Local Inference of Morse Indices Using Finite Point Samples

Ka Man Yim
St Anne's College, Oxford

A thesis presented for the degree of

Doctor of Philosophy in Mathematics

September 2021

To Mum and Dad, and Mr Chow Kan Hung.

Acknowledgements

I would like to thank my academic supervisors Prof Vidit Nanda, Prof Peter Grindrod CBE, and Dr Andrew Mellor, for generously imparting their time, knowledge and wisdom, and their thoughtful guidance throughout this DPhil. I would also like to thank my industry supervisors Dr Andrew Plume, Dr Kate Gasson, and Ms Rachel Herbert for their faith and support in this project, but above all being open-minded about applying novel mathematics to industrial problems. I am grateful for the experience of collaborating with my friend Jacob Leygonie, an ever patient collaborator from whom I have learnt a lot. I am also indebted to Prof Colin Please and Prof Chris Breward for setting up InFoMM, which has made this DPhil possible, and their continued personal support throughout the programme. Finally, I could not have done this DPhil without the love and support of my friends Diego and Lingyi, my brother Thomas, and my mum.

Contents

Abstract

We consider a smooth function $f : M \to \mathbb{R}$ on a Riemannian submanifold M embedded in an ambient Euclidean space $\mathbb{R}^d$. We provide theoretical guarantees for inferring the Morse index of a compact, dynamically isolated critical set of f using a pair of finite point clouds $\mathbb{X}$ and $\mathbb{X}_-$ in $\mathbb{R}^d$, and detail sufficient conditions which $\mathbb{X}$ and $\mathbb{X}_-$ have to satisfy. We also outline a realisable computational procedure for obtaining $\mathbb{X}$ and $\mathbb{X}_-$ and inferring the Morse index.

In a separate study with Jacob Leygonie, we discuss how persistent homology and wavelet signatures can be applied to graph classification problems in machine learning. We construct a supervised learning model which backpropagates through persistent homology to optimise over the space of wavelet signatures. We then apply this model to several benchmark graph classification problems and obtain accuracies in line with state of the art models.

The final portion of this thesis investigates how maximum diversity, a notion of effective cardinality in geometric measure theory, can be applied to industrial applications of interest to Elsevier, who sponsor this DPhil. In particular, we apply maximum diversity to assess the multi-disciplinarity of a collection of EPSRC grants, which takes relations between topics and grants into account. We also derive a generalised notion of Jaccard similarity such that we can measure the topical similarity between two grants, leveraging the pairwise similarities between their associated topics. We give recommendations on further industrial applications of maximum diversity in Elsevier's business.

Chapter 1

Introduction

This thesis documents three pieces of research undertaken during this degree. Across these projects, we apply ideas in *topological data analysis* to develop computational methods for real life data science problems.

The first project encompasses Chapters 2 to 4 and lends the thesis its title. In this project, we are interested in smooth functions $f : M \to \mathbb{R}$ on Riemannian manifolds. For the purposes of our discussion, we always assume a Riemannian manifold M is properly and smoothly embedded in some Euclidean space $\mathbb{R}^d$ of sufficiently high dimension. The main object of concern are *critical sets* of f, subsets of M where $\nabla f = 0$. We denote the critical sets of f as $\mathrm{Crit}(f)$. If $\nabla f(p) = 0$ for some $p \in f^{-1}(0)$, we say 0 is a *critical value* of f; else we say 0 is a *regular value* of f.

We focus on computing topological invariants associated with critical sets called *Morse indices* [Milnor et al., 1973, Conley, 1978, Mischaikow and Mrozek, 2002]. Morse indices are topological invariants in the following sense. First, in our setting, the Morse index of $S \subset \mathrm{Crit}(f)$ is the *relative homology* of special pairs of subsets of M called *index pairs*, which have to satisfy certain properties in relation to the gradient flow along $-\nabla f$. Second, the Morse index of S is an inherent property of S and independent of the choice of index pairs used to compute it.

The main obstacle in computing the Morse index of S lies in finding an appropriate index pair. We also require the index pair to be topologically 'tame' (e.g. homeomorphic to some finite simplicial complex), so that their relative homology can be inferred from finite data.

For a dynamically isolated critical set of f (see Definition 4.1.1), Gromoll and Meyer [1969] and Chang and Ghoussoub [1996] derived sufficient conditions for a pair of subsets to be an index pair. We refer to these instances of index pairs as *Gromoll-Meyer pairs* (Definition 4.1.3). In Section 4.2 of Chapter 4, we give an explicit construction of Gromoll-Meyer pairs for a compact dynamically isolated critical set S, relying only on bounds of $\|\nabla f\|$ in the vicinity of S in M.

We refer to the Gromoll-Meyer pairs that we construct as $(\mathscr{M}, \mathscr{M}_-)$. They are compact, and are constructed by intersecting sublevel and level sets of smooth functions of M, where the sublevel and level sets are taken at regular values (eq. (4.15)). We then argue that the relative homology of $(\mathscr{M}, \mathscr{M}_-)$ can be inferred from a pair of point clouds $(\mathbb{X}, \mathbb{X}_-)$ in $\mathbb{R}^d$. This relies on two key ideas. First, we prove that if a compact subset $A \subset \mathbb{R}^d$ can be written as a transverse intersections of sublevel sets at regular values on $M \subset \mathbb{R}^d$, then A is a subset with *positive reach* (Proposition 3.6.6). Thus, $(\mathscr{M}, \mathscr{M}_-)$ are subsets of $\mathbb{R}^d$ with positive reach. The *reach* (Definition 2.2.2) is a length scale τ_A that characterises how A is embedded in $\mathbb{R}^d$, and A is said to have positive reach if $\tau_A > 0$. The second idea is that the relative homology of pairs of compact subsets with positive reach can be inferred from finite point samples, if the Hausdorff distance between the pair and the point samples is sufficiently small relative to the reaches of the pair. (Proposition 2.3.20). In Proposition 3.6.9, we derive an explicit expression for a positive lower bound of the reach of A, where A is a compact transverse intersection of sublevel sets at regular values on a properly embedded submanifold. Thus, we know explicitly how close $(\mathbb{X}, \mathbb{X}_-)$ need to be to $(\mathscr{M}, \mathscr{M}_-)$ in their Hausdorff distances, so that $(\mathbb{X}, \mathbb{X}_-)$ provide sufficient data for computing the Morse index of S. We then describe how we can obtain $(\mathbb{X}, \mathbb{X}_-)$ in a realistic computational setting in Section 4.3.

The mathematical content pertaining to this project is organised in the following way. In Chapter 2, we recall concepts and properties related to subsets with positive reach, mainly introduced by Federer [1959]. We then generalise the homology inference techniques and theoretical guarantees for compact submanifolds in Niyogi et al. [2008] to subsets with positive reach. While we principally rely on arguments in Niyogi et al. [2008], there is a rich body of literature on estimating the homology of compact subsets from finite point samples. For example, Bobrowski et al. [2017] focused on the case where the point sample was obtained from a probability distribution supported on a compact manifold, and showed that as the num-

ber of points drawn from the distribution tends to infinity, the probability that the homology of the manifold can be recovered from the point sample tends to one. Other approaches to homology estimation include studies that use concepts related to the reach, such as the weak feature size [Chazal and Lieutier, 2005, Chazal et al., 2009c]. We also note that the reach is also a parameter of interest in many computational and statistical problems involving manifolds, for example in deciding whether a point sample is drawn from some underlying manifold [Fefferman et al., 2016], and deriving theoretical guarantees for approximating functions on manifolds using neural networks [Chen et al., 2019b].

The main aim of Chapter 3 is the derivation of Proposition 3.6.6 and Proposition 3.6.9 which we have mentioned above. As $(\mathscr{M}, \mathscr{M}_-)$ are embedded into $\mathbb{R}^d$ along the embedding of M, the geometry and indeed the reach of $(\mathscr{M}, \mathscr{M}_-)$ as a subset of $\mathbb{R}^d$ are dependent not only on the intrinsic properties of $(\mathscr{M}, \mathscr{M}_-)$ in M, but also on the extrinsic geometry of M in $\mathbb{R}^d$. Therefore, to build up to Proposition 3.6.6 and Proposition 3.6.9, we review some consequences of curvature on the local geometry of manifolds. Combining geometric knowledge about subsets with positive reach from Chapter 2 with those specific to manifolds that we acquire in Chapter 3, we are able to derive Proposition 3.6.6 and Proposition 3.6.9.

Finally, in Chapter 4, we review some theoretical facts about Morse indices from Conley index theory, and present our construction of Gromoll-Meyer pairs $(\mathscr{M}, \mathscr{M}_-)$. We then discuss a sampling strategy for $(\mathscr{M}, \mathscr{M}_-)$ such that we obtain finite point samples $(\mathbb{X}, \mathbb{X}_-)$ that help us infer the Morse index of a compact, dynamically isolated critical set S.

The other two projects are described in self contained chapters. We give a brief account here and defer their formal introductions to Section 5.1 and Section 6.1. In Chapter 5, we describe joint work undertaken with Jacob Leygonie, where we apply *persistent homology* [Zomorodian and Carlsson, 2005, Edelsbrunner and Harer, 2008] to classify graphs in data science applications. If we have a function $f : V \to \mathbb{R}$ on the vertex set of a simplicial graph $G = (V, E)$, we call (G, f) a *filtered graph* and f a *filter* of G. We can use persistent homology to map filtered graphs to the space of *persistence diagrams*, which is equipped with a notion of distance. Thus, we can compare filtered graphs using their persistence diagrams. Given graphs without filters, we can filter them using their *wavelet signatures* [Hammond et al., 2011, Li and Hamza, 2013], which are vertex functions whose function values are informed by the structure of the

graph itself. Wavelet transforms and signatures are derived from the field of harmonic analysis on graphs and manifolds; see for instance Antoine et al. [2010], Hammond et al. [2011] for a succinct summary of the field. We investigate whether the persistence diagrams of graphs filtered using wavelet signtures produce useful representations of graphs, in applications to graph classification problems in machine learning. We also make use of the differentiability of persistent homology to optimise the classification accuracy of graphs over the parameter space of wavelet signatures. We outline a graph classification pipeline, and discuss some experimental results on benchmark graph classification problems in Chapter 5.

Finally, in Chapter 6, we give an account of research undertaken in collaboration with Elsevier, a leading academic publisher who is a sponsor of my DPhil via Oxford University's Industrially Focussed Mathematical Modelling Centre for Doctoral Training. Elsevier is interested in measuring the multi-disciplinarity of research grants and assessing the topical similarity between different grants. Elsevier has catalogued how research grants awarded by major funding organisations are associated to research topics indexed by Elsevier [Klavans and Boyack, 2017]. While the number of topics associated to a grant can give a coarse measure of multi-disciplinarity, the actual topical diversity of the grant should be strictly less than the number of topics if the grant's topics are similar to each other. We account for the effect of similarities between topics using a concept in geometric measure theory called *maximum diversity*. The maximum diversity of the set of topics associated to a grant can be interpreted as the 'effective number of topics' of the set. We applied maximum diversity to a dataset of EPSRC grants to assess how often the number of topics associated to a grant gives a poor measure of its actual multi-disciplinarity. We also apply maximum diversity to derive a generalisation of the Jaccard similarity between sets, allowing us to take the similarity between topics into account in comparing the topical content of research grants. We assess the application of these concepts to the EPSRC dataset and give some recommendations for further downstream applications in Elsevier's business.

Chapter 2

Homology Inference

2.1 Introduction

In this chapter, we discuss how the homology of subsets of $\mathbb{R}^d$ with positive reach can be inferred from a finite collection of points $\mathbb{X} \subset \mathbb{R}^d$, which we refer to as a *point cloud*. Subsets with positive reach is a class of closed subsets (Lemma 2.2.4), which notably includes compact submanifolds of $\mathbb{R}^d$ (Proposition 3.1.2). Our primary contribution is the generalisation of the deformation retracts arguments of Niyogi et al. [2008] for compact Riemannian submanifolds to closed subsets with positive reach.

Unless otherwise stated, distances and lengths are Euclidean distances. We formally define the reach τ_A of $A \subset \mathbb{R}^d$ and describe its various geometric significance in Section 2.2. Intuitively, we can understand τ_A as a length scale, such that for $\epsilon < \tau_A$, the *ϵ-thickening* A^ϵ of A in the ambient Euclidean space is an open neighbourhood of A that deformation retracts to A (see Proposition 2.2.13 for formal statement). The ϵ-thickening of A here refers to the following union of open balls

$$A^\epsilon = \bigcup_{p \in A} B_\epsilon(p). \tag{2.1}$$

Here $B_\epsilon(x)$ refers to the open Euclidean ball of radius ϵ centered at x. We refer to the corresponding closed Euclidean ball with $B_\epsilon[x]$.

We set out conditions on the point cloud $\mathbb{X}$ such that $\mathbb{X}$ provides sufficient data to infer to homology of A. These conditions concern the relationship of τ_A with the following two other

length scales that relate $\mathbb{X}$ to A. The first property is density: we say $\mathbb{X}$ is δ-dense in A if the δ-thickening $\mathbb{X}^\delta$ of $\mathbb{X}$ contains A. The second property is closeness: we say $\mathbb{X}$ is r-close to A if conversely A is r-dense in $\mathbb{X}$. We recall that the *Hausdorff* distance $d_H(A, \mathbb{X})$ between A and $\mathbb{X}$ is the smallest ϵ such that $\mathbb{X}$ is both ϵ-dense in A and ϵ-close to A.

Definition 2.1.1. Let A and B be two subsets or metric space $\mathbb{R}^d$. The Hausdorff distance between A and B is

$$d_H(A, B) = \inf\{\epsilon \geq 0 \ : \ B \subseteq A^\epsilon \quad \text{and} \quad A \subseteq B^\epsilon\}.$$

Following the arguments by Niyogi et al. [2008] for the special case where A is a compact submanifold of $\mathbb{R}^d$, we show in Section 2.3 that if δ and r is sufficiently small, there is a range of thickening parameters ϵ such that $\mathbb{X}^\epsilon$ deformation retracts onto A, and hence has the same homotopy type as A. Furthermore, we show that a choice of (δ, r, ϵ) where the deformation retract condition are met can be determined if we are given τ_A or some positive lower bound of τ_A. We also derive sufficient conditions on $\mathbb{X}$ for inferring the homology of A from $\mathbb{X}$ using persistent homology, which are less stringent than those required for deformation retract.

We first review and prove some basic geometric properties of subsets A with positive local feature size or reach in Section 2.2, which help us establish the homology inference results in Section 2.3. In Section 2.2.3, we also recall from Federer [1959] a formula for computing the reach of subsets with with positive reach, involving generalisations of tangent and normal spaces of submanifolds to subsets with positive local feature size. This formula will be used in the subsequent chapter to derive a positive lower bound of the reach of compact intersections of sublevel and level sets of smooth functions on embedded submanifold.

Remark 2.1.2. While this thesis focuses on the implications of the reach on topological inference from point samples, we note that there are many similar approaches in literature which considers the implications of other geometric properties of compact subsets on topological inference. Different approaches are useful in different scenarios. While reach-based approaches recover the homotopy type using geometric data based on the ambient Euclidean metric, they are limited to subsets with positive reach, which constrain the neighbourhoods of points on A to resemble a convex cone (see Theorem 2.2.9). One alternative to the reach is the *weak feature size*, defined to be the infimum of positive critical values of the distance

function $d(\cdot, A) : \mathbb{R}^d \to [0, \infty)$. Chazal and Lieutier [2005] showed if $d_H(\mathbb{X}, A)$ is sufficiently small relative to the weak feature sizes of both A and $\mathbb{X}$, then the homotopy type of A^ϵ can be inferred for $\epsilon > 0$ sufficiently small. This approach can be applied to generalise homology inference guarantees to a class of reasonably tame subsets with zero reach. Compared to similar conditions placed on $d_H(\mathbb{X}, A)$ relative to the reach in Niyogi et al. [2008], the conditions arising from the weak feature size of the point cloud itself can be difficult to be simultaneously satisfied alongside conditions on A and $d_H(\mathbb{X}, A)$ in practical settings (we refer the reader to [Boissonnat et al., 2018, chapter 9, section 3] for a detailed discussion). Using a refinement of the weak feature size, called the μ-reach, Chazal et al. [2009c], the aforementioned issues afflicting the weak feature size were circumvented as the μ-reach implicitly encodes more information about A. There are other approaches that rely on quantities similar to the reach or weak feature size, that are derived from the intrinsic geodesic distances of the subset rather than distances in the ambient Euclidean space. Gao et al. [2008] showed that the homotopy type of a Lipschitz planar domain $A \subset \mathbb{R}^2$ can be inferred from the geodesic Delaunay triangulation on a sufficiently dense point sample of A; the sufficient density is relative to the *homotopy feature size* of A, which is the length of the shortest loop in A as measured with respect to the geodesic distance on A. In a slightly different context, Chazal et al. [2009a] uses the *strong convexity radius* of an underlying manifold M to derive theoretical guarantees for approximating the persistent diagrams of sublevel set filtrations of M from point samples, provided the filter function $f : M \to \mathbb{R}$ is sufficiently tame; the strong convexity radius of M is the radius of the smallest geodesic ball on M such that any two points in the ball are connected by a *unique* geodesic.

2.2 Subsets of Euclidean Space with Positive Local Feature Size

In his seminal work, Federer [1959] has shown that the condition of positive local feature size imposes geometric constraints the local geometry and topology of A. This section catalogues a few relevant geometric properties of subsets with positive local feature size, which include the class of subsets with positive reach.

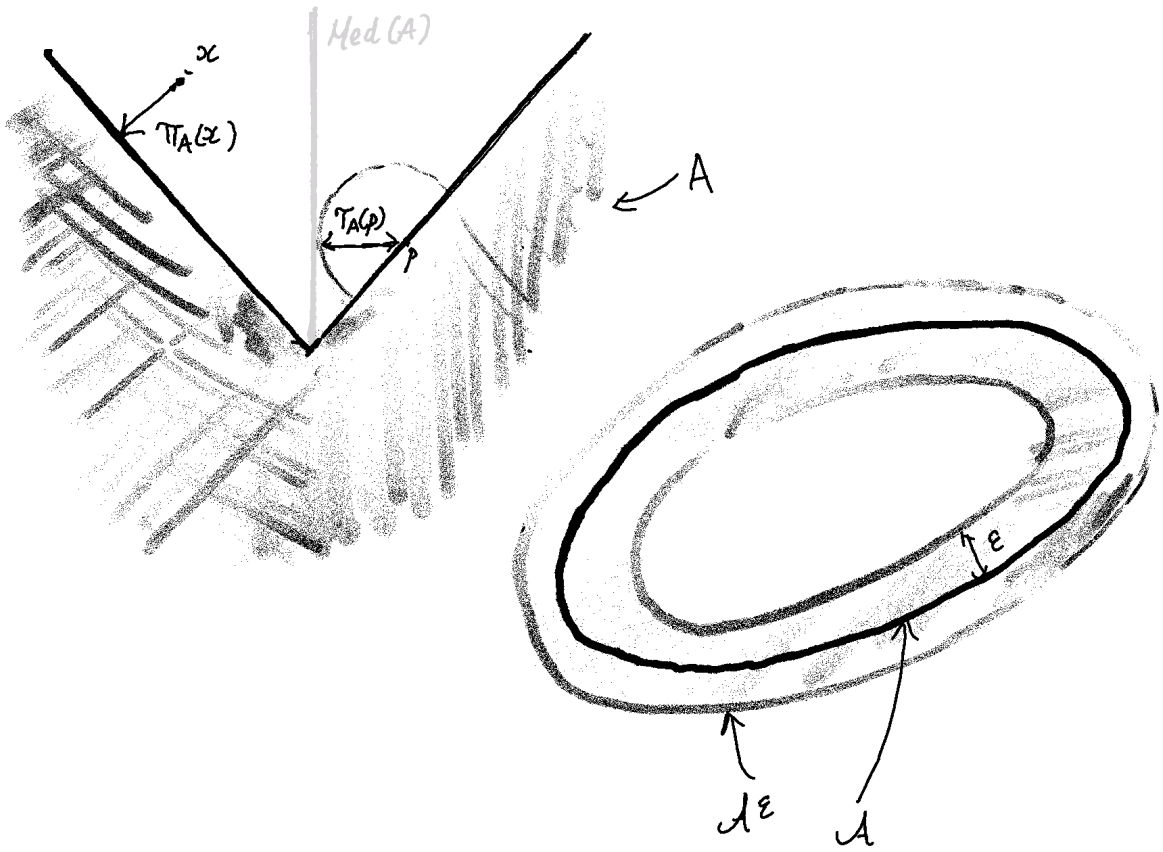

Figure 2.1: An illustration of some of the geometric concepts related to a subset $A \subset \mathbb{R}^d$ (where $d = 2$), such as the local feature size $\tau_A(p)$, the medial axis $\mathrm{Med}(A)$, the projection map π_A, and the ϵ-thickening of A^ϵ.

2.2.1 The Local Feature Size of Subsets of Euclidean Space

Federer [1959] describes the geometry of $A \subset \mathbb{R}^d$ in relation to other points in the ambient Euclidean space. Let $\mathrm{UP}(A)$ be the set of points $x \in \mathbb{R}^d$ for which there is a unique point of A nearest to x, and the projection map $\pi_A : \mathrm{UP}(A) \to A$ associates with $x \in \mathrm{UP}(A)$ the unique $p \in A$ such that $d_{\mathbb{R}^d}(x, A) = \|x - p\|$. We first recall the following elementary result about π_A.

Theorem 2.2.1 ([Federer, 1959, Theorem 4.8(2& 4)]). *If $A \subset \mathbb{R}^d$ is a closed subset, the projection map $\pi_A : \mathrm{UP}(A) \to A$ onto A is continuous, and the distance to A, given by $d_A = d_{\mathbb{R}^d}(\cdot, A) : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ is Lipschitz continuous.*

The local feature size of A at $p \in A$ is defined as the radius of the largest Euclidean open ball about p that one can fit in $\mathrm{UP}(A)$.

Definition 2.2.2 (Local Feature Size [Federer, 1959, Definition 4.1]). For $p \in A$, the *local feature size* is a function $\tau_A : A \to \mathbb{R}_{\geq 0}$ given by

$$\tau_A(p) = \sup\{r \geq 0 \ : \ B_r(p) \subset \mathrm{UP}(A)\}.$$

We say A has positive local feature size if $\tau_A(p) > 0$ for all $p \in A$. The *reach* of A, τ_A, is the infimum of the local feature size over all points in A:

$$\tau_A = \inf_{p \in A} \tau_A(p) \tag{2.2}$$

We say A has *positive reach* if $\tau_A > 0$. The *Medial Axis* of A, $\mathrm{Med}(A)$, is the set of points $x \in \mathrm{UP}(A)$ where the set of nearest neighbours of x in A has cardinality greater than one. If A is closed, then

$$\tau_A(p) = d_{\mathbb{R}^d}(p, \mathrm{Med}(A)) = \sup\{r \geq 0 \ : \ B_r(p) \cap \mathrm{Med}(A) = \emptyset\}$$
$$\tau_A = \inf_{x \in \mathrm{Med}(A)} d(x, A)$$

The following result about the continuity of the local feature size as a function on A was quoted as a remark in Federer [1959] without proof, which we prove here explicitly.

Lemma 2.2.3. *[Federer, 1959, Remark 4.2] The local feature size of $A \subset \mathbb{R}^d$ is Lipschitz continuous with respect to the Euclidean norm of the ambient space $\mathbb{R}^d$:*

$$|\tau_A(a) - \tau_A(b)| \leq \|a - b\|.$$

Proof. Consider points $a, b \in A$ and $c \in \mathbb{R}^d \setminus \mathrm{UP}(A)$. By the triangle inequality,

$$\|b - a\| + \|a - c\| \geq \|b - c\|.$$

Taking the infimum of the right hand side over $c \in \mathbb{R}^d \setminus \mathrm{UP}(A)$, we have

$$\|b - a\| + \|a - c\| \geq \inf_{c \in \mathbb{R}^d \setminus \mathrm{UP}(A)} \|b - c\| = \tau_A(b).$$

Similarly, we can take the infimum of the left hand side over $c \in \mathrm{Med}(M)$. Since the infimum is the greatest lower bound, we have

$$\|b - a\| + \tau_A(a) = \|b - a\| + \inf_{c \in \mathbb{R}^d \setminus \mathrm{UP}(A)} \|a - c\| \geq \tau_A(b).$$

Without loss of generality, we can choose a and b such that $\tau_A(a) \leq \tau_A(b)$, and deduce

$$\|b - a\| \geq \tau_A(b) - \tau_A(a) = |\tau_A(b) - \tau_A(a)|.$$

□

As a consequence of the continuity of the local feature size, we can show that only closed subsets of $\mathbb{R}^d$ have positive reach.

Lemma 2.2.4. *Only closed subsets of $\mathbb{R}^d$ have positive reach; For any $p \in A$ where A is not closed,*

$$\tau_A(p) \leq d_{\mathbb{R}^d}(p, \mathrm{cl}(A) \setminus A).$$

Proof. Let A be a subset of $\mathbb{R}^d$ and suppose $\mathrm{cl}(A) \setminus A \neq \emptyset$ i.e. A is not closed.

Suppose for some $q \in \mathrm{cl}(A) \setminus A$, and $p \in A$ that is distinct from q, we have $\tau_A(p) > \|p - q\|$. Then $q \in B_{\tau_A(p)}(p) \subset \mathrm{UP}(A)$. However, we reach a contradiction as $q \notin \mathrm{UP}(A)$, since for

every point in A we could always find another point that is closer to q. We therefore deduce that for any $p \in A$, $\tau_A(p) \leq \|p - q\|$. Thus, $\tau_A(p) \leq d_{\mathbb{R}^d}(p, \mathrm{cl}(A) \setminus A)$.

In particular, the continuity of local feature size (Lemma 2.2.3) implies for $q \in \mathrm{cl}(A) \setminus A$,

$$\tau_A(q) = \lim_{p \to q} \tau_A(p) = 0.$$

□

Another implication of the continuity of local feature size is that the positivity of $\tau_A(p)$ can be determined purely by the local geometry of p in A.

Lemma 2.2.5. *Consider $A \subset \mathbb{R}^d$ and $p \in A$.*

1. *If $\tau_A(p) > 0$, then $\tau_B(p) > 0$ for $B = A \cap B_\epsilon(p)$ and $\epsilon > 0$;*
2. *If there is some $\epsilon > 0$ such that $\tau_B(p) > 0$ for $B = A \cap B_\epsilon(p)$, then $\tau_A(p) > 0$.*

In other words $\tau_A(p) > 0$ if and only if $\tau_B(p) > 0$ for any $\epsilon > 0$.

Proof.

1. Suppose $\tau_A(p) > 0$ and $r \in (0, \tau_A(p))$. Then every $x \in B_r(p)$ has a unique projection onto some point in A that lies in $B_{2r}(p)$. Thus for $\epsilon > 2r > 0$, every $x \in B_r(p)$ has a unique projection onto $B = B_\epsilon(p) \cap A$. Therefore, $\tau_B(p) \geq r > 0$. Since r can be made arbitrarily small, the local feature size of p in $B_\epsilon(p) \cap A$ is positive for all $\epsilon > 0$.

2. Suppose there is some $\epsilon > 0$ such that $\tau_B(p) > 0$ for $B = A \cap B_\epsilon(p)$. Consider some $x \in B_r(p)$ where $0 < r < \min(\tau_B(p), \frac{\epsilon}{2})$. Since $\|x - p\| \leq r < \tau_B(p)$, from the definition of the local feature size Definition 2.2.2, we must have $x \in B_{\tau_B(p)}(p) \subset \mathrm{UP}(B)$. Therefore, there is a unique projection q of p onto B.

   Now suppose there is some $q' \in A$ where $q \neq q'$, such that $\|q' - x\| \leq \|q - x\|$; in other words, q is not the unique projection of x onto A. Then applying the triangle inequality in $\mathbb{R}^d$, we deduce that

$$\|q' - p\| \leq \|q' - x\| + \|x - p\| \leq \|q - x\| + \|x - p\| \leq 2\|x - p\| \leq 2r < \epsilon.$$

This implies that $q' \in B_\epsilon(p) \cap A \subset B$. Since q is the unique projection of x onto B and $q \neq q'$, the existence of q' contradicts the fact that $x \in \mathrm{UP}(B)$. Therefore, all $x \in B_r(p)$ must have a unique projection onto A that lies in B; in other words, $B_r(p) \subset \mathrm{UP}(A)$ and thus $\tau_A(p) \geq r > 0$.

□

2.2.2 Geometric and Topological Consequences of Positive Local Feature Size

We now catalogue a few useful geometric and topological properties of subsets with positive local feature size. Most of the geometric results derived in Federer [1959] are based on constructions of *tangent sets* and *normal cones* for points on the subset. These are subsets of $T_p\mathbb{R}^d$ that describe the local geometry of A at p, and its relation to surroundings points in the ambient Euclidean space respectively. If $\tau_A(p) > 0$, the tangent and normal cones are a pair of dual convex cones in $T_p\mathbb{R}^d$. In the case where A is a properly embedded submanifold, these tangent sets and normal cones are precisely thes tangent and normal spaces.

The language of tangent sets and normal cones give us the technical scaffolding to describe two results. The first one, Theorem 2.2.10, describes sufficient conditions under which the intersection of two subsets which positive local feature size also has positive local feature size. This result is the central plank in our subsequent derivation of a positive lower bound of the reach of compact intersections of sublevel sets of smooth functions. The second one is Federer's tubular neighbourhood theorem Theorem 2.2.12 for subsets with positive reach, which is the conceptual basis of the homology inference results using point clouds in the vein of Niyogi et al. [2008].

The tangent set also faciliates computing the reach τ_A of A if it is positive. Theorem 2.2.16 expresses the reach in terms of projections of A onto the tangent sets of individual points in A. Thus, we do not need to derive the explicit geometry of $\mathrm{UP}(A)$ to derive its reach, which was demanded in the definition of the reach in Definition 2.2.2.

Finally, we show in Lemma 2.2.15 that the intersections of A with open Euclidean balls is contractible if these balls are of sufficiently small and placed close to A. This result is applied in the derivation of lower bounds for the reach of sublevel sets of smooth functions on embedded

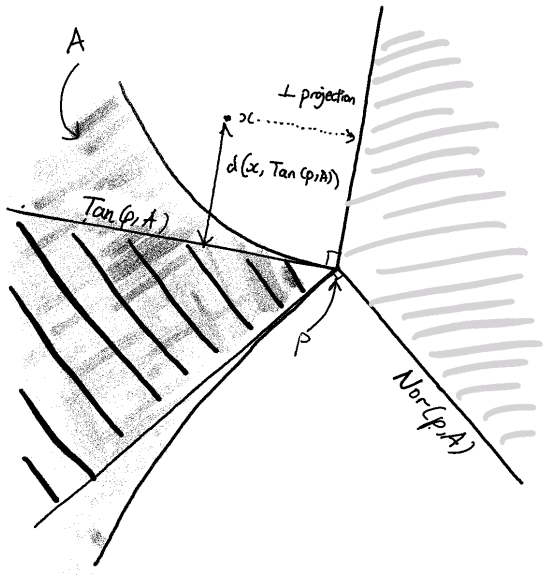

Figure 2.2: An illustration of the tangent set $\mathrm{Tan}(p, A)$ and normal cone $\mathrm{Nor}(p, A)$ of a $A \subset \mathbb{R}^2$.

submanifolds in the next chapter.

2.2.2.1 Tangent Sets and Normal Cones

In Federer [1959], the local geometry of A at a point p is modelled using tangent sets and normal cones, which are collections of vectors in $T_p\mathbb{R}^d$ that generalise notions of tangent and normal spaces for embedded submanifolds to arbitrary subsets of $\mathbb{R}^d$.

Definition 2.2.6 ([Federer, 1959, Definitions 4.3 and 4.4])**.** For any $A \subset \mathbb{R}^d$ and $p \in A$, the *tangent set* is the following set of points in $T_p\mathbb{R}^d$:

$$\mathrm{Tan}(p, A) = \left\{ u \in T_p\mathbb{R}^d \ : \ \forall \epsilon > 0, \exists v \in A - p \text{ s.t. } \|v\| \in (0, \epsilon) \ \& \ \left\| \frac{v}{\|v\|} - \frac{u}{\|u\|} \right\| < \epsilon \right\} \cup \{0\}; \tag{2.3}$$

and we denote the *normal cone* is the convex dual of $\mathrm{Tan}(p, A)$ in $T_p\mathbb{R}^d$,

$$\mathrm{Nor}(p, A) = \left\{ v \in T_p\mathbb{R}^d \mid \forall u \in \mathrm{Tan}(p, A), \ \langle v, u \rangle \leq 0 \right\}, \tag{2.4}$$

which by definition is a closed convex cone.

We illustrate an example of a tangent set and normal cone in Figure 2.2. The tangent set and

normal cone depends only on the local geometry of A at p, which we make explicit in the following lemma.

Lemma 2.2.7. *For any $p \in A$, and $\epsilon > 0$,*

$$\mathrm{Tan}(p, A) = \mathrm{Tan}(p, A \cap B_\epsilon(p)) \quad \textit{and} \quad \mathrm{Nor}(p, A) = \mathrm{Nor}(p, A \cap B_\epsilon(p)).$$

Furthermore, if there is some $\epsilon > 0$ such that $B_\epsilon(p) \cap A = B_\epsilon(p) \cap \mathrm{cl}(A)$,

$$\mathrm{Tan}(p, A) = \mathrm{Tan}(p, \mathrm{cl}(A)) \quad \textit{and} \quad \mathrm{Nor}(p, A) = \mathrm{Nor}(p, \mathrm{cl}(A)).$$

Proof. The first statement follows directly from the definition of the tangent and normal cones.

If $A \cap B_\epsilon(p) = \mathrm{cl}(A) \cap B_\epsilon(p)$, then

$$\mathrm{Tan}(p, A) = \mathrm{Tan}(p, A \cap B_\epsilon(p)) = \mathrm{Tan}(p, \mathrm{cl}(A) \cap B_\epsilon(p)) = \mathrm{Tan}(p, \mathrm{cl}(A)).$$

As the tangent sets are identical the normal cones defined as the convex dual of them must also be identical. □

Combining Lemma 2.2.5 and Lemma 2.2.7, we can deduce the following.

Corollary 2.2.8. *Consider $A \subset \mathbb{R}^d$ and $B \subset \mathbb{R}^d$. Suppose $p \in A \cap B$ and there is some open neighbourhood U of p in $\mathbb{R}^d$, such that $A \cap U = B \cap U$. Then $\tau_A(p) > 0$ if and only if $\tau_B(p) > 0$, and*

$$\mathrm{Tan}(p, A) = \mathrm{Tan}(p, B) \quad \textit{and} \quad \mathrm{Nor}(p, A) = \mathrm{Nor}(p, B).$$

Proof. Let $\epsilon > 0$ be sufficiently small, such that $B_\epsilon(p) \subset U$. Then Lemma 2.2.5 implies

$$\tau_A(p) > 0 \iff \tau_{B_\epsilon(x)0<\cap A}(p) = \tau_{B_\epsilon(x)\cap B}(p) > 0 \iff \tau_B(p) > 0.$$

Similarly, Lemma 2.2.7 implies

$$\mathrm{Tan}(p, A) = \mathrm{Tan}(p, A \cap B_\epsilon(p)) = \mathrm{Tan}(p, B \cap B_\epsilon(p)) = \mathrm{Tan}(p, B).$$

□

If $\tau_A(p) > 0$, then the tangent set is actually a closed convex cone, which is dual to the normal cone. Furthermore, the normal cone is the cone of rays emanating from p that pass through points that project onto p under π_A.

Theorem 2.2.9 ([Federer, 1959, Theorem 4.8(12)]). *Let A be a non-empty closed subset of $\mathbb{R}^d$. Suppose for some $a \in A$, $\tau_A(a) > 0$. Then $\mathrm{Nor}(a, A)$ is the closed convex cone*

$$\mathrm{Nor}(a, A) = \Big\{ \lambda v \ : \ v \in \mathbb{S}^{d-1}, \ \pi_A(a + \lambda v) = a, \text{ for any } 0 < \lambda < \tau_A(a) \Big\}. \tag{2.5}$$

Furthermore, $\mathrm{Tan}(a, A)$ is the convex cone dual to $\mathrm{Nor}(a, A)$, i.e.

$$\mathrm{Tan}(a, A) = \Big\{ u \in \mathbb{R}^d \ : \ \forall v \in \mathrm{Nor}(a, A), \ \langle v, u \rangle \leq 0 \Big\}. \tag{2.6}$$

The following theorem states a set of sufficient conditions under which the intersection of two subsets with positive local feature size also has positive local feature size.

Theorem 2.2.10 ([Federer, 1959, Theorem 4.10]). *For closed subsets A and B of $\mathbb{R}^d$, consider $p \in A \cap B$. If both A and B have positive local feature size at p, then*

$$\min\{\measuredangle(v, w) \ : \ v \in \mathrm{Nor}(p, A), \quad w \in \mathrm{Nor}(p, B)\} < \pi,$$

implies

$$\mathrm{Tan}(p, A \cap B) = \mathrm{Tan}(p, A) \cap \mathrm{Tan}(p, B)$$
$$\mathrm{Nor}(p, A \cap B) = \mathrm{Nor}(p, A) + \mathrm{Nor}(p, B) = \{v + w \ : \ v \in \mathrm{Nor}(p, A) \quad \text{and} \quad w \in \mathrm{Nor}(p, B)\}.$$

If there is some sufficiently small ϵ such that on $C = B_\epsilon[p] \cap A \cap B$, and we have

$$\min_{q \in C} \tau_A(q) > 0, \quad \min_{q \in C} \tau_B(q) > 0 \tag{2.7}$$

$$\max\{\measuredangle(v, w) \ : \ v \in \mathrm{Nor}(q, A), \quad w \in \mathrm{Nor}(q, B), \forall q \in C\} < \pi. \tag{2.8}$$

Then $\tau_{A \cap B}(p) > 0$.

In the case where A and B are submanifolds, the condition eq. (2.8) that the normal cones of either subsets cannot be opposed at q reduces to the statement that the $N_qA \perp N_qB$. If A and B are submanifolds with boundary and $q \in A \cap B$, we additionally require the outward pointing normals of A and B at q point in opposing directions.

2.2.2.2 Tubular Neighbourhoods of Subsets with Positive Reach

For $A \subset \mathbb{R}^d$ with positive reach $\tau_A > 0$, we refer to the ϵ-thickening of A in $\mathbb{R}^d$ as the *ϵ-tubular neighbourhood* whenever $\epsilon \in (0, \tau_A]$,

$$\mathrm{Tub}_\epsilon(A) = \bigcup_{p \in A} B_\epsilon(p).$$

As $\epsilon < \tau_A$ ensures that each open ball $B_\epsilon(p)$ is contained in $\mathrm{UP}(A)$, the tubular neighbourhood is an open neighbourhood of A in $\mathbb{R}^d$ contained in $\mathrm{UP}(A)$. We show explicitly that the fibres of the projection map π_A onto A partition $\mathrm{Tub}_\epsilon(A)$. As such, this enables us to construct a deformation retract of $\mathrm{Tub}_\epsilon(A)$ onto A.

Lemma 2.2.11. *Let A be a non-empty closed subset of $\mathbb{R}^d$. Suppose the local feature sizes $\tau_A(a)$ and $\tau_A(b)$ at two distinct points a and $b \in A$, are positive. For $x \in A$, let*

$$\mathrm{P^{-1}}_A(p) = (p + \mathrm{Nor}(p, A)) \cap B_{\tau_A(p)}(p).$$

Then we have

$$\mathrm{P^{-1}}_A(a) \cap \mathrm{P^{-1}}_A(b) = \emptyset.$$

Proof. Take $x \in \mathrm{P^{-1}}_A(a)$. From Theorem 2.2.9,

$$\mathrm{P^{-1}}_A(p) = (p + \mathrm{Nor}(p, A)) \cap B_{\tau_A(p)}(p) = \pi^{-1}{}_A(p) \cap B_{\tau_A(p)}(p).$$

Thus x has a unique nearest neighbour $\pi(x) = a$ in A. Thus, $x \notin \mathrm{P^{-1}}_A(b)$ where $b \neq a$, as this implies x has more than one neighbour in A. □

Theorem 2.2.12 (Federer's Tubular Neighbourhood Theorem)**.** *Suppose $A \subset \mathbb{R}^d$ has positive*

reach $\tau_A > 0$. *Then* $\mathrm{Tub}_\epsilon(A)$ *can be expressed as a disjoint union*

$$\mathrm{Tub}_\epsilon(A) = \bigsqcup_{p\in A} ((p + \mathrm{Nor}(p, A))) \cap B_\epsilon(p)) = \bigsqcup_{p\in A} \left(\pi^{-1}{}_A(p) \cap B_\epsilon(p)\right).$$

Proof. The second equality in the statement is a straightforward consequence of Theorem 2.2.9. Thus we focus on proving the first equality. As $\epsilon \in (0, \tau_A]$, it follows from the definition of the reach (Definition 2.2.2) that $A^\epsilon \cap \mathrm{Med}(A) = \emptyset$ and every point in $\mathrm{Tub}_\epsilon(A)$ has a unique projection onto A. Consider $x \in \mathrm{Tub}_\epsilon(A)$. As $B_\epsilon(x) \cap A \neq \emptyset$, the distance $\|x - \pi(x)\|$ between x and $\pi(x)$, its nearest neighbour in A, must be strictly less than ϵ. Therefore, $x \in \pi^{-1}(\pi(x)) \cap B_\epsilon(\pi(x))$, and thus, $x \in \mathrm{Nor}(\pi(x), A) \cap B_\epsilon(\pi(x))$. Since every point in $\mathrm{Tub}_\epsilon(A)$ has a unique projection onto A,

$$\mathrm{Tub}_\epsilon(A) = \bigcup_{p\in A} ((p + \mathrm{Nor}(p, A))) \cap B_\epsilon(p)).$$

As $((p + \mathrm{Nor}(p, A))) \cap B_\epsilon(p)) \subset \mathrm{P}^{-1}{}_A(p)$ and for any distinct a and b in A, the subsets $\mathrm{P}^{-1}{}_A(a)$ and $\mathrm{P}^{-1}{}_A(b)$ are disjoint (Lemma 2.2.11), the sets $(p + \mathrm{Nor}(p, A))) \cap B_\epsilon(p)$ must therefore be a partition of $\mathrm{Tub}_\epsilon(A)$. □

As a consequence of this result, we have the following well-known result about the robustness of the homotopy type of subsets with positive reach against thickenings.

Proposition 2.2.13. *Suppose* $A \subset \mathbb{R}^d$ *has positive reach* $\tau_A > 0$. *Then for all* $\epsilon \in [0, \tau_A)$, *we have a strong deformation retract* $F : [0, 1] \times \mathrm{Tub}_\epsilon(A) \to \mathrm{Tub}_\epsilon(A)$ *onto* A, *given by a straight line homotopy*

$$F(t, x) = (1 - t)x + t\pi(x) \tag{2.9}$$

where $\pi : \mathrm{Tub}_\epsilon(A) \to A$ *is the projection of* $x \in \mathrm{Tub}_\epsilon(A)$ *onto its nearest neighbour in* A *by the Euclidean metric.*

Proof. We first check that the image of F lies in $\mathrm{Tub}_\epsilon(A)$. Due to Theorem 2.2.12, any $x \in \mathrm{Tub}_\epsilon(A)$ belongs to a unique $\mathrm{Nor}(\pi(x), A) \cap B_\epsilon(\pi(x))$. It follows from the positive homogeneity of $\mathrm{Nor}(\pi(x), A)$ (Theorem 2.2.9) that the line segment $\overline{x\pi(x)}$ from x to $\pi(x)$ in Euclidean space is contained in $\mathrm{Tub}_\epsilon(A)$. The points along this line segment is precisely parametrised by $F(t, x)$ where $t \in [0, 1]$. Therefore, the image of F indeed lies in $\mathrm{Tub}_\epsilon(A)$. As π is continuous

(Theorem 2.2.1), F is thus a continuous map from $[0,1] \times \text{Tub}_\epsilon(A)$ to $\text{Tub}_\epsilon(A)$. One can then easily check that $F(t,x)$ is strong deformation retract onto A. □

2.2.2.3 Intersection with Euclidean Balls

We now show that a subset A with positive local feature size is locally contractible, in the sense that every $p \in A$ admits an open contractible neighbourhood in A.

Lemma 2.2.14. *Suppose for some $a \in A$, $\tau_A(a) > 0$. If A is closed in $\mathbb{R}^d$, then for any $r \in (0, \tau_A(a))$, $A \cap B_r(a)$ is contractible.*

Proof. This is a modification of the proof of [Federer, 1959, Remark 4.15]. Since $r < \tau_A(a)$, by definition $B_r(a) \subset \text{UP}(A)$. Therefore, any point in $x \in B_r(a)$ has a unique projection $\pi(x)$ onto A, and this projection is continuous. Furthermore, since $B_r(a)$ are convex subsets, the line segment $a(t,x) = (1-t)x + ta$ for $t \in [0,1]$ is contained in $B_r(a)$. Furthermore, we can show that for $t \in [0,1]$,

$$\pi(a(t,x)) \subset A \cap B_r(a).$$

Since $\pi(a(t,x))$ is the nearest neighbour of $a(t,x)$ in A, $\|\pi(a(t,x)) - a(t,x)\| \leq \|a(t,x) - x\|$. Along the line segment from x to a, $\|a(t,x) - a\| = t\|a - x\|$ and $\|a(t,x) - x\| = (1-t)\|a - x\|$. Therefore, we can bound $\|a - \pi(a(t,x))\|$ using the triangle inequality

$$\begin{aligned}
\|a - \pi(a(t,x))\| &\leq \|a(t,x) - a\| + \|a(t,x) - \pi(a(t,x))\| \\
&\leq \|a(t,x) - a\| + \|a(t,x) - x\| \\
&= (1-t)\|a - x\| + t\|a - x\| \\
&= \|a - x\| < r.
\end{aligned}$$

Therefore, we can define the continuous map $H : A \cap B_r(a) \times [0,1] \to A \cap B_r(a)$

$$H(x,t) = \pi_A(a(t,x))$$

which is a homotopy between $A \cap B_r(a)$ and the point $\pi(a)$. □

We can also generalise the aforementioned result for open balls based on $x \in A$ to balls based on points in the ambient Euclidean space that is within a bounded offset from A.

Lemma 2.2.15. *Suppose for some $p \in A$, $\tau_A(p) > 0$. Let $r \leq \max\left(\tau_M, \frac{\tau_A(p)}{3}\right)$ and consider $x \in B_r(p)$. If A is closed in $\mathbb{R}^d$, Then $B_r(x) \cap A$, if non-empty, is contractible.*

Proof. We first show that if $r \leq \max\left(\tau_M, \frac{\tau_A(p)}{3}\right)$, then for all $q \in B_r(x) \cap A$, we must have $r \leq \tau_A(q)$. This is self-evident in the case where $\tau_M \geq \tau_A(p)/3$, since the reach is the infimum of the local feature sizes of points on M. Consider the other case where $\tau_M < r < \tau_A(p)/3$, Since local feature size is 1-Lipschitz with respect to the Euclidean distance (Lemma 2.2.3), we must therefore have $\tau_A(q) \geq \tau_A(p) - \|q - p\|$. Since $\|q - p\| < 2r \leq \frac{2\tau_A(p)}{3}$, we therefore have $\tau_A(q) > \frac{\tau_A(p)}{3} \geq r$.

As such the line segment $\overline{qx}$ from q to x in Euclidean space, whose length is at most r, is contained in the open ball centered at q with radius $\tau_A(q)$. Therefore, $\overline{qx} \subset \mathrm{UP}(A)$ and any $y \in \overline{qx}$ has a unique projection onto A. Furthermore, we can show that $\pi_A(y) \in B_r(x)$. As $\|y - q\| \geq \|y - \pi_A(q)\|$, we can apply the triangle inequality and the fact that x, y and q are colinear, and deduce that

$$\begin{aligned}\|x - \pi_A(y)\| &\leq \|x - y\| + \|y - \pi_A(y)\| \\ &\leq \|x - y\| + \|y - q)\| = \|x - q\| < r.\end{aligned}$$

Therefore, we can define a continuous map $H : A \cap B_r(x) \times [0,1] \to A \cap B_r(x)$

$$H(q,t) = \pi_A(tx + (1-t)q)$$

which is a homotopy between $A \cap B_r(x)$ and $\pi_A(x)$. □

2.2.3 Computing the Reach

We discuss how the reach τ_A of $A \subset \mathbb{R}^d$ can be computed without explicit knowledge of $\mathrm{UP}(A)$, if A has positive local feature size. The key to this computation is a semi-local notion of quadratic curvature $\kappa_A(p)$ of A at p proposed by Federer [1959]. We first define $\kappa_A(p)$ in the theorem below and defer motivating $\kappa_A(p)$ as an analogue of curvature to Lemma 2.2.20.

Theorem 2.2.16 ([Federer, 1959, Theorem 4.18])**.** *Let A be a closed subset of $\mathbb{R}^d$ with positive local*

feature size. Then

$$\tau_A = \inf_{p \in A} \frac{1}{\kappa_A(p)} \quad \textit{where}$$
$$\kappa_A(p) = \sup_{q \in A \setminus p} \frac{2 d_{\mathbb{R}^d}(q-p, \mathrm{Tan}(p,A))}{\|p-q\|^2}. \tag{2.10}$$

Since the reach of A is the infimum of $\frac{1}{\kappa_A(p)}$, we can derive a lower bound on the reach of A from an upper bound on $\kappa_A(p)$ on A. The advantage of considering $\kappa_A(p)$ is that it is defined simply in terms of the quadratic deviation of points in A away from the tangent set $\mathrm{Tan}(p,A)$ at p. Thus, we can apply simpler geometric arguments to derive lower bounds on the reach, without requiring explicit knowledge of $\mathrm{UP}(A)$.

Remark 2.2.17. As A is assumed to have positive local feature size, the tangent set and normal cones at $p \in A$ are convex cones that are dual to each other in $T_p\mathbb{R}^d$. Hence, the distance to the tangent set $d_{\mathbb{R}^d}(q-p, \mathrm{Tan}(p,A))$ can be calculated from finding the closest orthogonal projection onto vectors in the normal cone, and vice versa:

$$d_{\mathbb{R}^d}(v, \mathrm{Tan}(p,A)) = \max_{\substack{n \in \mathrm{Nor}(p,A) \\ \|n\| \leq 1}} \langle v, n \rangle =: \langle v, \mathrm{Nor}(p,A) \rangle$$
$$d_{\mathbb{R}^d}(v, \mathrm{Nor}(p,A)) = \max_{\substack{t \in \mathrm{Tan}(p,A) \\ \|t\| \leq 1}} \langle v, t \rangle =: \langle v, \mathrm{Tan}(p,A) \rangle .$$

We show below that $\kappa_A(p)$ is well-defined; in fact, $\kappa_A(p)$ it is bounded above by the inverse of the local feature size at p, which is finite as if A has positive local feature size. Thus, the function $\frac{1}{\kappa_A(p)}$ on A is an upper bound of $\tau_A(p)$ with a common infimum that is the reach.

Theorem 2.2.18 ([Federer, 1959, Theorems 4.8(7)])**.** *Let A be a closed subset of $\mathbb{R}^d$ with positive local feature size. Consider two distinct points p and q in A, and $x \in \mathrm{UP}(A)$, where $\pi_A(x) = p$. Then*

$$\left\langle \frac{x-p}{\|x-p\|}, \frac{q-p}{\|q-p\|} \right\rangle \leq \frac{\|q-p\|}{2\tau_A(p)}. \tag{2.11}$$

Corollary 2.2.19. *Let A be a closed subset of $\mathbb{R}^d$ with positive local feature size. Consider two distinct points p and q in A. Then the projection of $q-p$ onto the normal cone of p is bounded above by*

$$d_{\mathbb{R}^d}(q-p, \mathrm{Tan}(p,A)) \leq \frac{\|q-p\|^2}{2\tau_A(p)}. \tag{2.12}$$

Thus, for $\kappa : A \to \mathbb{R}_{\geq 0}$ as defined in eq. (2.10),

$$\kappa_A(p) \leq \frac{1}{\tau_A(p)}.$$

Proof. Due to Theorem 2.2.9, for any vector $n \in \mathrm{Nor}(p, A)$ with $\|n\| \leq 1$, the nearest neighbour in A of point $x = p + \tau_A(p)n$ is simply p. Thus, for two distinct points p and q in A, Theorem 2.2.18 implies

$$\langle n, q - p \rangle = \|n\| \left\langle \frac{x - p}{\|x - p\|}, q - p \right\rangle \leq \frac{\|n\| \|q - p\|^2}{2\tau_A(p)} \leq \frac{\|q - p\|^2}{2\tau_A(p)}. \tag{2.13}$$

Therefor, the distance of $q - p$ to the tangent set of p is given by

$$d_{\mathbb{R}^d}(q - p, \mathrm{Tan}(p, A)) = \langle q - p, \mathrm{Nor}(p, A) \rangle = \max_{\substack{n \in \mathrm{Nor}(p,A) \\ \|n\| \leq 1}} \langle v, n \rangle \leq \frac{\|q - p\|^2}{2\tau_A(p)}$$

where we have applied the formula in Remark 2.2.17. Substituting this bound into the definition of $\kappa_A(p)$ (eq. (2.10)), we derive that $\kappa_A(p)$ is well-defined where p has positive local feature size: $\kappa_A(p) \leq \frac{1}{\tau_A(p)}$. □

We justify our claim that κ_A is analogous to curvature for submanifolds with the following observation. Any path in A through p cannot curve into a Euclidean ball of radius $\frac{1}{\kappa_A(p)}$ whose centre projects onto p, which is the content of the lemma below Lemma 2.2.20.

