

Graph Reconstruction by Discrete Morse Theory

Tamal K. Dey*, Jiayuan Wang*, Yusu Wang*

Abstract

Recovering hidden graph-like structures from potentially noisy data is a fundamental task in modern data analysis. Recently, a persistence-guided discrete Morse-based framework to extract a geometric graph from low-dimensional data has become popular. However, to date, there is very limited theoretical understanding of this framework in terms of graph reconstruction. This paper makes a first step towards closing this gap. Specifically, first, leveraging existing theoretical understanding of persistence-guided discrete Morse cancellation, we provide a simplified version of the existing discrete Morse-based graph reconstruction algorithm. We then introduce a simple and natural noise model and show that the aforementioned framework can correctly reconstruct a graph under this noise model, in the sense that it has the same loop structure as the hidden ground-truth graph, and is also geometrically close. We also provide some experimental results for our simplified graph-reconstruction algorithm.

1 Introduction

Recovering hidden structures from potentially noisy data is a fundamental task in modern data analysis. A particular type of structure often of interest is the geometric graph-like structure. For example, given a collection of GPS trajectories, recovering the hidden road network can be modeled as reconstructing a geometric graph embedded in the plane. Given the simulated density field of dark matters in universe, finding the hidden filamentary structures is essentially a problem of geometric graph reconstruction.

Different approaches have been developed for reconstructing a curve or a metric graph from input data. For example, in computer graphics, much work have been done in extracting 1D skeleton of geometric models using the medial axis or Reeb graphs [10, 28, 19, 15, 21, 7]. In computer vision and machine learning, a series of work has been developed based on the concept of *principal curves*, originally proposed by Hastie and Steutzle [17]. Extensions to graphs include the work in [18] for 2D images and in [24] for high dimensional point data.

In general, there is little theoretical guarantees for most approaches developed in practice to extract hidden graphs. One exception is some recent work in computational topology: Aanijaneya et al. [3] proposed the first algorithm to approximate a metric graph from an input metric space with guarantees. The authors of [8, 15] used Reeb-like structures to approximate a hidden (metric) graph with some theoretical guarantees. These work however

*Department of Computer Science and Engineering, The Ohio State University. tamaldey, yusu@cse.ohio-state.edu, wang.6195@osu.edu

only handles (Gromov-)Hausdorff-type of noise. When input points are embedded in an ambient space, they require the input points to lie within a small tubular neighborhood of the hidden graph. Empirically, these methods do not seem to be effective when the input contains ambient noise allowing some faraway points from the hidden graph.

Recently, a discrete Morse-based framework for recovering hidden structures was proposed and studied [9, 16, 25]. This line of work computes and simplifies a discrete analog of (un)stable manifolds of a Morse function by using the (Forman’s) discrete Morse theory coupled with persistent homology for 2D or 3D volumetric data. One of the main issues in such simplification is the inherent obstructions that may occur for cancelling critical pairs. The authors of [25] suggest sidestepping this and consider a combinatorial representation of critical pairs for further processing. The authors in [9] identify a restricted set of pairs called “cancellable close pairs” which are guaranteed to admit cancellation. This framework has been applied to, for example, extracting filament structures from simulated dark matter density fields [26] and reconstructing road networks from GPS traces [27].

This persistence-guided discrete Morse-based framework has shown to be very effective in recovering a hidden geometric graph from (non-Hausdorff type) noise and non-homogeneous data. The method draws upon the global topological structure hidden in the input scalar field and thus is particularly effective at identifying junction nodes which has been a challenge for previous approaches that rely mostly on local information. However, to date, theoretical understanding of such a framework remains limited. Simplification of a discrete Morse gradient vector field using persistence has been studied before. For example, the work of [9] clarifies the connection between persistence-pairing and the simplification of *discrete Morse chain complex* (which is closely related, but different from the cancellation in the discrete gradient vector field) for 2D and 3D domains. Bauer et al. [6] obtain several results on persistence guided discrete Morse simplification for combinatorial surfaces. The simplification of vertex-edge persistence pairing used in [6] has also been observed in [4] independently for simplifying Morse functions on surfaces. Leveraging these existing developments, we aim to provide a theoretical understanding of a persistence-guided discrete Morse based approach to reconstruct a hidden geometric graph.

Main contributions and organization of paper. In Section 3, we start with one version of the existing persistence-guided discrete Morse-based graph reconstruction algorithm (as employed in [26, 27, 11]). We show that this algorithm can be significantly simplified while still yielding the same output. To establish the theoretical guarantee of the reconstruction algorithm, we introduce a simple yet natural noise model in Section 4. Intuitively, this noise model assumes that we are given an input density field $\rho : \mathbb{R}^d \rightarrow \mathbb{R}$ where densities are significantly higher within a small neighborhood around a hidden graph than outside it. Under this noise model, we show that the reconstructed graph has the same loop structure as the hidden graph, and is also geometrically close to it; the technical details are in Sections 5 and 6 for the general case and the 2-dimensional case (with additional guarantees), respectively.

While our noise model is simple, our theoretical guarantees are first of a kind developed for a discrete Morse-based approach applied to graph reconstruction. In fact, prior to this, it was not clear whether a discrete Morse based approach can recover a graph even if there is no noise, that is, the density function has positive values *only* on the hidden graph. For

our specific noise model, it may be possible to develop thresholding strategies perhaps with theoretical guarantees. However, previous work (e.g. [26, 27]) have shown that discrete Morse approach succeeds in many cases handling non-homogeneous data where thresholding fails. We have implemented the proposed simplified algorithm and tested it on several data sets, which generally gives a speed-up of at least a factor of 2 over a state-of-the-art approach. We present more discussions and experimental results in the appendix.

2 Preliminaries

2.1 Morse theory

For simplicity, we consider only a smooth function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. See [12, 20] for more general discussions.

For a point $p \in \mathbb{R}^d$, the gradient vector of f at a point p is $\nabla f(p) = -[\frac{\partial f}{\partial x_1} \dots \frac{\partial f}{\partial x_d}]^T$, which represents the steepest descending direction of f at p , with its magnitude being the rate of change. An integral line of f is a path $\pi : (0, 1) \rightarrow \mathbb{R}^d$ such that the tangent vector at each point p of this path equals $\nabla f(p)$, which is intuitively a flow line following the steepest descending direction at any point. A point $p \in \mathbb{R}^d$ is *critical* if its gradient vector vanishes, i.e. $\nabla f(p) = [0 \dots 0]^T$. A *maximal* integral line necessarily “starts” and “ends” at critical points of f ; that is, $\lim_{t \rightarrow 0} \pi(t) = p$ with $\nabla f(p) = [0 \dots 0]^T$, and $\lim_{t \rightarrow 1} \pi(t) = q$ with $\nabla f(q) = [0 \dots 0]^T$. See Figure 1a where we show the graph of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, and there is an integral line from p' to the minimum v_1 .

For a critical point p , the union of p and all the points from integral lines flowing into p is referred to as the *stable manifold of p* . Similarly, for a critical point q , the union of q and all the points on integral lines starting from q is called the *unstable manifold of q* . The stable manifold of a minimum p intuitively corresponds the basin/valley around p in the terrain of f . The 1-stable manifolds of index $(d - 1)$ saddles consist of pieces of curves connecting $(d - 1)$ -saddles to maxima – These curves intuitively capture “mountain ridges” of the terrain (graph of the function f); see Figure 1a for an example. Symmetrically, the unstable manifold of a maximum q corresponds to the mountain around q . The 1-unstable manifolds consist of a collection of curves connecting 1-saddles to minima, corresponding intuitively to the “valley ridges”.

In this paper, we focus on a graph-reconstruction framework using Morse-theory (as in e.g. [16, 9, 26, 27]). Intuitively, the 1-stable manifolds of saddles (mountain ridges) of the density field ρ are used to capture the hidden graphs. To implement such an idea in practice, the *discrete Morse theory* is used for robustness and simplicity contributed by its combinatorial nature; and a simplification scheme guided by the persistence pairings is employed to remove noise. Below, we introduce some necessary background notions in these topics.

2.2 Discrete Morse theory

First we briefly describe some notions from discrete Morse theory (originally introduced by Forman [14]) in the simplicial setting.

A k -simplex $\tau = \{p_0, \dots, p_k\}$ is the convex hull of $k + 1$ affinely independent points; k is

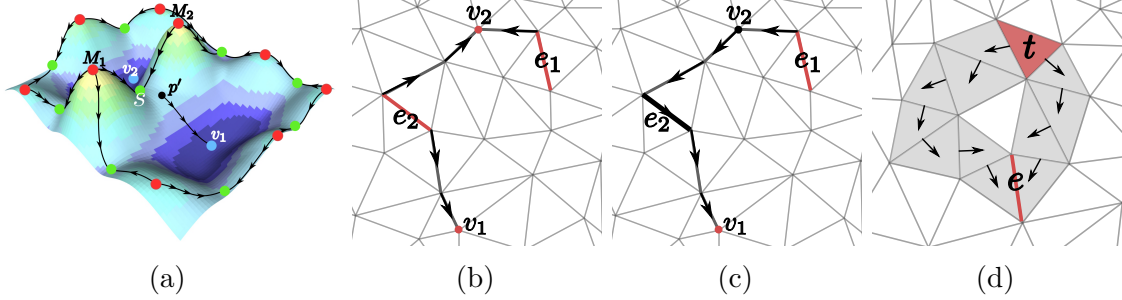


Figure 1: (a) M_1 and M_2 are maxima (red dots), v_1 and v_2 are minima (blue dots), s is a saddle (green dots) with its stable manifolds flowing to it from M_1 and M_2 . If we put a drop of water at p' it will flow to v_1 . If we put it on the other side of the mountain ridge it will flow to minimum v_2 . (b) Before cancellation of pair $\langle v_2, e_2 \rangle$. (c) After cancellation, the path from e_2 to v_2 is inverted, giving rise to a gradient path from e_1 to v_1 , making $\langle v_1, e_1 \rangle$ now potentially cancellable. (d) An edge-triangle pair $\langle e, t \rangle$ which is not cancellable as there are two gradient paths between them.

called the *dimension* of τ . A *face* σ of τ is a simplex spanned by a proper subset of vertices of τ ; σ is a *facet* of the k -simplex τ , denoted by $\sigma < \tau$, if its dimension is $k - 1$.

