

University of Southampton

Faculty of Social Sciences

School of Mathematical Sciences

Properties of Generalised Persistence Modules

Megan Palser

Thesis for the Degree of Doctor of Philosophy

September 2021

Abstract

The stability of persistent homology is rightly considered to be one of its most important properties, but persistence is still sensitive to choices of metrics, indexing sets, and methods of filtering. This thesis will expand upon previous discussions around stability, considering sources of invariance and symmetry, as well as potential sources of instability. While homology is a large-scale feature which is invariant under homotopy, transferring to the persistent setting does not preserve all of these properties. In this thesis, we show that there exists an excision property for persistent homology. This is a new result even for one-dimensional persistence, but we'll show that this result holds for persistence modules indexed by any partially ordered set.

We will then expand on the theory of generalised persistence modules, building on work by Lesnick (?) and Bubenik, de Silva and Scott (?). In this thesis we have adapted slightly the definition of a generalised interleaving as seen in (?) so that we can recover the usual definition of an interleaving of one-parameter persistence modules defined in (?) as a special case. We will give a new proof of the stability of this generalised distance, and describe how geometric symmetries in data result in interleavings between Vietoris-Rips filtrations. Finally, we will consider reparameterisations of persistence modules, which can be used to rescale persistence. We'll investigate the effect that reparameterisations have on the interleaving distance. We show that a reparameterisation by a Lipschitz order isomorphism is also a Lipschitz map, and we present several stability results for the operation of rescaling persistence.

Acknowledgements

I must first thank my supervisor, Jacek Brodzki, for four years of advice and support in producing this thesis. It has been a privilege to spend such this time researching in such an exciting field and I must thank Jacek for the belief he has shown in me in the course of this work. The last few years have been a strange and challenging time in which to work and I thank him for the efforts he has gone to in order to continue to support me through such upheaval.

Thank you to my examiners, Ran Levi and Jan Spakula for the delightful experience of being able to discuss my work with them, and also for their helpful comments and observations about some areas of my work. Having these perspectives were invaluable.

Thanks also to the other members of the TDA research group at Southampton, most especially Matt Burfitt, Ingrid Membrillo-Solis and Mariam Pirashvilli for the advice, the discussions and the great times at many conferences.

My enormous thanks and gratitude to my fellow students over the last few years. To Holly Paveling, Naomi Andrew, Tom Brown, Vlad Vankov, Sam Hughes, Lily Clements, Guy Boyde, Simon Rea, George Simmons, Matt Staniforth and so, so many others, I cannot thank you enough for making this experience so enjoyable. Your friendship is one of the best things I will take with me from Southampton.

In speaking of great things I take with me from Southampton, my enormous thanks next to Josh McCormick, who has supported me in ways I cannot articulate in these pages. Thank you for being there for every step of this thesis.

Thank you to the best friends anyone could ever hope to have, to Zoe Palmer, Ruby Nightingale, Tilly Kusmishko, Beth Wright, Ellie Atkinson, Robbie Diserens and Ryan Kirkness. For encouraging me and consoling me, thank you – I would not have made it here without you.

My thanks also to Pete Hibbert, Lorna Benson, Charlie Roskams, Chris Sloggett, Doug Larkam, Julie House, Rob Noye and Zak Ozeren. You remind me every day why I am proud to be a mathematician.

This thesis is dedicated to those whom I owe the most thanks. To my parents and my sister, Eleanor. For your unconditional help, support, and love, you are the greatest family I could ask for.

Contents

Introduction

Persistent homology is an algebraic method of describing a topological process which is evolving according to some parameters. This field is, within the scope of mathematics, a very young area of research. The beginnings of persistence can be seen independently in the doctoral work of Vanessa Robbins (?) and the work of Patrizio Frosini and Massimo Ferri and collaborators (?) in their work on size functions in the 1990s. Persistence as we know it now was then born with the publication of the seminal paper by Edelsbrunner, Letscher and Zomorodian (?) in 2002.

Consider the process in Figure ???. For each value of t , we have a simplicial complex. We can think of this diagram as showing a simplicial complex which is growing as the parameter t increases.




Figure 1: A filtration of a simplicial complex by parameter t .

This sequence of nested spaces shown in Figure ??? is an example of a filtered topological space. As the parameter increases, the topology of the subspaces may change. For example, in Figure ???, we begin with four connected components, which merge into two, and finally one component, as t increases. As well as the number of connected components, higher homology classes also appear – and, sometimes, disappear – as we progress through the filtration. Any classes which survive over multiple values of the parameter we define to be *persistent*. Rather than consider the individual homology groups, $H_n(X(t=i))$, for any particular $i \in \{1, \dots, 4\}$, it is more interesting to look for those features which *persist* from $t=i$ to $t=j$ for some $j \geq i$.

One of the reasons why persistence has become such a popular area of research is the way in which it has been utilised in the new field of topological data analysis. Gunnar Carlsson discusses how persistent homology can allow us to discover the “shape of data” (?). Take for example the first subspace, $X(t=1)$, in the filtration in Figure ???. The topology of four disconnected points is rather uninteresting, but by connecting them in a certain way, we create a sequence of subspaces whose topology is much more informative.

Vietoris-Rips and Čech filtrations provide methods of constructing sequences of simplicial complexes from collections of data points, which may be large in both cardinality and dimension. We may not be able to *see* any structure in a large cloud of points embedded in \mathbb{R}^d , if d is very large, but by using methods in persistent homology, that shape is there to

discover, and there to exploit. The successful uses of topological data analysis demonstrate its power: by comparing topological features of data sets, researchers have been able to discover hitherto unknown phenomena in fields such as medicine (?), network theory (?), shape classification (?) and many others. See (?) for a survey of some of the ways persistent homology has been used in practice. One of the most lauded examples of the power of persistent homology in detecting unseen shape in data is the discovery of a new subclass of breast cancers in (?).

The study of persistent homology is not only of interest for its potential applications, however. The algebraic objects which we construct in persistence are themselves worthy of study. Consider again the filtration in Figure ???. First fix an $n \geq 0$ and a field k . For each $t \in T = \{1, \dots, 4\}$, we have a homology group $H_n(X(t); k)$, and if we choose any $i, j \in T$ such that $i \leq j$, then this relation gives rise first to an inclusion of subspaces, $X(t = i) \hookrightarrow X(t = j)$, and subsequently to an induced linear map $H_n(X(t = i); k) \rightarrow H_n(X(t = j); k)$. Overall, the ordered set, $1 \leq 2 \leq 3 \leq 4$ gives rise to the diagram,

$$H_n(X(t = 1)) \rightarrow H_n(X(t = 2)) \rightarrow H_n(X(t = 3)) \rightarrow H_n(X(t = 4)). \quad (0.1)$$

This diagram is nothing more than the composition of two functors - the first from the ordered set T to the category of subcomplexes of $X(t = 4)$, and the second to the category of k -vector spaces. Such a diagram is an example of what we call a persistence module.

The study of persistence modules is therefore nothing more than the study of functors with certain properties. Many concepts in persistent homology are examples of concepts in category theory, but which have only been explored because of their applications to topological data analysis. For instance, without the motivation of computing distances between persistence modules, it is unlikely that the concept of an interleaving between functors would have been defined.

By considering the example in Figure ??? as a special case of a more general picture, we can widen the scope of persistence. Though originally defined in the one-dimensional case, where persistence modules are indexed by a single, linear parameter, there is by now considerable interest in more general cases (?). “Multiparameter” (?) persistence involves the study of persistence modules indexed by some subset of \mathbb{R}^n for $n \geq 2$. Consider, for example, the bifiltration in Figure ???, where our topological space evolves according to two parameters, γ_1 and γ_2 .

Though the one-parameter case is reasonably well-understood, multiparameter persistence is much less so (?). Many concepts which are ubiquitous in one-parameter persistence, such as a barcode, a persistence diagram and the matching distance, cannot be defined when we introduce additional parameters. It helps, in this case, to treat persistence in a more general way.

How, then, should we define a persistence module? According to Bubenik, de Silva and Scott (?), a persistence module is a functor,

$$\mathcal{M} : P \rightarrow \mathcal{C}, \quad (0.2)$$




Figure 2: A filtration indexed by parameters γ_1 and γ_2 .

where P is a partially ordered set, and \mathcal{C} is any category. Defining a persistence module in this way, where we do not specify a target category, means that concepts which crop up across the persistence literature – filtrations, diagrams of vector spaces such as the one in (?), and even Reeb graphs (?) – become examples of persistence modules. Allowing modules to be indexed by any poset means our generalised treatment encompasses one-parameter, multiparameter, and even zig-zag persistence (?). In this thesis, we’ll follow the conventions of (?). There are so many varying conventions, with authors often influenced by their expertise within particular fields of pure mathematics, that it is almost impossible to describe a standard treatment of persistence. This provides a significant motivation to study persistence in the most general way possible.

Outline of Thesis

This thesis consists of a preliminary section which contains the necessary definitions regarding categories, posets and metric spaces, followed by three sections, titled ‘Homology Theories and Persistence’, ‘Interleavings’, and ‘Poset Morphisms and the Interleaving Distance’.

Section One

In the first section we will define a generalised persistence module, and some well-known filtrations which are often used in topological data analysis. These include Vietoris-Rips and Čech filtrations on point cloud data, and sublevel-set filtrations.

In this section, we will investigate the extent to which persistent homology satisfies the properties of a homology theory. Instinctively, we know that there are many differences between the homology of a single space and the persistent homology of a filtered space. While homology is a large-scale feature, persistence depends not only on the space in question, but also on the filtration. It can be thought of as small-scale homology, with the fineness of the scale dictated by the indexing poset. Despite the obvious differences, it has been shown that many properties of a homology theory do extend – at least on some level – to the persistent setting. The existence of a long exact sequence of a pair, a Mayer-Vietoris sequence for a triple, and an excision property are all useful tools in calculating unknown homology groups (?), and so the existence of analogous properties for persistent homology could also be useful computationally.

In the first section, we’ll summarize the results of Di Fabio and Landi in (?), where it was shown that there exists a variant of the Mayer-Vietoris sequence in Persistent homology. Likewise, in (?), a similar method to that of (?) was used to prove the existence of a long exact sequence of persistence modules. However, in both of these cases, the sequences obtained are not on-the-nose generalisations of the homological properties - if we attempt to write a Mayer-Vietoris sequence in persistent homology by simply replacing homology groups with their persistent counterparts, with the appropriate restrictions of the connecting maps, then the resulting sequence fails to be exact in general. The same is true if we attempt to restrict a long exact sequence of a pair in homology to the persistent homology groups.

The prospect of generalising homological properties to the persistent setting is therefore not as straightforward as we may hope, and the results of (?) and (?) indicate that we do not have a set of Eilenberg-Steenrod axioms which hold for persistent homology groups. In contrast to the results of (?) and (?), however, we show in section 1 that there is an excision property both for persistence modules and for persistent homology groups. This result holds for modules indexed by any partially ordered set. In particular there exists an excision property for multiparameter persistence modules.

We can then use this result, together with the long exact sequence for persistence modules of

(?), to give an alternative proof of the Mayer-Vietoris-style sequence in persistent homology groups of (?). In doing so, we are able to generalise this result to modules indexed by arbitrary posets, and for a much wider range of filtrations. In particular, we can show that there exists a Mayer-Vietoris sequence for multiparameter as well as single-parameter persistence modules.

Section Two

In the second section, we'll introduce the definition of an interleaving between persistence modules, and the pseudometric which can be defined from this concept. The interleaving distance was originally defined only for one-parameter modules by Chazal et. al in (?), and was subsequently expanded to multiparameter modules, indexed by \mathbb{R}^n , by Lesnick (?). In (?), the definition was generalised to persistence modules indexed by arbitrary posets. In this thesis, we'll adapt the definition of (?). The new definition ensures that we recover the definition initially given in (?) as a special case of this more general one.

Calculating the proximity between persistence modules is clearly important in the context of applications in topological data analysis: if we can calculate the distance between the outcomes of a persistence process, then this gives a proxy for the distance between the data sets themselves. Stability theorems – first proved in (?) and (?) – give guarantees that the interleaving distance really does provide a good measure of similarity between input data, by showing that the interleaving distance between a pair of persistence modules is bounded above by the proximity between the inputs. We'll discuss the precise wording and implications of these stability results in greater detail in section ??.

A downside of the stability results of (?) and (?) is that they apply only to one-parameter persistence modules. In (?), after introducing their more general interleaving distance, Bubenik, de Silva and Scott present a stability result for this pseudometric, which is a generalisation of the result of (?). In section 2, we use a remark of Lesnick's in (?) to prove the stability of the generalised interleaving distance defined in this thesis. Namely, Lesnick suggests that we may consider the distance between \mathbb{R}^n -valued functions as an interleaving distance in its own right. We use this idea, to show that an interleaving of poset-valued filtering functions gives rise to an identical interleaving of the corresponding persistence modules. This stability theorem for the generalised interleaving distance is the main result of Section 2.

In Section 2, we consider some of our first examples of transformations in Persistent Homology. Vietoris-Rips filtrations are a method of constructing a filtered topological space from a finite metric space. This method is often used in topological data analysis to construct a filtration from a point cloud. However, with many examples of real world data, there is no canonical choice of metric to use. Our other main result in Section 2 shows that the Vietoris-Rips persistence modules of metric spaces which are Lipschitz or Coarsely equivalent are interleaved.

Section Three

A persistence module is simply a functor from a poset to a category, a process which can be described by the diagram in Figure ??.

$$\boxed{\text{poset } P} \longrightarrow \boxed{\text{category } \mathcal{C}} \quad (0.3)$$

In (?), the authors discuss the possibility of composing this process in (??) with additional functors, with a diagram similar to the one below (??).

$$\boxed{\text{poset } Q} \longrightarrow \boxed{\text{poset } P} \longrightarrow \boxed{\text{category } \mathcal{C}} \longrightarrow \boxed{\text{category } D} \quad (0.4)$$

In (?), they consider the effect of composing with the arrow on the right of the diagram in (??) on the interleaving distance. In Section 3, we'll consider the dual question – what is the effect of precomposing with a morphism of posets, the arrow on the *left* of the diagram in (??)?

This operation was discussed in (?) with a view towards discretizing modules, and in (?), where the motivation is to develop a relative version of the interleaving distance in cases where calculating interleavings in a given poset is not possible, due to the structure of the poset.

In Section 3 we show that Lipschitz properties of structure-preserving poset maps descend to identical properties for the pull-backs. This stability result for the operation of reparameterising persistence modules is our first main result of section 3. Our second main result in this section shows that the interleaving distance between two pull-back modules is bounded by the distance between the morphisms.

A further main result serves as a note of caution, however. We show that for Kan extensions of persistence modules along arbitrary poset maps, the interleaving distance between these modules can be distorted. This result shows that discretizing persistence modules for the purpose of simplifying computations should be done with caution, since the process is not necessarily stable.

Preliminaries

Categories and Posets

Definition 0.1. A **partial order** on a set, P is a relation \leq , defined on some subspace of $P \times P$ which is,

- Reflexive: for every $a \in P$, $a \leq a$.
- Anti-symmetric: if $a \leq b$ and $b \leq a$ then $a = b$.
- Transitive: if $a \leq b$ and $b \leq c$, then $a \leq c$.

A set P together with a partial order \leq is called a **partially ordered set**, and is denoted (P, \leq) , or simply P if \leq is clear.

Example 0.2. The real or natural numbers, together with the usual order \leq , is actually a totally ordered set, as any two elements a, b of \mathbb{R} are related by either $a \leq b$ or $b \leq a$.

Example 0.3. \mathbb{R}^n , together with the relation, $(a_1, \dots, a_n) \leq (b_1, \dots, b_n)$ whenever $a_i \leq b_i$ for every $i \in \{1, \dots, n\}$.

Throughout this thesis, we will be using the example of \mathbb{R}^n with the partial order described in Example ?? quite often. Any reference to the poset (\mathbb{R}^n, \leq) will be referring to the partial order of Example ??, unless stated otherwise. There are of course, multiple possible partial orders on \mathbb{R}^n , not all of which abide by any intuitive idea of ‘ordering’, as the following example shows.

Example 0.4. \mathbb{R}^n together with the relation \preceq is a partially ordered set, where $(a_1, \dots, a_n) \preceq (b_1, \dots, b_n)$ whenever $b_i \geq a_i$ for every $i \in \{1, \dots, n\}$, and \leq denotes the usual order on \mathbb{R} .

A poset P can also be defined as a directed graph. This graph describes the ‘shape’ of the poset, and it is often helpful to keep in mind when discussing posets, or, as we will later, structures indexed by posets.

Definition 0.5. The directed graph associated to a poset (P, \leq) is called the **Hasse diagram** of P , and we will denote it H_P . H_P has a vertex for each element $p_i \in P$, and a directed edge $p_1 \rightarrow p_2$ whenever $p_1 \leq p_2$, and there is no $p \in P$ such that $p_1 \leq p \leq p_2$.

For example, the Hasse diagram of the poset (\mathbb{N}, \leq) can be drawn as in Figure ??, where we only show the ‘generating’ arrows between consecutive integers.

$$1 \longrightarrow 2 \longrightarrow 3 \longrightarrow \dots$$

Figure 3: The Hasse diagram $H_{\mathbb{N}}$ for the poset of natural numbers.

We note that it is common to see Hasse diagrams drawn vertically – that is, if $p_i \geq p_j$ then p_i is drawn vertically above p_j . While this is sometimes informative, it is not necessary or, indeed, always practical.

Definition 0.6. A **category** \mathcal{C} is a collection of objects, $\text{Ob}(\mathcal{C})$, and morphisms $\text{Mor}(\mathcal{C})$, between the objects of \mathcal{C} , such that,

1. for each $A \in \text{Ob}(\mathcal{C})$, there exists an identity morphism, $id_A : A \rightarrow A$,
2. the composition of morphisms is associative. That is, given $f : A \rightarrow B$, $g : B \rightarrow C$, $h : C \rightarrow D$, then

$$h \circ (g \circ f) = (h \circ g) \circ f.$$

A **functor** between categories \mathcal{C} and \mathcal{D} is a map $F : \mathcal{C} \rightarrow \mathcal{D}$, such that for any $A \in \text{Ob}(\mathcal{C})$, $F(id_A) = id_{F(A)}$, and for any morphisms $f : A \rightarrow B$, $g : B \rightarrow C$ in $\text{Mor}(\mathcal{C})$,

$$F(g \circ f) = F(g) \circ F(f).$$

Examples of categories include the category $Vect_k$ of vector spaces over a field, k , where the morphisms are given by k -linear maps, or the category Top of topological spaces and continuous maps.

Remark 0.7. A poset (P, \leq) is also an example of a category, which has an object for each element of the set P , and a *unique* morphism $p_1 \rightarrow p_2$ whenever $p_1 \leq p_2$ in the partial order. We note that as the map $p_1 \rightarrow p_2$ is unique, a poset is therefore an example of a *thin* category. The upshot of this, as (?) points out, is that if we have a diagram,

$$\begin{array}{ccc} p_1 & \longrightarrow & p_2 \\ \downarrow & & \downarrow \\ p_3 & \longrightarrow & p_4 \end{array} \tag{0.5}$$

then this diagram must necessarily commute, since there is a unique map $p_1 \rightarrow p_4$. We will make use of this fact later.

Homology

One significant example of a functor is homology. Homology is a functor from the category of topological spaces and continuous maps, to the category of vector spaces and linear maps. We will give an overview of simplicial homology, which is more suited to use in applications than, say, singular homology, as it can be computed using matrix methods. See (?) for a description, accessible to non-mathematicians, of how simplicial homology can be computed using such methods.

Definition 0.8. An n -simplex is a generalisation of a tetrahedron in n dimensions.

Some low-dimensional simplices are shown in Figure ?? below. A 0-simplex is a single vertex, a 1-simplex is a line segment, a 2-simplex is a triangle and a 3-simplex is a tetrahedron. We denote a simplex on vertices x_0, \dots, x_n by $[x_0, \dots, x_n]$.



Figure 4: Left to right: A 0-simplex, a 1-simplex and a 2-simplex.

Definition 0.9. A simplex $\sigma_1 = [x_0, \dots, x_n]$ is a **face** of $\sigma_2 = [y_0, \dots, y_m]$ if $\{x_0, \dots, x_n\} \subseteq \{y_0, \dots, y_m\}$.

A simplicial complex is a union of simplices, closed under taking faces, such that the simplices are attached in a coherent way, meaning that the intersection of any two simplices σ_1 and σ_2 in the complex is another simplex, σ_3 , which is a face of both σ_1 and σ_2 .

Definition 0.10. A **simplicial n -chain** is a linear combination of n -simplices,

$$\sum_{i=1}^k a_i \sigma_i, \tag{0.6}$$

for some a_i in a commutative ring, R . Denote the set of all simplicial n -chains in a particular simplicial complex, X , with coefficients in R , by $C_n(X; R)$.

This set $C_n(X; R)$ is the free abelian group on the set of oriented n -simplices, where the operation is given by addition of chains:

$$\sum_{i=1}^k a_i \sigma_i + \sum_{i=1}^k b_i \sigma_i = \sum_{i=1}^k (a_i + b_i) \sigma_i. \tag{0.7}$$

Definition 0.11. The **boundary map**, $\partial_n : C_n(X; R) \rightarrow C_{n-1}(X; R)$ is given by,

$$\partial_n([x_0, \dots, x_n]) = \sum_{i=0}^n (-1)^i [x_0, \dots, \hat{x}_i, \dots, x_n], \tag{0.8}$$

for an oriented n -simplex, $[x_0, \dots, x_n]$, where \hat{x}_i denotes that vertex x_i is omitted from that term in the sum. The boundary map extends to an arbitrary n -chain via:

$$\partial_n \left(\sum_{i=1}^k a_i \sigma_i \right) = \sum_{i=1}^k a_i \partial_n(\sigma_i). \tag{0.9}$$




Figure 5: An oriented 2-simplex on vertices x_0, x_1 and x_2 .

If a chain, c , is in the kernel of the n^{th} boundary map – that is, $\partial_n(c) = 0$ – then c is a cycle of n -simplices. For example, the chain,

$$c = [x_0, x_1] + [x_1, x_2] + [x_2, x_0],$$

in Figure ?? is in the kernel of ∂_1 as

$$\partial_1(c) = \partial_1([x_0, x_1]) + \partial_1([x_1, x_2]) + \partial_1([x_2, x_0]) = x_1 - x_0 + x_2 - x_1 + x_0 - x_2 = 0.$$

If a chain is in the image of ∂_{n+1} then it is the boundary of some other chain of $(n+1)$ -simplices. For example, c is also in the image of ∂_2 , as it is the boundary of the 2-simplex, $[x_0, x_1, x_2]$:

$$\partial_2([x_0, x_1, x_2]) = [x_1, x_2] - [x_0, x_2] + [x_0, x_1] = [x_1, x_2] + [x_2, x_0] + [x_0, x_1] = c.$$

It is no coincidence that c is both in the image of ∂_2 and the kernel of ∂_1 . In fact we have that $\text{im}(\partial_{n+1}) \subseteq \ker(\partial_n)$ for every n . That is, the boundary of every $(n+1)$ -chain is an n -cycle (?).

The sequence of chain groups and boundary maps,

$$\cdots \rightarrow C_n(X; R) \xrightarrow{\partial_n} C_{n-1}(X; R) \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_1} C_0(X; R) \xrightarrow{\partial_0} 0, \quad (0.10)$$

is hence an example of a **chain complex** – a sequence of abelian groups and connecting group homomorphisms such that the image of each map is contained in the kernel of its successor. That is, for each $n \in \mathbb{N}$, $\text{im}(\partial_{n+1}) \subseteq \ker(\partial_n)$. If this inclusion is in fact an equality, then the sequence of groups and group homomorphisms is called **exact**.

The aim is then to find those cycles which do not bound any part of the simplicial complex, and so, must bound a *hole* of the complex, as the following definition clarifies.

Definition 0.12. The n^{th} **simplicial homology group** of a simplicial complex, X , $H_n(X; R)$, is given by the quotient,

$$H_n(X; R) = \frac{\ker(\partial_n) : C_n(X; R) \rightarrow C_{n-1}(X; R)}{\text{im}(\partial_{n+1}) : C_{n+1}(X; R) \rightarrow C_n(X; R)}. \quad (0.11)$$

The group $H_n(X; R)$ is generated by those chains which are cycles which do not form boundaries of any part of the complex, and so instead must bound an n -dimensional 'hole'.

If the chain complex (??) is an exact sequence, then $\text{im}(\partial_{n+1}) = \ker(\partial_n)$, and so the homology groups, H_n , are all zero. In general, the homology groups are a measure of the extent to which the chain complex (??) *fails* to be an exact sequence (?).

Definition 0.13. The n^{th} **Betti number** of a simplicial complex X is given by $\beta_n(X) = \text{rank}(H_n(X))$.

Often, we may be interested in the homology of only part of the simplicial complex, X , effectively ignoring the chains which lie outside of a particular subcomplex, A , of X . This concept is called **relative homology**.

Definition 0.14. Let A be a subcomplex of X . Define the **group of simplicial n -chains of X relative to A** as the quotient,

$$C_n(X, A) = \frac{C_n(X)}{C_n(A)}. \quad (0.12)$$

We can then form the relative chain complex,

$$\cdots \rightarrow C_n(X, A) \xrightarrow{\partial_n} C_{n-1}(X, A) \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_1} C_0(X, A) \xrightarrow{\partial_0} 0, \quad (0.13)$$

where the maps ∂_i are just the usual boundary maps, since the boundary of a relative chain must lie wholly in X .

Definition 0.15. The n^{th} **relative homology group** of the pair (X, A) is the n^{th} homology group of the chain complex,

$$H_n(X, A) = \frac{\ker(\partial_n) : C_n(X, A) \rightarrow C_{n-1}(X, A)}{\text{im}(\partial_{n+1}) : C_{n+1}(X, A) \rightarrow C_n(X, A)}. \quad (0.14)$$

Example 0.16. Let X be a wedge of two circles, as shown in Figure ???. Let $A \subset X$ be the right-hand circle in Figure ???. Then the relative homology group $H_1(X, A)$ is given by,

$$H_1(X, A; k) \cong k,$$

since this group is generated by the 1-cycle of $\frac{C_1(X)}{C_1(A)}$.




Figure 6: The pair of topological spaces (X, A) .

Metric Spaces

Definition 0.17. A **metric** on a set M is a function $d : M \times M \rightarrow \mathbb{R}_{\geq 0}$ such that for all $a, b, c \in M$,

1. $d(a, b) = 0$ if and only if $a = b$,
2. d is symmetric, so $d(a, b) = d(b, a)$,
3. d satisfies the triangle inequality, so $d(a, c) \leq d(a, b) + d(b, c)$.

A set M together with a metric d is called a **metric space**, denoted (M, d) .

Definition 0.18. A map $f : (M, d_1) \rightarrow (N, d_2)$ between metric spaces (M, d_1) and (N, d_2) is called **Lipschitz** if there exists a constant $C \geq 0$ such that for all $m_1, m_2 \in M$,

$$d_2(f(m_1), f(m_2)) \leq C d_1(m_1, m_2).$$

Such a map f is called **bi-Lipschitz** if there exists a constant $C \geq 1$ such that for all $m_1, m_2 \in M$,

$$\frac{1}{C} d_1(m_1, m_2) \leq d_2(f(m_1), f(m_2)) \leq C d_1(m_1, m_2).$$

The constant C is called the **Lipschitz constant** for f .

Definition 0.19. A **Lawvere metric space** is a set X together with a function $d : X \times X \rightarrow [0, \infty]$, called a Lawvere metric, such that,

1. $d(x, x) = 0$ for all $x \in X$,
2. for all $x, y, z \in X$, $d(x, z) \leq d(x, y) + d(y, z)$.

A Lawvere metric is therefore a function which has two properties of a metric, but which does not require symmetry (that $d(x, y) = d(y, x)$ for all $x, y \in X$), and does not require that $d(x, y)$ is non-zero for all distinct pairs x, y . Points in a Lawvere metric space can also be infinitely far from one another.

As well as a Lawvere metric, here we introduce a further variant of a metric, which will crop up several times in our discussion of ‘distances’ in persistent homology.

Definition 0.20. An **extended pseudometric space** is a set X together with a function $d : X \times X \rightarrow [0, \infty]$, called an extended pseudometric, such that,

1. $d(x, x) = 0$ for all $x \in X$,
2. $d(x, y) = d(y, x)$ for all $x, y \in X$.

3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$.

We end this preliminary section with some ways to describe the relationships between two metric spaces. First we have the concept of two metric spaces being coarsely equivalent.

Definition 0.21. A map $f : (M, d_1) \rightarrow (N, d_2)$ of metric spaces is called **coarse** if for every bounded subset B of N , the set $f^{-1}(B)$ is bounded in M , and for every $\varepsilon_i > 0$, there exists a real number $\delta_i > 0$ such that for every $m_1, m_2 \in M$ with $d(m_1, m_2) \leq \varepsilon_i$, we have $d(f(m_1), f(m_2)) \leq \delta_i$.