Lemma 2.2.20. *Let A be a closed subset of $\mathbb{R}^d$ with positive local feature size. Let $\lambda = \kappa_A(p)^{-1}$ where $\kappa_A(p)$ is as defined in eq. (2.10). Then or any $p \in A$, and any unit vector $n \in \mathrm{Nor}(p, A)$,*

$$B_\lambda(p + \lambda n) \cap A = \emptyset.$$

Proof. Consider two distinct points q and p in A and let $\|q - p\| = r$. Pick any $n \in \mathrm{Nor}(p, A)$ and let $\measuredangle(q - p, n) = \theta$. From the definition of $\kappa_A(p)$ in eq. (2.10),

$$\frac{2\langle q - p, n \rangle}{\|q - p\|^2} \leq \kappa_A(p) \implies r \geq 2\lambda \cos(\theta).$$

We now note that $r < 2\lambda \cos(\theta)$ is precisely the parametric expression for the open ball that is the intersection of $B_\lambda(p + \lambda n)$ with the two-dimensional plane containing n and $q - p$, with p

as its origin. Thus, for any $q \in A$, we conclude that $q \notin B_\lambda(p + \lambda n)$. □

Since $\kappa_A(p) \leq \frac{1}{\tau_A(p)}$ (Corollary 2.2.19), we obtain a less informative statement in terms of the local feature size.

Corollary 2.2.21 (See also [Boissonnat et al., 2019, Corollary 2]). *Let A be a non-empty closed subset of $\mathbb{R}^d$. Suppose for some $a \in A$, $\tau_A(a) > 0$. Then for any $n \in \mathrm{Nor}(a, A)$ where $r = \|n\| \leq \tau_A(a)$, the open Euclidean ball $B_r(a + n) \subset \mathbb{R}^d$ centered at $a + n$ does not intersect A.*

2.3 Homology Inference using Finite Point Samples

In this section, we will restrict our focus to compact subsets of Euclidean space with positive reach. The goal of this discussion is to establish theoretical guarantees on whether the homology of A can be reconstructed from a finite collection of points $\mathbb{X} = \{x_1, \ldots, x_N\} \subset \mathbb{R}^d$, given an explicit upper bound on the Hausdorff distance between $\mathbb{X}$ and A (see Definition 2.1.1). As A is compact, the bound on $d_H(A, \mathbb{X})$ implies $\mathbb{X}$ are on a compact domain. In particular, if the points are not sampled directly on A but rather offset away from A in the ambient Euclidean space by random noise, then the distribution of points must have compact support.

For compact Riemannian submanifolds M, Niyogi et al. [2008] showed that the ϵ-thickening of a point cloud $\mathbb{X}$ deformation retracts to M if $d_H(\mathbb{X}, M)$ is sufficiently small relative to τ_M. We generalise to deformation retract construction in Niyogi et al. [2008] to subsets with positive reach $\tau_A > 0$, and thus provide theoretical guarantees for the inference of homotopy types of compact subsets with positive reach from sufficiently Hausdorff close point samples. If we are only interested in the homology $H_\bullet(A; k) =: H_\bullet(A)$ of A with field coefficients k, then we can use persistent homology to infer $H_\bullet(A)$ from $\mathbb{X}$ under less stringent conditions on $d_H(\mathbb{X}, A)$.

2.3.1 The Deformation Retract Argument

We can regard the ϵ-thickening of $\mathbb{X} \subset A$ as an attempt to approximate a tubular neighbourhood of A. Since the ϵ-tubular neightoubrhood $\mathrm{Tub}_\epsilon(A)$ deformation retracts to A for $\epsilon \leq \tau_A$, one may ask whether $\mathbb{X}^\epsilon$ also deformation retracts to A. This has shown been to be the case for compact Riemannian submanifolds with and without boundary.

Theorem 2.3.1 ([Niyogi et al., 2008, Proposition 3.1]). *Let M be a compact Riemannian submanifold and suppose a finite point sample $\mathbb{X}$ of M is $\epsilon/2$-dense. Then for $\epsilon < \tau_M\sqrt{3/5}$, $\mathbb{X}^\epsilon$ deformation retracts to M.*

Theorem 2.3.2 ([Wang and Wang, 2020, Proposition 3.2]). *Let M be a compact Riemannian submanifold* with boundary *and suppose a finite point sample $\mathbb{X}$ of M is $\epsilon/2$-dense. Then for $\epsilon < \tau_M/2$, $\mathbb{X}^\epsilon$ deformation retracts to M.*

Here we show that the deformation retract argument of Niyogi et al. [2008] can be generalised to any subset of $\mathbb{R}$ with positive reach. The key insight is that the geometric arguments used in Niyogi et al. [2008] to derive the deformation retract result relies on geometric properties of subsets with positive reach, rather than properties derived from the differential structure of the manifold. The constructive argument in Niyogi et al. [2008] is motivated by placing conditions on a point cloud $\mathbb{X}$ such that the conditions of the lemma below are satsfied for $\mathbb{X}^\epsilon$ on an open interval for ϵ.

Lemma 2.3.3. *Let A be a subset of $\mathbb{R}^d$ with positive reach τ_A. Suppose $A \subset B \subset A^{\tau_A}$. If for every $p \in B$, the line segment $\overline{p\pi(p)}$ is also contained in B, then B deformation retracts to A.*

Proof. Let us consider the function $F : [0,1] \times B \to \mathbb{R}^d$

$$F(t,p) = (1-t)p + t\pi(p).$$

Since π is continuous (theorem 2.2.1), F is continuous. Furthermore, since we have assumed that $\overline{p\pi(p)}$ is contained in B, $F(t,p) \subset B$. We can easily check that $F(0,p) = p$, $F(1,p) = \pi(p) \in A$, and $F(t,p) = \pi(p) = p$. Therefore, F is a strong deformation retraction $F : [0,1] \times B \to B$ of B onto A. □

We consider three possible configurations of $\mathbb{X}$. First, where $\mathbb{X}$ is an error-free point sample directly on A. In this scenario, $\mathbb{X}$ is r-close to A for $r = 0$ and the only parameters of interest are the density δ and the thickening parrameter ϵ. In the second scenario, we relax the $r = 0$ condition and consider $\mathbb{X}$ a three dimensional parameter space (δ, r, ϵ). In the third scenario, we consider an ambient point cloud $\mathbb{X}$ that is δ-dense in $\mathbb{X}$, and consider thickening a point cloud $\mathbb{X}_r$ that is the collection of points in $\mathbb{X}$ that are r-close to A. In all three scenarios, we

show that there exists an open interval for ϵ depending on δ, r, and τ_A such that $\mathbb{X}^\epsilon$ or $\mathbb{X}^\epsilon_r$ deformation retracts to A.

2.3.1.1 Error-free Point Samples

Lemma 2.3.4. *Assume the conditions of proposition 2.3.5. Consider $x \in \mathbb{X}$ and $p \in B_\epsilon(x)$ where $\epsilon \in (0, \tau_A)$ and $a = \pi(p) \in A$. If $a \notin B_\epsilon(x)$, then there exists some q on the interior of the line segment $\overline{ap}$, such that*

$$\overline{qp} \subset B_\epsilon[x] \quad \textit{and} \quad \|q - a\| \leq \frac{\epsilon^2}{\tau_A(a)}.$$

Proof. Since $\epsilon < \tau_A$, and $p \in B_\epsilon(x)$, we know p has a unique projection onto A. Due to Theorem 2.2.12, the vector $p - a \in \mathrm{Nor}(a, A)$ lies in the normal cone of a. If $\pi(p) = a = x$, then the conditions are trivially satisfied as $\epsilon^2/\tau_A(x) < \epsilon$ and we can take q to be sufficiently close to x along $\overline{qx} \cap B_\epsilon(x)$.

We thus consider the case where $x \neq \pi(p)$. Let $y(t) = (1-t)a + tp$ for $t \in [0,1]$ parametrise the line segment $\overline{ap}$. Since $p - q$ lies in the normal cone of a, it follows from the definition of the normal cone that $y(t) - a = t(p - q)$ also lies in the normal cone $\mathrm{Nor}(a, A)$. Furthermore, since for all $t \in [0,1]$, $\|y(t) - a\| = t\|p - a\| < \tau_A$, we observe that $y(t) \subset \pi^{-1}(a)$.

We consider the continuous function $g(t) = \|y(t) - x\|$. Since $g(0) > \epsilon$ and $g(1) < \epsilon$, by continuity there must be some $t_* \in (0,1)$ such that $g(t_*) = \|y(t_*) - x\| = \epsilon$. Since p is in the interior of $B_\epsilon(x)$ and $B_\epsilon(x)$ is convex, t_* must be unique and $\overline{qp} \subset B_\epsilon[x]$ for $q = y(t_*)$.

Since $\pi(q) = a \in A$ and $x \in A$, we can now apply theorem 2.2.18: for

$$\left\langle \frac{q-a}{\|q-a\|}, \frac{x-a}{\|x-a\|} \right\rangle \leq \frac{\|x-a\|}{2\tau_A(a)}.$$

As such,

$$\begin{aligned}
\epsilon^2 = \|q-x\|^2 &= \|(q-a)+(x-a)\|^2 \\
&= \|q-a\|^2 + \|x-a\|^2 - 2\langle q-a, x-a\rangle \\
&\geq \|q-a\|^2 + \|x-a\|^2\left(1-\frac{\|q-a\|}{\tau_A(a)}\right) \\
&\geq \|q-a\|^2 + \epsilon^2\left(1-\frac{\|q-a\|}{\tau_A(a)}\right) \\
\implies \|q-a\|^2 &\leq \frac{\epsilon^2}{\tau_A(a)}.
\end{aligned}$$

where in the fourth line we applied our assumption that $a \notin B_\epsilon(x)$ and the fact that $\|q-a\| = \|y(t_*)-a\| < \tau_A$. □

Proposition 2.3.5. *Let A be a subset of $\mathbb{R}^d$ with positive reach τ_A. Suppose we have a finite point sample $\mathbb{X} \subset A$ is δ-dense in A where $\delta < \tau_A/4$. Then $\mathbb{X}^\epsilon$ deformation retracts to A if*

$$\frac{\epsilon}{\tau_A} \in \left(\frac{1}{2}-\sqrt{\frac{1}{4}-\frac{\delta}{\tau_A}}, \frac{1}{2}+\sqrt{\frac{1}{4}-\frac{\delta}{\tau_A}}\right).$$

Proof. Consider any point $p \in \mathbb{X}^\epsilon$ and let $\pi(p) = a$. By definition p is contained in at least one open Euclidean ball $B_\epsilon(x)$ for some $x \in \mathbb{X}$. If $a \in B_\epsilon(x)$ for one such x, then $\overline{ap}$ is also contained within $B_\epsilon(x)$ as Euclidean balls are convex. Therefore, $\overline{ap}$ is contained in $\mathbb{X}^\epsilon$.

Let us consider the case where a is not contained in any of the Euclidean balls containing p. From lemma 2.3.4, there exists some $q \in \overline{ap}$ such that $\overline{qp} \subset B_\epsilon[x]$, and $\|q-a\| \leq \epsilon^2/\tau_A(a)$. We can subdivide the line segment $\overline{ap}$ into two segments $\overline{aq}$ and $\overline{ap} \setminus \overline{aq} \subset B_\epsilon(x)$. If there is some $x' \in \mathbb{X}$ such that the closed line segment $\overline{aq}$ is contained in $B_\epsilon(x')$, then the entirety of $\overline{ap}$ is contained in $\mathbb{X}^\epsilon$.

This can be achieved if both a and q are contained in $B_\epsilon(x')$. By the δ-density assumption, we can pick a point $x' \in \mathbb{X}$ such that $a \in B_\delta(x')$. If we assume $\delta < \epsilon$, then $a \in B_\epsilon(x')$.

If q is to be contained in $B_\epsilon(x')$, we require $\|x'-q\| < \epsilon$. If we choose $\delta < \epsilon - \epsilon^2/\tau_A$, then by

the triangle inequality,

$$\begin{aligned}\|x' - q\| &\leq \|x' - a\| + \|a - q\| \\ &< \delta + \frac{\epsilon^2}{\tau_A(a)} \\ &< \epsilon - \frac{\epsilon^2}{\tau_A} + \frac{\epsilon^2}{\tau_A(a)} \\ &\leq \epsilon \qquad\qquad \because \tau_A(a) \geq \tau_A.\end{aligned}$$

Therefore for choices of ϵ that satisfy $\delta < \epsilon - \epsilon^2/\tau_A$, the line $\overline{ap}$ is contained in $\mathbb{X}^\epsilon$. Since this holds for any choice of p, by the lemma 2.3.3, this implies $\mathbb{X}^\epsilon$ deformation retracts to A.

Finally, we have $\delta < \epsilon - \epsilon^2/\tau_A$ if and only if

$$\frac{\epsilon}{\tau_A} \in \left(\frac{1}{2} - \sqrt{\frac{1}{4} - \frac{\delta}{\tau_A}}, \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{\delta}{\tau_A}}\right).$$

□

2.3.1.2 Point Samples with Bounded Offset

Lemma 2.3.6. *Suppose $0 \leq r < \epsilon < \frac{\tau_A}{2}$ and consider a point $x \in A^r$. Then any point $p \in B_\epsilon(x)$ has a unique projection $a = \pi(p) \in A$; if $a \notin B_\epsilon(x)$, then there exists some q in the interior of the line segment $\overline{ap}$, such that*

$$\overline{qp} \subset B_\epsilon[x] \quad \textit{and} \quad \|q - a\| \leq \frac{(\epsilon - r)^2}{2\tau_A}\left(1 + \sqrt{1 + \frac{16\epsilon r \tau_A^2}{(\epsilon - r)^4}}\right).$$

Proof. Since $d_{\mathbb{R}^d}(p, A) \leq \|p - x\| + d_{\mathbb{R}^d}(x, A) < \epsilon + r < \tau_A$, the point p belongs to the tubular neighbourhood of A, and has a unique projection $a = \pi(p) \in A$. If $a = \pi(x)$, then $a \in B_r(x)$, which implies $a \in B_\epsilon(x)$ as $\epsilon > r$, contradicting the assumption that $a \notin B_\epsilon(x)$. Thus, we can exclude the possibility that $\pi(x) = \pi(p)$ in this case.

Following an identical argument in the proof of lemma 2.3.4, we can establish that there is some $q \in \overline{pa}$ such that $\|q - x\| = \epsilon$. Unlike the case considered in lemma 2.3.4, $x \notin A$. However, since $x \in \mathbb{X} \subset A^r \subset A^{\tau_A}$, there is a unique $b \in A$ such that $\pi(x) = b$. Using the

triangle inequality, we can establish that

$$\|q-b\| \leq \|q-x\| + \|x-b\| < \epsilon + r; \quad \text{and}$$
$$\|a-b\| \geq |\|a-x\| - \|x-b\|| > |\epsilon - r|.$$

Since a and $b \in A$, we can now apply theorem 2.2.18:

$$\left\langle \frac{q-a}{\|q-a\|}, \frac{a-b}{\|a-b\|} \right\rangle \leq \frac{\|a-b\|}{2\tau_A}.$$

As such,

$$\begin{aligned}
(\epsilon + r)^2 \geq \|q-b\|^2 &= \|(q-a)+(a-b)\|^2 \\
&= \|q-a\|^2 + \|a-b\|^2 - 2\langle q-a, a-b\rangle \\
&\geq \|q-a\|^2 + \|a-b\|^2\left(1 - \frac{\|q-a\|}{\tau_A}\right) \\
&\geq \|q-a\|^2 + (\epsilon-r)^2\left(1 - \frac{\|q-a\|}{\tau_A}\right) \\
\implies 0 \leq \|q-a\| &\leq \frac{(\epsilon-r)^2}{2\tau_A}\left(1 + \sqrt{1 + \frac{16\epsilon r \tau_A^2}{(\epsilon-r)^4}}\right)
\end{aligned}$$

where in the fourth line we applied the inequality $\|a-b\| \geq |\epsilon - r|$ and the fact that $\|q-a\| < \|p-a\| < \tau_A$. □

Proposition 2.3.7. *Let A be a subset of $\mathbb{R}^d$ with positive reach $\tau_A > 0$ and $\mathbb{X} \subset A^r$ be a point cloud δ-dense in A. If $\delta/\tau_A < 1/4$,*

$$\frac{1}{2} - \sqrt{\frac{1}{4} - \frac{\delta}{\tau_A}} < \frac{\epsilon}{\tau_A} < \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{\delta}{\tau_A}} \quad \text{and} \quad 0 \leq r < \epsilon\left(1 - \frac{2\tau_A}{\epsilon - \delta}\left(1 - \sqrt{1 - \frac{\epsilon-\delta}{\tau_A} + \frac{(\epsilon-\delta)^3}{4\tau_A\epsilon^2}}\right)\right),$$

then $\mathbb{X}^\epsilon$ deformation retracts to A.

Proof. The proof follows that of proposition 2.3.5. We want the point cloud to be sufficiently dense in A (i.e. δ to be sufficient small) so that for every $p \in \mathbb{X}^\epsilon$, the straight line $\overline{p\pi(p)}$ is covered by $\mathbb{X}^\epsilon$. If p and $a = \pi(p)$ are contained in separate balls $B_\epsilon(x)$ and $B_\epsilon(x')$ respectively where $\|x'-a\| < \delta$, then we want δ to be sufficiently small so that $B_\epsilon(x')$ contains the distinguished point $q \in \overline{ap} \cap \partial B_\epsilon(x)$ described in lemma 2.3.6. According to lemma 2.3.6, the point

q is at most

$$\|q-a\| \leq \frac{(\epsilon-r)^2}{2\tau_A}\left(1+\sqrt{1+\frac{16\tau_A^2\epsilon r}{(\epsilon-r)^4}}\right)$$

away from a if $\epsilon + r < \tau_A$. Assuming that this is so, then $\overline{ap} \subset B_\epsilon(x) \cup B_\epsilon(x') \subset \mathbb{X}^\epsilon$. By the triangle inequality, we have

$$\begin{aligned}\|q-x'\| &\leq \|q-a\| + \|a-x'\| \\ &\leq \frac{(\epsilon-r)^2}{2\tau_A}\left(1+\sqrt{1+\frac{16\tau_A^2\epsilon r}{(\epsilon-r)^4}}\right) + \delta.\end{aligned}$$

If the right hand side of the inequality is strictly less than ϵ, then we achieve the sufficient condition for $q \in B_\epsilon(x')$. One can then show that the system inequalities

$$\frac{(\epsilon-r)^2}{2\tau_A}\left(1+\sqrt{1+\frac{16\tau_A^2\epsilon r}{(\epsilon-r)^4}}\right) + \delta < \epsilon \quad \text{and} \quad r+\epsilon < \tau_A \quad \text{and} \quad \epsilon > r > 0 \quad \text{and} \quad \delta > 0 \tag{2.14}$$

reduces to $\delta/\tau_A < 1/4$ and

$$\begin{aligned}&\frac{1}{2} - \sqrt{\frac{1}{4}-\frac{\delta}{\tau_A}} < \frac{\epsilon}{\tau_A} < \frac{1}{2} + \sqrt{\frac{1}{4}-\frac{\delta}{\tau_A}}; \quad \text{and} \\ &\qquad 0 \leq r < \epsilon\left(1-\frac{2\tau_A}{\epsilon-\delta}\left(1-\sqrt{1-\frac{\epsilon-\delta}{\tau_A}+\frac{(\epsilon-\delta)^3}{4\tau_A\epsilon^2}}\right)\right).\end{aligned}$$

Having shown that $\overline{ap} \subset B_\epsilon(x) \cup B_\epsilon(x') \subset \mathbb{X}^\epsilon$ for any $p \in \mathbb{X}^\epsilon$, we can apply lemma 2.3.3 which implies $\mathbb{X}^\epsilon$ deformation retracts to A. □

Proposition 2.3.8. *Let $\mathbb{X} \subset \mathbb{R}^d$ be δ-dense in A, a subset of $\mathbb{R}^d$ with positive reach $\tau_A > 0$. Let $\mathbb{X}_r = \mathbb{X} \cap A^r$. If $\delta < \frac{\sqrt{2}-1}{8}\tau_A$, there is an open interval of choices for ϵ*

$$\delta + \frac{\tau_A}{2} - \sqrt{\frac{\tau_A}{4} - 2(\sqrt{2}+1)\delta} < \epsilon < \delta + \frac{\tau_A}{2} + \sqrt{\frac{\tau_A}{4} - 2(\sqrt{2}+1)\delta}$$

and an open interval of choices for r

$$\delta \leq r < \epsilon - \frac{2\epsilon\tau_A}{\epsilon-\delta}\left(1-\sqrt{1-\frac{\epsilon-\delta}{\tau_A}+\frac{(\epsilon-\delta)^3}{4\tau_A\epsilon^2}}\right)$$

such that $\mathbb{X}_r$ is non-empty and $\mathbb{X}_r^\epsilon$ deformation retracts to A.

Proof. For concision, we express all quantities 'in units of τ_A': i.e. δ, ϵ, and r here refer to δ/τ_A, ϵ/τ_A, r/τ_A, without loss of generality.

We start with a simple observation: if $\mathbb{X}$ is δ-dense in A, then so is $\mathbb{X}_r = \mathbb{X} \cap A^r$ for $r \geq \delta$. $\mathbb{X}$ is δ-dense if and only if

$$A \subset \bigcup_{x \in \mathbb{X}} B_\delta(x) = \mathbb{X}^\delta.$$

We now observe that $B_\delta(x) \cap A = \emptyset$ if $d_{\mathbb{R}^d}(x, A) \geq \delta$. Therefore, we can deduce that for all $r \geq \delta$,

$$A = A \cap \mathbb{X}^\delta = A \cap \mathbb{X}_r^\delta \subset \mathbb{X}_r^\delta.$$

In other words, the subset $\mathbb{X}_r \subset \mathbb{X}$ of our original point cloud remains δ-dense in A if $r \geq \delta$.

Having obtained a point cloud $\mathbb{X}_r \subset A^r$ that is δ-dense in A, we now ask whether we can find choices of (ϵ, r) that are compatible with the conditions of proposition 2.3.7 and our restriction $r \geq \delta$, so that $\mathbb{X}_r^\epsilon$ deformation retracts to A. Specifically, we would like to pick some

$$\epsilon \in \left(\frac{1}{2} - \sqrt{\frac{1}{4} - \delta},\ \frac{1}{2} + \sqrt{\frac{1}{4} - \delta} \right) \tag{2.15}$$

so that the interval

$$r \in \left(\delta, \quad \epsilon - \frac{2\epsilon}{\epsilon - \delta} \left(1 - \sqrt{1 - (\epsilon - \delta) + \frac{(\epsilon - \delta)^3}{4\epsilon^2}} \right) \right) \tag{2.16}$$

is not empty. We now proceed to derive the conditions on ϵ for the interval to be non-empty. Since eq. (2.15) implies $\epsilon > \delta$, we can bound δ in terms of $\epsilon - \delta$:

$$\begin{aligned} & \delta < \epsilon - \frac{2\epsilon}{\epsilon - \delta} \left(1 - \sqrt{1 - (\epsilon - \delta) + \frac{(\epsilon - \delta)^3}{4\epsilon^2}} \right) \\ \iff & \delta < \frac{\epsilon - \delta}{2} \left(\sqrt{2 - (\epsilon - \delta)} - 1 \right). \end{aligned}$$

We can show that for $\epsilon - \delta \in (0,1)$ and $\delta \in (0, \frac{\sqrt{2}-1}{8})$,

$$\left(\epsilon - \delta - \frac{1}{2}\right)^2 < \frac{1}{4} - 2(\sqrt{2}+1)\delta \iff \delta < \frac{\sqrt{2}-1}{2}(\epsilon-\delta)(1-(\epsilon-\delta))$$
$$\implies \delta < \frac{\epsilon-\delta}{2}\left(\sqrt{2-(\epsilon-\delta)}-1\right),$$

and we therefore solve the quadratic equation to obtain a slightly stronger bound on ϵ

$$\epsilon \in \left(\delta + \frac{1}{2} - \sqrt{\frac{1}{4} - 2(\sqrt{2}+1)\delta}, \quad \delta + \frac{1}{2} + \sqrt{\frac{1}{4} - 2(\sqrt{2}+1)\delta}\right).$$

Choosing a thickening parameter ϵ in this interval ensures that there is a non-empty interval eq. (2.16) of the noise tolerance parameter r. Finally, we can check that for $\delta \in (0, \frac{\sqrt{2}-1}{8})$, the range of ϵ such that $\mathbb{X}^r_\epsilon$ deformation retracts to A is contained within the bounds on ϵ prescribed by proposition 2.3.7:

$$\frac{1}{2} - \sqrt{\frac{1}{4} - \delta} < \delta + \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{\sqrt{2}+1}{2}\delta} < \delta + \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{\sqrt{2}+1}{2}\delta} < \frac{1}{2} + \sqrt{\frac{1}{4} - \delta}.$$

□

Proposition 2.3.9. *Let A be a subset of $\mathbb{R}^d$ with positive reach $\tau_A > 0$ and $\mathbb{X}$ be a point cloud such that $d_H(\mathbb{X}, A) = \alpha\tau_A$, where $\alpha < \frac{82-19\sqrt{19}}{108}$. If*

$$\frac{\epsilon}{\tau_A} \in \left(20\frac{1846+431\sqrt{19}}{1503}\alpha, \frac{1}{2} - \left(\frac{79+18\sqrt{19}}{5}\right)\alpha\right),$$

then $\mathbb{X}^\epsilon$ deformation retracts to A.

Proof. The assumption that $d_H(\mathbb{X}, A) = \alpha$ implies $\mathbb{X}$ is α-dense and α-close to A. Note that this is a weaker assumption than having separate parameters for δ-density and r-closenesss. Thus, this places stronger constraints on the region of parameter space (ϵ, δ, r) given in eq. (2.14) where $\mathbb{X}^\epsilon$ deformation retracts to A. Setting $\delta = r = \alpha$ in eq. (2.14), and expressing quantities in units of τ_A, we obtain the following admissible region in the parameter space (ϵ, α) where $\mathbb{X}^\epsilon$ deformation retracts to A:

$$(\epsilon-\alpha)^2 + \sqrt{(\epsilon-\alpha)^4 + 16\epsilon\alpha} < \epsilon - \alpha \quad \text{and} \quad \epsilon + \alpha < 1 \quad \text{and} \quad \epsilon > \alpha > 0.$$

One can then show that the maximum α that satisfies these constraints is $\alpha < \frac{82-19\sqrt{19}}{108}$, and the triangle defined by

$$\epsilon \in \left(20\frac{1846+431\sqrt{19}}{1503}\alpha, \frac{\tau_A}{2} - \frac{79+18\sqrt{19}}{5}\alpha\right) \quad \text{and} \quad \alpha < \frac{82-19\sqrt{19}}{108}$$

is contained in the admissible region, and the interval for ϵ vanishes only when $\alpha \geq \frac{82-19\sqrt{19}}{108}$.

□

In Propositions 2.3.5, 2.3.7 and 2.3.8, we have assumed that we have a precise value of τ_A. We now assume that we only have some positive lower bound of the reach. Propositions 2.3.10 to 2.3.12 can be shown by substituting the stated bounds directly into Propositions 2.3.5, 2.3.7 and 2.3.8 respectively

Proposition 2.3.10. *Let $\mathbb{X} \subset A$ be an δ-dense sample of $A \subset \mathbb{R}^d$ and suppose we have a positive lower bound $\tau_- \in (0, \tau_A)$ on the reach. If $\delta \leq \tau_-/4$, then $\mathbb{X}^{2\delta}$ deformation retracts to A.*

Proposition 2.3.11. *Let $\mathbb{X} \subset A^r$ be an δ-dense sample of $A \subset \mathbb{R}^d$ and suppose we have a positive lower bound $\tau_- \in (0, \tau_A)$ on the reach. If*

$$16r \leq \delta \leq \frac{\tau_-}{8},$$

then $\mathbb{X}^{2\delta}$ deformation retracts to A.

Proposition 2.3.12. *Let $\mathbb{X} \subset \mathbb{R}^d$ be an δ-dense sample of $A \subset \mathbb{R}^d$ and suppose we have a positive lower bound $\tau_- \in (0, \tau_A)$ on the reach. If $\delta \leq \frac{\sqrt{2}-1}{8}\tau_-$, then for*

$$1 \leq \frac{r}{\delta} \leq \left(1+\sqrt{2}\right)\left(2\sqrt{6\left(5+2\sqrt{2}\right)} - 3\left(3+\sqrt{2}\right)\right) \quad \textit{and} \quad \epsilon = \left(4\sqrt{2}+5\right)\delta,$$

$\mathbb{X}^\epsilon$ deformation retracts to A.

2.3.2 Persistent Homology

If we are only interested in recovering the homology of A from $\mathbb{X}$ rather than the homotopy type, then we do not need to use the deformation retract argument, which requires stringent bounds on the density and closeness of $\mathbb{X}$ to A. Instead, we can use persistent homology.

2.3.2.1 Background: Persistence

A *filtration* here refers to a functor from the poset $(\mathbb{R}, \geq)$ to the category of topological spaces; explicitly, a filtration can be expressed as the following sequence

$$\mathsf{X}: \quad \cdots \to X_s \xrightarrow{\varphi(s \leq t)} X_t \to \cdots \tag{2.17}$$

where X_s are topological spaces indexed over $\mathbb{R}$, and we have a family of continuous maps $\varphi(t \geq s) : X_s \to X_t$ indexed over ordered pairs $(t \geq s)$ which respect composition $\varphi(r \geq s) = \varphi(r \geq t) \circ \varphi(t \geq s)$, and $\varphi(t \geq t)$ is the identity map on X_t.

Definition 2.3.13. Given a pair of filtrations X and Y, we say the following

$$(\mathsf{X}, \mathsf{Y}): \quad \cdots \to (X_s, Y_s) \xrightarrow{\varphi(s \leq t)} (X_t, Y_s) \to \cdots \tag{2.18}$$

is a *compatible filtration* of a pair if we have maps $Y_s \to X_s$ for every $s \in \mathbb{R}$, and for every $s \leq t$, the following square commutes:

$$\begin{array}{ccc} Y_s & \longrightarrow & Y_t \\ \downarrow & & \downarrow \\ X_s & \longrightarrow & X_t \end{array} . \tag{2.19}$$

Let $A^\epsilon = \bigcup_{p \in A} B_\epsilon(p)$ be the ϵ-thickening of A in the ambient Euclidean space. We refer the following sequence of inclusions

$$\check{\mathsf{C}}(A): \quad A \hookrightarrow \cdots \hookrightarrow A^s \xhookrightarrow{s \leq t} A^t \hookrightarrow \cdots \tag{2.20}$$

as the *Čech filtration* of A. If we have a pair of subsets $B \subset A$, then we refer to the following sequence of inclusions of pairs as the Čech filtration of (A, B):

$$\check{\mathsf{C}}(A, B) := \big(\check{\mathsf{C}}(A), \check{\mathsf{C}}(B)\big): \quad (A, B) \hookrightarrow \cdots \hookrightarrow (A^s, B^s) \xhookrightarrow{s \leq t} (A^t, B^t) \hookrightarrow \cdots \tag{2.21}$$

where one can check that this is a compatible filtration of a pair are satisfied since all the maps in eq. (2.19) are inclusions.

By applying the homology with field coefficients to these filtrations of topological spaces, we obtain a persistence module:

Definition 2.3.14 (Persistence Module)**.** A persistence module [Edelsbrunner et al., 2002, Zomorodian and Carlsson, 2005] is a a sequence of vector spaces V_t indexed over $t \in \mathbb{R}$, and maps $\theta(s \leq t) : V_s \to V_t$ for any $s \leq t$, such that

1. $\theta(t \leq t)$ is the identity on V_t; and
2. For $r \leq s \leq t$, we have $\theta(r \leq t) = \theta(s \leq t) \circ \theta(r \leq s)$.

We restrict our considerations to persistence modules that are 'tame', such as *pointwise finite dimensional* (p.f.d.) modules, where all of the vector spaces are finite dimensional. P.f.d. modules can be classified by their decomposition into constituent *interval modules*. For an interval $\langle b, d \rangle \subset \mathbb{R}$ (where $\langle b, d \rangle$ may be open, closed or half-closed), the interval module supported on $\langle b, d \rangle$ is a persistence module where a vector space V_t is one dimensional if and only if $t \in \langle b, d \rangle$, with isomorphisms between vector spaces indexed in $\langle b, d \rangle$ and zero maps elsewhere.

Theorem 2.3.15 (Interval Decomposition Gabriel [1972], Crawley-Boevey [2015])**.** *Any p.f.d. persistence module can be uniquely decomposed as a summand of interval modules up to isomorphism and permutation of indices.*

The multiset of the support of the interval modules of a p.f.d. persistence module V is called the module's *barcode* $|||V|||$.

For simplicity, let us consider singular homology with field coefficients. Applying homology as a functor from the category of topological spaces to that of vector spaces , we can send a filtration X to a persistence module which we call the *persistent homology* of the filtration X:

$$H_i(\mathsf{X}) : \quad \cdots \to H_i(X_s) \xrightarrow{\varphi(s \leq t)} H_i(X_t) \to \cdots . \tag{2.22}$$

We can similarly apply relative homology as a function to a filtration of a pair (X, Y), and obtain a persistence module that is the *persistent relative homology* of (X, Y):

$$H_i(\mathsf{X}, \mathsf{Y}) : \quad \cdots \to H_i(X_s, Y_s) \xrightarrow{\varphi(s \leq t)} H_i(X_t, Y_t) \to \cdots . \tag{2.23}$$

Here we are concerned with the persistent homology of the Čech filtrations of subsets or pairs

of subsets in Euclidean space:

$$\check{\mathsf{H}}_i(A) := H_i(\check{\mathsf{C}}(A)): \quad H_i(A) \to \cdots \to H_i(A^s) \xrightarrow{s\leq t} H_i(A^t) \to \cdots \tag{2.24}$$

$$\check{\mathsf{H}}_i(A,B) := H_i(\check{\mathsf{C}}(A,B)): \quad H_i(A,B) \to \cdots \to H_i(A^s,B^s) \xrightarrow{s\leq t} H_i(A^t,B^t) \to \cdots. \tag{2.25}$$

If we consider the barcode of the persistent homology of a filtration, we can interpret the interval $\langle b,d\rangle$ as a homological feature that is born at index b in the filtration and extinguished at d, persisting for a duration of $d-b$ in the filtration.

2.3.2.2 Background: Metrics

We recall two popular (and equivalent) metrics for comparing persistence modules and their corresponding barcodes respectively: the interleaving and bottleneck distances.

We say two persistence modules V and W are *δ-interleaved* [Chazal et al., 2009b] if for all $s \in \mathbb{R}$, all triangles and parallelograms in the diagram below commute:

$$\begin{array}{ccccc} V_s & \longrightarrow & V_{s+\delta} & \longrightarrow & V_{s+2\delta} \\ & \searrow \quad \nearrow & & \searrow \quad \nearrow & \\ W_s & \longrightarrow & W_{s+\delta} & \longrightarrow & W_{s+2\delta}. \end{array} \tag{2.26}$$

The interleaving distance $d_I(V,W)$ between persistence modules V and W is the infimum over all δ where V and W are δ-interleaved. We can similarly define an interleaving between filtrations, where the vector spaces and linear transformations in the diagam above are replaced by topological spaces and continuous maps that obey the same commutativity requirements .

If two filtrations are δ-interleaved, then the functoriality of homology implies we have an interleaving of their persistent homology modules. The following theorem states that interleaving on the level of individual filtrations extends to an interleaving of the relative persistent homology of the corresponding pairs if the filtrations are compatible.

Theorem 2.3.16 ([Skraba and Wang, 2012, Theorem 1]). *Suppose (X,Y) and $(\mathsf{X}',\mathsf{Y}')$ are compatible filtrations of pairs as defined in Definition 2.3.13. If X and X' are α-interleaved, and Y and Y' are β-interleaved, then the relative persistent homology modules $H_i(\mathsf{X},\mathsf{Y})$ and $H_i(\mathsf{X}',\mathsf{Y}')$ are $\max(\alpha,\beta)$-interleaved.*

From the definition of the Hausdorff distance in Definition 2.1.1, $d_H(A', A) = \alpha$ for two subsets A and A' of $\mathbb{R}^d$ if and only if $d_I(\check{\mathrm{C}}(A), \check{\mathrm{C}}(A')) = \alpha$. Since Čech filtrations of a pair are compatible filtrations, Theorem 2.3.16 implies the following for Čech filtrations.

Corollary 2.3.17. *If two subsets A and A' have Hausdorff distance $d_H(A, A') = \alpha$, then we have a δ-interleaving between their persistent homology $\check{\mathrm{H}}_i(A)$ and $\check{\mathrm{H}}_i(A')$ for all $\delta > \alpha$; thus $d_I(\check{\mathrm{H}}_i(A), \check{\mathrm{H}}_i(A')) \leq \alpha$. Furthermore, if we have $B \subset A$ that is β interleaved with $B' \subset A'$, then we have a δ-interleaving between $\check{\mathrm{H}}_i(A, B)$ and $\check{\mathrm{H}}_i(A', B')$ where $\delta > \max(\alpha, \beta)$; thus $d_I(\check{\mathrm{H}}_i(A, B), \check{\mathrm{H}}_i(A', B')) \leq \max(\alpha, \beta)$.*

Comparing persistence modules via their barcodes, we say there is a *δ-matching* [Bauer and Lesnick, 2015] between barcodes B and C if there is bijection between two subsets of intervals in $B_* \subset B$ and $C_* \subset C$, such that B_* and C_* contain intervals with persistence greater than 2δ in B and C respectively; furthermore, for a pair of matched intervals $\langle b, d \rangle$ and $\langle b', d' \rangle$,

$$\langle b, d \rangle \subseteq \langle b' - \delta, d' + \delta \rangle \quad \text{and} \quad \langle b', d' \rangle \subseteq \langle b - \delta, d + \delta \rangle .$$

The *bottleneck distance* [Bauer and Lesnick, 2015, Cohen-Steiner et al., 2007] $d_b(B, C)$ between barcodes B and C is the infimum over all δ where there is a δ-matching between B and C. The following theorem shows that comparing persistence modules with the interleaving distance is equivalent to comparing their barcodes using the bottleneck distance.

Theorem 2.3.18 (Isometry Theorem [Bauer and Lesnick, 2015])**.** *Two p.f.d. persistence modules V and W are δ-interleaved if and only if there exists a δ-matching between their corresponding barcodes $|||V|||$ and $|||W|||$. In particular, $d_b(|||V|||, |||W|||) = d_I(V, W)$.*

2.3.2.3 Persistence of Subsets with Positive Reach

We now consider the persistent homology of the Čech filtration of a subset A with positive reach $\tau_A > 0$, and pairs $A \supset B$ of subset where τ_B is also positive.

Lemma 2.3.19. *Suppose A is a subset of $\mathbb{R}^d$ with positive reach, and $H_i(A)$ is finite dimensional for all i. Consider $\check{\mathrm{H}}_i(A)$, the persistent homology of the Čech filtration of A. Then the only intervals in the barcode $|||\check{\mathrm{H}}_i(A)|||$ born during $[0, \tau_A)$ are those of the form $[0, d\rangle$, where $d \geq \tau_A$, and there are precisely $\dim H_i(A)$ such intervals.*

Similarly, if we have $B \subset A$ such that $H_i(B)$ is finite dimensional for all i, and $\tau_{AB} := \min(\tau_A, \tau_B) >$

0 , *then the only intervals in the barcode* $\left|\left|\left|\check{\mathsf{H}}_i(A,B)\right|\right|\right|$ *born during* $[0,\tau_{AB})$ *are those of the form* $[0,d\rangle$, *where* $d \geq \tau_{AB}$, *and there are precisely* dim $H_i(A,B)$ *such intervals.*

Proof. As a consequence of Proposition 2.2.13, the deformation retract of A^t onto A for $t < \tau_A$ implies the linear maps in the persistence module $\check{\mathsf{H}}_i(A)$ are isomorphisms for all $0 \leq s \leq t < \tau_A$:

$$H_i(A^s) \xrightarrow{\cong} H_i(A^t).$$

Thus, these isomorphisms imply the only intervals in $\left|\left|\left|\check{\mathsf{H}}_i(A)\right|\right|\right|$ born during $[0,\tau_A)$ can only be those of the form $[0,d\rangle$, which must die at some time $d \geq \tau_A$.

For the pair of subsets (A,B), the deformation retract of B^t onto B for $t < \tau_B$ also implies

$$H_i(B^s) \xrightarrow{\cong} H_i(B^t).$$

for all $0 \leq s \leq t < \tau_B$. Since we also have $H_i(A^s) \xrightarrow{\cong} H_i(A^t)$ for all $0 \leq s \leq t < \tau_A$:, these isomorphisms innduce an isomorphism on the level of relative homology for $t < \tau_{AB} := \min(\tau_A, \tau_B)$

$$H_i(A,B) \xrightarrow{\cong} H_i(A^t, B^t)$$

due to applying the five lemma between the homology long exact sequences of the pairs (A,B) and (A^t,B^t). Thus these isomorphisms imply the only intervals born during $[0,\tau_{AB})$ can only be those of the form $[0,d\rangle$, which must die at some time $d \geq \tau_{AB}$. □

This simple observation leads us to the following result that allows us to infer the rank of the homology of subsets with positive reach using the persistent homology of the Čech filtration of point clouds. Compared to the conditions placed on the point clouds for homology inference via deformation retract arguments, these conditions are far less stringent.

Proposition 2.3.20. *Suppose* $A \subset \mathbb{R}^d$ *and a point sample* $\mathbb{X} \subset \mathbb{R}^d$ *have Hausdorff distance* $d_H(A,\mathbb{X}) < \frac{\tau_A}{4}$. *If the homology of A has finite rank, Then*

$$\text{rank } H_i(A) = \#\left\{ \langle b,d\rangle \in \left|\left|\left|\check{\mathsf{H}}_i(\mathbb{X})\right|\right|\right| \ : \ b < \frac{\tau_A}{4} \quad \textit{and} \quad d > \frac{3\tau_A}{4} \right\}.$$

Furthermore, if we have $B \subset A$ *and a subsample* $\mathbb{Y} \subset \mathbb{X}$, *where* $d_H(B,\mathbb{Y}) < \frac{\tau_B}{4}$, *and the relative*

homology of (A, B) *has finite rank, then for* $\tau_{AB} = \min(\tau_A, \tau_B)$,

$$\text{rank } H_i(A,B) = \#\left\{ \langle b, d\rangle \in \left|\left|\left| \check{\mathsf{H}}_i(\mathbb{X}, \mathbb{Y}) \right|\right|\right| \ : \ b < \frac{\tau_{AB}}{4} \quad \textit{and} \quad d > \frac{3\tau_{AB}}{4} \right\}.$$

Proof. Let $d_H(A, \mathbb{X}) = \alpha$. We first apply Corollary 2.3.17 and deduce that $d_I(\check{\mathsf{H}}_i(A), \check{\mathsf{H}}_i(\mathbb{X})) \leq \alpha$. The isometry theorem (Theorem 2.3.18) implies for $\frac{\tau_A}{4} > \delta > \alpha$, there is a δ-matching between the barcodes of $\check{\mathsf{H}}_i(A)$ and $\check{\mathsf{H}}_i(\mathbb{X})$. Cosider the collection of intervals in the barcode of $\check{\mathsf{H}}_i(A)$, which are born at 0 with persistence greater than τ_A. These are the intervals that correspond to $H_i(A)$. As their persistence is greater than 2δ, they are thus matched to some set of intervals $\langle b', d'\rangle$ in the barcode of $\check{\mathsf{H}}_i(\mathbb{X})$, which satisfy

$$b' < \delta < \frac{\tau_A}{4} \quad \text{and} \quad d' > \tau_A - \delta > \frac{3}{4}\tau_A \implies d' - b' > \tau_A - 2\delta > 2\delta. \tag{2.27}$$

We now consider the converse, and consider intervals in the barcode of $\check{\mathsf{H}}_i(\mathbb{X})$ that satisfy the conditions above, and the possible intervals in the barcode of $\check{\mathsf{H}}_i(A)$ with whom they can be matched. As their persistence is greater than 2δ, all such intervals are matched with intervals $\langle b, d\rangle$ in the barcode of $\check{\mathsf{H}}_i(A)$, which satisfy

$$b < b' + \delta < 2\delta < \frac{\tau_A}{2} \quad \text{and} \quad d > \tau_A - 2\delta > \frac{\tau_A}{2}.$$

The only such intervals in the barcode of $\check{\mathsf{H}}_i(A)$ are the β_i intervals born at 0 with persistence greater than τ_A. As such, we have a bijection between intervals in $\check{\mathsf{H}}_i(\mathbb{X})$ satisfying eq. (2.27) and the intervals corresponding to the homology of A.

We can similarly apply the same arguments to deduce the analogous result in the relative case, as Corollary 2.3.17 and Theorem 2.3.18 also establish a δ-matching between the barcodes of $\left|\left|\left| \check{\mathsf{H}}_i(\mathbb{X}, \mathbb{Y}) \right|\right|\right|$ and $\left|\left|\left| \check{\mathsf{H}}_i(A, B) \right|\right|\right|$, the latter having no intervals born and destroyed on $(0, \tau_{AB})$. □

Remark 2.3.21. Compare the conditions on the Hausdorff distance between A and a point sample $\mathbb{X}$ in Proposition 2.3.20 to those in Proposition 2.3.9, where the latter is far more stringent. This is because in using persistent homology to infer the homology of A, we do not require the deformation retract of $\mathbb{X}^t$ onto A for any t, whereas Proposition 2.3.9 demands that there must be some open interval of thickening parameters where this is true. This illustrates

the advantages of using persistent homology where short intervals in the barcode of $\check{\mathrm{H}}_i(\mathbb{X})$ which cause $H_i(\mathbb{X}^t) \ncong H_i(A)$ for any t can be allowed without getting in the way of inferring topologically significant features.

Chapter 3

Submanifolds of Euclidean Space

In this chapter, we restrict our consideration to smooth, m-dimensional Riemannian submanifolds $M \subset \mathbb{R}^d$ that are *properly embedded* into $\mathbb{R}^d$. That is, the embedding $M \hookrightarrow \mathbb{R}^d$ is a proper map. An embedded submanifold is properly embedded in $\mathbb{R}^d$ if and only if it is a closed subset of $\mathbb{R}^d$ ([Lee, 2012, Proposition 5.5]).

Following on from the discussion about subsets of Euclidean space with positive local feature size, we first establish that M has positive local feature size. Thus, M inherits all of the geometric properties of subsets with positive local feature size that we have established in the previous chapter. Building on Niyogi et al. [2008] and Boissonnat et al. [2019], we then explore the consequences of curvature for the local geometry of the manifold, such as errors in approximating the geodesic metric of the manifold with the ambient Euclidean distance. We also show that for a radius r sufficiently small relative to the local feature size $\tau_M(p)$ of M at p, we have an open neighbourhood of p in M given by $B_r(p) \cap M$ that are diffeomorphic to its orthogonal projections onto the tangent plane $p + T_pM$.

We then turn our attention to smooth functions on M. Sublevel sets of smooth functions at regular values are *regular domains* of M, i.e. properly embedded codimension-0 submanifolds with boundary of M. We consider regular domains of M that are embedded into $\mathbb{R}^d$ along a proper embedding $M \hookrightarrow \mathbb{R}^d$. We derive an explicit lower bound of the reach of regular domains in terms of bounds on the first and second derivatives of functions that define them. We then discuss how we can recover the homology of compact regular domains with finite point

samples, which involves applying our findings in the discussion on orthogonal projections to results by Wang and Wang [2020].

Finally, we show that regular domains that intersect transversely in M have positive local feature size in the ambient Euclidean space. We also derive an explicit lower bound of their reach in terms of bounds on the first and second derivatives of the functions that define the intersecting regular domains.

As a matter of notation, we will let $d_M(p,q)$ denote the geodesic distance between two points p and q on the manifold M, and $\|x-y\|$ denote the Euclidean distance between two points x and y in $\mathbb{R}^d$. We also let

$$B_t^M(p) = \{q \in M \ : \ d_M(p,q) < s\}$$
$$B_t^M[p] = \{q \in M \ : \ d_M(p,q) \leq s\}$$

respectively denote open and closed geodesic balls of radius t on the manifold. Note that $B_r(x)$ and $B_r[x]$ denote the open and closed Euclidean balls respectively.

3.1 Local Feature Size

We show that such submanifolds have positive local feature size. This follows directly from the following classical result in differential geometry. We recall that the normal bundle $\mathcal{N}(M)$ of a submanifold M of $\mathbb{R}^d$ is a d-dimensional submanifold of the tangent bundle $\mathcal{T}(\mathbb{R}^d)$, defined as $\mathcal{N}(M) = \{(p,v) \ : \ p \in M \quad \text{and} \quad v \in T_p\mathbb{R}^d \quad \text{s.t.} \quad v \perp T_pM\}$.

Theorem 3.1.1 ([Lee, 2012, Theorem 6.24]). *For every properly embedded smooth m-dimensional submanifold $M \subset \mathbb{R}^d$, there exists a positive continuous function $\delta : M \to (0,\infty)$, and an open neighbourhood of M in $\mathbb{R}^d$ that is diffeomorphic to the image of the following open subset in the normal bundle of M*

$$\mathcal{U} = \{(x,v) \in \mathcal{N}(M) \ : \ \|v\| < \delta(x)\} \subset \mathcal{N}(M) \tag{3.1}$$

under the map $\mathcal{E} : \mathcal{N}(M) \to \mathbb{R}^d$ from the normal bundle $\mathcal{N}(M) \subset \mathcal{T}(\mathbb{R}^d)$ to $\mathbb{R}^d$, where for $p \in M$, the map

$$\mathcal{E}(p,v) = p + v$$

maps $v \in N_pM$ to an element in the affine plane centered at p, that is the orthogonal complement of the tangent plane to M at p.

Note that this theorem implies for any $x \in \mathcal{U}$, we invert the diffeomorphism $\mathcal{E}$ and associate to it a unique point $p \in M$ where $x - p \in N_pM$ and $\|x - p\| < \delta(p)$.

Proposition 3.1.2. *Every properly embedded m-dimensional smooth submanifold $M \subset \mathbb{R}^d$ has a positive local feature size.*

Proof. Suppose the contrary that M does not have positive local feature size. In other words, there is some $p \in M$ such that for every $r > 0$, there is some point $x_r \in \text{Med}(M)$ such that $\|p - x_r\| < r$.

Let $\mathcal{E}$ be the map from the normal bundle to $\mathbb{R}^d$ as defined in theorem 3.1.1 and $\mathcal{U}$ be as defined in eq. (3.1). As $\mathcal{E}(\mathcal{U})$ is an open neighbourhood of M, let r be sufficiently small such that $B_p(r) \subset \mathcal{E}(\mathcal{U})$. Thus $x_r \in \mathcal{E}(\mathcal{U})$. Let p_r be the point in M such that $\mathcal{E}(p_r, x_r - p_r) = x_r$. As $x_r \in \text{Med}(M)$, there is at least one nearest neighbour $q_r \in \text{NN}_M(x_r)$ of x_r in M that is distinct from p_r, such that

$$\|x_r - q_r\| \leq \|x_r - p_r\| < r$$

and $x_r - q_r \in N_{q_r}M$. Let $\delta : M \to (0, \infty)$ be the continuous, positive function as defined in Theorem 3.1.1. Since $\mathcal{E}$ is a diffeomorphism from $\mathcal{E}$ onto $\mathcal{E}(\mathcal{U})$, and $\mathcal{E}(p_r, x_r - p_r) = x_r$, for no $v \in N_{q_r}M$ where $\|v\| < \delta(q_r)$ such that $\mathcal{E}(q_r, v) = x_r$. Thus $\|x_r - q_r\| \geq \delta(q_r)$. Since $\|x_r - q_r\| < r$, this in turn implies $r > \delta(q_r)$.

We can bound the distance between q_r and p using the triangle inequality:

$$\|p - q_r\| \leq \|x_r - p\| + \|x - q_r\| < 2r.$$

Thus as $r \to 0$, the sequence $q_r \to p$. As δ is continuous,

$$\delta(p) = \lim_{r \to 0} \delta(q_r).$$

However, as $r > \delta(q_r)$, the limit on the right hand side evaluates to zero as we take $r \to 0$, implying $\delta(p) = 0$. This is a contradiction. □

We next show that the tangent and normal cones of Federer as defined in Definition 2.2.6 coincide with the usual notions of tangent and normal planes for manifolds.

Proposition 3.1.3. *If M is a properly embedded smooth submanifold of $\mathbb{R}^d$, then for any $p \in M$,*

$$\mathrm{Tan}(p, M) = T_pM \quad \textit{and} \quad \mathrm{Nor}(p, M) = N_pM.$$

Proof. We note that $p \in M$ is the projection $p = \pi(x)$ of $x \in B_{\tau_M(p)}(p) \subset \mathbb{R}^d \setminus \mathrm{Med}(M)$ if and only if p is a critical point of the restriction of the Euclidean distance to x to M, i.e. $x - p \in N_pM$. Therefore $\pi^{-1}(p) \cap B_{\tau_M(p)}(p) = \big(p + N_pM\big) \cap B_{\tau_M(p)}(p)$. Since M has a positive local feature size by Proposition 3.1.2, we can apply Theorem 2.2.9 and deduce that $\mathrm{Nor}(p, M) = N_pM$, and $\mathrm{Tan}(p, M)$ is the convex cone dual to N_pM, which is T_pM. □

Proposition 3.1.4. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and R be a regular domain of M, i.e. a properly embedded codimension-0 submanifold with boundary. Then R has positive local feature size, and in particular*

$$\tau_R \geq \min\left(\tau_{\partial R}, \inf_{p \in \mathring{R}} \tau_M(p)\right).$$

Proof. We proceed by comparing $\mathrm{Med}(R)$ with $\mathrm{Med}(M)$ and consider points $x \in \mathrm{Med}(R) \setminus \mathrm{Med}(M)$. Since $x \notin \mathrm{Med}(M)$, we know x has a unique nearest neighbour $\pi_M(x) \in M$. In restricting from M to R, either $\pi_M(x)$ remains in R or is absent in R. In the former case, $\pi_M(x)$ must still be the unique nearest neighbour of x in R, as R is strictly a subset of M and thus x cannot have more nearest neighbours in R than in M. Therefore, the additional points $x \in \mathrm{Med}(R) \setminus \mathrm{Med}(M)$ added to the medial axis must belong to the latter case, where $\pi_M(x) \in M \setminus R$. If $\tau_R < \tau_M$, this must be due to the existence of points in $\mathrm{Med}(R) \setminus \mathrm{Med}(M)$ being closer to R than those in $\mathrm{Med}(M)$.

