MORSE THEORY AND PERSISTENT HOMOLOGY FOR TOPOLOGICAL ANALYSIS OF 3D IMAGES OF COMPLEX MATERIALS

Olaf Delgado-Friedrichs, Vanessa Robins, Adrian Sheppard

Department of Applied Mathematics
Research School of Physics and Engineering
The Australian National University, Canberra ACT 0200.

ABSTRACT

We develop topologically accurate and compatible definitions for the skeleton and watershed segmentation of a 3D digital object that are computed by a single algorithm. These definitions are based on a discrete gradient vector field derived from a signed distance transform. This gradient vector field is amenable to topological analysis and simplification via Forman's discrete Morse theory and provides a filtration that can be used as input to persistent homology algorithms. Efficient implementations allow us to process large-scale x-ray micro-CT data of rock cores and other materials.

***Index Terms*—** Discrete Morse theory, Skeletonisation, Watershed transform, Topological data analysis

1. INTRODUCTION

3D imaging technology, notably X-ray micro-CT, is now able to routinely capture the structure inside complex materials such as porous rocks, bones and manufactured materials with excellent fidelity and resolution. The need to study these data in a quantitative manner requires efficient algorithms for processing large 3D images and for extracting topological and geometrical measures.

One of the most fundamental descriptions of structure is via *homology*: the mathematical characterisation of connectivity, including connected components, independent loops and enclosed voids. An essential ingredient in using homology in the study of experimental and computational data to is to build a filtration (a nested sequence of cell-complexes) that captures the topology of the data with respect to a parameter, usually a length-scale, that dictates the order in which cells are added. Topological features (such as a hole through an object) are born at some parameter value and are later merged or filled in at a larger value. Features that are created and destroyed almost simultaneously are considered noise while features that persist over longer parameter ranges are deemed to be more important. These lifetimes are encoded by persistent homology; see [1] for an overview.

Research supported in part by ARC Discovery Project DP110102888.

For a suitably well-behaved real-valued function f defined on a compact manifold, such a filtration can be derived from the lower level cuts (i.e. sub-level sets) of f: $\mathcal{L}_f(c) = \{x \mid f(x) \leq c\}$. Morse theory tells us that the topology of these sub-level sets can only change when c passes a critical value of f (i.e., at a local minimum, saddle point or maximum) [1]. The critical points and gradient flow-lines between them are used to define an abstract cell complex with a filtration ordering parameter (in this case c). The persistent homology of this filtration can be shown to be the same as the sequence of the sub-level sets.

There have been many applications of Morse theory to digital image analysis, see the review [2], for example. We model greyscale digital images as real-valued functions on the vertices of a *cubical complex* [3]. In our setting, the voxels $(i, j, k) \in \mathcal{D} \subset \mathbb{Z}^3$ are the vertices (0-cells) of the complex. Higher-dimensional cells are the unit edges (1-cells) between adjacent voxels, unit squares (2-cells), and unit cubes (3-cells). We initially extend a greyscale function $g : \mathcal{D} \to \mathbb{R}$ to the full cubical complex $\mathcal{K}$ by taking the maximal value from the vertices of a cell:

$$g(\beta) := \max\{g(\alpha) \mid \alpha < \beta, \alpha \in \mathcal{D}\}. \tag{1}$$

Here the notation $\alpha < \beta$ means that α is a *face* of β and that β is a *coface* of α. If the dimensions of α and β differ by one then we say that α is a *facet* of β.

Given any function defined on a complex $f : \mathcal{K} \to \mathbb{R}$ the *lower level subcomplex* is

$$\mathcal{K}_f(c) := \{\alpha \mid \alpha \leq \beta \text{ for } f(\beta) \leq c\}.$$

The connectivity of the cubical level subcomplex for the maximal-vertex function g is equivalent to the standard 6-neighbourhood of 3D digital topology [4].

An elegant and effective approach to extending Morse theory to functions on cell complexes is provided by Robin Forman's *discrete Morse theory* [5]. He defines a discrete Morse function on a cell complex so that a critical point of index i in the classical theory becomes a critical *cell* of dimension i. Gradient flow lines between critical points of index $(i + 1)$ and i become paths of (cell, cofacet) pairs in a discrete vector

978-1-4799-5751-4/14/$31.00 ©2014 IEEE

field, these are called *V-paths*. A *discrete vector field* pairs adjacent cells with dimensions i and $i+1$ so that no cell in the complex appears in more than one pair. Cells that are unpaired are critical. In a *gradient* vector field there are no loops. A discrete Morse function $f : \mathcal{K} \to \mathbb{R}$ is such that (α, β) form a (cell, cofacet) pair in a gradient vector field if and only if $f(\alpha) \geq f(\beta)$. The maximal-vertex function g defined in (1) does not fulfil the requirements for a discrete Morse function. Methods for extending the voxel-values to a valid discrete Morse function on the full cubical complex have recently emerged [6][7]. These methods build on work for simplicial meshes [8] and extend voxel classification schemes from digital topology[9].

