_C^\infty \frac{1}{y}\Big[\frac{M}{t^2+y^2} + \frac{M}{(t-y)^2+y^2}\Big]\,dy \\ &\le \int_C^\infty \frac{2M}{y^3}\,dy \\ &\le \frac{4M}{C^2}.\end{aligned}$$

From this we know we can interchange the partial derivative and integral in the expression for $I(t;\mu)$, giving us the formula

$$\frac{\partial^n}{\partial t^n}I(t;\mu) = \int_0^\infty \frac{1}{y}\Big[\frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t,y;\mu) - \frac{\partial^{n-1}}{\partial t^{n-1}}g_\sigma(t-y,y;\mu)\Big]\,dy,$$

completing the proof. □

With this Lemma we can easily show Theorem 21.

Theorem 21. *Let p be the topological morphology function associated to a neuron represented by a tree T. Let $d : N(T) \to [0, \infty)$ be the intrinsic distance to the soma. Let $F_\sigma(x, y)$ be the persistence surface corresponding to* $\mathrm{TMD}(T, d)$ *constructed with Gaussian functions of standard deviation σ. For any generic positive number t,*

$$p(t) = \lim_{\sigma \to 0} \int_{R_t} \frac{F_\sigma(x, y)}{y} \, dx \, dy. \tag{3.2}$$

Further the integral in the above expression converges and is an infinitely differentiable function of t for all $\sigma > 0$.

Proof of Theorem 21. Let $\mu_1, \ldots, \mu_N$ be the coordinates of the off-diagonal points in the transformed persistence diagram of $\mathrm{TMD}(T, f)$. The persistence surface is then given by the formula

$$F_\sigma(x, y) = \sum_{i=1}^{N} (\mu_i)_y \, g_\sigma(x, y; \mu_i). \tag{3.4}$$

Each function $g_\sigma(x, y; \mu_i)$ is a two dimensional Gaussian function, for which the requirements of Lemma 23 are elementary properties. Each statement of Theorem 21 is immediate from the fact that F_σ is a finite linear combination of functions satisfying the requirements of Lemma 23. Indeed, in light of Equation 3.4, the integral in Equation 3.2 is clearly infinitely differentiable from the previous lemma. When t is generic with respect to B, it is easily seen that each μ_i is not on the boundary of R_t. Applying Lemma 23 and Equation 3.4, we see that the limit

$$\lim_{\sigma \to 0} \int_{R_t} \frac{F_\sigma(x, y)}{y} \, dx \, dy$$

evaluates to the number of μ_i in R_t, which is equal to $p(t)$. □

Now we prove an analogous result for Sholl functions. Let T be a tree with an embedding P to $\mathbb{R}^n$ and f be the Euclidean distance of each point represented by v to the root r. Recall that we define the Sholl function $s(t)$ to be the number times the n-sphere of radius t centered about $P(r)$ intersects with $P(T)$. Let B be the barcode

$$\mathrm{EPH}_0(T, f) \sqcup -S(\mathrm{EPH}_0(T, -f)), \tag{3.5}$$

the disjoint union of barcodes given by persistent homology of the superlevel and sublevel sets of f. In this setting we say that t is generic if it is not an endpoint of an interval in B. Let D_B be the persistence diagram associated to B. If instead of

assuming the embedding of edges of T in $\mathbb{R}^n$ is linear, we assume that P is such that f is weakly increasing or decreasing on edges, Li et al. show in [59, Section 2.4] that the value $s(t)$ can be recovered from D_B for any generic t. For the remainder of this subsection, we only consider T with such an embedding. For generic t between 0 and $\max(f)$, Li et al. prove that the Sholl function $s(t)$ is the number of points above and to the left or below and to the right of (t, t) in D_B minus one. More formally written, if Q_t is the set of (x, y) such that ($x < t$ and $y > t$) or ($x > t$ and $y < t$), then

$$s(t) = \left|B \cap Q_t\right| - 1. \tag{3.6}$$

If t is generic, but t is not between 0 and $\max(f)$ it is easily seen that the left side of this equation is 0 while the right side of this equation is -1. We show how $s(t)$ may be calculated from D_B visually in Figure 3.7.

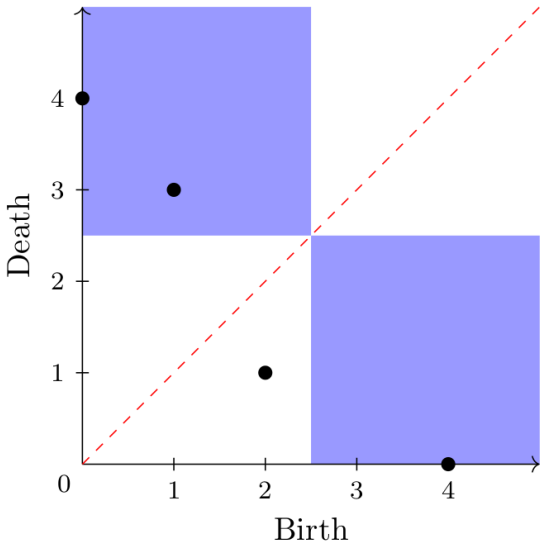

Figure 3.7: Calculating the $s(2.5)$ from the persistence diagram D_T. Shown in blue is the region $Q_{2.5}$, within which there are three points. Hence, $s(2.5) = 2$.

One may define a persistence surface F_σ for the barcode B defined in Equation 3.5 by adding a two dimensional Gaussian function for each point in D_B, with multiplicity, each weighted by the vertical distance of the point from the diagonal[2]. We can use our Lemma 23 to show that an approximate Sholl function can be recovered from this persistence surface.

[2]It is common to skip the transformation step when there are points below the diagonal in the initial barcode.

Theorem 24. *Let T be a rooted tree with Sholl function s and f be the associated Euclidean distance from the root function. Assume that between adjacent vertices, the embedding of T in $\mathbb{R}^2$ or $\mathbb{R}^3$ is such that f is either weakly increasing or decreasing along each edge. Let B be the barcode given by Equation 3.5 and D_B be the associated persistence diagram. Let $F_\sigma(x,y)$ be the persistence surface corresponding to D_B constructed with Gaussian functions of standard deviation σ, without the transformation step $(x,y) \mapsto (x, y-x)$. For any generic positive number t between 0 and $\max(f)$,*

$$s(t) = \lim_{\sigma \to 0} \int_{Q_t} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy - 1.$$

Otherwise, if t is generic, the left side of this equation is 0 and the right side of this expression is -1. Further, the integral in the above expression converges and is an infinitely differentiable function of t for all $\sigma > 0$.

Proof. Let $\mu_1, \ldots, \mu_N$ be the coordinates of the off-diagonal points in D_B, with x and y values $(\mu_i)_x$ and $(\mu_i)_y$ respectively. By definition the persistence surface F_σ is given by a weighted sum of Gaussians:

$$F_\sigma(x,y) = \sum_{i=1}^{N} \left|(\mu_i)_y - (\mu_i)_x\right| g_\sigma(x,y;\mu_i).$$

We denote by Q_t^+ the portion of Q_t satisfying $y > x$ and similarly denote by Q_t^- the portion of Q_t satisfying $x > y$. We have the equation

$$\int_{Q_t} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy = \int_{Q_t^+} \frac{F_\sigma(x,y)}{y-x}\, dx\, dy + \int_{Q_t^-} \frac{F_\sigma(x,y)}{x-y}\, dx\, dy.$$

Examining the first term, we have

$$\int_{Q_t^+} \frac{F_\sigma(x,y)}{y-x}\, dx\, dy = \int_0^\infty \int_{-\infty}^t \frac{F_\sigma(x,y)}{y-x}\, dx\, dy = \int_0^\infty \int_{t-v}^t \frac{F_\sigma(u,u+v)}{v}\, du\, dv$$

via the transform $(x,y) \mapsto (x, y-x) := (u,v)$. Noting that linear transformations of Gaussian density functions are still Gaussian density functions and that this last integral is over the region R_t in the uv plane, we have:

$$\int_0^\infty \int_{t-v}^t \frac{F_\sigma(u,u+v)}{v}\, du\, dv = \sum_{i=1}^{N} \left|(\mu_i)_y - (\mu_i)_x\right| \int_{R_t} \frac{g_\sigma(u,u+v;\mu_i)}{v}\, du\, dv.$$

We may now apply Lemma 23 as each g_σ is Gaussian with mean $((\mu_i)_x, (\mu_i)_y - (\mu_i)_x)$ in the uv plane and therefore upholds the requirements of the lemma. Thus the integrals on the right are well defined and infinitely differentiable. Consider the i^{th}

of these integrals on the right. Taking the limit $\sigma \to 0$ this integral evaluates to $((\mu_i)_y - (\mu_i)_x)$ if $((\mu_i)_x, (\mu_i)_y - (\mu_i)_x)$ is interior to R_t and 0 if it is exterior to R_t. Equivalently, the integral evaluates to $((\mu_i)_y - (\mu_i)_x)$ if μ_i is interior to Q_t^+ and 0 if it is exterior to Q_t^+.

No μ_i lies on the boundary of Q_t^+. Hence, the limit

$$\lim_{\sigma \to 0} \int_{Q_t^+} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy$$

evaluates to the number of μ_i in Q_t^+ for generic t.

Similarly, we may show via the transformation $(x, y) \mapsto (y, x - y)$ that the limit

$$\lim_{\sigma \to 0} \int_{Q_t^-} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy$$

is the number of μ_i in Q_t^- for generic t, and that the integral here is always well defined and infinitely differentiable regardless of t.

Hence the integral

$$\int_{Q_t} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy = \int_{Q_t^+} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy + \int_{Q_t^-} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy$$

is well defined and infinitely differentiable for any σ.

Now consider generic t. We have

$$\lim_{\sigma \to 0} \int_{Q_t} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy - 1 = \lim_{\sigma \to 0} \int_{Q_t^+} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy + \lim_{\sigma \to 0} \int_{Q_t^-} \frac{F_\sigma(x,y)}{|x-y|}\, dx\, dy - 1,$$

which is the number of points in the barcode that also lie in Q_t minus one. This value has been shown to be $s(t)$ in [59] when t is between 0 and $\max(f)$. When t is not between 0 and $\max(f)$, there can be no points in Q_t^+ or Q_t^-, and so this expression evaluates to -1. Clearly for such t, $s(t) = 0$.

□

3.4.3 1-Wasserstein Stability

Here we show that the TMD operation is stable for the 1-Wasserstein distance under certain perturbations one might consider in practice. As noted at the end of Section 2.2.2, persistence images are stable for L_1, L_2, and L_∞ metrics under the 1-Wasserstein distance of persistence diagrams. Hence the results presented here also lead to stability results for persistence images. To this end, we proceed in a fashion similar to [51, SI]. There are four different ways of perturbing data (T, f) to (T', f') that we will show are stable:

P1 *Vertex perturbation.* $T' = T$ and f' satisfies $|f'(v) - f(v)| \leq \varepsilon$ for all vertices v in $N(T)$.

P2 *Attaching an edge to a vertex.* $N(T') = N(T) \cup \{v'\}$, $E(T') = E(T) \cup \{(v, v')\}$ for some vertex $v \in N(T)$, $f' = f$ on $N(T)$, $|f'(v') - f(v)| \leq \varepsilon$, and $f(v) \leq f(l) + \varepsilon$ for at least one leaf l of T that is a descendant of v. Put another way, we attach a short edge to one of the vertices of T, one of whose descendant leaves does not have a much smaller f value than its own.

P3 *Attaching an edge to an edge.* There is an edge (n, n') in $E(T)$, $N(T') = N(T) \cup \{v, v'\}$, $E(T') = E(T) \cup \{(n, v), (v, n'), (v, v')\} - \{(n, n')\}$, $f = f'$ on $N(T)$, $|f(v) - f(v')| \leq \varepsilon$, and $f'(v) \leq f(l) + \varepsilon$ for at least one leaf l of T that is a descendant of v. In words, we add a vertex interior to an edge of T, ensure that one of the descendant leaves of this vertex does not have a much smaller f value than its own, and attach a short edge to this vertex.

P4 *Edge retraction.* (T, f) is obtained by a perturbation of type **P2** or **P3** from (T', f').

We leave the root vertex unchanged in all of these classes of perturbation.

We show visualisations of these transformations in Figure 3.8, along with a type of transformation we do not claim to be stable.

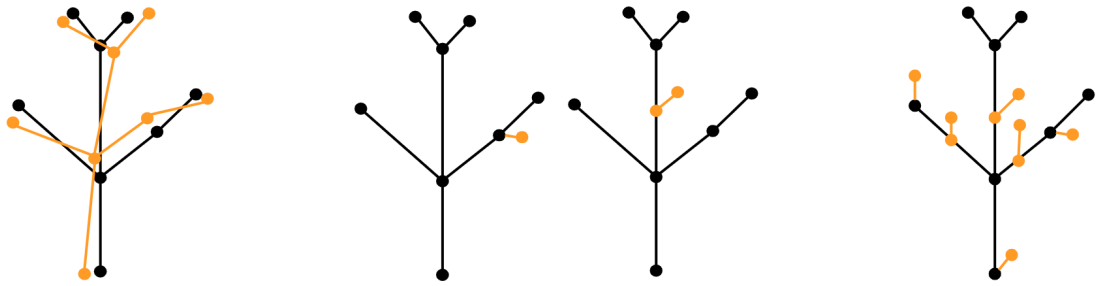

Figure 3.8: Tree perturbations we show are stable, and one type of perturbation that is unstable.

Now we state our main stability theorem.

Theorem 25. *Let T be a rooted tree and $f : T \to \mathbb{R}$. Let also T' be a rooted tree with $f' : T' \to \mathbb{R}$.*

1. *If (T', f') is obtained by a perturbation of type **P1** of (T, f), then*

$$d_1(\mathrm{TMD}(T, f), \mathrm{TMD}(T', f')) \leq |L(T)|\varepsilon.$$

2. *If (T', f') is obtained by a perturbation of type **P2** or **P3** of (T, f), then*

$$d_1(\mathrm{TMD}(T, f), \mathrm{TMD}(T', f')) \leq \frac{5}{2}\varepsilon.$$

*Furthermore, if the vertex v given by the perturbation of type **P2** or **P3** is such that $f'(v) \leq f(l)$ for some leaf l of T that is a descendant of v, then*

$$d_1(\mathrm{TMD}(T, f), \mathrm{TMD}(T', f')) \leq \varepsilon.$$

3. *If (T', f') is obtained by a perturbation of type **P4** of (T, f), then*

$$d_1(\mathrm{TMD}(T, f), \mathrm{TMD}(T', f')) \leq \frac{5}{2}\varepsilon.$$

*Furthermore, if the vertex v in T given by the perturbation of type **P2** or **P3** from (T', f') to (T, f) is such that $f(v) \leq f'(l)$ for some leaf l of T' that is a descendant of v, then*

$$d_1(\mathrm{TMD}(T, f), \mathrm{TMD}(T', f')) \leq \varepsilon.$$

Note that the 1-Wasserstein stability results are weaker than those obtained for the bottleneck distance in [51, SI, Theorem 1]. First, the Lipschitz constant we provide in statement 1 of the theorem is dependent on the number of leaves in T. Second, for statements 2 and 3, we are not allowing for arbitrarily intervals of a fixed length to be attached to or removed from T. With respect to the 1-Wasserstein distance, it is not hard to see that this kind of perturbation is unstable. For instance, if T is a single root vertex equipped with the zero map f, and we obtain (T', f') by attaching n vertices to this root, each with f' value ε, we can make the 1-Wasserstein distance between the two associated TMDs arbitrarily large by increasing n.

We will first prove statement 1 of this theorem, then statement 2. Statement 3 follows immediately from statement 2, since these statements are identical with the roles of (T, f) and (T', f') reversed. Our methods of proof are based on those of [51, SI, Section 4]. The following proposition is equivalent to statement 1.

Proposition 26. *Let T be a rooted tree with root r and $f, f' : N(T) \to \mathbb{R}$ such that $|f(v) - f'(v)| \leq \varepsilon$ for all vertices v in $N(T)$. Then*

$$d_1(\mathrm{TMDg}(T, f), \mathrm{TMDg}(T, f')) \leq |L(T)|\varepsilon.$$

Proof. If T is a single vertex, then T has one leaf, which is also its root, and the fact that

$$d_1(\mathrm{TMDg}(T, f), \mathrm{TMDg}(T, f')) \leq \varepsilon$$

is clear.

We proceed with the remaining cases by induction on the depth of T, each time proving that every point not in diag in each persistence diagram can simultaneously be matched with another point not in diag that is ε away in sup norm, provided T is of depth k. This will imply the desired result since the TMD algorithm produces exactly one interval for each leaf. We will assume inductively that the matching we get for trees of depth $k-1$ between functions f and f' sends a copy of $(f(r), L)$ to a point with first coordinate $f'(r)$, and whose inverse sends a copy of $(f'(r), L')$ to a point with first coordinate $f(r)$, where L and L' are the maximum values of f and f' on the leaves of T respectively. This additional constraint can clearly be satisfied for the base case where T has depth zero, and hence is a single vertex.

Suppose T is a tree of depth d with f and f' as in the statement of the theorem. Let $c_1, \ldots, c_n$ denote the children of the root of T. Each c_i is the root of a subtree T_i of T whose vertices are c_i, its children, its children's children, and so on. Let b_i be the leaf that takes the greatest value of f subject to the constraint that c_i is an ancestor of b_i. We define b_i' similarly for f'. Thus,

$$\mathrm{TMDg}(T, f) = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup \bigsqcup_{i=1}^{n} D_i,$$
$$\mathrm{TMDg}(T, f') = \mathrm{diag} \sqcup \{(f'(r), f'(b_i')) : 1 \leq i \leq n\} \sqcup \bigsqcup_{i=1}^{n} D_i',$$

where

$$D_i = \mathrm{TMDg}(T_i, f) - \{(f(c_i), f(b_i))\},$$
$$D_i' = \mathrm{TMDg}(T_i, f') - \{(f'(c_i), f'(b_i'))\}.$$

If we can show that there is a matching of D_i and D_i' that sends no point more than ε away with respect to the sup norm, then the matching that additionally sends

each $(f(r), f(b_i))$ to $(f'(r), f(b_i'))$ does the same for $\mathrm{TMDg}(T, f)$ and $\mathrm{TMDg}(T, f')$. Indeed, assuming without loss of generality that $f(b_i) \geq f'(b_i')$,

$$|f(b_i) - f'(b_i')| = f(b_i) - \max_{l \in L(T_i)} f'(l) \leq f(b_i) - f'(b_i) \leq \varepsilon$$

Further, this matching and its inverse also satisfy our additional inductive hypothesis regarding points with left coordinate equal to $f(r)$ and $f'(r)$. Since there are exactly $|L(T)|$ points in both persistence diagrams excluding diag we will be done once we can find a desired matching between D_i and D_i'.

By inductive hypothesis we already have such a matching between $\mathrm{TMDg}(T_i, f)$ and $\mathrm{TMDg}(T_i, f')$, and we will adjust this matching to produce a matching between D_i and D_i'. If our given matching sends $(f(c_i), f(b_i))$ to $(f'(c_i), f'(b_i'))$ then we immediately get a desired matching between D_i and D_i'.

Otherwise, let $\phi : \mathrm{TMDg}(T_i, f) \to \mathrm{TMDg}(T_i, f')$ be our matching and let $(x', y') = \phi((f(c_i), f(b_i)))$, $(x, y) = \phi^{-1}((f'(c_i), f'(b_i')))$. Thus we may define $\widetilde{\phi} : D_i \to D_i'$ by sending (x, y) to (x', y') and every other point to its image via ϕ. By our inductive hypothesis, $x = f(c_i)$ and $x' = f'(c_i)$. Thus $|x - x'| \leq \varepsilon$. Assume without loss of generality that $f'(b_i') \leq f(b_i)$. Now assume, *with* loss of generality, that $f'(b_i') \geq y$. Thus we have the equation.

$$|y - y'| = |y - f'(b_i') + f'(b_i') - y'|. \tag{3.7}$$

Both differences on the right side of the above equation are bounded by ε, the first due to the matching ϕ, and the last since by definition of b_i', $f'(b_i') \geq y'$ and so

$$|f'(b_i') - y'| = f'(b_i') - y' \leq f(b_i) - y' \leq \varepsilon.$$

Notice also that the last difference in Equation 3.7 is nonnegative. Meanwhile, $y - f'(b_i')$ is nonpositive since $f'(b_i') \geq y$. Thus, Equation 3.7 shows that $|y - y'|$ is a difference of nonnegative numbers each bounded by ε, and hence itself is less than or equal to ε. Hence $\widetilde{\phi}$ is a matching which alters each coordinate of every point by no more than ε.

Finally, still assuming that $f'(b_i') \leq f(b_i)$, we also assume $f'(b_i') \leq y$. Now the relevant equation is

$$|y - y'| = |y - f(b_i) + f(b_i) - y'|. \tag{3.8}$$

Once again, both differences on the right are bounded by ε, the second due to the matching ϕ, and the first since by definition of b_i, $y \leq f(b_i)$, and so

$$|y - f(b_i)| = f(b_i) - y \leq f(b_i) - f'(b_i') \leq \varepsilon.$$

As before, we have also just shown the first difference in the right side of Equation 3.8 is nonpositive, while the second is nonnegative since $f(b_i) \geq f'(b_i') \geq y'$. Hence $|y - y'|$ is the difference of two nonnegative numbers bounded by ε, and so is itself less than ε. Again, this implies that $\widetilde{\phi}$ is a matching which sends no point more than ε away with respect to the sup norm and the proof is complete. □

Remark 2. *Kanari et al. [51, SI] show perturbations of type **P1** are bottleneck stable by constructing a matching that sends points not in* diag *to each other. Proposition 26 follows from their matching and the fact that the TMD algorithm produces a diagram with $|L(T)|$ points excluding* diag*. The construction of the matching in [51] implicitly assumes that $b_i = b_i'$. Here, we use a different inductive hypothesis and therefore include our own full construction for completeness.*

In fact, for our above proposition to hold, the values of f and f' only need to be close on the root, branch points, and leaves- the following proposition shows we can remove all other vertices without changing TMDg.

Proposition 27. *Let T be a tree with root r and an edge (a, b). Let T' be the tree given by $N(T') = N(T) \cup \{v'\}$ and $E(T') = E(T) \cup \{(a, v'), (v', b)\} - \{(a, b)\}$, with root r. Let $f : N(T) \to \mathbb{R}$ and $f' : N(T') \to \mathbb{R}$ be such that f' restricts to f on $N(T)$. Then*

$$\mathrm{TMDg}(T, f) = \mathrm{TMDg}(T', f').$$

Proof. We prove this by induction on the depth of T. If T has depth zero, then T is a single vertex and the theorem is vacuously true since T has no edge.

Otherwise, assume the result holds for adding vertices to trees of depth less than k and T is of depth k. Then, first assume neither a nor b is the root of T. Thus, adopting the notation of c_i, b_i, T_i, and D_i from the previous proof, we observe that a and b must in some T_j. We define T_j' by $N(T_j') = N(T_j) \cup \{v'\}$ and $E(T_j') = E(T_j) \cup \{(a, v'), (v', b)\} - \{(a, b)\}$. We further define $D_j' = \mathrm{TMDg}(T_j', f') - \{(f(r), f(b_j))\}$. It follows that

$$\mathrm{TMDg}(T, f) = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup \bigsqcup_{i=1}^{n} D_i,$$
$$\mathrm{TMDg}(T', f') = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup D_j' \sqcup \bigsqcup_{i \neq j} D_i.$$

But by our inductive hypothesis these D_j and D_j' are identical, so the above persistence diagrams must be identical as well. It remains to check the case where either a

or b is the root. Without loss of generality, suppose a is the root of T. Then $b = c_j$ for some j. Let T'_j be given by $N(T'_j) = N(T_j) \cup \{v'\}$ and $E(T'_j) = E(T_j) \cup \{(v', c_j)\}$. Thus

$$\mathrm{TMDg}(T', f') = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup D'_j \sqcup \bigsqcup_{i \neq j} D_i,$$

where

$$D'_j := \mathrm{TMDg}(T'_j, f') - \{(f'(v'), f(b_i))\} = \mathrm{TMDg}(T_j, f) - \{(f(c_j), f(b_i))\} = D_j.$$

Hence $\mathrm{TMDg}(T, f) = \mathrm{TMDg}(T', f')$. □

The following proposition shows that statement 2 of Theorem 25 holds for perturbations of type **P2**. By applying the previous proposition as well, the statement also holds for perturbations of type **P3**.

Proposition 28. *Let T be a rooted tree with root r and $f : N(T) \to \mathbb{R}$. Let v be any vertex of T and T' be another tree given by $N(T') = N(T) \cup \{v'\}$ and $E(T') = E(T) \cup \{(v, v')\}$, with root r. Let $f' : N(T') \to \mathbb{R}$ be such that $f' = f$ on $N(T)$ and $|f'(v') - f'(v)| \leq \varepsilon$. If $f(v) \leq f(l) + \delta$ for at least one leaf l which is a descendant of v then*

$$d_1(\mathrm{TMDg}(T, f), \mathrm{TMDg}(T', f')) \leq \varepsilon + \frac{3}{2}\delta.$$

Proof. To avoid certain mild irritations during the proof, we will first prove the proposition when f' (and hence also f) is injective. We will then attain the general case by perturbing f' via Proposition 26.

For convenience, let $\eta = \varepsilon + \frac{3}{2}\delta$. In this proof we proceed by induction on the depth k of v instead of the depth of T. As our inductive hypothesis, we assume that when v has depth less than k in a tree T, and T' satisfies the hypothesis of the proposition the following is true. If l maximises f on the leaves of T and l' maximises f' on the leaves of T', then we can find a η-matching sending $(f(r), f'(l'))$ to $(f(r), f(l))$. Hence for our induction step we will need to construct a matching with this additional property.

For our base case $v = r$. As in the previous proofs, we may write

$$\mathrm{TMDg}(T, f) = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup \bigsqcup_{i=1}^{n} D_i$$

with D_i, T_i, b_i, and c_i defined as in the proof of Proposition 26. If we attach the vertex v' to r, then we have

$$\mathrm{TMDg}(T', f') = \mathrm{diag} \sqcup \{(f(r), f'(v'))\} \sqcup \{(f(r), f(b_i)) : 1 \leq i \leq n\} \sqcup \bigsqcup_{i=1}^{n} D_i$$

Hence by sending $(f(r), f'(v'))$ to $(f(r), f(r))$ on the diagonal, and everything else to itself we obtain a η-matching between $\mathrm{TMDg}(T', f')$ and $\mathrm{TMDg}(T, f)$. However, if v' maximises f' on the leaves of T' and the depth of T is nonzero then this matching does not satisfy the inductive hypothesis. Whenever this happens, we can instead choose the η-matching with the following mappings

$$(f(r), f'(v')) \longmapsto (f(r), f(b_m))$$
$$(f(r), f(b_m)) \longmapsto ((f(r) + f(b_m))/2, (f(r) + f(b_m))/2)$$

that sends everything else to itself, where b_m is the leaf maximizing f on the leaves of T. Notice that $|f'(v') - f(r)| \leq \varepsilon$, and further $f(r) - f(b_m) = f(v) - f(b_m) \leq \delta$, so if $f(b_m) < f(r)$ then this is a η-matching. Otherwise, $f(r) < f(b_m) < f(v')$, and

$$\begin{aligned} |f'(v') - f(b_m)| + |f(b_m) - (f(r) + f(b_m))/2| &\leq |f'(v') - f(b_m)| + |f(b_m) - f(r)| \\ &= f'(v') - f(b_m) + f(b_m) - f(r) \\ &= f'(v') - f(r) \\ &\leq \varepsilon, \end{aligned}$$

so the proposed matching is an ε-matching, and hence a η-matching.

Proceeding now by induction, we now assume the inductive hypothesis holds whenever the v has depth less than $k > 0$. Assume now that the depth of v is k. Hence $v \neq r$, so then v lies in T_j for some j. We define T'_j to be the tree with $N(T'_j) = N(T_j) \cup \{v'\}$ and $E(T'_j) = E(T_j) \cup (v, v')$, and b'_j to be the leaf in T'_j that maximises f' over the leaves of T'_j. Thus

$$\mathrm{TMDg}(T', f') = \mathrm{diag} \sqcup \{(f(r), f(b_i)) : i \neq j\} \sqcup \{(f(r), f'(b'_j))\} \sqcup D'_j \sqcup \bigsqcup_{i \neq j} D_i,$$

where

$$D'_j = \mathrm{TMDg}(T'_j, f') - \{(f(c_i), f'(b'_i))\}$$

For our convenience, we let γ denote the number $\eta - f'(b'_j) + f(b_j)$. Notice that v has depth $k-1$ in T_j. Consider the η-matching ϕ between $\mathrm{TMDg}(T'_j, f')$ and $\mathrm{TMDg}(T_j, f)$ guaranteed to exist by our inductive hypothesis. Since our inductive hypothesis also demands that $\phi((f(c_j), f(b'_j))) = (f(c_j), f(b_j))$, ϕ restricts to a γ-matching between D'_j and D_j.

We can use this matching to construct a η-matching between $\mathrm{TMDg}(T', f')$ to $\mathrm{TMDg}(T, f)$ by sending $(f(r), f'(b'_j))$ to $(f(r), f(b_j))$, D'_j to D_j via ϕ, and everything else to itself. However, this matching may not satisfy our additional inductive assumption that points with right coordinates corresponding to leaves maximising f

are sent to each other. If this is the case, then either b'_j is the leaf maximising f' or b_j is the leaf maximising f, but not both. Since $f'(b'_j) \geq f(b_j)$, it must be the case that b'_j maximises f', but b_j does not maximise f. Let b_m be the leaf that maximises f. In this case, consider the matching ψ:

$$\begin{aligned}(f(r), f(b'_j)) &\longmapsto (f(r), f(b_m))\\ (f(r), f(b_m)) &\longmapsto (f(r), f(b_j))\\ D'_j &\longmapsto D_j,\end{aligned}$$

that sends everything else to itself. Of course, we define the restriction of ψ to D'_j to be ϕ. Since $f'(b'_j) \geq f(b_m) \geq f(b_j)$, we observe

$$\gamma + |f'(b'_j) - f(b_m)| + |f(b_m) - f(b_j)| = \gamma + f'(b'_j) - f(b_j) = \eta.$$

As ϕ is a γ-matching, the fact that ψ is a η-matching satisfying our inductive hypothesis is immediate.

Now suppose f' on T' is not injective. In this case, we can always find an injective g' on $N(T')$ satisfying both that $\|f' - g'\|_\infty \leq \kappa$ for any fixed positive real κ and the hypotheses of the proposition with ε replaced by $\varepsilon + 2\kappa$ and δ replaced by $\delta + 2\kappa$. In fact, this will always be the case whenever we perturb f' at each vertex by a value no more than κ in such a way that the perturbed function is injective. We let g be the restriction of g' to $N(T)$. Hence by Proposition 26 and the fact that g' is injective,

$$\begin{aligned}d_1(\mathrm{TMDg}(T, f), \mathrm{TMDg}(T', f')) \leq\ & d_1(\mathrm{TMDg}(T, f), \mathrm{TMDg}(T, g))\\ &+ d_1(\mathrm{TMDg}(T, g), \mathrm{TMDg}(T', g'))\\ &+ d_1(\mathrm{TMDg}(T', g'), \mathrm{TMDg}(T', f')),\end{aligned}$$

and this is less than or equal to

$$(|L(T)| + |L(T')|)\kappa + 5\kappa + \eta.$$

Considering a sequence of values for κ approaching zero the result is immediate. □

The bound given by this proposition is sharp. We show this with an example in Figure 3.9. Indeed the tree on the right of this figure is produced by a perturbation of the tree on the left of the same type as in the hypothesis of the proposition. The tree on the left has TMDg equal to $(0, K - \delta) \sqcup \mathrm{diag}$ while the tree on the right has TMDg equal to $(0, K + \varepsilon) \sqcup (K, K - \delta) \sqcup \mathrm{diag}$. It is not difficult to see that for sufficiently large K, the optimal matching from the later diagram to the former is the matching which sends $(0, K + \varepsilon)$ to $(0, K - \delta)$, $(K, K - \delta)$ to the point $(K - \delta/2, K - \delta/2)$ on the diagonal, and everything else to itself.

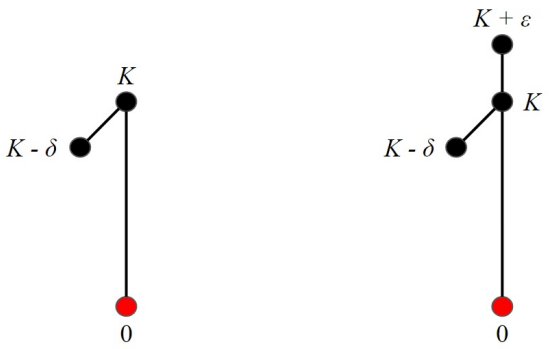

Figure 3.9: A pair of trees with roots in red showing that the bound in Proposition 28 is sharp.

Chapter 4

Inverse problems for persistent homology

In the last chapter we discussed two related ways to compute a barcode describing a function on a tree, the TMD and persistent homology. One reasonable question to ask is how much information is lost in representing tree data by a barcode. For the TMD, this question has been studied already in [52]. For persistent homology, this question fits into a general framework that is an area of active study: suppose we are given a topological space X, a space $\mathcal{F}$ of real valued functions on X, and B, a collection of barcodes in each dimension. What can be said about about the fiber of persistent homology $\mathrm{PH}^{-1}(B)$, the space of functions f in $\mathcal{F}$ such that $\mathrm{PH}(f) = B$? Before returning the the question at hand, we survey some ways in which this question has already been answered below.

For filter functions on a simplicial complex, it was observed in [58] that $\mathrm{PH}^{-1}(B)$ has the structure of a finite polyhedral complex. This polyhedral structure was exploited in [57] to design an algorithm for computing the homology groups of $\mathrm{PH}^{-1}(B)$, and this algorithm was demonstrated on a menagerie of small examples. When $\mathcal{F}$ is the subspace of filter functions determined by their values on vertices, it was shown that every connected component of $\mathrm{PH}^{-1}(B)$ is contractible when X is a simplicial decomposition of the unit interval [32], and homotopy equivalent to a circle when X is instead a simplicial decomposition of the circle [67].

Cases where X is non-discrete have also been investigated. For *Morse-like* continuous functions on the unit interval, the number of path components of $\mathrm{PH}^{-1}(B)$ was computed for generic barcodes [30]. For Morse functions on the 2-sphere $\mathbb{S}^2$ obtained by composing an embedding of $\mathbb{S}^2$ in $\mathbb{R}^3$ with the vertical projection, the tools developed in [19] motivated conjectures on the number of connected components of $\mathrm{PH}^{-1}(B)$. Large families of point clouds in $\mathbb{R}^n$ with identical one-dimensional

persistence under popular filtrations arising from point cloud data have also been identified [84].

In [67], it was observed that when X is a star-like tree and B is the specific barcode that has only one finite interval, then the path connected components of $\mathrm{PH}^{-1}(B)$ are wedges of circles. However, the authors require additionally that X has at least two nodes on each of of the outgoing spokes of X. This suggests that in general, when X is a discrete tree, the fiber is dependent on the chosen triangulation of X until the triangulation is sufficiently fine. To avoid complications of this sort, in Section 4.1 we analyze $\mathrm{PH}^{-1}(B)$ in the case of an arbitrary generic barcode B, for continuous functions on an arbitrary *geometric* tree. Our analysis relies upon the fact that for functions f on X, the persistence map factors in the following way:

$$\mathrm{PH} : f \xmapsto{\mathrm{MT}} \mathcal{T} \longmapsto B.$$

Hence, to characterise the fiber of persistent homology in this setting, we can instead characterise the space of functions with a given merge tree and the space of merge trees that map to B. The main contributions of this section are:

- In Theorem 29, we provide sufficient conditions for a merge tree derived from a function f on a topological space to be a cellular merge tree.
- In Theorem 37, we show that $\mathrm{MT}^{-1}(\mathcal{T})$ is homotopy equivalent to a constrained version of the configuration space of n points on X, denoted $\mathrm{Conf}(X, \mathcal{T})$, where the points must satisfy additional constraints determined by $\mathcal{T}$.
- In Theorem 42, we show that $\mathrm{Conf}(X, \mathcal{T})$, and hence $\mathrm{MT}^{-1}(\mathcal{T})$, is path connected when X has a branch point. We deduce a 1-1 correspondence between path connected components in the fiber $\mathrm{PH}^{-1}(B)$ and non-isomorphic merge trees with barcode B.
- We derive two important consequences of the above results for when X has at least one branch point: (i) in Corollary 53, we find a lower-bound on the distance between the path connected components in $\mathrm{PH}^{-1}(B)$, and (ii) in Corollary 52, we count the number of such components using existing work on merge trees [30, 52].

