A Scale-Space Approach to the Morphological Simplification of Scalar Fields

L. Rocca¹ and E. Puppo²

¹IGAG - CNR, Milan, Italy ²DIBRIS - University of Genoa, Italy

Abstract

We present a multi-scale morphological model of scalar fields based on the analysis of the spatial frequencies of the underlying function. Morphological models partition the domain of a function into homogeneous regions. The most popular tool in this field is the Morse-Smale complex, where each region is spanned by all integral lines that join a minimum to a maximum, with the integral lines departing from saddles as region boundaries. Morphological features usually occur at very different scales, from noise and high frequency details up to large trends at the lowest frequencies. Without some form of multi-scale analysis, only the morphology at the finest scale is visible and explicit in such a model. The most popular approach in the literature is the filtration provided by persistent homology, a method that combines the amplitude values of critical points with the topology of the sublevel sets of the function. We propose the adoption of an alternative filtration method, based on the analysis of the deep structure of the linear scale-space of the function. To retrieve an adequately fine-grained ranked sequence of pairs of critical points that vanish through the scales, we adopt a continuous representation of the Scale-space that overcomes the limits of discrete scale-space approaches. This sequence provides a progressive simplification of the Morse-Smale complex, resulting in a progressive multi-scale model of the morphology that always refers to the geometry of the original function, which is not changed by our model. We apply our method to digital elevation models, with results providing a multi-scale representation of the network of ridges and valley lines that joins peaks, pits and passes and divide the land into mountains and basins.

1. Introduction

According to Morse theory [Mil65], the morphology of a scalar function f can be characterized by its critical points together with two partitions of its domain into stable and unstable submanifolds that are centered at maxima and minima of f – mountains and basins in the bivariate case. Under proper conditions, the overlay of such partitions forms the *Morse-Smale complex*, whose cells have their vertices at the critical points of f and their edges at *separatrices* connecting them – ridges and valleys in the bivariate case.

Approaches based on Morse theory have found applications in several domains, such as geosciences, medical imaging, computational fluid dynamics, material sciences, chemistry, just to mention a few. However, real data are usually noisy, with a large number of spurious critical points. Several methods have been proposed in the literature, which try to analyze the input at different scales to extract the important features that characterize the underlying function.

One popular approach is based on *topological persistence*: the sublevel sets of the function are analyzed, and their topological changes, which occur at the critical points, are tracked. Eventually, a ranked sequence of pairs of critical points is obtained, which can be used to select the most relevant features, and also to simplify the Morse-Smale complex describing the morphology of the function

[ELZ02, EHZ03]. Broadly speaking, topological persistence operates on the amplitude of the signal, since it analyzes its sublevel sets. In spite of its popularity, this kind of analysis can be prone to noise; in particular, impulse noise consisting of isolated outliers with large amplitude may corrupt the results [RKG*11]. Here, we overcome the limitations of this approach by performing an alternative analysis, which works on the frequency of the signal.

The scale-space is an approach to the multi-scale analysis of signals, originally developed in the image processing literature and used in several low-level tasks in computer vision [Lin94]. In this case, the initial function f undergoes a diffusion process, providing a family of progressively smoother functions, called the scalespace. The deep structure of the scale-space encodes the morphological structure of the functions in this family at the different scales. In particular, tracking the critical points through the scales provides another ranked sequence of pairs, which is similar in nature, but different from the one generated with topological persistence. In this case, the analysis occurs in the frequency domain, resulting more robust to noise. However, the critical points drift across the domain while the function undergoes diffusion, hence tracking them may also be problematic [RKG*11]. Here, we overcome the limitations of the classical layered model of the scalespace, by adopting a continuous model.

In summary, we propose a method for simplifying the Morse-Smale complex of a scalar function, which is based on the analysis of the deep structure of its scale-space. We rely on a continuous, piecewise-linear model of the scale-space [RP13], which supports a robust tracking of the critical points through the scale, and a ranking of pairs of critical points that annihilate together. Such ranking is then used, akin to [EHZ03], to progressively simplify the Morse-Smale complex computed on the original function. This provides a fine-grained hierarchy of complexes that are faithful to the morphology of the original terrain, and can be efficiently queried at any desired scale. We present applications to topographic maps, providing a multi-scale representation of the network of ridges and valleys that join peaks, pits and passes. The method is rather general and can be directly applied to bivariate functions in other contexts. Extensions to higher dimensions are also possible.

2. Related work

Morse theory [Mil65] is a classical mathematical theory that relates the topology of the domain of a function to its critical points; see [Mat02] for a modern account. Morse and Morse-Smale complexes have been long used to describe the morphology of discrete scalar fields in a variety of applications; see [BDF*08] for a survey. There exist two main approaches to the computation of such complexes from sampled data: in the piecewise-linear setting, the underlying smooth function is approximated with a piecewise-linear signal defined on a simplicial complex, which discretizes the domain of the function [EH10]; in the discrete Morse theory [For98], the domain is also discretized with a simplicial complex, but just the incidence graph of cells is considered, and discrete gradients are associated to its arcs. Both approaches have advantages and disadvantages. In this work, we adopt the piecewise-linear setting.

Scale-space theory was proposed independently by Koenderink [Koe84] and Witkin [Wit83], thoroughly studied and formalized later on by Lindeberg [Lin94], and used in countless applications in image processing and computer vision; see [Pri23] for an annotated bibliography. The most recent developments of the scale-space theory address its application to the design of neural networks [Lin22].

