Min-Max Theory for Cell Complexes

Lacey Johnson   Kevin Knudson

Department of Mathematics, University of Florida, Gainesville, FL 32611, USA
E-mail: johnsonla@ufl.edu   kknudson@ufl.edu

Abstract. In the study of smooth functions on manifolds, min-max theory provides a mechanism for identifying critical values of a function. We introduce a discretized version of this theory associated to a discrete Morse function on a (regular) cell complex. As applications we prove a discrete version of the mountain pass lemma and give an alternate proof of a discrete Lusternik–Schnirelmann theorem.

1 Introduction

Given a smooth function $f : M \to \mathbb{R}$, where M is a smooth manifold, a collection of *min-max data* for f is a pair $(\mathcal{H}, \mathcal{S})$, where $\mathcal{H}$ is a collection of homeomorphisms of M containing functions that decrease the sublevel sets of f and $\mathcal{S}$ is a collection of subsets of M closed under the collection $\mathcal{H}$. Given such a pair, the min-max principle produces critical values of the function f.

In this paper we introduce a discretized version of this theory. While we state results for simplicial complexes, our results hold on the category of regular cell complexes. After reviewing the basics of discrete Morse theory in Section 2, we define in Section 3 the notion of min-max data associated to a discrete Morse function $f : K \to \mathbb{R}$ on a simplicial complex K. The collection $\mathcal{H}$ will not be a set of homeomorphisms; rather, it consists of functions defined on the power set of the set of cells of K. There are various operators associated to the function f which yield the requisite property of decreasing the level sets, and so we are able to obtain a discrete min-max principle to identify critical values of f.

As applications of this, we deduce the following. First, we obtain a discrete version of the mountain pass lemma: given two local minima for f, there is a critical edge e "between" them; see Section 4 (a proof of the corresponding smooth result may be found in [5, p. 89]). We also give an alternate proof of a discrete Lusternik–Schnirelmann theorem as first presented in [1]; this is the content of Section 5.

2 Discrete Morse Theory

Let K be any finite simplicial complex, where K need not be a triangulated manifold nor have any other special property [3]. When we write K we mean the set of simplices of K; by $|K|$ we mean the underlying topological space. Let $\alpha^{(p)} \in K$ denote a simplex of dimension p. Let $\alpha < \beta$ denote that the simplex α is a face of the simplex β. If $f \colon K \to \mathbb{R}$ is a function, we define $U(\alpha) = \{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\}$ and $L(\alpha) = \{\gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha)\}$. In other words, $U(\alpha)$ contains the immediate cofaces of α with lower (or equal) function values, while $L(\alpha)$ contains the immediate faces of α with higher (or equal) function values. Let $|U(\alpha)|$ and $|L(\alpha)|$ be their sizes.

Definition 2.1. A function $f \colon K \to \mathbb{R}$ is a *discrete Morse function* if for every $\alpha^{(p)} \in K$, (i) $|U(\alpha)| \leq 1$ and (ii) $|L(\alpha)| \leq 1$.

Forman showed that conditions (i) and (ii) are exclusive — if one of the sets $U(\alpha)$ and $L(\alpha)$ is nonempty, then the other one must be empty (see [3, Lemma 2.5]). Therefore, each simplex $\alpha \in K$ can be paired with at most one exception simplex: either a face γ with larger function value, or a coface β with smaller function value. We refer to conditions (i) and (ii) as the *Morse conditions*.

Definition 2.2. A simplex $\alpha^{(p)}$ is *critical* if (i) $|U(\alpha)| = 0$ and (ii) $|L(\alpha)| = 0$. A *critical value* of f is its value at a critical simplex.

Definition 2.3. A simplex $\alpha^{(p)}$ is *regular* if either of the following conditions holds: (i) $|U(\alpha)| = 1$; (ii) $|L(\alpha)| = 1$. As noted above, these conditions cannot be both true (see [3, Lemma 2.5]).