Section 4.1 is organised as follows.

In Section 4.1.1, we show that a continuous function on a compact connected Hausdorff space has a cellular merge tree if and only if it has finitely many local

minima. Then we show that the interleaving distance between merge trees is a true metric on the subspace of cellular merge trees. Section 4.1.2 is devoted to providing necessary and sufficient conditions for when a function on a geometric tree X has a given cellular merge tree.

In Section 4.1.3 we define the space $\mathrm{Conf}(X, \mathcal{T})$ and a few other intermediary configuration spaces constrained by rules determined by $\mathcal{T}$. By a series of consecutive homotopy equivalences between these configuration spaces, the section culminates in a proof of Theorem 37, showing that $\mathrm{MT}^{-1}(\mathcal{T})$, the space of continuous functions on X with merge tree $\mathcal{T}$, is homotopy equivalent to $\mathrm{Conf}(X, \mathcal{T})$.

Section 4.1.4 exploits this homotopy equivalence to deduce topological properties of $\mathrm{MT}^{-1}(\mathcal{T})$ and $\mathrm{PH}^{-1}(B)$ for generic barcodes B. The main result of this section, Theorem 42, says that $\mathrm{Conf}(X, \mathcal{T})$ and hence $\mathrm{MT}^{-1}(\mathcal{T})$ are connected when X has at least one branch point, i.e. X is not homeomorphic to an interval. This then allows us to provide a lower bound on the distance between any two path connected components in $\mathrm{PH}^{-1}(B)$ (Corollary 52), which depends only on the barcode B. In addition, combining Theorem 42 with existing work enumerating the number of merge trees with a given barcode [30, 52], we deduce in Corollary 53 that

$$\#\pi_0(\mathrm{PH}^{-1}(B)) = \prod_{[b,d)\in B} \#\big\{[b', d') \in B \mid [b, d) \subset [b', d')\big\},$$

whenever X has a branch point (the case where X is an interval has already been studied in [30]). We conclude Section 4.1 by computing the homotopy type of $\mathrm{PH}^{-1}(B)$ via $\mathrm{Conf}(X, \mathcal{T})$ when B has either one or two intervals. When B has one interval, we deduce that $\mathrm{PH}^{-1}(B)$ is contractible (Corollary 54). When B has two intervals, Corollary 55 shows that $\mathrm{PH}^{-1}(B)$ is homotopic to a wedge of

$$-1 + \sum_{v\in N(X)} (\eta(v) - 1)(\eta(v) - 2)$$

circles, where $N(X)$ is the set of vertices in any triangulation of X, and $\eta(v)$ is the degree of vertex v.

In Section 4.2 we show that when X is instead a circle, a different analysis allows us to study the fiber. Namely, when $\mathcal{F}$ is the space of continuous functions on X we show that path components in the fiber are homotopy equivalent to a circle. We attain this result by studying $\mathrm{End}_{\leq}(\mathbb{S}^1)$, the closure of the space of homeomorphisms from the circle to itself, with a given orientation. The result will follow from the observation that path components are exactly the orbits of certain continuous functions f under the action of $\mathrm{End}_{\leq}(\mathbb{S}^1)$ by pre-composition.

In Section 4.3, we see that an entirely analogous approach works for the case where X is a smooth manifold M with or without boundary, provided we choose a nicer space of functions $\mathcal{F}$. Namely we let $\mathcal{F}$ in this section denote the space of Morse functions on M. Suppose we fix a Morse function f such that $B = \mathrm{PH}(f)$ is the associated barcode. Acting on f by pre-composition with the group $\mathcal{D}_{\mathrm{Id}}(M)$ of diffeomorphisms which are isotopic to the identity, and from this action we get an orbit $\mathcal{O}_{\mathrm{Id}}(f)$ inside the space of Morse functions. Our core contribution of this section (Theorem 65) is the equality between this orbit and the path connected component $\mathrm{PH}_f^{-1}(B)$ containing f in the fiber of PH over B:

$$\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f).$$

This result crucially relies on Mather's fibration theorem for smooth mappings [64], which we slightly adapt to the case of Morse functions with equal critical values using results of Cerf [21].

We then put at work the abundant literature about the homotopy type of the orbit $\mathcal{O}_{\mathrm{Id}}(f)$, especially the work of Maksymenko [61]: the mapping $\phi \mapsto f \circ \phi$ in fact defines a locally trivial fibration from $\mathcal{D}_{\mathrm{Id}}(M)$ to the orbit $\mathcal{O}_{\mathrm{Id}}(f)$, with fiber $\mathcal{S}_{\mathrm{Id}}(f) \subseteq \mathcal{D}_{\mathrm{Id}}(M)$ the diffeomorphisms stabilising f, i.e. $f \circ \phi = f$. Hence a long exact sequence links the fiber $\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f)$ to the well-studied diffeomorphism groups of the manifold M. In particular for any compact surface M, we compute $\pi_n(\mathrm{PH}_f^{-1}(B))$ in terms of the homotopy groups of M for $n \geq 3$ (Proposition 69). In fact, if B has distinct interval endpoints we derive the complete homotopy type of $\mathrm{PH}_f^{-1}(B)$ for many compact surfaces (Propositions 70 and 71).

4.1 The fiber of persistent homology on trees

At the start of the chapter, we made the claim that, for functions on a geometric tree X, the persistence map factors into the merge tree operation composed with another operation. Now we explain why this is the case.

As observed in [68], the zero dimensional persistent homology of (X, f), for $f : X \to \mathbb{R}$ a pfd function, is also the zero dimensional persistent homology of $(\mathrm{MT}(f), \pi_f)$. Indeed, the dimension of $H_0(f^{-1}(-\infty, t])$ is exactly the number of path components of $f^{-1}(-\infty, t]$, which, being a subset of a tree, is the number of connected components of the same set. This is exactly $|\pi_f^{-1}(t)|$. However, the set $\pi_f^{-1}(-\infty, t]$ retracts onto $\pi_f^{-1}(t)$ via the homotopy

$$h_u : (x, s) \longmapsto (x, s(1-u) + tu).$$

Hence for each t the map $x \mapsto (x, t)$ induces pointwise isomorphisms between $H_0(f^{-1}(-\infty, t])$ and $H_0(\pi_f^{-1}(-\infty, t])$. Further, these isomorphisms commute with inclusions arising from inequalities $s \leq t$. Hence the persistence modules of (X, f) are completely determined by $\mathrm{MT}(f)$. In other words the map PH factors as a composite of maps

$$\mathrm{PH} : f \longmapsto \mathrm{MT}(f) \longmapsto D.$$

We thus refer to the *persistent homology* on merge trees as the second of these maps, and its value on a given merge tree $\mathrm{MT}(f)$ as the *barcode* of $\mathrm{MT}(f)$.

These observations naturally organise the problem of computing the fiber $\mathrm{PH}^{-1}(B)$ into two consecutive steps: we will first study which functions have a given merge tree, and second, which merge trees have a given barcode.

4.1.1 Tree structure and metric for merge trees

4.1.1.1 When merge trees are cellular merge trees

It is tempting to assume that $(\mathrm{MT}(f), \pi_f)$ is always a cellular merge tree, however this is not the case, even when X and f are very simple.

Example. *If $X = (-\infty, 0]$ and $f(x) = e^x$, $\mathrm{MT}(f)$ is an interval with two open endpoints. This cannot be a cellular merge tree since cellular merge trees have at most one open endpoint.*

The following result gives conditions under which $\mathrm{MT}(f)$ is indeed a cellular merge tree.

Theorem 29. *Let X be a compact connected Hausdorff space and $f : X \to \mathbb{R}$ be a continuous function. Then $(\mathrm{MT}(f), \pi_f)$ is a cellular merge tree if and only if f has finitely many local minima. More precisely, $(\mathrm{MT}(f), \pi_f)$ is isomorphic to the labelled cellular merge tree $(\mathcal{T}, \pi)$ with leaves $l_1, \ldots, l_n$ if and only if the following conditions on f are satisfied:*

1. *The function f has finitely many local minima $X_1, \ldots, X_n$, with values $\pi(l_1), \cdots, \pi(l_n)$.*
2. *For any $1 \leq i < j \leq n$, let t_{ij} denote the infimum of values t where X_i and X_j are in the same connected component of $f^{-1}(-\infty, t]$, i.e. $t_{ij} = \inf\{t \mid (X_i, t) \sim (X_j, t)\}$. Then*

$$t_{ij} = \mathcal{M}(\mathcal{T})_{ij}.$$

Note that the local minima of f are connected but do not need to be singleton sets (see Definition 4). For this reason we use the notation X_i rather than x_i to denote these local minima.

To prove this theorem, we fix the function $f : X \to \mathbb{R}$ and the cellular merge tree $(\mathcal{T}, \pi)$. We will use two lemmas.

Lemma 30. *Let $(y,t) \in \mathrm{MT}(f)$. There exists a local minimum M of f such that (x,t) is also a representative of (y,t) for any x in M. In particular, M can be chosen such that $f(M) \leq f(y)$. As a result, for any t, there are at most as many connected components of $f^{-1}(-\infty, t]$ as local minima of f.*

Proof. Fix y and t. Let Ω_y denote the connected component of y in $f^{-1}(-\infty, t]$. The set Ω_y is a closed subset of a compact set, and therefore is compact. Let $m := \min f|_{\Omega_y}$, and let $M \subseteq \Omega_y$ be a connected component of $f^{-1}(m) \cap \Omega_y$. We claim M is a local minimum of f in X.

Suppose $M' \subseteq X$ is a connected set containing M on which f is never greater than m. Since $M' \subseteq f^{-1}(-\infty, m] \subseteq f^{-1}(-\infty, t]$, we have $M' \subseteq \Omega_y$, and therefore $M' \subseteq f^{-1}(m)$ because $m = \min f|_{\Omega_y}$. Thus $M = M'$. It follows that M is a local minimum of f in X. Meanwhile, since M minimises f on Ω_y, it must be the case that $f(M) \leq f(y)$.

To each connected component $\Omega_y \subseteq f^{-1}(-\infty, t]$, we associate one of its local minima $M \subseteq \Omega_y$, and the last part of the lemma follows. □

Lemma 31. *For any $1 \leq i < j \leq n$, the local minima X_i and X_j are connected in $f^{-1}(-\infty, t_{ij}]$.*

Proof. Suppose the opposite. Let Ω_i be the connected component of X_i in $f^{-1}(-\infty, t_{ij}]$. By Lemma 30 there are finitely many connected components in $f^{-1}(-\infty, t_{ij}]$, so Ω_i is both open and closed in $f^{-1}(-\infty, t_{ij}]$. Thus Ω_i and $\Omega_j := f^{-1}(-\infty, t_{ij}] - \Omega_i$ are disjoint sets in X that are both open and closed in $f^{-1}(-\infty, t_{ij}]$. In particular, Ω_j is nonempty as it contains X_j. Since Ω_i and Ω_j are each closed, disjoint, and nonempty in a closed subspace, they are closed, disjoint, and nonempty subsets of X. Since X is compact and Hausdorff and hence normal, there exists disjoint open sets $U_i \supseteq \Omega_i$ and $U_j \supseteq \Omega_j$ in X. Therefore $X - U_i \cup U_j$ is closed and thus compact in X. If $X - U_i \cup U_j$ is empty, this would imply $X = U_i \cup U_j$ is disconnected, contradicting the hypotheses of the theorem. Hence $X - U_i \cup U_j$ is nonempty and compact and we may let $m := \min f|_{X - U_i \cup U_j}$.

The value m is greater than t_{ij} since $f^{-1}(-\infty, t_{ij}]$ is contained in $\Omega_i \cup \Omega_j$. Therefore by restricting the open cover consisting of U_i and U_j to $f^{-1}(-\infty, t]$ for any $t < m$, we see that Ω_i and Ω_j lie in separate connected components of $f^{-1}(-\infty, t]$. In particular, since $m > t_{ij}$ we may take $t > t_{ij}$, contradicting the hypotheses of the theorem. □

We next turn to the heart of the proof and find the actual tree structure of the merge tree $\mathrm{MT}(f)$.

Proof of Theorem 29. Let $x_1, \ldots, x_n$ be points in each of the local minima of X achieving the values $m_1, \cdots, m_n$ under f, and let:

$$\mathcal{T} := \left(\bigcup_{i=1}^{n} \{x_i\} \times [m_i, +\infty) \right) / \left\{ (x_i, t) \sim (x_j, t) \mid t \geq t_{ij} \right\}.$$

So $\mathcal{T}$ is a disjoint union of the n right-open intervals $\{x_i\} \times [m_i, +\infty)$, and the i^{th} interval is identified with the j^{th} one at and beyond the threshold t_{ij}. In particular, $\mathcal{T}$ is a cellular merge tree.

We have a continuous map $(x_i, t) \in \mathcal{T} \mapsto (x_i, t) \in \mathrm{MT}(f)$ which is well-defined by Lemma 31. We also have a continuous map in the other direction, namely $(x, t) \mapsto (x_i, t)$ where x_i is provided by Lemma 30 to ensure $(x, t) \sim (x_i, t)$ in $\mathrm{epi}(f)$. Therefore $\mathcal{T}$ and $\mathrm{MT}(f)$ are isomorphic. □

Remark 3. *If we defined merge trees using the equivalence relation $(x, t) \sim (y, t)$ whenever x and y are in the same path component of $f^{-1}(-\infty, t]$ instead of the same connected component, Theorem 29 would not hold. For a counterexample, consider the so-called topologist's sine curve*

$$S = \{(0, t) \in \mathbb{R}^2 : t \in [-1, 1]\} \cup \{(x, \sin \frac{1}{x}) \in \mathbb{R}^2 : x \in (0, 1]\}.$$

It is well known that S is closed and connected but has two path components. Taking B to be a closed disk covering S and $f : B \to \mathbb{R}$ the Euclidean distance from S, we see that $f^{-1}(-\infty, 0] = S$ is not path-connected. Thus $\mathrm{MT}(f)$ cannot be a cellular merge tree as it is not Hausdorff: any neighborhood of either point p with $\pi_f(p) = 0$ contains both such points.

Remark 4. *In the conditions of Theorem 29, implying that $\mathrm{MT}(f)$ is a cellular merge tree, a node v is a descendant of a node v' if and only if, when viewed as connected components of $f^{-1}(-\infty, \pi(v)]$ and $f^{-1}(-\infty, \pi(v')]$ respectively, v is a subset of v'.*

4.1.1.2 Interleaving distance on cellular merge trees

Here, we study how the interleaving distance behaves on cellular merge trees, and use our findings to show that two sufficiently close functions with the same finite barcode also have the same (cellular) merge tree.

Proposition 32. *The interleaving distance is a metric on the subspace of cellular merge trees (up to isomorphism).*

Proof. By [68, Lemma 1] the interleaving distance is an extended pseudometric on cellular merge trees, so it remains to show that the interleaving distance is real valued on cellular merge trees and that two cellular merge trees $(\mathcal{T}_1, \pi_1)$ and $(\mathcal{T}_2, \pi_2)$ are isomorphic if they have interleaving distance zero.

By Corollary 4.4 of [42] (which in turn depends upon [90, Theorem 1]), there exists two labellings (possibly with repetitions) of $\mathcal{T}_1$ and $\mathcal{T}_2$ with the same number of indices of leaves, such that:

$$d_I(\mathcal{T}_1, \mathcal{T}_2) = \min_{1 \leq i,j \leq N} |\mathcal{M}(\mathcal{T}_1)_{ij} - \mathcal{M}(\mathcal{T}_2)_{ij}|.$$

Therefore d_I is real-valued. In addition, by Lemma 2.9 of [42], if $\mathcal{M}(\mathcal{T}_1) = \mathcal{M}(\mathcal{T}_2)$ then $\mathcal{T}_1 \cong \mathcal{T}_2$, and so d_I is a metric on cellular merge trees up to isomorphism. □

Proposition 33. *Let X be a compact connected Hausdorff space, and let $f, g : X \to \mathbb{R}$ be continuous pfd functions with the same finite barcode B in dimension zero. Let δ_L be the minimum distance between pairs of non-equal interval left endpoints in B and δ_R be the minimum distance between pairs of non-equal right endpoints. If*

$$\|f - g\|_\infty < \min(\delta_L, \delta_R),$$

then $\mathrm{MT}(f)$ *and* $\mathrm{MT}(g)$ *are isomorphic.*

Proof. Suppose $(\mathcal{T}, \pi)$ is a cellular merge tree with barcode B. Then for any branch point v in $\mathcal{T}$, the map $H_0(\pi^{-1}(-\infty, \pi(v) - \varepsilon]) \to H_0(\pi^{-1}(-\infty, \pi(v)])$ must have a nontrivial kernel for all sufficiently small ε. Therefore, $\pi(v)$ must be the right endpoint of an interval of B. Similarly, for any leaf l in $\mathcal{T}$, the map $H_0(\pi^{-1}(-\infty, \pi(l) - \varepsilon]) \to H_0(\pi^{-1}(-\infty, \pi(l)])$ does not have full image for all sufficiently small ε, and so $\pi(l)$ must be the left endpoint of an interval of B.

By Lemma 5 and Theorem 29, $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are both cellular. By Corollary 4.4 of [42] (which in turn depends upon [90, Theorem 1]), there exist two labellings

(possibly with repetitions) of the cellular merge trees $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ with the same number of indices of leaves, such that:

$$d_I(\mathrm{MT}(f), \mathrm{MT}(g)) = \min_{1 \le i,j \le N} |\mathcal{M}(\mathrm{MT}(f))_{ij} - \mathcal{M}(\mathrm{MT}(g))_{ij}|.$$

The diagonal entries of $\mathcal{M}(\mathrm{MT}(f))$ are projected values of leaves of $\mathrm{MT}(f)$, and non-diagonal entries are projected values of branch points. Similarly for $\mathcal{M}(\mathrm{MT}(g))$. However, from the first paragraph, the projected values of leaves and branch points of $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are the values of left and right interval endpoints of B respectively. Thus if we assume $d_I(\mathrm{MT}(f), \mathrm{MT}(g))$ is positive, then in fact, by our above equation,

$$d_I(\mathrm{MT}(f), \mathrm{MT}(g)) \ge \min(\delta_L, \delta_R).$$

Hence the stability theorem for the interleaving distance [68, Theorem 2] gives us that

$$\min(\delta_L, \delta_R) \le d_I(\mathrm{MT}(f), \mathrm{MT}(g)) \le \|f - g\|_\infty < \min(\delta_L, \delta_R),$$

a contradiction. So $d_I(\mathrm{MT}(f), \mathrm{MT}(g)) = 0$. Since $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are both cellular, Proposition 32 implies that $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are isomorphic. □

4.1.2 Functions on a tree with a given merge tree

In this section we work with cellular merge trees whose branch points are non-degenerate:

Definition 34. *A cellular merge tree $(\mathcal{T}, \pi)$ is* generic *if it is a binary tree (each internal node has two children), and all leaves have distinct projection values.*

In particular if a generic $(\mathcal{T}, \pi)$ has n leaves then it has $(n-1)$ internal nodes. For the rest of the section we fix a generic labelled cellular merge tree $\mathcal{T}$, and for each internal node $v \in \mathcal{T}$, we fix an arbitrary labelling Lv and Rv of its two children. In particular $Lv \preceq v$, where as a reminder this notation means that Lv is a descendant of v, and likewise $Rv \preceq v$. For $1 \le i \le n$ we denote by $m_i := \pi(l_i)$ the value of a leaf. Let X be a geometric tree.

Proposition 35. *Let $f : X \to \mathbb{R}$ be a continuous function with finitely many local minima. Then $\mathrm{MT}(f)$ is isomorphic to $\mathcal{T}$ if and only if both the following conditions are satisfied:*

1. *The function f has n local minima $X_1, \cdots, X_n$ with values $m_1, \cdots, m_n$.*

2. *For any $i \neq j$, the maximum of f restricted on* $\mathrm{ShortPath}(X_i, X_j)$ *equals* $\mathcal{M}(\mathcal{T})_{ij}$.

Proof. By Theorem 29, $\mathrm{MT}(f)$ is a cellular merge tree, and it is isomorphic to $\mathcal{T}$ if and only if:

1. The function f has n local minima $X_1, \cdots, X_n$ with value $m_1, \cdots, m_n$.
2. For any $1 \leq i < j \leq n$, denoting $t_{ij} = \inf\{t \mid (X_i, t) \sim (X_j, t)\}$, we have $t_{ij} = \mathcal{M}(\mathcal{T})_{ij}$.

Furthermore by Lemma 31 we can replace the infimum by a minimum in the definition of t_{ij}. Since X is a geometric tree it is then clear that

$$\min\{t \mid X_i \text{ and } X_j \text{ are connected in } f^{-1}(-\infty, t]\} = \max f|_{\mathrm{ShortPath}(X_i, X_j)} \qquad \square$$

We arrive at our most useful characterisation of functions with a given merge tree. Recall that in this subsection we restrict our attention to cellular merge trees $\mathcal{T}$ which have distinct projection values on their leaves. If $X_1, \ldots, X_n$ are the local minima of a function f with merge tree $\mathcal{T}$, then we have the following two bijections:

$$f : \{X_1, \ldots, X_n\} \to \{m_1, \ldots, m_n\}$$
$$\pi : \{l_1, \ldots, l_n\} \to \{m_1, \ldots, m_n\},$$

where the first bijection is guaranteed by the above proposition (or Theorem 29). As a result we have a bijection between the local minima of f and the leaves of $\mathcal{T}$. We always label local minima so that X_i is mapped to l_i by this bijection (and hence f maps X_i to m_i).

Proposition 36. *Let $f : X \to \mathbb{R}$ be a continuous function with finitely many local minima. Then* $\mathrm{MT}(f)$ *is isomorphic to $\mathcal{T}$ if and only if both the following conditions are satisfied:*

1. *The function f has n local minima $X_1, \cdots, X_n$ with values $m_1, \cdots, m_n$.*
2. *Given a node v (possibly a leaf), let $X_f = (X_1, \ldots, X_m)$ and*

$$\mathrm{Conv}_{X_f}(v) := \mathrm{Conv}\left\{X_i \mid \text{ leaf } l_i \text{ is a descendent of } v \text{ in } \mathrm{MT}(f)\right\} \subseteq X.$$

Then, for any $1 \leq k \leq n-1$, we have:

$$\max\left\{f(x) \mid x \in \mathrm{ShortPath}(\mathrm{Conv}_{X_f}(Lv_k), \mathrm{Conv}_{X_f}(Rv_k))\right\} = \pi(v_k), \qquad (4.1)$$

and the maximum is attained at a unique connected subset Y_k of the shortest path.

Proof. Let f satisfy the conditions of the statement. One is condition 1 of Proposition 35, and an induction on nodes of $\mathcal{T}$ in increasing order of π-value immediately yields condition 2 as well, so $\mathrm{MT}(f) \sim \mathcal{T}$. Conversely, condition 2 in Proposition 35 is equivalent to

$$\forall v_k \in \mathcal{T}, \forall l_i \preceq Lv_k \text{ and } l_j \preceq Rv_k, \max f|_{\mathrm{ShortPath}(X_i,X_j)} = \pi(v_k).$$

From this, an immediate induction yields that, for each node $v_k \in \mathcal{T}$, the value $\max f|_{\mathrm{Conv}_{X_f}(v_k)}$ equals $\pi(v_k)$ and is attained on $\mathrm{ShortPath}(\mathrm{Conv}_{X_f}(Lv_k), \mathrm{Conv}_{X_f}(Rv_k))$.

Assume, seeking contradiction, that the maximum $\pi(v_k)$ is attained at two distinct connected subsets Y_1 and Y_2 of the shortest path. Then, inside $f^{-1}(-\infty, t]$ for $t < \pi(v_k)$, the connected component of elements between Y_1 and Y_2 is distinct from that of elements of $\mathrm{Conv}_{X_f}(Lv_k)$ and $\mathrm{Conv}_{X_f}(Rv_k)$, and at $t = \pi(v_k)$ we thus have three or more connected components of the sublevel-sets of f which are identified, contradicting that $\mathcal{T}$ is generic. □

4.1.3 Retraction of the fiber to configuration space on a tree

Let X be a geometric tree. We metrise spaces of functions on X via the supremum norm. In this section we fix a generic labelled cellular merge tree $\mathcal{T}$ and we analyze the subspace of functions f in the fiber:

$$\mathrm{MT}^{-1}(\mathcal{T}) = \Big\{ f : X \to \mathbb{R},\ \mathrm{MT}(f) = \mathcal{T} \Big\}.$$

We assume without loss of generality that $\mathcal{T}$ has only branch points and leaves as nodes. We will simplify $\mathrm{MT}^{-1}(\mathcal{T})$ by means of a series of homotopy equivalences

$$\mathrm{MT}^{-1}(\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}(X, \mathcal{T}),$$

where the spaces $\mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T})$, $\mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$, and $\mathrm{Conf}(X, \mathcal{T})$ are configuration spaces tracking the local minima and saddles of a function $f \in \mathrm{MT}^{-1}(\mathcal{T})$, detailed hereafter.

Consider $\widetilde{X} = (X_1, \ldots, X_n) \subseteq X^n$. Motivated by the definition we made in Proposition 36, for a node v of $\mathcal{T}$, possibly a leaf, we define:

$$\mathrm{Conv}_{\widetilde{X}}(v) := \mathrm{Conv}\Big\{ X_i \mid \text{ leaf } l_i \text{ is a descendent of } v \Big\} \subseteq X.$$

We define $\mathrm{Conv}_x(v)$ for $x = (x_1, \ldots, x_n) \in X^n$ similarly as

$$\mathrm{Conv}_x(v) := \mathrm{Conv}\Big\{ x_i \mid \text{ leaf } l_i \text{ is a descendent of } v \Big\} \subseteq X.$$

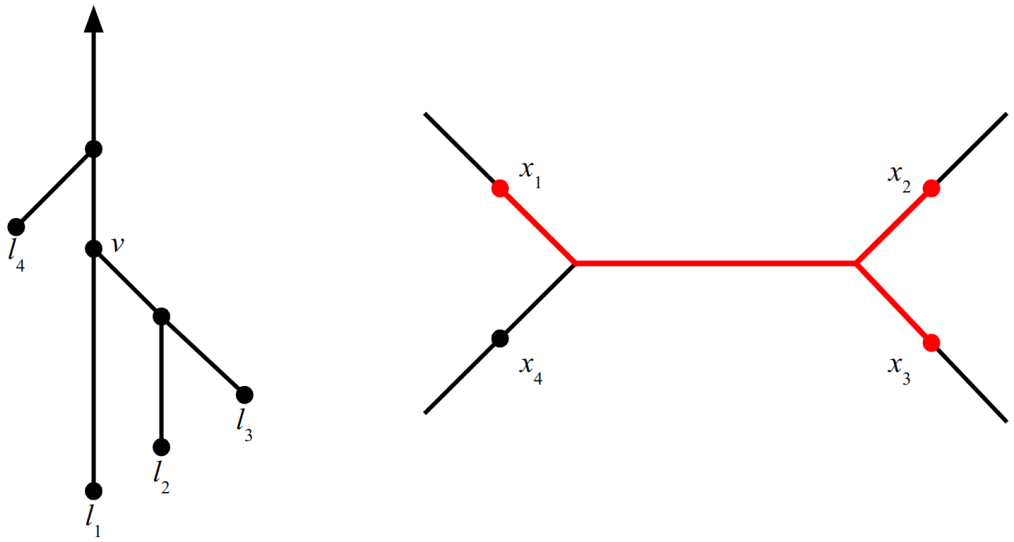

Figure 4.1: (Left) A cellular merge tree $\mathcal{T}$ with four leaves. (Right) An element $x = (x_1, x_2, x_3, x_4)$ of $\mathrm{Conf}_4(X)$, for X a tree. Here, the set $\mathrm{Conv}_x(v)$ associated to a node $v \in \mathcal{T}$ is highlighted.

We illustrate this construction in Figure 4.1.

Denote the usual ordered configuration space on n points by $\mathrm{Conf}_n(X)$. The space $\mathrm{Conf}(X, \mathcal{T})$ is given by

$$\mathrm{Conf}(X, \mathcal{T}) := \Big\{ x = (x_1, \cdots, x_n) \in \mathrm{Conf}_n(X) \mid \\ \mathrm{Conv}_x(v) \cap \mathrm{Conv}_x(v') \neq \emptyset \Rightarrow v \preceq v' \text{ or } v' \preceq v \Big\}.$$

A configuration $(x_1, \cdots, x_n) \in \mathrm{Conf}(X, \mathcal{T})$ should be thought of as the points where a function f with merge tree $\mathcal{T}$ achieves its local minima $m_1, \cdots, m_n$. Because in general a function f with merge tree $\mathcal{T}$ could achieve its minima on arbitrary closed sets rather than points, it is natural to extend $\mathrm{Conf}(X, \mathcal{T})$ to the following configuration space of closed sets:

$$\mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T}) := \Big\{ \widetilde{X} = (X_1, \cdots, X_n) \subseteq X^n \text{ disjoint connected closed sets} \mid \\ \mathrm{Conv}_{\widetilde{X}}(v) \cap \mathrm{Conv}_{\widetilde{X}}(v') \neq \emptyset \Rightarrow v \preceq v' \text{ or } v' \preceq v \Big\}.$$

For v an internal node with children Lv and Rv, we will also consider the following subset of X:

$$\mathrm{ShortPath}_{\widetilde{X}}(v) := \mathrm{ShortPath}\big(\mathrm{Conv}_{\widetilde{X}}(Lv), \mathrm{Conv}_{\widetilde{X}}(Rv)\big).$$

We are now ready introduce our last configuration space $\mathrm{Conf}_{\mathrm{Crit}}(X,\mathcal{T})$ of closed sets where we also record saddles of a function $g \in \mathrm{MT}^{-1}(\mathcal{T})$:

$$\mathrm{Conf}_{\mathrm{Crit}}(X,\mathcal{T}) := \Big\{ (\widetilde{X},\widetilde{Y}) = (X_1,\cdots,X_n,Y_1,\cdots,Y_{n-1}) \subseteq X^{2n-1} \\ \text{disjoint connected closed sets} \mid \\ (X_1,\cdots,X_n) \in \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T}), \\ \forall j,\, Y_j \cong [0,1],\, Y_j \text{ subset of the interior of } \mathrm{ShortPath}_{\widetilde{X}}(v_j) \Big\}.$$

Note that the Hausdorff distance inherited from the ground metric on X induces topologies on configuration spaces of closed sets, hence subset topologies for our spaces $\mathrm{Conf}_{\mathrm{Crit}}(X,\mathcal{T}), \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$ and $\mathrm{Conf}(X,\mathcal{T})$.

Theorem 37. *The spaces* $\mathrm{MT}^{-1}(\mathcal{T}), \mathrm{Conf}_{\mathrm{Crit}}(X,\mathcal{T}), \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$*, and* $\mathrm{Conf}(X,\mathcal{T})$ *are homotopy equivalent.*

The proof of Theorem 37 decomposes as the construction of three consecutive homotopy equivalences:

$$\mathrm{MT}^{-1}(\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Crit}}(X,\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}(X,\mathcal{T}).$$

Proof. For the proof we will fix a metric d on X.

Step 1: $\mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}(X,\mathcal{T})$

We choose an arbitrary leaf $\tau \in X$ as the root of the tree, hence for any connected closed subset $A \subseteq X$, the point $\tau(A)$ which is closest to the root is uniquely defined. We can continuously contract A to $\tau(A)$ with a family $(A_t)_{0\leq t\leq 1}$ of enclosed subsets:

$$A_t := \big\{ x \in A \mid d(a,\tau(A)) \leq (1-t)\mathrm{diam}(A) \big\}.$$

Given a configuration of connected closed sets $(X_1,\cdots,X_n)$, the continuous contraction of each X_i to $\tau(X_i)$ defines a deformation retract from $\mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$ to $\mathrm{Conf}(X,\mathcal{T})$:

$$\mathrm{H} : \big(t,(X_1,\cdots,X_n)\big) \in [0,1]\times\mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T}) \longmapsto ((X_1)_t,\cdots,(X_n)_t) \in \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$$

Indeed, under this map, any $\widetilde{X} \in \mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$ is mapped to an element $\mathrm{H}(t,\widetilde{X})$ that satisfies the condition for being in $\mathrm{Conf}_{\mathrm{Min}}(X,\mathcal{T})$: for any nodes $v,v' \in \mathcal{T}$, if $\mathrm{Conv}_{\mathrm{H}(t,\widetilde{X})}(v) \cap \mathrm{Conv}_{\mathrm{H}(t,\widetilde{X})}(v')$ is nonempty, since $\mathrm{Conv}_{\mathrm{H}(t,\widetilde{X})}(v) \subseteq \mathrm{Conv}_{\widetilde{X}}(v)$ and $\mathrm{Conv}_{\mathrm{H}(t,\widetilde{X})}(v') \subseteq \mathrm{Conv}_{\widetilde{X}}(v')$, then we also have $\mathrm{Conv}_{\widetilde{X}}(v) \cap \mathrm{Conv}_{\widetilde{X}}(v') \neq \emptyset$, and so $v \preceq v'$ or $v' \preceq v$, as desired.

Step 2: $\mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$

We continue to use the root $\tau \in X$. Let $\widetilde{X} = (X_1, \cdots, X_n) \in \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$ be a configuration of closed sets. Given $1 \leq j \leq (n-1)$, we have that $\mathrm{ShortPath}_{\widetilde{X}}(v_j)$ is a closed segment $[a_j, b_j]$ in X, where by convention the extreme a_j is the closest to the root τ.

Let (s, t) be an element of the open standard simplex $\Delta_2 := \{(s, t) \mid 0 < s \leq t < 1\}$. Define

$$Y_j^{s,t} := \left\{ x \in [a_j, b_j] \mid s \leq \frac{d(a_j, x)}{d(a_j, b_j)} \leq t \right\}.$$

Note that, for varying $1 \leq j \leq (n-1)$, the shortest paths $\mathrm{ShortPath}_{\widetilde{X}}(v_j)$ are disjoint from each other. Indeed, let's assume, seeking contradiction, that $\mathrm{ShortPath}_{\widetilde{X}}(v)$ intersects $\mathrm{ShortPath}_{\widetilde{X}}(v')$ for some distinct nodes $v, v' \in \mathcal{T}$. We then have $\mathrm{Conv}_{\widetilde{X}}(v) \cap \mathrm{Conv}_{\widetilde{X}}(v') \neq \emptyset$, and since $\widetilde{X} \in \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$, we can assume without loss of generality that v is a descendant of v'. But then $\mathrm{Conv}_{\widetilde{X}}(v)$ is disjoint from $\mathrm{ShortPath}_{\widetilde{X}}(v')$, and since $\mathrm{ShortPath}_{\widetilde{X}}(v) \subseteq \mathrm{Conv}_{\widetilde{X}}(v)$, we reach the contradiction $\mathrm{ShortPath}_{\widetilde{X}}(v) \cap \mathrm{ShortPath}_{\widetilde{X}}(v') = \emptyset$. Therefore the sets $Y_j^{s_j,t_j}$ are disjoints from each other and from the sets X_i, and we have the homeomorphism:

$$\begin{aligned}[(X_i)_{i=1}^n, (s_j, t_j)_{j=1}^{n-1}] \in \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T}) \times (\Delta_2)^{n-1} \\ \longmapsto [(X_i)_{i=1}^n, (Y_j^{s_j,t_j})_{j=1}^{n-1}] \in \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T}).\end{aligned}$$

The deformation retract of each copy of Δ_2 to a point gives us the homotopy equivalence from $\mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T})$ to $\mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$.

Step 3: $\mathrm{MT}^{-1}(\mathcal{T}) \xrightarrow{\sim} \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T})$

To show that there is a homotopy equivalence between these two spaces, we will need to define a map which sends a function f to its local minima $X_i(f)$ and saddles $Y_j(f)$, see Figure 4.2.

Let $f \in \mathrm{MT}^{-1}(\mathcal{T})$. For $1 \leq i \leq n$, let $X_i(f)$ denote the connected subset of X where f achieves the minimum m_i corresponding to leaf l_i of $\mathcal{T}$. The map $f \mapsto X_i(f)$ is continuous.

