DISCRETE MORSE THEORY

LINHAN SHEN

Abstract. In this paper, we will introduce discrete Morse theory, an analogue of Morse theory onto discrete surfaces called simplicial complexes. We will look into the applications of creating discrete Morse functions, which range from reducing simplicial complexes to analyzing their homology. Along the way, we will introduce Random Discrete Morse Theory, Vertex Refinement, and Morse Sequences.

1. Introduction

The concept of discrete Morse theory was first introduced by R. Forman in 1998 [5] as an analogue of (smooth) Morse theory onto analyzing discrete surfaces. Morse theory is used primarily with smooth manifolds. Specifically, it applies Morse functions—smooth functions with only non-degenerate critical points—onto manifolds in order to analyze their topology [11]. In this paper, we will primarily use Discrete Morse Theory for two goals: to reduce the complexity of a simplicial complexes (primarily through elementary collapse) and to make observations about their homotopy. These will come in the form of the Collapse Theorem and its corollaries, which allow us to view sequences of elementary collapses as generalized discrete vector fields, and the Discrete Morse Inequalities, which relate the critical points of a discrete Morse function with the $\mathbb{F}_2$-Betti numbers of a simplicial complex. Finally, we will introduce the concept of Morse sequences, which are an alternative way of utilizing elementary operations—specifically, elementary expansions and the inverse of elementary collapses—to represent simplicial complexes.

2. Background on Simplicial Complexes

We will begin by properly defining the discrete surfaces we are working with [10].

Definition 2.1. A *simplex* is the smallest possible Euclidean polytope in a given number of dimensions. Generally, we will define a simplex in i dimensions as $\tau^{(i)}$, or more abstractly, an i-simplex. For any simplex $\tau^{(i)}$, $\dim(\tau) := i$. The *codimension* of a simplex τ with respect to a simplex σ is defined as $\dim(\sigma) - \dim(\tau)$.

An i-simplex can be thought of as a set of vertices with cardinality $i + 1$. For example, a 0-simplex is a point, while a 2-simplex is a triangle. We will notate a simplex by the concatenation of its vertices; a 2-simplex defined over vertices v_1, v_2, v_3, and v_4 will be notated as $v_1v_2v_3v_4$. When the intention is clear, we will also write the last simplex as v_{1234} for convenience.

Definition 2.2. A simplicial complex K is defined as a subset of the power set of a collection of vertices $[v_n] := \{v_0, v_1, \ldots, v_n\}$, excluding $\emptyset$, satisfying the following conditions:

(a) If $\sigma \in K$, for every $\tau \subseteq \sigma$, $\tau \in K$;

Date: June 13, 2025.

(b) For every $v_i \in [v_n]$, $v_i \in K$.

We will denote the *vertex set* of K, or $[v_n]$, as $V(K)$. For any simplicial complex K in i dimensions, $\dim(K) := i$. The *c-vector* of K is defined as the vector $(c_0, c_1, ..., c_{dim(k)})$, where c_i is the number of i-simplices in K.

A natural thought that could arise from this definition is the question of the efficiency of writing out K. Consider the simplicial complex $K := \{v_1, v_2, v_3, v_{12}, v_{13}, v_{23}, v_{123}\}$. If the simplex v_0v_1 is contained in K, we also know that v_0 and v_1 are contained in K. In fact, the majority of the simplices in this depiction are unnecessary; all of then are contained in $v_0v_1v_2$! We will now attempt to construct a more efficient way of depicting a simplicial complex, only using the essential simplices. We start by defining these "essential simplices":

Definition 2.3. For every $\tau \subsetneq \sigma : \sigma \in K$, we say that τ is a face of σ and σ is a coface of τ. This relationship is denoted as $\tau < \sigma$.

Definition 2.4. A simplex $\sigma \in K$ is a facet of K if it is not the coface of any other simplex in K.

Lemma 2.5. *The set of facets of K properly contains every simplex of K.*

Proof. We will prove this claim by contradiction. Assume that there exists a simplex $\tau \in K$ such that for every facet σ of K, $\tau \not< \sigma$. If τ is not a face of any other simplex, Then, τ itself is a facet, and since τ contains itself, we have a contradiction. ■

Theorem 2.6. *Every simplicial complex K is uniquely represented by the set of its facets.*

Proof. We can rephrase this theorem as the following statement: two simplicial complexes K and L are identical—that is, the set of their simplices are equivalent—if and only if the set of their facets are equivalent.

We will start by proving the forward direction. Suppose that K and L are identical. Then, they both generate the same set of facets.

Now, we will prove the reverse direction through contradiction. Let S be defined as the set of the facets of K, which is equivalent to the set of the facets of L. If we assume that K and L are not identical, then at least one of the complexes will contain a simplex σ that the other does not. Without loss of generality, we will assume that K contains σ; we can make the same argument with L containing σ instead as needed. By Lemma 2.5, we know that S properly contains the simplices of both K and L, which includes σ. This implies that L, which does not contain σ, has a facet that does contain σ. This is a contradiction of Definition 2.2 ■

3. Simple Homotopy and Euler Characteristics

Now that we have a proper understanding of the construction of simplicial complexes, a logical next step is to define an equivalence relation between these structures. We will do this by creating a function that can transform a simplicial complex into a homotopy equivalent one.

Definition 3.1. A *free pair* in a simplicial complex K is a set of two simplices $\{\sigma^{(i-1)}, \tau^{(i)}\} \in K$ such that $\tau^{(i)}$ is a facet and $\tau^{(i)}$ is the only coface of $\sigma^{(i-1)}$.

Definition 3.2. We call the process of removing a free pair from a simplicial complex K an *elementary collapse*, and it will be notated as $K \searrow L$, where L is the resulting complex.

Similarly, an *elementary expansion* is the process of adding a free pair, which will be notated as $K \nearrow L$. Because of our definition of free pairs, L will always be a simplicial complex. If there exists a series of elementary expansions and collapses that transforms a complex K to a complex L, we say that K and L are homotopy equivalent, or $K \sim L$. Finally, we call a simplicial complex *collapsible* if there exists a sequence of collapses—no expansions allowed—that transforms the complex into a single 0-simplex, we call that complex *collapsible*.

There also exists an elementary operation called *elementary removal*, where a simplex (not necessarily part of a free pair) of the largest possible dimension is removed from the complex, but this operation will not be considered because it does not preserve the homology of a complex. As a shorthand rule, whenever we reference "elementary operations" from this point, we mean elementary expansions and elementary collapses.

Example 3.3. Consider the following complex K.