Suppose we are given a simplicial complex K which is simply a collection of simplices and all their faces so that if two simplices intersect, they do so in a common face. A *discrete (gradient) vector* is a pair of simplices (σ, τ) such that $\sigma < \tau$. A *Morse pairing* in K is a collection of discrete vectors $M(K) = \{(\sigma, \tau)\}$ where each simplex appears in *at most* one pair; simplices that are not in any pair are called *critical*.

Given a Morse pairing $M(K)$, a *V-path* is a sequence $\tau_0, \sigma_1, \tau_1, \dots, \sigma_\ell, \tau_\ell, \sigma_{\ell+1}$, where $(\sigma_i, \tau_i) \in M(K)$ for every $i = 1, \dots, \ell$, and each σ_{i+1} is a facet of τ_i for each $i = 0, \dots, \ell$. If $\ell = 0$, the V-path is *trivial*. This V-path is *cyclic* if $\ell > 0$ and $(\sigma_{\ell+1}, \tau_0) \in M(K)$; otherwise, it is *acyclic* in which case we call this V-path a *gradient path*. We say that a gradient path is a vertex-edge gradient path if $\text{dimension}(\sigma_i) = 0$, implying that $\text{dimension}(\tau_i) = 1$. Similarly, it is an edge-triangle gradient path if $\text{dimension}(\sigma_i) = 1$. A Morse pairing $M(K)$ becomes a *discrete gradient vector field* (or equivalently a *gradient Morse pairing*) if there is no cyclic V-path induced by $M(K)$.

Intuitively, given a discrete gradient vector field $M(K)$, a gradient path $\tau_0, \sigma_1, \dots, \tau_\ell, \sigma_{\ell+1}$ is the analog of an integral line in the smooth setting. But different from the smooth setting, a maximal gradient path may not start or end at critical simplices. However, those that do (i.e, when τ_0 and $\sigma_{\ell+1}$ are critical simplices) are analogous to maximal integral line in the smooth setting which “start” and “end” at critical points, and for convenience one can think of *critical k -simplices* in the discrete Morse setting as *index- k critical points* in the smooth setting. For example, for a function on \mathbb{R}^2 , critical 0-, 1- and 2-simplices in the discrete Morse setting correspond to minima, saddles and maxima in the smooth setting, respectively.

For a critical edge e , we define its *stable manifold* to be the union of edge-triangle gradient paths that ends at e . Its *unstable manifold* is defined to be the union of vertex-edge gradient paths that begins with e . While earlier we use “mountain ridges” (1-stable manifolds) to motivate the graph reconstruction framework, algorithmically (especially for the Morse cancellations below), vertex-edge gradient paths are simpler to handle. Hence in our algorithm

below, we in fact consider the function $g_\rho = -\rho$ (instead of the density field ρ itself) and the algorithm outputs (a subset of) the *1-unstable manifolds* (vertex-edge paths in the discrete setting) as the recovered hidden graph.

Morse cancellation / simplification. One can simplify a discrete gradient vector field $M(K)$ (i.e, reducing the number of critical simplices) by the following Morse cancellation operation: A pair of critical simplices $\langle \sigma, \tau \rangle$ with $\text{dimension}(\tau) = \text{dimension}(\sigma) + 1$ is *cancellable*, if there is a *unique* gradient path $\tau = \tau_0, \sigma_1, \dots, \tau_\ell, \sigma_{\ell+1} = \sigma$ starting at the $k + 1$ -simplex τ and ends at the k -simplex σ . The *Morse cancellation operation on $\langle \sigma, \tau \rangle$* then modifies the vector field $M(K)$ by removing all gradient vectors (σ_i, τ_i) , for $i = 1, \dots, \ell$, while adding new gradient vectors (σ_i, τ_{i-1}) , for $i = 1, \dots, \ell + 1$. Intuitively, the gradient path is inverted. Note that $\tau = \tau_0$ and $\sigma = \sigma_{\ell+1}$ are no longer critical after the cancellation as they now participate in discrete gradient vectors. If there is no gradient path, or more than one gradient path between this pair of critical simplices $\langle \sigma, \tau \rangle$, then this pair is *not cancellable* – the uniqueness condition is to ensure that no cyclic V-paths are formed after the cancellation operation. See Figure 1 (b) – (d) for examples.

2.3 Persistence pairing

The Morse cancellation can be applied to any sequence of critical simplices pairs as long as they are cancellable at the time of cancellation. There is no canonical cancellation sequence. To cancel features corresponding to “noise” w.r.t. an input piecewise-linear function $f : |K| \rightarrow \mathbb{R}$, a popular strategy is to guide the Morse cancellation by the persistent homology induced by the lower-star filtration [16, 26], which we introduce now.

Filtrations and lower-star filtration. Given a simplicial complex K , let S be an ordered sequence $\sigma_1, \dots, \sigma_N$ of all n simplices in K so that for any simplex $\sigma_i \in K$, all of its faces appear before it in S . Then S induces a (*simplex-wise*) *filtration* $F(K) : K_1 \subset K_2 \subset \dots \subset K_N = K$, where $K_i = \bigcup_{j \leq i} \sigma_j$ is the subcomplex formed by the prefix $\sigma_1, \dots, \sigma_i$ of S . Passing to homology groups, we have a *persistence module* $H_*(K_1) \rightarrow \dots \rightarrow H_*(K_N)$, which has a unique decomposition into the direct sum of a set of indecomposable summands that can be represented by the set of *persistence-pairing* $P(K)$ induced by $F(K)$: Each *persistence pair* $(\sigma_i, \sigma_j) \in P(K)$ indicates that a new k -th homological class, $k = \text{dimension}(\sigma_i)$, is created at K_i and destroyed at K_j ; σ_i is thus called a *positive simplex* as it creates, and σ_j a *negative simplex*. Assuming that there is a *simplex-wise function* $\bar{f} : K \rightarrow \mathbb{R}$ such that $\bar{f}(\sigma_i) \leq \bar{f}(\sigma_j)$ if $i < j$, then the *persistence* of the pair (σ, τ) is defined as $\text{pers}(\sigma) = \text{pers}(\tau) = \text{pers}(\sigma, \tau) = \bar{f}(\tau) - \bar{f}(\sigma)$. Some simplices σ_ℓ 's may be unpaired, meaning that homological features created at K_ℓ are never destroyed. We augment $P(K)$ by adding (σ_ℓ, ∞) for every unpaired simplex σ_ℓ to it, and set $\text{pers}(\sigma_\ell, \infty) = \infty$.

The persistent homology can be defined for any filtration of K . In our setting, there is an input function $f : V(K) \rightarrow \mathbb{R}$ defined at the vertices $V(K)$ of K whose linear extension leads to a piecewise-linear (PL) function still denoted by $f : |K| \rightarrow \mathbb{R}$. To reflect topological features of f , we use the lower-star filtration of K induced by f : Specifically, for any vertex $v \in V(K)$, its lower-star $\text{LowSt}(v)$ is the set of simplicies containing v where v has the highest

f value among their vertices. Now sort vertices of K in non-decreasing order of their f -values: v_1, \dots, v_n . An ordered sequence $S = \langle \sigma_1, \dots, \sigma_N \rangle$ induces a *lower-star filtration* $F_f(K)$ of K w.r.t. f if S can be partitioned to n consecutive pieces $\langle \sigma_1, \dots, \sigma_{I_1} \rangle, \langle \sigma_{I_1+1}, \dots, \sigma_{I_2} \rangle, \dots, \langle \sigma_{I_{n-1}+1}, \dots, \sigma_N \rangle$, such that the i -th piece $\langle \sigma_{I_{i-1}+1}, \dots, \sigma_{I_i} \rangle$ equals $\text{LowSt}(v_i)$.

Now let $P_f(K)$ be the resulting set of persistence pairs induced by the lower-star filtration $F_f(K)$. Extend the function $f : V(K) \rightarrow \mathbb{R}$ to a simplex-wise function $\bar{f} : K \rightarrow \mathbb{R}$ where $\bar{f}(\sigma) = \max_{v \in \sigma} f(v)$ (i.e, $\bar{f}(\sigma)$ is the highest f -value of any of its vertices). For each pair (σ, τ) , we measure its persistence by $\text{pers}(\sigma, \tau) = \bar{f}(\tau) - \bar{f}(\sigma)$. Every simplex in K contributes to a persistence pair in $P_f(K)$. However, assuming the value of f is distinct on all vertices, then those persistence pairs with zero-persistence are “trivial” in the sense they correspond to the local pairing of two simplices from the lower-star of the same vertex. A persistence pair (σ, τ) with positive persistence corresponds to a pair of (homological) critical points (p, q) for the PL-function $f : |K| \rightarrow \mathbb{R}$ [12] induced by the function f on $V(K)$, with $p \in \sigma$ and $q \in \tau$.

3 Reconstruction algorithm

Problem setup. Suppose we have a domain Ω (which will be a cube in \mathbb{R}^d in this paper) and a density function $\rho : \Omega \rightarrow \mathbb{R}$ (that “concentrates” around a hidden geometric graph $G \subset \Omega$). In the discrete setting, our input will be a triangulation K of Ω and a density function given as a PL-function $\rho : K \rightarrow \mathbb{R}$. Our goal is to compute a graph \widehat{G} approximating the hidden graph G . In Algorithm 1, we first present a *known* discrete Morse-based graph (1-skeleton) reconstruction framework, which is based on the approaches in [16, 9, 26, 27].

Intuitively, we wish to use “mountain ridges” of the density field to approximate the hidden graph, which are computed as the 1-unstable manifolds of $g_\rho = -\rho$, the negation of the density function. Specifically, after initializing the discrete gradient vector field M to be a trivial one, a persistence-guided Morse cancellation step is performed in Procedure `PerSimpVF()` to compute a new discrete gradient vector field M_δ so as to capture only important (high persistent) features of g_ρ . In particular, Morse-cancellation is performed for each pair of critical simplices from $P(K)$ (if possible) in increasing order of persistence values (for pairs with equal persistence, we use the nested order as in [6]). Finally, the union of the 1-unstable manifolds of all remaining high-persistence critical edges is taken as the output graph \widehat{G} , as outlined in Procedure `CollectOutputG()`.

