

Combinatorial Gradient Fields for 2D Images with Empirically Convergent Separatrices

Jan Reininghaus ^{*} David Günther [†] Ingrid Hotz [‡]
Tino Weinkauff [†] Hans-Peter Seidel [†]

September 3, 2012

Abstract

This paper proposes an efficient probabilistic method that computes combinatorial gradient fields for two dimensional image data. In contrast to existing algorithms, this approach yields a geometric Morse-Smale complex that converges almost surely to its continuous counterpart when the image resolution is increased. This approach is motivated using basic ideas from probability theory and builds upon an algorithm from discrete Morse theory with a strong mathematical foundation. While a formal proof is only hinted at, we do provide a thorough numerical evaluation of our method and compare it to established algorithms.

1 Introduction

Computer assisted analysis of two-dimensional image data has become an essential tool in scientific research and industrial applications. To deal with the growing amount of data, automated feature extraction methods are frequently employed in applications from, e.g., medical imaging, geosciences or computer vision. In particular, methods based on computational topology [EH10] are gaining traction due to their ability to robustly extract relevant features of the data.

In this paper, we propose a novel method that extracts the Morse-Smale complex (MS-complex) [Sma61] of a given two-dimensional image. The MS-complex consists of critical points and separatrices. In this setting, the critical points are the local minimum-, saddle-, and maximum points, while the separatrices are the paths of steep descent connecting the minima and maxima to the saddles [Cay59].

The MS-complex induces a segmentation of the image into regions of monotonic behavior [Mil65] and is strongly related [NS94] to the concept of the watershed transform [Max70]. In fact, the separatrices form a superset of the watersheds and water-courses [GC95, LLSV99].

There are three established methods to compute the MS-complex: The classical approach employs numerical methods. In this setting, the critical points are given by computing all zeros of the gradient. The separatrices are extracted by starting at

^{*}Institute for Science and Technology, Austria. jan.reininghaus@ist.ac.at

[†]MPI for Informatics, Germany. {dguenther, weinkauff, hpseidel}@mpi-inf.mpg.de

[‡]Zuse-Institute Berlin, Germany. hotz@zib.de

the saddle points and following the gradient in the direction of the eigenvectors of their Hessian [Wei08]. Using interval methods, the MS-complex can be extracted in a certified manner [Cha11]. The second approach works in a piecewise linear context. In this setting, the critical points are given by an analysis of the lower star of each vertex [Ban70]. The separatrices are typically approximated as a sequence of steepest edges in the triangulation [Zom01, BEHP04].

In this paper, we build upon a purely combinatorial approach [For98b, For01] to compute the MS-complex. Such an approach lends itself to computational purposes due to its discrete nature [Bau11, Gyu08, Lew05]. In contrast to the classical approach, it approximates the MS-complex directly on the grid defined by the image. In this setting, the critical points are defined by the topological changes in the sub-level sets of the data [Mil63]. These topological changes can be computed efficiently by constructing a combinatorial gradient [RWS11]. The separatrices are then computed by starting at the (combinatorial) saddle points and following the grid along this combinatorial gradient.

To analyze the approximation error of such a combinatorial algorithm, one can apply it to a sampling of an analytic function f using k^2 pixels and compute the distance of its result to the exact MS-complex of f . A natural expectation is that this distance should go to zero when k is increased. This is not the case for any established algorithm.

The main reason for this behavior is the combinatorial representation of the gradient direction in terms of the directions provided by the grid. In all established algorithms, the grid direction with the steepest descent is chosen to approximate the gradient. When one follows the combinatorial gradient, the quantization error can accumulate to a large value. Since this error only depends on the number of possible directions and the gradient of the data, it does not decrease when a finer grid is employed.

Based on the mathematical foundation presented in Section 2, we propose an efficient probabilistic method that computes a combinatorial gradient in Section 3. In contrast to all established algorithms, the approximation error of its induced MS-complex almost surely goes to zero when the image resolution is increased. While a formal proof is only hinted at, we do provide a thorough numerical evaluation of our method and compare it to established algorithms in Section 4. We conclude this paper with a discussion of possible future directions and extensions in Section 5.

2 Computational Discrete Morse Theory

This section introduces the main mathematical concepts and algorithms which our method builds upon. We will first give a brief introduction to discrete Morse theory [For98b] in a graph theoretical notation [RGH⁺10]. Using this notation, we will then describe the algorithm which we build upon in Section 3.

