Revisiting Accurate Geometry for Morse-Smale Complexes

Son Le Thanh* KTH Royal Institute of Technology
Michael Ankele† Indurad
Tino Weinkauf‡ KTH Royal Institute of Technology

ABSTRACT

The Morse-Smale complex is a standard tool in visual data analysis. The classic definition is based on a continuous view of the gradient of a scalar function where its zeros are the critical points. These points are connected via gradient curves and surfaces emanating from saddle points, known as separatrices. In a discrete setting, the Morse-Smale complex is commonly extracted by constructing a combinatorial gradient assuming the steepest descent direction. Previous works have shown that this method results in a geometric embedding of the separatrices that can be fundamentally different from those in the continuous case. To achieve a similar embedding, different approaches for constructing a combinatorial gradient were proposed. In this paper, we show that these approaches generate a different topology, i.e., the connectivity between critical points changes. Additionally, we demonstrate that the steepest descent method can compute topologically and geometrically accurate Morse-Smale complexes when applied to certain types of grids. Based on these observations, we suggest a method to attain both geometric and topological accuracy for the Morse-Smale complex of data sampled on a uniform grid.

Index Terms: Discrete Morse theory, Morse-Smale complex, Topology, Accurate geometry.

1 INTRODUCTION

The Morse-Smale complex [16, 20] has proven to be a powerful tool for topological data analysis and has found its applications in many fields such as in material science [11], fluid dynamics [7, 19], computer graphics [21], or in molecular biology [2].

Forman's discrete Morse theory [4] gives a formidable setting for computing the Morse-Smale complex in a combinatorial setting [5, 12, 18]. The essential component of these algorithms is the computation of the discrete gradient field, which implicitly encodes the structure of the Morse-Smale complex. The remarkably efficient method by Robins et al. [18] can be considered the *de facto* standard for the calculation of the discrete gradient field. The method aligns its discrete vectors locally with the steepest descent direction.

It has been shown by Gyulassy et al. [9] and Reininghaus et al. [17] that the geometric embedding of the discrete Morse-Smale complex does not coincide with its continuous counterpart when using the steepest descent direction in Robins' method – even when increasing the resolution of the discretization. This has been investigated for uniform grids, but not for other types of grids. Instead, Gyulassy et al. [9] and Reininghaus et al. [17] propose a new method to compute the discrete gradient field, which replaces the steepest descent direction in Robins et al. [18] with a probabilistic choice of descending directions. This achieves a better geometric embedding of the Morse-Smale complex.

However, as we will show in this paper, moving in directions other than the steepest descent will lead to a different topology. Specifically, the connectivity between critical points differs. We show that these differences can be found in many scenarios, from small data sets to large ones, from smooth data sets to noisy ones, from artificial data sets to real-world data sets. We provide a theoretical discussion that reveals which layers of the cell complex are affected by the probabilistic approach of Gyulassy et al. [9] and Reininghaus et al. [17].

To provide solutions, we investigate different types of grids for their susceptibility to distortions of the geometric embedding. Based on these observations, we propose a method to convert a uniform grid to a specific triangle grid that provides as good of a geometric embedding as the probabilistic approaches of Gyulassy et al. [9] and Reininghaus et al. [17], but can be used with the steepest descent method of Robins et al. [18], thereby inheriting its topological guarantees.

Our contributions are as follows:

- We provide a novel, systematic analysis of the differences in the topology of Morse-Smale complexes produced by various methods (Section 4). These changes happen across a wide range of scenarios and can not be removed by the means of topological simplification.
- We investigate different types of grids and identify empirically the properties that a grid needs to have such that the geometric embedding obtained using the steepest descent method aligns with the continuous case (Section 5).
- We propose a method to achieve both geometric accuracy and topological consistency for data sampled on a uniform grid by converting it to a specific type of triangle grid and applying the steepest descent method (Section 6).

Related work and theoretical background are discussed in the following two sections, with a focus on the topic of accurate geometry in Section 3.

2 BACKGROUND

We briefly go over the relevant concepts in the following.

2.1 Morse function and Morse-Smale complex

Morse theory [15] allows to analyze the topology of a manifold by examining a function defined on that manifold. Let $f : \mathbf{M} \to \mathbf{R}$ be a function defined on a d-manifold with boundary $\mathbf{M}$, $\mathbf{p} \in \mathbf{M}$ is a *critical point* of f if $\nabla f(\mathbf{p}) = 0$. If the Hessian of f at $\mathbf{p}$, $H(\mathbf{p})$, is non-singular, i.e. $|H(\mathbf{p})| \neq 0$, then $\mathbf{p}$ is called a *non-degenerate critical point*. The function f is called a *Morse function* if all of its critical points are non-degenerate. Given a Morse function f with a non-degenerate critical point $\mathbf{p}$, the result of the Morse lemma states there exists a local coordinate in the neighborhood of p such that f can take the quadratic form $f(\mathbf{x}) = f(\mathbf{p}) - x_1^2 - \ldots - x_\gamma^2 + x_{\gamma+1}^2 + \ldots + x_d^2$. The value of γ in this formula is called the *index* of f at $\mathbf{p}$ and can be used to characterize the type of the critical points of f. In the case of 2 dimensions, the indices 0, 1, and 2 correspond to the minima, saddle points, and maxima of the function, respectively. For volumetric data, minima have index 0, 1-saddles have index 1, 2-saddles have index 2, and maxima have index 3.

The tangent vector of a *tangent curve* $\phi : \mathbf{R} \to \mathbf{M}$ agrees with the gradient ∇f at every point along the line. In other words, ϕ is the

*e-mail: sonlt@kth.se
†e-mail: michael.ankele@indurad.com
‡e-mail: weinkauf@kth.se

solution of the equation $\frac{\partial}{\partial t}\phi(t) = \nabla f(\phi(t))$, where $\phi(0) = \mathbf{p}$. For each tangent curve, the point $\lim_{t\to-\infty}\phi(t)$ is called the *source* or *origin* of ϕ, whereas $\lim_{t\to\infty}\phi(t)$ is called the *destination* or *sink*. The sets of tangent curves having the same source or sink are called *ascending* and *descending* manifolds, respectively.

A function f fulfills the *Morse-Smale condition* if its ascending and descending manifolds intersect transversally for each pair of critical points. If f satisfies such condition, the intersection of the ascending and descending manifolds defines a structure known as the *Morse-Smale complex*. It segments $\mathbf{M}$ into regions with monotone gradient flow behaviors, i.e., all gradients inside a Morse-Smale cell share the same source and destination. The boundaries of the Morse-Smale regions are called *separatrices*, the unique tangent curve connecting a pair of critical points with consecutive indices.

2.2 Discrete Morse theory

Discrete Morse theory developed by Forman [4] aims to describe the concepts and properties of the original Morse theory in a combinatorial fashion. It is widely adopted as the foundation for efficient extraction schema for the Morse-Smale complex. We review only the necessary concepts that will be mentioned later on.

A d-cell is a topological space that is homeomorphic to the closed d-ball $\mathbf{B}^d = \{\mathbf{x} \in \mathbf{E}^d \,|\, \|\mathbf{x}\| \leq 1\}$, where $\mathbf{E}^d$ denotes the d-dimensional Euclidean space. For instance, a vertex is a 0-cell, an edge that connects two vertices is a 1-cell, and a polygon of any shape is a 2-cell. Given a d-cell $\alpha^{(d)}$, we will write this cell as α if the dimension is not necessary for the context. A cell $\alpha^{(d_1)}$ is called a *face* of another cell $\beta^{(d_2)}$, denoted by $\alpha^{(d_1)} \preceq \beta^{(d_2)}$, if $d_1 \leq d_2$ and the 0-cells of $\alpha^{(d_1)}$ are a subset of the 0-cells of $\beta^{(d_2)}$. In this case, we also say that $\beta^{(d_2)}$ is a *co-face* of $\alpha^{(d_1)}$.