Since we only need 1-unstable manifolds, K is assumed to be a 2-complex. It is pointed out in [11] that in fact, instead of performing Morse-cancellation for all critical pairs involving edges (i.e, vertex-edge pairs and edge-triangle pairs), one only needs to cancel vertex-edge pairs – This is because only vertex-edge gradient vectors will contribute to the 1-unstable manifolds, and also new vertex-edge vectors can only be generated while canceling other vertex-edge pairs. Hence in `PerSimpVF()`, we can consider *only vertex-edge pairs* $(\sigma, \tau) \in P$ in order. Furthermore, it is not necessary to check whether the cancellation is valid or not – it will always be valid as long as the pairs are processed in increasing orders of persistence [6]¹.

¹We remark that though [6] states that the cancellation is not valid in higher dimension or non-manifold 2-complexes, all cancellations in `PerSimpVF()` are for vertex-edge pairs in a spanning tree which can be viewed as a 1-complex, and thus are always valid.

However, we can further simplify the algorithm as follows: First, we replace procedure `PerSimpVF()` by procedure `PerSimpTree()` as shown in Algorithm 2, which is much simpler both conceptually and implementation speaking. Note that there is *no explicit cancellation operation* any more.

Algorithm 1: MorseRecon(K, ρ, δ)	
	Data: Triangulation K of Ω , density function $\rho : K \rightarrow \mathbb{R}$, threshold δ
	Result: Reconstructed graph \widehat{G}
	begin
1	Compute persistence pairings $P(K)$ by the lower-star filtration of K w.r.t $g_\rho = -\rho$
2	$M = \text{PerSimpVF}(P(K), \delta)$
3	$\widehat{G} = \text{CollectOutputG}(M)$
4	return \widehat{G}
	Procedure PerSimpVF($P(K), \delta$)
1	Set initial discrete gradient field M on K to be trivial
2	Rank all persistence pairs in $P(K)$ in increasing order of their persistence
3	for each $(\sigma, \tau) \in P(K)$ with $\text{pers}(\sigma, \tau) \leq \delta$ do
4	If possible, perform discrete-Morse cancellation of (σ, τ) and update the discrete gradient vector field M
5	return M
	Procedure CollectOutputG(M)
1	$\widehat{G} = \emptyset$
2	for each remaining critical edge e with $\text{pers}(e) > \delta$ do
3	$\widehat{G} = \widehat{G} \cup \{1\text{-unstable manifold of } e\}$
4	return \widehat{G}

The 1-dimensional simplicial complex \mathcal{T} output by procedure `PerSimpTree()` may have multiple connected components $\mathcal{T} = \{T_1, \dots, T_k\}$ – In fact, it is known that each T_i is a tree and \mathcal{T} is a forest (see results from [4, 6] as summarized in Lemma 3.2 below). For each component T_i , we define its *sink*, denoted by $\text{si}(T_i)$, as the vertex $v_i \in T_i$ with the lowest function $g_\rho = -\rho$ value. Lemma 3.2 also states that the sink of T_i would have been the *only critical simplex* among all simplices in T_i , if we had performed the δ -simplification as specified in procedure `PerSimpVF()`. Next, we replace procedure `CollectOutputG()` by procedure `Trebased-OutputG()` shown in Algorithm 2. We use `MorseReconSimp()` to denote our simplified version of Algorithm 1 (with `PerSimpVF()` replaced by `PerSimpTree()`, and `CollectOutputG()` replaced by `Trebased-OutputG()`). In summary, algorithm `MorseReconSimp(K, ρ, δ)` works by first computing all persistence pairs as before. It then collects all vertex-edge persistence pairs (v, e) with $\text{pers}(v, e) \leq \delta$. These edges along with the set of all vertices form a spanning forest \mathcal{T} . Then, for every edge $e = \langle u, v \rangle$ with $\text{pers}(e) > \delta$, it outputs the 1-unstable manifold of e , which is simply the union of e and the unique paths from u and v to the sink (root) of the tree containing them respectively. Its time complexity is stated below; note for the previous algorithm `MorseRecon()`, the cancellation step can take $\tilde{O}(n^2)$ time.

Algorithm 2: MorseReconSimp(K, ρ, δ)

```

Procedure PerSimpTree( $P(K), \delta$ ) /* This procedure replaces original PerSimpVF() */
1   $\Pi :=$  the set of vertex-edge persistence pairs from  $P(K)$ 
2  Set  $\Pi_{\leq \delta} \subseteq \Pi$  to be  $\Pi_{\leq \delta} = \{(v, e) \in \Pi \mid \text{pers}(v, e) \leq \delta\}$ 
3   $\mathcal{T} := \bigcup_{(v, \sigma) \in \Pi_{\leq \delta}} \{\sigma = \langle u_1, u_2 \rangle, u_1, u_2\}$ 
4  return  $\mathcal{T}$ 

Procedure Treebased-OutputG( $\mathcal{T}$ ) /* This procedure replaces CollectOutputG() */
1   $\widehat{G} = \emptyset$ 
2  for each edge  $e = \langle u, v \rangle$  with  $\text{pers}(e) > \delta$  do
3    Let  $\pi(u)$  be the unique path from  $u$  to the sink of the tree  $T_i$  containing  $u$ 
4    Define  $\pi(v)$  similarly; Set  $\widehat{G} = \widehat{G} \cup \pi(u) \cup \pi(v) \cup \{e\}$ 
5  return  $\widehat{G}$ 

```

Theorem 3.1. *The time complexity of our Algorithm PerSimpVF() is $O(\text{PerT}(K) + n)$, where $\text{PerT}(K)$ is the time to compute persistence pairings for K , and n is the total number of vertices and edges in K .*

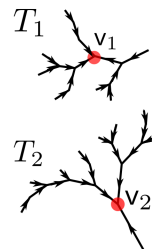
We remark that the $O(n)$ term is contributed by the step collecting all 1-unstable manifolds, which takes linear time if one avoids revisiting edges while tracing the paths.

Justification of the modified algorithm MorseReconSimp(). Let M_δ denote the resulting discrete gradient field after canceling all *vertex-edge* persistence pairs in $P(K)$ with persistence at most δ ; that is, M_δ is the output of the procedure PerSimpVF() (although we only compute the relevant part of the discrete gradient vector field). Using observations from [4, 6], we show that the output \mathcal{T} of procedure PerSimpTree() includes all information of M_δ . Furthermore, procedure Treebased-OutputG() computes the correct 1-unstable manifolds for all critical edges with persistence larger than δ . Indeed, observe that edges in Morse pairings from M_δ (for any $\delta \geq 0$) form a spanning forest of edges in K . Results of [6] imply that the output \mathcal{T} constructed by our modified procedure corresponds exactly to this spanning forest:

Lemma 3.2. *The following statements hold for the output \mathcal{T} of procedure PerSimpTree() w.r.t any $\delta \geq 0$:*

- (i) \mathcal{T} is a spanning forest consisting of potentially multiple trees $\{T_1, \dots, T_k\}$.
- (ii) For each tree T_i , its sink v_i is the only critical simplex in M_δ . The collection of v_i s corresponds exactly to those vertices whose persistence is bigger than δ .
- (iii) Any edge with $\text{pers}(e) > \delta$ remains critical in M_δ (and cannot be contained in \mathcal{T}).

Note that, (ii) above implies that for each T_i , any discrete gradient path of M_δ in T_i terminates at its sink v_i . See the right figure for an illustration. Hence for any vertex $v \in T_i$, the path $\pi(v)$ computed in procedure Treebased-OutputG() is the unique discrete gradient path starting at v . This immediately leads to the following result:

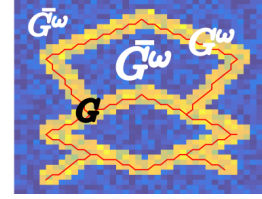


Corollary 3.3. *For each critical edge $e = \langle u, v \rangle$ with $\text{pers}(e) \geq \delta$, $\pi(u) \cup \pi(v) \cup \{e\}$ as computed in procedure `Treebased-OutputG()` is the 1-unstable manifold of e in M_δ . Hence the output of our simplified algorithm `MorseReconSimp()` equals that of the original algorithm `MorseRecon()`.*

4 Noise model

We first describe the noise model in the continuous setting where the domain is $\Omega = [0, 1]^d$. We then explain the setup in the discrete setting when the input is a triangulation K of Ω .

Given a connected “true graph” $G \subset \Omega$, consider a ω -neighborhood $G^\omega \subseteq \Omega$, meaning that (i) $G \subseteq G^\omega$, and (ii) for any $x \in G^\omega$, $d(x, G) \leq \omega$ (i.e. G^ω is sandwiched between G and its ω -offset). Given G^ω , we use $\text{cl}(\overline{G^\omega})$ to denote the closure of its complement $\text{cl}(\overline{G^\omega}) = \text{cl}(\Omega \setminus G^\omega)$. See the right figure, showing G (red graph) with its ω -neighborhood G^ω (orange).



Definition 4.1. *A density function $\rho : \Omega \rightarrow \mathbb{R}$ is a (β, ν, ω) -approximation of a connected graph G if the following holds:*

C-1 There is a ω -neighborhood G^ω of G such that G^ω deformation retracts to G .

C-2 $\rho(x) \in [\beta, \beta + \nu]$ for $x \in G^\omega$; and $\rho(x) \in [0, \nu]$ otherwise. Furthermore, $\beta > 2\nu$.

Intuitively, this noise model requires that the density ρ concentrates around the true graph G in the sense that the density is significantly higher inside G^ω than outside it; and the density fluctuation inside or outside G^ω is small compared to the density value in G^ω (condition C-2). Condition C-1 says that the neighborhood has the same topology of the hidden graph. Such a density field could for example be generated as follows: Imagine that there is an ideal density field $f_G : \Omega \rightarrow \mathbb{R}$ where $f_G(x) = \beta$ for $x \in G^\omega$ and 0 otherwise. There is a noisy perturbation $g : \Omega \rightarrow \mathbb{R}$ whose size is always bounded by $g(x) \in [0, \nu]$ for any $x \in \Omega$. The observed density field $\rho = f_G + g$ is an (β, ν, ω) -approximation of G .

