

Discrete Morse Theory & Persistent Homotopy

With Applications in Image Analysis

Kelly Maggs

A thesis submitted for the degree of Bachelor of Mathematics (Honours)
of the Australian National University



**Australian
National
University**

For Matthew Richardson, who kicked 800 career goals and the most of all time on the MCG.



Declaration

The work in this thesis is my own except where otherwise stated.

Kelly Maggs

Acknowledgements

First and foremost I would like to acknowledge my wonderful supervisors Vanessa Robins and Katherine Turner. My weekly meetings with Vanessa were something to look forward to throughout the year. Both were exceptionally generous with their time and support, providing invaluable guidance and mathematical insight. They gave me the freedom to experiment and pursue ideas, while recommending interesting literature and explaining things to me (slowly) whenever I got stuck.

Other academics I'd like to thank include Vidit Nanda (Oxford University), who kindly responded to my long emails about his papers, which feature heavily in my thesis. From ANU, I want to thank Vigleik Angeltveit for running excellent classes in algebraic topology, and Joan Licata for being a supportive and thoughtful honours co-ordinator. Both made the move to Canberra a more pleasant and worthwhile one.

I want to acknowledge a number of people in my personal life, without whom the thesis would not have been possible. My family for their support, both emotional and financial. My roommates, Anna, who made sure I looked after myself. My colleagues and good mates, Owen and Frederick, for their office banter and great work.

I am also fortunate to have a great network of friends in Melbourne, from all of whom I drew support throughout the year. Most importantly, I'd like to thank my partner Tamrin, whose understanding and constant support was indispensable.

It just wouldn't be right to begin this thesis without fully acknowledging the inspirational Richmond Football Club. Every time I was stuck on a particularly hard problem, I would re-watch their glorious 2017 AFL Premiership victory and the solution would immediately become obvious.

Abstract

In this thesis, we present new theoretical tools in topological data analysis with applications in image analysis.

We draw on a number of tools from Discrete Morse theory, presenting the relevant concepts from [7], [15] and [16] in Chapter 2. We prove an original result in Chapter 3 concerning the attaching maps of CW complex comprised of the critical cells of a discrete Morse function.

We introduce the concept of a *homotopy merge tree* in Chapter 4 as an algebraic tool to summarise homotopical changes over a filtered space. The definition is an extension of the work of [14], retaining the important properties of interleaving distance and stability.

We show that the results of Chapter 2 can be used to simplify calculations of the homotopy merge tree. We also provide original, provably correct algorithms in Chapters 3 and 5 to demonstrate the computational viability of our tools.

Contents

Acknowledgements	vii
Abstract	ix
Notation and terminology	xiii
1 Introduction	1
2 Discrete Morse Theory	5
2.1 Background	5
2.1.1 CW Complexes	5
2.1.2 Discrete Morse Functions	8
2.1.3 Partially Ordered Sets	10
2.2 Homology	12
2.2.1 Orientation of V-paths	12
2.2.2 Morse Chain Complex	13
2.3 Homotopy	14
2.3.1 Morse Homotopy Theorems	14
2.3.2 Induced Flow-lines	15
2.3.3 Flow-paths and Stable Subdivision	18
3 The Forman Complex	23
3.1 Construction	23
3.1.1 Main Theorem	24
3.2 Applications	35
3.2.1 Homology	35
3.2.2 Image Partitioning and Connectivity	37
3.3 Computing Attaching Maps	39
3.3.1 Embedding Functions	39

3.3.2	The 2-skeleton	40
3.3.3	Algorithms	45
4	Persistence	55
4.1	Persistent Homology	55
4.1.1	Persistence Modules	56
4.1.2	Interleaving Distance	57
4.2	Persistent Homotopy	58
4.2.1	Merge Trees	59
4.2.2	Canonical Paths	60
4.2.3	Homotopy Merge Tree	61
4.2.4	Comparing Homotopy Merge Trees	65
4.2.5	Stability Theorem	71
4.3	Application: Based Persistence	74
4.3.1	Based Persistence	74
4.3.2	Persistent Homology	77
5	Computing the Homotopy Merge Tree	79
5.1	Sub-level Complexes	79
5.2	Computing $\pi_1[T_{f_\circ}]$	81
5.2.1	The Fundamental Group	81
6	Discussion	85
6.1	Contributions	85
6.2	Future Work	85
A	Simplification	87
A.0.1	V-path Cancellation	87
A.0.2	Persistence Pairs	88
A.0.3	Theoretical Obstructions to Simplification	90
	Bibliography	93

Notation and terminology

In the following, X is a CW complex, $f : X \rightarrow \mathbb{R}$ is a real-valued function.

Notation

D^n	The Euclidean n -disk, $\{x \in \mathbb{R}^n \mid \ x\ \leq 1\}$.
S^n, S_-^n, S_+^n	The n -sphere $S^n = \{x \in \mathbb{R}^{n+1} \mid \ x\ = 1\}$, where S_-^n and S_+^n are the southern and northern hemispheres respectively.
$e_\lambda, \bar{e}_\lambda$	e_λ is an open cell in the CW complex X , and \bar{e}_λ is its closure.
$e_\alpha \prec e_\beta$	The cell e_α is a face of e_β in X .
$e_\alpha \triangleleft e_\beta$	The cell e_α is a facet of e_β in X ; in other words, a co-dimension one face.
$X_{f \leq k}$	The sub-level set $f^{-1}((-\infty, k]) \subseteq X$ of $k \in \mathbb{R}$.
$\bar{X}_{f \leq k}$	The sub-level complex $\bigcup_{\bar{e} \in X, f(\bar{e}) \leq k} \bigcup_{e \leq \bar{e}} e \subseteq X$ of $k \in \mathbb{R}$.
$X^{(n)}$	The n -skeleton $\{e \in X \mid \dim(e) \leq n\}$ of X .
$B_\varepsilon^{e_\lambda}(x)$	An ε -ball in the cell e_λ around the point $x \in e_\lambda$.
$\text{Face}(X)$	The face poset of cells in X .
$\mu : D \rightarrow U$	An acyclic partial matching on the cells of X , with $D \subset X$ and $U \subset X$.
$\mu_f : D \rightarrow U$	The induced acyclic partial matching of a discrete Morse function $f : X \rightarrow \mathbb{R}$.
$\text{Crit}(\mu), \text{Crit}(f)$	The set of critical cells of an acyclic partial matching $\mu : D \rightarrow U$, and of a discrete Morse function $f : X \rightarrow \mathbb{R}$.

$W^s(c)$	The stable subspace of a critical point $c \in \text{Crit}(\mu)$
$FP(\mu), FP'(\mu)$	The sets of flow-paths and reduced flow-paths of an acyclic partial matching $\mu : D \rightarrow U$.
$\text{Sd}_\mu(X)$	The stable subdivision of X with respect to μ .
$C_*^{CW}(X; A)$	The cellular chain complex of X with co-efficients in a ring A
$H_*^{CW}(X; A)$	The homology groups of X with co-efficients in a ring A
$\pi_n(X, x_0)$	The n -th homotopy group of X based at $x_0 \in X$. In other words, the group of homotopy classes of base-point preserving maps $[S^n, X]$.
$C_*(\mathcal{M}; A)$	The algebraic Morse complex with co-efficients in a ring A
$\text{Form}(X, \mu)$	The Forman complex of X over the acyclic partial matching $\mu : D \rightarrow U$
f_\diamond	The induced discrete Morse function on $\text{Form}(X, \mu)$.
\mathbb{V}	A real persistence module
$d_i(\mathbb{V}, \mathbb{U})$	Interleaving distance between two persistence modules
$\text{epi}(f)$	The epigraph of f
T_f, P_f	The merge tree T_f of f and its associated poset P_f
\hat{f}	The height function over the merge tree $\hat{f} : T_f \rightarrow \mathbb{R}$
$\pi_n[T_f]$	The n -th homotopy merge tree of f
$d_\pi(\pi_n[T_f], \pi_n[T_g])$	The interleaving distance between two n -th homotopy merge trees
n -connected	X is n -connected if $\pi_i(X) = 0$ for all $i \leq n$.
$i^\varepsilon, i_\sigma^\varepsilon$	The shift map $i^\varepsilon : T_f \rightarrow T_f$ by ε on the topological merge tree, and the induced group homomorphism $\pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(X_{i^\varepsilon(\sigma)}, e_{i^\varepsilon(\sigma)})$

Chapter 1

Introduction

In nonsense is strength.

-Kurt Vonnegut, Breakfast of
Champions

On the face of it, computer image analysis and algebraic topology share little in common. In the computer world, images are topologically meaningless - disconnected, discrete sets of rigid points. Algebraic topology* concerns itself with continuous entities, endlessly stretching and morphing between one another.

The story of this thesis is to present and develop existing theory which, miraculously, reconciles these disparate fields.

The over-arching problem is efficiently summarising space. For an image, three dimensional or otherwise, what are the important features? How can we compute and distill them from chaotic garbling? When we see an image, this happens naturally. We don't see a collection of disorganised pixels, but meaningful structures assembling themselves from a mess of information.

Algebraic topology has long had precise mathematical solutions to this problem. Homology and homotopy describe space by its voids, with both theories packaged neatly into computable algebraic structures.

Yet another approach is Morse theory, where real-valued functions reduce space to critical points and gradient flow-lines. Remarkably, this simple tool is

*affectionally regarded as abstract nonsense by friend and foe alike

enough to recover homology.

Near single-handedly invented by Robin Forman in the 1990s [7], *discrete* Morse theory recreates Morse theory in a combinatorial setting. For our purposes, it functions as a powerful preprocessing tool to reduce a space before calculating homology.

Forman himself had flirted with homotopy as well, showing discrete Morse functions induce a CW complex comprised of only critical cells that retains the homotopy type of a space. However, the explicit attaching maps to build this complex are not provided.

Recent publications have examined in depth the interrelation of discrete Morse theory and homotopy theory [15] [16]. In Chapter 2, we utilise the novel tools of [16] to formulate an original proof (Theorem 3.5) determining the attaching maps of the above complex.

There are number of simple geometric ideas which will, if the reader takes them to heart, make the somewhat technical proof more palatable. Figure 1.1 depicts a CW complex, where the arrows represent a *discrete gradient vector field* induced by a discrete Morse function. The cells in the left-most diagrams not paired by any arrow are the *critical cells*.




Figure 1.1: Reducing a CW complex to its critical cells via a discrete gradient vector field.

The arrows essentially provide a map, instructing us how to glue the critical cells together (rightmost diagram). The discrete pairing, however, is slightly too coarse. We use tools from [16] to induce continuous flow-lines that partition the space into stable regions, where the quotient of such flow-lines yields the attaching maps.

We contextualise our results in the work of [18], which describes how images can be reformulated as discrete Morse functions over CW complexes. The combinatorial flavour of discrete Morse theory proves a natural fit for computer calculations. We underscore the viability of our approach by outlining provable original algorithms.

It remains that the resolution homology and homotopy provide is wildly insufficient for our aims. Severely lacking is the ability to distinguish objects of different sizes.

The revelation of *persistent* homology is that one can utilise filtrations as a technique to track information about an underlying metric. Early geometric versions of persistence can be found under the guise of *size theory* in [8], and in the doctoral work [17] of my supervisor, Vanessa Robins.

Persistent homology evolved when people realised that it could be efficiently computed [20] [6], and again when it could be metrized and used statistically [3]. Modern persistence is a fully developed field, with theory at the most general level [1].

In Chapter 4, we use the work of [14] as a basis to extend the notion of persistence to homotopy groups. We introduce the original definition of a *homotopy merge tree*, and show that the main properties of persistent homology - interleaving and stability - follow easily.

There exist successful examples of work implementing this idea [9] [11]. However, our definitions are tailored to fit discrete Morse theoretic methods from Chapter 3. We describe this connection explicitly in Chapter 5, sketching an original algorithm for computing the first homotopy merge tree.

Throughout the work, we endeavour to provide illustrative geometric examples, and contextualise how theory can be applied to image analysis. However, we often present the theory in generality, with the hope that it may find application in other areas.

Chapter 2

Discrete Morse Theory

Discrete Morse theory [7] was formulated by Robin Forman in the 1990s as a combinatorial adaptation of Morse Theory for manifolds. By design, the main theorems of Morse theory all hold in the discrete setting.

In this chapter, we present the relevant background and results. Our approach follows Forman's original paper, but also incorporates the more modern work [16] relating the theory to homotopy.

2.1 Background

2.1.1 CW Complexes

The type of topological spaces on which discrete Morse theory operates are CW complexes. Many varied and complicated definitions of CW complexes exist, so we will firstly take some time to lay down notation that will be used throughout the paper. We follow the approach of [5].

By S^n and D^n , we denote the standard Euclidean n -sphere and n -disk, where $S^{-1} = \emptyset$ and D^0 is a point. We will use E^n to denote the interior of D^n , also written as $D^n \setminus S^{n-1}$ for $n \geq 1$.

Definition 2.1. [5] Let X be a topological space. An n -dimensional cell in X is a subset $e^n \subseteq X$ such that e^n is homeomorphic to E^n .

Different authors refer to cells as being either the above definition, or the closure of e^n , denoted \bar{e}^n , in X . Here, e^n will always be thought of as homeomorphic to the interior of an n -disk.

Definition 2.2. [5] A **CW complex** is a topological space X along with a collection of cells $\{e_\lambda \mid \lambda \in \Lambda\}$ satisfying the following:

1. For each n -cell e_λ^n there is a **characteristic map** $\Phi_\lambda : D^n = D_\lambda^n \rightarrow X$, inducing a homeomorphism $E^n \rightarrow e_\lambda^n$ and sending $\partial D^n = S^{n-1}$ into the union X^{n-1} of cells of dimension less than or equal to $n-1$.
2. The closure \bar{e}_λ of each cell $e_\lambda \in X$ intersects finitely many other cells.
3. X carries the colimit topology with respect to the set $\{e_\lambda \mid \lambda \in \Lambda\}$.

All CW complexes are Hausdorff topological spaces [10]. We refer to the restriction of Φ_λ to ∂D^n as the **attaching map**. The n -**skeleton** $X^{(n)}$ of X is the sub-complex of cells $\{e \in X \mid \dim(e) \leq n\}$.

Remark 2.3. In this thesis, we will assume all CW complexes have a finite number of cells, and all cells are finite dimensional.

In practice, the open sets in the colimit topology of a CW complex X are described via the following condition:

A subset $U \subset X$ is open if the restriction of U to \bar{e}_λ is open for every $e_\lambda \in X$.

Definition 2.4. Let X be a CW complex and $e_\alpha, e_\beta \in X$. If $e_\alpha \subseteq \bar{e}_\beta \setminus e_\beta$, we call e_α a **face** of e_β , writing $e_\alpha \prec e_\beta$. Dually, we call e_β a **coface** of e_α .

In the case that $\dim(e_\alpha) = \dim(e_\beta) - 1$, we call e_α a **facet** of e_β , and write $e_\alpha \triangleleft e_\beta$.

We are often interested in \mathbb{R} -valued functions on CW complexes, particularly their sub-level sets.

Definition 2.5. Let $f : X \rightarrow \mathbb{R}$ be a function on a CW complex. The **sub-level set** of f at $h \in \mathbb{R}$ is

$$X_{f \leq h} := f^{-1}((-\infty, h]).$$

The **sub-level complex** of f at h is the cellular closure of $X_{f \leq h}$, explicitly

$$\bar{X}_{f \leq h} := \bigcup_{\bar{e} \in X, f(\bar{e}) \leq h} \bigcup_{e \triangleleft \bar{e}} e$$

A special kind of cell complex that behaves nicely is given by the following condition from [12]:

Definition 2.6. [12] A CW-complex X is **regular** if each n -cell $\bar{e}_\lambda^n \in X$ is homeomorphic to D^n for all n .

An equivalent condition for being regular is that every characteristic map $\Phi_\lambda : D_\lambda^n \rightarrow X$ is a homeomorphism onto \bar{e}_λ . We can also define regularity locally for a pair $e_\alpha \triangleleft e_\beta$.

Definition 2.7. Let $e_\alpha \triangleleft e_\beta$ and $\Phi_\beta : D^n \rightarrow X$ be the characteristic map for e_β . Then e_α is a **regular face** of e_β if

1. Φ_β restricts to a homeomorphism on $\Phi_\beta^{-1}(e_\alpha)$,
2. The closure $\overline{\Phi_\beta^{-1}(e_\alpha)}$ is a closed n -ball.

Regular CW complexes satisfy an important property, where if a cell face differs in dimension by more than one, we can sandwich at least two cells in between.




Figure 2.1: Three different CW complexes. The first two are different decompositions of the same topological space. Only the middle is regular. The first and last are irregular in different ways; the boundary map of the 2-cell x_2 gets squashed to a point, whereas that of z_2 is wrapped around the 1-cell z_1 twice.

Theorem 2.8. [7] Suppose X is a regular CW complex, and there exists $e_\alpha^{n+k} \succ e_\beta^{n-1}$ such that $k \geq 1$. Then there exist two $(n+k-1)$ -cells e_λ and \tilde{e}_λ such that $e_\lambda \neq \tilde{e}_\lambda$ and

$$e_\alpha^{n+k} \triangleright e_\lambda \succ e_\beta \text{ and } e_\alpha^{n+k} \triangleright \tilde{e}_\lambda \succ e_\beta$$

An immediate corollary following by induction is that, in such a situation, there exist n -cells e_γ and \tilde{e}_γ such that $e_\gamma \neq \tilde{e}_\gamma$ such that

$$e_\alpha^{n+k} \succ e_\gamma \triangleright e_\beta \text{ and } e_\alpha^{n+k} \succ \tilde{e}_\gamma \triangleright e_\beta$$

The topology of a CW complex can be thought of as being inherited from the metric space topology on each disks. In some sense, we can glue these metrics together to get an equivalent one:

Lemma 2.9. [12] *A finite connected CW complex is metrizable*

2.1.2 Discrete Morse Functions

Definition 2.10. [7] Given a CW complex X , a **discrete Morse function** $f : X \rightarrow \mathbb{R}$ satisfies the following properties for all cells $e_\lambda \in X$:

1. If e_λ is an irregular facet of e_α then $f(e_\lambda) < f(e_\alpha)$. Further,

$$|\{e_\alpha \in X \mid e_\lambda \triangleleft e_\alpha \text{ and } f(e_\lambda) \geq f(e_\alpha)\}| \leq 1.$$

2. If e_β is an irregular facet of e_λ then $f(e_\beta) < f(e_\lambda)$. Further,

$$|\{e_\beta \in X \mid e_\lambda \triangleright e_\beta \text{ and } f(e_\lambda) \leq f(e_\beta)\}| \leq 1.$$

Remark 2.11. Here, f is constant on the cells, assigning a real number to each. A good way to think of the above definition is that, for each cell, higher (lower) dimensional cofaces (faces) take higher (lower) values under f , except for at most one exception.

Definition 2.12. [7] Given a discrete Morse function $f : X \rightarrow \mathbb{R}$ on a CW complex X , a cell $e_\lambda \in X$ with $\dim(e_\lambda) = n$ is an n -dimensional critical cell if both

1. $\{e_\alpha \in X \mid e_\lambda \triangleleft e_\alpha \text{ and } f(e_\lambda) \geq f(e_\alpha)\} = \emptyset$
2. $\{e_\beta \in X \mid e_\lambda \triangleright e_\beta \text{ and } f(e_\lambda) \leq f(e_\beta)\} = \emptyset$

For a given discrete Morse function $f : X \rightarrow \mathbb{R}$ we denote its set of critical cells $\text{Crit}(f)$.

Lemma 2.13. [7] *Given a CW complex X and discrete Morse function $f : X \rightarrow \mathbb{R}$, for any cell $e_\lambda \in X$, at least one of the following is satisfied:*

1. $\{e_\alpha \in X \mid e_\lambda \triangleleft e_\alpha \text{ and } f(e_\lambda) \geq f(e_\alpha)\} = \emptyset$
2. $\{e_\beta \in X \mid e_\lambda \triangleright e_\beta \text{ and } f(e_\lambda) \leq f(e_\beta)\} = \emptyset$

Notably, this lemma provides a partition of the cells of X into three separate classes; cells for which set 1 is empty, cells of which set 2 is empty and critical cells for which both sets are empty. This definition can often be quite difficult to work with. Fortunately, the language of discrete vector fields provide an intuitive procedure to generate discrete Morse functions over a cell complex.

Definition 2.14. [7] Let X be a CW-complex. A **discrete vector field** V is a collection of facet/cofacet pairs $\{(e_\lambda^n \triangleleft u_\lambda^{n+1})\}_\lambda$ such that each cell belongs to at most one pair.

Flow-lines of a Morse function in the smooth category have a discrete analogue in V -paths, where we piece together pairs of cells in the discrete vector field.

Definition 2.15. [7] Given a gradient vector field V , an n -dimensional **V-path** is a sequence of cells

$$(e_0 \triangleleft u_0 \triangleright e_1 \triangleleft u_1 \triangleright \dots \triangleright e_k \triangleleft u_k)$$

such that for all $0 \leq i \leq k$, we have $(e_i \triangleleft u_i) \in V$, $e_i \neq e_{i-1}$ and $\dim(e_i) = n$.

We say that a V -path is **closed** if there exists σ_i and σ_j such that $\sigma_i = \sigma_j$ and $i \neq j$.

Definition 2.16. [7] A discrete vector field V is called a **discrete gradient vector field** if there do not exist any closed V -paths.

That this is a *gradient* vector field refers to theorem proved in Forman, providing us with an alternate characterisation of discrete Morse functions:

Theorem 2.17. [7] *A cell complex admits a discrete gradient vector field if and only if it admits a discrete Morse function.*

The constructive proof given in Forman uses the pairings of the discrete vector field to stipulate when the inequalities in definition 1.2. Explicitly, the pairing $\{(e^n \triangleleft u^{n+1})\} \in V$ if and only if the given discrete Morse function $f : X \rightarrow \mathbb{R}$ satisfies $f(e^n) \geq f(u^{n+1})$. Consequentially, any given V -path

$$(e_0 \triangleleft u_0 \triangleright e_1 \triangleleft u_1 \triangleright \dots \triangleright e_k \triangleleft u_k)$$

in a discrete *gradient* vector field descends monotonically in f values

$$f(e_0) \geq f(u_0) > f(e_1) \geq f(u_1) > \dots > f(e_k) \geq f(u_k).$$

A closed V -path in a discrete vector field breaks the monotony of this sequence. The reader may perturb the values of f , say by adding a constant to every value, without altering the pairing; the important notion that the definition captures is that of transitivity.

2.1.3 Partially Ordered Sets

Since its conception, discrete Morse theory has evolved beyond its original form as a real-valued function. Many of the interesting theorems rely solely on properties of the discrete gradient vector field.

Definition 2.18. A **partially ordered set** (or poset) (P, \leq) is a set P with a binary relation \leq satisfying the following for all $p, q, r \in P$:

1. (Reflexivity) $p \leq p$,
2. (Antisymmetry) if $p \leq q$ and $q \leq p$ then $p = q$
3. (Transitivity) If $p \leq q$ and $q \leq r$ then $p \leq r$.

The faces of a CW complex comprise elements of a poset.

Definition 2.19. Let X be a CW complex. The **face poset** $\text{Face}(X)$ consists of the set of cells of X with order relation $e_\alpha \leq e_\beta$ iff $e_\alpha \preceq e_\beta$.

The only non-trivial condition to verify that $\text{Face}(X)$ is a poset is transitivity. If $e_\alpha \preceq e_\beta \preceq e_\gamma$ then $e_\alpha \subseteq \overline{e_\beta} \subseteq \overline{e_\gamma}$, so $e_\alpha \preceq e_\gamma$.

Recent reformulations of discrete Morse theory have recognised the centrality of posets to the definition. For example, the discrete gradient vector field of a discrete Morse function $f : X \rightarrow \mathbb{R}$ can be formulated via posets.

Definition 2.20. [16] An **acyclic partial matching** on X consists of a partition of the cells into three disjoint sets D, U and $\text{Crit}(\mu)$ along with a bijection $\mu : D \rightarrow U$ so that the following conditions hold:

1. **Incidence:** $e \triangleleft \mu(e)$ for $e \in D$ and
2. **Acyclicity:** the transitive closure of the binary relation

$$e \prec_\mu e' \text{ if and only if } e \prec \mu(e')$$

is a partially ordered set on D .

By transitive closure, we mean that the relations $e_i \prec_\mu e_j \prec_\mu e_k$ force the relation $e_i \prec_\mu e_k$.




