

Stratified Morse Theory for Cell Complexes

Vidit Nanda and Francesca Tombari

ABSTRACT. We develop a version of discrete Morse theory for finite regular CW complexes equipped with an auxiliary stratification. The key construction is the *halo* of a cell, which contains all those faces in the boundary that enter closed sublevelsets precisely when the threshold reaches that cell’s value. The complement of this halo in the boundary, called the *shadow*, is always a subcomplex. A stratified discrete Morse function requires Forman’s conditions on each stratum together with the requirement that closures of paired cells admit filtered collapses onto their shadows. We establish fundamental Morse lemmas: filtered collapses across regular intervals, and controlled attachments at critical values. For functions satisfying only the stratum-wise Forman condition, we construct an upper envelope on the barycentric subdivision whose local Morse data decomposes into horizontal and vertical components. This yields a simplicial analogue of the standard tangential-normal splitting of local Morse data in the sense of Goresky and MacPherson.

Introduction

Morse theory provides one of the finest and most successful mechanisms for translating local analytic information into global geometric insight. Beginning with Morse’s original work on counting geodesics [21], the core insight – that the topology of sublevelsets changes only at critical points, and does so in a controlled manner – has also found spectacular applications when extended to equivariant [2], piecewise-linear [4], symplectic [11], dynamical [7], and physical [31] contexts. In each case, the fundamental theorems survive: regular intervals produce no topological change, while critical values yield handle attachments whose nature can be determined by local data of the underlying function near critical points. We are concerned here with two such extensions — the first of these is the combinatorial adaptation developed by Forman for finite CW complexes [12], and the second is Goresky-MacPherson’s Morse theory for stratified spaces [14].

In Forman’s theory, a discrete Morse function f assigns real values to cells of a CW complex X such that each cell σ has at most one exceptional coface $\tau > \sigma$ with $f(\tau) \leq f(\sigma)$, or at most one exceptional face $\eta < \sigma$ with $f(\eta) \geq f(\sigma)$. Cells with no exceptional neighbours are deemed critical, while the remaining cells assemble into free-face pairs that can be collapsed away. This discrete avatar of Morse theory has also found substantial applications across diverse fields, including topological combinatorics [3], commutative algebra [15], algebraic topology [28, 10], and geometric group theory [1, 32]. It has also been modified in several directions — there are filtered [20], equivariant [13, 32], sheafy [9], noncompact [17] and even 2-categorical refinements [26, 24]. Our goal here is to develop a *stratified* discrete Morse theory that combines insights from Goresky and MacPherson’s pioneering work [14] while retaining the combinatorial essence and flexibility of Forman’s.

Therefore, the setting throughout is that of a finite regular CW complex X equipped with a filtration

$$X = X_n \supset X_{n-1} \supset \cdots \supset X_0 \supset X_{-1} = \emptyset$$

such that the successive differences $X_i - X_{i-1}$, whose connected components are called *strata*, satisfy the frontier axiom (see Section 1.1). This induces a partial order on the strata. Besides this order, we emphasise at the outset that we impose no manifold structure, no regularity conditions à la Whitney or Thom–Mather, and no requirement that strata admit conical neighbourhoods. We forsake these hypotheses deliberately, for two reasons: first, we do not need them. And second, imposing them would prevent our theory from recovering Forman’s results for the trivial stratification $X \supset \emptyset$.

The Challenge. Our primary target in this work is the pair of **fundamental Morse lemmas**. The first of these asserts that nothing of topological importance changes across regular intervals — sublevelsets are homotopy equivalent, and indeed, collapse onto one another in a filtration-preserving manner. The second one quantifies the change across a critical value: the sublevelset grows by a controlled handle attachment. Before we can establish (or even formulate) discrete analogues of these results, we must confront two obstacles.

The first of these is conceptual: *how should we even define a discrete stratified Morse function $f : X \rightarrow \mathbb{R}$?* One might naïvely require that the restriction $f|_S : S \rightarrow \mathbb{R}$ to every stratum $S \subset X$ satisfies the axioms of Forman. Such a strategy is antithetical to both the spirit and the letter of Goresky–MacPherson’s setting, where a crucial additional *normal nondegeneracy* requirement is imposed. The gradient of a stratified Morse function f at a critical point must not annihilate any limiting tangent space of a higher stratum at that point. This nondegeneracy condition ensures that gradient flow does not become trapped at stratum boundaries. *What is the combinatorial equivalent of this condition?* We have no tangent planes, no gradients in the differential-geometric sense, and no obvious way to define or detect transversality.

The second obstacle is technical, and we call it the *closure problem*. In Goresky and MacPherson’s setting, the sublevelset $f_{\leq c}$ is unambiguous — it is simply the preimage of the half-infinite interval $(-\infty, c]$ under a continuous function. In sharp contrast, sublevelsets of the form $f_{\leq c}$ need not form subcomplexes of the ambient CW complex X precisely because we allow cells $\sigma < \tau$ to satisfy $f(\sigma) \geq f(\tau)$. Following Forman, we are therefore compelled to work with closures: the relevant object is $\mathbf{cl}(f_{\leq c})$, the smallest subcomplex containing all cells with f -value at most c . The problem is that passing to the closure can introduce cells with arbitrarily large f -values — a cell σ with $f(\sigma) = 10000$ will appear in $\mathbf{cl}(f_{\leq 1})$ if σ has a coface $\tau > \sigma$ with $f(\tau) \leq 1$ (see Figure 1). There is no parallel to this phenomenon in the continuous setting, one might as well compare a volcano to a butterfly.

This Paper. Here we resolve both obstacles by introducing the *halo* $\mathbf{h}(\sigma; f)$ of a given cell $\sigma \in X$ along an injective map $f : X \hookrightarrow \mathbb{R}$. This is defined as the collection of all boundary cells $\eta < \sigma$ such that $f(\eta) \geq f(\sigma)$ and σ is the f -minimal coface of η . Informally, the halo consists of those faces that appear in the closed sublevelset $\mathbf{cl}(f_{\leq c})$ precisely when the threshold c reaches $f(\sigma)$. The complement of the halo in the boundary of σ , which we call the *shadow* of σ and denote $\mathbf{sh}(\sigma; f)$, forms a subcomplex of the closure $\mathbf{cl}(\sigma)$. The halo allows us to formulate our combinatorial nondegeneracy condition as follows. In addition to requiring our putative Morse function to be injective and satisfy the Forman axioms on each stratum, we further ask that for each cell pair $(\sigma > \tau)$ lying in the same stratum with

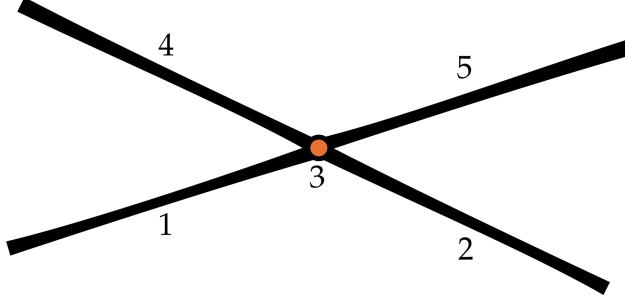


FIGURE 1. A piece of a cell complex containing one vertex and four edges, with f -values indicated on the cells. Note that the closure of $f_{\leq 1}$ already contains the central vertex, even though that vertex has f -value 3.

$f(\sigma) < f(\tau)$, the smaller cell τ lies in $\mathbf{h}(\sigma; f)$ and the closure $\mathbf{cl}(\sigma)$ admits a filtered collapse onto the shadow $\mathbf{sh}(\sigma; f)$.

This new condition ensures that when a cell σ is cancelled by its partner τ within a stratum, all the collateral faces in the halo can also be collapsed away without disrupting the stratification. The filtered collapse respects the filtration X_{\bullet} in the sense that it restricts to a homotopy equivalence at each level. As a result, critical cells now come in two flavours: there is the standard notion, where a cell σ lying in a stratum S is *critical* if it is unpaired within S . It is *s-critical* if it is critical and does not lie in the halo of any other cell. This distinction is genuinely new: a cell can be critical and contribute meaningfully to the topology of its stratum, and yet be invisible to the filtered topology of X because it admits cofaces with smaller f -values lying in higher strata. Only s-critical cells affect the filtered homotopy type of sublevelsets.

Main Results. Fix a stratified discrete Morse function $f : X \hookrightarrow \mathbb{R}$ and let σ be a cell of X . Our first main result establishes that there is no material change in the sublevelsets of f across intervals which contain no s-critical values. Here is an abridged version, see Theorem 4.1 for details.