2.1 Definitions

Let I denote a two dimensional image represented by real numbers defined on a rectangular grid Ω . In topological terms, Ω is called a cubical cell complex C [Hat02, KMM04]. This complex consists of cells with different dimensions (e.g., vertices, edges, cubes) and of boundary maps describing their neighborhood relation. For example, an edge is bounded by its two incident vertices, whereas a quad is bounded by its incident edges.




Figure 1: Illustration of a cell graph (a) of a 2×1 grid. The solid edges in (b) represent a combinatorial gradient containing critical points (black). A p -line is shown in (c), while (d) shows two separatrices (blue, green) connecting a saddle (yellow) to two minima (blue).

To define the essential concepts of discrete Morse theory, we consider the cell complex C in a graph theoretical setting: the cell graph $G = (N, E)$ encodes the essential combinatorial information of C . The nodes N represent the cells of C and each node u^p is labeled by the dimension p of the cell it represents. The edges E encode the neighborhood relation of the cells. If a cell u^p is in the boundary of a cell w^{p+1} , then $e^p = \{u^p, w^{p+1}\} \in E$. The edge e^p is said to be of index p .

The main task in computational discrete Morse theory is now to construct a combinatorial gradient such that the following combinatorial definitions of critical points and separatrices correspond to the input image I .

Formally, a combinatorial gradient field V is a subset of pairwise non-adjacent edges of G with a certain acyclic constraint [Cha00] defined below. Given such a combinatorial gradient field V , the critical points are the unmatched nodes of V . A critical point u^p that represents a cell of dimension p is a minimum ($p = 0$), saddle ($p = 1$), or maximum ($p = 2$). A combinatorial p -line is a path in the cell graph G whose edges are of index p and alternate between V and its complement $E \setminus V$. The above mentioned acyclic constraint is now specified as the non-existence of any closed p -line. A p -line connecting two critical points u^p and w^{p+1} is called a combinatorial p -separatrix. A 0-separatrix thereby connects a minimum with a saddle, while a 1-separatrix connects a saddle with a maximum.

Figure 1 shows a simple cell graph of a 2×1 grid and an arbitrary combinatorial gradient field with its critical points and separatrices.

2.2 Algorithm

As already mentioned, the main task in computational discrete Morse theory is to construct a combinatorial gradient that corresponds to the input data I . Many such algorithms have been proposed [BLW12, GBPH11, KKM05, Lew05]. In this paper, we make use of the algorithm *ProcessLowerStars* proposed by Robins et al. [RWS11]. The critical points of their combinatorial gradient provably correspond one-to-one to the topological changes of the lower-level sets of the input data in up to three dimen-

sions. Also, this algorithm is very efficient, since it has linear running time and a parallel implementation scales well [GRP⁺12].

In Section 3 we will propose an extension of this algorithm. Therefore, we now present it in detail in our graph theoretical notation.

We first propagate the input I from the 0-nodes to all nodes of the graph: each node u^p is assigned the maximum I -value of the vertices of the cell that u^p represents. We denote this extension of I by \hat{I} . Algorithm 1 decomposes then the cell graph G into the lower stars [Ban70] defined by \hat{I} (line 3, 4, 5 and 6). Note that this decomposition is disjoint, which allows for good parallel scalability. Each lower star is now grown from its vertex using *simple homotopic expansions* – the inverse of homotopic collapses [Coh73]. Such expansions are represented by the edges of the cell graph (line 7).

The combinatorial gradient V is now constructed iteratively (line 10): each time we expand a lower star using an edge $e \in E$ (lines 13, 14 and 15), we append e to V (line 16). An edge $e = \{u^p, w^{p+1}\} \in E$ is admissible for simple homotopic expansion if the following conditions hold:

1. u^p and w^{p+1} are not covered by an edge in the current V (line 13),
2. u^p and w^{p+1} have not been flagged previously (line 13),
3. there is no other edge $\{z^p, w^{p+1}\} \in E$ that fulfills 1. and 2. (line 14).

If the set of admissible edges L is empty (line 19), we flag an arbitrary node in the lower star (line 23) that is not covered by an edge in the current V (line 22). If no such node can be found, the expansion stops. If L is not empty (line 15), an admissible edge is chosen based on an order defined by \hat{I} and appended to V (line 16).

As shown by Robins et al. [RWS11], the way we choose an edge from L (line 16) does not affect the overall number nor the type of critical points in the resulting combinatorial gradient. Since the combinatorial gradient is supposed to correspond to the (continuous) gradient, a natural choice is the edge that represents locally the steepest descent.

3 Almost Surely Convergent Separatrices

Before we describe our method in Section 3.2, we give an explanation for the non-convergent behavior of the existing combinatorial gradient algorithms and motivate our almost surely convergent probabilistic approach.