Figure 2.2: A discrete gradient vector field over a CW complex, or equivalently, an acyclic partial matching $\mu : D \rightarrow U$. The set D comprises of the tails $\{e_i\}$ of each arrow. The heads consist of the set U , paired by the collection $\{\mu(e_i)\}$. The coloured cells are the unpaired critical cells.

Lemma 2.21. *Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function with discrete gradient vector field $V = \{(e_i \triangleleft u_i) \mid f(e_i) \geq f(u_i)\}$. Then the map*

$$\begin{aligned} \mu_f : D &\rightarrow U \\ e_i &\mapsto u_i \end{aligned}$$

is an acyclic partial matching, with $\text{Crit}(\mu_f) = \text{Crit}(f)$.

We call μ_f the induced acyclic partial matching of $f : X \rightarrow \mathbb{R}$.

Proof. The incidence requirement is satisfied immediately from the definition. To show that \prec_μ is a partial order, first note that reflexivity follows from the fact that $e \prec \mu(e)$. Further, transitivity is enforced by definition.

We just need to check that upon taking the transitive closure, we do not break the antisymmetry axiom. Suppose that $e_i \prec_{\mu_f} e_j$ and $e_j \prec_{\mu_f} e_i$ with $e_i \neq e_j$. So there exist sequences

$$e_j \triangleleft \mu(e_j) \succ e_{k_1} \triangleleft \mu(e_{k_1}) \succ \dots \succ e_i$$

and

$$e_i \triangleleft \mu(e_i) \succ e'_{k_1} \triangleleft \mu(e'_{k_1}) \succ \dots \succ e_j$$

such that each pair $(e \triangleleft \mu(e)) \in V$. The above paths imply that $\dim(e_j) \leq \dim(e_i)$ and $\dim(e_i) \leq \dim(e_j)$, so $\dim(e_i) = \dim(e_j)$. This forces the above paths to then be V -paths in the discrete gradient vector field. By concatenating the two, we get a V -path

$$e_j \triangleleft \mu(e_j) \triangleright e_{k_1} \triangleleft \dots \triangleright e_j$$

which is a closed loop, contradicting the fact V is a discrete gradient vector field. So $e_j \neq e_i$, and μ_f is a poset. \square

The results in this thesis will sometimes only require that we have an acyclic partial matching. However, as sub-level sets and sub-level complexes are of central importance, we will refer to the original real-valued function definition when we need it.

2.2 Homology

The central result of discrete Morse theory is that the homology of the space can be computed *solely* from information about the critical points and V -paths of a discrete Morse function.

For the reader unfamiliar with Morse theory, this should seem miraculous; after all, when we only retain these attributes, we often throw away most of the space!

2.2.1 Orientation of V -paths

Before delving into the above statement, we must discuss a technicality in how V -paths are oriented. First recall how orientation is defined for a CW complex:

Definition 2.22. [12] If $\varphi : D^n \rightarrow X$ is the characteristic map of an n -cell e^n , an **orientation of e^n** is a choice of generator for $\varphi_*(H_n(D^n, \partial D^n)) \subseteq H_n(X^n, X^{n-1})$. An **orientation for X** is a choice of generator for each cell.

More explicitly, note that $H_n(D^n, \partial D^n) \cong \mathbb{Z}$ and

$$H_n(X^n, X^{n-1}; \mathbb{Z}) = C_n^{CW}(X^n; \mathbb{Z}) \cong \bigoplus_{e^n \in X} \mathbb{Z}$$

in \mathbb{Z} -coefficients. An orientation picks whether to send the generator of $H_n(D^n, \partial D^n)$ to 1 or -1 in its corresponding summand of $H_n(X^n, X^{n-1})$. In order to define the Morse boundary operator, we need to extend this orientation to V -paths. Forman defines path multiplicity as a way of transferring orientation along a V -path.

Definition 2.23. Let $\mu : D \rightarrow U$ be an acyclic partial matching over a CW complex and

$$\gamma = (e_0 \triangleleft \mu(e_0) \triangleright e_1 \triangleleft \mu(e_1) \triangleright \dots \triangleright e_k)$$

be a V -path. The **multiplicity** of γ is:

$$m(\gamma) = \prod_{i=0, \mu(e_i) \neq 0}^{r-1} \langle \partial^{CW} \mu(e_i), e_i \rangle \langle \partial^{CW} \mu(e_i), e_{i+1} \rangle \in \{\pm 1\}$$

Here, ∂^{CW} is the cellular boundary operator on X , and $\langle _, _ \rangle$ is the ‘dot product’ over the free \mathbb{Z} -module $C_{\dim(e_i)}^{CW}(X; \mathbb{Z})$. That e_i is a regular face of $\mu(e_i)$ implies that this always returns ± 1 , depending only on whether the induced orientation of $\mu(e_i)$ under the boundary operator *agrees* with the choice of orientation on e_i .

2.2.2 Morse Chain Complex

The information about critical cells can be packaged together into an algebraic chain complex, which computes the homology of the space. Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function, and let $\Gamma(e_\alpha, e_\beta)$ denote the set of V -paths from e_α to e_β .

Definition 2.24. [7] The **algebraic Morse complex** \mathcal{M} of (X, f) is a chain complex given by the following data:

1. The chain groups

$$C_n(\mathcal{M}; \mathbb{Z}) := \bigoplus_{\sigma \in \text{Crit}(f)} \mathbb{Z}$$

2. The boundary operators $\partial_n^{\mathcal{M}} : C_n(\mathcal{M}; \mathbb{Z}) \rightarrow C_{n-1}(\mathcal{M}; \mathbb{Z})$ given by

$$\partial_n^{\mathcal{M}}(\sigma^{(n)}) = \sum_{\nu^{(n-1)} \in \text{Crit}(f)} \sum_{\gamma \in \Gamma(\sigma, \nu)} m(\gamma) \nu^{(n-1)}$$

The n -th homology of \mathcal{M} computed via $\text{Ker}(\partial_n^{\mathcal{M}})/\text{Im}(\partial_{n+1}^{\mathcal{M}})$, which we will refer to as $H_n(\mathcal{M}; \mathbb{Z})$.

Theorem 2.25. [7] (*Discrete Morse Homology Theorem*)

$$H_n(\mathcal{M}; \mathbb{Z}) \cong H_n(X; \mathbb{Z})$$

We defer the verification that the above construction is a chain complex, as well as the proof of Theorem 2.25 to a later section. There, we will use the main result of the next chapter to provide a different proof to that of the original appearing in [7].

2.3 Homotopy

V -paths and critical cells are sufficient to recover homology. A natural question is whether this approach works for homotopy.

As in homology, V -paths care about relationships between cells of co-dimension one. Homotopy, on the other hand, requires us to attach cells to cells of *all* lower dimensions.

Flow-paths, as introduced in [16], are a natural modification of V -paths accommodating this difference. In this section, we introduce the notation, tools and results of [16].

2.3.1 Morse Homotopy Theorems

There are two essential theorems about the relationship between critical points and homotopy. Recall that $\bar{X}_{f \leq h}$ is the cellular closure of the sub-level set $X_{f \leq h}$.

Theorem 2.26. [7] *If $a < b$ are real numbers such that $[a, b]$ contains no critical values of f , then*

$$\bar{X}_{f \leq a} \simeq \bar{X}_{f \leq b}$$

Theorem 2.27. [7] *If e is a critical n -cell such that*

- $f(e) \in [a, b]$ and;
- $f^{-1}[a, b]$ contains no other critical points.

Then

$$\bar{X}_{f \leq b} \simeq \bar{X}_{f \leq a} \sqcup_{\partial e} \bar{e}$$

The culmination of the above theorem results in the following corollary, which is of central interest in this thesis:

Corollary 2.28. [7] *M is homotopy equivalent to a CW complex with n -cells corresponding to the critical n -cells of f .*

In his original paper [7], Robin Forman provides an inductive proof of the above statement, without providing explicit attaching maps. As it will appear many times, we will give this canonical cell complex a name:

Definition 2.29. Given a regular cell complex X with an acyclic partial matching $\mu : D \rightarrow U$, the **Forman Complex of X given μ** , denoted $\text{Form}(X, \mu)$, is the cell complex prescribed by Corollary 2.28.

Corollary 2.28, by construction, can then be restated as $\text{Form}(X, \mu) \simeq X$. In the next chapter, we recover an explicit incarnation of the Forman complex, attaching maps included. As is the case for homology, critical points and flow-lines between them will be our main tools.

2.3.2 Induced Flow-lines

Discretization of the space destroys the uniqueness of gradient flow-lines in the smooth case. In [16], the authors present a clever way of reintroducing continuous flow-lines on a cell complex via its discrete morse pairing. This construction yields a surprising amount of information about the space.

The idea is to define linear flow over a general n -disk, then push it forward cell by cell through the characteristic maps. Let

$$\begin{aligned} \zeta : D^{n-1} \times [-1, 1] &\rightarrow D^n \\ (\mathbf{x}, t) &\mapsto (\mathbf{x}, t\sqrt{1 - \|\mathbf{x}\|^2}) \end{aligned}$$

be the function that surjectively smooths the n -cylinder into the n -disk. Then let

$$\begin{aligned} L : [-1, 1] \times [0, 1] &\rightarrow [-1, 1] \\ (x, t) &\mapsto (1 - t)x - t \end{aligned}$$

be the flow that linearly deforms the interval down into $\{-1\}$. Combining these together, we can then define a flow on D^n as $\text{pr}_1 : D^n \times [0, 1] \rightarrow D^n$, via the commutative diagram given below from [16]:

$$\begin{array}{ccc} D^{n-1} \times [-1, 1] \times [0, 1] & \xrightarrow{\zeta \times 1} & D^n \times [0, 1] \\ \downarrow 1 \times L & & \downarrow \text{pr}_1 \\ D^{n-1} \times [-1, 1] & \xrightarrow{\zeta} & D^n \end{array} \quad (2.1)$$




Figure 2.3: An illustration of the flow over D^2 . First, we build a flow on the 1-cylinder as a map $L : D^1 \times [-1, 1] \times [0, 1] \rightarrow D^1 \times [-1, 1]$. Then we use our smoothing function $\zeta : D^1 \times [-1, 1] \rightarrow D^2$ to push it forward to the 2-disk.

Having defined flow on the n -disk, we wish to extend it to a flow over a matched pair in the discrete gradient vector field. Suppose that (u, d) is a discrete morse pairing in (X, μ) , with $\dim(u) = n$. We can continuously deform the characteristic map $\varphi_u : D^n \rightarrow \bar{u}$ to a homeomorphism χ such that $\chi(S_+^n) = d \subseteq \partial u$, where S_+^n is the upper hemisphere on the boundary of D^n and d is the $(n - 1)$ -cell paired with u (See [16]). Lastly, we can extend diagram (1.1) to describe a flow L_u on u :

$$\begin{array}{ccccc}
 D^{n-1} \times [-1, 1] \times [0, 1] & \xrightarrow{\zeta \times 1} & D^n \times [0, 1] & \xrightarrow{\chi \times 1} & \bar{u} \times [0, 1] \\
 \downarrow 1 \times L & & \downarrow \text{pr}_1 & & \downarrow L_u \\
 D^{n-1} \times [-1, 1] & \xrightarrow{\zeta} & D^n & \xrightarrow{\chi} & \bar{u}
 \end{array} \tag{2.2}$$

Visually, a pairing induces a flow as per the figure below:




Figure 2.4: On the left is a cell paired by an acyclic partial matching. The flow defined on a 2-disk is transferred onto the pairing through its characteristic map.

Once we do this to every paired cell in the space, we are interested in seeing where points are sent under the flow:

Definition 2.30. [16] Let $x \in X$ and $\mu : X \rightarrow \mathbb{R}$ be a discrete Morse function. The **flowline** of x , denoted λ_x , is a real number h_x and continuous map $\lambda_x : [0, h_x] \rightarrow X$ such that:

- for all $u \in X$ intersecting the image of λ_x , there exists a point $y \in u$ and a real number $s \geq 0$ such that

$$\lambda_x(t) = L_u(y, t + s)$$

for all admissible $t \geq 0$.

In other words, every point in X has an associated continuous ‘gradient flow line’, whose trajectory agrees with the induced flow along each cell it intersects. We can always rescale h_x to be 1, and will often do so without comment. Some immediate observations to make about flow-lines:

1. The acyclicity of the discrete gradient vector field ensures that flow lines are aperiodic, and hence trace out a contractible image in the space;
2. Every flow-line intersects a critical cell uniquely at its end-point.

Example 2.31. Consider the cell complex with acyclic partial matching as pictured below, with four arbitrary points highlighted in colour:




Figure 2.5: The acyclic partial matching from Figure 2.2, with four points highlighted in colour.

In the Figure 2.6 below, we induce a continuous flow on the entire cell complex as described by the procedure above. The coloured lines then represent the flow-lines of their associated coloured points.

Notably, each flow-line ends at a critical (unpaired) cell, and is aperiodic. In particular, the purple point situated on a critical cell has a trivial flow-line.




Figure 2.6: The induced flow on Figure 2.5, with flow-lines of each coloured point.

2.3.3 Flow-paths and Stable Subdivision

A surprising result from [16] is that the behaviour of continuous flow-lines can be entirely summarised by combinatorial approximations.

However, unlike the case for homology, V -paths alone do not contain enough information to tell us everything. To see all of the *homotopical* information, we need a broader definition of combinatorial flow:

Definition 2.32. [16] Let $\mu : D \rightarrow U$ be an acyclic partial matching on X . A **flow path** of μ is a finite alternating sequence of cells in X of the form

$$\gamma = (e_1 \preceq u_1 \succ e_2 \preceq u_2 \succ \dots \succ e_n \preceq u_n \succ c)$$

such that

1. the cells u_i are distinct elements of U ;
2. either $e_i = u_i$ or $\mu(e_i) = u_i$ for $1 \leq i \leq n$ and
3. c is critical.

Continuing the notation from [16], we let $FP(\mu)$ denote the set of all flow-paths in (X, μ) , and let $\tau(\gamma)$ denote the critical cell at the end of γ .

Every V -path is a flow path, but the converse is not true. The key difference is that flow paths can *descend in dimension* to adjacent cells when we pass from a u_i to an e_{i+1} , whereas V -paths always alternate between co-dimension one cells.

Intuitively, a cell in a CW complex can be attached to cells of any lower dimension, so it's sensible to include information about *all* lower dimension cells in the combinatorial flow.

Lemma 2.33. *Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function on a regular CW complex, with induced acyclic partial matching μ_f . Let*

$$\gamma = (e_1 \preceq u_1 \succ e_2 \preceq u_2 \succ \dots \succ e_n \preceq u_n \succ c)$$

be a flow-path. Then f decreases monotonically along γ .

The regularity condition is essential, as we rely on the following theorem from [7] before we begin the proof:

Theorem 2.34. [7] *Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function over a regular CW complex. If e_λ is an n -cell and $e_\alpha \in X$ such that $e_\alpha \succ e_\lambda$, then there exists $e_\beta \in X$ such that $e_\lambda \triangleleft e_\beta \preceq e_\alpha$ and*

$$f(e_\beta) \leq f(e_\alpha)$$

Proof of Lemma 2.33. Let i be arbitrary and firstly note $f(u_i) = f(e_i)$ if $e_i = u_i$. Now consider the case $u_i = \mu_f(e_i)$. By the definition of μ_f as the acyclic partial matching induced by the discrete gradient vector field of f , we have $f(\mu_f(e_i)) \leq f(e_i)$.

We now show $f(u_i) \geq f(e_{i+1})$. Since $u_i \in U$, there exists some $e_i \in X$ such that $u_i = \mu_f(e_i)$, with $e_i \triangleleft \mu_f(e_i) \succ e_{i+1}$ satisfying $f(e_i) \leq f(\mu_f(e_i))$.

Case 1: Suppose $\dim(e_{i+1}) = \dim(e_i)$. Then, as $\mu_f(e_i)$ is already paired by e_i , the definition of a discrete Morse function implies that $f(e_{i+1}) < f(\mu_f(e_i))$ as required.

Case 2: Suppose that $\dim(e_{i+1}) < \dim(e_i)$. For contradiction's sake, suppose that $f(e_{i+1}) \geq f(\mu_f(e_i))$. By Theorem 2.34 and the fact that X is regular, there exists u_{i+1} such that

$$e_{i+1} \triangleleft u_{i+1} \prec \mu_f(e_i)$$

with $f(u_{i+1}) < f(\mu_f(e_i))$. By our assumption,

$$f(u_{i+1}) < f(\mu_f(e_i)) \leq f(e_{i+1}).$$

By regularity and Theorem 2.8, there exists $\tilde{e}_{i+1} \neq u_{i+1}$ such that

$$\mu_f(e_i) \succ \tilde{e}_{i+1} \triangleright e_{i+1}.$$

Since $f(e_{i+1}) \geq f(u_i)$ and $\tilde{e}_{i+1} \triangleright e_{i+1} \triangleleft u_{i+1}$, it follows from the definition of discrete Morse function that $f(\tilde{e}_{i+1}) > f(e_{i+1})$, else e_{i+1} would be paired with two cells in μ_f . So

$$f(\tilde{e}_{i+1}) > f(e_{i+1}) \geq f(\mu_f(e_i))$$

implying $e_i \triangleleft \mu_f(e_i) \succ \tilde{e}_{i+1}$ with $f(\tilde{e}_{i+1}) \geq f(\mu_f(e_i))$. Noting that $\dim(\tilde{e}_{i+1}) = \dim(e_{i+1}) + 1$, we can repeat the same argument inductively until we reach Case 1, in which we get the desired contradiction. □

For every flow-line in the *continuous* world, we can associate a flow-path in the *combinatorial* world.

Lemma 2.35. [16] *For each $x \in X$ there exists a unique reduced flow path γ_x such that the flow line originating at x is contained entirely within the union of cells in γ_x .*

Here, a flow-path is called **reduced** if it satisfies the algebraic condition

$$e_{i+1} \not\prec \mu^{-1}(u_i)$$

for all i .

Geometrically, once a reduced path leaves a paired cell $\mu^{-1}(u_i)$, it cannot return to its boundary immediately after. This corresponds to how flow-lines are defined; they can't flow back to the boundary of a cell in D that they have just left.

While each flow-line has a unique flow-path, many flow-lines can be assigned to the same flow-path. By partitioning our cell complex into regions who share the same flow-path, we get a well-behaved subdivision of the space.

Theorem 2.36. (Stable Subdivision) [16] *Let*

$$X = \bigcup_{\gamma} e_{\gamma}$$

be a stratification of X , where $e_{\gamma} \subseteq X$ is the set of points in X such that γ is the unique flow-path corresponding to x . Then this stratification is a regular cell decomposition.




Figure 2.7: The stable subdivision of a single cell.

We call the above decomposition of X the **stable subdivision** $Sd_\mu(X)$. For any critical cell $c \in X$, we call the set

$$W^s(c) = \bigcup_{\tau(\gamma)=c} e_\gamma$$

the **stable subspace** of c . An alternative expression for $W^s(c)$ which will often use is given by

$$W^s(c) = \{x \in X \mid \lambda_x(h_x) \in c\}$$

i.e. points whose flow-lines terminate at c . By gluing together all of the λ_x in the stable subspace, we can define a flow on the stable subspace of a critical cell c :

$$L_c : W^s(c) \times [0, 1] \rightarrow W^s(c)$$

such that $c = L_c(W^s(c), 1)$ (See [16] for details).

Example 2.37. To illuminate the flow-path definition, and its connection with stable subdivision, it is worthwhile to work through a simple example. Figure 2.7 shows the rightmost cell in the CW complex from example 2.5. Note that critical cells have been highlighted in colour in the left-most diagram.

The rightmost diagram illustrates how the space is partitioned into its stable subdivision. **Here, each cell corresponds to a reduced flow-path.** To demonstrate, we will identify each cell in the stable subdivision Sd_μ restricted to the stable subspace $W^s(a_2)$, highlighted in red, with its corresponding flow-path (Figure 2.8).

As the concepts of stable-subdivision and flow-paths feature heavily in the next section, the reader is encouraged to refer back to both Lemma 2.35 and Theorem 2.36, and make sure that they can be reconciled with the above example. Namely,




Figure 2.8: The stable subspace of a_2 inside the stable subdivision of Figure 2.7. Theorem 2.36 needs to be emphasised here; it states that the cells in the stable subdivision correspond bijectively with flow-paths.

one should see how the cells correspond to points whose *continuous* flow-lines share the same corresponding combinatorial flow-path.

It is worth noting the absence of the flow-path $e \preceq u \succ a_1 \preceq b_1 \succ a_2$. This flow-path is *reducible*, and hence, does not contribute to the stable subdivision. An easy way to see this is to note that unreduced flow-paths are those which don't correspond to a flow-line, and no flow-lines can flow out from e into u then return to a_1 .

Chapter 3

The Forman Complex

When Robin Forman invented discrete Morse theory, he provided a simple proof that for any discrete Morse function, there *exists* a homotopy equivalent cell complex consisting of only the critical cells (Corollary [2.28](#)).

However, how these cells are attached together is a more complicated story.

In this section, we define an explicit cell complex that realises the attaching maps one Forman predicted and christen it the **Forman Complex**.

While the construction is somewhat painful, knowing the cellular attaching maps turns out to be quite useful. We quickly recover the algebraic Morse complex and re-prove the discrete Morse homology theorem.

It also opens the door to interesting computations, and we provide an algorithm in the second half of this section for computing the 2-skeleton of the Forman Complex.

3.1 Construction

The language of stable subdivision, flow-lines and flow-paths is our tool for building the Forman complex.

We show that it can be constructed as the quotient space of flow-lines, which describe how critical cells eventually become glued to one another.

3.1.1 Main Theorem

Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function, $\mu_f : D \rightarrow U$ the induced acyclic partial matching over a regular CW complex X . Recall the stable subspace

$$W^s(c) = \{x \in X \mid \lambda_x(h_x) \in c\}$$

of $c \in \text{Crit}(\mu_f)$ is the subspace of X whose flow-lines end at some point in c . We have maps

$$W^s(c) \times I \begin{array}{c} \xrightarrow{L_c} \\ \xrightarrow{\pi_{W^s(c)}} \end{array} W^s(c)$$

where $\pi_{W^s(c)}$ is projection onto the first co-ordinate and L_c is the flow on $W^s(c)$ as per [16]. The points $x \in c$ parametrise the flow-lines in its stable subspace $W^s(c)$ in the sense that

$$W^s(c) = \coprod_{x \in c} \pi_{W^s(c)} \circ L_c^{-1}(x) \quad (3.1)$$

This partition is valid since every point in $W^s(c)$ intersects a unique point in c under L_c . There is a natural equivalence relation on $W^s(c)$ given by

$$a \sim b \text{ if there exists } x \in c \text{ such that } a, b \in \pi_{W^s(c)} \circ L_c^{-1}(x)$$

where elements of $W^s(c)/\sim$ can then be written as $[\pi_{W^s(c)} \circ L_c^{-1}(x)]$.

We show that the points in c and the space of flow-lines are the same topologically. Said more formally, the quotient space $W^s(c)/\sim$ of $W^s(c)$ by flow-lines is *homeomorphic* to c .

Quotient topologies can be delicate, and constructing the requisite homeomorphism makes the following lemma somewhat technical.

Lemma 3.1. *Let $W^s(c)$ and \sim be defined as above. Then $W^s(c)/\sim \cong c$.*

Proof. Define a map

$$f : c \xrightarrow{i} W^s(c) \xrightarrow{\pi} W^s(c)/\sim$$

which is continuous by definition, noting that $W^s(c)$ and $W^s(c)/\sim$ have the canonical subspace and quotient topologies. Define another map

$$\begin{aligned} g : W^s(c)/\sim &\rightarrow c \\ [\pi_{W^s(c)} \circ L_c^{-1}(x)] &\mapsto x \end{aligned}$$

which is well defined following from $W^s(c) = \coprod_{x \in c} \pi_{W^s(c)} \circ L_c^{-1}(x)$. By construction,

$$g \circ f = id_c \quad \text{and} \quad f \circ g = id_{W^s(c)/\sim}.$$

Thus, to show that $W^s(c) \cong c$, it is sufficient that g be continuous.