THEOREM (A). *Let $[c, d] \subset \mathbb{R}$ be an interval whose preimage $f^{-1}([c, d])$ is $\{\sigma\}$. If σ is not s-critical, then $\mathbf{cl}(f_{\leq d})$ admits a filtration-preserving collapse onto $\mathbf{cl}(f_{\leq c})$.*

Let us emphasise once again that the collapse promised by this result might occur even if σ is critical, for the reasons detailed above — if σ lies in the halo of another cell, then it will already be present in $\mathbf{cl}(f_{\leq c})$ and make no material contribution as we increase the threshold from c to d . Having confirmed that sublevelsets can only change when we cross an s-critical value, it remains to produce **local Morse data** by describing the precise handle attachment across such a value. This task becomes the purview of our next main result (see Theorem 4.2 below).

THEOREM (B). *Let $\sigma \in X$ be an s-critical cell of f with $c := f(\sigma)$. For all sufficiently small $\varepsilon > 0$, the sublevelset $\mathbf{cl}(f_{\leq c+\varepsilon})$ is the union of $\mathbf{cl}(\sigma)$ with $\mathbf{cl}(f_{\leq c-\varepsilon})$ along their common intersection $\mathbf{sh}(\sigma; f)$.*

In other words, the local Morse data at an s-critical cell σ is the pair $\mathbf{sh}(\sigma; f) \hookrightarrow \mathbf{cl}(\sigma)$. We should specify that this data is manifestly different from Goresky and MacPherson's tangential \times normal decomposition from [14]; their tangential data T occurs along the stratum S containing the critical point p and is identical to the usual handle attachment data for the

restriction $f|_S$. Their normal data N is generated by the *descending link* of S , which resides in the strata lying strictly above S . This disparity is both a feature – our attachment data is combinatorially explicit and requires no link computations – and a limitation, in that we obtain less disentangled geometric information than the smooth theory provides. Our third and final main result is therefore a peace offering to the reader who would much prefer a discrete theory that faithfully recreates the continuous one's $T \times N$ product decomposition of local Morse data.

In fact, here we consider a weaker setup. We still require $f : X \rightarrow \mathbb{R}$ to be injective and restrict to a discrete Morse function on each stratum, but no further halo-centric constraints are imposed. We then pass to the barycentric subdivision of X ; this is a simplicial complex $Sd(X)$ whose k -simplices are strictly ascending chains $\xi = [\sigma_0 < \dots < \sigma_k]$ of cells in X . The subdivision canonically inherits its stratification from X , where each stratum $S \subset X$ induces a unique stratum $Sd(S) \subset Sd(X)$. The map f extends to $Sd(X)$ via its *upper envelope* $\mathbf{U}f : Sd(X) \rightarrow \mathbb{R}$, given by

$$\mathbf{U}f(\xi) = \max \{f(\sigma_i) \mid 0 \leq i \leq k\}.$$

The upshot is that sublevelsets of $\mathbf{U}f$ are honest subcomplexes and there is no need to take closures. We establish that – for the sublevelsets of $\mathbf{U}f$ – the local Morse data at a vertex $[\sigma]$ in $Sd(X)$ corresponding to each cell $\sigma \in X$ admits a join decomposition of the form $H_\sigma \star V_\sigma$, where $H = H_\sigma$ and $V = V_\sigma$ are subcomplexes called the *horizontal* and *vertical* parts of the data. The terminology is justified by the following observations: if σ is critical and S is the stratum containing it, then H lies in the union of strata $\leq Sd(S)$ while V lies in the union of strata $> Sd(S)$.

THEOREM (C). *Let σ be a critical cell with $c := f(\sigma)$. For all sufficiently small $\varepsilon > 0$, the local Morse data for $\mathbf{U}f_{\leq c-\varepsilon} \hookrightarrow \mathbf{U}f_{\leq c+\varepsilon}$ is a pair*

$$(H \star V) \hookrightarrow (\text{cone}(H) \star V),$$

where H lies in the union of strata $\leq Sd(S)$ and V lies in the union of strata $> Sd(S)$.

(This is Theorem 5.9 below).

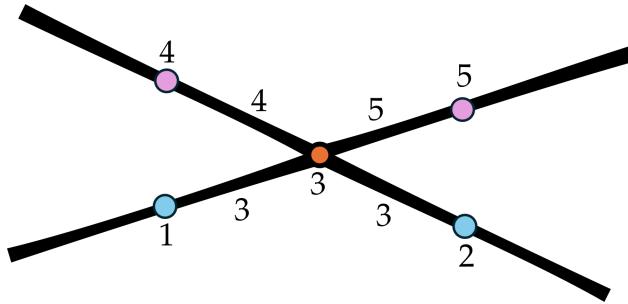


FIGURE 2. A subdivided version of the Example from Figure 1, this time with $\mathbf{U}f$ -values on simplices; here the central vertex does contribute to the change in topology because its lower link consists of two vertices. These lie at the barycenters of the original edges, and hence have $\mathbf{U}f$ -values 1 and 2.

This decomposition is our simplicial analogue of the tangential-normal splitting. The horizontal part H captures Morse data along the stratum, while the vertical part V records the directions in which higher strata impinge on σ . In Figure 2 we consider the barycentric

subdivision and upper envelope of Figure 1. The central vertex, which had no contribution in the original model of taking sublevelsets, now admits a lower link of the form $H \star V$, where H is empty but V has two connected components (these are barycenters of edges with f -values 1 and 2). Since this link is not contractible, crossing $Uf = 3$ is guaranteed to change sublevelset topology.

Related Work. This paper may be viewed as continuing three largely independent research programmes, all of which seek to generalise Forman’s original theory from [12]. The first of these is the filtered variant due to Mischaikow and Nanda [20]. The main difference between that setting and ours is that in our work, we allow for the existence of cells $\sigma < \tau$ lying in different strata with $f(\sigma) > f(\tau)$ — this is expressly prohibited in [20, Def 4.1].

The second overlapping research direction is the more recent output of Knudson and Wang [16], which also aims to build a stratified discrete Morse theory. Their functions are only required to restrict to Forman-style discrete Morse functions on each stratum, with no further requirements. As a consequence, they obtain weak Morse inequalities relating critical cell counts to Betti numbers, but are unable to extract local Morse data. Our Theorem (C) shows that even in their setting, one recovers a tangential-normal decomposition in the barycentric subdivision as long as one has injectivity.

Thirdly, we note the substantial literature on multivector fields, initiated by Mrozek and collaborators [22, 23, 18], which provides a framework for discrete Conley index theory. A multivector is a nonempty, convex, connected subset of the face poset; a multivector field is a partition of a regular CW complex X into multivectors. Stratifications satisfying the frontier axiom yield acyclic multivector fields, and the Conley–Morse spectral sequence assembles the homology of X from the Conley indices of its strata. We discuss this connection briefly in an appendix, but emphasise that multivector field theory imposes no analogue of our halo collapsibility condition, and therefore does not yield Morse lemmas of the kind we establish here.

Outline. In Section 1 we describe stratifications of regular CW complexes. Then Section 2 introduces halos and shadows, proving their basic properties. In Section 3 we define stratified discrete Morse functions and the two notions of criticality. Section 4 establishes Theorems (A) and (B). Finally, Section 5 constructs the upper envelope on the barycentric subdivision and proves Theorem C. Appendix A tersely summarises the relationship between our stratifications and acyclic multivector fields.

1. Stratified Cell Complexes

Let X be a regular CW complex. We denote by (X, \leq) the **face poset** of X , which contains all the cells of X ordered by the face relation¹. Given any nonempty subcollection S of cells lying in X , the *closure* of S is the subcomplex defined as a down-set

$$\mathbf{cl}(S) := \{\tau \in X \mid \tau \leq \sigma \text{ for some } \sigma \in S\};$$

and dually, the *(open) star* of S is the up-set

$$\mathbf{st}(S) := \{\tau \in X \mid \tau \geq \sigma \text{ for some } \sigma \in S\}.$$

¹Explicitly, the relation $\sigma < \tau$ holds if and only if the boundary of τ contains σ .

When $S = \{\sigma\}$ consists of a single cell, we write $\mathbf{cl}(\sigma)$ and $\mathbf{st}(\sigma)$ rather than $\mathbf{cl}(\{\sigma\})$ and $\mathbf{st}(\{\sigma\})$. We say that S is *connected* if any two cells σ, τ in S form endpoints of a finite zigzag:

$$\sigma \geq \eta_0 \leq \eta_1 \geq \cdots \geq \eta_{k-1} \leq \eta_k \geq \tau,$$

with all $\eta_i \in S$. Finally, we denote by $\dim S$ the maximal dimension encountered among the constituent cells of S .

1.1. Stratifications. A *filtration* of a finite regular CW complex X is a finite sequence of subcomplexes

$$X = X_n \supset X_{n-1} \supset \cdots \supset X_1 \supset X_0 \supset X_{-1} = \emptyset.$$

A **stratification** of X is a filtration such that the successive differences $\Delta_i := (X_i - X_{i-1})$ satisfy the following *frontier axiom* — given connected $S \subset \Delta_i$ and $T \subset \Delta_j$, if $\mathbf{cl}(S) \cap T$ is nonempty then $T \subset \mathbf{cl}(S)$. The connected components of Δ_i are called *i-strata* of the stratification X_\bullet . A standard example is the *skeletal* stratification, where the *i-strata* are precisely the *i*-dimensional cells; this is the finest possible stratification of X . At the other end of the spectrum, one has the coarsest (or canonical, or minimal) stratification of X into homology manifolds — see [25].