3.1 Motivation

As defined in Section 2.1, the separatrices of a combinatorial gradient V are alternating paths in the cell graph G with respect to V . Intuitively, the edges in V should therefore reflect the direction of the gradient of the input data. However, at a given vertex u , there are only a constant number of directions representable by the edges of G . This implies that the continuous gradient can only be represented in a quantized way. Loosely speaking, the gradient direction is snapped to the edges of the graph.

Therefore, the combinatorial gradient differs from the continuous gradient not only by a sampling error, but also by a quantization error. Note that the sampling error can be decreased using a denser sampling. However, this is not necessarily the case for the quantization error.

Algorithm 1 CombinatorialGradient(G, I)

Input: $G = (N, E), \hat{I} : N \rightarrow \mathbb{R}$ **Output:** $V \subset E$

```
1:  $V \leftarrow \emptyset$ 
2: for all  $v^0 \in N$  do
3:    $S \leftarrow v^0$ 
4:    $W \leftarrow \{w \in N : \hat{I}(w) \leq \hat{I}(v^0)\}$ 
5:   for  $p \leftarrow 0, \dots, d-1$  do
6:      $S \leftarrow S \cup \{w^{p+1} \in W : \exists \{u^p, w^{p+1}\} \in E, u^p \in S\}$ 
7:    $K \leftarrow E(S)$ 
8:    $C \leftarrow \emptyset$ 
9:    $abort \leftarrow false$ 
10:  while  $abort = false$  do
11:     $abort \leftarrow true$ 
12:    for  $p \leftarrow 0, \dots, d-1$  do
13:       $T \leftarrow \{\{u^p, w^{p+1}\} \in K : u^p, w^{p+1} \notin C \cup N(V)\}$ 
14:       $L \leftarrow T \setminus \{\{u^p, w^{p+1}\} \in T : \exists \{z^p, w^{p+1}\} \in T, u^p \neq z^p\}$ 
15:      if  $L \neq \emptyset$  then
16:         $V \leftarrow V \cup \text{ChooseEdge}(L)$ 
17:         $abort \leftarrow false$ 
18:        goto Line 10
19:      else
20:        for  $k \leftarrow 0, \dots, d$  do
21:           $\{u_0^k, u_1^k, \dots, u_m^k\} \leftarrow \{u^k \in S : u^k \notin C \cup N(V)\}$ 
22:          if  $\{u^k \in S : u^k \notin C \cup N(V)\} \neq \emptyset$  then
23:             $C \leftarrow C \cup u_0^k$ 
24:             $abort \leftarrow false$ 
25:            goto Line 10
```

Perhaps surprisingly, the ubiquitous steepest descent strategy for the edge selection in Line 16 of Algorithm 1 suffers from this quantization artifact. Suppose that the exact gradient is almost constant in a region K and points 'North-North-East'. At any given vertex in K the steepest descent direction is therefore always 'North' – independent of the resolution used to sample the exact gradient. Any exact separatrix passing through K is thereby approximated by a straight line going 'North'. This (resolution independent) behavior can be observed frequently in practice as can be seen in Figure 7.

To deal with the quantization error, we propose to choose the edges in Line 16 of Algorithm 1 adjacent to a vertex u in a probabilistic fashion. Since we cannot represent the (continuous) gradient direction exactly, we pick an edge according to a random variable X_u . The probability mass function g of X_u is defined by the image data I and the width and height of each pixel. The basic idea is to design g such that the expected value of X_u corresponds to the (continuous) gradient direction at u .

Note that these random variables are independent. Assuming that in this setting the law of large numbers [Ber13] is applicable, a path following this probabilistic combinatorial gradient will therefore almost surely proceed in the direction of the (continuous) gradient when the grid is refined. While this argument is far from a formal proof, we do provide a thorough numerical evaluation in Section 4 that substantiates this intuition.

3.2 Method

The main building block of our method consists of Algorithm 1. The only change is that we choose the direction (Line 16 of Algorithm 1) from L in a probabilistic fashion instead of choosing the locally steepest descent.

The index of the edges in L is either always 1 or always 0 (Line 14 of Algorithm 1). Perhaps surprisingly, it suffices to choose the edges of index 0 appropriately. The order in which the edges of index 1 are chosen has no effect. They are uniquely defined once the edges of index 0 and the saddle points have been selected. This fact motivated the original construction of a combinatorial gradient in [LLT03]. In the following, we therefore assume that L contains only edges of index 0.

For a given vertex $u^0 \in N$, the edge selection strategy is given by a random variable X_u . The value of this random variable is always an edge in L . We now define a probability mass function $g : L \rightarrow [0, 1]$ for X_u such that the expected value of X_u is collinear to the (continuous) gradient at u .

