Tracking features in image sequences using discrete Morse functions

Gregor Jerše and Neža Mramor Kosta

University of Ljubljana, Department of Computer and Information Science, Ljubljana, Slovenia
gregor.jerse@fri.uni-lj.si, neza.mramor@fri.uni-lj.si

Abstract The goal of this contribution is to present an application of discrete Morse theory to tracking features in image sequences. The proposed algorithm can be used for tracking moving figures in a filmed scene, for tracking moving particles, as well as for detecting canals in a CT scan of the head, or similar features in other types of data. The underlying idea which is used is the parametric discrete Morse theory presented in [13], where an algorithm for constructing the bifurcation diagram of a discrete family of discrete Morse functions was given. The original algorithm is improved here for the specific purpose of tracking features in images and other types of data, in order to produce more realistic results and eliminate irregularities which appear as a result of noise and excess details in the data.

1 Introduction

In this paper we would like to present an application of a parametric version of discrete Morse theory to image analysis. In [13], an application of discrete Morse theory was proposed, which results in an algorithm for constructing a discrete version of the bifurcation diagram of a1-parametric family of Morse functions. In the smooth case, a 1-parametric family $f_t\colon M \to \mathbb{R}$, $t \in I$, where $I \subset \mathbb{R}$ is an interval, of smooth functions on a smooth manifold M consists in the generic case of Morse function for all values of t, except maybe finitely many values $t_1, \ldots, t_k$, where bifurcations occur. The changes of the critical points of the functions f_t as t increases are traced in the bifurcation diagram, where merges and deaths of critical points, as well as births of new ones occur precisely in the bifurcation points $t_i, i = 1, \ldots, k$.

This phenomenon has potential applications in many fields, in particular it is useful in topological data analysis and image analysis. Critical points of the functions f_t often correspond to important features in the data which we would like to trace as t increases. For example in grayscale images, a figure corresponds to a local maximum of the grayscale function, a white feature within a dark area corresponds to a local minimum, and a saddle of the grayscale function can come for example from a joint connecting to bigger elements in a large construction or to a point where two lighted areas in a dark scene meet. In a sequence of images, figures and other features can thus be traced by following the critical points in the bifurcation diagram of the grayscale function. In the case of topological data analysis, values of the functions f_t are given only in a specific (usually finite) set of points from the domain M, or the domain itself is discrete, as in the case of digital images where the grayscale function associates a value to each pixel. In addition, only a finite sequence of function f_{t_i} is considered. The smooth phenomenon described above can therefore not be applied directly.

There are many approaches to extending Morse theory to the discrete setting, in particular for the purpose of topological data analysis, for example in [6], [5], [4], [2], or [17]. Among them the discrete Morse functions of Robin Forman [8] , [9] are well established and have been successfully used in many domains. A discrete Morse function is defined on a cellular complex. In the case of digital images, the natural cell structure is that of a cubical complex, but in the case of general data a cell decomposition of the domain (for example a triangulation) must first be constructed. The

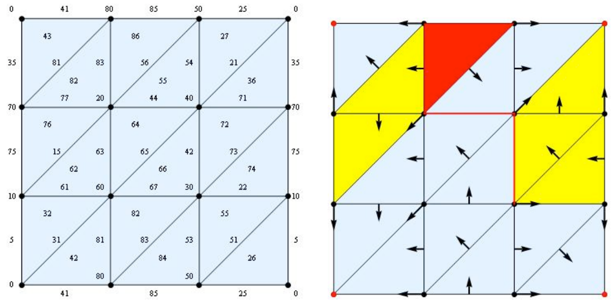

Figure 1: A discrete Morse function on the torus, its discrete vector field V, and a V-path

function values, which are initially given only in the vertices of the complex must then be extended to a discrete Morse function on the whole domain, which can be done using the algorithm of [12]. Morse functions obtained from data, for example from grayscale images, can have very many critical points corresponding to unwanted (or unimportant) details or to noise, and an additional advantage of Forman's discrete Morse theory is that it contains a simple procedure for discarding access critical points, similar to canceling pairs of critical points in smooth Morse theory. By introducing a persistence level p, and canceling pairs of critical points with values differing by less than p, unwanted details as well as noise can be reduced. As p increases, only the critical points which correspond to the relevant features persist.