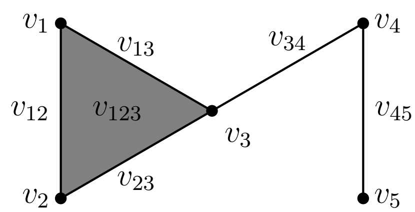

We can perform an elementary expansion, adding the free pair $\{v_{35}, v_{345}\}$, to get the following complex:

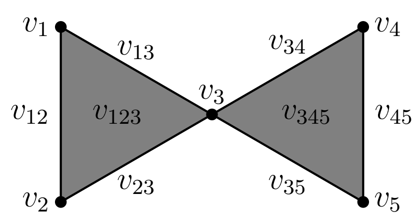

We can also collapse the following free pairs, in order, from the complex K: $\{v_{12}, v_{123}\}$, $\{v_1, v_{13}\}$, $\{v_2, v_{23}\}$, $\{v_3, v_{34}\}$, and $\{v_4, v_{45}\}$. Below is a visualization of these elementary collapses:

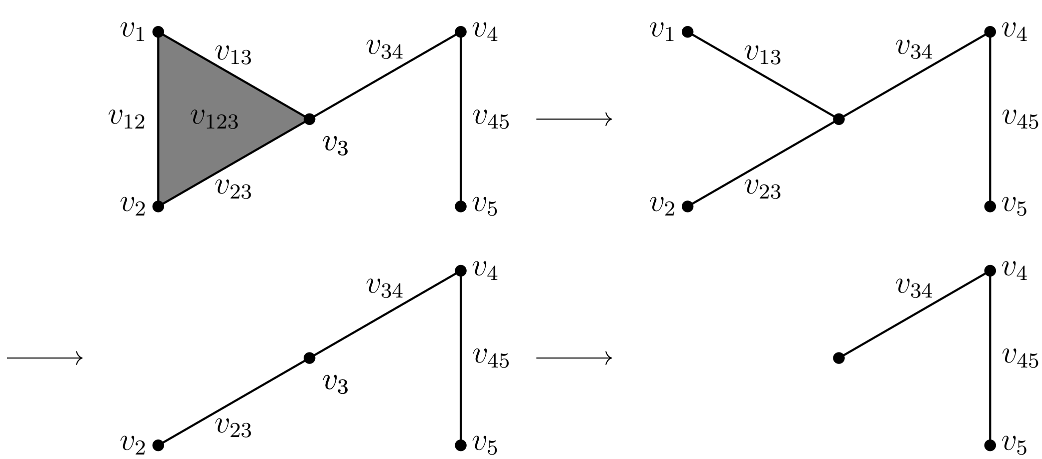

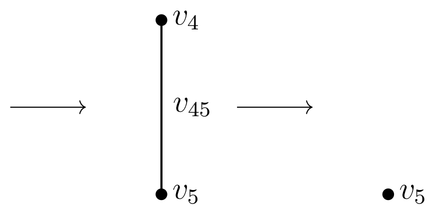

Since we are able to collapse K to a single 0-simplex, K is collapsible.

Remark 3.4. All collapsible simplicial complexes are homotopy equivalent.

Proof. Consider two collapsible simplicial complexes K and L. There exists a sequence of collapses that transforms K into a 0-simplex, and there exists a sequence of expansions that transforms a 0-simplex into L. Combining these sequences, we get a series of elementary operations that transforms K into L. Thus, $K \sim L$. ■

We can consider the Euler characteristics (denoted by χ) over these simplicial complexes to determine whether or not they are homotopy equivalent.

Remark 3.5. If we have simplicial complexes K and L such that $\chi(K) \neq \chi(L)$, then $K \nsim L$.

Proof. We will first show that the Euler characteristic of a free pair is 0. Remember that each free pair is a set of two simplices $\{\sigma^{(i-1)}, \tau^{(i)}\}$. The Euler characteristic of this pair is therefore $(-1)^{i-1} + (-1)^i = 0$. Because elementary expansions and collapses involves either adding or subtracting a free pair, it is impossible to change the Euler characteristic of a complex through elementary operations, and therefore a simplicial complex cannot be homotopy equivalent with another complex with a different Euler characteristic. ■

Example 3.6. Consider the following simplicial complex.

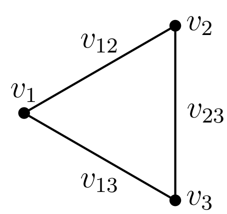

This simplicial complex is not homotopy equivalent to the complex described in Example 3.3 by the previous remark. By Remark 3.4, this also means that the above complex is not collapsible. (This result can be confirmed by the fact that there are no free pairs in the figure, so it cannot undergo elementary collapse.)

When we defined collapsible simplicial complexes, one might wonder about why we need a transformation without elementary expansions. The implication here is that there exist simplicial complexes that can be reduced to a 0-simplex through elementary collapses and expansions, yet are not collapsible. This was proved in [2] with the following example.

Example 3.7. The following simplicial complex, also known as the *Dunce hat*, is not collapsible because it has no free pairs. However, it is homotopy equivalent to the triangulation of a 3-sphere with 8 vertices and 19 facets, which is then collapsible.

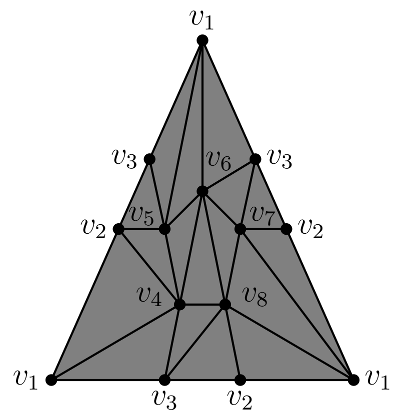

Remark 3.8. The order and types of collapses that are taken on a simplicial complex matter.

4. Discrete Morse Functions

Simplicial complexes can grow to be quite complex (pun not intended). In order to make these structures more practical, we can start by trying to find an equivalent simplicial complex that is smaller. We can do this through elementary operations, but visualizing these operations directly is time-consuming. In Example 3.3, it took us half a page to collapse a simplicial complex with only 11 elements. This is where discrete Morse functions come in.