Using the above framework, we can show that the closely related techniques of discrete Morse theory and persistent homology provide a rigorous framework for generalising and unifying various traditional tools from image analysis such as skeletonization by thinning [10, 11, 12, 13], and the watershed transform [14, 15, 16, 17]. The connections between discrete Morse theory (or simple homotopy), persistent homology, and the use of stable and unstable manifolds to partition scalar functions on cell complexes have also been explored recently by other authors [18, 19, 20, 21, 22, 7].

This paper is laid out as follows. In section 2 we provide a new definition for the homotopic skeleton of a 3D digital object, which emerges as a subcomplex of the Morse complex and can therefore be computed without explicit thinning. In Section 3 it is shown that one can also define a watershed segmentation (or partitioning) as part of the Morse extraction procedure. In Section 4 we discuss the application of persistent homology for the simplification of the resulting structures, while in Section 5 the techniques are applied to x-ray tomographic imaging data.

2. THE MORSE SKELETON

Throughout this section, let $\mathcal{K}$ be a cubical complex (e.g. as derived from a 3D digital image) with discrete Morse function $f\colon \mathcal{K} \to \mathbb{R}$ and associated gradient vector field V, defined by $V(\alpha) = \beta$ if α is the single facet of β such that $f(\alpha) \geq f(\beta)$. We also write $V(\gamma) = 0$ when γ is critical or when $\gamma = V(\alpha)$ for some facet α. We assume without loss of generality that the values of f are unique on the cells of $\mathcal{K}$.

Let α be a critical p-cell in $\mathcal{K}$. We use V-paths that end or start at α to define its stable or unstable sets respectively. Formally, the *stable set* of α, denoted $\mathcal{S}_\mathcal{K}(\alpha)$, is the smallest set of p-cells in $\mathcal{K}$ such that $\alpha \in \mathcal{S}_\mathcal{K}(\alpha)$ and

$$V(\delta^{(p)}) = \gamma^{(p+1)} > \beta^{(p)} \in \mathcal{S}_\mathcal{K}(\alpha) \text{ implies } \delta \in \mathcal{S}_\mathcal{K}(\alpha).$$

Similarly, the *unstable set* of α, denoted $\mathcal{U}_\mathcal{K}(\alpha)$, is the smallest set of p-cells such that $\alpha \in \mathcal{U}_\mathcal{K}(\alpha)$ and

$$V^{-1}(\delta^{(p)}) = \gamma^{(p-1)} < \beta^{(p)} \in \mathcal{U}_\mathcal{K}(\alpha) \text{ implies } \delta \in \mathcal{U}_\mathcal{K}(\alpha).$$

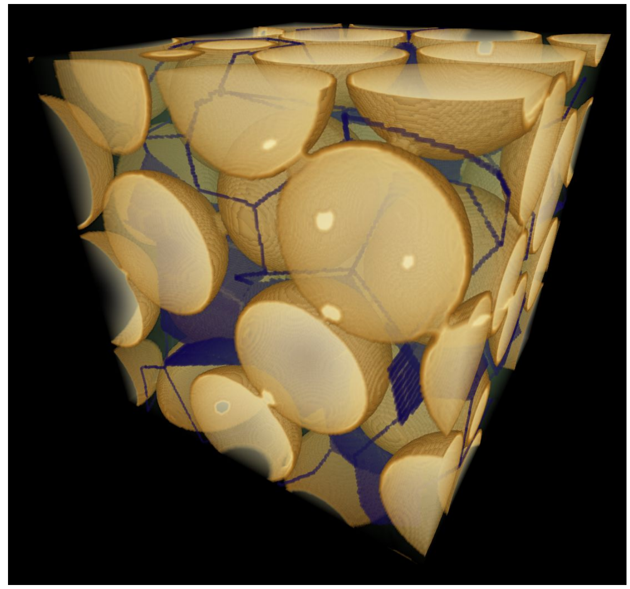

Fig. 1. The Morse skeleton (in blue) generated from the signed Euclidean distance transform of the binary image of a sphere pack, where the pore space is the foreground and takes negative image values. Where critical 2-cells are in the pore space their unstable sets are shown as blue skeleton patches; a line skeleton, lacking such patches, would poorly represent the geometry.