Definition 0.22. Metric spaces (M, d_1) and (N, d_2) are said to be **coarsely equivalent** if there exists a pair of coarse maps $f : (M, d_1) \rightarrow (N, d_2)$ and $g : (N, d_2) \rightarrow (M, d_1)$ such that $f \circ g$ and $g \circ f$ are close to the identity maps on N and M , respectively. That is, there exist real numbers $c_1, c_2 \geq 0$ such that for all $m \in M$,

$$d_1(m, g \circ f(m)) \leq c_1, \tag{0.15}$$

and for all $n \in N$,

$$d_2(n, f \circ g(n)) \leq c_2. \tag{0.16}$$

For more information about coarse maps and coarsely equivalent metric spaces, see (?). Finally, if we wish to consider the distance between two metric spaces, one way is to use the Gromov-Hausdorff distance.

Definition 0.23. The **Gromov-Hausdorff distance** between metric spaces (X, d_X) and (Y, d_Y) is given by,

$$d_{GH}(X, Y) = \inf_{Z, f, g} d_H(f(X), g(Y)), \tag{0.17}$$

where $f : X \rightarrow Z$ and $g : Y \rightarrow Z$ are isometric embeddings of X and Y into a third metric space, (Z, d) , and d_H denotes the Hausdorff distance,

$$d_H(f(X), g(Y)) = \max \left\{ \sup_{x \in X} \inf_{y \in Y} d(f(x), g(y)), \sup_{y \in Y} \inf_{x \in X} d(f(x), g(y)) \right\}. \tag{0.18}$$

1 Homology Theories and Persistence

This section is dedicated to persistence modules, which are the main algebraic objects of study in persistence. There is, by now, a vast literature devoted to persistent homology, with many works motivated by applications to topological data analysis. Different authors may have varying applications in mind, which lead to different conventions, and, often, different definitions of phrases including ‘persistence’, ‘persistence modules’ and ‘persistent homology’. Often, persistence will be defined as a process indexed by a subset of the real or natural numbers. This case, sometimes referred to as “one-parameter persistence” is very important - many of the most important examples of how persistent homology can be applied in data analysis make use of the one-parameter version and this is certainly the most well-understood case. See, for example, (?) for a survey of the use of persistent homology in topological data analysis. In Section 2, we will discuss the presentations of one-parameter persistence modules. The unique structure of these objects allow us to define complete, discrete invariants which prove extremely useful, especially if we wish to use persistence in practice. It is much easier to perform computations with these discrete invariants than with algebraic objects such as persistence modules. The same invariants cannot be defined for anything other than the one-parameter case. It is partly for these reasons that one-parameter persistence is the most studied, as well as the most used in applications.

There is, however, a growing interest in many more general cases. Persistence which is indexed by a subset of \mathbb{R}^n , for $n > 1$, is sometimes referred to as “multiparameter persistence” (?). As the name suggests, a multiparameter persistence module can be used to describe the homology of a space filtered by multiple parameters. This case is a generalisation of the one described above, and a complete understanding of the multiparameter case would be of huge significance - rarely can any real-world process be completely described by the influence of just one parameter. Zig-zag persistence (?) provides an even more interesting example, in which persistence modules may be indexed by diagrams of the form,

$$\begin{array}{ccccccc} & a & & c & & \dots & \\ & \swarrow & & \swarrow & & \swarrow & \\ & & b & & d & & \\ & \searrow & & \searrow & & \searrow & \end{array} \quad (1.1)$$

Going even further, Bubenik, de Silva and Scott (?) redefine a persistence module in the most general way possible, as a functor from any partially ordered set to any category. In Section 2, we’ll go on to build upon the work of (?) on generalised interleavings, so will follow their definitions in this section.

In a 2008 paper (?), Robert Ghrist describes persistent homology as “a homology theory for point cloud data”. The phrasing may have been intended innocuously enough, but it raises the question - to what extent is persistent homology a ‘homology theory’, in the strict sense of the term? In other words, which well-known properties of a homology theory continue to hold when we consider filtered spaces, and are working with persistent homology groups?

Firstly, homology with field coefficients is a functor from the category of topological spaces and continuous maps, to the category of vector spaces over that field, and linear maps. Persistent homology, again with coefficients in a field k , is a functor from a filtered topological space to the category of k -vector spaces and k -linear maps. But it is not a given that a continuous map of topological spaces induces a morphism in the persistent homology of those spaces. A recent paper (?) tackles the question “what is persistent homology a functor of?”. Specifically, given two finite sets X and Y , and two sets of functions, $\Phi = \{\phi : X \rightarrow \mathbb{R}\}$, and $\Psi = \{\psi : Y \rightarrow \mathbb{R}\}$, it is shown in (?) that a map $\alpha : \Phi \rightarrow \Psi$ only induces morphisms in persistent homology, $PH(\Phi) \rightarrow PH(\Psi)$, when there exists a map $f : X \rightarrow Y$ such that for every $\phi \in \Phi$, the following diagram commutes,

$$\begin{array}{ccc}
 X & & \\
 \downarrow f & \searrow \phi & \\
 & & \mathbb{R} \\
 & \nearrow \alpha(\phi) & \\
 Y & &
 \end{array} \tag{1.2}$$

As well as the property of being functorial, the Eilenberg-Steenrod axioms for a homology theory can provide many ways in which unknown homology groups can be calculated from known ones. A long exact sequence in homology for a pair (X, A) connects the relative homology $H_n(X, A)$ to the homology of the individual spaces, $H_n(X)$ and $H_n(A)$. The excision property says that $H_n(X, A)$ is preserved after removing a subspace from both X and A . Subsequently, using these axioms, we can also prove the existence of an additional exact sequence in homology. The Mayer-Vietoris sequence connects the homology of a space to that of a pair of subspaces, and can be thought of as an analogue of the Van Kampen theorem for homology.

Among the first to consider how homological properties transfer to the persistent setting was a 2011 paper by Di Fabio and Landi (?), which proved the existence of a form of a Mayer-Vietoris sequence for a triple, (X, A, B) , in the specific case that X, A and B are all filtered by sublevel set filtrations with respect to the same function. However, the authors note that their Mayer-Vietoris sequence,

$$\cdots \rightarrow H_{n+1}^{i,j}(X, A) \xrightarrow{\partial_n^{i,j}} H_n^{i,j}(A) \xrightarrow{\iota_n^{i,j}} H_n^{i,j}(X) \xrightarrow{\kappa_n^{i,j}} H_n^{i,j}(X, A) \rightarrow \cdots \tag{1.3}$$

written with persistent homology groups, is not, in general, an exact sequence, but is a chain complex. Despite this, they demonstrate that this property is sufficient to guarantee that parts of the persistence diagrams of the subspaces A and B appear as ‘signatures’ in the persistence diagram for X , suggesting that we may be able to infer the persistent homology of a space – at least to some degree – from that of a pair of subspaces.

Following this work, (?) showed that though the Mayer-Vietoris-style sequence in persistent homology groups fails to be exact in general, there does in fact exist an exact Mayer-Vietoris sequence of persistence *modules*. In a similar vein, they show that an attempt to construct

a long exact sequence for a pair in persistent homology groups results in a sequence,

$$\cdots \longrightarrow H_n^{i,j}(X) \xrightarrow{\kappa_X} H_n^{i,j}(X, A) \xrightarrow{\partial_X} H_{n-1}^{i,j}(A) \xrightarrow{\iota_A} H_{n-1}^{i,j}(X) \longrightarrow \cdots \quad (1.4)$$

which, like (??), is not exact in general, but is a chain complex. However, just as for the Mayer-Vietoris sequence, they show that there does in fact exist a long exact sequence of persistence *modules* for the pair (X, A) .

We will describe these sequences for persistent homology groups in greater detail, including the maps in the diagrams, in section ??.

The question of the existence of an excision theorem for persistent homology is an interesting one. Persistent homology can be thought of as a small scale version of homology. While homology is a large-scale feature, persistence describes a sequence of small changes. Removing a subspace may result in a very different filtration. It is certainly not a given that removing a subspace from a pair of spaces should preserve the persistent homology. It is important to note that a version of the property has been used implicitly in defining extended persistence (?) and in Bendich's (?) work defining a persistent version of local homology, though an excision theorem for persistent homology had not been formally stated or proved until the work of this thesis. This result also appeared in my earlier work (?).

In this section we prove that the excision property possessed by any homology theory descends to an identical property in persistent homology. We show that such a property holds both for persistence modules and for persistent homology groups. This is markedly different to the property of Mayer-Vietoris sequences, and long exact sequences for a pair, which only fully hold for persistence *modules* as opposed to persistent homology *groups*. Although the result for persistence modules has been used implicitly without proof in (?) and (?), the result for persistent homology groups is completely new.

We then use this excision result, together with the result of (?) which shows that when we restrict the long sequence for a pair to persistent homology groups, the resulting sequence is a chain complex, to give an alternative proof of the existence of a Mayer-Vietoris sequence in persistent homology groups, the result first seen in (?). In doing so, we are able to extend the result of (?) to a much wider class of filtered triples than the original result. In fact, given any filtered space X , and a pair of subspaces, A and B , we demonstrate a way to extend the filtration of X to one of the subspaces, in such a way that the conditions needed for the excision theorem, and the Mayer-Vietoris sequence, are guaranteed to hold.

1.1 Persistence Modules

Definition 1.1. Let P be a partially ordered set. A **P -persistence module**, or persistence module indexed by P , is a functor $\mathcal{M} : P \rightarrow \mathcal{C}$, where \mathcal{C} is any category. That is, to each $p \in P$, we associate an object $\mathcal{M}(p)$, and whenever $p_1 \leq p_2$, we have a morphism $\mathcal{M}(p_1) \rightarrow \mathcal{M}(p_2)$, which we denote by $\psi_{\mathcal{M}}(p_1, p_2)$. We call this set of morphisms $\Psi =$

$\{\psi_{\mathcal{M}}(p_i, p_j) | p_i \leq p_j\}$ the **transition maps** of \mathcal{M} . The definition of a functor between categories tells us that the transition maps must satisfy the following properties:

1. for every $p \in P$, the map $\psi_{\mathcal{M}}(p, p) : \mathcal{M}(p) \rightarrow \mathcal{M}(p)$ is the identity map on $\mathcal{M}(p)$,
2. for every $p_1 \leq p_2 \leq p_3$, the composite,

$$\psi_{\mathcal{M}}(p_2, p_3) \circ \psi_{\mathcal{M}}(p_1, p_2) : \mathcal{M}(p_1) \rightarrow \mathcal{M}(p_2) \rightarrow \mathcal{M}(p_3),$$

is equal to the map $\psi_{\mathcal{M}}(p_1, p_3) : \mathcal{M}(p_1) \rightarrow \mathcal{M}(p_3)$.

Example 1.2. Let (P, \leq) be the set of real numbers together with the usual partial order. Let $I = [a, b]$ be an interval, and let k be a field. The **interval module** $k[I] : \mathbb{R} \rightarrow Vect_k$ is the P -persistence module defined by,

$$k[I](p) = \begin{cases} k & a \leq p \leq b \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad \psi_{k[I]}(p, q) = \begin{cases} id_k & a \leq p \leq q \leq b \\ 0 & \text{otherwise.} \end{cases} \quad (1.5)$$

An interval module $k[I]$ therefore assigns a copy of the field k to each $p \in \mathbb{R}$ contained in the interval I , and a copy of the zero vector space to each real number outside of the interval I . The transition maps, $\psi_{k[I]}(p, q)$ between these vector spaces are then the identity maps in k whenever p and q belong to the interval I , and are the zero maps otherwise.

Example 1.3. Let P be the 6-element poset $\{a, b, c, d, e, f\}$ with partial order as determined by the Hasse diagram shown in Figure ??.




Figure 7: A Hasse diagram for the poset $P = \{a, b, c, d, e, f\}$.

The functor $\mathcal{M} : P \rightarrow Vect_k$ which assigns to each element of P a copy of the field k , and assigns the identity map on k to each relation is a persistence module. This persistence module is shown in Figure ??.




Figure 8: The persistence module $\mathcal{M} : P \rightarrow Vect_k$.

Further examples of persistence modules indexed by P are shown in Figure ??.




Figure 9: The zero module indexed by P (left) and another P -persistence module.

Definition 1.4. A **morphism**, $f : \mathcal{M} \rightarrow \mathcal{N}$ of P -persistence modules, \mathcal{M} and \mathcal{N} , is a family of morphisms, $\{f_p : \mathcal{M}(p) \rightarrow \mathcal{N}(p) | p \in P\}$, such that for each $p_1 \leq p_2$ in P , the diagram,

$$\begin{array}{ccc}
 \mathcal{M}(p_1) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(p_2) \\
 f_p \downarrow & & \downarrow f_q \\
 \mathcal{N}(p_1) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(p_2)
 \end{array} \tag{1.6}$$

commutes.

That is, a morphism $f : \mathcal{M} \rightarrow \mathcal{N}$ is a natural transformation of the functors \mathcal{M} and \mathcal{N} .

Definition 1.5. A morphism $f : \mathcal{M} \rightarrow \mathcal{N}$ is an **isomorphism** of P -persistence modules \mathcal{M} and \mathcal{N} if f_p is an isomorphism for each $p \in P$. Equivalently, f is an isomorphism if there exists a morphism $g : \mathcal{N} \rightarrow \mathcal{M}$ such that for every $p \in P$, $g_p \circ f_p = id_{\mathcal{M}(p)}$ and $f_p \circ g_p = id_{\mathcal{N}(p)}$.

1.2 Persistent Homology

1.2.1 Filtrations

Definition 1.6. Let (P, \leq) be a poset and let X be a topological space. A **filtration** of X is a persistence module $\mathcal{F} : (P, \leq) \rightarrow \text{Subsets}(X)$, where $\text{Subsets}(X)$ is the category of subsets of X , with morphisms given by inclusions of subsets.

That is, for each $p \in P$, we have a subset $\mathcal{F}(p)$ of X , and for $p_1 \leq p_2$, the transition map $\psi_{\mathcal{F}}(p_1, p_2) : \mathcal{F}(p_1) \rightarrow \mathcal{F}(p_2)$ is the inclusion of $\mathcal{F}(p_1)$ into $\mathcal{F}(p_2)$.

A topological space X together with a filtration \mathcal{F} is called a **filtered space**, which we denote (X, \mathcal{F}) .

Example 1.7. Let P be the three-element poset $\{1, 2, 3\}$ with partial order inherited from \mathbb{N} . Let X be the simplicial complex on five vertices as shown in Figure ??.




Figure 10: The simplicial complex, X .

Let $\mathcal{F} : P \rightarrow \text{Subsets}(X)$ be the filtration as shown in Figure ??.




Figure 11: A filtration of the simplicial complex X by P .

In Figure ??, we see that the relations $1 \leq 2$ and $2 \leq 3$ correspond to inclusions $X(1) \hookrightarrow X(2)$ and $X(2) \hookrightarrow X(3)$, and these inclusions are transitive.

The following are three of the most common examples of filtrations used in persistent homology which can be constructed from topological spaces.

Example 1.8. Let X be a topological space and $f : X \rightarrow \mathbb{R}^n$ a function. The **sublevel set filtration** of X by f is the functor $\mathcal{S} : (\mathbb{R}^n, \leq) \rightarrow \text{Subsets}(X)$ given by,

$$\mathcal{S}(a) = \{x \in X \mid f(x) \leq a\}, \quad (1.7)$$

where the transition maps $\psi_{\mathcal{S}}(a, b) : \mathcal{S}(a) \rightarrow \mathcal{S}(b)$ are the inclusions of sublevel sets,

$$\{x \in X \mid f(x) \leq a\} \rightarrow \{x \in X \mid f(x) \leq b\}.$$

Example 1.9. The diagram in Figure ?? shows a 2-sphere, X , filtered by a height filtration. To describe this height filtration, we take let $f : X \rightarrow \mathbb{R}$ to be a function where for $x \in X$, $f(x)$ is the distance from x to a fixed horizontal plane beneath X . The subspace $X(h_i)$ is then given by the sublevel set $X(h_i) = \{x \in X \mid f(x) \leq h_i\}$, or, in other words, $X(h_i)$ consists of all points of X whose vertical height above the chosen plane is no more than h_i .




Figure 12: A 2-sphere filtered by a height parameter, h .

For the following examples, let (X, d) be a finite metric space, and let \hat{X} denote the maximal simplicial complex on X . That is, the 0-simplices of \hat{X} correspond to the points of X , and \hat{X} contains an n -simplex $[x_0, \dots, x_n]$ for each n -tuple of points $\{x_0, \dots, x_n\}$ of X .

Example 1.10. A **Vietoris-Rips filtration** on (X, d) is a persistence module $VR : \mathbb{R}_{\geq 0} \rightarrow \text{Subsets}(\hat{X})$, where for $a \in \mathbb{R}_{\geq 0}$, $VR(a)$ is the simplicial complex with a simplex $[x_0, \dots, x_n]$ whenever $d(x_i, x_j) \leq 2a$ for each pair of points $x_i, x_j \in \{x_0, \dots, x_n\}$. When $a \leq b$, the set of n -tuples which are pairwise at distance at most a includes into the set of n -tuples which are pairwise at distance at most b , and so the transition maps $\psi_{VR}(a, b) : VR(a) \rightarrow VR(b)$ are the inclusions of simplicial complexes.

For any $a \in \mathbb{R}$, the simplicial complex $VR(a)$ is called a **Vietoris-Rips complex**.




Figure 13: An example of a Vietoris-Rips filtration on a set of points in \mathbb{R}^2 .

A related filtration is the Čech filtration.




Figure 15: A Vietoris-Rips complex, $VR(a)$ (left), and a Čech complex, $C(a)$ (right), constructed on X the set of vertices of an equilateral triangle, where a is half the length of a side of the triangle.

Example 1.11. A Čech filtration on (X, d) is a persistence module $C : \mathbb{R}_{\geq 0} \rightarrow \text{Subsets}(\hat{X})$, where for $a \in \mathbb{R}_{\geq 0}$, $C(a)$ is the simplicial complex with a simplex $[x_0, \dots, x_n]$ whenever $\bigcap_{i=1}^n B_a(x_i) \neq \emptyset$, where $B_a(x)$ denotes the closed ball of radius a centred at $x \in X$. As for the Vietoris-Rips filtration, if $a \leq b$, the set of n -tuples which are mutually at distance at most a includes into the set of n -tuples which are mutually at distance at most b , and so the transition maps $\psi_C(a, b) : C(a) \rightarrow C(b)$ are the inclusions of simplicial complexes.

For any $a \in \mathbb{R}$, the simplicial complex $C(a)$ is called a **Čech complex**.

The diagram in Figure ?? shows a Čech filtration on the same finite metric space (X, d) as in Figure ??. Notice that a 2-simplex on vertices x_i, x_j, x_k only appears in the subcomplex $C(a)$ when the balls $B_a(x_i), B_a(x_j)$ and $B_a(x_k)$ mutually intersect, whereas the 2-simplex $[x_i, x_j, x_k]$ would appear in $VR(a)$ when $B_a(x_i), B_a(x_j)$ and $B_a(x_k)$ pairwise intersect.




Figure 14: An example of a Čech filtration on a set of points in \mathbb{R}^2 .

There is a clear relationship between the Čech and Vietoris-Rips filtration on a given finite metric space. However, they are not equal. For a given $a \in \mathbb{R}$, the Čech complex $C(a)$ includes into $VR(a)$, as for a for an n -tuple of points, $x_1, \dots, x_n \in X$, if $\bigcap_{i=1}^n B_a(x_i) \neq \emptyset$, then we also have $d(x_i, x_j) \leq 2a$ for all pairs $i, j \in \{1, \dots, n\}$.

The example in Figure ??, however, shows that the converse is not true in general.

On the other hand, for a given $a \in \mathbb{R}$, we have an inclusion $VR_a(\mathbb{X}) \subseteq C_{\sqrt{2}a}(\mathbb{X})$. A proof of this result may be found in (?).

1.2.2 Persistent Homology Groups

Definition 1.12. Let X be a topological space and $\mathcal{F} : P \rightarrow \text{Subsets}(X)$ a filtration of X . The n^{th} **persistent homology of X with respect to \mathcal{F}** is the functor $\mathcal{PH} : P \rightarrow \text{Vect}_k$ given by,

$$\mathcal{PH}(p) = H_n(\mathcal{F}(p); k), \quad (1.8)$$

where the transition maps $\psi_{\mathcal{PH}(p_1, p_2)} : \mathcal{PH}(p_1) \rightarrow \mathcal{PH}(p_2)$ are given by the linear maps in homology induced by the inclusions of subsets $\mathcal{F}(p_1) \hookrightarrow \mathcal{F}(p_2)$.

If X is triangulable, and \mathcal{F} is consistent with the triangulation, then we may take H_i to be the simplicial homology functor.

Definition 1.13. Let γ be a generator of $H_n(\mathcal{F}(p_i))$.

1. If γ is such that $f(\gamma) \notin \text{im}\{f : H_n(\mathcal{F}(p)) \rightarrow H_n(\mathcal{F}(p_i))\}$ for any $p < p_i$, then γ is said to be **born** at p_i .
2. For $p_j \geq p_i$, γ is said to **persist** from p_i to p_j if $f(\gamma) \in \text{im}\{f : H_n(\mathcal{F}(p_i)) \rightarrow H_n(\mathcal{F}(p_j))\}$.
3. Finally, γ is said to **die** at p_k if $\gamma \in \text{im}\{H_n(\mathcal{F}(p_i)) \rightarrow H_n(\mathcal{F}(p))\}$ for every $p_i \leq p < p_k$, but $\gamma \notin \text{im}\{H_n(\mathcal{F}(p_i)) \rightarrow H_n(\mathcal{F}(p_k))\}$.

Definition 1.14. Let $\mathcal{F} : P \rightarrow \text{Top}$ be a filtration of a topological space X . The **persistent homology groups of X with respect to \mathcal{F}** are given by,

$$H_n^{i,j}(X) = \text{im}\{\mathcal{PH}(p_i) \rightarrow \mathcal{PH}(p_j)\} = \text{im}\{H_n(\mathcal{F}(p_i)) \rightarrow H_n(\mathcal{F}(p_j))\}. \quad (1.9)$$

That is, the persistent homology groups, $H_n^{i,j}(X)$ of X are given by the images of the maps in homology which are induced by the inclusions of subspaces, $\mathcal{F}(p_i) \hookrightarrow \mathcal{F}(p_j)$.

The persistent homology group, $H_n^{i,j}(X)$, is generated by non-bounding n -cycles in X which persist from $\mathcal{F}(p_i)$ to $\mathcal{F}(p_j)$.

Definition 1.15. The **persistent Betti numbers** of X with respect to a filtration \mathcal{F} are the ranks of the persistent homology groups as k -vector spaces. That is,

$$\beta_n^{i,j}(X) = \text{rank}\{H_n^{i,j}(X)\} = \text{rank}(\text{im}\{H_n(\mathcal{F}(p_i)) \rightarrow H_n(\mathcal{F}(p_j))\}). \quad (1.10)$$

Example 1.16. Consider the following filtration in Figure ?? indexed by the set $\{1, 2, 3, 4, 5\}$. At stage 1, there are three 0-dimensional cycles, or connected components. One of these dies at stage 2, when it merges with another component. Another dies at stage 3, when it merges with the other component. Hence $\beta_0^{1,4}(X) = 1$.




Figure 16: A filtration of a 2-simplex, X .

For each $n \geq 0$, we can form the persistence module $\mathcal{PH}_n(X)$ by composing the filtration in Figure ?? with a simplicial homology functor. For $n = 0$, the persistence module is of the form,

$$k \oplus k \oplus k \longrightarrow k \oplus k \longrightarrow k \longrightarrow k \longrightarrow k, \quad (1.11)$$

where the inclusions of subcomplexes in Figure ?? induce the transition maps in (??). Similarly, the persistence module $\mathcal{PH}_1(X)$ is of the form,

$$0 \longrightarrow 0 \longrightarrow 0 \longrightarrow k \longrightarrow 0. \quad (1.12)$$

1.3 Persistent Homology as a Homology Theory

1.3.1 Homology Theories

Simplicial homology is one example of a homology theory – it is a specific method of calculating homology groups of a certain kind of topological space. There are numerous other homology theories, such as singular, or Čech homology, which use a slightly different definition to assign a sequence of homology groups, to a topological space. There are certain properties, though, which are common to all homology theories. One which we have already discussed is that homology is functorial. That is, for any $n \in \mathbb{Z}$, given two topological spaces X and Y , and a continuous map, $f : X \rightarrow Y$, there is an induced map $f_* : H_n(X) \rightarrow H_n(Y)$.

The other properties which define an ordinary homology theory are as follows. For any $n \in \mathbb{Z}$,

1. If $f, g : X \rightarrow Y$ are homotopic, then the induced maps,

$$f_*, g_* : H_n(X) \rightarrow H_n(Y),$$

are isomorphic.

2. For each X and each $A \subset X$, there is a homomorphism, $\partial : H_n(X, A) \rightarrow H_n(A)$ such that the sequence of homology groups,

$$\cdots \rightarrow H_p(A) \xrightarrow{i_*} H_p(X) \xrightarrow{j_*} H_p(X, A) \xrightarrow{\partial} H_{p-1}(A) \rightarrow \cdots \quad (1.13)$$

is an exact sequence, where $i : A \rightarrow X$ is the inclusion of the subspace A into X , and $j : X \rightarrow (X, A)$ is the inclusion of the pair (X, \emptyset) into the pair (X, A) . This sequence is called the **long exact sequence associated to the pair (X, A)** .

3. For each X , and each pair of subspaces A, B such that the interiors of A and B cover X , the inclusion of pairs, $(B, A \cap B) \rightarrow (X, A)$, induces an isomorphism in homology,

$$H_n(X, A) \cong H_n(B, A \cap B).$$

4. If X is a one-point space, then

$$H_n(X) = \begin{cases} k & n = 0 \\ 0 & \text{otherwise,} \end{cases} \quad (1.14)$$

where k is the field of coefficients of the homology theory.

Persistent homology assigns a sequence of *persistent* homology groups to a topological space. As we see, to be correctly referred to as a homology theory, there are several properties which persistence would be required to satisfy. Various works (??) (??) (??) have explored the extent to which persistence can be said to exhibit some of these properties, including the functoriality property. What emerges is that persistent homology satisfies many of these properties *to a certain extent*, but not always in the same way that a homology theory such as simplicial homology does. For example, in (??), it was shown that we have a long exact sequence of relative persistence modules $\mathcal{PH}(X, A)$ for a pair, but not of persistent homology groups, $H_n^{i,j}(X, A)$. A similar statement is true for a Mayer-Vietoris sequence in persistent homology. We will discuss these results in more detail in section ??.

In section ??, we show that an excision property holds for persistent homology, both for persistence modules and, in contrast to the properties of a long exact sequence (??) and a Mayer-Vietoris sequence (??), also for persistent homology groups. Excision can be used to calculate unknown relative homology groups by relating a pair to one for which the group is known. With this result, we can do the same for persistent homology. Given the use of persistent homology in topological data analysis, it is a very useful tool.

1.3.2 Excision

Definition 1.17. Let A be a subspace of a topological space X . A **filtration of the pair** (X, A) is a functor from a poset, P , the category of pairs of subspaces, $(X_i, A_i) \subseteq (X, A)$.

pair of filtrations, $\mathcal{F}_A : P \rightarrow \text{Subsets}(A)$, $\mathcal{F}_X : P \rightarrow \text{Subsets}(X)$ such that for all $p \in P$, $\mathcal{F}_A(p) \subseteq \mathcal{F}_X(p)$.

Example 1.18. Let X be a sphere and let $A \subset X$ be the lower hemisphere as shown in Figure ??.

Let $\mathcal{F} : \mathbb{R} \rightarrow \text{Subsets}(X)$ be a sublevel set filtration of X by some function $f : X \rightarrow \mathbb{R}$. Define a filtration \mathcal{F}_A of A by setting,

$$\mathcal{F}_A(i) = \mathcal{F}(i) \cap A.$$




Figure 17: A sphere, X , and a subspace, A , of X .

Then $(\mathcal{F}, \mathcal{F}_A)$ is a filtration of the pair (X, A) .

The relative persistent homology of a pair (X, A) with respect to filtrations \mathcal{F}_X and \mathcal{F}_A is the persistence module $\mathcal{PH}(X, A)$, where for $p \in P$,

$$\mathcal{PH}(X, A)(p) = H_n(\mathcal{F}_X(p), \mathcal{F}_A(p)). \quad (1.15)$$

Given $\mathcal{PH}(X, A)$, and $p_i \leq p_j$ the relative persistent homology groups of (X, A) with respect to the filtrations \mathcal{F}_X and \mathcal{F}_A are given by,

$$\begin{aligned} H_n^{i,j}(X, A) &= \text{im}\{\psi_{\mathcal{PH}} : \mathcal{PH}(X, A)(p_i) \rightarrow \mathcal{PH}(X, A)(p_j)\} \\ &= \text{im}\{\psi_{\mathcal{PH}} : H_n(\mathcal{F}_X(p_i), \mathcal{F}_A(p_i)) \rightarrow H_n(\mathcal{F}_X(p_j), \mathcal{F}_A(p_j))\}. \end{aligned} \quad (1.16)$$

Definition 1.19. Let A be a subspace of X , and let $\mathcal{F}_X : P \rightarrow \text{Subsets}(X)$ be a filtration of X . A filtration $\mathcal{F}_A : P \rightarrow \text{Subsets}(A)$ of A is **induced** by \mathcal{F}_X if for all $p \in P$,

$$\mathcal{F}_A(p) = \mathcal{F}_X(p) \cap A. \quad (1.17)$$

Example 1.20. Let X be a topological space with subspace $A \subset X$. Given a function $f_X : X \rightarrow \mathbb{R}$, let $\mathcal{S} : \mathbb{R} \rightarrow \text{Subsets}(X)$ be the sublevel set filtration of X by f_X . Let f_A denote the restriction of f to A , and let $\mathcal{S}_A : \mathbb{R} \rightarrow \text{Subsets}(A)$ be the sublevel set filtration of A by f_A . That is, for $p \in \mathbb{R}$,

$$\mathcal{S}_A(p) = \{a \in A \mid f_A(a) \leq p\} = \{a \in X \mid f_X(a) \leq p\} \cap A = \mathcal{S}(p) \cap A. \quad (1.18)$$

Hence the restriction of a sublevel set filtration to a subspace is an example of an induced filtration.

