

Organized collapse: An introduction to discrete Morse theory, by Dmitry N. Kozlov,
Graduate Studies in Mathematics, Vol. 207, American Mathematical Society,
Providence, RI, 2020, xxiii+312 pp., hardcover, ISBN 978-1-4704-5701-3

1. TOPOLOGICAL SIMPLIFICATION

Topologists, and those who use topological methods in their work, like simplicial models for spaces. They are easy to describe in terms of a collection of vertices and which subsets form simplices, and this information may be stored readily in a computer. There are drawbacks, however. Minimal triangulations may be difficult to come by and may still contain a large number of simplices. For example, the minimal triangulation of the torus requires fourteen 2-simplices [3], and higher-dimensional spaces are more complicated still. Moreover, for the purposes of computing homology, simplicial complexes can be a bit unwieldy. For example, in data analysis applications one often works with complexes containing millions of simplices; the resulting boundary matrices are enormous, and the corresponding linear algebra can take quite a while for a computer to process.

On the other hand, one learns in a first topology course that spaces may also be modeled as cell complexes with a correspondingly simpler homology theory. The torus may be built from just four cells—one 0-cell, two 1-cells, and one 2-cell—and the resulting cellular chain complex has all boundary maps equal to 0, yielding an immediate calculation of the homology. The question then is how one bridges these two points of view.

Consider the triangulation of the torus shown in Figure 1. This is not minimal, but it is close. There are a few simplices that stand out: the corners (which form a single vertex after the sides of the square are identified), the two bold edges in the center of the figure, and the darker triangle at the top. Observe that there are many arrows in this diagram, pointing from vertices to edges, or from edges to triangles. Note further that there are no arrows emanating from the corners or the bold edges, nor are there arrows pointing into the bold edges or the dark triangle. Now imagine beginning with the dark triangle and deforming it into the other triangles by following the arrows emanating from its edges, following all the possible paths; in the end there is one large 2-cell with some edges running through it. Then collapse the marked edges by following the arrows from vertices into the edges to the other end. Note that the ends of the dark edges eventually merge with the single vertex, and the large 2-cell is glued to these. At the end of this process, what is left is a cell complex with one 0-cell, two 1-cells, and one 2-cell, and while it is not obvious that the corresponding boundary maps in the cellular chain complex are trivial, it is possible to prove this.

If this deformation process seems familiar, it is because it is similar to what happens when one examines the (negative) gradient vector field of a Morse function on a manifold. Recall that a smooth function $f : M \rightarrow \mathbb{R}$ is a *Morse function* on the compact n -manifold M if all the critical points of f are nondegenerate; that is, in local coordinates at a critical point p , the Hessian matrix of second partial derivatives of f at p is nonsingular. The *index* of p is then the number of negative




FIGURE 1. A triangulation of the torus

eigenvalues of the Hessian. The fundamental theorems of Morse theory are the following.

Theorem 1.1. Denote by M_a the sublevel set $f^{-1}(-\infty, a]$. If the interval $[a, b]$ does not contain a critical value of f , then M_a is a deformation retract of M_b .

Theorem 1.2. Suppose c is a critical value of f with $f^{-1}(c) = p$. Denote by i the index of p . Then for sufficiently small $\varepsilon > 0$ the space $M_{c+\varepsilon}$ is diffeomorphic to $M_{c-\varepsilon}$ with an i -handle attached: $M_{c+\varepsilon} \cong M_{c-\varepsilon} \cup_{\varphi} (D^i \times D^{n-i})$.

These theorems are proved by flowing along the negative gradient of f . The first theorem is fairly simple; the second is rather delicate (the proof consumes several pages in [6]). One consequence of these theorems is that the manifold has the homotopy type of a cell complex with one cell of dimension i for each critical point of index i . Indeed, the first theorem tells us that the topology of the sublevel sets can change only when passing a critical value, while the second theorem tells us that when we do pass a critical value, we attach an i -cell. (Note: the i -handle is a product of an i -disc called the core and an $(n - i)$ -disc called the co-core; in the proof of Theorem 1.2 one sees that $M_{c+\varepsilon}$ deformation retracts to $M_{c-\varepsilon} \cup D^i$.)

The deformation process outlined above for the simplicial torus is the idea underlying *discrete Morse theory*. It is based on the *simple homotopy theory* of J.H.C. Whitehead [9], in which one employs *elementary collapses* to deform one space to another. To perform a simplicial collapse we need a *free simplex*, a simplex τ which is a face of exactly one other simplex σ . If such a pair exists, we may remove

them from the simplicial complex without changing its homotopy type. Moreover, we may view this removal as a strong deformation retraction onto the remaining complex. For the torus example, however, we do not have any such free simplices, but the process we described may be thought of as a sequence of *internal collapses* transforming our simplicial torus into a simpler cell complex without changing homotopy type. The questions then are what sorts of collapses are allowed and how can we use them to simplify our complex.