Note that the stable and unstable sets of α contain cells only of dimension p. It is useful therefore, to also define an *unstable complex* of α as the closure of its unstable set.

$$\mathcal{W}_\mathcal{K}(\alpha) := \{\gamma \mid \gamma \leq \delta \in \mathcal{U}_\mathcal{K}(\alpha)\}.$$

We do not make a dual definition of stable complex as it would most naturally involve the cofaces of elements in $\mathcal{S}_\mathcal{K}(\alpha)$ (rather than faces) and therefore not be a subcomplex of $\mathcal{K}$.

The *Morse skeleton* $\mathcal{A}_\mathcal{K}$ is now defined to be the subcomplex built from unstable complexes of critical cells,

$$\mathcal{A}_\mathcal{K} := \bigcup_{\alpha \text{ critical}} \mathcal{W}_\mathcal{K}(\alpha).$$

When $\mathcal{K}$ and f are derived from a greyscale digital image on a rectangular domain, the skeleton is effectively the entire complex; it is when we examine the level *sub*complexes that the Morse skeleton gives us an interesting summary of structure. It can be shown that

$$\mathcal{A}_{\mathcal{K}(c)} = \bigcup_{\substack{\alpha \text{ critical} \\ f(\alpha) \leq c}} \mathcal{W}_{\mathcal{K}(c)}(\alpha) = \bigcup_{\substack{\alpha \text{ critical} \\ f(\alpha) \leq c}} \mathcal{W}_\mathcal{K}(\alpha).$$

This means we need only determine the unstable complexes of critical points with respect to the complex $\mathcal{K}$, there is no need to recompute for different level subcomplexes. The

Morse skeleton for any choice of threshold is then given by the unstable complexes whose critical value is less than the threshold.

We are also able to show that the Morse skeleton, $\mathcal{A}_{\mathcal{K}(c)}$ is homotopy equivalent to $\mathcal{K}(c)$ via a "regular collapse" along the V-paths. This regular collapse is analogous to a thinning process in digital image analysis [10]. The use of discrete Morse theory means that the collapse (or thinning) does not have to be performed explicitly and there are no ambiguities in its construction. Geometric properties of the skeleton (such as being centred with respect to the object boundary) depend on the initial function which is commonly a signed distance transform of a binary image [23]; see Figures 1 and 2 for illustration.

3. BASINS AND BRIDGES

Having used unstable sets to define the Morse skeleton, we now investigate the role of stable sets in partitioning the vertex set of the complex $\mathcal{K}$ with discrete Morse function f into regions similar to those derived by watershed algorithms. We have already noted that in general, the stable set of a critical cell cannot be extended to a subcomplex in a manner dual to the unstable complex. But for critical 0-cells, i.e. minima, we can define the *basin* $\mathcal{B}_{\mathcal{K}}(\alpha)$ as the maximal subcomplex of $\mathcal{K}$ that has a regular collapse onto α. It is not hard to show that each vertex of $\mathcal{K}$ belongs to the stable set (and therefore the basin) of exactly one local minimum. The same is not true of k-cells with $k \geq 1$. We call a 1-cell that is not contained in any basin a *bridge*. If β is a bridge, then we must have $V^{-1}(\beta) = 0$, so either β is critical (a 1-saddle), or $V(\beta) \neq 0$. When β is a bridge between the basins of two minima $\mathcal{B}(\alpha)$ and $\mathcal{B}(\gamma)$, the V-paths that start at each vertex of β form a path between α and γ that we call the *canonical path* defined by β. If β is a critical 1-cell, this path is the unstable complex of β and forms part of the Morse skeleton of $\mathcal{K}$, so we have established that when two minima are in the unstable complex of a 1-saddle, their basins are adjacent at that 1-saddle. The converse is not true: two basins may be adjacent when there is no critical 1-cell bridging their minima; the basins labelled 'A' and 'C' in Fig. 2 illustrate this. In this case the canonical path of a non-critical bridge captures the additional connectivity between minima. This is valuable for many applications where region adjacency is of physical relevance.