Note that regular simplices occur in pairs. Suppose there is a pair $\sigma^{(p)} < \tau^{(p+1)}$ of simplices in K such that σ has no other cofaces in K. Then $K \setminus \{\sigma, \tau\}$ is a simplicial complex called an *elementary collapse* of K. We say that K *collapses to* L if L can be obtained from K by a sequence of elementary collapses. We denote this by $K \searrow L$. A complex K is *collapsible* if $K \searrow \{v\}$ for some vertex v.

Two Morse functions $f, g \colon K \to \mathbb{R}$ are *equivalent* if for every pair $\sigma^{(p)} < \tau^{(p+1)}$ in K, $f(\sigma) < f(\tau)$ if and only if $g(\sigma) < g(\tau)$. A proof of the following is embedded in the proof of [3, Theorem 3.3].

Lemma 2.4. *Let $f \colon K \to \mathbb{R}$ be a discrete Morse function. Then there is an injective discrete Morse function $g \colon K \to \mathbb{R}$ equivalent to f having the same critical simplices as f.*

Main results of discrete Morse theory. Given $c \in \mathbb{R}$, we have the *level subcomplex* $K^c = \bigcup_{f(\alpha) \leq c} \bigcup_{\beta \leq \alpha} \beta$. That is, K^c contains all simplices α of K such that $f(\alpha) \leq c$ along with all of their faces. We have the following two combinatorial versions of the main results of classical Morse theory. The fundamental idea in their proofs is to collapse the sublevel complexes by removing pairs $\{\sigma < \tau\}$ in the discrete gradient vector field of the function f.

Theorem 2.5. [3, Theorem 3.3] *Suppose the interval (a, b) contains no critical value of f. Then K^b is homotopy equivalent to K^a. In fact, K^b collapses onto K^a.*

The next theorem explains how the topology of the sublevel complexes changes as one passes a critical value of a discrete Morse function. In what follows, $\dot{e}^{(p)}$ denotes the boundary of a p-cell $e^{(p)}$.

Theorem 2.6. [3, Theorem 3.4] *Suppose that $\sigma^{(p)}$ is a critical simplex with $f(\sigma)$ in $(a, b]$, and there are no other critical simplices with values in $(a, b]$. Then K^b is homotopy equivalent to attaching a p-cell $e^{(p)}$ along its entire boundary in K^a; that is, $K^b = K^a \bigcup_{\dot{e}^{(p)}} e^{(p)}$.*

The associated gradient vector field. Given a discrete Morse function $f \colon K \to \mathbb{R}$, we may associate a discrete gradient vector field as follows. Since any noncritical simplex $\alpha^{(p)}$ has at most one of the sets $U(\alpha)$ and $L(\alpha)$ nonempty, there is a unique face $\nu^{(p-1)} < \alpha$ with $f(\nu) \geq f(\alpha)$ or a unique coface $\beta^{(p+1)} > \alpha$ with $f(\beta) \leq f(\alpha)$. Denote by V the collection of all such pairs $\{\sigma < \tau\}$ (referred to as *Morse pairs*). Then every simplex in K is in at most one pair in V and the simplices not in any pair are precisely the critical simplices of the function f. We call V the *gradient vector field associated to f*. We visualize V by drawing an arrow from α to β for every pair $\{\alpha < \beta\} \in V$. Observe that V "points" in the direction of *decrease* for f; as such it corresponds to the *negative* gradient of the function and we may write $V = -\nabla f$. Theorems 2.5 and 2.6 may then be visualized in terms of V by collapsing the pairs in V using the arrows. Thus, a discrete gradient (or equivalently, a discrete Morse function) provides a collapsing order for the complex K, simplifying it to a complex L with potentially fewer cells but having the same homotopy type.

More generally, a *discrete vector field* on K is a collection V of pairs of simplices $\{\alpha^{(p)} < \beta^{(p+1)}\}$ such that each simplex in K is in at most one pair in V. By a *V-path* we mean a sequence

$$\alpha_0^{(p)} < \beta_0^{(p+1)} > \alpha_1^{(p)} < \beta_1^{(p+1)} > \cdots < \beta_r^{(p+1)} > \alpha_{r+1}^{(p)},$$

where each $\{\alpha_i < \beta_i\}$ is a pair in V. Such a path is *nontrivial* if $r > 0$, and *closed* if $\alpha_{r+1} = \alpha_0$. Forman proved the following result characterizing which discrete vector fields correspond to gradients of discrete Morse functions.