We partition points in $\mathrm{Med}(R) \setminus \mathrm{Med}(M)$ into two classes, where for $x \in \mathrm{Med}(R) \setminus \mathrm{Med}(M)$, either $\mathrm{NN}_R(x) \cap \mathring{R} = \emptyset$ or not. We first consider the latter case. For $x \in \mathrm{Med}(R)$, suppose $q \in \mathrm{NN}_R(x) \cap \mathring{R}$. Since q is a critical point of the distance function to M from x, the gradient of the distance function is annihilated by T_qM; thus, we infer that $x - q \in N_qM$. As M has positive local feature size, if $\|x - q\| < \tau_M(q)$, this would imply $x \in \mathrm{P}^{-1}{}_M(q)$, which implies $\pi_M(x) = q \in R$ and violates our assumption that $\pi_M(x) \in M \setminus R$. Thus, $\|x - q\| \geq \tau_M(q)$.

Let us consider the other case where $\mathrm{NN}_R(x) \cap \mathring{R} = \emptyset$; in other words, $\mathrm{NN}_R(x) = \mathrm{NN}_{\partial R}(x)$. Thus, $x \in \mathrm{Med}(\partial R)$. For any $q \in \mathrm{NN}_R(x)$, we must therefore have $\|x - q\| \geq \tau_{\partial R}(q)$.

Now suppose there is some point $p \in R$ such that for any $r > 0$, there is a point $x_r \in \mathrm{Med}(R)$ such that $\|x_r - p\| < r$. If we choose $0 < r < \tau_M(p)$, then $x_r \in \mathrm{Med}(R) \setminus \mathrm{Med}(R)$. Let q_r be a nearest neighbour of x_r in R. Then

$$\tau_M(p) > r \geq \|x_r - q_r\| \geq \tau_M(q_r) \quad \text{and} \quad \|p - q_r\| \leq 2r.$$

If $p \in \mathring{R}$, for sufficiently small r, we can force q_r to be also in $\mathring{R}$. If $p \in \partial R$, then we can choose $0 < r < \tau_{\partial R}(p)$ such that $x_r \notin \mathrm{Med}(\partial R)$. As such, $\mathrm{NN}_R(x_r) \cap \mathring{R} \neq \emptyset$ and there is always some $q_r \in \mathrm{NN}_R(x_r)$ such that $q_r \in \mathring{R}$. Since $\tau_M(p)$ is a continuous function (Lemma 2.2.3),

$$\tau_M(p) = \lim_{r \to 0} \tau_M(q_r).$$

However, as $r \geq \tau_M(q_r)$, the limit on the right hand side evalates to zero as we take $r \to 0$, implying $\tau_M(p) = 0$. This is a contradiction. Therefore, for some sufficiently small r, we must have that $B_r(p) \cap \mathrm{Med}(R) = \emptyset$, i.e. all points in R have positive local feature size. □

3.2 Curvature

The *second fundamental form* of an embedded submanifold M describes the extrinsic curvature of M. Before we describe the second fundamental form, let us first establish a few notational conventions. Let $\boldsymbol{\nabla} : \mathcal{T}(\mathbb{R}^d) \times \mathcal{T}(\mathbb{R}^d) \to \mathcal{T}(\mathbb{R}^d)$ be the Euclidean connection of the ambient Euclidean space, and $\nabla : \mathcal{T}(M) \times \mathcal{T}(M) \to \mathcal{T}(M)$ be the Levi-Civita connection of the metric on M induced by the embedding into Euclidean space.

The second fundamental form $\mathrm{II} : \mathcal{T}(M) \times \mathcal{T}(M) \to \mathcal{N}(M)$ is a symmetric bilinear form taking pairs of vector fields to sections of the normal bundle

$$\mathrm{II}(X, Y) = \boldsymbol{\nabla}_X Y - \nabla_X Y \tag{3.2}$$

where X and Y are arbitrarily extended from $\mathcal{T}(M)$ to $\mathcal{T}(\mathbb{R}^d)$ on the right hand side. Recall for any pair of vector fields X and Y in $\mathcal{T}(M)$, $\nabla_X Y$ is the tangential component of $\boldsymbol{\nabla}_X Y$ in

$\mathcal{T}(M) \subset \mathcal{T}(\mathbb{R}^d)$:

$$\nabla_X Y = (\boldsymbol{\nabla}_X Y)^\top. \tag{3.3}$$

Thus the second fundamental form is simply the normal component of $\boldsymbol{\nabla}_X Y$:

$$\mathrm{II}(X, Y) = (\boldsymbol{\nabla}_X Y)^\perp. \tag{3.4}$$

The second fundamental form encodes the extrinsic curvature of M due to the embedding into $\mathbb{R}^d$. For example, the *Weingarten equation* expresses the rate of change of the normal vector in the direction of Y evolved along X:

$$\langle \boldsymbol{\nabla}_X n, Y \rangle = - \langle n, \mathrm{II}(X, Y) \rangle . \tag{3.5}$$

We now recall some facts about smooth curves on M in relation to the second fundamental form. The extrinsic acceleration $\ddot{u}$ of a smooth curve $u : I \to M$ can be expressed as an orthogonal linear combination of the tangential acceleration $\nabla_{\dot{u}} \dot{u}$ along $T_{u(t)}M$ and a normal acceleration $\mathrm{II}(\dot{u}, \dot{u})$ along $N_{u(t)}M$:

$$\ddot{u} = \boldsymbol{\nabla}_{\dot{u}} \dot{u} = \nabla_{\dot{u}} \dot{u} + \mathrm{II}(\dot{u}, \dot{u}). \tag{3.6}$$

In particular, for a geodesic $\gamma(t)$ on M, the velocity vector $\dot{\gamma}$ is parallel transported along γ. Therefore a geodesic can be regarded as having constant speed:

$$\nabla_{\dot{\gamma}} \dot{\gamma} = 0 \iff \|\dot{\gamma}\|^2 = \langle \dot{\gamma}, \dot{\gamma} \rangle = \text{constant}.$$

Equation (3.6) implies that the extrinsic acceleration of geodesics are completely normal and given by the second fundamental form:

$$\ddot{\gamma} = \mathrm{II}(\dot{\gamma}, \dot{\gamma}).$$

Unless otherwise stated, we take geodesics to be parametrised by arc length i.e. $\|\dot{\gamma}\| = 1$. As distance minimising curves are geodesics, the distance between two points p and q on M is then the geodesic with the minimum arc length.

We make explicit the relationship between the extrinsic curvature and the second fundamental form by showing we can bound the angle variation of vectors parallel transported along geodesics in M in the ambient Euclidean space with the norm of the second fundamental form.

Lemma 3.2.1. *Consider a geodesic $\gamma : \mathbb{R} \to M$ where $\gamma(0) = p$ and $\|\dot{\gamma}\| = 1$. Suppose the norm of the second fundamental form is bounded above by $\frac{1}{\rho}$ along $\gamma([0, l])$:*

$$\|\mathrm{II}(X, X)\| \leq \frac{1}{\rho} \quad \forall X \in T_{\gamma(t)}M, \quad l \in [0, l]. \tag{3.7}$$

Let $X(t) \in T_{\gamma(t)}M$ be the parallel transport of a vector $X(0) \in T_{\gamma(0)}M$ along γ, where $\|X(0)\| = 1$. Then the extrinsic angle variation of the velocity vector $X : [0, \epsilon) \to \mathbb{S}^{d-1}$ along the geodesic is bounded above by

$$\measuredangle(X(l), X(0)) \leq \frac{l}{\rho}. \tag{3.8}$$

Proof. Since X is parallel transported, $\nabla_{\dot{\gamma}} X = 0$ implies $\|X\|$ is constant and equal to its intial value $\|X\| = \|X(0)\| = 1$. Therefore we can view $X(t)$ as a curve on $\mathbb{S}^{d-1}$. The angle (i.e. the geodesic distance on $\mathbb{S}^{d-1}$ between $X(0)$ and $X(l)$ is less than or equal to the arc length along $X(t)$ in $\mathbb{S}^{d-1}$. Therefore,

$$\measuredangle(X(l), X(0)) \leq \int_0^l \mathrm{d}t \left\| \frac{\mathrm{d}}{\mathrm{d}t} X(\gamma(t)) \right\| = \int_0^l \mathrm{d}t \, \|\boldsymbol{\nabla}_{\dot{\gamma}} X\| = \int_0^l \mathrm{d}t \, \|\mathrm{II}(X, \dot{\gamma})\| \leq \frac{l}{\rho}.$$

□

Niyogi et al. [2008] showed that the inverse of the norm of the second fundamental form at p, is bounded below by the local feature size of M at p.

Lemma 3.2.2 ([Niyogi et al., 2008, Proposition 6.1])**.** *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$. Then*

$$\|\mathrm{II}(X, X)\| \leq \frac{1}{\tau_M(p)} \quad \forall X \in T_p M.$$

If we consider the parallel transport of the velocity vector $\dot{\gamma}(t)$ of a geodesic $\gamma(t)$, and apply the inverse of the local feature size as a bound on the second fundamental form as shown

in Lemma 3.2.2, we obtain the following bounds on the variation of the angle of the velocity along a geodesic.

Corollary 3.2.3 ([Boissonnat et al., 2019, Lemma 5]). *Given the assumptions of Lemma 3.2.1, the extrinsic angle variation of the velocity vector $\dot{\gamma}(t) \in \mathbb{S}^{d-1}$ along the geodesic is bounded above by*

$$\measuredangle(\dot{\gamma}(l), \dot{\gamma}(0)) \leq \frac{l}{\rho}. \tag{3.9}$$

In particular, we can take $\rho = \min_{t\in[0,l]} \tau_M(\gamma(t))$.

In the case of a manifold with boundary M, a geodesic on M is a continuous, piecewise smooth path, where each piece is a geodesic on a manifold without boundary, that is either the interior or boundary of M.

Theorem 3.2.4 ([Alexander et al., 1987, P.169-170, Section 2]). *Let M be a smooth Riemannian manifold with a smooth boundary, and assume M is properly embedded into $\mathbb{R}^d$. Consider a geodesic $\gamma : \mathbb{R} \to M$. Then γ is C^1 and $\ddot{\gamma}$ is only undefined on a countable number of points. Where the acceleration $\ddot{\gamma}$ is defined, either $\ddot{\gamma} = \mathrm{II}_{\partial M}(\dot{\gamma}, \dot{\gamma})$, or $\ddot{\gamma} = \mathrm{II}_M(\dot{\gamma}, \dot{\gamma})$; the latter is the only possible case where $\gamma(t)$ is in the interior of M.*

We can thus generalises Corollary 3.2.3 to geodesics on submanifolds with boundary.

Proposition 3.2.5. *Consider a geodesic $\gamma : \mathbb{R} \to M$ where $\gamma(0) = p$ and $\|\dot{\gamma}\| = 1$ where M is a smooth m-dimensional Riemannian submanifold of $\mathbb{R}^d$ with or without boundary. Suppose*

$$\|\mathrm{II}_M(X, X)\| \leq \frac{1}{\rho} \quad \forall X \in T_pM, \quad p \in \gamma([0, \epsilon]); \quad \textit{and} \tag{3.10}$$

$$\|\mathrm{II}_{\partial M}(X, X)\| \leq \frac{1}{\rho} \quad \forall X \in T_p\partial M, \quad p \in \partial M \cap \gamma([0, \epsilon]) \quad \textit{if} \quad \partial M \neq \emptyset. \tag{3.11}$$

Then the extrinsic angle variation of the velocity vector $\dot{\gamma} : [0, \epsilon) \to \mathbb{S}^{d-1}$ along the geodesic is bounded above by

$$\measuredangle(\dot{\gamma}(l), \dot{\gamma}(0)) \leq \frac{l}{\rho}. \tag{3.12}$$

Proof. The case where $\partial M = \emptyset$ is accounted for in Corollary 3.2.3. We now consider the case where M is a manifold with boundary. Since γ is C^1 and $\ddot{\gamma}$ is only undefined on a countable number of points $t_1 < t_2 < \cdots \in [0, \epsilon]$ by Theorem 3.2.4, the path length of $\dot{\gamma}([0, l]) \subset \mathbb{S}^{d-1}$

is given by the integral $\sum_{i=0} \int_{t_i}^{t_{i+1}} \mathrm{d}t \, \|\ddot{\gamma}\|$ (where $t_0 = 0 \leq t_1$). We can thus apply the same argument in the proof of Corollary 3.2.3 and deduce

$$\measuredangle(\dot{\gamma}(l), \dot{\gamma}(0)) \leq \sum_{i=0} \int_{t_i}^{t_{i+1}} \mathrm{d}t \, \|\ddot{\gamma}\| = \sum_{i=0} \int_{t_i}^{t_{i+1}} \mathrm{d}t \, \|\mathrm{II}(\dot{\gamma}, \dot{\gamma})\| \leq \frac{l}{\rho}.$$

□

3.3 Metric Distortion

Having related the curvature and local feature sizes of M to the angular variation of the velocity of geodesics, we can demonstrate how curvature bounds local discrepancies between the geodesic distance on M and the Euclidean distance between two points on M. [Boissonnat et al., 2019, Lemma 12] showed that on a small neighbourhood of M of size τ_M, the Euclidean distance $\|p - q\|$ between two points p and q in M can be used to bound the distance $d_M(p, q)$ between them on the manifold. In the lemma below, we modify the proof of [Boissonnat et al., 2019, Lemma 12] so that the bound on the geodesic distance is parametrised by the local feature size rather than the reach.

Lemma 3.3.1. *Let M be a Riemannian submanifold with or without boundary, that is properly embedded in $\mathbb{R}^d$ with positive local feature size. Consider $p \in M$ and let $r \leq \max\left(\tau, \frac{\tau_M(p)}{\pi+1}\right)$. Thus, if we have a geodesic $\gamma(t)$ on M parametrised by arc length through $\gamma(0) = p$, for $t \in [0, \pi r]$ we have $\tau_M(\gamma(t)) \geq r$ and*

$$\|\gamma(t) - \gamma(0)\| \geq 2r \sin\left(\frac{t}{2r}\right). \tag{3.13}$$

Proof. Let us consider $t \in [0, \pi r]$. Since the Euclidean distance $\|\gamma(t) - p\|$ is bounded above by the arc length t along $\gamma([0, t])$, the geodesic $\gamma(t)$ is contained within the closed Euclidean ball or radius t: $\|\gamma(t) - p\| \leq t \leq \pi r$. If $r \leq \tau(p)/(\pi + 1)$, then $\|\gamma(t) - p\| \leq \pi\tau(p)/(\pi + 1)$. Since the local feature size is 1-Lipschitz (Lemma 2.2.3)

$$\tau(\gamma(t)) \geq \tau(p) - \|\gamma(t) - p\| \geq \frac{1}{\pi + 1}\tau(p) \geq r.$$

Therefore, we can establish that $\tau(\gamma(t)) \geq r$ for $r \leq \max(\tau, \tau(p)/(\pi + 1))$. For the segment of the geodesic $\gamma([s', s])$ where $s, s' \in [0, \pi r]$, the conditions of proposition 3.2.5 are satisfied ,

and thus the variation of the direction of $\dot{\gamma}$ is bounded above:

$$\measuredangle\big(\dot{\gamma}(s), \dot{\gamma}(s')\big) \leq \frac{|s - s'|}{r} \leq \frac{\pi r}{r} \leq \pi.$$

Therefore, in this regime of angle variation, cosine is a monotonically non-increasing function:

$$\big\langle \dot{\gamma}(s), \dot{\gamma}(s') \big\rangle = \cos\big(\measuredangle\big(\dot{\gamma}(s), \dot{\gamma}(s')\big)\big) \geq \cos\left(\frac{s - s'}{r}\right).$$

The rest of the proof is a modification of [Boissonnat et al., 2019, Lemma 12] where we replace τ with r. For $t \in [0, \pi r]$, we can apply proposition 3.2.5 again and deduce:

$$\begin{aligned}
\|\gamma(t) - \gamma(0)\|^2 &= \langle \gamma(t) - \gamma(0), \gamma(t) - \gamma(0) \rangle \\
&= \left\langle \int_0^t \mathrm{d}s\, \dot{\gamma}(s), \int_0^t \mathrm{d}s'\, \dot{\gamma}(s') \right\rangle \\
&= \int_0^t \mathrm{d}s \int_0^t \mathrm{d}s' \left\langle \dot{\gamma}(s), \dot{\gamma}(s') \right\rangle \\
&\geq \int_0^t \mathrm{d}s \int_0^t \mathrm{d}s' \cos\left(\frac{s - s'}{r}\right) \\
&= \left(2r \sin\left(\frac{t}{2r}\right)\right)^2.
\end{aligned}$$

□

The following corollary of Lemma 3.3.1 can be regarded as a special case of [Boissonnat et al., 2019, Theorem 1]. For two points on M sufficiently close relative to their local feature sizes, the geodesic distance between p and q can be bounded above using their Euclidean distance.

Corollary 3.3.2 (Also see [Boissonnat et al., 2019, Theorem 1])**.** *Let M be a Riemannian submanifold with or without boundary properly embedded in $\mathbb{R}^d$ with positive local feature size. Consider $p \in M$ and let $r \leq \max(\tau, \frac{\tau(p)}{\pi+1})$. Then for $q \in M \cap B_\epsilon[p]$ where $\epsilon \in (0, 2r)$*

$$2r \sin\left(\frac{d_M(p,q)}{2r}\right) \leq \|q - p\| \leq d_M(p,q) \leq \pi r.$$

Proof. Let $B^M_{\pi r}[p]$ be the set of points $x \in M$ with $d_M(x,p) \leq \pi r$. We first want to show that $M \cap B_\epsilon[p] = B^M_{\pi r}[p] \cap B_\epsilon[p]$.

Suppose the contrary: i.e. there is some point $q \in M \cap B_\epsilon[p]$ where $d_M(p,q) > \pi r$.

Since $\epsilon < \tau(p)$, the set $M \cap B_\epsilon[p]$ is contractible (and thus connected) by Lemma 2.2.14. Therefore, there exists some path $y(t)$ from $y(0) = p$ to $y(l) = q$ such that for all $t \in [0, l]$, $y(t) \in M \cap B_\epsilon[p]$.

By continuity, we have $d_M(p, y(t_*)) = \pi r$ at some point $t_* \in (0, l)$ along the path. Therefore, we can apply Lemma 3.3.1; since $y(t_*) \in B_\epsilon[p]$ and $\epsilon < 2r$,

$$2r \sin\left(\frac{d_M(p, y(t_*))}{2r}\right) \leq \|p - y(t_*)\| \leq \epsilon < 2r \implies 2r < 2r.$$

Thus, by contradiction, we have $M \cap B_\epsilon[p] = B^M_{\pi r}[p] \cap B_\epsilon[p]$. Since $M \cap B_\epsilon[p] = B^M_{\pi r}[p] \cap B_\epsilon[p]$, for any point $q \in M \cap B_\epsilon[p]$, the conditions of Lemma 3.3.1 apply, and we obtain the desired statement. □

Using a lower bound of $2r \sin\left(\frac{d_M(p,q)}{2r}\right)$ that is linear in $d_M(p, q)$, we can obtain a weaker but simpler statement of Corollary 3.3.2, which is also proved in [Boissonnat et al., 2019, Lemma 14] with different geometric arguments. We note that our proof of Corollary 3.3.2 and Corollary 3.3.3 is different to the arguments used in Boissonnat et al. [2019].

Corollary 3.3.3 (Also see [Boissonnat et al., 2019, Lemma 14])**.** *Let M be a Riemannian submanifold with or without boundary properly embedded in $\mathbb{R}^d$ with positive local feature size. Consider $p \in M$ and let $r \leq \max(\tau, \frac{\tau(p)}{\pi+1})$. If $\epsilon < \pi r$, then*

$$M \cap B_{\frac{2}{\pi}\epsilon}(p) \subset B^M_\epsilon(p) \subset M \cap B_\epsilon(p).$$

where $B^M_\epsilon(p)$ denotes the open geodesic ball on the manifold of radius ϵ about p.

Proof. First, since two points a geodesic distance ϵ apart can be at most ϵ Euclidean distance apart, $B^M_\epsilon(p) \subset M \cap B_\epsilon(p)$.

To show the other inclusion, consider $q \in M \cap B_{\frac{2}{\pi}\epsilon}(p)$. Since $\frac{2}{\pi}\epsilon < 2r$, we can apply Corollary 3.3.2 and deduce

$$\frac{2}{\pi}\epsilon > \|p - q\| \geq 2r \sin\left(\frac{d_M(p, q)}{2r}\right) \geq \frac{2}{\pi} d_M(p, q) \implies d_M(p, q) < \epsilon.$$

Thus, $M \cap B_{\frac{2}{\pi}\epsilon}(p) \subset B^M_\epsilon(p)$. □

3.4 Orthogonal Projections onto Tangent Spaces

Another geometric consequence of curvature is the deviation of the manifold away from the tangent space T_pM at some point $p \in M$. We can quantify this effect by considering the largest radius r such that the orthogonal projection of $B_r(p) \cap M$ onto T_pM is an embedding of $B_r(p) \cap M$ into T_pM. We obtain a lower bound of this radius in terms of the local feature size of M at p. We first improve on the bounds of [Niyogi et al., 2008, Proposition 6.2] which relates the radius where the the orthogonal projection of $B_r(p) \cap M$ onto T_pM is a *local* diffeomorphism to the reach.

Proposition 3.4.1. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and consider $p \in M$. Suppose $\tau(p) > 0$. Then for $r < \max(\tau(p)/(1+2/\pi), \pi\tau/2)$, the orthogonal projection $P : M \to T_pM$ onto the tangent plane $p + T_pM \subset \mathbb{R}^d$ restricted to the open geodesic ball $B_r^M(p)$ is a smooth immersion, i.e. $\mathrm{d}P$ has full rank on $B_r^M(p)$, and therefore a local diffeomorphism.*

Proof. The minimum of $\tau(q)$ for $q \in B_r^M(p)$ is lower bounded by the reach τ, or we use the fact that the local feature size is 1-Lipschitz with respect to the Euclidean distance (Lemma 2.2.3) and deduce

$$\tau(q) \geq \tau(p) - r > \frac{2}{\pi} r.$$

In either case,

$$\frac{r}{\min_q \tau(q)} < \frac{\pi}{2},$$

where the minimum is taken over $q \in B_r^M(p)$.

We now suppose $\mathrm{d}P$ is singular for some $q \in B_r^M(p)$. This implies there is at least some unit vector $u \in T_qM$ such that $u \in N_pM$. Since $q \in B_r^M(p)$, there is a geodesic $\gamma : [0, s] \to M$ where $\gamma(0) = q$ and $\gamma(s) = p$, such that $d_M(p, q) = s < r$ and $\gamma([0, s]) \subset B_r^M(p)$. Let us parallel transport u along γ to T_pM:

$$u \in T_qM \overset{\gamma}{\rightsquigarrow} u' \in T_pM.$$

Since $u \in N_pM$,

$$\measuredangle(u, u') = \pi/2.$$

However, this contradicts the upper bound on $\measuredangle(u, u')$ due to Lemma 3.2.1:

$$\frac{\pi}{2} > \frac{r}{\min_q \tau(q)} \geq \measuredangle(u, u') = \frac{\pi}{2}.$$

Thus, we reach a contradiction, and deduce that there are no points in $B_r^M(p)$ where $\mathrm{d}P$ is singular. In other words, P restricted to $B_r^M(p)$ is a smooth immersion. □

We now consider the case where we require the orthogonal projection to not only be locally injective, but globally injective when restricted to a neighbourhood of p in M. We first consider the range of radii where the projection of geodesic balls $B_r^M(p)$ onto the tangent space at p is injective.

Proposition 3.4.2. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and consider $p \in M$. Suppose $\tau(p) > 0$. Then for $r < \max(\tau(p)/2, \tau)$, the orthogonal projection $P : M \to T_pM$ onto the tangent plane $p + T_pM \subset \mathbb{R}^d$ restricted to $B_r^M(p)$ is injective.*

Proof. We first show that $r < \tau(q)$ for all $q \in B_r^M(p)$. By definition of the local feature size, $\tau(q) \geq \tau$. If $\tau(p)/2 \leq \tau$, then $r < \tau \leq \tau(q)$. If on the contrary $\tau(p)/2 > \tau$, we use the fact that the local feature size is 1-Lipschitz with respect to the Euclidean distance (Lemma 2.2.3) and deduce

$$\tau(q) \geq \tau(p) - \|p - q\| \geq \tau(p) - r > 2r - r = r.$$

We now proceed to prove our result by contradiction. Consider a pair of distinct points q and q' in $B_r^M(p)$, where we assume $P(q) = P(q')$. Since q and q' project onto the same point in the plane $p + T_pM$, the vector $q - q'$ is normal to T_pM, and parallel to $q - P(q)$. Since

$$P(q) = (q - p) - \langle q - p, n \rangle n$$

for some unit vector $n \propto q - q' \in N_pM \subseteq \mathrm{Nor}(p, M)$, we can apply Theorem 2.2.18 and deduce

$$\|(q - p) - P(q)\| = |\langle q - p, n \rangle| \leq \frac{\|q - p\|^2}{2\tau(p)} < \frac{r^2}{2\tau(p)}.$$

Therefore,

$$\|q - q'\| = \|(q - p) - P(q)\| \pm \|(q' - p) - P(q')\| < \frac{r^2}{\tau(p)}.$$

Let us now consider $u \in T_qM$, the unit vector of the projection of $q' - q$ onto T_qM. Once again applying Theorem 2.2.18, we can deduce

$$\sin\big(\measuredangle(u, q' - q)\big) \leq \frac{\|q - q'\|}{2\tau(q)} < \frac{r^2}{2\tau(p)\tau(q)} < \frac{1}{2} \implies \measuredangle(u, q' - q)) < \pi/6,$$

where we used the fact that $r < \tau(q)$ for all $q \in B_r^M(p)$ in the penultimate inequality.

As $q \in B_r^M(p)$, there is a geodesic $\gamma : [0, s] \to M$ where $\gamma(0) = q$ and $\gamma(s) = p$, such that $d_M(p, q) = s < r$ and $\gamma([0, s]) \subset B_r^M(p)$. Let us parallel transport u along γ to T_pM:

$$u \in T_qM \overset{\gamma}{\rightsquigarrow} u' \in T_pM.$$

We deduce from Lemma 3.2.1 that u' has undergone a rotation in $\mathbb{S}^{d-1}$ at most

$$\measuredangle(u, u') \leq r / \min \tau(\gamma(t)) < 1.$$

Since $u' \in T_pM$ and $q - q' \in N_pM$, we have

$$\measuredangle(u', q - q') = \pi/2.$$

Applying the triangle inequality on $\mathbb{S}^{d-1}$:

$$\measuredangle(u', u) + \measuredangle(u, q - q') \geq \measuredangle(u', q - q') \implies \frac{\pi}{6} + 1 > \frac{\pi}{2}$$

which is a contradiction. □

Combining Proposition 3.4.2 and Proposition 3.4.1 which prescribe the geodesic radii on which the projection is injective and a smooth immersion respectively, we thus derive the range of radii such that the projection of a geodesc ball is a diffeomorphism onto its image in the tangent space.

Proposition 3.4.3. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and consider $p \in M$. Suppose $\tau_M(p) > 0$. Then for $r < \max(\tau_M(p)/2, \tau)$, the orthogonal projection $P : M \to T_pM$ onto the tangent plane $p + T_pM \subset \mathbb{R}^d$ restricted to $B_r^M(p)$ is a diffeomorphism onto its image, i.e. a smooth embedding of $B_r^M(p)$ into T_pM.*

Proof. Let P restricted to $B_r^M(p)$ be $\tilde{P}$ for $r < \max(\tau(p)/2, \tau)$. Since Proposition 3.4.2 imply that $\tilde{P}$ is injective, we can apply invariance of domain (Theorem C.1) and deduce that $\tilde{P}$ is a homeomorphism onto its image. As Proposition 3.4.1 imply that $\tilde{P}$ is also a smooth immersion, we conclude that $\tilde{P}$ is a diffeomorphism onto its image. □

Having derived a range of radii for geodesic balls based at p to project injectively onto the tangent plane at p, we now use the bounds between the geodesic and Euclidean distances (Corollary 3.3.3) to derive a similar result for Euclidean balls intersecting the manifold.

Corollary 3.4.4. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and consider $p \in M$. Suppose $\tau(p) > 0$. Then for $\epsilon < \max(\tau(p)/\pi, 2\tau/\pi)$, the projection $P : M \to T_pM$ onto the tangent plane $p + T_pM \subset \mathbb{R}^d$ restricted to $B_\epsilon(p) \cap M$ is a diffeomorphism onto its image.*

Proof. Since the squared Euclidean distance away from p defines a smooth function on M, the set $B_\epsilon(p) \cap M$ is a submanifold of M. Since $\frac{\pi}{2}\epsilon < \max\left(\frac{\pi}{\pi+1}\tau(p), \pi\tau\right)$, we observe that

$$B_\epsilon(p) \cap M \subset B^M_{\frac{\pi}{2}\epsilon}(p).$$

due to Corollary 3.3.3. Because $\frac{\pi}{2}\epsilon < \max\left(\frac{\tau_M(p)}{2}, \tau\right)$, the projection map P restricted to $B^M_{\frac{\pi}{2}\epsilon}(p)$ is a diffeomorphism onto its image. As we have a smooth inclusion of $B_\epsilon(p) \cap M$ into $B^M_{\frac{\pi}{2}\epsilon}(p)$, the restriction of P to $B_\epsilon(p) \cap M$ is also a diffeomorphism onto its image. □

Corollary 3.4.5. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and R be a regular domain of M, i.e. a properly embedded codimension-0 submanifold with boundary. Consider $p \in N$ and $p + T_pM = p + T_pR$. Then for $\epsilon < \max\left(\frac{\tau_M(p)}{\pi}, \frac{2}{\pi}\tau_M\right)$, the projection $P : M \to T_pM$ restricted to $\mathring{R} \cap B_\epsilon(p)$ is a homeomorphism onto its image.*

Proof. We claim P is a continuous injection on $M \cap B_\epsilon(p)$ (which we show in Proposition 3.4.2). Since the inclusion $\iota : \mathring{R} \cap B_\epsilon(p) \hookrightarrow M \cap B_\epsilon(p)$ is a continuous injection, $P \circ \iota$ is a continuous injection on $\mathring{R} \cap B_\epsilon(p)$. Moreover, since R is a regular domain of M, the subset of interest $\mathring{R} \cap B_\epsilon(p)$ is open in M. Therefore, we can apply invariance of domain (Theorem C.1) and deduce that $P \circ \iota$ is a homeomorphism onto its image. □

3.5 Inferring the Homology of Regular Domains

We now consider the problem of inferring the homology of a regular domain $R \subset M$ from a finite point sample $\mathbb{X}$. Recall that a regular domain of M is a properly embedded codimension-0 submanifold with boundary of M. We have shown in Proposition 3.1.4 that R has positive local feature size. As such, if R is compact, then R has positive reach and we can recover its homology from a finite point sample $\mathbb{X}$ using techniques described in Section 2.3, provided $d_H(R, \mathbb{X})$ is sufficiently small relative to the reach of R.

For each regular domain, there exists a *defining function* $f \in C^\infty(M)$, such that zero is a regular value of f and $R = f^{-1}(-\infty, 0]$. We derive an explicit lower bound for the reach of regular domains and their boundaries in terms of the bounds on the norms of first and second derivatives of the defining function f. We then apply theorems from Niyogi et al. [2008] and Wang and Wang [2020] to show that a uniform i.i.d. random sample $\mathbb{X}$ on R can be made arbitrarily dense in R with high confidence if $\#\mathbb{X}$ is sufficiently large, where we have an explicit expression for $\#\mathbb{X}$ in terms of the reach of M and R. As such, we can infer the homology of R from a sufficiently large number of uniform i.i.d. random samples.

By a uniform i.i.d. random sample, we refer to points drawn from a uniform measure on the manifold. We recall a measurable set E of a smooth m-dimensional compact manifold M is one where for any chart $\varphi : U \to \mathbb{R}^m$ of M, $\varphi(E \cap U)$ is a Lebesgue measurable set in $\mathbb{R}^m$. The uniform measure is defined as the following integral.

Definition 3.5.1. Let M be a smooth m-dimensional Riemannian manifold with or without boundary with metric g. The volume of a measurable set E is a measure given by

$$\mathrm{vol}(E) = \sum_i \int_{\varphi(U_i)} (\psi_i \circ \varphi_i^{-1})(x) \sqrt{\det g(\varphi_i{}^{-1}(x))}\, \mathrm{d}x_1^i \cdots \mathrm{d}x_m^i$$

where $\{\psi_i\}$ is any smooth partition of unity subordinate to an open cover of M by smooth charts $\varphi_i : U_i \to \mathbb{R}^m$, and $\left\{x_j^i\right\}$ are a set of coordinates on U_i. The uniform probability measure of E is then the volume of E normalised by the volume of the manifold.

This definition is independent of the choice of $\{\psi_i\}$ ([Lee, 2012, Proposition 16.5]).

3.5.1 Lower Bounds of the Reach

Let us consider a sublevel set $f^{-1}(-\infty, 0]$ of a smooth function $f : M \to \mathbb{R}$ where 0 is a regular value of f. As $f^{-1}(-\infty, 0]$ has positive local feature size (Proposition 3.1.4), we can obtain a lower bound of the reach $\tau_{f^{-1}(-\infty,0]}$ using the result of Theorem 2.2.16, which states that the supremum of

$$\kappa_{f^{-1}(-\infty,0]}(p) = \sup_{\substack{v \in \mathbb{R}^d \setminus 0 \\ p+v \in f^{-1}(-\infty,0]}} \frac{2\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty,0]\right)\right\rangle}{\|v\|^2}$$

over all $p \in f^{-1}(-\infty, 0]$ is $1/\tau_{f^{-1}(-\infty,0]}$. Thus, we have a lower bound of $\tau_{f^{-1}(-\infty,0]}$ if we can obtain an upper bound of $\kappa_{f^{-1}(-\infty,0]}(p)$.

Lemma 3.5.2. *Consider $p \in f^{-1}(-\infty, 0]$, where $f : M \to \mathbb{R}$ is a smooth function on a m-dimensional Riemannian submanifold M properly embedded in $\mathbb{R}^d$ with positive reach, and $p \notin \mathrm{Crit}(f)$. If 0 is a regular value of f, then for $\kappa : M \to [0, \infty]$ given in Equation (2.10) in Theorem 2.2.16.*

$$\kappa_{f^{-1}(-\infty,0]}(p) \leq \kappa_M(p) \qquad \text{if} \quad f(p) < 0 \tag{3.14}$$

$$\kappa_{f^{-1}(-\infty,0]}(p) \leq \max\left(\kappa_{f^{-1}(0)}(p), \min\left(\frac{3}{\tau_M(p)}, \frac{1}{\tau_M}\right)\right) \qquad \text{if} \quad f(p) = 0. \tag{3.15}$$

Proof. From Proposition 3.6.6, we see that for p in the interior i.e. $f(p) < 0$, the tangent and normal cones for $f^{-1}(-\infty, 0]$ at p are the same as those for M at p. Since $f^{-1}(-\infty, 0] \subset M$, we can deduce from the definition of $\kappa_{f^{-1}(-\infty,0]}(p)$ in Corollary 2.2.19 that

$$\begin{aligned}
\kappa_{f^{-1}(-\infty,0]}(p) &= \sup_{\substack{v \in \mathbb{R}^d \setminus 0 \\ p+v \in f^{-1}(-\infty,0]}} 2\frac{\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty,0]\right)\right\rangle}{\|v\|^2} \\
&= \sup_{\substack{v \in \mathbb{R}^d \setminus 0 \\ p+v \in f^{-1}(-\infty,0]}} 2\frac{\langle v, \mathrm{Nor}(p, M)\rangle}{\|v\|^2} \\
&\leq \sup_{\substack{v \in \mathbb{R}^d \setminus 0 \\ p+v \in M}} 2\frac{\langle v, \mathrm{Nor}(p, M)\rangle}{\|v\|^2} = \kappa_M(p).
\end{aligned}$$

For $p \in f^{-1}(0)$, Proposition 3.6.5 implies

$$\mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right) = \left\{t\nabla f + n \ : \ n \in N_p M, \quad t \geq 0\right\}.$$

Thus for any $v \in \mathbb{R}^d$, we have

$$\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right)\right\rangle = \max_{t\geq 0,\ n\in N_pM,\ \|n\|^2\leq 1-t^2} t\left\langle v, n_f\right\rangle + \langle v, n\rangle \tag{3.16}$$

where n_F is the unit vector along ∇f. If $\langle \nabla f, v\rangle \leq 0$, then the right had side is maximised if we choose $t = 0$, i.e.

$$\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right)\right\rangle = \max_{n\in N_pM,\ \|n\|=1} \langle v, n\rangle = \langle v, \mathrm{Nor}(p, M)\rangle = \left\langle v, N_pM\right\rangle$$

Since $p + v \in f^{-1}(-\infty, 0] \subset M$, we can apply Theorem 2.2.16 and obtain an upper bound

$$\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right)\right\rangle = \left\langle v, N_pM\right\rangle \leq \frac{\kappa_M(p)}{2}\|v\|^2.$$

If $\langle \nabla f, v\rangle > 0$, the fact that $\mathrm{Nor}\big(p, f^{-1}(-\infty, 0]\big) \subset \mathrm{Nor}\big(p, f^{-1}(0)\big)$ implies

$$\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right)\right\rangle \leq \left\langle v, \mathrm{Nor}\left(p, f^{-1}(0)\right)\right\rangle.$$

If $p + v \in f^{-1}(0)$, then we can apply Theorem 2.2.16 and obtain an upper bound

$$\left\langle v, \mathrm{Nor}\left(p, f^{-1}(-\infty, 0]\right)\right\rangle \leq \left\langle v, N_pf^{-1}(0)\right\rangle \leq \frac{\kappa_{f^{-1}(0)}(p)}{2}\|v\|^2.$$

The only outstanding case left to consider are points $p + v \in f^{-1}(-\infty, 0)$, such that $\langle \nabla f, v\rangle > 0$. Let us consider $B_r(p + rn) \cap M$, where n is a unit vector in $\mathrm{Nor}\big(p, f^{-1}(0)\big)$, and r is a constant such that for $r_0 = \max\left(\frac{\tau_M(p)}{3}, \tau_M\right)$

$$r \leq \min\left(\kappa_{f^{-1}(0)}(p), r_0\right),$$

Since $r \leq r_0$, we can apply Lemma 2.2.15 and see that $B_r(p + rn) \cap M$ has one connected component. Furthermore, for any unit vector $n \in \mathrm{Nor}\big(p, f^{-1}(0)\big)$,

$$B_r(p + rn) \cap f^{-1}(0) = \emptyset.$$

This follows from Lemma 2.2.20 where we take $f^{-1}(0)$ to be an embedded submanifold of $\mathbb{R}^d$.

Assuming $B_r(p + rn) \cap M$ is non-empty, we now show that the sign of f is constant on

$B_r(p+rn) \cap M$. For suppose we have two points q and q' in $B_r(p+rn) \cap M$, such that $f(q) > 0$ and $f(q') < 0$. Since $B_r(p+rn) \cap M$ is connected, let us consider a path connecting q and q' contained in $B_r(p+rn) \cap M$. From the mean value theorem, there must be some point along the path where $F = 0$. However, this contradicts the fact that $B_r(p+rn) \cap f^{-1}(0) = \emptyset$. Thus, the sign of f is constant and non-zero on $B_r(p+rn) \cap M$.

We now return to deriving an upper bound for $\langle v, \mathrm{Nor}(p, f^{-1}(-\infty, 0]) \rangle \leq \langle v, N_p f^{-1}(0) \rangle$ where $p + v \in f^{-1}(-\infty, 0)$ and $\langle \nabla f, v \rangle > 0$. Let u be the unit vector in $N_p f^{-1}(0)$ parallel to the projection of v onto $N_p f^{-1}(0)$; then

$$\langle v, \nabla f \rangle = \langle v - \langle v, u \rangle u, \nabla f \rangle + \langle v, u \rangle \langle u, \nabla f \rangle = \langle v, u \rangle \langle u, \nabla f \rangle .$$

where $\langle v - \langle v, u \rangle u, \nabla f \rangle = 0$ because $v - \langle v, u \rangle$ is orthogonal to $N_p f^{-1}(0)$ and $\nabla f \in N_p f^{-1}(0)$. Since u points along the orthogonal projection of v onto $N_p f^{-1}(0)$, $\langle v, u \rangle > 0$. Since we have assumed $\langle v, \nabla f \rangle > 0$, we have $\langle u, \nabla f \rangle > 0$.

Let us consider $B_r(p+ru) \cap M$. Since $\langle u, \nabla f \rangle > 0$, by continuity, there is some point sufficiently close to p in $B_r(p+ru) \cap M$ on which $f > 0$. Since the sign of f is constant and non-zero on the ball, we deduce f is non-negative on the entirety of $B_r(p+ru) \cap M$. Thus, $p + v$ lies outside $B_r(p+ru) \cap M$.

If u and v are parallel, i.e. $v \in N_p f^{-1}(0)$, then the fact that $p + v$ lies outside $B_r(p+ru) \cap M$ implies $\|v\| \geq 2r$, and thus

$$\frac{\langle v, \mathrm{Nor}(p, f^{-1}(-\infty, 0]) \rangle}{\|v\|^2} = \frac{\langle v, N_p f^{-1}(0) \rangle}{\|v\|^2} = \frac{\langle u, v \rangle}{\|v\|^2} = \frac{1}{\|v\|} \leq \frac{1}{2r}.$$

For the case where u and v are not parallel, let us consider two dimensional plane spanned by u and v, with p at the origin. If we take u to point along the positive horizontal axis, the intersection of $B_r(p+ru) \cap M$ with the plane is described by the parametric equation in radial coordinates (ρ, θ)

$$\rho < 2r \cos(\theta),$$

where ρ is the distance from the origin, and θ is the angle away from u. Since $p + v$ lies outside

$B_r(p+ru) \cap M$, we have v must lie in the region $\rho \geq 2r\cos(\theta)$, and thus

$$\|v\|^2 \geq 2r\langle u, v\rangle \implies \frac{\langle v, N_p f^{-1}(0)\rangle}{\|v\|^2} = \frac{\langle u, v\rangle}{\|v\|^2} \leq \frac{1}{2r}.$$

Having considered all possibilities of $p+v \in f^{-1}(-\infty, 0]$ for $p \in f^{-1}(0)$, we conclude that

$$\kappa_{f^{-1}(-\infty,0]}(p) \leq \max\left(\kappa_M(p), \kappa_{f^{-1}(0)}(p), \frac{1}{r}\right) \leq \frac{1}{r}.$$

□

Corollary 3.5.3. *Let M be an m-dimensional Riemannian submanifold M properly embedded in $\mathbb{R}^d$ with positive reach. Consider $f \in C^\infty(M)$ and suppose zero is a regular value of f. If for all $p \in f^{-1}(0)$, we have $\kappa_{f^{-1}(0)}(p) \leq 1/r_0$, then*

$$\tau_{f^{-1}(0)} \geq r_0 \quad \textit{and}$$
$$\tau_{f^{-1}(-\infty,0]} \geq \min\left(r_0, \max\left(\tau_M, \frac{1}{3}\inf_{q \in f^{-1}(0)} \tau_M(q)\right)\right).$$

Proof. This follows directly from the application of the following two facts to the local result in Lemma 3.5.2: that firstly the condition number $1/\tau_A$ of any subset $A \subset \mathbb{R}^d$ is equal to the supremum of $\kappa_A(p)$ over $p \in A$ (Theorem 2.2.16), and secondly, that the local feature size $\tau_A(p)$ is a lower bound of $1/\kappa_A(p)$. (Corollary 2.2.19) □

Lemma 3.5.4. *Consider $p \in f^{-1}(0)$, where $f : M \to \mathbb{R}$ is a smooth function on a m-dimensional Riemannian submanifold M properly embedded in $\mathbb{R}^d$. Let $r \leq \max\left(\tau_M, \frac{\tau_M(p)}{\pi+1}\right)$. Suppose on $B_{\frac{\pi}{3}r}(p) \cap M$, the norm of the Hessian is bounded by above $\left\|H_f\right\| \leq \Lambda$ and $\|\nabla f(p)\| \geq \rho\Lambda > 0$. Then*

$$\kappa_{f^{-1}(0)}(p) \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\rho} + \frac{2}{r}$$
$$\kappa_{f^{-1}(-\infty,0]}(p) \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\rho} + \frac{2}{r}.$$

where $\kappa : M \to [0, \infty]$ is given in Equation (2.10) in Theorem 2.2.16.

Proof. Let us recall the definition of κ from Equation (2.10),

$$\frac{\kappa_{f^{-1}(0)}(p)}{2} = \sup_{\substack{v\in\mathbb{R}^d\setminus 0 \\ p+v\in f^{-1}(0)}} \frac{\langle v, N_p f^{-1}(0)\rangle}{\|v\|^2}.$$

We separate points $p+v \in f^{-1}(0)$ into those where $\|v\| \geq r$, and those where $\|v\| < r$. In the former category, we can trivially obtain a bound

$$\frac{\langle v, N_p f^{-1}(0)\rangle}{\|v\|^2} \leq \frac{\|v\|}{\|v\|^2} \leq \frac{1}{r}$$

which is implied by the bound in the lemma. We now turn our attention to points where $\|v\| < r$. We first recall Corollary 3.3.2: for $\|v\| < 2r$,

$$2r \sin\left(\frac{d_M(p, p+v)}{2r}\right) \leq \|v\|.$$

If $\|v\| < r$, this implies

$$s = d_M(p, p+v) < \frac{\pi}{3} r.$$

Thus, for any point along the shortest path geodesic $\gamma(t)$ between $\gamma(0) = p$ and $\gamma(s) = p+v$,

$$\|\gamma(t) - \gamma(0)\| \leq t \leq s < \frac{\pi}{3} r$$

and $\gamma(t)$ therefore contained entirely in the ball $B_{\frac{\pi}{3}r}(p) \cap M$, and the bounds on the Hessian of f apply on the geodesic. Furthermore, as $s < \pi r$, by Lemma 3.3.1, $\tau_M(\gamma(t)) > r$ along the geodesic.

Let us consider a unit vector $u \in N_p f^{-1}(0)$. Then let

$$C(t) = \langle u, \gamma(t) - \gamma(0)\rangle.$$

Since $C : \mathbb{R} \to \mathbb{R}$ is a smooth function, we can apply the mean value theorem, which implies for some $t_* \in (0, s)$,

$$\langle u, v\rangle = \langle u, \gamma(s) - \gamma(0)\rangle = C(s) = \cancelto{0}{C(0)} + sC'(t_*) = s\langle u, \dot\gamma(t_*)\rangle = s\cos(\measuredangle(u, \dot\gamma(t_*))).$$

Meanwhile, the triangle inequality on $\mathbb{S}^{d-1}$ implies

$$\measuredangle(u, \dot\gamma(t_*)) \geq |\measuredangle(u, \dot\gamma(0)) - \measuredangle(\dot\gamma(0), \dot\gamma(t_*))|.$$

Substituting this inequality into the expression for $\langle u, v\rangle$, we have an upper bound

$$\begin{aligned}
\langle u, v\rangle &= s\cos(\measuredangle(u, \dot\gamma(t_*))) \\
&\leq s\cos(\measuredangle(u, \dot\gamma(0)) - \measuredangle(\dot\gamma(0), \dot\gamma(t_*))) \\
&\leq s\cos(\measuredangle(u, \dot\gamma(0))) + s\sin(\measuredangle(\dot\gamma(0), \dot\gamma(t_*))) \\
&\leq s\,\langle u, \dot\gamma(0)\rangle + s\sin\left(\frac{t_*}{r}\right) && \text{Corollary 3.2.3} \\
&< s\,\langle u, \dot\gamma(0)\rangle + s\sin\left(\frac{s}{r}\right) && \because t_* < s < \frac{\pi}{2}r.
\end{aligned}$$

We now recall that $u \in N_p f^{-1}(0)$ and write $u = u_\top + u_\perp$, where $u_\top$ is the component of u along $\nabla f \in T_pM$ and $u_\perp$ is the component of u in N_pM and orthogonal to $\dot\gamma(0)$. Therefore,

$$\begin{aligned}
\langle u, v\rangle &< s\,\langle u, \dot\gamma(0)\rangle + s\sin\left(\frac{s}{r}\right) \\
&= s\cancelto{0}{\langle u_\perp, \dot\gamma(0)\rangle} + s\,\langle u_\top, \dot\gamma(0)\rangle + s\sin\left(\frac{s}{r}\right) \\
&< \frac{s}{\|\nabla f(p)\|}|\langle \nabla f(p), \dot\gamma(0)\rangle| + s\sin\left(\frac{s}{r}\right) \\
&< \frac{s}{\rho\Lambda}|\langle \nabla f(p), \dot\gamma(0)\rangle| + s\sin\left(\frac{s}{r}\right) && \because \|\nabla f(p)\| \geq \rho\Lambda
\end{aligned}$$

We now proceed to produce an upper bound for $|\langle \nabla f(p), \dot\gamma(0)\rangle|$ by considering the variation of the smooth function on t that is $\langle \nabla f(\gamma(t)), \dot\gamma(t)\rangle$. For $t \in (0, \frac{\pi}{3}r)$,

$$\begin{aligned}
\langle \nabla f(\gamma(t)), \dot\gamma(t)\rangle - \langle \nabla f(\gamma(0)), \dot\gamma(0)\rangle &= \int_0^t \frac{\mathrm{d}}{\mathrm{d}z}\langle \nabla f(\gamma(z)), \dot\gamma(t')\rangle\,\mathrm{d}z \\
&= \int_0^t \langle \boldsymbol{\nabla}_{\dot\gamma}\nabla f(\gamma(z)), \dot\gamma(z)\rangle + \left\langle \nabla f(\gamma(z)), \cancelto{0}{\ddot\gamma(z)} \right\rangle \mathrm{d}z \\
&= \int_0^t H_f(\dot\gamma, \dot\gamma)\,\mathrm{d}z \\
\implies |\langle \nabla f(\gamma(t)), \dot\gamma(t)\rangle - \langle \nabla f(\gamma(0)), \dot\gamma(0)\rangle| &\leq \int_0^t \left|H_f(\dot\gamma, \dot\gamma)\right|\mathrm{d}z \leq \Lambda t
\end{aligned}$$

where in the second line we used the fact that $\ddot\gamma = 0$ for geodesics; in the third line we noted that II maps tangent vectors to the normal space, and is thus orthogonal to $\dot\gamma$. In the final line, we applied our assumption on the bound on the norm of the Hessian. Applying the mean

value theorem to $F \circ \gamma$, we deduce that for some $s_* \in (0, s)$,

$$\langle \nabla f(\gamma(s_*)), \dot{\gamma}(s_*) \rangle = \frac{f(\gamma(s)) - f(\gamma(0))}{s} = 0.$$

Thus, we obtain an upper bound for $\langle u, v \rangle$:

$$|\langle \nabla f(p), \dot{\gamma}(0) \rangle| \leq \Lambda s_* < \Lambda s \implies \langle u, v \rangle < \frac{s^2}{\rho} + s \sin\Big(\frac{s}{r}\Big).$$

Since u is an arbitrary unit vector in $N_p f^{-1}(0)$, we have

$$\Big\langle v, N_p f^{-1}(0) \Big\rangle = \sup_{\substack{u \in N_p f^{-1}(0) \\ \|u\|=1}} \langle u, v \rangle < \frac{s^2}{\rho} + s \sin\Big(\frac{s}{r}\Big).$$

Thus, applying Corollary 3.3.2,

$$\begin{aligned}
\frac{\big\langle v, N_p f^{-1}(0) \big\rangle}{\|v\|^2} &< \frac{1}{\|v\|^2}\left(\frac{s^2}{\rho} + s \sin\Big(\frac{s}{r}\Big)\right) \\
&\leq \frac{s^2}{\big(2r \sin\big(\frac{s}{2r}\big)\big)^2}\frac{1}{\rho} + \frac{s \sin\big(\frac{s}{r}\big)}{\big(2r \sin\big(\frac{s}{2r}\big)\big)^2} \\
&= \frac{1}{\big(\operatorname{sinc}\big(\frac{s}{2r}\big)\big)^2}\frac{1}{\rho} + \frac{s}{2r^2}\cot\frac{s}{2r} \\
&< \Big(\frac{\pi}{3}\Big)^2 \frac{1}{\rho} + \frac{1}{r}
\end{aligned}$$

where we used the fact that on $[0, \pi/6]$, the functions $\operatorname{sinc}(x)$ is monotonically decreasing and $\cot x < 1/x$. Substituting this bound into the expression for $\kappa_{f^{-1}(0)}(p)$ yields the desired bound for $\kappa_{f^{-1}(0)}(p)$.

Finally, we can apply Lemma 3.5.2 to obtain the desired bound for $\kappa_{f^{-1}(-\infty,0]}(p)$. □

Proposition 3.5.5. *Let $f : M \to \mathbb{R}$ be a smooth function on an m-dimensional Riemannian submanifold M properly embedded in $\mathbb{R}^d$. Suppose 0 is a regular value of f and*

$$\inf_{p \in f^{-1}(0)} \|\nabla f(p)\| \geq \mu > 0.$$

Let $r > 0$ be a sufficiently small parameter, so that we have an upper bound on the norm of the Hessian

$\left\|H_f\right\| \leq \mu/\rho$ *on every* $B_{\frac{\pi}{3}r}(p) \cap M$ *where* $p \in f^{-1}(0)$. *Furthermore, if*

$$r \leq \max\left(\tau_M, \inf_{p \in f^{-1}(0)} \frac{\tau_M(p)}{\pi + 1}\right),$$

then the reach of the level and sublevel sets are bounded below by

$$\frac{1}{\tau_{f^{-1}(0)}} \quad \textit{and} \quad \frac{1}{\tau_{f^{-1}(-\infty,0]}} \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\rho} + \frac{2}{r}.$$

Proof. This follows from a combination of Corollary 3.5.3 and lemma 3.5.4 . □

3.5.2 Homology Inference from Uniform Random Samples

We consider an ideal scenario where $\mathbb{X}$ is a point sample directly on the regular domain $f^{-1}(-\infty, 0]$. Having obtained an explicit lower bound for the reach of sublevel sets in Proposition 3.5.5, we can apply Proposition 2.3.10 to derive the sufficient density of $\mathbb{X}$ so that the homology of $f^{-1}(-\infty, 0]$ can be recovered from $\mathbb{X}$. In particular, if $\mathbb{X}$ is drawn from a uniform random i.i.d. sample on on M, then for $\mathbb{X}$ a sufficiently large sample, the following theorems by Niyogi et al. [2008] and Wang and Wang [2020] ensure that $\mathbb{X}$ is ϵ-dense in M and $f^{-1}(-\infty, 0]$ with high confidence. Phrased in another way, for sufficiently many random samples, we can ensure that the set of Euclidean balls $\{B_\epsilon(x)\}_{x \in \mathbb{X}}$ covers M or $f^{-1}(-\infty, 0]$ for a given ϵ.