An algorithm for constructing the bifurcation diagram of a 1-parametric family of discrete Morse functions has been proposed in [13]. A drawback of this algorithm in the context of figure and feature tracking is that, since canceling is performed at each level t separately, the resulting paths in the bifurcation diagram tend to appear discontinuous. For example, a figure in an image is a dark patch on a lighter background with similar grayscale values, on which the grayscale function is close to constant. As a result, many of the cells which form the region are critical and, after canceling, they all collapse to a single critical point. As t varies and the figure moves this surviving critical cell may jump within the region.

In this contribution we first review the basics on Forman's discrete Morse functions, describe the parametric version introduced in [13] and shortly review the algorithm for reconstructing the bifurcation diagram of a 1-parametric family of discrete Morse functions (section 2). We then discuss a canceling strategy, suitable for feature tracking in image analysis (section 3). Finally, applications and further work are discussed (section 4).

2 Discrete Morse functions

A discrete Morse function, as defined by Robin Forman in [8], [9], is a function $F\colon M \to \mathbb{R}$ defined on a regular cell complex M a regular complex M that associates a value $F(\sigma) \in \mathbb{R}$ to each cell $\sigma \in M$ such that F increases with dimension, except possibly in one direction: for every $\sigma^k \in M$

- $F(\tau^{k-1}) \geq F(\sigma^k)$ for at most one face $\tau < \sigma$
- $F(\tau^{k+1}) \leq F(\sigma^k)$ for at most one coface $\tau > \sigma$.

The pairs (σ^k, τ^{k+1}), such that $\sigma < \tau$ and $F(\sigma) > F(\tau)$ contain the *regular cells.* The unpaired cells are *critical* with index equal to the dimension of the cell.

For the purpose of presentation, a regular pair (σ^k, τ^{k+1}) is often denoted by an arrow pointing in the direction of function descent, as in [10], and the collection of all regular pairs forms the *discrete gradient vector field* V associated to the function F. A sequence of regular pairs forming a path along which the function values decrease is a V-path. It is a well known result of Forman, that a partial pairing V on a finite regular cell complex M is the discrete gradient vector field of a

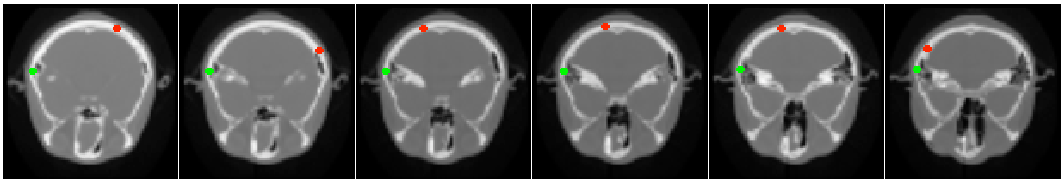

Figure 2: The remaining critical point corresponding to the top part of the scull after canceling, represented by a red point, jumps from slice to slice

discrete Morse function if and only if its V-paths contain no cycles. The discrete gradient vector field captures the qualitative properties of the function F.

A parametric discrete Morse function is a family F_{t_i}, $i = 0, \ldots, n$ of discrete Morse functions on a M. In [13] an algorithm for connecting the critical cells of the functions F_{t_i} among the slices t_i, and constructing the bifurcation diagram of such a family is given. The basic idea of the algorithm is to extend the discrete vector fields of the Morse functions F_{t_i} slice by slice to a discrete vector field on $M \times I$ in such a way, that all critical cells in slice t_i are paired backwards into the strip $M \times [t_{i-1}, t_i]$. In the extended discrete vector field there are no critical cells, except the initial ones in slice t_0. A critical cell τ in $M \times \{t_i\}$ is then connected to a critical cell τ' in $M \times \{t_{i-1}\}$ in the bifurcation diagram if there is a V-path in the extended discrete vector field connecting them. It can be shown [13] that if F_{t_i} are good enough cellular approximations of smooth functions on $|M|$, belonging to a smooth generic 1-parametric family, then the bifurcation diagram obtained in this way introduces no new unwanted connections (although some connections may be lost).