A *cell complex* K is a collection of cells such that the intersection between any two cells is either empty or a common face of both. To give some examples, uniform and curvilinear grids as well as triangle and tetrahedral meshes are cell complexes. The d-skeleton of K is the collection of all cells whose dimension is not larger than d. Let $f : K \to \mathbf{R}$ be a function defined on a cell complex K. For any face $\alpha \in K$, the *lower star* of α, denoted by $\mathrm{St}^-(\alpha) = \{\sigma \in K | \alpha \preceq \sigma \wedge f(\sigma) \leq f(\alpha)\}$, is the collection of co-faces of α whose values are not exceeding that of α.

Let $f : K \to \mathbf{R}$ be a function that assigns a value to every cell of K. If for every face $\alpha^{(d)} \in K$ the following holds

$$\left|\left\{\beta^{(d+1)} | \alpha^{(d)} \preceq \beta^{(d+1)} \wedge f(\alpha) \geq f(\beta)\right\}\right| \leq 1, \tag{1}$$

$$\left|\left\{\gamma^{(d-1)} | \gamma^{(d-1)} \preceq \alpha^{(d)} \wedge f(\gamma) \geq f(\alpha)\right\}\right| \leq 1, \tag{2}$$

then f is called a *discrete Morse function*. If a cell $\alpha^{(d)}$ satisfies

$$\left|\left\{\beta^{(d+1)} | \alpha^{(d)} \preceq \beta^{(d+1)} \wedge f(\alpha) \geq f(\beta)\right\}\right| = 0, \tag{3}$$

$$\left|\left\{\gamma^{(d-1)} | \gamma^{(d-1)} \preceq \alpha^{(d)} \wedge f(\gamma) \geq f(\alpha)\right\}\right| = 0, \tag{4}$$

then $\alpha^{(d)}$ is a critical cell with the index d. A *discrete vector* is a pair $\{\alpha^{(d)}, \beta^{(d+1)}\}$ such that $\alpha \preceq \beta$. The collection V of discrete vectors where each cell of K appears in at most in one vector is called a *discrete vector field*. Given a discrete vector field V, a *V-path* is an alternating sequence of cells

$$\alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \alpha_2^{(d)}, \ldots, \beta_\ell^{(d+1)}, \alpha_{\ell+1}^{(d)} \tag{5}$$

such that $\{\alpha, \beta\} \in V$ and $\alpha_i \neq \alpha_{i+1} \preceq \beta_i$. If $\ell \geq 0$ and $\alpha_0 = \alpha_{\ell+1}$ then the path is a *non-trival closed path*. If the discrete vector field V does not contain any such closed path, then V is called a *discrete gradient field* of f. This is analogous to the continuous gradient of a function defined on K. Thus, the V-paths are comparable to the tangent curves in the original Morse theory. A V-path of a discrete gradient field starts or terminates at critical cells. Hence, we can define the *discrete Morse-Smale complex* similarly to the continuous counterpart.

Using discrete Morse theory provides a framework to design algorithms to extract the Morse-Smale complex with higher robustness and reliability. The combinatorial nature of this theory enables handling degenerate cases such as plateau regions that can appear when dealing with data.

2.3 Topological Simplification

The goal of topological simplification is to remove small-scale, spurious features, thereby enabling a meaningful analysis of the data. In the context of Forman's discrete Morse theory, simplification of a discrete Morse function can be done by removing a pair of connected critical points and updating the connections of the neighbors of these two critical points. Alternatively, Forman [4] suggests simplification by reversing the gradient vector field along the path connecting two critical cells. This approach implicitly cancels the two mentioned critical cells while resulting in a new valid discrete Morse function. One can use the *height difference* between two critical cells, which is the absolute difference in the value of these cells, to guide the simplification. The use of the height difference is closely related to the concept of *persistence* [3]. See Günther et al. [6] for a detailed discussion.

3 Accurate Geometry for Morse-Smale Complexes

The term *accurate geometry* was coined by Gyulassy et al. [9] to refer to discrepancies between the geometric embeddings of the continuous and the discrete Morse-Smale complex. The underlying idea is that the geometric embedding of the discrete Morse-Smale complex should converge to its continuous counterpart, if the discretization of the domain becomes infinitely fine. This is not always the case [9, 17], but quite desirable in a number of applications, e.g., where separatrices are the main features of interest. We will discuss the previous work on this topic in the following.

Given a discrete gradient field, the Morse-Smale complex can be constructed straightforwardly using an alternating breadth-first search [18]. The main challenge is rather to compute the discrete gradient field itself from image or volumetric data such that the discrete Morse-Smale complex will align with its continuous counterpart. Several methods have been proposed on this subject [8, 14, 18]. The first provably correct method proposed by Robins et al. [18] solves the problem efficiently and is regarded as the *de facto* standard for this task: given an image or volume data that can be thought of as a complex, this algorithm grows from the vertices to every cell of the complex by partitioning the whole domain into disjoint lower stars of the vertex. For each lower star, the first gradient vector is chosen from the possible vectors based on the *steepest descent* direction. The rest of the vectors in this lower star will be constructed via a procedure called *simple homotopy expansion* by collecting cells with exactly one uncovered face and choosing the first possible pair according to an order. This procedure will be performed until such cells can no longer be found. If the expansion is not possible, a cell will be recorded as critical and the expansion continues to proceed from this cell until every cell in the lower star either belongs to a vector or is critical.

However, Gyulassy et al. [9] and Reininghaus et al. [17] show independently that the choice of the *steepest descent* direction leads to a non-convergent geometric embedding of the separatrices, i.e., inaccurate geometry. This is due to the local and greedy nature of the approach. Consider the following analytic function $f : [0,2]^2 \to \mathbf{R}$

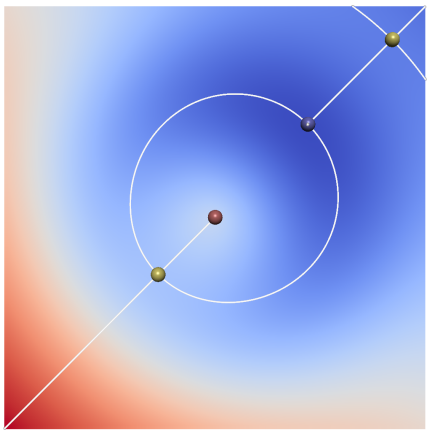

(a) Continuous topology.

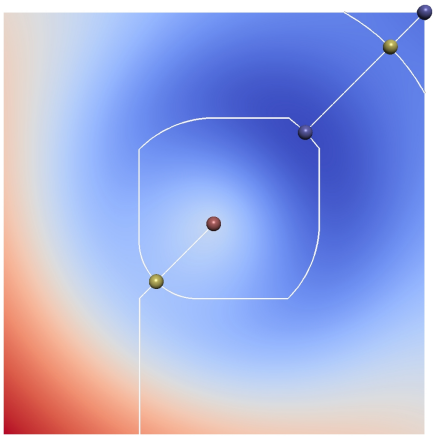

(b) Steepest descent method [18] on the uniform grid.

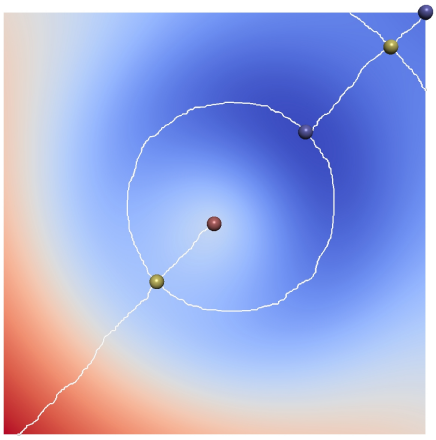

(c) Probabilistic method [9, 17] on the uniform grid.