In the discrete setting when we have a triangulation K of Ω , we define a ω -neighborhood G^ω to be a subcomplex of K , i.e. $G^\omega \subseteq K$, such that (i) G is contained in the underlying space of G^ω and (ii) for any vertex $v \in V(G^\omega)$, $d(v, G) \leq \omega$. The outside-region $\text{cl}(\overline{G^\omega}) \subseteq K$ is simply the smallest subcomplex of K that contains all simplices from $K \setminus G^\omega$ (i.e. all simplices **not** in G^ω and their faces). A PL-function $\rho : K \rightarrow \mathbb{R}$ (β, ν, ω) -approximation of G can be extended to this setting by requiring the underlying space of G^ω deformation retracts to G as in (C-1), and having those density conditions in (C-2) only at vertices of K .

We remark that the noise model is still limited – In particular, it does not allow significant non-uniform density distribution. However, this is the first time that theoretical guarantees are provided for a discrete Morse based reconstruction framework, despite that such a framework has been used for different applications before. We also give experiments and discussions in Appendix B that the algorithm works beyond this noise model empirical, where thresholding type approaches do not work.

5 Theoretical guarantee

In this section, we prove results that are applicable to any dimension. Recall that M_δ is the discrete gradient field after the δ -Morse cancellation process, where we perform Morse-cancellation for all *vertex-edge* persistence pairs from $P(K)$. (While our algorithm does not maintain M_δ explicitly, we use it for theoretical analysis.) At this point, all positive edges (i.e, those paired with triangles or unpaired in $P(K)$) remain critical in M_δ . Some negative edges (i.e, those paired with vertices in $P(K)$) are also critical in M_δ – these are exactly the negative edges with persistence bigger than δ . `Trebased-OutputG()` only takes the 1-unstable manifolds of those critical edges (positive or negative) with persistence bigger than δ ; so those positive edges whose persistence is $\leq \delta$ (if there is any) are ignored.

From now on, we use “*under our noise model*” to refer to (1) the input is a (β, ν, ω) -approximated density field w.r.t. G , and (2) $\delta \in [\nu, \beta - \nu)$. Let \hat{G} be the output of algorithm `MorseReconSimp(K, ρ, δ)`. The proof of the following result is in Appendix A.1.

Proposition 5.1. *Under our noise model, we have:*

- (i) *There is a single critical vertex left after `PerSimpVF()` which is in G^ω .*
- (ii) *Every critical edge considered by `Trebased-OutputG()` forms a persistence pair with a triangle.*
- (iii) *Every critical edge considered by `Trebased-OutputG()` is in G^ω .*

Theorem 5.2. *Under our noise model, the output graph satisfies $\hat{G} \subseteq G^\omega$.*

Proof. Recall that the output graph \hat{G} consists of the union of 1-unstable manifolds of all the edges e_1^*, \dots, e_g^* with persistence larger than δ – By Propositions 5.1 (ii) and (iii), they are all positive (paired with triangles), and contained inside G^ω .

Take any $i \in [1, g]$ and consider $e_i^* = \langle u, v \rangle$. Without loss of generality, consider the gradient path starting from u : $\pi : u = u_1, e_1, u_2, e_2, \dots, u_s, e_s, u_{s+1}$. By Lemma 3.2 and Proposition 5.1, u_{s+1} must be a critical vertex (a sink) and is necessarily the global minimum v_0 , which is also contained inside G^ω . We now argue that the entire path π (i.e, all simplices in it) is contained inside G^ω . In fact, we argue a stronger statement: First, we say that a gradient vector (v, e) is *crossing* if $v \in G^\omega$ and $e \notin G^\omega$ (i.e, $e \in \text{cl}(\overline{G^\omega})$) – Since v is an endpoint of e , this means that the other endpoint of e must lie in $K \setminus G^\omega$.

Claim 5.3. *During the δ -Morse cancellation, no crossing gradient vector is ever produced.*

Proof. Suppose the lemma is not true: Then let (v, e) be the *first* crossing gradient vector ever produced during the δ -Morse cancellation process. Since we start with a trivial discrete gradient vector field, the creation of (v, e) can only be caused by reversing of some gradient path π' connecting two critical simplices v' and e' while we are performing Morse-cancellation for the persistence pair (v', e') . Obviously, $\text{pers}(v', e') \leq \delta$. On the other hand, due to our (β, ν, ω) -noise model and the choice of δ , it must be that either both $v', e' \in G^\omega$ or both $v', e' \in K \setminus G^\omega$ – as otherwise, the persistence of this pair will be larger than $\beta - \nu > \delta$.

Now consider this gradient path π' connecting v' and e' in the current discrete gradient vector field M' . Since the pair (v, e) becomes a gradient vector after the inversion of this path,

it must be that (w, e) currently is a gradient vector where $e = \langle v, w \rangle$. Furthermore, since the path π' begins and ends with simplices either both in G^ω or both outside it, the path π' must contain a gradient vector (v'', e'') going in the opposite direction crossing inside/outside, that is, $v'' \in G^\omega$ and $e'' \notin G^\omega$. In other words, it must contain a crossing gradient vector. This however contradicts to our assumption that (v, e) would be the first crossing gradient vector. Hence the assumption is wrong and no crossing gradient vector can ever be created. \square

As there is no crossing gradient vector during and after δ -Morse cancellation, it follows that π , which is one piece of the 1-unstable manifold of the critical edge e_i^* , has to be contained inside G^ω . The same argument works for the other piece of 1-unstable manifold of e_i^* (starting from the other endpoint of e_i^*). Since this is for any $i \in [1, g]$, the theorem holds. \square

The previous theorem shows that \widehat{G} is close to G in geometry. Next we will show that they are also close in topology.

Proposition 5.4. *Under our noise model, \widehat{G} is homotopy equivalent to G .*

Proof. We show that \widehat{G} has the same first Betti number as that of G which implies the claim as any two graphs in \mathbb{R}^d with the same first Betti number are homotopy equivalent.

The underlying space of ω -neighborhood G^ω of G deformation retracts to G by definition. Observe that, by our noise model, G^ω is a sublevel set in the filtration that determines the persistence pairs. This sublevel set being homotopy equivalent to G must contain exactly g positive edges where g is the first Betti number of G . Each of these positive edges pairs with a triangle in $\overline{G^\omega}$. Therefore, $\text{pers}(e, t) > \delta$ for each of the g positive edges in G^ω . By our earlier results, these are exactly the edges that will be considered by procedure `Treebased-OutputG()`. Our algorithm constructs \widehat{G} by adding these g positive edges to the spanning tree each of which adds a new cycle. Thus, \widehat{G} has first Betti number g . \square

We have already proved that \widehat{G} is contained in G^ω . This fact along with Proposition 5.4 can be used to argue that any deformation retraction taking (underlying space) G^ω to G also takes \widehat{G} to a subset $G' \subseteq G$ where G' and G have the same first Betti number. In what follows, we use G^ω to denote also its underlying space.

Theorem 5.5. *Let $F : G^\omega \times [0, 1] \rightarrow G^\omega$ be any deformation retraction. Then, the restriction $F|_{\widehat{G}} : \widehat{G} \times [0, 1] \rightarrow G^\omega$ is a homotopy from the embedding \widehat{G} to $G' \subseteq G$ where G' is the minimal subset so that G and G' have the same first Betti number.*

Proof. The fact that $F|_{\widehat{G}}(\cdot, \ell)$ is continuous for any $\ell \in [0, 1]$ is obvious from the continuity of F . Only thing that needs to be shown is that $F|_{\widehat{G}}(\widehat{G}, 1) = G'$. Suppose not. Then, $G'' = F|_{\widehat{G}}(\widehat{G}, 1)$ is a proper subset of G which has a first Betti number less than that of G . We observe that the cycle in \widehat{G} created by a positive edge e along with the paths to the root of the spanning tree is also non-trivial in G^ω because this is a cycle created by adding the edge e during persistence filtration and the edge e is not killed in G^ω . Therefore, a cycle basis for \widehat{G} is also a homology basis for G^ω . Since the map $F(\cdot, 1) : G^\omega \rightarrow G$ is a homotopy equivalence, it induces an isomorphism in the respective homology groups; in particular, a homology basis in G^ω is mapped to a homology basis in G . Therefore, the image $G'' = F|_{\widehat{G}}(\widehat{G}, 1)$ must have a basis of cardinality g if \widehat{G} has first Betti number g . But, G'' cannot have a cycle basis of cardinality g if it is a proper subset of G' reaching a contradiction. \square

6 Additional guarantee for 2D

For \mathbb{R}^2 , we now show that G^ω actually deformation retracts to \widehat{G} , which is stronger than saying G and \widehat{G} are homotopy equivalent. We are unable to prove this result for dimensions higher than 2, as our current proof needs that the edge-triangle persistence pairs can always be canceled (even though our algorithm does not depend on edge-triangle cancellations at all). It would be interesting, as a future work, to see whether a different approach can be developed to avoid this obstruction for the special case under our noise model. The main result of this section is as follows.

Theorem 6.1. *Under our noise model, G^ω deformation retracts to G and \widehat{G} .*

This main result follows from Proposition 6.2 and Theorem 6.3 below. To prove them, we will show that there exists a partition $\mathcal{R} := \{R_i\}$ of the set of triangles in K for which Theorem 6.3 holds. (This theorem is our main tool in establishing the deformation retract.) We first state the results below before giving their proofs. Let $B_i = \partial R_i$ where ∂ is the boundary operator operating on the 2-chain R_i . We also abuse the notations R_i and B_i to denote the geometric space that is the point-wise union of simplices in the respective chains. Let t_i be a triangle in R_i whose choice will be explained later. In the following, let H be the maximal set of edges in \widehat{G} whose deletions do not eliminate a cycle (assume that a vertex is deleted only if all of its edges are deleted). Observe that H necessarily consists of negative edges forming “hairs” attached to the loops of \widehat{G} and hence to $\cup_i B_i$ because of the following proposition.