3 Canceling

In smooth Morse theory two critical cells of consecutive indices that are connected by only one gradient path may be canceled, producing a simpler Morse function with fewer critical points (this can be illustrated as straightening out unnecessary bumps in the graph). In discrete Morse theory, canceling is implemented by a very simple process: a pair of critical cells of consecutive dimensions, connected by only one V-path, is canceled simply by reversing all arrows along this single path. Since the new discrete gradient vector field has no cycles (due to the condition of a single V-path), it belongs to a discrete Morse function. Canceling is done up to a given persistence level p, that is, pairs of critical cells satisfying the necessary condition on V-paths can be canceled only in values differ by less than p.

In the case of image analysis, there are several reasons for canceling. The first is noise, which is generally present. Because of this, canceling up to a persistence level corresponding to the estimated noise level is a good idea. The second reason for canceling is the presence of larger regions with a uniform shade of gray, where the grayscale function is almost constant. Such regions contain many critical cells (in the extreme case all the cells can be critical), which all correspond to the same feature in the image.

For example, Figure 3 shows a sequence of CT scans of the head. As a result of canceling, only one critical point corresponds to the whole top part of the scull, but, since the canceling is performed separately in each slice, this critical point jumps across the arch of the scull from slice to slice.

This implies that in the task of feature tracking in image sequences it is reasonable to cancel critical cells in each slice separately, in order to reduce the number of critical cells, if possible, to the number of features that are of interest. As mentioned above, this introduces a new problem – the critical cell that is traced in the bifurcation diagram tends to jump around a uniform region as it moves from slice to slice, as in Figure 3. In order to eliminate this problem, we store for each critical cell τ which survives after canceling in some slice t_i also all the critical cells that have been canceled in the process and, in the end, belong to the descending disk of τ (that is,

Figure 3: Tracking the path of a skier, the smaller square denotes the original critical points after canceling, and the larger square the path resulting from the modified algorithm

are connected to τ by a V-path beginning in the boundary of τ, [11]). All these could just as well represent the same feature as τ if the order of canceling had been different. In other words, instead of representing a feature in the slice t_i by the one surviving critical cell τ, we represent it by a collection $V(\tau)$ of V-paths beginning in the boundary of τ.

After the bifurcation diagram is constructed and τ is connected to critical cells in other slices, a path $\tau = \tau_0, \tau_1, \tau_2, \ldots, \tau_n$ is optimized in sense that it seems to describe the smooth movement. A representative r_i for τ_i is chosen from the set of cells $V(\tau_i)$ for each time slice i using algorithm based on a local optimization. First we define $r_i = \tau_i$ for each i. Next we move through the path several times and at each step optimize the choice of the representative r_i as follows: fixing r_{i-1} and r_{i+1} the representative r_i is chosen such that $\max(D(r_{i-1}, r_i), D(r_i, r_{i+1}))$ is minimal. Here $D(x, y)$ denotes the Euclidean distance between points x and y. In other words the algorithm minimizes the maximal distance between the representative r_i and its neighbors, which creates a sense of a smooth movement.

4 Applications and further work

We have tested the algorithm on image sequences from three different domains. The original motivation for this work came from an attempt to use the algorithm of [13] for tracking a moving figure in a movie. The path produced by the algorithm of [13] was irregular, it tended to jump from point to point within the area representing the figure and even its surroundings. Figure 4 shows the results of the original and the modified algorithm applied to a movie showing a skier skiing down a mountain face. The original critical point corresponding to the figure of the skier produced by the original algorithm of [13] is framed by a smaller red square, and the path produced by the modified algorithm is framed by a larger white square. In general, the path new path is a much better approximation of the skiers path down the slope.

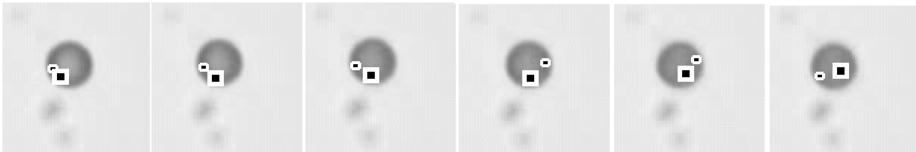

Figure 4: Tracking the path of a superparamagnetized colloid particles in water controlled by an external magnetic field, the smaller circle denotes the original critical points after canceling, and the larger square the path resulting from the modified algorithm

The second domain is a series of movies taken in experiments conducted at the at the Department of Complex Matter at the Josef Stefan Institute in Ljubljana, Slovenia, in the context of a research leading to simulating the dynamics of atoms on surfaces [14]. The algorithm of [13] was applied to a sequence of images taken by an optical microscope showing the interaction of superparamagnetized colloid particles in water controlled by an external magnetic field. On figure 4 the moving particle is tracked using the original (smaller white circle) and modified (big white square) algorithm.