So let $B_\varepsilon^c(x)$ be an open ball in c , with $x \in c$. The primary occupation of the next few pages is simply to show:

$$g^{-1}(B_\varepsilon^c(x)) \text{ is an open set in } W^s(c)/\sim$$

However, this turns out to be more complicated than first appears. By the universal property of quotient spaces, $g^{-1}(B_\varepsilon^c(x))$ is open in $W^s(c)/\sim$ iff $\pi^{-1} \circ g^{-1}(B_\varepsilon^c(x))$ is open in $W^s(c)$:

$$\begin{array}{ccc} W^s(c) & & \\ \downarrow \pi & \searrow^{g \circ \pi} & \\ W^s(c)/\sim & \xrightarrow{g} & c \end{array}$$

As we will deal with many different topologies, it is worthwhile to take the time to carefully describe the open sets of our space:

1. For a CW complex X , $A \subseteq X$ is open if and only if $A \cap \bar{e}$ is open in \bar{e} for all $e \in X$;
2. For a space X , with $V \subseteq W \subseteq X$, V is open in W if and only if there exists $U \in \mathcal{T}_X$ such that $V = W \cap U$; and
3. Together, for a CW complex X , with $V \subseteq W \subseteq X$, V is open in W if and only if there exists $U \subseteq X$ such that $V = W \cap U$ and $\bar{e} \cap U$ is open in \bar{e} for every $e \in X$.

In our context, let $V = \pi^{-1} \circ g^{-1}(B_\varepsilon^c(c))$ and $W = W^s(c)$ (See Fig. 3.1). Our updated goal is to now show that V is open. It is sufficient to show that

there exists $U \subseteq X$ such that $V = W \cap U$ and $\bar{e} \cap U$ is open for every $e \in X$.

Our goal is to construct such a U inductively by defining it over individual cells that intersect $W^s(c)$, making sure that it is open on each closed cell at each step.




Figure 3.1: A discrete Morse pairing on a CW complex. If c is the rightmost critical one cell, the coloured region is $W = W^s(c)$. The red region within is $V = \pi^{-1} \circ g^{-1}(B_\varepsilon(c))$; the pullback along L_c of an open 1-ball $B_\varepsilon(x) \subseteq c$. Note that neither region is open inside the 3-dimensional cell - yet.

First, we specify which cells we are going to do the induction over:

$$\{(c, \emptyset), (d_1, u_1), (d_2, u_2), \dots \mid (d_i, u_i) \cap W^s(c) \neq \emptyset\}$$

where $u_i = \mu_f(d_i)$ for all i .

We order the indices so that the stable subspace *grows* out of the critical cell, one pairing at a time. Explicitly, let the ordering of indices be such that

$$\bar{u}_n \cap \bigcup_{i < n} \bar{u}_i \neq \emptyset \quad \text{and} \quad d_j \prec e_k \Rightarrow j \leq k.$$

for all $e_k \in M$.

To show that such an ordering exists, first note that $\bigcup_i \bar{u}_i$ must be path-connected, as $W^s(c)$ is path-connected and intersects every \bar{u}_i . Suppose that we have chosen the first $(n - 1)$ -cells.

If we cannot chose u_n such that $\bar{u}_n \cap \bigcup_{i < n} \bar{u}_i \neq \emptyset$, then we can partition $\bigcup_i \bar{u}_i$ into disconnected sets $\bigcup_{i < n} \bar{u}_i$ and $\bigcup_{i \geq n} \bar{u}_i$, contradicting path connectedness.

From the finite set of eligible (d_n, u_n) , we choose a pair of cells minimal in dimension. This ensures that any cofaces of d_n lying in some pair (d_k, u_k) of greater dimension are added later on.

From this set, we additionally require that d_n be minimal with respect to the partial order \prec_μ in Definition 2.20. In other words, if $d_n \prec_\mu(d_k) = u_k$ for some (d_k, u_k) of the same dimension, then $d_n \prec_\mu d_k$ and $k \geq n$. The second condition is hence satisfied. In particular, this ensures that (d_n, u_n) is a free pair upon insertion.

Define a cover and partial cover of $W^s(c)$ as

$$M := \bigcup_i \bar{u}_i \quad \text{and} \quad M_n := \bigcup_{i \leq n} \bar{u}_i$$

respectively. We construct a subset $U \subseteq X$ satisfying

$$M \cap V = M \cap W^s(c) \cap U$$

with the restriction of U open on each closed cell in M . Our induction argument is as follows:

1. (Base Case) There exists $U_0 \subseteq X \cap M_0$ such that $M_0 \cap V = M_0 \cap W^s(c) \cap U_0$ and $\bar{e} \cap U_0$ is open in \bar{e} for every $e \in X \cap M_0$.
2. (Induction Step) If there exists $U_{i-1} \subseteq X \cap M_{i-1}$ such that

$$M_{i-1} \cap V = M_{i-1} \cap W^s(c) \cap U_{i-1}$$

and $\bar{e} \cap U_{i-1}$ is open in \bar{e} for every $e \in X \cap M_{i-1}$, then there exists $U_i \subseteq X \cap M_i$ such that

$$M_i \cap V = M_i \cap W^s(c) \cap U_i$$

and $\bar{e} \cap U_i$ is open in \bar{e} for every $e \in X \cap M_i$.

For the base case, note first that $M_0 = \bar{c}$. Then simply define

$$U_0 := V \cap M_0 = V \cap \bar{c} = B_\varepsilon^c(x).$$

Then $B_\varepsilon^c(x)$ is open in \bar{e} for each $e \in \bar{c}$ by definition, and

$$M_0 \cap V = B_\varepsilon^c(x) = M_0 \cap W^s(c) \cap U_0$$

as required.

For the induction step, assume there exists $U_{i-1} \subseteq X \cap M_{i-1}$ such that $M_{i-1} \cap V = M_{i-1} \cap W^s(c) \cap U_{i-1}$ and $\bar{e} \cap U_{i-1}$ is open in \bar{e} for every $e \in X \cap M_{i-1}$. We would like to construct $U_i \subseteq X \cap M_i$ such that $M_i \cap V = M_i \cap W^s(c) \cap U_i$ and $\bar{e} \cap U_i$ is open in \bar{e} for every $e \in X \cap M_i$.

Define $S := \bar{u}_i \cap U_{i-1}$. Our goal is to extend S to be open on all of the cells in \bar{u}_i . We first extend S to be open on the boundary $(\partial u_i - d_i)$. For example, if there exists $e \in (\partial u_i - d_i)$ such that $e \notin U_{i-1}$, $\bar{e} \cap V$ may not yet be open in \bar{e} .




Figure 3.2: The situation where we attach cells from $\partial u_i - d_i$. On the left, we have the set M_{i-1} . Here, $\bigcup_{e'} S_{e'}$ is the red set $V \cap M_{i-1}$, but extended to open balls on the new cells. We extend in such a way that restricting $\bigcup_{e'} S_{e'}$ to $W^s(c) \cap M_{i-1}$ (coloured region of left figure) agrees with $V \cap M_{i-1}$.

For each $\{e' \in \partial u_i \mid \partial e' \in M_{i-1}, e' \notin M_{i-1}\}$ and $x \in e' \cap U_{i-1}$, we can choose $\varepsilon_x \in \mathbb{R}_{>0}$ such that

$$S_{e'} := \bigcup_{x \in e' \cap U_{i-1}} B_{\varepsilon_x}^{e'}(x), \text{ where } S_{e'} \cap W^s(c) = e' \cap V$$

and $S_{e'}$ is open in \bar{e}' (See Fig. 3.2). Then $\bigcup_{e'} S_{e'}$ is open in \bar{e} for all $e \in (\partial u_i - d_i)$.

Having taken care of the cells in $\partial u_i - d_i$, we lift $\bigcup_{e'} S_{e'}$ to be an open set S_{u_i} on the entirety of \bar{u}_i , with the additional constraint that $S_{u_i} \cap W^s(c) = \bar{u}_i \cap V$. Consider Diagram 3.2, where the homeomorphism χ is the characteristic map of u_i .

$$\begin{array}{ccc}
 D^n & \xrightarrow{\chi} & \bar{u} \\
 \pi_{D^n} \uparrow & & \uparrow \pi_{\bar{u}} \\
 D^n \times I & \xrightarrow{\chi \times I} & \bar{u} \times I \\
 pr^{-1} \uparrow & & \uparrow L_u^{-1} \\
 S_-^{n-1} & \xrightarrow{\chi} & \partial u_i - d_i
 \end{array} \tag{3.2}$$

Now let $S_{u_i} = \pi_{\bar{u}} \circ L_u^{-1}(\bigcup_{e'} S_{e'})$. Then, by the definition of $V, W^s(c)$ and the flow L_u^{-1} , we have that $S_{u_i} \cap W^s(c) = \bar{u}_i \cap V$.




Figure 3.3: The red region $V \cap M_i$ from Figure 3.1 was not open in the 3-cell. Here, the red region is \tilde{U} , an extension of $V \cap M_i$ into an open set retaining the equality $\tilde{U} \cap W^s(c) \cap M_i = \tilde{U} \cap V$. \tilde{U} can be thought of as the pull back of $\bigcup_{e'} S_e$ under the flow L_{u_i} on u_i .

To see that S_{u_i} is open in \bar{u}_i , recall first that S_{u_i} is open in $\partial u_i - d_i$. As χ is a homeomorphism, we have that $\chi^{-1}(V_{u_i})$ is open in $\chi^{-1}(\partial u_i - d_i) = S_-^{n-1}$. The lift of an open set in S_-^{n-1} under $\pi_{D^n} \circ L_{D^n}^{-1}$ is the intersection of open cylinders in \mathbb{R}^n with D^n , and hence open in D^n .

The commutativity of Diagram 3.2 shows S_{u_i} is open in \bar{u} . Since d_i is a free face of u_i in M_i , $S_{u_i} \cup U_{i-1} \cap e$ is open in \bar{e} for all $e \in M_i$. By construction,

$$M_i \cap V = M_i \cap W^s(c) \cap (S_{u_i} \cup U_{i-1})$$

which concludes the induction step.

We now have a set $\tilde{U} \subset X$ such that $V = W^s(c) \cap \tilde{U}$, and $\tilde{U} \cap e$ is open in \bar{e} for all $e \in M$. However, in order to be an open set in X , \tilde{U} needs to be open on the closure of *all* cells.

There may exist some cells outside of M that intersect \tilde{U} only on their boundary. We then grow \tilde{U} to be open in all of X by growing ε -balls around the points who lie in closure of cells outside of M (See Fig. 3.4).

Now, \tilde{U} satisfies all the requirements to make $V = \pi \circ g^{-1}(B_\varepsilon(x))$ into an open set in $W^s(c)$. So our function g must be continuous, and hence,

$$W^s(c)/\sim \cong c.$$

□




Figure 3.4: The red region \tilde{U} needs to be open on the closure of *all* cells in X . We expand it slightly, so that it is also open on the right-hand cell, which was not a member of our cover M .

Now, we quotient the entire space all at once. The stable subspaces *become* the cells, binding together to yield the Forman complex. Before commencing the proof of 3.5, we state two key theorems we will use.

Theorem 3.2. [19] *Let $f : X \rightarrow Y$ be proper and surjective, where X and Y are 0-connected, locally compact, separable metric spaces, X is locally n -connected, and for each $y \in Y$, $f^{-1}(y)$ is $(n - 1)$ -connected and locally $(n - 1)$ -connected. Then*

1. *the induced map $f_* : \pi_k(X) \rightarrow \pi_k(Y)$ is an isomorphism for all $0 \leq k \leq n - 1$ and surjective for $k = n$.*

Remark 3.3. Here, by locally n -connected we mean that for all $x \in X$, there exists some open neighbourhood U such that $\pi_k(U, x) \cong 0$ for all $0 \leq k \leq n$.

Theorem 3.4. (Whitehead's Theorem) [10] *If $f : X \rightarrow Y$ is a map between CW complexes such that $f_* : \pi_n(X) \rightarrow \pi_n(Y)$ is an isomorphism for all $n \geq 0$, then f is a homotopy equivalence.*

Theorem 3.5. (Forman Complex Theorem) *Let X be a regular path-connected cell complex with acyclic partial matching $\mu : D \rightarrow U$. Let X / \sim_μ be the quotient space of X where $a \sim_\mu b$ if there exists $c \in \text{Crit}(\mu)$ and $x \in c$ such that $a, b \in L_c^{-1}(x)$. Then*

1. *X / \sim_μ has a cell complex structure, where the cells are given by $W^s(c) / \sim_\mu$ and the attaching maps are inherited from the quotient space, and*
2. *X / \sim_μ is homotopy equivalent to X .*

Proof. For (1), first note that the stable subspaces partition X , so proving that each $W^s(c)/\sim_\mu$ is an $\dim(c)$ -cell with valid characteristic map provides a CW decomposition of X/\sim_μ .

Assume that we have shown the result for the $(n-1)$ -skeleton of X/\sim_μ . Consider an n -cell c , with associated stable subspace $W^s(c)$. Let $\Phi : D^n \rightarrow X$ be the characteristic map of c . We claim that

$$D^n \xrightarrow{\Phi} X \xrightarrow{\pi} X/\sim_\mu$$

is a characteristic map for the n -cell $W^s(c)/\sim_\mu$. Continuity follows from the continuity of both maps in the composition. $\Phi(D^n - S^{n-1})$ maps homeomorphically onto c by the characteristic map Φ . Lemma 3.1 showed that

$$c \cong W^s(c)/\sim_\mu$$

via the projection map. It follows that $\pi \circ \Phi(D^n - S^{n-1})$ maps homeomorphically onto $W^s(c)/\sim_\mu$.

Now we show that $\pi \circ \Phi(S^{n-1})$ lies in the $(n-1)$ -skeleton of X/\sim_μ . This is equivalent to showing that boundary of c is contained within the stable subspaces of critical cells of dimension lower than $(n-1)$.

Recall that the stable subdivision stratifies X into regions e_γ , where $e_\gamma \subseteq X$ is the set of points $x \in X$ such that γ is the unique flow-path corresponding to x . Lemma 2.35 tells us that the flow-line of x ends at the critical cell $\tau(\gamma)$.

Hence, we need to show for the stable subdivision of the boundary of c , the associated flow paths of each subdivision terminate at a critical cell of dimension less than or equal to $(n-1)$. This is certainly true, as a flow-path leaving the boundary of c can only decrease in dimension. So $\pi \circ \Phi(S^{n-1})$ is attached to the $(n-1)$ -skeleton, and X/\sim_μ is a CW complex.

Now we show that X/\sim_μ is homotopy equivalent to X . Namely, we show that the projection map

$$\begin{aligned} \pi : X &\rightarrow X/\sim_\mu \\ x &\rightarrow [L_c(x, 1)] \end{aligned}$$

is a homotopy equivalence, where c is the unique critical cell such that $x \in W^s(c)$. First we show that π meets the conditions of Theorem 3.2.

All CW complexes are Hausdorff, separable metric spaces and locally contractible* [10]. Further, every continuous map from a compact space to a Hausdorff space is proper, and the projection map is surjective by definition.

Let $[y] \in X / \sim_\mu$, picking representative $y \in X$ such that it lies in a critical cell c . The fibres $\pi^{-1}([y])$ consist of all the flow-lines ending at y . By letting each point in $\pi^{-1}([y])$ move along its unique flow-line to y , we get a deformation retraction of $\pi^{-1}([y])$ onto y . So $\pi^{-1}([y])$ is contractible, and hence, n -connected for $n \geq 0$.

As X is locally contractible, Theorem 3.2 holds for arbitrarily high n , so the induced map $\pi_* : \pi_n(X) \rightarrow \pi_n(X / \sim_\mu)$ is an isomorphism for $n \geq 0$. Then π is a homotopy equivalence by Theorem 3.4. \square

The CW complex described by Proposition 3.5 has cells in bijection with the critical cells of (X, μ) , and is homotopy equivalent to X . Hence, this construction is the CW complex given in Corollary 2.28, namely is $\text{Form}(X, \mu)$, up to cellular equivalence.

Definition 3.6. Let X be a regular CW complex with acyclic partial matching $\mu : D \rightarrow U$. Then

$$\text{Form}(X, \mu) := X / \sim_\mu$$

Remark 3.7. The proof of part (1) shows that the stable subdivision of each critical cell's boundary determines the attaching map. Thus, knowing the flow-paths out of the boundary of each critical cell is enough to reconstruct $\text{Form}(X, \mu)$.

The following corollaries describe how $\text{Form}(X, \mu)$ retains the key properties of the original space and discrete Morse function. It inherits a canonical irregular discrete Morse function with the same critical values as before, the sub-level sets of which are homotopy equivalent to the original complex.

Accordingly, the Forman complex can be thought of as the minimal cell complex capturing all homotopical changes of a filtration by the discrete Morse function.

*Here, locally contractible means that a CW complex is locally n -connected for all n .

Corollary 3.8. *Let X be a regular CW complex and $f : X \rightarrow \mathbb{R}$ a discrete Morse function with induced acyclic partial matching μ_f . Then there exists a discrete Morse function $f_\diamond : \mathbf{Form}(X, \mu_f) \rightarrow \mathbb{R}$ such that*

$$\begin{array}{ccc} \mathbf{Crit}(f) & & \\ \downarrow \cong & \searrow f & \\ \mathbf{Form}(X, \mu_f) & \xrightarrow{f_\diamond} & \mathbb{R} \end{array}$$

commutes, where each cell in $\mathbf{Form}(X, \mu_f)$ is critical.

Here the bijection $\mathbf{Crit}(f)$ to $\mathbf{Form}(X, \mu_f)$ is a bijection on the set of cells given by sending a critical cell c of f to its corresponding cell $W^s(c)/\sim_{\mu_f}$.

Proof. Recall that an irregular discrete Morse function is an \mathbb{R} -valued function over the cells that satisfies the axioms of a discrete Morse function on regular cells, while making every irregular cell critical. We show the stronger statement that *every* cell is critical with respect to the induced map.

It is sufficient to show for each cell $x \in \mathbf{Form}(X, \mu_f)$, $f_\diamond(x) > f_\diamond(y)$ for any $y \in \partial x$. If $y \in \partial x$ in $\mathbf{Form}(X, \mu_f)$, then by construction there must have existed some flow-path down to y from the boundary of x in the original cell complex. Since the value of f decreases monotonically along flow paths, $f(x) > f(y)$ in (X, f) . By the definition of f_\diamond in the commutative diagram, $f_\diamond(x) > f_\diamond(y)$. \square

Recall that, for any $h \in \mathbb{R}$, $X_{f \leq h}$ denotes the sub-level set $f^{-1}((-\infty, h])$ and $\overline{X}_{f \leq h}$ its cellular closure.

Corollary 3.9. *For all $h \in \mathbb{R}$, $\mathbf{Form}(X, \mu_f)_{f_\diamond \leq h} \simeq \overline{X}_{f \leq h}$.*

Proof. Note first that $\mathbf{Form}(X, \mu_f)_{f_\diamond \leq h} = \overline{\mathbf{Form}(X, \mu_f)_{f_\diamond \leq h}}$. This follows from the fact that each cell is critical, so if a cell is in the sub-level set, its boundary must also be.

$\overline{X}_{f \leq h}$ is a regular cell complex with discrete Morse function given by restricting $f : X \rightarrow \mathbb{R}$ to sub-complex $f \circ i : \overline{X}_{f \leq h} \rightarrow \mathbb{R}$. Hence, we can apply Theorem 3.5 to $\overline{X}_{f \leq h}$, where we claim that $\mathbf{Form}(\overline{X}_{f \leq h}, \mu_{f \circ i}) = \mathbf{Form}(X, \mu_f)_{f_\diamond \leq h}$. We know that both $\mathbf{Form}(\overline{X}_{f \leq h}, \mu_{f \circ i})$ and $\mathbf{Form}(X, \mu_f)_{f_\diamond \leq h}$ have cells corresponding to critical cells in (X, f) with values below h , so we need only show that the attaching maps




Figure 3.5: A regular CW complex, homotopy equivalent to an annulus, with an acyclic partial matching, and its partition into stable subspaces

are the same.

For $\text{Form}(X, \mu_f)_{f \leq h}$, the attaching maps are determined by stably subdividing the boundary of each critical cell in X , which depends only on flow-paths out of the boundary. As flow-paths descend in f value by Lemma ??, for any critical cell $x \in X_{f \leq h}$, every flow-path out of x is contained in $\overline{X}_{f \leq h}$. Hence, these same flow-paths determine the attaching maps for $\text{Form}(\overline{X}_{f \leq h}, \mu_{f \circ i})$. Applying Theorem 3.5, we then have

$$\overline{X}_{f \leq h} \simeq \text{Form}(\overline{X}_{f \leq h}, \mu_{f \circ i}) = \text{Form}(X, \mu_f)_{f \leq h}.$$

□

Example 3.10. Consider the cell complex shown in Figure 3.5. The arrows indicate the discrete Morse pairing, with the critical cells labelled and highlighted with colour.

The reader is encouraged to first imagine how the pairing induces continuous flow-lines on the space. On the right, we show the partitioning of the space into its stable subspaces, whose colours correspond to their associated critical cells.

Remark 3.11. Note here that this partition is **not** the complete stable subdivision, which would contain more internal cell division in addition to those on the boundary of each stable subspace.

Now, we can depict $\text{Form}(X, \mu)$ as the quotient space in Figure 3.6. Here, the stable subspaces have been collapsed down onto the boundary of the critical cell e . The figure on the right is an equivalent space, with the regions representing




Figure 3.6: Two equivalent representations of the quotient space $\text{Form}(X, \mu)$. Note, the space is still a cell complex after the quotient operation, with cells corresponding to critical cells in the original matching.




Figure 3.7: When glued together, we can see that $\text{Form}(X, \mu)$ has the same homotopy type as X .

the stable subspace of a zero dimensional cell collapsed down to a point.

When we glue along the prescribed cells, we get the space depicted below. Note, $\text{Form}(X, \mu)$ is an *open* cylinder; hence, $\text{Form}(X, \mu)$ is homotopy equivalent to the annulus, which is homotopy equivalent to the original space X , as predicted.

3.2 Applications

3.2.1 Homology

One immediate application of Theorem 3.5 is that we can quickly reprove, and gain intuition about, the discrete Morse homology theorem. Recall from the previous section:

Theorem. [7] (Discrete Morse Homology Theorem)

$$H_n(\mathcal{M}; \mathbb{Z}) \cong H_n(X; \mathbb{Z})$$

The original proof given in [7] works through a somewhat technical construction of an intermediary chain complex, which is then shown to be chain isomorphic to both \mathcal{M} and $C_n^{\text{Sing}}(X, \mathbb{Z})$. In our case, let (X, μ) be an acyclic partial matching. Then

Observation 3.12. $C_*^{CW}(\text{Form}(X, \mu); \mathbb{Z}) = C_*(\mathcal{M}; \mathbb{Z})$

In other words, the cellular chain complex of the Forman complex is the algebraic Morse complex. Note that the cellular chains are both the free \mathbb{Z} -modules generated over $\text{Crit}(\mu)$.

For the boundary operators, observe that $\partial^{\mathcal{M}}(e_\alpha) = \pm e_\beta$ iff there exists a positively (or negatively) oriented V -path from e_α to e_β iff e_β is quotiented to the cellular boundary of e_α in the Forman complex. Summing this over all V -paths leads to the above observation.

Corollary 3.13. (*Discrete Morse Homology Theorem*)


$$H_n(\mathcal{M}; \mathbb{Z}) \cong H_n(X; \mathbb{Z})$$

Proof. From the main theorem, we know that $\text{Form}(X, \mu) \simeq X$, which implies that $H_n(\text{Form}(X, \mu); \mathbb{Z}) \cong H_n(X; \mathbb{Z})$. Then, as $C_*^{CW}(\text{Form}(X, \mu); \mathbb{Z}) = C_*(\mathcal{M}; \mathbb{Z})$, we have the equation:

$$H_n(\mathcal{M}; \mathbb{Z}) \cong H_n(\text{Form}(X, \mu); \mathbb{Z}) \cong H_n(X; \mathbb{Z}).$$

□

Example 3.14. Consider the example of (X, μ) and $\text{Form}(X, \mu)$ from the previous section:



Both chain groups are generated by critical cells:

$$\begin{aligned} C_2^{CW}(\text{Form}(X, \mu); \mathbb{Z}) &= C_2(\mathcal{M}; \mathbb{Z}) = \mathbb{Z}\{e\} \\ C_1^{CW}(\text{Form}(X, \mu); \mathbb{Z}) &= C_1(\mathcal{M}; \mathbb{Z}) = \mathbb{Z}\{a, b, c\} \\ C_0^{CW}(\text{Form}(X, \mu); \mathbb{Z}) &= C_0(\mathcal{M}; \mathbb{Z}) = \mathbb{Z}\{p, q\} \end{aligned}$$

For the boundary operators:

1. There are two V -paths from c to p , each with opposite orientation; the same is true for the two V -paths from b to q . There is a single V -path from a to q and also from a to p .
2. There are single V -paths from e to b and e to c . There are two V -paths of opposite orientation from e to a .