Definition 1.21. Let X be a topological space and A a subset of X . The **interior** of A inside X , denoted A_X° , is the largest subset of A which is open in the topology on X .

Lemma 1.22. *Suppose that A and B are subspaces of X such that $A^\circ \cup B^\circ = X$. Let $\mathcal{F}_X : P \rightarrow \text{Subsets}(X)$ be a filtration of X , and let $\mathcal{F}_A : P \rightarrow \text{Subsets}(A)$ and $\mathcal{F}_B : P \rightarrow \text{Subsets}(B)$ be filtrations on A and B respectively which are induced by \mathcal{F}_X . Then for all $p \in P$,*

$$\mathcal{F}_X(p) = (\mathcal{F}_A(p))^\circ \cup (\mathcal{F}_B(p))^\circ. \quad (1.19)$$

We will prove a slightly more general statement, that if U is a subspace of a topological space X , which is covered by the interiors of two subspaces A and B , then restricting A and B to U provides a cover of U . That is, if $X = A_X^\circ \cup B_X^\circ$, then

$$U \subseteq (U \cap A)_U^\circ \cup (U \cap B)_U^\circ. \quad (1.20)$$

Proof. By assumption, $X = A_X^\circ \cup B_X^\circ$.

Trivially, $U = U \cap X$, as $U \subset X$, so

$$U = U \cap X = U \cap (A_X^\circ \cup B_X^\circ) = (U \cap A_X^\circ) \cup (U \cap B_X^\circ). \quad (1.21)$$

Considering $(U \cap A_X^\circ)$, this set is open in the subspace topology for U , as by definition it is the intersection of U with an open set in X . As it is open in U , it equals its own interior, when the interior is taken inside U , so

$$(U \cap A_X^\circ) = (U \cap A_X^\circ)_U^\circ \subseteq (U \cap A)_U^\circ, \quad (1.22)$$

similarly for B , we have:

$$(U \cap B_X^\circ) = (U \cap B_X^\circ)_U^\circ \subseteq (U \cap B)_U^\circ. \quad (1.23)$$

Hence $U = (U \cap A_X^\circ) \cup (U \cap B_X^\circ) \subseteq (U \cap A)_U^\circ \cup (U \cap B)_U^\circ$. \square

Example 1.23. Let A and B be subsets of a topological space X such that $A \cap B \neq \emptyset$.

Let $f : X \rightarrow \mathbb{R}$ be a function which measures the height of a point $x \in X$ above some fixed horizontal plane, P , and let $\mathcal{F} : \mathbb{R} \rightarrow \text{Subsets}(X)$ be the filtration where for all $a \in \mathbb{R}$,

$$\mathcal{F}(a) = \{x \in X \mid f(x) \leq a\}. \quad (1.24)$$

For $a \leq b$, the transition maps $\psi_{\mathcal{F}}(a, b) : \mathcal{F}(a) \rightarrow \mathcal{F}(b)$ are then the inclusions of subsets. For any $a \in \mathbb{R}$, $\mathcal{F}(a)$ consists of all points of X which are of height at most a from P , and so, for any $a \in \mathbb{R}$, $\mathcal{F}_A(a) = \mathcal{F}|_A(a) = \mathcal{F}(a) \cap A$, and similarly for $\mathcal{F}_B = \mathcal{F}|_B(a)$. By Lemma ??, $a \in \mathbb{R}$,

$$\mathcal{F}(a) = \mathcal{F}_A(a)^\circ \cup \mathcal{F}_B(a)^\circ.$$

This example, with X a 2-sphere, is illustrated in Figure ??.




Figure 18: A 2-sphere, X , covered by subspaces A and B .

We now introduce our excision theorem for persistent homology.

Theorem 1.24. *Let A and B be subspaces of the topological space X such that $X = A^\circ \cup B^\circ$. Let $\mathcal{F}_X : P \rightarrow \text{Subsets}(X)$ be a filtration of X and let $\mathcal{F}_A : P \rightarrow \text{Subsets}(A)$ and $\mathcal{F}_B : P \rightarrow \text{Subsets}(B)$ be filtrations on A and B respectively such that for all $p \in P$,*

$$\mathcal{F}_X(p) = (\mathcal{F}_A(p))^\circ \cup (\mathcal{F}_B(p))^\circ. \quad (1.25)$$

Let $\mathcal{PH}_n(X, A)$ be the relative persistent homology of (X, A) in degree n with respect to the filtrations \mathcal{F}_X and \mathcal{F}_A , and let $\mathcal{PH}_n(B, A \cap B)$ be the relative persistent homology of $(B, A \cap B)$ in degree n relative to the filtrations \mathcal{F}_B and $\mathcal{F}_{A \cap B}$, where for $p \in P$,

$$\mathcal{F}_{A \cap B}(p) = \mathcal{F}_A(p) \cap \mathcal{F}_B(p). \quad (1.26)$$

Then for every $n \geq 0$, there is an isomorphism of persistence modules,

$$\mathcal{PH}_n(X, A) \cong \mathcal{PH}_n(B, A \cap B), \quad (1.27)$$

and for every $p_i \leq p_j$, there is an isomorphism of persistent homology groups,

$$H_n^{i,j}(X, A) \cong H_n^{i,j}(B, A \cap B). \quad (1.28)$$

Proof. For any $p_i \leq p_j$ in P , consider the diagram of inclusions,

$$\begin{array}{ccc} (\mathcal{F}_B(p_i), \mathcal{F}_{A \cap B}(p_i)) & \longrightarrow & (\mathcal{F}_B(p_j), \mathcal{F}_{A \cap B}(p_j)) \\ \downarrow & & \downarrow \\ (\mathcal{F}_X(p_i), \mathcal{F}_A(p_i)) & \longrightarrow & (\mathcal{F}_X(p_j), \mathcal{F}_A(p_j)) \end{array} \quad (1.29)$$

This diagram commutes, and so the diagram of induced maps,

$$\begin{array}{ccc}
\mathcal{PH}(B, A \cap B)(p_i) & \longrightarrow & \mathcal{PH}(B, A \cap B)(p_j) \\
\downarrow & & \downarrow \\
\mathcal{PH}(X, A)(p_i) & \longrightarrow & \mathcal{PH}(X, A)(p_j)
\end{array} \tag{1.30}$$

Additionally, the condition (??) means that each vertical map in (??) is an isomorphism, by the excision property for H_n . Since these vertical isomorphisms commute with the horizontal maps, we have an isomorphism of persistence modules,

$$\mathcal{PH}(B, A \cap B) \rightarrow \mathcal{PH}(X, A). \tag{1.31}$$

Note that for any $p_i \leq p_j$ in P , $\text{im}\{\mathcal{PH}(X, A)(p_i) \rightarrow \mathcal{PH}(X, A)(p_j)\}$ is the relative persistent homology group $H_n^{i,j}(X, A)$, and similarly,

$$H_n^{i,j}(B, A \cap B) = \text{im}\{\mathcal{PH}(B, A \cap B)(p_i) \rightarrow \mathcal{PH}(B, A \cap B)(p_j)\}. \tag{1.32}$$

We now show that the vertical isomorphisms in (??) restrict to isomorphisms on the images of the horizontal maps. For this we will consider the general picture, where we have a commutative diagram of vector spaces C, D, E and F ,

$$\begin{array}{ccc}
C & \xrightarrow{f} & D \\
i \downarrow & & \downarrow j \\
E & \xrightarrow{g} & F
\end{array} \tag{1.33}$$

Let i and j be isomorphisms. Then we can show that j restricts to an isomorphism $\bar{j} : \text{Im}(f) \rightarrow \text{Im}(g)$.

The proof involves a simple diagram chase. Let $y \in \text{Im}(f)$. Then there is some $c \in C$ such that $y = f(c)$. Define a map $\bar{j} : \text{Im}(f) \rightarrow \text{Im}(g)$ by $\bar{j}(y) = g(i(c))$.

First suppose $y \in \text{Im}(f)$ is such that $y = f(c)$ and $y = f(d)$ for $c, d \in C$. Then as (??) commutes, $g(i(c)) = j(f(c)) = j(y)$, and $g(i(d)) = j(f(d)) = j(y)$, and so $g(i(c)) = g(i(d))$, and \bar{j} is well-defined.

We now show that \bar{j} is both injective and surjective.

1. **Injectivity:** Suppose $\bar{j}(y) = \bar{j}(y')$ for $y, y' \in \text{Im}(f)$. Let $y = f(c)$ and $y' = f(c')$. Then by definition of \bar{j} ,

$$g(i(c)) = g(i(c')).$$

The square (??) commutes, so

$$j(f(c)) = j(f(c')).$$

But j is an isomorphism, and so $y = f(c) = f(c') = y'$.

2. **Surjectivity:** Let $x \in \text{Im}(g)$. We want to find a $y \in \text{Im}(f)$ such that $\bar{j}(y) = x$.

First, we have that as $x \in \text{Im}(g)$, there is some $z \in E$ such that $g(z) = x$. We also have that i is an isomorphism, so there is an $c \in C$ such that $i^{-1}(z) = c$. Let $y = f(c)$. Then $s(y) = x$. Hence we have an isomorphism $\text{Im}(f) \rightarrow \text{Im}(g)$.

Replacing the vector spaces C, D, E and F with $\mathcal{PH}(B, A \cap B)(p_i)$, $\mathcal{PH}(B, A \cap B)(p_j)$, $\mathcal{PH}(X, A)(p_i)$ and $\mathcal{PH}(X, A)(p_j)$, respectively, we see that the isomorphism,

$$\mathcal{PH}(B, A \cap B)(p_j) \rightarrow \mathcal{PH}(X, A)(p_j),$$

descends to an isomorphism on the images of the maps

$$\mathcal{PH}(B, A \cap B)(p_i) \rightarrow \mathcal{PH}(B, A \cap B)(p_j), \quad (1.34)$$

$$\mathcal{PH}(X, A)(p_i) \rightarrow \mathcal{PH}(X, A)(p_j). \quad (1.35)$$

In this case, these images are the persistent homology groups, and so we have an isomorphism,

$$H_n^{i,j}(X, A) \cong H_n^{i,j}(B, A \cap B), \quad (1.36)$$

for any $p_i \leq p_j$ and for any $n \geq 0$.

□

We note that Lemma ?? demonstrates that the conditions of Theorem ?? hold for the triple of filtered spaces (X, A, B) , whenever A and B are endowed with filtrations which are induced by that of X . This does not impose any restrictions on the filtration on X . Hence, Lemma ?? demonstrates a way in which any filtration on X can be extended to a pair of subspaces in such a way that the persistent excision theorem is guaranteed to hold on that triple.

1.3.3 Mayer-Vietoris Sequences in Persistent Homology

Another extremely useful property possessed by a homology theory is the existence of a long exact sequence known as the Mayer-Vietoris sequence. The existence of this sequence can be derived from the axioms of a homology theory (?) – namely, from the axiom stating the existence of long exact sequence for a pair, and the excision axiom.

Definition 1.25. Let X be a space with subspaces A and B such that $X = A_X^\circ \cup B_X^\circ$. The **Mayer-Vietoris sequence** associated to the triple (X, A, B) is the long exact sequence,

$$\cdots \rightarrow H_{n+1}(X) \xrightarrow{\partial_n} H_n(A \cap B) \xrightarrow{(\alpha_n, \beta_n)} H_n(A) \oplus H_n(B) \xrightarrow{\gamma_n} H_n(X) \rightarrow \cdots \quad (1.37)$$

where the maps α_n and β_n are induced by the respective inclusions,

$$(A \cap B) \hookrightarrow A, \text{ and } (A \cap B) \hookrightarrow B, \quad (1.38)$$

so that,

$$(\alpha_n, \beta_n)([z]) = ([z], [z]), \quad (1.39)$$

and $\gamma_n([z], [z']) = [z - z']$ is induced by the inclusions,

$$A \hookrightarrow X \text{ and } B \hookrightarrow X. \quad (1.40)$$

The map ∂_n is the usual boundary map – as $X = A_X^\circ \cup B_X^\circ$, every chain in X can be expressed as the sum of a chain in A and a chain in B , whose boundary lies in the intersection, $A \cap B$.

In (?), the authors derive a form of a Mayer-Vietoris sequence in persistent homology, in the special case that the persistence modules are the composition of a sublevel set filtration functor and a Čech homology functor. We will see that, unlike the sequence in (??), this Mayer-Vietoris-style sequence in persistent homology groups fails to be exact. Despite this, the authors show that the sequence obtained can be very useful when persistent homology is used in applications.

In this section, we derive the Mayer-Vietoris-style sequence of (?) in a new way, using the persistent excision theorem and the work of (?). In this paper (?), the authors construct a sequence for a pair in the style of a long exact sequence in homology. We note that, just like the Mayer-Vietoris-style sequence in persistent homology of (?), this sequence is not a long *exact* sequence of persistent homology groups, but is a chain complex. Our new derivation of the sequence of (?) shows that a Mayer-Vietoris sequence in persistent homology exists not only for sublevel set filtrations, as was initially proved, but for any triple of spaces which satisfies the conditions of Theorem ??.

The following Theorem shows the result of attempting to restrict a long exact sequence in homology to persistent homology groups.

Theorem 1.26. *Let $(\mathcal{F}_X, \mathcal{F}_A)$ be a filtration of the pair (X, A) . For any $n \geq 0$, the sequence of persistent homology groups,*

$$\cdots \rightarrow H_{n+1}^{i,j}(X, A) \xrightarrow{\partial_n^{i,j}} H_n^{i,j}(A) \xrightarrow{\iota_n^{i,j}} H_n^{i,j}(X) \xrightarrow{\kappa_n^{i,j}} H_n^{i,j}(X, A) \rightarrow \cdots \quad (1.41)$$

is a chain complex (?). We detail the maps in the course of the proof.

We give a brief overview of the proof, which can be seen in greater detail in (?). We note that it follows a similar structure to the proof of the existence of a Mayer-Vietoris-style sequence in persistent homology which appears in (?).

Proof. In the course of this proof and the next, we will use the shorthand $X_i := \mathcal{F}_X(p_i)$ and $A_i := \mathcal{F}_A(p_i)$.

For each $p_i \in P$, we have a pair of subspaces (X_i, A_i) , and for each such pair we have a long exact sequence in homology,

$$\cdots \rightarrow H_{n+1}(X_i, A_i) \xrightarrow{\partial_n^i} H_n(A_i) \xrightarrow{\iota_n^i} H_n(X_i) \xrightarrow{\kappa_n^i} H_n(X_i, A_i) \rightarrow \cdots \quad (1.42)$$

Let $p_i \leq p_j$. Then we have an identical sequence for the pair (X_j, A) , and we can form the commutative diagram,

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & H_{n+1}(X_i, A_i) & \xrightarrow{\partial_n^i} & H_n(A_i) & \xrightarrow{\iota_n^i} & H_n(X_i) & \xrightarrow{\kappa_n^i} & H_n(X_i, A_i) & \longrightarrow & \cdots \\
& & \downarrow h_{n+1} & & \downarrow f_n & & \downarrow g_n & & \downarrow h_n & & \\
\cdots & \longrightarrow & H_{n+1}(X_j, A_j) & \xrightarrow{\partial_n^j} & H_n(A_j) & \xrightarrow{\iota_n^j} & H_n(X_j) & \xrightarrow{\kappa_n^j} & H_n(X_j, A_j) & \longrightarrow & \cdots
\end{array} \tag{1.43}$$

where the horizontal sequences are the long exact sequences of the pairs as in (??), and the vertical maps are induced by the inclusions $X_i \hookrightarrow X_j$ and $A_i \hookrightarrow A_j$.

In general, given a sequence A_* of abelian groups, and a chain complex, B_* , and a map of sequences $f_* : A_* \rightarrow B_*$ such that each square,

$$\begin{array}{ccc}
A_i & \longrightarrow & A_j \\
f_i \downarrow & & \downarrow f_j \\
B_i & \longrightarrow & B_j
\end{array} \tag{1.44}$$

commutes, then the sequence of images, $\text{im}(f_*)$ is also a chain complex. The proof of this statement is omitted here, but can be shown via a simple diagram chase. In our case, we consider the restriction of the lower sequence of (??) to the images of the vertical maps. That is, we let

$$\partial_n^{i,j} = \partial_n^j|_{\text{im}(h_{n+1})}, \quad \iota_n^{i,j} = \iota_n^j|_{\text{im}(f_n)}, \quad \kappa_n^{i,j} = \kappa_n^j|_{\text{im}(g_n)}.$$

The inclusions $\text{im}(\partial_n^j) \subseteq \ker(\iota_n^j)$, $\text{im}(\iota_n^j) \subseteq \ker(\kappa_n^j)$ and $\text{im}(\kappa_n^j) \subseteq \ker(\partial_{n-1}^j)$ descend to inclusions,

$$\text{im}(\partial_n^{i,j}) \subseteq \ker(\iota_n^{i,j}), \quad \text{im}(\iota_n^{i,j}) \subseteq \ker(\kappa_n^{i,j}) \quad \text{and} \quad \text{im}(\kappa_n^{i,j}) \subseteq \ker(\partial_n^{i,j}),$$

and so the sequence of persistent homology groups,

$$\cdots \longrightarrow H_{n+1}^{i,j}(X, A) \xrightarrow{\partial_n^{i,j}} H_n^{i,j}(A) \xrightarrow{\iota_n^{i,j}} H_n^{i,j}(X) \xrightarrow{\kappa_n^{i,j}} H_n^{i,j}(X, A) \longrightarrow \cdots \tag{1.45}$$

is a chain complex. \square

The following result appeared in a less general form in (?), with a proof following a similar structure to that of Theorem ??. The result below is a generalisation of that of (?), and the proof follows a very different structure.

Theorem 1.27. *Let (X, A, B) be a triple with respective filtrations $\mathcal{F}_X : P \rightarrow \text{Subsets}(X)$, $\mathcal{F}_A : P \rightarrow \text{Subsets}(A)$ and $\mathcal{F}_B : P \rightarrow \text{Subsets}(B)$, such that for each $p \in P$,*

$$\mathcal{F}_X(p) = (\mathcal{F}_A(p))^\circ \cup (\mathcal{F}_B(p))^\circ. \tag{1.46}$$

Then for any $p_i \leq p_j$, the sequence of persistent homology groups,

$$\cdots \rightarrow H_{n+1}^{i,j}(X) \xrightarrow{\partial_n^{i,j}} H_n^{i,j}(A \cap B) \xrightarrow{(\alpha_n^{i,j}, \beta_n^{i,j})} H_n^{i,j}(A) \oplus H_n^{i,j}(B) \xrightarrow{\gamma_n^{i,j}} H_n^{i,j}(X) \rightarrow \cdots \quad (1.47)$$

is a chain complex. We detail the maps in the course of the proof.

Proof. In the course of the proof, we will use the shorthand $X_i = \mathcal{F}_X(p_i)$, $A_i = \mathcal{F}_A(p_i)$, $B_i = \mathcal{F}_B(p_i)$ and $(A \cap B)_i = \mathcal{F}_A(p_i) \cap \mathcal{F}_B(p_i) = \mathcal{F}_{A \cap B}(p_i)$.

For $p_i \in P$, consider the diagram,

$$\begin{array}{ccccccc} \cdots & \rightarrow & H_n(B_i) & \longrightarrow & H_n(B_i, (A \cap B)_i) & \xrightarrow{\partial} & H_{n-1}((A \cap B)_i) & \longrightarrow & H_{n-1}(B_i) & \rightarrow & \cdots \\ & & \downarrow & & \downarrow f_i & & \downarrow g_i & & \downarrow & & \\ \cdots & \rightarrow & H_n(X_i) & \longrightarrow & H_n(X_i, A_i) & \xrightarrow{\partial} & H_{n-1}(A_i) & \longrightarrow & H_{n-1}(X_i) & \rightarrow & \cdots \end{array} \quad (1.48)$$

Each horizontal row in this diagram is a long exact sequence, for the pairs $(B_i, (A \cap B)_i)$, and (X_i, A_i) . Hence each map, other than the boundary maps indicated, is induced by an inclusion. Moreover, each square in this diagram commutes, either because all four maps in the square are induced by an inclusion, or because the boundary map commutes with inclusions.

Let $p_i \leq p_j$, and consider the diagram,

$$\begin{array}{cccccccccccc} \cdots & \rightarrow & H_n(B_i) & \longrightarrow & H_n(B_i, (A \cap B)_i) & \xrightarrow{\partial} & H_{n-1}((A \cap B)_i) & \longrightarrow & H_{n-1}(B_i) & \rightarrow & \cdots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \cdots & \rightarrow & H_n(B_j) & \longrightarrow & H_n(B_j, (A \cap B)_j) & \xrightarrow{\partial} & H_{n-1}((A \cap B)_j) & \longrightarrow & H_{n-1}(B_j) & \rightarrow & \cdots \\ & & \downarrow & & \downarrow f_i & & \downarrow g_i & & \downarrow & & \\ \cdots & \rightarrow & H_n(X_i) & \longrightarrow & H_n(X_i, A_i) & \xrightarrow{\partial} & H_{n-1}(A_i) & \longrightarrow & H_{n-1}(X_i) & \rightarrow & \cdots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \cdots & \rightarrow & H_n(X_j) & \longrightarrow & H_n(X_j, A_j) & \xrightarrow{\partial} & H_{n-1}(A_j) & \longrightarrow & H_{n-1}(X_j) & \rightarrow & \cdots \end{array} \quad (1.49)$$

α_1 (red arrow from $H_n(B_i)$ to $H_n(B_j)$)
 α_2 (red arrow from $H_n(B_i, (A \cap B)_i)$ to $H_n(B_j, (A \cap B)_j)$)
 α_3 (red arrow from $H_n(X_i)$ to $H_n(X_j)$)
 α_4 (red arrow from $H_n(X_i, A_i)$ to $H_n(X_j, A_j)$)

The diagram is really the union of two diagrams of the form (??), the rear diagram for p_i , and the front diagram for p_j , with red connecting maps induced by the inclusions of subspaces,

$$A_i \hookrightarrow A_j, B_i \hookrightarrow B_j, \text{ and } X_i \hookrightarrow X_j.$$

If we restrict the front diagram to the images of the red maps, then we obtain the following diagram in persistent homology groups.

$$\begin{array}{ccccccccccc}
\cdots & \longrightarrow & H_n^{i,j}(B) & \xrightarrow{\kappa_B} & H_n^{i,j}(B, A \cap B) & \xrightarrow{\partial_B} & H_{n-1}^{i,j}(A \cap B) & \xrightarrow{\iota_B} & H_{n-1}^{i,j}(B) & \longrightarrow & \cdots \\
& & \downarrow & & \downarrow f & & \downarrow g & & \downarrow h & & \\
\cdots & \longrightarrow & H_n^{i,j}(X) & \xrightarrow{\kappa_X} & H_n^{i,j}(X, A) & \xrightarrow{\partial_X} & H_{n-1}^{i,j}(A) & \xrightarrow{\iota_A} & H_{n-1}^{i,j}(X) & \longrightarrow & \cdots
\end{array} \tag{1.50}$$

where the bottom row is the sequence in persistent homology for the pair (X, A) as defined in Theorem ??, the top row is the analogous sequence for the pair $(B, A \cap B)$, and the vertical maps are induced by the inclusions, $B \hookrightarrow X$, $A \cap B \hookrightarrow A$, and the inclusion of the pair, $(B, A \cap B) \hookrightarrow (X, A)$, which by Theorem ?? is an isomorphism. The left-hand and right-hand squares commute since they commute at the level of spaces — every map is induced from an inclusion and the inclusion maps themselves commute.

For the commutativity of the middle square of (??), we return to the diagram in (??). We will show that for any $\gamma_1 \in \text{Im}(\alpha_1)$, we have that,

$$\partial f_j(\gamma_1) = g_j \partial(\gamma_1) \in \text{Im}(\alpha_4).$$

We will first show that both $\partial f_j(\gamma_1)$ and $g_j \partial(\gamma_1)$ belong to $\text{Im}(\alpha_4)$, and then that the square commutes.

We first have that $f_j(\gamma_1) \in \text{Im}(\alpha_3)$ since $f_j(\gamma_1) = f_j(\alpha_1(x))$ for some $x \in H_n(B_i, (A \cap B)_i)$.

Since $f_j(\gamma_1) \in \text{Im}(\alpha_3)$, we have that $f_j(\gamma_1) = \alpha_3(z)$ for some $z \in H_n(X_i, A_i)$. Hence,

$$\partial f_j(\gamma_1) = \partial \alpha_3(z) = \alpha_4 \partial(z),$$

and so $\partial f_j(\gamma_1) \in \text{Im}(\alpha_4)$.

Next, $\partial(\gamma_1) = \partial \alpha_1(x) = \alpha_2 \partial(x)$ and so $\partial(\gamma_1) \in \text{Im}(\alpha_2)$.

As $\partial(\gamma_1) \in \text{Im}(\alpha_2)$, we have that $\partial(\gamma_1) = \alpha_2(y)$ for some $y \in H_{n-1}((A \cap B)_i)$, and so,

$$g_j \partial(\gamma_1) = g_j \alpha_2(y) = \alpha_4 g_j(y),$$

and so $g_j \partial(\gamma_1) \in \text{Im}(\alpha_4)$.

Since the square,

$$\begin{array}{ccc}
H_n(B_j, (A \cap B)_j) & \xrightarrow{\partial} & H_{n-1}((A \cap B)_j) \\
\downarrow f_j & & \downarrow g_j \\
H_n(X_j, A_j) & \xrightarrow{\partial} & H_{n-1}(A_j)
\end{array} \tag{1.51}$$

commutes, and given that both $\partial f_j(\text{Im}(\alpha_1)) \subseteq \text{Im}(\alpha_4)$ and $g_j \partial(\text{Im}(\alpha_1)) \subseteq \text{Im}(\alpha_4)$, we also have that the restriction of the square (??) to the persistent homology groups also commutes.

We'll now show that the condition that the upper and lower sequences in diagram (??) are chain complexes descends to a chain complex condition on the subsequence,

$$\cdots \rightarrow H_n^{i,j}(X) \xrightarrow{\partial_*} H_{n-1}^{i,j}(A \cap B) \xrightarrow{(g, \iota_B)} H_{n-1}^{i,j}(A) \oplus H_{n-1}^{i,j}(B) \xrightarrow{\gamma} H_{n-1}^{i,j}(X) \rightarrow \cdots \quad (1.52)$$

where $\gamma([z]) = h - \iota_A$. This proof is a standard way of proving the existence of the Mayer-Vietoris sequence for a homology theory from the excision and long exact sequences axioms for a homology theory in the non-persistent case, which can be found as an exercise without proof in (?).

- $\text{im}(\partial_*) \subseteq \ker(g, \iota_B)$:

Firstly, as ∂_* is the composition of maps,

$$\partial_* = \partial_B \circ f^{-1} \circ \kappa_X, \quad (1.53)$$

then any element of $\text{im}(\partial_*)$ is clearly also an element of $\text{im}(\partial_B)$. As $\text{im}(\partial_B) \subseteq \ker(\iota_B)$, we also have $\text{im}(\partial_*) \subseteq \ker(\iota_B)$. We also show that $\text{im}(\partial_*) \subseteq \ker(g)$.

Hence, using the fact that the composition of κ_X with ∂_X is zero, and that, given the claim, ∂_X is equal to the composition,

$$\partial_X = g \circ \partial_B \circ f^{-1}, \quad (1.54)$$

then we must also have that the composition $g \circ (\partial_B \circ f^{-1} \circ \kappa_X)$ is zero — that is, $\text{im}(\partial_*) \subseteq \ker(g)$.

- $\text{im}(g, \iota_B) \subseteq \ker(h - \iota_A)$:

Let $([Y], [Z]) \in \text{im}(g, \iota_B)$, so that there exists some $[X] \in H_{n-1}^{i,j}(A \cap B)$ such that $g([X]) = [Y]$ and $\iota_B([X]) = [Z]$. We have that the right-hand square in (??) commutes, so $h \circ \iota_B([X]) - \iota_A \circ g([X]) = 0$, and so $([Y], [Z]) = (g([X]), \iota_B([X])) \in \ker(h - \iota_A)$.