2. ACYCLIC MATCHINGS

An important property of gradient vector fields is that there are no closed orbits. That is, if one begins at a critical point for a smooth function and flows along the gradient field, one cannot return to the starting point. To pursue this analogy, let us call the unpaired simplices in the simplicial torus *critical*. Observe that if one begins at a face of a critical cell and follows paths of arrows, there are no closed loops; that is, no such path returns to the boundary of the critical cell. We are now ready for some definitions.

Definition 2.1. Let K be a simplicial complex. A *discrete vector field* on K is a collection V of pairs $\{\alpha < \beta\}$, where α is a codimension-1 face of β , such that each simplex in K is in at most one pair in V . If σ is not in a pair in V , we call σ *critical*. A V -*path* is a sequence

$$\alpha_0 < \beta_0 > \alpha_1 < \beta_1 > \cdots < \beta_{r-1} > \alpha_r$$

in which each pair $\alpha_i < \beta_i$ is an element of V and all the α_i and β_i are distinct for $i < r$. The path is *nontrivial* if $r > 0$, and it is *closed* if $\alpha_r = \alpha_0$. If V has no nontrivial closed paths, we call it a *discrete gradient vector field*.

A convenient way to visualize discrete vector fields is to use arrows as in Figure 1: draw an arrow from α to β for any pair $\{\alpha < \beta\} \in V$. Perhaps a better point of view is to consider the *Hasse diagram* of K . This is the directed graph whose vertices are the simplices of K , with a directed edge from β to α whenever α is a codimension-1 face of β . Now, if we are given a discrete vector field V on K , we reverse the arrows for each pair $\{\alpha < \beta\} \in V$. The vector field V is then a gradient if and only if this modified Hasse diagram is acyclic; that is, it has no directed loops. This point of view then allows one to use the extensive catalog of algorithms in directed graph theory to study these objects.

Summary: A discrete gradient vector field on K is the same as an acyclic partial matching on the Hasse diagram of K . We now have some theorems relating these to the topology of K . In what follows, we assume that K is a finite simplicial complex, although that condition may often be relaxed.

Theorem 2.2. *Suppose K' is a subcomplex of K . Then there is a sequence of elementary collapses from K to K' if and only if there is an acyclic complete matching on the set of all simplices of K not contained in K' .*

This is analogous to Theorem 1.1, even though we do not have functions and sublevel sets (more on this in Section 3 below). What we do have is a gradient, and since every simplex not in K' may be paired, this theorem says that we can remove them, preserving homotopy type.

What about Theorem 1.2? There is a direct analogue (again, more on this in Section 3), but for now we will content ourselves with the consequence that a

manifold equipped with a Morse function has the homotopy type of a cell complex whose cells correspond to the critical points of the function.

Theorem 2.3. *Suppose V is a discrete gradient vector field on K (i.e., an acyclic partial matching on the Hasse diagram of K). Then there exists a cell complex X such that*

- (1) *for each d , the number of d -cells in X is equal to the number of d -simplices in K that are critical with respect to V ; and*
- (2) *there is a homotopy equivalence $K \simeq X$.*

This theorem formalizes the ideas we discussed in the torus example above. In contrast to the proof of the manifold version of this theorem, the proof of Theorem 2.3 is relatively short and easily digestible. Moreover, if we consider the cellular chain complex $C_\bullet(X)$, the boundary maps are fairly easy to compute in terms of the sequence of internal collapses determined by V . If one is willing to work with $\mathbb{Z}/2$ coefficients for homology, then the boundary maps are given by counting the number of V -paths between critical cells of adjacent dimensions. Again, this is in contrast with the smooth case, where the associated Morse complex can be quite tricky to understand, at least at first glance.

3. DISCRETE MORSE FUNCTIONS

The language of acyclic matchings is particularly clean, but the original definition of discrete Morse theory given by Forman [2] begins with the function point of view.

Definition 3.1. A *discrete Morse function* on K is a map f on the simplices of K , which we denote by $f : K \rightarrow \mathbb{R}$, satisfying the following two conditions:

- (1) For each p -simplex α , the set $\{\nu^{(p-1)} < \alpha : f(\nu) > f(\alpha)\}$ has cardinality at most 1;
- (2) for each p -simplex α , the set $\{\beta^{(p+1)} > \alpha : f(\beta) < f(\alpha)\}$ has cardinality at most 1.