It is left to the reader to see that this coincides with the boundary operator $\partial : C_n^{CW}(\text{Form}(X, \mu); \mathbb{Z}) \rightarrow C_{n-1}^{CW}(\text{Form}(X, \mu); \mathbb{Z})$ for all n .

3.2.2 Image Partitioning and Connectivity

Recent work has applied discrete Morse theory to computer image analysis. Greyscale three-dimensional images can be represented as a collection of discrete *voxels* taking real values.

These voxels comprise the vertex set of a three-dimensional abstract cubical complex. To reason about homology and homotopy, we substitute this abstract representation for the original image.

In [18], the authors provide an algorithm extending voxel values to a discrete Morse function over the entire cubical complex. Later, in [4], the same authors use this algorithm to assist in partitioning images into regions based on the flow of the discrete gradient vector field. The key definitions from [4] are as follows:

Definition 3.15. Let e^0 be a critical 0-cell, X be a finite dimensional cubical complex and $f : X \rightarrow \mathbb{R}$ a discrete Morse function. The **basin of** e_0 $\mathcal{B}_X(e^0)$ is the maximal sub-complex of X that has a simple homotopy collapse onto e_0 along the discrete gradient vector field.

An important feature of basins is that they partition the vertices of X :

Lemma 3.16. [4] *Each vertex of X is in exactly one basin.*

This lemma does not, however, hold for higher dimensional cells. We give a special name to 1-cells that are not in any basins:

Definition 3.17. [4] A **bridge** is a 1-cell not contained in any basin.




Figure 3.8: Each coloured cell is a critical point of the overlaid discrete gradient vector field. The vertices are partitioned into the basins of critical zero cells. How should we define basins of 1-cells?

We wish to define higher dimensional analogues of basins. Basins partition the 0-cells is because each 0 has a unique V -path exiting it. Such a property does not hold for higher dimensional cells, so there is ambiguity about how we would partition the space (See Fig 3.8).

The stable subdivision solves this problem (See Fig. 3.9), dividing cells until the complex can be partitioned into stable subspaces. It is precisely these subspaces that become critical cells in the Forman Complex. The cells then have an important geometric interpretation, representing how higher dimensional basins are connected to one another.




Figure 3.9: The stable subdivision Sd_μ of the cubical complex from Figure 3.8 and the corresponding Forman complex.

By replacing the original cubical complex with the Forman complex, we accomplish the two aims:

1. We reduce the space to minimal number of meaningful cells.

2. We canonically partition the space into stable regions.

Further, since homotopy type is preserved, we can still faithfully analyse the homotopy type of the original complex.

3.3 Computing Attaching Maps

The usefulness of the above theorems as so far been mainly theoretical. In this section, we present an algorithm to compute the attaching maps of the 2-skeleton of the Forman complex.

This opens the door to pre-processing a space before calculating its fundamental group. Better still, each sub-level set of the induced filtration is homotopy equivalent to the original sub-level sets, a property we will exploit in Chapter 3.

3.3.1 Embedding Functions

An indispensable tool for computing attaching maps, as defined in [16], is the notion of an **embedding function**.

Let $\mu : D \rightarrow U$ be an acyclic partial matching over a cell complex X . Recall that $FP(\mu)$ is the set of flow-paths, and $FP'(\mu)$ the set of reduced flow-paths.

Definition 3.18. [16] Let $\gamma = (e_1 \geq u_1 \succ \dots < u_m \geq c)$ and $\gamma' = (e'_1 \preceq \dots \preceq u'_n \succ c)$ be flow-paths. An **embedding function** from γ to γ' is a strictly monotone function

$$\varphi : \{0, 1, \dots, k\} \rightarrow \{0, 1, \dots, n+1\}$$

where $1 \leq k \leq m+1$, such that the following are satisfied:

1. $\varphi(0) = 0$ and $\varphi(k) = n+1$,
2. $u_j = u'_{\varphi(j)}$ for $1 \leq j < k$,
3. $e_j \preceq e'_p$ for $i \leq j \leq k$ and $\varphi(j-1) < p \leq \varphi(j)$.

In this paper, whenever $\varphi(j) > \varphi(j-1) + 1$, we will say φ **lifts** at j . We can define a relation on flow-paths $\gamma \Rightarrow \gamma'$ whenever there exists an embedding function between them. We follow the convention in [16] that $e_{n+1} = c$.




Figure 3.10: Recall the sub-complex of the stable subdivision Sd_μ of Figure 2.7. Theorem 3.19 asserts that every for every face relation, there exists an embedding function between the corresponding flow-paths.

Two important properties of embedding functions given in [16] are:

1. The relation above defines a partial order on $FP(\mu)$, of which reduced flow-paths $FP'(\mu)$ are a sub-poset.
2. If there exists an embedding function between two reduced flow-paths, it is unique.

Recall that for any regular cell complex X , $\text{Face}(X)$ denotes its induced face poset. Further, if $\mu : D \rightarrow U$ is an acyclic partial matching on X , then $Sd_\mu(X)$ denotes the stable subdivision.

Theorem 3.19. [16] *The map $FP'(\mu) \rightarrow \text{Face}(Sd_\mu(X))$ from reduced flow paths to cells in the stable subdivision given by $\gamma \rightarrow e_\gamma$ is an isomorphism of posets.*

This implies that if flow-paths are related by an embedding function $\gamma \Rightarrow \delta$, then the cell represented by γ in the stable subdivision is in the boundary of the cell represented by δ . Thus, embedding functions allow us to reason combinatorially about the adjacency relations between faces in the stable subdivision.

3.3.2 The 2-skeleton

When we construct the Forman complex, we examine where the boundary of each critical cell is sent under flow to determine its attaching map. In particular, the stable subdivision of the boundary tells us which region should be sent to which critical cell, with its associated flow-path determining orientation.

In the case of a 1-cell, this is easy; unique flow-paths emanate from two boundary 0-cells, and each terminate at critical 0-cells. For the higher dimensional cells, the situation is more complicated; many different and complicated flow-paths can move out of the boundary.

For the 2-skeleton, our primary focus is working out the stable subdivision of 1-cells on the boundary of critical 2-cells, and where each region is mapped to under flow. At first glance searching for embedding functions seems computationally intractable. However, in the case of the 2-skeleton, a series of lemmas will make the task more manageable.

For the remainder of this section, let X be a regular cell complex and $\mu : D \rightarrow U$ be an acyclic partial matching over X . Let $e_1 \in X$ be a 1-cell on the boundary of a critical cell.

Corollary 3.20. *For a cell $e_\gamma \in Sd_\mu(e_1)$ with associated flow-path $\gamma \in FP'(\mu)$ and $e_\gamma \subset \bar{e}_1$, there can exist at most 2 embedding functions relating e_γ to other cells.*

Proof. The stable subdivision of a 1-cell can contain only 0-cells and 1-cells. Hence, if $\dim(e_\gamma) = 0$ it can be in at most the boundary of two 1-cells, and if $\dim(e_\gamma) = 1$ it can have at most two 0-cells in its boundary. The result then follows directly Theorem 3.19. \square

Lemma 3.21. *(Repetition Lemma) If $e_\gamma \subseteq \bar{e}_1$ with*

$$\gamma = (e_1 \preceq u_1 \succ \dots \preceq u_n \succ c) \in FP'(\mu)$$

then $e_i = u_i$ at most once and $\dim(u_i) = 1$.

Proof. Suppose i is the first integer such that $u_i = e_i$. We know that γ must start at e_1 in order for $e_\gamma \subseteq \bar{e}_1$. If $\dim(u_i) = 2$, there exists a nontrivial subsequence of γ given by $u_{i-1} \succ u_i \preceq u_i$. This implies that $2 = \dim(u_i) < \dim(u_{i-1})$, but each cell in the flow-path starting at a 1-cell can be at most dimension 2, so we get a contradiction.

If, on the other hand, $\dim(u_i) = 0$, then the subsequence of γ given by $u_i \preceq u_i \succ u_{i+1}$ tells us that $0 = \dim(u_i) > \dim(u_{i+1})$, which is also a contradiction. So we must have that $\dim(u_i) = 1$, which implies that $\dim(e_k) = 0$ for all $k > i$. By the above reasoning, there then cannot exist an e_k such that $e_k = u_k$ for $k > i$, so there can only be one instance of $e_i = u_i$. \square

Lemma 3.22. (*Lift Lemma*) *Suppose*

$$\gamma = (e_1 \preceq u_1 \succ e_2 \preceq u_2 \succ \dots \preceq u_n \succ c)$$

and

$$\gamma' = (e_1 \preceq u_1 \succ e'_2 \preceq u'_2 \succ \dots \preceq u'_m \succ c)$$

are reduced flow-paths such that $\gamma \Rightarrow \gamma'$ with embedding function φ lifting at i . Then

- (a) γ and γ' agree elsewhere,
- (b) $\dim(e_i) = 0$,
- (c) Either $e_{\varphi(i)} = u_{\varphi(i)}$ or $e_{\varphi(i)-1} = u_{\varphi(i)-1}$,
- (d) e_i is the first zero cell in γ .

This lemma describes the structure of embedding functions that lift ending at the same critical cell. (a) states that the lift at i is the only place γ and γ' differ; (b), (c) and (d) describe where lifts occur and how they end. Thus, the requirement that e_1 is a 1-cell turns out to impose quite a lot of structure on γ and γ' , providing shortcuts for computing the reduced flow-path poset $FP'(\mu)$.

Proof of (a). Suppose that i' is the first time φ lifts. We show that γ and γ' agree elsewhere, so $i' = i$. Lifting first at i' implies that, for $k < i'$, $\varphi(k) = k$ and $u_k = u'_{\varphi(k)} = u'_k$. Suppose additionally that there exists a minimal $j < i'$ such that $e_j \neq e'_j$. We then have two subsequences of γ and γ' given by $u_{j-1} \succ e_j \preceq u_j$ and $u_{j-1} \succ e'_j \preceq u_j$ respectively, and $e_j \preceq e'_j$. From the definition of flow-paths, we know that e_j and e'_j are both either u_j or $\mu^{-1}(u_j)$.

If $e_j = u_j$, then $e_j \preceq e'_j \preceq u_j$ implies that $e'_j = u_j = e_j$, contradicting our assumption that $e_j \neq e'_j$. So it must be that $e_j = \mu^{-1}(u_j)$ and $e'_j = u_j$. Then consider the flowpath δ derived from γ by changing e_j to u_j .

Now, the identity function $id : \{0, \dots, n\} \rightarrow \{0, \dots, n\}$ is a nontrivial embedding function from γ to δ , and φ is a nontrivial embedding function from δ to γ' lifting at i' . Then $e_\gamma \subset \overline{e_\delta}$ and $e_\delta \subset \overline{e_{\gamma'}}$, so $\dim(e_\gamma) < \dim(e_\delta) < \dim(e_{\gamma'})$. However, this contradicts that all cells are 0 and 1-cells lying in the stable subdivision of e_1 , so γ and γ' agree up to $i' - 1$. Now, define a flow-path

$$\delta' = (e_1 \preceq u_1 \succ e'_2 \preceq u'_2 \succ \dots \preceq u'_{\varphi(i)} = u_i \succ e_{i+1} \preceq u_{i+1} \dots).$$

In other words, δ' is the concatenation of the first $\varphi(i')$ pairs of γ' up to $u'_{\varphi(i')} = u_i$, followed by the rest of γ starting from u_i . We then have an embedding function from γ to δ given by

$$\psi(t) = \begin{cases} t & t < i' \\ \varphi(i') - i' + t & t \geq i' \end{cases}$$

lifting at i' and the identity shifted by $\varphi(i') - i'$ thereafter.

Now suppose that there exists $j > i'$ such that $e_j \neq e'_{\varphi(i')-i'+j}$ or $u_j \neq u'_{\varphi(i')-i'+j}$. We can define a second embedding function given by

$$\psi'(t) = \begin{cases} t & t \leq \varphi(i') \\ \varphi(t - (\varphi(i') - i')) & t > \varphi(i') \end{cases}$$

that maps δ' to γ' . We think of ψ' as mapping the first $\varphi(i')$ pairs of δ' identically to those of γ' , then applying φ (with shifted indices) to the second part of δ' which coincides with γ . We then have $\gamma \Rightarrow \delta' \Rightarrow \gamma'$, implying $\dim(e_\gamma) < \dim(e_{\delta'}) < \dim(e_{\gamma'})$. But this contradicts the fact that these cells lie in the stable subdivision of our 1-cell, e_1 . So γ must agree with γ' everywhere except where φ lifts at $i' = i$. \square

Proof of (b). First, we will show that $e_i = \mu^{-1}(u_i)$ rather than u_i . Suppose otherwise, that $e_i = u_i$. Then we can construct a flow-path δ such that δ agrees with γ everywhere except at e_i , where δ takes the cell $\mu^{-1}(u_i)$ instead of u_i . Then we can construct an embedding function from δ to γ that is the identity everywhere except e_i . Then we have $\delta \Rightarrow \gamma \Rightarrow \gamma'$, which contradicts that all cells lie in the stable subdivision of e_1 . So $e_i = \mu^{-1}(u_i)$.




Figure 3.11: A diagram of the embedding function φ from γ to γ' , where it lifts at i . The blue arrows are face adjacencies, and the purple lines are cell equalities.

Now, we show that $\dim(e_i) = 0$. We know that $e'_{\varphi(i-1)+1} = \mu^{-1}(u'_{\varphi(i-1)+1})$ is a strict face of $u'_{\varphi(i-1)+1}$, so 3.11 implies the relation

$$e_i \preceq e'_{\varphi(i-1)+1} \prec u'_{\varphi(i-1)+1}. \quad (3.3)$$

As $\dim(u'_{\varphi(i-1)}) \leq 2$ then $\dim(e_i) \leq 1$.

Suppose that $\dim(e_i) = 1$. By relation 3.3 and the fact that $\dim(u'_{\varphi(i-1)+1}) \leq 2$, this forces $\dim(e'_{\varphi(i-1)+1}) = 1$. Further, as $e_i \preceq e'_{\varphi(i-1)+1}$ and they share the same dimension, they must be equal $e_i = e'_{\varphi(i-1)+1}$. We then have two cases; either $e'_{\varphi(i-1)+1} = u'_{\varphi(i-1)+1}$ or $e'_{\varphi(i-1)+1} = \mu^{-1}(u'_{\varphi(i-1)+1})$, both of which lead to contradictions.

In the case that $e'_{\varphi(i-1)+1} = u'_{\varphi(i-1)+1}$, then $\dim(u'_{\varphi(i-1)+1}) = \dim(e'_{\varphi(i-1)+1}) = 1$. As $u'_{\varphi(i-1)+1} \succ e'_{\varphi(i-1)+2}$, $\dim(e'_{\varphi(i-1)+2}) = 0$. Then the rest of γ' after $\varphi(i-1) + 2$ has at most dimension 1, which contradicts $\dim(u_{\varphi(i)}) = \dim(u_i) = 2$, as lifting at i implies that $\varphi(i) \geq \varphi(i-1) + 2$.

In the case that $e'_{\varphi(i-1)+1} = \mu^{-1}(u'_{\varphi(i-1)+1})$, then

$$u_i = \mu(e_i) = \mu(e'_{\varphi(i-1)+1}) = u'_{\varphi(i-1)+1}.$$

φ lifting at i implies that $u_i = u'_{\varphi(i)}$ and $\varphi(i) > \varphi(i-1) + 1$. As flow-paths cannot return to a matched cell, we get a contradiction in $u'_{\varphi(i)} = u_i = u'_{\varphi(i-1)+1}$. Thus, $\dim(e_i) \neq 1$, so $\dim(e_i)$ must be 0. \square

Proof of (c). Suppose that $e'_{\varphi(i)} \neq u'_{\varphi(i)}$. Then, from the definition of flow-paths, we must have that $e'_{\varphi(i)} = \mu^{-1}(u'_{\varphi(i)})$. By the definition of embedding function, $e_i \preceq e_{\varphi(i)-1}$. If $e'_{\varphi(i)-1} = \mu^{-1}(u'_{\varphi(i)-1})$, then we have that $e_{\varphi(i)} \preceq e'_{\varphi(i)-1}$, which implies that γ' is a reducible flow-path, which is a contradiction. So if $e'_{\varphi(i)} \neq u'_{\varphi(i)}$, then $e'_{\varphi(i)-1} = u'_{\varphi(i)-1}$. \square

Proof of (d). As φ lifts at i , part (a) tells us that γ is identical to γ' outside of the lift. By part (c), either $e'_{\varphi(i)} = u'_{\varphi(i)}$ or $e'_{\varphi(i)-1} = u'_{\varphi(i)-1}$. By Lemma 3.21, we then have either $\dim(e'_{\varphi(i)}) = 1$ or $\dim(e'_{\varphi(i)-1}) = 1$.

In either case, this implies that $\dim(u'_{\varphi(i)-2}) = 2$. As, for any flow-path, the dimension of u_k decreases along k and $\dim(u_k) - 1 \leq \dim(e_k) \leq \dim(u_k)$, there cannot exist some 0-cell with $k < i$. So e_i is the first 0-cell in γ . \square

3.3.3 Algorithms

In this section, we present algorithms for computing the 2-skeleton of the Forman Complex. The attaching map of a 1-cells are easy to find. We follow the unique V -paths emanating from each boundary cell, attaching it to the end-points.

Finding the attaching map of a critical 2-cell is slightly more complicated. Remark 3.7 notes that this task is equivalent to finding the stable subdivision of its boundary. The central focus of our algorithms is to find the face adjacencies in these subdivisions.

Theorem 3.19 reframes this as a combinatorial problem in the world of embedding functions. Here, the useful insight of Lemma 3.22 is that embedding functions can be found by looking at predictable, local regions of flow-paths.

Suppose we have some flow-path γ and associated cell e_γ in the stable subdivision $\text{Sd}_\mu(e_1)$ of some 1-cell $e_1 \in X$. Aside from a special case where γ is a V -path, there are two dual situations to find the next adjacent cell e_δ :

- **Lifting:** If $\dim(e_\gamma) = 0$ then we need to find a flow-path δ such that $\gamma \Rightarrow \delta$, implying $e_\gamma \subseteq \bar{e}_\delta$
- **Colifting:** If $\dim(e_\gamma) = 1$ then we need to find a flow-path δ such that $\delta \Rightarrow \gamma$, implying $e_\delta \subseteq \bar{e}_\gamma$

For the purposes of the algorithm, we introduce the notation that a reduced flow-path

$$\gamma = (e_1 \preceq u_1 \succ \dots \preceq u_n \succ c) \in FP'(\mu)$$

may be written as a two dimensional array in the following way

$$\gamma = [[e_1, u_1], [e_2, u_2], \dots, [e_n, u_n], [c,]].$$

In this notation, each $\gamma[i] = [e_i, u_i]$ is a pair where either $e_i = u_i$ or $e_i = \mu^{-1}(u_i)$.

We retrieve individual elements by double indexing: for example, $\gamma[i][0] = e_i$ and $\gamma[k][1] = u_k$. We can also truncate the first i pairs of γ by the usual array notation $\gamma[i :]$, and use array methods such as concatenate, append and pop.

Lifting

We assume that there exist algorithms `COFACET` and `FACET` that compute the set of cofacets and facets of a cell in X . We also assume the existence of an algorithm `OPATH` for computing the unique V -path emanating from a given 0-cell.

Remark 3.23. In line 2, the ‘ $-$ ’ in the expression

$$\text{FACET}(\gamma[i-1][1]) \cap \text{COFACET}(\gamma[i][0]) - \gamma_{\text{prev}}[i][0]$$

signifies set minus, as in lines 6 and 10.

The while loop in Line 4 amounts to following a 1-dimensional V -path with the condition that each 1-cell is a cofacet of $\gamma[i][0]$, the anchoring cell we are lifting from. Once we reach a cell outside of D , we know that we are at the end of the lift and return the desired flow-path. Recall that, in our notation, the precise conditions of a flow-path

$$\gamma = \gamma[1][0] \preceq \gamma[1][1] \succ \dots \succ \gamma[n][0] \preceq \gamma[n][1] \succ \gamma[n+1][0]$$

are that the cells $\gamma[i][1]$ are distinct elements of U , either $\gamma[i][0] = \gamma[i][1]$ or $\mu(\gamma[i][0]) = \gamma[i][1]$ for $1 \leq i \leq n$ and c is critical. To be a reduced flow-path means that $\gamma[i+1][0] \not\preceq \mu^{-1}(\gamma[i][1])$.

Lemma 3.24. `LIFT`($X, \mu, \gamma, i, \gamma_{\text{prev}}$) returns a reduced flow-path δ such that $\delta \neq \gamma_{\text{prev}}$ and $\gamma \Rightarrow \delta$.

Proof. We must first show that δ is indeed a reduced a flow-path. The object path is built from successive pairs of the form $[e, e]$ for $e \in U$ or $[e, \mu(e)]$ for $e \in D$. Further, as each pair is appended, α is updated at line 6 to be a facet of the previous element of the path, as required.

In the concatenations at Lines 10 and 13, Line 1 ensures that the first element of `path` is adjacent to the final element of $\gamma[i-1][\cdot]$. In the case of line 10, `OPATH` returns a V -path, which must end at a critical cell. In the case of line 13, the returned path ends at $\alpha \in \text{Crit}(\mu)$. In either case, δ ends a critical cell, so δ must be a flow-path.

To see that $\delta \neq \gamma_{\text{prev}}$, note that in Line 1, the first α defined becomes $\delta[i][0]$ regardless of the rest of the algorithm. Since, $\delta[i][0] = \alpha \neq \gamma_{\text{prev}}[i][0]$, we have

Algorithm 1 Lift($X, \mu, \gamma, i, \gamma_{\text{prev}}$)

Input: X a regular cell complex, μ an acyclic partial matching, i the index of the first pair with a 0-cell, γ_{prev} the preceding flow-path in a bigger algorithm

Output: Returns a reduced flow-path lifted from γ .

```

1: function LIFT( $X, \mu, \gamma, i, \gamma_{\text{prev}}$ )
2:    $\alpha \leftarrow \text{FACET}(\gamma[i-1][1]) \cap \text{COFACET}(\gamma[i][0]) - \gamma_{\text{prev}}[i][0]$ 
3:   path  $\leftarrow []$ 
4:   while  $\alpha \in D$  do
5:     path.append( $[\alpha, \mu(\alpha)]$ )
6:      $\alpha \leftarrow \text{FACET}(\mu(\alpha)) \cap \text{COFACET}(\gamma[i][0]) - \alpha$ 
7:   end while
8:   if  $\alpha \in U$  then
9:     path.append( $[\alpha, \alpha]$ )
10:    return concatenate( $\gamma[:i-1]$ , path,  $\text{OPATH}(\text{FACET}(\alpha) - \mu^{-1}(\alpha))$ )
11:  else if  $\alpha \in \text{Crit}(\mu)$  then ▷ Truncation
12:    path.append( $[\alpha, ]$ )
13:    return concatenate( $\gamma[:i-1]$ , path)
14:  end if
15: end function

```

$\delta \neq \gamma_{\text{prev}}$ as required.

Now we show that $\gamma \Rightarrow \delta$. The claim is that there is an embedding function φ from γ to δ lifting at i . Regardless of how the algorithm ends, the first $i-1$ pairs of δ are the same as those of γ . There are three cases:

1. **Case 1:** The algorithm terminates on Line 10 and $(\text{FACET}(\alpha) - \mu^{-1}(\alpha)) = \gamma[i][0]$. If this is the case, then the lift takes $\gamma[i]$ to the final pair in **path**. Each time we adjoin a pair $[\alpha, \mu(\alpha)]$ to **path**, Line 6 ensures that it satisfies $\alpha \succeq \gamma[i][0]$. Once we reach $\gamma[i][0]$, the fact that it is a 0-cell uniquely determines the V-path it must be sent along, so γ and δ agree after this. Hence, we have a valid lift at i .
2. **Case 2:** The algorithm terminates on Line 10 and $(\text{FACET}(\alpha) - \mu^{-1}(\alpha)) \neq \gamma[i][0]$. Again, Line 6 ensures that each pair $[\alpha, \mu(\alpha)]$ we adjoin to **path** satisfies $\alpha \succeq \gamma[i][0]$. The final pair of 1-cells $[\alpha, \alpha]$ must be followed by $\gamma[i+1][0]$, as this is the only other possible cell in $\text{FACET}(\alpha)$. Thus, we have

a valid lift at i .