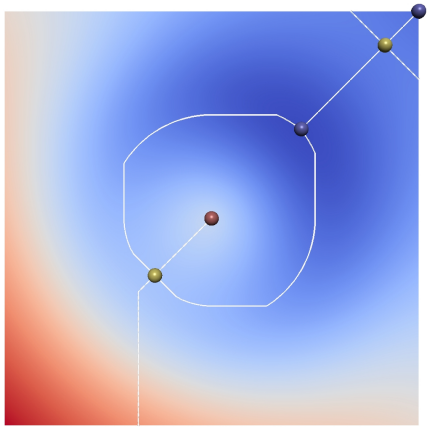

(d) Steepest descent method [18] on a triangle grid obtained by diagonally dividing each cell of a uniform grid once as shown in Figure 8b.

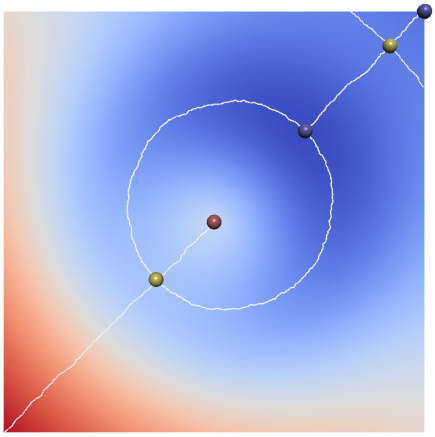

(e) Steepest descent method [18] on a triangle grid with random vertex positions and Delaunay triangulation as shown in Figure 8c.

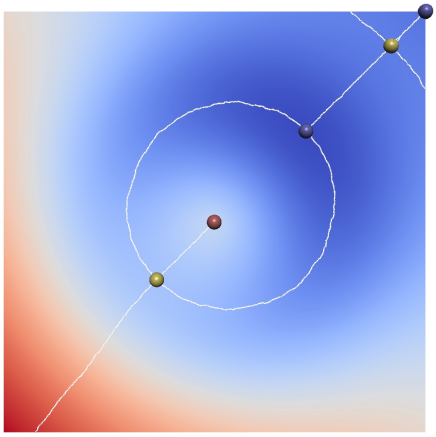

(f) Steepest descent method [18] on a triangle grid derived from the uniform grid following our suggestion, as illustrated in Figure 8d and discussed in Section 6.

Figure 1: The function from Equation (6) has been sampled on different types of grids and the Morse-Smale complex has been extracted using different methods. The goal of *accurate geometry* is to achieve a geometric embedding that is similar to the one from continuous topology. While it is well-known that the steepest descent method is not able to achieve this on uniform grids, we show that it can achieve *accurate geometry* on certain grid types and provide a suggestion of how to convert uniform grids accordingly.

from Reininghaus et al. [17]

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

It represents a circle engraved in a tilted plane. The continuous Morse-Smale complex reveals this shape in Figure 1a. To extract the discrete counterpart, we sampled this function on a 1024×1024 uniform lattice grid and applied the method of Robins et al. [18] using the steepest descent approach. As Figure 1b shows, the result does not capture the circular shape, neither for this nor any other sampling resolution.

Gyulassy et al. [9] and Reininghaus et al. [17] tackle this problem using a probabilistic approach: instead of always choosing the steepest-descent vector, any descending vector can be chosen following a probability that relates back to the continuous gradient. This means, the vectors in the discrete gradient field are more likely to be aligned with the continuous tangent curves. This rather small modification to the original method of Robins et al. [18] produces Morse-Smale complexes which will converge to the continuous version as the sampling resolution increases. The result can be seen in Figure 1c.

Gyulassy et al. [9] introduced a second approach to geometric accuracy that takes into account a larger region when constructing the gradient vector field. It is of global nature and comes with a substantial computational effort. Gyulassy et al. [10, 12] propose methods to create *conforming* Morse-Smale complexes: they are not just based on the input scalar data, but additional information can be supplied by the user to influence the output. Specifically, a discrete vector will be created between two cells, if these cells have equal value with respect to a map L. Using this map, one can encode extra information for the computation or even modification of the Morse-Smale complex, according to the needs of the application. Since a strict adherence to the original Morse-Smale complex is not desired when applying these methods, we will not discuss them further in this paper, but rather leave it to future work to investigate the similarities and differences to the steepest descent method.

4 Topological Inconsistencies

We can use the probabilistic methods [9, 17] from Section 3 to achieve geometrical accuracy. However, these methods come with a significant drawback that has not been reported before: the resulting Morse-Smale complex differs topologically from the one obtained using steepest descent. Most importantly, the connectivity of the separatrices differs, meaning, different critical points are connected

with each other. This holds true for 2D and 3D data sets, and the topological differences have a high persistence, i.e., they can be observed even after radical topological simplification. Besides these topological differences, the geometric embedding of the saddle points changes as well.

For simplicity, we limit our discussions to simplicial and cubical complexes, but they can be extended to the more general cell complex as well.

4.1 Identical Parts

First, we identify which parts of the Morse-Smale complex will *not* differ between different methods. Specifically, we show that any gradient vector field computed based on the lower star has a constant number of critical cells for each dimension and the positions of certain types of critical cells are fixed, independent of the order in which cells are chosen to proceed.

To simplify the discussion, we take advantage of a useful concept introduced by Robins et al. [18]. The *reduced lower star* of a vertex $\mathbf{v}$, denoted by $R(\mathbf{v})$, is the intersection of the lower star of $\mathbf{v}$ with the sphere centered at $\mathbf{v}$ with a small radius r, that is $R(\mathbf{v}) = \mathrm{St}^-(\mathbf{v}) \cap S(\mathbf{v}, r)$. The lower star can be thought of as projecting the cells in $\mathrm{St}^-(\mathbf{v})\backslash\mathbf{v}$ onto the sphere centered at $\mathbf{v}$. With this map, for $d > 0$, a d-cell in $\mathrm{St}^-(\mathbf{v})\backslash\mathbf{v}$ is mapped bijectively to a $(d-1)$-cell in $R(\mathbf{v})$. This map also transforms the algorithm by Robins et al. [18] to an analogous algorithm processing on the reduced lower stars, performing homotopy expansion starting from the steepest edge.

From the definition of the reduced lower star, we can see that it is closed, i.e. every face of a simplex σ of $R(\mathbf{v})$ is also in $R(\mathbf{v})$. Thus, the number of d-cells of $R(\mathbf{v})$ is constant. It was shown by Robins et al. [18] that every critical $(d-1)$-cells processed from $R(\mathbf{v})$ is a critical d-cell in the original complex, except for the initial 0-cell of $R(\mathbf{v})$. It follows that the number of critical d-cells produced by the lower star-based methods is constant for $d > 0$, independent of the order in the expansion step. Furthermore, for any vertices of a complex K, the number of elements in its lower star is constant, and a vertex $\mathbf{v}$ is marked as a minimum if $\mathrm{St}^-(\mathbf{v}) = \varnothing$. This implies that the number and the position of the minima are fixed.

Moving to higher dimensional cells, we break the results into smaller pieces. In the following, we say that a simplex is *covered* if it belongs to a discrete vector or is marked as critical during the homotopy expansion process, and *uncovered* otherwise. We stay in the language of reduced lower stars and show that using homotopy expansion in a 1-skeleton of this structure resulting in a same set of covered 0- and 1-simplices through the following lemmas.