Let $v, v' \in \mathcal{T}$ be two nodes such that $A := \mathrm{Conv}_{\widetilde{X}}(v) \cap \mathrm{Conv}_{\widetilde{X}}(v') \neq \emptyset$. We assume without loss of generality that $\pi(v) \leq \pi(v')$ and show that v is a descendant of v', i.e. $v \preceq v'$. The node v, viewed as a connected component in $f^{-1}(-\infty, \pi(v)]$, contains A. Since the connected component of $f^{-1}(-\infty, \pi(v')]$ represented by v' also contains A, viewing nodes once again as connected components of sublevel sets, v

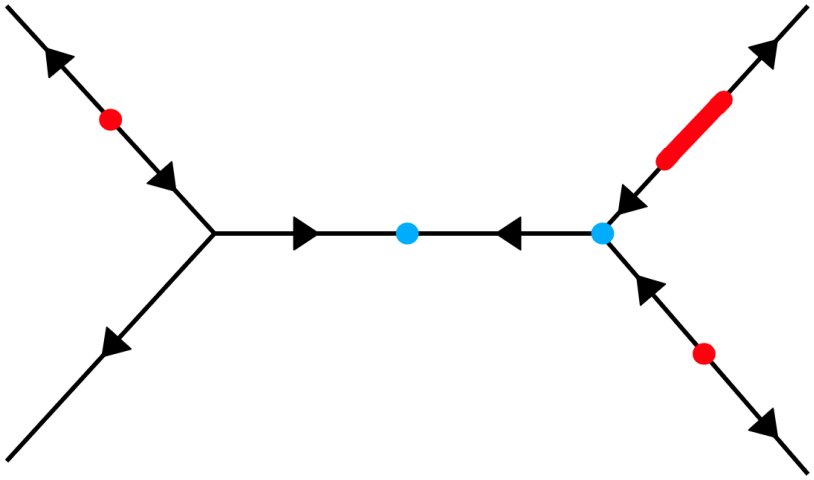

Figure 4.2: The local minima and saddles of a function on a tree. Directions of arrows on the depicted tree X indicate where the function is increasing. The three red regions indicate the locations of local minima $X_i(f)$ while the two blue regions indicate the locations of saddles $Y_j(f)$. Saddles do not need to be local maxima: in this example, one saddle is a local maximum while the other is not.

is a subset of v'. So $v \preceq v'$ (see Remark 4). Therefore $\widetilde{X} := (X_1(f), \ldots, X_n(f)) \in \mathrm{Conf}_{\mathrm{Min}}(X, \mathcal{T})$.

By Proposition 36, for $1 \leq j \leq n-1$, the restriction of f to $\mathrm{ShortPath}_{\widetilde{X}}(v_j)$ attains its maximum $\pi(v_j)$ on a unique connected closed set $Y_j(f)$. Since local minima $X_i(f)$ vary continuously with f, so do the convex hulls $\mathrm{Conv}_{\widetilde{X}}(v_j)$ and the shortest paths $\mathrm{ShortPath}_{\widetilde{X}}(v_j)$ between them, and therefore the maps $f \mapsto Y_j(f)$ are continuous, and so we have defined a continuous map:

$$F : f \in \mathrm{MT}^{-1}(\mathcal{T}) \longmapsto \big(X_1(f), \cdots, X_n(f), Y_1(f), \cdots, Y_{n-1}(f)\big) \in \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T}).$$

To show that F is a homotopy equivalence, we define a map in the other direction. Let $Z = (\widetilde{X}, \widetilde{Y}) = (X_1, \cdots, X_n, Y_1, \cdots, Y_{n-1}) \in \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T})$. We construct a function f_Z by induction on the nodes of $\mathcal{T}$. To begin with, we define $f_Z(X_i) := m_i$. Next, let $v_j \in \mathcal{T}$ be a node such that f_Z is already defined on $\mathrm{Conv}_{\widetilde{X}}(Lv_j)$ and $\mathrm{Conv}_{\widetilde{X}}(Rv_j)$. We extend f_Z on

$$\mathrm{Conv}_{\widetilde{X}}(v_j) = \mathrm{Conv}_{\widetilde{X}}(Lv_j) \cup \mathrm{Conv}_{\widetilde{X}}(Rv_j) \cup \mathrm{ShortPath}_{\widetilde{X}}(v_j),$$

by letting $f_Z(Y_j) := s_j$ and by linear interpolation on the rest of $\mathrm{ShortPath}(v_j)$. At the end of this process, f_Z is defined on the convex hull of all the X_i, outside of which

we let f_Z increase in all directions:

$$\forall x \in X \setminus \mathrm{Conv}(X_1, \cdots, X_n)$$
$$f_Z(x) := f_Z(\mathrm{proj}_{\mathrm{Conv}(X_1,\cdots,X_n)}(x)) + d(x, \mathrm{Conv}(X_1, \cdots, X_n)).$$

By Proposition 36 $\mathrm{MT}(f_Z) = T$ as desired.

This gives us a continuous map:

$$G : Z \in \mathrm{Conf}_{\mathrm{Crit}}(X, \mathcal{T}) \longmapsto f_Z \in \mathrm{MT}^{-1}(\mathcal{T}),$$

and clearly $F \circ G = \mathrm{Id}$. Conversely, by Proposition 36, the straight-line interpolation

$$(t, f) \in [0, 1] \times \mathrm{MT}^{-1}(\mathcal{T}) \longmapsto (1 - t)f + tG \circ F(f)$$

is valued in $\mathrm{MT}^{-1}(\mathcal{T})$, hence it defines a homotopy equivalence $G \circ F \sim \mathrm{Id}$. □

4.1.4 The topology of functions on a tree with a given barcode

Let X be a geometric tree. In this section we fix a barcode B and analyze the space $\mathrm{PH}^{-1}(B)$ of continuous pfd functions with barcode B. For $f \in \mathrm{PH}^{-1}(B)$ we let $\mathrm{PH}_f^{-1}(B) \subseteq \mathrm{PH}^{-1}(B)$ be the connected component of the fiber that contains f. In the rest of the section we assume that B is generic in the following sense:

Definition 38. *A barcode B is* generic *if it is finite and all its interval endpoints are distinct.*

Computing the fiber $\mathrm{PH}^{-1}(B)$ can be solved by computing the fibers of two composite maps:

$$\mathrm{PH} : f \xrightarrow{\mathrm{MT}} \mathcal{T} \longrightarrow B$$

The fiber of the second map is known from [30]: the number of merge trees giving rise to a generic barcode B is finite and computed computed in [30, Theorem 4.8]. Therefore it remains to analyze functions with a given merge tree $\mathcal{T}$.

We first identify $\mathrm{MT}^{-1}(\mathcal{T})$ with a connected component of $\mathrm{PH}^{-1}(B)$ in Subsection 4.1.4.1, and then, in Subsection 4.1.4.2, we derive some topological properties of these components. Before we begin we first record the following two useful properties.

Proposition 39. *If B is a generic barcode and $f \in \mathrm{PH}^{-1}(B)$, then the merge tree $\mathrm{MT}(f)$ is a generic cellular merge tree.*

Proof. Since B is finite, by Lemma 5 f has finitely many local minima, and in turn $(\mathrm{MT}(f), \pi_f)$ is cellular by Theorem 29. Since $\mathrm{PH}(\pi_f) = \mathrm{PH}(f)$ has distinct interval endpoints, no more than two connected subsets can merge at a time in the sublevel-sets of π_f, hence $\mathrm{MT}(f)$ is a binary tree. Similarly, no two leaves of $\mathrm{MT}(f)$ can have the same value of π_f, as this would force $\mathrm{PH}(f)$ to have a repeated left endpoint. □

Proposition 40. *Given a geometric tree $X \neq \emptyset$ and a cellular merge tree $\mathcal{T}$, the fiber $\mathrm{MT}^{-1}(\mathcal{T})$ and the configuration space $\mathrm{Conf}(X, \mathcal{T})$ are non-empty.*

Proof. A function on the unit interval (hence more generally on any non-empty tree X) with merge tree $\mathcal{T}$ can be constructed e.g. as in [30, Proposition 6.8]. Therefore $\mathrm{MT}^{-1}(\mathcal{T}) \neq \emptyset$, and by Theorem 37, $\mathrm{Conf}(X, \mathcal{T}) \neq \emptyset$ as well. □

4.1.4.1 Counting connected components in the fiber over binary trees

Recall that the barcode B is generic. To avoid trivial cases where $\mathrm{PH}^{-1}(B) = \emptyset$, we further assume that B has no interval in degree greater than 0, and only one unbounded interval $[b, \infty)$ in degree 0 that contains all other intervals.

Proposition 41. *Given $\mathcal{T}$ a generic cellular merge tree with barcode B, $\mathrm{MT}^{-1}(\mathcal{T})$ is a non-empty union of connected components in $\mathrm{PH}^{-1}(B)$.*

Proof. $\mathrm{MT}^{-1}(\mathcal{T}) \neq \emptyset$ (Proposition 40) and MT is locally constant on $\mathrm{PH}^{-1}(B)$ (Proposition 33). □

Theorem 42. *Let X be a geometric tree not homeomorphic to the unit interval. Given a generic cellular merge tree $\mathcal{T}$, the fiber $\mathrm{MT}^{-1}(\mathcal{T})$ is nonempty and path-connected. In particular, if $f \in \mathrm{MT}^{-1}(\mathcal{T})$ has barcode B, then $\mathrm{MT}^{-1}(\mathcal{T})$ equals $\mathrm{PH}_f^{-1}(B)$, the path connected component of the fiber $\mathrm{PH}^{-1}(B)$ containing f.*

The proof mainly relies on the following result.

Proposition 43. *Let X be a geometric tree not homeomorphic to the unit interval. Let $\mathcal{T}$ be a generic cellular merge tree. Then $\mathrm{Conf}(X, \mathcal{T})$ is path-connected.*

Before proving the proposition, let us see how it leads to the theorem.

Proof of Theorem 42. From Proposition 40, $\mathrm{MT}^{-1}(\mathcal{T}) \neq \emptyset$, and from Theorem 37, it is homotopy equivalent to $\mathrm{Conf}(X, \mathcal{T})$, which is path-connected by Proposition 43. If $\mathcal{T}$ has barcode B, by Proposition 41, $\mathrm{MT}^{-1}(\mathcal{T})$ is a non-empty union of connected components of $\mathrm{PH}^{-1}(B)$. Therefore it equals exactly one such connected component. □

To prove that $\mathrm{Conf}(X, \mathcal{T})$ is path-connected, we proceed in two steps, each relying on a key lemma. First, we show how to deform a configuration of points into one whose points all lie on a common edge of X. The main result to achieve this step is Lemma 46. Then, given two configurations whose points lie on an edge, we show how to connect them using a branch point of X. We achieve this step with Lemma 51. Throughout, we fix a labelling of the cellular merge tree $\mathcal{T}$.

The following two results, Lemma 44 and Lemma 45, will be used repeatedly during our argument.

Lemma 44. *Let $x = (x_1, \ldots, x_n) \in \mathrm{Conf}(X, \mathcal{T})$, and $x' = (x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n) \in \mathrm{Conf}_n(X)$. Fix P, the image of some path from x_i to y. If $\mathrm{Conv}_x(v)$ does not intersect P for any v that is not an ancestor of l_i, then $x' \in \mathrm{Conf}(X, \mathcal{T})$ and there is a path from x to x' in $\mathrm{Conf}(X, \mathcal{T})$.*

Proof. For $p \in P$, let $x(p) = (x_1, \ldots, x_{i-1}, p, x_{i+1}, \ldots, x_n)$. Let v, v' be nodes of T with neither $v \preceq v'$ nor $v' \preceq v$. Therefore it cannot be the case that both v and v' are ancestors of l_i as ancestors of l_i are totally ordered. If neither v nor v' are ancestors of l_i then

$$\mathrm{Conv}_{x(p)}(v) \cap \mathrm{Conv}_{x(p)}(v') = \mathrm{Conv}_x(v) \cap \mathrm{Conv}_x(v') = \emptyset,$$

with the final equality following from the fact that $x \in \mathrm{Conf}(X, \mathcal{T})$. Lastly, consider the case where either v or v' is an ancestor of l_i, but not both. Without loss of generality assume v' is an ancestor of l_i and v is not. Note that

$$\mathrm{Conv}_{x(p)}(v') \subseteq \mathrm{Conv}_x(v') \cup \mathrm{ShortPath}(x_i, p) \subseteq \mathrm{Conv}_x(v') \cup P.$$

Thus, since $\mathrm{Conv}_x(v) \cap \mathrm{Conv}_x(v') = \emptyset$, we have

$$\mathrm{Conv}_{x(p)}(v) \cap \mathrm{Conv}_{x(p)}(v') \subseteq \mathrm{Conv}_x(v) \cap \big(\mathrm{Conv}_x(v') \cup P\big) = \mathrm{Conv}_x(v) \cap P = \emptyset.$$

Thus the criteria for $x(p)$ to be an element of $\mathrm{Conf}(X, \mathcal{T})$ are satisfied for all $p \in P$. The result follows. □

Lemma 45. *Let b be a branch point or leaf in X with incident edge e, and $x = (x_1, \ldots, x_n) \in \mathrm{Conf}(X, \mathcal{T})$. If $x_i = b$ for some i, then sufficiently short paths from x_i to points y in the interior of e define paths from x to $(x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$ in $\mathrm{Conf}(X, \mathcal{T})$.*

Proof. Let y be sufficiently close to $x_i = b$ that there is no x_j, here $j \neq i$, and no additional branch point in the shortest path from y to x_i. Denote this path by P

and let $x' = (x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$. For $v \in \mathcal{T}$ a node not ancestral to l_i, the set $\mathrm{Conv}_x(v)$ does not intersect P, as doing so would mean it would contain x_i. The result follows by Lemma 44. □

Lemma 46. *Fix any $x = (x_1, \ldots, x_n) \in \mathrm{Conf}(X, \mathcal{T})$. There is an edge e in X such that there is a path from x to $x' = (x'_1, \ldots, x'_n)$ in $\mathrm{Conf}(X, \mathcal{T})$ where every x'_i is in the interior of e.*

Proof. Using Lemma 45, we assume without loss of generality that no point of x is on a branch point or leaf of X. Let $e \in X$ be an edge with some but not all entries of x in its interior. We will build a path to an x' with one more entry in e, giving us the lemma by induction.

By hypothesis, there is an entry of x in at least one of the two path components of X minus the interior of e. Let b denote the endpoint of e in this path component, which we will call C. Let x_j denote the entry of x in the interior of e that is closest to b. We will construct a path from b to some $x_i \in C$ and show that this path lifts to a path in $\mathrm{Conf}(X, \mathcal{T})$, using Lemma 44. The construction of the path is illustrated in Figure 4.3.

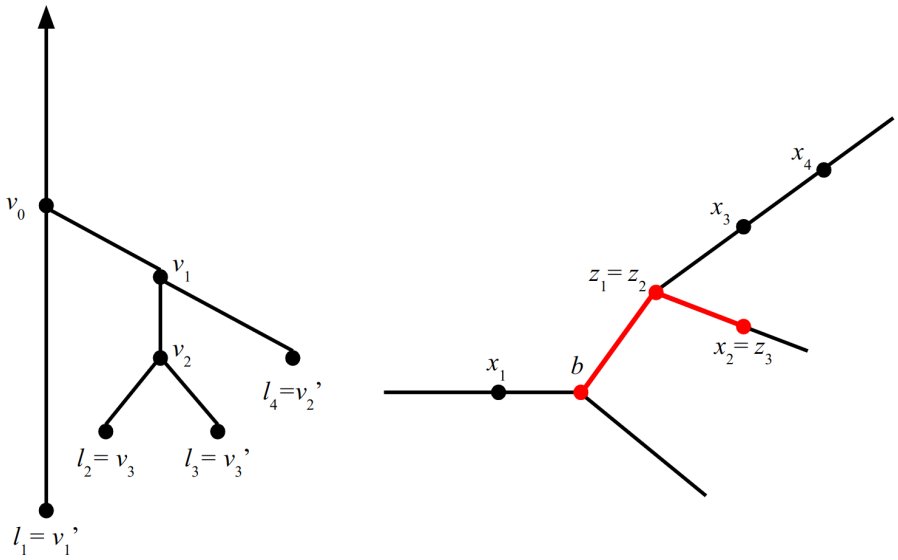

Figure 4.3: A merge tree $\mathcal{T}$ (left) and an example geometric tree X (right) illustrating the main construction of Lemma 46. Highlighted in red is the path constructed from from a point x_i in X to b, which lifts to a path in $\mathrm{Conf}(X, \mathcal{T})$. Here, x_1 plays the role of x_j in the proof. Note that in this example we could have also constructed a path from b to x_3, but not to x_4, because $\mathrm{Conv}_x(v_2)$ interrupts the path from b to $\mathrm{Conv}_x(v'_2)$.

Let $v_0 \in \mathcal{T}$ be the least ancestor of l_j such that $\mathrm{Conv}_x(v_0)$ contains b. Thus v_0 is not l_j itself, so v_0 has two children, v_1 and v_1', exactly one of which is an ancestor of l_j. Without loss of generality, suppose v_1' is an ancestor of l_j and v_1 is not. Let P denote the shortest path from b to $\mathrm{Conv}_x(v_1)$, and let z_1 denote the endpoint of P intersecting $\mathrm{Conv}_x(v_1)$. Since both b and z_1 are in $\mathrm{Conv}_x(v_0)$, so is P. If b is in $\mathrm{Conv}_x(v_1)$, then P consists of the singleton b and does not intersect $\mathrm{Conv}_x(v_1')$. Otherwise, b is in neither $\mathrm{Conv}_x(v_1)$ nor $\mathrm{Conv}_x(v_1')$. Thus b lies on the shortest path between these two sets. Hence P is a subset of this shortest path, and therefore does not intersect $\mathrm{Conv}_x(v_1')$. So P does not intersect $\mathrm{Conv}_x(v_1')$ in either case.

We will continue to augment P until it has reached some x_i. Inductively, assume we have already constructed a path P from b to a point z_{m-1} in $\mathrm{Conv}_x(v_{m-1})$. If v_{m-1} has no children, then $\mathrm{Conv}_x(v_{m-1})$ is a singleton containing x_i for some i, and we are done. Otherwise, let v_m and v_m' denote the children of v_{m-1}.

Either the shortest path from z_{m-1} to $\mathrm{Conv}_x(v_m)$ does not intersect $\mathrm{Conv}_x(v_m')$ or the shortest path from z_{m-1} to $\mathrm{Conv}_x(v_m')$ does not intersect $\mathrm{Conv}_x(v_m)$. Without loss of generality, assume we are in the first case. Since z_{m-1} is in $\mathrm{Conv}_x(v_{m-1})$, so is the shortest path from z_{m-1} to $\mathrm{Conv}_x(v_m)$. We augment P by this path and refer to its augmented endpoint as z_m. Since v_0 has finitely many descendants, this process is guaranteed to terminate eventually, and we will obtain a path P from b to some x_i. Further, since every augmented portion of P lies in $\mathrm{Conv}_x(v_m)$ for some m, all of P lies in $\mathrm{Conv}_x(v_0)$.

Let v be a node of $\mathcal{T}$ not ancestral to l_i. If $\mathrm{Conv}_x(v)$ intersects P then it intersects $\mathrm{Conv}_x(v_0)$, which contains P. Therefore either $v \preceq v_0$ or $v_0 \preceq v$. It cannot be that $v_0 \preceq v$ since l_i is not a descendant of v, so $v \preceq v_0$. Again since l_i is not a descendant of v, v cannot be v_m for any m, so v is a descendant of some v_m'. However, by construction, P does not intersect $\mathrm{Conv}_x(v_m')$ for any m, so it cannot intersect $\mathrm{Conv}_x(v)$ either. Applying Lemma 44 allows us to move x_i onto b. Then applying Lemma 45, we can further move x_i into the interior of e. □

Now that we know we can get all points onto one edge, we want to be able to move groups of points along curves. Let $x = (x_1, \ldots, x_n) \in \mathrm{Conf}_n(X)$, here $\mathrm{Conf}_n(X)$ is the usual ordered configuration space of n points on X, and suppose that there is a subset $Y \subseteq X$ homeomorphic via a map h to the unit interval $[0, 1]$, such that every x_i lies in Y. The coordinates of x thus inherit a total order via the total order of their images under h. Thus for some permutation σ of $\{1, \ldots, n\}$, the inherited total order has form $x_{\sigma(1)} \leq \ldots \leq x_{\sigma(n)}$. We refer to σ as the *h-permutation of x*. Notice that σ only depends on the orientation determined by h.

Lemma 47. *Let $x \in \mathrm{Conf}(X, \mathcal{T})$, $y \in \mathrm{Conf}_n(X)$ and suppose that there is a subset $Y \subseteq X$ homeomorphic via a map h to the unit interval $[0,1]$, such that every coordinate of x and y lies in Y. Then there is a path from x to y in $\mathrm{Conf}(X, \mathcal{T})$ if x and y have the same h-permutation.*

Proof. Let h and h^{-1} induce maps on $\mathrm{Conf}_n(Y)$ and $\mathrm{Conf}_n([0,1])$ by acting componentwise. Consider the path in $\mathrm{Conf}_n(X)$ from x to y

$$\gamma(t) := h^{-1}\big[(1-t)h(x) + th(y)\big]$$

whose image under h linearly interpolates between $h(x)$ and $h(y)$. Thus the h-permutation of $\gamma(t)$ is the same as that of x for all $t \in [0,1]$. Suppose $\mathrm{Conv}_{\gamma(t)}(v) \cap \mathrm{Conv}_{\gamma(t)}(v')$ is nonempty. Then either there exists i, j, and k such that l_i and l_k are descendants of v, l_j is a descendent of v' and $\gamma(t)_i \leq \gamma(t)_j \leq \gamma(t)_k$ or the same is true with the roles of v and v' reversed. Without loss of generality suppose we are in the first case. Thus $x_i \leq x_j \leq x_k$, since the h-permutation is constant along γ. So $\mathrm{Conv}_x(v) \cap \mathrm{Conv}_x(v')$ is also nonempty and $v \preceq v'$ or $v' \preceq v$. Hence γ is a path in $\mathrm{Conf}(X, \mathcal{T})$. □

Lemma 48. *Suppose $x \in \mathrm{Conf}(X, \mathcal{T})$ consists of points on a subset Y of X homeomorphic to the unit interval via a map h. Then there exists $1 \leq i < k \leq n$ such that $l_{\sigma(i)}, \ldots, l_{\sigma(k)}$ are the descendants of v.*

Proof. It suffices to show that the set $\{j \mid l_{\sigma(j)} \preceq v\}$ is convex as a subset of the integers. Consider two leaves $l_{\sigma(i)}$ and $l_{\sigma(k)}$, with $i < k$, that have v as ancestor, and let $j \in \{i, \cdots, k\}$. Since $x_{\sigma(i)} < x_{\sigma(j)} < x_{\sigma(k)}$, we have $x_{\sigma(j)} \in \mathrm{Conv}(x_{\sigma(i)}, x_{\sigma(k)}) \subseteq \mathrm{Conv}_x(v)$. Since $\mathrm{Conv}_x(l_{\sigma(j)}) = \{x_{\sigma(j)}\}$, $\mathrm{Conv}_x(l_{\sigma(j)})$ intersects $\mathrm{Conv}_x(v)$. Therefore either $l_{\sigma(j)} \preceq v$ or $v \preceq l_{\sigma(j)}$. But $l_{\sigma(j)}$ is a leaf so it must be the case that $l_{\sigma(j)} \preceq v$. □

This lemma implies that the homeomorphism h of the edge e where a configuration $x \in \mathrm{Conf}(X, \mathcal{T})$ lies determines, for each internal node $v \in \mathcal{T}$, its left child Lv and right child Rv. Explicitly, we can choose Lv and Rv to be such that $i < j$ whenever $l_{\sigma(i)} \preceq Lv$ and $l_{\sigma(j)} \preceq Rv$. Therefore, the configuration x induces the structure $\mathcal{T}_x$ of a *chiral merge tree* on $\mathcal{T}$, as defined in [30, Definition 5.3]:

Definition 49. *A* chiral merge tree *is a binary cellular merge tree where the two children of any internal node are labelled as either the* left *or* right *child.*

The next lemma will let us alter the chiral merge tree structure assigned to a configuration when X is especially simple.

Lemma 50. *Let X be a geometric starlike tree of degree 3, i.e. X is homeomorphic to three copies of $[0,1]$ identified at 0. Let $x \in \mathrm{Conf}(X,\mathcal{T})$ be a configuration lying on the interior of an edge e of X. Fix a homeomorphism h of e with $[0,1]$. Let $\mathcal{T}_c$ be any chiral merge tree structure on $\mathcal{T}$. Then there exists a path from x to some $y \in \mathrm{Conf}(X,\mathcal{T})$ lying on the interior of e, such that $\mathcal{T}_y = \mathcal{T}_c$.*

Proof. By induction, it is sufficient to prove the case where $\mathcal{T}_c$ differs from $\mathcal{T}_x$ by only inverting the left child Lv and right child Rv of a given node v.

For simplicity, assume that the permutation σ induced by x and h is the identity, thus $x_1 \leq \ldots \leq x_n$. Denote by $l_i, \ldots, l_k$ the descendants of v. There is some $i \leq j < k$ such that $l_i, \ldots, l_j$ are the descendants of Lv while $l_{j+1}, \ldots, l_k$ are the descendants of Rv. The goal is thus to build a path from x to some $y = (y_1, \ldots, y_n)$ in $\mathrm{Conf}(X,\mathcal{T})$, where each entry of y is interior to e, satisfying

$$y_1 \leq \ldots \leq y_{i-1} \leq y_{j+1} \leq \ldots \leq y_k \leq y_i \leq \ldots \leq y_j \leq y_{k+1} \leq \ldots \leq y_n.$$

Let $e_1 = e$, e_2, and e_3 be the edges of X, and b be the branch point of X. Without loss of generality assume that h is such that $h(b) = 0$. We will show that the following sequence of moves from x to y are allowable in $\mathrm{Conf}(X,\mathcal{T})$:

1. Move the 1^{st} through j^{th} coordinates from e_1 into e_2, in that order.
2. Move the $(j+1)^{\text{th}}$ through k^{th} coordinates from e_1 into e_3, in that order.
3. Move the i^{th} through j^{th} coordinates from e_2 into e_1, in reverse order.
4. Move the $(j+1)^{\text{th}}$ through k^{th} coordinates from e_3 into e_1, in reverse order.
5. Move the 1^{st} through i^{th} coordinates from e_2 into e_1, in reverse order.

Figure 4.4 shows a visualisation of this sequence of moves in an example where there are five coordinates.

Moves 1 and 5 can be realised by a path in $\mathrm{Conf}(X,\mathcal{T})$ via Lemma 47 with $Y = e_1 \cup e_2$. For move 2, suppose we have already moved coordinates $j+1$ through $m-1$ into e_3 to attain a configuration $z = (z_1, \ldots, z_n)$. Let P denote the shortest path from z_m to b and $z' = (z_1, \ldots, z_{m-1}, b, z_{m+1}, \ldots, z_n)$. Let v_0 be a node of $\mathcal{T}$ not ancestral to l_m.

Suppose that $\mathrm{Conv}_z(v_0)$ intersects P. Since $l_m \not\leq v_0$, $\mathrm{Conv}_z(v_0)$ can only intersect P at b. This means that $\mathrm{Conv}_{z'}(v_0)$ intersects the interiors of both e_2 and e_3, so there must be points on e_3 in z already, i.e. $m > j+1$. In particular, $\mathrm{Conv}_z(v_0)$

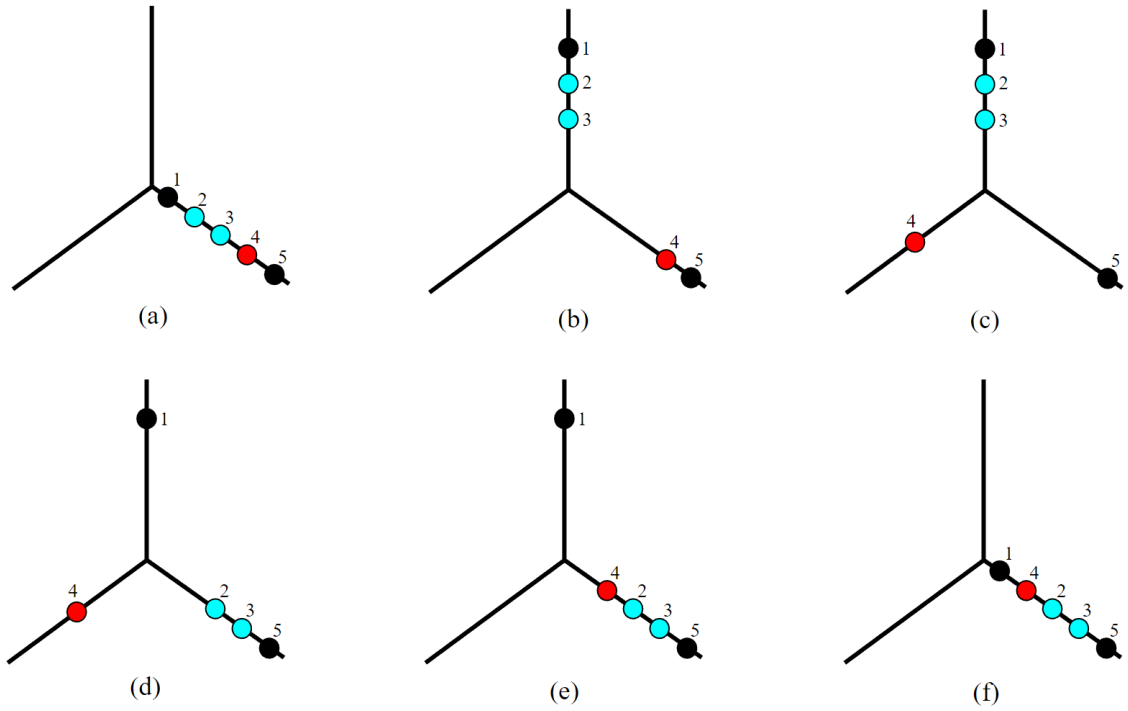

Figure 4.4: A visual representation of the proof of Lemma 50. In the depicted example, $\mathcal{T}$ has five leaves, three of which are descendant from v. One of these nodes, highlighted in red, is moreover a descendant of Lv. The other two descendants of v, highlighted in blue, are descendants of Rv. Panels (a) through (f) show the path used to reconfigure points in the proof.

contains both z_j and z_{m-1}. Since l_j and l_{m-1} are descendants of Lv and Rv respectively, v_0 must be an ancestor of v. Since l_m is a descendant of v, v_0 is an ancestor of l_m, a contradiction. Hence $\mathrm{Conv}_z(v)$ does not intersect intersect P, and so by Lemma 44 we can move the m^{th} coordinate of z to b. Then applying Lemma 45 we can move the m^{th} coordinate of z into the interior of e_3. Induction on m then allows us to complete move 2.

The cases of moves 3 and 4 are handled similarly. □

Lemma 51. *Let X be a geometric starlike tree of degree 3. Then $\mathrm{Conf}(X, \mathcal{T})$ is path-connected.*

Proof. Fix an edge $e \in X$ with an orientation h. By Proposition 40, $\mathrm{Conf}(X, \mathcal{T}) \neq \emptyset$. Let $x, y \in \mathrm{Conf}(X, \mathcal{T})$. Up to applying Lemma 46 and Lemma 47, we can assume that x lies in the interior of e. Similarly we may assume for y lies in the interior of e.

By Lemma 50, we can connect x to a configuration $x' \in \mathrm{Conf}(X, \mathcal{T})$ lying on e and such that $\mathcal{T}_{x'} = \mathcal{T}_y$. In particular x' and y induce the same h-permutation, and therefore, thanks to Lemma 47, there is a path between them in $\mathrm{Conf}(X, \mathcal{T})$. □

Finally we can prove the central proposition of the section, restated below for convenience.

Proposition 43. *Let X be a geometric tree not homeomorphic to the unit interval. Let $\mathcal{T}$ be a generic cellular merge tree. Then $\mathrm{Conf}(X, \mathcal{T})$ is path-connected.*

Proof of Proposition 43. Let l be a leaf of X, and e be the edge incident to l. Since X is connected and not homeomorphic to the unit interval, the other endpoint b of e must be a branch point. Hence, there is a subtree $Y \subseteq X$ which is a geometric starlike tree of degree 3 containing e. By Proposition 40, $\mathrm{Conf}(Y, \mathcal{T}) \neq \emptyset$. Let $y \in \mathrm{Conf}(Y, \mathcal{T}) \subseteq \mathrm{Conf}(X, \mathcal{T})$ be a fixed, target configuration on X.

Let $x = (x_1, \ldots, x_n) \in \mathrm{Conf}(X, \mathcal{T})$. Applying Lemma 46 and then Lemma 47, we find a path in $\mathrm{Conf}(X, \mathcal{T})$ from x to a configuration x' whose points lie in the interior of e. Viewing x' as a configuration in $\mathrm{Conf}(Y, \mathcal{T}) \subseteq \mathrm{Conf}(X, \mathcal{T})$, by Lemma 51 it can be joined to y via a path in $\mathrm{Conf}(Y, \mathcal{T})$, which also defines a path in $\mathrm{Conf}(X, \mathcal{T})$. □

Corollary 52. *Let X be a tree not homeomorphic to the unit interval. Let δ_L be the minimum distance between pairs of non-equal interval left endpoints in B and δ_R be the minimum distance between pairs of non-equal right endpoints. Then the path connected components of the fiber $\mathrm{PH}^{-1}(B)$ are at distance at least $\min(\delta_L, \delta_R)$ from each other. In particular, the path connected components of $\mathrm{PH}^{-1}(B)$ are the connected components of $\mathrm{PH}^{-1}(B)$.*

Proof. Let $f, g : X \to \mathbb{R}$ be functions in distinct path connected components of the fiber $\mathrm{PH}^{-1}(B)$. Then by Theorem 42, their merge trees $\mathrm{MT}(f)$ and $\mathrm{MT}(g)$ are non-isomorphic, and therefore by Proposition 33, we have that $\|f-g\|_\infty \geq \min(\delta_L, \delta_R)$. □

Corollary 53. *Let X be a tree not homeomorphic to the unit interval. The fiber $\mathrm{PH}^{-1}(B)$ has a finite number of connected components given by:*

$$\#\pi_0(\mathrm{PH}^{-1}(B)) = \prod_{[b,d)\in B} \#\big\{[b',d') \in B \mid [b,d) \subset [b',d')\big\}.$$

Proof. This is the number of distinct cellular merge trees with barcode B, see [30, Theorem 4.8]. From Theorem 42, such merge trees are in bijection with path connected components of $\mathrm{PH}^{-1}(B)$, which equal connected components of $\mathrm{PH}^{-1}(B)$ by Corollary 53. □

4.1.4.2 Topology of connected components in the fiber

When there are very few leaves in the cellular merge tree $\mathcal{T}$, and hence very few points in $\mathrm{Conf}(X, \mathcal{T})$, we are able to deduce the homotopy type of the connected components of $\mathrm{MT}^{-1}(\mathcal{T})$, and hence $\mathrm{PH}^{-1}(B)$ for simple barcodes B. The simplest case is when $\mathcal{T}$ has only one leaf.

Corollary 54. *Let X be the geometric realisation of a tree, $\mathcal{T}$ be a cellular merge tree with one leaf and B be the barcode associated to $\mathcal{T}$. Then $\mathrm{MT}^{-1}(\mathcal{T}) = \mathrm{PH}^{-1}(B)$ and both are contractible.*

Proof. The barcode B consists of one infinite interval $[a, \infty)$, where a is the value assigned to the one leaf of $\mathcal{T}$. The only merge tree that can give rise to B is $\mathcal{T}$, by [30, Theorem 4.8]. Hence $\mathrm{MT}^{-1}(\mathcal{T}) = \mathrm{PH}^{-1}(B)$. By Theorem 37,

$$\mathrm{MT}^{-1}(\mathcal{T}) \simeq \mathrm{Conf}(X, \mathcal{T}) = \mathrm{Conf}_1(X) = X,$$

and X is contractible, so $\mathrm{MT}^{-1}(\mathcal{T})$ is contractible. □

If there are only two points in $\mathrm{Conf}(X, \mathcal{T})$, its structure is still fairly simple, and has already been computed up to homotopy. We derive the following as an immediate consequence.

Corollary 55. *Let X be the geometric realisation of a tree with at least one vertex of degree ≥ 3 and $\mathcal{T}$ be a cellular merge tree with exactly two leaves l_1 and l_2. Suppose $\pi(l_1) \neq \pi(l_2)$. Then $\mathrm{MT}^{-1}(\mathcal{T})$ is homotopy equivalent to the wedge sum of*

$$-1 + \sum_{v \in N(X)} (\eta(v) - 1)(\eta(v) - 2)$$

circles, where $N(X)$ denotes the nodes in any cellular decomposition of X and $\eta(v)$ denotes the degree of node v in X. Moreover, denoting by B the barcode arising from the merge tree $\mathcal{T}$, we have $\mathrm{PH}^{-1}(B) = \mathrm{MT}^{-1}(\mathcal{T})$.