3. **Case 3:** The algorithm terminates on Line 11. Line 6 ensures that each pair $[\alpha, \mu(\alpha)]$ we adjoin to `path` satisfies $\alpha \succeq \gamma[i][0]$. Once $\alpha \in \text{Crit}(\mu)$ and $\alpha \in \text{COFACET}(\gamma[i][0])$ then we have a truncation at i , hence $\gamma \Rightarrow \delta$.

□

Colifting

The dual algorithm to LIFT is the COLIFT algorithm described below. The While loop in the LIFT algorithm inserted pairs into the flow-path γ to construct a flow-path $\delta \Leftarrow \gamma$. Dually, the While loop in the COLIFT algorithm *deletes* pairs from a flow-path γ to construct a flow-path $\delta \Rightarrow \gamma$.

Algorithm 2 `Colift`($X, \mu, \gamma, i, \gamma_{\text{prev}}$)

Input: X a regular cell complex, $\mu : D \rightarrow U$ an acyclic partial matching, γ a path, i the index of the unique pair such that $\gamma[i][0] = \gamma[i][1]$, γ_{prev} is the preceding path in a bigger algorithm

Output: Returns a modified flow-path γ and the index $k + 1$ of its first 0-cell

```

1: function COLIFT( $X, \mu, \gamma, i, \gamma_{\text{prev}}$ )
2:   if  $\gamma[i][1] \in \gamma_{\text{prev}}$  then
3:      $\gamma.\text{pop}(i)$ 
4:   else
5:      $\gamma[i][0] = \mu^{-1}(\gamma[i][1])$ 
6:   end if
7:    $k \leftarrow i - 1$ 
8:   while  $\gamma[k][0] \in \text{COFACET}(\gamma[i][0])$  do
9:      $\gamma.\text{pop}(k)$ 
10:     $k \leftarrow k - 1$ 
11:  end while
12:  return  $\gamma, k + 1$ 
13: end function

```

Lemma 3.25. `COLIFT`($X, \mu, \gamma, i, \gamma_{\text{prev}}$) returns a reduced flow path δ' such that $\delta' \neq \gamma_{\text{prev}}$ and $\delta' \Rightarrow \gamma$.

Proof. We show that δ' is indeed a reduced flow-path. Firstly, we show that γ satisfies the adjacency requirements after each alteration. In Line 3, we remove the i -pair, which is the first pair to repeat. In γ , we have

$$\gamma[i-1][1] \succ \gamma[i][0] \preceq \gamma[i][1] \succ \gamma[i+1][0]$$

so $\gamma[i-1][1] \succ \gamma[i+1][0]$ as required. In Line 5, γ gives us the face relation

$$\gamma[i-1][1] \succ \gamma[i][0] = \gamma[i][1] \succ \mu^{-1}(\gamma[i][1])$$

so $\gamma[i-1][1] \succ \gamma[i][0]$ after the alteration as required. Lastly, at each iteration of the While loop at Line 8, the condition $\gamma[k][0] \in \text{COFACET}(\gamma[i][0])$ ensures that $\gamma[k-1][1] \succ \gamma[k][0] \succ \gamma[i][0]$, so the requisite adjacencies hold for all of δ' .

To see that δ' is a *reduced* flow-path, we note that anytime

$$\gamma[k][0] \preceq \mu^{-1}(\gamma[k-1][1])$$

it is deleted by the While loop in Line 8.

To show $\delta' \neq \gamma_{\text{prev}}$, consider first Line 2 where $\gamma[i][1] \in \gamma_{\text{prev}}$. As, by the definition of flow-paths, each $\gamma[k][1]$ is distinct, $\gamma[i][1]$ only appears once in γ . It is deleted from γ in Line 3, so does not appear in δ' , thus $\delta' \neq \gamma_{\text{prev}}$. On the other hand, consider $\gamma[i][1] \notin \gamma_{\text{prev}}$. As it remains unchanged throughout the algorithm, $\gamma[i][1] = \delta'[i][1]$, so $\delta' \neq \gamma_{\text{prev}}$.

Lastly, we need to show that $\delta' \Rightarrow \gamma$. Let k_{final} be the final value of k upon termination of the algorithm. First consider the case where $\gamma[i][1] \in \gamma_{\text{prev}}$. We claim there is a lift at $k_{\text{final}} + 1$ mapping to the $(i+1)$ -th pair of γ . First note that $\delta'[k_{\text{final}} + 1] = \gamma[i+1]$. For each $k_{\text{final}} + 1 \leq j \leq i$, a pair is removed from γ to construct δ' such that each $\gamma[j][0]$ is a cofacet of $\delta'[k_{\text{final}} + 1][0] = \gamma[i+1][0]$. (Here, $\gamma[i][0]$ in Line 8 is the same as the original γ after removing the i -th pair in Line 3). As δ' and γ agree elsewhere by construction, there is a valid lift at $k_{\text{final}} + 1$, so $\delta' \Rightarrow \gamma$.

Now consider the case where $\gamma[i][1] \notin \gamma_{\text{prev}}$. In this case, note that $\delta'[k_{\text{final}}] = [\mu^{-1}(\gamma[i][1]), \gamma[i][1]]$. For each $k_{\text{final}} + 1 \leq j \leq i-1$, a pair is removed from γ to construct δ' such that each $\gamma[j][0]$ is a cofacet of $\delta'[k_{\text{final}} + 1][0] = \mu^{-1}(\gamma[i][1])$. (Here, $\gamma[i][0]$ in Line 8 is the same as $\mu^{-1}(\gamma[i][1])$ in the original γ in Line 3). As δ'




Figure 3.12: A regular cell complex X and an acyclic partial matching $\mu : D \rightarrow U$. The cell label colours show cells which lie on a 0-dimensional and 1-dimensional V -paths.

and γ agree elsewhere by construction, there is a valid lift at $k_{\text{final}} + 1$, so $\delta' \Rightarrow \gamma$.

□

Example 3.26. The following example will illustrate the geometric intuition behind the LIFT and COLIFT algorithms. Consider the flow-path in Figure 3.12 given by:

$$\gamma = (e_1 \preceq u_1 \succ x_3 \preceq y_3 \succ b_1 \preceq b_1 \succ c)$$

Assume that we are given $\gamma_{\text{prev}} = (e_1 \preceq u_1 \succ x_3 \preceq y_3 \succ c)$, with $\gamma_{\text{prev}} \Rightarrow \gamma$. Figure 3.13 depicts the situation where we apply the COLIFT algorithm. Firstly, we note that the first repeat is at the third pair $\gamma[3] = [b_1, b_1]$. As $b_1 \notin \gamma_{\text{prev}}$, we move to Line 5 of COLIFT and change $\gamma[i][0]$ to $\mu^{-1}(\gamma[i][1]) = a_1$.

We move onto Line 7 and set $k = 2$, whereby we verify that $\gamma[3][0] = a_1$ is a cofacet of $\gamma[2][0] = x_3$ and remove the pair $\gamma[2] = [x_3, y_3]$ from γ . Returning to the While loop, we see that $\gamma[1][0] = e_1$ is not a cofacet of a_1 . So we end the While loop and the algorithm returns δ , as pictured in Figure 3.13.

Figure 3.14 depicts the situation where we apply the LIFT algorithm. We take as input the last two paths from Figure 3.13 as an example of how LIFT and COLIFT are strung together.




Figure 3.13: An example of embedding functions between flow-paths in the COLIFT algorithm. Here, $\delta = \text{COLIFT}(X, \mu, \gamma, 3, \gamma_{\text{prev}})$, where the 3 indicates that the first repeat is the third pair. The arrows $x \rightarrow y$ represent the cellular face relations $x \preceq y$.




Figure 3.14: An example of embedding functions between flow-paths in the LIFT algorithm. Here, $\delta = \text{LIFT}(X, \mu, \gamma, 2, \gamma_{\text{prev}})$, where 2 indicates the index of the first pair containing a 0-cell $\gamma[2] = [a_1, b_1]$.

Firstly, we start at the minimum zero cell $\gamma[2][0] = a_1$. Our first value for α is the unique cell x_1 in $\text{FACET}(\gamma[i-1][1]) \cap \text{COFACET}(\gamma[i][0]) - \gamma_{\text{prev}}[2][0]$. Since $\mu(x_1) = y_1$, $x_1 \in D$ and we add the pair $[x_1, y_1]$ to **path**. Similarly, we then add the pairs $[x_2, y_2]$ and $[x_3, y_3]$ to **path**.

After $\alpha = x_3$, the variable α is assigned to b_1 . As $b_1 \in U$, we append $[b_1, b_1]$ to **path**. To end the algorithm, we concatenate $\gamma[:i-1]$, **path** and the V -path starting at $\text{FACET}(b_1) - a_1 = c$. Since c is a critical 0-cell, it is already a V -path, which leads us to the δ in Figure 3.14.

Remark 3.27. Geometrically, the outside flow-paths in Figure 3.14 ‘enclose’ the middle flow-path. For visual intuition, it is worthwhile referring back to Figure 3.12 to see how these flow-paths wrap around one another.

Subdivide

The SUBDIVIDE algorithm co-ordinates the LIFT and COLIFT algorithms. Starting at a boundary cell, we successively lift and co-lift until we reach the other side, recording each flow-path into a matrix.

We give a brief description of the algorithm. The first 4 lines initialize the algorithm, finding the facets of e_1 and creating an array $\text{Sd}_\mu(e_1)$ to store cells. Our starting flow-path, γ_0 , is the unique V -path emanating from a . The final path γ_1 emanates from the other end b of e_1 . Our goal is to compute all flow-paths in between, recovering the stable subdivision $\text{Sd}_\mu(e_1)$ in the process.

Algorithm 3 Subdivide(X, μ, e_1)

Input: X a regular cell complex, μ an acyclic partial matching, e_1 a 1-cell

Output: Returns an ordered array of cells $\text{Sd}_\mu(e_1)$ in the stable subdivision of e_1 .

```

1: function SUBDIVIDE( $X, \mu, e_1$ )
2:    $a, b \leftarrow \text{FACET}(e_1)$ 
3:    $\text{Sd}_\mu(e_1) \leftarrow [a]$  ▷ Array of Critical Cells in Subdivision
4:    $\gamma_0, \gamma_{\text{final}}, \gamma_1 \leftarrow \text{VPATH}(a, \mu), \text{VPATH}(b, \mu), [ ]$  ▷ Initialise the paths
5:    $i_0 \leftarrow 0$ 
6:   while  $\gamma_0 \neq \gamma_{\text{final}}$  do
7:      $\gamma_1 \leftarrow \text{LIFT}(X, \mu, \gamma_0, i_0, \gamma_1)$ 
8:     if  $\dim(\text{END}(\gamma_1)) == 1$  then ▷ Case  $\gamma_1$  is a  $V$ -path
9:        $\text{Sd}_\mu(e_1).\text{append}(\text{ORIENTATION}(\gamma_1) * \text{END}(\gamma_1))$ 
10:       $i_0 \leftarrow \text{len}(\gamma_1)$ 
11:       $\gamma_0 \leftarrow \text{concatenate}(\gamma_1[:i_0 - 1], \text{OPATH}(\text{FACET}(\text{END}(\gamma_1) - \gamma_0[i_0][0])))$ 
12:       $\text{Sd}_\mu(e_1).\text{append}(\text{END}(\gamma_0))$ 
13:     else
14:        $i_1 \leftarrow \text{FIRSTREPEAT}(\gamma_1)$ 
15:        $\gamma_0, i_0 \leftarrow \text{COLIFT}(X, \mu, \gamma_1, i_1, \gamma_0)$ 
16:     end if
17:   end while
18:   return  $\text{Sd}_\mu(e_1)$ 
19: end function

```

The While loop beginning on Line 6 iterates lifting and co-lifting, stopping




Figure 3.15: A partial stable subdivision of e_1 with flow-paths from Figures 3.12, 3.13 and 3.14. The SUBDIVIDE algorithm ties face adjacencies together by successively finding embedding functions with LIFT and COLIFT.

when we reach our end path γ_{final} . The variable i_0 records the position of the first 0-cell in γ_0 . We then lift to the next flow-path γ_1 , making sure to set the previous path to γ_1 as to not go backwards.

Here, there are two possibilities for γ_1 . If $\dim(\tau(\gamma_1)) = 1$, then γ_1 is a V -path, in which case we record its orientation and add it to $\mathbf{Sd}_\mu(e_1)$. In this situation, we know precisely how to find an adjacent flow-path. We simply look at the boundary of the final critical 1-cell, and follow the 0-dimensional V -path out of it.

The second possibility occurs at Line 14. If $\dim(\gamma_1) \neq 1$, it shares the same critical end-cell as the previous flow-path. By the Lifting Lemma 3.22, γ_1 must contain a repeat, which we set to i_1 . We then co-lift to find the next flow-path, returning to the beginning of the While Loop.

Iterations of LIFT and COLIFT remain *local* throughout the subdivide algorithm. We don't need to scan long flow-paths, nor search excessively for complicated embedding functions. The algorithm knows exactly where to look. From a computational standpoint, this property significantly increases the viability of implementation.

With the means to subdivide 1-cells, comes the power to compute the attaching map of any 2-cell. It should be noted that we have not provided an algorithm for first computing the 1-skeleton, the computation of which will be straightforward. Once implemented, it is then possible to compute the entire 2-skeleton. In Chapter 5, we use this construction to find the fundamental group.

Chapter 4

Persistence

Most algebraic invariants of topological spaces fail to distinguish between homotopy equivalent shapes of different sizes. The notion of persistence addresses this problem by embedding information about a metric in a filtration of the space, and measuring how the algebraic invariant changes.

When our algebraic invariant is homology, we get persistent homology. Recently, this field has been studied in detail, yielding many interesting applications. In the context of persistence, there is less study on homotopy groups as the algebraic invariant.

Some notable examples of such a notion can be found in [9] and [?]. However, the novel definition of a *homotopy merge tree* introduced in this thesis extends the work of [14], with an emphasis on computability and connections with discrete Morse theory.

4.1 Persistent Homology

Our presentation of persistent homology will focus on *persistence modules* [2]. This is somewhat backwards historically; persistent homology was first introduced in terms of *persistent betti numbers* with a strong geometric flavour (See [17] and [8]). However, this direction is taken to motivate the definition of a homotopy merge tree in a later section.

4.1.1 Persistence Modules

Definition 4.1. [2] A **real persistence module** \mathbb{V} over a field k is a family of k -vector spaces indexed by \mathbb{R}

$$\{V_t \mid t \in \mathbb{R}\}$$

and double indexed linear maps

$$\{v_s^t : V_s \rightarrow V_t\}$$

such that $v_s^t \circ v_t^r = v_s^r$ for all $s, t, r \in \mathbb{R}$.

Remark 4.2. If we consider \mathbb{R} as a poset under the standard ordering, a real persistence module is a functor from \mathbb{R} into the category of k -vector spaces.

Persistent homology is framed as a persistence module in the following situation. Suppose we have a real-valued function from a topological space $f : X \rightarrow \mathbb{R}$. Recall that $X_{f \leq h} = f^{-1}((-\infty, h])$. The sub-level set filtration X is given by

$$\{X_{f \leq h} \mid h \in \mathbb{R}\}.$$

This construction can be described as a functor from $X_f : \mathbb{R} \rightarrow \mathbf{Top}$, where $s \mapsto X_{f \leq s}$ and each order relation $s \leq t$ induces the inclusion

$$X_{f \leq s} \hookrightarrow X_{f \leq t}$$

Diagram 4.1 summarises this information.

$$\begin{array}{ccc}
 s & \xrightarrow{\quad} & r \\
 & \searrow & \nearrow \\
 & & t
 \end{array}
 \xrightarrow{X_f}
 \begin{array}{ccc}
 X_{f \leq s} & \xrightarrow{i} & X_{f \leq r} \\
 & \searrow i & \nearrow i \\
 & & X_{f \leq t}
 \end{array}
 \quad (4.1)$$

Fixing a field k , we can push forward this construction by applying the $H_n(_, k)$ functor, with the composition

$$\mathbb{R} \xrightarrow{X_f} \mathbf{Top} \xrightarrow{H_n(_, k)} k\mathbf{Vec}$$

denoted \mathbb{X}_f over k . This composition transforms the commutative Diagram 4.1 into the one below.

$$\begin{array}{ccc}
 H_n(X_{f \leq s}; k) & \xrightarrow{i_*} & H_n(X_{f \leq r}; k) \\
 & \searrow i_* & \nearrow i_* \\
 & & H_n(X_{f \leq t}; k)
 \end{array}$$

4.1.2 Interleaving Distance

There does exist a notion of two persistence modules being isomorphic. However, in practice, we wish to compare filtrations of spaces that almost never satisfy such an equivalence.

Interleavings are isomorphisms with a tolerance for error. They provide a sensible way to compare near-equivalent filtrations by allowing us to shift each filtration index when comparing persistence modules.

Definition 4.3. [2] Let $\delta \in \mathbb{R}_{\geq 0}$. Two persistence modules \mathbb{U}, \mathbb{V} are δ -interleaved if for all $t \in \mathbb{R}$ there exist maps

$$\phi_t : U_t \rightarrow V_{t+\delta} \text{ and } \psi_t : V_t \rightarrow U_{t+\delta}$$

such that the diagrams commute

$$\begin{array}{ccc} U_s & \xrightarrow{u_s^t} & U_t \\ \downarrow \phi_s & & \downarrow \phi_t \\ V_{s+\delta} & \xrightarrow{v_{s+\delta}^{t+\delta}} & V_{t+\delta} \end{array} \quad \begin{array}{ccc} U_{s-\delta} & \xrightarrow{u_{s-\delta}^{s+\delta}} & U_{s+\delta} \\ \searrow \phi_{s-\delta} & & \nearrow \psi_s \\ & & V_s \end{array}$$

$$\begin{array}{ccc} V_s & \xrightarrow{v_s^t} & V_t \\ \downarrow \psi_s & & \downarrow \psi_t \\ U_{s+\delta} & \xrightarrow{u_{s+\delta}^{t+\delta}} & U_{t+\delta} \end{array} \quad \begin{array}{ccc} V_{s-\delta} & \xrightarrow{v_{s-\delta}^{s+\delta}} & V_{s+\delta} \\ \searrow \psi_{s-\delta} & & \nearrow \phi_s \\ & & U_s \end{array}$$

Remark 4.4. We can think of an interleaving as a natural transformation between the persistence modules \mathbb{U} and \mathbb{V} , but shifted by δ .

We can define the **interleaving distance** $d_i(\mathbb{U}, \mathbb{V})$ to be

$$\inf\{\delta \mid \mathbb{U}, \mathbb{V} \text{ are } \delta\text{-interleaved}\}.$$

With these definitions, the interleaving distance is shown to be an extended pseudo-metric in [2].

A special case of interleaving is when we have two \mathbb{R} -valued functions $f, g : X \rightarrow \mathbb{R}$ over the same topological space X . [2] shows that if $\|f - g\|_\infty \leq \delta$ then the persistent homology modules \mathbb{X}_f and \mathbb{X}_g are δ -interleaved. In other words, $d_i(\mathbb{X}_f, \mathbb{X}_g) \leq \delta$. This property of *stability* is important, as it underscores that the interleaving metric responds appropriately to changes in the filtration of X .

4.2 Persistent Homotopy

Given the utility of persistent homology, a natural question is whether we can do the same thing with homotopy groups. Quickly, we realise there are two immediate problems. We now have to choose a base-point if the space has more than one path component, and the computations are much harder than for homology.


The aim of this section is to define a sensible notion of persistence in the context of homotopy groups, up to the point of proving the Stability Theorem.

The theory presented in this section closely follows the structure and proofs of [14], whose work provides the definitions and theorems in relation to the underlying topological merge trees.

Motivation

Persistence modules, as defined in the previous section, rely on \mathbb{R} as an indexing set for the homology functor H_* . In defining what a persistent fundamental group *should* be, a naive approach would be to apply $\pi_1(_, x)$ at the level of topological spaces for a given $x \in X$.


However, we must first decide which basepoint to pick. To see why this could be a problem, consider the following simple example of a filtration on a topological space:



Applying $\pi_1(_, x_1)$ and $\pi_1(_, x_2)$ to the filtration $X_1 \rightarrow X_2 \rightarrow X_3$ respectively yields two distinct sequences of groups:

$$\langle a \rangle \longrightarrow \langle a \rangle \longrightarrow \langle a, b \rangle \quad \text{and} \quad 0 \longrightarrow \langle b \rangle \longrightarrow \langle a, b \rangle$$

Neither sequence captures all the information about how generators of the fundamental group are born and die. A slight modification of the construction gives



where γ_p is a change of base-point homomorphism from x_2 to x_1 . The main modification here is occurs in the underlying indexing category, depicted above as a poset category with asterisks for objects and arrows for morphisms. In the following section, we will try to formalise this notion, and weave together $\pi_n(_, x)$ for many different $x \in X$.

4.2.1 Merge Trees

We need a more systematic way of constructing this indexing category as per the motivation above. As the information we need is related to base-points, we need something that measures path component changes in the filtration. Such an object is known as a topological **merge tree**.

We follow its construction in [14], where the merge tree is defined in terms of the sub-level set of an \mathbb{R} -valued function over a general topological space. Working with homotopy groups, we make a slight alteration to consider \mathbb{R} -valued functions over a CW complex, and discuss the necessary technical adjustments.

Let X be a CW complex and $f : X \rightarrow \mathbb{R}$ be a function on the cells. The **epigraph of f** , denoted $\text{epi}(f)$ is the set of points above f in its graph, $\{(x, h) \in X \times \mathbb{R} \mid f(x) \leq h\}$. The **merge tree** T_f is a quotient space of the epigraph of f under the following equivalence relation: $(x_1, h_1) \sim (x_2, h_2)$ if $h_1 = h_2$ and x_1 is in the same connected component as x_2 in the sub-level set $f^{-1}((-\infty, h_1])$.

Remark 4.5. Those familiar with **Reeb Graphs** will recognise that T_f could alternatively be defined as the Reeb graph of $\text{epi}(f)$.

Any merge tree T_f inherits two important continuous functions [14]:

1. A **shift function** $i^\varepsilon : T_f \rightarrow T_f$ for all ε . For any point $[x, h] \in T_f$, $i^\varepsilon([x, h]) = [(x', h')]$ is the unique point in T_f such that $h' = h + \varepsilon$ and x' belongs to the same connected component as x in $f^{-1}((-\infty, h'])$
2. A **height function** $\hat{f} : T_f \rightarrow \mathbb{R}$, such that $\hat{f}([x, h]) \mapsto h \in \mathbb{R}$.




Figure 4.1: An example of a CW complex X with piecewise-constant function $f : X \rightarrow \mathbb{R}$. The epigraph $\text{epi}(f)$ is pictured above, with regions coloured according to the cell they lie above. The merge tree T_f is pictured on the right. Notice f is not a discrete Morse function.

4.2.2 Canonical Paths

To deal merging path components across homotopy groups, we need to incorporate change of base-point homomorphisms. An important consideration is to consistently decide which of the two base points to select the homotopy class of path between them.

For completeness, we will briefly describe how the change of base-point homomorphism works for higher homotopy groups. Recall that for a space X , each higher homotopy group $\pi_n(X, x_0)$ can be thought of as homotopy classes of maps $(I^n, \partial I^n) \rightarrow (X, x_0)$.

Given a path $p : I \rightarrow X$ from $p(0) = x_0$ to $p(1) = x_1$, we can define (as per [10]) a **change of base-point homomorphism** $\gamma_p : \pi_n(X, x_0) \rightarrow \pi_n(X, x_1)$ in the following way. Let $[f] \in \pi_n(X, x_0)$, then $\gamma_p f$ is a map $(I^n, \partial I^n) \rightarrow (X, x_1)$ given by shrinking f into a concentric cube of I^n and radially expanding its boundary along the path p . (See Figure 4.2)

Minimal Spanning Trees

The 1-skeleton of any CW complex is a graph. Any function $f : X \rightarrow \mathbb{R}$ over a CW complex then restricts to a function over the 1-skeleton graph. For graphs, a **minimal spanning tree** of each path component is a path connected acyclic subgraph that minimises f -values. For any two points e_1, e_2 in a given minimal




Figure 4.2: The domain of a map $(I^n, \partial I^n) \rightarrow (X, x_1)$, cut up into different regions. The original map $f : (I^n, \partial I^n) \rightarrow (X, x_0)$ is shrunk, with the rest of the domain filled by maps $p : I \rightarrow X$ expanding radially outward.

spanning tree, there is a unique sequence of edges connecting them. We call this sequence of edges the **canonical path** p_{e_1, e_2} from e_1 to e_2 .