The scale-space is a one-parameter family of functions, built by applying a diffusion process to the input function. The *deep structure* of the scale-space encodes the evolution of the critical points of the function through the different scales, together with the topological structures that such evolution implies. See [GC95] for a summary of the morphological features encoded in the deep structure. In particular, the graph of saparatrices computed at each scale is equivalent to the Morse-Smale complex of the corresponding function, thus an analysis of the deep structure may provide an evolution of such complex through the scales.

In most cases, a discrete version of the scale-space is assumed, which consists of a sequence of layers, each filtered at a different scale. Finding the deep structure in such a model may be hard: since the diffusion process displaces the critical points, tracking each of them through the scales is prone to errors and noise, and often leads to false or broken trajectories [RKG*11]. Rocca and Puppo [RP13] proposed a continuous model, based on a piecewise-linear discretization of the whole scale-space domain, which supports the

robust tracking of critical points through the scales; this approach was applied successfully to the recognition of fiducial points on range images of faces [DRP15], and to the automatic placement of spot heights on topographic maps [RJP17]. In this work, we adopt the same approach to drive our morphological simplification.

Persistent homology was proposed independently by Edelsbrunner et al. [ELZ02] and by Frosini and Landi [FL99]. It provides a method to rank the importance of critical points of a Morse function, by studying the evolution of the topology of its sublevel sets. Similarly to tracking the critical points through the scale-space, persistent homology also provides a ranked sequence of pairs of critical points. A pair in the scale-space contains critical points that vanish together at some scale, and its ranking measures how long such a pair "survives" through the scales (i.e., it is resilient to filtering); while in persistent homology the pairing is related to the evolution of the sublevel sets and the score of a pair is the difference in amplitude of the function at the paired points. The two approaches thus differ both in forming pairs and in ranking them.

Morphological simplification can be performed by using the ranked sequence of pairs to simplify the Morse-Smale complex by progressively merging adjacent cells [EHZ03]. We build our multi-scale morphological model in a similar fashion, by relying on the sequence generated from the scale-space. Several other works follow a different approach, using topological persistence as a guide to modify the input function directly, in order to obtain a function with a simpler morphology [BS98, DDM^{*}03, DS18, FKM20, ID17, SFID21]. However, as the function is modified, the separatrices describing its morphology drift across the domain, thus describing a morphology that is no longer referred to the original function in a metric sense. We rather prefer to maintain the simplified morphology strictly referred to the original function: our model is a hierarchy, in which a simplified complex is bounded by a subgraph of the detailed morphology of the original function; we just progressively discard less important critical points and separatrices, by preserving the most representative features of the original function through the scales.

3. Background notions

Morse-Smale theory. Let \mathcal{M} be a smooth manifold of dimension d. A smooth function $f : \mathcal{M} \longrightarrow \mathbb{R}$ is said to be a Morse function if all its critical points are isolated; this is equivalent to say that its Riemannian Hessian does not vanish at critical points. In the following, we will stick to d = 2; the theory holds for higher dimensions too, but this is out of the scope of this work.

Let $p \in \mathcal{M}$ be a minimum of f; we define the *unstable submanifold* (a.k.a, *basin*) of p as the locus of points of \mathcal{M} that lie on integral curves of f emanating from p; for d = 2, each such region is bounded by a set of *separatrices* that are integral curves connecting maxima and saddles. The unstable manifolds form a partition of \mathcal{M} . Similarly, the *stable submanifold* (a.k.a, *mountain*) of a maximum q is the locus of points that lie on integral curves converging at q. The stable submanifolds form another partition of \mathcal{M} , and each of them is bounded by separatrices that connect minima to saddles. If the two partitions intersect transversally, then their overlay

is called a Morse-Smale complex. Hereafter, we will assume f satisfies this property. In the bivariate case, the Morse-Smale complex is described by a planar graph, whose edges are the separatices that connect saddles to maxima and saddles to minima. See [Mat02] for a more thorough formal treatment of this subject.

Scale-space. Let *f* be a Morse-Smale function defined on a twomanifold \mathcal{M} as above. The *linear scale-space* $F_f(p,t)$ of *f* is defined as the solution of the *heat equation*

$$\frac{\partial}{\partial t}F_f = \lambda \Delta F_f,$$

with initial condition $F_f(p,0) = f(p)$, where Δ denotes the Laplace-Beltrami operator with respect to the space variable p, and λ is a constant term tuning the speed of diffusion. So, the scale-space is defined on a three dimensional domain $\mathcal{M}t = \mathcal{M} \times [0, t_{\max}]$, the first two dimensions referring to space. We will use interchangeably the words *scale* and *time* referring to the third dimension. In general, the scale-space F_f is obtained through a diffusion process starting at f. If \mathcal{M} is Euclidean, i.e., $\mathcal{M} \subset \mathbb{R}^2$, the scale-space can be obtained equivalently by convolving f with Gaussian kernels of increasing variance.