A simplex is *critical* if both of these sets are empty. A simple example of a discrete Morse function is given by $f(\alpha) = \dim \alpha$; every cell is critical in this case. It is easy to see that the two conditions above are exclusive in the sense that if one set is nonempty the other must be empty. From this, one can then define the associated gradient vector field: every simplex in K has at most one other simplex to pair with. In the other direction, given an acyclic partial matching V on the Hasse diagram of K (what we had been calling a discrete gradient above), a classical theorem in graph theory asserts that the modified Hasse diagram supports a function which decreases along any directed path. Any such function is a discrete Morse function on K whose associated gradient is V .

With this language, Forman [2] proved direct analogues of Theorems 1.1 and 1.2. If a is a real number, denote by K_a the sublevel complex

$$K_a = \bigcup_{\substack{\sigma \in K \\ f(\sigma) \leq a}} \bigcup_{\tau \leq \sigma} \tau.$$

Note the second union here; we must include it because a given simplex α with $f(\alpha) \leq a$ may have a face with a larger function value.

Theorem 3.2. *Suppose the interval $[a, b]$ contains no critical values for the discrete Morse function f . Then K_b collapses to K_a .*

Theorem 3.3. *Suppose the interval $[a, b]$ contains exactly one critical value c and that there is a unique critical d -simplex σ with $f(\sigma) = c$. Then K_b is homotopy equivalent to $K_a \cup_{\partial e^d} e^d$, where e^d is a d -cell with boundary ∂e^d .*

Theorems 2.2 and 2.3 are direct consequences of these, of course.

Since the topological information is completely captured by a discrete gradient vector field, it is natural to ask why one would consider functions at all. The short answer is that vector fields are qualitative while functions are quantitative. In some applications, it may be necessary to keep some quantitative information intact. For example, suppose one is given a function on the vertices of a simplicial complex. The picture to have in mind here is a sampling of some smooth function on a manifold with the sample points taken as the vertices of a triangulation, a standard situation in data analysis or physical modeling. The smooth function has associated dynamics. Is it possible to construct a discrete Morse function that captures this information? The answer is yes; see [4].

4. SOME APPLICATIONS

We present here a simple application, namely the *complex of not-connected graphs*. Denote by K_n the complete graph on n vertices. A *spanning subgraph* is a subgraph $G \subseteq K_n$ containing all n vertices. Let N_n be the collection of all spanning subgraphs which are not connected. Observe that if $G_1 \subset G_2$ and if $G_2 \in N_n$, then $G_1 \in N_n$ as well. From N_n we construct a simplicial complex X_n as follows. The k -simplices of X_n are the graphs having $k + 1$ edges. If G is such a graph, then the faces of G are all the nontrivial spanning subgraphs of G .

Theorem 4.1. *The space X_n has the homotopy type of a wedge of $(n - 1)!$ copies of the sphere S^{n-3} .*

Forman proved this theorem by constructing a discrete gradient on X_n . Note that we need only consider $n \geq 3$ since $N_1 = \emptyset$ and N_2 consists of a single element. Label the vertices of any $G \in N_n$ by $\{1, 2, \dots, n\}$ and denote by e_{ij} the edge in K_n joining vertex i with vertex j (a particular G may or may not contain e_{ij}). Let V_{12} be the vector field consisting of all pairs $\{G < G + e_{12}\}$ for graphs $G \in N_n$ with $e_{12} \notin G$ and $G + e_{12} \in N_n$. Note that $G_0 = e_{12}$ remains unpaired. The graphs in X_n other than G_0 which are not paired in V_{12} are those G not containing e_{12} with $G + e_{12}$ connected; such a G has exactly two connected components, one containing vertex 1 and one containing vertex 2. Denote these components by G_1 and G_2 . In such a G , vertex 3 lies in either G_1 or G_2 . If it lies in G_1 and G does not contain e_{13} , then $G + e_{13}$ is unpaired in V_{12} ; add the pair $\{G < G + e_{13}\}$. If G contains e_{13} , then G remains unpaired if and only if $G - e_{13}$ is the union of three connected components, each containing one of vertices 1, 2, or 3. If vertex 3 is in G_2 and $e_{23} \notin G$, then add the pair $\{G < G + e_{23}\}$. Denote the resulting vector field by V_3 . The unpaired graphs in V_3 are G_0 and those G that either

- (a) contain e_{13} and are such that $G - e_{13}$ is the union of three connected components, each containing one of the vertices 1, 2, or 3; or
- (b) contain e_{23} and are such that $G - e_{23}$ is the union of three connected components, each containing one of the vertices 1, 2, or 3.

Now consider vertex 4 and pair any G which is unpaired in V_3 with $G + e_{14}$, $G + e_{24}$, or $G + e_{34}$, if possible (at most one of these graphs is unpaired in V_3). This gives the vector field V_4 . Continue in this manner to generate $V = V_n$. It is not

difficult to see that V is a discrete gradient and that the only unpaired graphs in V are G_0 and those graphs that are the union of two connected trees, one containing vertex 1 and one containing vertex 2. There are $(n-1)!$ such graphs, each having $(n-2)$ edges; they therefore correspond to $(n-3)$ -simplices in X_n . Theorem 2.3 then gives the desired result.