Definition 4.1. A *discrete Morse function* f is a function $f : K \to \mathbb{R}$ that satisfies the following properties:

$$|\{\tau^{(i-1)} < \sigma : f(\tau) \geq f(\sigma)\}| \leq 1$$

and

$$|\{\tau^{(i+1)} > \sigma : f(\tau) \leq f(\sigma)\}| \leq 1$$

for every $\sigma^{(i)} \in K$. A *critical simplex* is a simplex $\sigma^i \in K$ such that

$$|\{\tau^{(i-1)} < \sigma : f(\tau) \geq f(\sigma)\}| = 0$$

and

$$|\{\tau^{(i+1)} > \sigma : f(\tau) \leq f(\sigma)\}| = 0,$$

and $f(\sigma)$ is a *critical value*. We define the *discrete Morse vector* of f as $\vec{f} := (m_0^f, m_1^f, m_2^f, \ldots, m_{\dim(K)}^f)$, where m_i^f is the number of i-dimensional critical simplices of f (also denoted as m_i if the function is clear). Every simplex τ that is not critical is a *regular simplex*, and $f(\tau)$ is a *regular value*. The discrete Morse functions g on K are said to be *Forman equivalent* to f if for every pair of simplices $\sigma^{(i)} < \tau^{(i+1)}$ in K, $f(\sigma) < f(\tau) \iff g(\sigma) < g(\tau)$.

We can think of a discrete Morse function as a function that, generally, assigns higher values to higher-dimension simplices. Each simplex is given an "exception" to this rule, allowing us to create pairs of simplices that are each others' "exceptions".

Definition 4.2. A discrete Morse function $f : K \to \mathbb{R}$ is called *excellent* if it is 1-1 on the critical simplices of K.

Lemma 4.3. *Let $f : K \to \mathbb{R}$ be a discrete Morse function. Then there is an excellent discrete Morse function $g : K \to \mathbb{R}$ which is Forman equivalent.*

Proof. Let $\sigma, \tau \in K$ be critical simplices such that $f(\sigma) = f(\tau)$. Then, we construct a discrete Morse function $f' : K \to \mathbb{R}$ such that $f'(\gamma) = f(\gamma)$ for all $\gamma \neq \tau$ and $f'(\tau) = f(\tau) + \epsilon$, where $f(\tau) < f(\tau) + \epsilon$ and $f(\tau) + \epsilon$ is strictly less than the smallest value in f greater than $f(\tau)$. Repeating this process, we will eventually create a discrete Morse function g that is 1-1 on the critical simplices of K. ∎

Example 4.4. Here is an example of a discrete Morse function on the simplicial complex from Example 3.3.

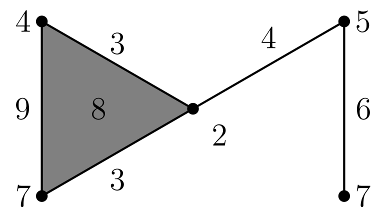

Using these functions, we can create gradient vector fields that outline pairs of regular simplices.

Definition 4.5. The *induced gradient vector field* V_f for a discrete Morse function f on K is defined as follows:

$$V_f := \{(\sigma^{(i)}, \tau^{(i+1)}) : \sigma < \tau, f(\sigma) \geq f(\tau)\}.$$

Each $(\sigma^{(i)}, \tau^{(i+1)}) \in V_f$ is called a *vector* or a *pair*, with σ being the *tail* and τ being the *head*.

Lemma 4.6. *Let K be a simplicial complex, with $\sigma \in K$ being a simplex and f a discrete Morse function on K. Then, exactly one of the following holds:*

(a) σ is the tail of exactly one vector in V_f;
(b) σ is the head of exactly one vector in V_f;
(c) σ is a critical simplex.

Proof. First, note that set of all regular simplices from f is equivalent to the set of simplices contained in the vectors of V_f; this proves the uniqueness of (3).

Now, writing that σ is a regular simplex defined by $\sigma = v_1 v_2 v_3 \dots v_{i-1} v_i$ and renaming the elements if necessary, suppose by contradiction that $\tau = v_1 v_2 v_3 \dots v_i v_{i+1} > \sigma$ and $\nu = v_1 v_2 v_3 \dots v_{i-2} v_{i-1} < \sigma$ satisfy $f(\tau) \leq f(\sigma) \leq f(\nu)$. In other words, suppose that both conditions (1) and (2) apply to σ. Next, observe that $\tilde{\sigma} := v_1 v_2 v_3 \dots v_{i-1} v_{i+1}$ satisfies $\nu < \tilde{\sigma} < \tau$. Because $\nu < \sigma$ and $f(\nu) \geq f(\tau)$, $v < \tilde{\sigma}$ tells us that $f(\nu) < f(\tilde{\sigma})$, and $f(\tilde{\sigma}) < f(\tau)$ follows similarly. We can then write that

$$f(\tau) \leq f(\sigma) \leq f(\nu) < f(\tilde{\sigma}) < f(\tau),$$

which is a contradiction. Therefore, if σ is a regular simplex, exactly one of (1) and (2) must be true. ∎

Theorem 4.7. *Two discrete Morse functions f and g defined on a simplicial complex K are Forman equivalent if and only if f and g induce the same gradient vector field.*

Proof. We will start by proving the forward direction. If we let f and g be Forman equivalent, then for any $\sigma^{(i)} < \tau^{(i+1)} \in K$, $f(\sigma) < f(\tau) \iff g(\sigma) < g(\tau)$. This implies that $f(\sigma) \geq f(\tau) \iff g(\sigma) \geq g(\tau)$, meaning that $(\sigma, \tau) \in V_f \iff (\sigma, \tau) \in V_g$.

Now, we will prove the reverse direction. Suppose that $V := V_f = V_g$. By Lemma 4.6, any simplex in K is either critical or in exactly one pair in V. We want to prove that any two simplices $\sigma^{(i)} < \tau^{(i+1)} \in K$ must satisfy $f(\sigma) \geq f(\tau) \iff g(\sigma) \geq g(\tau)$. This gives us the following cases:

(a) Suppose that $(\sigma, \tau) \in V$. This implies that $f(\sigma) \geq f(\tau) \iff g(\sigma) \geq g(\tau)$;
(b) Without loss of generality, suppose that σ is not in a pair in V while τ is in a pair in V. Because σ is not in a pair in V, it is a critical simplex for both f and g. As a result, $f(\sigma) < f(\tau)$ and $g(\sigma) < g(\tau)$;
(c) Suppose that σ and τ are in different pairs in V. Then $f(\sigma) < f(\tau)$ and $g(\sigma) < g(\tau)$;
(d) Suppose that neither σ nor τ are in a pair in V. Then they are both critical, and $f(\sigma) < f(\tau)$ and $g(\sigma) < g(\tau)$.

In all of these cases, $f(\sigma) \geq f(\tau) \iff g(\sigma) \geq g(\tau)$. ■

Corollary 4.8. *If two discrete Morse functions f, g defined on a simplicial complex K are Forman equivalent, then they share the same critical simplices.*

Proof. From Theorem 4.7, we have that V_f and V_g are equivalent. This implies that the set of simplices that are not in a pair in V_f—that is, are critical in f—are also not in a pair in V_g, and thus critical in g. ■

Example 4.9. The induced gradient vector field of the discrete Morse function described in Example 4.4 is as follows:

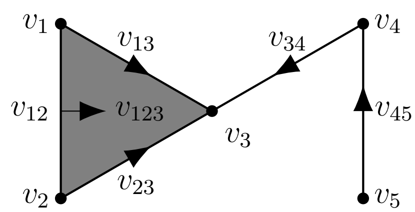

Note that if we think of these vectors as free pairs to perform elementary collapses on, the induced gradient vector field described a series of (unordered) collapses that allows us to simplify the complex. Specifically, it becomes the sequence of collapses described in Example 3.3.