As we will see in the next section, the reason for picking paths in minimal spanning trees is that they behave well under change of base-point homomorphisms and inclusions. The following lemma shows that the path homotopy class of composing canonical paths depends only on the endpoints:

Lemma 4.6. *Let $e_1, e_2, e_3 \in X$. Then $p_{e_1, e_2} \cdot p_{e_2, e_3} \simeq p_{e_1, e_3}$.*

Proof. First note that all canonical paths lie in the minimal spanning tree, which is contractible. If $p_{e_1, e_2} \cdot p_{e_2, e_3}$ is not path homotopic to p_{e_1, e_3} , then the composition $[p_{e_1, e_2} \cdot p_{e_2, e_3} \cdot p_{e_1, e_3}^{-1}]$ forms a non-trivial loop class based at e_1 in the fundamental group of the minimal spanning tree, which is a contradiction. \square

4.2.3 Homotopy Merge Tree

Now we define the main object of this section, the homotopy merge tree of an \mathbb{R} -valued function. Intuitively, we can think of it as assigning a homotopy group to every point on the merge tree by looking at the sub-level set.

For homotopy groups, the base-points matter. The general class of functions we will consider are functions $f : X \rightarrow \mathbb{R}$ over a CW complex X satisfying the following condition:

$$\text{for each } e \in X, \text{ there exists a } 0\text{-cell } e_0 \in \bar{e} \text{ such that } f(e_0) \leq f(e) \quad (\star)$$

For practical purposes we will always want to choose our base-point to be the minimal 0-cell of a sub-level set.

Remark 4.7. It should be noted that all discrete Morse functions satisfy (\star) , and in practice, are the functions we will mainly work with. However, the definitions and theorems in the following section hold for the more general class of functions satisfying (\star) .

Definition 4.8. Let $f : X \rightarrow \mathbb{R}$ be a function over a CW complex X satisfying (\star) . Let T_f be its merge tree. The n -**homotopy merge tree** $\pi_n[T_f]$ is a family of groups $\{\pi_n(X_\sigma, e_\sigma)\}_{\sigma \in T_f}$ indexed by T_f such that

1. X_σ is $f^{-1}((-\infty, \hat{f}(\sigma)])$.
2. e_σ is the minimal zero cell in X_σ

That each sub-level set X_σ has a minimal 0-cell follows directly from the fact that f satisfies the \star condition.

Definition 4.9. Let $\pi_n[T_f]$ be a homotopy merge tree. The ε -**shift morphism** $i^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_f]$ is an ε -shift map $i^\varepsilon : T_f \rightarrow T_f$ with family of group homomorphisms given by

$$\{i_\sigma^\varepsilon = \gamma_p \circ i_* : \pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(X_{i^\varepsilon(\sigma)}, e_{i^\varepsilon(\sigma)})\}_{\sigma \in T_f}$$

where i_* is the induced map of the inclusion from $X_\sigma \rightarrow X_{i^\varepsilon(\sigma)}$ and γ_p is the change of base-point homomorphism along the canonical path from e_σ to $e_{i^\varepsilon(\sigma)}$. Since the canonical path is implied and unique, when appropriate we will often drop the p and just write γ .

Lemma 4.10. For $\sigma \in T_f$, the following diagram

$$\begin{array}{ccccc} \pi_n(X_\sigma, e_\sigma) & \xrightarrow{\gamma_{p_1} \circ i_*} & \pi_n(X_{i^\varepsilon(\sigma)}, e_{i^\varepsilon(\sigma)}) & \xrightarrow{\gamma_{p_2} \circ i_*} & \pi_n(X_{i^{\varepsilon+\delta}(\sigma)}, e_{i^{\varepsilon+\delta}(\sigma)}) \\ & & & \searrow & \nearrow \\ & & & \gamma_{p_3} \circ i_* & \end{array}$$


commutes, where the paths are canonical paths and the i_* are the maps induced by inclusions $X_\sigma \subseteq X_{i^\varepsilon(\sigma)} \subseteq X_{i^{\varepsilon+\delta}(\sigma)}$.

Proof. We should first check that these maps are well defined by verifying the canonical paths between points are contained in the correct sub-level sets. More formally, if $\varepsilon' \in \mathbb{R}_{\geq 0}$ then we want to check the canonical path p from e_σ to $e_{i^{\varepsilon'}(\sigma)}$

lies inside $X_{i^{\varepsilon'}(\sigma)}$.

Here, we note that e_σ and $e_{i^{\varepsilon'}(\sigma)}$ are in the same path component in $X_{i^{\varepsilon'}(\sigma)}$ by the definition of the shift map. Hence, there is path p' of 1-cells connecting the two points in $X_{i^{\varepsilon'}(\sigma)}$. The canonical path p lies in the minimal spanning tree, so it must take values below that of p' everywhere. Hence, p is contained in the same sub-level set $X_{i^{\varepsilon'}(\sigma)}$ as p' , as required.

In checking commutativity, the inclusions i_* serve only to put the change of base-point homomorphisms in the right space. Thus, it is sufficient to check that $\gamma_{p_2}\gamma_{p_1}[s] = \gamma_{p_3}[s]$ for all $[s]$ in $\pi_n(X_{i^{\varepsilon+\delta}(\sigma)}, e_{i^{\varepsilon+\delta}(\sigma)})$. By Lemma 4.6, we know that $p_1 \cdot p_2 \simeq p_3$. Hence, we can homotope the paths in the change of base-point map as shown below:



so $\gamma_{p_2}\gamma_{p_1}[s] = \gamma_{p_3}[s]$ as required. □

Functorial Construction

These definitions can be re-interpreted as a more general construction. Generalised persistence modules, as defined in [1], retain many of the nice properties of persistence. They provide a basis for our framework for doing persistence calculations with the fundamental group later on:

Definition 4.11. [1] A **generalised persistence module** consists of the following data:

1. A partially ordered set P , called the indexing category*;
2. A category D , called the target category and

*The original paper uses pre-ordered sets, of which partially ordered sets are an example

3. A functor $F : P \rightarrow D$ called the generalised persistence module in D over P

In our context, let $f : X \rightarrow \mathbb{R}$ be a function on a cell complex X . The merge tree T_f can be made into a poset P_f by letting each $\sigma \in T_f$ be an element of P_f , and letting $\sigma \leq \tau$ whenever we can reach τ from σ by flowing up along T_f . We can describe this more formally using the shift functions $i^\varepsilon : T_f \rightarrow T_f$: the order relation $\sigma \leq \tau$ holds iff there exists $\varepsilon \in \mathbb{R}_{\geq 0}$ such that $i^\varepsilon(\sigma) = \tau$.

Lemma 4.12. *P_f is a poset. Recall that a poset satisfies the following conditions:*

1. *if $\sigma \in P_f$ then $\sigma \leq \sigma$.*
2. *if $\sigma, \tau \in P_f$ and $\sigma \leq \tau$ and $\tau \leq \sigma$ then $\sigma = \tau$*
3. *if $\sigma_1, \sigma_2, \sigma_3 \in P_f$, with $\sigma_1 \leq \sigma_2$ and $\sigma_2 \leq \sigma_3$ then $\sigma_1 \leq \sigma_3$.*

Proof. Considering $id = i^0 : T_f \rightarrow T_f$ shows that $\sigma \leq \sigma$ for all $\sigma \in T_f$. Now consider the case where we would like to show $\sigma \leq \tau$ and $\tau \leq \sigma$ implies $\sigma = \tau$. Here, there exist shift functions i^ε and i^δ such that $i^\varepsilon(\sigma) = \tau$ and $i^\delta(\tau) = \sigma$. This means $\hat{f}(\sigma) + \varepsilon = \hat{f}(\tau) = \hat{f}(\sigma) - \delta$, implying that $\varepsilon = -\delta = 0$. So $\tau = i^0(\sigma) = id(\sigma) = \sigma$.

The last condition to verify P_f is a poset is transitivity. Let $\sigma_1 \leq \sigma_2 \leq \sigma_3$. By construction, there exists $\varepsilon, \delta \in \mathbb{R}_{\geq 0}$ such that $i^\varepsilon(\sigma_1) = \sigma_2$ and $i^\delta(\sigma_2) = \sigma_3$. We then have $i^{\delta+\varepsilon}(\sigma_1) = i^\delta(i^\varepsilon(\sigma_1)) = i^\delta(\sigma_2) = \sigma_3$, so $\sigma_1 \leq \sigma_3$ as required. \square

Any poset P_f can be thought of as category whose objects are the elements of P_f , and whose morphisms $\sigma \rightarrow \tau$ are given by order relations $\sigma \leq \tau$ in P_f . That morphisms in this category are composable follows from the transitivity condition on the order relations of P_f . Given our reformulation of T_f as a poset category P_f , we can define the **induced functor of a homotopy merge tree** $\pi_n[T_f] : P_f \rightarrow \mathbf{Grp}$ in the following way:

1. For $\sigma \in P_f$, $\pi_n[T_f](\sigma) = \pi_n(X_\sigma, e_\sigma)$.
2. For a morphism $\sigma \leq \tau$ with $i^\varepsilon(\sigma) = \tau$, $\pi_n[T_f](\sigma \leq \tau)$ is given by the map

$$i_\sigma^\varepsilon = \gamma \circ i_* : \pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(X_\tau, e_\tau)$$

where γ is the change of base-point homomorphism along the canonical path from σ to τ .

Lemma 4.13. $\pi_n[T_f]$ is a functor.

Proof. The identity morphism in P_f given by $\sigma \leq \sigma$ is sent to $\pi_n[T_f](\sigma \leq \sigma) = \beta_h \circ i_* = i_* = id : \pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(X_\sigma, e_\sigma)$ as h is the trivial path.

Now suppose that we have $\sigma_1 \leq \sigma_2 \leq \sigma_3$ in P_f with canonical paths h_{12} from σ_1 to σ_2 , h_{23} from σ_2 to σ_3 and h_{13} from σ_1 to σ_3 . So there exists $\varepsilon_1, \varepsilon_2 \in \mathbb{R}_{\geq 0}$ such that $i^{\varepsilon_1}(\sigma_1) = \sigma_2$ and $i^{\varepsilon_2}(\sigma_2) = \sigma_3$. We need to check that

$$(\pi_n[T_f])(\sigma_1 \leq \sigma_2) \circ (\pi_n[T_f])(\sigma_2 \leq \sigma_3) = (\pi_n[T_f])(\sigma_1 \leq \sigma_3).$$

Equivalently, we need to check the following diagram commutes.

$$\begin{array}{ccccc} \pi_n(X_\sigma, e_\sigma) & \xrightarrow{\gamma \circ i_*} & \pi_n(X_{i^{\varepsilon_1}(\sigma)}, e_{i^{\varepsilon_1}(\sigma)}) & \xrightarrow{\gamma \circ i_*} & \pi_n(X_{i^{\varepsilon_1+\varepsilon_2}(\sigma)}, e_{i^{\varepsilon_1+\varepsilon_2}(\sigma)}) \\ & \searrow & & \nearrow & \\ & & \gamma \circ i_* & & \end{array}$$

This follows precisely from Lemma 4.10. □

Since every homotopy merge tree has shift morphisms, the functorial and non-functorial definition contain the same information

4.2.4 Comparing Homotopy Merge Trees

Given that we would like to use homotopy merge trees to compare the properties of spaces, we need a way of comparing them as objects. Namely, we need to define a map between them that respects their structure.

Definition 4.14. Let $f : X \rightarrow \mathbb{R}$ and $g : X \rightarrow \mathbb{R}$ be functions over cell complexes, with respective homotopy merge trees $\pi_n[T_f]$ and $\pi_n[T_g]$. A **morphism** between them $\phi : \pi_n[T_f] \rightarrow \pi_n[T_g]$ is a continuous map $\phi : T_f \rightarrow T_g$ with a family of group homomorphisms $\{\phi_\sigma : \pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(Y_{\phi(\sigma)}, e_{\phi(\sigma)})\}_{\sigma \in T_f}$ such that for all $\sigma \in T_f$ and $\varepsilon, \delta \in \mathbb{R}$ we have

$$\begin{array}{ccc} \pi_n(X_\sigma, e_\sigma) & \xrightarrow{i_\sigma^\delta} & \pi_n(X_{i^\delta(\sigma)}, e_{i^\delta(\sigma)}) \\ \downarrow \phi_\sigma & & \downarrow \phi_{i^\delta(\sigma)} \\ \pi_n(Y_{\phi(\sigma)}, e_{\phi(\sigma)}) & \xrightarrow{i_{\phi(\sigma)}^\delta} & \pi_n(Y_{i^\delta \circ \phi(\sigma)}, e_{i^\delta \circ \phi(\sigma)}) \end{array}$$

In this situation, we call the continuous map ϕ between merge trees the **indexing map**. The commuting square in the definition is inspired by the corresponding definition of a morphism of persistence modules [2].

While this may seem arbitrary at first, the reader should adopt the mantra that the shift morphisms are central to the structure of a homotopy merge tree, so morphisms should naturally respect them.

Remark 4.15. Having the same underlying indexing maps $\phi, \varphi : T_f \rightarrow T_g$ between spaces does not guarantee the same group homomorphisms ϕ_σ and φ_σ for all $\sigma \in T_f$. These two pieces of information are defined independently of one another.

Together, these definitions form a category with objects being homotopy merge trees and morphisms described as above. Composition of morphisms is given by composing the continuous functions between their merge trees, and composing group homomorphisms point-wise at their indexing point.

Remark 4.16. For a category whose objects are functors, natural transformations are often the ‘correct’ category-theoretic morphisms. While homotopy merge trees have a functorial interpretation, they should **not** be thought of as a functor category. Unlike persistence modules (which are all indexed by \mathbb{R}), the indexing categories P_f will often be different, meaning that no natural transformation exists.

Example 4.17. The ε -shift map over a given merge tree T_f can itself be thought of as a morphism $i^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_f]$. For all $\delta \in \mathbb{R}_{>0}$, the commutativity of the diagram below:

$$\begin{array}{ccc} \pi_n(X_\sigma, e_\sigma) & \xrightarrow{i_\sigma^\delta} & \pi_n(X_{i^\delta(\sigma)}, e_{i^\delta(\sigma)}) \\ \downarrow i_\sigma^\varepsilon & & \downarrow i_{i^\delta(\sigma)}^\varepsilon \\ \pi_n(X_{i^\varepsilon(\sigma)}, e_{i^\varepsilon(\sigma)}) & \xrightarrow{i_{i^\varepsilon(\sigma)}^\delta} & \pi_n(X_{i^{\varepsilon+\delta}(\sigma)}, e_{i^{\varepsilon+\delta}(\sigma)}) \end{array} \quad (4.2)$$

follows from the functorial construction of $\pi_n[T_f]$. Making this precise, in the associated poset category P_f of T_f , we have the equality of morphisms below

$$(i^\delta(\sigma) \leq i^{\varepsilon+\delta}(\sigma)) \circ (\sigma \leq i^\delta(\sigma)) = (\sigma \leq i^{\varepsilon+\delta}(\sigma)) = (i^\varepsilon(\sigma) \leq i^{\varepsilon+\delta}(\sigma)) \circ (\sigma \leq i^\varepsilon(\sigma))$$

This equality follows from the property of poset categories having at most one morphism between objects: any pair of morphisms $f : X \rightarrow Y$, $g : Y \rightarrow Z$

must compose to the unique morphism $h : X \rightarrow Z$. By the functoriality of the homotopy merge tree, applying the functor $\pi_n[T_f]$ to the above equality gives:

$$i_{i^\delta(\sigma)}^\varepsilon \circ i_\sigma^\delta = i_\sigma^{\varepsilon+\delta} = i_{i^\varepsilon(\sigma)}^\delta \circ i_\sigma^\varepsilon$$

ε -Compatibility

Definition 4.18. Let $f : X \rightarrow Y$ and $g : Y \rightarrow X$ satisfy (\star) . Two morphisms of homotopy merge trees $\alpha^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_g]$ and $\beta^\varepsilon : T_g \rightarrow T_f$ are said to be ε -compatible for some $\varepsilon \in \mathbb{R}_{>0}$ if

1. (Shifting) $\hat{g}(\alpha^\varepsilon(x)) = \hat{f}(x) + \varepsilon$ and $\hat{f}(\beta^\varepsilon(y)) = \hat{g}(y) + \varepsilon$ for all $x \in X$ and $Y \in Y$.
2. (Index Map Compatibility) The diagrams below commute:

$$\begin{array}{ccc} T_f & \xrightarrow{i^{2\varepsilon}} & T_f \\ & \searrow \alpha^\varepsilon & \nearrow \beta^\varepsilon \\ & T_g & \end{array} \quad \begin{array}{ccc} & & T_f \\ & \nearrow \beta^\varepsilon & \searrow \alpha^\varepsilon \\ T_g & \xrightarrow{i^{2\varepsilon}} & T_g \end{array}$$

3. (Group Map Compatibility) For all $\sigma \in T_f$ and $\tau \in T_g$, the diagrams below commute:

$$\begin{array}{ccc} \pi_n(X_\sigma, e_\sigma) & \xrightarrow{i_\sigma^{2\varepsilon}} & \pi_n(X_{i^{2\varepsilon}(\sigma)}, e_{i^{2\varepsilon}(\sigma)}) \\ & \searrow \alpha_\sigma^\varepsilon & \nearrow \beta_{\alpha^\varepsilon(\sigma)}^\varepsilon \\ & \pi_n(Y_{\alpha^\varepsilon(\sigma)}, e_{\alpha^\varepsilon(\sigma)}) & \end{array}$$

$$\begin{array}{ccc} & \pi_n(X_{\beta^\varepsilon(\tau)}, e_{\beta^\varepsilon(\tau)}) & \\ \nearrow \beta_\tau^\varepsilon & & \searrow \alpha_{\beta^\varepsilon(\tau)}^\varepsilon \\ \pi_n(Y_\tau, e_\tau) & \xrightarrow{i_\tau^{2\varepsilon}} & \pi_n(Y_{i^{2\varepsilon}(\tau)}, e_{i^{2\varepsilon}(\tau)}) \end{array}$$

Lemma 4.19. [14] Given two ε -compatible indexing maps of merge trees $\alpha^\varepsilon : T_f \rightarrow T_g$ and $\beta^\varepsilon : T_g \rightarrow T_f$, we have that for all $\delta \in \mathbb{R}_{>0}$

1. $\beta^\varepsilon \circ i^\delta \circ \alpha^\varepsilon = i^{\delta+2\varepsilon}$
2. $\alpha^\varepsilon \circ i^\delta \circ \beta^\varepsilon = i^{\delta+2\varepsilon}$

The following simple proof extends this lemma to the group structure overlaid on the indexing merge tree. While following almost immediately from the above definitions, it is worthwhile to help digest the notation.

Lemma 4.20. *Given two ε -compatible morphisms of merge tress $\alpha^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_g]$ and $\beta^\varepsilon : \pi_n[T_g] \rightarrow \pi_n[T_f]$, for all $\delta \in \mathbb{R}_{>0}$, $\sigma \in T_f$ and $\tau \in T_g$*

1. $\beta_{i_{\delta \circ \alpha^\varepsilon}^\varepsilon}^\varepsilon \circ i_{\alpha^\varepsilon}^\delta \circ \alpha_\sigma^\varepsilon = i_\sigma^{\delta+2\varepsilon}$
2. $\alpha_{i_{\delta \circ \beta^\varepsilon}^\varepsilon}^\varepsilon \circ i_{\beta^\varepsilon}^\delta \circ \beta_\sigma^\varepsilon = i_\sigma^{\delta+2\varepsilon}$

Proof. Consider the commutative diagram:

$$\begin{array}{ccccc}
 \pi_n(X_\sigma, e_\sigma) & \xrightarrow{i_\sigma^{2\varepsilon}} & \pi_n(X_{i^{2\varepsilon}(\sigma)}, e_{i^{2\varepsilon}(\sigma)}) & \xrightarrow{i_{i^{2\varepsilon}(\sigma)}^\delta} & \pi_n(X_{i^{\delta+2\varepsilon}(\sigma)}, e_{i^{\delta+2\varepsilon}(\sigma)}) \\
 \searrow \alpha_\sigma^\varepsilon & & \nearrow \beta_{\alpha^\varepsilon}^\varepsilon & & \nearrow \beta_{i_{\delta \circ \alpha^\varepsilon}^\varepsilon}^\varepsilon \\
 \pi_n(Y_{\alpha^\varepsilon(\sigma)}, e_{\alpha^\varepsilon(\sigma)}) & \xrightarrow{i_{\alpha^\varepsilon(\sigma)}^\delta} & \pi_n(Y_{i_{\delta \circ \alpha^\varepsilon}^\varepsilon(\sigma)}, e_{i_{\delta \circ \alpha^\varepsilon}^\varepsilon(\sigma)}) & &
 \end{array}$$

The commutativity of the triangle follows from the definition of ε -compatible morphisms, whereas the commutativity of the square follows from the definition of a morphism that says the group maps commute with the shift maps. Following the outside of the diagram we get:

$$\beta_{i_{\delta \circ \alpha^\varepsilon}^\varepsilon}^\varepsilon \circ i_{\alpha^\varepsilon}^\delta \circ \alpha_\sigma^\varepsilon = i_{i^{2\varepsilon}(\sigma)}^\delta \circ i_\sigma^{2\varepsilon} = i_\sigma^{2\varepsilon+\delta}$$

where the second relation follows from the functorial construction of the homotopy merge tree. The second claim in the lemma follows from repeating the argument, with α^ε and β^ε interchanged. \square

Interleaving Distance

Definition 4.21. Given two homotopy merge trees $\pi_n[T_f], \pi_n[T_g]$, the **interleaving distance** $d_\pi(\pi_n[T_f], \pi_n[T_g])$ between them is

$$\inf\{\varepsilon \in \mathbb{R}_{>0} \mid \text{there exist } \varepsilon\text{-compatible maps } \alpha^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_g], \beta^\varepsilon : \pi_n[T_g] \rightarrow \pi_n[T_f]\}$$

Given the existence of ε -compatible morphisms between $\pi_n[T_f]$ and $\pi_n[T_g]$, we will say that $\pi_n[T_f]$ and $\pi_n[T_g]$ are ε -**interleaved**.

Proposition 4.22. *The interleaving distance is an extended pseudo-metric. In other words, for all CW complexes X, Y, Z and maps $f : X \rightarrow \mathbb{R}$, $g : Y \rightarrow \mathbb{R}$, $h : Z \rightarrow \mathbb{R}$ satisfying (\star) , it satisfies the following:*

1. (*Identity*) $d_\pi(\pi_n[T_f], \pi_n[T_f]) = 0$
2. (*Symmetry*) $d_\pi(\pi_n[T_f], \pi_n[T_g]) = d_\pi(\pi_n[T_g], \pi_n[T_f])$
3. (*Triangle Inequality*) $d_\pi(\pi_n[T_f], \pi_n[T_h]) \leq d_\pi(\pi_n[T_f], \pi_n[T_g]) + d_\pi(\pi_n[T_g], \pi_n[T_h])$

Proof. We essentially follow the proof given in [14], extending it to the group structure.

For the identity axiom, we let the indexing map be the identity $id : T_f \rightarrow T_f$, as well as the group homomorphisms $id_\sigma : \pi_n(X_\sigma, e_\sigma) \rightarrow \pi_n(X_\sigma, e_\sigma)$. Taking this map in both directions, we get a pair of 0-compatible morphisms; so the interleaving distance $d_\pi(\pi_n[T_f], \pi_n[T_f])$ must be 0.

For the symmetry axiom, if $d_\pi(\pi_n[T_f], \pi_n[T_g]) \leq \varepsilon$, the ε -compatible morphisms α^ε and β^ε can be interchanged to show that $d_\pi(\pi_n[T_g], \pi_n[T_f]) \leq \varepsilon$ for all ε . Hence, $d_\pi(\pi_n[T_f], \pi_n[T_g]) = d_\pi(\pi_n[T_g], \pi_n[T_f])$.