We assume (without loss of generality) that the (continuous) gradient points North-East, i.e., $\nabla I(u) = (I_x, I_y)$ with $I_x, I_y \geq 0$. Furthermore, the width of the current pixel is denoted by w and its height by h . The set L thereby consists of the directions $(0, h)$ and $(w, 0)$.

To simplify notation, we refer to $g((w, 0))$ by λ . Since g is a probability mass function, we have $g((0, h)) = 1 - \lambda$. The expected direction $E(X_u)$ is now given by

$$E(X_u) = (1 - \lambda) \begin{pmatrix} 0 \\ h \end{pmatrix} + \lambda \begin{pmatrix} w \\ 0 \end{pmatrix} = \begin{pmatrix} \lambda w \\ (1 - \lambda)h \end{pmatrix}. \quad (1)$$

Since $E(X_u)$ should be collinear to $\nabla I(u) = (I_x, I_y)$, the following condition must hold

$$\det \begin{pmatrix} \lambda w & I_x \\ (1 - \lambda)h & I_y \end{pmatrix} = 0. \quad (2)$$

This yields

$$g((w, 0)) = \frac{hI_x}{wI_y + hI_x}, \text{ and } g((0, h)) = \frac{wI_y}{wI_y + hI_x}. \quad (3)$$

Since $\nabla I(u)$ is not directly available, we approximate it using finite differences:

$$I_x \approx \frac{I((u + (w, 0))) - I(u)}{w}, \text{ and } I_y \approx \frac{I((u + (0, h))) - I(u)}{h}. \quad (4)$$

Denoting the height difference $I((u + (w, 0))) - I(u)$ by W and $I((u + (0, h))) - I(u)$ by H , and inserting (4) into (3) yields the final probability mass function g in terms of I , w and h :

$$g((w, 0)) = \frac{h^2W}{w^2H + h^2W}, \text{ and } g((0, h)) = \frac{w^2H}{w^2H + h^2W}. \quad (5)$$

Note that, in practice, L may consist of more than 2 edges due to the sampling of I . Each edge in L is therefore assigned the height difference weighted by the squared length of the dual edge. Its probability is then given by the normalized value with respect to the other edges in L .

4 Evaluation and Comparison

In the following, we evaluate our probabilistic method and compare it to established algorithms. All experiments were performed on a machine with two Intel Xeon E5645 CPUs. We applied the linear Algorithm 1 in parallel. For an image of resolution 4096^2 , Algorithm 1 needed about 6 seconds. The probabilistic and the steepest descent version take the same time, since the amount of time needed to choose the edges is negligible. Using the implicit representation of the cell graph proposed in [GRP⁺12], the memory requirement is very low. To process an image of k pixels we need $2k$ bytes of main memory.

An analytic function. Let $\Omega = [-2, 2]^2$ and $\alpha \in \mathbb{R}^+$. The function $f : \Omega \rightarrow \mathbb{R}$ is given as

$$f(x, y) = -e^{-\alpha(\sqrt{x^2+y^2}-1)^2} - 0.3(x+y). \quad (6)$$

The function f describes a circle engraved on a tilted plane. The sharpness of this circle is defined by the parameter α . For $\alpha \rightarrow \infty$, the circle becomes arbitrary sharp. For $\alpha \rightarrow 0$, f gets flattened. Varying α allows us to simulate smooth as well as sharp features appearing in many applications. An illustration of f sampled on a 2048^2 grid for different choices of α is given in the first row of Figure 6. Integral lines of the continuous gradient ∇f are depicted by black lines using the dual streamline seeding technique [RPP⁺09]. Converging integral lines indicate thereby the existence of a separatrix. In the following paragraphs, we choose the engraved circle as a reference feature. For illustration, it is shown as a white circle line in the first row of Figure 6.

A qualitative comparison. We applied Algorithm 1 using the steepest descent version as well as our probabilistic version to construct a combinatorial gradient of f for different choices of α . The resulting MS-complexes of the steepest descent version are shown in the second row of Figure 6. Our reference feature – the white circle – is visually well recovered for large values of α , which confirms also the extraction results of recently proposed methods [CCL03, KRHH11, WG09]. In many applications, the desired features are sufficiently sharp.

However, deviations to the reference circle become visible if the feature gets smooth, i.e., for small choices of α . The steepest descent version of Algorithm 1 is not able to recover the circle for $\alpha = 2^0$ and $\alpha = 2^1$. The probabilistic approach, in contrast, is able to recover the circle for all choices of α . The resulting MS-complexes are shown in the third row of Figure 6.