Proposition 3.5.6 ([Niyogi et al., 2008, Proposition 3.2]). *Let* $M \subset \mathbb{R}^d$ *be a properly embedded compact Riemannian manifold with positive reach* $\tau > 0$, *and* $\mathbb{X}$ *be drawn by sampling* M *in i.i.d. fashion according to the uniform probability measure on* M *as defined in Definition 3.5.1. For* $\delta < \tau/4$, *the point sample* $\mathbb{X}$ *is* δ*-dense in* M *with probability greater than* $1 - v$, *if*

$$\#\mathbb{X} \geq C_{\tilde{\tau}}(\delta)\left(\log\left(C_{\tilde{\tau}}\left(\frac{\delta}{2}\right)\right) + \log\left(\frac{1}{v}\right)\right) \quad \textit{where} \tag{3.17}$$

$$C_{\tilde{\tau}}(x) = \left(1 - \left(\frac{x}{4\tau}\right)^2\right)^{-m/2} \frac{\mathrm{vol}(M)}{\mathrm{vol}\left(B^m_{x/2}\right)}. \tag{3.18}$$

where $\mathrm{vol}\left(B^m_{x/2}\right)$ *is the volume of an* m*-dimensional closed ball with radius* $x/2$, *where* m *is the dimension* M.

The analogous statement for compact submanifolds with boundary from Wang and Wang

[2020] requires an extra geometric condition, which is the projection diffeomorphism radius r such that $\mathring{M} \cap B_r(p)$ is diffeomorphic onto its image under the projection map $P : M \to T_pM$ onto the tangent space at p, for all $p \in M$.

Theorem 3.5.7 ([Wang and Wang, 2020, Theorem 3.3]). *Let $M \subset \mathbb{R}^d$ be a properly embedded m-dimensional compact Riemannian manifold with boundary and positive reach $\tau > 0$. Suppose for all $p \in M$, the projection map $P : M \to T_pM$ is a diffeomorphism onto its image when restricted to $\mathring{M} \cap B_r(p)$ for all $r < \tilde{\tau} < \tau$. Consider a finite point sample $\mathbb{X} \subset M$ drawn by sampling M in i.i.d. fashion according to the uniform probability measure on M, as defined in Definition 3.5.1. Then with probability greater than $1 - v$, the set $\mathbb{X}$ is δ-dense in M for $\delta < \tilde{\tau}/4$ if*

$$\#\mathbb{X} > D_{\tilde{\tau}}(\delta)\left(\log\left(D_{\tilde{\tau}}\left(\frac{\delta}{2}\right)\right) + \log\left(\frac{1}{v}\right)\right) \tag{3.19}$$

$$\text{and} \quad D_{\tilde{\tau}}(x) = \frac{\mathrm{vol}(M)}{\mathrm{vol}(B_x^m)}\left(\frac{1}{2}\left(1 - \left(\frac{x}{2\tilde{\tau}}\right)^2\right)^m I_{1-\left(\frac{x}{2\tilde{\tau}}\right)^2+\left(\frac{x}{2\tilde{\tau}}\right)^4}\left(\frac{m+1}{2}, \frac{1}{2}\right)\right)^{-1}, \tag{3.20}$$

where $I_x(a, b)$ is the regularized incomplete beta function.

If a submanifold with boundary is a compact regular domain of a properly embedded Riemannian submanifold $M \subset \mathbb{R}^d$, we have shown in Corollary 3.4.5 that the projection diffeomorphism radius is inherited from that of M, which we know to be $\frac{2}{\pi}\tau_M$. Thus, we can characterise the sufficient conditions for recovering the homology of R from random finite point samples in terms of the reaches of R and M. We remark that the boundary of R is a co-dimension one subset of M and is thus measure zero.

Proposition 3.5.8. *Let M be a smooth Riemannian manifold properly embedded in $\mathbb{R}^d$ and R be a compact regular domain of M, i.e. a properly embedded codimension-0 submanifold with boundary. Suppose M and R have positive reach τ_M and τ_R respectively. Consider a finite point sample $\mathbb{X} \subset R$ drawn by sampling R in i.i.d. fashion according to the uniform probability measure on R as defined in Definition 3.5.1, where the volume form is inherited from M. Then with probability exceeding $1 - v$, the set $\mathbb{X}$ is δ-dense in R for $0 < \delta < \tilde{\tau} = \min\left(\tau_R/4, \inf_{q \in R} \tau_M(q)/\pi\right)$ if*

$$\#\mathbb{X} > D_{\tilde{\tau}}(\delta)\left(\log\left(D_{\tilde{\tau}}\left(\frac{\delta}{2}\right)\right) + \log\left(\frac{1}{v}\right)\right) \tag{3.21}$$

where $D_{\tilde{\tau}}(x)$ is given in eq. (3.20).

Proof. This follows from a combination of Theorem 3.5.7, a sampling guarantee that requires $\delta < \tau_R/4$ to be sufficiently small such that the projection of $\mathring{R} \cap B_r(p)$ onto the tangent plane at p is a homeomorphism onto its image, and Corollary 3.4.5, which states that for regular domains, the projection is a homeomorphism onto its image at any $p \in R$ if $\delta < \tau_M(p)/\pi$. □

3.6 Intersections of Level and Sublevel Sets

We focus on subsets $A \subset \mathbb{R}^d$ that can be expressed as a finite intersections of sublevel and level sets of smooth functions on a properly embedded smooth Riemannian submanifold M, i.e.

$$A = \left(\bigcap_{\alpha=1}^{k} {f_\alpha}^{-1}(0) \right) \cap \left(\bigcap_{\beta=1}^{n} {h_\beta}^{-1}(-\infty, 0] \right). \tag{3.22}$$

In this section, we show that A has positive local feature size in the ambient Euclidean space if the (sub)level sets intersect transversely as subsets of M. As such, Theorem 2.2.16 implies we can derive explicit positive lower bounds on τ_A if A is compact, using upper bounds of κ_A.

Our argument that A has positive reach relies on Theorem 2.2.10, which provides sufficient conditions for the intersection of two closed subsets of $\mathbb{R}^d$ having positive local feature size. To apply this theorem, we need to show that the normal cone of one subset is not 'opposed' to the normal cone of the other, and have positive local feature sizes where they intersect. This prohibits us from naïvely apply this theorem to the intersection of sublevel sets of embedded submanifolds, as the normal cones of any sublevel sets of smooth functions on M contain N_pM, violating the opposedness condition. Thus, we address this problem by extending the functions to the tubular neighbourhood of M and considering intersections of the corresponding sublevel sets on the tubular neighbourhood. The extension operation is formally defined in Definition 3.6.2. We then establish in Lemma 3.6.4 that these sublevel sets have positive local feature size on a neighbourhood of the original (sub)level sets and compute their corresponding normal cones. Finally, we recover the original intersections of sublevel sets on M by intersecting the intersections of the extended sublevel sets with M, and show that this intersection satisfies the conditions of Theorem 2.2.10 in Proposition 3.6.6.

Having applied Theorem 2.2.10 in such a way to show that the intersection of finitely many (sub)level sets has positive reach and obtained an explicit expression for the normal cone the

intersection at each point, we show in Proposition 3.6.7 that the quadratic curvature κ of the intersection is bounded above by sublevel sets of functions composed of convex combinations of the functions involved in the intersection. We subsequently show how to derive bound the quadratic curvature of single sublevel sets and thus derive the bound for intersections.

Let us establish some assumptions and notation to facilitate the technical proofs below. For a subset $\mathfrak{I} \subset \{1, \ldots, k\}$ that indexes the sublevel set functions h_β in eq. (3.22), we let $F_\mathfrak{I} : M \to \mathbb{R}^{\#\mathfrak{I}+k}$ denote the function

$$F_\mathfrak{I} = (h_{j_1}, \ldots, h_{j_{\#\mathfrak{I}}}, f_1, \ldots, f_k) \quad \text{s.t.} \quad j_l \in \mathfrak{I} \quad \text{and} \quad j_1 < j_2 < \ldots < j_{\#\mathfrak{I}} \quad \text{and} \tag{3.23}$$

$$A_\mathfrak{I} = A \cap {F_\mathfrak{I}}^{-1}(0). \tag{3.24}$$

For every $F_\mathfrak{I}$, we will find it useful to consider the set of smooth functions

$$\psi = \sum_\alpha t_\alpha f_\alpha + \sum_{\beta \in \mathfrak{I}} s_\beta h_\beta \tag{3.25}$$

where $t_\alpha \in \mathbb{R}$ and $s_\beta \geq 0$. We refer to this set of functions as $\Delta_\mathfrak{I}$.

We say $\mathfrak{I} \subset \{1, \ldots, k\}$ is the *maximal index set* for $p \in A$ if

$$h_\beta(p) = 0 \iff \beta \in \mathfrak{I}. \tag{3.26}$$

Finally, we focus on sublevel and level sets that intersect transversely in M.

Definition 3.6.1. We say A as defined in eq. (3.22) is *regular* at $p \in A$ if zero is a regular value for all f_α and h_β, and $\mathrm{d}F_\mathfrak{I}$ is injective, where $\mathfrak{I}$ is the maximal index set of p. In other words, the level sets ${f_\alpha}^{-1}(0)$ and ${h_\beta}^{-1}(0)$ that contain p in A intersect transversely as subsets of M.

3.6.1 Extension Argument

We first recall the notion of a tubular neighbourhood of a properly embedded submanifold $M \subset \mathbb{R}^d$. Theorem 3.1.1 states that for every such M, there exists some continuous, positive function $\delta : M \to (0, \infty)$ such that the following subset of the normal bundle

$$\mathcal{U} = \{(x, v) \in \mathcal{N}(M) \ : \ \|v\| < \delta(x)\}$$

is diffeomorphic onto its image $\mathcal{E}(\mathcal{U}) \subset \mathbb{R}^d$ under the map $\mathcal{E} : \mathcal{N}(M) \to \mathbb{R}^d$ (theorem 3.1.1)

$$\mathcal{E}(x, v) = x + v.$$

Consider the continuous function $T : \mathcal{U} \to [0, 1)$ given by

$$T(x, v) = \frac{\|v\|}{\delta(x)}.$$

Since $\mathcal{E}$ is a diffeomorphism from between $\mathcal{U} \subset \mathcal{N}(M)$ and $\mathcal{E}(\mathcal{U}) \subset \mathbb{R}^d$, we can also consider the pushforward $T_* : \mathcal{E}(\mathcal{U}) \to [0, 1)$ given by $T_* = T \circ \mathcal{E}^{-1}$. Sublevel sets of T_* parametrise a family of open neighbourhoods for M in $\mathbb{R}^d$ which we denote, for $t \in [0, 1]$,

$$M(t) = {T_*}^{-1}[0, t) \quad \text{and} \quad M[t] = {T_*}^{-1}[0, t]$$

We can denote $\mathcal{E}(\mathcal{U}) = M(1)$.

Definition 3.6.2 (Extension of Functions to Tubular Neighbourhood)**.** Let M be an m-dimensional smooth Riemannian manifold properly embedded in $\mathbb{R}^d$. We say a smooth function $\check{f} : M(1) \to \mathbb{R}$ on the δ-tubular neighbourhood is the extension of a smooth function $f : M \to \mathbb{R}$ if

$$\check{f} = f \circ \mathcal{E} \circ \xi \circ \mathcal{E}^{-1}.$$

where $\xi : \mathcal{N}(M) \to M$ be the projection $\xi(x, v) = x$ of the normal bundle onto the manifold.

We first show that a point x in the tubular neighbourhood is a critical point of an extended function only if the projection of x onto M is a critical point of the original function on M. Thus, regular values of a smooth function on the manifold are also regular values of its extension on the tubular neighbourhood

Lemma 3.6.3. *Consider some $x \in M(1)$ and $p = \mathcal{E} \circ \xi \circ \mathcal{E}^{-1}(x) \in M$. For any extension $\check{f} \in C^\infty(M(1))$ of $f \in C^\infty(M)$ to the δ-tubular neighbourhood of M, if $\nabla f(p) \in T_pM$ is non-zero, then $\boldsymbol{\nabla} \check{f}(x) \in T_x\mathbb{R}^d$ is also non-zero.*

Proof. Consider the vector subspace $W \subset T_x\mathbb{R}^d$ that is the parallel transport of $T_pM \subset T_p\mathbb{R}^d$ to $T_x\mathbb{R}^d$ in the ambident Euclidean space. Since $\mathrm{d}E$ and $\mathrm{d}\xi$ have full rank when restricted to

W, we see that $\mathrm{d}\left(\mathcal{E} \circ \xi \circ \mathcal{E}^{-1}\right)(W) \cong T_pM$. Thus for w in this subspace, there is some $u \in T_pM$ where $w = \mathrm{d}\left(\mathcal{E} \circ \xi \circ \mathcal{E}^{-1}\right)(u)$, such that by the chain rule, $\left\langle w, \boldsymbol{\nabla}\check{f}(x)\right\rangle = \langle u, \nabla f(p)\rangle$. Since one can always choose some $u \in T_pM$ not orthogonal to ∇f, we deduce that $\boldsymbol{\nabla}\check{f}(x) \neq 0$ if $\nabla f(p) \neq 0$. □

For any $x \in M(1)$, $\check{f}(x) = 0$ if and only if f is zero on the projection of x onto M. Thus, Lemma 3.6.3 implies if 0 is a regular value of f, then it is also a regular value of $\check{f}$. Therefore, a regular domain $f^{-1}(-\infty, 0] \subset M$ can be extended to a regular domain on $\check{f}^{-1}(-\infty, 0] \subset M(1)$.

For the subsequent lemma, let us denote the intersections of these regular domains with $M[t]$, where $t \in (0, 1)$ is taken to be fixed, as

$$F = \check{f}^{-1}(-\infty, 0] \cap M[t] \tag{3.27}$$

We show that F has positive local feature size where it intersects M, and that the normal cones of such points point along ∇f where $f = 0$.

Lemma 3.6.4. *Let $\check{f} : M(1) \to \mathbb{R}$ be the extension of a smooth function f on M to the $\mathcal{E}(\mathcal{U})$ (see eq. (3.1) and theorem 3.1.1), and let F be as defined in eq. (3.27) for some fixed $t \in (0, 1)$. Then F is a closed subsets of $\mathbb{R}^d$. Furthermore, if $\nabla f \neq 0$ at $p \in f^{-1}(-\infty, 0] = F \cap M$, F has positive local feature size at p, and*

$$\mathrm{Nor}(p, F) = \begin{cases} \{\lambda \nabla f(p) \ : \ \lambda \geq 0\} & \text{if} \quad f(p) = 0 \\ \{0\} & \text{if} \quad f(p) < 0. \end{cases}$$

Proof. We first show that F is a closed subsets of $\mathbb{R}^d$. Consider ${T_*}^{-1}(t, \infty)$, which is open in $M(1)$ as $T_* : M(1) \to [0, 1)$ is continuous. Since $M(1)$ is an open subset of $\mathbb{R}^d$, we deduce that ${T_*}^{-1}(t, \infty)$ is also open in $\mathbb{R}^d$. Since $M[t]$ is precisely the complement of ${T_*}^{-1}(t, \infty)$ in $M(1)$ and $\mathbb{R}^d$, it is thus closed in $\mathbb{R}^d$. As F is the sublevel set of $\check{f}$ restricted to $M[t]$, F is a closed subset of $M[t]$ and therefore a closed subset of $\mathbb{R}^d$ as well.

We want to prove that F has positive local feature size where it intersects M. Consider $p \in M \cap F = f^{-1}(-\infty, 0]$.

Suppose p is in the interior of F. As F is closed in $\mathbb{R}^d$, the interior of F is open in $\mathbb{R}^d$. Thus,

there is some sufficiently small open Euclidean ball around p that is contained in $F \subset \mathrm{UP}(F)$, which implies p has positive local feature size.

Therefore, we only need to consider the case where p lies on the boundary of F. Since $p \in M$, it must therefore lie on $f^{-1}(0)$.

We show that the medial axis of F is of bounded distance away from p. Because F is closed, the local feature size is the distance between p and the closest point in the medial axis of F. If F does not have positive local feature size at p, then for any $r > 0$, we can find some $x \in \mathrm{Med}(F)$ such that $\|x - p\| < r$. Then for any r there is at least a pair of distinct points q_1 and q_2 in F such that $\|x - q_i\| \leq \|x - p\| < r$. Let r be sufficiently small such that $B_{2r}(p) \subset M(t)$. In this case, the triangle inequality implies q_i are also in $M(t)$. Since $B_r(x)$ intersects $F = \check{f}^{-1}(-\infty, 0]$ at q_i within $M(t)$, we surmise that $\check{f}(q_i) = 0$.

As $B_r(x) \subset M(t)$, the extended function $\check{f}$ is defined on $B_r(x)$. If we choose r to be sufficiently small, then the continuity of $\nabla\check{f}$ ensures the gradients $\nabla\check{f}$ at q_i are non-zero (Lemma 3.6.3) and bounded below by some positive number μ. As $B_r(x)$ must intersect $\check{f}^{-1}(-\infty, 0]$ tangentially at q_i

$$x - q_1 \propto \boldsymbol{\nabla}\check{f}(q_1) \quad \text{and} \quad x - q_2 \propto \boldsymbol{\nabla}\check{f}(q_2)$$

As q_1, q_2 and x form an isosceles triangle in Euclidean space with two internal angles equal to θ, we can write $\|q_1 - q_2\| = 2r\cos(\theta)$ and consider the straight line from q_1 to q_2, which is contained in $B_r(x)$ by convexity. Thus $\check{f}$ is defined on the line segment. Let n be the unit vector along $q - p$ and Λ be the upper bound on the norm of the Hessian of $\check{f}$ on $B_r(x)$. Then in either case of sublevel or level sets, we have

$$2\mu\cos(\theta) \leq \left\| \left\langle \boldsymbol{\nabla}\check{f}(q_1), n \right\rangle - \left\langle \boldsymbol{\nabla}\check{f}(q_2), n \right\rangle \right\| \leq \Lambda\|q_1 - q_2\| = \Lambda 2r\cos(\theta)$$

This implies $r \geq \frac{\mu}{\Lambda}$. Thus, we reach a contradiction if we had taken r sufficiently small and x sufficiently close to p, such that $r < \frac{\mu}{\Lambda}$. Therefore, $x \in \mathrm{Med}(F)$ cannot be arbitrarily close to p and p must have positive local feature size.

As F is closed with positive local feature size at p, we can apply Theorem 2.2.9 and deduce that the normal cone of F at p is the set of rays based in $\mathbb{R}^d$ based at p, along which exists a point which has p as its nearest neighbour in F. As $F = M[t] \cap \check{f}^{-1}(-\infty, 0]$, and p is in the

interior of $M[t]$, the normal cone of F at p is the ray along $\boldsymbol{\nabla}\check{f}(p) = \nabla f(p)$ if $f(p) = 0$, and the zero vector if $f(p) < 0$. □

We now compute the normal cones of regular domains. Since $f^{-1}(-\infty, 0] = M \cap F_{\leq}$, we can apply Theorem 2.2.10 and construct the normal cone of $f^{-1}(-\infty, 0]$ from the normal cones of M and $F_{\leq}$.

Proposition 3.6.5. *Let 0 be a regular value of a smooth function f on an m-dimensional Riemannian manifold M, and consider the regular domain $f^{-1}(-\infty, 0]$. If M is properly embedded in $\mathbb{R}^d$, then the normal cone $\mathrm{Nor}\big(p, f^{-1}(-\infty, 0]\big) \subset T_p\mathbb{R}^d$ of $f^{-1}(-\infty, 0]$ at p is given by*

$$\mathrm{Nor}\Big(p, f^{-1}(-\infty, 0]\Big) = \{t\nabla f + n \ : \ t \geq 0 \quad \textit{and} \quad n \in N_pM\} \qquad \textit{if} \quad p \in f^{-1}(0)$$
$$\mathrm{Nor}\Big(p, f^{-1}(-\infty, 0]\Big) = N_pM \qquad \textit{if} \quad p \in f^{-1}(-\infty, 0).$$

Proof. Consider the intersection of M and the extension F of $f^{-1}(-\infty, 0]$ to the tubular neighbourhood. Since $f^{-1}(-\infty, 0] = M \cap F$, if the conditions of Theorem 2.2.10 are satisfied, then $\mathrm{Nor}\big(p, f^{-1}(-\infty, 0]\big) = N_pM + \mathrm{Nor}(p, F)$. Then the expression for $\mathrm{Nor}(p, F)$ in Lemma 3.6.4 implies

$$\mathrm{Nor}\Big(p, f^{-1}(-\infty, 0]\Big) = N_pM + \{\lambda\nabla f(p) \ : \ \lambda \geq 0\} \qquad \text{if} \quad p \in f^{-1}(0)$$
$$\mathrm{Nor}\Big(p, f^{-1}(-\infty, 0]\Big) = N_pM \qquad \text{if} \quad p \in f^{-1}(-\infty, 0).$$

Now we verify that M and F intersect in a way that satisfies the conditions of Theorem 2.2.10. As the normal cone of F is a subset of T_pM, and the normal cone of M is simply N_pM, any pair of vectors drawn from $\mathrm{Nor}(p, M)$ and $\mathrm{Nor}(p, F)$ respectively are orthogonal. Thus, we satisfy the condition demanded by Theorem 2.2.10 to compose the normal cone of $f^{-1}(-\infty, 0] = M \cap F$ as a vector sum of N_pM and $\mathrm{Nor}(p, F)$. □

In the proposition below, we consider transverse intersections of sublevel sets and level sets of functions on M, such as A as defined in eq. (3.22). Similar to the approach in the single function setting, we consider the extension of regular domains and level sets on M to regular domains and level sets on the tubular neighbourhood. We then proceed to recover A through successive operations of intersections, such that the conditions of Theorem 2.2.10 are satisfied

at each operation. If the conditions of Theorem 2.2.10 are satisfied at all successive intersections, then A has positive local feature size and we obtain the normal cone of A from the vector sum of the normal cones of all the subsets that participate in the successive intersections.

Proposition 3.6.6. *Suppose $A \subset \mathbb{R}^d$ can be written in the form of eq. (3.22), such that A is regular at p in the sense of Definition 3.6.1. Then $\tau_A(p) > 0$, and*

$$\mathrm{Nor}(p, A) = \left\{ n + \sum_{\alpha} t_\alpha \nabla f_\alpha + \sum_{\beta \in \mathfrak{I}} s_\beta \nabla h_\beta \ : \ n \in N_p M, \quad t_\alpha \in \mathbb{R} \quad \textit{and} \quad s_\beta \geq 0 \right\}$$

$$\mathrm{Tan}(p, A) = \left\{ v \in T_p M \ : \ \langle v, \nabla f_\alpha \rangle = 0 \quad \textit{and} \quad \langle v, \nabla h_\beta \rangle \leq 0 \quad \forall \beta \in \mathfrak{I} \right\}$$

where $\mathfrak{I}$ is the maximal index set of p as defined in Equation (3.26).

Proof. We first deal with the level sets. The gradients of f_α are non-zero and linearly independent by the assumption of A being regular at p Definition 3.6.1. Thus, the pre-image of the origin in $\mathbb{R}^k$ under the map $\varphi = (f_1, \ldots, f_k) : M \to \mathbb{R}^k$ is a properly embedded submanifold of M (without boundary). Thus by Proposition 3.1.2 and Proposition 3.1.3, $\varphi^{-1}(0)$ has positive local feature size, and the normal cone of $\varphi^{-1}(0)$ is simply the normal space of $\varphi^{-1}(0)$, which is the vector space of spanned by $N_p M$ and ∇f_α.

Let $\mathfrak{I}$ index the functions h_β that evalute to zero on p. Suppose without loss of generality $\mathfrak{I} = \{1, \ldots, j\}$ We then argue that we can ignore the functions h_β for $\beta > j$ when we consider whether p has positive local feature size, and our determination of the normal cone. As $h_\beta(p) < 0$ for $\beta > j$, by continuity of these functions, there is a sufficiently small open neighbourhood U of p in $\mathbb{R}^d$, such that $U \cap A = U \cap \tilde{A}_\mathfrak{I}$ for

$$\tilde{A}_\mathfrak{I} = \varphi^{-1}(0) \cap \bigcap_{\beta \in \mathfrak{I}} {h_\beta}^{-1}(-\infty, 0].$$

Thus we can apply Corollary 2.2.8, and deduce

$$\tau_A(p) > 0 \iff \tau_{\tilde{A}_\mathfrak{I}}(p) > 0 \quad \text{and}$$
$$\mathrm{Nor}(p, A) = \mathrm{Nor}(p, \tilde{A}_\mathfrak{I}).$$

We now focus on $\tilde{A}_\mathfrak{I}$. Let $H_\beta = \breve{h_\beta}^{-1}(-\infty, 0] \cap M[t]$ for some fixed $t \in (0, 1)$ where $\beta \in \mathfrak{I}$. By

assumption of regularity, the gradients of h_β are all non-zero and linearly independent from each other and f_α at p. By continuity, the linear independence of the gradient vectors extend to $B_\epsilon[p] \cap M$ for $\epsilon > 0$ sufficiently small. Thus, the conditions of Lemma 3.6.4 are satisfied for $q \in B_\epsilon[p] \cap H_\beta$ for any H_β. Thus, H_β has positive local feature size at all such $q \in B_\epsilon[p] \cap H_\beta$, and the normal cone is either 0 if $f(q) = 0$, or a ray along $\nabla h(q)$.

We construct $\tilde{A}_\mathfrak{I}$ by sequentially taking intersections of $\varphi^{-1}(0)$ with $H_1, \ldots, H_j$. Consider the base case where we first take the intersection between $\varphi^{-1}(0)$ and H_1. As we have noted above, $\varphi^{-1}(0)$ and H_i have positive local feature sizes on $C_1 = B_\epsilon[p] \cap \varphi^{-1}(0) \cap H_1$. For $q \in C_1$, the vectors in $N_q\varphi^{-1}(0)$ are not antiparallel to those in $\mathrm{Nor}(q, H_1)$, due to the linear independence of the gradient vectors on C_1.

Thus the conditions of Theorem 2.2.10 are met: $\varphi^{-1}(0) \cap H_1$ has positive feature size at p, and for $q \in C_1$,

$$\begin{aligned} \mathrm{Nor}\Big(q, \varphi^{-1}(0) \cap H_1\Big) &= N_q\Big(\varphi^{-1}(0)\Big) + \{\lambda \nabla h_1(q) \ : \ \lambda \geq 0\} && \text{if} \quad h_1(q) = 0 \\ \mathrm{Nor}\Big(q, \varphi^{-1}(0) \cap H_1\Big) &= N_q\Big(\varphi^{-1}(0)\Big) && \text{if} \quad h_1(q) < 0. \end{aligned}$$

We once again note that the positivity of local feature size of $\varphi^{-1}(0) \cap H_1$ at p extends to points in $B_{\epsilon_1}[p] \cap \varphi^{-1}(0) \cap H_1 \subset C_1$ for some sufficiently small $\epsilon_1 \in (0, \epsilon]$.

Let $\tilde{A}_l = \varphi^{-1}(0) \cap H_1 \cap \cdots \cap H_l$. Having established a base case for $\tilde{A}_1$, let us assume, $\tilde{A}_l$ has positive local feature size at p; and for $q \in C_l = B_{\epsilon_{l-1}}[p] \cap \tilde{A}_l$,

$$\mathrm{Nor}\big(q, \cap \tilde{A}_l\big) = N_q\Big(\varphi^{-1}(0)\Big) + \left\{ \sum_{\beta=1}^{l} s_\beta \nabla h_\beta(q) \ : \ s_\beta \geq 0, \quad \text{and} \quad s_\beta = 0 \quad \text{if} \quad h_\beta(q) = 0 \right\}$$

where we have taken $\epsilon_0 = \epsilon$. Now let us consider the inductive step where we intersect $\tilde{A}_l$ with H_{l+1} to yield $\tilde{A}_{l+1}$, and check whether the conditions of Theorem 2.2.10 are being met. Let us choose $\epsilon_l \in (0, \epsilon_{l-1})$ sufficiently small, so that the local feature size of $\tilde{A}_l$ is positive on $B_{\epsilon_l}(p) \cap C_l$ by continuity from $\tau_{\tilde{A}_l}(p) > 0$. As $\epsilon_l \leq \epsilon$ and the local feature size of H_{l+1} is assumed to be positive on $B_\epsilon(p) \cap M$, the local feature size of H_{l+1} is also positive on $B_{\epsilon_l}(x) \cap C_l$.

As ∇h_{l+1} is linearly independent from ∇h_β for $\beta = 1, \ldots, l$ on $B_\epsilon[p] \supset B_{\epsilon_l}(x) \cap C_l$, we can

combine the normal cones of $\tilde{A}_l$ and H_{l+1} at all $q \in C_{l+1}$,

$$\mathrm{Nor}(q, \cap \tilde{A}_l) = N_q\Big(\varphi^{-1}(0)\Big) + \left\{ \sum_{\beta=1}^{l} s_\beta \nabla h_\beta(q) \ : \ s_\beta \geq 0, \quad \text{and} \quad s_\beta = 0 \quad \text{if} \quad h_\beta(q) = 0 \right\}$$

and deduce that $\tilde{A}_{l+1}$ has positive feature size at p. We have thus shown that our claim for l inductively implies the corresponding claim for $l+1$.

Taking l up to j to encompass all functions h_β indexed in $\mathfrak{I}$, we thus prove the statement in the proposition for the positivity of the local feature size of A and the expression of the normal cone at p. As $\tau_A(p) > 0$, we can deduce the form of $\mathrm{Tan}(p, A)$ by taking the convex dual of $\mathrm{Nor}(p, A)$ by Theorem 2.2.9. □

3.6.2 Lower Bounds of the Reach

Having established an expression for the normal cone and showed that A has positive local feature size where it is regular, we can now compute a lower bound to the curvature $\kappa_A(p)$ of A at p. We use the insight below that boils down the computation of $\kappa_A(p)$ to that of a sublevel set.

Proposition 3.6.7. *Suppose $A \subset \mathbb{R}^d$ that can be written as eq. (3.22), such that A is regular at p in the sense of Definition 3.6.1. Then*

$$\kappa_A(p) \leq \sup_{\psi \in \Delta_{\mathfrak{I}}} \kappa_{\psi^{-1}(-\infty, 0]}(p).$$

where $\mathfrak{I}$ is the maximal index set of $p \in A$ (Equation (3.26)), and $\Delta_{\mathfrak{I}}$ is the set of smooth functions parametrised over a cone as defined in eq. (3.25).

Proof. We first recall the definition of κ (Equation (2.10)) and Remark 2.2.17:

$$\frac{\kappa_A(p)}{2} = \sup_{\substack{q \in A \\ q \neq p}} \frac{d_{\mathbb{R}^d}(q - p, \mathrm{Tan}(p, A))}{\|q - p\|^2} = \sup_{\substack{q \in A \\ q \neq p}} \frac{\langle q - p, \mathrm{Nor}(p, A) \rangle}{\|q - p\|^2}$$

We first observe for any $q \in A$, and any function $\psi \in \Delta_{\mathfrak{I}}$, that $\psi(q) \leq 0$:

$$\psi(q) = \sum_{\alpha=1}^{k} t_\alpha f_\alpha(q) + \sum_{\beta \in \mathfrak{I}(p)} s_\beta h_\beta(q) \leq 0 \tag{3.28}$$

$$\because \quad f_\alpha(q) = 0, \quad h_\beta(q) \leq 0 \quad \text{and} \quad s_\beta \geq 0. \tag{3.29}$$

Similarly, we can show that $\psi(p) = 0$. Hence both p and q lie in the sublevel set $\psi^{-1}(-\infty, 0]$. Next, let us consider the projection onto the normal cone of A at p

$$\langle q - p, \mathrm{Nor}(p, A) \rangle := \sup \{ \langle q - p, u \rangle \ : \ n \in \mathrm{Nor}(p, A) \quad \text{and} \quad \|u\| \leq 1 \}.$$

Unless $q - p$ lies in the tangent set $\mathrm{Tan}(p, A)$ (which is the convex dual of $\mathrm{Nor}(p, A)$ by Theorem 2.2.9), the vector $q - p$ has a positive orthogonal projection onto $\mathrm{Nor}(p, A)$, and $\langle q - p, \mathrm{Nor}(p, A) \rangle = \langle q - p, u \rangle$ for some unit vector $u \in \mathrm{Nor}(p, A)$. We recall from Proposition 3.6.6 that

$$\mathrm{Nor}(p, A) = \left\{ t_0 n + \sum_\alpha t_\alpha \nabla f_\alpha + \sum_{\beta \in \mathfrak{I}(p)} s_\beta \nabla h_\beta \ : \ n \in N_p M, \quad t_\alpha \in \mathbb{R} \quad \text{and} \quad s_\beta \geq 0 \right\}.$$

Therefore we can write u as some

$$\begin{aligned} u &= n + \sum_{\alpha=1}^{k} t_\alpha^* \nabla f_\alpha(p) + \sum_{\beta \in \mathfrak{I}(p)} s_\beta^* \nabla h_\beta(p) \\ &= n + \nabla \psi \in \mathrm{Nor}\Big(p, \psi^{-1}(-\infty, 0]\Big) \end{aligned}$$

where $n \in N_p M$, and $\psi : M \to \mathbb{R}$ is a smooth function

$$\psi = \sum_{\alpha=1}^{k} t_\alpha^* f_\alpha + \sum_{\beta \in \mathfrak{I}(p)} s_\beta^* h_\beta.$$

Since we have assumed that the gradients of f_α and h_β for $\beta \in \mathfrak{I}$ are non-zero and linearly independent, the gradient of ψ is non-zero at p. Hence

$$\langle q - p, \mathrm{Nor}(p, A) \rangle = \langle q - p, u \rangle \leq \left\langle q - p, \mathrm{Nor}\Big(p, \psi^{-1}(-\infty, 0]\Big) \right\rangle.$$

Since we have shown that both p and q lie in $\psi^{-1}(-\infty, 0]$ (eq. (3.29)), we can apply the defini-

tion of κ (eq. (2.10)) to $\psi^{-1}(-\infty, 0]$, and show that

$$\langle q - p, \mathrm{Nor}(p, A)\rangle \leq \left\langle q - p, \mathrm{Nor}\left(p, \psi^{-1}(-\infty, 0]\right)\right\rangle \leq \frac{\kappa_{\psi^{-1}(-\infty,0]}(p)}{2}\|p - q\|^2.$$

Thus, considering all $q \in A$, we deduce that

$$\kappa_A(p) \leq \sup_{\psi \in \Delta_{\mathfrak{I}}} \kappa_{\psi^{-1}(-\infty,0]}(p).$$

□

We can then obtain an upper bound for $\kappa_A(p)$ by applying Lemma 3.5.4 to bound $\kappa_{\psi^{-1}(-\infty,0]}(p)$ acrosss all functions $\psi \in \Delta_{\mathfrak{I}}$.

Proposition 3.6.8. *Suppose $A \subset \mathbb{R}^d$ can be written as eq. (3.22),*

$$A = \left(\bigcap_{\beta=1}^{k} {f_\alpha}^{-1}(0)\right) \cap \left(\bigcap_{\alpha=1}^{n} {h_\beta}^{-1}(-\infty, 0]\right).$$

for smooth functions f and h on an embedded submanifold $M \subset \mathbb{R}^d$, such that A is regular at p in the sense of Definition 3.6.1. Let $\mathfrak{I}$ be the maximal index set of p in A (Equation (3.26)) and consider the Jacobian $\mathrm{d}F_{\mathfrak{I}}$ of the map $F_{\mathfrak{I}}$ as defined in (eq. (3.23)). Let μ be the $k + \#\mathfrak{I}(p)^{th}$ smallest singular value of $\mathrm{d}F_{\mathfrak{I}(p)}$ evaluated at p, and suppose it is positive $\mu > 0$. Suppose also for $r = \max\left(\tau_M, \frac{\tau_M(p)}{3}\right)$, the norms of the Hessians of f_α and h_β are bounded by Λ_α and Ω_β respectively on $B_{\frac{\pi}{3}r}(p) \cap M$, then

$$\kappa_A(p) \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\rho} + \frac{2}{r}$$
$$\text{where} \quad \rho = \frac{1}{\mu}\sqrt{\sum_{\alpha=1}^{k} \Lambda_\alpha^2 + \sum_{\beta \in \mathfrak{I}} \Lambda_\alpha^2}.$$

Proof. For brevity, let $\mathfrak{I} = \mathfrak{I}(p)$. We begin with Proposition 3.6.7, which states that

$$\kappa_A(p) \leq \sup_{\psi \in \Delta_{\mathfrak{I}}} \kappa_{F^{-1}(-\infty,0]}(p)$$

where $\Delta_{\mathfrak{I}}$ is the family of functions defined in eq. (3.25). Recall that we can write $\psi \in \Delta_{\mathfrak{I}}$ as

$$\psi = \sum_{\alpha} t_\alpha f_\alpha + \sum_{\beta \in \mathfrak{I}} s_\beta h_\beta$$

where $s_\beta \geq 0$. For the special case $\psi = 0 \in \Delta_\mathfrak{I}$, the sublevel set $F^{-1}(-\infty, 0]$ is simply the whole manifold, and we can bound $\kappa_{\psi^{-1}(-\infty,0]}(p)$ from above using $\frac{1}{\tau_M(p)}$ (Corollary 2.2.19), which is below the stated bound. Thus let us consider the case where there at least one of the coefficients is non-zero. We have shown in Lemma 3.5.4 that $\kappa_{\psi^{-1}(-\infty,0]}(p)$ on the right hand side can be bounded by

$$\kappa_{\psi^{-1}(-\infty,0]}(p) \leq \left(\frac{\pi}{3}\right)^2 \frac{2\Lambda_\psi}{\mu_\psi} + \frac{2}{r}.$$

if $0 < \mu_\psi \leq \|\nabla\psi(p)\|$ and Λ_ψ is greater or equal to norm on the Hessian of ψ on $B_{\frac{\pi}{3}r}(p) \cap M$ for $r \leq \max\left(\tau_M, \frac{\tau_M(p)}{3}\right)$.

We would thus like to obtain lower bounds for $\|\nabla\psi(p)\|$ and the norm of the Hessian of $\psi \in \Delta_\mathfrak{I}$ at p. If we let θ be the column vector with entries consisting of coefficients t_α and s_β, then we can write

$$\mathrm{d}\psi = \mathrm{d}F_\mathfrak{I}\,\theta.$$

As the column vectors ∇f_{α_i} of the Jacobian matrix $\mathrm{d}F_\mathfrak{I}$ are non-zero and linearly independent at p,

$$\|\mathrm{d}\psi\|^2 = \theta^\intercal\,\mathrm{d}F_\mathfrak{I}{}^\intercal\,\mathrm{d}F_\mathfrak{I}\,\theta > 0 \quad \text{for} \quad \theta \neq 0.$$

As the sublevel set $\psi^{-1}(-\infty, 0]$ is defined up to a rescaling of F by positive real numbers, we can rescale ψ such that $\mu_\psi = \|\mathrm{d}\psi\,(p)\| = 1$.

We turn our attention to bounding the norm of the Hessian of ψ. For $v \in T_qM$, we can write

$$H_\psi(v, v) = \sum_{\alpha\in\mathfrak{I}} t_\alpha H_{f_\alpha}(v, v).$$

Thus the norm of the Hessian of ψ on $B_{\frac{\pi}{3}r}(p) \cap M$ is bounded above by

$$\left\|H_\psi\right\| \leq \sum_\alpha |t_\alpha|\Lambda_\alpha + \sum_{\beta\in\mathfrak{I}} |s_\beta|\Omega_\beta$$

where $\left\|H_{f_\alpha}\right\| \leq \Lambda_\alpha$ and $\|H_{h_\alpha}\| \leq \Omega_\beta$ on $B_{\frac{\pi}{3}r}(p) \cap M$. Hence, we can write down an upper

bound for the functional curvature term of $\kappa_{\psi^{-1}(-\infty,0]}(p)$ for all $\psi \in \Delta_{\mathfrak{I}}$:

$$\frac{\Lambda_\psi}{\mu_\psi} \leq \max_\theta \lambda^\intercal \theta \quad \text{s.t.} \quad \|\mathrm{d}F_{\mathfrak{I}}\,\theta\| = 1 \quad \text{and} \quad \theta \geq 0$$

where λ is a column vector with entries consisting of the Hessian norm bounds Λ_α and Ω_β. As $\lambda \geq 0$, the solution to this maximisation problem is

$$\max_\theta \lambda^\intercal \theta = \sqrt{\lambda^\intercal (\mathrm{d}F_{\mathfrak{I}}^\intercal\, \mathrm{d}F_{\mathfrak{I}})^{-1} \lambda} \leq \frac{1}{\mu}\|\lambda\| = \frac{1}{\mu}\sqrt{\sum_\alpha \Lambda_\alpha^2 + \sum_{\beta\in\mathfrak{I}} \Omega_\beta^2}$$

where μ is the smallest non-zero singular value of $\mathrm{d}F_{\mathfrak{I}}$, the Jacobian matrix which we recall we have assumed to have non-zero and linearly indepedent column vectors. Hence

$$\kappa_A(p) \leq \sup_{\psi\in\Delta_{\mathfrak{I}}} \kappa_{\psi^{-1}(-\infty,0]}(p) \leq \left(\frac{\pi}{3}\right)^2 \frac{2\Lambda_\psi}{\mu_\psi} + \frac{2}{r} \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\mu}\sqrt{\sum_\alpha \Lambda_\alpha^2 + \sum_{\beta\in\mathfrak{I}} \Omega_\beta^2} + \frac{2}{r},$$

which completes the proof. □

Proposition 3.6.9. *Suppose $A \subset \mathbb{R}^d$ can be written as eq. (3.22),*

$$A = \left(\bigcap_{\beta=1}^{k} {f_\alpha}^{-1}(0)\right) \cap \left(\bigcap_{\alpha=1}^{n} {h_\beta}^{-1}(-\infty, 0]\right).$$

for smooth functions f and h on an embedded submanifold $M \subset \mathbb{R}^d$, such that A is regular at in the sense of Definition 3.6.1 at every point. Let

$$\mu_{\mathfrak{I}} = \inf_{p\in A_{\mathfrak{I}}} \sigma_{k+\#\mathfrak{I}}(\mathrm{d}F_{\mathfrak{I}}\,(p))$$

where for an index set $\mathfrak{I}$, the subset $A_{\mathfrak{I}}$ of A as defined in eq. (3.24), the function $F_{\mathfrak{I}}$ is as defined in eq. (3.23), and $\sigma_i(X)$ denotes the i^{th} singular value of a matrix X. In addition, for $0 < r < \max(\frac{1}{\pi+1}\inf_{p\in A}\tau_M(p), \tau_M)$,

$$\begin{aligned}
\Lambda_\alpha &= \sup\left\{\left\|H_{f_\alpha}(p)\right\| \;:\; p \in M, \quad d_{\mathbb{R}^d}(p, A) < \frac{\pi}{3}r\right\} \\
\Omega_\beta &= \sup\left\{\left\|H_{h_\beta}(p)\right\| \;:\; p \in M, \quad d_{\mathbb{R}^d}(p, A) < \frac{\pi}{3}r\right\} \\
\textit{and} \quad \Lambda_{\mathfrak{I}} &= \sum_\alpha \Lambda_\alpha^2 + \sum_{\beta\in\mathfrak{I}} \Omega_\beta^2.
\end{aligned}$$

Let

$$\rho = \min\left\{ \frac{\mu_{\mathfrak{I}}}{\Lambda_{\mathfrak{I}}} \ : \ A_{\mathfrak{I}} \neq \emptyset \right\}$$

If $\rho > 0$ and $r > 0$, then the reach of A has a lower bound given by

$$\frac{1}{\tau_A} \leq \left(\frac{\pi}{3}\right)^2 \frac{2}{\rho} + \frac{2}{r}.$$

Proof. This follows straightforwardly from applying Proposition 3.6.8 to every point in $p \in A$ to derive an upper bound of $\kappa_A(p)$, and then applying Theorem 2.2.16 to bound $\frac{1}{\tau_A}$ with an upper bound of κ_A. □

Chapter 4

Morse Indices and Gromoll-Meyer Pairs

In this chapter, we consider critical sets $\mathrm{Crit}(f)$ of smooth functions f on manifolds M. A subset $S \subset M$ is critical if $\nabla f = 0$ on S. In particular, we are interested in the *Morse indices* Conley [1978] of compact, dynamically isolated critical sets, and methods of computing them from finite local data near S. The required local data is called a Gromoll-Meyer pair; it consists of $\mathscr{W}$, a compact neighbourhood of S, along with a subset of its boundary $\mathscr{W}_-$, which we refer to as the *exit set* of $\mathscr{W}$. The relative homology of the pair, referred to as the *critical group* of S in Chang and Ghoussoub [1996],

$$C_\bullet(S, f) = H_\bullet(\mathscr{W}, \mathscr{W}_-).$$

is essentially the Morse index of S; we make explicit the connections in Section 4.1.1. In Section 4.1, we give the definitions of a dynamically isolated critical set S, its critical group, and Gromoll-Meyer pairs. We then show an explicit construction of Gromoll-Meyer pairs for any compact dynamically isolated critical set, using no more assumptions than a lower bound of $\|\nabla f\|$ on an annular neighbourhood of S. We denote our construction of this special instance of a Gromoll-Meyer pair $(\mathscr{M}, \mathscr{M}_-)$.

In our construction, both $(\mathscr{M}, \mathscr{M}_-)$ are compact, transverse intersections of level sets and sublevel sets of smooth functions on the manifold M. Thus, if M is smoothly and properly

embedded in some Euclidean space $\mathbb{R}^d$, the images of both $\mathscr{M}$ and $\mathscr{M}_-$ under the embedding map have positive reach, as we have shown in the previous chapter. As such, we can bring the techniques of homological inference in Chapter 2 to bear to infer $H_\bullet(\mathscr{M}, \mathscr{M}_-)$ using a finite set of point samples. In Section 4.3, we discuss conditions under which a point cloud $\mathbb{X}$ provides sufficient data for inferring $H_\bullet(\mathscr{M}, \mathscr{M}_-)$, as well as a simple procedure for obtaining such a point cloud.

4.1 Gromoll-Meyer Theory

Let M be a smooth m-dimensional Riemannian manifold. A *flow* on M is a continuous map $u : \mathbb{R} \times M \to M$, such that

$$u(0, x) = x \quad \forall x \in M \quad \text{and} \quad u(t + s, x) = u(t, u(s, x)). \tag{4.1}$$

In particular, we focus on flows generated by gradients of smooth functions. For $f \in C^2(M)$ continuously differentiable, the *gradient flow* $\sigma : \mathbb{R} \times M \to M$ refers to the flow that solves the following Cauchy problem:

$$\frac{\partial \sigma}{\partial t} = -\nabla f, \quad \text{and} \quad \sigma(0, x) = x \quad \forall x \in M. \tag{4.2}$$

For a gradient flow σ, the critical points $\mathrm{Crit}(f) = \left\{ x \in M \mid \mathrm{d}f\big|_x = 0 \right\}$ of f are precisely the fixed points of σ. We focus on the following class of critical sets of f.

Definition 4.1.1 (Dynamically Isolated Critical Set [Chang and Ghoussoub, 1996, Definition I.10])**.** A subset $S \subset \mathrm{Crit}(f)$ is said to be *dynamically isolated* if there exists a closed neighbourhood $\mathcal{O}$ of S and regular values $\alpha < \beta$ of f such that

(K1) $\mathcal{O} \subset f^{-1}[\alpha, \beta]$

(K2) $\mathrm{cl}(\sigma(\mathbb{R}, \mathcal{O})) \cap \mathrm{Crit}(f) \cap f^{-1}[\alpha, \beta] = S$

where the flow σ is given by eq. (4.2). We say $(\mathcal{O}, \alpha, \beta)$ is an *isolating triplet* for S.

We elucidate this definition with a counter example. Consider two subsets S_1 and S_2 of $\mathrm{Crit}(f)$ with two closed disjoints neighbourhoods $\mathcal{O}_1, \mathcal{O}_2 \subset f^{-1}[a, b]$ respectively. Furthermore,

assume they are linked by gradient flow: $\sigma(\mathbb{R}_{>0}, \mathcal{O}_1) \cap \mathcal{O}_2 \neq \emptyset$ and $\sigma(\mathbb{R}_{<0}, \mathcal{O}_2) \cap \mathcal{O}_1 \neq \emptyset$. While they are isolated from each other by disjoint neighbourhoods, they are *not* dynamically isolated. However, their union $S_1 \cup S_2$ can be if it is only linked with other critical sets along the gradient flow that are outside $f^{-1}[a, b]$ and isolated from other critical sets.

Example 4.1.2. Let p be an isolated critical point of $f \in C^2(M)$. Suppose a, b are regular values of f. If p is the only critical point in $f^{-1}[a, b]$, then we can choose $\mathcal{O} = f^{-1}[a, b]$.

To each dynamically isolated critical set S, we can associate a pair of $(\mathscr{W}, \mathscr{W}_-)$ such that their relative homology is an invariant associated with S.

Definition 4.1.3 (Gromoll-Meyer Pair)**.** Suppose f is a C^2 function and let S be a dynamically isolated critical set of f. Let σ be the gradient flow defined by eq. (4.2). A subset $\mathscr{W}$ together with its *exit set*

$$\mathscr{W}_- = \{x \in \mathscr{W} \ : \ \sigma(t, x) \notin \mathscr{W} \forall t > 0\} \tag{4.3}$$

are a *Gromoll-Meyer pair* for S if

(GM1) $\mathscr{W}$ is a closed neighbourhood of S

(GM2) $\mathscr{W} \cap \mathrm{Crit}(f) = S$;

(GM3) $\mathscr{W} \cap f^{-1}(-\infty, a) = \emptyset$ for some a ;

(GM4) $\mathscr{W}$ satisfies the *mean value property*:

$$\sigma(t_1, x), \sigma(t_2, x) \in \mathscr{W} \implies \sigma([t_1, t_2], x) \subset \mathscr{W}; \tag{4.4}$$

(GM5) $\mathscr{W}_-$ is closed and is a union of finitely many submanifolds transverse to the flow σ.

The conditions placed on Gromoll-Meyer pairs ensure any such pair $(\mathscr{W}, \mathscr{W}_-)$ encodes consistent information about the role of S in the gradient flow σ on M. In particular, **(GM4)** implies the flow is a one way street in $\mathscr{W}$: once we exit the set, we cannot return. In particular, this implies an element of the boundary $\partial\mathscr{W}$ belongs to the exit set defined above (eq. (4.3)) as where the flow leaves $\mathscr{W}$ at the boundary, or the entry set $\mathscr{W}_+$ where the flow enters. The transversality condition **(GM5)** ensures the flow cannot exit $\mathscr{W}$ tangentially at any point.

The following result from Chang and Ghoussoub [1996] shows that the relative homology of Gromoll-Meyer pairs is an invariant of the critical set.

Theorem 4.1.4 ([Chang and Ghoussoub, 1996, Proposition II.2 and Theorem III.3]). *Given $f \in C^2(M)$, suppose $S \subset \mathrm{Crit}(f)$ is a dynamically isolated critical set of f with isolating triplet $(\mathcal{O}, \alpha, \beta)$. Let $\mathcal{O}^+$ be the closure of $\sigma(\mathbb{R}_{\geq 0}, \mathcal{O})$ for the gradient flow σ given in eq. (4.2). Then the* critical group *of S, defined as*

$$C_\bullet(f, S) := H_\bullet\left(f^{-1}(-\infty, \beta] \cap \mathcal{O}^+, f^{-1}(-\infty, \alpha] \cap \mathcal{O}^+\right) \tag{4.5}$$

is independent of the choice of isolating triplets: there is an isomorphism

$$C_\bullet(f, S) \cong H_\bullet(\mathscr{W}, \mathscr{W}_-). \tag{4.6}$$

for any given Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$.

Example 4.1.5. Let p be an isolated critical point of $f \in C^2(M)$. Suppose α, β are regular values of f. If $\left(f^{-1}[\alpha, \beta], \alpha, \beta\right)$ is an isolating triplet for p, then we can choose $\left(f^{-1}[\alpha, \beta], f^{-1}(\alpha)\right)$ to be a Gromoll-Meyer pair, and

$$C_\bullet(f, p) \cong H_\bullet(f^{-1}[\alpha, \beta], f^{-1}(\alpha)).$$

Furthermore, if p is non-degenerate, $H_\bullet(f^{-1}[\alpha, \beta], f^{-1}(\alpha))$ is the Morse index of p.

Proposition 4.1.6. *Let p be a non-degenerate critical point of $f \in C^2(M)$ with index $\mathscr{I}(f, p)$. If $(\mathscr{W}, \mathscr{W}_-)$ is a Gromoll-Meyer pair for p, then*

$$H_\bullet(\mathscr{W}, \mathscr{W}_-) \cong \tilde{H}_\bullet\left(\mathbb{S}^{\mathscr{I}(f,p)}\right), \tag{4.7}$$

where $\tilde{H}_\bullet$ denotes reduced homology.