Proposition 6.2. *Under our noise model, $\widehat{G} = \cup B_i \cup H$.*

Theorem 6.3. *Under our noise model, there exists a partition $\{R_i\}$ of triangles in K such that, there is a deformation retraction of $\cup_i (R_i \setminus t_i)$ to $\cup B_i$ that comprises of two deformation retractions, one from $\cup_i (R_i \setminus t_i)$ to G^ω and another one from G^ω to $\cup_i B_i \cup H$ which is \widehat{G} .*

Now we describe the construction of a partition \mathcal{R} of the triangles in K to prove Proposition 6.2 and Theorem 6.3. For technicality we assume that K is augmented to a triangulation of a sphere by putting a vertex v at infinity and joining it to the boundary of K with edges and triangles all of whom have function value ∞ . Let $P(K)$ be the collection of persistence pairs of the form either (σ, τ) or (σ, ∞) generated from the lower-star filtration $F(K)$ as described before. Since K is 2-dimensional, each pair (σ, τ) is either a vertex-edge pair or an edge-triangle pair. We order persistence pairs in $P(K)$ by their persistence, where ties are broken via the nested order in the filtration $F(K)$, and obtain:

$$P(K) = \{(\sigma_1, \tau_1), \dots, (\sigma_n, \tau_n), (\alpha_1, \infty), \dots, (\alpha_s, \infty)\}. \quad (1)$$

Starting with a trivial discrete gradient vector field M_0 where all simplices in K are critical, the algorithm `PerSimpVF()` performs Morse cancellations for the first $m \leq n$ persistence pairs $(\sigma_1, \tau_1), \dots, (\sigma_m, \tau_m)$ in order where $\text{pers}(\sigma_m, \tau_m) \leq \delta$ but $(\sigma_{m+1}, \tau_{m+1}) > \delta$. Let M_i denote the gradient vector field after canceling (σ_i, τ_i) . Recall that in the implementation of the algorithm we do not need to perform Morse cancellation for any edge-triangle pairs. However in this section, for the theoretical analysis, we will cancel edge-triangle pairs as well. Recall

that a positive edge is one that creates a 1-cycle, namely, it is either paired with a triangle or unpaired; while a negative edge is one that destroys a 0-cycle (i.e, paired with a vertex).

Consider the ordered sequence of edge-triangle persistence pairs, $(e_1, t_1), \dots, (e_n, t_n)$, which is a subsequence of the one in (1). Consider the sequence t_1, \dots, t_n of triangles in K ordered by the above sequence. Recall the standard persistence algorithm [13]. It implicitly associates a 2-chain with a triangle t when searching for the edge it is about to pair with. This 2-chain is non-empty if t is a destructor, and is empty otherwise. Let D_i denote this 2-chain associated with t_i for $i \in [1, n]$. Initially, the algorithm asserts $D_i = t_i$. At any stage, if D_i is not empty, the persistence algorithm identifies the edge e in the boundary ∂D_i that has been inserted into the filtration $F(K)$ most recently. If e has not been paired with anyone, the algorithm creates the persistence pair (e, t_i) . Otherwise, if e has already been paired with a triangle, say $t_{i'}$, then D_i is updated with $D_i = D_i + D_{i'}$ and the search continues. Given an index $j \in [1, n]$, we define a modified set of chains C_i^j inductively as follows. For $j = 1$, $C_i^1 = t_i$. Assume that C_i^{j-1} has been already defined. To define C_i^j , similar to the persistence algorithm, check if the most recent edge e_{j-1} in the persistence order (not the filtration order) is on the boundary ∂C_i^{j-1} . If so, define $C_i^j := C_i^{j-1} + C_{j-1}^{j-1}$ and $C_i^j := C_i^{j-1}$ otherwise. The proof of the following result is in Appendix A.3.

Proposition 6.4. *For $i \in [1, n]$, e_i is in ∂C_i^i . Furthermore, e_i is the most recent edge in ∂C_i^i according to the filtration order $F(K)$.*

Procedure $\text{PerSimpVF}()$ also implicitly maintains a 2-chain R_i^* with each triangle t_i . Initially, $R_i^* = t_i$ as in the case of D_i . Then, inductively assume that R_i^* is the 2-chain implicitly associated with t_i when a persistence pair $(e_{i'}, t_{i'})$ is about to be considered by $\text{PerSimpVF}()$ and the boundary ∂R_i^* contains $e_{i'}$. By reversing a gradient path between $t_{i'}$ and $e_{i'}$, it implicitly updates the 2-chain R_i^* as $R_i^* := R_i^* + R_{i'}^*$. We observe that R_i^* is identical with C_i^i . Proposition 6.5 below establishes this fact along with some other inductive properties useful to prove Theorem 6.3. The proof can be found in Appendix A.4.

Proposition 6.5. *Let (e_j, t_j) be the edge-triangle persistence pair $\text{PerSimpVF}()$ is about to consider and let C_i^j be the 2-chains defined as above. Then, the following statements hold:*

- (a) *For each triangle t_i , $i = 1, \dots, n$, in the persistence order, the 2-chain R_i^* satisfies the following conditions: (a.i) $R_i^* = C_i^i$, (a.ii) interpreting R_i^* as a set of triangles, one has that the sets R_i^* , $i = j, \dots, n$, partition the set of all triangles in K .*
- (b) *There is a gradient path from t_i to all edges of the triangles in R_i^* , and (b.i) the path is unique if the edge is in the boundary ∂R_i^* for every $i = j, \dots, n$; (b.ii) if there is more than one gradient path from t_i to an edge e , then e must be a negative edge.*

We are now ready to setup the regions R_i s needed for Theorem 6.3 and Proposition 6.2. Suppose the first m edge-triangle pairs have persistence less than or equal to δ , the parameter supplied to $\text{PerSimpVF}()$. Then, we set $R_i = R_i^*$ as in Proposition 6.5 for $i \geq m + 1, \dots, n$. The proof for Proposition 6.2 is in Appendix A.2.

Finally, similar to the vertex-edge gradient vectors, we say that a gradient vector (e, t) is *crossing* if $e \in G^\omega$ and $t \notin G^\omega$. The following claim can be proved similarly as Claim 5.3.

Claim 6.6. *During the δ -Morse cancellation of edge-triangle pairs, no crossing gradient vector is ever produced.*

Proof of Theorem 6.3. Set $\hat{R}_i = R_i \setminus t_i$. Let \mathcal{T} be the spanning tree formed by all negative edges and their vertices. Let L_i be the set of edges in R_i that has more than one gradient path from t_i to them; $L_i \subset \mathcal{T}$ by Proposition 6.5 (b.ii). First, we want to establish a deformation retraction from $\cup(R_i \setminus t_i)$ to G^ω . To do this, for $k = 0, 1, \dots, s$, we will define \hat{R}_i^k inductively where \hat{R}_i^{k-1} deformation retracts to \hat{R}_i^k and $\hat{R}_i^s \subseteq G^\omega \cup L_i$. Let $R_i^0 = \hat{R}_i$. For $k = 1, \dots, s$, consider a *positive* edge e in \hat{R}_i^{k-1} where (a) e is not in G^ω and (b) there is a unique gradient path in R_i from t_i to e that passes through triangles all of which are in $R_i \setminus \hat{R}_i^{k-1}$. If such an edge e exists, then e is necessarily incident to a single triangle, say t , in \hat{R}_i^{k-1} . We collapse the pair (e, t) , which is necessarily an edge-triangle gradient vector pair because e is positive. We take \hat{R}_i^k to be $\hat{R}_i^{k-1} \setminus \{e, t\}$. If no such e exists, then either (A) there is no positive edge in $\hat{R}_i^{k-1} \setminus G^\omega$ any more; or (B) for each positive edge $e' \in \hat{R}_i^{k-1} \setminus G^\omega$, (B-1) there is a unique gradient path from t_i to e' but this path passes through some triangle in \hat{R}_i^{k-1} ; or (B-2) there are two gradient paths from t_i to e' that passes through only triangles from $R_i \setminus \hat{R}_i^{k-1}$.

If there is no positive edge in $\hat{R}_i^{k-1} \setminus G^\omega$ any more, then $\hat{R}_i^{k-1} \subseteq G^\omega \cup L_i$, as otherwise, there will be at least some triangle from $\hat{R}_i^{k-1} \setminus G^\omega \cup L_i$ with at least one boundary edge of it being positive. The induction then terminates; we set $s = k - 1$ and reach our goal.

We now show that case (B-1) is not possible. Suppose it happens, that is, $e' \in \hat{R}_i^{k-1} \setminus L_i$ is an edge not in G^ω for which the unique gradient path from t_i passes through triangles in \hat{R}_i^{k-1} . Let e'' be the first edge in this path that is in $\hat{R}_i^{k-1} \setminus L_i$. Then, if $e'' \notin G^\omega$, it qualifies for the conditions (a) and (b) required for e reaching a contradiction. So, assume $e'' \in G^\omega$. But, in that case, we have a gradient path that goes into G^ω and then comes out to reach $e' \notin G^\omega$. There has to be a gradient pair in this path where the edge is in G^ω and the triangle is not in G^ω . This contradicts Claim 6.6. Thus, case (B-1) is not possible.

Now consider (B-2). By Proposition 6.5 (b.ii), e' must be negative, and cannot be positive. Hence case (B-2) is not possible either.

To summarize, the induction terminates in case (A), at which time we would have that $\hat{R}_i^s \subseteq G^\omega \cup L_i$. Furthermore, this process also establishes a deformation retraction from \hat{R}_i to \hat{R}_i^s realized by successive collapses of edge-triangle pairs. Furthermore, by construction, each collapsed pair (e, t) must be from $\text{cl}(\overline{G^\omega})$, hence $\cup_i \hat{R}_i^s$ contains all simplices in G^ω . Combined with that $\hat{R}_i^s \subseteq G^\omega \cup L_i$, we have that $\cup_i \hat{R}_i^s = G^\omega \cup L$, where $L = \cup_i L_i$ is a subset of the spanning tree \mathcal{T} . The edges in L being part of a spanning tree cannot form a cycle and thus can be retracted along the tree to G^ω , which gives rise to a deformation retraction from $\cup_i (R_i \setminus t_i)$ to $G^\omega \cup L$ and then to G^ω , establishing the first part of Theorem 6.3.