A quantitative comparison. To quantify the approximation error, we measured the Hausdorff distance [Hau14] of the reference circle to the approximated circle. Figure 6 g) shows the approximation error for Algorithm 1 using the steepest descent strategy in a log-log plot. Although the extraction result looked visually reasonable for large α -values, there is no convergence for any α . The Hausdorff distance to the reference circle does not decrease when a finer grid is employed.

For the probabilistic approach, we did 200 runs of Algorithm 1 using the method presented in Section 3.2. Figure 6 h) shows the mean value of the Hausdorff distance and its standard deviation. Both quantities are converging to zero. Hence, the sampling as well as the quantization error mentioned in Section 3.2 are reduced when a finer grid is employed. However, it needs to be noted that the approximation error for very sharp




Figure 2: Distribution of Hausdorff distance error. The blue, black and red curves show the estimated probability density function of the Hausdorff distance between the center circle and the reference circle (as shown in Figure 6) for different resolutions.

features ($\alpha = 2^5$) is slightly larger compared to the steepest descent strategy when coarse grids are used.

In Figure 2, we plotted a kernel density estimate [BGK10] of the probability density of the Hausdorff distance error for different resolutions. We applied Algorithm 1 using the probabilistic approach 1000 times to obtain a sufficient number of samples for the density estimate. The overall shape of the densities suggests a sequence of binomial error distributions that converge to a normal distribution.

A segmentation comparison. In Figure 3 we compared ourselves (f) with different established segmentation approaches: steepest descent [RWS11] (b), Meyer’s watershed algorithm [Mey94] using a 4- (c) and 8-connectivity (d) provided by the image processing toolbox of Matlab [MAT09], and a fitted spanning forest approach [CCL03] (e). Since the implementation of [CCL03] that was available to us requires a triangulation we subdivided each pixel into two triangles.

The objective is the segmentation of the inner part of our reference circle of f with $\alpha = 1$ sampled on a 2048^2 grid. It can be observed that the steepest descent strategy yields a similar result as the watershed algorithm using an 8-connectivity, while the result of the spanning forest approach looks similar to the watershed result using a 4-connectivity. However, none of these approaches is able to segment the circle, in contrast, to the probabilistic approach presented in this paper.

Anisotropic grids. In certain applications, the image data is not provided on uniform grids. Based on the imaging technique anisotropic grids might be employed. In Figure 4, we illustrate the behavior of the steepest descent and our probabilistic strategy applied to f with $\alpha = 1$ using two different anisotropic grids.

It can be seen in Figure 4 a) and c) that the result of the steepest descent approach heavily depends on the kind of anisotropy. The resulting segmentation is distorted when the anisotropy becomes large. However, this is not the case for our probabilistic approach, see Figure 4 b) and d).




Figure 3: Segmentation comparison. The analytic function f as defined in (6) with $\alpha = 1$ is visualized as in Figure 6 – the black lines depict integral lines of the gradient. Converging integral lines indicate the separatrices/watersheds. (b) shows the result using Algorithm 1 with the steepest descent strategy. (c) and (d) show the result using the Matlab implementation of the classic watershed algorithm [Mey94]. (e) shows the result using a fitted spanning forest approach [CCL03]. (f) shows the segmentation using our probabilistic approach based on Algorithm 1. The red circle shows the reference separatrix/watershed of the continuous function.

Grid dependency. To evaluate the grid dependency of our approach, we rotated f by different angles $\theta \in [-\pi/2, \pi/2]$. Since this is a rigid transformation of f , a rotation by $-\theta$ of the extracted separatrices should coincide with the original separatrices of f . The rotated function $f_\theta : \Omega \times [-\pi/2, \pi/2] \rightarrow \mathbb{R}$ is given by

$$f_\theta(x, y) = f(R(x, y, \theta)) \quad (7)$$

with

$$R(x, y, \theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

We concentrate ourselves on the reference feature of Figure 6. Since the center circle is rotated and back-rotated by θ , the ground truth is a cylinder. We extracted the center circle using the steepest descent and our probabilistic strategy for each rotation step. The spatial domain was discretized using a 1024^2 grid and a full rotation was sampled using 1024 steps.




Figure 4: Segmentation comparison on anisotropic grids. The left column shows the segmentation of the analytic function f (6) with $\alpha = 1$ using Algorithm 1 with the steepest descent approach on two anisotropic grids. The right column shows the corresponding results for our probabilistic approach. The red circle depicts the continuous separatrix/watershed of f .