Now we prove the triangle inequality. Here, we suppose $d_\pi(\pi_n[T_f], \pi_n[T_g]) = \varepsilon_1$ and $d_\pi(\pi_n[T_g], \pi_n[T_h]) = \varepsilon_2$. By the properties of the infimum, for all $\delta > 0$, there are $(\varepsilon_1 + \delta)$ -compatible maps

- $(\alpha_{fg})^{\varepsilon_1 + \delta} : \pi_n[T_f] \rightarrow \pi_n[T_g]$ and $(\beta_{gf})^{\varepsilon_1 + \delta} : \pi_n[T_g] \rightarrow \pi_n[T_f]$
- $(\alpha_{gh})^{\varepsilon_2 + \delta} : \pi_n[T_g] \rightarrow \pi_n[T_h]$ and $(\beta_{hg})^{\varepsilon_2 + \delta} : \pi_n[T_h] \rightarrow \pi_n[T_g]$

We wish to show that there $d_\pi(\pi_n[T_f], \pi_n[T_h]) \leq \varepsilon_3 := \varepsilon_1 + \varepsilon_2$. We can do this by defining $(\varepsilon_3 + \delta)$ -compatible morphisms $(\alpha_{fh})^{\varepsilon_3 + \delta} : \pi_n[T_f] \rightarrow \pi_n[T_h]$ and $(\beta_{hf})^{\varepsilon_3 + \delta} : \pi_n[T_h] \rightarrow \pi_n[T_f]$ for all $\delta \in \mathbb{R}_{>0}$ as follows:

1. the indexing maps are given by compositions

$$(\alpha_{fh})^{\varepsilon_3 + \delta} = (\alpha_{gh})^{\varepsilon_2 + \delta/2} \circ (\alpha_{fg})^{\varepsilon_1 + \delta/2} \quad \text{and} \quad (\beta_{hf})^{\varepsilon_3 + \delta} = (\beta_{gf})^{\varepsilon_1 + \delta/2} \circ (\beta_{hg})^{\varepsilon_2 + \delta/2}$$

To prove that these are $(\varepsilon_3 + \delta)$ -compatible as indexing maps, we follow exactly the working in [14]:

$$\begin{aligned} i^{2(\varepsilon_3 + \delta)} &= i^{2(\varepsilon_1 + \varepsilon_2 + \delta)} \\ &= (\beta_{gf})^{\varepsilon_1 + \delta/2} \circ i^{2(\varepsilon_2 + \delta/2)} \circ (\alpha_{fg})^{\varepsilon_1 + \delta/2} && \text{Lemma 4.19} \\ &= (\beta_{gf})^{\varepsilon_1 + \delta/2} \circ (\beta_{hg})^{\varepsilon_2 + \delta/2} \circ (\alpha_{gh})^{\varepsilon_2 + \delta/2} \circ (\alpha_{fg})^{\varepsilon_1 + \delta/2} && (\varepsilon_2 + \delta/2)\text{-compatibility} \\ &= (\beta_{hf})^{\varepsilon_3 + \delta} \circ (\alpha_{fh})^{\varepsilon_3 + \delta} && \text{by construction} \end{aligned}$$

with $i^{2(\varepsilon_3 + \delta)} = (\alpha_{fh})^{\varepsilon_3 + \delta} \circ (\beta_{hf})^{\varepsilon_3 + \delta}$ following by symmetry.

2. for $\sigma \in T_f$ and $\tau \in T_h$, the group maps are defined by compositions

$$(\alpha_{fh})_{\sigma}^{\varepsilon_3+\delta} = (\alpha_{gh})_{(\alpha_{fg})^{\varepsilon_1+\delta/2}(\sigma)}^{\varepsilon_2+\delta/2} \circ (\alpha_{fg})_{\sigma}^{\varepsilon_1+\delta/2} \quad \text{and} \quad (\beta_{hf})_{\tau}^{\varepsilon_3+\delta} = (\beta_{gf})_{(\beta_{hg})^{\varepsilon_2+\delta/2}(\tau)}^{\varepsilon_2+\delta/2} \circ (\beta_{hg})_{\tau}^{\varepsilon_2+\delta/2}$$

To show these are $(\varepsilon_3 + \delta)$ -compatible morphisms, we can use the exact same working as before, albeit with slightly more indexing and using Lemma 4.20 instead of 4.19.

$$\begin{aligned} i_{\sigma}^{2(\varepsilon_3+\delta)} &= i_{\sigma}^{2(\varepsilon_1+\varepsilon_2+\delta)} \\ &= ((\beta_{gf})^{\varepsilon_1+\delta/2} \circ i^{2(\varepsilon_2+\delta/2)} \circ (\alpha_{fg})^{\varepsilon_1+\delta/2})_{\sigma} && \text{by Lemma 4.20} \\ &= ((\beta_{gf})^{\varepsilon_1+\delta/2} \circ (\beta_{hg})^{\varepsilon_2+\delta/2} \circ (\alpha_{gh})^{\varepsilon_2+\delta/2} \circ (\alpha_{fg})^{\varepsilon_1+\delta/2})_{\sigma} && \text{by } (\varepsilon_2 + \delta/2)\text{-compatibility} \\ &= ((\beta_{hf})^{\varepsilon_3+\delta} \circ (\alpha_{fh})^{\varepsilon_3+\delta})_{\sigma} && \text{by construction} \end{aligned}$$

We have proven that for all $\delta \in \mathbb{R}_{>0}$, there exists an $(\varepsilon_3 + \delta)$ -interleaving between $\pi_n[T_f]$ and $\pi_n[T_h]$. Thus, following from the properties of infimums,

$$d_{\pi}(\pi_n[T_f], \pi_n[T_h]) \leq \varepsilon_3 = \varepsilon_1 + \varepsilon_2 = d_{\pi}(\pi_n[T_f], \pi_n[T_g]) + d_{\pi}(\pi_n[T_g], \pi_n[T_h])$$

as required. \square

Example 4.23. In Figure 4.3, consider the two CW complexes X and Y in blue and orange respectively. The numbers denote functions $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ over the cells. Denote by e_{x_i} the cell in X with $f(e_{x_i}) = i$ and e_{y_j} a cell in Y with $g(e_{y_j}) = j$. For simplicity, both functions are discrete Morse functions, where every cell is critical.




Figure 4.3: Two functions $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ over CW complexes X and Y

Throughout the filtration, three loops appear in X and one in Y . In X , these are described by

$$a = e_{x_4}e_{x_3}, \quad b = e_{x_3}e_{x_6}e_{x_3}^{-1}, \quad c = e_{x_7}e_{x_8}$$

all based at e_{x_1} . In Y , we have the loop $d = e_{y_7}$ based at e_{y_2} .

Figure 4.4 depicts an ε -interleaving between homotopy merge trees $\pi_1[T_f]$ and $\pi_1[T_g]$, with fundamental group overlaid at each point in underlying topological merge trees. Arrows between the spaces represent the maps α^ε and β^ε , where $\varepsilon = 3$. At the point $\sigma \in T_f$, the group maps are defined to take

$$\begin{array}{ccc} a & \xrightarrow{\alpha^\varepsilon(\sigma)} & d & \xrightarrow{\beta_{\alpha^\varepsilon(\sigma)}} & a \\ & \searrow & & \nearrow & \\ & & i_\sigma^{2\varepsilon} & & \end{array}$$

4.2.5 Stability Theorem

Arguably, the most crucial notion in Topological Data Analysis is that of *stability*. For any statistical invariant of a mathematical object, stability measures how perturbations in the object bound perturbations in the invariant.

In our context, the mathematical objects we study are \mathbb{R} -valued functions over cells complexes $f : X \rightarrow \mathbb{R}$; the algebraic invariants are homotopy merge trees $\pi_n[T_f]$.

Theorem 4.24. (*Stability of Homotopy Merge Trees*) *Let X be a CW complex, with $f, g : X \rightarrow \mathbb{R}$. Then*

$$d_\pi(\pi_n[T_f], \pi_n[T_g]) \leq \|f - g\|_\infty$$

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

Proof. Assume $\|f - g\|_\infty = \varepsilon$. Our aim to prove the existence of ε -compatible maps $\alpha^\varepsilon : \pi_n[T_f] \rightarrow \pi_n[T_g]$ and $\beta^\varepsilon : \pi_n[T_g] \rightarrow \pi_n[T_f]$.

Firstly, we find the indexing maps, following identically the method from [14]. Since $\|f - g\|_\infty = \varepsilon$, for all $\sigma \in T_f$ and $\tau \in T_g$ we have

$$f^{-1}((-\infty, \hat{f}(\sigma)]) \subseteq g^{-1}((-\infty, \hat{g}(\sigma) + \varepsilon]) \subseteq f^{-1}((-\infty, \hat{f}(\sigma) + 2\varepsilon])$$




Figure 4.4: ($\varepsilon = 3$) An ε -interleaving between the homotopy merge trees of filtered CW complexes described in Figure 4.3. The maps α^ε and β^ε are the ε -compatible morphisms. For any smaller ε , the compatible morphisms at σ factor through zero and cannot commute with the shift morphism $i^{2\varepsilon}$. Hence, this interleaving is minimal and $d_\pi(\pi_1[T_f], \pi_1[T_g]) = 3$

and

$$g^{-1}((-\infty, \hat{f}(\tau)]) \subseteq f^{-1}((-\infty, \hat{f}(\tau) + \varepsilon]) \subseteq g^{-1}((-\infty, \hat{g}(\tau) + 2\varepsilon])$$

Letting $X_\sigma^f = f^{-1}((-\infty, \hat{f}(\sigma)])$ and $X_\tau^g = g^{-1}((-\infty, \hat{g}(\tau)])$, we can write these inclusions more succinctly as

$$X_\sigma^f \subseteq X_{i^\varepsilon(\sigma)}^g \subseteq X_{i^{2\varepsilon}(\sigma)}^f \quad \text{and} \quad X_\tau^g \subseteq X_{i^\varepsilon(\tau)}^f \subseteq X_{i^{2\varepsilon}(\tau)}^g$$

Recall the definition of merge tree as a quotient space $T_f = (X \times \mathbb{R}) / \sim$. For each $\sigma = ([x, h]) \in T_f$, the inclusion of $x_1 \in X_\sigma^f$ into $X_{i^\varepsilon(\sigma)}^g$ induces a map on the merge trees $\iota_1 : T_f \rightarrow T_g$ sending $([x, h])$ to itself. Define our left indexing map α^ε by composition $i^\varepsilon \circ \iota_1$, the inclusion followed by the ε -shift map in T_g . Similarly, we can define a map induced by inclusion $\iota_2 : T_g \rightarrow T_f$ and right indexing map β^ε by $i^\varepsilon \circ \iota_2$.

Then, for all $\sigma = ([x, h]) \in T_f$, we have

$$\begin{aligned} \beta^\varepsilon \circ \alpha^\varepsilon(\sigma) &= i^\varepsilon \circ \iota_2 \circ i^\varepsilon \circ \iota_1([x, h]) \\ &= i^\varepsilon \circ \iota_2 \circ [(x', h + \varepsilon)] \\ &= [(x'', h + 2\varepsilon)] \\ &= i^{2\varepsilon}([x, h]) \end{aligned}$$

by construction, as x'' must lie in the same path component as x at height $h + 2\varepsilon$. Similarly, $\alpha^\varepsilon \circ \beta^\varepsilon = i^{2\varepsilon}$, so the indexing maps are ε -compatible.

Now we need to define the ε -compatible group maps lying on top of our indexing maps. Define these maps as follows:

$$\{ \alpha_\sigma^\varepsilon = \gamma \circ i_* : \pi_n(X_\sigma^f, e_\sigma) \rightarrow \pi_n(X_{\alpha^\varepsilon(\sigma)}^g, e_{\alpha^\varepsilon(\sigma)}) \}_{\sigma \in T_f}$$

where h is the canonical path from σ to $\alpha^\varepsilon(\sigma)$ and i_* is induced by the inclusion $i : X_\sigma^f \hookrightarrow X_{\alpha^\varepsilon(\sigma)}^g$. Similarly, define the right group maps as

$$\{ \beta_\tau^\varepsilon = \gamma \circ i_* : \pi_n(X_\tau^g, e_\tau) \rightarrow \pi_n(X_{\beta^\varepsilon(\tau)}^f, e_{\beta^\varepsilon(\tau)}) \}_{\tau \in T_g}$$

The last step to verify ε -compatibility of our defined maps is then to check the group map compatibility. Said diagrammatically, we must check the following

$$\begin{array}{ccc}
\pi_n(X_\sigma^f, e_\sigma) & \xrightarrow{i_\sigma^{2\varepsilon}} & \pi_n(X_{i^{2\varepsilon}(\sigma)}^f, e_{i^{2\varepsilon}(\sigma)}) \\
& \searrow \alpha_\sigma^\varepsilon & \nearrow \beta_{\alpha^\varepsilon(\sigma)}^\varepsilon \\
& & \pi_n(X_{\alpha^\varepsilon(\sigma)}^g, e_{\alpha^\varepsilon(\sigma)})
\end{array}$$

$$\begin{array}{ccc}
& \pi_n(X_{\beta^\varepsilon(\tau)}^f, e_{\beta^\varepsilon(\tau)}) & \\
& \nearrow \beta_\tau^\varepsilon & \searrow \alpha_{\beta^\varepsilon(\tau)}^\varepsilon \\
\pi_n(X_\tau^g, e_\tau) & \xrightarrow{i_\tau^{2\varepsilon}} & \pi_n(X_{i^{2\varepsilon}(\tau)}^g, e_{i^{2\varepsilon}(\tau)})
\end{array}$$

commute for all $\sigma \in T_f$ and $\tau \in T_g$. The maps in the first diagram can be rewritten as

$$\begin{array}{ccc}
\pi_n(X_\sigma^f, e_\sigma) & \xrightarrow{\gamma \circ i_*} & \pi_n(X_{i^{2\varepsilon}(\sigma)}^f, e_{i^{2\varepsilon}(\sigma)}) \\
& \searrow \gamma \circ i_* & \nearrow \gamma \circ i_* \\
& & \pi_n(X_{\alpha^\varepsilon(\sigma)}^g, e_{\alpha^\varepsilon(\sigma)})
\end{array}$$

where each γ is the change of basepoint homomorphism along canonical paths. Here, commutativity follows directly from lemma 2.9. Hence, α^ε is an ε -compatible morphism. Similarly, it follows that β^ε is an ε -compatible morphism. So $\pi_n[T_f]$ and $\pi_n[T_g]$ are ε -compatible, and

$$d_\pi(\pi_n[T_f], \pi_n[T_g]) \leq \varepsilon = \|f - g\|_\infty$$

as required. □

4.3 Application: Based Persistence

The homotopy merge tree of a filtration is universal in the sense that we can recover the information about filtrations from every base-point. The following section makes this connection explicit, and shows further that we can recover the one-dimensional persistent homology module.

4.3.1 Based Persistence

The definitions above provide a natural way to examine the persistence of a path-component in a given filtration, starting from any 0-cell e_α .

Suppose we have a CW complex X with $f : X \rightarrow \mathbb{R}$ satisfying (\star) and n -homotopy merge tree $\pi_n[T_f]$. Recall that the persistent homology over a field k can be defined functorially as

$$\mathbb{R} \xrightarrow{X_f} \mathbf{Top} \xrightarrow{H_n(_, k)} k\mathbf{Vec}$$

where $X_f : h \mapsto X_{f \leq h}$. For a locally minimal 0-cell $e_\alpha \in X$, replacing $H_n(_, k)$ with $\pi_n(_, e_\alpha)$ yields

$$\mathbb{R}_{\geq f(e_\alpha)} \xrightarrow{X_f} \mathbf{Top} \xrightarrow{\pi_n(_, e_\alpha)} \mathbf{Grp}.$$

where we restrict \mathbb{R} to $\mathbb{R}_{\geq f(e_\alpha)}$ to make sure $\pi_n(_, e_\alpha)$ is well-defined.

Geometrically, this records how the homotopy groups of the path-component containing e_α change with respect to the filtration.

For a cell $e_\lambda \in X$, let t_λ be minimal such that $X_{f \leq t_\lambda}$ contains e_α and e_λ in the same path component, with $t_\lambda = \infty$ if this never occurs. An equivalent construction replaces f with the function.

$$\begin{aligned} f_{(e_\alpha)} : X &\rightarrow \mathbb{R} \cup \{\infty\} \\ e_\lambda &\mapsto \max\{f(e_\lambda), t_\lambda\} \end{aligned}$$

Remark 4.25. $f_{(e_\alpha)}$ only *sees* a cell once it is in the same path component as e_α , ignoring the history of other path-components up to that point. Also note that $f_{(e_\alpha)}$ satisfies (\star) , so we can take its homotopy merge tree $\pi_n[T_{f_{e_\alpha}}]$.

These constructions are equivalent in the sense that we have the following commutative diagram

$$\mathbb{R}_{\geq f(e_\alpha)} \begin{array}{c} \xrightarrow{X_f} \\ \xrightarrow{X_{f_{(e_\alpha)}}} \end{array} \mathbf{Top} \xrightarrow{\pi_n(_, e_\alpha)} \mathbf{Grp} \quad (4.3)$$

This means that we can track how the homotopy groups evolve with either f or $f_{(e_\alpha)}$.

Recall that we can think of $\pi_n[T_f]$ as a functor from the poset P_f over T_f into the category \mathbf{Grp} . Define the sub-poset (P_f, e_0) of P_f given by

$$(P_f, e_\alpha) = \{x \in P_f \mid e_\alpha \leq x\}.$$

The poset (P_f, e_α) is isomorphic to $\mathbb{R}_{\geq f(e_\alpha)}$ via the isomorphism $\sigma \mapsto \hat{f}(\sigma)$. We then extend Diagram 4.3 to

$$\begin{array}{ccccc}
 & & \pi_n[T_{f(e_\alpha)}] & & \\
 & \searrow & \text{---} & \searrow & \\
 \mathbb{R}_{\geq f(e_\alpha)} & \xrightarrow{X_f} & \text{Top} & \xrightarrow{\pi_n(_, e_\alpha)} & \text{Grp} \\
 & \xrightarrow{X_{f(e_\alpha)}} & & & \\
 \downarrow \cong & & & & \\
 (P_f, e_\alpha) & \xrightarrow{i} & P_f & \xrightarrow{\pi_n[T_f]} & \text{Grp}
 \end{array} \tag{4.4}$$

where i is the inclusion functor of poset categories.

There are two important takeaways from all of this abstract nonsense. Suppose we have computed $\pi_n[T_f]$.

1. We get the induced filtration on $\pi_n(_, e_\alpha)$ for every minimal base-point e_α **for free**. In other words, we can skip the construction of $f(e_\alpha)$, and instead restrict $\pi_n[T_f]$ to the poset (P_f, e_α) .
2. We can use our interleaving metric to compare filtrations of different base-points. Namely, for two minimal 0-cells $e_\alpha, e_\beta \in X$, the quantity

$$d_\pi(\pi_n[T_{f(e_\alpha)}], \pi_n[T_{f(e_\beta)}])$$

is well-defined.

Remark 4.26. Alongside the information in Diagram 4.4, the second result can be loosely interpreted as a metric on the functors $\pi_n(_, e_\alpha)$ and $\pi_n(_, e_\beta)$.

This discussion yields a natural application in image analysis. Let e_α be a critical 0-cell, X be a CW complex and $f : X \rightarrow \mathbb{R}$ a discrete Morse function.

Persistent homology, traditionally, gives a *global* description of the image. If we want to look at how a particular region of image changes in the filtration, the above method describes how this can be achieved. Further, we can take measures of similarity between different filtered regions with the interleaving distance, leading to the following definition.

Definition 4.27. For f -critical 0-cells e_α, e_β of a discrete Morse function $f : X \rightarrow \mathbb{R}$, the n -th component **similarity number** is given by

$$d_\pi(\pi_n[T_{f(e_\alpha)}], \pi_n[T_{f(e_\beta)}])$$




Figure 4.5: Two functions $f : X \rightarrow \mathbb{R}$ and $g : Y \rightarrow \mathbb{R}$ over CW complexes X and Y

Example 4.28. Suppose there is a bridge connecting the two CW complexes in Figure 4.3. We get a new CW complex Z and discrete Morse function $h : Z \rightarrow \mathbb{R}$ (See Figure 4.5).

The two 0-cells, $e_x, e_y \in Z$, are the previous global minima of X and Y respectively. The two based homotopy merge trees

$$\pi_1[T_{h(e_x)}] \quad \text{and} \quad \pi_1[T_{h(e_y)}]$$

are respectively the same as $\pi_1[T_f]$ and $\pi_1[T_g]$ in Figure 4.4. Hence, the 1-st component similarity number is

$$d_\pi(\pi_1[T_{h(e_x)}], \pi_1[T_{h(e_y)}]) = d_\pi(\pi_1[T_f], \pi_1[T_g]) = 3$$

4.3.2 Persistent Homology

Within the homotopy merge tree lies information about the homological filtration. For any path connected topological space X and $n \in \mathbb{Z}_{\geq 1}$, there exists a group homomorphism

$$h_* : \pi_n(X, x_0) \rightarrow H_n(X; \mathbb{Z})$$

called the **Hurewicz homomorphism**. This homomorphism is natural in the sense that for any map $f : X \rightarrow Y$, the diagram

$$\begin{array}{ccc}
\pi_n(X, x_0) & \xrightarrow{f_*} & \pi_n(Y, f(x_0)) \\
\downarrow h_* & & \downarrow h_* \\
H_n(X; \mathbb{Z}) & \xrightarrow{f_*} & H_n(Y; \mathbb{Z})
\end{array}$$

commutes.

Theorem 4.29. [10] *When X is path connected $h_* : \pi_1(X, x_0) \rightarrow H_1(X; \mathbb{Z})$ is the abelianization map and surjective.*

It follows that we have a group isomorphism

$$h_* : \pi_1(X, x_0) / [\pi_1(X, x_0), \pi_1(X, x_0)] \rightarrow H_1(X; \mathbb{Z})$$

where $[\pi_1(X, x_0), \pi_1(X, x_0)]$ is the commutator subgroup of $\pi_1(X, x_0)$.

Suppose we have $\pi_1[T_f]$. Let

$$\{ \sigma_i \in T_f \mid \hat{f}(\sigma_i) = h \}$$

be representatives for the path components A_i of the sub-level set of X at height h . We then take the abelianization map split over the direct sum of path components

$$\oplus h_* : \bigoplus_{\sigma_i} \pi_1(X_{\sigma_i}, e_{\sigma_i}) \rightarrow \bigoplus_i H_1(A_i; \mathbb{Z})$$

to recover the homology groups at time h .

We can also use the Hurewicz map to naturally define based persistence for homology. Let e_α be a globally minimal 0-cell. Then we simply push forward our previous construction

$$\mathbb{R}_{\geq f(e_\alpha)} \xrightarrow{X_f} \mathbf{Top} \xrightarrow{\pi_1(_, e_\alpha)} \mathbf{Grp} \xrightarrow{h_*} \mathbf{AbGrp}$$

to recover a notion of the homological persistence of a path component.

Persistent homology, in practice, is usually only computed as a global phenomena. The theory presented above allows one to examine and compare how local regions of data change with a filtration.

Chapter 5

Computing the Homotopy Merge Tree

Computations of homotopy merge trees should thus far seem completely and hopelessly intractable. For a given T_f , there are uncountably many groups. Worse still, when we wish to compare two of these complicated objects, we need to try out uncountably many families of uncountably many group homomorphisms between them.

Discrete Morse theory provides the tools toward making these calculations possible. In this section, we unify the two main ideas of the previous section, describing how our beloved Forman Complex can assist in computing homotopy merge trees.

5.1 Sub-level Complexes

The homotopy merge tree examines sub-level *sets* of a function at each real value. This is essential to agree with the standard literature concerning the underlying topological merge tree as per [14].

In dealing with CW complexes, a stronger but more natural notion of a sub-level set is that of a sub-level complex. Recall that for a CW complex X , a sub-level complex $\bar{X}_{f \leq h}$ at $h \in \mathbb{R}$ is given by the cellular closure of its sub-level set $X_{f \leq h}$.

Conjecture A. *Let X be a regular CW complex and $f : X \rightarrow \mathbb{R}$ be a discrete*

Morse function. Then

$$X_{f \leq h} \simeq \overline{X}_{f \leq h}$$

for all $h \in \mathbb{R}$.

This is not necessarily true for general functions over CW complexes. Consider the example given in Figure 5.1. The function $f : X \rightarrow \mathbb{R}$ over the CW complex X is not a discrete Morse function; the top vertex is trying to be paired with both 1-cells. When we examine the sub-level set $X_{f \leq 4}$, we get a non-path connected space. Taking the cellular closure $\overline{X}_{f \leq 4}$ makes the sub-level complex path connected.