A *layer* of the scale-space for a given time \bar{t} is the restriction $f_{\bar{t}} = F_f|_{t=\bar{t}}$. Let p be a critical point of $f_{\bar{t}}$. There exist a maximal smooth *trajectory* $\gamma_p : [t_p^c, t_p^a] \longrightarrow \mathcal{M}t$, with $\overline{t} \in [t_p^c, t_p^a]$, such that $\gamma_p(\bar{t}) = (p, \bar{t})$ and for all other values $\gamma_p(t)$ is a critical point of f_t of the same type of p. Trajectory γ_p describes the evolution of critical point p through the scales; with abuse of notation, when the scale is clear we will refer to p and γ_p interchangeably. The values t_p^c and t_p^a are called the times of *creation* and *annihilation* of critical point p, respectively. If $t_p^c = 0$, then p is an original critical point of the input function, otherwise it is called a newborn. Similarly, if $t_p^a < t_{\text{max}}$, then p vanishes at time t_p^a , annihilating with another critical point; otherwise, it is a survivor at the largest scale. A creation or annihilation of (pairs of) critical points is called a catastrophic event. Each catastrophic event always involves a saddle and either a minimum or a maximum. If we consider the Morse-Smale complexes for all layers of F_f , we find that also separatrices sweep surfaces through the scales; and each separatrix will collapse and/or originate at catastrophic events involving its endpoints. This evolution of the morphological structure of f through the scale is called the deep structure of the scale-space. See, e.g., [FK00] for a more thorough formal treatment of this subject.

4. Morphological simplification

We first describe the piecewise-linear scale-space together with the algorithm to find catastrophic events in its deep structure, and their use to pair and rank critical points (\$4.1). Next, we describe the algorithm to extract the Morse-Smale complex under the piecewise-linear approach (\$4.2). Finally, we describe how to use the ranked set of pairs of critical points to obtain a progressive simplification of the morphology of the input function (\$4.3).

4.1. The piecewise-linear scale space

Let f_0 be a Morse function sampled at a finite set of points in \mathcal{M} , and let $f_1, \dots, f_{t_{\text{max}}}$ be progressively smoothed versions of f_0 , ob-

tained as described in §3. The collection $(f_0, f_1, \dots, f_{t_{\text{max}}})$ provides a *discrete* representation of the scale-space F_f .

Following [RP13], we build a continuous piecewise-linear approximation of F_f . Let M be a triangular mesh approximating \mathcal{M} , having its vertices at the points where f_0 is sampled. The construction of M is out of the scope of this work; in our experiments, we assume f_0 sampled at the nodes of a regular grid discretizing rectangle $\mathcal{M} = [0, w] \times [0, h]$ and the mesh M is built trivially by diagonally splitting each cell of the grid into two triangles; as a consequence, each internal vertex of M has exactly six neighbors.

We take $t_{\text{max}} + 1$ copies of M, namely $M_0, \ldots, M_{t_{\text{max}}}$ and for all $i = 0, \ldots, t_{\text{max}}$ we associate the values of f_i to the corresponding vertices of M_i . We extend f_i to the triangles of M_i by linear interpolation, thus obtaining a piecewise-linear bivariate function over \mathcal{M} interpolating f_i at the sampled points. Such a function constitutes a layer in the scale-space, which is now continuous in the spatial domain, yet discrete in the time domain. See Fig.1 left.

In order to achieve continuity in time, too, we connect corresponding triangles of consecutive layers to form triangular prisms, in which we approximate F_f by bilinear interpolation (linear in the space and time variables, respectively). See Fig.1 center.

Let τ be a triangle of M having vertices a, b, c, and let τ_i, τ_{i+1} be the copies of τ in M_i, M_{i+1} , respectively. Let f_i^a, f_i^b, f_i^c and $f_{i+1}^a, f_{i+1}^b, f_{i+1}^{c}$ be the values of the functions at a, b, c in layers i and i + 1, respectively. We approximate F_f within the triangular prism between τ_i and τ_{i+1} with the following bilinear function:

$$f^{\tau}(\alpha,\beta,t) = (1-t)(\alpha f_i^a + \beta f_i^b + (1-\alpha-\beta)f_i^c) + t(\alpha f_{i+1}^a + \beta f_{i+1}^b + (1-\alpha-\beta)f_{i+1}^c)$$

where (α, β) are the barycentric coordinates of a generic point in triangle τ and $t \in [t_i, t_{i+1}]$ is the time variable. See Fig.1 right. Note that for any *i* and for any time $\bar{t} \in [t_i, t_{i+1}]$ the collection of functions $f^{\cdot}(\cdot, \cdot, \bar{t})$ form a piecewise linear function $f_{\bar{t}}$ on the triangles of M, which provides a virtual layer of the scale-space F_f at time \bar{t} .

Analysis of the deep structure. Given any layer f_t , with $t \in [0, t_{max}]$, its critical points are located at vertices of M and characterized as follows:

- A point *p* is a maximum if $f_t(p)$ is larger than the value of f_t at all neighbors of *p*;
- A point p is a minimum if $f_t(p)$ is smaller than the value of f_t at all neighbors of p;
- A point *p* is a *k*-saddle if, when the neighbors of *p* are traversed in cyclic radial order, the number *s* of times that the values of f_t at them alternate between smaller and larger values with respect to $f_t(p)$ is larger than two. The index of the saddle is k = s/2 - 1(in our experiments, since each sampling point has exactly six neighbors, only 1- and 2-saddles may exist).

In all other cases, *p* is said to be a *regular* point. We are interested in finding the time of creation and annihilation of all critical points, and to pair those that are cerated/annihilated together.