Theorem 4.1 is a typical application of discrete Morse theory. Many simplicial complexes arise in combinatorial applications and analyzing their topology is often possible via these techniques. Examples include the order complex of the partition lattice (Section 11.5.2 in the book under review), explicit calculations of homology generators, and Hom-complexes of graphs (Section 14.1.3 in the book under review).

Discrete Morse theory has also found a home in applied algebraic topology via its utility in reducing homology calculations. Indeed, one may use acyclic matchings to reduce the simplicial chain complex of K to a smaller complex with basis the critical cells of a discrete gradient. While finding an optimal gradient (i.e., one with the minimum number of critical simplices) is an NP-complete problem, one can still reduce the problem significantly via clever reductions (or even not-so-clever ones). Many software packages for computing persistent homology take advantage of discrete Morse theory to improve efficiency (e.g., [1], [7]).

5. THE BOOK UNDER REVIEW

Kozlov's text is the third book-length treatment of discrete Morse theory. The reviewer's book [5] introduces classical Morse theory on manifolds alongside Forman's discrete Morse theory, with a focus on more computational and algorithmic aspects of the latter. Scoville's text [8] is aimed at advanced undergraduates and includes applications in which he has made contributions (e.g., the theory of strong collapses, discrete Lusternik–Schnirelmann category, etc.). The text under review is quite different and reflects Kozlov's combinatorial point of view.

Parts 1 and 2 are about the basics of homology theory, beginning with the simplicial theory, moving through the abstraction of arbitrary chain complexes of modules and the corresponding homological algebra, and then ending with cellular homology. These 140 pages could serve as a thorough introduction to the subject for beginning topology students. Part 3 concerns basic discrete Morse theory, given completely from the point of view of acyclic matchings. Indeed, discrete Morse functions receive mention only in Section 14.2, which is a mere two pages long. This is not a criticism; from a combinatorial and algorithmic point of view, one is most often concerned with constructing discrete gradient vector fields, and so the author's choice to present the theory in this way is entirely appropriate.

Part 4 discusses extensions of discrete Morse theory, which include many of Kozlov's own contributions to the subject. One such topic is *algebraic discrete Morse theory*, in which one produces acyclic matchings in an arbitrary chain complex and reduces its homology to that of a smaller "Morse" complex. Discrete Morse theory for posets makes an appearance in this part as well; there is a rich theory here which is only just hinted at in this book. The final chapter includes a brief introduction to the use of discrete Morse theory for computing the persistent homology of filtered complexes; as mentioned above this is an area of active research in applied topology.

Kozlov has given us a book that fills a gap in the literature. Most mathematicians who use discrete Morse theory have settled on the acyclic matching point of view, and until now these ideas were scattered in various papers. He has brought it all

together into a unified discussion that provides a self-contained introduction to the subject.

REFERENCES

- [1] U. Bauer, *Ripsner: efficient computation of Vietoris-Rips barcodes*, preprint (2019), [arXiv:1908.02518](https://arxiv.org/abs/1908.02518).
- [2] R. Forman, *Morse theory for cell complexes*, Adv. Math. **134** (1998), no. 1, 90–145, DOI 10.1006/aima.1997.1650. MR1612391
- [3] M. Jungerman and G. Ringel, *Minimal triangulations on orientable surfaces*, Acta Math. **145** (1980), no. 1-2, 121–154, DOI 10.1007/BF02414187. MR586595
- [4] H. King, K. Knudson, and N. Mramor, *Generating discrete Morse functions from point data*, Experiment. Math. **14** (2005), no. 4, 435–444. MR2193806
- [5] K. P. Knudson, *Morse theory: Smooth and discrete*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2015. MR3379451
- [6] J. Milnor, *Morse theory*, Based on lecture notes by M. Spivak and R. Wells. Annals of Mathematics Studies, No. 51, Princeton University Press, Princeton, N.J., 1963. MR0163331
- [7] V. Nanda, *Perseus, the persistent homology software*, <http://www.sas.upenn.edu/~vnanda/perseus>, accessed 02/01/2021.
- [8] N. A. Scoville, *Discrete Morse theory*, Student Mathematical Library, vol. 90, American Mathematical Society, Providence, RI, 2019. MR3970274
- [9] J. H. C. Whitehead, *Simple homotopy types*, Amer. J. Math. **72** (1950), 1–57, DOI 10.2307/2372133. MR35437

KEVIN P. KNUDSON

DEPARTMENT OF MATHEMATICS

UNIVERSITY OF FLORIDA

P.O. BOX 118105

GAINESVILLE, FLORIDA 32611

Email address: kknudson@ufl.edu