4. SIMPLIFICATION

A naive application of the watershed partition to a real image (even one with small-amplitude noise) generally creates more basins than necessary to effectively summarize the structures present. Many strategies have been proposed to reduce this effect, with marker-based watersheds [24, 17] proving highly successful. Medial axis skeletons are also typically pruned

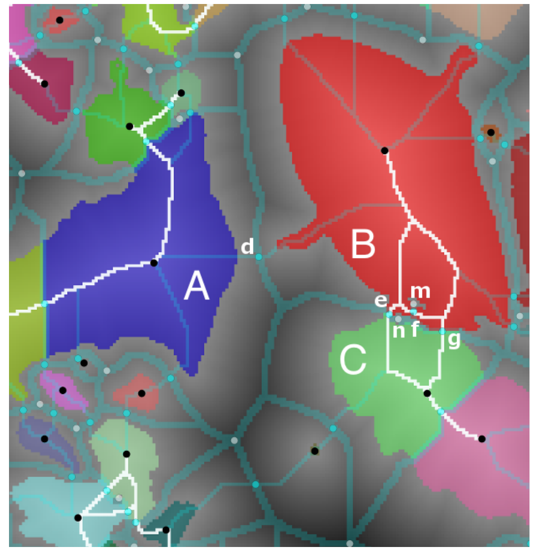

Fig. 2. A small section of Morse skeleton and partition derived from the signed Euclidean distance transform of a 2D binary image. The discrete Morse gradient vector field is calculated, simplified to a persistence threshold of 1 voxel unit and used to simultaneously find the basins and skeleton of the foreground, i.e. the lower level cut at threshold 0. The coloured regions show a partitioning of the foreground via basins of the local minima (black dots). Basin boundaries are marked by thick blue-grey lines that lie along V-paths between maxima (grey dots) and saddles (blue dots). Basin **B** is connected to basin **A** by a critical bridge (the saddle point **d**) and to basin **C** by two critical bridges (saddles **e** and **g**), but **A** and **C** have only non-critical bridges between them. Note the intricate structure associated with saddle **f** and maximum **m**. The basin **B** and the coloured pore are both topological disks because the V-path between **f** and **m** contains bridges that make a cut.

to remove artefacts that result from noise or surface roughness [23].

Naturally these issues also occur when we construct the Morse basins and skeleton of a real image, resulting in features that have a low importance. The Morse theory approach, which yields a complete filtration, has two advantages when undertaking this simplification exercise. Firstly, we have additional information, from the fact that the complex contains cells of all dimensions from 0 to 3, and the fact that we have a compatible skeletonisation and partitioning. Secondly, we can use concepts from persistent homology that can be applied when a complete filtration is present [25]. As we saw in the previous section, adjacent basins may or may not be connected via a 1-saddle, and identifying this is crucial to merging basins in a way that preserves important topological information as measured by persistent homology. Discrete Morse theory has built-in techniques for cancelling critical points by reversing pairs along a V-path and the natural connection with homology means we can simplify the topology of the level sets in a controlled way [19]. Morse theory also incorporates higher-order topological changes such as filling in extraneous loops and voids with the same techniques. We find that when working with distance transforms, cancellation of critical point pairs with a persistence of less than 1 voxel unit provides a sufficient level of simplification while maintaining a high level of detail; see Fig. 2.

Simplification strategies are the subject of ongoing research. For example, it is possible to incorporate measures of feature size similar to those in [17] within the persistence framework. The features we can measure go beyond basins of local minima to include higher-dimensional structures defined by the 1- and 2-saddles. This allows us to discriminate between convex and concave shapes of various classes. For example the distance transform of a barbell shape is characterised by two local minima joined by a 1-saddle while that of a squashed ball (similar to a red blood cell) must contain a 2-saddle and one or more local minima and 1-saddles. By examining the persistence of the 1-cycle and the size of the unstable set of the 2-saddle that it is paired with, we can decide whether to cancel or preserve this feature.

5. APPLICATION TO POROUS MATERIALS

Starting with a grayscale 3D image of a porous material, we first classify voxels as either 'pore' (foreground) or 'grain' (background) thus obtaining a binary image. To study the geometry defined by this image it is natural to apply the signed Euclidean distance transform (SEDT), so that the greyscale function g encodes the distance of each voxel from the boundary between foreground and background voxels. By making the distance negative in the foreground and positive in the background, we see that the Morse skeleton at threshold $c = 0$ defines a medial surface of the foreground and that the basins partition the foreground in a dual fashion. The Morse skeleton and basins now provide *simultaneous* and *compatible* methods for describing the foreground geometry. Figure 1 shows the Morse skeleton of the pore space of a packing of glass spheres, computed from a micron-scale x-ray micro-CT image as described above. To illustrate both the skeleton and basins, Fig. 2 shows a small subset of a 2D dataset: a single slice taken from a micro-CT image of a Mt Gambier limestone core.