Lemma 1. *The homotopy expansion of the* 1*-skeleton of each* $R(\mathbf{v})$ *covers all of its* 0*-simplices.*

Proof. Assuming that there exists one uncovered 0-simplex α during the expansion process. Then this uncovered simplex connects to a covered 0-simplex via a 1-simplex β. The homotopy expansion will pair $\{\alpha, \beta\}$ during the process. □

Lemma 2. *Given* $R(\mathbf{v})$*, the set of all* 1*-simplices covered by the homotopy expansion in* $R(\mathbf{v})$ *is the same, independent of the starting point of the process.*

Proof. A 1-simplex would not be paired during the homotopy expansion if its two 0-simplices have been paired with other 1-simplices. This behavior can only happen if that 1-simplex is a part of a cycle. Two intersecting cycles would prevent the pairing of their last 1-simplices, marking them as critical cells. All other 1-simplices will be paired during this homotopy expansion as it covers all of the 0-simiplices owing to the result of Lemma 1. It implies that a 1-simplex is unpaired if and only if it is maximal in a cycle. This defines the fixed set of covered 1-simplex by homotopy expansion of the 1-skeleton of $R(\mathbf{v})$. □

The result of Lemma 2 concerns only up to the coverage of the 1-simplices. The following lemma ensures that after performing the homotopy expansion on the 1-skeleton of a complex, the result gradient field of that complex is fixed.

Lemma 3. *For each* $R(\mathbf{v})$*, the homotopy expansion of its* 1*-skeleton can be done before expanding into* 2*-simplices without changing the result gradient field.*

Proof. Supposing that we completed expanding in the 1-skeleton of $R(\mathbf{v})$. Assigning the first vector $\{\alpha^{(1)}, \beta^{(2)}\}$ requires that $\alpha^{(1)}$ is the only face of $\beta^{(2)}$ that is unpaired during the expansion of the 1-skeleton. Thus, no vector $\{\gamma^{(0)}, \alpha^{(1)}\}$ is chosen during the first expansion since this would imply the expansion had not been finished. □

From Lemma 3, mapping to the original lower star, we get the following result.

Theorem 1. *For any vertex* $\mathbf{v}$ *of a complex* K*, the gradient vector fields between the* d*- and* $(d+1)$*-cells in* $St^-(\mathbf{v})$ *are fixed for* $d \geq 2$*, independent of the order of processing the cells.*

In the arguments for Lemma 2, we can see that the critical 1-simplices which are unpaired during the homotopy expansion of the 1-skeleton will also not be paired during the expansion of 2-skeleton of $R(\mathbf{v})$, and their positions will also be fixed. From Theorem 1, we can conclude the same for higher dimensional cells. The following results summarize our observations.

Corollary 1. *For* $d \geq 2$*, the position of critical* d*-cells in* $St^-(\mathbf{v})$*, for every vertex* $\mathbf{v}$*, are independent of the chosen order in the computing process of the gradient vector field.*

Corollary 2. *For* $d \geq 2$*, each separatrix connecting a pair of* d*- and* $(d+1)$*-critical cells retains its connectivity, independent of the chosen order in the computing process of the gradient vector field.*

We conclude that the positions of the minima and maxima in 2D and 3D data sets are identical between the different methods. The same holds for the 2-saddle points in 3D data sets. We further conclude that the connectivity of the separatrices between 2-saddles and maxima in 3D data sets is identical between the different methods.

4.2 Different Parts

The saddle points (1-saddle points in the case of three dimensions) did not appear in the results of Section 4.1. Indeed, with different ways of choosing the first vector in the computation (steepest descent or probabilistically), they can have different positions. We demonstrate this behavior by applying the steepest descent method and the probabilistic method on a scalar field given by a small randomly generated 4×4 matrix

$$A = \begin{pmatrix} 9 & 8 & 7 & 13 \\ 1 & 6 & 10 & 0 \\ 12 & 14 & 2 & 4 \\ 3 & 11 & 5 & 15 \end{pmatrix}. \tag{7}$$

The differences in the positions of the saddle points are shown in Figure 2. We can see from the two figures that the positions of the minima and the maximum are fixed whereas three saddle points changed their positions. This is in agreement with the theoretical results stated earlier. It should be noted that while the saddle points can move in different iterations of the probabilistic approach, their movements are restricted to their respective lower stars.

The changes in the position of the saddle points further cause separatrices to connect to different critical points as illustrated in Figure 3. We can easily see the top-most saddle connects to a

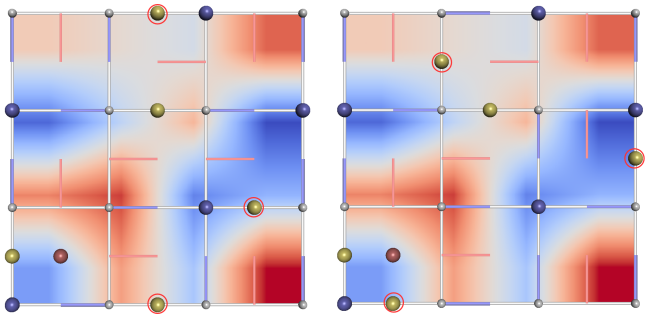

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 2: Given the same input data from Equation (7), the two different methods for computing the discrete gradient field can result in different positions for saddle points as indicated by the red circles. The white spheres, white lines, and squares represent the 0-cells, 1-cells, and 2-cells. The blue and red segments illustrate the pairings between the 0- and 1-cells as well as the 1- and 2-cells, respectively.

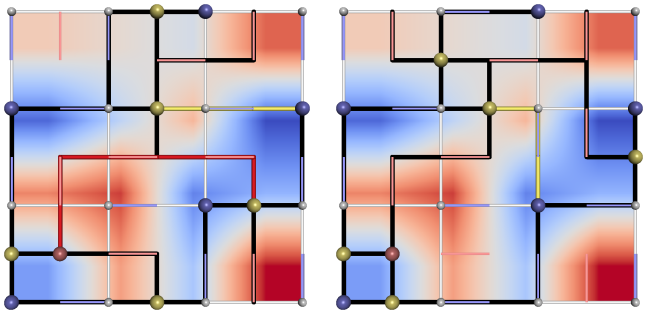

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 3: Given the same input data from Equation (7), the two different methods for computing the discrete gradient field can result in a different connectivity of the separatrices. The separatrices connecting to the same critical points are colored black. The separatrices connecting to different minima are highlighted in yellow. The red separatrix only presents for the steepest descent method, and is missing from the output of the probabilistic method.

different minimum, and that another saddle looses its connection to a maximum. This observation can be explained by the difference in the gradient vector fields. Notably, the result of Corollary 2 only holds for separatrices in higher layers, whereas this example shows separatrices between 0-cells, 1-cells, and 2-cells, for which Corollary 2 does not hold.

As the changes already happen in two dimensions, it is straightforward that these changes can also appear in three dimensions. Indeed, consider the scalar field given by the small randomly generated tensor

$$B = \left(\begin{pmatrix} 14 & 18 & 9 \\ 6 & 24 & 15 \\ 23 & 4 & 7 \end{pmatrix}, \begin{pmatrix} 10 & 20 & 0 \\ 21 & 1 & 11 \\ 2 & 26 & 25 \end{pmatrix}, \begin{pmatrix} 12 & 13 & 8 \\ 19 & 22 & 3 \\ 17 & 5 & 16 \end{pmatrix} \right). \tag{8}$$

The inconsistencies of the positions of the 1-saddles (green) together with the connectivity of the separatrices are illustrated in Figure 4.