Example 4.10. There are many scenarios, however, when the collapses we want to perform are obstructed. Consider the following induced gradient vector field.

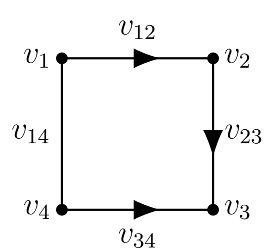

Because there are no free pairs, we are unable to perform any collapses. We are "obstructed" by the simplex v_{14}, since if it were removed, the resulting simplicial complex would be collapsible. We will see a way to view these obstructions in Section 6.

5. Random Discrete Morse Theory

Definition 5.1. A *optimal discrete Morse function* is a discrete Morse function f on a simplicial complex K such that for any discrete Morse function g on K,

$$m_i^f \leq m_i^g : 0 \leq i \leq \dim(K).$$

For example, the discrete Morse function described in Example 4.4 is optimal.

For the purpose of reducing simplicial complexes, our goal is to do as many elementary collapses as possible; thus, we prefer to have optimal discrete Morse functions, since they have the most regular pairs. [1] documents an algorithm that outputs the discrete Morse vector of a random discrete Morse function on a simplicial complex K: The idea of this

---

Algorithm 1 Random Discrete Morse Algorithm

INPUT: An i-dimensional simplicial complex K, given by its list of facets.

(0) Initialize $m_0 = m_1 = m_2 = \ldots m_i = 0$.
(1) Is the complex empty? If yes, then STOP; otherwise, go to (2).
(2) Are there free pairs? If yes, go to (3); if no, go to (4).
(3) *Elementary Collapse*: Pick one free pair uniformly at random and delete it. Go back to (1).
(4) *Critical Face*: Pick one of the top-dimensional faces uniformly at random and delete it from the complex. If n is the dimension of the face just deleted, increment m_n by 1 unit. Go back to (1).

OUTPUT: The resulting discrete Morse vector $(m_0, m_1, m_2, \ldots, m_i)$.

---

algorithm is simple: we randomly "remove" free pairs by turning them into regular pairs until there are no more; then, we take a random simplex of the highest dimension, label it as a critical simplex, and repeat.

This algorithm is not very computationally complex, and often returns optimal discrete Morse vectors in general; however, there exist simplicial complexes that pose an issue, as seen in the following example from [10].

Proposition 5.2. *For every $\epsilon > 0$, there exists a simplicial complex G_ϵ such that the probability that the Random Discrete Morse Algorithm yields an optimal discrete Morse vector of G_ϵ is less than ϵ.*

Proof. For any $\epsilon > 0$, choose $n \in \mathbb{N}$ such that $\frac{6}{n+6} < \epsilon$. Consider the simplicial complex

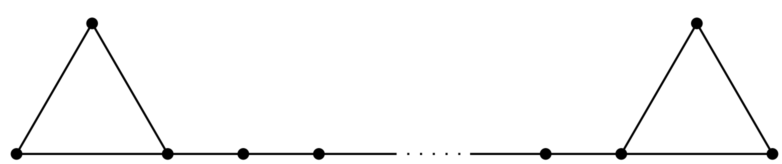

with at least n edges between the two cycles. Note that there are no free pairs in the simplicial complex.

In order for a random discrete Morse vector by the Random Discrete Morse Algorithm to be optimal, the first removed edge must be one from the two cycles of length 3. (This results in a discrete Morse vector of $(1, 2)$.) The probability of this happening is $\frac{6}{n+6}$; thus, our chances of obtaining an optimal discrete Morse vector is less than ϵ. ∎

This observation tells us that discrete Morse functions can still be unreliable for finding elementary collapses. In the next section, we will discuss ways to improve our toolbox for reducing simplicial complexes.

6. Generalized Discrete Morse Functions and the Collapse Theorem

Tabling the discussion of collapses for the moment, we will note that the discrete Morse functions are somewhat limited in the information that they provide; they only partition a simplicial complex into groups of one (critical simplices) and two (pairs of regular simplices). From this idea, we get the concept of generalized discrete Morse functions. These functions were first introduced by R. Freij in 2009 [6].

Definition 6.1. Let K be a simplicial complex. Then, for $\alpha, \beta \in K$, the *interval* $[\alpha, \beta]$ is defined as

$$[\alpha, \beta] := \{\gamma \in K : \alpha \subseteq \gamma \subseteq \beta\}$$

Note that the interval $[\alpha, \beta] \neq \emptyset \iff \alpha < \beta$.

Definition 6.2. Let $g : K \to \mathbb{R}$ be a function on a simplicial complex K that satisfies the following conditions:

(a) $g(\sigma) \leq g(\tau)$ for all $\sigma < \tau \in K$;
(b) A partition W of K into intervals exists such that for any simplices σ, τ in the same interval, $g(\sigma) = g(\tau)$.

Then, we call g a *generalized discrete Morse function* and W its *generalized discrete vector field*. Finally, any interval that only contains one simplex σ is called *singular*, σ is a *critical simplex*, and $g(\sigma)$ is a *critical value*. Any interval that is not singular is *regular* and its simplices are *regular simplices*.

Example 6.3. This simplicial complex is partitioned into into 8 different intervals, and its critical values are 1, 2, 5, and 6.

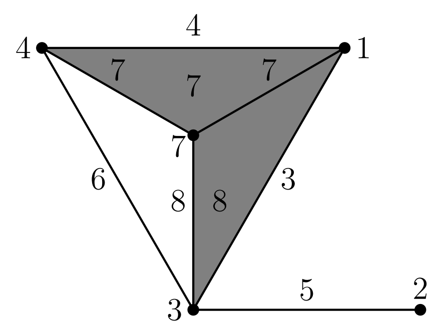

Generalized discrete Morse functions are useful because they provide an easier way to visualize collapses. We will provide a way to do so in Corollary 6.9.

Definition 6.4. A *subcomplex* L of a simplicial complex K is a simplicial complex such that $L \in K$.