- $\text{im}(h - \iota_A) \subseteq \ker(\partial_*)$:

Let $[X] \in \text{im}(h - \iota_A)$. Then there exists $[X_1] \in H_{n-1}^{i,j}(B)$, and $[X_2] \in H_{n-1}^{i,j}(A)$ such that $h(X_1) - \iota_A(X_2) = X$. Then we have that $\partial_*(X) = \partial_*(h(X_1) - \iota_A(X_2)) = \partial_*(h(X_1)) - \partial_*(\iota_A(X_2))$. Given that $\partial_* = \partial_B \circ f^{-1} \circ \kappa_X$, and the composition $\kappa_X \circ \iota_A$ is zero, then we must have that $\partial_*(\iota_A(X_2)) = 0$. We also have that $f^{-1} \circ \kappa_X \circ h = \kappa_B$. Hence the composition,

$$\partial_* \circ h = \partial_B \circ f^{-1} \circ \kappa_X \circ h = \partial_B \circ \kappa_B. \quad (1.55)$$

We now use the fact that the top and bottom sequences in (??) are chain complexes, to conclude that $\partial_* \circ \beta(X_1) = \partial_B \circ \kappa_B(X_1) = 0$.

□

We have hence demonstrated that the Mayer-Vietoris sequence of (?) exists for any filtered triple satisfying the condition of (?). In particular, this sequence exists for multiparameter filtrations satisfying this condition. The existence of this sequence - despite the fact that it is not exact - was used in (?) to show that for such a filtered triple, the points of the persistence diagrams of the subspaces A and B appear as *signatures* in the persistence diagram of either X or that of the intersection, $A \cap B$.

We could ask when the sequence of persistent homology groups in (??) and (??) are exact, as opposed to merely chain complexes. We thank Ran Levi and Jan Spakula for the following observation, which could be used as a source of further work to investigate this question. As well as considering the images of the vertical maps in (??), we can consider the sequence of kernels, K_* , which is given by,

$$\cdots \rightarrow \ker(h_{n+1}) \rightarrow \ker(f_n) \rightarrow \ker(g_n) \rightarrow \ker(h_n) \rightarrow \cdots \quad (1.56)$$

with h_i, f_i and g_i as in (?). There is a left exact sequence of chain complexes,

$$0 \rightarrow K_* \rightarrow A_* \rightarrow B_*, \quad (1.57)$$

where A_* is the upper sequence in the diagram in (??) and B_* is the lower sequence in the diagram in (??). Then the quotient sequence, A_*/K_* is isomorphic to the sequence of images in (?). Hence (??) is exact precisely when the quotient sequence, A_*/K_* is exact. Given that (??) is used in the construction of the Mayer-Vietoris sequence in persistent homology (??), this would also be of use in determining when (??) is an exact sequence.

We have also shown that an excision theorem exists for persistent homology, both on the level of persistence modules and persistent homology groups. Lemma ?? also gives us conditions under which a filtered triple will satisfy the conditions of Theorem ?? and Theorem ?. In particular, given any filtered space X , we can define a way to restrict the filtration to a pair of subsets in such a way that the conditions of Theorems ?? and ?? are guaranteed to hold.

2 Interleavings

In the first section, we defined a persistence module and looked at some of the properties of persistent homology. In the remainder of this thesis, we look at relationships between persistence modules.

If we want to compare persistence modules, then we currently have the notions of morphisms or isomorphisms. But what does each tell us about the relationship between modules? On the one hand, an isomorphism is a very strict relationship, and is certainly too much to hope for in general. On the other hand, the existence of a morphism between two modules is a rather weak relationship – as an extreme example, there exists a morphism from any module to the zero module, but the existence of such a morphism doesn't tell us anything about the non-zero module. In this section we will define the concept of an interleaving between persistence modules, and show how we can use this relationship to define a form of distance between modules.

An interleaving was first defined in (?). Prior to this definition in (?), the standard way of comparing persistence processes was to compare the discrete invariants called barcodes. A barcode is a discrete invariant, which can be defined for persistence modules indexed by \mathbb{R} , but – crucially – not all modules (?). In (?) it was first suggested that we can define relationships between the algebraic objects – the persistence modules – as opposed to looking for similarities between barcodes, which was one of the important first steps to generalise persistence to more than the initial special cases for which it was first defined.

If we want to define a relationship between P -persistence modules, where P is an arbitrary poset, then the definition due to (?) is no longer sufficient. Interleavings of \mathbb{R} -modules were originally defined with respect to maps $\mathbb{R} \rightarrow \mathbb{R}$ of the form $a \mapsto a + \varepsilon$. For a general poset, a map of the form $a \mapsto a + \varepsilon$ may not be defined. Even when the above definition *can* be defined for our given poset, the relationship is rather rigid, and may not describe the relationship between the persistence modules in the most useful way. Consider for example the relationship between the Čech and Vietoris-Rips filtrations on a given finite metric space, X . As we discussed in section ??, for a given $a \in \mathbb{R}$, the Čech complex, $C_X(a)$, includes into the Vietoris-Rips complex, $VR_X(\sqrt{2}a)$. Hence, in this case, it could be more useful to describe the relationship between the persistence modules C_X and VR_X via a map $f : \mathbb{R} \rightarrow \mathbb{R}$ given by,

$$f(a) = \sqrt{2}a.$$

In this section, we will define an interleaving, and the associated metric on the space of P -persistence modules, in the most general way possible. The concept of an interleaving has been generalised to a greater range of poset morphisms before, in (?) and (?). There is one important distinction, however, between the definition of an interleaving *distance* we'll see in section ?? and the one in (?): to recover the original definition of the interleaving distance as a special case of the generalised one, we specify an allowable set of translations over which we wish to minimize. The original interleaving distance we saw above allows translations of the form $a \mapsto a + \varepsilon$ for some positive ε , and no others.

One source of examples of these more general types of interleavings between \mathbb{R} -persistence modules is to consider maps between metric spaces. For example, suppose we'd like to carry out a Vietoris-Rips filtration on a finite set, M , of $n \times n$ matrices. There are many examples of metrics which we could define on M , as we'll discuss in section ??, but there is no canonical choice. In section ??, we'll show that coarse and Lipschitz maps between metric spaces give rise to interleavings between Vietoris-Rips filtrations, and demonstrate that the interleavings defined in (?) provide a much more natural way of describing the relationship between such modules than the rather rigid original definition of an ε -interleaving.

One of the most fundamental properties of persistent homology is the fact that the common methods of filtering spaces - namely by Vietoris-Rips or sublevel-set filtrations - are *stable*. This property essentially means that if we have similar inputs to a persistence process, then the outputs — that is, the persistence modules — should also be similar. In particular, if a persistence module is the composition of a sublevel set filtration and a homology functor, then the distance between any two such modules is bounded by the distance between the filtering functions. Stability is such a fundamental property that any generalisation of the totally ordered case, to arbitrary posets, should also satisfy some form of stability properties. When the so-called “generalised” interleaving distance was defined in (?), a version of a stability theorem for sublevel set filtrations was proved. However, this proof involved a rather technical assumption about the poset.

In this section, we offer a new perspective on the stability properties of this generalised form of the interleaving distance. In (?), Lesnick observes that the supremum distance, d_∞ , on the space of \mathbb{R}^n -valued functions can itself be described as an interleaving distance. We use this observation, and the generalisation of the concept of an interleaving seen in (?), to show that if a pair of P -persistence modules are obtained from the sublevel set filtrations of a pair of P -valued functions, then an interleaving between the filtering functions implies the existence of an identical interleaving between the persistence modules.

2.1 Translations

In this section, we will explore how we can *shift* a persistence module by a translation. A translation is simply an increasing, order-preserving self-map of a poset.

Definition 2.1. A **translation** of a poset (P, \leq) is a map $s : P \rightarrow P$ such that for all $p_1, p_2 \in P$, $p_1 \leq p_2$ implies that $s(p_1) \leq s(p_2)$, and $s(p) \geq p$ for all $p \in P$.

If $s, t \in \text{Trans}(P, \leq)$ then clearly $s \circ t$ is also increasing. $\text{Trans}(P, \leq)$ is thus a monoid with respect to the operation of composition, where the identity in $\text{Trans}(P, \leq)$ is the identity translation, $\text{id}_P(p) = p$.

We note that $\text{Trans}(P, \leq)$ is also a partially ordered set itself, with the partial order,

$$s \leq t \text{ if } s(p) \leq t(p) \text{ for all } p \in P. \tag{2.1}$$

Example 2.2. Consider the poset \mathbb{R}^n , with the partial order given by,

$$a \leq b \text{ if and only if } a_i \leq b_i \text{ for all } 1 \leq i \leq n.$$

Given the partial order on the poset \mathbb{R}^n , a map $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a translation of \mathbb{R}^n is a translation of each coordinate. Hence,

$$\text{Trans}(\mathbb{R}^n) \cong \text{Trans}(\mathbb{R} \times \cdots \times \text{Trans}(\mathbb{R})).$$

Let $\varepsilon \geq 0$, and let $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be given by $s(a) = a + \varepsilon$, where $\varepsilon = (\varepsilon, \dots, \varepsilon)$. Then s is a translation of the poset \mathbb{R}^n .

This first example is a very important one. When an interleaving was first defined, it was only with respect to \mathbb{R} -persistence modules, and only in relation to translations of the form $a \mapsto a + \varepsilon$. The definitions in this section, which coincide with those of (?), are generalisations of this special case.

Example 2.3. Let $\delta \geq 1$, and let $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be given by $s(a) = \delta \cdot a$. Then, again, s is a translation of the poset \mathbb{R}^n .

Definition 2.4. Given a translation $s \in \text{Trans}(P, \leq)$, and a P -persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$, define the **s-shift of \mathcal{M}** to be the persistence module $\mathcal{M}(s) : P \rightarrow \mathcal{C}$ where for $p \in P$,

$$\mathcal{M}(s)(p) = \mathcal{M}(s(p)),$$

and whose transition maps are given by,

$$\psi_{\mathcal{M}(s)}(p_1, p_2) = \psi_{\mathcal{M}}(s(p_1), s(p_2)).$$

$$\begin{array}{ccccccc} \cdots & \longrightarrow & \mathcal{M}(a) & \xrightarrow{\quad} & \mathcal{M}(b) & \xrightarrow{\quad} & \mathcal{M}(c) & \longrightarrow & \cdots \\ & & \searrow f(a) & & \searrow f(b) & & \searrow f(c) & & \\ & & \mathcal{M}(s)(a) & \xrightarrow{\quad} & \mathcal{M}(s)(b) & \xrightarrow{\quad} & \mathcal{M}(s)(c) & \longrightarrow & \cdots \end{array}$$

Figure 19: The shift of a persistence module, \mathcal{M} by a translation, s .

We note that for any $s \in \text{Trans}(P, \leq)$, and any P -persistence module \mathcal{M} , we have a morphism $f : \mathcal{M} \rightarrow \mathcal{M}(s)$, called the **shift map**, where for each $p \in P$, $f(p) : \mathcal{M}(p) \rightarrow \mathcal{M}(s)(p)$ is the transition map $\psi_{\mathcal{M}}(p, s(p))$. This map is shown in Figure ??.

Additionally, we note that if we have a morphism $f : \mathcal{M} \rightarrow \mathcal{N}$ of P -persistence modules, then for any $s \in \text{Trans}(P, \leq)$ we have a morphism $f(s) : \mathcal{M} \rightarrow \mathcal{N}(s)$ given by,

$$f(s)(p) = \psi_{\mathcal{N}}(p, s(p)) \circ f(p). \tag{2.2}$$

2.2 Interleavings

Definition 2.5. For $s, t \in \text{Trans}(P, \leq)$, two P -persistence modules $\mathcal{M} : P \rightarrow \mathcal{C}$ and $\mathcal{N} : P \rightarrow \mathcal{C}$ are said to be (s, t) -interleaved if for all $p \in P$, there exist morphisms,

$$\begin{aligned} f(p) : \mathcal{M}(p) &\rightarrow \mathcal{N}(s(p)), \\ g(p) : \mathcal{N}(p) &\rightarrow \mathcal{M}(t(p)), \end{aligned}$$

such that for every $p_1 \leq p_2$ in P , the diagrams,

$$\begin{array}{ccc} \mathcal{M}(p_1) & \xrightarrow{\psi_{\mathcal{M}(p_1, p_2)}} & \mathcal{M}(p_2) \\ f(p_1) \downarrow & & \downarrow f(p_2) \\ \mathcal{N}(s(p_1)) & \xrightarrow{\psi_{\mathcal{N}(s(p_1), s(p_2))}} & \mathcal{N}(s(p_2)) \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{M}(t(p_1)) & \xrightarrow{\psi_{\mathcal{M}(t(p_1), t(p_2))}} & \mathcal{M}(t(p_2)) \\ g(p_1) \uparrow & & \uparrow g(p_2) \\ \mathcal{N}(p_1) & \xrightarrow{\psi_{\mathcal{N}(p_1, p_2)}} & \mathcal{N}(p_2) \end{array} \quad (2.3)$$

commute, and for every $p \in P$,

$$\begin{aligned} g(s(p)) \circ f(p) &= \psi_{\mathcal{M}(a, ts(p))}, \\ f(t(p)) \circ g(p) &= \psi_{\mathcal{N}(a, st(p))}. \end{aligned}$$

Or, equivalently, for every $p \in P$, the diagrams,

$$\begin{array}{ccc} \mathcal{M}(p) & \xrightarrow{\psi_{\mathcal{M}(p, ts(p))}} & \mathcal{M}(ts(p)) \\ f(p) \searrow & & \nearrow g(s(p)) \\ & \mathcal{N}(s(p)) & \end{array} \quad \text{and} \quad \begin{array}{ccc} & \mathcal{M}(t(p)) & \\ g(p) \nearrow & & \searrow f(t(p)) \\ \mathcal{N}(p) & \xrightarrow{\psi_{\mathcal{N}(p, st(p))}} & \mathcal{N}(st(p)) \end{array}$$

commute.

Notice how if both s and t are the identity translation on P , then two P -persistence modules are (s, t) -interleaved if and only if they are isomorphic. For this reason, an interleaving between persistence modules can be thought of as a generalisation of an isomorphism, or an isomorphism up to s and t .

The following example appears as Lemma 3.5(a) in (?) without proof. The example shows how interval modules, introduced in section ??, are interleaved.

Example 2.6. For an interval $I \subset \mathbb{R}$, and $\varepsilon > 0$ define $Ex^\varepsilon(I)$ to be the set,

$$Ex^\varepsilon(I) = \{a \in \mathbb{R} \mid \exists b \in I \text{ such that } |a - b| \leq \varepsilon\}. \quad (2.4)$$

Suppose that I and J are intervals such that $I \subseteq Ex^\varepsilon(J)$ and $J \subseteq Ex^\varepsilon(I)$. Let $s, t \in \text{Trans}(\mathbb{R}, \leq)$ be given by $s(a) = t(a) = a + \varepsilon$.

Let $k[I]$ and $k[J]$ be the interval modules of Example ?? Define morphisms $f : k[I] \rightarrow k[J](s)$ and $g : k[J] \rightarrow k[I](t)$ by,

$$f(a) = \begin{cases} id_k & a \in I, a + \varepsilon \in J \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad g(a) = \begin{cases} id_k & a \in J, a + \varepsilon \in I \\ 0 & \text{otherwise.} \end{cases} \quad (2.5)$$

That is, $f(a) : k[I](a) \rightarrow k[J](s)(a)$ is the identity map whenever both $k[I](a)$ and $k[J](s)(a)$ are a copy of k , and is the zero map otherwise. Similarly, $g(a) : k[J](a) \rightarrow k[I](t)(a)$ is the identity map whenever both $k[J](a)$ and $k[I](t)(a)$ are a copy of k , and is the zero map otherwise.

We leave it for the interested reader to check that these morphisms are well-defined. This amounts to checking that the diagrams,

$$\begin{array}{ccc} k[I](a) & \xrightarrow{\psi_{k[I](a,b)}} & k[I](b) & & k[I](t)(a) & \xrightarrow{\psi_{k[I](t(a),t(b))}} & k[I](t)(b) \\ f(a) \downarrow & & \downarrow f(b) & \text{and} & g(a) \uparrow & & \uparrow g(b) \\ k[J](s)(a) & \xrightarrow{\psi_{k[J](s(a),s(b))}} & k[J](s)(b) & & k[J](a) & \xrightarrow{\psi_{k[J](a,b)}} & k[J](b) \end{array} \quad (2.6)$$

commute for all $a, b \in \mathbb{R}$. There are 16 cases to check for each diagram. For instance, for the first diagram, we much check all cases depending on whether a, b are elements of I , or not, and whether $a + \varepsilon, b + \varepsilon$ are elements of J , or not. We will not include this proof here.

Once we are suitably convinced that f and g are well defined persistence module morphisms, we can now consider how they compose. Our aim is to show that for all $a \in \mathbb{R}$,

$$f(a + \varepsilon) \circ g(a) = \psi_{k[J](a, a + 2\varepsilon)}, \quad \text{and} \quad g(a + \varepsilon) \circ f(a) = \psi_{k[I](a, a + 2\varepsilon)} \quad (2.7)$$

All maps in (??) are either the identity map on k , or the zero map, depending on whether $a, a + \varepsilon, a + 2\varepsilon$ belong to I and J . We now need to show that the maps compose to coincide for all values of a . We note that $f(a + \varepsilon) \circ g(a)$ is the identity map when both $g(a)$ and $f(a + \varepsilon)$ are the identity map. That is,

$$f(a + \varepsilon) \circ g(a) = \begin{cases} id_k & a \in J, a + \varepsilon \in I \text{ and } a + 2\varepsilon \in J \\ 0 & \text{otherwise} \end{cases}$$

Meanwhile, the transition map $\psi_{k[J](a, a + 2\varepsilon)}$ is the identity map when $[a, a + 2\varepsilon] \subseteq J$, or since J is an interval, when $a, a + 2\varepsilon \in J$.

We note that for any $a \in \mathbb{R}$, if $a \in J$ and $a + 2\varepsilon \in J$ then $a + \varepsilon \in I$. We see why this must be true by supposing otherwise — that there is some $a \in \mathbb{R}$ such that $a \in J, a + 2\varepsilon \in J$, but that $a + \varepsilon \notin I$.

As $J \subset Ex^\varepsilon(I)$, then for every $j \in J$, there exists an $i \in I$ such that $|i - j| \leq \varepsilon$. So, in particular, there exists $i_1 \in I$ such that,

$$|a - i_1| \leq \varepsilon. \quad (2.8)$$

If $a + \varepsilon \notin I$, then this i_1 must be strictly less than $a + \varepsilon$, as any i_1 greater than $a + \varepsilon$ violates (??). Hence, there is some $i_1 < a + \varepsilon$ such that $i_1 \in I$. Similarly, there exists $i_2 \in I$ such that,

$$|(a + 2\varepsilon) - i_2| \leq \varepsilon. \quad (2.9)$$

Again, given that $a + \varepsilon \notin I$, then this i_2 must be strictly greater than $a + \varepsilon$, since any i_2 less than $a + \varepsilon$ violates (??). Hence there is some $i_2 > a + \varepsilon$ such that $i_2 \in I$.

As I is an interval, it must be that $a + \varepsilon \in I$.

Hence,

$$f(a + \varepsilon) \circ g(a) = \begin{cases} id_k & a \in J, a + 2\varepsilon \in J \\ 0 & \text{otherwise} \end{cases} = \psi_{k[J]}(a, a + 2\varepsilon). \quad (2.10)$$

A similar argument shows that $g(a + \varepsilon) \circ f(a) = \psi_{k[I]}(a, a + 2\varepsilon)$ for all $a \in \mathbb{R}$.

Hence $k[I]$ and $k[J]$ are (s, t) -interleaved.

Remark 2.7. If persistence modules $\mathcal{M} : P \rightarrow \mathcal{C}$ and $\mathcal{N} : P \rightarrow \mathcal{C}$ are (s, t) -interleaved, and s' and t' are such that $s' \geq s, t' \geq t$, then \mathcal{M} and \mathcal{N} are also (s', t') -interleaved.

Conversely, if \mathcal{M} and \mathcal{N} are *not* (s', t') -interleaved, then they are also not (s, t) -interleaved.

2.3 Metric Spaces and the Interleaving Distance

In section ??, we presented the definition of an interleaving in its most general form. But, as we have discussed, this is not the definition we will see in most places in the literature. The original definition — see, for example, (?) or (?) — allows only translations of the poset (\mathbb{R}, \leq) of the form,

$$a \mapsto a + \varepsilon.$$

Clearly Definition ?? seems appropriate if we are to define interleavings for arbitrary poset modules, but is there an advantage to using this more general definition even for \mathbb{R} -modules? In this section, we explore some of the relationships between data sets which the generalised definition of an interleaving can be used to describe.

The following examples arise as the result of a simple question - how does the choice of metric affect a Vietoris Rips filtration of a metric space? It may be the case that our data lives in a space on which we can define more than one metric. For example, suppose that we have a set of points, $\mathbb{X} = \{X_1, \dots, X_m\}$ such that each X_i is an $n \times n$ matrix with entries in \mathbb{R} . Then each X_i can be thought of as a point of the metric space (\mathbb{R}^{n^2}, d_1) , where we

think of a matrix $X \in M_{n \times n}(\mathbb{R})$ with elements,

$$X = \begin{pmatrix} X_{1,1} & X_{1,2} & \cdots & X_{1,n} \\ X_{2,1} & X_{2,2} & \cdots & X_{2,n} \\ \vdots & \ddots & & \\ X_{n,1} & & & X_{n,n} \end{pmatrix} \quad (2.11)$$

as the vector $(X_{1,1}, X_{1,2}, \dots, X_{1,n}, X_{2,1}, \dots, X_{n,n})$ in \mathbb{R}^{n^2} , and $d_1(X_i, X_j)$ is simply the Euclidean distance between X_i and X_j as vectors in \mathbb{R}^{n^2} . On the other hand, we could also consider the metric space $(M_{n \times n}(\mathbb{R}), d_2)$, where for $X, Y \in M_{n \times n}(\mathbb{R})$,

$$d_2(X, Y) = \text{rank}(Y - X). \quad (2.12)$$

In this section, we will investigate how the choice of metric, d , on a space, M , affects a Vietoris-Rips filtration on (M, d) .

Proposition 2.8. *Let (M, d_1) and (N, d_2) be finite metric spaces. Suppose there exists a bi-Lipschitz map $f : (M, d_1) \rightarrow (N, d_2)$ with Lipschitz constant $C \geq 1$. Let \mathcal{M} and \mathcal{N} denote the Vietoris-Rips filtrations on (M, d_1) and $(f(M), d_2)$ respectively. Then \mathcal{M} and \mathcal{N} are (s, s) -interleaved, where $s : \mathbb{R} \rightarrow \mathbb{R}$ is given by $s(a) = C \cdot a$.*

Proof. If $f : (M, d_1) \rightarrow (N, d_2)$ is a bi-Lipschitz map with constant $C \geq 1$, then for any $m_i, m_j \in M$,

$$\frac{1}{C}d_1(m_i, m_j) \leq d_2(f(m_i), f(m_j)) \leq Cd_1(m_i, m_j). \quad (2.13)$$

The Vietoris-Rips filtrations on (M, d_1) and $(f(M), d_2)$ are completely determined by distances between pairs of points. That is, for a given $a \in \mathbb{R}$, the complex $\mathcal{M}(a)$ is determined by pairs m_i, m_j such that $d_1(m_i, m_j) \leq a$. We know that for every such pair, given (??), the images of m_i and m_j in $(f(M), d_2)$ satisfy,

$$\frac{a}{C} \leq d_2(f(m_i), f(m_j)) \leq Ca.$$

Hence, for every $a \in \mathbb{R}$, there is an inclusion of Vietoris-Rips complexes,

$$\mathcal{M}(a) \hookrightarrow \mathcal{N}(Ca). \quad (2.14)$$

For any $n_i = f(m_i), n_j = f(m_j)$,

$$d_1(f^{-1}(n_i), f^{-1}(n_j)) \leq Cd_2(ff^{-1}(n_i), ff^{-1}(n_j)) = Cd_2(n_i, n_j).$$

Hence f^{-1} is also Lipschitz with constant C , and there is an inclusion of Vietoris-Rips complexes,

$$\mathcal{N}(a) \hookrightarrow \mathcal{M}(Ca). \quad (2.15)$$

Together (??) and (??) give rise to an (s, s) -interleaving of \mathcal{M} and \mathcal{N} . \square




Figure 20: Coarse maps f and g between metric spaces (M, d_1) and (N, d_2) .

We note that this result would also hold if we took \mathcal{M} and \mathcal{N} to be the homology of the Vietoris-Rips complexes, that is, if we took \mathcal{M} to be the persistence module where for $a \in \mathbb{R}$,

$$\mathcal{M}(a) = H_n(\text{VR}(M, a); k),$$

for some $n \in \mathbb{Z}$ and some field k .

Proposition 2.9. *Suppose that (M, d_1) and (N, d_2) are coarsely equivalent metric spaces, where $f : (M, d_1) \rightarrow (N, d_2)$ and $g : (N, d_2) \rightarrow (M, d_1)$ denote the coarse maps, as in Figure ???. Let $\mathbb{M} = \{m_1, \dots, m_l\}$ be a finite set of points sampled from M . Let \mathcal{M} and \mathcal{N} denote the Vietoris-Rips filtrations on (\mathbb{M}, d_1) and $(f(\mathbb{M}), d_2)$ respectively. Then there exists an interleaving of \mathcal{M} and \mathcal{N} .*

Proof. The condition that for every $\varepsilon_i > 0$, there exists a real number $\delta_i > 0$ such that for every $m_i, m_j \in \mathbb{M}$ with $d_1(m_i, m_j) \leq \varepsilon_i$, we have $d_2(f(m_i), f(m_j)) \leq \delta_i$, means that for every $\varepsilon_i \in \mathbb{R}$, we have an inclusion of Vietoris-Rips complexes,

$$\mathcal{M}(\varepsilon_i) \hookrightarrow \mathcal{N}(\delta_i). \quad (2.16)$$

For an opposite inclusion, we first note that as g is a coarse map, then for any $\varepsilon_i \in \mathbb{R}$, there exists a $\gamma_i \geq 0$ such that for any $n_i, n_j \in f(\mathbb{M})$, whenever $d_2(n_i, n_j) = d_2(f(m_i), f(m_j)) \leq \varepsilon_i$, then $d_1(g \circ f(m_i), g \circ f(m_j)) \leq \gamma_i$. However, the points $g \circ f(m_i)$ and $g \circ f(m_j)$ in M do not necessarily belong to the set \mathbb{M} . Therefore the inequality above does not necessarily give an inclusion,

$$\mathcal{N}(\varepsilon_i) \hookrightarrow \mathcal{M}(\gamma_i), \quad (2.17)$$

but instead an inclusion,

$$\mathcal{N}(\varepsilon_i) \hookrightarrow \widetilde{\mathcal{M}}(\gamma_i), \quad (2.18)$$

where $\widetilde{\mathcal{M}}$ denotes a Vietoris-Rips filtration on $\widetilde{\mathbb{M}} = \{g \circ f(m_1), \dots, g \circ f(m_n)\} = g \circ f(\mathbb{M})$. However, we use the fact that since f and g are coarse maps, there exists c_1, c_2 such that for all $m \in M$,

$$d_1(m, g \circ f(m)) \leq c_1.$$

and for all $n \in N$,

$$d_2(n, f \circ g(n)) \leq c_2.$$

Together with the triangle inequality for d_1 , we have that whenever $d_1(g \circ f(m_i), g \circ f(m_j)) \leq \gamma_i$,

$$d_1(m_i, m_j) \leq d_1(m_i, g \circ f(m_i)) + d_1(g \circ f(m_i), g \circ f(m_j)) + d_1(g \circ f(m_j), m_j) = \gamma_i + 2c_1, \quad (2.19)$$

hence we have an inclusion of Vietoris-Rips complexes,

$$\mathcal{N}(\varepsilon_i) \hookrightarrow \mathcal{M}(\gamma_i + 2c_1). \quad (2.20)$$

Let $s : \mathbb{R} \rightarrow \mathbb{R}$ be the map given by,

$$s(\varepsilon_i) = \min\{\varepsilon_i, \inf\{\delta_i \mid d_1(m_i, m_j) \leq \varepsilon_i \Rightarrow d_2(f(m_i), f(m_j)) \leq \delta_i\}\}, \quad (2.21)$$

and let $t : \mathbb{R} \rightarrow \mathbb{R}$ be the map given by,

$$t(\varepsilon_i) = \min\{\varepsilon_i, \inf\{\gamma_i \mid d_2(n_i, n_j) \leq \varepsilon_i \Rightarrow d_1(g(n_i), g(n_j)) \leq \gamma_i\} + 2c_1\}. \quad (2.22)$$

Then s and t are translations such that \mathcal{M} and \mathcal{N} are (s, t) -interleaved. \square

2.4 The Interleaving Distance

The relationship of an interleaving can be to define a distance on the space of P -persistence modules (?). First, we need a measure of the quality of the relationship which an interleaving describes. Clearly, an (id_P, id_P) -interleaving describes the closest possible relationship between P -modules. In order to describe the proximity between a pair of (s, t) -interleaved modules, we require that P has certain extra structure, as we now detail.

