

This is a pre print version of the following article:

Bapat, Asilata, Robyn, Brooks, Celia, Hacker, Claudia, Landi e Barbara I., Mahler. "Morse-based Fiberings of the Persistence Rank Invariant" Working paper, ASSOCIATION FOR WOMEN IN MATHEMATICS SERIES

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

26/05/2026 11:35

(Article begins on next page)

Morse-based Fiberings of the Persistence Rank Invariant

Asilata Bapat*
asilata.bapat@anu.edu.au

Robyn Brooks †
robyn.brooks@bc.edu

Celia Hacker‡
celia.hacker@epfl.ch

Claudia Landi§
claudia.landi@unimore.it

Barbara I. Mahler ¶
mahler@maths.ox.ac.uk

April 15, 2021

Abstract

Although there is no doubt that multi-parameter persistent homology is a useful tool to analyse multi-variate data, efficient ways to compute these modules are still lacking in the available topological data analysis toolboxes. Other issues such as interpretation and visualization of the output remain difficult to solve. Software visualizing multi-parameter persistence diagrams is currently only available for 2-dimensional persistence modules. One of the simplest invariants for a multi-parameter persistence module is its rank invariant, defined as the function that counts the number of linearly independent homology classes that live in the filtration through a given pair of values of the multi-parameter. We propose a step towards interpretation and visualization of the rank invariant for persistence modules for any given number of parameters. We show how discrete Morse theory may be used to compute the rank invariant, proving that it is completely determined by its values at points whose coordinates are critical with respect to a discrete Morse gradient vector field. These critical points partition the set of all lines of positive slope in the parameter space into equivalence classes, such that the rank invariant along lines in the same class are also equivalent. We show that we can deduce all persistence diagrams of the restrictions to the lines in a given class from the persistence diagram of the restriction to a representative in that class¹.

*Australian National University, Canberra, Australia

†Boston College, Massachusetts, United States

‡EPFL, Lausanne, Switzerland

§Università di Modena e Reggio Emilia, DISMI, Italy

¶University of Oxford, United Kingdom

¹*MSC: 55N31, 57Q70, Keywords:* persistence module, persistence diagram, discrete gradient vector field, critical value

Introduction

Digital data are being produced at a constantly increasing pace, and their availability is changing the approach to science and technology. The fundamental hypothesis of Topological Data Analysis is that data come as samples taken from an underlying shape, and unveiling such shape is important for understanding the studied phenomenon. Topological shape analysis amounts to determining non-trivial topological holes in any dimension. Computational Topology provides tools to derive specific signatures – topological invariants – which depend only on topological features of the shape of data and are robust to local noise [15]. Among them, *persistent homology* [19, 6, 26, 17, 29] stands out as most useful, having already found numerous applications in a diverse range of fields [7, 8, 20, 23, 28].

The first step in the persistence pipeline is to build a family, called a filtration, of nested simplicial complexes that model the data at various scales by varying one or more parameters. The second step focuses on the maps induced in homology by the simplicial inclusions to extract invariants such as the *persistence module*. The third step is to use persistence invariants as a source of feature vectors in machine learning contexts, the final goal being to use the acquired topological information to improve the understanding of the underlying data. An important feature of this pipeline is its robustness to noise in the input data [15].

Some systems warrant analysis across multiple parameters, so it is important to focus on multi-parameter persistence [11] where the filtration may depend on any number of real-valued parameters, and n -tuples of parameter values have an inherited partial order. Unlike single-parameter persistent homology, which is completely described by a *persistence diagram*, multi-parameter persistence modules contain more information than it is possible to handle, understand, and visualize easily. Therefore, it is convenient to summarize them by simpler invariants.