Figure 5: Grid dependence. (a) shows the ground truth of the rotating function f_θ as defined in (7). The gray surface depicts the evolution of the spatial reference feature shown in Figure 6. (b) and (c) show the discrete result of Algorithm 1 using the steepest descent edge selection and our probabilistic strategy, respectively.

In order to obtain a segmentation of this region over the the range of the rotation, we stacked the spatial separatrices on top of each other and connected them to a surface. The result is shown in Figure 5.

The result of the steepest descent strategy shown in Figure 5 b) clearly reflects the structure of the grid. The surface is snapped to the grid, no cylindrical shape can be observed. In contrast, the probabilistic approach shown in Figure 5 c) is a rough approximation of the ground truth, and the cylindrical shape is well recovered.

Generic functions. To demonstrate that our approach is convergent for general smooth functions, we considered a set of smooth functions generated by the expression

$$\sum_{m,n=1}^2 \left(X_{m,n}^{(1)} \sin(mx) + X_{m,n}^{(2)} \cos(mx) \right) \left(X_{m,n}^{(3)} \sin(ny) + X_{m,n}^{(4)} \cos(ny) \right), \quad (8)$$

where the $X_{m,n}^{(j)}$'s are random variables uniformly distributed in $[-1, 1]$. This expression is now evaluated on the domain $[-\pi, \pi]^2$ discretized using two different grid resolutions. We selected two representatives of this set of functions and applied the steepest descent and our probabilistic version of Algorithm 1. The resulting MS-complexes are shown in Figure 7.

It is apparent that the steepest descent approach does not yield the correct MS-complex. Our probabilistic edge selection strategy, however, visually converges to the correct solution. Note that we applied our method to a much larger number of such functions. We observed the above behavior in every case.

5 Conclusion and Future Work

We presented a probabilistic approach that computes a combinatorial gradient for a given two dimensional image data set. Since our method is a modification of a previously proposed method [RWS11], it shares its main properties. It has a linear running time and can be applied efficiently in a parallel setting. Furthermore, its critical points correspond one-to-one with the topological changes of the sublevel sets of the input data.



Our probabilistic extension of this algorithm yields combinatorial gradients whose separatrices converge to their continuous counterpart when the grid resolution is increased. We also demonstrated that our method works for anisotropic grids and does not exhibit a grid dependency. While we only hinted at a formal proof of this property, we provided a thorough numerical evaluation and compared it to existing algorithms.

We believe that this approach could be extended and improved in several directions:

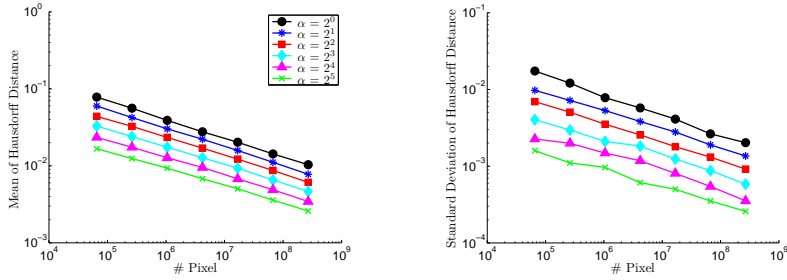
- It seems that an generalization to scalar data defined on triangulated surfaces is feasible. The main challenge is thereby the generalization of the derivation of the edge selection strategy in Section 3.2. This may also yield some insight on the representation of the discrete metric provided by the triangulation.
- Extending our approach to higher dimensional image data could also be interesting and useful. This may be quite challenging since combinatorial separatrices that connect the saddle points can merge and split in 3D. This stems from a fundamentally different structure of the cell graph.
- The efficiency of the method may be increased by developing an adaptive grid refinement approach. This may be particularly effective since a high resolution is only needed in the vicinity of the separatrices.
- It could also be possible to extend this idea to the combinatorial vector field context [For98a, RLH11].