Definition 2.10. A **weight function** on a poset (P, \leq) is a function $w : P \times P \rightarrow [0, \infty]$ satisfying,

- $w(p_i, p_i) = 0$ for all $p_i \in P$,
- for any p_i, p_j, p_k such that $p_i \leq p_j \leq p_k$, $w(p_i, p_k) \leq w(p_i, p_j) + w(p_j, p_k)$.

A poset (P, \leq) together with a weight function, w , is called a **weighted poset**.

This condition states that P is what is known in (?) as a **weighted category**, also referred to as a normed category in (?) or an additive category in (?).

Example 2.11. Let H_P denote the Hasse diagram of the poset (P, \leq) . For any $p_i, p_j \in P$, let $w(p_i, p_j)$ be the length of the shortest directed path in H_P from p_i to p_j , where the length of a path between vertices in H_P is the number of edges in the path. Note that $w(p_i, p_j) = \infty$ if there is no directed path between p_i and p_j - that is, if $p_i \not\leq p_j$.

Example 2.12. Let P be the 5-element poset $\{a, b, c, d, e\}$ with partial order as shown in the Hasse diagram, H_P , in Figure ???. For any $p_i, p_j \in P$, let $w(p_i, p_j)$ be the length of the shortest directed path in H_P .

Then, for example, $w(a, c) = 2$, but $w(c, a) = \infty$ as there is no directed path from c to a . Similarly, $w(c, e) = \infty$.




Figure 21: A Hasse diagram, H_P for a 5-element poset.

Definition 2.13. Given a poset P and a weight function, w , set,

$$d_P(p_1, p_2) = \begin{cases} w(p_i, p_j) & p_i \leq p_j \\ \infty & \text{otherwise.} \end{cases}$$

Then (P, d_P) is an example of a Lawvere metric space. Hence any weighted poset defines a Lawvere metric space (?).

Definition 2.14. Let $\omega : Trans(P, \leq) \rightarrow \mathbb{R}_{\leq 0}$ be given by,

$$\omega(s) = \sup_{p \in P} \{d_P(p, s(p))\}. \quad (2.23)$$

A translation $s \in Trans(P, \leq)$ is an ε -translation if $\omega(s) \leq \varepsilon$.

Definition 2.15. Let T be a submonoid of $Trans(P, \leq)$, and let $\varepsilon \geq 0$. Two P -persistence modules $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$ are (ε, T) -interleaved if there exists a pair of ε -translations $s, t \in T$ such that \mathcal{M} and \mathcal{N} are (s, t) -interleaved.

Definition 2.16. Let P be a weighted poset and let T be a submonoid of $Trans(P, \leq)$. The **generalised interleaving distance** between P -persistence modules with respect to T is,

$$d_T^T(\mathcal{M}, \mathcal{N}) = \inf\{\varepsilon \mid \mathcal{M} \text{ and } \mathcal{N} \text{ are } (\varepsilon, T)\text{-interleaved}\}. \quad (2.24)$$

Lemma 2.17. *The generalised interleaving distance is an extended pseudometric.*

Proof. We need to check that the three conditions of Definition ?? are satisfied for d_T^T . Firstly, as T is a submonoid of $Trans(P, \leq)$, the identity translation id_P must belong to T . For any P -persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$, there exists an (id_P, id_P) -interleaving of \mathcal{M} and \mathcal{M} . Hence for any $\mathcal{M} : P \rightarrow \mathcal{C}$, $d_T^T(\mathcal{M}, \mathcal{M}) = 0$ since \mathcal{M} and \mathcal{M} are $(0, T)$ -interleaved.

Secondly, we have that for any $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$, if \mathcal{M} and \mathcal{N} are (s, t) -interleaved for $s, t \in T$, then \mathcal{N} and \mathcal{M} are (t, s) -interleaved. Hence $d_I^T(\mathcal{M}, \mathcal{N}) = d_I^T(\mathcal{N}, \mathcal{M})$ for any $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$ and for any $T < Trans(P, \leq)$.

To show that the triangle inequality holds for d_I^T , we will use a result of (?), where it is proved that for any $\mathcal{M}, \mathcal{N}, \mathcal{L} : P \rightarrow \mathcal{C}$, if \mathcal{M} and \mathcal{N} are (s_1, t_1) -interleaved, and \mathcal{N} and \mathcal{L} are (s_2, t_2) -interleaved, then \mathcal{M} and \mathcal{L} are $(s_1 s_2, t_1 t_2)$ -interleaved. Since T is a submonoid of $Trans(P, \leq)$, it is closed under the monoid operation of composition. Hence if $s_1, t_1 \in T$ and $s_2, t_2 \in T$, then $s_1 s_2, t_1 t_2 \in T$.

Finally, if s_1, t_1 are ε_1 -translations, and s_2, t_2 are ε_2 -translations, then,

$$\sup_{p \in P} d_P(p, s_1 s_2(p)) \leq \sup_{p \in P} d_P(p, s_2(p)) + \sup_{p \in P} d_P(s_2(p), s_1 s_2(p)) \leq \varepsilon_2 + \varepsilon_1,$$

and similarly for $t_1 t_2$.

Hence if \mathcal{M} and \mathcal{N} are (ε_1, T) -interleaved, and \mathcal{N} and \mathcal{L} are (ε_2, T) -interleaved, then \mathcal{M} and \mathcal{L} are $(\varepsilon_1 + \varepsilon_2, T)$ -interleaved. \square

Here we see why it is necessary to take T to be a submonoid of $Trans(P, \leq)$ as opposed to just a subset. It is the property of being closed under composition of translations that allows us to demonstrate that a triangle inequality holds for the interleaving distance, d_I^T .

Example 2.18. Consider the poset (\mathbb{R}^n, \leq) where \leq is the usual partial order on \mathbb{R}^n . Let T_E be the submonoid of $Trans(\mathbb{R}^n)$ given by,

$$T_E = \{t \in Trans(\mathbb{R}^n, \leq) \mid t(a) = a + \varepsilon\}, \quad (2.25)$$

where $\varepsilon = (\varepsilon \dots, \varepsilon)$, for some $\varepsilon \geq 0$. Then for any $t \in T_E$, $\omega(t) = \varepsilon$.

This is the usual interleaving distance first defined in (?) for single parameter \mathbb{R} -modules.

Example 2.19. For any poset (P, \leq) , setting $T = Trans(P)$ gives the interleaving distance seen in (?).

Remark 2.20. If a pair of P -persistence modules are (ε, T_1) -interleaved, for some $T_1 < Trans(P, \leq)$, then they are also (ε, T_2) -interleaved for any $T_1 < T_2 < Trans(P, \leq)$. Hence for any $T_1 < T_2 < Trans(P, \leq)$, we have that $d_I^{T_2} \leq d_I^{T_1}$.

The following Remark, noted in (?), will be particularly useful when we want to determine whether a pair of \mathbb{R}^n -persistence modules are (s, t) -interleaved for a given pair of ε -translations s and t .

Remark 2.21. If $s, t \in Trans(\mathbb{R}^n, \leq)$ are ε -translations for some $\varepsilon \geq 0$, then

$$\omega(s), \omega(t) \leq \varepsilon = \omega(s_\varepsilon),$$

where $s_\varepsilon \in \text{Trans}(\mathbb{R}^n, \leq)$ is given by,

$$s_\varepsilon(a) = a + \varepsilon.$$

Given Remark ??, we have that if a pair of \mathbb{R}^n -persistence modules are **not** $(s_\varepsilon, s_\varepsilon)$ -interleaved, then they are not (s, t) -interleaved for any ε -translations $s, t \in \text{Trans}(P, \leq)$.

The interleaving distance provides a way to measure the proximity between persistence modules, but is not particularly easy to compute. As Lesnick (?) observes, the problem of determining whether a pair of persistence modules $\mathcal{M}, \mathcal{N} : \mathbb{R}^n \rightarrow \text{Vect}_k$ are (ε, T_E) -interleaved for some $\varepsilon \geq 0$ is NP-complete. For $n = 1$, we can define a discrete invariant which completely describes the persistence module, known as a barcode or persistence diagram. The interleaving distance between \mathbb{R} -modules is then equivalent to a combinatorial matching distance between the modules' barcodes (?). One of the greatest challenges in multiparameter, or generalised persistent homology is the fact that these invariants can no longer be defined. As (?) and (?) have shown, the challenge is not in defining some sensible invariant which mimics the barcode for more general modules, but in defining an invariant which is complete – that is, contains all of the information in a persistence module – and which is as useful as the barcode.

2.5 The Stability of Persistent Homology

If persistence is to be of any use in applications, and especially if it is to be used as a means of distinguishing spaces and data sets, we need some guarantee that it is suitably well-behaved, in the sense that filtrations of similar spaces – or datasets – are similar as persistence modules. We would also like the guarantee that comparing persistence modules is at least as useful as other standard methods we have of comparing the filtered spaces. These properties are collectively referred to as ‘stability’ properties.

Here we detail two of the ways in which persistent homology can be said to be stable.

Firstly, suppose X and Y are two finite metric spaces. In section 1, we described the Vietoris-Rips filtrations of finite metric spaces. Let $VR(X) : \mathbb{R}_{\geq 0} \rightarrow \text{Subsets}(X)$ and $VR(Y) : \mathbb{R}_{\geq 0} \rightarrow \text{Subsets}(Y)$ be Vietoris-Rips filtrations of X and Y , respectively, and let $\mathcal{M}, \mathcal{N} : \mathbb{R}_{\geq 0} \rightarrow \text{Vect}_k$ be the compositions of $VR(X)$ and $VR(Y)$ with a simplicial homology functor, respectively.

The first stability theorem is due to (?).

Theorem 2.22. *Let $T_E < \text{Trans}(\mathbb{R}, \leq)$ be the submonoid given by,*

$$T_E = \{s \in \text{Trans}(\mathbb{R}, \leq) \mid s(a) = a + \varepsilon \text{ for some } \varepsilon \geq 0\}.$$

Then,

$$d_I^{T_E}(\mathcal{M}, \mathcal{N}) \leq d_{GH}(X, Y), \tag{2.26}$$

where d_{GH} denotes the Gromov-Hausdorff distance between X and Y .

We can think of Theorem ?? as a guarantee that if we move the points of a finite metric space a small amount, then the resulting persistence module will change by a similarly small amount. That is, we will not create a significant persistent homology class which was not present in the original module, or destroy one which was. This is clearly a desirable and important quality for persistent homology to possess – it tells us that similar point clouds really do have similar Vietoris-Rips filtrations, making persistent homology a good measure of similarity between such objects. Proofs of similar statements involving Čech and witness complexes can also be found in (?).

The second stability theorem we acknowledge in this section generalises those seen in (?) and (?). The precise wording is due to (?). This second result concerns sublevel set filtrations as opposed to Vietoris-Rips ones.

Let **n-Fun** be the category whose objects are functions of the form $\gamma_X : X \rightarrow \mathbb{R}^n$ for some topological space X . A morphism in this category, or an element of $hom(\gamma_X, \gamma_Y)$ for some $\gamma_X : X \rightarrow \mathbb{R}^n, \gamma_Y : Y \rightarrow \mathbb{R}^n$, is a continuous function $f : X \rightarrow Y$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(f(x)). \quad (2.27)$$

For such a $\gamma_X \in \mathbf{n-Fun}$, let $\mathcal{M} : \mathbb{R}^n \rightarrow \mathbf{Vect}$ be the persistence module defined by,

$$\mathcal{M}(a) = H_i(\{x \in X | \gamma_X(x) \leq a\}), \quad (2.28)$$

with linear maps $\psi_{\mathcal{M}}(a, b) : \mathcal{M}(a) \rightarrow \mathcal{M}(b)$ induced by the inclusions of sub-level sets

$$\{x \in X | \gamma_X(x) \leq a\} \hookrightarrow \{x \in X | \gamma_X(x) \leq b\},$$

whenever $a \leq b$. Similarly, for $\gamma_Y \in \mathbf{n-Fun}$, let $\mathcal{N} : \mathbb{R}^n \rightarrow \mathbf{Vect}$ be the persistence module defined by,

$$\mathcal{N}(a) = H_i(\{y \in Y | \gamma_Y(y) \leq a\}), \quad (2.29)$$

with linear maps $\psi_{\mathcal{N}}(a, b) : \mathcal{N}(a) \rightarrow \mathcal{N}(b)$ induced by the inclusions of sub-level sets,

$$\{y \in Y | \gamma_Y(y) \leq a\} \hookrightarrow \{y \in Y | \gamma_Y(y) \leq b\},$$

whenever $a \leq b$.

Then the interleaving distance, $d_I^{TE}(\mathcal{M}, \mathcal{N})$, provides one means of measuring the proximity between γ_X and γ_Y . Another way is to use the pseudometric,

$$d_\infty(\gamma_X, \gamma_Y) = \inf_{h \in Homeo(X, Y)} \|\gamma_X - \gamma_Y \circ h\|_\infty, \quad (2.30)$$

where for $a \in \mathbb{R}^n$,

$$\|a\|_\infty = \max\{|a_i| \mid i = 1, \dots, n\}. \quad (2.31)$$

Theorem 2.23. *Let X be a topological space. For any $\gamma_X, \gamma_Y : X \rightarrow \mathbb{R}^n$,*

$$d_I^{TE}(\mathcal{M}, \mathcal{N}) \leq d_\infty(\gamma_X, \gamma_Y). \quad (2.32)$$

This second stability result tells us that if two persistence modules are constructed from sublevel set filtrations of \mathbb{R}^n -valued functions, then similar functions give rise to similar persistence modules. Again, this stability property tells us that the interleaving distance between persistence modules is a good measure of similarity between filtering functions. A proof for this second result can be found in (?).

2.6 Stability of the Generalised Interleaving Distance

The previous section covers the stability of the interleaving distance d_I^{TE} . But what about the stability of the generalised interleaving distance d_I^T for other choices of T ? In this section we define a measure of similarity between P -valued functions, which provides an upper bound for the generalised interleaving distance between sublevel set filtrations of such functions. We note that a dual problem was considered by Frosini in (?). Given a pair of \mathbb{R} -valued functions $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$, for some topological spaces X and Y , the stability theorem for persistent homology tells us that if we construct a pair of persistence modules from sublevel set filtrations by f and g , then the interleaving distance between the modules provides a lower bound for the supremum distance $d_\infty(f, g)$, where the supremum distance is defined by,

$$d_\infty(f, g) = \inf_{h \in \text{Homeo}(X, Y)} \|f - g \circ h\|_\infty. \quad (2.33)$$

Frosini's work finds a lower bound for a variant d_∞^G of the supremum distance, which for G a proper subgroup of $\text{Homeo}(X, Y)$, is defined as,

$$d_\infty^G(f, g) = \inf_{h \in G < \text{Homeo}(X, Y)} \|f - g \circ h\|_\infty, \quad (2.34)$$

The appropriate lower bound for d_∞^G can be thought of as a variant of the interleaving distance. The dual problem we consider in this section considers the opposite: by adapting the interleaving distance, to specify a particular monoid of poset translations, we find an appropriate variant of d_∞ which is an upper bound for this interleaving distance.

The question of the stability of d_I^T where $T = \text{Trans}(P, \leq)$ for some poset (P, \leq) was considered in (?). This result is given for a partially ordered set $P \subseteq \text{Subsets}(L)$, where $\text{Subsets}(L)$ is the set of subsets of a Lawvere metric space (L, d_L) , and the partial order on P is given by the inclusion of subsets. Take two functions, $\gamma_1, \gamma_2 : X \rightarrow L$, for some topological space, X , and let the persistence modules \mathcal{M} and \mathcal{N} be given by,

$$\mathcal{M}(A) = \gamma_1^{-1}(A), \text{ and } \mathcal{N}(A) = \gamma_2^{-1}(A), \quad (2.35)$$

for $A \in P$. That is, \mathcal{M} and \mathcal{N} are persistence modules $\mathcal{M}, \mathcal{N} : P \rightarrow \mathbf{Top}$, with transition maps $\psi_{\mathcal{M}}(A, B)$ and $\psi_{\mathcal{N}}(A, B)$ given by the inclusion of subspaces,

$$\gamma_1^{-1}(A) \hookrightarrow \gamma_1^{-1}(B) \text{ and } \gamma_2^{-1}(A) \hookrightarrow \gamma_2^{-1}(B), \quad (2.36)$$

respectively, for any $A \subseteq B$.

Definition 2.24. A poset $P \subseteq \text{Subsets}(L)$ has **enough translations** if for every $0 \leq \varepsilon \leq \eta$ there exists $s = s_{\varepsilon, \eta} \in \text{Trans}(P)$ with $\omega(s) \leq \eta$, such that for every $A \in P$ we have an inclusion,

$$A^\varepsilon = \{l \in L \mid \text{there exists } a \in A \text{ such that } d_L(a, l) \leq \varepsilon\} \subseteq s(A). \quad (2.37)$$

Theorem 2.25. *If P has enough translations, then the generalised interleaving distance between \mathcal{M} and \mathcal{N} satisfies,*

$$d_I(\mathcal{M}, \mathcal{N}) \leq d(\gamma_1, \gamma_2), \quad (2.38)$$

where

$$d(\gamma_1, \gamma_2) = \max\{\hat{d}(\gamma_1, \gamma_2), \hat{d}(\gamma_2, \gamma_1)\}, \quad (2.39)$$

and

$$\hat{d}(\gamma_1, \gamma_2) = \sup_{x \in X} \{d_L(\gamma_1(x), \gamma_2(x))\}. \quad (2.40)$$

In this section, we provide an alternative statement of the stability properties of the generalised interleaving distance, which avoids the technical assumption that P has enough translations. We will do this by reformulating the pseudometric d_∞ as an interleaving distance in the category of P -valued functions, **P-Fun**, following an intriguing remark by Lesnick in (?).

2.6.1 d_∞ as an Interleaving Distance

Earlier in section ?? we saw the category **n-Fun**, whose objects are functions $\gamma_X : X \rightarrow \mathbb{R}^n$, and whose set of morphisms $\text{hom}(\gamma_X, \gamma_Y)$ consists of continuous functions $f : X \rightarrow Y$ such that for all $x \in X$, $\gamma_X(x) \geq \gamma_Y(f(x))$. In this section we will explain how Lesnick's Remark 5.1 in (?) can be used to define an interleaving distance in the category **n-Fun**. For now, we'll consider the monoid $T_E < \text{Trans}(\mathbb{R}^n, \leq)$.

Definition 2.26. Given $\gamma_X, \gamma_Y \in \text{Ob}(\mathbf{n}\text{-Fun})$, $f \in \text{hom}(\gamma_X, \gamma_Y)$ and $s \in T_E$, we define the shift of γ_X by s , $\gamma_X(s) \in \text{Ob}(\mathbf{n}\text{-Fun})$, the shift map, $f(s) \in \text{hom}(\gamma_X, \gamma_Y(s))$, and the transition map, $\psi_{\gamma_X}(\gamma_X, \gamma_X(s))$.

1. Suppose $s \in T_E$ is given by $s(a) = a + \varepsilon$ for some $\varepsilon = (\varepsilon, \dots, \varepsilon)$, where $\varepsilon \geq 0$. Given $\gamma_X \in \text{Ob}(\mathbf{n}\text{-Fun})$, we set $\gamma_X(s)$ to be,

$$\gamma_X(s)(x) := \gamma_X(x) - \varepsilon. \quad (2.41)$$

2. Note that an element of $\text{hom}(\gamma_X, \gamma_Y(s))$ is a continuous map $f(s) : X \rightarrow Y$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(s)(f(s)(x)), \quad (2.42)$$

or equivalently, $\gamma_X(x) \geq \gamma_Y(f(s)(x)) - \varepsilon$.

If we are given $f \in \text{hom}(\gamma_X, \gamma_Y)$, then we have that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(f(x)),$$

and so $\gamma_X(x) \geq \gamma_Y(f(x)) \geq \gamma_Y(f(x)) - \varepsilon$. Hence, given $f \in \text{hom}(\gamma_X, \gamma_Y)$, we let $f(s) \in \text{hom}(\gamma_X, \gamma_Y(s))$ be given by $f(s) = f$.

3. To define transition maps $\psi_{\gamma_X}(\gamma_X, \gamma_X(s))$ for a translation $s \in T_E$, we note that an element of $\text{hom}(\gamma_X, \gamma_X(s))$ is a continuous map $\psi_{\gamma_X} : X \rightarrow X$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_X(s)(\psi_{\gamma_X}(x)) = \gamma_X(\psi_{\gamma_X}(x)) - \varepsilon.$$

If we let $\psi_{\gamma_X} = \text{id}_X$, then this is true for any $\varepsilon \geq 0$.

Definition 2.27. For $\gamma_X, \gamma_Y \in \mathbf{n}\text{-Fun}$, we say that $\gamma_X : X \rightarrow \mathbb{R}^n$ and $\gamma_Y : Y \rightarrow \mathbb{R}^n$ are (ε, T_E) -interleaved if there exist $s, t \in T_E$, $f \in \text{hom}(\gamma_X, \gamma_Y(s))$, and $g \in \text{hom}(\gamma_Y, \gamma_X(t))$ such that

$$f(t) \circ g = \psi_{\gamma_X}(\gamma_X, \gamma_X(ts)) \text{ and } g(s) \circ f = \psi_{\gamma_Y}(\gamma_Y, \gamma_Y(st)). \quad (2.43)$$

The interleaving distance, $d_I^{T_E}(\gamma_X, \gamma_Y)$ is then given by,

$$d_I^{T_E}(\gamma_X, \gamma_Y) = \inf\{\varepsilon \mid \gamma_X \text{ and } \gamma_Y \text{ are } (\varepsilon, T_E)\text{-interleaved}\}. \quad (2.44)$$

The following statement was given without proof in (?). We give a proof here for completeness, but also to illuminate the generalisation which will follow in section ??.

Proposition 2.28. Let $T_E = \{s \in \text{Trans}(\mathbb{R}^n, \leq) \mid s(a) = a + \varepsilon\}$, where $\varepsilon = (\varepsilon, \dots, \varepsilon)$ for some $\varepsilon \geq 0$. Then,

$$d_I^{T_E}(\gamma_X, \gamma_Y) = d_\infty(\gamma_X, \gamma_Y). \quad (2.45)$$

Proof. Given Definitions ?? and ??, we have that γ_X and γ_Y are (ε, T_E) -interleaved if, for $s \in T_E$ given by $s(a) = a + \varepsilon$, there exist morphisms $f \in \text{hom}(\gamma_X, \gamma_Y(s))$ and $g \in \text{hom}(\gamma_Y, \gamma_X(s))$ such that $f(s) \circ g = \psi_{\gamma_X}$ and $g(s) \circ f = \psi_{\gamma_Y}$. Or, in other words, γ_X and γ_Y are (ε, T_E) -interleaved if there exist continuous maps,

$$\tilde{f} : X \rightarrow Y \text{ and } \tilde{g} : Y \rightarrow X,$$

such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(\tilde{f}(x)) - \varepsilon,$$

and for all $y \in Y$,

$$\gamma_Y(y) \geq \gamma_X(\tilde{g}(y)) - \varepsilon,$$

and $\tilde{g} \circ \tilde{f} = \text{id}_X$, $\tilde{f} \circ \tilde{g} = \text{id}_Y$.

That is, γ_X and γ_Y are (ε, T_E) -interleaved if and only if there exists a homeomorphism $\tilde{f} : X \rightarrow Y$ such that $\|\gamma_X - \gamma_Y \circ \tilde{f}\|_\infty \leq \varepsilon$. \square

2.6.2 Generalising d_∞

Now that we have identified d_∞ with an interleaving distance between objects of **n-Fun**, in this section we generalise this concept to an interleaving distance between objects of **P-Fun**, where **P-Fun** is the category with objects given by functions $\gamma_X : X \rightarrow P$ for some poset P and some topological space X . The set of morphisms in this category, that is, the set $\text{hom}(\gamma_X, \gamma_Y)$ for two objects γ_X and γ_Y , consists of continuous maps $f : X \rightarrow Y$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(f(x)). \quad (2.46)$$

We generalise the interleaving distance for **P-Fun** as follows.

Definition 2.29. Let X and Y be two topological spaces, let $\gamma_X : X \rightarrow P, \gamma_Y : Y \rightarrow P$ be elements of **P-Fun**, and let $T < \text{Trans}(P)$.

The maps γ_X and γ_Y are (ε, T) -**interleaved** if there exist ε -translations $s, t \in T$ and a homeomorphism $f : X \rightarrow Y$ such that for all $x \in X, y \in Y$,

$$s(\gamma_X(x)) \geq \gamma_Y(f(x)), \text{ and } t(\gamma_Y(y)) \geq \gamma_X(f^{-1}(y)). \quad (2.47)$$

As in section ??, for $\varepsilon \geq 0, s \in \text{Trans}(P, \leq)$ is an ε -translation if $\sup_{p \in P} d_P(p, s(p)) \leq \varepsilon$.

Definition 2.30. Let X and Y be two topological spaces, and let $\gamma_X : X \rightarrow P, \gamma_Y : Y \rightarrow P \in \mathbf{P-Fun}$. Define the distance d_∞^T between γ_X and γ_Y with respect to $T < \text{Trans}(P, \leq)$ to be given by,

$$d_\infty^T(\gamma_X, \gamma_Y) = \inf\{\varepsilon \mid \gamma_X \text{ and } \gamma_Y \text{ are } (\varepsilon, T)\text{-interleaved}\}. \quad (2.48)$$

Proposition 2.31. *Definition ?? is an interleaving in the category **P-Fun**.*

Proof. We first need to define the concepts of the shifts, $\gamma_X(s)$ and $f(s)$, of maps γ_X and f , respectively, and the transition maps $\psi_{\gamma_X}(\gamma_X, \gamma_X(s))$.

For an object, γ_X , of **n-Fun**, and a translation $s \in T_E$ of the form $s(a) = a + \varepsilon$, we defined the shift, $\gamma_X(s)$, of γ_X by s to be given by $\gamma_X(s)(x) = \gamma_X(x) - \varepsilon$.

That is, we defined $\gamma_X(s)(x) = s^{-1}(\gamma_X(x))$. Note that the inverse, s^{-1} , is not a member of T_E , as it is not increasing, and therefore not a translation.

In a similar vein, given $\gamma_X \in \text{Ob}(\mathbf{P-Fun}), t \in \text{Trans}(P, \leq)$, we set $\gamma_X(t)$ to be,

$$\gamma_X(t)(x) := t^{-1}(\gamma_X(x)). \quad (2.49)$$

As with an element $s \in T_E$, this inverse t^{-1} will not be an element of $\text{Trans}(P, \leq)$. However, as we see in Definition ??, we will not require this construction, or t to be invertible, in the definition of the interleaving. This construction is merely to illustrate the parallels between Definition ?? and Definition ??.

An element $f \in \text{hom}(\gamma_X, \gamma_Y)$ is a continuous function $f : X \rightarrow Y$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(f(x)), \quad (2.50)$$

For $s \in \text{Trans}(P, \leq)$, an element of $\text{hom}(\gamma_X, \gamma_Y(s))$ is a continuous map $f(s) : X \rightarrow Y$ such that for all $x \in X$,

$$\gamma_X(x) \geq \gamma_Y(s)(f(s)(x)), \quad (2.51)$$

or equivalently, given (??), $s(\gamma_X(x)) \geq \gamma_Y(f(s)(x))$. Let $f(s) = f$. Then,

$$s(\gamma_X(x)) \geq \gamma_X(x) \geq \gamma_Y(f(x)),$$

as required, since s is a translation.

Finally, we define transition maps $\psi_{\gamma_X} : \gamma_X \rightarrow \gamma_X(s)$ for a translation $s \in \text{Trans}(P, \leq)$. An element of $\text{hom}(\gamma_X, \gamma_X(s))$ is a continuous map $\psi_{\gamma_X} : X \rightarrow X$ such that for all $x \in X$,

$$s(\gamma_X(x)) \geq \gamma_X(\psi_{\gamma_X}(x)). \quad (2.52)$$

Setting $\psi_{\gamma_X} = \text{id}_X$ for all γ_X gives (??) as required, since s is an increasing map .

Hence, with these definitions of shifts of γ_X, γ_Y and f , and of transition maps ψ , we have that $\gamma_X : X \rightarrow P$ and $\gamma_Y : Y \rightarrow P$ are (s, t) -interleaved for some $s, t \in \text{Trans}(P)$ if there exist $f \in \text{hom}(\gamma_X, \gamma_Y(s))$, $g \in \text{hom}(\gamma_Y, \gamma_X(t))$ such that,

$$f(t) \circ g = \psi_{\gamma_Y}(\gamma_Y, \gamma_Y(ts)) \text{ and } g(s) \circ f = \psi_{\gamma_X}(\gamma_X, \gamma_X(st)) = \text{id}_X. \quad (2.53)$$

Since $f(t) = f, g(t) = g$ for any $s, t \in \text{Trans}(P)$, and since $\psi_{\gamma_X} = \text{id}_X, \psi_{\gamma_Y} = \text{id}_Y$ for any $\gamma_X, \gamma_Y \in \text{Ob}(\mathbf{P-Fun})$, the condition (??) is equivalent to saying that γ_X, γ_Y are (s, t) -interleaved if there exist continuous maps $f : X \rightarrow Y, g : Y \rightarrow X$ such that,

$$f \circ g = \text{id}_Y \text{ and } g \circ f = \text{id}_X,$$

such that for all $x \in X, y \in Y$,

$$s(\gamma_X(x)) \geq \gamma_Y(f(x)), \text{ and } t(\gamma_Y(y)) \geq \gamma_X(f^{-1}(y)), \quad (2.54)$$

and so the interleaving distance obtained from these definitions of shifts and transition maps is equivalent to Definition ??.