Proof. Let $i = \mathscr{I}(f, p)$. We consider a sufficiently small neighbourhood U of p such that by Morse's lemma [Milnor et al., 1973], we have a set of local coordinates $x_1, \dots x_m$ such that $p \in f^{-1}(0)$ is at the origin,

$$f(x) = -x_1^2 - \cdots - x_i^2 + x_{i+1}^2 + \cdots + x_m^2.$$

We choose a sufficiently small $a > 0$, such that the cuboid

$$K = \left\{ x \in U \mid \max_{l=1,\dots,m} |x_l| \leq a \right\}$$

lies in the domain of coordinates x_i. We can also define

$$L = \{x \in K \mid |x_l| = a,\ l = 1, \dots i\},$$

such that pair (K, L) is a Gromoll-Meyer pair for p. Let us consider the i-dimensional disk K_0 and its boundary

$$K_0 = \{x \in K \mid |x_l| \leq a,\ l = 1, \dots i\} \cap \{x \in K \mid x_j = 0,\ j = i+1, \dots, m\} \approx D^i$$
$$\partial K_0 = \{x \in K \mid |x_l| = a,\ l = 1, \dots i\} \cap \{x \in K \mid x_j = 0,\ j = i+1, \dots, m\} \approx \partial D^i.$$

We now note that $K = K_0 \times [-a, a]^{m-i}$ and $L = \partial K_0 \times [-a, a]^{m-i}$, and construct a deformation retract of (K, L) to $(K_0, \partial K_0)$ by sending the coordinates $x_{i+1}, \dots, x_m$ to zero. Thus, we have

$$H_\bullet(K, L) \cong H_\bullet(K_0, \partial K_0) \cong H_\bullet\left(D^i, \partial D^i\right) \cong \tilde{H}_\bullet(\mathbb{S}^i).$$

Since the relative homology of Gromoll-Meyer pairs is invariant by Theorem 4.1.4,

$$H_\bullet(\mathscr{W}, \mathscr{W}_-) \cong H_\bullet(K, L) \cong \tilde{H}_\bullet(\mathbb{S}^i).$$

□

4.1.1 Correspondence with Conley Index Theory

The critical group of $S \in \mathrm{Crit}(f)$ is essentially the *Conley index* of the minimal invariant set of the gradient flow σ containing S. We restrict our attention to critical sets admitting a compact Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$, and show that a Gromoll-Meyer pair for S is a Conley index pair for the maximal invariant set of the gradient flow contained in $\mathscr{W}$. We refer the reader to Conley [1978], Mischaikow and Mrozek [2002] for a more complete exposition about Conley index theory and Chang and Ghoussoub [1996] for further technical details of the correspondence between Gromoll-Meyer theory and Conley index theory where Gromoll-

Meyer pairs are not necessarily compact.

Definition 4.1.7 (Isolated Invariant Set)**.** An invariant set $S \subset M$ of flow u, i.e. a subset satisfying $u(\mathbb{R}, S) = S$, is said to be an *isolated invariant set* if there exists a compact subset $N \subset M$ such that $S = \mathrm{Inv}(N) \subset \mathrm{int}(N)$; here $\mathrm{int}(N)$ denotes the interior of N, and $\mathrm{Inv}(N)$ is the maximal closed invariant set contained in N:

$$\mathrm{Inv}(N) = \{x \in N \ : \ u(\mathbb{R}, x) \subset N\} = \bigcap_{t \in \mathbb{R}} u(t, N). \tag{4.8}$$

Such a neighbourhood N is said to be an *isolating neighbourhood* of S.

The homological Conley index of an isolated invariant set is the relative homology of a pair of spaces called index pairs.

Definition 4.1.8 (Index Pair)**.** A pair of compact subsets $N \supset N_-$ are said to be an *index pair* for an isolated invariant set S if

(IP1) $S = \mathrm{Inv}(\mathrm{cl}(N \setminus N_-))$ and $S \subset \mathrm{int}(N \setminus N_-)$;

(IP2) N_- is an *exit set*: for any $x \in N$ such that $u(t, x) \notin N$ for some $t > 0$, there exists some $t_0 \in [0, t]$, such that $u(t_0, x) \in N_-$ and $u([0, t_0], x) \subset N$.

(IP3) N_- is *positively invariant* in N: for $x \in N_-$,

$$u([0, t], x) \in N \implies u([0, t], x) \in N_-.$$

Theorem 4.1.9. *Suppose (N, N_-) and (N', N'_-) are index pairs for an isolated invariant set S. Then*

$$H_\bullet(N, N_-) \cong H_\bullet(N', N'_-).$$

We refer to $H_\bullet(N, N_-)$ as the (homological) Conley index of S.

We recall the correspondence between Conley index pairs and Gromoll-Meyer pairs for critical sets that admit a compact Gromoll-Meyer pair.

Lemma 4.1.10 ([Chang and Ghoussoub, 1996, Lemma I.8(1)])**.** *If W is a compact neighbourhood of a critical set $S \subset \mathrm{Crit}(f)$ satisfying the mean value property* ***(GM4)*** *with respect to the gradient*

flow σ eq. (4.2) and $W \cap \mathrm{Crit}(f) = S$ (GM2), then W is an isolating neighbourhood of $\mathrm{Inv}(W)$*, i.e.*

$$\mathrm{Inv}(W) \subset \mathrm{int}(W).$$

Lemma 4.1.11. *For a Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$, the closure $\mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$ equals $\mathscr{W}$.*

Proof. Suppose $\mathscr{W} \not\subset \mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$. Then there is some point $x \in \mathscr{W}$ such that $x \notin \mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$. Therefore, $x \in \mathscr{W}_-$ and there is some neighbourhood U of x such that $U \cap \mathscr{W} \setminus \mathscr{W}_- = \emptyset$. Since $x \in \mathscr{W}_-$, the definition of $\mathscr{W}_-$ in eq. (4.3) along with the transversality condition **(GM5)** imply that for the gradient flow σ intersecting $\mathscr{W}_-$ at x, we have $\sigma((-\infty, 0], x) \cap \mathscr{W} \setminus \mathscr{W}_- \neq \emptyset$. This contradicts the assertion that $U \cap (\mathscr{W} \setminus \mathscr{W}_-) = \emptyset$. Thus, we deduce that $\mathscr{W} \subset \mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$.

Since $\mathscr{W}$ is closed, $\mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-) \subset \mathrm{cl}(\mathscr{W}) = \mathscr{W}$. Therefore, $\mathscr{W} = \mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$ □

Lemma 4.1.12. *If $(\mathscr{W}, \mathscr{W}_-)$ is a compact Gromoll-Meyer pair for $S \subset \mathrm{Crit}(f)$, then $\mathscr{W}$ is an isolating neighbourhood of* $\mathrm{Inv}(\mathscr{W})$*; futhermore, $(\mathscr{W}, \mathscr{W}_-)$ satisfies (IP1), i.e.*

$$\mathrm{Inv}(\mathscr{W}) = \mathrm{Inv}(\mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)) \quad \textit{and} \quad \mathrm{Inv}(\mathscr{W}) \subset \mathrm{int}(\mathscr{W} \setminus \mathscr{W}_-) = \mathrm{int}(\mathscr{W}).$$

Proof. Since we have $\mathscr{W} = \mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)$ from Lemma 4.1.11, we trivially satisfy the first condition of **(IP1)**

$$\mathrm{Inv}(\mathscr{W}) = \mathrm{Inv}(\mathrm{cl}(\mathscr{W} \setminus \mathscr{W}_-)).$$

Since $\mathscr{W}$ is a compact neighbourhood of $S \subset \mathrm{Crit}(f)$ that satisfies **(GM2)** and **(GM5)** with respect to the gradient flow σ eq. (4.2), Lemma 4.1.10 implies $\mathscr{W}$ is an isolating neighbourhood of $\mathrm{Inv}(\mathscr{W})$, i.e.

$$\mathrm{Inv}(\mathscr{W}) \subset \mathrm{int}(\mathscr{W}).$$

As $\mathscr{W}_- \subset \partial\mathscr{W}$,

$$\mathrm{int}(\mathscr{W}) \supset \mathrm{int}(\mathscr{W} \setminus \mathscr{W}_-) \supset \mathrm{int}(\mathscr{W} \setminus \partial\mathscr{W}) = \mathrm{int}(\mathscr{W})$$

where the final equality is due to the fact that $\mathscr{W}$ is closed, and thus $\mathscr{W} \setminus \partial\mathscr{W} = \mathrm{int}(\mathscr{W})$. Hence, we satisfy the second condition of **(IP1)**:

$$\mathrm{Inv}(\mathscr{W}) \subset \mathrm{int}(\mathscr{W}) = \mathrm{int}(\mathscr{W} \setminus \mathscr{W}_-).$$

□

Lemma 4.1.13. *If $(\mathscr{W}, \mathscr{W}_-)$ is a compact Gromoll-Meyer pair for $S \subset \mathrm{Crit}(f)$, then $(\mathscr{W}, \mathscr{W}_-)$ satisfies* ***(IP2)****, i.e. $\mathscr{W}_-$ is an exit set of $\mathscr{W}$ with respect to the gradient flow σ as defined in eq. (4.2). In other words, for any $x \in \mathscr{W}$ such that $\sigma(t,x) \notin \mathscr{W}$ for some $t > 0$, there exists some $t_0 \in [0,t]$, such that $\sigma(t_0,x) \in \mathscr{W}_-$ and $\sigma([0,t_0],x) \subset \mathscr{W}$.*

Proof. Consder $x \in \mathscr{W}$ such that $\sigma(t,x) \notin \mathscr{W}$ for some $t > 0$. Let $t_0 = \max\{s \in [0,t] \ : \ \sigma(s,x) \in \mathscr{W}\}$. By definition, $\sigma(t_0,x) \in \partial W$. Let $\epsilon = \sup\{s \in [0,\infty) \ : \ \sigma(s,\sigma(t_0,x)) \in \mathscr{W}\}$. Note that since $\mathscr{W}$ is compact the supremum is attained. Suppose $\epsilon > 0$. From the definition of t_0, we must have $t_0 + \epsilon > t$, else $\epsilon = 0$. If $t_0 + \epsilon > t$, then the mean value property of $\mathscr{W}$ **(GM4)** implies $\sigma([t_0, t_0+\epsilon],x) \subset \mathscr{W}$. Since $t \in (t_0, t_0+\epsilon)$, we have $\sigma(t,x) \in \mathscr{W}$, which is a contradiction. Therefore, $\epsilon = 0$ and $\sigma(t_0,x) \in \mathscr{W}_-$. Furthermore, from the mean value property of $\mathscr{W}$ **(GM4)**, $\sigma([0,t_0],x) \in \mathscr{W}$. Therefore, $\mathscr{W}_-$ is an exit set for $\mathscr{W}$ with respect to the gradent flow σ. □

Lemma 4.1.14. *If $(\mathscr{W}, \mathscr{W}_-)$ is a compact Gromoll-Meyer pair for $S \subset \mathrm{Crit}(f)$, then $(\mathscr{W}, \mathscr{W}_-)$ satisfies* ***(IP3)****, i.e. for any $x \in \mathscr{W}_-$,*

$$\sigma([0,t],x) \in \mathscr{W} \implies \sigma([0,t],x) \in \mathscr{W}_-.$$

Proof. As $\mathscr{W}_- \cap \mathrm{Crit}(f) = \emptyset$ (**(GM1)** and **(GM2)**) and σ intersects $\mathscr{W}_-$ transversely **(GM5)**, for any $x \in \mathscr{W}_-$, we have:

$$\sigma([0,t],x) \in \mathscr{W} \iff t = 0 \implies \sigma([0,t],x) \in \mathscr{W}_-.$$

□

Proposition 4.1.15. *Let $(\mathscr{W}, \mathscr{W}_-)$ be a compact Gromoll-Meyer pair for $S \subset \mathrm{Crit}(f)$. Then $(\mathscr{W}, \mathscr{W}_-)$ are an index pair for the maximal isolated invariant set $\mathrm{Inv}(\mathscr{W})$ contained in $\mathscr{W}$ of the gradient flow σ as defined in eq. (4.2).*

Proof. We have shown that a compact Gromoll-Meyer pair for S satisfies the conditions **(IP1)** to **(IP3)** for being an index pair for S in Lemmas 4.1.12 to 4.1.14 respectively. □

4.2 Constructing Gromoll-Meyer Pairs

We now focus on a compact dynamically isolated critical set S with dynamically isolating triplet $(\mathcal{O}, \alpha, \beta)$ where $\mathcal{O}$ is compact. We present a construction of Gromoll-Meyer pairs for S, which we denote as $(\mathscr{M}, \mathscr{M}_-)$ to distinguish them from a general Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$. While the pair $\left(f^{-1}[\alpha, \beta], f^{-1}(\alpha)\right)$ is a convenient choice for a Gromoll-Meyer pair, neither of these sets are compact if M is not compact; for example, take $M = \mathbb{R}^2$ and the saddle point of $f = x^2 - y^2$ at the origin.

4.2.1 Constraints on the Structure of Gromoll-Meyer Pairs

Before we proceed with our construction of $(\mathscr{M}, \mathscr{M}_-)$, we briefly discuss constraints on any Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$ that follow from the condition of its definition **(GM1)** to **(GM5)**, where $\mathscr{W}$ has a smooth structure. We consider the case where $\mathscr{W}$ is a *regular domain*, i.e. a codimension-0 submanifold with boundary closed in M. In particular, we show that the insistence on the transversality of the exit set $\mathscr{W}_-$ with respect to the gradient flow **(GM5)** implies we cannot model $\mathscr{W}$ as a regular domain homeomorphic to a closed ball. If we relax the transversality constraint on C^1 data of f and consider more general index pairs (N, N_-) that admit points in the exit set tangent to the flow, then N cannot be a regular domain homeomorphic to a closed ball without explicit C^2 constraints involving the Hessian of f.

Proposition 4.2.1. *Let (N, N_-) be an index pair for an isolated invariant set of the gradient flow σ, where σ is as defined in eq. (4.2). Suppose N is a regular domain with a defining function g, and $N_- \cap \mathrm{Crit}(f) = \emptyset$. If $N_- \cap \partial N$ is transverse to the gradient flow and ∂N is connected, then $N_- = \emptyset$ or $N_- \supseteq \partial N$.*

Proof. Let $\varphi : \partial N \to \mathbb{R}$ be the smooth function $\varphi(x) = \langle \nabla f(x), \nabla g(x) \rangle$ and $L = N_- \cap \partial N$. The exit set condition **(IP2)** implies

$$\varphi^{-1}(-\infty, 0] \supseteq L \supseteq \varphi^{-1}(-\infty, 0).$$

As $N_- \cap \mathrm{Crit}(f) = \emptyset$, we observe that $\varphi^{-1}(0)$ are precisely the points on L that are tangent to

the flow. If the exit set is transverse to the flow, i.e. $L \cap \varphi^{-1}(0) = \emptyset$, then

$$\varphi^{-1}(-\infty, 0) \supseteq L \supseteq \varphi^{-1}(-\infty, 0) \implies L = \varphi^{-1}(-\infty, 0).$$

Since $L = \varphi^{-1}(-\infty, 0)$ and φ is continuous, L is open. As N_- is compact, $L = N_- \cap \partial N$ is also closed in ∂N. Hence L is a clopen subset of ∂N.

We recall this elementary fact about clopen subsets: ∂N is connected if and only if the only clopen sets are the $\emptyset$ and ∂N. Thus, if ∂N is connected, then either N_- is empty or $L = N_- \cap \partial N = \partial N_-$. □

As a Gromoll-Meyer pair $(\mathscr{W}, \mathscr{W}_-)$ is an index pair where $\mathscr{W}_- \subseteq \partial\mathscr{W}_-$ and $\mathscr{W}_-$ is transverse to gradient flow, we obtain the following corollary.

Corollary 4.2.2. *Suppose $(\mathscr{W}, \mathscr{W}_-)$ is a Gromoll-Meyer pair for $S \subset \mathrm{Crit}(f)$. Suppose $\mathscr{W}$ is a regular domain. If $\partial\mathscr{W}$ is connected, then $\mathscr{W}_-$ is either empty or the entirety of $\partial\mathscr{W}$.*

Corollary 4.2.3. *Suppose a dynamically isolated critical set $S \subset \mathrm{Crit}(f)$ does not contain a local maximum or minimum of f. If $(\mathscr{W}, \mathscr{W}_-)$ is a Gromoll-Meyer pair for S, then $\mathscr{W}$ cannot be a regular domain with a connected boundary.*

This rules out modelling $\mathscr{W}$ as a *regular domain* with a connected boundary for any critical set. In particular, these constraints prevent any submanifold with boundary homeomorphic to a closed ball as a candidate for $\mathscr{W}$ for manifold dimension > 1.

One might argue that the constraints of Gromoll-Meyer pairs are too restrictive and we can consider constructing index pairs outside the Gromoll-Meyer framework that only need to satisfy the weaker conditions of **(IP1)** to **(IP3)**. In particular, index pairs do not require the exit set to be necessarily transverse to the flow (unlike **(GM5)** for Gromoll-Meyer pairs). We recall that for any regular domain N, there exists a *defining function* $g \in C^\infty(M)$ and a regular value r of g such that $N = g^{-1}(-\infty, r]$ (see [Lee, 2012, Theorem 5.48]).

Proposition 4.2.4. *Let (N, N_-) be an index pair for an isolated invariant set of the gradient flow σ of $f \in C^2(M)$, where σ is as defined in eq. (4.2). Suppose N is a regular domain with a defining function g, and $N_- \cap \mathrm{Crit}(f) = \emptyset$. If $x \in N_- \cap \partial N$ is tangent to $\sigma(\mathbb{R}, x)$, then there exists some*

point $y \in N_- \cap \sigma(\mathbb{R}_{\geq 0}, x)$ *where*

$$H_f(\nabla f, \nabla g) + H_g(\nabla f, \nabla f) \geq 0, \tag{4.9}$$

at y *and the inequality is strict if* $y \in \text{int}(N)$.

Proof. We follow the notation adopted in Proposition 4.2.1. Let us consider the flow line $\sigma(t, x)$ through $x \in L$. Let $G = g(\sigma(\cdot, x)) : \mathbb{R} \to \mathbb{R}$. As $L \cap \text{Crit}(f) = \emptyset$, the flow is tangent to ∂N if and only if

$$\left.\frac{\mathrm{d}G}{\mathrm{d}t}\right|_{t=0} = \varphi(x) = \langle \nabla f(x), \nabla g(x) \rangle = 0.$$

If $G([0, \epsilon)) \geq r$ for some sufficiently small $\epsilon > 0$, the positive invariance property **(IP3)** is satsfied where σ is tangent to L only if

$$\left.\frac{\mathrm{d}^2 G}{\mathrm{d}t^2}\right|_{t=0} = \langle \nabla f(x), \nabla \varphi(x) \rangle = H_f(\nabla f, \nabla g)\Big|_x + H_g(\nabla f, \nabla f)\Big|_x \geq 0. \tag{4.10}$$

where H_f and H_g are the Hessians of f and g respectively. In the case where the inequality is strict, $G(0) = r$ is a local minimum of G, i.e. $\sigma((-\epsilon, \epsilon), 0) \cap N = x$ for a sufficiently small $\epsilon > 0$. If $G(0) = r$ is not a local minimum of G, then x must be an inflection point where $\left.\frac{\mathrm{d}^2 G}{\mathrm{d}t^2}\right|_{t=0} = 0$.

We now consider the alternative case where $G([0, \epsilon]) < r$ for some sufficiently small $\epsilon > 0$, i.e. $\sigma([0, \epsilon), x) \subset \text{int}(N)$. As N_- is compact and we have assumed there are no critical points of f in N_-, there is some t such that $\sigma(t, x) \notin N$: all that flows into N_- must leave N_-. Furthermore, positive invariance **(IP3)** implies there is some $t_0 \in [0, t]$ such that $\sigma(t_0, x) \in \partial N$ and $\sigma([0, t_0], x) \subset N_-$. Since $G(t_0) = G(0) = r$ and $G([0, \epsilon]) < r$, there is at least one local minimum of G at some $t_1 \in (0, t_0)$ where $\sigma(t_1, x) \subset \text{int}(N) \cap N_-$. Since $\sigma(t_1, x)$ is a local minimum of G, we must have $\frac{\mathrm{d}^2 G}{\mathrm{d}t^2} \geq 0$. □

Corollary 4.2.5. *Suppose a dynamically isolated critical set* $S \subset \text{Crit}(f)$ *does not contain a local maximum or minimum of* f. *Let* (N, N_-) *be an index pair for the minimal isolated invariant set of the gradient flow containing* S. *If* N *is a regular domain with a connected boundary, and* $N_- \cap \text{Crit}(f) = \emptyset$, *then* N *must satisfy the constraints as specified in eq.* (4.10).

Proof. The notation below follows that employed in the proof of Proposition 4.2.4. If $S \subset$

$\mathrm{Crit}(f)$ does not contain a local maximum or minimum of f, then N_- is neither empty, nor does it contain ∂N. If N is a regular domain with a connected boundary, then it follows from Proposition 4.2.1 that $N_- \cap \partial N$ contains a point x tangent to the gradient flow. Proposition 4.2.4 then implies there must be some point in N_- that satisfies the constraint eq. (4.9) on the Hessians of f and g. □

4.2.2 Constructing Gromoll-Meyer Pairs

We first recall a fundamental result of Morse theory. For a non-degenerate critical point $p = \mathrm{Crit}(f) \cap f^{-1}[-\eta, \eta]$ with index $i = \mathscr{I}(f, p)$, the sublevel set $f^{-1}(-\infty, \eta]$ is homeomorphic to $f^{-1}(-\infty, -\eta]$ with an i-handle D^i attached. In particular, $f^{-1}(-\infty, -\eta] \sqcup D^i$ can be smoothed into a smooth manifold with boundary $H \supset f^{-1}(-\infty, -\eta]$. In the construction of H in standard Morse theory texts such as [Milnor et al., 1973, p.14-19] or [Nicolaescu, 2011, p. 34-42], $H \setminus f^{-1}(-\infty, -\eta]$ can be bound in an arbitrarily small neighbourhood of p. One can show that $\mathrm{cl}\big(H \setminus f^{-1}(-\infty, -\eta]\big)$ and $\mathrm{cl}\big(H \setminus f^{-1}(-\infty, -\eta]\big) \cap f^{-1}(-\eta)$ is a Gromoll-Meyer pair for p.

However, the standard construction of H relies on coordinates that diagonalise the Hessian of f at p, which may be hard to obtain in applied contexts. Furthermore, it is specfic to non-degenerate critical points. Our construction of Gromoll-Meyer pairs $(\mathscr{M}, \mathscr{M}_-)$ generalises the standard construction of H to a coordinate-agnostic procedure that can be applied to compact critical sets. Furthermore, our procedure does not require the knowledge of second derivatives, only a lower bound on the norm of the gradients of f on an annular neighbourhood of the compact critical set S.

We now describe the conditions under which our construction is a valid Gromoll-Meyer pair. For a compact critical set $S \subset \mathrm{Crit}(f)$ with an isolating triplet $(\mathcal{O}, \alpha, \beta)$, we require the existence of a smooth function $g \in C^\infty(M)$ and constants $r_0 < \frac{r_0+r_1}{2} \leq r < r_1$ in the image of g that satisfy the conditions below:

(G1) $g^{-1}(-\infty, r_1]$ is compact;

(G2) $\mathcal{O} \subset g^{-1}(-\infty, r_0]$;

(G3) $g^{-1}(-\infty, r_0] \cap \mathrm{Crit}(f) = g^{-1}(-\infty, r_1] \cap \mathrm{Crit}(f) = S$;

(G4) g is smooth and $\mathrm{Crit}(g) \cap g^{-1}[r_0, r_1] = \emptyset$;

(G5) $\beta - \alpha$ can be taken to be sufficiently small, such that

$$2(\beta - \alpha) < \frac{r_1 - r_0}{2} \inf_{x \in g^{-1}[r_0, r_1]} \frac{\|\nabla f\|}{\|\nabla g\|}$$

Note that $\|\nabla g\|$ is bounded away from zero by **(G4)**.

(G6) η is a regular value of f and r is a regular value of the restriction of g to $f^{-1}(\eta)$.

These conditions can be satisfied easily in common scenarios. If M is smoothly embedded in $\mathbb{R}^d$, we can take g to be $\|x - p\|$, the Euclidean distance from $p \in M$ to some $x \in \mathbb{R}^d \setminus M$. In this case, g is smooth everywhere on M, and $\nabla g(p)$ is simply the cosine of the angle between $T_p M$ and $x - p$. Thus, $\|\nabla g\| \leq 1$ and the following condition

$$2(\beta - \alpha) < \frac{r_1 - r_0}{2} \inf_{y \in B_{r_1}[x] \setminus B_{r_0}(x)} \|\nabla f\|$$

satisfies **(G5)**. If S is sufficiently small, such that **(G2)** is satisfied for

$$0 < r_0 < r_1 < \inf_{y \in B_{r_1}(x) \cap M} \tau_M(y),$$

then there is only one critical point of g in $B_{r_1}(x)$ which is the minimum of g by Lemma 2.2.20. Since **(G2)** implies $B_{r_0}(p)$ contains a non-empty open set of M, the value of the minimum must be strictly less than r_0. Therefore, **(G4)** is satisfied. Furthermore, if S is connected, then f is constant on S, and we can take $\beta - \alpha$ to be an arbitrarily small positive number so that **(G5)** is satisfied. Finally, **(G6)** is satisfied for r on a dense subset of (r_0, r_1) due to Sard's theorem: since η is a regular value of f, the level set $f^{-1}(\eta)$ is a submanifold of M. As the restriction of g to $f^{-1}(\eta)$ is a smooth function, we have a dense subset of regular values for the restriction of g by Sard's theorem, satisfying **(G6)**.

We construct a compact $\mathscr{M}$ by considering a smooth, local perturbation $h : M \to \mathbb{R}$ of f so that

$$\mathscr{M} = h^{-1}(-\infty, \eta_0] \cap f^{-1}[\eta, \infty) \quad \text{and} \quad \mathscr{M}_- = f^{-1}(\eta) \cap g^{-1}(-\infty, r]$$

are a Gromoll-Meyer pair for S. In our construction, we ensure that the diameter of $\mathscr{M}$ is

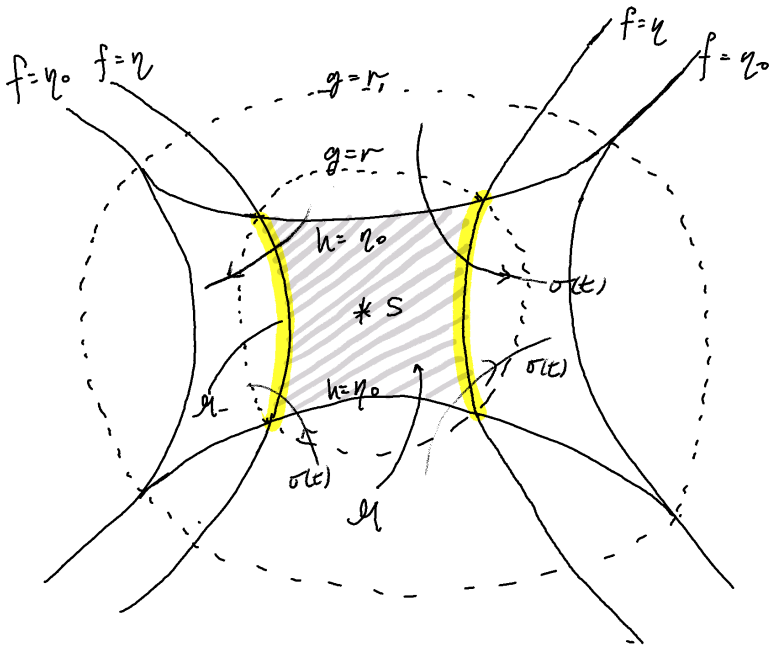

Figure 4.1: An illustration of our construction of Gromoll-Meyer pairs $(\mathscr{M}, \mathscr{M}_-)$. The grey thatched region is $\mathscr{M}$, and the yellow highlighted portion of the level set $f^{-1}(\eta)$ is the exit set $\mathscr{M}_-$.

finite and explicitly bounded above. We now use g and the constants r_1 and r_0 to construct h. First, consider the smoothed step function $q : \mathbb{R} \to \mathbb{R}_{\geq 0}$ which we show in Figure 4.2

$$q(t) = \begin{cases} q_0 & t < r_0 \\ q_0\left(1 + \exp\left(\frac{r_1 - r_0}{r_1 - t} + \frac{r_1 - r_0}{r_0 - t}\right)\right)^{-1} & t \in [r_0, r_1] \\ 0 & t > r_1 \end{cases}. \tag{4.11}$$

Choosing q_0 to be in a positive constant, we observe that:

(Q1) q is compactly supported on $[0, r_1]$;

(Q2) q constant and maximal on $[0, r_0]$; and

(Q3) q is strictly monotonically decreasing on (r_0, r_1).

(Q4) $q'(t)$ is most negative at $t = \frac{r_1 + r_0}{2}$ where $q'\left(\frac{r_1 + r_0}{2}\right) = -\frac{2q_0}{r_1 - r_0}$.

(Q5) $q\left(\frac{r_1 + r_0}{2}\right) = \frac{q_0}{2}$.

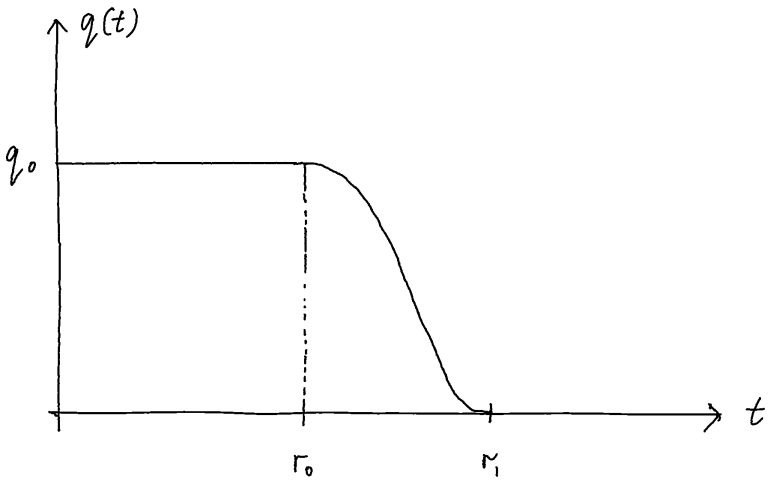

Figure 4.2: A graph of the function $q(t)$ in eq. (4.11).

As we have assumed **(G5)**, we henceforth take

$$2(\beta - \alpha) < q_0 < \frac{r_1 - r_0}{2} \inf_{x \in g^{-1}[r_0, r_1]} \frac{\|\nabla f\|}{\|\nabla g\|}. \tag{4.12}$$

Lemma 4.2.6. *Consider f and g in $C^\infty(M)$ such that for $S \in \mathrm{Crit}(f)$ with an isolating triplet $(\mathcal{O}, \alpha, \beta)$, g satisfies conditions **(G1)** to **(G6)**. For*

$$h(x) = f(g) - q(g(x)) \tag{4.13}$$

where q is as defined in eqs. (4.11) and (4.12), we have

(H1) $f(x) \geq h(x)$;

(H2) $g(x) \leq r_0 \iff h(x) = f(x) - q(0)$;

(H3) $g(x) \geq r_1 \iff h(x) = f(x)$; *and*

(H4) $\langle \nabla h, \nabla f \rangle \geq 0$ *with equality only attained on* $\mathrm{Crit}(f)$.

(H5) $\mathrm{Crit}(f) = \mathrm{Crit}(h)$.

Proof. The desired verifications are below.

(H1): Since q is non-negative, $h(x) = f(x) - q(g(x)) \leq f(x)$;

(H2): If $g(x) \leq r_0$, then $q(g(x)) = q(0)$ and $h(x) = f(x) - q(0)$, due to **(Q2)**.

Conversely, if $f(x) - h(x) = q(0)$, then x can only belong in $g^{-1}[0, r_0]$, since $f(x) - h(x) = q(g(x)) < q(0)$ elsewhere by **(Q2)**.

(H3): If $g(x) \geq r_1$, then $q(g(x)) = 0$ and $h(x) = f(x)$.

Conversely, if $h(x) = f(x)$, then x can only belong in $g^{-1}[r_1, \infty)$, since $f(x) - h(x) = q(g(x)) > 0$ elsewhere by **(Q1)**.

(H4): Outside the annulus $g^{-1}[r_0, r_1]$, **(Q1)** and **(Q2)** imply that $q' = 0$, and therefore $\nabla h = \nabla f$.

We now consider $x \in g^{-1}[r_0, r_1]$. Substituting the definition of h eq. (4.13),

$$\begin{aligned}
\langle \mathrm{d}h\,, \nabla f \rangle &= \langle \mathrm{d}f\,, \nabla f \rangle - q' \langle \mathrm{d}g\,, \nabla f \rangle \\
&\geq \|\nabla f\| \|\nabla g\| \left(\frac{\|\nabla f\|}{\|\nabla g\|} - |q'| \right) && \text{Cauchy-Schwarz, } \textbf{(G4)} \\
&\geq \|\nabla f\| \|\nabla g\| \left(\frac{\|\nabla f\|}{\|\nabla g\|} - \frac{2q_0}{r_1 - r_0} \right) && \textbf{(Q4)} \\
&> 0 && \textbf{(G3)} \text{ and } \textbf{(G4)} \text{ and eq. (4.12)}
\end{aligned}$$

(H5): **(H4)** implies $\mathrm{d}h \neq 0$ on the annulus $g^{-1}[r_0, r_1]$ of the annulus. As $\mathrm{d}h = \mathrm{d}f$ where $g \leq r_0$ or $g \geq r_1$, we establish that $\mathrm{Crit}(f) = \mathrm{Crit}(h)$

□

We now explicitly construct $\mathscr{M}$ in terms of a sublevel and superlevel set of h and f respectively.

Lemma 4.2.7. *For a critical set $S \subset \mathrm{Crit}(f)$ with a dynamically isolating triplet $(\mathcal{O}, \alpha, \beta)$, Assume the conditions of Lemma 4.2.6 are satisfied. For choices of constants*

$$\eta \in (\beta - \frac{q_0}{2}, \alpha) \quad \textit{and} \quad \eta_0 = \eta - q(r) \tag{4.14}$$

and h as constructed in eq. (4.13) of Lemma 4.2.6, the subset

$$\mathscr{M} = h^{-1}(-\infty, \eta_0] \cap f^{-1}[\eta, \infty) \tag{4.15}$$

is non-empty and closed; furthermore, we have

$$S \subset f^{-1}(\eta, \infty) \cap h^{-1}(-\infty, \eta_0) \subset \mathscr{M} \subset g^{-1}(-\infty, r].$$

Proof. We first confirm that $\alpha > \beta - \frac{q_0}{2}$: this follows from eq. (4.12). As $\mathscr{M}$ is the intersection of two closed sets it is itself closed. From the definition of $\mathscr{M}$ (by eq. (4.15)), for any $x \in \mathscr{M}$,

$$\eta - q(g(x)) \leq f(x) - q(g(x)) = h(x) \leq \eta_0 \implies q(g(x)) \geq \eta - \eta_0 = q(r).$$

Since q is monotonically decreasing on $(r_0, r_1) \ni r$ (**(Q3)**), we deduce that $\mathscr{M} \subset g^{-1}(-\infty, r]$.

Because $(\mathcal{O}, \alpha, \beta)$ is a dynamically isolating triplet for S, we have $S \subset f^{-1}(\alpha, \infty)$. As we have chosen η such that $\eta < \alpha$ (eq. (4.14)), we see that $S \subset f^{-1}(\eta, \infty)$.

Next, since we have assumed $S \subset g^{-1}(-\infty, r_0]$ **(G2)**, we can apply **(H2)** and deduce that $S \subset h^{-1}(-\infty, \beta - q_0)$. Since we have assumed $r \geq \frac{r_1 + r_0}{2}$, **(Q3)** and **(Q5)** imply $\eta_0 = \eta - q(r) \geq \eta - \frac{q_0}{2}$. Since we have taken $\eta > \beta - \frac{q_0}{2}$, we deduce that $\eta_0 > \beta - q_0$. Therefore, $S \subset h^{-1}(-\infty, \eta_0)$.

Thus, we conclude that $S \in h^{-1}(-\infty, \eta_0) \cap f^{-1}(\eta, \infty)$. □

Lemma 4.2.7 shows that $\mathscr{M}$ satisfies condition **(GM1)** for $(\mathscr{M}, \mathscr{M}_-)$ being a Gromoll-Meyer pair for S. The following two lemmas show that $\mathscr{M}$ is compact and satisfies **(GM2)**.

Lemma 4.2.8. *Assume the conditions of Lemma 4.2.7 are satisfied and consider h as constructed in eq. (4.13) of Lemma 4.2.6. Then*

$$\mathscr{M} \subset g^{-1}(-\infty, r] \tag{4.16}$$

$$\mathscr{M}_- = h^{-1}(-\infty, \eta_0] \cap f^{-1}(\eta) = f^{-1}(\eta) \cap g^{-1}(-\infty, r] \tag{4.17}$$

$$\mathscr{M}_+ = h^{-1}(\eta_0) \cap f^{-1}[\eta, \infty) = h^{-1}(\eta_0) \cap g^{-1}(-\infty, r] \tag{4.18}$$

Proof. Since q is a montonically non-increasing function

$$q(g(x)) \geq q(r) \implies g(x) \leq r.$$

Therefore, $\mathscr{M} \subset g^{-1}(-\infty, r]$.

We now prove eq. (4.17). For $x \in f^{-1}(\eta)$, we have $h(x) = \eta - q(g(x))$. Since q is a montonically non-increasing function

$$g(x) \leq r \implies h(x) \leq \eta - (\eta - \eta_0) = \eta_0.$$

Thus, $f^{-1}(\eta) \cap g^{-1}(-\infty, r] \subset f^{-1}(\eta) \cap h^{-1}(-\infty, \eta_0]$. Since we have shown Lemma 4.2.7 that $f^{-1}(\eta) \cap h^{-1}(-\infty, \eta_0] \subset \mathscr{M} \subset g^{-1}(-\infty, r]$, we have

$$f^{-1}(\eta) \cap h^{-1}(-\infty, \eta_0] = f^{-1}(\eta) \cap g^{-1}(-\infty, r].$$

The same argument can be applied to prove eq. (4.18). □

Lemma 4.2.9. *Assume the conditions of Lemma 4.2.7 are satisfied. Then* $\mathscr{M} \cap \mathrm{Crit}(f) = S$.

Proof. We have shown in lemma 4.2.7 that $S \subset \mathscr{M}$. Furthermore, we have shown in lemma 4.2.8 that $\mathscr{M} \subset g^{-1}(-\infty, r]$. Thus,

$$S \subset \mathscr{M} \cap \mathrm{Crit}(f) \subset g^{-1}(-\infty, r] \cap \mathrm{Crit}(f).$$

Since we have assumed in **(G3)** that $g^{-1}(-\infty, r] \cap \mathrm{Crit}(f) = S$, we conclude that $S = \mathscr{M} \cap \mathrm{Crit}(f)$. □

Lemma 4.2.10. *Assume the conditions of Lemma 4.2.7 are satisfied. Then* $\mathscr{M} \subset f^{-1}[\eta, \eta_0]]$, *where* $\mathscr{M}$ *and constants* η *and* η_0 *are as defined in eq. (4.14) and eq. (4.15) in Lemma 4.2.7 respectively.*

Proof. By definition eq. (4.15), $\mathscr{M} \subset f^{-1}[\eta, \infty)$. Since $h \leq f$ **(H1)**, the fact that $\mathscr{M} \subset h^{-1}(-\infty, \eta_0]$ implies $\mathscr{M} \subset f^{-1}(-\infty, \eta_0]$. □

Lemma 4.2.11. *Let* $\sigma : \mathbb{R} \times M \to M$ *be the gradient flow along* $\dot{\sigma} = -\nabla f$ *as described in eq. (4.2). Then* $\mathscr{M}$, *as defined in eq. (4.15) in Lemma 4.2.8, satisfies the mean value property* ***(GM4)****:*

$$\sigma(t_1, x), \sigma(t_2, x) \in \mathscr{M} \implies \sigma(t, x) \in \mathscr{M} \ \forall\ t \in [t_1, t_2].$$

Proof. Suppose $\sigma(t_1, x), \sigma(t_2, x) \in \mathscr{M}$, where $t_2 > t_1$. Recall from eq. (4.2) that σ flows along

the gradient $-\nabla f$. Hence, along the flow $\sigma(t,x)$,

$$\frac{\mathrm{d}f}{\mathrm{d}t} = -\|\mathrm{d}f\|^2 \leq 0.$$

Thus, f is non-increasing along the flow. Then for all $t \in [t_1, t_2]$,

$$f(\sigma(t_1,x)) \geq f(\sigma(t,x)) \geq f(\sigma(t_2,x)) \geq \eta.$$

Since $\langle \mathrm{d}h\,, -\nabla f\rangle \leq 0$ **(H4)**, h is also non-increasing along the flow:

$$\frac{\mathrm{d}h}{\mathrm{d}t} = \langle \mathrm{d}h\,, -\nabla f\rangle \leq 0$$

Thus for all $t \in [t_1, t_2]$,

$$\eta_0 \geq h(\sigma(t_1,x)) \geq h(\sigma(t,x)) \geq h(\sigma(t_2,x)),$$

and $\sigma(t,x) \in f^{-1}[\eta,\infty) \cap h^{-1}(-\infty,\eta_0] = \mathscr{M}$ by Equation (4.13). □

Lemma 4.2.12. *Assume the conditions of Lemma 4.2.7 are satisfied. Then $\mathscr{M}_\pm$ as defined in eq. (4.17) and eq. (4.18) in Lemma 4.2.8 respectively are smooth submanifolds of M with a common boundary $f^{-1}(\eta) \cap g^{-1}(r)$. Furthermore, $\mathscr{M}_\pm$ are transverse to the gradient flow σ as defined in eq. (4.2).*

Proof. Since we have assumed in **(G6)** that η is a regular value of f, the level set $f^{-1}(\eta)$ is a submanifold of M. Furthermore, we have assumed in **(G6)** that r is a regular value of the restriction of g to $f^{-1}(\eta)$. Thus, the sublevel set of g restriction to $f^{-1}(\eta)$, which is precisely $\mathscr{M}_-$, is a submanifold with boundary $g^{-1}(r) \cap f^{-1}(\eta)$.

For $\mathscr{M}_+$, let us consider h restricted to submanifold with boundary $g^{-1}(-\infty, r]$. If η_0 is a regular value of both h restricted to the boundary $g^{-1}(r)$ and interior $g^{-1}(-\infty, r)$ respectively, then $h^{-1}(\eta_0) \cap g^{-1}(-\infty, r]$ is a submanifold with boundary. Since $\mathrm{Crit}(f) = \mathrm{Crit}(h)$ **(H5)**, and the only critical set of f which is S (lemma 4.2.9) is a subset of $h^{-1}(-\infty, \eta_0)$ (lemma 4.2.7), we deduce that $\mathrm{Crit}(h) \cap \mathscr{M}_+ = \emptyset$ and η_0 is a regular value of h restricted to $g^{-1}(-\infty, r)$. We can also deduce that η_0 can only be a critical value of h restricted to the boundary $g^{-1}(r)$ if and only if

$$\mathrm{d}h = \mathrm{d}f - q'\,\mathrm{d}g = \lambda\,\mathrm{d}g \iff \mathrm{d}f = (q' + \lambda)\,\mathrm{d}g$$

in T_pM on $p \in h^{-1}(\eta_0) \cap g^{-1}(r)$ for some $\lambda \in \mathbb{R} \setminus 0$. However, by eq. (4.11),

$$h = \eta_0 \quad \text{and} \quad g = r \implies f = h + q(r) = \eta_0 + \eta - \eta_0 = \eta$$
$$f = \eta_0 \quad \text{and} \quad g = r \implies h = f - q(r) = \eta - (\eta - \eta_0) = \eta_0$$

we have $h^{-1}(\eta_0) \cap g^{-1}(r) = f^{-1}(\eta_0) \cap g^{-1}(r)$, which is the common boundary of $\mathscr{M}_\pm$. Since f is a regular value restricted to $g^{-1}(r)$, this implies $\mathrm{d}f$ and $\mathrm{d}g$ are linearly independent on the common boundary and $\mathrm{d}f \neq (q' + \lambda)\,\mathrm{d}g$. Thus, η_0 cannot be a critical value of h restricted to $g^{-1}(r)$. Therefore η_0 is a regular value of h restricted to $g^{-1}(-\infty, r]$ and $\mathscr{M}_+ = h^{-1}(-\infty, \eta_0] \cap g^{-1}(-\infty, r]$ is a submanifold with boundary.

We now note that the tangent spaces of $\mathscr{M}_-$ and $\mathscr{M}_+$ are orthogonal complements of ∇f and ∇h respectively. Thus, for a flow σ to be transverse to $\mathscr{M}_-$ and $\mathscr{M}_+$, the velocity vector $\dot{\sigma}$ of the flow must have a non-zero projection onto ∇f and ∇h respectively. As the velocity vector of the gradient flow σ as defined in eq. (4.2) is aligned with $-\nabla f$, the flow σ is patently transverse to $\mathscr{M}_-$. Since $\mathscr{M}_+ \cap \mathrm{Crit}(h) = \emptyset$, and $\langle \mathrm{d}h, \nabla f \rangle > 0$ where f and therefore h are not critical **(H4)**, σ is also transverse to $\mathscr{M}_+$. □

Proposition 4.2.13. *Assuming **(G1)** to **(G6)**, then $(\mathscr{M}, \mathscr{M}_-)$ and $(\mathscr{M}, \mathscr{M}_+)$ are Gromoll-Meyer pairs for $S \subset \mathrm{Crit}(f)$ and $S \subset \mathrm{Crit}(-f)$ respectively, where $\mathscr{M}$ is as defined in eq. (4.15) and $\mathscr{M}_\pm$ are as defined in eqs. (4.17) and (4.18) respectively.*

Proof. The preceding lemmas show that $(\mathscr{M}, \mathscr{M}_-)$ and $(\mathscr{M}, \mathscr{M}_+)$ satisfy the conditions **(GM1)** to **(GM5)** for being Gromoll-Meyer pairs. **(GM1)** is satisfied by lemma 4.2.7, **(GM2)** is satisfied by lemma 4.2.9, and **(GM3)** for $(\mathscr{M}, \mathscr{M}_-)$ and $(\mathscr{M}, \mathscr{M}_+)$ is satisfied by the lower and upper bounds of f on $\mathscr{M}$ in lemma 4.2.10. As swapping f for $-f$ simply reverses the direction of gradient flow σ, **(GM4)** holds for both $(\mathscr{M}, \mathscr{M}_-)$ and $(\mathscr{M}, \mathscr{M}_+)$ due to lemma 4.2.11. As we have shown that $\mathscr{M}_\pm$ are submanifolds with boundary transverse to the (anti)gradient flow in lemma 4.2.12, **(GM5)** are both satisfied. □

4.3 Inference of Morse Indices from Finite Data

We now turn to the practical considerations of inferring the Morse indices of compact critical sets. Recall we are concerned with critical sets of smooth functions on smooth Riemannian

manifold M, which we properly embed in some Euclidean space $\mathbb{R}^d$.

In Section 4.2, we have shown that we can compute the Morse index of a critical set S of $f \in C^\infty(M)$ using the relative homology of a Gromoll-Meyer pair $(\mathscr{M}, \mathscr{M}_-)$. Furthermore, a Gromoll-Meyer pair $(\mathscr{M}, \mathscr{M}_-)$ of S can be modelled as intersections of sublevel and level sets. Recall we have explicitly constructed Gromoll-Meyer pairs for compact critical sets that take the form of

$$\mathscr{M} = f^{-1}[0,\infty) \cap h^{-1}(-\infty, \eta_0] \quad \text{and} \quad \mathscr{M}_- = f^{-1}(0) \cap h^{-1}(-\infty, \eta_0].$$

where we have set η in eq. (4.15) to zero without loss of generality. In Section 3.6, we have shown that if $(\mathscr{M}, \mathscr{M}_-)$ are compact and 0 is a regular value of $F : M \to \mathbb{R}^2$ defined by $F(x) = (f(x), h(x))$, then $(\mathscr{M}, \mathscr{M}_-)$ are subsets with positive reach in the ambient Euclidean space. We show that our construction satisfies these conditions.

Proposition 4.3.1. *Assume the conditions of Lemma 4.2.7 are satisfied. Then $\mathscr{M}$ and $\mathscr{M}_\pm$ as defined in eq. (4.15), eq. (4.17) and eq. (4.18) in Lemmas 4.2.7 and 4.2.8 respectively have positive reach.*

Proof. We show that $\mathscr{M}$ and $\mathscr{M}_\pm$ are regular in the sense of Definition 3.6.1. As such, we can apply Proposition 3.6.6 which implies regular intersections of level sets and sublevel sets have postive reach. We first partition $\mathscr{M} = f^{-1}[0,\infty) \cap h^{-1}(-\infty, \eta_0]$ into the following subsets

$$f^{-1}(0,\infty) \cap h^{-1}(-\infty, \eta_0), \quad f^{-1}(0,\infty) \cap h^{-1}(\eta_0), \quad f^{-1}(0) \cap h^{-1}(-\infty, \eta_0) \quad \text{and} \quad f^{-1}(0) \cap h^{-1}(\eta_0).$$

To show that $\mathscr{M}$ and $\mathscr{M}_\pm$ are regular, we need to check that $\nabla h \neq 0$ and $\nabla f \neq 0$ on $f^{-1}(0,\infty) \cap h^{-1}(\eta_0)$, as well as $F = (f,h) : M \to \mathbb{R}^2$ being regular on $f^{-1}(0) \cap h^{-1}(\eta_0)$. This is shown to be true in the proof of Lemma 4.2.12 where $\mathscr{M}_\pm$ are shown to be submanifolds with boundary. □

As such, the homology of $\mathscr{M}$ and $\mathscr{M}_\pm$ can be reliably using finite point samples $(\mathbb{X}, \mathbb{X}_-) \subset \mathbb{R}^d$. Suppose we have $(\mathbb{X}, \mathbb{X}_-)$ satisfy

$$\max\left(d_H(\mathbb{X}, \mathscr{M}), d_H(\mathbb{X}_-, \mathscr{M}_-)\right) \leq \alpha \quad \text{and} \quad \min\left(\tau_{\mathscr{M}}, \tau_{\mathscr{M}_-}\right) \geq \tau > 0 \tag{4.19}$$

where d_H denotes the Hausdorff distance between subsets of $\mathbb{R}^d$. We defer the discussion

of how one might generate a point sample satisfying these conditions in Section 4.3.2. If $\alpha > 0$ is sufficiently small relative to τ, then there is some open interval (t_-, t_+) of thickening parameter t such that there is a deformation retract of $\mathbb{X}^t$ and $\mathbb{X}^t_-$ to $\mathscr{M}$ and $\mathscr{M}_-$ respectively for all $t \in (t_-, t_+)$. Using the long exact sequences of the pairs $(\mathscr{M}, \mathscr{M}_-)$ and $(\mathbb{X}^t, \mathbb{X}^t_-)$, we can show that the inclusions of $\mathscr{M}$ and $\mathscr{M}_-$ into $\mathbb{X}^t$ and $\mathbb{X}^t_-$ induce an isomorphism

$$H_\bullet(\mathscr{M}, \mathscr{M}_-) \cong H_\bullet(\mathbb{X}^t, \mathbb{X}^t_-)$$

for all $t \in (t_-, t_+)$. Thus, given some lower bound of the reach, we can recover the Morse index of a compact critical set from some non-trivial measure of the family of pairs $(\mathbb{X}^t, \mathbb{X}^t_-)$.

4.3.1 Persistent Relative Homology Modules

Given a pair of point clouds $\mathbb{X}_- \subset \mathbb{X}$ such that $d_H(\mathscr{M}, \mathbb{X}) \leq \alpha$ and $d_H(\mathscr{M}_-, \mathbb{X}_-) \leq \alpha$, we would like to compare the persistent relative homology module $\check{\mathrm{H}}_i(\mathbb{X}, \mathbb{X}_-)$ derived from thickening $\mathbb{X}$ and $\mathbb{X}_-$, with that of the Gromoll-Meyer pairs $\check{\mathrm{H}}_i(\mathscr{M}, \mathscr{M}_-)$. In Proposition 2.3.20, we have shown that so long as α is less than a quarter of the minimum of the reaches of $\mathscr{M}$ and $\mathscr{M}_-$, we can recover the rank of the Morse index from $\check{\mathrm{H}}_i(\mathbb{X}, \mathbb{X}_-)$.

4.3.2 Sampling

We now assume we have a finite, random uniform i.i.d sample $\mathbb{Y}$ of some sufficiently large compact regular domain N containing $\mathscr{M}$. Since we have chosen N to be a regular domain, Proposition 3.5.8 implies that as long as we draw sufficiently many uniform i.i.d. samples on N, we can guarantee these samples are δ-dense in N with high confidence. We show how we can choose subsets $\mathbb{X}$ and $\mathbb{X}_-$ of $\mathbb{Y}$ using function evaluations of f and h, such that $\mathbb{X}$ and $\mathbb{X}_-$ has bounded Hausdorff distances away from $\mathscr{M}$ and $\mathscr{M}_-$ respectively, under some mild assumptions on f and h.

We recall from Definition 2.1.1 that two subsets A and B of $\mathbb{R}^d$ have a Hausdorff distance of α if A is both dense in, and close to B. We first construct subsets $\mathscr{M}' \supset \mathscr{M}$ and $\mathscr{M}'_- \supset \mathscr{M}_-$ of M that can be made θ-close to $\mathscr{M}$ and $\mathscr{M}_-$ respectively, where θ can be made sufficiently small by construction. We defer the construction to and prove this point in Lemma 4.3.3. We then restrict the ambient point cloud $\mathbb{Y}$ to $\mathscr{M}'$ and $\mathscr{M}'_-$ to obtain $\mathbb{X}$ and $\mathbb{X}_-$. As $\mathbb{X}$ and $\mathbb{X}_-$ are

subsets of $\mathscr{M}'$ and $\mathscr{M}'_-$ respectively, they must also be θ-close to $\mathscr{M}$ and $\mathscr{M}_-$ respectively.

We then show in Lemma 4.3.2 that the $\mathscr{M}'$ and $\mathscr{M}'_-$ can be constructed such that $\mathbb{X}$ and $\mathbb{X}_-$ — the restriction of a δ-dense point sample $\mathbb{Y}$ to $\mathscr{M}'$ and $\mathscr{M}'_-$ — is δ-dense if δ is sufficiently small relative to parameters that define $\mathscr{M}'$ and $\mathscr{M}'_-$.