References

- [Ban70] T. F. Banchoff, *Critical points and curvature for embedded polyhedral surfaces*, The American Mathematical Monthly **77** (1970), no. 5, pp. 475–485 (English).
- [Bau11] Ulrich Bauer, *Persistence in discrete Morse theory*, Ph.D. thesis, University of Göttingen, 2011.
- [BEHP04] Peer-Timo Bremer, Herbert Edelsbrunner, Bernd Hamann, and Valerio Pascucci, *A topological hierarchy for functions on triangulated surfaces*, IEEE Transactions on Visualization and Computer Graphics **10** (2004), 385–396.
- [Ber13] Jakob Bernoulli, *Ars conjectandi*, impensis Thurnisiorum, fratrum, 1713.
- [BGK10] Z. I. Botev, J. F. Grotowski, and D. P. Kroese, *Kernel density estimation via diffusion*, Annals of Statistics **38** (2010), no. 5, 2916–2957.
- [BLW12] Ulrich Bauer, Carsten Lange, and Max Wardetzky, *Optimal topological simplification of discrete functions on surfaces*, Discrete and Computational Geometry **47** (2012), 347–377.
- [Cay59] A. Cayley, *On contour and slope lines*, The London, Edinburg and Dublin Philosophical Magazine and Journal of Science **18** (1859), 264–268.
- [CCL03] F. Cazals, F. Chazal, and T. Lewiner, *Molecular shape analysis based upon the Morse-Smale complex and the Connolly function*, SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry (New York, NY, USA), ACM Press, 2003, pp. 351–360.



(g) Steepest descent error



(h) Statistical characteristics of probabilistic error (this paper)

Figure 6: The analytic example f (6). The first row of (a)-(f) shows the sampled function f color-coded for different choices of α . Red denotes a high function value, while blue denotes a low value. Black lines depict integral lines of the gradient ∇f seeded using the dual streamline technique [RPP⁺09]. The second and third row show the MS-complex based on Algorithm 1 using the steepest descent and our probabilistic edge selection strategy, respectively. Minima, saddles and maxima are shown as blue, yellow and red spheres, whereas the 0- and 1-separatrices are shown as blue and red lines, respectively. Since separatrices are integral lines of the gradient, the blue and red lines should follow the black lines. (g) and (h) show the Hausdorff distance of the approximated center circle (blue) to the reference circle (white, first row) for different choices of α and increasing resolution. (g) shows the evolution of the error using the steepest descent approach. (h) shows two statistical characteristics of the error of the probabilistic approach proposed in this paper.




Figure 7: Comparison of the steepest descent and the probabilistic approach using generic functions. The two rows show two different functions randomly generated using expression (8). Red denotes a high function value, while blue denotes a low value. Black lines depict integral lines of their gradient. The left two columns show the extracted MS-complexes for each function sampled on a 512^2 grid using the steepest descent and our probabilistic edge selection strategy, respectively. The right two columns show the result on a 4096^2 grid. Minima, saddles and maxima are shown as blue, yellow and red spheres, whereas the 0- and 1-separatrices are shown as blue and red lines, respectively. Since separatrices are integral lines of the gradient, the blue and red lines should follow the black lines. White arrows indicate regions where the difference of the two approaches is visually apparent.

[Cha00] Manoj K. Chari, *On discrete Morse functions and combinatorial decompositions*, Discrete Math. **217** (2000), no. 1-3, 101–113.

[Cha11] A. Chattopadhyay, *Certified geometric computation: Radial basis function based isosurfaces and Morse-Smale complexes*, Ph.D. thesis, University of Groningen, 2011.

- [Coh73] M.M. Cohen, *A course in simple-homotopy theory*, Springer-Verlag, 1973.
- [EH10] H. Edelsbrunner and J. Harer, *Computational topology. an introduction*, American Mathematical Society, January 2010.
- [For98a] Robin Forman, *Combinatorial vector fields and dynamical systems*, Mathematische Zeitschrift **228** (1998), 629–681.
- [For98b] Robin Forman, *Morse theory for cell complexes*, Advances in Mathematics **134** (1998), 90–145.
- [For01] Robin Forman, *A user's guide to discrete Morse theory*, Proceedings of the 2001 Internat. Conf. on Formal Power Series and Algebraic Combinatorics, Advances in Applied Mathematics, 2001.
- [GBPH11] Attila Gyulassy, Peer-Timo Bremer, Valerio Pascucci, and Bernd Hamann, *Practical considerations in Morse-Smale complex computation*, Topological Methods in Data Analysis and Visualization (Valerio Pascucci, Xavier Tricoche, Hans Hagen, and Julien Tierny, eds.), Mathematics and Visualization, Springer Berlin Heidelberg, 2011, pp. 67–78.
- [GC95] Lewis D Griffin and Alan CF Colchester, *Superficial and deep structure in linear diffusion scale space: isophotes, critical points and separatrices*, Image & Vision Comp. **13** (1995), no. 7, 543 – 557.
- [GRP⁺12] D. Günther, J. Reininghaus, S. Prohaska, T. Weinkauff, and H.-C. Hege, *Efficient computation of a hierarchy of discrete 3d gradient vector fields*, Topological Methods in Data Analysis and Visualization II (R. Peikert, H. Hauser, H. Carr, and R. Fuchs, eds.), Mathematics and Visualization, Springer, 2012, TopoInVis 2011, Zürich, Switzerland, April 4 - 6, pp. 15–30.
- [Gyu08] Attila Gyulassy, *Combinatorial construction of Morse-Smale complexes for data analysis and visualization.*, Ph.D. thesis, University of California, Davis, 2008.
- [Hat02] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, U.K., 2002.
- [Hau14] F. Hausdorff, *Grundzüge der Mengenlehre*, Chelsea Publishing Series, Chelsea Publishing Company, 1914.
- [KKM05] Henry King, Kevin Knudson, and Neza Mramor, *Generating discrete Morse functions from point data*, Experimental Mathematics **14** (2005), no. 4, 435–444.
- [KMM04] T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational homology*, Applied Math. Sciences, vol. 157, Springer-Verlag, 2004.
- [KRHH11] Jens Kasten, Jan Reininghaus, Ingrid Hotz, and Hans-Christian Hege, *Two-dimensional time-dependent vortex regions based on the acceleration magnitude.*, IEEE Trans. Vis. Comput. Graph. **17** (2011), no. 12, 2080–2087.