Among the various invariants considered for a multi-parameter persistence module, such as the blockcodes of [16] and the multi-graded Betti numbers of [25], one of the simplest invariants is the *rank invariant* [11], defined as the function that counts the number of linearly independent homology classes that live in the filtration through any given pair of values of the multi-parameter. Theoretically, computation of the rank invariant can be carried out by fibering it along lines with positive slope in an n -dimensional space, where n is the number of parameters [12]. Indeed, the rank invariant of the restriction of a persistence module to an increasing line is completely described by a persistence diagram. The union of all such persistence diagrams forms a compact object called the persistence space [14]. Practically though, one needs to restrict the number of relevant lines to a reasonable number. For example, the state-of-the-art tool for rank invariant visualization RIVET [24] achieves reduction to some template lines using Betti tables, takes $O(m^3)$ runtime (with m the number of simplices), and is limited to two parameters.

Given the dimensional, computational, and interpretability limits of the cur-

rently available methods, multi-parameter persistence is not yet a viable option for data analysis. It is thus worth exploring different approaches that could enable faster computation and better understanding of the rank invariant in a dimension-agnostic way. To this end, we propose to exploit the information contained in the critical cells of a discrete gradient vector field consistent with the given multi-filtration as a means to enhancing geometrical understanding and computational efficiency of multi-parameter persistence.

Evidence for the usefulness of Morse theory in multi-parameter persistence is given in [4, 5, 27]. In these papers, discrete Morse theory is used to reduce the multi-parameter persistence input data size by substituting the original simplicial complex with a Morse complex containing only critical cells but having the same persistence as the initial complex. This approach takes advantage of the fact that number of critical cells is very small in comparison with m , the total number of simplices. Computations are sped up by reducing the actual value of m , while the runtime complexity of obtaining a discrete gradient vector field compatible with the filtration is $m \cdot s^2$, where s denotes the maximum number of simplices in a vertex star. Although in the worst case s may be as large as m , in many applications it is negligible in comparison to m , so that the complexity may be considered linear in m . Moreover, as shown by tests in [27], this global reduction also avoids repeating the retrieval of the same null persistence pairs when fibering the rank invariant along different lines, improving time performances proportionally to the number of lines: the larger the number of lines, the more convenient the reduction to the Morse complex.

Supported by such empirical evidence, in this paper we investigate the theoretical connection between critical cells of discrete Morse theory and rank invariant computation along lines, with the goal of improving the available methods for fibering the rank invariant of multi-parameter persistence in a way that is computationally efficient, geometrically interpretable, and readily visualizable. To this end, we exploit the correspondence between critical cells of a discrete gradient vector field and topological changes in filtrations compatible with it to:

1. Show that the rank invariant for n -parameter persistence modules can be computed by selecting a small number of values in the parameter space \mathbb{R}^n determined by the critical cells of the discrete gradient vector field, and using these values to reconstruct the rank invariant for all other possible values in the parameter space (Theorem 2.4).
2. Use such critical values to define an equivalence relation among lines so that representatives from each equivalence class form a convenient selection of lines for fibering the rank invariant (Definition 3.8).
3. Show that fibering the rank invariant along lines in the same equivalence class yields persistence diagrams that can be obtained one from the other by simple bijections between critical values (Theorem 3.17).
4. Present a method based on the previous results for computing any fiber of the rank invariant using only finitely many template fibers (Section 4).

We emphasize that all of our results hold for any number of parameters, thus improving the state-of-the-art methods that allow for only two. Moreover, our

method’s input requires only the critical cells of a discrete gradient vector field, so that the required pre-processing can be achieved in just linear time according to various available algorithms [4, 5, 27]. Finally, the connection to Morse theory allows for a more immediate geometric interpretation and visualization of the topological features detected by persistence as pairs of critical cells in the given simplicial complex.

The outline of this paper is as follows: in Section 1, we outline background information which will be necessary for the reader to understand the rest of the text. Section 2 focuses on computing the rank invariant, supported by lemmas and diagram-chasing. Section 3 explains the computation of the persistence space along with delving into fibering the rank invariant along equivalent lines. We conclude with Section 4 by laying out a method for fibering the rank invariant of persistence by lines, and suggesting possible applications. Finally, in Appendix A we discuss a result that allows us to develop an algorithm which chooses a representative line for each equivalence class in the case of 2-parameter persistence modules. The generalization of this algorithm to higher dimensional persistence modules remains future work.