Theorem 2.7. [3, Theorem 9.3] *V is a gradient vector field if and only if it has no nontrival closed V-paths.*

Ascending and descending regions. If $g \colon M \to \mathbb{R}$ is a Morse function on a smooth compact manifold M and if p is a critical point of g, there are two subspaces associated to p: the ascending and descending discs. The former is the set of points q in M such that the integral curve of $-\nabla g$ passing through q converges to p as $t \to \infty$, while the latter is the set of points whose integral curves converge to p as $t \to -\infty$. If p has index λ, then the ascending disc is topologically a disc of dimension $n - \lambda$ and the descending disc has dimension λ.

The discrete analogues of these objects were constructed by Jerše and Mramor Kosta in [4]. The idea is that descending discs should be unions of V-paths, just as in the smooth case (where descending discs are foliated by integral curves of a gradient-like vector field). Ascending discs are then constructed as descending discs for the dual vector field V^* on the dual cell complex K^*.

For our purposes, however, we need only consider ascending discs of local minima and for this it suffices to consider the *basin* of a minimum v, defined in [2]. This is the maximal subcomplex of K that collapses to v; we denote it by $A(v)$.

3 Min-max data associated to a discrete Morse function

Let K be a finite simplicial complex and let $f : K \to \mathbb{R}$ be a discrete Morse function (which we may assume injective). Recall that for $a \in \mathbb{R}$, K^a denotes the sublevel complex associated to a. We denote by L^a the *sublevel set* $L^a = \{\sigma \in K : f(\sigma) \le a\}$. Note that K^a is the closure of L^a. Denote by $\mathcal{P}(K)$ the power set of K.

Definition 3.1. A collection of *min-max data* for the discrete Morse function $f : K \to \mathbb{R}$ is a pair $(\mathcal{H}, \mathcal{S})$ satisfying the following conditions:
  (i) $\mathcal{H}$ is a collection of maps $\mathcal{P}(K) \to \mathcal{P}(K)$ such that for every regular value a of f there exist $\varepsilon > 0$ and $h \in \mathcal{H}$ such that $h(L^{a+\varepsilon}) \subset L^{a-\varepsilon}$.
  (ii) $\mathcal{S}$ is a collection of subsets of K such that for all $h \in \mathcal{H}$ and all $S \in \mathcal{S}$, $h(S) \in \mathcal{S}$.

Theorem 3.2. *If $(\mathcal{H}, \mathcal{S})$ is a collection of min-max data for the discrete Morse function f on K, then the number $c = \min_{S \in \mathcal{S}} \max_{\sigma \in S} f(\sigma)$ is a critical value for f.*

Proof. Suppose the contrary; that is, assume that c is a regular value for f. Then there exist $\varepsilon > 0$ and $h \in \mathcal{H}$ such that $h(L^{c+\varepsilon}) \subset L^{c-\varepsilon}$. By the definition of c, there is an $S \in \mathcal{S}$ such that $\max_{\sigma \in S} f(\sigma) < c + \varepsilon$; that is, $S \subset L^{c+\varepsilon}$. Thus, $S' = h(S) \in \mathcal{S}$ and $h(S) \subset L^{c-\varepsilon}$. It follows that $\max_{\sigma \in S'} f(\sigma) \le c - \varepsilon$, which implies $\min_{S' \in \mathcal{S}} \max_{\sigma \in S'} f(\sigma) \le c - \varepsilon$, contrary to the choice of c. $\square$

4 Mountain Pass Lemma

Our goal in this section is to prove the mountain pass lemma in the setting of an arbitrary simplicial complex. Recall the classical smooth situation. Given a Morse function $f : M \to \mathbb{R}$ on the compact manifold M, suppose f has two local minima, say at x_0 and x_1 with $f(x_0) < f(x_1)$. Assuming M to be connected, we see that there is a path from x_0 to x_1. The mountain pass lemma asserts that the infimum over all such paths γ of the maximum value $\sup_t \{\gamma(t)\}$ is a critical value of f.