Since the classification of a critical point p depends just on the relation between its value and the values of its neighbors, relevant events can occur only when one of such relations change. We say

submitted to STAG: Smart Tools and Applications in Graphics (2023)



Figure 1: Scheme of the piecewise-linear scale-space. We approximate the domain with a triangle mesh, which is replicated at all layers of the discrete model, obtaining a piecewise-linear representation at each layer (left); Triangles of subsequent layers are connected to form prisms, approximating the volume of the scale-space domain (center); within each prism, the scale space F_f is approximated with a bilinear function, which is linear in the space and time dimensions, respectively (right).

Displacement	Annihilation	Creation			
$\begin{array}{c} (m,r) \rightarrow (r,m) \\ (M,r) \rightarrow (r,M) \\ (s1,r) \rightarrow (r,s1) \\ (s2,r) \rightarrow (r,s2) \\ (s2,r) \rightarrow (s1,s1) \\ (s1,s1) \rightarrow (s2,r) \end{array}$	$ \begin{array}{c} (m,s1) \rightarrow (r,r) \\ (M,s1) \rightarrow (r,r) \\ (m,s2) \rightarrow (r,s1) \\ (M,s2) \rightarrow (r,s1) \end{array} $	$ \begin{array}{c} (r,r) \rightarrow (m,s1) \\ (r,r) \rightarrow (M,s1) \\ (r,s1) \rightarrow (m,s2) \\ (r,s1) \rightarrow (M,s2) \end{array} $			
$(s2,s1) \rightarrow (s1,s2)$					

Table 1: Possible transitions in the state of a pair of vertices connected by a flipping edge, after [RP13]. r: a regular point; M: a maximum; m: a minimum; s1: a 1-saddle; s2: a 2-saddle. Note that for every event a specular one is also possible, for a total of 32 possible events. An example would be: $(M, s2) \rightarrow (r, s1)$ is equivalent to $(s2, M) \rightarrow (s1, r)$.

that an edge pq of M flips at time $t \in [t_i, t_{i+1}]$ if the relation between $f_i(p)$ and $f_i(q)$ is different from the relation between $f_{i+1}(p)$ and $f_{i+1}(q)$. Since the function f_t evolves linearly along pq between t_i and t_{i+1} , we can compute exactly the time of flip as

$$t_{\text{flip}} = \frac{f_i(q) - f_i(p)}{f_i(q) - f_i(p) + f_{i+1}(p) - f_{i+1}(p)} + i.$$
 (1)

We identify all flips that occur within our piecewise-linear scalespace and sort them by time. Note that a flip is a relevant event only if it changes the labels of the endpoints of the edge; in this case, we have three possible types of flip: *displacements*, which allow us to trace the evolution of a critical point through the scales; *annihilations*, in which two critical points end their trajectories; and *creations*, in which two newborn critical points start their trajectories. Table 1 reports the relevant transitions caused by edge flips.

The analysis of flips generates a set of critical points (either original or newborn), where for each point p we record: the times t_p^c of its creation (possibly $t_p^c = 0$ if p is original) and t_p^a of its annihilation (possibly $t_p^a = t_{\text{max}}$ if p is a survivor); and the critical points q_c and q_a that are created and annihilated together with p, respectively (possibly empty if p is original and/or a survivor).



Figure 2: Critical points m, s are original while m', s' are newborn together at time t; m annihilates with s' at time t' while m' annihilates with s at time t'': the life of m is extended to [0,t''] and it is paired with s. Similarly, s_1 is original while s'', m'' are newborn together at time t'; m'' annihilates with s_1 at time t'' while s'' is a survivor at time t_{max} : the life of s_1 is this extended to $[0, t_{max}]$ making it a survivor.

Dealing with newborns. Eventually, we want to find a ranked sequence of pairs, each consisting of two original critical points, which are annihilated together; the ranking will be provided by the time of annihilation. In this perspective, the birth of new critical points represents a perturbation in the usual flow of displacements and collapses, which should results in a steadily decreasing number of critical points as the scale parameter increases. Most newborns are ephemeral and can be safely discarded. However, a small but sizable fraction of them does not disappear and survives arbitrarily long through the scales. Moreover, some original critical points in fact are annihilated together with newborn ones. It turns out that, in most cases, these long-lived newborns in fact extend the life span of pre-existing critical points of the newborn ones in their proximity [RJP17].

Consider the example in the left side of Fig.2. A minimum m that was present in the original data is annihilated together with a saddle s' at time t'. If s' is a newborn saddle that was born at time t < t' together with a minimum m', which lives longer than t, we interpret the collapse of m and s' as a transition of the minimum from m to m'. We thus extend the life of m until m' is annihilated. The life of m' could later be extended by the same mechanism, thus prolonging the life of m further, and so on. Fig.2 shows also another example of an original saddle that becomes a survivor because of this life extension.

The recovered life span is computed by following the sequence of annihilation and creation events, checking and propagating the timestamps along the way. The procedure stops when reaching either is a survivor, or a critical point that annihilates together with an original critical point. Eventually, we are left with two sets of original critical points: the survivors that have a life span of $[0, t_{max}]$; and a sequence of pairs of points that are annihilated together at some $t < t_{max}$; the latter sequence is ranked by this time *t*.