Definition 6.5. Consider a discrete Morse function $f : K \to \mathbb{R}$. For any $c \in \mathbb{R}$, the *level subcomplex* $K(c)$ is the subcomplex consisting of any simplices τ along with its faces such that $f(\tau) \leq c$. In other words,

$$K(c) = \bigcup_{f(\tau) \leq c} \bigcup_{\sigma \leq \tau} \sigma.$$

The following lemma tells us that we can slightly alter any discrete Morse function f to make it 1-1 without changing some specific level subcomplexes.

Lemma 6.6. *Let $f : K \to \mathbb{R}$ be a discrete Morse function and $[\alpha, \beta] \subseteq \mathbb{R}$ be an interval that contains no critical values. Then there is a discrete Morse function f' on K that satisfies the following conditions:*

(a) f' is 1-1 on $[\alpha, \beta]$;
(b) f' has no critical values in $[\alpha, \beta]$;
(c) The level subcomplexes $K_f(\gamma)$ and $K_{f'}(\gamma)$ are equivalent for $\gamma \in \{\alpha, \beta\}$, where $K_f(\gamma)$ refers to the level subcomplex of K defined by γ with respect to f;
(d) $f = f'$ outside of $[\alpha, \beta]$.

Proof. We will attempt to construct such a function f'.

We start by copying over all values outside of $[\alpha, \beta]$, since they are identical from *(d)*. Next, we will consider every value κ in f in $[\alpha, \beta)$. Let the next largest value in f be λ. For all simplices with a value of κ, we will reassign each of their values for f' in such a way that one of the highest-dimension simplices has a value of κ, all of the simplices have a different value in the interval $[\kappa, \lambda]$, and all $i+1$-simplices have a lesser value than all i-simplices. Finally, we will define δ as the greatest value in f' less than β. We will then perform a similar process for the simplices with a value of β in f, with one of the lowest-dimension simplices being assigned a value of β in f'. This preserves all regular pairs, and contains all of the values within the range $[\alpha, \beta]$, so it satisfies all of the conditions. ∎

The following theorem is known as the *Collapse Theorem*.

Theorem 6.7. *Let $f : K \to \mathbb{R}$ be a discrete Morse function and $[\alpha, \beta] \subseteq \mathbb{R}$ be an interval that contains no critical values. Then $K(\beta) \searrow K(\alpha)$.*

Proof. Applying Lemma 6.6 and with an abuse of notation, we assume that f is 1-1. Because of this, we can split $[\alpha, \beta]$ into subintervals such that each subinterval contains exactly one regular value. Consider the simplex $\sigma^{(i)} \in K$ such that $f(\sigma)$ is the largest regular value in $[\alpha, \beta]$. Then, Lemma 4.6 tells us that exactly one of the following conditions holds:

(a) There exists $\tau^{(i+1)} > \sigma$ such that $f(\tau) \leq f(\sigma)$;
(b) There exists $\nu^{(i-1)} < \sigma$ such that $f(\nu) \geq f(\sigma)$.

For the first case, suppose that there exists $\tau^{(i+1)} < \sigma$ such that $f(\tau) \leq f(\sigma)$. We will prove that $\{\tau, \sigma\}$ is a free pair in $K(\beta)$ through contradiction.

Suppose that there exists a second coface $\tilde{\tau}^{(i+1)} > \sigma$ with $\tilde{\tau} \in K(\beta)$. Because $f(\tau) \leq f(\sigma)$, we know that either $f(\tilde{\tau}) \geq f(\sigma)$. However, this is impossible, since $f(\sigma)$ is the largest value in the interval $[\alpha, \beta]$. Thus, $\{\tau, \sigma\}$ is a free pair, and $K(b) \searrow K(b) - \{\tau, \sigma\}$. The same argument can be applied to the second case, so we can apply this argument to every subinterval in $[\alpha, \beta]$, removing each subinterval as we "collapse" them. ∎

This theorem tells us that for the purpose of collapsing simplicial complexes, we only have to consider the collapsability of critical subcomplexes.

Lemma 6.8. *For every generalized discrete vector field, there is a standard discrete vector field that further partitions every non-singular, non-empty interval into pairs.*

Proof. We will prove this lemma by showing an algorithm known as *vertex refinement* that partitions an interval into pairs. ∎

Algorithm 2 Vertex Refinement

INPUT: A non-singular, non-empty interval $[\alpha, \beta]$ (which implies $\alpha < \beta$).

(0) Initialize a set v with cardinality 0.
(1) Choose a random 0-simplex $\sigma \in \beta - \alpha$.
(2) Is the interval empty? If yes, then STOP; otherwise, go to (3).
(3) Choose a random simplex $\gamma \in [\alpha, \beta]$ and add the pair $\{\gamma - \{\sigma\}, \gamma \cup \{\sigma\}\}$ to v. Go back to (2).

OUTPUT: The resulting set of pairs v that fully partition the interval $[\alpha, \beta]$.

The following result is immediate from Theorem 6.7 and Lemma 6.8.

Corollary 6.9. *Let K be a simplicial complex with a generalized discrete vector field W, and let $K' \subseteq K$ be a subcomplex of K. If $K - K'$ is a union of non-singular, non-empty intervals in W, then $K \searrow K'$.*

Example 6.10. Let us visualize our results with the following generalized discrete Morse function over a simplicial complex K. [10]

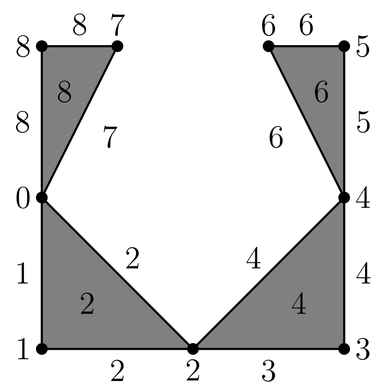

The sole critical value is 0; indeed, if we remove the critical simplex with a value of 0, the remaining simplices can be expressed as the union of the intervals with values 1, 2, 3, 4, 5, 6, 7, and 8. Therefore, K can be collapsed to the 0-simplex with a value of 0.

7. Simplicial Unreduced Homology

Another characteristic that remains constant across elementary operations are a simplicial complex's homology; we could start by thinking about ways to count the number of holes in a complex. An example of a 1-dimensional hole would be the sequence of 1-simplices $\{v_{34}, v_{45}, v_{35}\}$ from Example 3.6; we have a boundary formed by 1-simplices but not 2-simplices to "fill in the gaps". In order to calculate the number of these types of holes, we first need to find all of these possible "boundaries".