We now show that $(\cup_i \hat{R}_i^s) \cap G^\omega = G^\omega$ deformation retracts to $\cup B_i \cup H$. Let \hat{L}_i be the edges in $\hat{R}_i^s \cap G^\omega$ that have more than one gradient path from t_i to them. These edges are negative by Proposition 6.5 (b.ii). Replacing \hat{R}_i with $\hat{R}_i^s \cap G^\omega$ and edges in $B_i \cup \hat{L}_i$ playing the role of edges in $G^\omega \cup L_i$ in the above inductive procedure, we can obtain that $\hat{R}_i^s \cap G^\omega$ deformation retracts to $B_i \cup \hat{L}_i$. To this end, we also observe that $\cup \hat{L}_i = \mathcal{T} \cap G^\omega$ where \mathcal{T} is the spanning tree formed by all negative edges, as we only collapse edge-triangle pairs that are gradient pairs (hence the participating edges are always positive). This implies that $H \subset \cup \hat{L}_i$. Again, edges in \hat{L}_i (being part of a spanning tree) can be retracted along the spanning tree till one reaches B_i or edges in H . Performing this for each i , we thus obtain a

deformation retraction from $(\cup_i \hat{R}_i^s) \cap G^\omega = G^\omega$ to $\cup B_i \cup \cup \hat{L}_i$ and further to $\cup B_i \cup H = \hat{G}$. This finishes the proof of Theorem 6.3.

In Appendix B, we also provide some experiments demonstrating the efficiency of the simplified algorithm, as well as discussion on thresholding strategies.

Acknowledgments We thank Suyi Wang for generously helping with the software. The Enzo dataset used in our experiments is obtained from [1]. The Bone dataset is obtained from [2]. This work is supported by National Science Foundation under grants CCF-1526513, CCF-1618247, and CCF-1740761, and by National Institute of Health under grant R01EB022899.

References

- [1] Enzo code is developed by the Laboratory for Computational Astrophysics at the University of California in San Diego (<http://lca.ucsd.edu>).
- [2] CIBC CT dataset archive. <http://www.sci.utah.edu/cibc-software/ctdata.html>.
- [3] M. Aanjaneya, F. Chazal, D. Chen, M. Glisse, L. Guibas, and D. Morozov. Metric graph reconstruction from noisy data. *International Journal of Computational Geometry & Applications*, 22(04):305–325, 2012.
- [4] D. Attali, M. Glisse, S. Hornus, F. Lazarus, and D. Morozov. Persistence-sensitive simplification of functions on surfaces in linear time. *Presented at TOPOINVIS*, 9:23–24, 2009.
- [5] U. Bauer, M. Kerber, J. Reininghaus, and H. Wagner. Phat–persistent homology algorithms toolbox. *Journal of Symbolic Computation*, 78:76–90, 2017.
- [6] U. Bauer, C. Lange, and M. Wardetzky. Optimal topological simplification of discrete functions on surfaces. *Discr. Comput. Geom.*, 47(2):347–377, 2012.
- [7] S. Biasotti, D. Giorgi, M. Spagnuolo, and B. Falcidieno. Reeb graphs for shape analysis and applications. *Theoretical Computer Science*, 392(1-3):5–22, 2008.
- [8] F. Chazal, R. Huang, and J. Sun. Gromov–hausdorff approximation of filamentary structures using reeb-type graphs. *Discr. Comput. Geom.*, 53(3):621–649, 2015.
- [9] O. Delgado-Friedrichs, V. Robins, and A. Sheppard. Skeletonization and partitioning of digital images using discrete morse theory. *IEEE Trans. Pattern Anal. Machine Intelligence*, 37(3):654–666, March 2015.
- [10] T. Dey and J. Sun. Defining and computing curve-skeletons with medial geodesic function. In *Sympos. Geom. Proc.*, volume 6, pages 143–152, 2006.
- [11] T. Dey, J. Wang, and Y. Wang. Improved road network reconstruction using discrete morse theory. In *Proc. 25th ACM SIGSPATIAL*. ACM, 2017.

- [12] H. Edelsbrunner and J. Harer. *Computational Topology - an Introduction*. American Mathematical Soc., 2010.
- [13] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discr. Comput. Geom.*, 28:511–533, 2002.
- [14] R. Forman. Morse theory for cell complexes. *Advances in mathematics*, 134(1):90–145, 1998.
- [15] X. Ge, I. I Safa, M. Belkin, and Y. Wang. Data skeletonization via reeb graphs. In *Advances in Neural Info. Proc. Sys.*, pages 837–845, 2011.
- [16] A. Gyulassy, M. Duchaineau, V. Natarajan, V. Pascucci, E. Bringa, A. Higginbotham, and B. Hamann. Topologically clean distance fields. *IEEE Trans. Visualization Computer Graphics*, 13(6):1432–1439, Nov 2007.
- [17] T. Hastie and W. Stuetzle. Principal curves. *Journal of the American Statistical Association*, 84(406):502–516, 1989.
- [18] B. Kégl and A. Krzyzak. Piecewise linear skeletonization using principal curves. *IEEE Trans. Pattern Anal. Machine Intelligence*, 24(1):59–74, 2002.
- [19] L. Liu, E. W Chambers, D. Letscher, and T. Ju. Extended grassfire transform on medial axes of 2d shapes. *Computer-Aided Design*, 43(11):1496–1505, 2011.
- [20] J. Milnor. *Morse Theory*. Princeton Univ. Press, New Jersey, 1963.
- [21] M. Natali, S. Biasotti, G. Patanè, and B. Falcidieno. Graph-based representations of point clouds. *Graphical Models*, 73(5):151–164, 2011.
- [22] M. L. Norman, G. L. Bryan, R. Harkness, J. Bordner, D. Reynolds, B. O’Shea, and R. Wagner. Simulating Cosmological Evolution with Enzo. *ArXiv e-prints*, May 2007. [arXiv:0705.1556](https://arxiv.org/abs/0705.1556).
- [23] B. W. O’Shea, G. Bryan, J. Bordner, M. L. Norman, T. Abel, R. Harkness, and A. Kritsuk. Introducing Enzo, an AMR Cosmology Application. *ArXiv Astrophysics e-prints*, March 2004. [arXiv:astro-ph/0403044](https://arxiv.org/abs/astro-ph/0403044).
- [24] U. Ozertem and D. Erdogmus. Locally defined principal curves and surfaces. *Journal of Machine learning research*, 12(Apr):1249–1286, 2011.
- [25] V. Robins, P. J. Wood, and A. P. Sheppard. Theory and algorithms for constructing discrete morse complexes from grayscale digital images. *IEEE Trans. Pattern Anal. Machine Intelligence*, 33(8):1646–1658, Aug 2011.
- [26] T. Sousbie. The persistent cosmic web and its filamentary structure - I. Theory and implementation. 414:350–383, June 2011. [arXiv:1009.4015](https://arxiv.org/abs/1009.4015).
- [27] S. Wang, Y. Wang, and Y. Li. Efficient map reconstruction and augmentation via topological methods. In *Proc. 23rd ACM SIGSPATIAL*, page 25. ACM, 2015.

- [28] Y. Yan, K. Sykes, E. Chambers, D. Letscher, and T. Ju. Erosion thickness on medial axes of 3d shapes. *ACM Trans. on Graphics*, 35(4):38, 2016.

A Missing proofs

A.1 Proof of Proposition 5.1

Proof of claim (i): We first show that there can be no critical vertex of M_δ from the interior of outside-region $\text{cl}(\overline{G^\omega})$. Indeed, suppose there is a critical vertex $v \notin G^\omega$, then it forms a persistence pair either with ∞ or with a critical edge $e \in K$ in $P(K)$. The former cannot happen as the only vertex unpaired by the persistence algorithm is the global minimum of the input function f , which necessarily lies in the neighborhood G^ω . So suppose (v, e) is the persistent pair containing v . It then follows that $e \in \text{cl}(\overline{G^\omega})$ as it must come after v in the lower-star filtration $F(K)$ induced by f . Hence $\text{pers}(v, e)$ is necessarily smaller than δ under our noise model. In other words, the persistence pairing (v, e) should have already been canceled during the δ -Morse simplification process. As a result, there cannot be any critical vertex left in $\text{cl}(\overline{G^\omega})$.

Next, we argue that there is exactly one critical vertex in the neighborhood G^ω in the final discrete gradient vector field M_δ . First, note that the global minimum v_0 of the PL-function f must remain critical, as v_0 will be paired with ∞ by the persistent homology induced by the lower-star filtration, and the persistence pairing (v_0, ∞) remains after the δ -Morse cancellation by Lemma 3.2 (ii). We now prove that there cannot be any other critical vertex in G^ω .

Assume on the contrary that there is another critical vertex $u \in G^\omega$. This means we have a persistence pairing (u, e') with $\text{pers}(u, e') > \delta$. It then follows that $e' \in \text{cl}(\overline{G^\omega})$ by our assumption on the noise model for $f = -\rho$. Recall that the low-star filtration $F(K)$ is induced by adding the lower-star of v_i s in order where v_0, v_1, \dots, v_n are sorted in increasing order of f -values. Specifically, $F(K)$ contains the following sequence: $\emptyset \subset \bar{K}_1 \subset \bar{K}_2 \subset \dots \subset \bar{K}_n$, where $\bar{K}_i = \bigcup_{j \leq i} \text{LowSt}(v_j)$. Assume that $e' = \langle v_a, v_b \rangle$ such that $a < b$ (i.e, $f(v_a) < f(v_b)$). Then \bar{K}_b is the subcomplex in the filtration $F(K)$ when the edge e' is first included. It follows from the persistence algorithm [13] that, there are two connected components $C_1, C_2 \subset \bar{K}_{b-1}$ that is merged with the addition of e' ; that is, $v_a \in C_1$ and $v_b \in C_2$. Furthermore, since e' forms a persistence pair with u , this means that u must be the global minimum of one of the components, say C_1 w.o.l.g, and $f(u) > f(z)$ where z is the global minimum of the other component C_2 . See Figure 2.

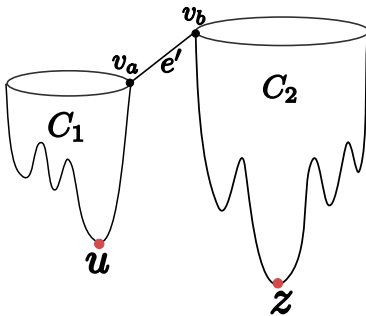


Figure 2: Illustration for proof of Proposition 5.1.