Next, we show that these issues also appear in larger and smooth data sets. We constructed a data set of size 40×40 from Equation (7) using bilinear interpolation. We can see in Figure 5 that the connectivity between critical points is different as highlighted in the figure.

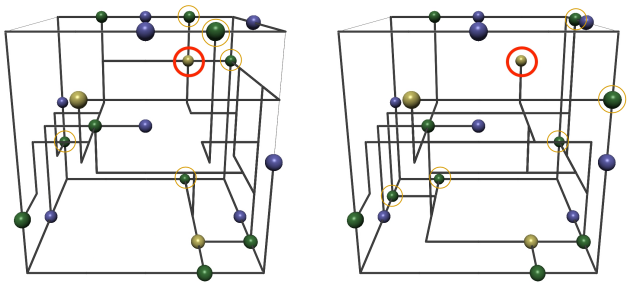

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 4: Morse-Smale complexes of the simple 3D data set from Equation (8) computed using different methods. Many 1-saddle points (green) have different positions (marked by the orange circles). This also leads to a different connectivity of the separatrices (marked by the red circle).

Data set	# cases with positional change	# cases with connectivity change
Random 4×4	603	235
Random 8×8	996	942
Random 16×16	1000	1000
Interpolated 8×8	567	533
Interpolated 16×16	581	555

Table 1: Number of cases with differences in the positions of the saddle points and in the connectivity of the separatrices between the two methods when applied to 1000 randomly generated 2D data sets.

This shows that topological changes appear in smooth data sets as well.

These differences do not only happen in a few selected scalar fields. We randomly generated 1000 2D data sets of size 4×4, 8×8, and 16×16. Then we counted for how many of those data sets we can observe a positional change of at least one saddle, and a connectivity change of at least one separatrix. The results are summarized in Table 1. We can see that these changes appear in almost all data sets of size 8×8 and larger, and in many smaller data sets as well, even after interpolating the 4×4 data sets to higher resolutions. Automatically detecting a connectivity change of a separatrix can be done as follows: for each separatrix in a Morse-Smale complex, identify the two lower stars where it originates/ends. Then try to find a separatrix in the other Morse-Smale complex which originates/ends in the same pair of lower stars. If it cannot be found, we record a connectivity change. This procedure is rooted in the fact that the positional change of a saddle point is restricted to the respective lower star and hence its separatrices need to originate/end there.

We repeat this experiment on a noisy data set, which we obtained by adding noise to the previous data set. We find 909 critical points in this data set. We apply a topological simplification with a persistence threshold of 45% of the data range. The simplified Morse-Smale complexes then have 17 critical points. We can see in Figure 6 that connectivity changes still persist under these conditions – even at this very high level of persistence.

To show that these differences also exist in real-life data sets, we performed our experiments on the Hurricane Isabel data set.[1] The data set represents different atmospheric variables over different

[1]This data set can be downloaded at https://www.earthsystemgrid.org/dataset/isabeldata.html

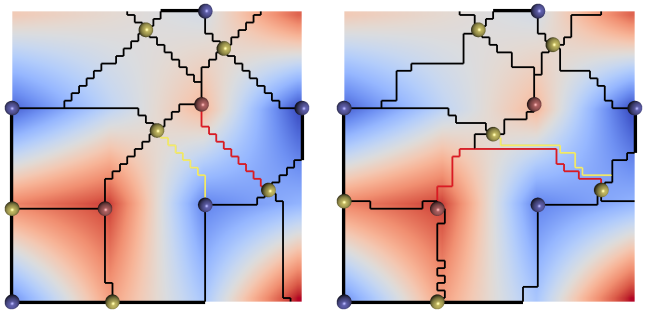

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 5: The data set from Equation (7) has been bilinearly interpolated on a 40×40 uniform grid. The Morse-Smale complexes have been extracted using the different methods and we observe connectivity changes for some separatrices (highlighted in red and yellow).

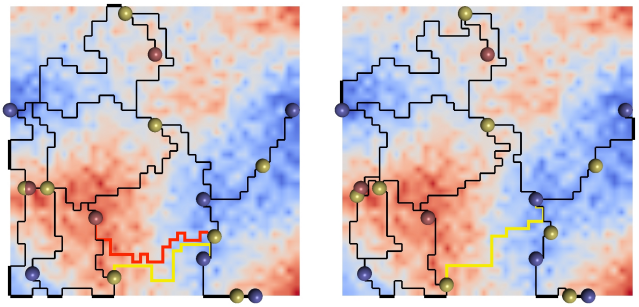

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 6: Noise has been added to the data set from Figure 5. The Morse-Smale complexes produced by the two methods show differences in separatrix connectivity (red and yellow) even at a very high level of topological simplification corresponding to 45% of the data range.

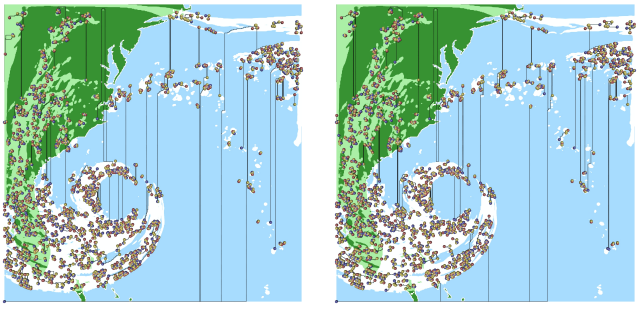

(a) Steepest descent [18] (b) Probabilistically chosen [9, 17]

Figure 7: The Morse-Smale complex of the Hurricane Isabel data set has been extracted using the two different methods. Only the separatrices with differences in connectivity are shown together with their endpoints. This is about 10% of all separatrices.

time steps. We chose the 41-st time step and the *total cloud* variable. Further, we extracted the 50-th slice in z-direction and use it as a 2D data set with a 500×500 uniform grid. The differences in the results of the two methods can be seen in Figure 7. A total of 2087 out of 20159 separatrices with different connectivity were found. Additionally, 10135 critical points were detected, among them, 983 out of 5067 saddle points have different positions.

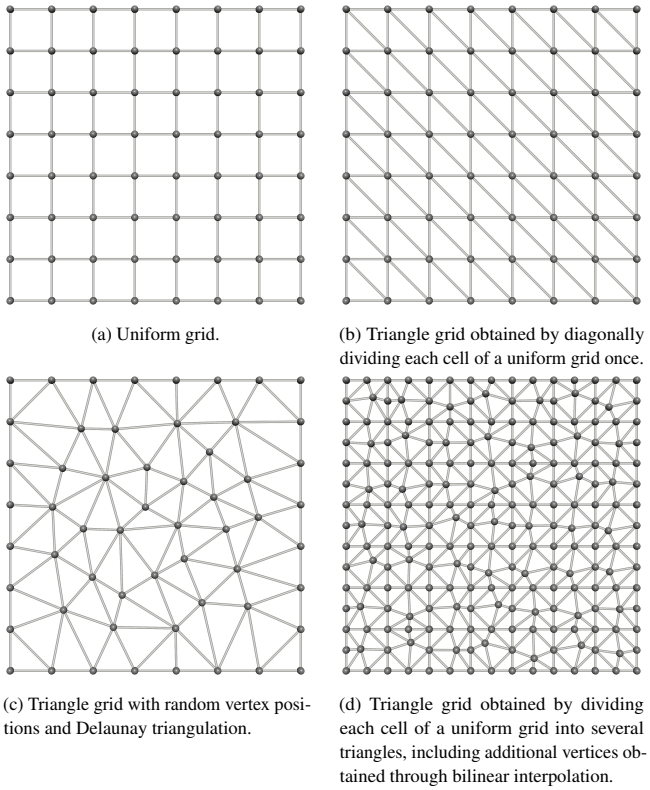

(a) Uniform grid. (b) Triangle grid obtained by diagonally dividing each cell of a uniform grid once.