We now generalise the stability result in (??) to arbitrary posets P , and arbitrary choices of $T < \text{Trans}(P, \leq)$.

Theorem 2.32. *Let $\gamma_X, \gamma_Y \in \mathbf{P-Fun}$, and let $\mathcal{M}, \mathcal{N} : P \rightarrow \text{Vect}_k$ be given by,*

$$\mathcal{M}(p) = H_n(\{x \in X | \gamma_X(x) \leq p\}; k) \text{ and } \mathcal{N}(p) = H_n(\{y \in Y | \gamma_Y(y) \leq p\}; k), \quad (2.55)$$

respectively, with linear maps $\psi_{\mathcal{M}}$ and $\psi_{\mathcal{N}}$ induced by the inclusions of sublevel sets. Then for any $T < \text{Trans}(P, \leq)$, the generalised interleaving distance between \mathcal{M} and \mathcal{N} satisfies,

$$d_T^T(\mathcal{M}, \mathcal{N}) \leq d_\infty^T(\gamma_X, \gamma_Y). \quad (2.56)$$

Proof. Suppose $d_\infty^T(\gamma_X, \gamma_Y) = \varepsilon$ for some $\varepsilon < \infty$, or else the result is trivially true. Then there exists a homeomorphism $\bar{f} : X \rightarrow Y$ and some ε -translations $s, t \in T$ such that for all $x \in X$,

$$s(\gamma_X(x)) \geq \gamma_Y(\bar{f}(x)), \quad (2.57)$$

and for all $y \in Y$,

$$t(\gamma_Y(y)) \geq \gamma_X(\bar{f}^{-1}(y)). \quad (2.58)$$

If $x \in X$ is such that $\gamma_X(x) \leq a$, then $y = \bar{f}(x)$ is such that,

$$\gamma_Y(y) \leq s(\gamma_X(x)) \leq s(a).$$

Hence there is an inclusion of sublevel sets,

$$\{x \in X \mid \gamma_X(x) \leq a\} \hookrightarrow \{x \in X \mid \gamma_Y(\bar{f}(x)) \leq s(a)\}. \quad (2.59)$$

which gives rise to a linear map,

$$\mathcal{M}(a) \rightarrow \mathcal{N}(s(a)), \quad (2.60)$$

for every $a \in P$.

Similarly, if $y \in Y$ is such that $\gamma_Y(y) \leq b$, then $x = \bar{f}^{-1}(y)$ is such that,

$$\gamma_X(\bar{f}^{-1}(y)) \leq t(\gamma_Y(y)) \leq t(b).$$

Hence there is an inclusion of sublevel sets,

$$\{y \in Y \mid \gamma_Y(y) \leq b\} \hookrightarrow \{y \in Y \mid \gamma_X(\bar{f}^{-1}(y)) \leq t(b)\}, \quad (2.61)$$

which gives rise to a linear map,

$$\mathcal{N}(b) \rightarrow \mathcal{M}(t(b)), \quad (2.62)$$

for every $b \in P$.

The composition,

$$\mathcal{M}(a) \rightarrow \mathcal{N}(s(a)) \rightarrow \mathcal{M}(ts(a)), \quad (2.63)$$

is equal to the transition map $\psi_M(a, ts(a))$ for each $a \in P$, since all maps are induced by the inclusions of sub-level sets. Similarly, the composition,

$$\mathcal{N}(a) \rightarrow \mathcal{M}(t(a)) \rightarrow \mathcal{N}(st(a)), \quad (2.64)$$

is equal to the transition map $\psi_N(a, st(a))$ for any $a \in P$. Hence, we have that if γ_X and γ_Y are (s, t) -interleaved, then \mathcal{M} and \mathcal{N} are (s, t) -interleaved. □

We have therefore shown that we have a stability result for sublevel-set persistence modules indexed by arbitrary posets P , which demonstrates precisely how a relationship between P -valued functions extends to a relationship between their persistence modules, with no assumptions about the poset.

2.6.3 Universality of the Interleaving Distance

In section ??, we stated a number of stability results which hold for the interleaving distance. In particular, let us recall Theorem ??, proved by Lesnick in (?). This result stated that for $\gamma_X, \gamma_Y \in \mathbf{n-Fun}$, if \mathcal{M} and \mathcal{N} are the persistence modules corresponding to sublevel-set filtrations by γ_X and γ_Y , respectively, then d_I^{TE} satisfies,

$$d_I^{TE}(\mathcal{M}, \mathcal{N}) \leq d_\infty(\gamma_X, \gamma_Y), \quad (2.65)$$

where $d_\infty(\gamma_X, \gamma_Y) = \inf_{h \in \text{Homeo}(X, Y)} \|f - g \circ h\|_\infty$.

As well as this stability result, Lesnick also proves (?) that d_I^{TE} is the universal distance for \mathbb{R}^n -persistence modules with respect to d_∞ .

Theorem 2.33. *Let $\gamma_X, \gamma_Y, \mathcal{M}$ and \mathcal{N} be as above. If d is another pseudometric on the space of \mathbb{R}^n -persistence modules, which satisfies,*

$$d(\mathcal{M}, \mathcal{N}) \leq d_\infty(\gamma_X, \gamma_Y),$$

then $d \leq d_I^{TE}$.

This means that for all distances on the space of \mathbb{R}^n -persistence modules which are stable with respect to d_∞ , then d_I^{TE} is the most discerning distance. In the same paper (?), Lesnick acknowledges the generalised interleaving distance of (?), and wonders whether a similar universality result holds for the generalised d_I . It is clear that if the generalised interleaving distance satisfies a universality condition, then it is with respect to some distance other than d_∞ . We conjecture that the appropriate analogue of d_∞ for the generalised interleaving distance is precisely our d_∞^T of Definition ??.

Conjecture 2.34. *Let $\gamma_X, \gamma_Y \in \mathbf{P-Fun}$. Let \mathcal{M}, \mathcal{N} be the P -persistence modules corresponding to the sublevel-set filtrations by γ_X and γ_Y , respectively. If d is another distance on the space of P -persistence modules which satisfies,*

$$d(\mathcal{M}, \mathcal{N}) \leq d_\infty^T(\gamma_X, \gamma_Y),$$

then $d \leq d_I^T$.

Universality of d_I^{TE} for \mathbb{R}^n -persistence modules was proved by Lesnick (?) by showing that for any \mathbb{R}^n -persistence modules \mathcal{M} and \mathcal{N} which are ϵ -interleaved, there exists a pair of topological spaces X and Y , and a pair of continuous functions $f : X \rightarrow \mathbb{R}^n$, $g : Y \rightarrow \mathbb{R}^n$ such that the persistent homology of the sublevel set filtrations by f and g are isomorphic to \mathcal{M} and \mathcal{N} . It is our conjecture that a similar method could be used to prove conjecture ??. This is a source of further work.

3 Poset Morphisms and the Interleaving Distance

We have seen that persistent homology can be described as a functor from a partially ordered set, P , to a category, \mathcal{C} , where \mathcal{C} is commonly the category of vector spaces and linear maps, or of topological spaces and continuous maps. In (?), the authors showed that composing this process with an additional functor, $F : \mathcal{C} \rightarrow \mathcal{D}$, is 1-Lipschitz with respect to the interleaving distance. That is, for any $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$,

$$d_I(F \circ \mathcal{M}, F \circ \mathcal{N}) \leq d_I(\mathcal{M}, \mathcal{N}). \quad (3.1)$$

This suggests that the interleaving distance does not depend on the choice of target category. We have seen this inequality implicitly in action in the previous section – if a pair of filtrations $f : P \rightarrow \text{Subsets}(X)$, and $g : P \rightarrow \text{Subsets}(Y)$ are (s, t) -interleaved for some pair of translations $s, t \in \text{Trans}(P, \leq)$, then the persistence modules $H \circ f$ and $H \circ g$, where H is a homology functor, are also (s, t) -interleaved.

Bubenik, de Silva and Scott (?) use the following diagram to describe the persistence process.

$$\text{poset } Q \xrightarrow{\varphi} \text{poset } P \xrightarrow{\mathcal{M}} \text{category } \mathcal{C} \xrightarrow{F} \text{category } \mathcal{D} \quad (3.2)$$

Their result (?), outlined in (??), describes the effect of composing with the functor F . In this section, it is the arrow on the left of their diagram, the poset morphism $\varphi : Q \rightarrow P$, which will concern us. Specifically, we will consider the effect of precomposing with a morphism $\varphi : Q \rightarrow P$ of partially ordered sets on the interleaving distance. We saw in the previous section that interleavings are defined with respect to the structure of the poset, since they depend on the set of translations, $\text{Trans}(P, \leq)$, and the weights of these translations. We will consider how we can bound the interleaving distance between pull-back modules with respect to the distance between the original modules.

There are many reasons for which we may be interested in considering the pull-backs and push-forwards along poset morphisms, and order isomorphisms. We can use a pullback along an inclusion $\varphi : \mathbb{Z}^n \rightarrow \mathbb{R}^n$ in order to discretise a continuous persistence module. If we are to perform any calculations with persistence modules, then discrete structures are easier to store. See (?) for an in-depth discussion about the merits of discretising persistence modules for applications, which also links discretisations to the concept of tameness. Discretising in this way, however, should be done with caution. Given that interleavings depend on the poset structure, pulling back along a poset morphism may distort distances. Hence the distance between two discretised modules may not be the same as the distance between the original ones. In section ??, we show that, outside of some special cases, reparameterising does not preserve distances between modules, and in fact this action can be unstable.

Another reason for considering pull-backs and push-forwards along poset morphisms is the possibility that we may have a persistence module indexed by a poset P , in which it is impossible to calculate the interleaving distance. We have seen that interleavings, and

calculating the interleaving distance, require the existence of non-trivial translations in P . The work of (?) discusses a way to proceed in the case that $Trans(P, \leq)$ is the trivial monoid. One example, seen in (?), of such a poset is a ‘zig-zag’ poset - one with a Hasse diagram resembling that of Figure ??.




Figure 22: A ‘zig-zag’ poset which has no non-trivial translations.

This is the motivation for the relative interleaving distance introduced in (?). If we have a morphism $\varphi : P \rightarrow Q$, where Q is a poset with non-trivial $Trans(Q, \leq)$, then even if we cannot calculate the interleaving distance in P , we can push forward a pair of P modules along φ , and consider interleavings in Q . The relative interleaving distance is then a measure of the distance between the P -persistence modules relative to the morphism $\varphi : P \rightarrow Q$.

There is one more motivation for considering the reparameterisations of persistence modules by poset morphisms, and that is the idea of rescaling modules. Scale is such a central part of persistent homology, and yet the way it is treated is rather rigid. An ε -interleaving described modules whose features arise at scales which differ by no more than ε , and sets their distance to be ε accordingly. However, this is not ideal if we wish to look for relationships between data sets which are identical – or suitably close – up to some change of scale.

By considering pull-backs and push-forwards of persistence modules along order isomorphisms, we can describe rescalings of persistence modules. We will then describe how bounding the interleaving distance between pull-back modules allows us to describe interleavings of rescaled modules.

In this section, we explore the question of when the operation of pushing forward or pulling back a pair of persistence modules along a poset morphism is an isometry with respect to the interleaving distance, and how we can use properties of the morphism to bound the distance between a pair of push-forward or pull-back modules. This section could therefore be said to be an investigation into the stability of the interleaving distance with respect to morphisms of posets. In particular, we show that Lipschitz properties of an order isomorphism $\varphi : P \rightarrow Q$ descend to identical properties with respect to the interleaving distance between the push-forwards and pull-backs of modules along φ .

In (?), it is stated that pushing forward along any poset morphism is always an isometry with respect to the interleaving distance. This work is done in the context of situations in which it is impossible to calculate the interleaving distance in a particular poset P , but it is possible in another poset Q for which there exists a morphism $\varphi : P \rightarrow Q$. In this section, we’ll show that the statement that “the push-forward operation ... is an isometry onto its image” (?) should be taken with care. In situations where it is possible to calculate

interleavings in either P or in Q , we show that the two concepts – and the distances obtained from both – are not equal in general. We will also show that if φ is not an order isomorphism, any Lipschitz properties of φ are in general not preserved by a Kan extension along φ . This shows that reparameterising persistence modules, or discretising modules, can be a source of instability with respect to the interleaving distance.

3.1 Poset Morphisms

Recall that a poset (P, \leq) can be viewed as a category, with an object for each element of the set and a morphism between a pair of objects p_1, p_2 whenever $p_1 \leq p_2$. Thinking of posets in this way, a morphism between two posets is really just a functor between categories.

Definition 3.1. Let (P, \leq_P) and (Q, \leq_Q) be partially ordered sets. A **poset morphism** $\varphi : (P, \leq_P) \rightarrow (Q, \leq_Q)$ is a map $\varphi : P \rightarrow Q$ such that for all $p_1, p_2 \in P$, if $p_1 \leq_P p_2$ then $\varphi(p_1) \leq_Q \varphi(p_2)$.

Such maps are sometimes called order-preserving or monotone.

Example 3.2. The map $\varphi : \mathbb{R} \rightarrow \mathbb{R}^n$ given by $\varphi(a) = \mathbf{a}$, where \mathbf{a} is the constant vector $\mathbf{a} = (a, \dots, a)$, is a morphism of posets $(\mathbb{R}, \leq) \rightarrow (\mathbb{R}^n, \leq)$, where the partial order on \mathbb{R}^n is given by

$$(a_1, \dots, a_n) \leq (b_1, \dots, b_n) \text{ if and only if } a_i \leq b_i \text{ for all } 1 \leq i \leq n. \quad (3.3)$$

Example 3.3. Let $\varphi : \mathbb{Z} \rightarrow \mathbb{R}$ be any monotone function. Then φ is order-preserving.

Example 3.4. Let $\varphi_c : \mathbb{R} \rightarrow \mathbb{Z}$ be given by $\varphi_c(a) = \lceil a \rceil$ where $\lceil a \rceil$ denotes the **ceiling** of a ,

$$\lceil a \rceil = \min\{z \in \mathbb{Z} \mid z \geq a\}. \quad (3.4)$$

Then φ_c is an order-preserving map, as is the map $\varphi_f : \mathbb{R} \rightarrow \mathbb{Z}$ given by $\varphi_f(a) = \lfloor a \rfloor$, where $\lfloor a \rfloor$ denotes the **floor** of a ,

$$\lfloor a \rfloor = \max\{z \in \mathbb{Z} \mid z \leq a\}. \quad (3.5)$$

Definition 3.5. An **order isomorphism** $\varphi : P \rightarrow Q$ is a surjective **order-embedding**. That is, an order isomorphism is a surjective function such that for all $p_1, p_2 \in P$, $\varphi(p_1) \leq \varphi(p_2)$ if and only if $p_1 \leq p_2$. We write $P \cong Q$ if there exists an order isomorphism of posets $f : P \rightarrow Q$.

We do not have to check that φ is injective in order to conclude that φ is an isomorphism since if $\varphi : P \rightarrow Q$ is a surjective order-embedding, then $\varphi(p_1) = \varphi(p_2)$ if and only if both $\varphi(p_1) \geq \varphi(p_2)$ and $\varphi(p_1) \leq \varphi(p_2)$, which, as φ is an order-embedding, means that $p_1 \geq p_2$ and $p_1 \leq p_2$. That is, $\varphi(p_1) = \varphi(p_2)$ if and only if $p_1 = p_2$.

Example 3.6. Let P and Q be subsets of the poset (\mathbb{R}, \leq) given by $P = [0, n]$, $Q = [0, m \cdot n]$ for some $m > 0$. Then the map $\varphi : P \rightarrow Q$,

$$\varphi(a) = m \cdot a, \quad (3.6)$$

is a surjective order-embedding, so is an order isomorphism between P and Q .

Example 3.7. As in Example ??, let P be the subset of (\mathbb{R}, \leq) given by $P = [0, n]$, and let Q be the subset of (\mathbb{R}, \leq) given by $Q = [m, n + m]$. Then the map

$$\varphi(a) = m + a, \quad (3.7)$$

is a surjective order-embedding, so is an order isomorphism between P and Q .

3.2 Pull-back Modules

Given any categories \mathcal{A}, \mathcal{B} and \mathcal{C} , with functors $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ and $F : \mathcal{B} \rightarrow \mathcal{C}$, there is a functor, $\varphi^*F : \mathcal{A} \rightarrow \mathcal{C}$ given by $\varphi^*F(a) = F(\varphi(a))$. When \mathcal{A} and \mathcal{B} are posets, and F is a persistence module, this gives us a way of pulling back the domain of F to \mathcal{A} .

Definition 3.8. Given any poset morphism $\varphi : P \rightarrow Q$, and any Q -persistence module $\mathcal{M} : Q \rightarrow \mathcal{C}$, the **pull-back of \mathcal{M} along φ** is the P -persistence module, $\varphi^*\mathcal{M} : P \rightarrow \mathcal{C}$ with,

$$\varphi^*\mathcal{M}(p) = \mathcal{M}(\varphi(p)), \quad (3.8)$$

for any $p \in P$, and with linear maps $\psi_{\varphi^*\mathcal{M}(p_1, p_2)} : \varphi^*\mathcal{M}(p_1) \rightarrow \varphi^*\mathcal{M}(p_2)$ given by,

$$\psi_{\varphi^*\mathcal{M}(p_1, p_2)} = \psi_{\mathcal{M}(\varphi(p_1), \varphi(p_2))} : \mathcal{M}(\varphi(p_1)) \rightarrow \mathcal{M}(\varphi(p_2)), \quad (3.9)$$

for every $p_1 \leq p_2$.

We note that the map $\psi_{\varphi^*\mathcal{M}(p_1, p_2)}$ exists whenever the map $\psi_{\mathcal{M}(p_1, p_2)}$ exists, since $\varphi(p_1) \leq \varphi(p_2)$ whenever $p_1 \leq p_2$.

We name $\varphi^*\mathcal{M}$ a pull-back module to show that we pull back the domain of $\varphi : Q \rightarrow \mathcal{C}$ to P along φ , yet this should not be confused with the commonly used term of a pull-back in category theory.

Example 3.9. Consider the map $\varphi : \mathbb{Z} \rightarrow \mathbb{R}$ in Example ??, given by $\varphi(a) = a$. If \mathcal{M} is any \mathbb{R} -module, then the pull-back $\varphi^*\mathcal{M}$ is a \mathbb{Z} -module, which is the discretisation of \mathcal{M} .

Example 3.10. Let $\delta > 0$, and let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be the order automorphism given by

$$\varphi(a) = \delta \cdot a. \quad (3.10)$$

Then the pull-back, $\varphi^*\mathcal{M}$, of an \mathbb{R} -persistence module $\mathcal{M} : \mathbb{R} \rightarrow \mathcal{C}$ along φ is rescaling of \mathcal{M} by a factor of δ .

For example, let \mathcal{M} be the interval module, $\mathcal{M} = k[0, n]$, for some $n > 0$. Then

$$\mathcal{M}(a) = \begin{cases} k & a \in [0, n] \\ 0 & \text{otherwise,} \end{cases}$$

with transition maps $\psi_{\mathcal{M}}$ given by,

$$\psi_{\mathcal{M}}(a, b) = \begin{cases} id_k & (a, b) \subseteq [0, n] \\ 0 & \text{otherwise.} \end{cases}$$

Then $\varphi^*\mathcal{M}(a) = \mathcal{M}(\varphi(a))$, and so,

$$\varphi^*\mathcal{M}(a) = \begin{cases} k & a \in [0, \frac{n}{\delta}] \\ 0 & \text{otherwise,} \end{cases}$$

with transition maps $\psi_{\mathcal{M}}$ given by,

$$\psi_{\mathcal{M}}(a, b) = \psi_{\mathcal{M}}(\varphi(a), \varphi(b)) = \begin{cases} id_k & (a, b) \subseteq [0, \frac{n}{\delta}] \\ 0 & \text{otherwise.} \end{cases}$$

Hence $\varphi^*\mathcal{M}$ is the interval module $k[0, \frac{n}{\delta}]$. We can think of $\varphi^*\mathcal{M}$ as a shrinking of \mathcal{M} by a factor of δ .

In a similar vein, as well as rescaling modules we can also translate modules, as the following example shows.

Example 3.11. Let $\mathcal{M} : \mathbb{R} \rightarrow Vect_k$ be the same interval module as in Example ??, $\mathcal{M} = k[0, n]$, and let $\delta \in \mathbb{R}$. This time, let φ be the order automorphism of \mathbb{R} given by,

$$\varphi(a) = a + \delta. \tag{3.11}$$

Then $\varphi^*\mathcal{M}$ is a translation of the module \mathcal{M} by δ .

We have,

$$\varphi^*\mathcal{M}(a) = \mathcal{M}(\varphi(a)) = \begin{cases} k & a \in [-\delta, n - \delta] \\ 0 & \text{otherwise,} \end{cases} \text{ with } \psi_{\varphi^*\mathcal{M}}(a, b) = \begin{cases} id_k & (a, b) \subseteq [-\delta, n - \delta] \\ 0 & \text{otherwise.} \end{cases}$$

Hence the pull-back module $\varphi^*\mathcal{M}$ is given by $\varphi^*\mathcal{M} = k[-\delta, n - \delta]$.

Let's look at some specific examples of how pull-backs can be used to describe rescalings in practice.

Example 3.12. Consider the following point clouds, which we can think of as two separate finite metric spaces, which are both subsets of \mathbb{R}^2 . The left hand set, X , is a rescaling of the right-hand set, Y , which can also be stated as saying that we have a bijection $f : X \rightarrow Y$, and some $m \geq 0$ such that for any $x_i, x_j \in X$, we have

$$d(f(x_i), f(x_j)) = md(x_i, x_j). \tag{3.12}$$




Figure 23: Two point clouds \mathbb{X} (left) and \mathbb{Y} (right), which are topologically identical up to a change of scale.

Let $\mathcal{M} : \mathbb{R} \rightarrow Vect_k$ be the persistence module given by,

$$\mathcal{M}(a) = H_n(VR_X(a); k). \quad (3.13)$$

That is, \mathcal{M} is the persistence module where each vector space, $\mathcal{M}(a)$ is given by the homology in degree $n \geq 0$ of the Vietoris-Rips complex on X at scale a . The transition maps, $\psi_{\mathcal{M}}(a, b)$, are the linear maps in homology induced by the inclusions of Vietoris-Rips complexes $VR_X(a) \hookrightarrow VR_X(b)$. Similarly, let $\mathcal{N} : \mathbb{R} \rightarrow Vect_k$ be the persistence module with,

$$\mathcal{N}(a) = H_n(VR_Y(a); k), \quad (3.14)$$

and with transition maps $\psi_{\mathcal{N}}$ induced by inclusions of Vietoris-Rips complexes. Consider the poset morphism $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ given by $\varphi(a) = ma$. Then (??) tells us that

$$\mathcal{M} \cong \varphi^* \mathcal{N}, \quad (3.15)$$

since for every $a \in \mathbb{R}$, we have an isomorphism, $\mathcal{M}(a) \cong \varphi^* \mathcal{N}(a) = \mathcal{N}(ma)$.

Example 3.13. Let X be a torus. Let P_1 be the plane with equation $z = z_1$. Let $f : X \rightarrow \mathbb{R}$ measure the height of a point $x \in X$ from the horizontal plane P_1 as shown in Figure ??.

The sublevel set,

$$\mathcal{F}(a) = \{x \in X \mid f(x) \leq a\},$$

consists of all points in the torus which are at height at most a above P_1 . We have marked the points a_1, a_2, a_3 and a_4 as it is at these points that the topology of the sublevel sets changes. For $a \leq a_1$, the sublevel set $\mathcal{F}(a)$ is empty. For $a_1 < a < a_2$, the sublevel set $\mathcal{F}(a)$ is a disk. For $a_2 < a < a_3$, it is homotopic to a cylinder. For $a_3 < a < a_4$ it is homotopic to a punctured torus, and for any $a > a_4$, the sublevel set $\mathcal{F}(a)$ is the whole torus. We can see this filtration by different values of a in Figure ??.

Let P_2 be another horizontal plane below P_1 , with equation $z = z_2$. Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be given by $\varphi(a) = a + (z_1 - z_2)$.




Figure 24: The function f measures the height of a point of the torus above the plane P_1 .

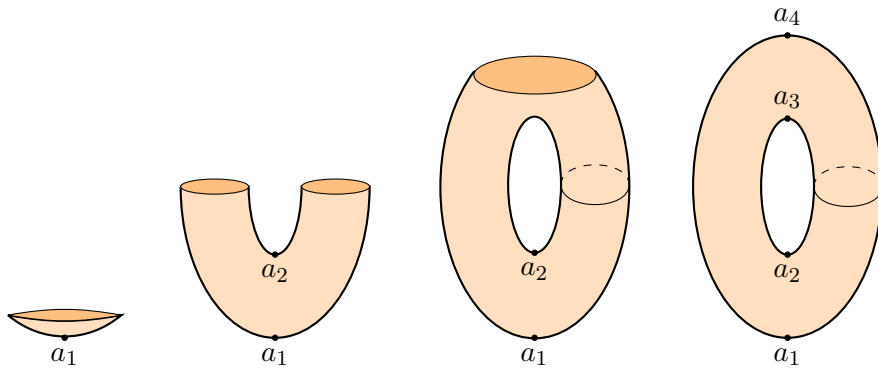


Figure 25: A height filtration, \mathcal{F} , of a torus.

Then the pull-back of \mathcal{F} along φ corresponds to a height filtration of the torus, X , by $f_2 : X \rightarrow \mathbb{R}$, where f_2 measures the height of a point $x \in X$ above the plane P_2 . If P_i is any other plane below P_1 with equation $z = z_i$, then the height filtration corresponding to the plane P_i is related to \mathcal{F} via a pull-back along the order isomorphism $\varphi_i : \mathbb{R} \rightarrow \mathbb{R}$ given by $\varphi_i(a) = a + (z_1 - z_i)$.




Figure 26: Various choices of plane P_i lead to different – but related – filtrations.

Remark 3.14. For any different choice of plane P_i below P_1 , we can obtain the height filtration relative to P_i as a pull-back of \mathcal{F} . However, a plane which lies above the torus, or between any a_i and a_i , gives a filtration which cannot be obtained as a pull-back of \mathcal{F} . This simple observation leads to the possibility of considering equivalence classes of filtrations under the relation,

$$\mathcal{F}_1 \sim_{\varphi} \mathcal{F}_2, \tag{3.16}$$

if \mathcal{F}_1 can be obtained from \mathcal{F}_2 via a pull-back along the order isomorphism φ . This is a source of potential further work.

We can consider other posets and other order-automorphisms.

Example 3.15. Consider the action of the permutation group S_n on \mathbb{R}^n . That is, for $\sigma \in S_n$, let $\varphi_{\sigma} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be given by,

$$\varphi_{\sigma}(a_1, \dots, a_n) = (a_{\sigma(1)}, \dots, a_{\sigma(n)}).$$

For instance, consider the two-parameter filtration, \mathcal{F} , seen in Figure ??.




Figure 27: A symmetric bifiltration, \mathcal{F} .

Let $\varphi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the order automorphism given by,

$$\varphi(a_1, a_2) = (a_2, a_1). \tag{3.17}$$

The pull-back of \mathcal{F} along φ amounts to a reflection in the line $\gamma_1 = \gamma_2$. We may see that in this case, the reflection $\varphi^*\mathcal{F}$ is identical to the original filtration \mathcal{F} . This reflects the fact that births of simplices occur along the diagonal, $\gamma_1 = \gamma_2$.

3.3 Pull-backs and the Interleaving Distance

In this section we will consider the question of when the pull-back operation along a poset morphism $\varphi : P \rightarrow Q$ is an isometry with respect to the interleaving distance. Let $\mathcal{M} : Q \rightarrow \mathcal{C}$ and $\mathcal{N} : Q \rightarrow \mathcal{C}$ be a pair of Q -persistence modules. We will consider the cases $d_I(\mathcal{M}, \mathcal{N}) = 0$ and $d_I(\mathcal{M}, \mathcal{N}) > 0$ separately. Within this section, d_I will be used to denote the interleaving distance, $d_I^{Trans(P)}$.