Remark 1.1. As we have mentioned already in the Introduction, it is customary when defining stratified spaces to impose additional regularity conditions on the strata. For instance: one often requires *i-strata* to be (smooth, piecewise-linear, or homology) manifolds of dimension *i* [30, 19, 29, 25]. **We make no such assumptions here.** Similarly, it is customary to assume that strata admit conical neighbourhoods in X . Namely, for each stratum S there exists a stratified space LS , called the **link** of S , such that any sufficiently small open neighbourhood $US \subset X$ of S admits a fiber bundle structure $US \twoheadrightarrow S$ with fiber $\text{Cone}(LS)$. Our stratifications are less rigid in the sense that the strata are not required to admit well-defined links or homogeneous neighbourhoods.

Nevertheless, we impose the frontier axiom in order to avail of two direct benefits. The first of these is the existence of a well-beloved and oft-used partial order on the strata.

Proposition 1.2. *In any stratification X_\bullet of X , the relation*

$$T \leq S \text{ if and only if } T \cap \mathbf{cl}(S) \neq \emptyset$$

is a partial order on the set of strata.

PROOF. Reflexivity of the proposed relation is immediate. If $T \leq S$ and $S \leq R$, by definition of the relation and the frontier axiom, we have $T \subset \mathbf{cl}(S)$ and $S \subset \mathbf{cl}(R)$. This implies $T \subset \mathbf{cl}(S) \subset \mathbf{cl}(R)$, and hence $T \leq R$, showing transitivity. It remains to establish antisymmetry. Proceeding by contradiction, assume there exist strata $S \neq T$ such that both $S \cap \mathbf{cl}(T)$ and $T \cap \mathbf{cl}(S)$ are nonempty. Let us assume that $S \subset \Delta_i$ and $T \subset \Delta_j$ where $\Delta_i = (X_i - X_{i-1})$ and similarly for Δ_j . We now have containments

$$T \subset \mathbf{cl}(S) \subset X_i,$$

where the first inclusion follows from the frontier axiom while the second follows because X_i is a subcomplex, and hence closed in X . Thus, T lies in $\Delta_j \cap X_i$, which forces $i \geq j$. The same argument gives $j \geq i$ if we exchange the roles of S and T , whence $i = j$. Thus, both S and T must be connected components of the same Δ_i . Finally, since $T \subset \mathbf{cl}(S)$, there exists cells $\tau \in T$ and $\sigma \in S$ with $\tau \leq \sigma$. This forces S and T to lie in the same connected component of Δ_i , and hence $S = T$ as desired. \square

For the skeletal stratification, the frontier partial order evidently reduces to the familiar face relation between cells. The second benefit of the frontier axiom is that it guarantees *convexity* of all strata, as described below.

Proposition 1.3. *Given any triple of cells $\sigma \leq \tau \leq \sigma'$ in a regular CW complex X , if both σ and σ' lie in the same stratum S of a stratification X_\bullet , then so does τ .*

PROOF. The only nontrivial case occurs when both inequalities among cells are strict, otherwise $\tau \in S$ is automatic. By construction, the strata of X_\bullet partition the cells of X ; we may therefore assume for the sake of contradiction that τ lies in a different stratum $T \neq S$. Now T intersects $\text{cl}(S)$ because $\tau < \sigma'$ whereas S intersects $\text{cl}(T)$ because $\sigma < \tau$. By the frontier axiom, we therefore get both $S \leq T$ and $T \leq S$, which immediately forces the desired contradiction $S = T$. \square

1.2. Filtered Collapses. We say that distinct cells (σ, τ) constitute a *free-face pair* in a regular CW complex X whenever

$$\mathbf{st}(\tau) = \{\sigma, \tau\}.$$

In this case, the regularity of X forces $\dim \tau = \dim \sigma - 1$ and it is well-known (see eg [6]) that removing both σ and τ from X produces a subcomplex which remains homotopy equivalent to X . This free-face pair removal operation is called an **elementary collapse**; we say that X *collapses* to a subcomplex Y if there is a finite sequence of intermediate subcomplexes Z^i

$$X = Z^1 \supset Z^2 \supset \cdots \supset Z^k = Y$$

where each Z^i has been obtained from the preceding Z^{i-1} by performing a single elementary collapse. We denote this relationship between X and Y by $X \searrow Y$.

Remark 1.4. The reader unfamiliar with this machinery is warned that collapsing X onto Y is not simply a matter of finding a collection of disjoint free-face pairs in X . When a free-face pair (σ, τ) is removed from X to produce the first subcomplex Z^1 , the excision of σ might create a new free-face pair (σ', τ') in Z^1 for some $\tau' < \tau$. Crucially, this new (σ', τ') is *not* a free-face pair in X because $\mathbf{st}(\tau') \supset \{\sigma', \sigma, \tau'\}$ has more than two elements.

Let us assume now that X is equipped with a stratification X_\bullet ; every subcomplex $Y \subset X$ inherits this stratification via $Y_i := Y \cap X_i$ for all i . We now seek a more refined collapsing mechanism from X to Y which produces a *filtered* homotopy equivalence. Thus, the goal is to not only obtain a homotopy equivalence $\phi_i : X_i \xrightarrow{\sim} Y_i$ at each level i , but also to require compatibility with the natural inclusion maps which relate adjacent levels. Explicitly, the diagram

$$\begin{array}{ccc} X_i & \hookrightarrow & X_{i+1} \\ \phi_i \downarrow \sim & & \sim \downarrow \phi_{i+1} \\ Y_i & \hookrightarrow & Y_{i+1} \end{array}$$

should commute for all i . This task is accomplished by allowing the stratification data to constrain the class of acceptable elementary collapses — we only permit the removal of a free-face pair (σ, τ) when both σ and τ lie in the same stratum S . We will denote the existence of such a **filtered collapse** as $X \searrow Y$ in order to distinguish it from the ordinary unfiltered collapses described above.

Remark 1.5. Assume that the regular CW subcomplex $Y \subset X$ has been obtained via a filtered elementary collapse. Denote by (σ, τ) the excised pair, let $S \subset (X_j - X_{j-1})$ be the stratum containing both cells. For each dimension i , we seek to relate the induced stratification $Y_i := Y \cap X_i$ of Y to the original stratification of X in order to make clear why a filtered homotopy equivalence $\phi_i : X_i \rightarrow Y_i$ exists. For this purpose, note first that $Y_i = X_i$ holds whenever $i < j$ simply because neither σ nor τ are present in any such X_j , so we let ϕ_i be the identity for such i . Next, we note that (σ, τ) is a free-face pair in X_i for all $i \geq j$ — we have $\mathbf{st}(\tau) = \{\sigma, \tau\}$ in X , and since both σ and τ lie in X_j this pair remains free in X_i for all $i \geq j$. It follows that one can choose the global homotopy equivalence $\phi : X \rightarrow Y$ to restrict to homotopy equivalences $X_i \rightarrow Y_i$ as expected, and hence the natural square commutes.

2. Halos and Shadows

Let X be a finite regular CW-complex. Throughout this section, we consider a function $f : X \rightarrow \mathbb{R}$ that sends each cell σ of X to a real number $f(\sigma)$. The *lower star* of σ along f is

$$\mathbf{st}^-(\sigma; f) := \{\tau \in X \mid \tau > \sigma \text{ and } f(\tau) \leq f(\sigma)\}; \quad (1)$$

this is evidently a subset of $\mathbf{st}(\sigma)$, but — unlike the usual star — it never contains σ itself. We will be particularly interested here in the case where f is injective. For each collection of cells $S \subset X$, let $\text{argmin}[f/S]$ denote the subset of all cells $\tau \in S$ which attain the minimal f -value in the image $f(S) \subset \mathbb{R}$, with the explicit understanding that $\text{argmin}[f/\emptyset] = \emptyset$. For injective f and nonempty S , the set $\text{argmin}[f/S]$ is guaranteed to be a singleton.

Definition 2.1. The **halo** of a cell $\sigma \in X$ along f is defined as

$$\mathbf{h}(\sigma; f) := \{\tau \in X \mid \sigma \in \text{argmin}[f/\mathbf{st}^-(\tau; f)]\},$$

while the *augmented halo* $\widehat{\mathbf{h}}(\sigma; f)$ is the (necessarily disjoint) union of $\mathbf{h}(\sigma; f)$ with $\{\sigma\}$.