(c) Triangle grid with random vertex positions and Delaunay triangulation. (d) Triangle grid obtained by dividing each cell of a uniform grid into several triangles, including additional vertices obtained through bilinear interpolation.

Figure 8: Different grid types as used in this paper.

Summary We have shown that the probabilistic methods of Gyulassy et al. [9] and Reininghaus et al. [17] are successful in producing accurate geometry, but fail at faithfully reproducing the topology of the method by Robins et al. [18], which follows the steepest descent and is regarded as the *de facto* standard for computing discrete gradient fields. These discrepancies arise in many different scenarios and may affect the conclusions drawn from topological analyses of data sets.

This leads us in the following to investigating alternatives to the probabilistic methods in order to find solutions that provide accurate geometry and consistent topology at the same time.

5 Effect of Sampling Methods

In this section, we delve into the improvement in geometry that can be achieved by employing the steepest descent method of Robins et al. [18] on different grid types. To this end, we refer the reader to Figure 8 where we illustrate the types of grids that will be discussed in this and the following section.

The pre-assumption of the probabilistic methods [9, 17] is that the data is given on a uniform grid. With a uniform grid, a cell has a limited number of higher dimensional cells to pair with during the calculation of the gradient vector field: a maximum of 4 choices for the 2D case, and a maximum of 6 choices for the 3D case. These numbers remain constant even when we increase the resolution of the grid. Furthermore, due to the nature of the grid, the directions of the grid edges are either vertical or horizontal, which limits the geometric expression of the separatrices drastically. Hence, the steepest descent method [18] is not able to capture the geometry of the tangent curves on such grid.

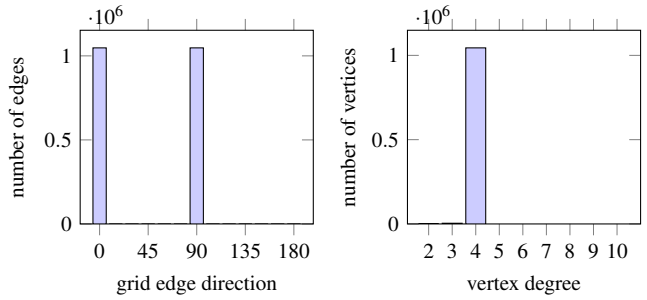

(a) Uniform grid.

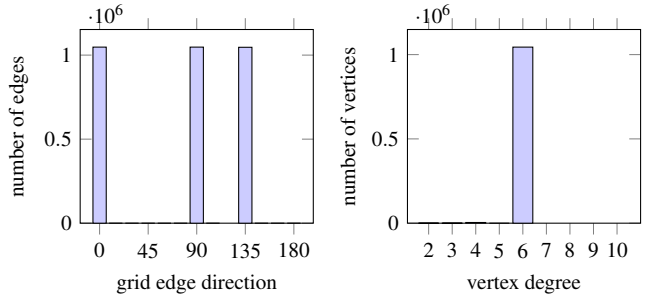

(b) Triangle grid obtained by diagonally dividing each cell of a uniform grid once as shown in Figure 8b.

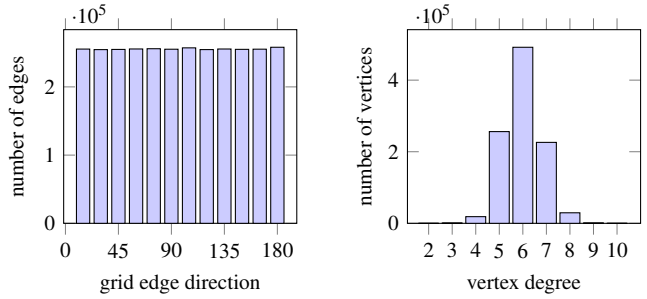

(c) Triangle grid with random vertex positions and Delaunay triangulation as shown in Figure 8c.

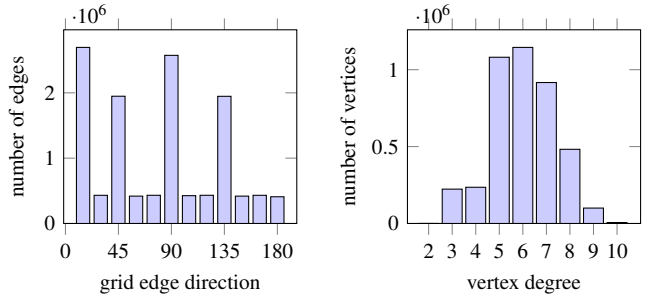

(d) Triangle grid derived from the uniform grid following our suggestion, as illustrated in Figure 8d and discussed in Section 6.

Figure 9: Statistics of different grids. Shown are the histogram of the grid edge directions (left) and the histogram of the vertex degrees (right).

It is interesting to investigate the output of the steepest descent method on other types of grids. Specifically, we want to investigate triangle grids as they allow for a wider variety in the directions of the grid edges as well as a higher degree of the vertices.

A straightforward approach is to divide each cell of a uniform grid into two triangles. This fairly common choice of grid is illustrated in Figure 8b. It allows all non-boundary vertices to have 6 possible choices of direction instead of just 4. Furthermore, it introduces diagonal grid edges. Figure 9b shows the distributions of the vertex degrees and of the grid edge directions of a 1024×1024 grid of this kind. The grid edge directions are measured as the angle to the x-axis. We used such a triangle grid to sample the function of Equation (6) and applied the steepest descent method. The result is shown in Figure 1d. We observe a slightly rounder shape for the circular part, but overall the geometric embedding is still quite different compared to the continuous case from Figure 1a.

Next, we investigate a triangle grid with randomly placed vertices as shown in Figure 8c. It is built by uniformly placing vertices at the boundary and then randomly placing the remaining vertices inside using a Poisson disc sampling. A Delaunay triangulation is then applied to obtain the mesh. The statistics of this grid are shown in Figure 9c and reveal a high diversity in terms of grid edge directions and vertex degrees. Again, we sampled Equation (6) onto a 1024×1024 version of this grid and computed the discrete gradient using the steepest descent method. The resulting Morse-Smale complex (Figure 1e) comes remarkably close to the continuous version (Figure 1a). In fact, in all our experiments, this type of grid exhibited the most accurate geometry of all grid types.

We repeated our experiments with another function g sampled in the domain $[-2,2]^2$

$$g(x,y) = \sin(xy)\cos(x+y). \tag{9}$$

The results are shown in Figure 11. Our previous observation is confirmed: the steepest descent method does produce accurate geometry on grids with high diversity of the grid edge directions and vertex degrees.

We conclude that when extracting the Morse-Smale complex in certain applications, if possible, a grid should be chosen that allows for a higher "degree of freedom" for the edge directions and vertex degrees. The method by Robins et al. [18] will then be able to extract accurate geometry.

6 Suggestion for Uniform Grids

It is not always possible to choose a type of grid for an application. Most applications come with pre-defined grids and some of those will be uniform grids, or very similar to uniform grids. In this section, we propose a simple solution to achieve acceptable geometry for the Morse-Smale complex given data sampled on a uniform grid while preserving its topology through the steepest descent direction.

Our proposal draws inspiration from the observations in Section 5: better geometry can be achieved by allowing a higher number of options to choose from during the computation of the discrete vector. Thus, given a data set sampled on a uniform grid, we create a triangle mesh using well-chosen auxiliary vertices to obtain the desired grid characteristics.