Lemma 3.16. *If $\varphi : P \rightarrow Q$ is order-preserving, then the map $\varphi^{-1} : Q \rightarrow P$ is also order-preserving. That is, for any $p_1, p_2 \in P$, if $\varphi(p_1) \leq \varphi(p_2)$ whenever $p_1 \leq p_2$, then for any $q_1, q_2 \in Q$, $\varphi^{-1}(q_1) \leq \varphi^{-1}(q_2)$ whenever $q_1 \leq q_2$.*

Proof. Since φ is an order isomorphism, it is also an order embedding. Hence for any $p_1, p_2 \in P$, if $\varphi(p_1) \leq \varphi(p_2)$ then $p_1 \leq p_2$, if and only if $\varphi^{-1}\varphi(p_1) \leq \varphi^{-1}\varphi(p_2)$. Since φ is surjective, then every $q_1, q_2 \in Q$ is given by $q_1 = \varphi(p_1), q_2 = \varphi(p_2)$ for some $p_1, p_2 \in P$. Hence if $q_1 \leq q_2$ then $\varphi^{-1}(q_1) \leq \varphi^{-1}(q_2)$. \square

Lemma 3.17. *The pull-back of Q -persistence modules along an order isomorphism $\varphi : P \rightarrow Q$ preserves isomorphism classes of persistence modules.*

Proof. Suppose $\mathcal{M} \cong \mathcal{N}$. Then for every $q \in Q$, we have morphisms $f(q) : \mathcal{M}(q) \rightarrow \mathcal{N}(q)$, $g(q) : \mathcal{N}(q) \rightarrow \mathcal{M}(q)$ such that $g(q) \circ f(q) = id_{\mathcal{M}(q)}$ and $f(q) \circ g(q) = id_{\mathcal{N}(q)}$. From now on, we will simply refer to these morphisms $f(q) : \mathcal{M}(q) \rightarrow \mathcal{N}(q)$ and $g(q) : \mathcal{N}(q) \rightarrow \mathcal{M}(q)$ as f and g , regardless of $q \in Q$.

In particular, for every $p \in P$, there exist maps,

$$f : \mathcal{M}(\varphi(p)) \rightarrow \mathcal{N}(\varphi(p)), \quad g : \mathcal{N}(\varphi(p)) \rightarrow \mathcal{M}(\varphi(p)),$$

or, in other terms, maps,

$$f : \varphi^*\mathcal{M}(p) \rightarrow \varphi^*\mathcal{N}(p), \quad g : \varphi^*\mathcal{N}(p) \rightarrow \varphi^*\mathcal{M}(p),$$

such that $f \circ g = id_{\varphi^*\mathcal{N}(a)}$, and $g \circ f = id_{\varphi^*\mathcal{M}(a)}$. Hence $\varphi^*\mathcal{M} \cong \varphi^*\mathcal{N}$. Hence if $d_I(\mathcal{M}, \mathcal{N}) = 0$, then the pull-back along any poset morphism induces an isometry with respect to the interleaving distance.

Conversely, as φ is surjective, then every $q \in Q$ is equal to $\varphi(p)$ for some $p \in P$. Hence if there exists an isomorphism $\mathcal{M}(\varphi(p)) \rightarrow \mathcal{N}(\varphi(p))$ for each $p \in P$, then there exists an isomorphism $\mathcal{M}(q) \rightarrow \mathcal{N}(q)$ for every $q \in Q$. Hence if φ is surjective, then $d_I(\mathcal{M}, \mathcal{N}) = 0$ if and only if $d_I(\varphi^*\mathcal{M}, \varphi^*\mathcal{N}) = 0$. \square

We will now consider when $d_I(\mathcal{M}, \mathcal{N}) = d_I(\varphi^*\mathcal{M}, \varphi^*\mathcal{N})$, supposing that $d_I(\mathcal{M}, \mathcal{N}) > 0$.

Lemma 3.18. *For any order isomorphism $\varphi : P \rightarrow Q$, and any $T < Trans(P)$, the set T' given by,*

$$T' = \{\varphi t \varphi^{-1} | t \in T\},$$

is a submonoid of $Trans(Q)$.

Proof. Clearly, $\varphi t \varphi^{-1}$ is order-preserving, since it is the composition of three order-preserving morphisms, and by Lemma ??, the composition of order-preserving morphisms is order-preserving.

Secondly, for any $q \in Q$, we have that $t\varphi^{-1}(q) \geq \varphi^{-1}(q)$ as $t \in Trans(P)$. As φ is order-preserving, we then have that $\varphi t \varphi^{-1}(q) \geq \varphi \varphi^{-1}(q) = q$. Hence T' is a subset of $Trans(Q)$.

T' is closed under composition, since for any $s, t \in T$,

$$(\varphi s \varphi^{-1})(\varphi t \varphi^{-1}) = \varphi s t \varphi^{-1},$$

which is an element of T' since T is a monoid and hence is closed under composition.

Finally, the identity translation id_P is an element of T since T is a monoid, and so $\varphi id_P \varphi^{-1} = id_Q$ is an element of T' . \square

Proposition 3.19. *For any order isomorphism $\varphi : P \rightarrow Q$, any two Q -persistence modules \mathcal{M} and \mathcal{N} are (s, t) -interleaved if and only if the P -persistence modules $\varphi^* \mathcal{M}$ and $\varphi^* \mathcal{N}$ are (s', t') -interleaved, where,*

$$s' = \varphi^{-1} s \varphi \quad \text{and} \quad t' = \varphi^{-1} t \varphi. \quad (3.18)$$

Proof. If \mathcal{M} and \mathcal{N} are (s, t) -interleaved for some $s, t \in \text{Trans}(Q)$, then there exist morphisms of persistence modules $f : \mathcal{M} \rightarrow \mathcal{N}(s)$ and $g : \mathcal{N} \rightarrow \mathcal{M}(t)$. Therefore, for each $q_1 \leq q_2 \in Q$, there exists a pair of commutative diagrams,

$$\begin{array}{ccc} \mathcal{M}(q_1) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(q_2) \\ \downarrow & & \downarrow \\ \mathcal{N}(s(q_1)) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(s(q_2)) \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{M}(t(q_1)) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(t(q_2)) \\ \uparrow & & \uparrow \\ \mathcal{N}(q_1) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(q_2) \end{array} \quad (3.19)$$

and we have that for every $q \in Q$, $g(s(q)) \circ f(q) = \psi_{\mathcal{M}}(q, ts(q))$ and $f(t(q)) \circ g(q) = \psi_{\mathcal{N}}(q, st(q))$. Hence for every $q \in Q$, there exists a pair of commutative diagrams,

$$\begin{array}{ccc} \mathcal{M}(q) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(ts(q)) \\ & \searrow & \nearrow \\ & \mathcal{N}(s(q)) & \end{array} \quad \begin{array}{ccc} & \mathcal{M}(t(q)) & \\ \nearrow & & \searrow \\ \mathcal{N}(q) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(st(q)) \end{array} \quad (3.20)$$

In particular, for each $p_1 \leq p_2$ in P , we have a pair of commutative diagrams,

$$\begin{array}{ccc} \mathcal{M}(\varphi(p_1)) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(\varphi(p_2)) \\ \downarrow & & \downarrow \\ \mathcal{N}(s\varphi(p_1)) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(s\varphi(p_2)) \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathcal{M}(t\varphi(p_1)) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(t\varphi(p_2)) \\ \uparrow & & \uparrow \\ \mathcal{N}(\varphi(p_1)) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(\varphi(p_2)) \end{array} \quad (3.21)$$

and for each $p_i \in P$, there exists a pair of commutative diagrams,

$$\begin{array}{ccc} \mathcal{M}(\varphi(p_i)) & \xrightarrow{\psi_{\mathcal{M}}} & \mathcal{M}(ts\varphi(p_i)) \\ & \searrow & \nearrow \\ & \mathcal{N}(s\varphi(p_i)) & \end{array} \quad \begin{array}{ccc} & \mathcal{M}(t\varphi(p_i)) & \\ \nearrow & & \searrow \\ \mathcal{N}(\varphi(p_i)) & \xrightarrow{\psi_{\mathcal{N}}} & \mathcal{N}(st\varphi(p_i)) \end{array} \quad (3.22)$$

Since $\mathcal{M}(\varphi(p_i)) = \varphi^*\mathcal{M}(p_i)$, we can rewrite the diagrams in (??) as,

$$\begin{array}{ccc}
\varphi^*\mathcal{M}(p_1) & \xrightarrow{\psi_{\varphi^*\mathcal{M}}} & \varphi^*\mathcal{M}(p_2) \\
\downarrow & & \downarrow \\
\varphi^*\mathcal{N}(\varphi^{-1}s\varphi(p_1)) & \xrightarrow{\psi_{\varphi^*\mathcal{N}}} & \varphi^*\mathcal{N}(\varphi^{-1}s\varphi(p_2))
\end{array}
\quad \text{and} \quad
\begin{array}{ccc}
\varphi^*\mathcal{M}(\varphi^{-1}t\varphi(p_1)) & \xrightarrow{\psi_{\varphi^*\mathcal{M}}} & \varphi^*\mathcal{M}(\varphi^{-1}t\varphi(p_2)) \\
\uparrow & & \uparrow \\
\varphi^*\mathcal{N}(p_1) & \xrightarrow{\psi_{\varphi^*\mathcal{N}}} & \varphi^*\mathcal{N}(p_2)
\end{array}
\tag{3.23}$$

and we can rewrite the diagrams in ?? as,

$$\begin{array}{ccc}
\varphi^*\mathcal{M}(p) & \xrightarrow{\psi_{\varphi^*\mathcal{M}}} & \varphi^*\mathcal{M}(\varphi^{-1}t\varphi(p)) \\
\searrow & & \nearrow \\
& \varphi^*\mathcal{N}(\varphi^{-1}s\varphi(p)) &
\end{array}
\quad
\begin{array}{ccc}
& \varphi^*\mathcal{M}(\varphi^{-1}t\varphi(p)) & \\
\nearrow & \xrightarrow{\psi_{\varphi^*\mathcal{N}}} & \searrow \\
\varphi^*\mathcal{N}(p) & & \varphi^*\mathcal{N}(\varphi^{-1}st\varphi(p))
\end{array}
\tag{3.24}$$

We note that $\varphi^{-1}st\varphi = (\varphi^{-1}s\varphi)(\varphi^{-1}t\varphi) = s't'$, and so the existence of the above diagrams means that, by definition, $\varphi^*\mathcal{M}$ and $\varphi^*\mathcal{N}$ are (s', t') -interleaved. The converse statement follows from the surjectivity of φ . \square

Remark 3.20. We could consider this proof in a more categorical way. Namely, if we consider the functor categories $Func(P, \mathcal{C})$ and $Func(Q, \mathcal{C})$, whose objects are respectively P - and Q -persistence modules, and whose morphisms are natural transformations between functors, then an isomorphism $\varphi : P \rightarrow Q$ gives rise to an induced isomorphism of Functor categories. Hence, for every pair of P -persistence modules $\mathcal{M}, \mathcal{N} \in Func(P, \mathcal{C})$, we have a corresponding pair of persistence modules $\varphi^*\mathcal{M}, \varphi^*\mathcal{N} \in Func(Q, \mathcal{C})$. Likewise, for $s \in Trans(P)$, for every morphism $\mathcal{M} \rightarrow \mathcal{N}(s)$, we have a corresponding morphism $\varphi^*\mathcal{M} \rightarrow \varphi^*\mathcal{N}(\varphi^{-1}s\varphi)$, and for $t \in Trans(P)$, for every morphism $\mathcal{N} \rightarrow \mathcal{M}(t)$, we have a corresponding morphism $\varphi^*\mathcal{N} \rightarrow \varphi^*\mathcal{M}(\varphi^{-1}t\varphi)$, which must commute as needed. Hence for any (s, t) -interleaving in $Func(P, \mathcal{C})$, we have a $(\varphi^{-1}s\varphi, \varphi^{-1}t\varphi)$ -interleaving between the corresponding objects of $Func(Q, \mathcal{C})$.

We may also interpret this result in a categorical way - we have discussed that an isomorphism of categories $P \rightarrow Q$ gives rise to an induced isomorphism between the Functor categories, which preserves natural transformations between objects of the functor categories. Interleavings have often been described as ‘‘approximate isomorphisms’’ (?). In Proposition ??, we see that the isomorphism $P \rightarrow Q$ gives rise to a one-to-one correspondence between these approximate isomorphisms of objects of the functor categories.

In the statement of the following theorem d_P and d_Q are the pseudometrics defined in Section 2.4, which are defined with respect to chosen weight functions, $w_P : P \times P \rightarrow [0, \infty]$ and $w_Q : Q \times Q \rightarrow [0, \infty]$, respectively. The interleaving distances between P - and Q -persistence modules, meanwhile, are $d_I^{Trans(P)}$ and $d_I^{Trans(Q)}$, respectively.

Theorem 3.21. *If $\varphi : P \rightarrow Q$ is an order isomorphism such that φ^{-1} is Lipschitz with constant $C \geq 0$, then the pull-back of Q -persistence modules along φ is also a Lipschitz map*

with constant C . That is, if φ is such that for any $q_1, q_2 \in Q$, there exists a constant $C > 0$ such that,

$$d_P(\varphi^{-1}(q_1), \varphi^{-1}(q_2)) \leq C d_Q(q_1, q_2), \quad (3.25)$$

then for any Q -persistence modules \mathcal{M} and \mathcal{N} , $T < \text{Trans}(Q)$,

$$d_I^{T'}(\varphi^* \mathcal{M}, \varphi^* \mathcal{N}) \leq C d_I^T(\mathcal{M}, \mathcal{N}), \quad (3.26)$$

where $T' = \varphi^{-1} T \varphi$.

Proof. Suppose $d_I(\mathcal{M}, \mathcal{N}) < \infty$, or else the result is trivially true. Then there exists a pair of translations $s, t \in \text{Trans}(Q, \leq)$ such that \mathcal{M} and \mathcal{N} are (s, t) -interleaved, where $\omega(s), \omega(t) \leq \varepsilon$ for some $\varepsilon < \infty$. That is,

$$\sup_{q \in Q} \{d_Q(q, s(q))\} \leq \varepsilon \quad \text{and} \quad \sup_{q \in Q} \{d_Q(q, t(q))\} \leq \varepsilon. \quad (3.27)$$

Using Proposition ??, $\varphi^* \mathcal{M}$ and $\varphi^* \mathcal{N}$ are (s', t') -interleaved, where $s' = \varphi^{-1} s \varphi$, and $t' = \varphi^{-1} t \varphi$.

For any $p \in P$, the condition (??) tells us that,

$$d_P(p, \varphi^{-1} s \varphi(p)) = d_P(\varphi^{-1} \varphi(p), \varphi^{-1} s \varphi(p)) \leq C \cdot d_Q(\varphi(p), s \varphi(p)) \leq C \cdot \varepsilon. \quad (3.28)$$

Hence $\omega(s') = \omega(\varphi^{-1} s \varphi) = \sup_{q \in Q} \{d_Q(q, \varphi^{-1} s \varphi(q))\} \leq C \cdot \varepsilon$.

Similarly, $\omega(t') \leq C \cdot \varepsilon$, and so $d_I(\varphi^* \mathcal{M}, \varphi^* \mathcal{N}) \leq C d_I(\mathcal{M}, \mathcal{N})$.

□

Definition 3.22. Poset morphisms $\varphi, \gamma : P \rightarrow Q$ are **comparable** if either $\varphi(p) \geq \gamma(p)$ for all $p \in P$, or $\varphi(p) < \gamma(p)$ for all $p \in P$.

In particular, if Q is a totally ordered set, then any two morphisms $\varphi, \gamma : P \rightarrow Q$ are comparable.

Lemma 3.23. For any order isomorphisms $\varphi, \gamma : P \rightarrow Q$, φ and γ are comparable if and only if φ^{-1} and γ^{-1} are comparable.

Proof. Suppose φ, γ are comparable. Then for every $p \in P$ we have either $\varphi(p) \geq \gamma(p)$ or $\varphi(p) < \gamma(p)$. We have that $\varphi(p) \geq \gamma(p)$ if and only if $\gamma^{-1} \varphi(p) \geq p$, since γ^{-1} is order-preserving, if and only if $\gamma^{-1}(q) \leq \varphi^{-1}(q)$, where $q = \varphi(p)$.

Similarly, $\varphi(p) < \gamma(p)$ if and only if $\gamma^{-1} \varphi(p) \leq p$, if and only if $\gamma^{-1}(q) \leq \varphi^{-1}(q)$, where $q = \varphi(p)$. We note that as every $q \in Q$ is equal to $\varphi(p)$ for some $p \in P$, then for every $q \in Q$ we have that either $\varphi^{-1}(q) \leq \gamma^{-1}(q)$ or $\varphi^{-1}(q) \geq \gamma^{-1}(q)$.

□

We note that the distances d_P and d_Q are not necessarily symmetric, since the weight functions w_P and w_Q are not. Recall that if $p_i \leq p_j$, then $w_P(p_i, p_j)$ is the weight of the relation $p_i \leq p_j$, and $w_P(p_i, p_j) = \infty$ if $p_i \not\leq p_j$. Hence $w_P(p_i, p_j) = w_P(p_j, p_i)$ if and only if either $p_i = p_j$ or there is no relation between p_i and p_j . We can, however, define a symmetric version of both d_P and d_Q , which allows us to define a generalisation of d_∞ .

Definition 3.24. Given $\varphi, \gamma : P \rightarrow Q$, let $\hat{d}_Q(\varphi, \gamma)$ be given by,

$$\hat{d}_Q(\varphi, \gamma) = \min \left\{ \sup_{p \in P} \{d_Q(\varphi(p), \gamma(p))\}, \sup_{p \in P} \{d_Q(\gamma(p), \varphi(p))\} \right\} \quad (3.29)$$

Theorem 3.25. Let $\varphi, \gamma : P \rightarrow Q$ be comparable order isomorphisms. If φ and γ are such that φ^{-1} and γ^{-1} are Lipschitz maps with constants C_φ, C_γ , respectively, then for any Q -module \mathcal{M} ,

$$d_I(\varphi^* \mathcal{M}, \gamma^* \mathcal{M}) \leq C \hat{d}_Q(\varphi, \gamma), \quad (3.30)$$

where $C = \max\{C_\varphi, C_\gamma\}$.

Proof. First suppose that $\hat{d}(\varphi, \gamma) < \infty$, or else the result is trivially true. Let $\hat{d}(\varphi, \gamma) = \varepsilon$.

Let $s : P \rightarrow P$ be given by,

$$s(p) = \begin{cases} \gamma^{-1}\varphi(p) & \gamma^{-1}\varphi(p) \geq p \\ p & \text{otherwise,} \end{cases} \quad (3.31)$$

and let $t : P \rightarrow P$ be given by,

$$t(p) = \begin{cases} \varphi^{-1}\gamma(p) & \varphi^{-1}\gamma(p) \geq p \\ p & \text{otherwise} \end{cases}. \quad (3.32)$$

We will show that $\varphi^* \mathcal{M}$ and $\gamma^* \mathcal{M}$ are (s, t) -interleaved.

We first check that $s, t \in \text{Trans}(P, \leq)$. Note that the condition that φ and γ are comparable means that for all $p \in P$, either $\varphi(p) \geq \gamma(p)$, in which case $\gamma^{-1}\varphi(p) \geq p$ and $\varphi^{-1}\gamma(p) < p$, or $\varphi(p) < \gamma(p)$, in which case $\gamma^{-1}\varphi(p) < p$, and $\varphi^{-1}\gamma(p) > p$. In either case, we have that s and t satisfy,

$$s(p) \geq p \text{ and } t(p) \geq p \text{ for all } p \in P. \quad (3.33)$$

We note that since φ, γ and their inverses are order-preserving, then so are their compositions. Hence for all $p_1, p_2 \in P$, if $p_1 \leq p_2$ then $s(p_1) \leq s(p_2)$ and $t(p_1) \leq t(p_2)$

We now show that $\varphi^* \mathcal{M}$ and $\gamma^* \mathcal{M}$ are (s, t) -interleaved.

If we look back to how s and t were defined, we see that for all $p \in P$, we have that $s(p) \geq \gamma^{-1}\varphi(p)$, and so, since γ is order-preserving, $\gamma s(p) \geq \varphi(p)$ for all $p \in P$. Hence for every $p \in P$ there exist transition maps in the module \mathcal{M} ,

$$\psi_{\mathcal{M}}(\varphi(p), \gamma s(p)) : \mathcal{M}(\varphi(p)) \rightarrow \mathcal{M}(\gamma s(p)), \quad (3.34)$$

or in other words, for every $p \in P$, there exists a map $\varphi^*\mathcal{M}(p) \rightarrow \gamma^*\mathcal{M}(s(p))$. Similarly, $\varphi t(p) \geq \gamma(p)$ for all $p \in P$, and so there exist transition maps in the module \mathcal{M} ,

$$\psi_{\mathcal{M}}(\gamma(p), \varphi t(p)) : \mathcal{M}(\gamma(p)) \rightarrow \mathcal{M}(\varphi t(p)), \quad (3.35)$$

or in other words, for every $p \in P$, there exists a map $\gamma^*\mathcal{M}(p) \rightarrow \varphi^*\mathcal{M}(t(p))$. Let $p_1, p_2 \in P$ be such that $p_1 \leq p_2$. We can assemble the transition maps in (??) and (??) into diagrams,

$$\begin{array}{ccc} \varphi^*\mathcal{M}(p_1) & \longrightarrow & \varphi^*\mathcal{M}(p_2) & \varphi^*\mathcal{M}(t(p_1)) & \longrightarrow & \varphi^*\mathcal{M}(t(p_2)) \\ \downarrow & & \downarrow & \uparrow & & \uparrow \\ \gamma^*\mathcal{M}(s(p_1)) & \longrightarrow & \gamma^*\mathcal{M}(s(p_2)) & \gamma^*\mathcal{M}(p_1) & \longrightarrow & \gamma^*\mathcal{M}(p_2) \end{array} \quad (3.36)$$

and for any $p \in P$, we can assemble these transition maps into a pair of diagrams,

$$\begin{array}{ccc} \varphi^*\mathcal{M}(p) & \longrightarrow & \varphi^*\mathcal{M}(ts(p)) \\ & \searrow & \nearrow \\ & \gamma^*\mathcal{M}(s(p)) & \end{array} \quad \begin{array}{ccc} & \varphi^*\mathcal{M}(t(p)) & \\ \nearrow & & \searrow \\ \gamma^*\mathcal{M}(p) & \longrightarrow & \gamma^*\mathcal{M}(st(p)) \end{array} \quad (3.37)$$

Since all maps in these diagrams are given by transition maps in \mathcal{M} , they must necessarily commute. Which, by definition, means that the pull-back modules $\varphi^*\mathcal{M}$ and $\gamma^*\mathcal{M}$ are (s, t) -interleaved. Finally, we have that,

$$\omega(s) = \sup_{p \in P} \{d_P(p, s(p))\} = \sup_{p \leq \gamma^{-1}\varphi(p)} \{d_P(p, \gamma^{-1}\varphi(p))\}. \quad (3.38)$$

Given the assumption that γ^{-1} is Lipschitz with constant C_γ , we have that for any $p \in P$,

$$d_P(p, \gamma^{-1}\varphi(p)) = d_P(\gamma^{-1}\gamma(p), \gamma^{-1}\varphi(p)) \leq C_\gamma d_Q(\gamma(p), \varphi(p)) \leq C_\gamma \varepsilon \leq C\varepsilon,$$

where $C = \max\{C_\gamma, C_\varphi\}$.

Similarly,

$$\omega(t) = \sup_{p \in P} \{d_P(p, t(p))\} = \sup_{p \leq \varphi^{-1}\gamma(p)} \{d_P(p, \varphi^{-1}\gamma(p))\}, \quad (3.39)$$

and assuming φ^{-1} is Lipschitz with constant C_φ , then for any $p \in P$,

$$d_P(p, \varphi^{-1}\gamma(p)) = d_P(\varphi^{-1}\varphi(p), \varphi^{-1}\gamma(p)) \leq C_\varphi d_Q(\varphi(p), \gamma(p)) \leq C_\varphi \varepsilon \leq C\varepsilon,$$

where $C = \max\{C_\gamma, C_\varphi\}$.

Hence $\varphi^*\mathcal{M}$ and $\gamma^*\mathcal{M}$ are $C\varepsilon$ -interleaved. \square

We will now consider a special case of Theorem ??, in which we assume φ and γ are automorphisms of a poset, P , and in which γ is the identity map on P . We will see that in this case, we do not need a Lipschitz condition in order to bound $d_I(\varphi^*\mathcal{M}, \mathcal{M})$.

Proposition 3.26. *Let P be any partially ordered set, and let $\varphi : P \rightarrow P$ be a poset automorphism such that φ is comparable to the identity map on P . Then for any P -persistence module \mathcal{M} ,*

$$d_I(\varphi^* \mathcal{M}, \mathcal{M}) \leq \sup_{p \in P} \{d_P(p, \varphi(p))\}. \quad (3.40)$$

Proof. Let

$$s(p) = \begin{cases} \varphi(p) & \varphi(p) \geq p \\ p & \text{otherwise,} \end{cases} \quad \text{and} \quad t(p) = \begin{cases} \varphi^{-1}(p) & \varphi^{-1}(p) \geq p \\ p & \text{otherwise.} \end{cases} \quad (3.41)$$

Then s and t are both translations of P . The condition that φ is comparable with the identity map means that for every $p \in P$, we have either $\varphi(p) \geq p$ or $\varphi(p) < p$. Hence in all cases we have that $s(p) \geq \varphi(p)$, and $t(p) \geq \varphi^{-1}(p)$ for all $p \in P$. Hence for every $p \in P$, we have a transition map in the module \mathcal{M} ,

$$\psi_{\mathcal{M}}(\varphi(p), s(p)) : \mathcal{M}(\varphi(p)) \rightarrow \mathcal{M}(s(p)). \quad (3.42)$$

Since for every $p \in P$, $\mathcal{M}(\varphi(p)) = \varphi^* \mathcal{M}(p)$, this is the same as saying that for every $p \in P$, we have a map,

$$f_p := \psi_{\mathcal{M}}(\varphi(p), s(p)) : \varphi^* \mathcal{M}(p) \rightarrow \mathcal{M}(s(p)). \quad (3.43)$$

Similarly, as $t(p) \geq \varphi^{-1}(p)$ for all $p \in P$, then in particular, $\varphi t(p) \geq p$ for all $p \in P$, and so we have a map,

$$g_p := \psi_{\mathcal{M}}(p, \varphi t(p)) : \mathcal{M}(p) \rightarrow \varphi^* \mathcal{M}(t(p)), \quad (3.44)$$

and moreover the composition of these maps must satisfy,

$$f_{t(p)} \circ g_p = \psi_{\mathcal{M}}(p, ts(p)) \quad \text{and} \quad g_{s(p)} \circ f_p = \psi_{\varphi^* \mathcal{M}}(p, st(p)) = \psi_{\mathcal{M}}(\varphi(p), \varphi st(p)), \quad (3.45)$$

since both f and g are themselves transition maps, which commute as we require by definition. Hence $\varphi^* \mathcal{M}$ and \mathcal{M} are (s, t) -interleaved.

Finally, we note that

$$\omega(s) = \sup_{p \in P} \{d_P(p, s(p))\} = \sup_{p \leq \varphi(p)} \{d_P(p, \varphi(p))\} \leq \sup_{p \in P} \{d_P(p, \varphi(p))\}, \quad (3.46)$$

and

$$\omega(t) = \sup_{p \in P} \{d_P(p, t(p))\} = \sup_{p \leq \varphi^{-1}(p)} \{d_P(p, \varphi^{-1}(p))\} \leq \sup_{p \in P} \{d_P(p, \varphi(p))\}. \quad (3.47)$$

where the final inequality in (??) follows as φ is surjective. \square

3.4 Kan Extensions of Persistence Modules

So far in this section, we have considered the pull-back of a Q -persistence module along a morphism of posets $\varphi : P \rightarrow Q$. We now consider the dual concept of a push-forward of a persistence module. That is, given a P -persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$, and a poset morphism $\varphi : P \rightarrow Q$, we consider how to extend the domain of \mathcal{M} to Q . We have seen that a persistence module is nothing more than a functor between categories, where the source category is a partially ordered set. Thinking of persistence modules in this way means that we can make use of concepts in category theory to solve problems in persistence. Since \mathcal{M} is simply a functor from P to \mathcal{C} , and a morphism of posets is really just a functor of thin categories, then we may recognise that the problem of extending the domain of \mathcal{M} to Q amounts to finding the *Kan extension* of \mathcal{M} along φ .

In this section, we'll give an overview of the definition of a left Kan extension of a persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$ along a morphism $\varphi : P \rightarrow Q$, of posets P and Q , as seen in (?). We will then show explicit examples of how the Left Kan extension is constructed for given persistence modules and poset morphisms. Beyond this, we will give a detailed description of the structure of a left Kan extension along an order isomorphism. This will lead us to dual results to those of section ?? for the push-forward of a module along a morphism of posets.

A Kan extension can, however, be defined for arbitrary maps, and was described in (?) as a way to discretise continuous persistence modules. Our final main result in this section shows that we should be cautious when using Kan extensions to discretise a persistence process. We show that discretising in this way is, in general, not a stable process.

We note that there is a dual concept of a right Kan extension of a persistence module along a morphism of posets, which can be seen in (?). Since this section is intended to illustrate the potential instability of the construction, we only detail a left Kan extension here, rather than describe two very similar constructions.

A reader unfamiliar with Kan extensions who wishes to see the general definition may consult (?).