Note that if τ lies in $\mathbf{h}(\sigma; f)$ then we must have $\sigma \in \mathbf{st}^-(\tau; f)$. Thus, both $\mathbf{h}(\sigma; f)$ and $\widehat{\mathbf{h}}(\sigma; f)$ are always subsets (but generally not subcomplexes) of the closure $\mathbf{cl}(\sigma)$. For injective f , a cell τ lies in $\mathbf{h}(\sigma; f)$ if and only if σ is the unique cell in $\mathbf{st}^-(\tau; f)$ with minimal f -value. This forces the halos of distinct cells along injective f to be disjoint.

Proposition 2.2. Let $f : X \hookrightarrow \mathbb{R}$ be an injective function. If $\sigma \neq \sigma'$ are two distinct cells of X , then we have $\mathbf{h}(\sigma; f) \cap \mathbf{h}(\sigma'; f) = \emptyset$.

PROOF. If a cell τ is contained in the intersection $\mathbf{h}(\sigma; f) \cap \mathbf{h}(\sigma'; f)$, then by definition both σ and σ' must lie in $\mathbf{st}^-(\tau; f)$ and attain the minimal value of f over this lower star. This violates the injectivity of f . \square

The halo of a cell σ along an injective f is also well-behaved as a subposet of $\mathbf{cl}(\sigma)$.

Proposition 2.3. Assume $f : X \hookrightarrow \mathbb{R}$ is injective. For each cell σ of X , the augmented halo $\widehat{\mathbf{h}}(\sigma; f)$ is an up-set of $\mathbf{cl}(\sigma)$.

PROOF. We seek to show that given two cells $\beta \geq \alpha$ in $\mathbf{cl}(\sigma)$ with α contained in $\widehat{\mathbf{h}}(\sigma; f)$, we also have $\beta \in \widehat{\mathbf{h}}(\sigma; f)$. If any of the inequalities in $\sigma \geq \beta \geq \alpha$ is an equality, then the result is immediate. Therefore, we may safely assume $\sigma > \beta > \alpha$. We now claim that $f(\beta) > f(\sigma)$; to see why, note by injectivity of f that we either have

- $f(\beta) > f(\alpha)$, in which case $f(\beta) > f(\sigma)$ holds because $\sigma \in \mathbf{st}^-(\alpha; f)$; or,
- $f(\beta) < f(\alpha)$, whence $\beta \in \mathbf{st}^-(\alpha; f)$ and so $f(\beta) > f(\sigma)$ by definition of argmin .

Thus, in both cases we obtain $\sigma \in \mathbf{st}^-(\beta; f)$. Assume, for the sake of contradiction, that $\beta \notin \mathbf{h}(\sigma; f)$; then there must exist another cell $\sigma' \neq \sigma$ in $\mathbf{st}^-(\beta; f)$ satisfying $f(\sigma') < f(\sigma)$. But any such σ' would automatically lie in the lower star of α and contradict the fact that σ has the smallest f -value amongst lower-star cofaces of α . This gives the desired containment $\beta \in \mathbf{h}(\sigma; f)$ and establishes upward closure of the augmented halo. \square

Since the removal of any up-set of cells from a regular CW complex leaves behind a subcomplex, the preceding result guarantees that

$$\mathbf{sh}(\sigma; f) := \mathbf{cl}(\sigma) - \widehat{\mathbf{h}}(\sigma; f), \quad (2)$$

is a (possibly empty) subcomplex of $\mathbf{cl}(\sigma)$, and hence, of X whenever f is injective. We will call this subcomplex the **shadow** of σ along f . Define, for each real number c , the sub- and super-levelset of f at c as

$$\begin{aligned} f_{\leq c} &:= \{\sigma \in X \mid f(\sigma) \leq c\}, \\ f_{\geq c} &:= \{\sigma \in X \mid f(\sigma) \geq c\}. \end{aligned}$$

Our interest in the halo stems from the fact that it completely determines the differences between closures of certain sublevelsets.

Lemma 2.4. *Let $f : X \hookrightarrow \mathbb{R}$ be an injective function, and let $[c, d] \subset \mathbb{R}$ be an interval such that $f_{\geq c} \cap f_{\leq d}$ consists of a single cell $\sigma \in X$. For any such interval, the difference*

$$\Delta := \mathbf{cl}(f_{\leq d}) - \mathbf{cl}(f_{\leq c})$$

of sublevelsets has the following structure:

- (1) if $\sigma \in \mathbf{cl}(f_{\leq c})$, then $\Delta = \emptyset$;
- (2) if $\sigma \notin \mathbf{cl}(f_{\leq c})$, then $\Delta = \widehat{\mathbf{h}}(\sigma; f)$.

PROOF. A cell γ lies outside $\mathbf{cl}(f_{\leq c})$ if and only if $f(\tau) > c$ holds for every coface $\tau \geq \gamma$. Conversely, γ lies inside $\mathbf{cl}(f_{\leq d})$ if and only if there exists some coface $\tau_* \geq \gamma$ with $f(\tau_*) \leq d$. Thus, in order for γ to lie in Δ , two properties must hold simultaneously. First, there must exist some $\tau_* \geq \gamma$ with $f(\tau_*) \in (c, d]$. And second, for every other $\tau \neq \tau_*$ in $\mathbf{st}(\gamma)$ we must have $f(\tau) > c$. The hypothesis requiring $f_{\geq c} \cap f_{\leq d}$ to equal $\{\sigma\}$ affects both of these properties and yields the desired conclusions. Explicitly,

- (1) The first property simplifies to the requirement that $\sigma = \tau_*$ lies in $\mathbf{st}(\gamma)$ and satisfies $f(\sigma) \neq c$. Now if $\sigma \in \mathbf{cl}(f_{\leq c})$ then all of its faces also lie in $\mathbf{cl}(f_{\leq c})$ by the subcomplex property. But since Δ can only contain faces of σ , it must be empty in this case.
- (2) The second property strengthens to $f(\tau) > d$ for all $\tau \neq \sigma$ in $\mathbf{st}(\gamma)$. Thus, a cell γ lies in Δ if and only if $\{\sigma\} = \operatorname{argmin}[f / \mathbf{st}(\gamma)]$. There are now two cases to consider. Either $\gamma = \sigma$, or $\gamma < \sigma$ and $f(\gamma) > d \geq f(\sigma)$. In the latter case, note that σ lies in $\mathbf{st}^-(\gamma; f)$ and is the argmin of f over this lower star.

To conclude the proof, we note from the strengthened version of the second property above that $\gamma \in \Delta$ holds if and only if γ lies in $\{\sigma\} \sqcup \mathbf{h}(\sigma; f)$, as desired. \square

3. Stratified Discrete Morse Functions

Fix a finite regular CW complex X and a function $f : X \rightarrow \mathbb{R}$. The *upper closure* of each cell σ along f is defined as:

$$\mathbf{cl}^+(\sigma; f) := \{\tau \in X \mid \tau < \sigma \text{ and } f(\tau) \geq f(\sigma)\}. \quad (3)$$

Upper closures are dual to lower stars from (1) in the sense that $\tau \in \mathbf{cl}^+(\sigma; f)$ holds if and only if $\sigma \in \mathbf{st}^-(\tau; f)$ holds. We call f a *discrete Morse function* on X in the sense of Forman whenever the inequality

$$|\mathbf{cl}^+(\sigma; f)| + |\mathbf{st}^-(\sigma; f)| \leq 1$$

holds for every cell $\sigma \in X$, with $|\cdot|$ denoting cardinality — see Definition 2.2 and Lemma 2.5 of [12] for details. Since X is finite, we can always perturb f slightly so that it becomes injective while preserving all relevant cardinalities. Let us now impose a stratification X_\bullet on X , as in Section 1.1.

Definition 3.1. A **stratified discrete Morse function** on X is any injective map $f : X \hookrightarrow \mathbb{R}$ such that for every cell σ lying in a stratum $S \subset X$, we have

- (1) $|\mathbf{cl}^+(\sigma; f) \cap S| + |\mathbf{st}^-(\sigma; f) \cap S| \leq 1$, and
- (2) if $\mathbf{cl}^+(\sigma; f) \cap S = \{\tau\}$, then $\tau \in \mathbf{h}(\sigma; f)$ and $\mathbf{cl}(\sigma) \searrow \mathbf{sh}(\sigma; f)$.

The first condition of this definition requires the restriction $f|_S : S \rightarrow \mathbb{R}$ to be a standard (albeit injective) discrete Morse function on each stratum S in the sense of [12]. The second condition is new to the best of our knowledge. Whenever a cell $\sigma \in S$ admits a unique τ in $\mathbf{cl}^+(\sigma; f) \cap S$, this condition requires the closure of σ to collapse onto the shadow subcomplex from (2). In other words, all cells in the augmented halo $\widehat{\mathbf{h}}(\sigma; f)$ must disappear along a sequence of elementary collapses². Another novel feature engendered by our definition is that critical cells come in two flavours.