4.2. Extracting the Morse-Smale complex

We compute the morphology of the input signal at its most refined scale, by considering the piecewise-linear function defined by (M, f_0) , as described in the previous section. Broadly speaking, a Morse-Smale complex is computed by tracing ascending and descending paths of maximum steepness, starting at the saddles. From a standard 1-saddle there exist two ascending and two descending paths, and on a smooth Morse function all paths are separated, except at their crossings, which occur just at critical points.

Ad discussed in [EHZ03, EH10], piecewise-linear functions do not easily satisfy the Morse requirements for the presence of multiple saddles and collapsed paths. Since we bound the degree of vertices to valence six in our meshes, there can indeed exist at most 2-saddles. While such saddles could be decoupled into pairs of 1saddles by small perturbations, we prefer to process our input unchanged: we address 2-saddles as special cases, tracing all six paths emanating from each such saddle.

Different ascending and descending paths leading to some maximum or minimum may collapse to the same path for a portion of their route. This fact implies the presence of regular points where multiple paths either split or join, which we call *fork* points. A notable example of fork point on a terrain is the location where a tributary from a side valley joins a river in a main valley.

We trace ascending and descending paths from saddles by following the edges of M. Although this is not exact in terms of gradients of the piecewise-linear function, it is simple to compute, it provides an equivalent topology, and the numerical error is negligible if the input is at high resolution (i.e., the size of triangles of M is tiny with respect to the morphology of the function). Akin to [EHZ03], each time an ascending or descending path hits a fork point (i.e., a vertex traversed by a previously computed path), or a saddle, we stop tracing it at that point. After this first step, each fork point q has several incoming paths π_1^i, \ldots, π_k^i and just one outgoing path π^o , which is either ascending or descending. In a second step, we extend each path π_j^i incident at q, constrained by the other paths crossing q: if π_i^i is of the same type of π^o , then it is extended with

submitted to STAG: Smart Tools and Applications in Graphics (2023)

 π^{o} ; otherwise, π^{i}_{j} has just one incoming path π^{i}_{k} of opposite type beside it in the radial order about q and we extend π^{i}_{j} with π^{i}_{k} taken in reverse order. Different paths may travel together through multiple fork points; the procedure above is repeated at all fork points until a maximum or minimum is reached. Paths hitting saddles are treated similarly. See [EHZ03] for further details. Whenever possible, we follow the approach of [BEHP03] to disambiguate ascending and descending paths that overlap by running them parallel, but separated by at least a strip of triangles of M.

Eventually, we obtain complete paths that join each saddle to its related maxima and minima. Such paths never cross, although they may share some of their route with other paths.

4.3. Morphological simplification

We use the ranked list of pairs computed as in §4.1 to simplify the Morse-Smale complex obtained as described in §4.2. The simplified complex at time $t \in [0, t_{\text{max}}]$ is obtained by deleting all pairs of critical points that annihilate at a time $\leq t$ and updating the complex accordingly.

Let (s,m) be a pair of critical points, where *s* is a saddle point and *m* is either a maximum or a minimum that annihilates with *s* at time *t*. Note that only 1-saddles can annihilate, while a 2-saddle must undergo a transition into a pair of 1-saddles before it can hit a catastrophic event; thus, *s* will have exactly two ascending and two descending incident paths. Let MS_t be the complex simplified by deleting all pairs that annihilate before time *t*. The annihilation of (s,m) modifies MS_t as follows:

- The two paths joining *s* with the extrema of type opposite to *m* are deleted and their incident regions are merged;
- Saddle s is turned into a regular point, extremum m is turned into a fork point, and direction of the path joining them is reversed;
- As a consequence, all paths leading to *m* are now extended to paths that lead to the other extrema joined to *s* having the same type of *m*.

5. Experimental results

We tested our method on two synthetic datasets and a real world terrain. Statistics about the datasets and the construction of our multiscale model on each of them are presented in Table 2. Datasets are all regular grids, meshed with a regular pattern of triangles, of sizes between about 66K and 1M vertices. Following a customary approach in the scale-space literature, the levels of the discrete scale-space have been computed by convolving the input function with Gaussian kernels of variance increasing exponentially at each level. All our graphs refer to the scale t in the scale-space, which corresponds to the variance of the filter applied at such scale. Scales between levels, where catastrophic events occur, are estimated by rescaling the linear value of Eq.1 to the exponential scale.

The construction of our multi-scale morphological model requires between about one second for the smallest datasets to over twenty seconds for the largest one. Most of the time is spent in building the scale-space, while the final construction of the multiscale model is negligible even for the largest dataset. The time needed to build the scale-space seems to be dependent only on the

Dataset			Model		Times (ms)					
Name Grid size (Vertices)		Levels in Regions in		Merge Scale-spa		-space	Base M-S	Multi-scale	Total	
			discrete s-s	base M-S	operations	discrete	PL	complex	model	
CosCos	256×256	(65536)	14	96	79	293.7	546.8	3.6	1.6	845.7
CosCos-Spikes	256×256	(65536)	14	572	422	324.0	606.4	58.5	20.7	1009.5
CosCos-Bumps	256×256	(65536)	14	1044	1004	306.0	747.5	178.7	53.5	1285.6
GaussHills-1	400×400	(160K)	14	12	11	898.3	1464.8	0.9	0.2	2364.2
GaussHills-2	400×400	(160K)	14	56	51	916.0	1421.7	4.2	0.9	2347.4
GaussHills-3	400×400	(160K)	14	226	209	920.6	1476.4	15.7	5.9	2418.5
Graian Alps	1200×900	(1080K)	16	5139	4867	8300.9	11467.1	3428.7	153.0	23349.7