Proof. Theorem 37 tells us that $\mathrm{MT}^{-1}(\mathcal{T})$ is homotopy equivalent to $\mathrm{Conf}(X, \mathcal{T})$, here $\mathrm{Conf}(X, \mathcal{T}) = \mathrm{Conf}_2(X)$ is the configuration space of two points on X, whose homotopy type is computed in [38, Theorem 11.1]. The last statement follows from [30, Theorem 4.8]: B is the barcode with two intervals, one finite contained by the other, infinite interval, and therefore $\mathcal{T}$ is the only merge tree giving rise to the barcode B. □

Remark 5. *Let X be the star-like tree made of n edges joined at one vertex, and let B be a barcode as in Corollary 55. Then the corollary tells us that* $\mathrm{PH}^{-1}(B)$ *is homotopy equivalent to a wedge of $n^2 - 3n + 1$ circles.*

In [67], the authors consider the same barcode B and the discrete tree Y obtained from X by inserting an additional vertex in the middle of each edge. Consider the space of functions f on the vertices and edges of Y, where $f(e) = \max(f(v), f(w))$ for any edge $e = (v, w)$. In this scenario, the authors find that $\mathrm{PH}^{-1}(B)$ *is also homotopy equivalent to a wedge of $n^2 - 3n + 1$ circles [67, Theorem 8.1]. This suggests that there may be a general relationship between the fiber of persistent homology of geometric trees and discrete trees with sufficiently fine triangulations.*

4.2 The fiber of persistent homology on manifolds of dimension one

In this section the domain of the persistence map instead consists of continuous maps on the circle, equipped with the supremum norm:

$$\mathrm{PH} : C^0(\mathbb{S}^1, \mathbb{R}) \longrightarrow \mathrm{Bar}^2.$$

Note that in the codomain we record the two barcodes with non-trivial homology, those in degree 0 and 1. In fact the second barcode contains a unique unbounded interval starting at the maximum of the function on the circle. Note that the barcode B does not depend on the underlying choice of field $\mathbb{F}$.

We fix a barcode B with finitely many intervals. When $f = \underline{\mathrm{cst}}$ is constant it forms the fiber by itself over the trivial barcode $B = \mathrm{PH}(f)$ with only two infinite bars $(\underline{\mathrm{cst}}, +\infty)$, one in each degree 0 and 1. Other finite barcodes such

that $\mathrm{PH}^{-1}(B) \neq \emptyset$ have one infinite interval $[b_0, +\infty)$ in degree 0, one infinite interval $[b_1, +\infty)$ with $b_0 < b_1$ in degree 1, finitely many bounded intervals $[a, b)$ in degree 0 with endpoints in $[b_0, b_1]$, and no other intervals. In the rest of this section we assume that B is non-trivial and denote by $(n-1)$, for some $n \geq 1$, its number of bounded intervals in degree 0.

Let $\mathrm{Aut}_{\leq}(\mathbb{S}^1)$ be the space of orientation-preserving homeomorphisms of the circle, and $\mathrm{End}_{\leq}(\mathbb{S}^1)$ be the closure of $\mathrm{Aut}_{\leq}(\mathbb{S}^1)$ in $C^0(\mathbb{S}^1, \mathbb{R})$ in the compact-open topology. Via the identification of the $\mathbb{S}^1$ with $\mathbb{R}/\mathbb{Z}$, $\mathrm{End}_{\leq}(\mathbb{S}^1)$ can be viewed as the space of nondecreasing continuous functions $\phi : \mathbb{R} \to \mathbb{R}$ satisfying $\phi(x) + 1 = \phi(x+1)$. Then $\mathrm{Aut}_{\leq}(\mathbb{S}^1)$ can be viewed as the subspace of functions that are strictly increasing, and hence have a continuous inverse. Given $f \in C^0(\mathbb{S}^1, \mathbb{R})$ we have the pre-composition map $\phi \in \mathrm{End}_{\leq}(\mathbb{S}^1) \mapsto f \circ \phi \in C^0(\mathbb{S}^1, \mathbb{R})$; we denote by $\mathcal{S}_{\mathrm{Id}}(f)$ the stabiliser of f and by $\mathcal{O}_{\mathrm{Id}}(f)$ its orbit under this action, i.e.

$$\mathcal{S}_{\mathrm{Id}}(f) := \{\phi \in \mathrm{End}_{\leq}(\mathbb{S}^1) : f \circ \phi = f\},$$
$$\mathcal{O}_{\mathrm{Id}}(f) := \{g \in C^0(\mathbb{S}^1, \mathbb{R}) : \text{there exists } \phi \in \mathrm{End}_{\leq}(\mathbb{S}^1) \text{ such that } g = f \circ \phi\}.$$

Proposition 56. *The fiber $\mathrm{PH}^{-1}(B)$ has finitely many path connected components. In each such component $\Omega(B)$ there exists some $f_{\Omega(B)} : \mathbb{S}^1 \to \mathbb{R}$ such that:*

$$\Omega(B) = \mathcal{O}_{\mathrm{Id}}(f_{\Omega(B)}),$$

and then $\Omega(B)$ is homeomorphic to the quotient $\mathrm{End}_{\leq}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$, and in particular is homotopy equivalent to $\mathbb{S}^1$.

The component $\Omega(B) \subseteq \mathrm{PH}^{-1}(B)$ in the fiber equals the orbit of a function only for a careful choice of function $f_{\Omega(B)}$: the requirement will be that $f_{\Omega(B)}$ is injective between its consecutive extrema. Nevertheless the fact that the pre-composition map induces a homeomorphism from $\mathrm{End}_{\leq}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ to the orbit $\mathcal{O}_{\mathrm{Id}}(f_{\Omega(B)})$ is reminiscent of the smooth case, and in fact with slightly more work it can be shown that it defines a $\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$-principal bundle. We state without proof the analogous and simpler result for the unit interval $[0, 1]$, which works with or without fixed values on the boundary points 0 and 1. We remark that the results of the previous section can also be used for the case of the interval with no fixed boundary values.

Proposition 57. *For any finite barcode B the fiber $\mathrm{PH}^{-1}(B) \subseteq C^0([0,1], \mathbb{R})$ has finitely many path connected components, each of which is contractible.*

Using a fixed orientation on $\mathbb{S}^1$ and going around starting from the north pole we can order the n minima and n maxima of a non-constant $f \in C^0(\mathbb{S}^1, \mathbb{R})$ into a sequence $\mathrm{Val}(f)$ which we view as an element in $\mathbb{R}^{2n}$:

$$\mathrm{Val}(f) := m_1(f) < M_1(f) > \cdots > m_n(f) < M_n(f).$$

Associated to this sequence we have the sequence of critical sets of f:

$$\mathrm{Seq}(f) : c_1(f),\, d_1(f), \cdots,\, c_n(f),\, d_n(f).$$

Explicitly, each $c_i(f)$ (resp. $d_i(f)$) is a connected component of $f^{-1}(m_i(f))$ (resp. of $f^{-1}(M_i(f))$).

Proposition 58. *Let $f \in \mathrm{PH}^{-1}(B)$. Then f has $2n$ extrema, i.e. $\mathrm{Val}(f) \in \mathbb{R}^{2n}$. In addition, let Γ_n be the group of cyclic permutations on n elements, which acts on $\mathbb{R}^{2n}$ by cyclically permuting the n pairs of entries. Then the path component $\Omega(B)$ in the fiber containing f is made of functions g whose sequence of extrema is the same as that of f up to cyclic permutations, that is:*

$$\Omega(B) = \left\{ g \in C^0(\mathbb{S}^1, \mathbb{R}) \,|\, \mathrm{Val}(g) \in \Gamma_n.\mathrm{Val}(f) \right\} \tag{4.2}$$

We omit the proof of this statement, which can be obtained for example by a simpler version of the argument in the third step of the proof of Theorem 37. So if $\Omega(B)$ is a path component in the fiber, we can pick the following simple function $f_{\Omega(B)}$ in $\Omega(B)$, whose critical sets and extrema are denoted by c_i, d_i, m_i, M_i for simplicity: the critical sets c_i and d_i are singletons arranged on the regular $2n$-gon in $\mathbb{S}^1$ and on each circular arc $[c_i, d_i]$, $f_{\Omega(B)}$ restricts to the linear homeomorphism to $[m_i, M_i]$.

Proposition 59. *Let $\Omega(B) \subseteq \mathrm{PH}^{-1}(B)$ be a path component in the fiber. Then the pre-composition map $\phi \mapsto f_{\Omega(B)} \circ \phi$ induces a homeomorphism from $\mathrm{End}_{\leq}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ to $\Omega(B)$.*

Proof. The map $\phi \in \mathrm{End}_{\leq}(\mathbb{S}^1) \mapsto f_{\Omega(B)} \circ \phi \in \Omega(B)$ is well-defined, i.e. $\mathrm{PH}(f \circ \phi) = \mathrm{PH}(f) = B$. This is because a homeomorphism $\phi \in \mathrm{Aut}_{\leq}(\mathbb{S}^1)$ restricts to a homeomorphism between the sub level-sets of $f_{\Omega(B)} \circ \phi$ and those of $f_{\Omega(B)}$, hence it induces an isomorphism of persistent homology modules and the equality of barcodes $\mathrm{PH}(f_{\Omega(B)} \circ \phi) = \mathrm{PH}(f_{\Omega(B)})$, which holds as well for any $\phi \in \mathrm{End}_{\leq}(\mathbb{S}^1)$ by the stability theorem.

Let $f \in \Omega(B)$. From Proposition 58 there are cyclic permutations $\pi \in \Gamma_n$ such that $\mathrm{Val}(f) = \pi.\mathrm{Val}(f_{\Omega(B)})$. For each such permutation π there is a unique map $\phi^{f,\pi}$

satisfying both $f_{\Omega(B)} \circ \phi^{f,\pi} = f$ and $\phi^{f,\pi}(c_i(f)) = c_{\pi(i)}$ (and $\phi^{f,\pi}(d_i(f)) = d_{\pi(i)}$): It is defined on each circular arc $[c_i(f), d_i(f)]$ by

$$\phi^{f,\pi}|_{[c_i(f),d_i(f)]} := [(f_{\Omega(B)})|_{[c_{\pi(i)},d_{\pi(i)}]}]^{-1} \circ f|_{[c_i(f),d_i(f)]}, \tag{4.3}$$

and similarly on circular arcs $[d_{i-1}(f), c_i(f)]$. In particular for $f = f_{\Omega(B)}$ the set of such $\phi^{f,\pi}$ equals the group $\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ of stabilisers. Therefore $\phi \mapsto f_{\Omega(B)} \circ \phi$ descends to a continuous bijection from $\mathrm{End}_{\leq}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ to $\Omega(B)$.

Finally we show that the inverse is continuous. Let $f \in \Omega(B)$ and $\phi^{f,\pi}$ as in Equation (4.3). Up to pre-composing f by a suitable homeomorphism the north pole does not belong to any extremal set $c_i(f), d_i(f)$. Consequently, for g in a small neighborhood $\mathcal{U} \subseteq \Omega(B)$ around f, we also have $\mathrm{Val}(g) = \pi.\mathrm{Val}(f_{\Omega(B)})$, hence we can define $\phi^{g,\pi} \in \mathrm{End}_{\leq}(\mathbb{S}^1)$ like in Equation (4.3) and then $f_{\Omega(B)} \circ \phi^{g,\pi} = g$. Hence the map $g \in \mathcal{U} \longmapsto \phi^{g,\pi} \in \mathrm{End}_{\leq}(\mathbb{S}^1)$ is a local section, whose continuity is a consequence of the fact that on each circular arc $[c_i, d_i]$ the linear restriction $(f_{\Omega(B)})|_{[c_i,d_i]}$ and its inverse are Lipschitz, and of the fact that the maximal distance from points in the critical sets $c_i(g), d_i(g)$ to the critical sets $c_i(f), d_i(f)$ of f can be continuously tracked in a sufficiently small neighborhood $\mathcal{U}$ of f. □

Proof of Proposition 56. From Proposition 59 the pre-composition map $\phi \mapsto f_{\Omega(B)} \circ \phi$ induces a homeomorphism from $\mathrm{End}_{\leq}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ to the path connected component $\Omega(B)$. Besides it is well-known that $\mathrm{End}_{\leq}(\mathbb{S}^1)$ deformation retracts to the group $\mathrm{SO}(2) \cong \mathbb{S}^1$ of orientation preserving rotations.[1] Recall that $f_{\Omega(B)}$ is a piecewise linear interpolation between extremal values arranged on a regular $2n$-gon, therefore its stabiliser $\mathcal{S}_{\mathrm{Id}}(f_{\Omega(B)})$ is a finite subgroup of $\mathrm{SO}(2)$ which is preserved under the deformation retraction. Hence $\Omega(B)$ is homotopy equivalent to the quotient of $\mathrm{SO}(2) \cong \mathbb{S}^1$ by a finite subgroup, so it is in fact homotopy equivalent to $\mathbb{S}^1$. □

4.3 The fiber for Morse functions

In this section we let the underlying space X be rather more complicated, a smooth compact manifold M with or without boundary. It turns out that at a high level, the same arguments used to study the space $\mathrm{PH}^{-1}(B)$ for functions on one-dimensional manifolds generalise to higher dimensional manifolds, provided we restrict ourselves to a nicer space of functions on M. Namely, instead of continuous functions, we restrict our attention to the space of Morse functions on M, which we denote by

[1] For instance the deformation retract of $\mathrm{Aut}_{\leq}(\mathbb{S}^1)$ of [46, Theorem 1.1.2] extends to $\mathrm{End}_{\leq}(\mathbb{S}^1)$.

Morse(M). Given another smooth manifold N, we denote by $C^\infty(M,N)$ the space of smooth maps from M to N equipped with the weak C^∞ topology (see for example Chapter 2.1 of [48]). This induces a topology on Morse(M) via the inclusion Morse(M) $\to C^\infty(M,\mathbb{R})$. We let $\mathcal{D}(M)$ denote the subspace of $C^\infty(M,M)$ consisting of diffeomorphisms, and let $\mathcal{D}_{\mathrm{Id}}(M) \subseteq \mathcal{D}(M)$ denote the subspace of diffeomorphisms homotopic through $\mathcal{D}(M)$ to the identity map, i.e. the path component of $\mathcal{D}(M)$ containing the identity map. As in the previous section we have a precomposition map

$$\begin{aligned} C^\infty(M,\mathbb{R}) \times \mathcal{D}(M) &\to C^\infty(M,\mathbb{R}) \\ (f,\phi) &\mapsto f \circ \phi. \end{aligned}$$

Also as in the previous section, the precomposition map can be used to define several spaces of interest:

$$\begin{aligned} \mathcal{S}_{\mathrm{Id}}(f) &:= \{\phi \in \mathcal{D}_{\mathrm{Id}}(M) : f \circ \phi = f\}, \\ \mathcal{O}_{\mathrm{Id}}(f) &:= \{g \in C^\infty(M,\mathbb{R}) : \text{there exists } \phi \in \mathcal{D}_{\mathrm{Id}}(M) \text{ such that } g = f \circ \phi\} \\ \mathcal{O}(f) &:= \{g \in C^\infty(M,\mathbb{R}) : \text{there exists } \phi \in \mathcal{D}(M) \text{ such that } g = f \circ \phi\}. \end{aligned}$$

We remark here that for a Morse function f, the choice of field $\mathbb{F}$ underlying the persistence map PH has an effect on the resulting barcode B. However, in this section the primary object of study will be $\mathrm{PH}_f^{-1}(B)$, the path component of the space of functions with barcode B containing f. The main theorem of this section, Theorem 65, says that $\mathrm{PH}_f^{-1}(B)$ is exactly the space $\mathcal{O}_{\mathrm{Id}}(f)$. As a consequence the space $\mathrm{PH}_f^{-1}(B)$ does not depend on the choice of $\mathbb{F}$.

4.3.1 Stability of Morse Functions

The stability result we wish to show here is that if f is a Morse function, then any sufficiently close Morse function with the same critical values can be written as $f \circ \phi$ for some $\phi \in \mathcal{D}_{\mathrm{Id}}(M)$, with ϕ continuously determined. To express this key result, we rely on the notion of local cross-sections defined below.

Definition 60. *Let $\mathcal{G}$ be a topological group acting on a topological space X. Given $x_0 \in X$, a* local cross-section *for the action of $\mathcal{G}$ on X at x_0 is a continuous map $s : \mathcal{U} \to \mathcal{G}$ defined on an open neighborhood $\mathcal{U}$ of x_0 satisfying:*

$$\forall x \in \mathcal{U},\ s(x) \cdot x_0 = x.$$

We say that the action of $\mathcal{G}$ on X admits local cross-sections *if it does so at each point.*

Up to replacing $s(x)$ by $s(x) \cdot s(x_0)^{-1}$ in the above definition, we can assume that $s(x_0)$ is the identity element.

Remark 6. *It is well-known that, if X admits local cross-sections, then any $\mathcal{G}$-equivariant map from a $\mathcal{G}$-space to X is locally trivial, see e.g. [72, Theorem A].*

The main result of this subsection adapts Mather's stability of smooth mappings [64] to the case of Morse functions with equal critical points by combining results of Cerf:

Proposition 61. *Let $f \in \mathrm{Morse}(M)$ and $\mathrm{Morse}_f(M)$ be the subspace of Morse functions with the same critical values as f. Then the action of $\mathcal{D}_{\mathrm{Id}}(M)$ on $\mathrm{Morse}_f(M)$ admits local cross-sections.*

The first result of Cerf we use is essentially a version of Proposition 61 restricted to the space $\mathrm{Morse}_f(M; \mathrm{Crit}(f), \partial M)$ of Morse functions with the same critical points and critical values as f and the same value and derivatives of any order as f on ∂M.

Proposition 62 ([21], Appendix, §1, Proposition 1)**.** *Let $f : M \to \mathbb{R}$ be a Morse function. Let $\mathcal{G}$ be the subspace of diffeomorphisms fixing the critical points of f and ∂M, and which have the same value and derivatives of every order as the identity on ∂M. Then the action of $\mathcal{G}$ on $\mathrm{Morse}_f(M; \mathrm{Crit}(f), \partial M)$ admits local cross-sections.*

To enhance this result to Morse functions with critical points and derivatives at the boundary allowed to vary, we need the reformulation of a result of Cerf [20, Theorem 5] for the C^∞ case given in [49].

Proposition 63. *Let $N \subseteq M$ be a compact submanifold. Suppose that:*

(a) either N has no boundary and does not intersect the boundary ∂M;

(b) or N is a closed collar neighborhood of a boundary component ∂M_j.

Let L be a compact neighborhood of N. Let $\mathrm{Emb}_L(N, M)$ be the space of embeddings $j : N \hookrightarrow M$ such that $j(N) \subseteq L$, $j^{-1}(\partial M) = N \cap \partial M$, and that restricts to the identity on $N \cap \partial M$. Denote by $\mathcal{D}_L(M)$ the space of diffeomorphisms inducing the identity on $(M \setminus L)$. Then the action of $\mathcal{D}_L(M)$ on $\mathrm{Emb}_L(N, M)$ admits local cross-sections.

Proof. This is a direct consequence of Theorem 3.1 in [49]. Note that, strictly speaking, embeddings j as defined in [49, Definition 2.5] are required to admit an extension to a diffeomorphism of M. However this assumption is unnecessary for their Theorem 3.1 and so we omit it. □

Proof of Proposition 61. Since $g \in \mathrm{Morse}_f(M)$ implies that $\mathrm{Morse}_g(M) = \mathrm{Morse}_f(M)$, it is enough to construct a local cross-section at f. By Proposition 62, there exists a local cross-section $g \in \mathcal{U}_f \mapsto \phi_g \in \mathcal{D}_{\mathrm{Id}}(M)$ defined on a neighborhood $\mathcal{U}_f$ of f in $\mathrm{Morse}_f(M; \mathrm{Crit}(f), \partial M)$, the space of functions with the same critical points $p_1, \cdots, p_n$ and critical values as f, and with the same value and derivatives $\partial^k f$ as f on the boundary ∂M. Therefore

$$\forall g \in \mathcal{U}_f, \;\; g = f \circ \phi_g. \tag{4.4}$$

For the general case where critical points and derivatives on the boundary are allowed to vary we simply find a diffeomorphism sending them back to $p_1, \cdots, p_n$ and $\partial^k f$ and apply the above result.

Namely, from item (a) of Proposition 63, we can find disjoint compact neighborhoods $U_1, \cdots, U_n$ of $p_1, \cdots, p_n$ and continuously associate to $(p'_1, \cdots, p'_n) \in U_1 \times \cdots \times U_n$ a diffeomorphism $\psi_{(p'_1, \cdots, p'_n)} \in \mathcal{D}_{\mathrm{Id}}(M)$ such that $\psi_{(p_1, \cdots, p_n)} = \mathrm{Id}$ and:

$$\forall (p'_1, \cdots, p'_n) \in U_1 \times \cdots \times U_n, \;\; \forall 1 \leqslant i \leqslant n, \;\; \psi_{(p'_1, \cdots, p'_n)}(p_i) = p'_i.$$

Let $\mathcal{U} \subseteq \mathrm{Morse}_f(M)$ be a neighborhood of f for which any $g \in \mathcal{U}$ has critical points $\mathrm{Crit}(g)$ in $U_1, \cdots, U_n$. In particular in this case $g \circ \psi_{\mathrm{Crit}(g)}$ has the same critical points $p_1, \cdots, p_n$ as f, so it remains to deal with the boundary ∂M.

Let ∂M_j be a boundary component. By flowing along the normalised gradient of f (or its inverse) from the boundary ∂M_j we get a compact collar $V_j \cong \partial M_j \times [0, \alpha]$ that is adapted to f in the sense that $f(x,t) = a_j \pm t$, w.l.o.g. $f(x,t) = a_j + t$. For g in a small neighborhood $\mathcal{V} \subseteq \mathrm{Morse}_f(M)$ of f, we have $g(\partial M_j \times [0, \frac{\alpha}{2}]) \subseteq [a_j, a_j + \alpha]$. Therefore, after potentially shrinking $\mathcal{V}$, we can continuously associate to $g \in \mathcal{V}$ the embedding $\iota_g : (x,t) \in \partial M_j \times [0, \frac{\alpha}{2}] \mapsto (x, g(x,t) - a_j) \in V_j$ that preserves ∂M_j and satisfies $g = f \circ \iota_g$ on $\partial M_j \times [0, \frac{\alpha}{2}]$. Hence by using item (b) of Proposition 63, up to shrinking $\mathcal{V}$, we can extend ι_g to a diffeomorphism χ_g that induces the identity outside V_j. By repeating this process for each boundary component ∂M_j, we can continuously associate to $g \in \mathcal{V}$ a diffeomorphism χ_g such that g and $f \circ \chi_g$ agrees on a closed collar neighborhood of the boundary.

By reducing the neighorhoods U_i and V_j to avoid overlaps, we have that for any g in $\mathcal{U} \cap \mathcal{V}$ the Morse function $g \circ \psi_{\mathrm{Crit}(g)} \circ \chi^{-1}_{g \circ \psi_{\mathrm{Crit}(g)}}$ has the same critical points as f and agrees with f on a neighborhood of the boundary ∂M, in particular it belongs to $\mathrm{Morse}_f(M; \mathrm{Crit}(f), \partial M)$. Hence by Eq. (4.4), up to shrinking $\mathcal{U} \cap \mathcal{V}$, we have that

for all $g \in \mathcal{U} \cap \mathcal{V}$,

$$\begin{aligned} g &= (g \circ \psi_{\mathrm{Crit}(g)} \circ \chi^{-1}_{g\circ\psi_{\mathrm{Crit}(g)}}) \circ \chi_{g\circ\psi_{\mathrm{Crit}(g)}} \circ \psi^{-1}_{\mathrm{Crit}(g)} \\ &= f \circ \phi_{g\circ\psi_{\mathrm{Crit}(g)}\circ\chi^{-1}_{g\circ\psi_{\mathrm{Crit}(g)}}} \circ \chi_{g\circ\psi_{\mathrm{Crit}(g)}} \circ \psi^{-1}_{\mathrm{Crit}(g)}, \end{aligned}$$

hence the local cross-section:

$$s : g \in \mathcal{U} \cap \mathcal{V} \longmapsto \phi_{g\circ\psi_{\mathrm{Crit}(g)}\circ\chi^{-1}_{g\circ\psi_{\mathrm{Crit}(g)}}} \circ \chi_{g\circ\psi_{\mathrm{Crit}(g)}} \circ \psi^{-1}_{\mathrm{Crit}(g)} \in \mathcal{D}_{\mathrm{Id}}(M). \qquad \square$$

Corollary 64. *Let $(f_t)_{t\in[0,1]} \subseteq \mathrm{Morse}(M)$ be a path of Morse functions with the same critical values. Then there exists $\phi \in \mathcal{D}_{\mathrm{Id}}(M)$ an* Id*-isotopic diffeomorphism such that $f_1 = f_0 \circ \phi$.*

Proof. By Proposition 61, each $t \in [0,1]$ has a neighborhood $I_t \subseteq [0,1]$ such that f_h can be written $f_h = f_t \circ \phi_h$ whenever $h \in I_t$. By compactness $[0,1]$ is covered by finitely many such intervals. Therefore f_1 equals $f_0 \circ \phi$, where ϕ is a finite composition of diffeomorphisms in $\mathcal{D}_{\mathrm{Id}}(M)$ hence is itself in $\mathcal{D}_{\mathrm{Id}}(M)$. $\square$

4.3.2 Covering the fiber with isotopies

As usual, a Morse function determines a sublevel set filtration and hence a barcode for every dimension. Here, if we let d be the dimension of the underlying manifold M, then there are $d+1$ barcodes of interest. Hence in this setting we may view the persistence map in the following way:

$$\mathrm{PH} : f \in \mathrm{Morse}(M) \longmapsto [\mathrm{PH}_0(f), \cdots, \mathrm{PH}_d(f)] \in \mathrm{Bar}^{d+1}.$$

We assume that Bar is equipped with its natural bottleneck distance which turns PH into a continuous map by the stability theorem. Given a barcode B and a Morse function $f \in \mathrm{Morse}(M)$ such that $\mathrm{PH}(f) = B$, we denote by $\mathrm{PH}_f^{-1}(B)$ the path connected component of the fiber $\mathrm{PH}^{-1}(B) \subseteq \mathrm{Morse}(M)$ containing f.

Theorem 65. *Let B be a barcode and $f \in \mathrm{PH}^{-1}(B)$. Then $\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f)$.*

Proof. Let $(\phi_t)_{0\le t\le 1}$ be a path in $\mathcal{D}_{\mathrm{Id}}(M)$. Each ϕ_t restricts to a homeomorphism between the sub level-sets of $f \circ \phi_t$ and f, hence it induces an isomorphism between the associated persistent homology modules. In turn $\mathrm{PH}(f \circ \phi_t) = \mathrm{PH}(f)$, so that $(f \circ \phi_t)_{0\le t\le 1}$ is a path in the fiber $\mathrm{PH}_f^{-1}(B)$, which implies $\mathcal{O}_{\mathrm{Id}}(f) \subseteq \mathrm{PH}_f^{-1}(B)$.

Conversely let $g \in \mathrm{PH}_f^{-1}(B)$ and let $(f_t)_{0\le t\le 1}$ be a path in the fiber $\mathrm{PH}_f^{-1}(B)$ joining f to g, thus $\mathrm{PH}(f_t) = B$ for each t. By Proposition 11, each f_t has the same critical values as f, because the barcode $\mathrm{PH}(f_t) = B$ is constant. By Corollary 64 there exists an Id-isotopic diffeomorphism ϕ such that $g = f \circ \phi$. Consequently $\mathrm{PH}_f^{-1}(B) \subseteq \mathcal{O}_{\mathrm{Id}}(f)$. $\square$

4.3.3 Topological properties of the fiber

We derive direct consequences of Theorem 65 combined with the extensive study of $\mathcal{O}_{\mathrm{Id}}(f)$ by Maksymenko [61]. Strictly speaking, it is $\mathcal{O}_f(f)$, the path component of $\mathcal{O}(f)$ containing f, whose properties are studied in [61][2]. However, there is an obvious inclusion $\mathcal{O}_{\mathrm{Id}}(f) \subseteq \mathcal{O}_f(f)$, and the reverse inclusion holds as well by Corollary 64. Therefore $\mathcal{O}_f(f) = \mathcal{O}_{\mathrm{Id}}(f)$.

Proposition 66. *Assume that M is connected. Let B be a barcode and $f \in \mathrm{PH}^{-1}(B)$. Then the action of $\mathcal{D}_{\mathrm{Id}}(M)$ on $\mathrm{PH}_f^{-1}(B)$ defines a locally trivial principal $\mathcal{S}_{\mathrm{Id}}(f)$-fibration.*

Proof. From [61, Theorem 2.1, (2)] the action of diffeomorphisms $\mathcal{D}(M)$ on $\mathcal{O}(f)$ defines a locally trivial principal fibration with fiber the diffeomorphisms ϕ satisfying $f \circ \phi = f$. Restricting to the action of $\mathcal{D}_{\mathrm{Id}}(M)$ on $\mathcal{O}_{\mathrm{Id}}(f)$ defines a locally trivial principal $\mathcal{S}_{\mathrm{Id}}(f)$-fibration, and $\mathrm{PH}_f^{-1}(B)$ equals $\mathcal{O}_{\mathrm{Id}}(f)$ by Theorem 65. □

Remark 7. *The principal bundle $\mathcal{S}_{\mathrm{Id}}(f) \to \mathcal{D}_{\mathrm{Id}}(M) \to \mathrm{PH}_f^{-1}(B)$ has computationally useful and direct implications. First, it is a locally trivial fibration hence it induces a homotopy long exact sequence:*

$$\cdots \to \pi_n(\mathcal{S}_{\mathrm{Id}}(f)) \to \pi_n(\mathcal{D}_{\mathrm{Id}}(M)) \to \pi_n(\mathrm{PH}_f^{-1}(B)) \to \pi_{n-1}(\mathcal{S}_{\mathrm{Id}}(f)) \to \cdots \to \pi_0(\mathcal{D}_{\mathrm{Id}}(M)).$$

Second, we have the homeomorphism:

$$\mathrm{PH}_f^{-1}(B) \cong \mathcal{D}_{\mathrm{Id}}(M)/\mathcal{S}_{\mathrm{Id}}(f).$$

We apply this result to compute the path components of the fiber $\mathrm{PH}^{-1}(D)$ when M is a circle:

Proposition 67. *Assume $M = \mathbb{S}^1$. Let B be a barcode and $f \in \mathrm{PH}^{-1}(D)$. Then $\mathrm{PH}_f^{-1}(B)$ is homotopy equivalent to $\mathbb{S}^1$.*

Proof. From Proposition 66 $\mathrm{PH}_f^{-1}(B)$ is homeomorphic to $\mathcal{D}_{\mathrm{Id}}(\mathbb{S}^1)/\mathcal{S}_{\mathrm{Id}}(f)$. Let n be the number of minima of f, which is then also the number of maxima of f because $\chi(\mathbb{S}^1) = 0$. Without loss of generality we assume that the associated $2n$ critical points of f are evenly spaced on $\mathbb{S}^1$. The space $\mathcal{D}_{\mathrm{Id}}(\mathbb{S}^1)$ of Id-isotopic diffeomorphisms of the circle deformation retracts to $\mathbb{S}^1$, i.e. the rotations of the circle. The

[2]We further warn the interested reader that in [61] the notation $\mathcal{S}_{\mathrm{Id}}(f)$ rather stands for the space of diffeomorphisms ϕ preserving f that are isotopic to Id_M *though* maps preserving f, thus it is the path connected component of Id_M in our own $\mathcal{S}_{\mathrm{Id}}(f)$.

subgroup $\mathcal{S}_{\mathrm{Id}}(f)$ of Id-isotopic diffeomorphisms ϕ preserving f, that is $f \circ \phi = f$, is then (isomorphic to) the subgroup of rotations consisting of the $2n$-th roots of unity that preserve the sequence of extremal values of f. The result follows since the quotient of $\mathbb{S}^1$ by a finite subgroup is again $\mathbb{S}^1$. □

When $M = [0, 1]$ recall that Morse functions have prescribed values a_0 and a_1 on the boundary points 0 and 1.

Proposition 68. *Assume $M = [0, 1]$. Let B be a barcode and $f \in \mathrm{PH}^{-1}(D)$. Then $\mathrm{PH}_f^{-1}(B)$ is contractible.*

Proof. From Proposition 66 $\mathrm{PH}_f^{-1}(B)$ is homeomorphic to $\mathcal{D}_{\mathrm{Id}}([0, 1])/\mathcal{S}_{\mathrm{Id}}(f)$. However $\mathcal{D}_{\mathrm{Id}}([0, 1])$ deformation retracts on the identity diffeomorphism $\mathrm{Id}_{[0,1]}$ by straight-line interpolations, and $\mathcal{S}_{\mathrm{Id}}(f) = \{\mathrm{Id}_{[0,1]}\}$. □

Propositions 67 and 68 can be viewed as Morse versions of the main results of the previous section for continuous functions. The analogues for lower-star filtrations on the subdivided interval and circle have been studied by [33] and [67] respectively.

Remark 8. *When $M = M_1 \sqcup M_2$ has more than one connected component, the path component $\mathrm{PH}_f^{-1}(B)$ in the fiber over $B = \mathrm{PH}(f)$ can be retrieved as the product of the path components of the fibers over $B_1 := \mathrm{PH}(f|_{M_1})$ and $B_2 := \mathrm{PH}(f|_{M_2})$ containing the restrictions $f|_{M_1}$ and $f|_{M_2}$ respectively:*

$$\mathrm{PH}_f^{-1}(B) = \mathrm{PH}_{f|_{M_1}}^{-1}(B_1) \times \mathrm{PH}_{f|_{M_2}}^{-1}(B_2).$$

For this reason we focus our analysis to the interesting case where M is connected.

For the rest of the section we fix a compact connected surface M and a function f with barcode B, whose number of critical points of index 1 is denoted by c_1. We make use of the analysis of the orbit $\mathcal{O}_{\mathrm{Id}}(f)$ by [61].

Proposition 69. *Assume that $c_1 > 0$. Then $\pi_2(\mathrm{PH}_f^{-1}(B)) = 0$ and $\pi_n(\mathrm{PH}_f^{-1}(B)) = \pi_n(M)$ for $n \geq 3$.*

Proof. $\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f)$ by Theorem 65, and by [61, (2), Theorem 1.5] we have $\pi_2(\mathcal{O}_{\mathrm{Id}}(f)) = 0$ and $\pi_n(\mathcal{O}_{\mathrm{Id}}(f)) = \pi_n(M)$ for $n \geq 3$. □

Remark 9. *From [61, (2), Theorem 1.5] we can also derive a short exact sequence $0 \to \pi_1(\mathcal{D}_{\mathrm{Id}}(M)) \oplus \mathbb{Z}^{k_f} \to \pi_1(\mathrm{PH}_f^{-1}(B)) \to G \to 0$ where G is a finite group and the integer $k_f \geq 0$ depends on the component $\mathrm{PH}_f^{-1}(B)$ in the fiber, on the number c_1 of saddles and the surface M.*

Proposition 70. *Assume that $c_1 = 0$. Then the homotopy type of the fiber $\mathrm{PH}_f^{-1}(B)$ is classified as follows:*

Surface M	$\mathbb{S}^2$	$\mathbb{S}^1 \times I$	$\mathbb{D}^2$
Fiber $\mathrm{PH}_f^{-1}(B)$	$\mathbb{S}^2$	$\{*\}$	$\{*\}$

Proof. $\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f)$ by Theorem 65, and the homotopy type of $\mathcal{O}_{\mathrm{Id}}(f)$ is computed by [61, Theorem 1.9]. □

Proposition 71. *Assume that B has pairwise distinct bounded interval endpoints, and that $c_1 > 0$. Then we have the following homotopy types for the fiber $\mathrm{PH}_f^{-1}(B)$:*

Surface M	$\mathbb{S}^2$	*Projective Plane*	*Torus*	$\mathbb{S}^1 \times I$	$\mathbb{D}^2$
Fiber $\mathrm{PH}_f^{-1}(B)$	$\mathrm{SO}(3) \times (\mathbb{S}^1)^{c_1-1}$	$\mathrm{SO}(3) \times (\mathbb{S}^1)^{c_1-1}$	$(\mathbb{S}^1)^{c_1+1}$	$(\mathbb{S}^1)^{c_1}$	$(\mathbb{S}^1)^{c_1}$

When M is obtained from the surfaces in the above tables by removing finitely many 2-disks, then $\mathrm{PH}_f^{-1}(B) \sim (\mathbb{S}^1)^{c_1-1}$. If M is the Möbius strip, then $\mathrm{PH}_f^{-1}(B) \sim (\mathbb{S}^1)^{c_1}$. For other orientable surfaces M, we have $\mathrm{PH}_f^{-1}(B) \sim (\mathbb{S}^1)^{c_1+\chi(M)}$. For the remaining non-orientable surfaces, we have $\mathrm{PH}_f^{-1}(B) \sim (\mathbb{S}^1)^{k_f}$ for some integer $k_f \leq c_1 + \chi(M)$, unless M is the Klein bottle in which case $k_f \leq c_1 + 1$.