Figure 5.1: A non-discrete Morse function $f : X \rightarrow \mathbb{R}$ over a CW complex X , with its sub-level set $X_{f \leq 4}$ and sub-level complex $\overline{X}_{f \leq 4}$

If Conjecture A holds, an easy and natural corollary equates the homotopy merge tree of a discrete Morse function $f : X \rightarrow \mathbb{R}$ with that of the induced filtration $f_\diamond : \text{Form}(X, \mu_f) \rightarrow \mathbb{R}$ over the Forman complex.

Conjecture B. Let X be a regular cell complex with a discrete Morse function $f : X \rightarrow \mathbb{R}$ and induced acyclic partial matching $\mu_f : D \rightarrow U$. Let

$$f_\diamond : \text{Form}(X, \mu_f) \rightarrow \mathbb{R}$$

be the induced discrete Morse function on the Forman complex. Then

$$\pi_n[T_f] = \pi_n[T_{f_\diamond}]$$

Proof. Recall Corollary 3.9 states that for all $h \in \mathbb{R}$, $\text{Form}(X, \mu_f)_{f_\diamond \leq h} \simeq \overline{X}_{f \leq h}$. If Conjecture A holds, then

$$\text{Form}(X, \mu_f)_{f_\diamond \leq h} \simeq X_{f \leq h}.$$

This implies that they have the same topological merge trees $T_{f_\diamond} = T_f$. Further,

$$\pi_n(\text{Form}(X, \mu_f)_{f_\diamond \leq h}, e_\sigma) = \pi_n(X_{f \leq h}, e'_\sigma)$$

where $e_\sigma = e'_\sigma$ follows from the fact that spaces share the same minimal 0-cells. \square

5.2 Computing $\pi_1[T_{f_\diamond}]$

To compute $\pi_1[T_f]$ directly, we would require some method of computing the fundamental group in terms of all the cells in the complex.

If Conjecture *B* holds, computing $\pi_1[T_{f_\diamond}]$ is equivalent, with the advantage of using the SUBDIVIDE algorithm for acquiring the attaching maps. Here, we sketch an algorithm for finding both $\pi_1[T_{f_\diamond}]$ and the 2-skeleton $\text{Form}(X, \mu_f)^{(2)}$.

5.2.1 The Fundamental Group

Firstly, we discuss some standard results concerning the fundamental group of a CW complex. We will need the following theorem from [10].

Proposition 5.1. [10] *If a CW pair (X, A) is r -connected and A is s -connected, with $r, s \geq 0$, then the map $\pi_i(X, A) \rightarrow \pi_i(X/A)$ induced by the quotient map $X \rightarrow X/A$ is an isomorphism for $i \leq r + s$ and a surjection for $i = r + s + 1$.*

Lemma 5.2. *Let X be a path-connected CW complex and M the spanning tree of its 1-skeleton. Then $X \simeq X/M$.*

Proof. Given any CW complex X , its spanning tree M is CW sub-complex, so (X, M) form a CW pair. Since M is contractible, it is s -connected for arbitrarily high s . It follows from Proposition 5.1 that

$$\pi_i(X, M) \rightarrow \pi_i(X/M)$$

is an isomorphism for all i . Taking the homotopy long exact sequence of the pair (X, M) yields

$$\dots \longrightarrow \pi_i(M) = 0 \longrightarrow \pi_i(X) \xrightarrow{\cong} \pi_i(X, M) \longrightarrow \pi_{i-1}(M) = 0 \longrightarrow \dots$$

for all $i \geq 0$, implying $\pi_i(X, x_0) \simeq \pi_i(X/M, [x_0])$. \square

Let X be a path connected CW complex and $X^{(n)}$ be the n -skeleton, and pick a base-point x_0 . For each 2-cell e_α , let $\varphi_\alpha : S^1 \rightarrow X$ be its attaching map, where γ_α is a path from $\varphi_\alpha(s_0)$ to x_0 . Then $\gamma_\alpha \varphi_\alpha \gamma_\alpha^{-1} \in \pi_1(X, x_0)$. Let N be the normal subgroup generated by all such $\gamma_\alpha \varphi_\alpha \gamma_\alpha^{-1}$.

Proposition 5.3. [10]

1. The inclusion $i : X^{(1)} \hookrightarrow X^{(2)}$ induces a surjection

$$i_* : \pi_1(X^{(1)}, x_0) \rightarrow \pi_1(X^{(2)}, x_0)$$

with kernel N , so $\pi_1(X^{(2)}, x_0) \cong \pi_1(X^{(1)}, x_0)/N$.

2. If X is path-connected then the inclusion $i : X^{(2)} \hookrightarrow X$ induces an isomorphism $i_* : \pi_1(X^{(2)}, x_0) \rightarrow \pi_1(X, x_0)$.

In particular, when we quotient a CW complex X by its minimal spanning tree M , we have that

- $\pi_1(X^{(1)}, x_0) \cong \pi_1(X^{(1)}/M, x_0)$ is the free group over edges outside of the minimal spanning tree.
- $\pi_1(X, x_0)$ is $\pi_1(X^{(1)}, x_0)$, quotient the attaching maps of the 2-cells adjusted for base-points.

It follows that edges in X are partitioned into elements of M and free generators of $\pi_1(X^{(1)}, x_0)$.

Computation

Remark 5.4. In many applied situations, a discrete Morse function $f : X \rightarrow \mathbb{R}$ will be perturbed or constructed so that its induced acyclic partial matching is the same, while f is injective. The consequence is that a single cell is inserted at a time in a filtration of sub-level complexes. We make this assumption throughout this section.

Let X be a regular CW complex with injective discrete Morse function $f : X \rightarrow \mathbb{R}$ and induced acyclic partial matching μ_f . Let $\text{Form}(X, \mu_f)$ be the associated Forman complex, and $f_\diamond : \text{Form}(X, \mu_f) \rightarrow \mathbb{R}$ be the induced discrete Morse function.

As stated previously, any discrete Morse function satisfies (\star) , so the homotopy merge tree is well-defined. Further, every cell in $\text{Form}(X, \mu_f)$ is f_\diamond -critical, so each sub-level set is a sub-level complex.

This section discusses computing $\pi_1[T_{f_\diamond}]$, the 1-homotopy merge tree of f_\diamond , and $\text{Form}(X, \mu_f)^{(2)}$, the 2-skeleton of the Forman complex. We will sketch an

inductive algorithm by successively adding cells from $\text{Crit}(\mu_f)^{(2)}$, the critical cells in the 2-skeleton.

We know that a discrete Morse function must always have a global minimum, so the first cell in the filtration of $\text{Form}(X, \mu_f)^{(2)}$ is a critical 0-cell.

Each cells attaching map $\text{Form}(X, \mu_f)^{(2)}$ depends on flow-paths from its boundary ending at critical points. As flow-paths decrease monotonically in f , it is attached to critical cells *lower* in f value.

Thus, if we give $\text{Crit}(\mu_f)^{(2)}$ an order by f -values and attach the first $k - 1$ critical cells, the k -th will have a well-defined attaching map.

Suppose we have attached the first $k - 1$ critical cells. In other words, we have computed $\text{Form}(X^{(2)}, \mu_f)_{f_\circ \leq \lambda_{k-1}}$, where λ_{k-1} is the $(k - 1)$ -th critical value. Let e_k be the k -th critical cell, and let $f(e_k) = \lambda_k$ its critical value. The cell e_k can 0, 1 or 2 dimensional, and we consider each case below.

1. When $\dim(e_k) = 0$, the attaching map must be trivial, so we add a new path component to $\text{Form}(X^{(2)}, \mu_f)_{f_\circ \leq \lambda_{k-1}}$ containing only e_k . In other words,

$$\pi_1(\text{Form}(X, \mu_f)_{f_\circ \leq \lambda_k}^{(2)}, e_k) = 0.$$

2. When $\dim(e_k) = 1$, we have a further two cases. Suppose both ends of e_k are attached to the same path component, with minimal zero cell e_{σ_k} at time λ_k . Then e_k becomes a free generator for $\pi_1(X_{f \leq \lambda_k}, e_{\sigma_k})$. As e_k is the only cell we are attaching, we have that

$$\pi_1(\text{Form}(X, \mu_f)_{f_\circ \leq \lambda_k}^{(2)}, e_{\sigma_k}) \cong \pi_1(\text{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}, e_{\sigma_k}) * \langle e_k \rangle.$$

Now suppose that both ends map into different path components A, B of $\text{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}$. Suppose e_α and e_β are the minimal zero cells of A and B respectively and, without loss of generality, $f(e_\alpha) < f(e_\beta)$.

As e_k is the first cell to connect these path components, we add it to the minimal spanning tree M . Further, we can deformation retract e_k to a point without altering the homotopy type of $\text{Form}(X^{(2)}, \mu_f)_{f_\circ \leq \lambda_k}$, implying that

$$A \cup B \cup e_k \simeq A \vee B.$$

By Van Kampen's theorem, $\pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_k}^{(2)}, e_\alpha)$ is isomorphic to the free product

$$\pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}, e_\alpha) * \pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}, e_\beta).$$

Remark 5.5. We have a neat way to compute the shift map across changing base-points in this case. Recall that the shift map

$$\pi_1(\mathbf{Form}(X^{(2)}, \mu_f)_{f_\circ \leq \lambda_{k-1}}, e_\beta) \rightarrow \pi_1(\mathbf{Form}(X^{(2)}, \mu_f)_{f_\circ \leq \lambda_k}, e_\alpha)$$

is given by $\gamma_p \circ i_*$, where γ_p is the change of base-point homomorphism. Since the path p lies in the minimal spanning tree, and $\pi_1(X, x_0) \simeq \pi_1(X/M, x_0)$, the change of base-point homomorphism γ_p can be ignored.

3. When $\dim(e_k) = 2$, the continuity of the attaching map ensures that it maps into a single path component of $\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}$.

The attaching map $\varphi_{e_k} : S^1 \rightarrow \mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}$ can be computed by the SUBDIVIDE algorithm, which returns an array of 1-cells in $\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}$.

By removing the cells that lie in the minimal spanning tree, the attaching map can be expressed as a sequence of generators

$$e_{i_1}, e_{i_2}, \dots, e_{i_n} \in \pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}, e_{\sigma_k})$$

such that all e_{i_n} are 1-cells in $\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)} - M$. Let N be the normal subgroup generated by $\{e_{i_1}, e_{i_2}, \dots, e_{i_n}\}$. Then

$$\pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_k}^{(2)}, e_{\sigma_k}) \cong \pi_1(\mathbf{Form}(X, \mu_f)_{f_\circ \leq \lambda_{k-1}}^{(2)}, e_{\sigma_k})/N$$

by Proposition 5.3.

In any case, we can compute the Forman complex incrementally, cell by cell. This not only provides a way of computing $\pi_1[T_f]$ in its entirety, but also computes

$$\pi_1(\mathbf{Form}(X, \mu_f), e_0) \cong \pi_1(X, e_0)$$

once the final cell is added.

Chapter 6

Discussion

6.1 Contributions

The main original contributions in this thesis are the following.

1. Provide an original constructive proof that the Forman Complex exists, specifying its attaching maps. We also showed that this construction inherits a natural, well-behaved filtration.
2. Extend the theory and definitions of [14] to include information about homotopy groups, introducing the concept of the homotopy merge tree.
3. Detail original algorithms for constructing the 2-skeleton of the Forman Complex. We contextualise this with the homotopy merge tree of a space, and sketch an algorithm for computing it in the first dimension.

While we have contextualised the above results with applications in image analysis, they could potentially be useful in other areas of applied topology, particularly involving CW or simplicial complexes.

6.2 Future Work

An immediate future goal is implementation and testing on real data sets. This includes optimising and developing the algorithms in Section 3.3.3. We also need to provide a proof of Conjecture A.

Underlying the usefulness of homotopy merge trees is the interleaving distance for comparing them. Creating an efficient algorithm for computing interleaving

distance is a subsidiary goal.

When examining three dimensional images, the fundamental group is a small fraction of the complete homotopical information. A more ambitious goal is to compute the second homotopy merge tree $\pi_2[T_f]$ of a discrete Morse function.

On the discrete Morse theory side, one urgent matter is simplifying the heretical point-set topology proof of Lemma 3.1. It is suspected that an easier proof exists, using that L_c is a deformation retract, but couldn't be found prior to the submission deadline.

The work of [13] proves that persistent homology can be computed by filtering the algebraic Morse complex. It is possible that these results could be reformulated and re-proven using the Forman complex.

Toward the end of the thesis, the work of [9] regarding persistent homotopy was discovered. On the face of it, the definitions within seem related to the notion of a homotopy merge tree, but further investigation would be required to make the connection explicit.

Simplification is an important tool both to reduce calculations and to remove homological or homotopical noise. The work on homotopical simplification in Appendix A is promising and requires further study. The section was omitted primarily to length constraints of the thesis.

Methods already described in [4] could be employed to increase computability of the homotopy merge tree. The end goal, however, is to devise formal simplification methods suited to our abstract homotopical representations. We hope that these might transcend simplification obstructions in the underlying space.

Appendix A

Simplification

The rationale of discrete Morse theory is that we can reduce spaces to smaller ones based on critical cells of a function. However, if the function has many critical cells, this reduction makes little difference to computation.

Simplification is the process of removing critical cells of a given discrete Morse function, while retaining as much information as possible.

In this section, we describe how simplification can be reformulated in terms of the Forman Complex.

A.0.1 V-path Cancellation

In [7], Forman provides a convenient method to cancel critical cells based on the acyclic partial matching.

Lemma A.1. [7] *Let $f : X \rightarrow \mathbb{R}$ be a discrete Morse function over a CW complex X with induced acyclic partial matching μ_f and $e_\alpha, e_\beta \in \text{Crit}(\mu_f)$. Suppose there exists a unique V -path from the boundary of e_α to e_β ,*

$$\gamma = (e_0 \triangleleft u_0 \triangleright e_1 \triangleleft u_2 \triangleright \dots \triangleright e_k \triangleleft u_k \triangleright e_\beta)$$

Then there exists a discrete Morse function $\tilde{f} : X \rightarrow \mathbb{R}$ such that

$$\text{Crit}(\mu_f) = \text{Crit}(\mu_{\tilde{f}}) \cup \{e_\alpha, e_\beta\}$$

Moreover, $\mu_f = \mu_{\tilde{f}}$ except along the gradient path γ .

The work of [7] provides a simple method for constructing $\mu_{\tilde{f}}$. We can reverse the pairing in γ , by letting $\mu_{\tilde{f}}(e_i) = u_{i-1}$. We call this method **V -path reversal**.

A.0.2 Persistence Pairs

Often, when merge trees are not involved, we will take the sub-level *complex* filtration of a discrete Morse function $f : X \rightarrow \mathbb{R}$ over a CW complex X . Namely, the filtration

$$\begin{aligned} \bar{X}_f : \mathbb{R} &\rightarrow \text{Top} \\ h &\mapsto \bar{X}_{f \leq h}. \end{aligned}$$

In the situation that a discrete Morse function is injective, we add one cell at a time.

For a pair $(e, \mu_f(e)) \in X$ paired by $\mu_f : D \rightarrow U$, both cells are added to the filtration at time $f(\mu_f(e))$. At this time, e must be a free face of $\mu(e)$, so their effect on homology is negligible.

On the other hand, for an n -cell $e_\alpha \in \text{Crit}(\mu_f)$, its entire boundary must have already been added. This means that either

- e_α destroys the $(n - 1)$ -cycle represented by its boundary or;
- e_α creates a new n -cycle.

By pairing cells that destroy and create a given cycle, we recover the **persistence pairing**. For each pair, the difference in entry time with respect to the filtration is its **persistence**.

Remark A.2. Pairs of non-critical cells $(e_\lambda, \mu_f(e_\lambda))$ are persistence pairs with persistence equal to 0.

Remark A.3. One of many reasons for the prevalence of the persistence pairing is the vast array of algorithms for its computation. An illuminating early example is given in [6]; for a more modern, efficient algorithm see [20].

From a data analysis perspective, pairs of low persistence represent ‘homological noise’. Pairs with large persistence represent features that last longer in the filtration. **Essential cycles** are cycles that never die, and represent elements of the final homology of the space.

V -path cancellation provides a method for cancelling persistence pairs within a certain threshold, and eliminating homological noise. However, it is not the case

that all persistence pairs can be cancelled. We discuss the framework from [4] that examines persistence pair cancellation, re-examining definitions using the Forman Complex.

Let \mathcal{M} be the algebraic Morse complex of an acyclic partial matching μ_f induced by a discrete Morse function $f : X \rightarrow \mathbb{R}$. Let $\partial_{\mathcal{M}}$ be the Morse boundary operator.

Definition A.4. [4] A **close pair** in X is a pair $(e_{\alpha}^{(n)}, e_{\beta}^{(n+1)}) \in \text{Crit}(\mu_f)$ such that

- $e_{\alpha} \in \partial_{\mathcal{M}}(e_{\beta})$
- $f(e_{\alpha}) = \max\{f(e_{\gamma}) \mid e_{\gamma} \in \partial_{\mathcal{M}}(e_{\beta})\}$
- $f(e_{\beta}) = \min\{f(e_{\tau}) \mid e_{\alpha} \in \partial_{\mathcal{M}}(e_{\tau})\}$

From Chapter 2, we know that $C_*^{CW}(\text{Form}(X, \mu_f); \mathbb{Z}) = C_*(\mathcal{M}; \mathbb{Z})$. Thus, we can formulate an equivalent definition of close pair in terms of face adjacencies within the Forman complex.

Definition A.5. A **close pair** in X is a pair $e_{\alpha}, e_{\beta} \in \text{Form}(X, \mu_f)$ such that


- $e_{\alpha} \triangleleft e_{\beta}$
- $f(e_{\alpha}) = \max\{f(e_{\gamma}) \mid e_{\gamma} \triangleleft e_{\beta}\}$
- $f(e_{\beta}) = \min\{f(e_{\tau}) \mid e_{\alpha} \triangleleft e_{\tau}\}$

From this viewpoint, a close pair is a persistence pair of adjacent faces in $\text{Form}(X, \mu_f)$.

Lemma A.6. [4] *If $\alpha \in \partial_{\mathcal{M}}(\beta)$ form a close pair in \mathcal{M} , then α must create a homological cycle which is destroyed by β .*

Given Definition A.5 of close pair, this follows immediately. A consequence of Lemma A.6 is that all close pairs are persistent pairs.

Example A.7. The converse of Lemma A.6 is not true. Consider the simplicial complex with function values defined on the cells via the following filtration:



This filtration implies discrete Morse function over the space in which every cell is critical. The persistence pairing will produce (e_α, e_β) as a pair. However, if every cell is critical, then there does not exist a V -path between them, so they cannot be a close pair.

Definition A.8. [4] A **cancellable close pair** is a close pair (α, β) such that there exists exactly one V -path from β to α in the gradient vector field.

For cancellable close pairs, we can reverse the chosen V -path to attain a new gradient vector field V . A natural question is how much can a discrete Morse function can be simplified via the path cancellation operation. The following definition describes the optimal discrete Morse function.

Definition A.9. A discrete Morse function has a **δ -perfect path simplification** if a sequence of V -path reversals removes all non-essential persistence pairs with persistence less than δ .

An ∞ -perfect path simplification of f is then a sequence of V -path reversals that completely remove all non-essential critical points.

A.0.3 Theoretical Obstructions to Simplification

We know that $\text{Form}(X, \mu_f)$ has the same homology as its underlying space X . Given that $C_*^{CW}(\text{Form}(X, \mu_f); \mathbb{Z})$ is generated by $\text{Crit}(\mu_f)$, the p -th Betti number

$$\beta_p := \text{Rank}(H_p(X; \mathbb{Z}))$$

must provide a lower bound for c_p , the number of p -dimension cells in $\text{Crit}(\mu_f)$

Definition A.10. (Perfect Discrete Morse Function) A discrete morse function is **perfect** if

$$c_p = \beta_p$$

for all p .

The existence of a discrete Morse function can be obstructed by the shape of the underlying space.

Proposition A.11. *There exist spaces which do not admit any perfect discrete morse functions*

In general, any complex that is collapsible but not contractable will contain non-cancellable persistence pairs. To see this connection, we need to prove the following lemma:

Lemma A.12. *A complex is collapsible if and only if there exists a perfect discrete morse function with one critical 0-cell.*

Proof. Suppose a complex is collapsible onto a point x . So there exists a sequence of free pairs

$$(e_1, u_1), (e_2, u_2), \dots, (e_n, u_n)$$

by which the complex collapses onto x , with the collection of pairs covering $X - x$. Define an acyclic partial matching by

$$\begin{aligned} \mu : D &\rightarrow U \\ e_i &\rightarrow u_i. \end{aligned}$$


By construction, every V -path is then a sequence of simple homotopy expansions

$$\sigma_{i_1} > \tau_{i_1} < \sigma_{i_2} > \dots < \sigma_{i_m}$$

with $i_1 > i_2 > \dots > i_m$. As simple homotopy expansions can only grow into free faces, there cannot be any non-trivial closed V -paths; hence, V is a discrete gradient vector field and defines a discrete Morse function. Further, x is the only unpaired cell and, hence, a critical point.

Now suppose that a space admits a perfect discrete Morse function with only one-critical 0-cell. By the Morse homology theorem, the space must be connected. As the space has only one connected component, we can grow it from the critical zero cell by simple homotopy expanding along the V -paths. As there is only one critical cell, all other cells lie in a vector field pair, so these expansions cover the whole space. As there are no non-trivial closed V -paths, each V -path ends with a free face and can be collapsed back down. Hence, the space is collapsible. \square

We can provide an explicit example confirming proposition 4.4 in the form of the dunce hat. The figure below shows the delta complex construction of the dunce hat and its regular, barycentric subdivision:



As the attaching map of the two cells in Figure (a) is homotopic to the identity map, the dunce hat is homotopy equivalent to the disk and, hence, contractible. However, after the quotient operation, the space has no free faces, and cannot be collapsible.

This means that the dunce hat has homotopy type of a point. A perfect discrete Morse function would then generate only one essential 0-cycle in its persistence diagram.

Question A.13. What does the Forman Complex re-formulation of close pairs tell us about simplification?

Bibliography

- [1] P. Bubenik, V. Silva, and J. Scott. Metrics for generalized persistence modules. *Found. Comput. Math.*, 15(6):1501–1531, Dec. 2015.
- [2] F. Chazal, V. De Silva, M. Glisse, and S. Oudot. *The structure and stability of persistence modules*. Springer, 2016.
- [3] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, Jan 2007.
- [4] O. Delgado-Friedrichs, V. Robins, and A. Sheppard. Skeletonization and partitioning of digital images using discrete morse theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37:654–666, 2 2015.
- [5] T. Dieck. *Algebraic Topology*. EMS textbooks in mathematics. European Mathematical Society, 2008.
- [6] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, Nov 2002.
- [7] R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134(1):90 – 145, 1998.
- [8] P. Frosini and C. Landi. Size theory as a topological tool for computer vision. *Pattern Recognition and Image Analysis*, 9(4):596–603, 1999.
- [9] P. Frosini, C. Landi, and F. Mémoli. The persistent homotopy type distance. *CoRR*, abs/1702.07893, 2017.
- [10] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [11] D. Letscher. On persistent homotopy, knotted complexes and the alexander module. In *Proceedings of the 3rd Innovations in Theoretical Computer*

- Science Conference*, ITCS '12, pages 428–441, New York, NY, USA, 2012. ACM.
- [12] A. T. Lundell and S. Weingram. *The topology of CW complexes*. Springer Science & Business Media, 2012.
- [13] K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete and Computational Geometry*, 50:330–353, 2013.
- [14] D. Morozov, K. Beketayev, and G. Weber. Interleaving distance between merge trees. *Discrete and Computational Geometry*, 49:22–45, 01 2013.
- [15] V. Nanda. Discrete morse theory and localization. *Journal of Pure and Applied Algebra*, 223(2):459 – 488, 2019.
- [16] V. Nanda, D. Tamaki, and K. Tanaka. Discrete morse theory and classifying spaces. *Advances in Mathematics*, 340:723 – 790, 2018.
- [17] V. Robins. Towards computing homology from finite approximations. In *Topology proceedings*, volume 24, pages 503–532, 1999.
- [18] V. Robins, P. J. Wood, and A. P. Sheppard. Theory and algorithms for constructing discrete morse complexes from grayscale digital images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33:1646–1658, 2011.
- [19] S. Smale. A vietoris mapping theorem for homotopy. *Proc. Amer. Math. Soc.*, 8:604–610, 1957.
- [20] A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, Feb 2005.