1 Notation and Definitions

These definitions are partially based on the definitions in [24, 21] and [18].

Let \mathbb{K} be a field. For computational purposes, \mathbb{K} is often taken to be finite. Define the following partial order on \mathbb{R}^n : for $u = (u_i), v = (v_i) \in \mathbb{R}^n$, we say that $u \preceq v$ (resp. $u \prec v$) if and only if $u_i \leq v_i$ (resp. $u_i < v_i$) for all i . The poset (\mathbb{R}^n, \preceq) will be our *parameter space*.

Definition 1.1 (Persistence module). A *persistence module* \mathbb{V} over the parameter space \mathbb{R}^n is an assignment of a \mathbb{K} -vector space V_u to each $u \in \mathbb{R}^n$, and transition maps $i^{u,v} : V_u \rightarrow V_v$ to pairs of points $u \preceq v \in \mathbb{R}^n$, satisfying the following properties:

- $i^{u,u}$ is the identity map for all $u \in \mathbb{R}^n$.
- $i^{v,w} \circ i^{u,v} = i^{u,w}$ for all $u \preceq v \preceq w \in \mathbb{R}^n$.

A persistence module over \mathbb{R}^n is also known as an n -parameter persistence module or an n D persistence module.

Definition 1.2 (Rank invariant). For $n \geq 1$, let $\mathcal{H}^n \subset \mathbb{R}^n \times \mathbb{R}^n$ be the subset of pairs (u, v) such that $u \preceq v$. Let \mathbb{V} be an n D persistence module. Then the *rank invariant of \mathbb{V}* is a function $\rho_{\mathbb{V}} : \mathcal{H}^n \rightarrow \mathbb{Z}$, defined as

$$\rho_{\mathbb{V}}(u, v) = \text{rank}(i^{u,v}).$$

Persistence modules most often arise from filtered simplicial complexes.

Definition 1.3 (Finite simplicial complex). A *finite simplicial complex* K is a collection of subsets of a finite set V_0 , such that:

- all singletons of V_0 are in K , and
- If $\beta \in K$ and $\alpha \subset \beta$, then $\alpha \in K$.

The elements of V_0 are called *vertices*, and the elements of K are called *simplices*. If $\alpha \in K$ contains $p + 1$ vertices, then we say α has *dimension* p , sometimes denoted $\alpha^{(p)}$. If $\beta \in K$ and $\alpha \subset \beta$, then we say α is a *face* of β and β is a *coface* of α , and denote this by $\alpha < \beta$. If α is a codimension one face of β , we say that α is a *facet* of β and β is a *cofacet* of α . If α is a face of a dimension p simplex β , and is not a face of any other p dimensional simplex, then we say that α is a *free face* of β .

Definition 1.4 (Filtration). Let K be a finite simplicial complex. An n -parameter filtration of K is a collection of subcomplexes $\mathcal{K} = \{K^u\}_{u \in \mathbb{R}^n}$ of K , such that $K^u \subset K^v$ whenever $u \preceq v$, and moreover that

$$K = \bigcup_{u \in \mathbb{R}^n} K^u.$$

A complex with such a filtration is called a *multi-filtered simplicial complex* or an *n-filtered simplicial complex*.

Remark. Since K is finite, there is some $u \in \mathbb{R}^n$ such that $K = K^u$.

Definition 1.5. We say that a simplex σ in K has *entrance value* $u \in \mathbb{R}^n$ if

$$\sigma \in K^u - \bigcup_{w \preceq u, w \neq u} K^w$$

This value tells us where σ has entered the filtration. Note that if the entrance value of σ is u , then $\sigma \in K^v$ for all $u \preceq v$, and $\sigma \notin K^w$ for all $w \preceq u$ such that $u \neq w$. Therefore, if u is an entrance value of σ , then u is a (possibly non-unique) minimal value of the set $\{v \in \mathbb{R}^n \mid \sigma \in K^v\}$.