Our goal is to convert each uniform grid cell into more than two triangles. To do so, we add one auxiliary vertex in the middle of each grid edge, and another auxiliary vertex randomly into the interior of each grid cell. The data values for these new vertices are obtained using bilinear interpolation. A Delaunay triangulation within each grid cell then yields the final triangle mesh. An example is shown in Figure 8d. Due to the random nature of the new points, spurious features can be created. We apply an ε-simplification to cancel these unwanted features. In our experiments, a persistence threshold of at most 0.001% of the data range is applied.

Figure 9d shows the characteristics of this grid. We can see that this grid covers all possible grid directions, albeit to a varying degree. Some directions clearly dominate, which stems from the fact that this triangle grid inherits the large number of vertical and horizontal grid

edges from the original uniform grid. Nevertheless, the histograms show a vast improvement over the uniform grid.

We applied this procedure to the uniform grids of size 1024×1024 that had been originally used to sample Equations (6) and (9). The resulting triangle meshes have then be used to compute the Morse-Smale complexes via the steepest descent method. The results are shown in Figures 1f and 11f, respectively. Note how the geometric embedding is much more aligned with the continuous case in comparison to the steepest descent method applied to the uniform grid. In particular, the circle in Figure 1f is well represented. However, we can also see that the diagonal separatrix in the lower-left corner of the same figure does not end at the corner of the domain, but only nearby.

It remains to see whether this new type of grid gives rise to a Morse-Smale complex that is equivalent to the one obtained from the uniform grid – at least after an ε-simplification. We have reason to believe that this is true, since we apply the steepest descent method in both cases. Further, the original vertices are kept in the new grid, and the bilinearly interpolated values of the new vertices will always fall within the range of the original values. Thus, the original minima remain the same, the saddle points and maxima are in the now refined lower star of the original vertex. Although spurious minima, saddles, or maxima can be created due to the randomness of our suggestion, they tend to be short-lived and can be canceled at a very low persistence threshold as mentioned. With more carefully chosen and optimized methods, new points can be added without creating new criticals.

An experimental validation is shown in Figure 10, where we show the Morse-Smale complexes obtained using our suggested method and using the continuous approach. These should be compared also to Figure 5 where the steepest descent and the probabilistic method have been applied to the original uniform grid. We can clearly see that our method avoids the topological issues of the probabilistic method, while achieving a good geometric embedding.

Our discussion in this section is limited to 2D cases. Given a cubical grid, we suspect that a *face-centered subdivision*, similar to the one by Carr et al. [1], could achieve similar results to ours. We left this for further investigations.

Increased computational effort Our suggestion leads to increased memory requirements and computation times. Suppose that we add a points at each edge and b points inside each cell. For a grid of size $n \times n$, we added $2a(n-1)n$ points at the edges and $b(n-1)^2$ points in the cells, a total of $(n-1)((2a+b)n-b)$ points in addition to the original n^2 points, which is a substantial memory cost, especially considering that the newly created unstructured grid requires explicit handling of the connectivity whereas this could be handled implicitly for the original structured grid. The computation time required to create such a triangulation is also non-negligible: it takes in the order of tens of seconds for larger grids such as 2048^2.

Despite having higher computation time and memory requirements, our method is able to provide Morse-Smale complexes with accurate geometry and topology while being algorithmically less involved than the probabilistic variants.

7 Conclusion

We discussed the different aspects that can affect the quality of the Morse-Smale complex in terms of geometry and topology. We have shown that the undesired shape of the Morse-Smale complex using the traditional methods comes from the low number of choices for the vector paring during the computation. By allowing a higher number of options for the vectors through a more optimized grid structure, we can attain a Morse-Smale complex that closely resembles its continuous version. It is thus encouraged, if possible, to take this into account when sampling the data.

However, not all applications allow such sampling methods and we need to employ methods for accurate geometry. While it is known that the existing methods produce geometrically different embeddings for the separatrices, we have shown that there are also unintended consequences in terms of topological changes. We have shown that the 1-saddles can take different positions, which can lead to significant alterations in the overall structure of the Morse-Smale complex. However, we did not investigate all methods that fall under the category of *accurate geometry*, e.g., the work by Gyulassy et al. [10]. We leave this to future work.

We also proposed a method to achieve better geometry for the Morse-Smale complex given data sampled on a uniform grid, drawing inspiration from our observations. Compared to the other methods aimed at tackling the geometry problem, our suggestion is significantly less complicated. This method required no modification on the algorithm side, as we only added more information to the given data. Therefore, our suggestion does not increase the complexity of these methods, albeit it does increase the memory requirements and computation times. Despite its simplicity, our method produced comparable results in terms of geometry while retaining the topological features inherited from the steepest descent method.

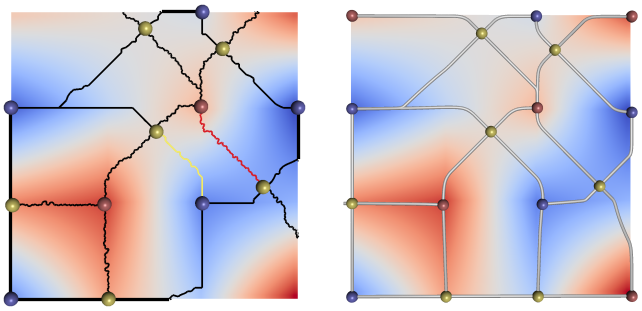

(a) Steepest descent method applied to our suggested triangle grid.

(b) Continuous Morse-Smale complex.

Figure 10: The discrete Morse-Smale complex obtained using the steepest descent method applied to our suggested triangle grid in comparison to the continuous Morse-Smale complex. The latter has a slightly different handling of the cases at the boundary, but otherwise the topologies coincide. Compare these results to Figure 5.

Acknowledgments

This work was supported through grants from The Swedish Research Council (Vetenskapsrådet, project 2020-05461), and the Swedish e-Science Research Centre (SeRC). The visualizations were created using the Inviwo framework [13]. The Hurricane Isabel data set was produced by the Weather Research and Forecast (WRF) model, courtesy of NCAR, and the U.S. National Science Foundation (NSF).

References

[1] H. Carr, T. Moller, and J. Snoeyink. Artifacts caused by simplicial subdivision. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):231–242, 2006. doi: 10.1109/TVCG.2006.22 8

[2] F. Cazals, F. Chazal, and T. Lewiner. Molecular shape analysis based upon the Morse-Smale complex and the connolly function. In *Proceedings of the nineteenth annual symposium on Computational geometry*, pp. 351–360, 2003. 1

[3] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511 – 533, 2002. 2

[4] R. Forman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire*, 48:B48c–35, 2002. 1, 2

[5] D. Günther, J. Reininghaus, S. Prohaska, T. Weinkauf, and H.-C. Hege. Efficient computation of a hierarchy of discrete 3d gradient vector fields. In *Topological Methods in Data Analysis and Visualization II*, pp. 15–30. Springer, 2012. 1