Persistent homology of SEDTs can help describe the morphology of a porous material in a variety of ways including pore- and grain-size distributions. As a further example, in Fig. 3 we show the 1-dimensional persistence diagrams [1] for the pore-space of four granular samples that are publicly available (xct.anu.edu.au/network_comparison). In general, 1D homology encodes the independent cycles (loops) in a cell complex. In the Morse filtration of an SEDT, a cycle is born when a 1-saddle is added between two minima that already

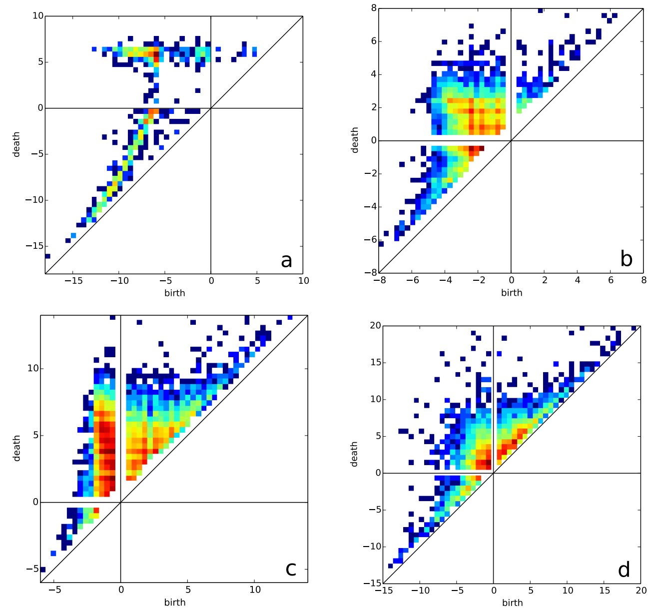

Fig. 3. Persistence diagrams (represented as 2D histograms of persistence pairs) for 1D cycles in the pore-space of four samples: (a) mono-disperse spherical bead pack, (b) polydisperse unconsolidated sand, (c) well-consolidated Castlegate sandstone, (d) fossiliferous Mt Gambier limestone.

belong to the same connected component and that cycle dies when a 2-saddle is added turning the 1-cycle into a boundary. Cycles that have birth values < 0 are loops in the pore space; those that also have death value < 0 correspond to 2D patches in the Morse skeleton, cf. Fig. 1. From the data in Fig. 3, we can see that only the pore-space of the sandstone (c) will be well-represented by a linear skeleton, the other samples have a large proportion of persistent 2-saddles lying in the pore-space. The consolidation of grains can also be seen in these diagrams as the proportion of pairs that have birth and death values above 0.

Aside from their theoretical interest, we see the results of this paper as being of practical value for two reasons. Firstly, a consistent skeleton and partition, along with tools from the emerging field of persistent homology, provide a rich set of measures for characterization. Secondly, simultaneous partitioning and skeletonisation provides just the information required to construct a *representative pore-network*. Pore-network models are used in the simulation of fluid transport in porous media, and particularly the flow of multiple immiscible fluids. The algorithms and objects defined here will allow the generation of such network models with greater geometrical and topological fidelity. We emphasise that the algorithms to perform this Morse analysis on 3D images, based on that in [6], are of genuine practical value since they can be implemented efficiently and on distributed architectures.

6. REFERENCES

[1] H. Edelsbrunner and J. Harer, *Computational Topology: An introduction*, American Mathematical Society, Providence, Rhode Island, 2010.

[2] S. Biasotti, L. De Floriani, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, L. Papaleo, and M. Spagnuolo, "Describing shapes by geometrical-topological properties of real functions," *ACM Comput. Surv.*, vol. 40, no. 4, pp. 187, 2008.

[3] V. A. Kovalevsky, "Finite topology as applied to image analysis," *Comput. Vision Graph. Image Process.*, vol. 46, no. 2, pp. 141–161, 1989.

[4] T.Y. Kong and A. Rosenfeld, "Digital topology: introduction and survey," *CVGIP*, vol. 48, no. 3, pp. 357–393, 1989.