Table 2: Statistics for the results presented in Figures 3, 4, and 5. From the left: name and size of the dataset; number of levels in the initial discrete scale-space, number of regions in the Morse-Smale complex at the finest scale, and total number of merge operations in the multi-scale model; time to build the discrete and the piecewise-linear scale-space, including time to find and sort the sequence of catastrophic events; time to compute the Morse-Smale complex at the finest scale; time to perform the morphological simplification and build the multi-scale morphological model; total time to build the model.

size of the input, for a given number of levels of the discrete model; while the (comparatively cheaper) following steps are also affected by the complexity of the Morse-Smale complex, which we measure by the number of its regions. While the current implementation is non-optimized prototype code, the construction of the scalespace could be optimized and also parallelized, possibly achieving a speedup of one or two orders of magnitude. After the multi-scale model has been built off-line, querying it on-line at a desired scale is almost immediate; we do not report times for online queries, as they are always compatible with interaction.

5.1. Datasets and analysis of results

Dataset GaussHills is a collection of three stochastic terrains: in each of them, the final function is obtained by summing multiple Gaussian functions with means, variances and positions randomly sampled from different uniform distributions. In Fig.3, we present experiments on the base dataset, created using only a few functions at large means and variances; and on two additional datasets obtained by adding on top of the previous dataset other sets of Gaussian functions, created in greater quantities and sampled from progressively smaller ranges. The main structure, which is evident in the filtered version of the clean dataset, also occurs roughly at the same scale in every dataset, despite all the additional finer-scale features. The geometric paths of the separatrix lines are more convoluted, but this is to be expected, as they have to navigate through a rougher terrain. The main structure is well preserved at large scale even in the most detailed dataset, which contains many features organized in a highly complicated morphology at the finest scale.

At the bottom of the figure, we report signature graphs, which are built akin to the diagrams used in the literature on topological persistence: each bullet represents a critical point; its horizontal position denotes the intensity of the signal at that point; while its vertical position denotes the extent of its life in the scale-space. The signature of the signal is given by the most representative points, once those ones with a short life have been discarded. It can be clearly seen that the three signatures are very close to each other by filtering spurious points at a scale of about 2^9 , preserving most of the features of the base signal while discarding most of the noise given by additional details.

In Figure 6 we show graphs of the number of regions in the Morse-Smale complex as a function of the scale in our multi-scale model. For the GaussHills datasets (left), it is clear how noise (or, rather, high frequency detail) is filtered fast before scale 2^8 , while the clean dataset is resilient to filtering up to that scale; after a nearplateau spanning scales about $[2^7, 2^9]$, a few features disappear at progressively near scales up to 2^{10} , where the number of regions in the three models becomes almost identical.

Dataset CosCos is a collection of three synthetic terrains: a clean one, sampled from a sinusoidal function dampened with a Gaussian centered at the origin, and two noisy versions of it. The original function is symmetric with respect to both coordinate axes, and all its critical points occur at the same frequency while having different intensities because of the damping. In Fig.4, we present results on the clean dataset, on a version perturbed with full-range impulsive noise on 0.3% of the data, and on a version perturbed with bumps, representing Gaussian noise at a frequency lower than the sampling rate. All the main features - critical points and separatrices - of the clean function are resilient to filtering in the scale space and are well represented in the deep structure, as tracked by the multiscale Morse-Smale complex. The critical points at the boundary are all connected to a virtual minimum outside the square domain and are the first to be filtered out; then saddles that form an inner frame close to the boundary are filtered together in a symmetric pattern, as it can be seen in the filtered version of the clean dataset. Both noisy versions of the dataset initially contain many more critical points, connected by a fine-scale network. Despite their presence, the structure of the original function clearly emerges in the filtered versions, confirming its resilience in the multi-scale model. All main critical points and their connecting lines are present and the morphological structure is very well preserved in both noise models. The geometrical position of features is mostly unaffected by impulse noise while it shows an occasional tendency for a slight displacement when perturbed with Gaussian bumps.

The signature graphs are less clear in this case. Note that, given the high symmetry of the clean dataset, darker bullets represent



Figure 3: Morse-Smale complexes of a collection of stochastic terrains. Top to bottom: input function; complex at the finest scale; complex at a larger scale; signature plot. Left to right: low frequency only (dataset (dataset GaussHills-1); same with medium-scale features added (dataset GaussHills-2); same with further addition of high frequency details on top of previous functions (dataset GaussHills-3). Red/green/blue bullets represent maxima/saddle/minima; bullet size and line width in the M-S complex denote life extent in the scale-space.



Figure 4: Morse-Smale complex of function $e^{-\frac{x^2+y^2}{2\pi}} \cos x \cos y$. Top to bottom: input function; complex at the finest scale; complex at a larger scale; signature plot. Left to right: lean function (dataset CosCos); perturbed with impulse noise (dataset CosCos-Spikes); perturbed with Gaussian noise at a frequency lower than the sampling rate (dataset CosCos-Bumps). Red/green/blue bullets represent maxima/saddle/minima; bullet size and line width in the M-S complex denote life extent in the scale-space.

more than one critical points, since symmetric points occur at the same range and disappear at the same scale. Overall, the signatures of the clean dataset and the dataset perturbed with impulse noise are very similar from about scale 2^9 . The signature of the dataset perturbed with bumps is shifted a bit to the right because the bumps were all added with a positive sign, causing an overall shift of the range towards higher values. Except for that, the points present in the signature beyond scale 2^{10} are about the same as those in the clean dataset, albeit sometimes with a different extent in the scale space.