Proof. $\mathrm{PH}_f^{-1}(B) = \mathcal{O}_{\mathrm{Id}}(f)$ by Theorem 65. Since B has distinct bounded interval endpoints, f has distinct critical points, and then the homotopy type of $\mathcal{O}_{\mathrm{Id}}(f)$ is computed by [61, (2)&(3), Theorem 1.5]. □

Remark 10. *When M has no boundary, $\partial M = \emptyset$, the number c_1 of saddles of Morse functions f in the fiber $\mathrm{PH}^{-1}(B)$ can be directly inferred from the barcode B. Namely, if we denote by k_B the number of intervals in B, then the quantity $k_B - \beta_0 - \beta_2$ counts (i) all the intervals $[b, d)$ of B in degree 1, which correspond by their birth value b to saddle points of f whose attaching handle increases the 1-dimensional homology of the sub level-set $f^{-1}(-\infty, b]$, and (ii) all the bounded intervals $[b, d)$ of B in degree 0, which correspond by their death value $d < \infty$ to saddle points of f whose attaching handle decreases the 0-dimensional homology of the sub level-set $f^{-1}(-\infty, d]$. Hence $c_1 = k_B - \beta_0 - \beta_2$. When $\partial M = \bigsqcup_j \partial M_j \neq \emptyset$, we can partition the boundary components ∂M_j into the sets $\partial M^{\min}$ (resp. $\partial M^{\max}$) of components ∂M_j that are local minimum (resp. maximum) of one (hence any) function f in the component of $\mathrm{PH}^{-1}(B)$ at stake. Since M is a surface each ∂M_j is a circle, therefore if $\partial M_j \subseteq \partial M^{\min}$, then it corresponds in the barcode B to the births of one interval in degree 0 and one interval in degree 1. Otherwise $\partial M_j \subseteq \partial M^{\max}$ induces no topological change when entering the sub level-sets of f. Consequently the correspondence between critical points and interval endpoints adapts and yields $c_1 = k_B - \beta_0 - \beta_2 - \#\partial M^{\min}$.*

Chapter 5

The Euler characteristic transform for shapes of dimension one

In the last chapter, we studied the amount of information that is lost when using persistent homology as a geometric descriptor. In this chapter we shift focus to another shape descriptor arising from TDA, the Euler characteristic transform (ECT) [91], which as we remarked in Chapter 2 converts shapes into functions and loses no information at all: the ECT provides a sufficient statistic for broad classes of shapes [31, 44] (e.g., compact semi-algebraic sets) in Euclidean space. A large array of methods for classifying shapes exist [13, 36, 41, 94], however among them, the ECT is theoretically well-motivated, interpretable and has been successfully applied in practice [2, 28, 63, 69, 94]. Further, the ECT signature lies in a vector space of functions and is thus well-suited for further statistical analysis. While the ECT is fast to compute, small perturbations in the input shape can lead to large differences in the output signature [24]. By contrast to other TDA methods, such as persistent homology and the persistent homology transform [91], we are not aware of any general stability results for the ECT which are independent of the triangulation of a shape.

In this chapter, we propose a new metric on the embeddings of a finite one-dimensional CW complex that is sensitive to changes in arc-length. Next, we introduce a norm on Euler characteristic transforms, in a similar vein to the norm introduced in Meng et al. [65, Equation 3.1], defined by taking first the 1-norm over the $\mathbb{R}$ component, and then the ∞-norm over the $\mathbb{S}^{d-1}$ component. We then prove a novel stability result for the ECT, showing that the ECT is continuous in our metric of embedded spaces (Theorem 73). In other words, if two embeddings of the same one-dimensional CW complex are sufficiently close in our metric, their corresponding ECTs are also close. To the best of our knowledge, our result is the first stability result for the ECT which is independent of the triangulation of a shape. Using similar

ideas, we also show that the ECT of a smooth underlying shape can be approximated using sufficiently fine triangulations (Theorem 85). Further, we propose a smoothing method for embeddings of one-dimensional CW complexes that were perturbed by independent Gaussian noise in ambient space. We use our stability result to prove that our smoothing method does not only yield stability but also provides a consistent statistical estimator for the ECT of a noisy data set (Theorem 93), i.e. the ECT of the smoothed shape converges to the ECT of the underlying shape in probability as we increase the number of noisy observations.

The well-known stability results in applied topology for Čech and Vietoris-Rips filtrations of point clouds are stated in terms of the Hausdorff distance [22]. Proving stability results for the ECT is complicated by the fact that this metric is too coarse for the ECT to be continuous. Crucially, it is straightforward to construct examples of two shapes embedded in Euclidean space which are close in Hausdorff distance but whose ECTs are far apart. We loosely classify such instabilities into two categories.

The first type of instability arises when two shapes are close in Hausdorff distance, yet not homeomorphic to each other. Counterexamples can be constructed by adding a single point to a shape at an arbitrarily close distance, as visualised in Figure 5.1 for the case of an embedded simplicial complex. We point out that classical persistent homology and the persistent homology transform (PHT) [91] suffer from the same instability. However, extended persistence [26] and the extended persistent homology transform [92] can be used to partially overcome this type of instability. In this paper, we resolve the described type of instability by restricting ourselves to shapes that are homeomorphic. Restricting an ECT analysis to a homeomorphism class of shapes is common in applications, see for example [2, 63, 69, 88].

Secondly, the ECT can suffer from instability through excessive curvature. For example, in the case of shapes homeomorphic to $\mathbb{S}^1$ or $I = [0, 1]$, which can be parameterised as curves, this type of instability occurs when two curves are close in the embedded space, but one curve changes curvature much more rapidly. An example of such curves is given in Figure 5.2. Such instability is expected to occur if a shape is approximated based on points which are perturbed *independently* of each other by ambient noise.

Our work resolves instabilities of the second type for one-dimensional shapes by proposing a new metric which is sensitive to curvature. We also provide a statistical estimator of the ECT which is consistent under perturbations by independently distributed Gaussian noise. These perturbations are likely to produce changes in curvature when we naively interpolate points. While the PHT and extended PHT

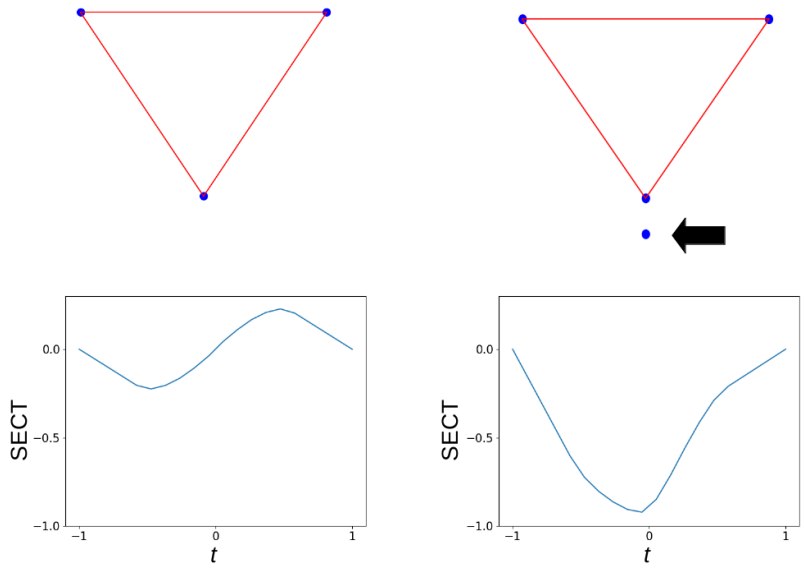

Figure 5.1: We visualise two embedded simplicial complexes (top) which differ by a single vertex. Their SECTs (bottom), visualised for a filtration in the bottom-to-top direction, are significantly different. The illustrated behaviour persists when we move the disconnected vertex in the top right panel (indicated by arrow) arbitrarily close to the larger connected component.

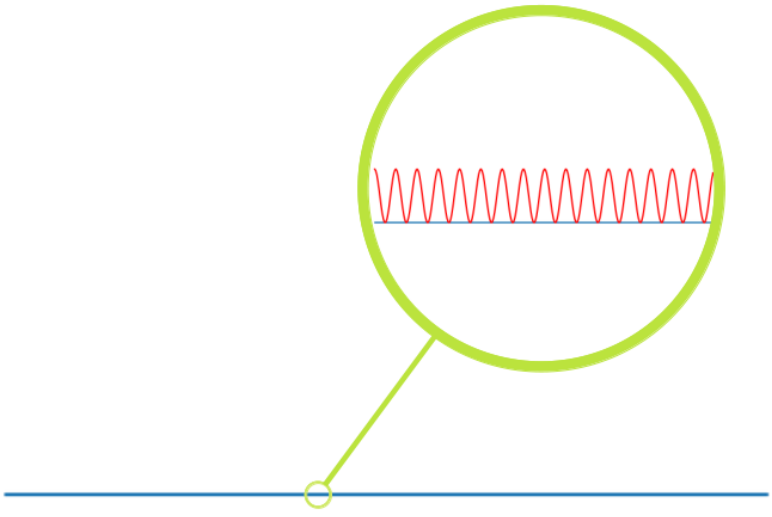

Figure 5.2: Two shapes homeomorphic to $[0,1]$ embedded into $\mathbb{R}^2$: A straight line (blue) and a wave (red) closely following the straight line with a small amplitude ε and high frequency. As long as the frequency is high enough for the wave to go through $n := \lceil 1/\varepsilon \rceil$ amplitudes, the distance of the ECTs of the two curves is at least 1 (fix the $\mathbb{S}^1$-component of the ECTs to (0,1) and compute the 1-norm over $\mathbb{R}$), while the Hausdorff distance between the curves is ε.

do not suffer from instabilities as illustrated in Figure 5.2, neither method provides a consistent estimator. Furthermore, the ECT arguably provides signatures more amenable to the application of further statistical and machine learning methods and are, by themselves as well as in conjunction with our new method, faster to compute.

5.0.1 Related Work

Already the work of Berkouk [10] shows that there is no metric on Euler curves that is stable against the interleaving distance of underlying persistence modules and satisfies a few mild, desirable conditions. We note that the stability of the Wasserstein distance proved by Skraba and Turner [83] provides a straightforward stability result for the ECT. Further, Dłotko and Gurnari [35] prove a similar result for the Euler characteristic curve. Nadimpalli et al. [69] prove a stability result for the ECT on binary image data, which is linear in the number of voxels at which two images differ. However, these stability results depend on the number of simplices in the underlying simplicial complex and the bound on the ECT becomes increasingly loose as the number of data points increases. Meng et al. [65] provide results that imply stability of the ECT when a shape is perturbed by rotations and translations but not for more general perturbations. Tameness assumptions, which are not needed in our results, are required for the stability they prove to hold. They also provide a statistical inference pipeline for shapes using the SECT. However, their pipeline considers parameterised families of shapes and random perturbations only happen in parameter space. As a result, the perturbations of points in shape space are correlated. By contrast, our results on the estimation of the ECT and SECT allow independent perturbations in ambient space.

5.1 Deterministic stability

5.1.1 Stability for smooth curves

As we observed at the start of the chapter, controlling the proximity of two different one-dimensional shapes is not enough to control the difference between their ECTs. This motivates the definition of a metric between such shapes which is also concerned with perturbations to length.

Definition 72. *Let Z be a finite one-dimensional CW complex with a fixed CW structure $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$. Fix $r \geq 1$. For $X, Y \in \mathcal{G}^r(Z^*, d)$, we define $d_{Z^*}(X, Y)$ to*

be the infimum of all ε such that there exists $h_X, h_Y \in \mathcal{E}^r(Z^*, d)$, whose images are X and Y respectively, satisfying:

1. *The difference of arc lengths between $h_X \circ \Phi_\lambda$ and $h_Y \circ \Phi_\lambda$ is less than or equal to ε for each λ.*

2. *Both $h_X \circ \Phi_\lambda$ and $h_Y \circ \Phi_\lambda$ are curves of constant velocity for each λ.*

3. *$\|h_X - h_Y\|_\infty \leq \varepsilon$.*

By using the compactness of Z it is not difficult to show that d_{Z^*} is a metric on $\mathcal{G}^r(Z^*, d)$. For the remainder of this paper, we endow $\mathcal{G}^r(Z^*, d)$ with this metric and the topology arising from it. A key goal of this paper is to show that the ECT is a continuous map on $\mathcal{G}^r(Z^*, d)$ for $r \geq 2$.

Theorem 73. *Let Z be a finite one-dimensional CW complex with a fixed CW structure $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda \in \Lambda})$. The map $X \mapsto \mathrm{ECT}_X$ is continuous on $\mathcal{G}^r(Z^*, d)$ for $r \geq 2$.*

In particular, if X has curvature bounded by M, and the image of the λ^{th} 1-cell in X has arc length L_λ, then whenever $d_{Z^*}(X, Y) < \varepsilon$, we have

$$\|\mathrm{ECT}_X - \mathrm{ECT}_Y\| \leq |Z_0|\varepsilon + \sum_{\lambda \in \Lambda} G_\lambda(\varepsilon),$$

where

$$G_\lambda(\varepsilon) := \begin{cases} 8\sqrt{L_\lambda n_\lambda \varepsilon} + n_\lambda \varepsilon & L_\lambda / n_\lambda > 2\varepsilon \\ 11 n_\lambda \varepsilon & L_\lambda / n_\lambda \leq 2\varepsilon \end{cases}$$

and

$$n_\lambda := \max\left(\left\lceil \left(\frac{M^2 L_\lambda^3}{24\varepsilon} \right)^{1/3} \right\rceil, \left\lceil \frac{L_\lambda M}{\pi} \right\rceil \right).$$

We prove this theorem via the following proposition, which will be useful when considering functional information in Section 5.2.

Proposition 74. *Let Z be a finite one-dimensional CW complex with a fixed CW structure $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda \in \Lambda})$. Let $f, g \in \mathcal{F}^r(Z^*, d)$, with $r \geq 2$, and suppose that:*

1. *The curves $f \circ \Phi_\lambda$ and $g \circ \Phi_\lambda$ have arc lengths that differ by at most ε for each λ.*

2. *The curves $f \circ \Phi_\lambda$ and $g \circ \Phi_\lambda$ have constant velocity for each λ.*

3. *$\|f - g\|_\infty \leq \varepsilon$.*

Then if f has curvature bounded by M, and $f \circ \Phi_\lambda$ has arc length L_λ, we have

$$\left\|\mathrm{ECT}_f - \mathrm{ECT}_g\right\| \leq |Z_0|\varepsilon + \sum_{\lambda \in \Lambda} G_\lambda(\varepsilon),$$

where G_λ and n_λ are defined as in Theorem 73.

The idea of the proof of this proposition is as follows. We show that the norm of the ECT of a curve can be controlled using its differential properties. Using this observation, we can bound the difference in ECT of two curves that are nearly straight lines. We can also refine the structure of a finite one-dimensional CW complex Z^* with a map $f : Z \to \mathbb{R}^d$ into enough pieces that the image of every 1-cell is a nearly linear curve. The above proposition then follows from a glueing argument.

In this chapter we let I denote the unit interval $[0, 1]$ and say that $f : I \to \mathbb{R}^d$ is piece-wise C^1 if it is continuous and there exists a collection $T_1, \ldots, T_k$ of closed intervals covering I, on the interiors of which f is C^1.

Proposition 75. *Let $\gamma : I \to \mathbb{R}^d$ piece-wise C^1 map, with X being the image of γ. Fix any $v \in \mathbb{S}^{d-1}$. Let a and b be the minimal and maximal values of $f : x \mapsto \langle v, \gamma(x)\rangle$ on I respectively. Then*

$$\int_a^b |\mathrm{ECT}_\gamma(v, t)| \, \mathrm{d}t \leq V(f),$$

where $V(f)$ denotes the variation of f:

$$V(f) := \int_0^1 |f'| \, \mathrm{d}t.$$

For $t \geq b$ we have $\mathrm{ECT}_\gamma(v, t) = 1$. For $t < a$ we have $\mathrm{ECT}_\gamma(v, t) = 0$. In particular, $\mathrm{ECT}_\gamma(v, t)$ is defined almost everywhere.

Proof. The last two statements follow from the fact that I has Euler characteristic one and the empty set has Euler characteristic zero. We now prove the remainder of the proposition.

The main goal of the proof is to establish the following two equalities:

$$\int_a^b |\mathrm{ECT}_\gamma(v, t)| \, \mathrm{d}t = \int_a^b \left|\pi_0\left[f^{-1}(-\infty, t]\right]\right| \, \mathrm{d}t$$
$$V(f) = \int_a^b \left|\pi_0\left[f^{-1}(t)\right]\right| \, \mathrm{d}t.$$

The first of these equalities follows from the fact that every subset of the unit interval is component-wise contractible. Hence, the Euler characteristic of any subset

of I is the number of path-components it has. The second of these equalities is more difficult to show, and its establishment is the bulk of the proof. Once both equalities are shown, demonstrating that the first integrand on the right is less than or equal to the second integrand on the right completes the proof.

To begin, notice that f is piece-wise C^1 since γ is. For the proof, we let $G(f)$ be the set of points where f' is defined and positive, $D(f)$ be the set of points where f' is defined and negative, and $C(f)$ be the set of points in I neither in $G(f)$ or $D(f)$. Since f is piece-wise C^1, $G(f)$ and $D(f)$ are both open. Meanwhile, all but finitely many points of $C(f)$ satisfy $f'(x) = 0$. Clearly, $G(f)$, $D(f)$, and $C(f)$ partition I. Hence,

$$\begin{aligned} V(f) &= \int_{G(f)} f'(x)\,\mathrm{d}x + \int_{D(f)} -f'(x)\,\mathrm{d}x + \int_{C(f)} |f'(x)|\,\mathrm{d}x \\ &= \int_{G(f)} f'(x)\,\mathrm{d}x + \int_{D(f)} -f'(x)\,\mathrm{d}x, \end{aligned}$$

since $f' = 0$ almost everywhere on $C(f)$. Since both $G(f)$ and $D(f)$ are open, each is a countable union of open subintervals of I, which we denote by $\{I_k\}_{k\in\Xi}$ and $\{J_l\}_{l\in\Theta}$ respectively. On each I_k, f is increasing and on each J_l, f is decreasing. Hence, we get

$$\begin{aligned} V(f) &= \int_{G(f)} f'(x)\,\mathrm{d}x + \int_{D(f)} -f'(x)\,\mathrm{d}x \\ &= \sum_{k\in\Xi} \int_{I_k} f'(x)\,\mathrm{d}x + \sum_{l\in\Theta} \int_{J_l} -f'(x)\,\mathrm{d}x \\ &= \sum_{k\in\Xi} \int_{\mathbb{R}} \left|\pi_0\big[(f|_{I_k})^{-1}(t)\big]\right| \mathrm{d}t + \sum_{l\in\Theta} \int_{\mathbb{R}} \left|\pi_0\big[(f|_{J_l})^{-1}(t)\big]\right| \mathrm{d}t \\ &= \int_{\mathbb{R}} \left|\pi_0\big[(f|_{G(f)})^{-1}(t)\big]\right| \mathrm{d}t + \int_{\mathbb{R}} \left|\pi_0\big[(f|_{D(f)})^{-1}(t)\big]\right| \mathrm{d}t \\ &= \int_{\mathbb{R}} \left|\pi_0\big[(f|_{G(f)\cup D(f)})^{-1}(t)\big]\right| \mathrm{d}t. \end{aligned}$$

Here, the last line follows from the fact that if x and y have the same f value, each point in $G(f)$ or $D(f)$, then by the definition of these sets there must be a point between them that obtains either a smaller or larger value of f. The line before follows from similar reasoning. If it is granted that $f(C(f))$ has measure zero we

then have

$$\begin{aligned} V(f) &= \int_{\mathbb{R}} \left| \pi_0\left[(f|_{G(f)\cup D(f)})^{-1}(t) \right] \right| \, \mathrm{d}t \\ &= \int_{\mathbb{R}} \left| \pi_0\left[(f|_{G(f)\cup D(f)})^{-1}(t) \right] \right| \, \mathrm{d}t + \int_{\mathbb{R}} \left| \pi_0\left[(f|_{C(f)})^{-1}(t) \right] \right| \, \mathrm{d}t \\ &= \int_{\mathbb{R}} \left| \pi_0\left[f^{-1}(t) \right] \right| \, \mathrm{d}t \\ &= \int_a^b \left| \pi_0\left[f^{-1}(t) \right] \right| \, \mathrm{d}t. \end{aligned}$$

Here, the second to last equality follows from the fact that if x is in $C(f)$ and y is in $G(f)$ or $D(f)$, then by definition of $G(f)$ and $D(f)$ there is a point between x and y with an f value not equal to $f(y)$. This equation implies the integrand at the end of this equation must be finite almost everywhere.

Now we show that indeed $f(C(f))$ has measure zero. Since f is piece-wise C^1, let $T_1, \ldots, T_k$ be closed intervals covering I on the interiors of which f is C^1. Consider Z_i, the subset of the interior of T_i with $f' = 0$. By Sard's Theorem (see for example [78, Theorem 7.2]), $f(Z_i)$ has measure zero. We have that $C(f)$ is a subset of $Z_1 \cup \ldots \cup Z_k \cup \partial T_1 \cup \ldots \cup \partial T_k$, whose image under f has measure zero. Thus $f(C(f))$ has measure zero.

Hence, the desired result follows once we establish that for any $t \in [a, b]$,

$$\left| \pi_0\left[f^{-1}(-\infty, t] \right] \right| \leq \left| \pi_0\left[f^{-1}(t) \right] \right|.$$

This inequality implies that $\mathrm{ECT}_f(v, t)$ is defined for almost all t since $\mathrm{ECT}_f(v, t)$ is just the left-hand side of this inequality, which is positive-valued, and bounded by a function that is finite for almost all t.

To prove this statement, it suffices to show that any path-component of $f^{-1}(-\infty, t]$ contains a point with f value t. Suppose otherwise, that there is a path-component C of $f^{-1}(-\infty, t]$ with $f(C) < t$. By continuity of f, C must also be a path-component of $f^{-1}(-\infty, b]$. Indeed, suppose α is a path from the complement of C to C. Thus, every neighborhood of $\alpha^{-1}(C)$ must intersect $(f \circ \alpha)^{-1}(t, b]$. But this produces a contradiction of the intermediate value theorem. So C is a path-component of $I = f^{-1}(-\infty, b]$ and hence is I. The fact that $f(C) < t \leq b$ thus contradicts the definition of b as the maximum of f, completing the proof. □

Remark 11. *Suppose $\gamma : I \to \mathbb{R}^d$ is continuous and definable with respect to an o-minimal structure on $\mathbb{R}$. By [93, Chapter 7, Theorem 3.2]), γ is piece-wise C^1 and hence the above result applies.*

Remark 12. *In the case where f is tame, this result is implied by a stronger result of [11, Corollary 4.6] for tame functions. The main contribution of this proposition is that the stated bound still holds when f is not tame (or even pfd).*

With the previous result in mind, we establish a bound on the variation of a curve that is approximately straight.

Lemma 76. *Suppose $\gamma : I \to \mathbb{R}^d$ is a piece-wise differentiable path with length L and the first coordinate γ_1 satisfies $|\gamma_1(1) - \gamma_1(0)| = L_x$. Then the variation of any other coordinate function γ_n of γ is bounded by*

$$V(\gamma_n) \leq \sqrt{L^2 - L_x^2}. \tag{5.1}$$

Proof. Without loss of generality, we can assume that $\gamma(0) = 0$, $\gamma_1(1) = L_x$, and we can show this bound holds for γ_2 only.

From γ we can construct another function $\bar{\gamma}$ by

$$\bar{\gamma}(t) := \int_0^t (\gamma_1'(t),\ |\gamma_2'(t)|,\ \gamma_3'(t), \ldots,\ \gamma_d'(t))\ \mathrm{d}t.$$

Put differently, $\bar{\gamma}$ has the same coordinate functions as γ except in the second coordinate, where $\bar{\gamma}_2$ has the same absolute value of its derivative as γ, but is never decreasing. It is immediate that γ and $\bar{\gamma}$ have the same length, value of the first coordinate at 1, and variation in the second coordinate. Hence, it suffices to show that the lemma holds for γ_2 for curves γ with length L, $\gamma(0) = 0$, $\gamma_1(1) = L_x$, and $\gamma_2'(t) \geq 0$ for all t.

The fact that γ has length L implies that $\gamma(1)$ lies in the closed d-disk of radius L centred at the origin. The fact that $\gamma_1(x) = L_x$ implies that $\gamma(1)$ lies on the hyperplane of points with the first coordinate L_x. Elementary geometry shows that the intersection of the disk and the hyperplane is

$$\{(y_1, y_2, \ldots, y_n) \in \mathbb{R}^d : y_1 = L_x,\ y_2^2 + \ldots + y_d^2 \leq L^2 - L_x^2\}.$$

It is easily seen that the greatest value of y_2 on this set is $\sqrt{L^2 - L_x^2}$. So $\gamma_2(1)$ is bounded above by this value. Hence,

$$V(\gamma_2) = \int_0^1 \left|\gamma_2'(t)\right| \mathrm{d}t = \int_0^1 \gamma_2'(t)\ \mathrm{d}t = \gamma_2(1) - \gamma_2(0) = \gamma_2(1) \leq \sqrt{L^2 - L_x^2}.$$

□

We can now bound the L_1 distance between Euler characteristic transforms of nearby curves, assuming one of them is approximately straight.

Proposition 77. *Let $\alpha : I \to \mathbb{R}^d$ be a piece-wise C^1 map such that the distance of $\alpha(0)$ to $\alpha(1)$ is L, with arc length no greater than $L + \varepsilon$. Let $\beta : I \to \mathbb{R}^d$ be another piece-wise C^1 map with arc length no greater than $L + 2\varepsilon$, and endpoints within ε of the corresponding endpoints of α. Then*

$$\left\|\mathrm{ECT}_\alpha - \mathrm{ECT}_\beta\right\| \leq \begin{cases} 8\sqrt{L\varepsilon} & L > 2\varepsilon \\ 10\varepsilon & L \leq 2\varepsilon. \end{cases}$$

Proof. Let v be an arbitrary unit vector, $w = \alpha(1) - \alpha(0)$, and θ be the angle between w and the hyperplane normal to v. After potentially applying a rotation to α and β, we may assume that $v = (0, 1, 0, \ldots, 0)$. By applying another rotation we may assume also that w is only non-zero in the first two coordinates.

Let f denote the inner-product with v and let a and b be the minimum and maximum of $f \circ \alpha$ and c and d be the minimum and maximum of $f \circ \beta$. Throughout the proof, we use the fact that if both $\mathrm{ECT}_\alpha(v, t)$ and $\mathrm{ECT}_\beta(v, t)$ are non-zero, then

$$|\mathrm{ECT}_\alpha(v, t) - \mathrm{ECT}_\alpha(v, t)| \leq \mathrm{ECT}_\alpha(v, t) + \mathrm{ECT}_\beta(v, t) - 2,$$

as both Euler characteristic transforms must be positive (since subsets of the interval are component-wise contractible) and greater than 1.

First, suppose $\max(a, c) \leq \min(b, d)$. Then

$$\begin{aligned} \int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v, t) - \mathrm{ECT}_\beta(v, t)| \, \mathrm{d}t = &\int_{\min(a,c)}^{\max(a,c)} |\mathrm{ECT}_\alpha(v, t) - \mathrm{ECT}_\beta(v, t)| \, \mathrm{d}t \\ &+ \int_{\max(a,c)}^{\min(b,d)} |\mathrm{ECT}_\alpha(v, t) - \mathrm{ECT}_\beta(v, t)| \, \mathrm{d}t \\ &+ \int_{\min(b,d)}^{\max(b,d)} |\mathrm{ECT}_\alpha(v, t) - \mathrm{ECT}_\beta(v, t)| \, \mathrm{d}t. \end{aligned}$$

Suppose also that $b \leq d$. Then the above expression is bounded by

$$\int_{\min(a,c)}^{\max(a,c)} \mathrm{ECT}_\alpha(v, t) + \mathrm{ECT}_\beta(v, t) \, \mathrm{d}t + \int_{\max(a,c)}^{b} \mathrm{ECT}_\alpha(v, t) + \mathrm{ECT}_\beta(v, t)$$

$$-2 \, \mathrm{d}t + \int_b^d \mathrm{ECT}_\beta(v, t) - 1 \, \mathrm{d}t,$$

where we have only had to approximate the middle term. If we additionally suppose $a \leq c$, then our bound is equal to

$$\int_a^c \mathrm{ECT}_\alpha(v,t)\ \mathrm{d}t + \int_c^b \mathrm{ECT}_\alpha(v,t) + \mathrm{ECT}_\beta(v,t) - 2\ \mathrm{d}t + \int_b^d \mathrm{ECT}_\beta(v,t) - 1\ \mathrm{d}t.$$

Rearranging and by the linearity of the integral, the above is equal to

$$\int_c^d \mathrm{ECT}_\beta(v,t)\ \mathrm{d}t + \int_a^b \mathrm{ECT}_\alpha(v,t)\mathrm{d}t - 2(b-c) - (d-b).$$

Similar analysis when $a > c$ and/or $b > d$ shows that when $\max(a,c) \leq \min(b,d)$,

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t \leq \int_c^d \mathrm{ECT}_\beta(v,t)\ \mathrm{d}t + \int_a^b \mathrm{ECT}_\alpha(v,t) \\ - 2|\min(b,d) - \max(a,c)| - |\max(b,d) - \min(b,d)|.\end{aligned} \tag{5.2}$$

Note that $|b-a| = L|\sin\theta|$. By hypothesis $\|\alpha(0) - \beta(0)\| \leq \varepsilon$, so $|a-c| \leq \varepsilon$. Similarly $|b-d| \leq \varepsilon$. Hence,

$$|\min(b,d) - \max(a,c)| \geq L|\sin\theta| - 2\varepsilon$$

by the triangle inequality. Of course, the quantity on the left is also positive, so

$$|\min(b,d) - \max(a,c)| \geq \max(0, L|\sin\theta| - 2\varepsilon).$$

Trivially, we also have $|\max(b,d) - \min(b,d)| \geq 0$. Applying these inequalities, Proposition 75, and Lemma 76 to Equation (5.2), we have

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t \leq \sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} \\ + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} \\ - 2\max(0, L|\sin\theta| - 2\varepsilon).\end{aligned} \tag{5.3}$$

In the application of Lemma 76 for the first term, we use that $|\beta_1(1) - \beta_1(0)| \leq \max(0, L|\cos\theta| - 2\varepsilon)$.

Otherwise, $\max(a,c) \geq \min(b,d)$, so either $a \leq b \leq c \leq d$ or $c \leq d \leq a \leq b$. In either of these cases, we must have that $L|\sin\theta| \leq \varepsilon$. Consider the first of these cases. We observe

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t = \int_a^b |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t \\ + \int_b^c |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t \\ + \int_c^d |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\ \mathrm{d}t.\end{aligned}$$

This quantity is equal to

$$\int_a^b \mathrm{ECT}_\alpha(v,t)\,\mathrm{d}t + \int_b^c 1\,\mathrm{d}t + \int_c^d \mathrm{ECT}_\beta(v,t) - 1\,\mathrm{d}t,$$

Bounding from above, we have

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\,\mathrm{d}t &\leq \int_a^b \mathrm{ECT}_\alpha(v,t)\,\mathrm{d}t + \int_c^d \mathrm{ECT}_\beta(v,t)\,\mathrm{d}t \\ &\quad + (c-b) - (d-c) \\ &\leq \int_a^b \mathrm{ECT}_\alpha(v,t)\,\mathrm{d}t + \int_c^d \mathrm{ECT}_\beta(v,t)\,\mathrm{d}t + (c-b).\end{aligned}$$

Similar analysis when $c \leq d \leq a \leq b$ shows that in general, if $\max(a,c) \leq \min(b,d)$, then

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\,\mathrm{d}t &\leq \int_a^b \mathrm{ECT}_\alpha(v,t)\,\mathrm{d}t \\ &\quad + \int_c^d \mathrm{ECT}_\beta(v,t)\,\mathrm{d}t + \min(|c-b|, |d-a|).\end{aligned}$$

By the triangle inequality, $\min(|c-b|, |d-a|) \leq \varepsilon - L\sin\theta$. Applying Proposition 75 and Lemma 76 once again, we see

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\,\mathrm{d}t &\leq \sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} \\ &\quad + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} \\ &\quad + \max(0, \varepsilon - L|\sin\theta|).\end{aligned}$$

In summary, $\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\,\mathrm{d}t$ is bounded by

$$\sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} - 2\max(0, L|\sin\theta| - 2\varepsilon)$$

whenever $L|\sin\theta| \geq \varepsilon$. Otherwise, either we still have $\max(a,c) \leq \min(b,d)$ and the above bound still holds or $\max(a,c) \geq \min(b,d)$ and we instead have the bound

$$\sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} + \max(0, \varepsilon - L|\sin\theta|).$$

Hence, in general,

$$\begin{aligned}\int_{\mathbb{R}} |\mathrm{ECT}_\alpha(v,t) - \mathrm{ECT}_\beta(v,t)|\,\mathrm{d}t &= \sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} \\ &\quad + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} \\ &\quad - 2\max(0, L|\sin\theta| - 2\varepsilon) + \max(0, \varepsilon - L|\sin\theta|).\end{aligned}$$

The proof is complete once we have established the following tedious lemma. □

Lemma 78. *The function*

$$f(\theta) = \sqrt{(L+2\varepsilon)^2 - \max(0, L|\cos\theta| - 2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2|\cos^2\theta|} \\ - 2\max(0, L|\sin\theta| - 2\varepsilon) + \max(0, \varepsilon - L|\sin\theta|)$$

is bounded above by

$$f(\theta) \leq \begin{cases} 8\sqrt{L\varepsilon} & L > 2\varepsilon \\ 10\varepsilon & L \leq 2\varepsilon. \end{cases}$$

Proof. Thanks to the symmetries of the sine and cosine functions and the absolute values present in the formula for f, we have that $f(-\theta) = f(\theta)$ and $f(\pi/2-\theta) = f(\theta)$. So $f(\theta) = f(\pi/2+\theta)$. Therefore it suffices to bound f on the interval $[0, \pi/2]$. On this interval, we can remove the absolute values in the formula for f, giving

$$f(\theta) = \sqrt{(L+2\varepsilon)^2 - \max(0, L\cos\theta - 2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2\cos^2\theta} \\ - 2\max(0, L\sin\theta - 2\varepsilon) + \max(0, \varepsilon - L\sin\theta).$$

With the exception of finitely many values of θ, the derivative of f exists and is equal to

$$\mathbb{I}(L\cos\theta > 2\varepsilon)\frac{L\sin\theta(L\cos\theta - 2\varepsilon)}{\sqrt{(L+2\varepsilon)^2 - (L\cos\theta - 2\varepsilon)^2}} + \frac{L^2\sin\theta\cos\theta}{\sqrt{(L+\varepsilon)^2 - L^2\cos^2\theta}} \\ -\mathbb{I}(L\sin\theta > 2\varepsilon)2L\cos(\theta) - \mathbb{I}(L\sin\theta < \varepsilon)L\cos\theta,$$

where $\mathbb{I}$ denotes the indicator function.

Notice that

$$\frac{L^2\sin\theta\cos\theta}{\sqrt{(L+\varepsilon)^2 - L^2\cos^2\theta}} \leq \frac{L^2\sin\theta\cos\theta}{\sqrt{L^2 - L^2\cos^2\theta}} = L\cos\theta,$$

and similarly

$$\frac{L\sin\theta(L\cos\theta - 2\varepsilon)}{\sqrt{(L+2\varepsilon)^2 - (L\cos\theta - 2\varepsilon)^2}} \leq \frac{L^2\sin\theta\cos\theta}{\sqrt{L^2 - L^2\cos^2\theta}} = L\cos\theta.$$

Using these identities, we get that

$$f'(\theta) \leq \mathbb{I}(L\cos\theta > 2\varepsilon)L\cos\theta + L\cos\theta - \mathbb{I}(L\sin\theta > 2\varepsilon)2L\cos(\theta) - \mathbb{I}(L\sin\theta < \varepsilon)L\cos\theta.$$

Hence f is weakly decreasing whenever $L\sin\theta > 2\varepsilon$. When $\varepsilon < L\sin\theta < 2\varepsilon$, every non-zero term in f' is positive, and so f is increasing. We now bound f' in absolute value when $L\sin\theta < \varepsilon$. In this case,

$$\begin{aligned} |f'(\theta)| &\leq \mathbb{I}(L\cos\theta > 2\varepsilon)\frac{L\sin\theta(L\cos\theta - 2\varepsilon)}{\sqrt{(L+2\varepsilon)^2 - (L\cos\theta - 2\varepsilon)^2}} \\ &\quad + \frac{L^2\sin\theta\cos\theta}{\sqrt{(L+\varepsilon)^2 - L^2\cos^2\theta}} + L\cos\theta \\ &\leq 3L\cos\theta \\ &\leq 3L, \end{aligned}$$

using our approximations for the first and second terms from earlier.