- [Lew05] Thomas Lewiner, *Geometric discrete Morse complexes*, Ph.D. thesis, Department of Mathematics, PUC-Rio, 2005.
- [LLSV99] A. M. Lopez, F. Lumbreras, J. Serrat, and J. J. Villanueva, *Evaluation of methods for ridge and valley detection*, IEEE Trans. on PAMI **21** (1999), no. 4, 327–335.
- [LLT03] Thomas Lewiner, Hlio Lopes, and Geovan Tavares, *Optimal discrete Morse functions for 2-manifolds*, Computational Geometry **26** (2003), no. 3, 221 – 233.
- [MAT09] MATLAB, *version (r2009a)*, The MathWorks Inc., Natick, Massachusetts, 2009.
- [Max70] J. C. Maxwell, *On hills and dales*, The London, Edinburgh and Dublin Philosophical Magazine and Journal of Science **40** (1870), 421–425.
- [Mey94] Fernand Meyer, *Topographic distance and watershed lines*, Signal Processing **38** (1994), no. 1, 113 – 125, Mathematical Morphology and its Applications to Signal Processing.
- [Mil63] John Milnor, *Morse theory*, Princeton University Press, 1963.
- [Mil65] J.W. Milnor, *Topology from the differentiable viewpoint*, Univ. Press Virginia, 1965.
- [NS94] Laurent Najman and Michel Schmitt, *Watershed of a continuous function*, Signal Processing **38** (1994), no. 1, 99 – 112, Mathematical Morphology and its Applications to Signal Processing.
- [RGH⁺10] Jan Reininghaus, David Günther, Ingrid Hotz, Steffen Prohaska, and Hans-Christian Hege, *TADD: A computational framework for data analysis using discrete Morse theory*, Mathematical Software – ICMS 2010 (Komei Fukuda, Joris van der Hoeven, Michael Joswig, and Nobuki Takayama, eds.), Lecture Notes in Computer Science, vol. 6327, Springer, 2010, pp. 198–208.
- [RLH11] Jan Reininghaus, Christian Löwen, and Ingrid Hotz, *Fast combinatorial vector field topology*, IEEE Transactions on Visualization and Computer Graphics **17** (2011), 1433–1443.
- [RPP⁺09] Olufemi Rosanwo, Christoph Petz, Steffen Prohaska, Ingrid Hotz, and Hans-Christian Hege, *Dual streamline seeding*, Proceedings of the IEEE Pacific Visualization Symposium (Peter Eades, Thomas Ertl, and Han-Wei Shen, eds.), 2009, pp. 9 – 16.
- [RWS11] Vanessa Robins, Peter John Wood, and Adrian P. Sheppard, *Theory and algorithms for constructing discrete Morse complexes from grayscale digital images*, IEEE Trans. on PAMI **33** (2011), no. 8, 1646–1658.
- [Sma61] Stephen Smale, *On gradient dynamical systems*, The Annals of Mathematics **74** (1961), 199–206.
- [Wei08] T. Weinkauff, *Extraction of topological structures in 2d and 3d vector fields*, Ph.D. thesis, University Magdeburg, 2008.

- [WG09] T. Weinkauff and D. Günther, *Separatrix persistence: Extraction of salient edges on surfaces using topological methods*, Computer Graphics Forum (Proc. SGP '09) **28** (2009), no. 5, 1519–1528.
- [Zom01] Afra Zomorodian, *Computing and comprehending topology: Persistence and hierarchical morse complexes*, Ph.D. thesis, Urbana, Illinois, October 2001.