The graph of the number of regions as a function of scale (Fig.6, center) shows that this number is nearly coincident at scales between 2^6 and 2^9 , suggesting that the deep structure might represent the same features at those scales. Beyond 2^9 , catastrophic events indeed involve the same points in the three datasets, but occur at different scales: earlier in the version perturbed with bumps, later in the version perturbed with spikes and in the clean version, which has the longest lives for the most prominent features.

Dataset Graian Alps is a real terrain, derived from the SRTM Nasa Mission[dF12]. The digital elevation model represents the relief of the Graian Alps, roughly centered on the Mont Blanc Massif. Results are presented in Fig.5. Natural data often presents features at all scales and frequencies and no comparison with a "clean" version is possible, nor meaningful. Our aim is to show that the multiscale morphological model obtained by refining the Morse-Smale complex using the weighted sequence provided by our scale-space analysis captures and correctly represents the relative importance of ridge and valley lines in the terrain morphology, preserving the most relevant ones as the scale parameter grows. It should be noted that our analysis can consider only lines departing from saddles. While all separatrices are either ridges or valleys, there may exist local valley-like or ridge-like lines that do not end at critical points and not correspond to separatrices. In rugged terrains, such lines may be visually evident. This is a limitation of the Morse-Smale complex, though, which is not related to our specific method.

The Graian Alps dataset contains a much higher number of critical points compared to the other datasets. Its graph of the number of regions as a function of scale (Fig.6, right) exhibits an almost perfectly linear decay of the number of regions through the scale, which is clearly representative of the presence of features at all scales.

5.2. Discussion

While we did not perform yet any direct comparison with methods based on persistence, or methods based on surface simplification, we may draw some conclusion from the reported results:

- Our method preserves the location of critical points, as well as the shape of regions in the Morse-Smale complex, always referring to the graph of the input function. This is not possible with methods based on surface simplification, which change the graph of the input function and extract the Morse-Smale complex from the simplified surface.
- Our method is robust to noise, including impulse noise, which is likely to be problematic for methods based on persistence. In

fact, outliers added from input noise annihilate early in our scalespace, while their persistence is likely to be high, possibly at the same scale of the most relevant critical points in the clean function.

- Signatures obtained by preserving just the critical points with longer lives in the scale-space are representative of the main features of the signal, and the Morse-Smale complex at the corresponding scale provides a robust morphological characterization of its graph.
- Application to terrain data shows that even a drastically simplified Morse-Smale complex is able to capture the major ridges and valleys of a terrain, and to partition it into its constituent land masses.

6. Concluding remarks

We have presented a progressive multi-scale morphological model of scalar fields, based on scale-space analysis, and on a consequent simplification of the Morse-Smale complex of the underlying function. This represents an alternative approach to the more popular analysis based on topological persistence and to more common approaches that simplify the geometry of the input signal. The main difference derives from the initial analysis occurring in the frequency domain instead of the amplitude of the signal.

Preliminary results on synthetic data, with and without noise, as well as on real terrain data suggest that our approach is promising, it is resilient to noise, it and produces meaningful results. Further investigation is needed to compare it to techniques based on persistent homology. Both methods provide a ranked sequence of pairs of critical points that guide the simplification of the Morse-Smale complex starting at the finest scale. A theoretical analysis, though, suggests that the two rankings might be intrinsically different. A direct comparison could offer insights on their complementarity and on their respective strengths and weaknesses.

Several other aspects will be the subject of our future work. Concerning scale-space analysis, although our treatment of newborn critical points seems to work well in practice, providing trajectories in the scale space that are much better than those obtained with discrete analyses, the interaction between original and newborn critical points is still unclear, and more investigation is needed to better understand and better justify why some newborn points can be considered as extensions of original points. Moreover, the ranked sequence provided by the continuous scale-space analysis seems to be always topologically correct for the datasets we have tried so far (i.e., it never suggests the annihilation of two critical points that are not directly connected in the Morse-Smale complex computed up until that event). More work is needed to understand if this is a guarantee or not, and if situations where this is not the case might arise, and how they might be described. Concerning the computation of the Morse-Smale complex, the piecewise-linear approach has several corner cases that often occur with real data, adding to the complexity of the underlying algorithms. An alternative approach based on the discrete Morse theory by Forman is worth investigating, as it might give more stable results.



Figure 5: Morse-Smale complexes of the Graian Alps dataset: input data (upper left); complex at the finest scale (upper right); and complexes filtered at progressively larger scales (lower row). Red/green/blue bullets represent maxima/saddle/minima; bullet size and line width represents life extent in the scale-space.



Figure 6: Number of regions in the Morse-Smale complex as a function of the scale: GaussHills dataset, three variants (left); CosCos dataset, three variants (center); Graian Alps dataset (right).