The picture to have in mind is shown in Figure 1. Thinking of f as a height function, the minima at x_0 and x_1 represent depressions in the landscape. To get from one low point to another, one must cross over a ridge; the lowest point z on the ridge will be a saddle point, i.e., a critical point of f.
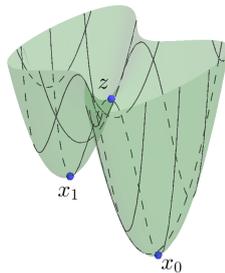
Figure 1. Visualizing mountain pass points

While this result may seem obvious, it is not at all clear how one could prove it on a general n-manifold, let alone on a general simplicial complex. Our plan is to construct an appropriate collection of min-max data and then apply Theorem 3.2.

Suppose V is a discrete gradient vector field on K, say $V = -\nabla f$ for some discrete Morse function f. Denote by $\langle -, - \rangle$ the inner product on the chains $C_p(K, \mathbb{Z})$ defined by making all cells orthogonal, and define a map $V : C_p(K, \mathbb{Z}) \to C_{p+1}(K, \mathbb{Z})$ as follows. Suppose that σ is a p-cell with fixed orientation. If there is a $(p+1)$-cell τ with $\{\sigma < \tau\} \in V$, then set $V(\sigma) = -\langle \partial \tau, \sigma \rangle \tau$, where ∂ is the boundary map in K. If there is no such τ (for example, if σ is critical), then set $V(\sigma) = 0$. We extend this linearly to a map on all of $C_p(K)$.

Definition 4.1. The *discrete flow operator* $\phi : C_p(K) \to C_p(K)$ is defined by

$$\phi = \text{Id} + \partial V + V \partial.$$

The intuition here is that if v is a vertex, then either (a) v is critical and should be fixed by the flow, or (b) v is not critical and it should flow to the other vertex that the edge e is paired with by the gradient V; that is, $V(v) = \pm e = v + \partial V(v)$.

The following is [3, Theorem 6.4].

Proposition 4.2. *Let $\sigma_1, \ldots, \sigma_r$ denote the p-cells of K, each with a chosen orientation. Write $\phi(\sigma_i) = \sum_j a_{ij} \sigma_j$. Then the following statements hold:*
  (i) *For every i, $a_{ii} = 0$ or 1, and $a_{ii} = 1$ if and only if σ_i is critical.*
  (ii) *If $i \neq j$, then $a_{ij} \in \mathbb{Z}$. If $i \neq j$ and $a_{ij} \neq 0$, then $f(\sigma_j) < f(\sigma_i)$.*

For our purpose, we need to consider the collection of cells that appear in the image of ϕ. Note that we may apply ϕ to sets of p-cells in K to obtain a map, denoted by $\Phi : \mathcal{P}(K) \to \mathcal{P}(K)$, defined by $\Phi(A) = \bigcup_{\sigma \in A} \phi(\sigma)$. That is, for a cell σ we are using the notation $\Phi(\sigma)$ to mean the collection of cells appearing in the chain $\phi(\sigma)$. Define also a second map $\overline{\Phi}$ on $\mathcal{P}(K)$ by

$$\overline{\Phi}(A) = \text{subcomplex of } K \text{ generated by } \Phi(A).$$

We give examples of Φ and $\overline{\Phi}$ in Figure 2.
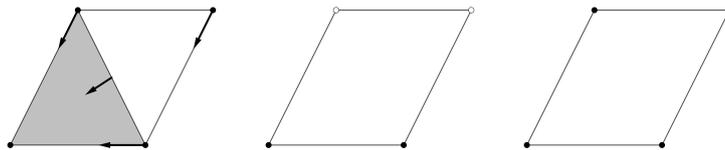
Figure 2. A complex K (left), $\Phi(K)$ (center), $\overline{\Phi}(K)$ (right)

The maps Φ and $\overline{\Phi}$ will be useful for constructing min-max data. Consider $\mathcal{H} = \{\Phi\}$. By Proposition 4.2, condition (i) in the definition of min-max data for a discrete Morse function is satisfied. In fact, for any a, regular or not, $\Phi(K^a) \subseteq K^a$ (see, e.g., [2, Lemma 1]), but Φ only decreases the sublevel set for regular values. The same is true for the map $\overline{\Phi}$. In fact, we have the following, which is a special case of [2, Theorem 4].