Hence $z \in G^\omega$ as $u \in G^\omega$ and $f(u) > f(z)$. This however is not possible. Indeed, note

that $G^\omega \subseteq \bar{K}_{b-1}$: This is because that $e' \in \text{cl}(\overline{G^\omega})$ (and thus $\bar{f}(e') > \delta > \nu$), implying \bar{K}_{b-1} will contain all vertices (thus as well as simplices they span) whose function value is lower than δ , which contains G^ω (whose vertices have f values $< \nu < \delta$). Since both u and z are in G^ω (which is connected), they are already connected in \bar{K}_{b-1} , contradictory to that $u \in C_1$ and $z \in C_2$ are two different connected components in \bar{K}_{b-1} . Hence the persistence pairing (u, e) cannot exist, and there is no critical vertex other than the global minimum v_0 .

Proof of claim (ii). A critical edge forms a persistence pair with either a vertex, or a triangle, or remains unpaired. The last case is not possible as the input domain is simply connected. Any remaining critical edge cannot be paired with a (critical) vertex either by claim (i). Statement (ii) then follows.

Proof of claim (iii). We now prove Proposition (iii): Let e be a critical edge that is considered by `Treebased-OutputG()`. By Proposition 5.1, e forms a persistence pair with a triangle t . Since e is considered by `Treebased-OutputG()`, $\text{pers}(e, t) > \delta$. Under our noise model, this means that e and t cannot be both in G^ω or both in its complement $\overline{G^\omega}$. Then, we have two possibilities: (i) $e \in G^\omega$ and $t \notin G^\omega$ or (ii) $t \in G^\omega$ and $e \notin G^\omega$. In case (i) there is nothing to prove because $e \in G^\omega$; case (ii) is impossible because the function value of t will be less than that of its pairing edge e .

A.2 Proof of Proposition 6.2

First, we need the following obvious claim.

Claim A.1. *Let G be any graph and $T \subseteq G$ be a spanning tree. Let $E \subseteq G \setminus T$ be a set of edges. Let $G_{E,T} \subseteq G$ be the subgraph where $G_{E,T} = (T \cup E) \setminus H$ where H is the largest set of edges in $T \cup E$ whose deletions do not eliminate any cycle. Given T and E , the set H and hence $G_{E,T}$ is unique.*

We call $G_{E,T}$ in the above claim to be the minimal subgraph with respect to the edge set E and the spanning tree T . Now, consider the spanning tree T of the 1-skeleton of K consisting of all negative edges. Addition of positive edges creates cycle. Let E be the set of all positive edges with persistence more than δ . The graph consisting of edges $T \cup E$ has cycles. Consider the minimal subgraph $G_{E,T}$. The graph \hat{G} computed by the algorithm is the graph $G_{E,T}$ plus a maximal set of edges in T whose deletions do not eliminate any cycle in \hat{G} . Denoting this set of edges as H , one has $\hat{G} = G_{E,T} \cup H$.

On the other hand, the union of the boundaries $\cup_i B_i$ of the regions R_i consist of only negative edges or positive edges that are critical. This is because otherwise the edges have to be positive and non-critical, but then the two regions containing such edges are already merged, eliminating them from boundaries. Therefore, $\cup_i B_i$ can be formed by taking the spanning tree T consisting of negative edges, adding the positive edges E to it and then eliminating all edges that cannot eliminate any cycle. Therefore, $\cup_i B_i$ also forms a subgraph of the 1-skeleton of K that is minimal with respect to the same set of positive edges E and the spanning tree T . Hence, by claim $G_{E,T} = \cup_i B_i$ and thus $\hat{G} = \cup_i B_i \cup H$.

A.3 Proof of Proposition 6.4

First we show that the chain D_i constructed by the persistence algorithm [13] is a subchain of C_i^i . The chain D_i itself is constructed by adding chains as the algorithm searches for the edge e_i to be paired with t_i . Let $D_{i_1} = t_i, D_{i_2}, \dots, D_{i_k}$ be the ordered sequence of chains that are added to construct D_i .

We use double induction. First, assume inductively that, for $j < i$, D_j is a subchain of C_j^j . It is true initially for $j = 1$ because $C_1^1 = D_1 = t_1$. To prove the hypothesis for C_i^i , assume in a nested induction that $D_{i_1}, \dots, D_{i_{k-1}}$ are subchains of $C_{i_k}^{i_k}$. Initially, t_i and hence D_{i_1} is a subchain of $C_{i_k}^{i_k}$. For the nested induction consider the edge e_{i_k} which is in the boundary $\partial(D_{i_1} + \dots + D_{i_{k-1}})$. The persistence pair (e_{i_k}, t_{i_k}) appears before the pair (e_i, t_i) . Therefore, the chain $C_{i_k}^{i_k}$ is added to $C_{i_k}^{i_k}$ and thus becomes a subchain in $C_{i_k}^{i_k}$. But, the chain $C_{i_k}^{i_k}$ contains D_{i_k} as a subchain which is also added to C_i^i as a result. This establishes that D_i is a subchain of C_i^i .

If the edge $e_i \in \partial D_i$ is not in the boundary ∂C_i^i , it must be the case that, for some $i' < i$, $C_{i'}^{i'}$ has been added to C_i^i where $\partial C_{i'}^{i'}$ contains e_i . In that case, $(e_i, t_{i'})$ must be a persistence pair according to the construction of C_j^j 's. This is impossible because (e_i, t_i) is a persistence pair and $i > i'$.

To show that e_i is the most recent edge in ∂C_i^i , assume inductively that e_j is the most recent edge in ∂C_j^j for $j < i$. If e_i is not the most recent edge in ∂C_i^i , let $e_{i'}$ be the most recent one. Let $C_{j'}^{j'}$ be the chain that was added to C_i^i because of which $e_{i'}$ was included in the boundary C_i^i . Clearly, $j' < i$. Then, $e_{i'}$ was in the boundary $\partial C_{j'}^{j'}$ and by inductive hypothesis $i' \leq j'$. It follows that $i' \leq j' < i$ reaching a contradiction that e_i is not the most recent edge. The proposition then follows.

A.4 Proof of Proposition 6.5

We induct on j . When $j = 1$, we have $R_i^* = t_i$ for every $i = 1, \dots, n$. Then, (a) & (b) are satisfied trivially. Assume that they hold for j . Since $R_i^* = C_i^j$ by inductive hypothesis, Proposition 6.4 ensures that ∂R_j^* contains e_j , the edge with which t_j pairs with. Then, by inductive hypothesis (b), there is a unique gradient path π_j from t_j to e_j . The algorithm $\text{PerSimpVF}()$ only reverses π_j .

The set R_j^* is a constituent of the sets $\{R_i^*\}, i = j, \dots, n$ that partition the set of triangles in K . Hence, the edge $e_j \in \partial R_j^*$ necessarily belongs to another boundary $\partial R_{j'}^*$ for some $j' > j$. Observe that since each edge can be incident to at most two triangles, there is a unique such j' . We update $R_{j'}^* := R_{j'}^* + R_j^*$. Clearly, the sets $R_i^*, i = j + 1, \dots, n$ partitions the set of triangles in K , proving claim (a.ii). We argue that $R_{j'}^* = C_{j'}^{j+1}$ completing the proof that (a) holds after $\text{PerSimpVF}()$ processes t_j . Before the update it holds inductively that $R_{j'}^* = C_{j'}^j$. After $\text{PerSimpVF}()$ processes (e_j, t_j) , the only 2-chain that gets updated is $R_{j'}^*$ because e_j is only in $\partial R_{j'}^*$ where $j' \neq j$. The new 2-chain $R_{j'}^*$ after the update exactly satisfies the definition of $C_{j'}^{j+1}$ because $R_{j'}^* := R_{j'}^* + R_j^* = C_{j'}^j + C_j^j$. This proves claim (a.i).

To show that claim (b) holds as well, we need to consider the only updated 2-chain $R_{j'}^*$. We say that an edge is in a 2-chain if one of its triangles contains the edge in its boundary. Let e be any edge in $R_{j'}^*$. If e is in $R_{j'}^*$ before the update, then we already have a gradient

path from $t_{j'}$ to e by inductive hypothesis. If e is in R_j^* but not in $R_{j'}^*$ before the update, we have a gradient path from $t_{j'}$ to e in the updated $R_{j'}^*$. This gradient path is obtained by concatenating three sequences, say π_1, π_2 and π_3 . The sequence π_1 is the gradient path from $t_{j'}$ to e_j in $R_{j'}^*$ before the update. Let t be the last triangle shared by a path from t_j to e and the unique path from t_j to e_j before the update. The sequence π_2 is the subsequence of the reversed path π_j from e_j to t with e_j and t removed, and the sequence π_3 is the gradient path from t to e that existed in R_j^* . This establishes that there is a gradient path from $t_{j'}$ to all edges in $R_{j'}^*$ after the update. Now, to prove (b.i), if e is a boundary edge of $R_{j'}^*$ after the update, it must be the case that e is in the boundary of either R_j^* or $R_{j'}^*$ but not in both before the update. In that case, uniqueness of the gradient path from either t_j or $t_{j'}$ to e before the update implies the uniqueness of the path after the update. Hence, (b.i) holds for updated $R_{j'}^*$ and hence for all $i > j$.

To prove claim (b.ii), suppose $e \in R_{j'}^*$ has two gradient paths from $t_{j'}$ to it after the update. If e is from $R_{j'}^*$ before the update, then the gradient path from $t_{j'}$ to it will not change. Hence it must be negative in this case. So now suppose e is from R_j^* but not from $R_{j'}^*$ before the update. By the argument in the previous paragraph, we now that new gradient path from $t_{j'}$ to e consists of three portions, and it is easy to see that we can have two paths from $t_{j'}$ to e only if there were two paths from t_j to e in R_j^* before the update. Hence by induction hypothesis, e must be negative as well. This proves claim (b.ii), and finishes the proof of Proposition 6.5.

B Experiments

In this section, we perform our algorithm on 2D and 3D density fields. The 2D density fields are generated from the GPS trajectories and the goal is to extract the hidden road network behind [27]. There are three 3D density fields: A synthetic dataset where the ground truth is known; and two real-life datasets where the noise model assumptions may or may not be satisfied: Specifically, we have the *Enzo dataset* [22, 23], which comes from the simulations of cosmological structure formation in university and where the goal is to extract the filament structure behind; and the *Bone dataset* which are Micro CT images of bones from the CT Dataset Archive from CIBC [2]. For the *Bone dataset*, we crop a portion of one bone since the triangulation is huge. The sizes of these data sets are in Table 1, where “Athens”, “Beijing”, and “Berlin” represents the three 2D density fields obtained from large collection of noisy GPS traces in the three respective cities.