Further, if $L > 2\varepsilon$,

$$\begin{aligned} f(0) &= \sqrt{(L+2\varepsilon)^2 - \max(0, L-2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2} - \varepsilon \\ &= \sqrt{8L\varepsilon} + \sqrt{2L\varepsilon + \varepsilon^2} - \varepsilon \\ &\leq (\sqrt{8} + \sqrt{3})\sqrt{L\varepsilon} - \varepsilon. \end{aligned}$$

Otherwise $L \leq 2\varepsilon$ and,

$$\begin{aligned} f(0) &= \sqrt{(L+2\varepsilon)^2 - \max(0, L-2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2} - \varepsilon \\ &= L + 2\varepsilon + \sqrt{2L\varepsilon + \varepsilon^2} - \varepsilon \\ &= L + \varepsilon + \sqrt{2L\varepsilon + \varepsilon^2} \\ &\leq (3 + \sqrt{5})\varepsilon. \end{aligned}$$

Hence, we can bound $f(\theta)$ for θ on the interval $[0, \sin^{-1}(\varepsilon/L)]$ (or $[0, \pi/2]$ if $\varepsilon > L$) by using our upper bounds for $f(0)$ and $|f'(\theta)|$ on this interval. By additionally using the inequality $\sin^{-1}(x) \leq \pi x/2$ for positive x, we obtain that when $L \sin\theta \leq \varepsilon$,

$$\begin{aligned} f(\theta) &\leq \begin{cases} (\sqrt{8} + \sqrt{3})\sqrt{L\varepsilon} + (3\pi/2 - 1)\varepsilon & L > 2\varepsilon \\ (3 + \sqrt{5} + 3\pi/2)\varepsilon & \varepsilon \leq L \leq 2\varepsilon \\ (3 + \sqrt{5})\varepsilon + (3\pi/2)L & L < \varepsilon \end{cases} \\ &\leq \begin{cases} (\sqrt{8} + \sqrt{3})\sqrt{L\varepsilon} + (3\pi/2 - 1)\varepsilon & L > 2\varepsilon \\ (3 + \sqrt{5} + 3\pi/2)\varepsilon & L \leq 2\varepsilon \end{cases} \\ &\leq \begin{cases} (\sqrt{8} + \sqrt{3} + (3\pi/4 - 1/2)\sqrt{2})\sqrt{L\varepsilon} & L > 2\varepsilon \\ (3 + \sqrt{5} + 3\pi/2)\varepsilon & L \leq 2\varepsilon \end{cases} \\ &\leq \begin{cases} 8\sqrt{L\varepsilon} & L > 2\varepsilon \\ 10\varepsilon & L \leq 2\varepsilon. \end{cases} \end{aligned}$$

Otherwise, we know f is weakly increasing until $L \sin\theta > 2\varepsilon$, after which point it is weakly decreasing. Hence, if f is not maximised where $L \sin\theta \leq \varepsilon$, it must attain its maximum when $L \sin\theta = 2\varepsilon$, or equivalently $\theta = \sin^{-1}(2\varepsilon/L)$. Note that this implies

$2\varepsilon \leq L$. We compute

$$\begin{aligned}
f(\sin^{-1}(2\varepsilon/L)) &= \sqrt{(L+2\varepsilon)^2 - \max(0, \sqrt{L^2-4\varepsilon^2} - 2\varepsilon)^2} + \sqrt{(L+\varepsilon)^2 - L^2 - 4\varepsilon^2} \\
&= \sqrt{(L+2\varepsilon)^2 - \max(0, \sqrt{L^2-4\varepsilon^2} - 2\varepsilon)^2} + \sqrt{2L\varepsilon - 3\varepsilon^2} \\
&= \begin{cases} \sqrt{4L\varepsilon + 4\varepsilon^2 + 2\varepsilon\sqrt{L^2-4\varepsilon^2}} + \sqrt{2L\varepsilon - 3\varepsilon^2} & L > 2\sqrt{2}\varepsilon \\ L + 2\varepsilon + \sqrt{2L\varepsilon - 3\varepsilon^2} & 2\varepsilon \leq L \leq 2\sqrt{2}\varepsilon \end{cases} \\
&\leq \begin{cases} \sqrt{6L\varepsilon + 4\varepsilon^2} + \sqrt{2L\varepsilon} & L > 2\sqrt{2}\varepsilon \\ (2+2\sqrt{2})\varepsilon + \sqrt{2L\varepsilon} & 2\varepsilon \leq L \leq 2\sqrt{2}\varepsilon \end{cases} \\
&\leq \begin{cases} (\sqrt{6+\sqrt{2}} + \sqrt{2})\sqrt{L\varepsilon} & L > 2\sqrt{2}\varepsilon \\ (2+2\sqrt{2})\sqrt{L\varepsilon} & 2\varepsilon \leq L \leq 2\sqrt{2}\varepsilon \end{cases} \\
&\leq 5\sqrt{L\varepsilon}.
\end{aligned}$$

Hence, to totally bound $f(\theta)$ on the interval $[0, \pi/2]$ we need only use our earlier bound, namely

$$f(\theta) \leq \begin{cases} 8\sqrt{L\varepsilon} & L > 2\varepsilon \\ 10\varepsilon & L \leq 2\varepsilon. \end{cases}$$

□

The goal of the following proposition is to bound from below the chord length of a short segment of a curve given that it has bounded curvature.

Proposition 79. *Suppose $\gamma : [0, L] \to \mathbb{R}^d$ is a twice differentiable curve parametrised by arc length with curvature κ bounded in norm by M. Let $0 < \varepsilon < \pi/M$. Then for any $t \in [0, L - \varepsilon]$,*

$$\varepsilon \geq \|\gamma(t+\varepsilon) - \gamma(t)\|_2 \geq \frac{2}{M} \sin\left(\frac{M}{2}\varepsilon\right).$$

In particular,

$$\|\gamma(t+\varepsilon) - \gamma(t)\|_2 \geq \varepsilon - \frac{M^2}{24}\varepsilon^3.$$

To prove this we make use of the following theorem of Schwarz, which we cite from [23]:

Theorem 80 (Schwarz)**.** *Let C be an arc joining two given points A and B with curvature $\kappa(s) \leq 1/R$, such that $R \geq \frac{1}{2}\delta$, where δ is the distance between A and B. Let S be a circle of radius R through A and B. Then the length of C is either less than, or equal to, the shorter arc AB or greater than, or equal to, the longer arc AB on S.*

Proof of Proposition 79. The first inequality is clear since γ is parameterised by arc length. Now fix t. For the second inequality, consider the optimisation problem of minimising $\|\alpha(t+\varepsilon)-\alpha(t)\|_2$ subject to the constraints that $\|\alpha'\|_2 = 1$ and $\|\alpha''\|_2 \leq M$. Consider an arc of length ε on the circle of curvature M, which has radius $1/M$. Elementary geometry shows that the distance between the endpoints of such an arc is $\frac{2}{M}\sin(M\varepsilon/2)$. We claim that this arc provides an optimal solution. Indeed, let γ be any curve that performs at least as well as this arc in the sense that

$$\|\gamma(t+\varepsilon) - \gamma(t)\|_2 \leq \frac{2}{M}\sin\left(\frac{M}{2}\varepsilon\right),$$

while $\|\gamma'\|_2 = 1$, $\|\gamma''\|_2 \leq M$.

Let S be a circle of radius $1/M$ crossing both $\gamma(t)$ and $\gamma(t+\varepsilon)$. Such a circle must exist since $\|\gamma(t+\varepsilon) - \gamma(t)\|_2 \leq 2/M$. The curve γ on the interval $[t, t+\varepsilon]$ is of length $\varepsilon < \pi/M$ while the longer arc on S connecting $\gamma(t)$ and $\gamma(t+\varepsilon)$ has length greater than π/M. Hence, by the theorem of Schwarz, ε is less than or equal to the length of the shorter arc on S from $\gamma(t)$ to $\gamma(t+\varepsilon)$. If this inequality is strict, we may take a shorter portion of the circular arc with length ε, which has a shorter distance between endpoints. This proves that an arc of length ε on a circle of radius $1/M$ is an optimal solution of the optimisation problem.

Thus for potentially suboptimal γ, we have

$$\|\gamma(t+\varepsilon) - \gamma(t)\|_2 \geq \frac{2}{M}\sin\left(\frac{M}{2}\varepsilon\right).$$

For the last statement of the proposition, by the Lagrange remainder theorem

$$\sin x - x + x^3/6 = \int_0^x \cos t \frac{(x-t)^5}{5!}\,\mathrm{d}t.$$

The right side is clearly positive provided that $0 < x \leq \pi/2$. Since $0 < \frac{M}{2}\varepsilon < \pi/2$, we have

$$\frac{2}{M}\sin\left(\frac{M}{2}\varepsilon\right) \geq \frac{2}{M}\left[\frac{M}{2}\varepsilon - \frac{M^3}{48}\varepsilon^3\right] = \varepsilon - \frac{M^2}{24}\varepsilon^3.$$

□

We now use the results we have already proven about curves that are approximately straight to obtain a stability result for the Euler characteristic transform of more general shapes. To do this, we prove a lemma that allows us to glue together Euler characteristic transforms of functions restricted to different regions of a domain.

Definition 81. *Let $V^* = (V, V_0, \{\Phi_\lambda\}_{\lambda\in\Lambda_V})$ and $W^* = (W, W_0, \{\Phi_\lambda\}_{\lambda\in\Lambda_W})$ be finite one-dimensional CW complexes, each with a fixed CW structure. Suppose there exist maps $f_V \in \mathcal{F}^r(V^*, d)$ and $f_W \in \mathcal{F}^r(W^*, d)$, a subset $S \subseteq V_0$ and an injective map $m : S \to W_0$ such that $f_V = f_W \circ m$ on S. We define the glue of V^* and W^* under m to be a finite complex with structure:*

$$Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda_Z}) := \big((V \sqcup W)/m,\ \ (V_0 \sqcup W_0)/m,\ \ \{\Phi_\lambda\}_{\lambda\in\Lambda_V\sqcup\Lambda_W}\big).$$

We define the glue of f_V and f_W under m to be the map $f_Z : Z \to \mathbb{R}^d$ which restricts to f_V on V and f_W of W. This map is well defined (since $f_V = f_W \circ m$ on S) and is an element of $\mathcal{F}^r(Z^*, d)$.

Lemma 82. *Using the notation of the previous definition, suppose $\mathrm{ECT}_{f_V}(v,t)$ and $\mathrm{ECT}_{f_W}(v,t)$ are defined for almost all t for any fixed v. Then*

$$\mathrm{ECT}_{f_Z}(v,t) = \mathrm{ECT}_{f_V}(v,t) + \mathrm{ECT}_{f_W}(v,t) - \mathrm{ECT}_{f_S}(v,t) \tag{5.4}$$

for almost all t when v is fixed, where f_S is the restriction of f_V to S.

Proof. Fix a unit vector v in $\mathbb{R}^d$. Let $p_1, \ldots, p_k$ be the points of S. We denote by $V(v,t)$ the subset of points x in V satisfying that $\langle v, f_V(x)\rangle \leq t$. We define $W(v,t)$ and $Z(v,t)$ analogously. We let $S(v,t)$ denote the intersection of S and $V(v,t)$.

Via the inclusions of V and W into Z, we can view Z as the union of V and W, with V and W intersecting in Z at S. Similarly, we can view $Z(v,t)$ as the union of $V(v,t)$ and $W(v,t)$, with these two subsets intersecting at $S(v,t)$. For almost all $t \in \mathbb{R}$, $\langle v, f_V(p_i)\rangle \neq t$ for all i. Fix any such t. Hence, we have that the interiors of $V(v,t)$ and $W(v,t)$ cover their intersection $S(v,t)$, by continuity of f_V and f_W. Therefore, we have a Mayer-Vietoris exact sequence of homology groups:

$$\ldots \to H_i\big(S(v,t)\big) \to H_i\big(V(v,t)\big) \oplus H_i\big(W(v,t)\big) \to H_i\big(Z(v,t)\big) \to \ldots.$$

A routine argument then deduces the identity

$$\chi\big(Z(v,t)\big) = \chi\big(V(v,t)\big) + \chi\big(W(v,t)\big) - \chi\big(S(v,t)\big),$$

whenever all Euler characteristics on the right-hand side are defined. This happens for almost all t and is another way of writing the identity of Equation (5.4). □

We now have the prerequisites to prove Proposition 74.

Proof of Proposition 74. Let $\alpha_\lambda := f \circ \Phi_\lambda$ and $\beta_\lambda := g \circ \Phi_\lambda$. Since the index set Λ is finite, we let $\Lambda = \{1, \ldots, k\}$. We define $Z^0 := Z_0$, and inductively, $Z^\lambda = Z^{\lambda-1} \cup \operatorname{im} \Phi_\lambda$ for $\lambda \in \Lambda$. We then let $f_\lambda = f|_{Z^\lambda}$ and $g_\lambda = g|_{Z^\lambda}$.

Inductively, we assume that

$$\left\|\mathrm{ECT}_{f_{\lambda-1}} - \mathrm{ECT}_{g_{\lambda-1}}\right\| \leq |Z_0|\varepsilon + \sum_{k=1}^{\lambda-1} G_k(\varepsilon).$$

Indeed, as a base case, it is easily observed that

$$\left\|\mathrm{ECT}_{f_0} - \mathrm{ECT}_{g_0}\right\| \leq |Z_0|\varepsilon.$$

We can split α_λ into n pieces by restricting $\alpha_{\lambda,i} : [\frac{i-1}{n}, \frac{i}{n}] \to \mathbb{R}^d$. Analogously we can split β_λ into curves $\beta_{\lambda,i}$. By Proposition 79, the arc length of each $\alpha_{\lambda,i}$ is at most $M^2 L_\lambda^3 / 24n^3$ greater than the distance between its endpoints, if $n > L_\lambda M/\pi$. We now apply Proposition 77. Thus, provided

$$\frac{M^2 L_\lambda^3}{24n^3} \leq \varepsilon, \text{ or equivalently, } n \geq \left(\frac{M^2 L_\lambda^3}{24\varepsilon}\right)^{1/3},$$

we observe

$$\left\|\mathrm{ECT}_{\alpha_{\lambda,i}} - \mathrm{ECT}_{\beta_{\lambda,i}}\right\| \leq \begin{cases} 8\sqrt{L_\lambda \varepsilon / n} & L_\lambda/n > 2\varepsilon \\ 10\varepsilon & L_\lambda/n \leq 2\varepsilon. \end{cases}$$

Let $m_\lambda \in \{1, 2\}$ be the number of 0-cells (i.e. elements of Z_0) in the image of Φ_λ. The number m_λ will be 1 whenever the edge indexed by λ forms a self-loop, and will be equal to 2 otherwise. By repeatedly applying Lemma 82 to the curves $\alpha_{\lambda,i}$ (keeping λ fixed) we have that, since each $\alpha_{\lambda,i}$ has closed endpoints,

$$\mathrm{ECT}_{\alpha_\lambda}(v,t) = (m_\lambda - 2)\mathrm{ECT}_{\alpha_\lambda(0)}(v,t) + \sum_{i=1}^{n} \mathrm{ECT}_{\alpha_{\lambda,i}}(v,t) - \sum_{i=1}^{n-1} \mathrm{ECT}_{\alpha_{\lambda,i}(i/n)}(v,t),$$

for almost all t when v is fixed. The first term on the right side corresponds to the endpoint(s) of α_λ. The second term corresponds to the curves $\alpha_{\lambda,i}$, each of which has closed endpoints on either side. The final term corresponds to the endpoints of the $\alpha_{\lambda,i}$ which are not endpoints of α_λ. By the same argument, a similar equality holds for $\mathrm{ECT}_{\beta_\lambda}$. Hence, by the triangle inequality, we deduce that $\|\mathrm{ECT}_{\alpha_\lambda} - \mathrm{ECT}_{\beta_\lambda}\|$ is bounded above by

$$\begin{aligned} (m_\lambda - 2)\left\|\mathrm{ECT}_{\alpha_\lambda(0)} - \mathrm{ECT}_{\beta_\lambda(0)}\right\| + \sum_{i=1}^{n} \left\|\mathrm{ECT}_{\alpha_{\lambda,i}} - \mathrm{ECT}_{\beta_{\lambda,i}}\right\| + \\ \sum_{i=1}^{n-1} \left\|\mathrm{ECT}_{\alpha_{\lambda,i}(i/n)} - \mathrm{ECT}_{\beta_{\lambda,i}(i/n)}\right\| \\ \leq \begin{cases} 8\sqrt{L_\lambda n \varepsilon} + (n + m_\lambda - 3)\varepsilon & L_\lambda/n > 2\varepsilon \\ (11n + m_\lambda - 3)\varepsilon & L_\lambda/n \leq 2\varepsilon. \end{cases} \end{aligned}$$

In particular, this bound holds when we let

$$n = n_\lambda := \max\left(\left\lceil\left(\frac{M^2 L_\lambda^3}{24\varepsilon}\right)^{1/3}\right\rceil, \left\lceil\frac{L_\lambda M}{\pi}\right\rceil\right).$$

Applying Lemma 82 again, we have

$$\mathrm{ECT}_{f_\lambda}(v,t) = \mathrm{ECT}_{f_{\lambda-1}}(v,t) + \mathrm{ECT}_{\alpha_\lambda}(v,t) - \mathrm{ECT}_{\alpha_\lambda(0)}(v,t) - (m_\lambda - 1)\mathrm{ECT}_{\alpha_\lambda(1)}(v,t),$$

for almost all t when v is fixed. Similarly, such an equation holds involving g_λ, $g_{\lambda-1}$, and β_λ.

Applying the triangle inequality as before, along with our bound for $\|\mathrm{ECT}_{\alpha_\lambda} - \mathrm{ECT}_{\beta_\lambda}\|$, we deduce

$$\left\|\mathrm{ECT}_{f_\lambda} - \mathrm{ECT}_{g_\lambda}\right\| \leq \left\|\mathrm{ECT}_{f_{\lambda-1}} - \mathrm{ECT}_{g_{\lambda-1}}\right\| + \left\|\mathrm{ECT}_{\alpha_\lambda} - \mathrm{ECT}_{\beta_\lambda}\right\| + (2 - m_\lambda)\varepsilon.$$

The last two terms sum to $G_\lambda(\varepsilon) - \varepsilon$, so in particular we have the bound

$$\left\|\mathrm{ECT}_{f_\lambda} - \mathrm{ECT}_{g_\lambda}\right\| \leq |Z_0|\varepsilon + \sum_{k=1}^{\lambda} G_k(\varepsilon).$$

Induction then proves the proposition. □

From Proposition 74, Theorem 73 follows easily.

Proof of Theorem 73. Fix some $X, Y \in \mathcal{G}^r(Z^*, d)$ and suppose $d_{Z^*}(X,Y) < \varepsilon$. Hence, we may choose $h_X, h_Y \in \mathcal{E}^r(Z^*, d)$ with the properties given in Definition 72. Suppose that X has curvature bounded by M under Z^*. It follows that h_X also has curvature bounded by M. Proposition 74 gives that

$$\|\mathrm{ECT}_{h_X} - \mathrm{ECT}_{h_Y}\| \leq |Z_0|\varepsilon + \sum_{\lambda\in\Lambda} G_\lambda(\varepsilon),$$

but $\mathrm{ECT}_X = \mathrm{ECT}_{h_X}$ and $\mathrm{ECT}_Y = \mathrm{ECT}_{h_Y}$ since h_X and h_Y are homeomorphisms. The second statement of the theorem follows.

For the first statement, note that every $X \in \mathcal{G}^r(Z^*, d)$ has a bound M on its curvature and that $G_\lambda(\varepsilon) \to 0$ as $\varepsilon \to 0$ for all $\lambda \in \Lambda$. □

5.1.2 Stability of Piece-wise Linear Interpolation

If we are given $X \subseteq \mathbb{R}^d$, the C^2-image of a homeomorphism h from some one-dimensional CW complex Z, it may not be easy to exactly compute ECT_X. The main goal of this section is to show that a dense subset of Z can be used to approximate $\mathrm{ECT}_h = \mathrm{ECT}_X$. First, we make precise the kind of dense subset we need to properly estimate the ECT_h.

Definition 83. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a connected finite one-dimensional CW complex with some fixed cellular decomposition and f be a C^2 map $f : Z \to \mathbb{R}^d$. We say that $A = \{a_1, \ldots, a_n\} \subseteq Z$ is a* compatible subset *of Z^* if the following hold:*

1. *$Z_0 \subseteq A$ and*
2. *$A - Z_0$ contains a point in each 1-cell of Z.*

*These requirements ensure that $Z - A$ is a union of disjoint open intervals. If additionally, the length of the image of each of these intervals under f is less than ε, we say that A is an ε-*dense subset *for f. An infinite subset of Z is* compatible and dense *for f if it contains an ε-dense subset for all positive ε.*

Definition 84. *Let f be as in the previous definition and $A = \{a_1, \ldots, a_n\}$ be a compatible subset of Z^*. Let a_{ij} denote the line segment from a_i to a_j. We define a multiset E with elements in the set of unordered pairs in $\{1, \ldots, n\}$. E contains a copy of (i, j) for each open curve in $Z - A$ whose endpoints are a_i and a_j. We define*

$$\begin{aligned}\mathrm{ECT}_f^A(v,t) =& \#\{1 \le i \le n : \langle f(a_i), v\rangle \le t\} \\ &- \#\{(i,j) \in E : \max(\langle f(a_i), v\rangle, \langle f(a_j), v\rangle) \le t\}.\end{aligned}$$

The main theorem of the section says that we can use dense subsets to approximate the Euler characteristic transform of a one-dimensional CW complex:

Theorem 85. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a connected finite one-dimensional CW complex with some fixed cellular decomposition and f be a C^2 map $f : Z \to X \subseteq \mathbb{R}^d$. Suppose that f has curvature bounded by M on Z^* and let A be an ε-dense subset of Z^*, where $0 < \varepsilon < \pi/M$. Let L be the sum of the arc lengths of the images of 1-cells of Z under f. Then*

$$\left\|\mathrm{ECT}_f - \mathrm{ECT}_f^A\right\| \le \frac{1}{\sqrt{12}} M L \varepsilon.$$

In practice, the Theorem 85 implies that the ECT of a function on a one-dimensional CW complex can be computed approximately via a dense subset. The proof of this theorem is similar to the proof of Theorem 73, but requires two additional lemmas.

Lemma 86. *Using the notation Definition 84, let $b_i := f(a_i)$, and b_{ij} denote the line segment from b_i to b_j, and let c_i denote the number of pairs in E containing i. Then*

$$\mathrm{ECT}_f^A = \sum_{(i,j)\in E} \mathrm{ECT}_{b_{ij}} - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}.$$

Proof. Fix some v and t. If $\max(\langle b_i, v\rangle, \langle b_j, v\rangle) \leq t$, then $\mathrm{ECT}_{b_{ij}}(v,t) = \chi(b_{ij}) = 1$. If instead $\min(\langle b_i, v\rangle, \langle b_j, v\rangle) > t$, then $\mathrm{ECT}_{b_{ij}}(v,t) = \chi(\emptyset) = 0$. Otherwise, without loss of generality, suppose $\langle b_i, v\rangle \leq t$ and $\langle b_j, v\rangle > t$. Again, we have that $\mathrm{ECT}_{b_{ij}(v,t)}$ is equal to the Euler characteristic of a line segment, which is equal to 1.

Define the submultisets

$$\begin{aligned}
E_{\mathrm{up}} &= \{(i,j) \in E : \min(\langle b_i, v\rangle, \langle b_j, v\rangle) > t\},\\
E_{\mathrm{down}} &= \{(i,j) \in E : \max(\langle b_i, v\rangle, \langle b_j, v\rangle) \leq t\},\\
E_{\mathrm{mid}} &= \{(i,j) \in E : \max(\langle b_i, v\rangle, \langle b_j, v\rangle) > t,\ \min(\langle b_i, v\rangle, \langle b_j, v\rangle) \leq t\}.
\end{aligned}$$

Note $E = E_{\mathrm{up}} \sqcup E_{\mathrm{down}} \sqcup E_{\mathrm{mid}}$. Therefore,

$$\begin{aligned}
&\sum_{(i,j)\in E} \mathrm{ECT}_{b_{ij}}(v,t) - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t)\\
&= \sum_{(i,j)\in E_{\mathrm{up}}} \mathrm{ECT}_{b_{ij}}(v,t) + \sum_{(i,j)\in E_{\mathrm{down}}} \mathrm{ECT}_{b_{ij}}(v,t)\\
&\qquad + \sum_{(i,j)\in E_{\mathrm{mid}}} \mathrm{ECT}_{b_{ij}}(v,t) - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t)\\
&= \sum_{(i,j)\in E_{\mathrm{down}}} 1 + \sum_{(i,j)\in E_{\mathrm{mid}}} 1 - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t)\\
&= \sum_{(i,j)\in E_{\mathrm{down}}} (2-1) + \sum_{(i,j)\in E_{\mathrm{mid}}} 1 - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t)\\
&= \sum_{(i,j)\in E_{\mathrm{down}}} 2 + \sum_{(i,j)\in E_{\mathrm{mid}}} 1 - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t) - \sum_{(i,j)\in E_{\mathrm{down}}} 1\\
&= \sum_{i=1}^{n} c_i \mathrm{ECT}_{b_i}(v,t) - \sum_{i=1}^{n} (c_i - 1)\mathrm{ECT}_{b_i}(v,t) - \sum_{(i,j)\in E_{\mathrm{down}}} 1\\
&= \sum_{i=1}^{n} \mathrm{ECT}_{b_i}(v,t) - \sum_{(i,j)\in E_{\mathrm{down}}} 1 = \mathrm{ECT}_f^A(v,t).
\end{aligned}$$

□

Lemma 87. *Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be any differentiable function with increasing positive derivative satisfying $f(0) = 0$. For positive numbers L and ε consider the set*

$$S(L) = \left\{ (x_1, \ldots, x_k) \in \mathbb{R}^k : 0 \leq x_i \leq \varepsilon,\ \sum_{i=0}^{k} x_i = L \right\}.$$

If $S(L)$ is non-empty, then

$$\sum_{i=1}^{k} f(x_i) \leq L f(\varepsilon)/\varepsilon$$

on $S(L)$ (note that $S(L)$ is always non-empty if $k > L/\varepsilon$).

Proof. Let $a \leq c$ and $b \geq 0$. We have

$$f(b+c) - f(a+b) - \big(f(c) - f(a)\big) = \int_a^c f'(b+t) - f'(t)\, \mathrm{d}t \geq 0$$

and so

$$f(a+b) + f(c) \leq f(a) + f(b+c) \quad \text{when} \quad a \leq c \text{ and } b \geq 0. \tag{5.5}$$

Suppose $S(L)$ is non-empty. Since $S(L)$ is compact, f must attain a maximum on $S(L)$. Pick any such maximiser $x = (x_1, \ldots, x_k) \in S(L)$. By potentially reordering entries, we may assume the x_i are in decreasing order without affecting the value of $\sum_i f(x_i)$. Let j be the smallest index with $x_j \neq \varepsilon$ and l be the largest index with x_l not equal to zero.

If $l > j$, let $m = \min(x_l, \varepsilon - x_j)$. Equation (5.5) shows that if we replace x_j with $x_j + m$ and replace x_l with $x_l - m$, the value of $\sum_i f(x_i)$ does not decrease. Therefore, by applying this replacement procedure several times, we can always find a maximiser of $\sum_i f(x_i)$ on $S(L)$ with $j \geq l$. This condition forces the value of $\sum_i f(x_i)$ to be

$$\lfloor L/\varepsilon \rfloor f(\varepsilon) + f\big(L - \lfloor L/\varepsilon \rfloor \varepsilon\big).$$

If L is divisible by ε, the result is immediate. Otherwise, since f is convex,

$$\begin{aligned} f(L - \lfloor L/\varepsilon \rfloor \varepsilon) &= f\Big(\big(\lceil L/\varepsilon \rceil - L/\varepsilon\big)0 + \big(L/\varepsilon - \lfloor L/\varepsilon \rfloor\big)\varepsilon\Big) \\ &\leq \big(\lceil L/\varepsilon \rceil - L/\varepsilon\big) f(0) + \big(L/\varepsilon - \lfloor L/\varepsilon \rfloor\big) f(\varepsilon) \\ &= \big(L/\varepsilon - \lfloor L/\varepsilon \rfloor\big) f(\varepsilon). \end{aligned}$$

Thus

$$\lfloor L/\varepsilon \rfloor f(\varepsilon) + f\big(L - \lfloor L/\varepsilon \rfloor \varepsilon\big) \leq \lfloor L/\varepsilon \rfloor f(\varepsilon) + \big(L/\varepsilon - \lfloor L/\varepsilon \rfloor\big) f(\varepsilon) = L f(\varepsilon)/\varepsilon.$$

Since the value on the left is the maximum of $\sum_i f(x_i)$ on $S(L)$, we are done. □

Proof of Theorem 85. Adopting the notation of Lemma 86, for each $e = (i, j) \in E$, we let $b_e = b_{ij}$. Each $e \in E$ corresponds to some open interval in $Z - A$. We can always fix a finer CW structure $Z^\dagger = (Z, A, \{\Phi_e\}_{e\in E})$ of Z, and still have that $f \in \mathcal{F}^r(Z^\dagger, d)$. Let $\gamma_e = f \circ \Phi_e$. Induction on the number of elements in E with Lemma 82 applied to $Z^\dagger$ gives that for fixed v,

$$\mathrm{ECT}_f(v,t) = \sum_{e\in E} \mathrm{ECT}_{\gamma_e}(v,t) - \sum_{i=1}^{n}(c_i - 1)\mathrm{ECT}_{b_i}(v,t),$$

for almost all t.

Therefore, again fixing v and using Lemma 86,

$$\begin{aligned}
&\int_{\mathbb{R}} \left|\mathrm{ECT}_f(v,t) - \mathrm{ECT}_f^A(v,t)\right| \,\mathrm{d}t \\
=&\int_{\mathbb{R}} \Bigg| \sum_{e\in E} \mathrm{ECT}_{\gamma_e}(v,t) - \sum_{i=1}^{n}(c_i-1)\mathrm{ECT}_{b_i}(v,t) - \\
&\qquad \Big[\sum_{e\in E} \mathrm{ECT}_{b_e}(v,t) - \sum_{i=1}^{n}(c_i-1)\mathrm{ECT}_{b_i}(v,t)\Big]\Bigg| \,\mathrm{d}t \\
=&\int_{\mathbb{R}} \left| \sum_{e\in E} \mathrm{ECT}_{\gamma_e}(v,t) - \sum_{e\in E} \mathrm{ECT}_{b_e}(v,t)\right| \,\mathrm{d}t \\
\leq&\sum_{e\in E} \int_{\mathbb{R}} \left|\mathrm{ECT}_{\gamma_e}(v,t) - \mathrm{ECT}_{b_e}(v,t)\right| \,\mathrm{d}t.
\end{aligned}$$

Focusing on any particular $e = (i, j) \in E$, let d_1 be the minimum of $\langle \gamma_e(s), v\rangle$ over s, and d_4 be the maximum of the same function over s. Let $d_2 = \min(\langle b_i, v\rangle, \langle b_j, v\rangle))$ and $d_3 = \max(\langle b_i, v\rangle), \langle b_j, v\rangle))$. It follows that $d_1 \leq d_2 \leq d_3 \leq d_4$.

Since subsets of I and b_e always consist of contractible components, ECT_{γ_e} and ECT_{b_e} never have negative values. We have

$$\begin{aligned}
t \geq d_1 \quad &\Longrightarrow \quad \mathrm{ECT}_{\gamma_e}(v,t) \geq 1, \\
t \geq d_4 \quad &\Longrightarrow \quad \mathrm{ECT}_{\gamma_e}(v,t) = 1, \\
t < d_1 \quad &\Longrightarrow \quad \mathrm{ECT}_{\gamma_e}(v,t) = 0, \\
t \geq d_2 \quad &\Longrightarrow \quad \mathrm{ECT}_{b_e}(v,t) = 1, \\
t < d_2 \quad &\Longrightarrow \quad \mathrm{ECT}_{b_e}(v,t) = 0.
\end{aligned}$$

Combining these observations, we see

$$\begin{aligned}
\int_{\mathbb{R}} \left|\mathrm{ECT}_{\gamma_e}(v,t) - \mathrm{ECT}_{b_e}(v,t)\right| \,\mathrm{d}t &= \int_{d_1}^{d_4} \mathrm{ECT}_{\gamma_e}(v,t) - \mathrm{ECT}_{b_e}(v,t) \,\mathrm{d}t \\
&\leq \int_{d_1}^{d_4} \mathrm{ECT}_{\gamma_e}(v,t) \,\mathrm{d}t - (d_3 - d_2).
\end{aligned}$$

After applying a rotation, we may assume that $v = (0, 1, 0, \ldots, 0)$. After applying another rotation about v we may assume that b_e is parallel to the plane spanned by the first two coordinates. Let l_e be the arc length of γ_e. By Proposition 79, the length of b_e is at least $l_e - M^2 l_e^3/24$. Suppose that the line segment b_e meets the hyperplane perpendicular to v at an angle $\theta \in [0, \pi/2]$.

Applying Proposition 75 and Lemma 76 to this scenario, we observe

$$\int_{d_1}^{d_4} \mathrm{ECT}_{\gamma_e}(v,t)\,\mathrm{d}t - (d_3 - d_2) \leq \sqrt{l_e^2 - \left(l_e - \frac{M^2}{24} l_e^3\right)^2 \cos^2\theta} - \left(l_e - \frac{M^2}{24} l_e^3\right) \sin\theta.$$

We refer to the right side of this inequality as $f(\theta)$. Let $G = l_e - M^2 l_e^3/24$. G is positive since $l_e < \varepsilon < \pi/M < \sqrt{24}/M$. We have

$$f'(\theta) = \frac{G^2 \sin\theta\cos\theta}{\sqrt{l_e^2 - G^2\cos^2\theta}} - G\cos\theta.$$

A routine calculation shows that f' is either zero only when $\theta = \pi/2$ or for every θ. Meanwhile $f'(0) = -G$. Since this value is negative, f must be maximised at $\theta = 0$. Hence,

$$\begin{aligned}
\int_{\mathbb{R}} |\mathrm{ECT}_{\gamma_e}(v,t) - \mathrm{ECT}_{b_e}(v,t)|\;\mathrm{d}t &\leq \int_{d_1}^{d_4} \mathrm{ECT}_{\gamma_e}(v,t)\,\mathrm{d}t - (d_3 - d_2) \\
&\leq \sqrt{l_e^2 - \left(l_e - \frac{M^2}{24} l_e^3\right)^2} \\
&= \sqrt{\frac{M^2}{12} l_e^4 - \frac{M^4}{24^2} l_e^6} \\
&\leq \frac{M}{\sqrt{12}} l_e^2.
\end{aligned} \tag{5.6}$$

For $\lambda \in \Lambda$, let L_λ denote the arc length of $f \circ \Phi_\lambda$. Also for $\lambda \in \Lambda$, denote by $\Gamma(\lambda)$ the submultiset of $e \in E$ such that $\operatorname{im}\Phi_e$ is a subset of $\operatorname{im}\Phi_\lambda$. By Equations (5.6) and (5.6), along with Lemma 87, we get that

$$\begin{aligned}
\int_{\mathbb{R}} \left|\mathrm{ECT}_f(v,t) - \mathrm{ECT}_f^A(v,t)\right| \mathrm{d}t &\leq \frac{M}{\sqrt{12}} \sum_{e\in E} l_e^2 \\
&= \frac{M}{\sqrt{12}} \sum_{\lambda\in\Lambda} \sum_{e\in\Gamma(\lambda)} l_e^2 \\
&\leq \frac{M}{\sqrt{12}} \sum_{\lambda\in\Lambda} L_\lambda \varepsilon^2/\varepsilon \\
&= \frac{ML\varepsilon}{\sqrt{12}}.
\end{aligned}$$

Since this bound holds for any v, we are done. □

5.2 ECT stability for random data

In this section, we consider observations taken from an embedded finite one-dimensional CW complex Z which are perturbed by ambient Gaussian noise. We show that the Gaussian smoothing of these observations converges to satisfy the assumptions of Proposition 74. In particular, we show that the ECT and SECT of the Gaussian smoothing give consistent estimators of the ECT and SECT of Z, respectively. To provide the theorems, we first need to introduce technical conditions on the kernel we use in the Gaussian smoothing:

Definition 88 (Definition 5 in [55])**.** *Let Z be a topological space and $k : Z \times Z \to \mathbb{R}$ be a continuous kernel. Define*

$$d_k(t,s) = \sqrt{k(t,t) + k(s,s) - 2k(t,s)}.$$

For any $\varepsilon > 0$ let $N(Z, \varepsilon, d_k)$ be the minimal numbers of d_k-balls with radius ε needed to cover Z. Then define

$$J(Z, d_k) = \int_0^\infty \sqrt{\log N(Z, \varepsilon, d_k)}\, \mathrm{d}\varepsilon.$$

The standard metric on $[0,1]$ induces a metric d on any connected finite one-dimensional CW complexes via geodesics. The number $J(Z, d_k)$ is finite whenever d_k is sufficiently well-behaved with respect to d:

Proposition 89. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a connected finite one-dimensional CW complex. Let $k : Z \times Z \to \mathbb{R}$ be a continuous kernel. Suppose there are some positive constants C and δ such that whenever $d(t,s) < \delta$, we have*

$$k(t,t) - k(t,s) \leq C\big[d(t,s)\big]^2.$$

Then $J(Z, d_k) < \infty$.