To do this, we will create vector spaces $\mathbb{U}_n^{c_i}$ over $\mathbb{F}_2$ for every $i : 1 \leq i \leq \dim(k)$ in a simplicial complex K that represent the linear combinations of n vectors, where each of i-simplices contained in K are represented as a vector. (The vector space being over $\mathbb{F}_2$ means that the coefficients of the vectors will only be the numbers 0 or 1, with $0 + 0 = 0$, $0 + 1 = 1$, and $1 + 1 = 0$.) For example, the vector space over K over 1-simplices would be $\mathbb{U}_1^3 := \{\vec{0}, v_{12}, v_{13}, v_{23}, v_{12} + v_{13}, v_{12} + v_{23}, v_{13} + v_{23}, v_{12} + v_{13} + v_{23}\}$, where $\vec{0}$ is the zero vector.

In order to create a chain complex utilizing these vector spaces, we need to define a linear transformation $\partial_i : \mathbb{U}^{c_i} \to \mathbb{U}^{c_{i-1}}$. For each simplex, this linear transformation will associate its boundary.

Definition 7.1. Let $\sigma \in K_i$ and write $\sigma = \sigma_{n_0 n_1 n_2 \ldots n_i}$. For $i = 0$, define $\partial_o : \mathbb{U}^{c_0} \to 0$ by $\partial_0 \to 0$, a zero matrix of the appropriate size. For $i \geq 0$, define the *boundary operator* $\partial_i : \mathbb{U}^{c_i} \to \mathbb{U}^{c_{i-1}}$ by

$$\partial_i(\sigma) := \sum_{0 \leq j \leq i} (\sigma - \sigma_{n_j}).$$

Because our vector spaces are over $\mathbb{F}_2$, $\partial\partial = 0$.

Example 7.2. We will give an example of this by calculating $\partial\partial(v_{123})$:

$$\begin{aligned}
\partial\partial(v_{123}) &= \partial(v_{12} + v_{23} + v_{13}) \\
&= \partial(v_{12}) + \partial(v_{23}) + \partial(v_{13}) \\
&= v_1 + v_2 + v_2 + v_3 + v_0 + v_3 \\
&= 2v_1 + 2v_2 + 2v_3 \\
&= \vec{0}
\end{aligned}$$

Thus, we have defined a chain complex.

Definition 7.3. The ith *unreduced* $\mathbb{F}_2$*-homology* of K is defined as the vector space

$$H_i(K; \mathbb{F}_2) := \mathbb{U}^{\operatorname{null} \partial_i - \operatorname{rank} \partial_{i+1}}.$$

The ith $\mathbb{F}_2$-*Betti number* of K is defined as

$$b_i(K; \mathbb{F}_2) := \operatorname{null} \partial_i - \operatorname{rank} \partial_{i+1}.$$

The ith Betti number of a simplicial complex K refers to the number of i-dimensional holes in K. For example, in Example 3.6, there is one 1-dimensional hole (since v_{123} is "missing") and one 0-dimensional hole (since we only have one distinct "object"), so $b_1 = 1$ and $b_0 = 1$.

We can calculate b_n using a well-known theorem of linear algebra called the *Rank-Nullity Theorem*:

Theorem 7.4. *Let $\partial : \mathbb{U} \to \mathbb{V}$ be a linear transformation between two finite vector spaces. Then,* $\operatorname{rank}(\partial) + \operatorname{null}(\partial) = \dim(\mathbb{U})$.

For the sake of convenience, we will refer to the $\mathbb{F}_2$-Betti numbers as Betti numbers from here on. However, these differ from standard Betti numbers, which contain more information.

The following lemma is given by [4]:

Lemma 7.5. *If simplicial complexes K and L are homotopy equivalent, then for every nonnegative integer i, $b_i(K) = b_i(L)$.*

Before we move on to talking about applying Discrete Morse Theory, we must first prove one lemma: the addition of a p-simplex will either increase b_p by 1 or decrease b_{p-1} by 1.

Lemma 7.6. *Let K be a simplicial complex and $\sigma^{(p)} \in K$ a facet of K, where $p \geq 1$. If $K' := K - \{\sigma\}$ is a simplicial complex, then one of the following holds:*

(a) One hole is added to the complex, or $b_p(K) = b_p(K') + 1$ and $b_{p-1}(K) = b_{p-1}(K')$;
(b) One hole is removed from the complex, or $b_p(K) = b_p(K')$ and $b_{p-1}(K) + 1 = b_{p-1}(K')$.

Furthermore, $b_i(K) = b_i(K')$ for all $i \neq p-1, p$.

Proof. Let $(\mathbb{U}_*, \partial_*)$ and $(\mathbb{U}'_*, \partial'_*)$ be the associated chain complexes for K and K', respectively. Since σ is a facet, it follows that $\partial_i = \partial'_i$ for all $i \neq p$. Thus, it is generally true that $b_i(K) = b_i(K')$ for all i, with the only possible exceptions being $b_p = \text{null}(\partial_p) - \text{rank}(b_{p+1})$ and $b_p - 1 = \text{null}(\partial_{p-1}) - \text{rank}(\partial_p)$. We will consider the cases where $\sigma \in \ker(\partial_p)$ and $\sigma \notin \ker(\partial_p)$.

For the first case, suppose that $\sigma \in \ker(\partial_p)$. Then $\text{Im}(\partial_p) = \text{Im}(\partial'_p)$, and $\text{rank}(\partial_p) = \text{rank}(\partial'_p)$ follows as a result. Because $\sigma \notin \mathbb{U}'_*$, we know that $\sigma \notin \ker(\partial'_p)$ and $\text{null}(\partial_p) = \text{null}(\partial'_p) + 1$. Finally, we reach

$$\begin{aligned} b_p(K) &= \text{null}(\partial_p) - \text{rank}(\partial_{p+1}) \\ &= \text{null}(\partial'_p) - \text{rank}(\partial_{p+1}) + 1 \\ &= b_p(K') + 1 \end{aligned}$$

and $b_{p-1}(K) = \text{null}(\partial_{p-1}) - \text{rank}(\partial_p) = \text{null}(\partial'_p) - \text{rank}(\partial'_p) = b_{p-1}(K')$.

For the second case, suppose that $\sigma \notin \ker(\partial_p)$, so that $\partial_p(\sigma)$ is nonzero and contained within $\text{Im}(\partial_p)$. Then $\ker(\partial_p) = \ker(\partial'_p)$, which gives us $b_p(K) = \text{null}(\partial_p) - \text{rank}(\partial_{p+1}) = \text{null}(\partial'_p) - \text{rank}(\partial'_{p+1}) = b_p(K')$. Next, since $\partial_p(\sigma)$ is nonzero and σ is a basis element, $\text{rank}(\partial_p) = \text{rank}(\partial'_p) + 1$. We reach

$$\begin{aligned} b_{p-1}(K) &= \text{null}(\partial_{p-1}) - \text{rank}(\partial_{p-1}) \\ &= \text{null}(\partial'_{p-1}) - \text{rank}(\partial'_{p-1}) - 1 \\ &= b_{p-1}(K') - 1. \end{aligned}$$

■

8. Discrete Morse Theory and Homology

We can now apply discrete Morse functions to our discussion on homology. The following theorem is known as the *Weak discrete Morse inequalities*.