Lemma 4.3. *For each $a \in \mathbb{R}$, $K^a \searrow \overline{\Phi}(K^a)$.*

Proof. Note that any critical simplex in K^a is also in $\overline{\Phi}(K^a)$, and hence so are its faces. Thus, a simplex σ in $K^a \backslash \overline{\Phi}(K^a)$ must be regular and paired with some $\tau \in K^a$. Order such pairs by decreasing f-value; the maximal pair must be of the form $\{\alpha < \beta\}$, where α is a free face of β (note that if a is a regular value, then this pair consists of the unique simplex in $f^{-1}(a)$ and either a coface of lower value or a face of higher value). Removing the pairs in decreasing order yields a collapsing sequence $K^a \searrow \overline{\Phi}(K^a)$. $\qquad\square$

Remark 4.4. Note that the collapse $K^a \searrow \overline{\Phi}(K^a)$ consists of elementary collapses of V-pairs $\{\sigma < \tau\}$. This will be important in Section 5 below.

Now suppose that v_0 and v_1 are critical vertices for the discrete Morse function f and assume $f(v_0) < f(v_1)$. Let $\mathcal{S}$ be the collection of subsets of K consisting of sets of the form $E = \{v_1, e_1, e_2, \ldots, e_r\}$, where the e_i form an edge path beginning at v_1 and ending in the basin $A(v_0)$. Writing $e_i = \langle u_{i-1}, u_i \rangle$, we insist on the following condition: if u_j belongs to $A(v_0)$, then $f(e_j) > f(e_{j+1}) > \cdots > f(e_r)$. This condition is not strictly necessary, but it does prevent path bifurcation when acted upon by Φ. We may assume that such paths do not pass through another critical vertex $v \neq v_0, v_1$. Note that if $E \in \mathcal{S}$, we have $\Phi(E) \neq \emptyset$. Indeed, since v_1 is a critical vertex, $\Phi(v_1) = v_1$.

Theorem 4.5. *If $E \in \mathcal{S}$, then $\Phi(E) \in \mathcal{S}$.*

Assuming this, we see that $(\mathcal{H}, \mathcal{S})$ is a collection of min-max data on K. By Theorem 3.2 we see that $c = \min_{E \in \mathcal{S}} \max_{e \in E} f(e)$ is a critical value of f. This critical value must be greater than $f(v_1)$ since any set E in $\mathcal{S}$ has the form $E = \{v_1, e_1, \ldots, e_r\}$ with $r \geq 1$ and $f(e_1) > f(v_1)$. Thus, among all the edge paths from v_1 to v_0, there is one passing through a critical edge e and the value of $f(e)$ is minimal among all such paths. Hence, e may be thought of as the "saddle point" along the "ridge" between v_0 and v_1.

Modulo the proof of Theorem 4.5 we have thus proved the discrete version of the mountain pass lemma and we refer to the resulting critical edges as mountain pass edges. The mountain pass lemma implies that if a discrete Morse function has two vertices that are strict local minima, then it must admit a critical edge.

Another explanation of why there should exist a critical edge with function value greater than $f(v_1)$ is to observe that the sublevel complex $K^{f(v_1)}$ is disconnected while the cell complex K is connected. The change of topological type in going from $K^{f(v_1)}$ to K can be explained by the presence of a critical edge with function value greater than $f(v_1)$.

Remark 4.6. A slightly different version of this result was proved in [2], where the authors did not employ min-max theory to do so.

Proof of Theorem 4.5. We begin by computing the action of Φ on a single edge $e = \langle v, w \rangle$. There are three cases.

(i) Suppose that e is critical. Then $\phi(e) = e + V\partial e = e - V(v) + V(w)$. It follows that $\Phi(e) = \{e, V(v), V(w)\}$.

(ii) If e is regular and paired with one of its vertices (say v), then we have $\phi(e) = e + V\partial e = V(w)$. Thus, $\Phi(e) = \{V(w)\}$.