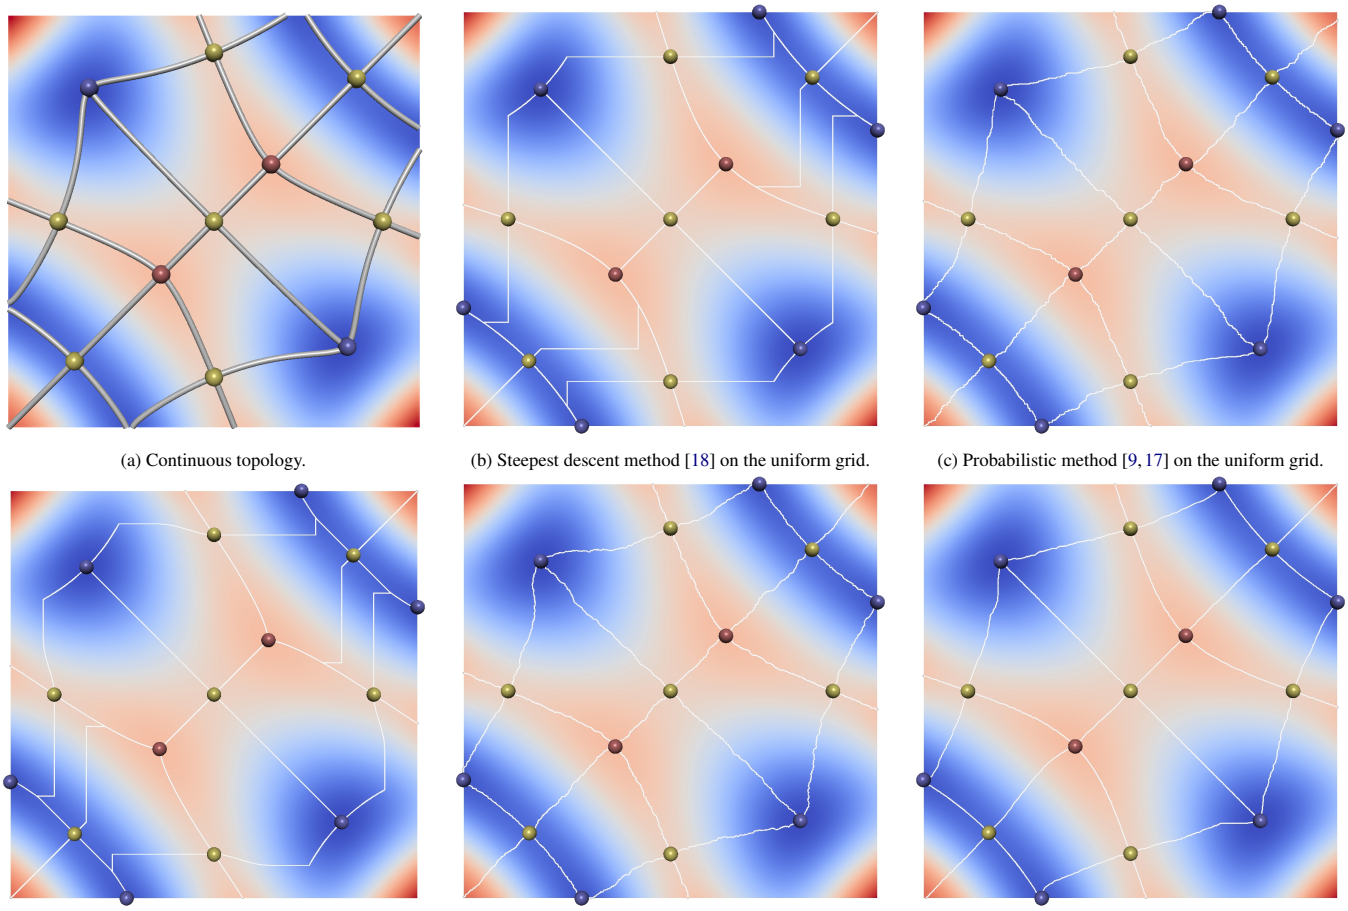

(a) Continuous topology.

(b) Steepest descent method [18] on the uniform grid.

(c) Probabilistic method [9, 17] on the uniform grid.

(d) Steepest descent method [18] on a triangle grid obtained by diagonally dividing each cell of a uniform grid once as shown in Figure 8b.

(e) Steepest descent method [18] on a triangle grid with random vertex positions and Delaunay triangulation as shown in Figure 8c.

(f) Steepest descent method [18] on a triangle grid derived from the uniform grid following our suggestion, as illustrated in Figure 8d and discussed in Section 6.

Figure 11: The function from Equation (9) has been sampled on different types of grids and the Morse-Smale complex has been extracted using different methods. Qualitatively, we see the same results here as in Figure 1, namely that *accurate geometry* can be achieved with different methods, most notably also with the steepest descent method if the right kind of grid is provided.

[6] D. Günther, J. Reininghaus, H.-P. Seidel, and T. Weinkauf. Notes on the simplification of the Morse-Smale complex. In P.-T. Bremer, I. Hotz, V. Pascucci, and R. Peikert, eds., *Topological Methods in Data Analysis and Visualization III*, Mathematics and Visualization, pp. 135–150. Springer, 2014. 2

[7] D. Günther, H.-P. Seidel, and T. Weinkauf. Extraction of dominant extremal structures in volumetric data using separatrix persistence. *Computer Graphics Forum*, 31(8):2554–2566, December 2012. 1

[8] 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*, 14:1619–1626, 2008. doi: 10.1109/TVCG.2008.110 2

[9] A. Gyulassy, P.-T. Bremer, and V. Pascucci. Computing Morse-Smale complexes with accurate geometry. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2014–2022, 2012. doi: 10.1109/TVCG.2012.209 1, 2, 3, 5, 6, 9

[10] A. Gyulassy, P.-T. Bremer, and V. Pascucci. Shared-memory parallel computation of Morse-Smale complexes with improved accuracy. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1183–1192, 2019. doi: 10.1109/TVCG.2018.2864848 3, 8

[11] A. Gyulassy, M. Duchaineau, V. Natarajan, V. Pascucci, E. Bringa, A. Higginbotham, and B. Hamann. Topologically clean distance fields. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1432–1439, 2007. doi: 10.1109/TVCG.2007.70603 1

[12] A. Gyulassy, D. Günther, J. A. Levine, J. Tierny, and V. Pascucci. Conforming Morse-Smale complexes. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2595–2603, 2014. doi: 10.1109/TVCG.2014.2346434 1, 3

[13] D. Jönsson, P. Steneteg, E. Sundén, R. Englund, S. Kottravel, M. Falk, A. Ynnerman, I. Hotz, and T. Ropinski. Inviwo - a visualization system with usage abstraction levels. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3241–3254, 2019. doi: 10.1109/TVCG.2019.2920639 8

[14] T. Lewiner, H. Lopes, and G. Tavares. Optimal discrete Morse functions for 2-manifolds. *Comput. Geom. Theory Appl.*, 26:221–233, November 2003. doi: 10.1016/S0925-7721(03)00014-2 2

[15] J. Milnor. *Morse Theory*. Princeton University Press, 1963. 1

[16] M. Morse. *The calculus of variations in the large*. Colloquium publications, American mathematical society 18. The American mathematical society, New York, 1934. 1

[17] J. Reininghaus, D. Günther, I. Hotz, T. Weinkauf, and H. P. Seidel. Combinatorial gradient fields for 2D images with empirically convergent separatrices. *arXiv preprint arXiv:1208.6523*, 2012. 1, 2, 3, 5, 6, 9

[18] V. Robins, P. Wood, and A. Sheppard. Theory and algorithms for constructing discrete Morse complexes from grayscale digital images. *Pattern Analysis and Machine Intelligence, IEEE Trans. on*, 33(8):1646–1658, aug. 2011. 1, 2, 3, 4, 5, 6, 7, 9

[19] J. Sahner, T. Weinkauf, N. Teuber, and H.-C. Hege. Vortex and strain skeletons in eulerian and lagrangian frames. *IEEE Transactions on Visualization and Computer Graphics*, 13(5):980–990, September - October 2007. 1

[20] S. Smale. On gradient dynamical systems. *Annals of Mathematics*, pp. 199–206, 1961. 1

[21] T. Weinkauf and D. Günther. Separatrix Persistence: Extraction of salient edges on surfaces using topological methods. *Computer Graphics Forum (Proc. SGP '09)*, 28(5):1519–1528, July 2009. 1