The input points of the datasets are actually vertices from a 2D or 3D grid, so we obtain a triangulation of input points by simply triangulating each 2D/3D cubic cell. It is possible to use a cubical complex directly, but using a triangulation allows us to threshold on the input density function to remove noise and reduce the size of input simplicial complex.

There are two parameters used in our experiments: the threshold δ which is the input parameter of `MorseRecon()` and `MorseReconSimp()` used for persistence-based simplification; and the parameter t to threshold the density function as used by the thresholding method (not needed for our algorithm). Both parameters are chosen empirically.

Below, we first show the efficiency of our simplified algorithm `MorseReconSimp()` versus the original algorithm `MorseRecon()`. Then, we present some experimental results showing

that, while the practical data does not fall under our noise model, empirically the algorithm still works well (as also demonstrated earlier in work such as [26, 27]). Furthermore, as we mentioned earlier, thresholding may be able to produce a graph with theoretical guarantee for input density fields under our noise model, when it is combined with say, medial axis type approaches, to extract the graph after thresholding (even for that it is not yet straightforward to obtain such theoretical result). However, we will present experiments that show that simple thresholding does not work well empirically for non-uniformly sampled input.

Name	#vertex	#edge	#triangle
Athens	444,600	1,331,111	886,512
Beijing	3,754,580	11,255,893	7,501,314
Berlin	80,741	241,084	160,344
ENZO	262,144	1,536,192	2,524,284
Bone	5,351,976	31,829,419	52,768,394

Table 1: Size of the triangulations of the datasets.

Running time. We implemented our simplified algorithm `MorseReconSimp()` and now compare its running time with the original algorithm `MorseRecon()`. As we mentioned in Section 3 in [11], this algorithm has already been simplified so that Morse-cancellation is **only done** for vertex-edge critical pairs. Hence we implemented this improved version of `MorseRecon()`, which we refer to as `MorseRecon+()` (the version of `PerSimpVF()` where only vertex-edge critical pairs are canceled is referred to as `PerSimpVF+()`). Note that a speedup of a factor of at least 2 has already been reported for `MorseRecon+()` over `MorseRecon()` on GPS data [11]. The step of computing the persistence pairing for the negation of density field (both in `MorseRecon+()` and our `MorseReconSimp()`) is done by PHAT software package [5]. The comparison of their running time is shown in Table 2. Specifically, note the step of computing persistence pairing is common to both the original `MorseRecon()` algorithm and our simplified `MorseReconSimp()` version. Furthermore, for 3D data, this step is currently the bottleneck (although it may be improved by using persistence algorithm more suitable for volumetric data, such as DiPha). Hence we report the time for this step separately in the 3rd column of the table. The 4th and 5th columns of Table 2 show the running time (in seconds) of algorithm `MorseRecon+()` (without persistence computation) and of algorithm `MorseReconSimp()` (without persistence computation). As we can see, our simplified algorithm is more efficient, generally we see at least a factor of 2 speed-up.

Reconstructions results. We do not show the reconstruction results for 2D data sets, as these data sets are originally used in both [27, 11] and our output is the same (Corollary 3.3). We now show the reconstruction results for 3D data sets. Since it is hard to obtain ground truth for real life datasets, we first show the result of a synthetic dataset where we know the ground truth in Figure 3a. This dataset is generated as follows: First we create an arbitrary graph (the black lines), then we diffuse this graph by convolution with a Gaussian kernel to obtain a 3D density field. The band width of the Gaussian kernel is 4, comparing

Name	δ	Pre-process	PerSimpVF ⁺ +CollectOutputG	PerSimpTree+Treebased-OutputG
Athens	0.01	12.3	1.2	0.5
Beijing	0.1	97.8	13.1	5.4
Berlin	10	2.0	0.25	0.17

Name	δ	Pre-process	PerSimpVF ⁺ +CollectOutputG	PerSimpTree+Treebased-OutputG
ENZO	50	26.5	1.0	0.38
Bone	40	869	21.6	8.2

Table 2: Running time (in seconds) of pre-process (column 3, including filtration and persistence computation), algorithm MorseRecon⁺() w/o persistence part (column 4) and our simplified algorithm MorseReconSimp() w/o persistence part (column 5).

to the radius 50 of the input data. As we can see in Figure 3a, our output graph captures the two loops in the ground truth graph, which follows our theoretical results.

The reconstruction results for real life data sets Enzo and Bone are given in Figure 3b and in Figure 3c, respectively, overlaid with the original density field. There is no single good threshold t as the density has a rather non-homogeneous distribution, and structures can exist at different level of thresholds (which we will show more shortly below). Hence we provide a volume rendering of the input density field overlapped with our output reconstruction.

We also provide reconstructions of the bone dataset at different threshold δ level in Figure 4. As we increase the threshold δ , we capture fewer features of the data.




Figure 3: In the three figures above, the red lines are the reconstructions, the volumes show the volume renderings of the input density functions. (a) Synthetic dataset. The black lines are the ground truth. (b) Enzo dataset. (c) Part of a bone dataset.

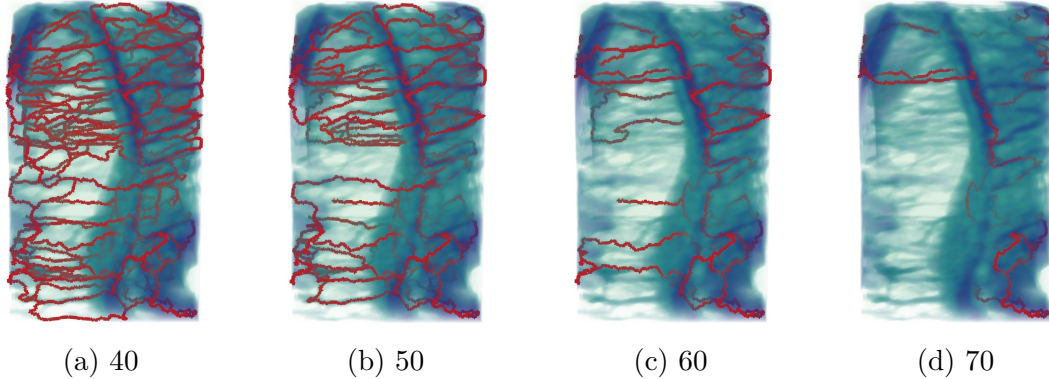


Figure 4: (a)-(d): Reconstructions of the bone dataset at different threshold δ level (show in the sub caption).

Comparing with the thresholding method. A natural way to denoise is simply thresholding. However, for many data, due to the non-homogeneous, no single threshold t can capture all features. For example, as we show in Figures 5 and 6, branching / loop features appear at different time as we vary the density threshold t . Furthermore, features that appear at high density threshold t may actually be destroyed at low threshold t . Hence there is no single good threshold t to capture all features. As an example, see the two loops circled with blue in Figure 5a for a high density threshold t , which is filled in a lower threshold t in 5c and 5d respectively. However, we need to lower the threshold t as many new features, such as the loops circled with green shown in 5c only appears in a lower threshold t . Similarly, for the bone dataset, we note that most of the vertical fibers in the lower part can only be captured at a lower threshold t . However, at that point, the features in the high density regions (say the top part) are already merged.

On the other hand, while we do not yet have theoretical guarantees for the discrete Morse based graph reconstruction algorithm for such non-homogeneous data sets, we note that it performs very well empirically, captures these features of different density scales. (The red graphs in both figures are the output reconstruction by our algorithm `MorseReconSimp`.)

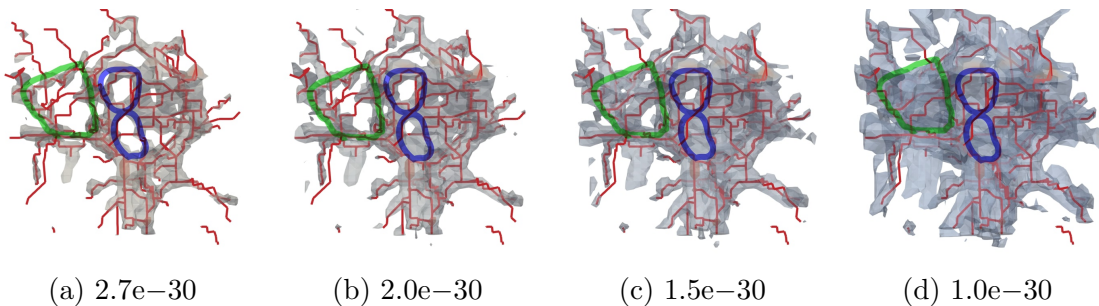


Figure 5: (a)-(d): Part of the Enzo dataset. The gray volumes are isovolumes with increasing lower bounds (show in the sub caption), the red lines are the reconstructions. Reliable features existing in high thresholds got killed when we lower it.

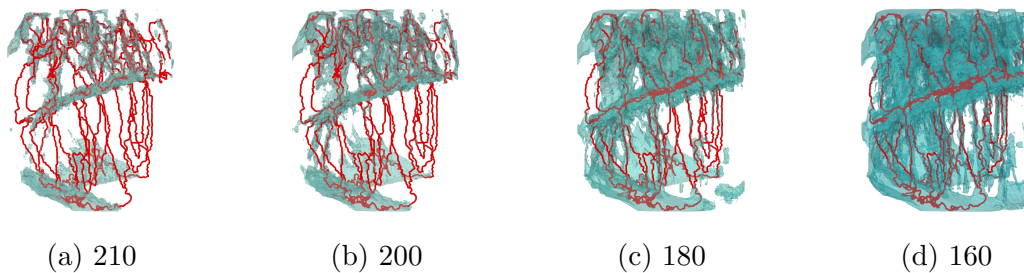


Figure 6: (a)-(d): Bone dataset. The blue volumes are isovolumes with increasing lower bounds (show in the sub caption), the red lines are the reconstructions. Low threshold t captures features in lower part, high threshold t captures loops in upper part.