Definition 3.2. Let $f : X \hookrightarrow \mathbb{R}$ be an injective function satisfying Definition 3.1.1. A cell σ lying in a stratum $S \subset X$ is called

- **critical** for f if both $\mathbf{cl}^+(\sigma; f) \cap S$ and $\mathbf{st}^-(\sigma; f) \cap S$ are empty; and,
- **s-critical** for f if it is critical in the sense above and $\mathbf{st}^-(\sigma; f)$ is empty.

(The requirement that $\mathbf{st}^-(\sigma; f) = \emptyset$ is equivalent to $\sigma \notin \mathbf{h}(\tau; f)$ for any cell $\tau \in X$).

From [12, Definition 2.2] it follows that σ is critical if and only if it is critical in Forman's sense for the restriction $f|_S : S \rightarrow \mathbb{R}$. If the stratification is trivial, then the only stratum in sight is $S = X$; in this case, s-criticality coincides with criticality because the first requirement of the above definition already forces $\mathbf{st}^-(\sigma; f)$ to be empty.

4. Local Morse Data

We now turn towards the task of establishing the fundamental results pertaining to sublevelsets of a stratified discrete Morse function $f : X \hookrightarrow \mathbb{R}$. The first of these pertains to the case where an interval $[c, d] \subset \mathbb{R}$ contains no s-critical values. The goal, as mentioned in the preamble to Lemma 2.4, is to relate $\mathbf{cl}(f_{\leq c})$ and $\mathbf{cl}(f_{\leq d})$.

THEOREM 4.1. Let $f : X \hookrightarrow \mathbb{R}$ be a stratified discrete Morse function and consider an interval $[c, d] \subset \mathbb{R}$ for which $f_{\geq c} \cap f_{\leq d}$ consists of a single cell σ , which is not s-critical. Then there exists a filtered collapse $\mathbf{cl}(f_{\leq d}) \searrow \mathbf{cl}(f_{\leq c})$.

PROOF. If $f(\sigma) = c$ then by the uniqueness of σ we get $f_{\leq c} = f_{\leq d}$ and there is nothing to prove; thus we assume $f(\sigma) \in (c, d]$. The argument now decomposes into two cases: either the lower star $\mathbf{st}^-(\sigma; f)$ is empty, or it is not. Of these, the nonempty case is far simpler.

²Moreover, the first pair in any such sequence is necessarily (τ, σ) , because σ lies in the open star of every cell in $\mathbf{cl}(\sigma)$ and the only other cell in S that it can possibly be paired with σ is τ .

Case 1: If $\mathbf{st}^-(\sigma; f)$ contains some cell τ , then by (1) we have both $\sigma < \tau$ and $f(\sigma) \geq f(\tau)$. By the injectivity of f and the uniqueness of σ in $f_{\geq c} \cap f_{\leq d}$, we get $f(\tau) < c$, which forces $\sigma \in \mathbf{cl}(f_{\leq c})$. Lemma 2.4(1) now gives $\mathbf{cl}(f_{\leq d}) = \mathbf{cl}(f_{\leq c})$, so the trivial collapse suffices.

Case 2: If $\mathbf{st}^-(\sigma; f)$ is empty, then since $f(\sigma) > c$ we have $\sigma \notin \mathbf{cl}(f_{\leq c})$. Therefore Lemma 2.4(2) yields

$$\mathbf{cl}(f_{\leq d}) - \mathbf{cl}(f_{\leq c}) = \widehat{\mathbf{h}}(\sigma; f).$$

Since σ is not s-critical and $\mathbf{st}^-(\sigma; f)$ is empty, Definition 3.1(2) forces $\mathbf{cl}(\sigma) \setminus \mathbf{sh}(\sigma; f)$. We now claim that the same sequence of elementary collapses which reduces $\mathbf{cl}(\sigma)$ to $\mathbf{sh}(\sigma; f)$ also reduces $\mathbf{cl}(f_{\leq d})$ to $\mathbf{cl}(f_{\leq c})$. First, note from (2) that

$$\mathbf{cl}(\sigma) - \mathbf{sh}(\sigma; f) = \widehat{\mathbf{h}}(\sigma; f),$$

so in both scenarios one must excise only those cells which lie in the augmented halo of σ . It remains to confirm that the free-face pairs removed from $\mathbf{cl}(\sigma)$ are also free-face pairs in $\mathbf{cl}(f_{\leq d})$. This amounts to checking open stars, as described in Section 1.2. Explicitly, for each cell $\gamma \in \widehat{\mathbf{h}}(\sigma; f)$, we must establish

$$[\mathbf{st}(\gamma) \cap \mathbf{cl}(\sigma)] \supset [\mathbf{st}(\gamma) \cap \mathbf{cl}(f_{\leq d})],$$

since the reverse containment is guaranteed by the fact that $f(\sigma) \leq d$. We will proceed by contradiction, assuming the existence of some $\rho > \gamma$ lying in $\mathbf{cl}(f_{\leq d}) - \mathbf{cl}(\sigma)$. By definition of $\mathbf{cl}(f_{\leq d})$ there exists a coface $\rho' \geq \rho$ with $f(\rho') \leq d$. Now $\rho \notin \mathbf{cl}(\sigma)$ implies $\rho' \neq \sigma$; and since σ is the only cell valued in $[c, d]$, we get $f(\rho') < c \leq f(\sigma)$. This violates the argmin property required by $\gamma \in \widehat{\mathbf{h}}(\sigma; f)$ and hence concludes the proof. \square

One strange consequence of the above result is that critical cells of f can only affect the filtered topology of sublevelsets if they also happen to be s-critical. We now quantify the change in topology from $\mathbf{cl}(f_{\leq c})$ to $\mathbf{cl}(f_{\leq d})$ whenever the interior (c, d) contains a single s-critical value. This is the purview of the second fundamental theorem of Morse theory. Since our f is injective, we may safely restrict attention to the case where $f_{\geq c} \cap f_{\leq d}$ is a singleton.

THEOREM 4.2. *Let $f : K \hookrightarrow \mathbb{R}$ be a stratified discrete Morse function and $\sigma \in X$ an s-critical cell with $c := f(\sigma)$. For all sufficiently small $\varepsilon > 0$, the subcomplex $\mathbf{cl}(f_{\leq c+\varepsilon})$ is the union of $\mathbf{cl}(\sigma)$ with $\mathbf{cl}(f_{\leq c-\varepsilon})$ along the intersection $\mathbf{sh}(\sigma; f)$.*

PROOF. By finiteness of X and injectivity of f , we are free to choose $\varepsilon > 0$ such that σ is the only cell of X lying in $f_{\geq c-\varepsilon} \cap f_{\leq c+\varepsilon}$. And since σ is s-critical, we know from Definition 3.2 that $\mathbf{st}^-(\sigma; f)$ is empty. Consequently, σ is not in $\mathbf{cl}(f_{\leq c-\varepsilon})$, so Lemma 2.4(2) implies that $\mathbf{cl}(f_{\leq c+\varepsilon})$ equals the union $\mathbf{cl}(f_{\leq c-\varepsilon}) \cup \widehat{\mathbf{h}}(\sigma; f)$. Since $\widehat{\mathbf{h}}(\sigma; f) \subset \mathbf{cl}(\sigma)$ holds by Definition 2.1, we get the containment

$$\mathbf{cl}(f_{\leq c+\varepsilon}) \subset [\mathbf{cl}(f_{\leq c-\varepsilon}) \cup \mathbf{cl}(\sigma)].$$

Conversely, $\mathbf{cl}(\sigma) \subset \mathbf{cl}(f_{\leq c+\varepsilon})$ since $f(\sigma) = c \leq c + \varepsilon$, so the reverse containment also holds. This proves that $\mathbf{cl}(f_{\leq c+\varepsilon})$ equals the union displayed above. It remains to establish the pushout property by showing that

$$\mathbf{sh}(\sigma; f) = \mathbf{cl}(f_{\leq c-\varepsilon}) \cap \mathbf{cl}(\sigma).$$

Below we will show that the left side is contained in the right side and vice-versa.

* **left \subset right** : By (2) and Definition 2.1, if a cell γ lies in $\mathbf{sh}(\sigma; f)$ then $\gamma < \sigma$ and

$$\sigma \neq \operatorname{argmin}[f / \mathbf{st}^-(\gamma; f)].$$

We claim that this inequality forces $\gamma \in \mathbf{cl}(f_{\leq c-\varepsilon})$. If $\mathbf{st}^-(\gamma; f)$ is empty, then $f(\gamma) < f(\sigma)$ and hence $f(\gamma) \leq c - \varepsilon$ and we get the desired $\gamma \in f_{\leq c-\varepsilon} \subset \mathbf{cl}(f_{\leq c-\varepsilon})$. Alternately, if some $\rho \in \mathbf{st}^-(\gamma; f)$ satisfies $f(\rho) < f(\sigma)$, then $f(\rho) \leq c - \varepsilon$ and hence $\gamma \in \mathbf{cl}(f_{\leq c-\varepsilon})$.