Thus, by enforcing that sufficiently many points are randomly sampled on N, and restricting the resultant set of points to subsets $\mathscr{M}'$ and $\mathscr{M}'_-$, we obtain point clouds $\mathbb{X}$ and $\mathbb{X}_-$ that we can make sufficiently close to $\mathscr{M}$ and $\mathscr{M}_-$ in the Hausdorff distance for the purposes of inferring the relative homology of the pair $(\mathscr{M}, \mathscr{M}_-)$.

We construct $\mathscr{M}'$ and $\mathscr{M}'_-$ by first recalling the definitions of $\mathscr{M}$ and $\mathscr{M}'$. In eq. (4.15), we have defined $\mathscr{M}$ to be the following intersection of sublevel sets

$$\mathscr{M} = h^{-1}(-\infty, \eta_0] \cap f^{-1}[0, \infty),$$

where $\eta_0 = -q(r)$, and $r \in \left(\frac{r_0+r_1}{2}, r_1\right)$. In eq. (4.17), we have defined $\mathscr{M}_- = h^{-1}(-\infty, \eta_0] \cap f^{-1}(0)$. Now let us consider constants $r' \in (r, r_1)$, as well as η' and η'', such that $\beta - \frac{q_0}{2} < -\eta' < \eta = 0 < \eta'' < \alpha$ for β and α corresponding to those in eq. (4.14). We also choose $\eta' > 0$ to be sufficiently small, such that $\eta'_0 := -\eta' - q(r') > \eta_0 = -q(r)$. As there are no critical points of f and h on the annulus $g^{-1}[r_0, r_1]$ due to assumptions **(G3)** and **(G4)** and **(H5)**, the conditions above imply $\mathscr{M} \subset g^{-1}(-\infty, r']$, and

$$\mathscr{M}' = h^{-1}(-\infty, \eta'_0] \cap f^{-1}[-\eta', \infty) \supset \mathscr{M} \quad \text{and} \tag{4.20}$$

$$\mathscr{M}'_- = h^{-1}(-\infty, \eta'_0] \cap f^{-1}\left[-\eta', \eta''\right] \supset \mathscr{M}_-. \tag{4.21}$$

We illustrate $\mathscr{M}'$ and $\mathscr{M}'_-$ in relation to $\mathscr{M}$ and $\mathscr{M}_-$ in Figure 4.3.

To keep the notation concise in the subsequent discussion, we shift $h \mapsto h - q(r)$ by a negative constant $-q(r)$ where r is the radius that defines $\mathscr{M}$ and q is the function eq. (4.11). Doing so we can write $\mathscr{M}$ as

$$\mathscr{M} = h^{-1}(-\infty, 0] \cap f^{-1}[0, \infty)$$

We also shift the constant $\eta'_0 \mapsto \eta'_0 - q(r)$ by the same amount so that the definition of $\mathscr{M}$ and $\mathscr{M}'$ are consistent.

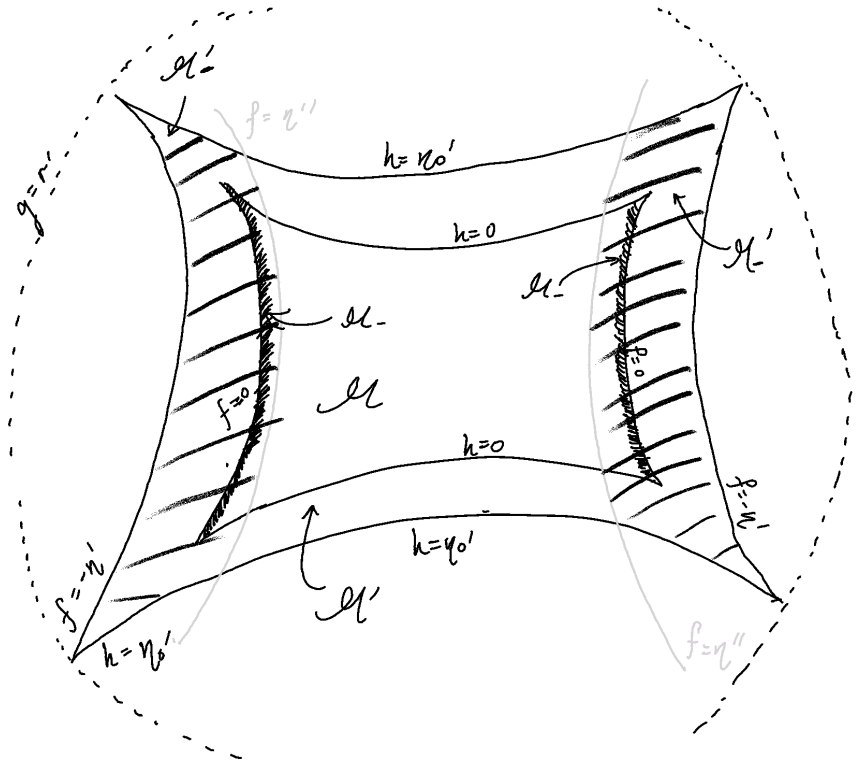

Figure 4.3: An illustration of $\mathscr{M}'$ and $\mathscr{M}'_-$ as defined in eq. (4.20).

Let us now consider the point clouds

$$\mathbb{X} = \mathbb{Y} \cap \mathscr{M}' \tag{4.22}$$

$$\mathbb{X}_- = \mathbb{Y} \cap \mathscr{M}'_- \subset \mathbb{X}. \tag{4.23}$$

We now derive sufficient conditions for a point cloud $(\mathbb{X}, \mathbb{X}_-)$ to be Hausdorff-close to $(\mathscr{M}, \mathscr{M}_-)$. We first address the issue of density of $\mathbb{X}$ and $\mathbb{X}_-$ in $\mathscr{M}$ and $\mathscr{M}_-$ respectively.

Lemma 4.3.2. *If $\mathbb{Y}$ is δ-dense in $N \supset \mathscr{M}' \supset \mathscr{M}$, then for δ sufficiently small, $\mathbb{X} = \mathscr{M}' \cap \mathbb{Y}$ is δ-dense in $\mathscr{M}$. Similarly, $\mathbb{X}_- = \mathscr{M}'_- \cap \mathbb{Y}$ is also δ-dense in $\mathscr{M}_-$*

Proof. Suppose f and h have a Lipschitz constant bounded above by K on N and let τ be the minimum of the local feature size of M on the compact subset N. Let us consider some point $p \in \mathscr{M}$. If $\mathbb{Y}$ is δ-dense in $N \supset \mathscr{M}$, then there exists some $y \in \mathbb{Y}$ such that $\|y - p\| < \delta$. If $y \notin \mathbb{X}$, then either $f(y) < -\eta'$ or $h(y) > \eta'_0$. Suppose the former is true and $\delta < \tau$. Then due to Lipschitz continuity and Corollary 3.3.2, $|f(p) - f(y)| \leq K d_M(p, y) \leq 2K\tau \operatorname{asin}\left(\frac{\delta}{2\tau}\right)$. This

implies

$$2K\tau \operatorname{asin}\left(\frac{\delta}{2\tau}\right) \geq f(p) - f(y) \geq \eta'. \tag{4.24}$$

Similarly we can perform the same analysis for h and deduce that if $h(y) > \eta_0'$ implies

$$2K\tau \operatorname{asin}\left(\frac{\delta}{2\tau}\right) \geq h(p) - h(y) \geq \eta_0'. \tag{4.25}$$

Thus for a sufficiently small δ that violates both of these conditions, we reach a contradiction and therefore guarantee $y \in \mathbb{X} = \mathscr{M}' \cap \mathbb{Y}$. Thus, for any $p \in \mathscr{M}$, we can always find a point $x \in \mathbb{X}$ such that $\|x - p\| \leq \delta$ if the point sample $\mathbb{Y}$ is sufficiently dense so that despite removing points with $h > \eta_0'$ or $f < -\eta'$ from $\mathbb{Y}$, the remaining point cloud $\mathbb{X}$ is still dense in $\mathscr{M}$.

For $\mathscr{M}'_-$, we need to ensure that for any $p \in \mathscr{M}_-$ there is some $y \in \mathbb{X}_- \subset \mathbb{Y}$ such that $\|p - y\| < \delta$. Suppose $y \in \mathbb{Y}$ and $\|p - y\| < \delta$, but $y \notin \mathbb{X}_-$, i.e. either $f(y) < -\eta'$ or $h(y) > \eta_0'$, which are the condtions for $\mathbb{X}$ we have discussed above, or additionally, $f(y) > \eta''$. The latter condition implies

$$2K\tau \operatorname{asin}\left(\frac{\delta}{2\tau}\right) \geq \eta''.$$

Thus taking δ to be sufficiently small so that this is ruled out alongside eqs. (4.24) and (4.25) ensures that y remains in the point cloud after restricting from $\mathbb{Y}$ to $\mathbb{X}_-$. □

We now consider the issue of closeness. Recall we say $\mathscr{M}'$ is θ-close to $\mathscr{M}$ if for any $q \in \mathscr{M}'$ there is some $p \in \mathscr{M}'$ such that $\|p - q\| < \theta$. Here we argue $\mathscr{M}'$ can be made θ-close for arbitrarily small θ if our construction of Gromoll-Meyer pairs is sufficiently robust against choices of parameters for the function values of f and h. By robustness, we mean our construction and in particular choice of g admits a deformation retract of $\mathscr{M}'$ onto $\mathscr{M}$ for small enough values of η' and η_0'. As such a deformation retract defines an explicit path that sends $q \in \mathscr{M}'$ to some $p \in \mathscr{M}'$, the arc length of such a path provides an upper bound for the geodesic distance along M between p and q, which in turn is an upper bound for the Euclidean distance between p and q. By setting parameters η' and η_0' to be sufficiently small, we can bound the arc length to be at most θ. In the lemma below, when we say a positive quantity is $\mathcal{O}(t)$, we imply its variation with t is bounded above by Mt, where M is a sufficiently large constant independent of t.

Lemma 4.3.3. *If $[r, r']$ are regular values of g restricted to $f^{-1}(-\eta')$ and $h^{-1}(0)$, then $\mathscr{M}'$ is $(\mathcal{O}(\eta' + \eta_0'))$-close to $\mathscr{M}$, where $\eta_0' = \eta' + q(r) - q(r')$ and q is the monotonically decreasing function as defined in eq. (4.11). In addition, $\mathscr{M}'_-$ is $(\mathcal{O}(\eta' + \eta_0' + \eta''))$-close to $\mathscr{M}_-$.*

Proof. Let us consider the gradient flow σ along $-\nabla f$ as defined in eq. (4.2) on $\mathscr{M}'$. Due to Lemma 4.2.7 and **(G3)**, all critical points of f and h in $g^{-1}(-\infty, r] \supsetneq \mathscr{M}'$ are contained in $\mathscr{M}$, and hence we have non-vanishing gradient flows along ∇f or ∇h on an open neighbourhood of $\mathscr{M}'$. As $\mathscr{M}'$ satisfies the same conditions as $\mathscr{M}$, the pair $\mathscr{M}'$ is one half of a Gromoll-Meyer pair for σ in its own right, with $\mathscr{M}' \cap h^{-1}(\eta_0')$ as its entry set and $\mathscr{M}' \cap f^{-1}(-\eta')$ as its the exit set. For points $x \in \mathscr{M}' \cap h^{-1}[0, \eta_0']$, the gradient flow $\sigma(t, x)$ either intersects $h^{-1}(-\infty, 0] \cap \mathscr{M}' = h^{-1}(-\infty, 0] \cap f^{-1}[\eta', \infty) \supset \mathscr{M}$ in finite time or exits $\mathscr{M}'$ at $\mathscr{M}' \cap f^{-1}(-\eta') \cap h^{-1}[0, \eta_0'] = f^{-1}(\eta') \cap h^{-1}[0, \eta_0']$ without intersecting $h^{-1}(-\infty, 0] \cap \mathscr{M}'$.

In either case, let s_1 be the arc length of $\sigma(t, x)$ from x to some point p where $h(p) = 0$, which may be in $\mathscr{M}' \cap h^{-1}(-\infty, 0]$ or $f^{-1}(-\eta') \cap h^{-1}(-\infty, 0]$. Then integrating along $\sigma(t, x)$, we can obtain the following bound

$$\mu s_1 \leq h(x) - h(p) \leq \eta_0'. \tag{4.26}$$

where μ is the lower bound of ∇h projected onto the direction of ∇f on $\mathscr{M}'$, which is positive due to **(H5)** and the absence of critical points of h and f on $\mathscr{M}' \setminus \mathscr{M}$.

We now show that if $p \in f^{-1}(\eta') \cap h^{-1}[0, \eta_0']$, then our assumptions allow us to flow p onto some point $q \in f^{-1}(\eta') \cap h^{-1}(0) \subset \mathscr{M}' \cap h^{-1}(-\infty, 0]$ along a flow that decreases h. From the definition of h as $h(p) = f - q(g(p)) - q(r)$ and the monotonicity of q on (r_0, r_1) **(Q3)**, one can verify that $g^{-1}[r, r'] \cap f^{-1}(-\eta') \supset f^{-1}(\eta') \cap h^{-1}[0, \eta_0']$. As we have assumed there are no critical values of g restricted to $f^{-1}(-\eta')$ in the interval $[r, r']$, there are no critical points of h restricted to $f^{-1}(\eta')$ on $h^{-1}[0, \eta_0']$. Therefore, as $f^{-1}(\eta') \cap h^{-1}[0, \eta_0']$ is closed in a compact set $g^{-1}[r, r'] \cap f^{-1}(-\eta')$, the norm of the gradient of h restricted to $f^{-1}(\eta')$ can be bounded below by $\nu > 0$ on $f^{-1}(\eta') \cap h^{-1}[0, \eta_0']$. Let s_2 be the arc length of the gradient flow line decreasing h on $f^{-1}(\eta')$, connecting p and q. Then s_2 is bounded by

$$\nu s_2 \leq h(p) - h(q) \leq \eta_0'. \tag{4.27}$$

Combining eqs. (4.26) and (4.27), we deduce that for any point $x \in \mathscr{M}'$, there is a path from x

to another point $q \in \mathscr{M}' \cap h^{-1}(-\infty, 0]$ such that

$$\|q - x\| \leq s_1 + s_2 \leq \eta_0' \left(\frac{1}{\mu} + \frac{1}{\nu} \right).$$

In a completely analogous way, we can construct a path from any point $q \in \mathscr{M}' \cap h^{-1}(-\infty, 0] = f^{-1}[-\eta', \infty) \cap h^{-1}(-\infty, 0]$ to some point in $f^{-1}(0) \cap h^{-1}(-\infty, 0] = \mathscr{M}_- \subset \mathscr{M}$. Flowing q in $f^{-1}[-\eta', \infty) \cap h^{-1}(-\infty, 0]$ along ∇f, the path it takes either intersects $\mathscr{M}$ or $f^{-1}[-\eta', 0] \cap h^{-1}(0)$ in finite time. If it intersects $f^{-1}[-\eta', 0] \cap h^{-1}(0)$, due to our assumptions that there are no critical values of g restricted to $h^{-1}(0)$ in the interval $[r, r']$, we can deduce along analogous arguments as above that there is a path from q to $\mathscr{M}_- \subset \mathscr{M}$ that is order η'.

Comibining the two paths from x to p and then from p to q, we deduce that for any $x \in \mathscr{M}'$, the distance to $\mathscr{M}$ is at most $\mathcal{O}(\eta') + \mathcal{O}(\eta_0')$.

We now turn our attention to $\mathscr{M}_-' = h^{-1}(-\infty, \eta_0'] \cap f^{-1}[-\eta', \eta'']$. We partition points in $\mathscr{M}_-'$ into those that lie in $\mathscr{M}' \setminus \mathscr{M}$, and those that lie in $\mathscr{M}$. For points in $\mathscr{M}_-'$ that lie in $\mathscr{M}' \setminus \mathscr{M}$, consider the flow along $-\nabla f$; these points either intersect $\mathscr{M}' \cap f^{-1}(-\eta')$ or $\mathscr{M}$ in finite time.

For points that intersect $\mathscr{M}' \cap f^{-1}(-\eta')$, we have established a path between them and points of $\mathscr{M}_-$ with $\mathcal{O}(\eta_0') + \mathcal{O}(\eta')$ bounded path length in the discussion above regarding the closeness of $\mathscr{M}'$ to $\mathscr{M}$. Thus, we only have to concern ourselves with points that intersect $\mathscr{M}$ along the flow. As f decreases along the flow, such points must lie in $h^{-1}(-\infty, \eta_0'] \cap f^{-1}[0, \eta'']$, as otherwise they would not intersect $\mathscr{M} = h^{-1}(-\infty, 0] \cap f^{-1}[0, \infty)$ along the flow. Furthermore, the flow away from these points along $-\nabla f$ must intersect $\mathscr{M}$ in a region contained in $\mathscr{M} \cap f^{-1}[0, \eta'']$ since f cannot increase along the flow. As we have chosen $\eta'' < \alpha$ and the critical points of f in $\mathscr{M}$ have $f > \alpha$, the gradient field ∇f on $\mathscr{M} \cap f^{-1}[0, \eta'']$ is non-zero. As $\mathscr{M}_-$ is the exit set of the gradient flow on $\mathscr{M}$, we can deformation retract $\mathscr{M} \cap f^{-1}[0, \eta'']$ onto $\mathscr{M}_-$. Thus this deformation retract provides a path of length $\mathcal{O}(\eta'')$ between points in $\mathscr{M} \cap f^{-1}[0, \eta'']$ to $\mathscr{M}'$. □

Chapter 5

Optimisation of Spectral Wavelets for Persistence-Based Graph Classification

The material of this chapter is reproduced from the paper Yim and Leygonie [2021] jointly written with Jacob Leygonie. The overall framework was jointly conceived by both authors. Yim was responsible for proposing wavelet signatures as a vertex function parameterization framework, along with the experimental design and analysis. The proof of the differentiability of extended persistence is due to Jacob Leygonie.

5.1 Introduction

5.1.1 Background

Graph classification is a challenging problem in machine learning. Unlike data represented in Euclidean space, there is no easily computable notion of distance or similarity between graphs. As such, graph classification requires techniques that lie beyond mainstream machine learning techniques focused on Euclidean data. Much research has been conducted on methods such as graph neural networks (GNNs) [Xu et al., 2019] and graph kernels [Vishwanathan et al., 2010, Shervashidze et al., 2009] that embed graphs in Euclidean space in a consistent manner.

Recently, *persistent homology* [Zomorodian and Carlsson, 2005, Edelsbrunner and Harer, 2008] has been applied as a feature map that explicitly represents topological and geometric features

of a graph as a set of *persistence diagrams* (a.k.a. *barcodes*). In the context of our discussion, the persistent homology of a graph $G = (V, E)$ depends on a vertex function $f : V \to \mathbb{R}$.

In the case where a vertex function is not given with the data, several schemes have been proposed in the literature to assign vertex functions to graphs in a consistent way. For example, vertex functions can be constructed using local geometric descriptions of vertex neighbourhoods, such as discrete curvature [Zhao and Wang, 2019] and Weisfeiler–Lehman graph kernels [Rieck et al., 2019]. There are also many approaches based on graph Laplacians, such as heat kernel signatures Sun et al. [2009], wave kernel signatures Aubry et al. [2011], and further generalisations Litman and Bronstein [2013]. .

However, it is often difficult to know *a priori* whether a heuristic vertex assignment scheme will perform well in addressing different data science problems. For a single graph, we can optimise the vertex function over $|V|$ many degrees of freedom in $\mathbb{R}^V$. In recent years, there have been many other examples of persistence optimisation in data science applications. The first two examples of persistence optimisation are the computation of Fréchet mean of barcodes using gradients on Alexandrov spaces [Turner et al., 2014a], and that of point cloud inference [Gameiro et al., 2016], where a point cloud is optimised so that its barcode fits a target fixed barcode. The latter is an instance of topological inverse problems (see [Oudot and Solomon, 2020] for a recent overview of such). Another inverse problem is that of surface reconstruction [Brüel-Gabrielsson et al., 2020]. Besides, in the context of shape matching [Poulenard et al., 2018], persistence optimisation is used in order to learn an adequate function between shapes. Finally, there are also many recent applications of persistence optimisation in Machine Learning, such as the incorporation of topological information in Generative Modelling [Moor et al., 2020, Hofer et al., 2019, Gabrielsson et al., 2020] or in Image Segmentation [Hu et al., 2019, Clough et al., 2019], the design of topological losses for regularization in supervised learning [Chen et al., 2019a] or for dimension reduction [Kachan, 2020].

Each of these applications can be thought of as minimising a certain *loss* function over a manifold $\mathcal{M}$ of parameters:

$$\min_{\theta \in \mathcal{M}} \mathcal{L}(\theta),$$

where $\mathcal{L} : \mathcal{M} \to \mathbf{Bar}^N \to \mathbb{R}$ factors through the space $\mathbf{Bar}^N$ of N-tuples of barcodes. The aim is to find the parameter θ that best fits the application at hand. Gradient descent is a very

popular approach in minimisation, but it requires the ability to differentiate the loss function.

[Leygonie et al., 2019] provide notions of differentiability for maps in and out **Bar** that are compatible with smooth calculus, and show that the loss functions $\mathcal{L}$ corresponding to the applications cited in the above paragraph are generically differentiable. The use of (stochastic) gradient descent is further legitimated by Carriere et al. [2020], where convergence guarantees on persistence optimisation problems are devised, using a recent study of stratified non-smooth optimisation problems [Davis et al., 2020]. In practice, the minimisation of $\mathcal{L}$ can be unstable due to its non-convexity and partial non-differentiability. Some research has been conducted in order to smooth and regularise the optimisation procedure [Solomon et al., 2020, Corcoran and Deng, 2020].

In a supervised learning setting, we want to optimise our vertex function assignment scheme over many individual graphs in a dataset. Since graphs may not share the same vertex set and come in different sizes, optimising over the $|V|$ degrees of freedom of any one graph is not conducive to learning a vertex function assignment scheme that can generalise to another graph. The degrees of freedom in any practical vertex assignment scheme should be independent of the number of vertices of a graph. However, a framework for parametrising and optimising the vertex functions of many graphs over a common parameter space $\mathcal{M}$ is not immediately apparent.

The first instance of a graph persistence optimisation framework (`GFL`) [Hofer et al., 2020] uses a one layer graph isomorphism network (GIN) [Xu et al., 2019] to parametrise vertex functions. The GIN learns a vertex function by exploiting the local topology around each vertex. In this paper, we propose a different framework for assigning and parametrising vertex functions, based on a graph's normalised Laplacian operator. Using the normalised Laplacian, we can explicitly take both local and global structures of the graph into consideration in an interpretable and transparent manner.

5.1.2 Outline and Contributions

We address the issue of vertex function parametrisation and optimisation using *wavelet signatures*. Wavelet signatures are vertex functions derived from the eigenvalues and eigenvectors

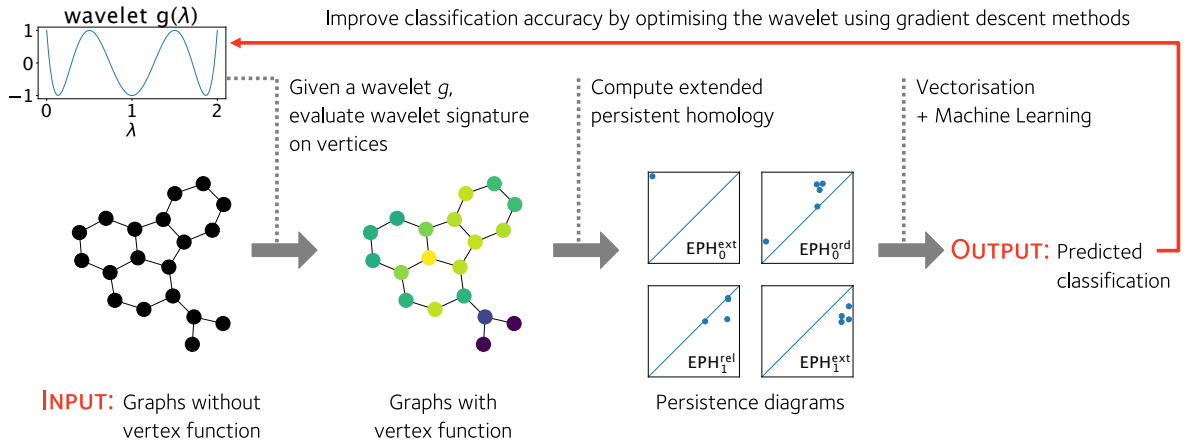

Figure 5.1: Given a wavelet $g : \mathbb{R} \to \mathbb{R}$, we can equip any graph with a non-trivial vertex function. This allows us to compute the extended persistence diagrams of a graph and use the diagrams as features of the graph to predict a graph's classification in some real world setting. The wavelet g can be optimised to improve the classification accuracy of a graph classification pipeline based on the extended persistence diagrams of a graph's vertex function.

of the graph Laplacian and encode multiscale geometric information about the graph [Li and Hamza, 2013]. The wavelet signature of a graph is dependent on a choice of wavelet $g : \mathbb{R} \to \mathbb{R}$, a function on the eigenvalues of the graph's Laplacian matrix. We can thus obtain a parametrisation of vertex functions for any graph $\mathrm{F} : \mathcal{M} \to \mathbb{R}^V$ by parametrising g. Consequently, the *extended persistence* of a graph – which has only four non-trivial persistence diagrams – can be varied over the parameter space $\mathcal{M}$. If we have a function $\mathrm{Out} : \mathbf{Bar}^4 \to \mathbb{R}$ on persistence diagrams that we wish to minimise, we can optimise over $\mathcal{M}$ to minimise the loss function

$$\mathcal{L} : \mathcal{M} \xrightarrow{\mathrm{F}} \mathbb{R}^V \xrightarrow{\mathrm{EPH}} \mathbf{Bar}^4 \xrightarrow{\mathrm{Out}} \mathbb{R}. \tag{5.1}$$

If $\mathcal{L}$ is generically differentiable, we can optimise the wavelet signature parameters $\theta \in \mathcal{M}$ using gradient descent methods. We illustrate an application of this framework to a graph classification problem in Figure 5.1, where the loss function $\mathcal{L}$ is the classification error of a graph classification prediction model based on the graph's extended persistence diagrams.

In Section 5.2, we describe the assignment of vertex functions $\mathrm{F} : \mathcal{M} \to \mathbb{R}^V$ by reviewing the definition of wavelet signatures. While spectral wavelets have been used in graph neural network architectures that predict vertex features [Xu et al., 2019] and compress vertex functions [Rustamov and Guibas, 2019], they have not been considered in a persistent homology framework for graph classification. We describe several ways to parametrise wavelets. We also

show in Proposition 5.2.2 that wavelet signatures are independent of the choice of eigenbasis of the graph Laplacian from which it is derived, ensuring that it is well-defined.

Finally, in Section 5.4, we apply our framework to graph classification problems on several benchmark datasets. We show that our model is competitive with state-of-the-art persistence-based models. In particular, optimising the vertex function appreciably improves the prediction accuracy on some datasets.

5.2 Filter Function Parametrization

We describe our recipe for assigning vertex functions to any simple graph $G = (V, E)$. Recall a simple graph is an undirected graph with no self-edges or multiple edges between vertices. Our assignment is based on a parametrised spectral wavelet, the first part F of the loss function

$$\mathcal{L} : \mathcal{M} \xrightarrow{\mathrm{F}} \mathbb{R}^V \xrightarrow{\mathsf{EPH}} \mathbf{Bar}^4 \xrightarrow{\mathsf{Out}} \mathbb{R} \ . \qquad \text{(Equation (5.1) recalled)}$$

Our recipe is based on a graph's wavelet signature, a vertex function derived from the graph's Laplacian. The wavelet signature also depends on a so-called 'wavelet function' in $g : \mathbb{R} \to \mathbb{R}$, which is independent of the graph. By modulating the wavelet, we can jointly vary the wavelet signature across many graphs. We parametrise the wavelet using a finite linear combination of basis functions, such that the wavelet signature can be manipulated in a computationally tractable way. In the following section, we define the wavelet signature and describe our linear approach to wavelet parametrisation.

5.2.1 Wavelet Signatures

The wavelet signature is a vertex function initially derived from wavelet transforms of vertex functions on graphs Hammond et al. [2011], a generalisation of wavelet transforms for square integrable functions on Euclidean space Graps [1995], Chui [2016] for signal analysis Akansu et al. [2001]. Wavelet signatures for graphs have been applied to encode geometric information about meshes of 3D shapes Akansu et al. [2001], Li and Hamza [2013]. Special cases of wavelets signatures, such as the heat kernel signature Sun et al. [2009] and wave kernel signature Aubry et al. [2011], have also been applied to describe graphs and 3D shapesBronstein

and Kokkinos [2010], Hu et al. [2014].

The wavelet signature of a graph is constructed from the graph's Laplacian operator. Let k_u denote the degree of vertex u. A graph's normalised Laplacian $L \in \mathbb{R}^{V \times V}$ is a symmetric positive semi-definite matrix, whose entries are given by

$$L_{uv} = \begin{cases} 1 & u = v \quad \text{and} \quad k_u \neq 0 \\ -\frac{1}{\sqrt{k_u k_v}} & (u,v) \in E \\ 0 & \text{otherwise} \end{cases} . \tag{5.2}$$

The Laplacian's eigenvalues λ and eigenvectors $\boldsymbol{\varphi}$ are known to encode various topological and geometric information about the graph Chung and Graham [1997], Biyikoglu et al. [2007]; for example, the number of zero eigenvalues corresponds to the number of connected components of the graph. The spectrum of the normalised Laplacian have real eigenvalues in $[0,2]$ Chung and Graham [1997]. As such, any function $g : \mathbb{R} \to \mathbb{R}$ evaluated on the eigenvalues need only be defined on $[0,2]$. Moreover, functions on a compact domain are easily parametrised using convenient bases.

Definition 5.2.1. [Wavelet Signature Li and Hamza [2013]] Let $L \in \mathbb{R}^{V \times V}$ be the normalised Laplacian of a simplical graph $G = (V, E)$. Let $\boldsymbol{\varphi}_1, \ldots, \boldsymbol{\varphi}_{|V|}$ be an orthonormal eigenbasis for L and $\lambda_1, \ldots, \lambda_{|V|}$ be their corresponding eigenvalues. The wavelet signature $W : \mathbb{R}^{[0,2]} \to \mathbb{R}^V$ maps a function $g : [0,2] \to \mathbb{R}$, which we refer to as a *wavelet*, to a vertex function $W(g) \in \mathbb{R}^V$ linearly, where the value of $W(g)$ on vertex v is given by

$$W(g)_v = \sum_{i=1}^{|V|} g(\lambda_i)(\boldsymbol{\varphi}_i)_v^2, \tag{5.3}$$

and $(\boldsymbol{\varphi}_i)_v$ denotes the component of eigenvector $\boldsymbol{\varphi}_i$ corresponding to vertex v.

If the eigenvalues of L have geometric multiplicity one (i.e. their eigenspaces are one dimensional), then the orthonormal eigenvectors are uniquely defined up to a choice of sign. It is then apparent from Equation (5.3) that the wavelet signature is independent of the choice of sign. However, if some eigenvalues have geometric multiplicity greater than one, then the orthonormal eigenvectors of L are uniquely defined up to orthonormal transformations in the individual eigenspaces. However, the wavelet signature is well-defined even when the

multiplicities of eigenvalues are greater than one. This is the content of the next Proposition.

Proposition 5.2.2. *The wavelet signature of a graph is independent of the choice of orthonormal eigenbasis for the Laplacian.*

Proof. Let $\mathsf{Spec}(L) \subset \mathbb{R}$ denote the spectrum of L and $\boldsymbol{\varphi}_1, \dots \boldsymbol{\varphi}_{|V|}$ be a set of orthonormal eigenvectors of L. Let us denote $\Phi(\lambda)$ to be a $|V| \times m$ matrix where m corresponds to the geometric multiplicity of λ, and the m column vectors of $\Phi(\lambda)$ correspond to eigenvectors $\boldsymbol{\varphi}_{i_1}, \dots, \boldsymbol{\varphi}_{i_m}$ with eigenvalue λ. Then we can rewrite the wavelet signature eq. (5.3) as

$$W(g)_v = \sum_{i=1}^{|V|} g(\lambda_i)(\boldsymbol{\varphi}_i)_v^2 = \sum_{\lambda \in \mathsf{Spec}(L)} g(\lambda)(\Phi(\lambda)\Phi(\lambda)^\intercal)_{vv} \tag{5.4}$$

Suppose we have another choice of eigenbasis of L. Without loss of generality for $\lambda \in \mathsf{Spec}(L)$, the new basis $\boldsymbol{\varphi}'_{i_1}, \dots, \boldsymbol{\varphi}'_{i_m}$ for $\mathsf{eig}(\lambda)$ is related to the previous eigenbasis $\boldsymbol{\varphi}_{i_1}, \dots, \boldsymbol{\varphi}_{i_m}$ by an orthonormal transformation transformation $U(\lambda) \in \mathbb{R}^{m \times m}$ on $\Phi(\lambda)$:

$$\Phi'(\lambda) = \begin{bmatrix} \boldsymbol{\varphi}'_1 & \cdots & \boldsymbol{\varphi}'_m \end{bmatrix} = \Phi(\lambda)U(\lambda).$$

As $U(\lambda)$ is an orthonormal transformation with $U(\lambda)U(\lambda)^\intercal = 1$,

$$\begin{aligned} \Phi'(\lambda)\Phi'(\lambda)^\intercal &= \Phi(\lambda)U(\lambda)(\Phi(\lambda)U(\lambda))^\intercal \\ &= \Phi(\lambda)U(\lambda)U(\lambda)^\intercal\Phi(\lambda)^\intercal \\ &= \Phi(\lambda)\Phi(\lambda)^\intercal. \end{aligned}$$

Since the $V \times V$ matrix $\Phi(\lambda)\Phi(\lambda)^\intercal$ is independent of the choice of eigenbasis, the wavelet signature given on the right hand side of eq. (5.4) must also be independent of the choice of eigenbasis. □

Remark 5.2.3. In addition to the traditional view of wavelets from a spectral signal processing perspective Hammond et al. [2011], we can also relate the wavelet signature of a vertex v to the degrees of vertices in some neighbourhood of v prescribed by g. Consider a wavelet $g : [0, 2] \to \mathbb{R}$. On a finite graph G with no isolated vertices, the normalised Laplacian L has at most $|V|$ many distinct eigenvalues. As such, there exists a polynomial $\hat{g}(x) = \sum_{n=0}^{p} a_n x^n$ of finite order that interpolates g at the eigenvalues $g(\lambda_i) = \hat{g}(\lambda_i)$. Therefore, $W(g) = W(\hat{g})$. Moreover, the

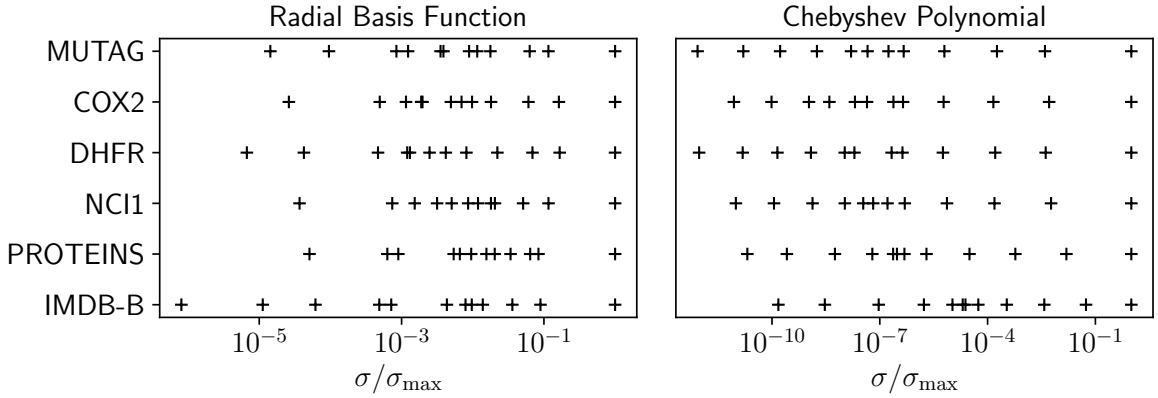

Figure 5.2: We consider the parametrisations of wavelet signatures on some datasets of graphs in machine learning, namely `MUTAG`, `COX2`, `DHFR`, `NCI1`, `PROTEINS` and `IMDB-B`, using coefficients of 12 radial basis functions (see eq. (5.27)) and a degree 13 Chebyshev polynomial respectively. For each dataset, we plot the distribution of the singular values σ of the map F in eq. (5.17) from the basis function coefficients $\theta \in \mathbb{R}^{12}$ to the wavelet signature on the whole dataset of graphs, as a fraction of the largest singular value $\sigma_{\max}$ of F. We can observe that for both parametrisations, the singular values span many orders of magnitudes across different datasets. Note that the singular values of F not only depend on the choice of basis but also on the dataset of graphs.

vertex values assigned by $W(\hat{g})$ are the diagonal entries of the matrix polynomial $\hat{g}(L)$:

$$\hat{g}(L)_{vv} = \sum_{n=0}^{p} a_n (L^n)_{vv} = \sum_{i=1}^{|V|} \hat{g}(\lambda_i)(\boldsymbol{\varphi}_i)_v^2 = \sum_{i=1}^{|V|} g(\lambda_i)(\boldsymbol{\varphi}_i)_v^2 = W(g)_{vv}. \tag{5.5}$$

Furthermore, we can also write the matrix polynomial $\hat{g}(L)$ as a matrix polynomial in $A = I - L$, the *normalised adjacency matrix*. If G has no isolated vertices, then the diagonal entry of the matrix monomial A^r corresponding to vertex v is an inverse degree weighted count of paths (a path being sequences of vertices that are connected to the next vertex in the sequence) $[v_0, v_1, \ldots, v_r]$ on the graph, which begin and end on vertex $v = v_0 = v_r$ Newman [2018]:

$$(A^r)_{vv} = \frac{1}{k_v} \sum_{[v,v_1,\ldots,v_{r-1},v]} \left(\prod_{l=1}^{r-1} \frac{1}{k_{v_l}} \right). \tag{5.6}$$

By expressing the wavelet signature as a matrix polynomial in A, we see that g controls how information at different length scales of the graph contribute to the wavelet signature. For instance, if g were an order p polynomial, then $W(g)_v$ only takes the degrees of vertices that are at most half the length of loops of length p based at v, i.e. vertices $\lfloor p/2 \rfloor$ separated from v. As the degree of the furthermost vertex in the loop also accounts for the connectivity of the vertices it is connected to, the subgraph of G that is within $\lfloor (p-1)/2 \rfloor + 1 = \lfloor (p+1)/2 \rfloor$

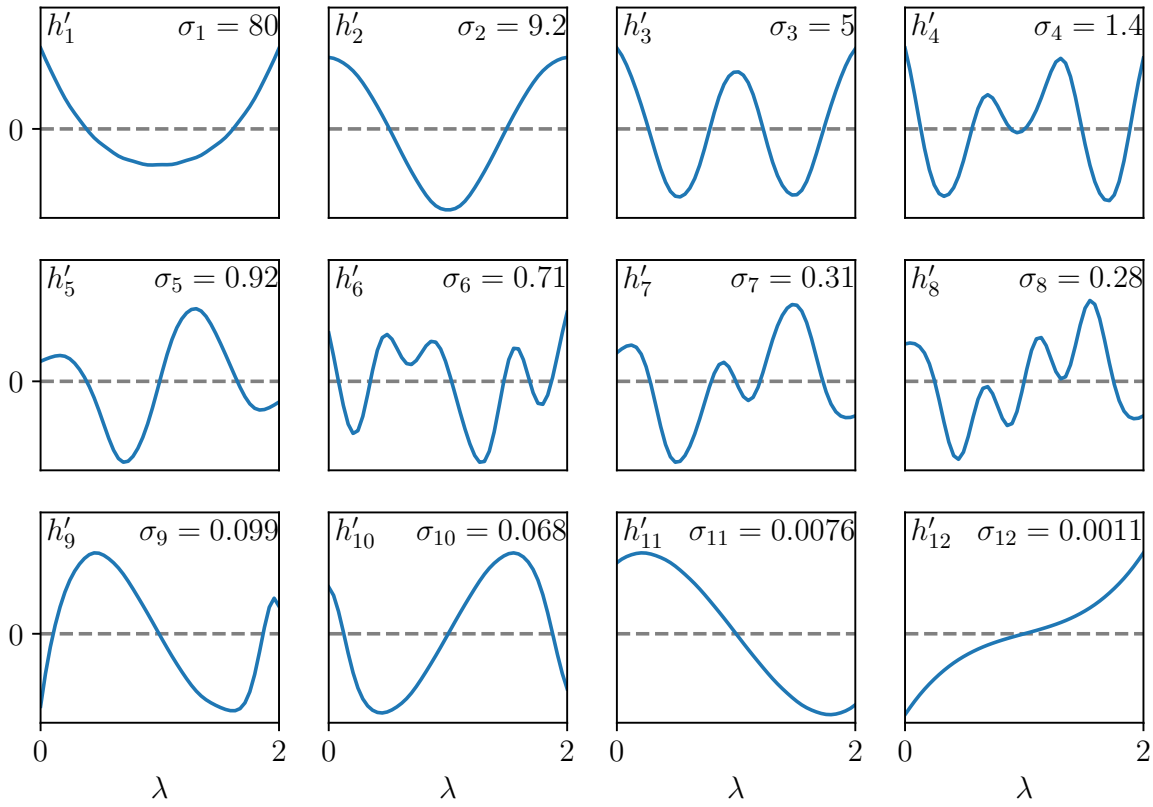

Figure 5.3: The functions shown are the new, stable wavelet basis $h'_1, \ldots, h'_{12}$ (eq. (5.14)) for the MUTAG dataset, derived from an initial numerically unstable parametrisation using twelve inverse multiquadric radial basis functions (eq. (5.27)). We parametrise the wavelet as a linear combination of these basis functions.

steps away from v, the wavelet signature at a vertex is only dependent on the subgraph of G that is within $\lfloor (p+1)/2 \rfloor$ steps away from v.

Since $W(g)$ can be specified by replacing g with a polynomial $\hat{g}$ of order at most $|V| - 1$, the wavelet signature at a vertex is only dependent on the subgraph of G that is within $\lfloor (|V|+1)/2 \rfloor$ steps away from v.

5.2.2 Parametrising the Wavelet

We see from Remark 5.2.3 that the choice of wavelet g determines how the topology and geometry of the graph is reflected in the vertex function. Though the space of wavelets is potentially infinite dimensional, here we only consider wavelets $g_\theta(x)$ that are parametrised by parameters θ in a finite dimensional manifold, so that we can easily optimise them using computational methods. In particular, we focus on wavelets written as a linear combination

of m *basis functions* $h_1, \ldots, h_m : [0,2] \to \mathbb{R}$

$$g_\theta(x) := \sum_{j=1}^{m} \theta_j h_j(x) \tag{5.7}$$

This parametrisation of wavelets in turn defines a parametrisation of vertex functions $\mathrm{F} : \mathbb{R}^m \to \mathbb{R}^V$ for our optimisation pipeline in eq. (5.1)

$$\mathrm{F}: \quad \theta \in \mathbb{R}^m \quad \longmapsto \quad \mathrm{F}(\theta) := W(g_\theta) \quad \in \mathbb{R}^V. \tag{5.8}$$

Since $W(g)$ is a linear function of the wavelet g, F is a linear transformation:

$$\mathrm{F}(\theta) = W\left(\sum_{j=1}^{m} \theta_j h_j(x)\right) = \sum_{j=1}^{m} \theta_j W(h_j). \tag{5.9}$$

We can write F as a $|V| \times m$ matrix acting on a vector $[\theta_1, \ldots \theta_m]^\intercal \in \mathbb{R}^m$, whose columns are the vertex functions $W(h_j)$.

Example 5.2.4 (Chebyshev Polynomials)**.** Any Lipschitz continuous function on an interval can be well approximated by truncating its Chebyshev series at some finite order Trefethen and Bau III [1997]. The Chebyshev polynomials $T_n : [-1,1] \to \mathbb{R}$

$$T_n(x) = \cos(n \arccos(x)) \quad n \in \mathbb{N}_{\geq 0}. \tag{5.10}$$

form an orthonormal set of functions. We can thus consider $h_j(\lambda) = T_j(\lambda - 1),\ j = 0, 2, \ldots, m$ as a naïve basis for wavelets. We exclude $T_1(x) = x$ in the linear combination as $W(T_1(1 - x)) = 0$ for graphs without self loops.

Example 5.2.5 (Radial Basis Functions)**.** In the machine learning community, a *radial function* refers loosely to a continuous monotonically decreasing function $\rho : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. There are many possible choices for ρ, for example, the inverse multiquadric

$$\rho(r) = \left(\left(\frac{r}{\epsilon}\right)^2 + 1\right)^{-\frac{1}{2}} \tag{5.11}$$

or the Gaussian

$$\rho(r) = \exp\left(-\frac{r^2}{2\epsilon}\right). \tag{5.12}$$

In both cases, the parameter $\epsilon > 0$ controls the width of the function. We can obtain a naïve

wavelet basis $h_j(x) = \rho(\|x - x_j\|)$ using copies of ρ offset by a collection of centroids $x_j \in \mathbb{R}$ along $\mathbb{R}$. In general, the centroids are parameters that could be optimised, but we fix them in this study. This parametrisation can be considered as a *radial basis function neural network*. RBNNs are well-studied in function approximation and subsequently machine learning; we refer readers to Chen et al. [1991], Park and Sandberg [1991] for further details.

5.2.3 The Choice of Wavelet Basis

The choice of basis functions determines the space of wavelet signatures and also the numerical stability of the basis function coefficients which serve as the wavelet signature parameters. The stability of the parametrisation depends on the graphs as much as the choice of wavelet basis $h_1, \ldots, h_m$. We can analyse the stability of a parametrisation F by its the singular value decomposition

$$\mathrm{F} = \sum_{k=1}^{r} \sigma_k \boldsymbol{u}_k \boldsymbol{v}_k^{\mathsf{T}} \tag{5.13}$$

where $\sigma_1, \ldots, \sigma_r$ are the non-zero singular values of the matrix, and $\boldsymbol{u}_k \in \mathbb{R}^{|V|}$ and $\boldsymbol{v}_k \in \mathbb{R}^m$ are orthonormal sets of vectors respectively. If the distribution of singular values span many orders of magnitude, we say the parametrisation is *ill-conditioned*. An ill-conditioned parametrisation interferes with the convergence of gradient descent algorithms on a loss function evaluated on wavelet signatures. We discuss the relationship between the conditioning of F and the stability of gradient descent in detail in Remark 5.2.7.

We empirically observe that the coefficients of a naïve choice of basis functions, such as Chebyshev polynomials or radial basis functions, are numerically ill-conditioned. In Figure 5.2, we can see that the singular values of radial basis function and Chebyshev polynomial parametrisations respectively are distributed across a large range on the logarithmic scale for some datasets of graphs in machine learning. We address this problem by picking out a new wavelet basis

$$h'_k(x) = \frac{1}{\sigma_k} \sum_{j=1}^{m} (\boldsymbol{v}_k)_j h_j(x), \quad k = 1, \ldots, r, \tag{5.14}$$

where σ_k are the singular values of F and $\boldsymbol{v}_k$ are the associated vectors in $\mathbb{R}^m$ from the singular value decomposition of matrix F in eq. (5.13). Then the parametrisation $\mathrm{F}' : \mathbb{R}^r \to \mathbb{R}^V$

$$\mathrm{F}'(\theta') = \sum_{k=1}^{r} \theta'_k W(h'_k). \tag{5.15}$$

have singular values equal to one, since this is a linear combination of orthonormal vectors $\boldsymbol{u}_k \in \mathbb{R}^V$:

$$W(h'_k) = \sum_{j=1}^{m} \frac{1}{\sigma_k} (\boldsymbol{v}_k)_j W(h_j) = \frac{1}{\sigma_k} \mathrm{F}\boldsymbol{v}_k = \boldsymbol{u}_k. \tag{5.16}$$

As an example, we plot the new wavelet basis h'_k derived from a twelve parameter radial basis function parametrisation for the `MUTAG` dataset in Figure 5.3.

Remark 5.2.6 (Learning a Wavelet Basis for Wavelet Signatures on Multiple Graphs)**.** In the case where the wavelet coefficients parametrise the wavelet signatures over graphs $G_1, \ldots, G_N$, we can view the maps $F_1, \ldots, F_N$ that map wavelet basis coefficients to vertex functions of graphs $G_1, \ldots, G_N$ respectively as a parametrisation for the disjoint union $\bigsqcup_i G_i$:

$$f = \begin{bmatrix} f_1 \\ \vdots \\ f_N \end{bmatrix} = \begin{bmatrix} \mathrm{F}_1 \\ \vdots \\ \mathrm{F}_N \end{bmatrix} \theta =: \mathrm{F}\theta. \tag{5.17}$$

We can then perform a singular value decomposition of the parametrisation F on $\bigsqcup_i G_i$ and derive a new, well-conditioned basis.

Remark 5.2.7 (Why the Conditioning of F Matters)**.** Let us optimise a loss function $\mathcal{L}$ on the parameter space of wavelet coefficients θ using a gradient descent algorithm. In a gradient descent step of step size s, the wavelet coefficients are updated to $\theta \mapsto \theta - s\boldsymbol{\nabla}_\theta \mathcal{L}$. Using the singular value decomposition of F (eq. (5.13)), we can write

$$\boldsymbol{\nabla}_\theta \mathcal{L} = \boldsymbol{\nabla}_\theta f^\intercal \boldsymbol{\nabla}_f \mathcal{L} = \mathrm{F}^\intercal \boldsymbol{\nabla}_f \mathcal{L} = \sum_{k=1}^{r} \sigma_k \left\langle \boldsymbol{\nabla}_f \mathcal{L}, \boldsymbol{u}_k \right\rangle \boldsymbol{v}_k. \tag{5.18}$$

The change in the vertex function is simply the matrix F applied to the change in wavelet parameters. Hence the vertex function is updated to $f \mapsto f - s\mathrm{F}\boldsymbol{\nabla}_\theta \mathcal{L}$, where

$$\mathrm{F}\boldsymbol{\nabla}_\theta \mathcal{L} = \sum_{k=1}^{r} \sigma_k \left\langle \boldsymbol{\nabla}_f \mathcal{L}, \boldsymbol{u}_k \right\rangle \mathrm{F}\boldsymbol{v}_k = \sum_{k=1}^{r} \sigma_k^2 \left\langle \boldsymbol{\nabla}_f \mathcal{L}, \boldsymbol{u}_k \right\rangle \boldsymbol{u}_k. \tag{5.19}$$

If the loss function $\mathcal{L}$ has large second derivatives– for example, due to nonlinearities in the function on persistence diagrams $\mathrm{Out} : \mathbf{Bar}^4 \to \mathbb{R}$ – the projections $\left\langle \boldsymbol{\nabla}_f \mathcal{L},\ \boldsymbol{u}_k \right\rangle$ in eqs. (5.18) and (5.19) may change dramatically from one gradient descent update to another. If the smallest singular value is much smaller than the largest, then updates to the wavelet signature can be especially unstable throughout the optimisation process. This source of instability can

be removed if we choose a parametrisation with singular values $\sigma_k = 1$ for all k. In this case, the update to f is simply the projection of $\nabla_f \mathcal{L}$ onto the space of wavelet signatures spanned by $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_r$, without any distortion introduced by non-uniform singular values:

$$f \mapsto f - s \sum_{k=1}^{r} \left\langle \boldsymbol{u}_k, \nabla_f \mathcal{L} \right\rangle \boldsymbol{u}_k. \tag{5.20}$$

5.3 Extended Persistent Homology

In this section, we begin by summarising some basic facts about the extended persistence of filtered graphs. We refer the reader to Cohen-Steiner et al. [2007] for a full treatment of the theory of extended persistence. Compared to *ordinary persistence*, extended persistence is a more informative and convenient feature map for graphs. For example, we can obtain the ordinary persistence diagram of the filtered graph from the corresponding extended diagram, by taking points in the 'extended' portion of the extended diagram, and replacing the finite death coordinates with infinity. In order to vectorise ordinary persistence diagrams, practitioners also have to introduce arbitrary finite cutoffs to remove the infinite death coordinates. These arbitrary choices are unnecessary for extended persistence diagrams as all points have finite birth and death coordinates.

5.3.1 Extended Persistent Homology

Let $G = (V, E)$ be a finite simple graph. For the purposes of this paper, the associated *extended persistent homology* is a map

$$\mathsf{EPH} : \mathbb{R}^V \to \mathbf{Bar}^4$$

from functions $f \in \mathbb{R}^V$ on its vertices to the space of four *persistence diagrams* or *barcodes*, which we define below. The map arises from a *filtration* of the graph, a sequential attachment of vertices and edges in ascending or descending order of f. We extend f on each edge $e = (v, v')$ by the maximal value of f over the vertices v and v', and we then let $G_t \subset G$ be the sub graph induced by vertices taking value less than t. Then we have the following sequence of inclusions:

$$\emptyset \longrightarrow \cdots \longrightarrow G_s \xrightarrow{s \leq t} G_t \longrightarrow \cdots \longrightarrow G. \tag{5.21}$$

Similarly, the sub graphs $G^t \subset G$ induced by vertices taking value greater than t assemble into a sequence of inclusions:

$$G \longleftarrow \cdots \longleftarrow G^s \xleftarrow{s \leq t} G^t \longleftarrow \cdots \longleftarrow \emptyset \ . \tag{5.22}$$

The changes in the topology of the graph along the filtration in ascending and descending order of f can be detected by its *extended persistence module*, indexed over the poset $\mathbb{R} \cup \{\infty\} \cup \mathbb{R}^{\mathrm{op}}$:

$$\mathsf{V}_p(f): \begin{array}{ccccccccc} 0 = H_p(\emptyset) & \longrightarrow \cdots \longrightarrow & H_p(G_s) & \xrightarrow{s \leq t} & H_p(G_t) & \longrightarrow \cdots \longrightarrow & H_p(G) \\ & & & & & & \downarrow \cong \\ 0 = H_p(G, G) & \longleftarrow \cdots \longleftarrow & H_p(G, G^s) & \xleftarrow{s \leq t} & H_p(G, G^t) & \longleftarrow \cdots \longleftarrow & H_p(G, \emptyset) \end{array}, \tag{5.23}$$

where H_p is the singular (relative) homology functor in degre $p \in 0, 1$ with coefficients in a fixed field, chosen to be $\mathbb{Z}/2\mathbb{Z}$ in pratice. The modules $\mathsf{V}_0(f)$ and $\mathsf{V}_1(f)$ together capture the evolution of the connected components and loops in the sub graphs of G induced by the function f.