(iii) If e is regular and paired with the 2-simplex with edges e, e', e'', then we have $\phi(e) = e + V\partial e + \partial V(e)$. Several things could happen in this case depending on how the vertices v and w are paired (or not) by V. We will evaluate this scenario as needed in what follows.

We proceed by induction on r, the number of edges in a set E. The case $r = 1$ is essentially described above. Here we have $E = \{v_1, e\}$ for some edge e. The other vertex u of e lies in the ascending disc of v_0. If e is critical, then we have $\phi(e) = e + V(u)$, and so $\Phi(E) = \{v_1, e, V(u)\}$ is another path ending in $A(v_0)$. Since u is in $A(v_0)$, it cannot be paired with e, and so the second case above cannot occur. This leaves the final possibility that e is regular and paired with a 2-simplex $\langle v_1, u, w \rangle$. In this case we see that $\phi(e) = \langle v_1, w \rangle + \langle u, w \rangle + V(u)$. It could be that $V(u) = -\langle u, w \rangle$, which implies that w also belongs to $A(v_0)$ and hence $\Phi(E) = \{v_1, \langle v_1, w \rangle\}$ is another element of $\mathcal{S}$. Otherwise, $V(u)$ is another edge ending in $A(v_0)$ and so $\Phi(E) = \{v_1, \langle v_1, w \rangle, \langle u, w \rangle, V(u)\} \in \mathcal{S}$.

Assume the result holds for all paths of length smaller than r and suppose $E = \{v_1, e_1, \ldots, e_r\} \in \mathcal{S}$. If $\{v_1, e_1, \ldots, e_{r-1}\}$ ends in $A(v_0)$, then $\phi(e_r)$ also lands in $A(v_0)$; by the induction hypothesis, $\Phi(E) \in \mathcal{S}$. Suppose that $\{v_1, e_1, \ldots, e_{r-1}\}$ does not end in $A(v_0)$. It means there is no gradient path

$$u_{r-1} < \epsilon_0 > w_1 < \epsilon_1 > \cdots < \epsilon_s > v_0.$$

However, there is a path $\eta_0 > u_r < \eta_1 > \nu_1 < \eta_2 > \cdots < \eta_t > v_0$. It follows that $e_r \neq \eta_0$ and $f(e_r) > f(u_r)$. Note that this implies that u_{r-1} could not be paired with e_r (or else $f(u_r) < f(e_r)$, which would imply $u_{r-1} \in A(v_0)$). Thus, e_r is paired with a 2-simplex σ (whose faces are e_r, e', and e''), or e_r is critical. In the first case, $\phi(e_r) = e' + e'' + V(u_r) + V(u_{r-1})$ simply diverts the path E around the edges of σ, and so $\Phi(E) \in \mathcal{S}$. If e_r is critical, then $\phi(e_r) = e_r + V(u_{r-1}) + V(u_r)$ and so $\Phi(E)$ still lies in $\mathcal{S}$. This completes the proof. $\qquad\square$

Remark 4.7. The reader might wonder why we used the map Φ and edge paths that do not include the vertices as opposed to the map $\overline{\Phi}$ and 1-dimensional subcomplexes. The issue is that arbitrary edge paths emanating from v_1 might contain vertices that are paired with edges not in the path. Applying Φ or $\overline{\Phi}$ to such an object leads to subcomplexes that have extraneous edges jutting from the path. We would then be forced to expand the collection $\mathcal{S}$ to accommodate this, causing some technical difficulties.

5 A Discrete Lusternik–Schnirelmann Theorem

Suppose that K is a simplicial complex and that L is a subcomplex. Recall that L is *collapsible* if it collapses to a vertex. The following definition appears in [1].

Definition 5.1. Suppose that L is a subcomplex of K. Then L has *geometric precategory less than or equal to m in K*, denoted by $\widetilde{\mathrm{dgcat}}_K(L) \leq m$ if there exist $m + 1$ subcomplexes $U_0, \ldots, U_m$ in K, each of which is collapsible in K, such that $L \subseteq \bigcup_{i=0}^m U_i$. If $\widetilde{\mathrm{dgcat}}_K(L) \not< m$, we say $\widetilde{\mathrm{dgcat}}_K(L) = m$. Then the *discrete (geometric) category* of L in K is $\mathrm{dgcat}_K(L) = \min\{\widetilde{\mathrm{dgcat}}_K(L') : L \searrow L'\}$.