*** right \subset left:** We show that any cell γ in $\mathbf{cl}(\sigma) \cap \mathbf{cl}(f_{\leq c-\varepsilon})$ must lie outside $\widehat{\mathbf{h}}(\sigma; f)$. If this γ lies in $f_{\leq c-\varepsilon}$ then we have $\sigma \notin \mathbf{st}^-(\gamma; f)$ since σ has a higher f -value. Thus, σ can not be the argmin of f over the lower star of γ , which means that γ lies outside the augmented halo of σ . On the other hand, if γ lies in $\mathbf{cl}(f_{\leq c-\varepsilon})$ because some cell $\rho > \gamma$ satisfies $f(\rho) \leq c - \varepsilon$, then the presence of this ρ once again rules out the possibility that $\sigma = \operatorname{argmin}[f / \mathbf{st}^-(\sigma; f)]$ because we have

$$f(\rho) \leq c - \varepsilon < c = f(\sigma).$$

Consequently, we see that γ must lie in $\mathbf{cl}(\sigma) - \widehat{\mathbf{h}}(\sigma; f)$ and hence in $\mathbf{sh}(\sigma; f)$ by (2). \square

5. Local Morse Data via the Upper Envelope

The **barycentric subdivision** of a finite regular CW complex X is the simplicial complex $\mathbf{Sd}(X)$ whose k -simplices, for all $k \geq 0$, are all strictly ascending chains of the form

$$\xi := [\sigma_0 < \sigma_1 < \cdots < \sigma_k]$$

in the face poset (X, \leq) . The faces $\eta < \xi$ are obtained by deleting one or more cells from the underlying chain $\sigma_0 < \cdots < \sigma_k$. There is a natural order-preserving *last cell* map $\omega : \mathbf{Sd}(X) \rightarrow X$ that sends each such ξ to its terminal cell σ_k . It is relatively straightforward to check that $\mathbf{Sd}(X)$ canonically inherits a stratification from X via ω — namely, for each stratum $S \subset X$ there exists a unique stratum $S' \subset \mathbf{Sd}(X)$ such that the simplex ξ lies in S' if and only if its last cell $\omega(\xi)$ lies in S . It is readily seen that S' consists exclusively of barycentrically subdivided S -cells, so we will denote it by $\mathbf{Sd}(S)$ rather than S' . We assume throughout this section that X — and hence, $\mathbf{Sd}(X)$ — is stratified.

Definition 5.1. The **upper envelope** of an injective function $f : X \hookrightarrow \mathbb{R}$ is the function $\mathbf{U}f : \mathbf{Sd}(X) \rightarrow \mathbb{R}$ given by

$$\mathbf{U}f([\sigma_0 < \cdots < \sigma_k]) := \max \{f(\sigma_i) \mid 0 \leq i \leq k\}.$$

(One may safely define upper envelopes for arbitrary $X \rightarrow \mathbb{R}$, but the only case of interest to us here occurs when f is injective).

The reader is warned that $\mathbf{U}f$ is almost never injective even though the original f is injective. Consequently, upper envelopes of stratified discrete Morse functions on X rarely yield stratified discrete Morse functions on $\mathbf{Sd}(X)$. The advantage of passing to the upper envelope is that $\mathbf{U}f$ is order-preserving — given $\xi \leq \eta$ in $\mathbf{Sd}(X)$, we always have $\mathbf{U}f(\xi) \leq \mathbf{U}f(\eta)$ in \mathbb{R} . This monotonicity fuels the following elementary result.

Proposition 5.2. *Given any function $f : X \rightarrow \mathbb{R}$ and value $c \in \mathbb{R}$, the sublevelset*

$$\mathbf{U}f_{\leq c} := \{\xi \in \mathbf{Sd}(X) \mid \mathbf{U}f(\xi) \leq c\}$$

is a subcomplex of $\mathbf{Sd}(X)$.

PROOF. Assume a simplex $\xi \in \mathbf{Sd}(X)$ lies in $\mathbf{U}f_{\leq c}$ and consider a face $\eta < \xi$. We see from Definition 5.1 that $\mathbf{U}f(\eta)$ is the maximum of f over a smaller set, and therefore we obtain $\mathbf{U}f(\eta) \leq \mathbf{U}f(\xi) \leq c$. Therefore, we have $\eta \in \mathbf{U}f_{\leq c}$ as desired. \square

We now seek to explicitly describe the local Morse data of a stratified Morse function $f : X \hookrightarrow \mathbb{R}$ at a cell $\sigma \in X$ in terms of the sublevelsets of $\mathbf{U}f$ in $\text{Sd}(X)$ rather than (closures of) sublevelsets of f in X . For this purpose, it is convenient to better understand the neighbourhood of the corresponding vertex $[\sigma] \in \text{Sd}(X)$. The open star of $[\sigma]$ is evidently the collection of all barycentric simplices whose underlying chains contain σ in some position:

$$\mathbf{st}[\sigma] = \{[\sigma_0 < \cdots < \sigma_k] \in \text{Sd}(X) \mid \sigma_i = \sigma \text{ for some } 0 \leq i \leq k\}.$$

It follows that the **link** of $[\sigma]$, given by

$$\mathbf{lk}[\sigma] := \mathbf{cl}(\mathbf{st}[\sigma]) - \mathbf{st}[\sigma],$$

is the subcomplex of $\text{Sd}(X)$ comprising all ξ whose underlying chain $\sigma_0 < \cdots < \sigma_k$ can be augmented by σ .

Remark 5.3. Specifically, given any $\xi = [\sigma_0 < \cdots < \sigma_k]$ in $\mathbf{lk}[\sigma]$, we must have either

- (1) $\sigma < \sigma_0$, so σ fits at the beginning of the chain; or,
- (2) $\sigma_k < \sigma$, so σ fits at the end of the chain; or,
- (3) $\sigma_i < \sigma < \sigma_{i+1}$ for some i in $\{0, 1, \dots, k-1\}$.

In every case these inequalities are strict, so $\sigma \neq \sigma_i$ for any i and therefore $\xi \notin \mathbf{st}[\sigma]$.

To confirm that $\mathbf{lk}[\sigma]$ is a subcomplex of $\text{Sd}(X)$ as asserted above, one can either use the fact that $\mathbf{st}[\sigma]$ is an up-set in its closure, or simply note that the removal of any σ_i from $\sigma_0 < \cdots < \sigma_k$ preserves σ -augmentability. Set $c := f(\sigma)$, and define the **lower link** of $[\sigma]$ along $\mathbf{U}f$ as

$$\mathbf{lk}_f^-[\sigma] := \mathbf{lk}[\sigma] \cap \mathbf{U}f_{<c}, \quad (4)$$

where $\mathbf{U}f_{<c}$ is defined analogously to $\mathbf{U}f_{\leq c}$. Since X is finite, there exists $\varepsilon > 0$ such that $\mathbf{U}f_{<c}$ coincides with $\mathbf{U}f_{\leq c-\varepsilon}$, and hence forms a subcomplex of $\text{Sd}(X)$ by Proposition 5.2. It follows that the lower link is an intersection of subcomplexes, and hence a subcomplex of $\text{Sd}(X)$. Lower links play a crucial role in the study of sublevelsets of injective real-valued maps defined on the vertices of a simplicial complex. By far most prominent instance of this phenomenon is found in the piecewise linear Morse theory of Bestvina and Brady [4].

Definition 5.4. Let us denote by K_i the set of i -dimensional simplices in a simplicial complex K . Given another simplicial complex L , we recall that the **join** $K \star L$ is the new complex whose n -simplices are given by

$$(K \star L)_n = K_n \sqcup L_n \sqcup \coprod_{i+j=n-1} K_i \times L_j.$$

The **cone** over L is the special case of this construction where K consists of a single vertex.

The cone $[\sigma] \star \mathbf{lk}_f^-[\sigma]$ is evidently a subcomplex of $\text{Sd}(X)$; combined with the lower link itself, it completely characterises the inclusion of sublevelsets $\mathbf{U}f_{\leq c-\varepsilon} \hookrightarrow \mathbf{U}f_{\leq c+\varepsilon}$ when f is injective on X .

Proposition 5.5. Let $\mathbf{U}f : \text{Sd}(X) \rightarrow \mathbb{R}$ be the upper envelope of an injective map $f : X \hookrightarrow \mathbb{R}$. For every cell $\sigma \in X$ with $c := f(\sigma)$, there exists sufficiently small $\varepsilon > 0$ such that $\mathbf{U}f_{\leq c+\varepsilon}$ is the union of $\mathbf{U}f_{\leq c-\varepsilon}$ with $[\sigma] \star \mathbf{lk}_f^-[\sigma]$ along their intersection $\mathbf{lk}_f^-[\sigma]$.