We recall a persistence module is completely characterised by its decomposition into interval modules representing the birth and death of homological features. As a shorthand, we denote $\mathsf{EPH}_p(f)$ to be the barcode of $\mathsf{V}_p(f)$. Cohen-Steiner et al. [2009] partitions the intervals in $\mathsf{EPH}_p(f)$ into the *ordinary* part of the diagram $\mathsf{EPH}_p^{\mathsf{ord}}(f)$, the *extended* part of the diagram $\mathsf{EPH}_p^{\mathsf{ext}}(f)$, and the *relative* part of the diagram $\mathsf{EPH}_p^{\mathsf{rel}}(f)$, based on whether the birth and death of the interval occur in $\mathbb{R}$ or $\mathbb{R}^{\mathrm{op}}$:

$$\mathsf{EPH}_p(f) = \underbrace{\{\langle b, d\rangle \mid b \in \mathbb{R},\ d \in \mathbb{R}\}}_{=\mathsf{EPH}_p^{\mathsf{ord}}(f)} \sqcup \underbrace{\{\langle b, d\rangle \mid b \in \mathbb{R},\ d \in \mathbb{R}^{\mathrm{op}}\}}_{=\mathsf{EPH}_p^{\mathsf{ext}}(f)} \sqcup \underbrace{\{\langle b, d\rangle \mid b \in \mathbb{R}^{\mathrm{op}},\ d \in \mathbb{R}^{\mathrm{op}}\}}_{=\mathsf{EPH}_p^{\mathsf{rel}}(f)}. \tag{5.24}$$

As $\mathbb{R}$ and $\mathbb{R}^{\mathrm{op}}$ only differ in their order relations, individual elements of $\mathbb{R}$ and $\mathbb{R}^{\mathrm{op}}$ have identical numerical representations on a computer. To disambiguate the different parts of

$\mathsf{EPH}_p(f)$ in our numerical representation, we split $\mathsf{EPH}_p(f)$ into three *separate* diagrams along the partition into ordinary, extended, and relative parts:

$$\mathsf{EPH}_p(f) = \left[\mathsf{EPH}_p^{\text{ord}}(f), \mathsf{EPH}_p^{\text{ext}}(f), \mathsf{EPH}_p^{\text{rel}}(f)\right] \in \mathbf{Bar}^3.$$

We now focus on the non-trivial persistence diagrams of the extended persistence of a filtered graph. As the filtered space is a graph where the highest order simplices are edges, the highest order non-trivial extended persistence diagram is $\mathsf{EPH}_1(f)$. Since the rank of the zeroth order relative homology of G with respect to any subset is at most the rank of $H_0(G)$, there can be no new intervals in the $\mathbb{R}^{\text{op}}$ part of $\mathsf{EPH}_0(f)$ and thus $\mathsf{EPH}_0^{\text{rel}}(f) = \emptyset$. On the other hand, $\mathsf{EPH}_1^{\text{ord}}(f) = \emptyset$ because the rank of $H_1(G_t)$ can only increase as we add edges to the graph along the $\mathbb{R}$ part of the filtration . Thus the extended persistent homology of a filtered graph only has four non-trivial diagrams, and we denote the map from a filtered graph to its four non-trivial diagrams as the *extended persistence map* $\mathsf{EPH} : \mathbb{R}^V \to \mathbf{Bar}^4$

$$\mathsf{EPH}(f) = \left[\mathsf{EPH}_0^{\text{ord}}(f), \mathsf{EPH}_0^{\text{ext}}(f), \mathsf{EPH}_1^{\text{ext}}(f), \mathsf{EPH}_1^{\text{rel}}(f)\right] \in \mathbf{Bar}^4. \tag{5.25}$$

Remark 5.3.1. As the extended diagrams $\mathsf{EPH}_p^{\text{ext}}(f)$ represent precisely the interval modules that factor through $H_p(G)$ in the extended filtration eq. (5.23), the number of intervals in $\mathsf{EPH}_p^{\text{ext}}(f)$ are the p^{th} Betti numbers of the graph.

Remark 5.3.2. If we only apply homology to the filtration of eq. (5.21), we get an *ordinary persistence module* indexed over the real line, which is essentially the first row in Eq. (5.23). This module is characterised by a unique barcode $\mathsf{PH}_p(f) \in \mathbf{Bar}$. We refer to the map

$$\mathsf{PH} : f \in \mathbb{R}^V \longmapsto [\mathsf{PH}_0(f), \mathsf{PH}_1(f)] \in \mathbf{Bar}^2 \tag{5.26}$$

as the *ordinary persistence map*.

We now dicuss the issue of differentiability for our loss function

$$\mathcal{L} : \mathcal{M} \xrightarrow{\mathrm{F}} \mathbb{R}^V \xrightarrow{\mathsf{EPH}} \mathbf{Bar}^4 \xrightarrow{\mathsf{Out}} \mathbb{R} \ . \qquad \text{(Equation (5.1) recalled)}$$

We call a map $\mathrm{F} : \mathcal{M} \to \mathbb{R}^V$ a *parametrisation*, as it corresponds to a selection of filter functions

over G parametrised by the manifold $\mathcal{M}$. Then $B := \mathsf{EPH} \circ \mathrm{F}$ is the barcode valued map whose differentiability properties are of interest in applications. We recall from Leygonie et al. [2019] the definition of differentiability for maps to barcodes.

Definition 5.3.3. A map $B : \mathcal{M} \to \mathbf{Bar}$ on a smooth manifold $\mathcal{M}$ is said to be differentiable at $\theta \in \mathcal{M}$ if for some neighbourhood U of θ, there exists a finite collection of differentiable maps $b_i, d_i : U \to \mathbb{R} \cup \{\infty\}$, called a *local coordinate system* for B at θ, such that

$$\forall \theta' \in U,\ B(\theta') = \left\{ \left\langle b_i(\theta'), d_i(\theta') \right\rangle \mid b_i(\theta') \neq d_i(\theta') \right\}.$$

For $N \in \mathbb{N}$, we say that a map $B : \mathcal{M} \to \mathbf{Bar}^N$ is differentiable at θ if all its components are so.

In Yim and Leygonie [2021], Jacob Leygonie showed that EPH is *generically* differentiable, i.e. differentiable on an open dense subset. This property guarantees that local gradients yield reliable descent directions in a neighbourhood of the current iterate. Thus, our loss function is smooth so long as Out and F are also smooth. We recall the statement of the proposition from Yim and Leygonie [2021] below.

Proposition 5.3.4 ([Yim and Leygonie, 2021, Proposition 3.3])**.** *Let* $\mathrm{F} : \mathcal{M} \to \mathbb{R}^V$ *be a parametrisation differentiable on an open dense subset of* $\mathcal{M}$*. Then the composition* $\mathsf{EPH} \circ \mathrm{F}$ *is also differentiable on an open dense subset of* $\mathcal{M}$*.*

5.4 Binary graph classification

We investigate whether optimising the extended persistence of wavelet signatures can be usefully applied to graph classification problems, where persistence diagrams are used as features to predict discrete, real life attributes of networks. In this setting, we aim to learn $\theta \in \mathcal{M}$ that minimise the classification error of graphs over a training dataset.

We apply our wavelet optimisation framework to classification problems on the graph datasets `MUTAG` Debnath et al. [1991], `COX2` Sutherland et al. [2003], `DHFR` Sutherland et al. [2003], `NCI1` Wale et al. [2008], Shervashidze et al. [2011], `PROTEINS` Borgwardt et al. [2005], Dobson and Doig [2003] and `IMDB-B` Yanardag and Vishwanathan [2015]. The former five datasets are biochemical molecules and `IMDB-B` is a collection of social ego networks. In our models,

we use persistence images Adams et al. [2017] as a fixed vectorisation method and use a feed forward neural network to map the persistence images to classification labels. We also include the eigenvalues of the graph Laplacian as additional features; model particulars are described in the sections below.

To illustrate the effect of wavelet optimisation on different classification problems, we also perform a set of *control experiments* where for the same model architecture, we fix the wavelet and only optimise the parameters of the neural network. The control experiment functions as a baseline against which we assess the efficacy of wavelet optimisation.

We benchmark our results with two existing persistence based architectures, `PersLay` Carrière et al. [2020] and `GFL` Hofer et al. [2020]. `Perslay` optimises the vectorisation parameters and use two heat kernel signatures as fixed rather than optimisable vertex functions for computing extended persistence. `GFL` optimises and parametrises vertex functions using a graph isomorphism network Xu et al. [2019], and computes *ordinary* sublevel and superlevel set persistence instead of extended persistence.

5.4.1 Model Architecture

We give a high level description of our model and relegate details and hyperparameter choices of the vectorisation method and neural network architecture to appendix B. In our setting, the extended persistence diagrams of the optimisable wavelet signatures for each graph are vectorised as persistence images. We also include the static persistence images of a *fixed* heat kernel signature, $W(e^{-0.1x})$, as an additional set of features, alongside some non-persistence features. Both the optimised and static persistence diagrams are transformed into the persistence images using identical hyperparameters. We feed the optimisable and static persistence images into two separate convolutional neural networks (CNNs) with the same architecture. Similarly, we feed the non-persistence features as a vector into a separate multilayer perceptron. The outputs of the CNNs are concatenated with the outputs of the multi-layer perceptron. Finally, an affine transformation sends the concatenated vector to a real number whose sign determines the binary classification.

5.4.1.1 Wavelet Parametrisation

We choose a space of wavelets spanned by 12 inverse multiquadric radial basis functions

$$h_j(x) = \left(\left(\frac{x - x_j}{\epsilon} \right)^2 + 1 \right)^{-\frac{1}{2}} \tag{5.27}$$

whose centroids x_j are located at $x_j = 2(j-1)/9,\ j = 0, \ldots, 11$. The width parameter is chosen to be the distance between the centroids, $\epsilon = 2/9$. We note that there are many different radial basis functions to choose from; both Gaussians and inverse multiquadrics perform similarly well in empirical studies Mongillo et al. [2011], Chen et al. [2005]. We refer the reader to the [Chen et al., 2005] for details on practical considerations in choosing radial basis function candidates.

On each dataset, we derive a numerically stable parametrisation using the procedure described in Section 5.2.2; the parameters we optimise are the coefficients of the new basis given by eq. (5.14). We initialise the parameters by fitting them via least squares to the heat kernel signature $W(e^{-10x})$ on the whole dataset of graphs.

5.4.1.2 Non-Persistence Features

We also incorporate the eigenvalues of the normalised Laplacian as additional, fixed features of the graph. Since the number of eigenvalues for a given graph is equal to the number of vertices, it differs between graphs in the same dataset. To encode the information represented in the eigenvalues as a fixed length vector, we first sort the eigenvalues into a time-series; we then compute the log path signature of the time series up to level four, which is a fixed length vector in $\mathbb{R}^8$. The log-signature captures the geometric features of the path; we refer readers to Chevyrev and Kormilitzin [2016] for details about path signatures. For `IMDB-B` in particular, we also include the maxima and minima of the heat kernel signatures $W(e^{-10x})$ and $W(e^{-0.1x})$ respectively of each graph.

5.4.2 Experimental set up

We employ a *10 ten-fold* test-train split scheme on each dataset to measure the accuracy of our model. Each ten-fold is a set of ten experiments, corresponding to a random partition of the dataset into ten portions. In each experiment, a different portion is selected as the test set while the model is trained on the remaining nine portions. We perform 10 ten-folds to obtain a total of 10×10 experiments, and report the accuracy of the classifier as the average accuracy over 100 such experiments. The epochs at which the accuracies were measured are specified in Table 5.3.

Across all experiments, we use binary cross entropy as the loss function. We use the `Adam` optimiser Kingma and Ba [2014] with learning rate `lr = 1e-3` to optimise the parameters of the neural network. The wavelet parameters are updated using stochastic gradient descent with learning rate `lr = 1e-2`, for all datasets except for `IMDB-B`, where the learning rate is set to $\texttt{lr} = 1e - 1$. The batch sizes for each experiment are shown in Table 5.4. In all experiments, we stop the optimisation of wavelet parameters at epoch 50 while the neural network parameters continue to be optimised.

We use the `GUDHI` library to compute persistence, and make use of the optimisation and machine learning library `PyTorch` for the construction of the graph classifications models.

5.4.3 Results and Discussion

In Table 5.2, we present the classification accuracies of our model. For each dataset, we perform four experiments using our model, varying whether the wavelet parameter is optimised and whether additional features are included. In Table 5.1, we show the test accuracy of our model alongside two persistence-based graph classification architectures, `Perslay` and `GFL`, as well as other state-of-the-art graph classification architectures. Following the presentation in Carrière et al. [2020], we record the standard deviation of the mean accuracies of the ten-folds as the error on the reported accuracy.

We first compare the performances of our model between cases where we optimise and fix the wavelets. In Table 5.2, we see that when only persistence features are used, the wavelet

optimised models perform better than the control model in MUTAG, DHFR, and NCI1, where the gain in mean accuracy exceeds one standard deviation of the control model. On other datasets, the differences in performance after optimising the wavelet lie within one standard deviation of the control model performance. When we include non-persistence features in the model, the performances of the optimised and control models is only statistically significant for MUTAG, suggesting that the non-persistence features in other datasets can offset the deficit in predictive accuracy due to an unoptimised wavelet. This also indicates that the initial wavelet signature – the heat kernel signature $W(e^{-10x})$ – may be a locally optimal choice of wavelet for our neural network classifier.

We now compare our architecture to other persistence based architectures, Perslay and GFL, where node attributes are excluded from their vertex function models. Except on PROTEINS, our wavelet optimised model matches or exceeds Perslay. While our model architecture and choice of wavelet initialisation is similar to that of Perslay, we differ in two important respects. Perslay fixes the vertex functions but optimises the weights assigned to points on the persistence diagrams, as well as the parameters of the persistence images. The mean accuracy of our model is higher than that of Perslay for MUTAG, DHFR, NCI1 1, and IMDB-B; in particular the improvements for NCI1 1 and IMDB-B exceed the one standard deviation of both models. These results indicate that vertex function optimisation yields statistically significant improvements for some datasets that cannot be obtained through vectorisation optimisation alone.

Compared to GFL (without node attributes), both Perslay and our architecture achieves similar or higher classification accuracies on PROTEINS and NCI1. In particular, the difference in performance on PROTEINS is greater than one standard deviation of our model's performance. This supports wavelet signatures being viable models for vertex functions on those datasets. On the other hand, both Perslay and our model lag behind GFL on IMDB-B. We attribute this to the fact that IMDB-B, unlike the other bionformatics datasets, consists of densely connected graphs. The graphs in IMDB-B have diameter at most two and 14% of the graphs are cliques. This fact has two consequences. First, we expect the one-layer GIN used in GFL – a local topology summary – to be highly effective in optimising for the salient features of a graph with small diameter. Second, the extended persistence modules for cliques have zero persistence, since all vertices are assigned the same function value due to symmetry. In contrast, ordinary

persistence used in GFL is able to capture the cycles in a complete graph as points with infinite persistence.

Compared to non-persistence state-of-the-art architectures in Table 5.2, our model achieves competitive accuracies on MUTAG, COX2, and DHFR. For NCI1 and PROTEINS, all persistence architectures listed that exclude additional node attributes perform poorly in comparison, though PWL was able to achieve leading results with node attributes.

All in all, we observe that wavelet signatures can be an effective parametrisation of vertex functions when we use extended persistence as features for graph classification. In particular, on some bioinformatics datasets, we show that optimising the wavelet signature can lead to improvements in classification accuracy. The wavelet signature approach is complementary to the GFL approach to vertex function parametrisation as they show strengths on different datasets.

5.5 Conclusion

We have presented a framework for equipping any graph G with a set of extended persistence diagrams $\mathsf{EPH} \circ \mathrm{F} : \mathcal{M} \to \mathbf{Bar}^4$ parametrised over a manifold $\mathcal{M}$, a parameter space for the graph's wavelet signature. We described how wavelet signatures can be parametrised and interpreted. Given a function on extended persistence diagrams $\mathsf{Out} : \mathbf{Bar}^4 \to \mathbb{R}$ that is differentiable, we have shown how a loss function $\mathcal{L} = \mathsf{Out} \circ \mathsf{EPH} \circ \mathrm{F}$ can be generically differentiable with respect to $\theta \in \mathcal{M}$ as $\mathcal{L}$. Thus, we can apply gradient descent methods to optimise the extended persistence diagrams of a graph to minimise $\mathcal{L}$.

We applied this framework to a graph classification architecture where the wavelet signature is optimised for classification accuracy. We are able to demonstrate an increase in accuracy on several benchmark datasets where the wavelet is optimised, and perform competitively with state-of-the-art persistence based graph classification architectures.

	Non-Persistence State-of-the-Art				Persistence Based					
	P-SAN	RetGK	GIN	FGSD	PWL	GFL		Perslay	Control	Wavelet Opt.
									This model	
Node attr.	Yes		No		Yes		No			
MUTAG	92.6	90.3±1.1	89.4	92.1	90.5±1.3	–	–	89.8±0.9	89.0±0.6	**90.4**±1.3
COX2	–	81.4±0.6	–	–	–	–	–	**80.9**±1.0	80.8±0.4	80.8±1.0
DHFR	–	82.5±0.8	–	–	–	–	–	80.3±0.8	80.3±0.9	**81.0**±0.9
NCI1	78.6	84.5±0.2	82.7	79.8	85.6±0.3	77.2	71.2	73.5±0.3	74.3±0.3	**74.4**±0.3
PROTEINS	75.9	78.0±0.3	76.2	73.4	75.9±0.8	73.4	74.1	**74.8**±0.3	74.5±0.4	74.6±0.6
IMDB-B	71.0	72.3±0.6	75.1	73.6	73.0±1.0	–	**74.5**	71.2±0.7	71.6±0.9	72.0±0.7
# Ten-folds	10	10	1	1	10	1	1	10	10	10

Table 5.1: Binary classification accuracy on datasets of graphs. The best accuracy of persistence-based models *without* using node attributes is made bold for each dataset. The performance of our model is reported in the column Wavelet Opt. on the right hand side. The accuracies of the control model, where the wavelet parameters are fixed to the initial values, are shown in the column Control. Both these models use additional features (see Section 5.4.1.2). The accuracies of our model are the means over 10 ten-folds, recorded at epochs reported in Table 5.3. We also provide the standard deviations of the 10 mean accuracies of each ten-fold. For other architectures, we indicate whether their accuracies were reported as averages over 1 ten-fold or 10 ten-fold in the bottom row of the table. To avoid confusion, we leave out the errors reported for P-SAN Niepert et al. [2016], GINXu et al. [2019] and GFL Hofer et al. [2019] and refer the reader to the original sources, as they were calculated using a different formula. Errors were not reported in Verma and Zhang [2017] for FGSD. We refer the reader to Zhang et al. [2018], Rieck et al. [2019], and Carrière et al. [2020] for further details regarding the results of RetGK, PWL, and PersLay respectively.

	Persistence Only		Non-Persistence Features incl.	
	Control	Wavelet Opt.	Control	Wavelet Opt.
MUTAG	89.2±0.6	89.8±0.8	89.0±0.6	90.4±0.4
COX2	79.6±1.0	79.4±0.7	80.8±1.0	80.8±1.0
DHFR	79.9±0.4	80.4±0.4	80.3±0.9	81.0±0.9
NCI1	73.7±0.2	74.3±0.5	74.3±0.3	74.4±0.3
PROTEINS	72.9±0.3	73.0±0.4	74.5±0.4	74.6±0.6
IMDB-B	68.3±0.5	68.6±0.7	71.6±0.9	72.0±0.7

Table 5.2: Binary classification accuracy of our model where we vary whether non-Persistence features are included and whether the wavelet is optimised. The reported accuracies are the mean over 10 ten-folds, recorded at epochs reported in Table 5.3. We also provide standard deviations of the 10 mean accuracies of each ten-fold. See Section 5.4.1.2 for the particulars about the non-persistence features.

	Persistence Only		Non-Persistencee Features incl.	
	Control	Wavelet Opt.	Control	Wavelet Opt.
MUTAG	25	125	25	75
COX2	50	50	25	25
DHFR	125	250	125	45
NCI1	270	270	500	370
PROTEINS	50	50	125	125
IMDB-B	100	25	75	50

Table 5.3: Epochs at which accuracies in Table 5.2 are recorded.

	MUTAG	COX2	DHFR	NCI1	IMDB-B
# graphs	188	467	756	4110	1000
batch size	10	9	11	20	50

Table 5.4: Batch sizes in the graph classification experiments for different datasets described in section 5.4.

Chapter 6

Assessing the Topical Diversity of Research Grants

6.1 Introduction

In this chapter, we describe a method for measuring *multi-disciplinarity* in academic research, as a part of research supervised by Elsevier. Elsevier is interested in helping academic stakeholders from individual researchers to major institutions better understand their research output within their context in academia, either through providing data-analytics tools that interface with their academic database Scopus, or professional advice proffered by Elsevier's consultants.

Within this wider mission, Elsevier is interested in developing methods of assessing the multi-disciplinarity of a body of research. As interactions between previously disparate scientific disciplines has led to fruitful new areas of scientific discovery in recent decades, multi-disciplinarity has become a relevant metric for profiling the research portfolios of institutions, such as universities and funding bodies. If Elsevier were to advise institutions on such matters, they need a mathematically principled way of comparing the multi-disciplinarity of one body of research to another. In this study, we focus on assessing the multi-disciplinarity of research grants. This is of keen interest to institutions in academia. For example, a measure of multi-disciplinarity of research grants allows universities and funding agencies to assess the ratio of multi-disciplinarity versus 'pure' research that is undertaken in the work they sup-

port. Our method of assessing multi-disciplinarity is designed to make use of the available data attached to research funds in Elsevier's Scopus and SciVal databases. Elsevier catalogues research grants referenced by published articles and associates a set of topics to each grant. These topics are generated by Elsevier's SciVal analytics software [Klavans and Boyack, 2017], based primarily on citation analysis of research articles indexed in Scopus. SciVal also constructs a graph X between topics where two topics are linked by an edge if the citation analysis deems them to be related. To each topic, SciVal associates a set of keywords which we use to interpret the topics learnt by SciVal's algorithms. We note that while SciVal focuses on using citation links to infer topic similarity, there are other attributes of research articles, such as coauthor-ship that can help inform topic similarity.

A naïve measure of the multi-disciplinarity of a grant is the number of topics associated to the grant. However, this fails to account for the similarities between these topics. Heuristically, if the topics associated to an grant are all closely related, then they should have a lower topical diversity than grants involving unrelated topics. The similarity data between topics associated to grant α is encoded as a subgraph $X_\alpha \subset X$ where the vertices of X_α are the topics associated to the grant α, and the edges are the all the edges between such topics in the topic graph X. A multi-disciplinarity measure would take the graph X_α and output a real number between 1 and the number of vertices that can be interpreted as the effective number of topics associated to the grant.

Our solution is to apply an invariant in geometric measure theory called *maximum diversity* [Leinster and Meckes, 2016], which can be interpreted as a measure of the effective cardinality of a set with similarity. A set of topics has a lower maximum diversity than that of another set of topics of the same cardinality that are less similar to each other. In Section 6.2, we will recall some properties of maximum diversity that conform to a number of common sense conditions that one would expect for a proxy for cardinality.

We have also applied maximum diversity to derive a new way of assessing the similarity of two grants. In our approach to this problem, we use the interpretation of maximum diversity as an effective cardinality to generalise the Jaccard index, a notion of similarity between sets based on enumerating the elements shared between them. Compared to the Jaccard index, our generalisation takes the similarities between the underlying objects in the sets into account.

We prove that our generalised Jaccard index satisfy a set of common sense conditions for a similarity measure in Section 6.3.

We applied our methods to a collection of EPSRC grants catalogued by Elsevier. Our primary aim is to assess the significance of topical relations in practice in affecting our assessment of multi-disciplinarity of grants, as well as the comparison of grants. We do so by comparing the maximum diversity against the number of topics for each grant, and the Jaccard index against its generalisation. We discuss our findings in Section 6.4.

6.2 Maximum Diversity

We first define the maximum diversity of a similarity matrix and note a few properties that make maximum diversity a desirable proxy for the effective cardinality of a set with similarity. We then discuss its relation to generalised notions of entropy and computational considerations. We refer the reader to Leinster and Meckes [2016], Leinster and Roff [2019] for further background on the basic definitions and the proofs for the theoretical theorems and lemmas quoted here.

6.2.1 Basic Definitions

We restrict our notion of similarity between two topics to those that satisfy the following condition.

Definition 6.2.1 (Similarity)**.** A similarity matrix Z on a finite set X is a real symmetric matrix where

1. $Z(i,i) = 1$
2. $Z(i,j) \in [0,1]$.

We call (X, Z) a finite set equipped with a similarity matrix Z. For a fixed underlying set X, we have a partial order $\preceq$ over all similarity matrices on X, where $Z \succeq S$ if all entries of $S - Z$ are non-negative.

We build up to the concept of the maximum diversity of a finite finite set equipped with a similarity matrix Z by considering a special class of functions $\omega : X \to \mathbb{R}$.

Definition 6.2.2 (Balanced Distribution [Leinster and Roff, 2019, Remark 4.7])**.** Let (X, Z) be a finite set equipped with a similarity matrix Z. Then a *balanced distribution* of X is a vector $\omega \in \mathbb{R}^X$ such that $(Z\omega)_i = 1$ for all $i \in X$. We say ω is a balanced, positive distribution $\omega > 0$ if $\omega_i > 0$ for all $i \in X$.

We can interpret $(Z\omega)_i$ as the amount of mass of the weight ω close to i. If ω is a balanced distribution, while the mass ω_i on each i may be different, the amount of mass close to i is the same for every $i \in X$. Balanced distributions enjoy the following invariant property.

Lemma 6.2.3 ([Leinster and Roff, 2019, Lemma 5.3])**.** *For any two balanced distributions μ and ν of X,*

$$\sum_{i \in X} \mu = \sum_{i \in X} \nu.$$

Definition 6.2.4 (Magnitude)**.** If a finite set equipped with a similarity matrix Z (X, Z) admits a balanced distribution ω, the magnitude of X is

$$|X| = \sum_{i \in X} \omega_i$$

If Z is invertible, then one can show that (X, Z) has a unique balanced distribution $\omega_i = \sum_j \left(Z^{-1}\right)_{ij}$, and the magnitude of (X, Z) is $\sum_{ij} \left(Z^{-1}\right)_{ij}$. This conforms with our intuition of a good proxy for effective size as a quantity that scales inversely with the similarities between elements. However magnitude is not defined for all spaces with similarity, as the set of balanced distributions of X can be empty. Moreover, magnitude can be negative, or even increase with increasing similarity. Maximum diversity is an adaptation of magnitude which eliminates the properties of magnitude that goes against our intuition of effective size.

Definition 6.2.5 (Maximum Diversity (See Leinster and Meckes [2016]))**.** Let $\Omega(X, Z) \subset \mathbb{R}^X_{\geq 0}$ be the set of non-negative distributions on X that are positive and balanced on some subset of X. The maximum diversity $|X|_+$ of (X, Z) is then the largest magnitude over all subsets of X that admit a positive and balanced distribution:

$$|(X, Z)|_+ = \sup_{\omega \in \Omega(X,Z)} \underbrace{\sum_i \omega_i}_{=|\text{supp } \omega|}.$$

We define $\Omega_+(X, Z) = \left\{\omega \in \Omega(X, Z) \mid \sum \omega = |(X, Z)|_+\right\}$ to be the weight distributions in

$\Omega(X, Z)$ that realise the maximum diversity of X.

Unlike magnitude, the maximum diversity of a finite set equipped with a similarity matrix Z is always well-defined. Since the magnitude of any point is simply 1, there is at least a subset of X that has well-defined magnitude.

If we interpret maximum diversity as a proxy for the effective size of X, for a given weight $\omega \in \Omega(X, Z)$, each entry ω_i can be viewed as a contribution from element i of X to the effective size of X under ω. The following lemma shows that the individual contributions of each element cannot be greater than the maximum diversity of a single point.

Lemma 6.2.6. *Let (X, Z) be a finite set equipped with a similarity matrix Z. For $\omega \in \Omega_+(X, Z)$, we have $\omega_i \in [0, 1]$ for every $i \in X$.*

Proof. By definition, $\omega \geq 0$. Suppose i is in the support of ω. From Definition 6.2.5, ω is a positive balanced distribution on its support, and it follows from the definition of a balanced distribution Definition 6.2.2 that $(Z\omega)_i = 1$. Because we have constrained $Z_{ii} = 1$, and all Z_{ij} and ω_j are non-negative,

$$1 = \sum_j Z_{ij}\omega_j = \omega_i + \sum_{j \neq i} Z_{ij}\omega_j \implies \omega_i = 1 - \sum_{j \neq i} Z_{ij}\omega_j \leq 1.$$

□

Corollary 6.2.7. *Let (X, Z) be a finite set equipped with a similarity matrix Z. Then*

$$|X|_+ \leq \inf_{\omega \in \Omega_+(X,Z)} \#(\operatorname{supp} \omega).$$

Proof. Choose a weight $\omega \in \Omega_+(X, Z)$ with the fewest elements in its support among all weights in $\Omega_+(X, Z)$, the set of weights that realise the maximum diversity of X . Since $\omega_i \leq 1$,

$$|X|_+ = \sum_i \omega_i = \sum_{i \in \operatorname{supp}\ \omega} \omega_i \leq \#(\operatorname{supp} \omega).$$

□

We would like to understand why some objects $i \in X$ are accorded zero weight $\omega_i = 0$ for

$\omega \in \Omega_+(X, Z)$ and $|X|_+ = |\text{supp}\omega|_+$ for $\text{supp}\omega \subsetneq X$. We start with the following observation.

Lemma 6.2.8. *Let X be a finite set equipped with a similarity matrix Z. If $\omega_i = 0$ for a weight $\omega \in \Omega_+(X, Z)$ that realises the maximum diversity of (X, Z), then $(Z\omega)_i \geq 1$.*

Proof. Suppose $\omega_i = 0$ and $1 > (Z\omega)_i$. Then let $\hat{\omega}_j = \omega_j$ for $j \in \text{supp}\omega$ and $\hat{\omega}_i = 1 - (Z\omega)_i > 0$. Then $\hat{\omega}$ is balanced on $\text{supp}\omega \cup \{i\}$ and $\sum \hat{\omega} > \sum \omega$, which implies ω is not a set of weights that realise the maximum diversity of X. Hence we reach a contradiction and conclude that $\omega_i = 0$ implies $(Z\omega)_i \geq 1$. □

Given a set of weights ω on X, we can regard $(Z\omega)_i = \sum_j Z_{ij}\omega_j$ as a heuristic measure of the concentration of mass around i. Thus, the observation that $(Z\omega)_i \geq 1$ indicates that $i \notin \text{supp}\omega$ has an excessive concentration of mass. The balance constraint $(Z\omega)_i = 1$ on the support of ω encourages allocating weight to elements dissimilar to each other in the process of maximising the sum of weights. The weights can thus be interpreted as a way of spreading out mass on X in a 'maximally diverse' manner. If $\omega_i = 0$ for $\omega \in \Omega_+(X, Z)$, then we can intuitively understand i to be so similar to other elements in X that despite spreading out the mass in a maximally diverse way, the concentration of mass around i still exceeds those in the support. In our application to measuring the multi-disciplinarity of a research grant, if we view the weights on topics as the contribution of each topic to the effective topical content of the grant, then the zero weight topics can be interpreted as those that are so similar to other topics in the grant that its inclusion does not meaningfully increase the grant's topical content.

This intuition is codified more formally in the connection between maximum diversity and generalised notions of entropy called diversity indices.

6.2.2 Computing Maximum Diversity via Diversty Indices

The combinatorial definition of maximum diversity in Definition 6.2.5 makes it a rather unwieldy quantity to compute on large sets, as we have to attempt to compute the magnitude of in principle all subsets in X. As it turns out, the maximum diversity of spaces with similarity are related to generalised notions of entropy called diversity indices; not only does this offer a new perspective on maximum diversity, it allows us to compute maximum diversity us-

ing convex quadratic programming in the case where the similarity matrix is positive definite, thus allowing us to bypass computing the magnitude of a combinatorially prohibitive number of subsets.

Definition 6.2.9 (Diversity indices)**.** For $q \in [0,\infty]$, the q-diversity $D_X^q : \Delta^X \to \mathbb{R}$ of a probability distribution $p \in \Delta^X$ on a finite set equipped with a similarity matrix Z (X, Z) is given by

$$D_Z^q(p) = \begin{cases} \left\langle (Zp)^{q-1} \right\rangle^{1/1-q} & q \neq 1, \infty \\ \exp(-\langle \log(Zp) \rangle) & q = 1 \\ 1/\max_{i \in \operatorname{supp} p} (Zp)_i & q = \infty. \end{cases}$$

where $\langle \cdot \rangle$ denotes the expected value over p.

We refer the reader to Leinster and Meckes [2016] on how the logarithm of diversity indices for different choices of q coincide with various notions of entropy if $Z = I$. For example, if we set $q = 1$, we recover the Boltzmann entropy $\log\big(D_I^1(p)\big) = -\sum_i p_i \log(p_i)$. The following theorem relates the maximisation of diversity indices of probabilities on Δ^X to the maximum diversity of the space itself.

Theorem 6.2.10 ([Leinster and Meckes, 2016, Theorem 1 & 2] and [Leinster and Roff, 2019, Corollary 7.4])**.** *Let (X, Z) be a finite set equipped with a similarity matrix Z. There exists a probability distribution $\hat{p} \in \Delta^X$ that maximises the diversity indices for all $q \in [0,\infty]$,*

$$\sup_{p \in \Delta^X} D_Z^q(p) = D_Z^q(\hat{p}),$$

and the maximal value of the diversity indices $D_Z^q(\hat{p})$ are the same for all $q \in [0,\infty]$, and equal to the maximum diversity of X:

$$D_Z^q(\hat{p}) = |X|_+.$$

Furthermore, $\hat{\omega} = |X|_+ \hat{p} \in \Omega_+(X, Z)$, i.e. $\hat{\omega}$ is a set of weights that realises the maximum diversity of X.

This theorem implies we can compute maximum diversity and obtain a set of weights $\omega \in$

$\Omega_+(X, Z)$ by maximising the D_Z^2 diversity index, which is a quadratic programming problem:

$$\frac{1}{|X|_+} = \inf_{p \in \Delta^X} p^\intercal Z p. \tag{6.1}$$

Intuitively, a probability distribution $\hat{p}$ that maximises D_Z^2 is one such that the similarity between any randomly pair of objects independently drawn from $\hat{p}$ with replacement is minimised. We can regard $\hat{p}$ as the maximally 2-diverse distribution. Moreover, Theorem 6.2.10 implies that no other probability distribution is more q-diverse than $\hat{p}$.

If Z is positive definite, minimising D_Z^2 is a convex optimisation problem which has a unique solution. While not all similarity matrices are positive definite, we can always rescale off-diagonal elements of any similarity matrix Z such that Z is diagonally dominant and therefore positive definite in practice.

Remark 6.2.11. If X is a set of topics of interest to a funding agency, we can view a solution to the maximum diversity problem $\hat{p}$ as a way of distributing resources that corrects for the 'knock-on' effect of funding on supporting other similar topics. However, we note that knock-on effects between resources allocated to a pair of topics are difficult to mathematically model and quantify. For example, they are likely to be asymmetric and conditional on case by case factors. Thus the maximum diversity distribution computed from a symmetric similarity matrix may yield misleading results as the data encoded in thee matrix is a poor approximation of knock-on effects.

6.2.3 Theoretical Properties

We now recall two monotonicity properties of maximum diversity from Leinster and Meckes [2016], and show that maximum diversity is additive with respect to the union of sets if two subsets A and B share no elements and are totally dissimilar. Using these three properties, we show that maximum diversity is subadditive with respect to the union of sets, and derive necessary and sufficient conditions for maximum diversity to be additive.

Theorem 6.2.12 ([Leinster and Roff, 2019, Corollary 7.5]). *Let (X, Z) be a finite set equipped with a similarity matrix Z. Then maximum diversity is monotonic with respect to the partial order induced by inclusion of subsets:*

$$Y \subseteq X \implies |Y|_+ \leq |X|_+.$$

Theorem 6.2.13 ([Leinster and Roff, 2019, Proposition 7.6]). *Let (X, Z) and (X, S) be two spaces with similarity with the same underlying set X. Then*

$$Z \preceq S \implies |(X, Z)|_+ \geq |(X, S)|_+.$$

Corollary 6.2.14. *For any finite set equipped with a similarity matrix Z (X, Z),*

$$1 \leq |(X, Z)|_+ \leq |(X, I)|_+ = \#X.$$

Lemma 6.2.15. *Let A and B be subsets of (X, Z) where $A \cap B = \emptyset$ and $Z_{ij} = 0$ for all pairs of elements $i \in A$ and $j \in B$. Then*

$$|A \cup B|_+ = |A|_+ + |B|_+.$$

Proof. This can be shown by straightforward computation of $|A \cup B|_+$ as the maximum of the diversity index $D_Z^2(p)$, using the fact that Z restricted to $A \cup B$ can be decomposed as a block diagonal matrix $Z^A \oplus Z^B$ where the block Z^A corresponds to entries between elements in A and similarly for block Z^B. Making the substitution $p = \sigma p_A + (1 - \sigma) p_B$ where $\sigma \in [0, 1]$, with $p_A \in \Delta^A$, and $p_B \in \Delta^B$:

$$\begin{aligned}
\frac{1}{|A \cup B|_+} &= \inf_{p \in \Delta^{A \cup B}} p^\intercal Z p \\
&= \inf_{\sigma \in [0,1]} \inf_{\substack{p_A \in \Delta^A \\ p_B \in \Delta^B}} \left(\sigma^2 p_A^\intercal Z p_A + (1 - \sigma)^2 p_B^\intercal Z p_B \right) \\
&= \left(\left(\inf_{p_A \in \Delta^A} p_A^\intercal Z p_A \right)^{-1} + \left(\inf_{p_B \in \Delta^B} p_B^\intercal Z p_B \right)^{-1} \right)^{-1} \\
&= \frac{1}{|A|_+ + |B|_+}.
\end{aligned}$$

□

Proposition 6.2.16. *Maximum diversity is subadditive. Let A and B be subsets of (X, Z). Then*

$$|A \cup B|_+ \leq |A|_+ + |B|_+.$$

Proof. We use the fact that maximum diversity is monotonic with respect to inclusion (Theo-

rem 6.2.12) and deduce

$$|A|_+ + |B|_+ \geq |A \setminus B|_+ + |B|_+.$$

Let $Z^{A\setminus B}$ and Z^B be the restriction of Z to $A \setminus B$ and B respectively. Since $A \setminus B$ and B are disjoint, we can apply Lemma 6.2.15:

$$|A \setminus B|_+ + |B|_+ = \left|\left(A \cup B, Z^{A\setminus B} \oplus Z^B\right)\right|_+.$$

Because $Z^{A\setminus B} \oplus Z^B$ is the matrix Z with all entries outside diagonal blocks corresponding to $A \setminus B \times A \setminus B$ and $B \times B$ sent to zero, $Z \succeq Z^{A\setminus B} \oplus Z^B$. By the monotonicity of maximum diversity with respect to the partial order on similarity matrices (Theorem 6.2.13),

$$\left|\left(A \cup B, Z^{A\setminus B} \oplus Z^B\right)\right|_+ \geq |(A \cup B, Z)|_+ = |A \cup B|_+.$$

Linking all the inequalities above yields the desired result. □

Proposition 6.2.17. *Let A and B be subsets of (X, Z). We have*

$$|A \cup B|_+ = |A|_+ + |B|_+$$

if and only if $|A|_+ = |A \setminus B|_+$ and $|B|_+ = |B \setminus A|_+$, and

$$|A \cup B|_+ = |A \setminus B|_+ + |B \setminus A|_+.$$

Proof. Assume $|A|_+ + |B|_+ = |A \cup B|_+$. Since maximum diversity is subadditive (Proposition 6.2.16),

$$|A|_+ + |B|_+ = |A \cup B|_+ \leq |A \setminus B|_+ + |B|_+ \implies |A|_+ \leq |A \setminus B|_+.$$

Since maximum diversity is monotonic with respect to inclusion (Theorem 6.2.12),

$$|A \setminus B|_+ \leq |A|_+ \leq |A \setminus B|_+ \implies |A|_+ = |A \setminus B|_+.$$

Exchanging A and B, we can also show that $|B|_+ = |B \setminus A|_+$.

The converse is obvious. □

Example 6.2.18. Consider the path graph of 4 vertices $P_4 = (\bullet - \bullet - \bullet - \bullet)$ where vertices connected by an edge have similarity $\sigma \in [0,1]$ and vertices not connected by an edge have zero similarity. Then the similarity matrix of this space with similarity is

$$Z = \begin{bmatrix} 1 & \sigma & 0 & 0 \\ \sigma & 1 & \sigma & 0 \\ 0 & \sigma & 1 & \sigma \\ 0 & 0 & \sigma & 1 \end{bmatrix}$$

where $\sigma \in [0,1]$. The maximum diversity of P_4 is obtained by solving the quadratic programme eq. (6.1)

$$\frac{1}{|P_4|_+} = \inf_{p \in \Delta^4} p^\intercal Z p.$$

Let us consider the different cases where p is supported on different subsets of P_4. Theorems 6.2.12 and 6.2.13 define a partial order between the subgraphs of P_4, and one can check that the maximal elements of the poset of subgraphs ordered by increasing maximum diversity are

$$(\bullet - \bullet - \bullet - \bullet), \qquad (\bullet - \bullet \quad \bullet), \qquad (\bullet \quad \bullet - \bullet), \quad \text{and} \qquad (\bullet \quad \bullet).$$

One can explicitly compute that the distributions that minimise $p^\intercal Z p$ supported on the four subgraphs above are respectively

$$\begin{aligned} p &= \begin{bmatrix} \frac{1}{4-2\sigma} & \frac{1-\sigma}{4-2\sigma} & \frac{1-\sigma}{4-2\sigma} & \frac{1}{4-2\sigma} \end{bmatrix} & p^\intercal Z p &= \frac{1+\sigma-\sigma^2}{4-2\sigma}, \text{ or} \\ p &= \begin{bmatrix} \frac{1}{3+\sigma} & \frac{1}{3+\sigma} & 0 & \frac{1+\sigma}{3+\sigma} \end{bmatrix} & p^\intercal Z p &= \frac{2+\sigma}{3+\sigma}, \text{ or} \\ p &= \begin{bmatrix} \frac{1+\sigma}{3+\sigma} & 0 & \frac{1}{3+\sigma} & \frac{1}{3+\sigma} \end{bmatrix} & p^\intercal Z p &= \frac{2+\sigma}{3+\sigma}, \text{ or} \\ p &= \begin{bmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{bmatrix} & p^\intercal Z p &= \frac{1}{2}. \end{aligned}$$

We thus have three candidate expressions for the maximum diversity of P_4. From the definition of maximum diversity ,

$$|P_4|_+ = \inf_{p \in \Delta^4} p^\intercal Z p = 1/\min\left\{\frac{1+\sigma-\sigma^2}{4-2\sigma}, \frac{2+\sigma}{3+\sigma}, \frac{1}{2}\right\} = \begin{cases} \frac{4-2\sigma}{1+\sigma-\sigma^2} & \sigma \in \left[0, \frac{\sqrt{5}-1}{2}\right] \\ \frac{3+\sigma}{1+\sigma} & \sigma \in \left[\frac{\sqrt{5}-1}{2}, 1\right] \end{cases}. \tag{6.2}$$

We note that at $\sigma = \frac{\sqrt{5}-1}{2}$, the probability distribution p that minimising $p^\intercal Z p$ discontinuously changes from one with full support to two possible distributions supported on three vertices.

6.3 Jaccard Similarity

We now use maximum diversity as a means of computing a notion of similarity between subsets of a space with similarity. We take inspiration from the traditional Jaccard similarity.

Definition 6.3.1 (Jaccard Similarity)**.** The Jaccard similarity between two finite non-empty sets X and Y is given by

$$J(X,Y) = \frac{\#(X \cap Y)}{\#(X \cup Y)} = \frac{\#X + \#Y - \#(X \cup Y)}{\#(X \cup Y)} = 1 - \frac{\#(X \vartriangle Y)}{\#(X \cup Y)}.$$

wherre $X \vartriangle Y$ denotes the symmetric difference between the two sets.

The Jaccard index does not take similarities between objects into account. If X and Y are subsets of some space with similarity, we now replace the cardinality of each set in the definition with its maximum diversity. However, we note that unlike cardinality,

$$|X \cap Y|_+ \neq |X|_+ + |Y|_+ - |X \cup Y|_+ \neq |X \vartriangle Y|_+ \neq |X \cup Y|_+ - |X \vartriangle Y|_+.$$

We observe that $|X \cap Y|_+$ does not take similarities between pairs $i \in X$ and $j \in X$ that are not both in $X \cap Y$ into account. On the other hand, the symmetric difference definition fails the common sense check

$$1 - \frac{|X \vartriangle Y|_+}{|X \cup Y|_+} = \frac{|X|_+ - |X \setminus Y|_+}{|X|_+} \neq \frac{|Y|_+}{|X|_+}$$

while one can check that the following definition does.

Definition 6.3.2 (Generalised Jaccard Similarity)**.** Let (X, Z) be a space with similarity. Let A and B be non-empty subsets of X. Then we define the generalised Jaccard similarity between A and B to be

$$J_Z(A,B) = \frac{|A|_+ + |B|_+ - |A \cup B|_+}{|A \cup B|_+}.$$

Intuitively, two subsets A and B of a finite set equipped with similarity are similar under this notion of similarity if the combination of A and B does not increase the effective size relative

to A and B considered individually. We can interpret this as A and B having similar content and thus their resultant increase in effective size is less than the sum of its parts.

Proposition 6.3.3. *For non-empty subsets A and B of X, we have*

1. $J_Z(A,B) \in [0,1]$;
2. $J_Z(A,B) = 1$ *iff* $|A|_+ = |B|_+ = |A \cup B|_+$.
3. $J_Z(A,B) = 0$ *iff* $|A|_+ = |A \setminus B|_+$, *and* $|B|_+ = |B \setminus A|_+$, *as well as*

$$|A \cup B|_+ = |A|_+ + |B|_+ = |A \setminus B|_+ + |B \setminus A|_+.$$

Proof.

1. The generalised Jaccard similarity is non-negative due to the fact that maximum diversity is subadditive (Proposition 6.2.16). We proceed to derive the upper bound on $J_Z(A,B)$. Suppose, without loss of generality, that $|A|_+ \geq |B|_+$. Then

$$\begin{aligned} J_Z(A,B) &= \frac{|A|_+ + |B|_+ - |A \cup B|_+}{|A \cup B|_+} \\ &\leq 2\frac{|A|_+}{|A \cup B|_+} - 1 \\ &\leq 2\frac{|A|_+}{|A|_+} - 1 = 1. \end{aligned}$$

   where we apply the monotonicity of maximum diversity with respect to inclusion of subsets (Theorem 6.2.12) to obtain the final inequality.

2. The implication $|A|_+ = |B|_+ = |A \cup B|_+ \implies J_Z(A,B) = 1$ follows directly from computation. We now consider the converse. If $J_Z(A,B) = 1$, then

$$\left(|A|_+ - |A \cup B|_+\right) + \left(|B|_+ - |A \cup B|_+\right) = 0.$$

   By monotonicity of maximum diversity with respect to inclusions (Theorem 6.2.12), each bracketed term on the left hand side is negative or zero. Since the right hand side is zero, the bracketed terms must be zero, i.e. $|A|_+ = |B|_+ = |A \cup B|_+$.

3. This is a consequence of Proposition 6.2.17.

□

We can interpret different limits of the generalised Jaccard index as thus. Two subsets A and B are deemed maximally similar $J_Z(A,B) = 1$ if their union does not increase the effective size of either subsets. From the definition of maximum diversity (Definition 6.2.5), there is at least one diversity maximising weight measure for $A \cup B$ that is supported on some subset of A, and another that is supported on some subset of B. If the diversity maximising weight measure on $A \cup B$ is unique, then this implies all this weight measure is supported on a subset of $A \cap B$; if the diversity maximising weight measure on $A \cup B$ is unique and has full support, then $J_Z(A,B) = 1$ if and only if $A = B$.

On the other hand, A and B are deemed maximally dissimilar if the combination of A and B yields the maximum increase in effective size. The effective size of $A \cup B$ is equivalent to that of the disjoint union of $A \setminus B$ and $B \setminus A$ where no pair of elements $i \in A \setminus B$ and $j \in B \setminus A$ are similar. Furthermore, $A \cap B$ has no contribution to the effective size of A nor B, a consequence of elements in $A \cap B$ being too similar to those in $A \setminus B$ and $B \setminus A$ or empty. Thus, A and B can be maximally dissimilar if $A \setminus B$ and $B \setminus A$ dissimilar despite being similar to elements in $A \cap B$.

6.4 Case study: EPSRC Grant Data

We apply maximum diversity and our generalised Jaccard index to analyse Elsevier's proprietary dataset of 10206 grants awarded by EPSRC. For each grant, we have an associated set of 14489 topics referenced by these grants; we are also given a set of relations between topics, which can be expressed as an unweighted graph between topics. The majority of grants reference few topics; we can see in Figure 6.1 that about 74% of grants reference fewer than 5 topics and 79% of grants reference fewer than 6 topics.4

6.4.1 The Multi-Discipliniarity of Research Grants

We first examine the impact of topical relations on our assessment of multi-disciplinarity of grants. We would like to compare multi-disciplinarity as measured by maximum diversity with the naïve approach of counting the number of topics. We model the similarity between

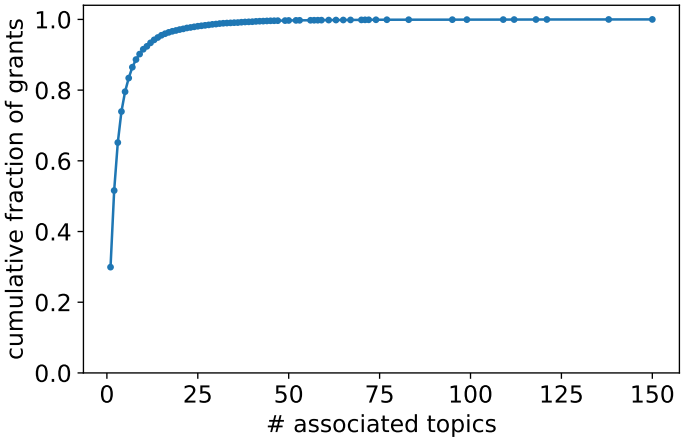

Figure 6.1: The cumulative distribution of the number of topics associated to EPSRC grants in Elsevier's catalogue.

two topics with a global similarity parameter $\sigma \in [0,1]$, where $\sigma = 0$ corresponds to no similarity and $\sigma = 1$ corresponds to maximum similarity. Thus, the similarity matrix between topics is given by

$$Z(\sigma) = I + \sigma A \tag{6.3}$$

where A is the adjacency matrix of the topic graph and I is the identity matrix on the set of topics. As a significant majority of grants have fewer than five associated topics, the similarity matrices associated to these grants are sufficiently small such that we can obtain explicit expressions of the maximum diversity of such grants as a function of σ. We plot the the maximum diversity of the 18 possible graphs (up to isomorphism) of size less than or equal to four as a function of σ in fig. 6.2. We have labelled the graphs in this figure by their degree distribution as this uniquely identifies them for graphs with fewer than five vertices. As we can see, for large enough σ, the relative size of the maximum diversity between two graphs is not determined by the number of vertices or the number of edges; for example, for large σ, the Y shaped graph with degrees $(1,1,1,3)$ has a larger maximum diversity than that of the $|\ |$-shaped graph with degrees $(1,1,1,1)$, despite the former having an extra edge. The complete graphs all decrease with σ until they attain a maximum diversity of a single vertex at $\sigma = 1$.

We estimate the chance of cardinality being a practically misleading method of computing the probability eq. (6.4) that a grant associated with four topics has a lower or equivalent

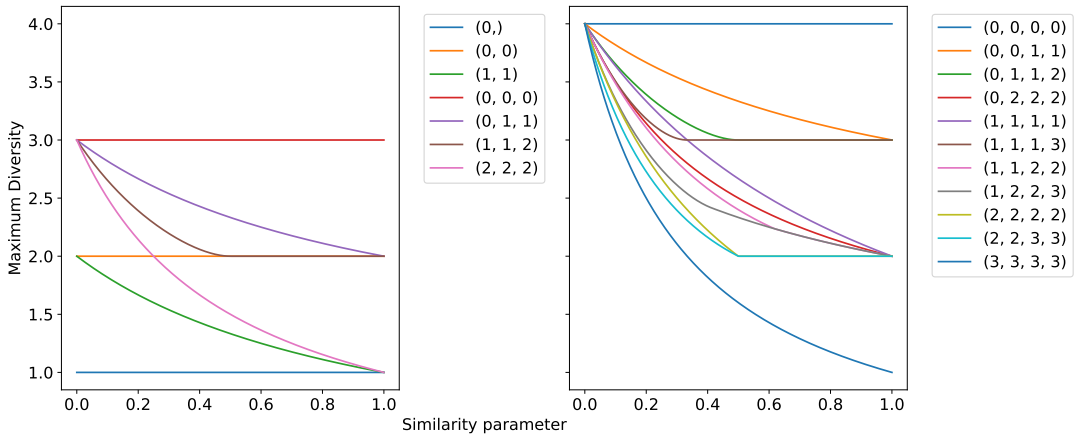

Figure 6.2: The maximum diversity of all graphs with fewer than five vertices, where where two vertices connected by an edge is assumed to have similarity σ, and two vertices not connected by an edge is assumed to have zero similarity. We label these 18 graphs by their degree sequence, which uniquely identifies graphs with fewer than five vertices.

maximum diversity than a grant associated with three topics

$$\Pr\left(|X_\alpha|_+ \leq |X_\beta|_+ \quad | \quad \#X_\alpha = 4 \quad \text{and} \quad \#X_\beta = 3\right) \tag{6.4}$$

as a function of the similarity parameter σ. This probability depends on the relative abundance of the different types of graphs in the full collection, which we show in Figure 6.3. We plot the probability eq. (6.4) against σ in Figure 6.4, and note that for a large range of σ, the probability of cardinality being a misleading indication of the relative sizes of the multi-disciplinarity of grants associated with three and four topics is about 10%.