Suppose that $f\colon K \to \mathbb{R}$ is a discrete Morse function with gradient V. When L and L' are subcomplexes of K, we say that L *V-collapses to* L', denoted by $L \searrow_V L'$, if $L \searrow L'$ and each elementary collapse in the sequence is a pair in V. For each k, let $\Gamma_k = \{L \subset K : \exists\, K^a \searrow_V L \text{ with } \mathrm{dgcat}_K(K^a) \geq k-1\}$. Recall that if $A \in \mathcal{P}(K)$, we define $\overline{\Phi}(A)$ to be the subcomplex of K generated by $\Phi(A)$.

Theorem 5.2. *For each k, $(\{\overline{\Phi}\}, \Gamma_k)$ is a collection of min-max data.*

Proof. Suppose that $L \in \Gamma_k$. Then there is some $a \in \mathbb{R}$ with $K^a \searrow_V L$. We claim that $\overline{\Phi}(K^a) = \overline{\Phi}(L)$. First note that if L is obtained from K^a by a single elementary collapse ($L = K^a - \{\sigma, \tau\}$ with σ a free face of τ and $\{\sigma < \tau\} \in V$), then since $\Phi(\tau) = \emptyset$ and $\Phi(\sigma) \subset \partial\tau \subset L$, we see that $\Phi(K^a) = \Phi(L)$, and hence $\overline{\Phi}(K^a) = \overline{\Phi}(L)$. The general case follows by induction on the number of V-collapses $K^a \searrow_V L$. Since Lemma 4.3 implies $K^a \searrow_V \overline{\Phi}(K^a)$ (Remark 4.4), we see that $\overline{\Phi}(L) \in \Gamma_k$. $\quad\square$

Corollary 5.3. *For each k, the real number $c_k = \min_{L \in \Gamma_k} \max_{\sigma \in L} f(\sigma)$ is a critical value of f.*

As a consequence of this, we obtain the following Lusternik–Schnirelmann theorem, originally proved in [1, Theorem 26].

Corollary 5.4. *Let $f\colon K \to \mathbb{R}$ be a discrete Morse function with m critical values. Then $\mathrm{dgcat}(K) + 1 \leq m$.*

Proof. By Corollary 5.3, each c_k is a critical value of f. We may assume that the global minimum of f, which occurs at some vertex v, is $f(v) = 0$. Then $c_1 = 0$ and K^{c_1} contains one critical simplex. Proceeding inductively, we suppose that $K^{c_{k-1}}$ has at least $k-1$ critical simplices and consider K^{c_k}. Since by Lemma 2.4 we may assume that f is injective, $c_{k-1} < c_k$ and so $f^{-1}(c_k)$ contains at least one new critical simplex. Thus, K^{c_k} contains at least k critical simplices. So if $c_1 < c_2 < \cdots < c_{\mathrm{dgcat}(K)+1}$ are the distinct critical values, then $K^{c_{\mathrm{dgcat}(K)+1}} \subseteq K$ contains at least $\mathrm{dgcat}(K) + 1$ critical simplices. Therefore, $\mathrm{dgcat}(K) + 1 \leq m$. $\quad\square$

References

[1] S. Aaronson, N. Scoville, Lusternik–Schnirelmann category for simplicial complexes, *Illinois J. Math.* **57** (2013) 743–753.

[2] O. Delgado-Friedrichs, V. Robins, A. Sheppard, Skeletonization and partitioning of digital images using discrete Morse theory, *IEEE Trans. Pattern Anal. Mach. Intell.* **37** (2015) 654–666.

[3] R. Forman, Morse theory for cell complexes, *Adv. Math.* **134** (1998) 90–145.

[4] G. Jerše, N. Mramor Kosta, Ascending and descending regions of a discrete Morse function, *Comput. Geom.* **42** (2009) 639–651.

[5] L. Nicolaescu, *An Invitation to Morse Theory*, 2nd ed., Springer, New York, 2011.