The third image domain the sequence of CT scans of the head, a part of which is shown on Figure 3. The algorithm was used to detect and follow canals and cavities within the scull. The green dot in the figure shows the path of the single remaining critical point representing the top part of the scull, produced by the modified algorithm.

The algorithm was also used in artificial intelligence in machine learning where a robot was taught the concept of occlusion by learning from bifurcation diagrams arising from images taken by an onboard camera [15],[16].

We have not addressed the question of the persistence level p systematically. In particular, persistence barcodes or diagrams [1], [7], could be used as a mechanism for determining a suitable choice of persistence level. In this context, 1-parametric families of persistence diagrams have been introduced [3] which offer a framework for constructing discrete bifurcation diagrams.

References

[1] G. Carlsson, A. Collins, L. J. Guibas, A. Zomorodian, *Persistence Barcodes for Shapes*. International Journal of Shape Modeling, 2005. (ps.gz) doi:10.1142/S0218654305000761 Symposium on Geometry Processing, Nice, France, 2004. (ps.gz)

[2] F. Cazals, F. Chazal, T. Lewiner, *Molecular shape analysis based upon Morse-Smale complex and the Connolly function*, in: 19th ACM Symposium on Computational Geometry, 2003, pp. 351–360.

[3] D. Cohen-Steiner, H.Edelsbrunner, D. Morozov, *Vines and vineyards by updating persistence in linear time*, In Proc. 22nd Ann. Sympos. Comput. Geom., 2006, 119–126.

[4] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, B. Hamann, *Volumetric analysis using Morse-Smale complexes*, in: Intl. Conference on Shape Modeling and Applications (SMI), 2005, pp. 320–325.

[5] H. Edelsbrunner, J. Harer, V. Natarajan, V. Pascucci, *Morse-Smale complexes for piecewise linear 3-manifolds*, in: 19th ACM Symposium on Computational Geometry, 2003, pp. 361–370.

[6] H. Edelsbrunner, J. Harer, A. Zomorodian, *Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds*, Discrete Comput. Geom. 30 (2003) 87–107.

[7] H. Edelsbrunner, D. Letscher, A. Zomorodian, *Topological Persistence and Simplification*, Discrete Comput. Geom. 28 (2002), 511–533.

[8] R. Forman, *A Discrete Morse Theory for Cell Complexes Geometry*, Topology and Physics for Raoul Bott, S.T. Yau, ed., Int'l Press, 1995.

[9] R. Forman, *Morse Theory for Cell Complexes*, Advances in Math., vol. 134, pp. 90-145, 1998.

[10] R. Forman, *A user's guide to discrete Morse theory*, Sém. Lothar. Combin. 48 (2002) Art. B48c, 35 pp. (electronic).

[11] G. Jerše and N. Mramor Kosta, *Ascending and descending regions of a discrete morse function*, Computational Geometry 42 (2009), no. 6-7, 639 651.

[12] H. King, K. Knudson, N. Mramor Kosta, Generating discrete Morse functions from point data, Exp. math. 14 (4) (2005) 435–444.

[13] H. King, K. Knudson, N. Mramor Kosta, *Birth and death in discrete Morse theory*, available at arXiv:0808.0051v1 (2007).

[14] M. Vilfan, A. Poto?nik, B. Kavčič, N. Osterman, I. Poberaj, A. Vilfan, and D. Babič, Self-assembled artificial cilia, Proc. Natl. Acad. Sci. (PNAS) 107, 1844 (2010).

[15] J. Žabkar, G. Jerše, N. Mramor Kosta, I. Bratko, *Induction of qualitative models using discrete Morse theory*, in: C. Price (Ed.), Proceedings of the 21st Annual Workshop on Qualitative Reasoning, 2007, pp. 203–208.

[16] J. Žabkar, I. Bratko, G. Jerše, J. Prankl, M. Schlemmer, *Learning qualitative models from image sequences*, in: E. Bradley, L. Travé-Massuyès (Eds.), 22nd International Workshop on Qualitative Reasoning, 2008, pp. 146–149.

[17] Afra Zomorodian, *Topology for Computing*, Cambridge Monographs on Applied and Computational Mathematics (No. 16)