References

- [BDF*08] BIASOTTI S., DE FLORIANI L., FALCIDIENO B., FROSINI P., GIORGI D., LANDI C., PAPALEO L., SPAGNUOLO M.: Describing shapes by geometrical-topological properties of real functions. ACM Comput. Surv. 40, 4 (2008), 1–87. 2
- [BEHP03] BREMER P.-T., EDELSBRUNNER H., HAMANN B., PAS-CUCCI V.: A multi-resolution data structure for two-dimensional morsesmale functions. In *IEEE Visualization*, 2003 (2003), pp. 139–146. 5
- [BS98] BAJAJ C. L., SCHIKORE D. R.: Topology preserving data simplification with error bounds. *Computers & Graphics* 22, 1 (1998), 3–12.
- [DDM*03] DANOVARO E., DE FLORIANI L., MAGILLO P., MES-MOUDI M. M., PUPPO E.: Morphology-driven simplification and multiresolution modeling of terrains. In Proc. 11th ACM Int. Symp. on Advances in Geographic Information Systems (2003), p. 63–70. 2
- [dF12] DE FERRANTI J.: Viewfinder panoramas, 2012. URL: http: //www.viewfinderpanoramas.org/dem3.html.9
- [DRP15] DE GIORGIS N., ROCCA L., PUPPO E.: Scale-space techniques for fiducial points extraction from 3d faces. In *Image Analysis and Processing — ICIAP 2015* (2015), Murino V., Puppo E., (Eds.), Springer International Publishing, pp. 421–431. 2
- [DS18] DEY T., SLECHTA R.: Edge contraction in persistence-generated discrete morse vector fields. *Computers & Graphics* 74 (2018), 33–43.
- [EH10] EDELSBRUNNER H., HARER J.: Computational Topology: An Introduction. American Mathematical Society, 2010. 2, 5
- [EHZ03] EDELSBRUNNER H., HARER J., ZOMORODIAN A.: Hierarchical Morse–Smale Complexes for Piecewise Linear 2-Manifolds. *Discrete* and Computational Geometry 30, 1 (2003), 87–107. 1, 2, 5
- [ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete and Computational Ge*ometry 28, 4 (2002), 511–533. 1, 2
- [FK00] FLORACK L., KUIJPER A.: The Topological Structure of Scale-Space Images. *Jou. of Math. Imaging and Vision 12*, 1 (2000), 65–79.
- [FKM20] FUGACCI U., KERBER M., MANET H.: Topology-preserving terrain simplification. p. 36–47. 2
- [FL99] FROSINI P., LANDI C.: Size theory as a topological tool for computer vision. Pattern Recognition and Image Analysis 9 (1999), 596–603.
- [For98] FORMAN R.: Morse theory for cell complexes. Advances in Mathematics 134, 1 (1998), 90–145. 2
- [GC95] GRIFFIN L. D., COLCHESTER A. C. F.: Superficial and deep structure in linear diffusion scale space: isophotes, critical points and separatrices. *Image and Vision Computing 13*, 7 (1995), 543–557. 2
- [ID17] IURICICH F., DE FLORIANI L.: Hierarchical forman triangulation: A multiscale model for scalar field analysis. *Computers & Graphics* 66 (2017), 113–123. 2
- [Koe84] KOENDERINK J. J.: The structure of images. Biological Cybernetics 50 (1984), 363–370. 2
- [Lin94] LINDEBERG T.: Scale-Space Theory in Computer Vision. Kluwer Academic Publishers, 1994. 1, 2
- [Lin22] LINDEBERG T.: Scale-Covariant and Scale-Invariant Gaussian Derivative Networks. *Journal of Mathematical Imaging and Vision 64*, 3 (2022), 223–242. 2
- [Mat02] MATSUMOTO Y.: An Introduction to Morse Theory, vol. 208 of Translations of Mathematical Monographs. American Mathematical Society, 2002. 2, 3
- [Mil65] MILNOR J.: Topology from the differentiable viewpoint. Univ. Press Virginia, 1965. 1, 2

- [Pri23] PRICE K.: Annotated computer vision bibliography, 2023. URL: http://www.visionbib.com/bibliography/contents. html. 2
- [RJP17] ROCCA L., JENNY B., PUPPO E.: A continuous scale-space method for the automated placement of spot heights on maps. *Computers* & *Geosciences* (2017). 2, 4
- [RKG*11] REININGHAUS J., KOTAVA N., GUENTHER D., KASTEN J., HAGEN H., HOTZ I.: A scale space based persistence measure for critical points in 2d scalar fields. *IEEE Trans. Vis. and Comp. Graph.* 17, 12 (2011), 2045–2052. 1, 2
- [RP13] ROCCA L., PUPPO E.: A virtually continuous representation of the deep structure of scale-space. In *Image Analysis and Processing -ICIAP 2013*, Petrosino A., (Ed.), vol. 8157 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 522–531. 2, 3, 4
- [SFID21] SONG Y., FELLEGARA R., IURICICH F., DE FLORIANI L.: Efficient topology-aware simplification of large triangulated terrains. In *Proc. 29th Int. Cont. on Advances in Geographic Information Systems* (New York, NY, USA, 2021), Association for Computing Machinery, p. 576–587. 2
- [Wit83] WITKIN A. P.: Scale-space filtering. In Proc. 8th Int. Joint Conf. Art. Intell. (Karlsruhe, Germany, August 1983), pp. 1019–1022. 2