Theorem 8.1. *Let $f : K \to \mathbb{R}$ be a discrete Morse function, and let $n := \dim(K)$. Then, we have the following inequalities:*

(a) For all $i = 0, 1, 2, \ldots, n$, $m_i \geq b_i$, and

(b) $\sum_{i=0}^{n} (-1)^i m_i = \chi(K)$.

Proof. We will start by proving the first part of the Weak Discrete Morse Inequalities. Since it does not affect the values of m_i, we will use Lemma 4.3 to assume that f is excellent. Now, we will use strong induction on c, the number of simplices of K. For $c = 1$, the only simplex is a 0-simplex, giving us $m_0 = b_0 = 0$. This proves the base case.

Assume the inductive hypothesis that there is a $c \geq 1$ such that for every simplicial complex with $1 \leq n \leq c$ simplices, any discrete Morse function satisfies $m_i \geq b_i$. Now, suppose that K is any simplicial complex with $c+1$ simplices and $f : K \to \mathbb{R}$ is an excellent discrete Morse function. Now, consider the value $m = \max\{f\}$. If m is a critical value with a corresponding critical p-simplex σ, we can now consider the subcomplex $K' := K - \{\sigma\}$ and the function $f' = f|_{K'} : K' \to \mathbb{R}$; in other words, f' is equivalent to f on the values of the simplices of K'. f' is a discrete Morse function with one less critical p-simplex than f, and by Lemma 7.6, the removal of σ results in either $b_p(K) = b_p(K') + 1$ or $b_{p-1}(K) = b_{p-1}(K') - 1$, while $b_i(K) = b_i(K')$ for $i \neq p-1, p$. Plugging in the inductive hypothesis, we obtain the equation

$$b_p(K) - 1 = b_p(K') \leq m_p(K') = m_p(K) - 1$$

for the first case and

$$b_{p-1}(K) + 1 = b_{p-1}(K') \leq m_{p-1}(K') = m_{p-1}(K)$$

for the second case.

Otherwise, if σ is a regular simplex, it must be part of a free pair. By Lemma 7.5, removing this free pair through elementary collapse preserves the Betti numbers of the complex. By the inductive hypothesis, we obtain the equation

$$b_i(K) = b_i(K') \leq m_i(K') = m_i(K)$$

for all i.

Now, we will prove the second part of the Weak Discrete Morse Inequalities. Since every regular pair consists of two simplices $\sigma^{(i)}$ and $\tau^{(i+1)}$, its Euler characteristic is $(-1)^i + (-1)^{i+1} = 0$. Thus, if we take the sum $\chi(K) = \sum_{i=0}^{n}(-1)^i c_i$ and cancel out every regular pair, we will be left with the alternating sum $\sum_{i=0}^{n}(-1)^i m_i$. ■

There also exists the *Strong Discrete Morse Inequalities*, the proof of which requires the following three results. These are by [5], [9], and [8], respectively.

Lemma 8.2. *Let K be a simplicial complex of dimension n and $f : K \to \mathbb{R}$ be a discrete Morse function. Then*

(a) K is homotopy equivalent to a CW complex X where the p-cells of X are in bijective correspondence with the set of critical p-simplices of f;
(b) $b_i(X) = b_i(K)$ for all nonnegative integers i;
(c) For each nonnegative integer p, we have

$$b_p(X) - b_{p-1}(X) + b_{p-2}(X) - \ldots + (-1)^p b_0(X) \leq c_p - c_{p-1} + c_{p-2} - \ldots + (-1)^p c_0.$$

This lemma essentially tells us that for any simplicial complex K, there exists a homotopy equivalent CW complex X such that $c_i = m_i$ for all i. (A definition of CW complexes can be found in [7].)

We can now state the *Strong Discrete Morse Inequalities.*

Theorem 8.3. *Let K be a simplicial complex of dimension n and $f : K \to \mathbb{R}$ be a discrete Morse function. Then, for each $p = 0, 1, 2, \ldots, n, n+1$, we have the following equation:*

$$b_p - b_{p-1} + b_{p-2} - \ldots + (-1)^p b_0 \leq m_p - m_{p-1} + m_{p-2} - \ldots + (-1)^p m_0.$$

Proof. From Lemma 8.2, we know there exists a CW complex X with p-cells in bijective correspondence with the critical p-simplices of K. Lemma 8.2 and Theorem 8.1 yield the following:

$$\begin{aligned} & b_p(K) - b_{p-1}(K) + b_{p-2}(K) - \ldots + (-1)^p b_0(K) \\ =& b_p(X) - b_{p-1}(X) + b_{p-2}(X) - \ldots + (-1)^p b_0(X) \\ =& c_p - c_{p-1} + c_{p-2} - \ldots + (-1)^p c_0 \\ =& m_p - m_{p-1} + m_{p-2} - \ldots + (-1)^p m_0. \end{aligned}$$

■

Example 8.4. We will give an example of the Discrete Morse Inequalities with the following simplicial complex K and discrete Morse function $f : K \to \mathbb{R}$.

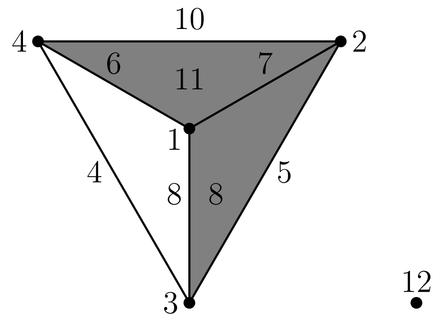

We can observe that the c-vector of K is $\{5, 6, 2\}$, $b_0 = 2$, $b_1 = 1$, $b_2 = 0$, $b_3 = 0$, $m_0 = 4$, $m_1 = 4$, $m_2 = 1$ and $m_3 = 0$. Then, we obtain

$$m_0 = 4 \geq 2 = b_0$$
$$m_1 = 4 \geq 1 = b_1$$
$$m_2 = 1 \geq 0 = b_2$$

and

$$m_0 - m_1 + m_2 = 4 - 4 + 1 = 1 = 2 - 1 + 0 = b_0 - b_1 + b_2$$

for the Weak Discrete Morse Inequalities and

$$m_0 = 4 \geq 2 = b_0$$
$$m_1 - m_2 = 4 - 4 = 0 \geq -1 = 1 - 2 = b_1 - b_0$$
$$m_2 - m_1 + m_0 = 1 - 4 + 4 = 1 \geq 1 = 0 - 1 + 2 = b_2 - b_1 + b_0$$
$$m_3 - m_2 + m_1 - m_0 = 0 - 4 + 4 - 1 = -1 \geq -1 = 0 - 0 + 1 - 2 = b_3 - b_2 + b_1 - b_0$$

for the Strong Discrete Morse Inequalities.