[5] R. Forman, "Morse theory for cell complexes," *Advances in Mathematics*, vol. 134, pp. 90–145, 1998.

[6] V. Robins, P. Wood, and A. Sheppard, "Theory and algorithms for constructing discrete morse complexes from grayscale digital images," *IEEE T. Pattern Anal.*, vol. 33, no. 8, pp. 16461658, 2011.

[7] A. Gyulassy, P. T Bremer, B. Hamann, and V. Pascucci, "A practical approach to morse-smale complex computation: Scalability and generality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 16191626, 2008.

[8] T. Lewiner, H. Lopes, and G. Tavares, "Applications of forman's discrete morse theory to topology visualization and mesh compression," *IEEE Transactions on Visualization and Computer Graphics*, vol. 10, no. 5, pp. 499–508, Sept. 2004.

[9] M. Couprie and G. Bertrand, "New characterizations of simple points in 2D, 3D and 4D discrete spaces," *IEEE T. Pattern Anal.*, vol. 31, no. 4, pp. 637648, 2009.

[10] D. G. Morgenthaler, *Three-dimensional simple points: serial erosion, parallel thinning, and skeletonization*, University of Maryland, 1981.

[11] G. Malandain, G. Bertrand, and N. Ayache, "Topological segmentation of discrete surfaces," *Int. J. Comp. Vis.*, vol. 10, no. 2, pp. 183–197, Apr. 1993.

[12] K. Siddiqi and S. Pizer, *Medial Representations*, Springer, Berlin Heidelberg, 2008.

[13] S. Biasotti, D. Attali, J.-D. Boissonnat, H. Edelsbrunner, G. Elber, M. Mortara, G. Sanniti di Baja, M. Spagnuolo, M. Tanase, and R. Veltkamp, "Skeletal structures," in *Shape Analysis and Structuring*, pp. 145–183. Springer Berlin Heidelberg, 2008.

[14] S. Beucher and C. Lantejoul, "Use of watersheds in contour detection," in *International Workshop on Image Processing: Real-time Edge and Motion Detection/Estimation, Rennes, France.*, 1979.

[15] L. Vincent and P. Soille, "Watersheds in digital spaces: an efficient algorithm based on immersion simulations," *IEEE T. Pattern Anal.*, vol. 13, no. 6, pp. 583598, 1991.

[16] L. Najman and M. Schmitt, "Watershed of a continuous function," *Signal Processing*, vol. 38, no. 1, pp. 99–112, July 1994.

[17] M. Couprie and G. Bertrand, "Topological gray-scale watershed transformation," *Optical Science, Engineering and Instrumentation '97*, pp. 136–146, Oct. 1997.

[18] J. Cousty, G. Bertrand, M. Couprie, and L. Najman, "Collapses and watersheds in pseudomanifolds of arbitrary dimension," *J. Math. Imaging Vis.*, pp. 1–25, Feb. 2014.

[19] U. Bauer, C. Lange, and M. Wardetzky, "Optimal topological simplification of discrete functions on surfaces," *Discrete Comput. Geom.*, vol. 47, no. 2, pp. 347377, Mar. 2012.

[20] S. Guenther, J. Reininghausen, H. Wagner, and I. Hotz, "Efficient computation of 3D Morse-Smale complexes and persistent homology using discrete morse theory," *The Visual Computer*, vol. 28, pp. 959–969, 2012.

[21] L. Čomić, M. Mesmoudi, and L. De Floriani, "Smale-like decomposition and forman theory for discrete scalar fields," in *Discrete Geometry for Computer Imagery*, number 6607 in Lecture Notes in Computer Science, pp. 477–488. Springer Berlin Heidelberg, Jan. 2011.

[22] G. Jerše and N. Mramor Kosta, "Ascending and descending regions of a discrete morse function," *Computational Geometry*, vol. 42, no. 67, pp. 639–651, Aug. 2009.

[23] C. Arcelli, G. Sanniti di Baja, and L. Serino, "Distance-driven skeletonization in voxel images," *IEEE T. Pattern Anal.*, vol. 33, no. 4, pp. 709720, 2011.

[24] F. Meyer and S. Beucher, "Morphological segmentation," *Journal of Visual Communication and Image Representation*, vol. 1, no. 1, pp. 21–46, Sept. 1990.

[25] H. Edelsbrunner, J. Harer, and A. Zomorodian, "Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds," *Discrete and Computational Geometry*, vol. 30, pp. 98107, 2003.