The proof of this fact is deferred to Appendix B.

Definition 90. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a connected finite one-dimensional CW complex with some fixed cellular decomposition. Let $k : Z \times Z \to \mathbb{R}$ be a continuous kernel. We say k is r-times differentiable on Z^* if*

1. *for each $\lambda \in \Lambda$ the map $k^\lambda : I \times I \to \mathbb{R}$ given by $(s,t) \mapsto k(\Phi_\lambda(s), \Phi_\lambda(t))$ is r-times continuously differentiable and*

2. *for each* $\lambda \in \Lambda$ *and* $z \in Z$ *the map* $k^{\lambda,z} : I \to \mathbb{R}$ *given by* $s \mapsto k(\Phi_\lambda(s), z)$ *is* r*-times continuously differentiable.*

Differentiability is defined by one-sided limits at the boundaries of $I \times I$ *and* I*.*

Remark 13. *For a given connected finite 1 one-dimensional CW complex* $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ *with fixed cellular composition there is a straightforward way to construct an* r*-times differentiable kernel on* Z^**: let* $f : Z \to \mathbb{R}^d$ *be a continuous function such that* $f \circ \Phi_\lambda$ *is* r*-times differentiable for each* $\lambda \in \Lambda$*. Then if* k *is an* r*-times differentiable kernel on* $\mathbb{R}^d$*, it follows that* $k'(s,t) := k(f(s), f(t))$ *is an* r*-times differentiable kernel on* Z *by the chain rule.*

While it might be tempting to define a geodesic distance on Z *and then apply a stationary kernel (such as the Gaussian kernel) to this distance, it should be noted that, even in the case of* Z *being a manifold, the resulting function does not give a positive-definite kernel in general [39].*

We can now state the first theorem of this section:

Theorem 91. *Let* $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ *be a connected finite one-dimensional CW complex with some fixed cellular structure. Let* $k : Z \times Z \to \mathbb{R}$ *be a continuous, four-times differentiable kernel on* Z^**. Assume* k *satisfies* $J(Z, d_k) < \infty$*.*

Let $f : Z \to \mathbb{R}$ *be a function in the RKHS of* k*. Let* $\mathbf{a} \subset Z$ *be a sequence which is dense. Denote by* $\mathbf{a}_n$ *the first* n *terms of* $\mathbf{a}$ *and by* a_n *the* n*-th term of* $\mathbf{a}$*. Let* $\hat{f}_n$ *denote the Gaussian smoothing of* f *based on observations* $y_i = f(a_i) + \zeta_i$ *using kernel* k*, where* $i = 1, ..., n$ *and* $\zeta_i \sim \mathcal{N}(0, \sigma)$ *i.i.d. for some* $\sigma > 0$*. Then*

$$\mathbb{E}\left[\left\|\hat{f}_n(t, \mathbf{a}_n, f) - f(t)\right\|_\infty\right] \to 0$$

as $n \to \infty$*. Moreover, for each* $\lambda \in \Lambda$ *define* $\hat{f}_{n,\lambda}(t) = \hat{f}_n(\Phi_\lambda(t))$ *and* $f_\lambda(t) = f(\Phi_\lambda(t))$*. Then*

$$\mathbb{E}\left[\left|V\left(\hat{f}_{n,\lambda}\right) - V\left(f_\lambda\right)\right|^2\right] \to 0$$

on each 1-cell of Z^* *as* $n \to \infty$*; i.e. the variation of* $\hat{f}_{n,\lambda}$ *converges to the variation of* f_λ *in mean square.*

When proving the above result, we write k_x and k_y for the partial derivatives in the first and second components, respectively, and K_x and K_y for their corresponding Gram matrices. In particular, for fixed $t \in I$, $\lambda \in \Lambda$, and $\mathbf{a}$, we write $K_x^\lambda(t, \mathbf{a}_n) = [k_x^{\lambda,a_1}(t), ..., k_x^{\lambda,a_n}(t)]$ and K_y^λ for its transpose. A repeated subscript indicates repeated differentiation in that variable.

Let $g : X \to \mathbb{R}$ be a GP with kernel k and a deterministic function $h : X' \to X$. Then $g \circ h$ is a GP with kernel $k'(x, y) := k(h(x), h(y))$ for all $x, y \in X'$. This insight immediately follows from the definition of a GP in Definition 19. In particular, for the GP f in the statement of this theorem and any $\lambda \in \Lambda$, the composition $f \circ \Phi_\lambda$ is a GP for any number of observations n.

The derivative of a Gaussian process on I with a differentiable kernel is almost surely differentiable. As differentiation is a linear operator, the derivative of a Gaussian process is again a Gaussian process in such a case [75]. In particular, this derivative GP has kernel k_{xy} and for any $t \in I$ we have the joint distribution

$$\begin{bmatrix} g(t) \\ g'(t) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu(t) \\ \mu'(t) \end{bmatrix}, \begin{bmatrix} k(t,t) & k_y(t,t) \\ k_x(t,t) & k_{xy}(t,t) \end{bmatrix} \right). \tag{5.7}$$

In Theorem 91, we consider the GP regression of f_λ based on observations at $\mathbf{a}$ for fixed λ. Even if not all elements in the sequence $\mathbf{a}$ need to be in the image of Φ_λ, the GP posterior pre-composed with Φ_λ defines a GP on I. We are interested in the convergence of the derivative of this GP.

For fixed $\lambda \in \Lambda$, we denote the variance of this derivative GP at $t \in I$ by $v'_{n,\lambda}(t)$ (which is not the same as the derivative of $v_{n,\lambda}(t)$ in t). In particular, we have

$$v'_{n,\lambda}(t) = k^\lambda_{xy}(t,t) - K^\lambda_x(t, \mathbf{a}_n)(K(\mathbf{a}_n, \mathbf{a}_n) + \sigma^2 I)^{-1} K^\lambda_y(\mathbf{a}_n, t).$$

Lemma 92. *Given the Gaussian processes of Theorem 91, we get that for each $\lambda \in \Lambda$ the $v'_{n,\lambda}$ satisfy*

$$v'_{n,\lambda}(t) = \mathbb{E}\left[\left| \hat{f}'_{n,\lambda}(t, \mathbf{a}_n, f) - f'_\lambda(t) \right|^2 \right].$$

Furthermore, $v'_{n,\lambda}(t)$ is monotonically decreasing in n for all $t \in I$.

Proof. The first statement follows from Lemma 11 of [55]. In particular,

$$\begin{aligned}
\mathbb{E}_{f'_\lambda}\left[\left| \hat{f}'_{n,\lambda}(t) - f'_\lambda(t) \right|^2 \right] &= \mathbb{E}_{\boldsymbol{\zeta}_n}\left[\mathbb{E}_{f'_\lambda}\left[\left| \hat{f}'_{n,\lambda}(t) - f'_\lambda(t) \right|^2 \right] \middle| \mathbf{f}(\mathbf{a}_n) + \boldsymbol{\zeta}_n = \mathbf{y}_n \right] \\
&= \mathbb{E}_{\boldsymbol{\zeta}_n}\left[\mathbb{E}_{f'_\lambda}\left[\left| f'_\lambda(t) - \mathbb{E}_{f'_\lambda}[f'_\lambda(t) | \mathbf{f}(\mathbf{a}_n) + \boldsymbol{\zeta}_n] \right|^2 \right] \middle| \mathbf{f}(\mathbf{a}_n) + \boldsymbol{\zeta}_n = \mathbf{y}_n \right] \\
&= \mathbb{E}_{\boldsymbol{\zeta}_n}[\mathrm{Var}(f'_\lambda(t) | \mathbf{f}(\mathbf{a}_n) + \boldsymbol{\zeta}_n)] \\
&= \mathbb{E}_{\boldsymbol{\zeta}_n}[v'_{n,\lambda}(t)] = v'_{n,\lambda}(t).
\end{aligned}$$

For the second statement, we can write

$$v'_{n,\lambda}(t) = k^\lambda_{xy}(t,t) - K^\lambda_x(t, \mathbf{a}_{n+1}) \begin{bmatrix} B_n^{-1} & \mathbf{0}_{n\times 1} \\ \mathbf{0}_{1\times n} & 0 \end{bmatrix} K^\lambda_y(\mathbf{a}_{n+1}, t),$$

where $B_n := (K(\mathbf{a}_n, \mathbf{a}_n) + \sigma^2 I_n)$. Using the bordering method to obtain an expression for B_{n+1}^{-1} in terms of B_n, we get

$$
\begin{aligned}
&v'_{n,\lambda}(t) - v'_{n+1,\lambda}(t) \\
&= \nu^{-1} K_x^\lambda(t, \mathbf{a}_{n+1}) \begin{bmatrix} B_n^{-1} K(\mathbf{a}_n, a_{n+1}) K(a_{n+1}, \mathbf{a}_n) B_n^{-1} & -B_n^{-1} K(\mathbf{a}_n, a_{n+1}) \\ -K(a_{n+1}, \mathbf{a}_n) B_n^{-1} & 1 \end{bmatrix} K_y^\lambda(\mathbf{a}_{n+1}, t) \\
&= \nu^{-1}((K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, a_{n+1}))^2 - 2K_x^\lambda(t, a_{n+1}) K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, a_{n+1}) \\
&\qquad\qquad + k_x^\lambda(t, a_{n+1})^2) \\
&= \nu^{-1}(K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, a_{n+1}) - k_x^\lambda(t, a_{n+1}))^2,
\end{aligned}
$$

where $\nu := k(a_{n+1}, a_{n+1}) + \sigma - K(a_{n+1}, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, a_{n+1}) = v_n(a_{n+1}) + \sigma$ is the Schur complement of B_n inside B_{n+1}. As the ν is the Schur complement of a positive-definite matrix inside a positive-definite matrix, it is positive. As the second factor in the final line above is a square and thus positive too, we conclude that the sequence of functions $v'_{n,\lambda}(t)$ is monotonically decreasing. □

Proof of Theorem 91. The first statement follows from Theorem 8 in [55].

To prove the remainder of the theorem, we recall from Equation (5.7) that the covariance matrix of the distribution of $(f_\lambda(t), f'_\lambda(t))^T$ given n noisy observations of f is

$$
\begin{bmatrix} k(t,t) - K(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t) & k_y^\lambda(t,t) - K(t, \mathbf{a}_n) B_n^{-1} K_y^\lambda(\mathbf{a}_n, t) \\ k_x^\lambda(t,t) - K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t) & k_{xy}^\lambda(t,t) - K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K_y^\lambda(\mathbf{a}_n, t) \end{bmatrix}.
$$

As this matrix needs to be positive-definite, by taking the determinant and using the symmetry of k we get

$$
(k(t,t) - K(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t))(k_{xy}^\lambda(t,t) - K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K_y^\lambda(\mathbf{a}_n, t)) \tag{5.8}
$$

$$
\geq (k_x^\lambda(t,t) - K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t))^2 \geq 0. \tag{5.9}
$$

Thus, $k_x^\lambda(t,t) - K_x^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t) \to 0$ uniformly on I as the first factor of (5.8) converges uniformly by Proposition 10 of [55] and the second factor of (5.8) is bounded by the monotonicity established in Lemma 92 and the compactness of I. Repeating the same procedure with $\hat{f}''_{n,\lambda}$ in place of $\hat{f}'_{n,\lambda}$ gives $k_{xx}^\lambda(t,t) - K_{xx}^\lambda(t, \mathbf{a}_n) B_n^{-1} K(\mathbf{a}_n, t) \to 0$ uniformly: in this case, the second factor is $k_{xxyy}^\lambda(t,t) - K_{xx}^\lambda(t, \mathbf{a}_n) B_n^{-1} K_{yy}^\lambda(\mathbf{a}_n, t)$, which equals $v''_{n,\lambda}(t)$, the variance of the second derivative of the GP f_λ. We can show that $v''_{n,\lambda}(t)$ monotonically decreases by a proof analogous to the case $v'_{n,\lambda}(t)$ given in Lemma 92. For $v''_{n,\lambda}(t)$ to be well-defined we require k to be four times differentiable.

Then, by Jensen's inequality and Lemma 92, we can bound the expected value of the squared difference $V\left(\hat{f}_{n,\lambda}\right) - V(f_\lambda)$:

$$\mathbb{E}\left[\left|\int_0^1 |f'_\lambda(t)| - \left|\hat{f}'_{n,\lambda}(t,\mathbf{a}_n,f)\right| \,\mathrm{d}t\right|^2\right] \leq \mathbb{E}\left[\left(\int_0^1 \left||f'_\lambda(t)| - \left|\hat{f}'_{n,\lambda}(t,\mathbf{a}_n,f)\right|\right| \,\mathrm{d}t\right)^2\right]$$

$$\leq \mathbb{E}\left[\int_0^1 \left||f'_\lambda(t)| - \left|\hat{f}'_{n,\lambda}(t,\mathbf{a}_n,f)\right|\right|^2 \,\mathrm{d}t\right] = \int_0^1 \mathbb{E}\left[\left||f'_\lambda(t)| - \left|\hat{f}'_{n,\lambda}(t,\mathbf{a}_n,f)\right|\right|^2\right] \,\mathrm{d}t$$

$$= \int_0^1 v'_{n,\lambda}(t)\,\mathrm{d}t$$

$$= \left[k^\lambda_x(t,t) - K^\lambda_x(t,\mathbf{a}_n)B_n^{-1}K(\mathbf{a}_n,t) - \int_0^t k^\lambda_{xx}(s,s) - K^\lambda_{xx}(s,\mathbf{a}_n)B_n^{-1}K(\mathbf{a}_n,s)\,\mathrm{d}s\right]_0^1.$$

The final equation above converges to 0 as $n \to \infty$, as both the left-hand term and the function under in the integral of the right-hand term in the above difference converge uniformly to 0 by Equation (5.9) and its analogue for $v''_{n,\lambda}(t)$. □

It follows that the ECT of the Gaussian smoothing of f, denoted $\mathrm{ECT}_{\hat{f}_n}$, is a consistent estimator of the ECT of X:

Theorem 93. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a finite one-dimensional CW complex with some fixed cellular structure and $f : Z \to X \subseteq \mathbb{R}^d$ be a C^2 homeomorphism with bounded curvature. Further, assume that all components of f are functions in the RKHS of k, where k is a kernel satisfying the assumptions of Theorem 91. Moreover, assume that $\|f'_\lambda(t)\|_2 = L_\lambda$ is constant on all 1-cells $\lambda \in \Lambda$. Let $\mathbf{a}$ be a sequence in Z which is compatible with Z^* and dense for f. Let*

$$f(t) := \left(f^1(t), ..., f^d(t)\right)^T, \qquad \hat{f}_n := \left(\hat{f}^1_n, ..., \hat{f}^d_n\right)^T,$$

where for $j = 1, ..., d$ and $i = 1, ..., n$ the function $\hat{f}^j_n$ is the Gaussian smoothing of f^j given observations $y_{ij} = f^j(a_i) + \zeta_{ij}$ using kernel k and $\zeta_{ij} \sim \mathcal{N}(0, \sigma_j)$ i.i.d for some $\sigma_j > 0$. Then for each $\varepsilon > 0$

$$\lim_{n\to\infty} \mathbb{P}\left(\left\|\mathrm{ECT}_{\hat{f}_n} - \mathrm{ECT}_f\right\| < \varepsilon\right) \to 1.$$

Note that as f is a homeomorphism, $\mathrm{ECT}_f = \mathrm{ECT}_{\mathrm{im}\, f} = \mathrm{ECT}_X$ and we thus have constructed a consistent estimator for ECT_X. If for given observations $\mathbf{y}_n$ the curvature of $\hat{f}_n$ is bounded on each 1-cell, we can approximate $\mathrm{ECT}_{\hat{f}_n}$ by $\mathrm{ECT}^{\mathbf{a}_m}_{\hat{f}_n}$ arbitrarily closely for a sufficiently large m by Theorem 85. We conjecture that for sufficiently well-behaved kernels k, $\mathrm{ECT}^{\mathbf{a}_m}_{\hat{f}_n}$ converges to $\mathrm{ECT}_{\hat{f}_n}$ in probability, when m is some function in n. Proving this conjecture will involve bounding the curvature with high probability and is beyond the scope of this work.

Lemma 94. *Let f and $\hat{f}_n$ be as in the statement of Theorem 93. Denote the arc-lengths of $\hat{f}_{\lambda,n} := \hat{f}_n \circ \Phi_\lambda$ and $f_\lambda := f \circ \Phi_\lambda$ by $L_{n,\lambda}$ and L_λ respectively for each $\lambda \in \Lambda$. Then $L_{n,\lambda} \to L_\lambda$ and*

$$\int_0^1 \left| \left\| f'_{n,\lambda}(t) \right\|_2 - \left\| f'_\lambda(t) \right\|_2 \right| \mathrm{d}t \to 0$$

in probability.

Proof. First, note that $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$ and $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x-y|}$ for all $x, y \geq 0$. We have

$$\begin{aligned}
\left| \int_0^1 \left\| f'_{n,\lambda}(t) \right\|_2 - \left\| f'_\lambda(t) \right\|_2 \mathrm{d}t \right| &\leq \int_0^1 \left| \left\| f'_{n,\lambda}(t) \right\|_2 - \left\| f'_\lambda(t) \right\|_2 \right| \mathrm{d}t \\
&\leq \int_0^1 \sqrt{\left| \left\| f'_{n,\lambda}(t) \right\|_2^2 - \left\| f'_\lambda(t) \right\|_2^2 \right|} \mathrm{d}t \\
&\leq \sum_{j=1}^d \int_0^1 \sqrt{\left| |(f^j_{n,\lambda})'(t)|^2 - |(f^j_\lambda)'(t)|^2 \right|} \mathrm{d}t \\
&\leq \sum_{j=1}^d \left[\int_0^1 |(f^j_{n,\lambda})'(t)| + |(f^j_\lambda)'(t)| \mathrm{d}t \right] \\
&\qquad \cdot \left[\int_0^1 \left| |(f^j_{n,\lambda})'(t)| - |(f^j_\lambda)'(t)| \right| \mathrm{d}t \right].
\end{aligned}$$

The first inequality follows from the triangle inequality for integrals. The second and third inequalities follow from the inequalities for square roots introduced at the start of the proof. The final inequality is the Cauchy-Schwarz inequality for integrals.

The first factor in the final line converges to $2V(f_\lambda)$ and the second factor converges to 0 for each $j = 1, ..., d$ by Theorem 91 in probability. □

Proof of Theorem 93. We recall that convergence in mean implies convergence in probability. By applying Theorem 91 to each component of f, we find that $\hat{f}_n$ converges to f in mean in the ∞-norm. Note that $\hat{f}_{n,\lambda}$ need not be parameterised to constant velocity. Denote the arc length of $\hat{f}_{n,\lambda}$ by $L_{n,\lambda}$ and the arc length of f_λ by L_λ. By Lemma 94, $L_{n,\lambda} \to L_\lambda$ in probability for each $\lambda \in \Lambda$. Let $s_{n,\lambda}$ be the re-parametrisation of $\hat{f}_{n,\lambda}$ to constant-velocity on I, which is given by

$$s_{n,\lambda}(t) = \frac{1}{L_{n,\lambda}} \int_0^t \left\| \hat{f}'_{n,\lambda}(x) \right\|_2 \mathrm{d}x$$

and satisfies $\left\| \left(\hat{f}_{n,\lambda} \circ s^{-1} \right)'(x) \right\|_2 = 1$ on I. Thus,

$$|s_{n,\lambda}(t) - t| = \left| \int_0^t \left\| \hat{f}'_{n,\lambda}(x) \right\|_2 / L_{n,\lambda} - 1 \, \mathrm{d}x \right|$$

$$\leq \left| \int_0^t \left| \left\| \hat{f}'_{n,\lambda}(x) \right\|_2 / L_{n,\lambda} - \left\| f'_\lambda(x) \right\|_2 / L_\lambda \right| \mathrm{d}x \right|$$

$$\leq \frac{1}{L_{n,\lambda}} \int_0^1 \left| \left\| f'_{n,\lambda}(x) \right\|_2 - \left\| f'_\lambda(x) \right\|_2 \right| \mathrm{d}x + \left\| f'_\lambda(x) \right\|_2 \left| \frac{1}{L_{\lambda,n}} - \frac{1}{L_\lambda} \right|. \tag{5.10}$$

As both terms in Equation (5.10) converge to 0 in probability independently of t by Lemma 94, we get $\|s_{n,\lambda}(t) - t\|_\infty \xrightarrow{p.} 0$.

We then define $s_n : Z \to Z$ as

$$s_n(z) = \begin{cases} z & \text{if } z \in Z_0, \\ \left(\Phi_\lambda \circ s_{n,\lambda}^{-1} \circ \Phi_\lambda^{-1}\right)(z) & \text{if } z \in \Phi_\lambda((0,1)). \end{cases}$$

The map s_n is continuous as each $s_{n,\lambda}^{-1}$ is continuous, $s_{n,\lambda}^{-1}(0) = 0$ and $s_{n,\lambda}^{-1}(1) = 1$.

The result of the theorem then follows from Proposition 74:

$$\left\| \mathrm{ECT}_{\hat{f}_n} - \mathrm{ECT}_f \right\| \leq \left\| \mathrm{ECT}_{\hat{f}_n} - \mathrm{ECT}_{\hat{f}_n \circ s_n} \right\| + \left\| \mathrm{ECT}_{\hat{f}_n \circ s_n} - \mathrm{ECT}_f \right\|. \tag{5.11}$$

Note that the first term is 0 as re-parametrisation does not change the image of a function. For the second term, we find that $\hat{f}_n \circ s_n$ converges to satisfy the conditions such that Proposition 74 yields increasingly tight bounds: the arc lengths of $\hat{f}_n \circ s_n \circ \Phi_\lambda = \hat{f}_{n,\lambda} \circ s_{n,\lambda}^{-1}$ converge to those of $f \circ \Phi_\lambda = f_\lambda$ by Lemma 94 (the composition of f_λ with $s_{n,\lambda}^{-1}$ does not change its arc length). Further, both aforementioned functions have constant velocity and

$$\left\| \hat{f}_n \circ s_n - f \right\|_\infty \leq \left\| \hat{f}_n \circ s_n - \hat{f}_n \right\|_\infty + \left\| \hat{f}_n - f \right\|_\infty \xrightarrow{p.} 0.$$

In the above, the second term converges in probability by Theorem 91. The first term converges in probability as

$$\left\| \hat{f}_{n,\lambda} \circ s_{n,\lambda}^{-1} - \hat{f}_{n,\lambda} \right\|_\infty \leq \left\| \hat{f}_{n,\lambda} \circ s_{n,\lambda}^{-1} - f_\lambda \circ s_{n,\lambda}^{-1} \right\|_\infty + \left\| f_\lambda \circ s_{n,\lambda}^{-1} - f_\lambda \right\|_\infty + \left\| f_\lambda - \hat{f}_{n,\lambda} \right\|_\infty \xrightarrow{p.} 0$$

on each 1-cell $\lambda \in \Lambda$. The first term converges in probability by Theorem 91 (as re-parametrisation does not change the ∞-norm). Note that f_λ is continuous on I, which is compact, and therefore uniformly continuous. The second term equals $\left\| f_\lambda \circ s_{n,\lambda} - f_\lambda \right\|_\infty$ by pre-composition with $s_{n,\lambda}(t)$ and thus converges by Equation (5.10) and the uniform continuity of f_λ. The last term converges in probability by Theorem 91. □

Furthermore, our consistency result extends to the SECT of X.

Lemma 95. *Define the* ECT *on some interval* $[-a, a]$. *Assume the distance between the* ECT*s of two shapes* X *and* Y *is* δ. *Then the distance between their* SECT*s is at most* $(2a+1)\delta$.

Proof. Fix $v \in \mathbb{S}^{d-1}$. Then

$$\|\mathrm{SECT}_X(v,\,\cdot\,) - \mathrm{SECT}_Y(v,\,\cdot\,)\|_1$$

$$= \int_{-a}^{a} \left| \int_{-a}^{t} \mathrm{ECT}_X(v,x) - \mathrm{ECT}_Y(v,x)\mathrm{d}x - \frac{t+a}{2a} \int_{-a}^{a} \mathrm{ECT}_X(v,x) - \mathrm{ECT}_Y(v,x)\mathrm{d}x \right| \mathrm{d}t$$

$$\leq \int_{-a}^{a} \int_{-a}^{t} |\mathrm{ECT}_X(v,x) - \mathrm{ECT}_Y(v,x)|\,\mathrm{d}x + \frac{t+a}{2a} \int_{-a}^{a} |\mathrm{ECT}_X(v,x) - \mathrm{ECT}_Y(v,x)|\,\mathrm{d}x\mathrm{d}t$$

$$\leq 2a\delta + \delta = (2a+1)\delta.$$

Since the above is independent of v, we are done. □

The main limitation of our results is that the topology of our embedded space is assumed to be known and the results only work for a restricted class of CW complexes. Extending our statistical estimator and the related results to perturbations in the topology of the underlying shape remains future work.

5.3 Examples

We now illustrate our methods by means of a simulated example. In our simulation, we focus on a single simple closed curve in $\mathbb{R}^2$ and sample different numbers of noisy points from the curve. Our curve has been constructed by judiciously choosing complex Fourier coefficients. The samples are then taken by evenly spaced evaluations of our curve and are corrupted by adding independent multivariate Gaussian noise with mean 0 and covariance $(0.002)^2 I_2$. The curve, together with the noisy samples, is visualised in Figure 5.3.

As a kernel in our Gaussian smoothing, we pick the *sine-squared exponential kernel*. Assuming our curve is parameterised by $\gamma : [0, 2\pi] \to \mathbb{R}^2$ with $\gamma(0) = \gamma(2\pi)$, it is given by

$$k(s,t) = \exp\left(-2\sin\left(\frac{s-t}{2}\right)^2\right).$$

It satisfies the conditions of Theorems 91 and 93 (see Lemma 96; it is infinitely differentiable as it is the composition of infinitely differentiable functions). Its RKHS contains the curve we generated (see Lemma 97).

In Figure 5.4, we visualise the SECT of our true curve (in a fixed direction) and compare it to the SECT of curves sampled from Gaussian process regression (GPR) posterior distributions based on 20, 50 and 100 noisy evaluations of our original curves, respectively. In addition, we plot the distributions of the distance (given by the norm

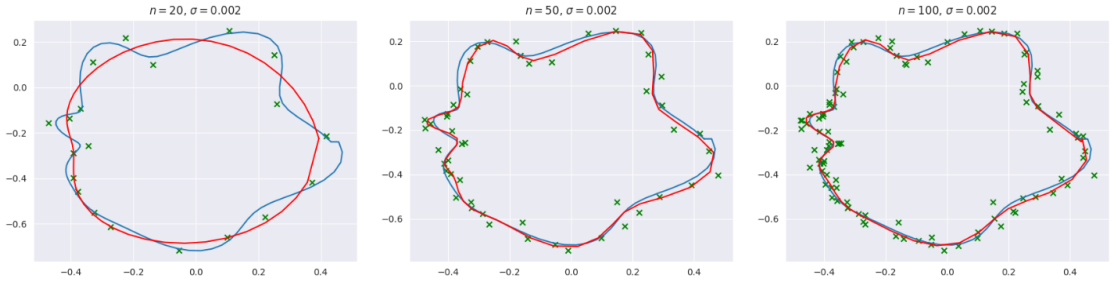

Figure 5.3: The Gaussian smoothings (red lines) of a simple closed curve (blue line) based on noisy samples (green crosses). The number of points is 20 on the left panel, 50 in the middle panel and 100 in the right panel. All points have been independently corrupted with mean zero Gaussian noise with standard deviation $\sigma = 0.002$ in each component.

introduced in Equation (2.4)) of the SECTs of the posterior samples with the SECT of the true curve.

In both types of plots, we see the posterior curves' mass moving closer to the true SECT, thereby illustrating the results of our theorems. We furthermore report that the distance between the SECT of the true curve and the SECTs of Gaussian smoothings are approximately 0.0627 ($n = 20$), 0.0366 ($n = 50$) and 0.0214 ($n = 100$), respectively. However, while Figure 5.4 illustrates that our results provide a consistent estimator of the SECT, the estimator need not be unbiased. In particular, in the upper left plot we see that the posterior curves' mass appears to be distributed near a curve which is distinct from that of the underlying SECT. This distinction between the posterior curves and underlying SECT suggests that our estimator is indeed biased.

The example presented here section illustrates how our estimator naturally gives rise to a posterior distribution over the space of SECT curves. We believe that there is potential to use this posterior distribution in a statistical inference or classification pipeline. Proving convergence rates for estimators like ours would help with quantifying the confidence of statistical ECT analyses.

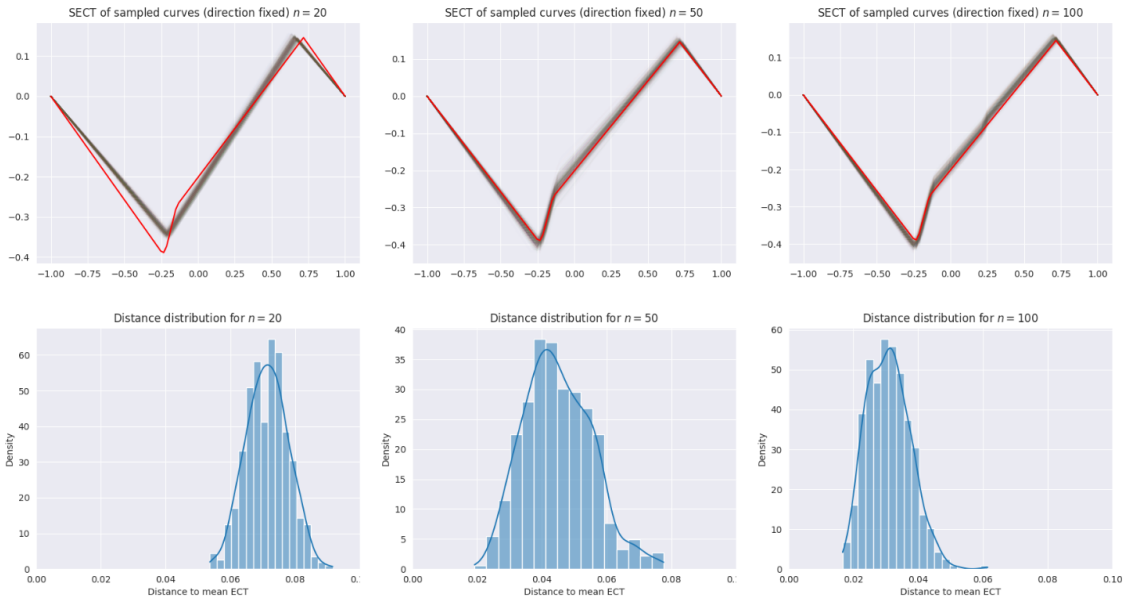

Figure 5.4: Top: The SECT in a fixed direction of the true shape (in red) compared to the SECTs of GPR posterior samples (in opaque blue) based on 20 (left), 50 (middle) and 100 (right) samples. The fixed direction corresponds to left-to-right in Figure 5.3. The SECTs are based on interpolations on the samples. Bottom: The distribution of the distance between the SECT of the true curve to SECTs of GPR posterior curves based on 20 (left), 50 (middle) and 100 (right) noisy samples from the underlying curve.

Chapter 6

Conclusions

We have studied several situations in TDA where by restricting to spaces of dimension one we can simplify complex algorithms and obtain results that are presently out of reach for more general shapes. In Chapter 3, where we study the TMD, this simplification is conceptual: instead of using persistent homology, whose definition is based on a classification theorem for persistence modules, we can instead imagine tearing apart a tree in a manner prescribed by simple rules, and in many cases get the same result. In cases where the TMD differs from persistent homology, the output of the TMD nevertheless has a biologically relevant interpretation which can be used to understand neuronal data. We provided a proof that the TMD can be recast as an application of persistent homology for trees with functions that increase along paths away from the root. We saw that one such function, the path distance from the soma, can be used to generate persistence images which identify key differences in healthy versus diseased neurons, and proved that persistence images arising in this way encode a variant of a neuron's Sholl function. We concluded the chapter with a novel stability result showing that persistence images are stable against a class of perturbations one might expect to result from measurement error. The results presented here help connect the TMD to the more widely used method of persistent homology, and provide theoretical and experimental justifications of its use in the practical setting of neuronal morphology.

In Chapters 4 and 5 we instead focused on low dimensional shapes to make progress on challenging problems in the theory of TDA. For Chapter 4, we studied the fiber of persistent homology, which can be viewed as a space measuring the amount of information lost by applying persistent homology to a function. We developed two different strategies for studying this space. The first of these strategies involves studying the mapping which sends a function to its critical sets. This methodology was effective for studying the fiber for continuous functions on geometric trees in Section 4.1, since we

were able to deduce the number of path components in the fiber of generic barcodes. In Section 4.2 we saw that when the underlying space is a manifold of dimension one instead of a tree, we also have the option of studying the fiber through the action of automorphisms of the underlying space on it. This observation proved extremely fruitful, as we were able to use conceptually the same methods to study the fiber for functions on any manifold, by restricting our function space to the more well-behaved space of Morse functions. In Section 4.1, critical sets were explicitly related to our methodology for studying the fiber of persistent homology. In the latter two sections, critical sets still play an important role as we must pay special attention to the action of endomorphisms (or diffeomorphisms) on critical points. This suggests that viewing future inverse problems in terms of permissible configurations of critical sets may be a useful conceptual approach for inverse problems which have yet to be studied.

Chapter 5 was motivated by the need to stably approximate the ECT of a shape in the presence of a noisy sample. To this end, in Section 5.1 we defined a novel metric on the space of embeddings of a fixed one dimensional CW complex into Euclidean space and showed that the ECT is a continuous map for this metric. This metric does not rely upon any simplicial structure of embeddings into Euclidean space, but instead on the differential properties of these embeddings. Then, in Section 5.2 we leveraged our stable metric to show that Gaussian processes can be used to regress noisy data in such a way that our regression is likely to have have the same ECT as the underlying embedding, and that the same is true for the SECT. We then verified our results on an example simple closed curve in Section 5.3. This chapter of the thesis suggests that in general it may be advisable to regress rather than interpolate noisy point cloud data in situations where we may expect error to accumulate from interpolation.

We conclude this chapter by mentioning a few problems that follow naturally from this work, ordered roughly by the structure of the thesis itself.

- **The TMD for branch order:** We have largely restricted our attention in Chapter 3 to the TMD for functions which encode geometric information regarding an underlying neuron. We could also consider the significance of the TMD for a combinatorial function, such as *branch order*. The branch order of a vertex v in a rooted tree is the number of branch points along the path from the root to v, excluding v if v is itself a branch point. Another relevant combinatorial function is the *subtree persistence* of a node, the maximum value of the branch order function on its child branches [47, 79]. We conjecture that

relevant quantities involving subtree persistence can be related by formulae to a tree's branch order TMD.