PROOF. Since f is injective and X is finite, we may safely choose an $\varepsilon > 0$ such that σ is the unique cell of X lying in the intersection $f_{\geq c-\varepsilon} \cap f_{\leq c+\varepsilon}$. By Definition 5.1, the upper envelope $\mathbf{U}f$ attains the same values as f in \mathbb{R} . Therefore, any cell

$$\eta = [\sigma_0 < \dots < \sigma_k] \text{ lying in } (\mathbf{U}f_{\leq c+\varepsilon} - \mathbf{U}f_{\leq c-\varepsilon})$$

satisfies $\mathbf{U}f(\eta) = c$. It now follows from Definition 5.1 and the injectivity of f that $\sigma = \sigma_i$ for some $0 \leq i \leq k$ and $f(\sigma_j) < c$ for all the other $j \neq i$. Therefore, using the uniqueness of σ guaranteed by our choice of ε , we know that the face $\xi < \eta$ obtained by removing σ_i from the underlying chain satisfies $f(\xi) \leq c - \varepsilon$, and hence $\xi \in \mathbf{lk}_f^-(\sigma)$. Thus, our η has the form $\xi \cup [\sigma]$ for some $\xi \in \mathbf{lk}_f^-(\sigma)$, as desired. \square

The preceding result may disappoint the reader who seeks a more faithful discretisation of f 's local Morse data at $[\sigma]$ à la Goresky and MacPherson [14]. The central discrepancy here is that the lower link L of $[\sigma]$ along f does not admit an obvious factorisation of the form $TL \times NL$, where TL is a *tangential* component which lies entirely within the unique stratum $\text{Sd}(S) \subset X$ containing $[\sigma]$, while NL is the *normal* component that resides in strictly higher strata $\text{Sd}(T) > \text{Sd}(S)$. There are two immediate obstacles — first, simplicial complexes are not closed under products, so it is unreasonable to expect the desired TL and NL to form subcomplexes of $\text{Sd}(X)$. Secondly, the simplicial setting frustrates all naïve attempts to take arbitrarily small neighbourhoods around $[\sigma]$ — the smallest available option is $\mathbf{cl}(\mathbf{st}[\sigma])$, which will necessarily intersect *lower* strata $\text{Sd}(T) < \text{Sd}(S)$ whenever there exists a face $\tau < \sigma$ which lies in T .

Nevertheless, we persist. Consider a (not necessarily critical) cell σ of X and fix an injective function $f : X \hookrightarrow \mathbb{R}$. For brevity, we write L for the lower link $\mathbf{lk}_f^-(\sigma)$ from (4). The first step is to isolate two relevant subcomplexes of L , which correspond to the first two cases of Remark 5.3.

Definition 5.6. A simplex $\xi = [\sigma_0 < \dots < \sigma_k]$ in L is said to lie in

- (1) the **horizontal part** H of L whenever $\sigma_k < \sigma$, and in
- (2) the **vertical part** V of L whenever $\sigma < \sigma_0$.

(Both defining properties are preserved when passing to faces of ξ , whence H and V form subcomplexes of L).

We catalogue the following elementary and helpful properties of H and V .

Proposition 5.7. *Let H and V be the horizontal and vertical parts of the lower link L of σ along f . Then,*

- (1) *the intersection $H \cap V$ is empty;*
- (2) *the join $H \star V$ is L .*

PROOF. For (1), note that if a simplex $\xi = [\sigma_0 < \dots < \sigma_k]$ lies in $H \cap V$ then we arrive at the absurdity $\sigma < \sigma_0 < \sigma_k < \sigma$, where the first inequality comes from the definition of V and the last inequality comes from the definition of H . Moving to (2), assume that ξ lies in the difference $L - (H \cup V)$. By Remark 5.3 there is some i in $\{0, 1, \dots, k-1\}$ with $\sigma_i < \sigma < \sigma_{i+1}$. But now, the front face $[\sigma_0 < \dots < \sigma_i]$ of ξ lies in H while the back face $(\sigma_{i+1} < \dots < \sigma_k)$ lies in V , so ξ is indeed the disjoint union of simplex of H and a simplex of V as required by Definition 5.4. \square

Combining the preceding result with Proposition 5.5, we see that the local Morse data for $\mathbf{U}f$ at σ is $(H \star V) \hookrightarrow ([\sigma] \star (H \star V))$. Since the join operation is manifestly associative on simplicial complexes, we may safely rewrite this as

$$(H \star V) \hookrightarrow (([\sigma] \star H) \star V). \quad (5)$$

Our next task is to examine how H and V interact with the stratification of $\text{Sd}(X)$ inherited from X . From here onwards, we will assume that $f : X \hookrightarrow \mathbb{R}$ satisfies Definition 3.1 and that σ is a critical cell of f , but do not insist on s-criticality.

Proposition 5.8. *Let H and V be the horizontal and vertical part of the link L of $[\sigma] \in \text{Sd}(X)$. If $S \subset X$ is the stratum containing σ , then*

- (1) *H lies in the union $\bigcup_{T \leq S} \text{Sd}(T)$ of (subdivided) strata lying at or below the level of S in the frontier partial order, while*
- (2) *V lies in the union $\bigcup_{T > S} \text{Sd}(T)$ of (subdivided) strata lying strictly above S with respect to the frontier partial order.*

(In fact the first assertion holds for all σ , whereas the second assertion requires σ to be critical).

PROOF. To establish (1), we note from Definition 5.6 that a simplex $\xi = [\sigma_0 < \dots < \sigma_k]$ of L lies in H if and only if the last cell $\omega(\xi) = \sigma_k$ is a strict face of σ in X . Thus if $T \subset X$ is the stratum containing σ_k , we get a nonempty $\text{cl}(S) \cap T$ and hence $T \leq S$. Passing to the subdivision, ξ must lie in $\text{Sd}(T)$ as desired. For (2), given $\xi \in V$, we know that σ is a strict face of the first cell σ_0 by definition of V , and that $f(\sigma_0) < f(\sigma)$ holds because ξ lies in the lower link L of σ . Therefore, we have $\sigma_0 \in \text{st}^-(\sigma; f)$, and the fact that σ is critical now forces $\sigma_0 \notin S$ (as per Definition 3.2). Thus, σ_0 lies in some $T_0 > S$ along the frontier order. Proceeding similarly from left to right along the chain $\sigma_0 < \dots < \sigma_k = \omega(\xi)$, we see that the stratum T_i containing σ_i must satisfy $T_i \geq T_{i-1}$ for all $i > 0$. In particular, it follows that the stratum T_k whose subdivision contains ξ satisfies $T_k \geq T_0 > S$, as desired. \square

Finally, we combine (5) with Propositions 5.5 and 5.8 to arrive at the main result of this Section.

THEOREM 5.9. *Let σ be a critical cell for an injective function $f : X \hookrightarrow \mathbb{R}$ which satisfies the first requirement of Definition 3.1. Set $c := f(\sigma)$ and choose any $\varepsilon > 0$ so that σ is the unique cell of X with $f(\sigma) \in [c - \varepsilon, c + \varepsilon]$. Let H and V be the horizontal and vertical parts of the lower link of $[\sigma]$ along $\mathbf{U}f$. Then,*

- (1) *$\mathbf{U}f_{\leq c+\varepsilon}$ is the union of $\mathbf{U}f_{\leq c-\varepsilon}$ with $([\sigma] \star H) \star V$ along their intersection $H \star V$; and,*
- (2) *if $S \subset X$ is the stratum containing σ , then $[\sigma] \star H$ lies in the union of strata $\leq \text{Sd}(S)$ while V lies in the union of strata $> \text{Sd}(S)$.*

PROOF. The only assertion which has not been explicitly addressed in previous results is the claim that $[\sigma] \star H$ lies in the union U of strata $\leq \text{Sd}(S)$. We already know that every simplex $\xi \in H$ lies in U from Proposition 5.8, so it suffices to check that $[\sigma] \cup \xi$ is also in U . By Definition 5.6, we see that σ must be the last cell in the underlying chain of $[\sigma] \cup \xi$, whence $[\sigma] \cup \xi$ always lies in S and hence in U , as desired. \square

References

- [1] K. Adiprasito and B. Benedetti. Collapsibility of CAT(0) spaces. *Geometriae Dedicata*, 206:181–199, 2020.
- [2] M. F. Atiyah and R. Bott. The Yang–Mills equations over Riemann surfaces. *Philosophical Transactions of the Royal Society of London. Series A*, 308(1505):523–615, 1983.