9. Morse Sequences

Morse sequences are a way of visualizing discrete Morse functions introduced by [3].

Definition 9.1. Let K and L be simplicial complexes. If $\sigma \in K$ is a facet of K and if $L = K \setminus \{\sigma\}$, we call L an *elementary perforation* of K, and K an *elementary filling* of L.

Definition 9.2. Let K be a simplicial complex. A *Morse sequence* on K is a sequence $\vec{W} = (\emptyset = K_0, K_1, K_2, \ldots, K_n = K)$ of simplicial complexes such that, for every $i = 1, 2, 3, \ldots, n$, K_i is either an elementary expansion or an elementary filling of K_{i-1}. We also write $\vec{W}(K)$ for a Morse sequence $\vec{W}$ on K.

Below are two potential algorithms for building a Morse sequence for a simplicial complex K.

Algorithm 3 Maximal Increasing Morse Sequence Algorithm

INPUT: A simplicial complex K.

(0) Initialize $\vec{W} = (\emptyset)$.
(1) Consider the last simplicial complex K_n in $\vec{W}$. Is it identical to K? If yes, then STOP; otherwise, go to (2).
(2) Can we perform elementary expansions on K_n that create structures identical to those in K? If yes, go to (3); if no, go to (4).
(3) *Elementary Expansion*: Perform one such elementary expansion to get a new simplicial complex K_{n+1}, and add it to $\vec{W}$. Go back to (1).
(4) *Elementary Filling*: Use an elementary filling to add a simplex in K to K_n to get a new simplicial complex K_{n+1}, and add it to $\vec{W}$. Go back to (1).

OUTPUT: The Morse sequence $\vec{W}$.

Algorithm 4 Maximal Decreasing Morse Sequence Algorithm

INPUT: A simplicial complex K.

(0) Initialize $\vec{W} = (K)$.
(1) Consider the first simplicial complex K_n in $\vec{W}$. Is it identical to $\emptyset$? If yes, then STOP; otherwise, go to (2).
(2) Can we perform elementary collapses on K_n? If yes, go to (3); if no, go to (4).
(3) *Elementary Collapse*: Perform one such elementary collapse to get a new simplicial complex K_{n-1}, and add it to $\vec{W}$. Go back to (1).
(4) *Elementary Perforation*: Use an elementary perforation to remove a facet from K_n to get a new simplicial complex K_{n-1}, and add it to $\vec{W}$. Go back to (1).

OUTPUT: The Morse sequence $\vec{W}$.

Definition 9.3. Let $\vec{W} = (K_0 m =, K_1, K_2, \dots, K_n)$ be a Morse sequence. For each $i = 1, 2, 3, \dots, n$, let κ_i be defined as such:

(a) If K_i is an elementary filling of K_{i-1} and $K_i = K_{i-1} \cup \{\sigma\}$, then $\kappa_i = \sigma$. We say that κ_i is *critical* for $\vec{W}$.
(b) If K_i is an elementary expansion of K_{i-1} and $K_i = K_{i-1} \cup \{\sigma^{(p-1)}, \tau^{(p)}\}$, then $\kappa_i = \{\sigma, \tau\}$. We say that κ_i is *regular* for $\vec{W}$, σ is a *lower regular* for $\vec{W}$, and τ is a *higher regular* for $\vec{W}$.

We will use $C(\vec{W})$ to denote the set of critical simplices of $\vec{W}$ in K, $R(\vec{W})$ to denote the set of regular pairs, $U(\vec{W})$ to denote the set of upper regular simplices, and $L(\vec{W})$ to denote the set of lower regular simplices.

An important concept in Morse sequences are reference and coreference maps, which assign a set of critical p-simplices to each p-simplex of a simplicial complex; this provides an efficient method to label each simplex of a Morse sequence with a set of critical simplices. Specifically, critical simplices map to themselves, upper regular simplices map to the mappings of the coboundaries of their respective lower regular simplices, and lower regular simplices map to the mappings of the boundaries of their respective upper regular simplices. Just like with discrete Morse functions, Morse sequences can be used to analyze the topological structure of a simplicial complex; however, Morse sequences provide a more connected and dynamic, albeit complicated representation of simplicial complexes.

10. Acknowledgments

This paper was written under the supervision of the Euler Circle Independent Paper and Research Writing program. I would like to thank Simon Rubinstein-Salzedo for providing helpful insight into discrete Morse theory and mathematical thinking as a whole. I would also like to thank my mentor, Ethan Martirosyan, for his advice on the process of paper-writing and his feedback on the initial stages of my paper.

References

[1] Bruno Benedetti and Frank H. Lutz. Random discrete Morse theory and a new library of triangulations. *Exp. Math.*, 23(1):66–94, 2014.

[2] Bruno Benedetti and Frank Hagen Lutz. The dunce hat in a minimal non-extendably collapsible 3-ball. Preprint, arXiv:0912.3723 [math.AT] (2009), 2009.

[3] Gilles Bertrand. Morse sequences: a simple approach to discrete Morse theory. *J. Math. Imaging Vis.*, 67(2):22, 2025. Id/No 16.

[4] Anton Dochtermann and Andrew Newman. Random subcomplexes and Betti numbers of random edge ideals. *Int. Math. Res. Not.*, 2023(10):8832–8871, 2023.

[5] Robin Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

[6] Ragnar Freij. Equivariant discrete Morse theory. *Discrete Math.*, 309(12):3821–3829, 2009.

[7] A. T. Lundell and S. Weingram. The topology of CW complexes. The University Series in Higher Mathematics. New York etc.: Van Nostrand Reinhold Company. VIII, 216 p. (1969)., 1969.

[8] John W. Milnor. *Morse theory. Based on lecture notes by M. Spivak and R. Wells*, volume 66 of *Texts Read. Math.* New Delhi: Hindustan Book Agency, reprint of the 1963 original published by Princeton University Press edition, 2013.

[9] J. Rotman. *An Introduction to Algebraic Topology.* Graduate Texts in Mathematics. Springer New York, 1998.

[10] Nicholas A. Scoville. *Discrete Morse theory*, volume 90 of *Stud. Math. Libr.* Providence, RI: American Mathematical Society (AMS), 2019.

[11] Anna Song, Ka Man Yim, and Anthea Monod. Generalized Morse theory of distance functions to surfaces for persistent homology. *Adv. Appl. Math.*, 166:79, 2025. Id/No 102857.