- **Generalising the TMD:** The TMD can be thought of as a variant of persistent homology for pointed spaces, where we require that the underlying space is a tree. Is there a construction generalizing the TMD to more complicated spaces? One possible obstruction to such a construction is that rooted trees induce orientations on every edge, but this is not true even for arbitrary graphs (consider for example the discretisation of the circle). Hence instead of trying to generalise the TMD to pointed spaces it may be more fruitful to instead generalise the construction to spaces equipped with a notion of flow.

- **The fiber of persistent homology for non-generic barcodes on trees:** How does the fiber change for trees when barcode endpoints can be equal? We suspect the answer changes depending on whether we allow left versus right endpoints to coincide. When left endpoints are allowed to be equal, we need to study merge trees with identical values on leaves. This causes an identifiability issue for local minima: we can no longer say that a particular local minimum is *the* local minimum corresponding to a certain leaf of a merge tree. We suspect that in these degenerate cases $\mathrm{MT}^{-1}(\mathcal{T})$ is homotopy equivalent to a quotient of $\mathrm{Conf}(X, \mathcal{T})$ by the action by a subgroup of the symmetric group corresponding to certain symmetries of $\mathcal{T}$. When right endpoints are allowed to be equal in a barcode B, it means that some merge trees with barcode B may have a node v with three or more children. For such a merge tree $\mathcal{T}$, we can still define a notion of a saddle corresponding to v, but it no longer needs to be connected in general[1], complicating our analysis considerably. We still suspect that the map $\mathrm{MT}^{-1}(\mathcal{T}) \to \mathrm{Conf}(X, \mathcal{T})$ induces an isomorphism on H_0, but it is no longer clear that this map is a homotopy equivalence. When both left and right endpoints are allowed to be equal, both complications must be handled simultaneously to study the fiber.

- **The fiber of persistent homology for point clouds:** One of the most common applications of persistent homology is to point cloud data in $\mathbb{R}^d$ via the Čech or Vietoris-Rips filtrations. It is not difficult to show that these filtrations produce barcodes that are invariant under the action of translations and rotations on point cloud data, an often desirable property. It is reasonable to

[1] In fact, in this more general case two functions can have the same local minima and the same merge tree and still differ in the number of path components in their saddle corresponding to v.

ask how many other ways there are to perturb a particular point cloud without affecting the resulting barcode: what is the local dimension of the fiber of persistent homology at a particular point cloud? The local dimension can be viewed as a measure of the ability of the persistence map to differentiate a point cloud from its neighbors.

- **Integral formulas for extended persistence:** In Proposition 75 we saw that the topology of the sublevel sets of a smooth function on the interval are constrained by its variation. Shortly thereafter it was remarked that for tame functions, the same result follows from the work of Biswas et al. [11]. In this work, the authors show that for any tame function f on the interval (or more generally a geometric tree), the sum of the lengths of the intervals in its extended persistence is exactly the variation of f. For a function on a higher dimensional manifold M, is there a formula relating an integral of a function of f over M to the sum of the lengths of intervals in the extended persistence barcode of f? Should such a formula not be possible, is it possible to bound the summed lengths of intervals in the extended persistence of f by an integral involving the derivatives of f? An affirmative answer to this question in any dimension greater than one would be the first step to extending the results of Chapter 5 to higher dimensions. Further, such a result might have fruitful applications in the stochastic topology of random functions.

Appendix A

Details of neuronal data acquisition

A.1 Tissue harvest and processing

Tau35 mice were generated by targeted knock-in to the *Hprt* locus under the control of the human tau promoter as described previously [12]. Mice, Tau35 and wild-type (WT) controls, were sacrificed at two time-points, presymptomatic (4 month old) and post-symptomatic (10-12 month old), using terminal anaesthesia, perfused with 1x PBS and post-fixed in 4% (w/v) PFA overnight at 4°C. Brain sections (200 μm) were prepared using a VT1000 S Vibrating blade microtome (Leica Biosystems) and stored free-floating in cryoprotectant (30% (v/v) ethylene glycol, 15% (w/v) sucrose in PBS) at -20°C. All procedures were carried out in accordance with the Animals (Scientific Procedures) Act, 1986, following approval by the local ethical review committee.

A.2 Tissue staining and image acquisition

To quantify dendritic morphology, Golgi-stained neurons [76] were imaged using a 60x objective (NA=1.4) on a Nikon microscope. Using the live acquisition feature images were collected at a depth starting at 20μm below the surface of the specimen. Z-stack images (30-90μm total on the Z-axis; Z-stack step size=0.3μm; 90-270 images per stack) were acquired. Each image stack was extracted using NIS-Elements (Nikon) software and imported to Neurolucida (MBF Bioscience) software for analysis. Neurolucida was used to identify cell bodies and their outgrowth for each of the sections in the z-stack for stacks below 50μm. For stacks above 50μm, due to the image size, an average intensity projection image was generated by calculating the average intensity (AIP) values of each pixel along the z-axis for all the layers and combining them into the 2D AIP image. The resulting AIP images were then used to identify cell bodies and outgrowth with the Neurolucida software.

Appendix B

Characterisation of the sine-squared exponential kernel

Before focusing on the sine-squared kernel, we provide general conditions on a kernel for the number $J(Z, d_k)$ to be finite.

Proposition 89. *Let $Z^* = (Z, Z_0, \{\Phi_\lambda\}_{\lambda\in\Lambda})$ be a connected finite one-dimensional CW complex. Let $k : Z \times Z \to \mathbb{R}$ be a continuous kernel. Suppose there are some positive constants C and δ such that whenever $d(t, s) < \delta$, we have*

$$k(t, t) - k(t, s) \leq C\big[d(t, s)\big]^2.$$

Then $J(Z, d_k) < \infty$.

Proof. Let $B^d_\varepsilon(t)$ denote an ε ball around $t \in Z$ with respect to d. Analogously define $B^{d_k}_\varepsilon(t)$. We have that if $d(t, s) \leq \delta$

$$\begin{aligned} d_k(t, s) &= \sqrt{k(t, t) + k(s, s) - 2k(t, s)} \\ &= \sqrt{k(t, t) - k(t, s) + k(s, s) - k(t, s)} \\ &\leq \sqrt{2C}d(t, s). \end{aligned}$$

For convenience, we let $C_0 = \frac{1}{\sqrt{2C}}$. The inequality $d_k(t, s) \leq \frac{1}{C_0}d(t, s)$ implies the following for $\varepsilon < \delta/C_0$:

$$\begin{aligned} d(t, s) < C_0\varepsilon \quad &\Rightarrow \quad d_k(t, s) < \varepsilon \\ B^d_{C_0\varepsilon}(t) \quad &\subseteq \quad B^{d_k}_\varepsilon(t) \\ N(Z, \varepsilon, d_k) \quad &\leq \quad N(Z, C_0\varepsilon, d). \end{aligned}$$

Since Z is compact, the function $k(t,s)$ is bounded above by some M on $Z \times Z$. Hence the function d_k is bounded above by $\sqrt{2M}$. Therefore, $N(Z,\varepsilon,d_k) = 1$ for $\varepsilon > \sqrt{2M}$. By potentially increasing M so that $\sqrt{2M} \geq \delta/C_0$ we have

$$\begin{aligned}\int_0^\infty \sqrt{\log N(Z,\varepsilon,d_k)}\,\mathrm{d}\varepsilon &= \int_0^{\sqrt{2M}} \sqrt{\log N(Z,\varepsilon,d_k)}\,\mathrm{d}\varepsilon \\ &= \int_0^{\delta/C_0} + \int_{\delta/C_0}^{\sqrt{2M}}.\end{aligned}$$

Note that since k is continuous, d_k is continuous. Hence, d_k-balls are open in the topology on Z. A standard compactness argument shows that $N(Z,\delta/C_0,d_k)$ is finite. Therefore, the integral

$$\int_{\delta/C_0}^{\sqrt{2M}} \sqrt{\log N(Z,\varepsilon,d_k)}\,\mathrm{d}\varepsilon \leq N(Z,\delta/C_0,d_k)(\sqrt{2M} - \delta/C_0)$$

is finite.

It only remains to show that the integral

$$\int_0^{\delta/C_0} \sqrt{\log N(Z,\varepsilon,d_k)}\,\mathrm{d}\varepsilon \leq \int_0^{\delta/C_0} \sqrt{\log N(Z,C_0\varepsilon,d)}\,\mathrm{d}\varepsilon$$

is finite.

However,

$$N(Z,\varepsilon,d) \leq |\Lambda| N([0,1],\varepsilon,d) + |Z_0|,$$

and

$$N([0,1],\varepsilon,d) \leq \left\lceil \frac{1}{\varepsilon} \right\rceil \leq \frac{1}{\varepsilon} + 1,$$

so

$$N(Z,\varepsilon,d) \leq \frac{|\Lambda|}{\varepsilon} + |\Lambda| + |Z_0|.$$

Therefore

$$\begin{aligned}\int_0^{\delta/C_0} \sqrt{\log N(Z,C_0\varepsilon,d)}\,\mathrm{d}\varepsilon &\leq \int_0^{\delta/C_0} \sqrt{\log\left[\frac{|\Lambda|}{\varepsilon} + |\Lambda| + |Z_0|\right]}\,\mathrm{d}\varepsilon \\ &= \frac{1}{|\Lambda|}\int_{|\Lambda|C_0/\delta}^\infty \frac{\sqrt{\log\left[x + |\Lambda| + |Z_0|\right]}}{x^2}\,\mathrm{d}x\end{aligned}$$

However for x sufficiently large, $\log(x + |\Lambda| + |Z_0|) \leq x^{1/2}$, and so the above integral converges as the integral

$$\int_b^\infty \frac{1}{x^{3/2}}\,\mathrm{d}x$$

converges for all positive b. □

Remark 14. *A similar result could also be proven using Proposition 7 from [55].*

Now we focus on the sine-squared kernel.

Lemma 96. *For the sine-squared kernel, we have* $J(\mathbb{S}^1, d_k) < \infty$.

Proof. Let k denote the sine-squared kernel. For $s - t$ approaching zero, we have

$$k(s,t) = \exp\left(-2\sin\left(\frac{s-t}{2}\right)^2\right) = 1 - \frac{(s-t)^2}{2} + O\big((s-t)^4\big).$$

Thus when $s - t$ is sufficiently close to zero we have

$$k(t,t) - k(s,t) \leq (s-t)^2$$

and so by Proposition 89 $J(\mathbb{S}^1, d_k) < \infty$. □

Lemma 97. *Define the Hilbert space* $\mathcal{H}'$ *of sequences* $w_{ab} \in \mathbb{R}$, $a, b \in \mathbb{N}_0$, *satisfying*

$$\sum_{n=0}^{\infty} n! \sum_{\substack{a\geq 0,\, b\geq 0 \\ a+b=n}} \frac{w_{ab}^2}{C_a^n} < \infty,$$

where C_a^n *denotes* n *choose* a. *For* $\{w_{ab}\}, \{v_{ab}\} \in \mathcal{H}'$, *the inner-product of* $\mathcal{H}'$ *is given by*

$$\langle \{w_{ab}\}, \{v_{ab}\} \rangle_{\mathcal{H}'} := \gamma \sum_{n=0}^{\infty} n! \sum_{\substack{a\geq 0,\, b\geq 0 \\ a+b=n}} \frac{w_{ab} v_{ab}}{C_a^n}$$

where $\gamma > 0$ *is a constant. Define* V *to be the closed subspace of sequences* $\{w_{ab}\} \in \mathcal{H}'$ *such that*

$$\sum_{(a,b)\in\mathbb{N}^2} w_{ab} \cos^a(t) \sin^b(t) = 0 \tag{B.1}$$

for all $t \in [0, 2\pi)$. *Then the Hilbert space* $\mathcal{H}$ *given by the functions*

$$f(t) = \sum_{(a,b)\in\mathbb{N}^2} w_{ab} \cos^a(t) \sin^b(t) \tag{B.2}$$

with $\{w_{ab}\} \in V^{\perp}$ *and inner-product induced from* $\mathcal{H}'$ *is isomorphic to the RKHS of the sine-squared-exponential kernel, denoted by* $\mathcal{H}_k$.

Proof. By using standard trigonometric identities, we see that the sine-squared kernel is proportional (by a positive constant) to the kernel

$$k(s,t) = \exp\left(\cos(s)\cos(t) + \sin(s)\sin(t)\right).$$

Thus, by using the Taylor expansion of exp, $k(\cdot, t) \in \mathcal{H}$ for all t with coefficients

$$w_{ab} = \frac{C_a^{a+b}}{(a+b)!} \cos^a(t) \sin^b(t).$$

Moreover, for $f \in \mathcal{H}$ with coefficients v_{ab} and fixed t, we get

$$\begin{aligned} \langle k(\cdot, t), f \rangle_{\mathcal{H}} &= \sum_{n=0}^{\infty} n! \sum_{\substack{a \geq 0, b \geq 0 \\ a+b=n}} \frac{C_a^n \cos^a(t) \sin^b(t) v_{ab}}{n! C_a^n} \\ &= \sum_{(a,b) \in \mathbb{N}^2} \cos^a(t) \sin^b(t) v_{ab} = f(t). \end{aligned} \tag{B.3}$$

Thus, the inner-product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ has the reproducing property and coincides with the inner-product induced by the kernel k (i.e. the inner-product of $\mathcal{H}_k$) given in Equation (2.6). Further, the coefficients of $k(\cdot, t)$ lie in $V^{\perp}$: let $\{v_{ab}\} \in V$ and let $\{w_{ab}\}$ be the coefficients of $k(\cdot, t)$. Then by Equation (B.3), $\langle \{v_{ab}\}, \{w_{ab}\} \rangle_{\mathcal{H}'} = 0$. As $V^{\perp}$ is closed as it is perpendicular to V, so is $\mathcal{H}$, implying that $\mathcal{H}$ is a Hilbert space. We have that $\mathcal{H}_k \subseteq \mathcal{H}$. As $\mathcal{H}_k$ is complete by definition, we get $\mathcal{H} = \mathcal{H}_k \oplus W$ for some closed subspace W. Let $f \in W$. Then $\langle g, f \rangle_{\mathcal{H}} = 0$ for all $g \in \mathcal{H}_k$ and in particular $f(t) = \langle k(\cdot, t), f \rangle_{\mathcal{H}} = 0$ for all $t \in [0, 2\pi)$. Thus, $W = 0$ and $\mathcal{H} \cong \mathcal{H}_k$. □

Lemma 98. *Every $f \in \mathcal{H}$ is continuous and the inclusion $\mathcal{H} \hookrightarrow C(\mathbb{S}^1, d_\infty)$ is continuous, where $C(\mathbb{S}^1, d_\infty)$ is the space of continuous real-valued functions on S^1 endowed with the ∞-norm. Further, $\cos(nt)$ and $\sin(nt)$ are elements of $\mathcal{H}$ for all $n \in \mathbb{N}$.*

Proof. Note that $\|k(\cdot, t)\|_{\mathcal{H}} = 1$ for all t. Thus, by the reproducing property of k and the Cauchy-Schwarz inequality, for all $f \in \mathcal{H}$ and any $t \in \mathbb{S}^1$ we get

$$|f(t)| = |\langle k(\cdot, t), f \rangle_{\mathcal{H}}| \leq \|f\|_{\mathcal{H}}.$$

Hence, convergence in the $\mathcal{H}$-norm implies convergence in the ∞-norm. As f can be written as a series of continuous functions converging in the $\mathcal{H}$-norm (c.f. Equation (B.2)), it follows that f is continuous. As for any $f \in \mathcal{H}$ with $\|f\|_{\mathcal{H}} \leq \varepsilon$ and any $t \in \mathbb{S}^1$ we have $|f(t)| \leq \varepsilon$, we get that the inclusion $\mathcal{H} \hookrightarrow C(\mathbb{S}^1, d_\infty)$ is continuous.

Further, we can expand

$$\cos(nt) = \sum_{k \text{ even}}^{n} (-1)^{\frac{k}{2}} \binom{n}{k} \cos^{n-k}(t) \sin^k(t),$$

$$\sin(nt) = \sum_{k \text{ odd}}^{n} (-1)^{\frac{k-1}{2}} \binom{n}{k} \cos^{n-k}(t) \sin^k(t).$$

Thus, $\cos(nt)$ and $\sin(nt)$ can be expanded as powers of cos and sin with coefficients in $\{w_{ab}\} \in \mathcal{H}'$ (as in Lemma 97). We can project these coefficients into $V^{\perp}$ without changing the value of our series at any $t \in \mathbb{S}^1$: the difference in the series we observe by subtracting from elements of V from $\{w_{ab}\}$ is 0 for all t (c.f. Equation (B.1)). Thus, $\cos(nt), \sin(nt) \in \mathcal{H}$ for all $n \in \mathbb{N}$. $\square$

Bibliography

[1] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: A stable vector representation of persistent homology. *The Journal of Machine Learning Research*, 18(1):218–252, 2017.

[2] Erik J Amezquita, Michelle Quigley, Tim Ophelders, Jacob Landis, Elizabeth Munch, Daniel Chitwood, and Daniel Koenig. Quantifying barley morphology using the euler characteristic transform. In *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*, 2020.

[3] Erik J Amézquita, Michelle Y Quigley, Tim Ophelders, Jacob B Landis, Daniel Koenig, Elizabeth Munch, and Daniel H Chitwood. Measuring hidden phenotype: Quantifying the shape of barley seeds using the euler characteristic transform. *in silico Plants*, 4(1):diab033, 2022.

[4] Giorgio A Ascoli. Mobilizing the base of neuroscience data: the case of neuronal morphologies. *Nature Reviews Neuroscience*, 7(4):318–324, 2006.

[5] Giorgio A. Ascoli, Lidia Alonso-Nanclares, Stewart A. Anderson, German Barrionuevo, Ruth Benavides-Piccione, Andreas Burkhalter, György Buzsáki, Bruno Cauli, Javier DeFelipe, Alfonso Fairén, Dirk Feldmeyer, Gord Fishell, Yves Fregnac, Tamas F. Freund, Daniel Gardner, Esther P. Gardner, Jesse H. Goldberg, Moritz Helmstaedter, Shaul Hestrin, Fuyuki Karube, Zoltán F. Kisvárday, Bertrand Lambolez, David A. Lewis, Oscar Marin, Henry Markram, Alberto Muñoz, Adam Packer, Carl C. H. Petersen, Kathleen S. Rockland, Jean Rossier, Bernardo Rudy, Peter Somogyi, Jochen F. Staiger, Gabor Tamas, Alex M. Thomson, Maria Toledo-Rodriguez, Yun Wang, David C. West, and Rafael Yuste. Petilla terminology: nomenclature of features of gabaergic interneurons of the cerebral cortex. *Nature Reviews Neuroscience*, 9(7):557–568, 2008.

[6] Serguei Barannikov. The framed Morse complex and its invariants. *American Mathematical Society Translations, Series 2*, 1994.

[7] David Beers, Despoina Goniotaki, Diane P Hanger, Alain Goriely, and Heather A Harrington. Barcodes distinguish morphology of neuronal tauopathy. *arXiv preprint arXiv:2204.03348*, 2022.

[8] David Beers, Heather A Harrington, and Alain Goriely. Stability of topological descriptors for neuronal morphology. *arXiv preprint arXiv:2211.09058*, 2022.

[9] David Beers and Jacob Leygonie. The fiber of persistent homology for trees. *arXiv preprint arXiv:2303.16176*, 2023.

[10] Nicolas Berkouk. Persistence and the sheaf-function correspondence. *arXiv preprint arXiv:2207.06335*, 2022.

[11] Ranita Biswas, Sebastiano Cultrera di Montesano, Herbert Edelsbrunner, and Morteza Saghafian. A window to the persistence of 1d maps. i: Geometric characterization of critical point pairs. *LIPIcs*, 2022.

[12] Marie K Bondulich, Tong Guo, Christopher Meehan, John Manion, Teresa Rodriguez Martin, Jacqueline C Mitchell, Tibor Hortobagyi, Natalia Yankova, Virginie Stygelbout, Jean-Pierre Brion, et al. Tauopathy induced by low level expression of a human brain-derived tau fragment in mice is rescued by phenylbutyrate. *Brain*, 139(8):2290–2306, 2016.

[13] Doug M Boyer, Yaron Lipman, Elizabeth St. Clair, Jesus Puente, Biren A Patel, Thomas Funkhouser, Jukka Jernvall, and Ingrid Daubechies. Algorithms to automatically quantify the geometric similarity of anatomical surfaces. *Proceedings of the National Academy of Sciences*, 108(45):18221–18226, 2011.

[14] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015.

[15] Peter Bubenik and Jonathan A Scott. Categorification of persistent homology. *Discrete & Computational Geometry*, 51(3):600–627, 2014.

[16] Chen Cai, Woojin Kim, Facundo Mémoli, and Yusu Wang. Elder-rule-staircodes for augmented metric spaces. *SIAM Journal on Applied Algebra and Geometry*, 5(3):417–454, 2021.

[17] Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.

[18] Mathieu Carrière, Steve Y Oudot, and Maks Ovsjanikov. Stable topological signatures for points on 3d shapes. In *Computer Graphics Forum*, volume 34, pages 1–12. Wiley Online Library, 2015.

[19] Michael J Catanzaro, Justin M Curry, Brittany Terese Fasy, Jānis Lazovskis, Greg Malen, Hans Riess, Bei Wang, and Matthew Zabka. Moduli spaces of morse functions for persistence. *Journal of Applied and Computational Topology*, 4(3):353–385, 2020.

[20] Jean Cerf. Topologie de certains espaces de plongements. *Bulletin de la Société Mathématique de France*, 89:227–380, 1961.

[21] Jean Cerf. La stratification naturelle des espaces de fonctions différentiables réelles et le théorème de la pseudo-isotopie. *Inst. Hautes Études Sci. Publ. Math.*, (39):5–173, 1970.

[22] Frédéric Chazal, Vin De Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014.

[23] Shiing Shen Chern. Curves and surfaces in euclidean space. *Studies in global geometry and analysis*, 4(1):967, 1967.

[24] Ilya Chevyrev, Vidit Nanda, and Harald Oberhauser. Persistence paths and signature features in topological data analysis. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):192–202, 2018.

[25] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.

[26] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using poincare and lefschetz duality. *Foundations of Computational Mathematics*, 9:pages79–103, 2008.

[27] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using poincaré and lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.

[28] Lorin Crawford, Anthea Monod, Andrew X Chen, Sayan Mukherjee, and Raúl Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.

[29] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.

[30] Justin Curry. The fiber of the persistence map for functions on the interval. *Journal of Applied and Computational Topology*, 2(3):301–321, 2018.

[31] Justin Curry, Sayan Mukherjee, and Katharine Turner. How many directions determine a shape and other sufficiency results for two topological transforms. *arXiv preprint arXiv:1805.09782*, 2018.

[32] Jacek Cyranka, Konstantin Mischaikow, and Charles Weibel. Contractibility of a persistence map preimage. *Journal of Applied and Computational Topology*, 4(4):509–523, 2020.

[33] Jacek Cyranka, Konstantin Mischaikow, and Charles Weibel. Contractibility of a persistence map preimage. *Journal of Applied and Computational Topology*, 4(4):509–523, 2020.

[34] Javier DeFelipe, Pedro L López-Cruz, Ruth Benavides-Piccione, Concha Bielza, Pedro Larrañaga, Stewart Anderson, Andreas Burkhalter, Bruno Cauli, Alfonso Fairén, Dirk Feldmeyer, et al. New insights into the classification and nomenclature of cortical gabaergic interneurons. *Nature Reviews Neuroscience*, 14(3):202–216, 2013.

[35] Paweł Dłotko and Davide Gurnari. Euler characteristic curves and profiles: a stable shape invariant for big data problems. *arXiv preprint arXiv:2212.01666*, 2022.

[36] Claire Donnat, Axel Levy, Frederic Poitevin, and Nina Miolane. Deep generative modeling for volume reconstruction in cryo-electron microscopy. *arXiv preprint arXiv:2201.02867*, 2022.

[37] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.

[38] Michael Farber. Configuration spaces and robot motion planning algorithms. In *Combinatorial And Toric Homotopy: Introductory Lectures*, pages 263–303. World Scientific, 2018.

[39] Aasa Feragen, Francois Lauze, and Soren Hauberg. Geodesic exponential kernels: When curvature and linearity conflict. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3032–3042, 2015.

[40] Eberhard E Fetz and Bengt Gustafsson. Relation between shapes of post-synaptic potentials and changes in firing probability of cat motoneurones. *The Journal of physiology*, 341(1):387–410, 1983.

[41] Tingran Gao, Shahar Z Kovalsky, and Ingrid Daubechies. Gaussian process landmarking on manifolds. *SIAM Journal on Mathematics of Data Science*, 1(1):208–236, 2019.

[42] Ellen Gasparovic, Elizabeth Munch, Steve Oudot, Katharine Turner, Bei Wang, and Yusu Wang. Intrinsic interleaving distance for merge trees. *arXiv preprint arXiv:1908.00063*, 2019.

[43] Robert Ghrist. Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1):61–75, 2008.

[44] Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and euler integral transforms. *Journal of Applied and Computational Topology*, 2(1-2):55–60, 2018.

[45] Wesley B Grueber, Lily Y Jan, and Yuh Nung Jan. Tiling of the drosophila epidermis by multidendritic sensory neurons. *Development*, 129, 2002.

[46] Mary-Elizabeth Hamstrom et al. Homotopy in homeomorphism spaces, TOP and PL. *Bulletin of the American Mathematical Society*, 80(2):207–230, 1974.

[47] Edouard Hannezo, Colinda LGJ Scheele, Mohammad Moad, Nicholas Drogo, Rakesh Heer, Rosemary V Sampogna, Jacco Van Rheenen, and Benjamin D Simons. A unifying theory of branching morphogenesis. *Cell*, 171(1):242–255, 2017.

[48] Morris W Hirsch. *Differential topology*, volume 33. Springer Science & Business Media, 2012.

[49] Sungbok Hong, John Kalliongis, Darryl McCullough, and J Hyam Rubinstein. *Diffeomorphisms of elliptic 3-manifolds*, volume 2055. Springer, 2012.

[50] Gregory SXE Jefferis, Christopher J Potter, Alexander M Chan, Elizabeth C Marin, Torsten Rohlfing, Calvin R Maurer Jr, and Liqun Luo. Comprehensive maps of drosophila higher olfactory centers: spatially segregated fruit and pheromone representation. *Cell*, 128(6):1187–1203, 2007.

[51] Lida Kanari, Paweł Dłotko, Martina Scolamiero, Ran Levi, Julian Shillcock, Kathryn Hess, and Henry Markram. A topological representation of branching neuronal morphologies. *Neuroinformatics*, 16(1):3–13, 2018.

[52] Lida Kanari, Adélie Garin, and Kathryn Hess. From trees to barcodes and back again: theoretical and statistical perspectives. *Algorithms*, 13(12):335, 2020.

[53] Lida Kanari, Srikanth Ramaswamy, Ying Shi, Sebastien Morand, Julie Meystre, Rodrigo Perin, Marwan Abdellah, Yun Wang, Kathryn Hess, and Henry Markram. Objective morphological classification of neocortical pyramidal cells. *Cerebral Cortex*, 29(4):1719–1735, 2019.

[54] Reem Khalil, Sadok Kallel, Ahmad Farhat, and Pawel Dlotko. Topological sholl descriptors for neuronal clustering and classification. *PLOS Computational Biology*, 18(6):e1010229, 2022.

[55] Peter Koepernik and Florian Pfaff. Consistency of gaussian process regression in metric spaces. *J. Mach. Learn. Res.*, 22:244–1, 2021.

[56] Sophie Laturnus, Dmitry Kobak, and Philipp Berens. A systematic evaluation of interneuron morphology representations for cell type discrimination. *Neuroinformatics*, 18(4):591–609, 2020.

[57] Jacob Leygonie and Gregory Henselman-Petrusek. Algorithmic reconstruction of the fiber of persistent homology on cell complexes. *arXiv preprint arXiv:2110.14676*, 2021.

[58] Jacob Leygonie and Ulrike Tillmann. The fiber of persistent homology for simplicial complexes. *Journal of Pure and Applied Algebra*, page 107099, 2022.

[59] Yanjie Li, Dingkang Wang, Giorgio A Ascoli, Partha Mitra, and Yusu Wang. Metrics for comparing neuronal tree shapes based on persistent homology. *PloS one*, 12(8):e0182184, 2017.

[60] Michael London and Michael Häusser. Dendritic computation. *Annu. Rev. Neurosci.*, 28:503–532, 2005.

[61] Sergiy Maksymenko. Homotopy types of stabilizers and orbits of Morse functions on surfaces. *Annals of Global Analysis and Geometry*, 29(3):241–285, 2006.

[62] Lewis Marsh and David Beers. Stability and inference of the euler characteristic transform. *arXiv preprint arXiv:2303.13200*, 2023.

[63] Lewis Marsh, Felix Y Zhou, Xiao Qin, Xin Lu, Helen M Byrne, and Heather A Harrington. Detecting temporal shape changes with the euler characteristic transform. *arXiv preprint arXiv:2212.10883*, 2022.

[64] John N Mather. Stability of C^∞ mappings: II. infinitesimal stability implies stability. *Annals of Mathematics*, pages 254–291, 1969.

[65] Kun Meng, Lorin Crawford, and Ani Eloyan. Randomness and statistical inference of shapes via the smooth euler characteristic transform. *arXiv preprint arXiv:2204.12699*, 2022.

[66] John Willard Milnor, Michael Spivak, Robert Wells, and Robert Wells. *Morse theory*. Number 51. Princeton university press, 1963.

[67] Konstantin Mischaikow and Charles Weibel. Persistent homology with non-contractible preimages. *arXiv preprint arXiv:2105.08130*, 2021.

[68] Dmitriy Morozov, Kenes Beketayev, and Gunther Weber. Interleaving distance between merge trees. *Discrete and Computational Geometry*, 49(22-45):52, 2013.

[69] Kalyan Varma Nadimpalli, Amit Chattopadhyay, and Bastian Rieck. Euler characteristic transform based topological loss for reconstructing 3d images from single 2d slices. 2023.

[70] Monica Nicolau, Arnold J Levine, and Gunnar Carlsson. Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival. *Proceedings of the National Academy of Sciences*, 108(17):7265–7270, 2011.

[71] Ippei Obayashi and Michio Yoshiwaki. Field choice problem in persistent homology. *arXiv preprint arXiv:1911.11350*, 2019.

[72] Richard S Palais. Local triviality of the restriction map for embeddings. *Commentarii Mathematici Helvetici*, 34(1):305–312, 1960.

[73] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.

[74] WILFRID Rall, RE Burke, TG Smith, Po G Nelson, and K Frank. Dendritic location of synapses and possible mechanisms for the monosynaptic epsp in motoneurons. *Journal of Neurophysiology*, 30(5):1169–1193, 1967.

[75] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer, 2003.

[76] W Christopher Risher, Tuna Ustunkaya, Jonnathan Singh Alvarado, and Cagla Eroglu. Rapid golgi analysis method for efficient and unbiased classification of dendritic spines. *PloS one*, 9(9):e107591, 2014.

[77] Vanessa Robins and Katharine Turner. Principal component analysis of persistent homology rank functions with case studies of spatial point patterns, sphere packing and colloids. *Physica D: Nonlinear Phenomena*, 334:99–117, 2016.

[78] Arthur Sard. The measure of the critical values of differentiable maps. *Bulletin of the American Mathematical Society*, 48(12):883–890, 1942.

[79] Colinda LGJ Scheele, Edouard Hannezo, Mauro J Muraro, Anoek Zomer, Nathalia SM Langedijk, Alexander Van Oudenaarden, Benjamin D Simons, and Jacco Van Rheenen. Identity and dynamics of mammary stem cells during branching morphogenesis. *Nature*, 542(7641):313–317, 2017.

[80] Ruggero Scorcioni, Sridevi Polavaram, and Giorgio A Ascoli. L-measure: a web-accessible tool for the analysis, comparison and search of digital reconstructions of neuronal morphologies. *Nature protocols*, 3(5):866–876, 2008.

[81] Donald A Sholl. Dendritic organization in the neurons of the visual and motor cortices of the cat. *Journal of anatomy*, 87(Pt 4):387, 1953.

[82] Gurjeet Singh, Facundo Mémoli, Gunnar E Carlsson, et al. Topological methods for the analysis of high dimensional data sets and 3d object recognition. *PBG@ Eurographics*, 2, 2007.

[83] Primoz Skraba and Katharine Turner. Wasserstein stability for persistence diagrams. *arXiv preprint arXiv:2006.16824*, 2020.

[84] Philip Smith and Vitaliy Kurlin. Families of point sets with identical 1d persistence. *arXiv preprint arXiv:2202.00577*, 2022.

[85] Joseph Snider, Andrea Pillai, and Charles F Stevens. A universal property of axonal and dendritic arbors. *Neuron*, 66(1):45–56, 2010.

[86] EW Stockley, HM Cole, AD Brown, and HV Wheal. A system for quantitative morphological measurement and electrotonic modelling of neurons: three-dimensional reconstruction. *Journal of neuroscience methods*, 47(1-2):39–51, 1993.

[87] Mikolaj J Sulkowski, Srividya Chandramouli Iyer, Mathieu S Kurosawa, Eswar Prasad R Iyer, and Daniel N Cox. Turtle functions downstream of cut in differentially regulating class specific dendrite morphogenesis in drosophila. *PloS one*, 6(7):e22611, 2011.

[88] Wai Shing Tang, Gabriel Monteiro da Silva, Henry Kirveslahti, Erin Skeens, Bibo Feng, Timothy Sudijono, Kevin K Yang, Sayan Mukherjee, Brenda Rubenstein, and Lorin Crawford. A topological data analytic approach for discovering biophysical signatures in protein dynamics. *PLoS computational biology*, 18(5):e1010045, 2022.

[89] Corinne M Teeter and Charles F Stevens. A general principle of neural arbor branch density. *Current Biology*, 21(24):2105–2108, 2011.

[90] Elena Farahbakhsh Touli and Yusu Wang. Fpt-algorithms for computing gromov-hausdorff and interleaving distances between trees. *arXiv preprint arXiv:1811.02425*, 2018.

[91] Katharine Turner, Sayan Mukherjee, and Doug M Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 2014.

[92] Katharine Turner, Vanessa Robins, and James Morgan. The extended persistent homology transform of manifolds with boundary. *arXiv preprint arXiv:2208.14583*, 2022.

[93] Lou Van den Dries et al. *Tame topology and o-minimal structures*, volume 248. Cambridge University Press, 1998.

[94] Bruce Wang, Timothy Sudijono, Henry Kirveslahti, Tingran Gao, Douglas M Boyer, Sayan Mukherjee, and Lorin Crawford. A statistical pipeline for identifying physical features that differentiate classes of 3d shapes. *The Annals of Applied Statistics*, 15(2):638–661, 2021.

[95] Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. *Advances in neural information processing systems*, 13, 2000.

[96] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

Acknowledgements

I would firstly like to thank my supervisors Alain Goriely and Heather Harrington for offering their guidance to me for the last four years. Thank you both for your patience and encouragement during my time in Oxford. Your open doors proved an invaluable resource, and I feel that both of you knew how to put my academic worries to rest. My time as a graduate student has been about as stress-free as one could wish as a result.

Second I would like to thank my fellow mathematicians in the TDA group. Jacob, collaborating with you has been a pleasure. What started as a quick chat about a poorly thought out idea I had turned into two novel approaches to a challenging problem, and years of work together. Lewis, thank you for taking interest in my research on stability, and for your insight in seeing that our research was in fact related. Together we were able to write a paper that was greater than the sum of its parts. And there are many others with whom I had the joy of discussing research with, often amidst other more nonsensical conversations. Prama, Pavan, Kate, Otto, Thomas, Hadrian, Georgia, Matthew, Ambrose, Christoffer, Thomas, Adam, Gill, Abhinav, Tommi, Agnese, Travis, Vidit, and Ulrike, you have all made the Mathematical Institute a wonderful place to work.

Another great pleasure of my time in Oxford has been meeting a breadth interesting people outside of mathematics who I can now call my friends. Ben, Connor, Lisanne, Sarah, Julia, and Priss, I am grateful to have met you in my very first year here and to have stuck together with you until this very day. Jaś, Nathan, Dan, Carina, Sarah, Adam, Magnus, Sharm, Honoka, Isabela, Cat, Hannah, and Theo, my only wish is that I had met you all sooner.

I also have to thank my many friends from Boston whose online presence kept me more or less sane during the pandemic. Sebastian, Ryan, David ($\times 2$), Vincent, Alex, Gabe, Brian, Sean, Malcolm, Francis, Tyrone, and Freddy, I am grateful that I had an opportunity that I had the chance to see each of you every time I went home. Every one of you made my cry with laughter at some point during our calls across the

pond. Finally, I must thank my family for their support during these years, especially my brother Louis whose friendship and driver's license have been a great boon to me.