- [3] E. Babson and P. Hersh. Discrete Morse functions from lexicographic orders. *Transactions of the American Mathematical Society*, 357(2):509–534, 2005.
- [4] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. *Inventiones Mathematicae*, 129:445–470, 1997.
- [5] M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Mathematics*, 217(1):101–113, 2000.
- [6] M. M. Cohen. *A course in simple homotopy theory*. Springer-Verlag, 1973.
- [7] C. C. Conley. *Isolated Invariant Sets and the Morse Index*, volume 38 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1978.
- [8] O. Cornea, K. A. De Rezende, and M. R. Da Silveira. Spectral sequences in Conley’s theory. *Ergodic Theory and Dynamical Systems*, 30(4):1009–1054, 2010.
- [9] J. Curry, R. Ghrist, and V. Nanda. Discrete Morse theory for computing cellular sheaf cohomology. *Foundations of Computational Mathematics*, 16(4):875–897, 2016.
- [10] D. Farley and L. Sabalka. Discrete Morse theory and graph braid groups. *Algebraic & Geometric Topology*, 5:1075–1109, 2005.
- [11] A. Floer. Morse theory for Lagrangian intersections. *Journal of Differential Geometry*, 28(3):513–547, 1988.
- [12] R. Forman. Morse theory for cell complexes. *Advances in mathematics*, 134(1):90–145, 1998.
- [13] R. Freij. Equivariant discrete Morse theory. *Discrete Mathematics*, 309(12):3821–3829, 2009.
- [14] M. Goresky and R. MacPherson. *Stratified Morse theory*. Springer, 1988.
- [15] M. Jöllenbeck and V. Welker. *Minimal Resolutions via Algebraic Discrete Morse Theory*, volume 197 of *Memoirs of the American Mathematical Society*. AMS, 2009.
- [16] K. Knudson and B. Wang. Discrete stratified Morse theory: Algorithms and a user’s guide. *Discrete & Computational Geometry*, 67(4):1023–1052, 2022.
- [17] M. Kukieła. The main theorem of discrete Morse theory for Morse matchings with finitely many rays. *Topology and its Applications*, 160(9):1074–1082.
- [18] M. Lipiński, J. Kubica, M. Mrozek, and T. Wanner. Conley–Morse–Forman theory for generalized combinatorial multivector fields on finite topological spaces. *Journal of Applied and Computational Topology*, 7:139–184, 2023.
- [19] J. N. Mather. Notes on topological stability. *Bulletin of the American Mathematical Society*, 49(4):475–506, 2012.
- [20] K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, 50(2):330–353, 2013.
- [21] M. Morse. The calculus of variations in the large. *Annals of Mathematics*, 26(3):213–234, 1925.
- [22] M. Mrozek. Conley–Morse–Forman theory for combinatorial multivector fields on lefschetz complexes. *Foundations of Computational Mathematics*, 17:1585–1633, 2017.
- [23] M. Mrozek and T. Wanner. Dynamics of combinatorial multivector fields. In *Connection Matrices in Combinatorial Topological Dynamics*, SpringerBriefs in Mathematics, pages 105–118. Springer, Cham, 2015.
- [24] V. Nanda. Discrete Morse theory and localization. *Journal of Pure and Applied Algebra*, 223(2):459–488, 2019.
- [25] V. Nanda. Local cohomology and stratification. *Foundations of Computational Mathematics*, 20(2):195–222, 2020.
- [26] V. Nanda, D. Tamaki, and K. Tanaka. Discrete Morse theory and classifying spaces. *Advances in Mathematics*, 340:723–790, 2018.
- [27] D. Salamon. Connected simple systems and the Conley index of isolated invariant sets. *Bulletin of the American Mathematical Society*, 291:1–41, 1985.
- [28] M. Salvetti. Topology of the complement of real hyperplanes in \mathbb{R}^n . *Inventiones Mathematicae*, 88(3):603–618, 1987.
- [29] S. Weinberger. *The Topological Characterization of Stratified Spaces*. The University of Chicago Press, 1994.
- [30] H. Whitney. Tangents to an analytic variety. *Annals of Mathematics*, 81(3):496–549, 1965.
- [31] E. Witten. Supersymmetry and Morse theory. *Journal of Differential Geometry*, 17(4):661–692, 1982.
- [32] N. Yerolemou and V. Nanda. Morse theory for complexes of groups. *Journal of Pure and Applied Algebra*, 228(6):107606, 2024.

Appendix A. Multivector Fields and the Conley Index

Let X be a finite regular CW complex. A **multivector** of X is any nonempty, convex and connected subset M of the poset (X, \leq) . A partition $X = \bigsqcup_i M_i$ into multivectors is a **multivector field**; these have found substantial use in the combinatorial study of dynamical systems [22, 23, 18]. Here are four examples of multivector fields, in approximately ascending order of generality:

- (1) The most trivial case occurs whenever each multivector is a single cell of X , so the multivector field is simply the given partition of X into its constituent cells.
- (2) Given a discrete Morse function $f : X \rightarrow \mathbb{R}$ in the sense of Forman, each multivector is either a single f -critical cell or a pair $\{\sigma, \tau\}$ of cells where $\sigma < \tau$, $\dim \tau = \dim \sigma + 1$ and $f(\sigma) \geq f(\tau)$. The resulting multivector field, called an *acyclic partial matching*, is fundamental in discrete Morse theory [5].
- (3) In [13, Sec 3.2], Freij considers a partition of X into *intervals*; each interval I_σ^τ consists of all cells $\gamma \in X$ of the form $\sigma \leq \gamma \leq \tau$ for some fixed $\sigma \leq \tau$ in X , not necessarily of codimension ≤ 1 . Such interval partitions are also multivector fields.
- (4) Let X_\bullet be a stratification of X as described in Section 1.1 above. The collection of strata constitutes a multivector field on X — connectedness holds by definition and convexity follows from Proposition 1.3.

Conversely, not every multivector field produces a stratification in our sense because the frontier axiom may not hold.

The *exit set* of a multivector $M \subset X$ is $\mathbf{ex}(M) := \mathbf{cl}(M) - M$; since M is convex, it forms an up-set in its closure, whence $\mathbf{ex}(M)$ is always a (possibly empty) subcomplex of $\mathbf{cl}(M)$. The following definition forms the starting point in a discrete version of Conley's celebrated topological study [7] of isolated invariant sets — see for instance [23, Eq (7.4)].

Definition A.1. The **homological Conley index** of a multivector $M \subset X$ is defined as the relative homology group

$$\mathbf{Con}_k(M) := H_k(\mathbf{cl}(M), \mathbf{ex}(M); \mathbb{Z})$$

for each dimension $k \geq 0$. We say that M is **critical** whenever $\mathbf{Con}_k(M) \neq 0$ for some k .

It is readily checked (by writing out the relevant chain complexes) that when $M = \{\sigma\}$ is a single cell, then $\mathbf{Con}_k(M)$ is nontrivial iff $k = \dim \sigma$. On the other hand, if $M = \{\sigma, \tau\}$ contains two cells with $\sigma < \tau$ and $\dim \tau = \dim \sigma + 1$, then $\mathbf{Con}_k(M)$ is trivial for all $k \geq 0$.

Consider the binary relation \blacktriangleright on a multivector field $\mathcal{M} = \{M_i \subset X\}$ given by

$$M_i \blacktriangleright M_j \text{ if there exist } \sigma_i \in M_i \text{ and } \sigma_j \in M_j \text{ with } \sigma_i \geq \sigma_j,$$

(This is clearly reflexive.) We say that \mathcal{M} is **acyclic** whenever the transitive closure \triangleright of \blacktriangleright forms a partial order on the set of multivectors. Acyclic multivector fields therefore generalise the acyclic partial matchings induced by discrete Morse functions [5]. In fact, it is easily checked that acyclic multivector fields also subsume the stratifications from Section 1.1.

Proposition A.2. *The strata in any stratification of X form an acyclic multivector field, and \triangleright coincides with the frontier partial order on strata.*

PROOF. By definition, $S \blacktriangleright T$ holds for a pair of strata $S \neq T$ if and only if there exist cells $\sigma \in S$ and $\tau \in T$ with $\sigma \geq \tau$, which in turn occurs if and only if $\mathbf{cl}(S) \cap T$ is nonempty. The desired conclusion now follows from Proposition 1.2. \square

The reader is warned that the preceding result does not imply that the frontier axiom holds amongst the multivectors of an acyclic multivector field; we only recover the partial order implied by the frontier axiom via Proposition 1.2. On the other hand, acyclicity does allow us to recover the global homology of X from the Conley indices of its multivectors. Explicitly, since \triangleright is a partial order, we may enumerate the multivectors $\{M_1, \dots, M_n\}$ such that $M_i \triangleright M_j$ forces $i \geq j$. Therefore,

$$F_q := \bigcup_{i \leq q} M_i$$

constitutes a filtration of X by subcomplexes. We may now examine the associated **Conley-Morse spectral sequence** of this filtration [27, 8], which is populated by relative homology groups

$$E_{p,q}^1 := H_{p+q}(F_q, F_{q-1}; \mathbb{Z}),$$

and converges to $E_{p,q}^\infty \simeq H_{p+q}(X)$. Crucially, it follows from excision that $E_{p,q}^1$ is isomorphic to the Conley index $\mathbf{Con}_{p+q}(M_q)$. Proposition A.2 guarantees a similar mechanism for assembling the global homology of X from the Conley indices of its strata.