We then computed the maximum diversity of all the grants in our collection for $\sigma = \frac{1}{8}$ and $\sigma = \frac{1}{3}$. We plot the maximum diversity of against their number of topics in Figure 6.5, and note that there are some grants where the number of topics overestimates its effective sizes by potentially two-fold. In particular, we observe that the number of topics seem is more likely to be a poor predictor for the multi-disciplinarity of grants that reference many topics. This is hardly surprising, as it is practically unlikely for a grant to encompass a wide range of unrelated research avenues.

As a sanity check, we examined the three grants with the highest topic to maximum diversity ratio for $\sigma = \frac{1}{3}$ and $\frac{1}{8}$ and display their associated topic keywords in Table 6.1. The top three grants for $\sigma = \frac{1}{3}$ and $\frac{1}{8}$ in fact are the same. Individually, we can see that they reference a

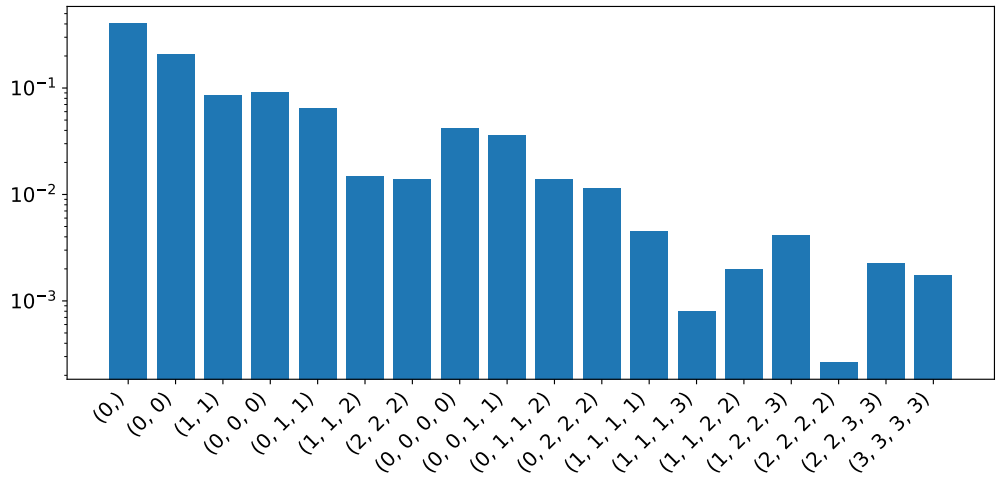

Figure 6.3: Relative abundance of topic subgraphs associated to grants in our collection of EPSRC grants with fewer than five associated topics. For these small graphs, we can uniquely identify them with their degree sequence which are the labels in the x-axis .

collection of mutually related topics. Despite the many topics associated with the grant in the first row, we can interpret from the keywords associated with the topics that the purpose of the grant concerns machine learning models for speech recognition and language comprehension. The grant in the second row reference the primary themes of modern mathematical physics. Lastly, the third grant concerns topics that belong in quantum information and quantum condensed matter physics, two intimately related fields in physics. Using maximum diversity, we are able to detect from the topical relations alone that these grants are topically far less diverse than the number of topics they reference suggests.

6.4.2 Generalised Jaccard Similarity between Grants

We compare the Jaccard similarity of pairs of grants to our generalisation using maximum diversity. To reduce our computational burden, we restrict our dataset of 10206 grants to the 1039 grants whose topic graphs have one connected component. Of all the $\binom{1039}{2} = 539241$ possible pairs of grants in this subset, there are 1381 pairs of grants which have a non-empty intersection of topics, and 9778 pairs that have no shared topics but whose topics are related. Thus, we already see that the usual Jaccard index is only useful for comparing a very small subset of grants, and there are seven times as many pairs of grants that may potentially be topically related where the Jaccard index would yield a zero similarity.

We compute the generalised Jaccard similarity across these 11159 pairs of grants for $\sigma = \frac{1}{8}$

Cardinality/Max Div	Topics and Topic Keywords
2.74 $\left(\sigma = \frac{1}{3}\right)$ 1.85 $\left(\sigma = \frac{1}{8}\right)$	Acoustic Model — Speaker Adaptation — Automatic Speech Recognition News Broadcast — Telephone Conversation — Speaker Speech Processing — Noise Type — Voice Activity Detection Speech Coding — Unvoiced — Musical Instrument Creaky Voice — Vowel — Speech Synthesis Spoken Dialog System — Partially Observable Markov Decision Process — Dialog Act Syllabification — Phoneme — Grapheme Volleyball — Soccer — Tennis Named Entity Recognition — Sentiment Classification — Entailment Keyword Search — Subword — Metrorrhagia Punctuation — Filled Pause — Disfluency N-Gram — Language Model — Automatic Speech Recognition Speech Enhancement — Linear Prediction — Reverberation Boltzmann Machine — Generative — Belief Network Natural Speech — Intonation — Uygur (people) Speaker Verification — Speech Enhancement — Vocoder Parallel Corpus — Translation Memory — Statistical Machine Translation Speech Intelligibility — Cochlear Nerve — Listener Code-Switching — Dialect — Language Identification Dysarthria — Motor Speech Disorder — Automatic Speech Recognition Speech Enhancement — Microphone — Public Speaking Acoustic Model — Text-To-Speech — Czech Language Speech Synthesis — Prosody — Prosodic Boundary Speech Intelligibility — Lombard — Talker Speech Signal — Robust Speech Recognition — Speaker Identification Voice Conversion — Text-To-Speech — Speech Synthesis
2.41 $\left(\sigma = \frac{1}{3}\right)$ 1.95 $\left(\sigma = \frac{1}{8}\right)$	Gauge Theory — Wilson Loop — Supergravity Nilpotency — Gauge Theory — Superspace Gauge Theory — Instanton — Brane Supersymmetry — Supergravity — Gravitino Complete Integrability — Brane — Einstein Space Spinor — Twistor — Superstring M-Theory — Supergravity — Brane String Theory — Dilaton — Conformal Field Theory Scattering Amplitude — Sigma Model — Duality String Theory — Matrix Model — Brane Ads/Cft — Holography — Brane
2.30 $\left(\sigma = \frac{1}{3}\right)$ 1.82 $\left(\sigma = \frac{1}{8}\right)$	Quantum Information Processing — Quantum State — Qubit Optical Vortex — Gaussian Beam — Angular Momentum Interpretation Of Quantum Mechanic — Wave Function — Decoherence Quantum Phase Transition — Critical Dynamic — Bose-Einstein Condensate Lobachevskian Geometry — 4-Manifold — Hyperbolic Knot Bell Inequality — Nonlocality — Contextuality Boseum — Vortex — Bose-Einstein Condensate Quantum Phase Transition — XY Model — Spin Chain Qubit — Monogamy — Entangled State Density Matrix — Fermion — Entanglement Quantum Walk — Evolution Operator — Coin Operated Equipment Mechanical Oscillator — Resonator — Optomechanic Coherent State — Squeezed State (quantum Theory) — Quantum Teleportation Counterfactual — Quantum System — Zeno Quantum State — Superposition — Decoherence Qubit — Quantum Measurement — Entanglement Josephson Junction — Superconducting Resonator — Qubit Quantum State — Cloning — Teleportation Quantum Fisher Information — Qubit — Metrology Qubit — Casimir Effect — Non-Stationary Qubit — Master Equation — Open Quantum System Qubit — Quantum Control — Speed Limit Quantum State — Qubit — Spin Chain Qubit — Discord — Decoherence Entropy Production — Maxwell Demon — Fluctuation Theorem Qubit — Ion Trap (instrumentation) — Trapped Ion Quantum Phase Transition — Dissipative — Jaynes-Cumming Model

Table 6.1: Grants with the top three topic cardinality to maximum diversity over-estimation ratio, shown here with the keywords associated to the topics of the award. Each line of keywords corresponds to a single topic.

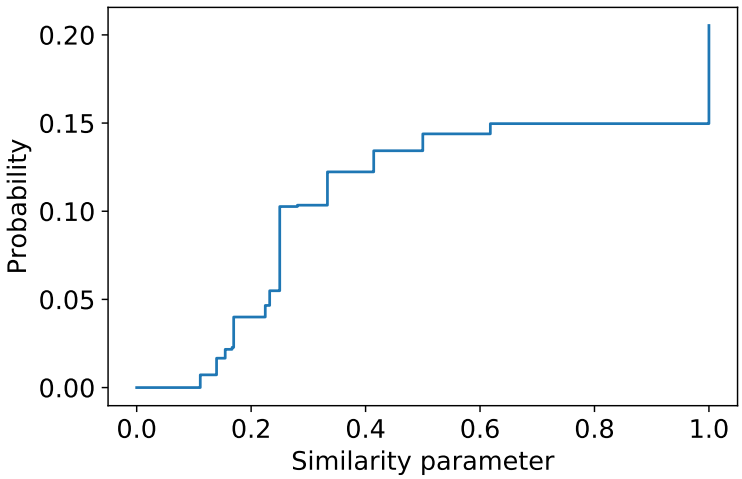

Figure 6.4: The probability eq. (6.4) that a grant associated with four topics has a lower or equivalent maximum diversity than a grant associated with three topics, as a function of the similarity parameter σ.

and $\frac{1}{3}$. In Figure 6.6, we plot the cumulative distributions of pairs across grants which have a non-empty intersection of topics, and those which have related topics but do not share topics separately. We can note from Figure 6.6 that while pairs with non-empty intersection tend to have a higher similarity, this is not always true: a small fraction of pairs of grants with related, non-overlapping topics have a similarity exceeding the median of pairs of grants with shared topics.

In Figure 6.6, we plot the generalised Jaccard similarity of the pairs against their Jaccard similarity, which only takes the intersection of topics into account and neglects topical relations. In the case where $\sigma = \frac{1}{3}$, we note there there are a few pairs of grants whose generalised Jaccard similarity is lower than the usual Jaccard similarity. Intuitively, we can attribute this to topical relations being denser in the individual grants in comparison with the union, such that the relatively large number of dissimilar topics in the union outweigh the effect of shared topics towards increasing the similarity between two grants. This is a more pronounced effect as the similarity parameter σ increases. Counter-intuitively at first glance, we recall the generalised Jaccard index when $\sigma = 0$ is simply the usual Jaccard index where no topics are similar. As we increase σ, the absence of topical relations become more keenly felt and plays a relatively larger role in determining multi-disciplinarity of and similarity between grants.

As an example, we consider the pair of grants which has the highest generalised Jaccard similarity to Jaccard similarity ratio at $\sigma = \frac{1}{3}$. For both grants, their topic graphs are the

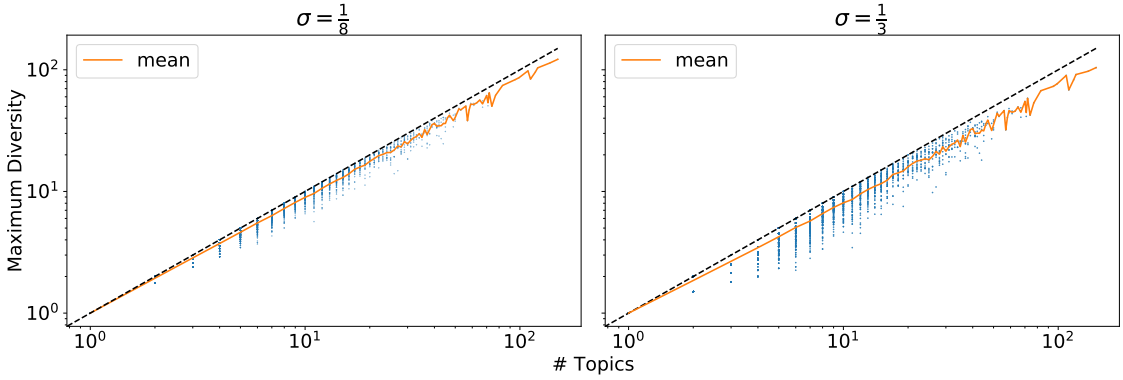

Figure 6.5: For each EPSRC grant in our collection, we plot the maximum diversity of the topics associated to the grantfor similarity parameters set to $\sigma = \frac{1}{8}$ and $\sigma = \frac{1}{3}$, against the cardinality of topics. We also plot the mean maximum diversity of grants binned by their number of topics.

path graphs with three vertices. Their union is the X shaped graph with five vertices and the middle vertex of the path graphs is the intersection of the two paths. On such a small graph, we can explicitly compute an analytic expression for the generalised Jaccard similarity as a function of σ, which we plot in Figure 6.8. In the small σ regime, we observe that the fact that topics not present in the overlap are related to the topic in the overlap leads to a higher generalised Jaccard similarity than the case where these relations are absent in the $\sigma = 0$ limit. As σ increases, the consequences of there being no relations between topics across the two grants apart from the topic in the overlap becomes more pronounced. For $\sigma \geq \frac{1}{2}$, the central vertex of the path graph is accorded zero weight in its maximum diversity realising distribution, and only the ends of the path have non-zero contribution to the topical content of the grant. In that regime of σ, we are effectively comparing two disjoint sets of vertices, and thus the generalised Jaccard similarity is zero.

6.5 Discussion and Industrial Recommendations

We have shown that relations between topics are important factors in evaluating the multi-disciplinarity of research grants. When the similarity between different topics is strong, then it is plausible that the number of topics is a poor proxy for the multi-disciplinarity of grants.

On our dataset of EPSRC grants, we are able to use maximum diversity as a proxy for multi-disciplinarity to evaluate how topical relations can cause the true multi-disciplinarity of a

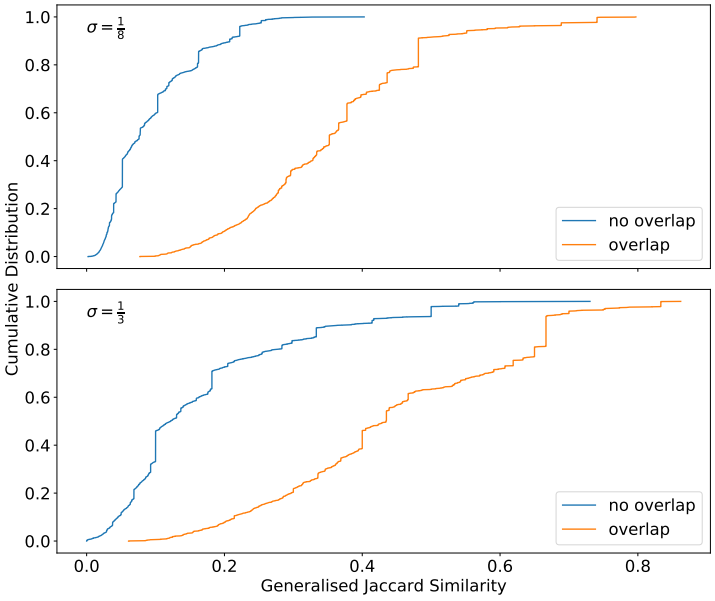

Figure 6.6: On the subset of grants where the associated topic subgraph of each graph is connected, we plot the cumulative distributions of pairs of grants which have a non-empty intersection of topics, and those which have related topics but do not share topics separately. The plot above and below shows the case where the topical similarity parameters are fixed at $\sigma = \frac{1}{8}$ and $\frac{1}{3}$ respectively.

grant to deviate from the number of topics. We show that the impact of topical relations are significant on the multi-disciplinarity of grants that reference topics that are mutually similar, which is more common in our dataset for grants associated a large number of topics. For grants that are associated with a few topics, the number of topics is more often than not a reliable indication of multi-disciplinarity (see Figure 6.5), though this is a reflection of the fact that a significant majority of grants that are only associated with few topics reference topics that are dissimilar (see Figure 6.3). All the while the chance that the number of topics is an inaccurate tool for comparing the multi-disciplinarity of two grants is not negligible (see Figure 6.4).

We note that for grants that are only associated with a few topics, the addition or subtraction of a topic relation can have a large impact on the assessment of multi-disciplinarity via maximum diversity (see Figure 6.2). As such, our assessment of multi-disciplinarity can be sensitive to errors and depend on the assumptions and biases in deriving these topical relations.

Having applied maximum diversity to derive a generalisation of the Jaccard similarity, we are able to compare grants with related topics that do not necessarily share associated top-

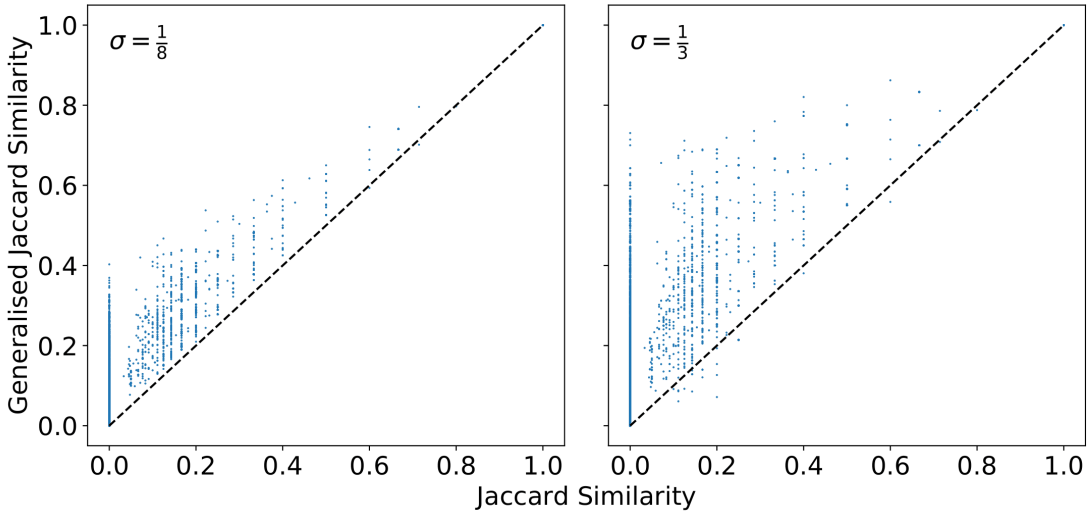

Figure 6.7: On the subset of grants where the associated topic subgraph of each graph is connected, we plot the Jaccard similarity between pairs of grants against its generalised Jaccard similarity, for similarity parameters $\sigma = \frac{1}{8}$ and $\frac{1}{3}$. The black dashed diagonal indicates where the Jaccard and generalised Jaccard indices coincide.

ics. In cases where topics are deemed to be highly related (i.e. large σ), it is possible for grants that share topics to be maximally dissimilar under our generalised notion of Jaccard similarity, if topics in the overlap make no effective contribution to the multi-disciplinarity of either grants, and topics not shared are dissimilar. Our investigation draws attention to the tension between weighing up the contributions of topical overlaps and topical relationships in assessing the similarity between grants. While our generalisation of the Jaccard similarity does not definitively resolve the tension between these competing factors, we note that a more nuanced assessment of the similarity between two grants can be gleaned by varying the similarity parameter σ, where for larger σ the topical relations play a more prominent role and may potentially reduce the influence of topical overlaps.

While maximum diversity has a strong theoretical foundation and justification for its interpretation as a proxy for the effective cardinality of a set equipped with a similarity matrix, it is computationally intensive to evaluate for large sets as we have to perform computations over potentially all subsets of the set in consideration (see definition 6.2.5). In practice, we recommend choosing a sufficiently small σ such that the similarity matrix between topics is positive definite. In this case, the maximum diversity of grant can be evaluated by solving a convex quadratic programming problem (eq. (6.1)) instead, which has many efficient software imple-

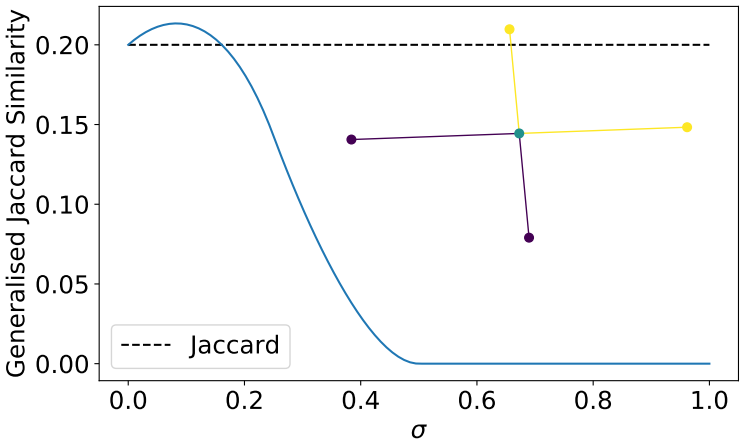

Figure 6.8: We plot the generalised Jaccard similarity between two path graphs (coloured purple and yellow respectively) with three vertices intersecting at the middle vertex (coloured green), as a function of the similarity σ between vertices connected by an edge. The black dashed line indicates the value of the usual Jaccard similarity between the vertex sets of the two graphs for comparison.

mentations. For grants associated with few topics, it is possible to derive exact expressions for the maximum diversity of grants by hand, as we have shown in Figure 6.2. As the topic graphs associated with 95% of the EPSRC grants can be expressed as a disjoint union of connected components with fewer than six vertices, in principle we can derive exact expressions by hand for the maximum diversity of the vast majority grants in the EPSRC dataset across all values of σ. As such, while maximum diversity might be computationally too costly to apply at first glance, there are several feasible strategies to overcome its computational limitations for applying it to analysing the multi-disciplinarity of grants.

Finally, we also note that our approach is not limited to analysing multi-disciplinarity of grants; for instance, if we have an embedding of research articles into some metric space using a topic model algorithm such as `doc2vec`, we can compute a notion of multi-disciplinarity of an author of a research group using the distance matrix between papers associated to the author or research group.

Chapter 7

Conclusion and Future Work

In this thesis, we built on recent advances in applied algebraic topology to develop new tools for computation of algebraic invariants of critical points and data science applications. In Chapters two to four, we considered critical points of smooth function of submanifolds embedded in Euclidean space, and showed that the Morse index of a critical point can be inferred from point samples in a neighbourhood of the critical point. We drew on two mathematical themes. The first theme is Gromoll-Meyer theory, which offers a recipe for constructing pairs of compact subsets called Gromoll-Meyer pairs, whose relative homology is isomorphic to the Morse index. Gromoll-Meyer pairs enable us to restrict our computational input to a neighbourhood of the critical point, rather than having to use global sublevel sets of the function. The second theme is homology estimation from finite point samples. We illustrated how the relative homology of compact subsets of Euclidean space can be inferred from finite point samples, if the subsets have positive reach and the Hausdorff distance between the point samples and the compact subsets is sufficiently small relative to the reach of the compact subsets. We subsequently showed that Gromoll-Meyer pairs can be constructed with an explicit positive lower bound of their reach if f satisfies reasonable bounds on the gradient and Hessian on a neighbourhood of the critical point. Bringing the two themes together, we outlined a method of obtaining point samples that are sufficiently close to the Gromoll-Meyer pair in the Hausdorff distance, and would thus provide sufficient data to infer the Morse index of the critical point.

We illustrate two possible directions for future work. First, we conjecture that the Conley

indices of isolated invariant sets of flows generated by smooth vector fields on a manifold can also be inferred from finitely many point samples on the isolating neighbourhood of the isolated invariant set. This conjecture is motivated by the existence of a smooth function on the isolating neighbourhood that is constant on the isolated invariant set and strictly decreasing along the flow elsewhere [Conley, 1978]. Thus, if such a function can be inferred from finitely many point samples of the vector field on the isolating neighbourhood, our theoretical guarantees for inferring the Morse index of gradient flows can be extended to this wider class of dynamical systems. Second, for a smooth function f on a compact submanifold M, we may employ the methods used in this thesis to deduce the sufficient density of point samples so that we can recover the persistent homology of the filtration of M by sublevel set filtrations of f. This is of practical interests where we do not have explicit expressions for f, and can only perform finitely many function evaluations. We posit that this approach may be useful in providing theoretical guarantees for recocering the the persistent homology transform of M [Turner et al., 2014b, Ghrist et al., 2018, Curry et al., 2018] from a sufficiently dense point sample of M.

We note that there are some obstacles in applying this work in computational settings. Since current homology estimation techniques from point samples rely on building Čech or Vietoris-Rips simplicial complexes on the point sample vertex set, our approach is computationally intensive for high dimensional manifolds, as the number of simplices in the simplicial complexes increases exponentially with the dimension of the manifold. As future work, we would investigate whether the construction of the simplicial complexes can be bypassed by using polynomials to approximate the function near the critical point, and determining the conditions upon which the index analytically computed from the polynomial approximation is isomorphic to the true index. Since polynomial approximation only requires pointwise function evaluation, this approach may bypass the computational burden due to the combinatorial explosion of building higher order simplices.

In Chapter 5, we studied how the filtration of graphs by wavelet signatures can be optimised to improve the classification accuracy of neural networks that use the extended persistence diagrams of the filtered graphs as predictive features. Optimising features of graphs for classification and regression problems is an area of active and intense research, and our study adds to the literature that demonstrates persistence homology is a competitive feature map.

While the improvement in classification accuracy due to wavelet optimisation may not be statistically significant on all datasets, many avenues of further investigation can be pursued to yield improvements. First, while smooth functions composed with the persistent homology map are generically differentiable, it is important to emphasise that the continuity of derivatives of the composed map only holds on dense open subsets of the domain called strata, and the derivatives may change wildly as we traverse from one stratum to another [Leygonie et al., 2019]. As such, gradient descent methods may not be the most efficient way of optimising loss functions that factor through persistence diagrams. Currently, methods for effectively optimising maps composed with persistent homology remains an open area of research [Leygonie et al., 2021, Carriere et al., 2021]. Second, more work can be done on how to select an appropriate parametrisation of wavelet signatures. For all the datasets that we had examined, we empirically observed that the map from the space of wavelet basis coefficients to wavelet signatures is often numerically ill-conditioned and require additional numerical procedures to stabilise the gradient descent computations. Furthermore, the degree of ill-conditioning is dependent on the choice of wavelet basis functions. Thus, further investigation can be made into how a judicious choice of basis functions can be made to improve numerical stability and classification accuracy.

In Chapter 6, we demonstrated how maximum diversity, a measure of the effective size of sets with similarity, can be applied to analyse the multi-disciplinarity of research grants with topic associations and derive a notion of topical similarity between grants using a generalisation of the Jaccard index. In our empirical case study on a collection of EPSRC grants, where the set of topic associations to those grants and topic similarities is given by Elsevier's proprietary SciVal topic generator, we demonstrate how topic similarities can grants with the same number of associated topics can have wildly different multi-disciplinarity as measured by maximum diversity. Using our generalised Jaccard similarity, we are able to similarly quantify how topic similarity can skew the assessment of topic closeness between grants compared to simply accounting for common topics between grants in the standard Jaccard index. Maximum diversity and the generalised Jaccard index have the advantage of being interpretable in applications, and can be efficiently computed using convex quadratic programming algorithms if the similarity between elements is sufficiently small. However, if the similarity between elements is sufficiently large, then the computational cost of computing these quantities increases

exponentially with the number of elements in the set. Given these computational obstacles, further theoretical work can be conducted on formulating computable and informative lower and upper bounds on maximum diversity, as well as asymptotic behaviours for similarity matrices with random entries as the number of elements tends to infinity. These approximations of maximum diversity may prove helpful to practitioners in data science that wish to estimate the effective size of sets with similarity, without having to expend undue computational resources to compute an exact value of maximum diversity.

Bibliography

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. Publisher: JMLR. org.

Ali N Akansu, Paul A Haddad, Richard A Haddad, and Paul R Haddad. *Multiresolution signal decomposition: transforms, subbands, and wavelets*. Academic press, 2001.

Stephanie B. Alexander, I. David Berg, and Richard L. Bishop. The Riemannian obstacle problem. *Illinois Journal of Mathematics*, 31(1), March 1987. ISSN 0019-2082. doi: 10.1215/ijm/1255989406. URL https://projecteuclid.org/journals/illinois-journal-of-mathematics/volume-31/issue-1/The-Riemannian-obstacle-problem/10.1215/ijm/1255989406.full.

Jean-Pierre Antoine, Daniela Roşca, and Pierre Vandergheynst. Wavelet transform on manifolds: Old and new approaches. *Applied and Computational Harmonic Analysis*, 28(2):189–202, 2010.

Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 1626–1633. IEEE, 2011.

Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *Journal of Computational Geometry*, pages 162–191 Pages, March 2015. doi: 10.20382/JOCG.V6I2A9. URL https://jocg.org/index.php/jocg/article/view/2983. Artwork Size: 162–191 Pages Publisher: Journal of Computational Geometry.

Türker Biyikoglu, Josef Leydold, and Peter F Stadler. *Laplacian eigenvectors of graphs: Perron-Frobenius and Faber-Krahn type theorems*. Springer, 2007.

Omer Bobrowski, Sayan Mukherjee, and Jonathan E Taylor. Topological consistency via kernel estimation. *Bernoulli*, 23(1):288–328, 2017.

Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and topological inference*, volume 57. Cambridge University Press, 2018.

Jean-Daniel Boissonnat, André Lieutier, and Mathijs Wintraecken. The reach, metric distortion, geodesic convexity and the variation of tangent spaces. *Journal of Applied and Computational Topology*, 3(1-2):29–58, June 2019. ISSN 2367-1726, 2367-1734. doi: 10.1007/s41468-019-00029-8. URL http://link.springer.com/10.1007/s41468-019-00029-8.

Karsten M Borgwardt, Cheng Soon Ong, Stefan Schönauer, SVN Vishwanathan, Alex J Smola, and Hans-Peter Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 21 (suppl_1):i47–i56, 2005. Publisher: Oxford University Press.

Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1704–1711. IEEE, 2010.

Rickard Brüel-Gabrielsson, Vignesh Ganapathi-Subramanian, Primoz Skraba, and Leonidas J Guibas. Topology-Aware Surface Reconstruction for Point Clouds. In *Computer Graphics Forum*, volume 39 No. 5, pages 197–207. Wiley Online Library, 2020.

Mathieu Carriere, Frédéric Chazal, Marc Glisse, Yuichi Ike, and Hariprasad Kannan. A note on stochastic subgradient descent for persistence-based functionals: convergence and practical aspects. *arXiv preprint arXiv:2010.08356*, 2020.

Mathieu Carriere, Frédéric Chazal, Marc Glisse, Yuichi Ike, Hariprasad Kannan, and Yuhei Umeda. Optimizing persistent homology based functions. In *International Conference on Machine Learning*, pages 1294–1303. PMLR, 2021.

Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. Perslay: A neural network layer for persistence diagrams and new graph topological signatures. In *International Conference on Artificial Intelligence and Statistics*, pages 2786–2796. PMLR, 2020.

Kung Ching Chang and N. Ghoussoub. The Conley index and the critical groups via an extension of Gromoll-Meyer theory. *Topological Methods in Nonlinear Analysis*, 7(1):77, March 1996. ISSN 1230-3429. doi: 10.12775/TMNA.1996.003. URL https://www.tmna.ncu.pl/static/files/v07n1-03.pdf.

Frédéric Chazal and André Lieutier. Weak feature size and persistent homology: computing homology of solids in rn from noisy data samples. In *Proceedings of the twenty-first annual symposium on Computational geometry*, pages 255–262, 2005.

Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Analysis of scalar fields over point cloud data. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 1021–1030. SIAM, 2009a.

Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the 25th annual symposium on Computational geometry - SCG '09*, page 237, Aarhus, Denmark, 2009b. ACM Press. ISBN 978-1-60558-501-7. doi: 10.1145/1542362.1542407. URL http://portal.acm.org/citation.cfm?doid=1542362.1542407.

Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A Sampling Theory for Compact Sets in Euclidean Space. *Discrete & Computational Geometry*, 41(3):461–479, April 2009c. ISSN 0179-5376, 1432-0444. doi: 10.1007/s00454-009-9144-8. URL http://link.springer.com/10.1007/s00454-009-9144-8.

Chao Chen, Xiuyan Ni, Qinxun Bai, and Yusu Wang. A Topological Regularizer for Classifiers via Persistent Homology. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2573–2582, 2019a.

CS Chen, YC Hon, and RA Schaback. Scientific computing with radial basis functions. *Department of Mathematics, University of Southern Mississippi, Hattiesburg, MS*, 39406, 2005.

Minshuo Chen, Haoming Jiang, Wenjing Liao, and Tuo Zhao. Efficient approximation of deep relu networks for functions on low dimensional manifolds. *Advances in neural information processing systems*, 32:8174–8184, 2019b.

Sheng Chen, Colin FN Cowan, and Peter M Grant. Orthogonal least squares learning algo-

rithm for radial basis function networks. *IEEE Transactions on neural networks*, 2(2):302–309, 1991.

Ilya Chevyrev and Andrey Kormilitzin. A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*, 2016.

Charles K Chui. *An introduction to wavelets*. Elsevier, 2016.

Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.

James R Clough, Ilkay Oksuz, Nicholas Byrne, Julia A Schnabel, and Andrew P King. Explicit topological priors for deep-learning based image segmentation using persistent homology. In *International Conference on Information Processing in Medical Imaging*, pages 16–28. Springer, 2019.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of Persistence Diagrams. *Discrete & Computational Geometry*, 37(1):103–120, January 2007. ISSN 0179-5376, 1432-0444. doi: 10.1007/s00454-006-1276-5. URL http://link.springer.com/10.1007/s00454-006-1276-5.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using Poincaré and Lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009. Publisher: Springer.

Charles Conley. *Isolated invariant sets and the morse index: expository lectures*. Number 38 in Regional conference series in mathematics. American Mathematical Society, Providence, RI, 1978. ISBN 978-0-8218-1688-2.

Padraig Corcoran and Bailin Deng. Regularization of Persistent Homology Gradient Computation. *arXiv preprint arXiv:2011.05804*, 2020.

William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and Its Applications*, 14(05):1550066, June 2015. ISSN 0219-4988, 1793-6829. doi: 10.1142/S0219498815500668. URL https://www.worldscientific.com/doi/abs/10.1142/S0219498815500668.

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.

Damek Davis, Dmitriy Drusvyatskiy, Sham Kakade, and Jason D Lee. Stochastic subgradient method converges on tame functions. *Foundations of computational mathematics*, 20(1):119–154, 2020. Publisher: Springer.

Asim Kumar Debnath, Rosa L Lopez de Compadre, Gargi Debnath, Alan J Shusterman, and Corwin Hansch. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of Medicinal Chemistry*, 34(2):786–797, 1991. Publisher: ACS Publications.

Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of Molecular Biology*, 330(4):771–783, 2003. Publisher: Elsevier.

Edelsbrunner, Letscher, and Zomorodian. Topological Persistence and Simplification. *Discrete & Computational Geometry*, 28(4):511–533, November 2002. ISSN 0179-5376, 1432-0444. doi: 10.1007/s00454-002-2885-2. URL http://link.springer.com/10.1007/s00454-002-2885-2.

Herbert Edelsbrunner and John Harer. Persistent homology-a survey. *Contemporary mathematics*, 453:257–282, 2008. Publisher: Providence, RI: American Mathematical Society.

Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93 (3):418–418, March 1959. ISSN 0002-9947. doi: 10.1090/S0002-9947-1959-0110078-1. URL http://www.ams.org/jourcgi/jour-getitem?pii=S0002-9947-1959-0110078-1.

Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016.

Peter Gabriel. Unzerlegbare Darstellungen I. *manuscripta mathematica*, 6(1):71–103, March 1972. ISSN 0025-2611, 1432-1785. doi: 10.1007/BF01298413. URL http://link.springer.com/10.1007/BF01298413.

Rickard Brüel Gabrielsson, Bradley J Nelson, Anjan Dwaraknath, and Primoz Skraba. A topology layer for machine learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1553–1563. PMLR, 2020.

Marcio Gameiro, Yasuaki Hiraoka, and Ippei Obayashi. Continuation of point clouds via persistence diagrams. *Physica D: Nonlinear Phenomena*, 334:118–132, 2016. Publisher: Elsevier.

Jie Gao, Leonidas J Guibas, Steve Y Oudot, and Yue Wang. Geodesic delaunay triangulation and witness complex in the plane. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 571–580, 2008.

Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and euler integral transforms. *Journal of Applied and Computational Topology*, 2(1):55–60, 2018.

Amara Graps. An introduction to wavelets. *IEEE computational science and engineering*, 2(2): 50–61, 1995. Publisher: IEEE.

Detlef Gromoll and Wolfgang Meyer. On differentiable functions with isolated critical points. *Topology*, 8(4):361–369, September 1969. ISSN 00409383. doi: 10.1016/0040-9383(69)90022-6. URL https://linkinghub.elsevier.com/retrieve/pii/0040938369900226.

David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011. Publisher: Elsevier.

Christoph Hofer, Roland Kwitt, Marc Niethammer, and Mandar Dixit. Connectivity-optimized representation learning via persistent homology. In *International Conference on Machine Learning*, pages 2751–2760. PMLR, 2019.

Christoph Hofer, Florian Graf, Bastian Rieck, Marc Niethammer, and Roland Kwitt. Graph filtration learning. In *International Conference on Machine Learning*, pages 4314–4323. PMLR, 2020.

Nan Hu, Raif M Rustamov, and Leonidas Guibas. Stable and informative spectral signatures for graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2305–2312, 2014.

Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen. Topology-preserving deep image segmentation. In *Advances in Neural Information Processing Systems*, pages 5658–5669, 2019.

Oleg Kachan. Persistent Homology-Based Projection Pursuit. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 856–857, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Richard Klavans and Kevin W. Boyack. Research portfolio analysis and topic prominence. *Journal of Informetrics*, 11(4):1158–1174, November 2017. ISSN 17511577. doi: 10.1016/j.joi.2017.10.002. URL https://linkinghub.elsevier.com/retrieve/pii/S1751157717302110.

John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 2012. ISBN 978-1-4419-9981-8 978-1-4419-9982-5. doi: 10.1007/978-1-4419-9982-5. URL http://link.springer.com/10.1007/978-1-4419-9982-5.

Tom Leinster and Mark W. Meckes. Maximizing Diversity in Biology and Beyond. *Entropy*, 18 (3):88, March 2016. doi: 10.3390/e18030088. URL https://www.mdpi.com/1099-4300/18/3/88. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Tom Leinster and Emily Roff. The maximum entropy of a metric space. *arXiv:1908.11184 [cs, math]*, October 2019. URL http://arxiv.org/abs/1908.11184. arXiv: 1908.11184.

Jacob Leygonie, Steve Oudot, and Ulrike Tillmann. A Framework for Differential Calculus on Persistence Barcodes. *arXiv preprint arXiv:1910.00960*, 2019.

Jacob Leygonie, Mathieu Carrière, Théo Lacombe, and Steve Oudot. A gradient sampling algorithm for stratified maps with applications to topological data analysis. *arXiv preprint arXiv:2109.00530*, 2021.

Chunyuan Li and A Ben Hamza. A multiresolution descriptor for deformable 3D shape retrieval. *The Visual Computer*, 29(6-8):513–524, 2013. Publisher: Springer.

Roee Litman and Alexander M Bronstein. Learning spectral descriptors for deformable shape correspondence. *IEEE transactions on pattern analysis and machine intelligence*, 36(1):171–180, 2013.

John Willard Milnor, Michael Spivak, and Raymond O Wells. *Morse theory*. 1973. ISBN 978-1-4008-8180-2. URL https://doi.org/10.1515/9781400881802. OCLC: 1165548913.

Konstantin Mischaikow and Marian Mrozek. Conley Index. In *Handbook of Dynamical Systems*, volume 2, pages 393–460. Elsevier, 2002. ISBN 978-0-444-50168-4. doi: 10.1016/S1874-575X(02)80030-3. URL https://linkinghub.elsevier.com/retrieve/pii/S1874575X02800303.

Michael Mongillo et al. Choosing basis functions and shape parameters for radial basis function methods. *SIAM undergraduate research online*, 4(190-209):2–6, 2011.

Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In *International Conference on Machine Learning*, pages 7045–7054. PMLR, 2020.

Mark Newman. *Networks*. Oxford university press, 2018.

Liviu I. Nicolaescu. *An invitation to morse theory*. Universitext. Springer Science+Business Media, LLC Springer e-books, New York, NY, 2nd edition edition, 2011. ISBN 978-1-4614-1105-5.

Mathias Niepert, Mohamed Ahmed, and Konstantin Kutzkov. Learning Convolutional Neural Networks for Graphs. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2014–2023, New York, New York, USA, June 2016. PMLR. URL http://proceedings.mlr.press/v48/niepert16.html.

Partha Niyogi, Stephen Smale, and Shmuel Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete & Computational Geometry*, 39(1):419–441, March 2008. ISSN 1432-0444. doi: 10.1007/s00454-008-9053-2. URL https://doi.org/10.1007/s00454-008-9053-2.

Steve Oudot and Elchanan Solomon. Inverse problems in topological persistence. In *Topological Data Analysis*, pages 405–433. Springer, 2020.

Jooyoung Park and Irwin W Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 3(2):246–257, 1991. Publisher: MIT Press.

Adrien Poulenard, Primoz Skraba, and Maks Ovsjanikov. Topological function optimization for continuous shape matching. In *Computer Graphics Forum*, volume 37 No. 5, pages 13–25. Wiley Online Library, 2018.

Bastian Rieck, Christian Bock, and Karsten Borgwardt. A persistent weisfeiler-lehman procedure for graph classification. In *International Conference on Machine Learning*, pages 5448–5458. PMLR, 2019.

Raif M Rustamov and Leonidas J Guibas. Wavelets on graphs via deep learning. In *Vertex-Frequency Analysis of Graph Signals*, pages 207–222. Springer, 2019.

Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. Efficient graphlet kernels for large graph comparison. In *Artificial Intelligence and Statistics*, pages 488–495, 2009.

Nino Shervashidze, Pascal Schweitzer, Erik Jan Van Leeuwen, Kurt Mehlhorn, and Karsten M Borgwardt. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9), 2011.

Primoz Skraba and Bei Wang. Approximating Local Homology from Samples. *arXiv:1206.0834 [cs, math]*, June 2012. URL http://arxiv.org/abs/1206.0834. arXiv: 1206.0834.

Elchanan Solomon, Alexander Wagner, and Paul Bendich. A Fast and Robust Method for Global Topological Functional Optimization. *arXiv preprint arXiv:2009.08496*, 2020.

Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, volume 28 No. 5, pages 1383–1392. Wiley Online Library, 2009.

Jeffrey J Sutherland, Lee A O'brien, and Donald F Weaver. Spline-fitting with a genetic algorithm: A method for developing classification structure- activity relationships. *Journal of Chemical Information and Computer Sciences*, 43(6):1906–1915, 2003. Publisher: ACS Publications.

Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.

Katharine Turner, Yuriy Mileyko, Sayan Mukherjee, and John Harer. Fréchet means for distributions of persistence diagrams. *Discrete & Computational Geometry*, 52(1):44–70, 2014a. Publisher: Springer.

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, 2014b.

Saurabh Verma and Zhi-Li Zhang. Hunt for the Unique, Stable, Sparse and Fast Feature Learning on Graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pages 87–97, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 978-1-5108-6096-4. event-place: Long Beach, California, USA.

S Vichy N Vishwanathan, Nicol N Schraudolph, Risi Kondor, and Karsten M Borgwardt. Graph kernels. *The Journal of Machine Learning Research*, 11:1201–1242, 2010. Publisher: JMLR. org.

Nikil Wale, Ian A Watson, and George Karypis. Comparison of descriptor spaces for chemical compound retrieval and classification. *Knowledge and Information Systems*, 14(3):347–375, 2008. Publisher: Springer.

Yuan Wang and Bei Wang. Topological inference of manifolds with boundary. *Computational Geometry*, 88:101606, June 2020. ISSN 09257721. doi: 10.1016/j.comgeo.2019.101606. URL https://linkinghub.elsevier.com/retrieve/pii/S0925772119301476.

Bingbing Xu, Huawei Shen, Qi Cao, Yunqi Qiu, and Xueqi Cheng. Graph wavelet neural network. *arXiv preprint arXiv:1904.07785*, 2019.

Pinar Yanardag and SVN Vishwanathan. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1365–1374, 2015.

Ka Man Yim and Jacob Leygonie. Optimization of Spectral Wavelets for Persistence-Based Graph Classification. *Frontiers in Applied Mathematics and Statistics*, 7:651467, April 2021. ISSN 2297-4687. doi: 10.3389/fams.2021.651467. URL https://www.frontiersin.org/articles/10.3389/fams.2021.651467/full.

Zhen Zhang, Mianzhi Wang, Yijian Xiang, Yan Huang, and Arye Nehorai. RetGK: Graph Kernels Based on Return Probabilities of Random Walks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pages 3968–3978, Red Hook, NY, USA, 2018. Curran Associates Inc. event-place: Montréal, Canada.

Qi Zhao and Yusu Wang. Learning metrics for persistence-based summaries and applications for graph classification. In *Advances in Neural Information Processing Systems*, pages 9859–9870, 2019.

Afra Zomorodian and Gunnar Carlsson. Computing Persistent Homology. *Discrete & Computational Geometry*, 33(2):249–274, February 2005. ISSN 0179-5376, 1432-0444. doi: 10.1007/s00454-004-1146-y. URL https://link.springer.com/article/10.1007/s00454-004-1146-y.

Appendices

A Selected List of Notation

τ_A	The reach of a subset $A \subset \mathbb{R}^d$.
$\tau_A(p)$	The local feature size of a subset $A \subset \mathbb{R}^d$ at $p \in A$.
$\mathrm{Med}(A)$	The medial axis of a subset $A \subset \mathbb{R}^d$.
$\mathrm{UP}(A)$	The set of points in $\mathbb{R}^d$ which has a unique nearest neighbour in $A \subset \mathbb{R}^d$.
$\mathrm{NN}_A(x)$	The set of nearest neighbours of $x \in \mathbb{R}^d$ in a closed subset $A \subset \mathbb{R}^d$.
$\mathrm{Tan}(A, p)$	The tangent set of a subset A at $p \in A$.
$\mathrm{Nor}(A, p)$	The normal cone of a subset A at $p \in A$.
$B_r(x)$	An open Euclidean ball with radius r, centered at $x \in \mathbb{R}^d$.
$B_r[x]$	A closed Euclidean ball with radius r, centered at $x \in \mathbb{R}^d$.
$B_r^M(p)$	An open geodesic ball with radius r on a manifold M, centred at $p \in M$.
$B_r^M[p]$	A closed geodesic ball with radius r on a manifold M, centred at $p \in M$.
$\overline{xp}$	A straight line segment in Euclidean space between points x and p.
∇	The Levi-Civita connection of the metric on a Riemannian manifold.
$\boldsymbol{\nabla}$	The Euclidean connection.
$\mathrm{II}(X, Y)$	The second fundamental form of a Riemmannian manifold.
$\mathcal{T}(M)$	The tangent bundle of a manifold.
$\mathcal{N}(M)$	The normal bundle of a manifold embedded in $\mathbb{R}^d$.
$\measuredangle(u, v)$	The angle between two non-zero vectors on u and v in $\mathbb{R}^d$.
$\mathrm{Crit}(f)$	The critical points of a smooth function.
$\mathrm{Inv}(N)$	The maximal invariant set of a dynamical system in a compact subset N.
$\check{\mathsf{C}}(A)$	The Čech filtration of $A \subset \mathbb{R}^d$.
$\check{\mathsf{H}}_i(A)$	The persistent homology of a Čech filtration
$\lVert\!\lVert V \rVert\!\rVert$	The barcode of a p.f.d. persistence module V.
Bar	The space of persistence diagrams.
$\mathsf{EPH}_p^{\mathsf{ord}}(f)$	The ordinary part of the extended persistence diagram of a filtered graph (G, f) in p^{th} degree homology.
$\mathrm{supp}(p)$	The support of a probability distribution p.
$\lvert X \rvert$	The magnitude of a set X equipped with a similarity matrix.
$\lvert X \rvert_+$	The maximum diversity of a set X equipped with a similarity matrix.

B Computational Particulars of Experiments Described in Chapter 5

B.1 Persistence Images Parameters

We vectorised each of the three persistence diagrams EPH_0, $\mathsf{EPH}_1^{\mathsf{rel}}$ and $\mathsf{EPH}_1^{\mathsf{ext}}$ as a persistence image. Prior to vectorising the persistence diagrams, we apply a fixed and identical affine transformation to the values of the vertex functions across all graphs in the dataset concerned, such that the maximum and minimum values taken across all graphs in the dataset of the initial vertex function prior to optimisation are scaled to 1 and 0 respectively. The persistence image is sampled on a 20×20 grid, whose grid points are equidistantly placed $\sigma = 1/17$ apart on the square $[-\sigma, 1+\sigma]^2$ of the persistence diagrams, where σ is the width of the Gaussian. The Gaussian centred on the birth and persistence coordinates $\langle b, p \rangle$ of each point is weighted according to its persistence

$$\omega(p) = \sin^2\Big(\frac{\pi}{2} \min\Big(\frac{p}{\sigma}, 1\Big)\Big).$$

Points with persistence $p \geq \sigma$ are assigned a uniform weight $\omega = 1$, else assigned a weight that diminishes to zero as $p \to 0$.

B.2 Convolutional Neural Network Architecture for Persistence Images

We feed each set of three persistence images belonging to either the optimisable or static persistence diagrams as a three channel image into the following convolutional neural network to obtain a 22×22 image:

$$\mathsf{CNN} : \mathbb{R}^{3\times 20\times 20} \xrightarrow{\mathsf{ReLU} \circ \mathsf{Conv} \circ \mathsf{BN2D}} \mathbb{R}^{20\times 21\times 21} \xrightarrow{\mathsf{ReLU} \circ \mathsf{Conv} \circ \mathsf{DO} \circ \mathsf{BN2D}} \mathbb{R}^{22\times 22}. \tag{1}$$

The function Conv denotes a convolutional layer with kernel size 2, stride 1, padding 1; BN2D denotes a 2D batch normalisation layer; and DO denotes a dropout layer with dropout probability $p = 0.5$.

B.3 Multilayer Perceptron for Non-Persistence Features

We feed non-persistence features as a vector of length $n = \#\text{features}$ into the following multilayer perceptron:

$$\mathsf{MLP} : \mathbb{R}^n \xrightarrow{\mathsf{BN} \circ \mathsf{ReLU} \circ \mathsf{Aff} \circ \mathsf{BN}} \mathbb{R}^n \tag{2}$$

where $\mathsf{Aff} : \mathbb{R}^n \to \mathbb{R}^n$ denotes an affine transformation, and BN denotes a batch normalisation layer.

B.4 Path Encoding of Laplacian Eigenvalues

For `MUTAG`, `COX2`, `DHFR`, and `NCI1`, we sort the Laplacian eigenvalues in ascending order and transform the one-dimensional sequence into a two-dimensional time series via a *delay embedding*

$$\begin{bmatrix} \lambda_1 & \lambda_2 & \cdots & \lambda_{N-1} & \lambda_N \end{bmatrix} \mapsto \left[\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} \quad \begin{bmatrix} \lambda_2 \\ \lambda_3 \end{bmatrix} \quad \cdots \quad \begin{bmatrix} \lambda_{N-2} \\ \lambda_{N-1} \end{bmatrix} \quad \begin{bmatrix} \lambda_{N-1} \\ \lambda_N \end{bmatrix} \right]. \tag{3}$$

For `IMDB-B`, we incorporate a fictitious time coordinate $t_j = 2(j-1)/(N-1)$ for $j = 1, \dots N$ as the second coordinate instead of a 'delayed' eigenvalue:

$$\begin{bmatrix} \lambda_1 & \lambda_2 & \cdots & \lambda_{N-1} & \lambda_N \end{bmatrix} \mapsto \left[\begin{bmatrix} \lambda_1 \\ t_1 \end{bmatrix} \quad \begin{bmatrix} \lambda_2 \\ t_2 \end{bmatrix} \quad \cdots \quad \begin{bmatrix} \lambda_{N-1} \\ t_{N-1} \end{bmatrix} \quad \begin{bmatrix} \lambda_N \\ t_N \end{bmatrix} \right]. \tag{4}$$

C Maps between Manifolds

We recall a smooth map $f : M \to N$ between manifolds with or without boundary is

- A *smooth immersion* if at each point p, the rank of the Jacobian $\mathrm{d}f_p : T_pM \to T_pN$ is equal to the dimension of the domain M;
- A *smooth submersion* if at each point p, the rank of $\mathrm{d}f_p$ is equal to the dimension of the codomain N;
- A *smooth embedding* if f is a smooth immersion, and f is a topological embedding, i.e. a homeomorphism onto its image.

A smooth map $f : M \to N$ between manifolds with or without boundary is a *local diffeomorphism* if for every point $p \in M$ there is a neighbourhood $U \ni p$ such that $f(U)$ is open in N and $f|_U : U \to f(U)$ is a diffeomorphism.

We recall the following two theorems for maps between manifolds without boundary of the same dimension.

Theorem C.1 (Invariance of Domain)**.** . *Consider two topological manifolds M and N of the same dimension. If U is an open set of M and $f : U \to N$ is an injective, continuous map, then $f(U)$ is open in N and f is a homeomorphism onto $f(U)$.*

Theorem C.2 (Inverse Function Theorem)**.** . *Let $f : M \to N$ be a smooth map between smooth manifolds M and N without boundary of the same dimension. If the Jacobian $\mathrm{d}f_p$ is invertible at some point $p \in M$, then there are connected neighbourhoods $U \ni p$ and $V \ni f(p)$ such that $f|_U : U \to V$ is a diffeomorphism.*

Theorem C.3. *Let $f : M \to N$ be a smooth map between smooth manifolds M and N without boundary of the same dimension. f is a local diffeomorphism if and only if it is a smooth submersion and smooth immersion. Furthermore, if f is bijective, then f is a diffeomorphism.*