Definition 3.27. The left Kan extension of $\mathcal{M} : P \rightarrow \mathcal{C}$ along $\varphi : P \rightarrow Q$ is a persistence module $Lan_\varphi \mathcal{M} : Q \rightarrow \mathcal{C}$,

$$\begin{array}{ccc}
 & Q & \\
 \varphi \nearrow & & \dashrightarrow Lan_\varphi \mathcal{M} \\
 P & \xrightarrow{\mathcal{M}} & C
 \end{array} \tag{3.48}$$

which can be thought of as the best approximation of \mathcal{M} with a change of domain from $P \rightarrow Q$, in the sense that there is a morphism of persistence modules, $\eta : \mathcal{M} \rightarrow Lan_\varphi \mathcal{M} \circ \varphi$ such that if $X : Q \rightarrow \mathcal{C}$ is any other persistence module for which there exists a morphism, $\alpha : \mathcal{M} \rightarrow X \circ \varphi$, there is a unique morphism of persistence modules $\sigma : Lan_\varphi \mathcal{M} \rightarrow X$ such that $\sigma \circ \eta = \alpha$.

In certain circumstances, we can define the left Kan extension on a module in a simpler way.

Definition 3.28. Given a functor $F : A \rightarrow B$, a **co-cone** of F is an object $b \in Ob(B)$, together with morphisms,

$$\alpha_i : F(a_i) \rightarrow b,$$

for each $a_i \in Ob(A)$, such that whenever we have a morphism $f : a_i \rightarrow a_j$ between a pair of objects of A , the diagram,

$$\begin{array}{ccc}
 F(a_i) & & \\
 \downarrow F(f) & \searrow \alpha_i & \\
 & & b \\
 & \nearrow \alpha_j & \\
 F(a_j) & &
 \end{array}
 \tag{3.49}$$

commutes.

A co-cone (b, α) is a **colimit** of F if for any other co-cone (b', α') , there exists a unique morphism $b \rightarrow b'$ such that the following diagram commutes.

$$\begin{array}{ccccc}
 F(a_i) & & \xrightarrow{\alpha'_i} & & \\
 \downarrow f & \searrow \alpha_i & & \searrow & \\
 & & b & \dashrightarrow & b' \\
 & \nearrow \alpha_j & & \nearrow & \\
 F(a_j) & & \xrightarrow{\alpha'_j} & &
 \end{array}
 \tag{3.50}$$

That is, the colimit of F is the universal co-cone of F .

Definition 3.29. A category C is **cocomplete** if the colimit of all functors $A \rightarrow C$ exists, for all small categories, A .

If C is cocomplete then, firstly, we have that the left Kan extension $Lan_\varphi \mathcal{M}$ exists, and secondly, it can be defined object-wise as the colimit of a particular functor (?). Cocomplete categories are not rare – the category Top of topological spaces, the category R-mod of modules over a commutative ring R , and the category $Vect_k$ of vector spaces over a field k are all cocomplete. From now on we will assume that our persistence modules take values in a cocomplete category, \mathcal{C} .

Definition 3.30. Given a poset morphism $\varphi : P \rightarrow Q$, and $q \in Q$, define the **comma category**, $(\varphi \downarrow q)$ to be the category with objects given by the morphisms of Q of the form $\varphi(p) \leq q$, where $p \in P$. For $p_1 \leq p_2$, then the morphisms in the comma category,

$$\varphi(p_1) \leq q \rightarrow \varphi(p_2) \leq q,$$

are simply given by the morphisms $p_1 \leq p_2$ in P .

There is a forgetful functor, $\pi : (\varphi \downarrow q) \rightarrow P$, which sends an object $\varphi(p) \leq q$ of the comma category to $p \in P$.

Definition 3.31. The **left Kan extension**, $Lan_\varphi \mathcal{M}$, of $\mathcal{M} : P \rightarrow \mathcal{C}$ along $\varphi : P \rightarrow Q$ is the colimit of the composite functor,

$$(\varphi \downarrow q) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}. \quad (3.51)$$

For $q_i \leq q_j$, a co-cone of $(\varphi \downarrow q_j) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}$ is also a co-cone of $(\varphi \downarrow q_i) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}$. Hence the colimit, $Lan_\varphi \mathcal{M}(q_j)$, of $(\varphi \downarrow q_j) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}$ is a co-cone of $(\varphi \downarrow q_i) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}$, and so there is a unique map $\psi(i, j) : Lan_\varphi \mathcal{M}(q_i) \rightarrow Lan_\varphi \mathcal{M}(q_j)$, by definition of a colimit. We define the transition maps $\psi_{Lan_\varphi \mathcal{M}}$ to be given by,

$$\psi_{Lan_\varphi \mathcal{M}}(q_i, q_j) := \psi(i, j). \quad (3.52)$$

When the poset morphism $\varphi : P \rightarrow Q$ is an order isomorphism, then the left Kan extension of a persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$ along $\varphi : P \rightarrow Q$ has a much simpler description.

Proposition 3.32. *If $\varphi : P \rightarrow Q$ is an order isomorphism, then the left Kan extension of a P -persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$ along φ is given by,*

$$Lan_\varphi \mathcal{M}(q) = \mathcal{M}(\varphi^{-1}(q)). \quad (3.53)$$

and the transition maps in $Lan_\varphi \mathcal{M}$ are equal to corresponding transition maps in \mathcal{M} ,

$$\psi_{Lan_\varphi \mathcal{M}}(q_1, q_2) = \psi_{\mathcal{M}}(\varphi^{-1}(q_1), \varphi^{-1}(q_2)). \quad (3.54)$$

Proof. We first note that $(\mathcal{M}(\varphi^{-1}(q)), \psi_{\mathcal{M}})$ is a co-cone of the composite functor,

$$(\varphi \downarrow q) \xrightarrow{\pi} P \xrightarrow{\mathcal{M}} \mathcal{C}, \quad (3.55)$$

since for all $\varphi(p_i) \leq q$, there exists a transition map $\psi_{\mathcal{M}} : \mathcal{M}(p_i) \rightarrow \mathcal{M}(\varphi^{-1}(q))$, and for any $\varphi(p_i) \leq \varphi(p_j) \leq q$, the diagram of transition maps,

$$\begin{array}{ccc} \mathcal{M}(p_i) & & \\ \downarrow \psi_{\mathcal{M}(p_i, p_j)} & \searrow \psi_{\mathcal{M}(p_i, \varphi^{-1}(q))} & \\ & & \mathcal{M}(\varphi^{-1}(q)) \\ & \nearrow \psi_{\mathcal{M}(p_j, \varphi^{-1}(q))} & \\ \mathcal{M}(p_j) & & \end{array} \quad (3.56)$$

must necessarily commute, since all transition maps in \mathcal{M} commute.

If (N, f) is another co-cone of $(\varphi \downarrow q_i) \rightarrow P \rightarrow \mathcal{C}$, then since $\varphi(\varphi^{-1}(q)) \leq q$, there exists a morphism, $f_q : \mathcal{M}(\varphi^{-1}(q)) \rightarrow N$. Hence $\mathcal{M}(\varphi^{-1}(q))$ satisfies the universal property and is hence the Left Kan Extension of $\mathcal{M} : P \rightarrow \mathcal{C}$ along φ . \square

Hence, in the case that $\varphi : P \rightarrow Q$ is an order isomorphism, the left Kan extension of a P -persistence module $\mathcal{M} : P \rightarrow \mathcal{C}$ along φ is the same as the pull-back of \mathcal{M} along the inverse map $\varphi^{-1} : Q \rightarrow P$. We will refer to this module as the **push-forward of $\mathcal{M} : P \rightarrow \mathcal{C}$ along φ** , and denote it by $\varphi_*\mathcal{M}$.

We can use a similar method to find the left Kan extension of a P -persistence module along a general morphism of posets.

Example 3.33. Let P and Q be the posets with Hasse diagrams as shown in Figure ??, and let $\varphi : P \rightarrow Q$ be given by,

$$\varphi(a) = e, \varphi(b) = f \text{ and } \varphi(c) = g. \quad (3.57)$$




Figure 28: Hasse diagrams for posets P (left) and Q (right).

Let $\mathcal{M} : P \rightarrow Vect_k$ be the P -persistence module given by,

$$\begin{array}{c} k \\ \uparrow id_k \\ k \\ \uparrow id_k \\ k \end{array} \quad (3.58)$$

Just as in the proof of Proposition ??, if $q \in Q$ is such that $\varphi(p) = q$, then,

$$Lan_{\varphi}\mathcal{M}(q) = \mathcal{M}(p).$$

Hence,

$$Lan_{\varphi}\mathcal{M}(e) = \mathcal{M}(a) = k,$$

$$Lan_{\varphi}\mathcal{M}(f) = \mathcal{M}(b) = k,$$

$$Lan_{\varphi}\mathcal{M}(g) = \mathcal{M}(c) = k,$$

and the transition maps between these vector spaces are precisely those inherited from the persistence module \mathcal{M} . Explicitly, $\psi_{Lan_{\varphi}\mathcal{M}}(e, f) = \psi_{Lan_{\varphi}\mathcal{M}}(f, g) = id_k$.

We still have to consider $Lan_\varphi\mathcal{M}(d)$ and $Lan_\varphi\mathcal{M}(h)$.

Firstly, $(\mathcal{M}(a), \psi_{\mathcal{M}}(-, a))$ is a co-cone of the diagram $(\varphi \downarrow d) \rightarrow P \rightarrow Vect_k$ since for every $p \in P$ with $\varphi(p) \leq d$, we have a morphism $\psi_{\mathcal{M}} : \mathcal{M}(p) \rightarrow \mathcal{M}(a)$, and this is the universal co-cone by an identical argument to the one in the proof of Proposition ?? – if (N, f) is any other co-cone of $(\varphi \downarrow d) \rightarrow P \rightarrow Vect_k$, then there exists a unique morphism $f_a : \mathcal{M}(a) \rightarrow N$, since $\varphi(a) \leq d$. Hence $Lan_\varphi\mathcal{M}(d) = \mathcal{M}(a) = k$.

Finally, $Lan_\varphi\mathcal{M}(h)$ is the colimit of $(\varphi \downarrow h) \rightarrow P \rightarrow Vect_k$, but in this case, the comma category $(\varphi \downarrow h)$ is empty, and so $Lan_\varphi\mathcal{M}(h)$ is the colimit of the empty diagram, that is, the zero vector space (?). The Q -persistence module $Lan_\varphi\mathcal{M}$ is then shown in Figure ??.




Figure 29: The persistence module $Lan_\varphi\mathcal{M}$.

3.5 Push-forwards along Order Isomorphisms

Given Proposition ??, the push-forward of a P -persistence module \mathcal{M} along an order isomorphism $\varphi : P \rightarrow Q$ is the same as the pull-back of \mathcal{M} along the inverse, $\varphi^{-1} : Q \rightarrow P$. Hence all of the results from section ?? have a dual statement, for push-forwards. All of the following results hold due to Proposition ??, and the corresponding results in section ??.

Lemma 3.34. *The push-forward of P -persistence modules along an order isomorphism $\varphi : P \rightarrow Q$ preserves isomorphism classes of persistence modules.*

Lemma 3.35. *For any poset morphism $\varphi : P \rightarrow Q$, and any $s \in Trans(Q, \leq)$, $\varphi^{-1}s\varphi \in Trans(P, \leq)$.*

Proposition 3.36. *For any order isomorphism $\varphi : P \rightarrow Q$, any P -persistence modules \mathcal{M} and \mathcal{N} are (s, t) -interleaved if and only if the Q -persistence modules $\varphi_*\mathcal{M}$ and $\varphi_*\mathcal{N}$ are (\tilde{s}, \tilde{t}) -interleaved, where,*

$$\tilde{s} = \varphi s \varphi^{-1} \quad \text{and} \quad \tilde{t} = \varphi t \varphi^{-1}. \quad (3.59)$$

Theorem 3.37. *If $\varphi : P \rightarrow Q$ is an order isomorphism such that φ is Lipschitz with constant $C \geq 0$, then the push-forward of P -persistence modules along φ is also a Lipschitz map with constant C . That is, if φ is such that for any $p_1, p_2 \in P$,*

$$d_Q(\varphi(p_1), \varphi(p_2)) \leq C d_P(p_1, p_2), \quad (3.60)$$

then for any P -persistence modules \mathcal{M} and \mathcal{N} , and any $T < \text{Trans}(P)$,

$$d_I^{T'}(\varphi_*\mathcal{M}, \varphi_*\mathcal{N}) \leq C d_I^T(\mathcal{M}, \mathcal{N}), \quad (3.61)$$

where $T' = \varphi T \varphi^{-1}$.

We can combine this result with that of Theorem ?? to get the following result.

Corollary 3.38. *If $\varphi : P \rightarrow Q$ is an order isomorphism which is bi-Lipschitz with constant $C \geq 0$, then the push-forward of P -persistence modules along φ , and the pull-back of Q -persistence modules along φ , are also bi-Lipschitz maps with constant C . That is, if φ is such that for any $p_1, p_2 \in P$, $T < \text{Trans}(P)$,*

$$\frac{1}{C} d_P(p_1, p_2) \leq d_Q(\varphi(p_1), \varphi(p_2)) \leq C d_P(p_1, p_2), \quad (3.62)$$

then for any P -persistence modules \mathcal{M} and \mathcal{N} ,

$$\frac{1}{C} d_I^T(\mathcal{M}, \mathcal{N}) \leq d_I^{T'}(\varphi_*\mathcal{M}, \varphi_*\mathcal{N}) \leq C d_I^T(\mathcal{M}, \mathcal{N}), \quad (3.63)$$

and for any Q -persistence modules \mathcal{M}' and \mathcal{N}' ,

$$\frac{1}{C} d_I^{T'}(\mathcal{M}', \mathcal{N}') \leq d_I^T(\varphi^*\mathcal{M}', \varphi^*\mathcal{N}') \leq C d_I^{T'}(\mathcal{M}', \mathcal{N}'), \quad (3.64)$$

where $T' = \varphi T \varphi^{-1}$.

Example 3.39. Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be given by $\varphi(a) = ma + c$, for some $m \geq 0$. Then φ is bi-Lipschitz with respect to the Euclidean metric on \mathbb{R} , with constant m . Let \mathcal{M} and \mathcal{N} be the interval modules,

$$\mathcal{M} = k[0, 1], \quad \mathcal{N} = k[0, 2].$$

Then by Example ?? \mathcal{M} and \mathcal{N} are 1-interleaved. Using Corollary ??, the interleaving distance between the push-forward modules along φ satisfies,

$$\frac{1}{m} \leq d_I(\varphi_*\mathcal{M}, \varphi_*\mathcal{N}) \leq m.$$

Combining the results of Theorem ?? and Theorem ??, we also obtain the following corollary.

Corollary 3.40. *If $\varphi : P \rightarrow Q$ is an order isomorphism such that there exists a $C \geq 0$, such that for every $p_1, p_2 \in P$,*

$$d_Q(\varphi(p_1), \varphi(p_2)) = Cd_P(p_1, p_2), \quad (3.65)$$

then for any P -persistence modules \mathcal{M} and \mathcal{N} , and any $T < \text{Trans}(P)$,

$$d_I^{T'}(\varphi_*\mathcal{M}, \varphi_*\mathcal{N}) = Cd_I^T(\mathcal{M}, \mathcal{N}). \quad (3.66)$$

and for any Q -persistence modules \mathcal{M}' and \mathcal{N}' ,

$$d_I^T(\varphi^*\mathcal{M}', \varphi^*\mathcal{N}') = \frac{1}{C}d_I^{T'}(\mathcal{M}', \mathcal{N}'), \quad (3.67)$$

where $T' = \varphi T \varphi^{-1}$.

Example 3.41. Let P be a subset of the real line consisting of a set of 100 equally-spaced real numbers, $P = \{p_i | i = 1, \dots, 100\}$, with common difference $|p_i - p_{i+1}| = g$. Let the partial order on P be that inherited from \mathbb{R} , and let the metric d_P be the Euclidean metric, again inherited from \mathbb{R} . Let Q be the subset of the integers given by $Q = \{1, \dots, 100\}$, with the usual partial order inherited from \mathbb{R} , and with metric d_Q given by $d_Q(q_i, q_j) = |q_j - q_i|$.

Let $\varphi : P \rightarrow Q$ be given by $\varphi(p_i) = i$. Then φ is an order isomorphism between posets P and Q , and satisfies,

$$d_Q(\varphi(p_i), \varphi(p_j)) = gd_P(p_i, p_j).$$

Then Corollary ?? states that for any P -persistence modules $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$,

$$d_I(\varphi^*\mathcal{M}, \varphi^*\mathcal{N}) = gd_I(\mathcal{M}, \mathcal{N}),$$

and for any Q -persistence modules \mathcal{M}' and \mathcal{N}' ,

$$d_I(\varphi_*\mathcal{M}', \varphi_*\mathcal{N}') = \frac{1}{g}d_I(\mathcal{M}', \mathcal{N}'). \quad (3.68)$$

In particular, Corollary ?? tells us that if $C = 1$, we have the following result.

Corollary 3.42. *If $\varphi : P \rightarrow Q$ is an isometry with respect to the chosen metrics on P and Q , then both pulling back a pair of Q -persistence modules, and pushing forward a pair of P -persistence modules along φ induce an isometry with respect to the interleaving distance.*

The following Theorem is dual to Theorem ?. The proof follows a near-identical structure, so is not included here.

Theorem 3.43. *Let $\varphi, \gamma : P \rightarrow Q$ be comparable order isomorphisms. If φ and γ are Lipschitz maps with constants C_φ, C_γ , respectively, then for any P -modules \mathcal{M} and \mathcal{N} ,*

$$d_I(\varphi_*\mathcal{M}, \gamma_*\mathcal{M}) \leq C\hat{d}(\varphi^{-1}, \gamma^{-1}), \quad (3.69)$$

where $C = \max\{C_\varphi, C_\gamma\}$.

3.6 The Instability of Kan Extensions of Persistence Modules

We have shown that if the map $\varphi : P \rightarrow Q$ preserves the poset structure — that is, if φ is an order isomorphism — then pushing forward and pulling back along φ causes Lipschitz properties of φ to descend to identical properties in the interleaving distance between the pull-back and push-forward modules. But a pull-back module can be defined even in the case that φ is not an order isomorphism, as can the left Kan extension of a module. In this section, we demonstrate that for general $\varphi : P \rightarrow Q$, and modules $\mathcal{M}, \mathcal{N} : P \rightarrow \mathcal{C}$, when we can calculate either $d_I^{Trans(P)}(\mathcal{M}, \mathcal{N})$ or $d_I^{Trans(Q)}(\varphi_*\mathcal{M}, \varphi_*\mathcal{N})$, the two are, in general, not equal.

We will also show that a Kan extension along a non-isomorphism $\varphi : P \rightarrow Q$ does not preserve the interleaving distance even in the case that φ is an isometry of posets. This demonstrates how important the structure of the poset, dictated by the order, is in determining interleaving distances between modules. It also provides a strong note of caution when it comes to reparameterising, or discretizing persistence modules. If the morphism is not an order isomorphism, then reparameterising along such a morphism may distort the distance between persistence modules. This section continues to build upon the work of (?) in considering Kan extensions of persistence modules and the interleaving distance between push-forward modules.

The following example demonstrates this.

Example 3.44. Let P and Q be the posets with Hasse diagrams H_P and H_Q as shown in Figure ??, and define Lawvere metrics on P and Q to be given by the directed path length metrics. That is, for $p_1, p_2 \in P$, $d_P(p_1, p_2)$ is the number of edges in the shortest directed path in H_P between p_1 and p_2 , and similarly for d_Q .




Figure 30: Hasse diagrams for posets P (left) and Q (right).

Let $\varphi : P \rightarrow Q$ be given by,

$$\varphi(a) = f, \varphi(b) = g \text{ and } \varphi(c) = h. \tag{3.70}$$

Notice that φ is an isometry with respect to d_P and d_Q – for any $p_i, p_j \in P$, $d_Q(\varphi(p_i), \varphi(p_j)) = d_P(p_i, p_j)$. In particular, this means that φ is Lipschitz with constant 1.

Consider the following pair of P -persistence modules, $\mathcal{M}, \mathcal{N} : P \rightarrow Vect_k$.

$$\begin{array}{ccc}
 k & & 0 \\
 id_k \uparrow & & 0 \uparrow \\
 k & & k \\
 id_k \uparrow & & id_k \uparrow \\
 k & & k
 \end{array}$$

Figure 31: The P -persistence modules \mathcal{M} (left) and \mathcal{N} (right).

Then \mathcal{M} and \mathcal{N} are (s, t) -interleaved, where $s, t \in Trans(P, \leq)$ are given by $s = id_P$, and t is defined by,

$$t(a) = a, t(b) = a, t(c) = b.$$

Since $\omega(s) = \sup_{p \in P} \{d_P(p, s(p))\} \leq 1$ and $\omega(t) = \sup_{p \in P} \{d_P(p, t(p))\} \leq 1$, we have that \mathcal{M} and \mathcal{N} are 1-interleaved.

Now consider the left Kan extensions of \mathcal{M} and \mathcal{N} along $\varphi : P \rightarrow Q$. Using a method similar to that in Example ??, we have that $Lan_\varphi \mathcal{M}$ and $Lan_\varphi \mathcal{N}$ are given by the Q -persistence modules,

$$\begin{array}{ccc}
 k & & 0 \\
 id_k \uparrow & & 0 \uparrow \\
 k & & 0 \\
 id_k \uparrow & & 0 \uparrow \\
 k & \text{and} & 0 \\
 id_k \uparrow & & 0 \uparrow \\
 k & & k \\
 id_k \uparrow & & id_k \uparrow \\
 k & & k
 \end{array} \tag{3.71}$$

respectively.

To see that left Kan extensions do not preserve Lipschitz properties of φ , it is enough to see that $Lan_\varphi \mathcal{M}$ and $Lan_\varphi \mathcal{N}$ are not 1-interleaved. This can be seen by noting that there is no pair of maps $Lan_\varphi \mathcal{M}(g) \rightarrow Lan_\varphi \mathcal{N}(f)$ and $Lan_\varphi \mathcal{N}(f) \rightarrow Lan_\varphi \mathcal{M}(e)$ which compose to the transition map $\psi_{Lan_\varphi \mathcal{M}}(g, e) = id_k$.

The results of this section show that when the morphism of posets preserves the structure given by the order, then Lipschitz properties of the map pass to identical properties of

the pull-back and push-forward operations. However, if the morphism of posets does not preserve the order, then these Lipschitz properties are not necessarily preserved. This shows us that when we rescale modules via an order-preserving map, then the distance between modules scales accordingly. The operation of pushing a persistence module forward along a non-isomorphism, however, is not stable with respect to the interleaving distance, and so reparameterising persistence modules should be done with care.

Further Work

There is a rich selection of further work which could follow from the results of this thesis. We have considered generalised interleavings of arbitrary poset modules, but exploring generalised interleavings of the more familiar \mathbb{R} -modules is also of interest. We have focused on the interleaving distance, but for \mathbb{R} -modules, we could also look at their barcodes or persistence diagrams. As Lesnick notes (?), the relationship between the barcodes of (s, t) -interleaved \mathbb{R} -modules for general $s, t \in \text{Trans}(\mathbb{R})$ is still not fully understood in the way that it is for $s, t \in T_E$.

There are also numerous opportunities to investigate in more detail the effect of changing the choice of submodule $T < \text{Trans}(P)$ on the interleaving distance, d_I^T , or indeed, the effect of changing the choice of pseudometric, d_P on the poset, P .

Within this thesis we have commented on various other sources of further work inspired by our results. One of which is the possibility of considering the universality of the interleaving distance d_I^T , in light of our generalisation of the supremum distance, d_∞ .

One other source of further work which we have commented on in section 3 was the possibility of considering equivalence classes of persistence modules, under the relation of being isomorphic under a pull-back along a morphism of posets. This idea came about from the simple observation that many persistence modules which we know to be related can be described as being identical up to a pull-back along some poset morphism.

One final, related idea would be to investigate whether we can use relationships between pull-back modules to determine whether one persistence module is obtained as a rescaling of another. That is, if we know that the pull-backs of modules display certain symmetries, it would be interesting to see whether we can infer that there is some symmetry in the original process.

References

- [1] Bauer, U., Landi, C. and Memoli, F. The Reeb Graph Edit Distance is Universal. *arXiv preprint arXiv:1801.01866*; 2018.
- [2] Bendich, P. Analyzing Stratified Spaces Using Persistent Versions of Intersection and Local Homology. PhD Thesis, Duke University, 2009.
- [3] Botnan, M. B., Curry, J. and Munch, E. A Relative Theory of Interleavings. *arXiv preprint*. 2020.
- [4] Bubenik, P., De Silva, V. and Scott, J. Metrics for Generalized Persistence Modules. *Foundations of Computational Mathematics*. 2015; 15(6): 1501-1531.
- [5] Bubenik, P., De Silva, V. and Scott, J. Interleaving and Gromov-Hausdorff distance. *arXiv preprint arXiv:1707.06288*; 2017.
- [6] Carlsson, G. and de Silva, V. Zigzag Persistence. *arXiv preprint arXiv:0812.0197*; 2008.
- [7] Carlsson, G. and Zomorodian, A. The Theory of Multidimensional Persistence. *Discrete & Computational Geometry*. 2009; 42(1): 71-93.
- [8] Carlsson, G. *On the Shape of Data*. 2012 SIAM Annual Meeting, July 9, Minneapolis.
- [9] Chacholski, W., De Gregorio, A., Quercioli, N. and Tombari, F. Landscapes of data sets and functoriality of persistent homology. *arXiv preprint arXiv:2002.05972*; 2020.
- [10] Chazal, F., Cohen-Steiner, D., Glisse, M., Guibas, L.J. and Oudot, S.Y. Proximity of persistence modules and their diagrams. *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*. 2009; 237-246.
- [11] Chazal, F., De Silva, V., Glisse, M. and Oudot, S. *The Structure and Stability of Persistence Modules*. Springer; 2016.
- [12] Cohen-Steiner, D., Edelsbrunner, H. and Harer, J. Stability of Persistence Diagrams. *Discrete & Computational Geometry*. 2007; 37(1):103-120.
- [13] Cohen-Steiner, D., Edelsbrunner, H. and Harer, J. Extending Persistence using Poincaré and Lefschetz Duality. *Foundations of Computational Mathematics*. 2009; 9(1): 79-103.
- [14] De Silva, V. and Ghrist, R. Coverage in Sensor Networks via Persistent Homology. *Algebraic & Geometric Topology*. 2007; 7(1): 339-358.
- [15] Di Fabio, B. and Landi, C. A Mayer Vietoris Formula for Persistent Homology with an Application to Shape Recognition in the Presence of Occlusions. *Foundations of Computational Mathematics*. 2011; 11(5): 499.
- [16] Edelsbrunner, H., Letscher, D. and Zomorodian, A. Topological Persistence and Simplification. *Discrete Computational Geometry*. 2002; 28:511-533.
- [17] Edelsbrunner, H. and Harer, J. L. *Computational Topology: An Introduction*. Providence, RI: AMS; 2010.
- [18] Frosini, P. G-invariant Persistent Homology. *Mathematical Methods in the Applied Sciences*. 2015; 38(6):1190-1199.

- [19] Ghrist, R. Barcodes: The Persistent Topology of Data. *Bulletin of the American Mathematical Society New Series*. 2008; 45(1): 61-75.
- [20] Ghrist, R. Three Examples of Applied and Computational Homology. *Lab Papers (GRASP)*. 2008; p18.
- [21] Grandis, M. Categories, Norms and Weights. *Journal of Homotopy and Related Structures*. 2007; 2(2):171-186.
- [22] Hatcher, A. *Algebraic Topology*. New York: Cambridge University Press; 2001.
- [23] Lawvere, F.W. Metric Spaces, Generalized Logic, and Closed Categories. *Rendiconti del seminario matematico e fisico di Milano*, 1973; 43(1), 135-166.
- [24] Lesnick, M. The Theory of the Interleaving Distance on Multidimensional Persistence Modules. *Foundations of Computational Mathematics*. 2015; 15(3): 613-650.
- [25] Lesnick, M. Multidimensional Interleavings and Applications to Topological Inference. PhD Thesis, Stanford University, 2012.
- [26] MacLane, S. *Categories for the Working Mathematician. Graduate Texts in Mathematics*. 5 (2nd ed.). New York, NY: Springer-Verlag; 1998.
- [27] Miller, E. Modules Over Posets: Commutative and Homological Algebra. *arXiv preprint arXiv:1908.09750*; 2019.
- [28] Nicolau, M., Levine, A.J. and Carlsson, G. Topology Based Data Analysis Identifies a Subgroup of Breast Cancers with a Unique Mutational Profile and Excellent Survival. *Proceedings of the National Academy of Sciences*. 2011; 108(17): 7265-7270.
- [29] Palser, M. An Excision Theorem for Persistent Homology. *arXiv preprint*; 2019.
- [30] Robins, V. Towards Computing Homology from Finite Approximations. *Topology Proceedings*. 1999; 24(1):503-532
- [31] Roe, J. What is... A Coarse Space? *Notices of the American Mathematical Society*, 2006; 53(6): 668-669.
- [32] Varli, H., Yilmaz, Y. and Pamuk, M. Homological Properties for Persistent Homology. *arXiv preprint*; 2018.
- [33] Verri, A., Uras, C., Frosini, P. and Ferri, M. On the use of Size Functions for Shape Analysis. *Biological Cybernetics*. 1993; 70(2):99-107.
- [34] Zomorodian, A. and Carlsson, G. Computing Persistent Homology. *Discrete & Computational Geometry*. 2005; 33(2):249-274.