According to the definition of filtration, if a simplex σ enters the filtration at u , all of its faces have to be in K^u as well. The faces can enter jointly with σ or at earlier values of the filtration. However, entrance values are not guaranteed to exist. To avoid pathological situations in which a simplex does not have any entrance values, such as in the 1-parameter filtration of K defined by $K^u = \emptyset$ for each $u \leq 0$ and $K^u = K$ for each $u > 0$, we will only consider filtrations as follows.

Definition 1.6. (Tameness) A filtration of a simplicial complex K is said to be *tame* if each simplex of K has at least one entrance value.

Remark. For a tame filtration, it is guaranteed that there exists some $u \in \mathbb{R}^n$ such that $K^u = \emptyset$.

Example 1.7. Consider a function $f: K \rightarrow \mathbb{R}^n$ that is monotonic with respect to the face relation. That is, $\alpha < \beta \in K$ implies that $f(\alpha) \preceq f(\beta)$. Every such function gives rise to a tame filtration, by defining the filtered pieces of K to be sublevel sets as follows:

$$K^v = f^{-1}(\{u \in \mathbb{R}^n : u \preceq v\}).$$

On the other hand, not all filtrations of K as in Definition 1.4 define a function $f : K \rightarrow \mathbb{R}^n$, even under the tameness assumptions, because the entrance value of a simplex may not be unique. This motivates the following definition that can be found in [10].

Definition 1.8 (One-criticality). An n -parameter filtration is said to be *one-critical* if every simplex of \mathcal{K} has a unique entrance value in the filtration.

Proposition 1.9. *Each tame one-critical filtration of a simplicial complex K is the sublevel set filtration of a monotonic function $f : K \rightarrow \mathbb{R}^n$. Conversely, if $f : K \rightarrow \mathbb{R}^n$ is monotonic, then its sublevel sets form a tame and one-critical filtration of K .*

Proof. Let K be a simplicial complex and \mathcal{K} a tame and one-critical filtration of K . We construct a monotonic function $f_{\mathcal{K}} : K \rightarrow \mathbb{R}$ whose sublevel set filtration is precisely \mathcal{K} . For each simplex $\sigma \in \mathcal{K}$, set $f_{\mathcal{K}}(\sigma)$ to be the unique entrance value of σ in K with respect to \mathcal{K} , which exists and is unique because the filtration is tame and one-critical. Let $\alpha, \beta \in K$ such that $\alpha < \beta$. Recall that if $\beta \in K^u$ for some $u \in \mathbb{R}^n$, then $\alpha \in K^u$ as well. In particular, $\alpha \in K^{f_{\mathcal{K}}(\beta)}$. Since $f_{\mathcal{K}}(\alpha)$ is the unique minimum value u such that $\alpha \in K^u$, it must be the case that $f_{\mathcal{K}}(\alpha) \preceq f_{\mathcal{K}}(\beta)$, and therefore $f_{\mathcal{K}}$ is monotonic.

Conversely, let $f : K \rightarrow \mathbb{R}^n$ be a monotonic function. As before, set \mathcal{K} to be the set of all sublevel sets

$$K^v = f^{-1}(\{u \in \mathbb{R}^n : u \preceq v\}).$$

Let σ be any simplex of K . It is clear that $\sigma \in K^{f(\sigma)}$. We prove that $f(\sigma)$ is the unique entrance value of σ . Note that it is indeed an entrance value: if $u \preceq f(\sigma)$ and $u \neq f(\sigma)$, then $f(\sigma) \notin K^u$. Suppose that v is any entrance value of σ , which means that $\sigma \in K^v$ and $\sigma \notin K^u$ for any $u \neq v$ such that $u \preceq v$. Since $\sigma \in K^v$, we must have $f(\sigma) \preceq v$. Since $\sigma \notin K^u$ for any $u \neq v$ such that $u \preceq v$, it must be the case that $f(\sigma) \neq u$ for any such u . Therefore $v = f(\sigma)$. Since every simplex $\sigma \in K$ has the unique entrance value $f(\sigma)$, we obtain a tame and one-critical filtration of K from the sublevel sets of f . \square

To avoid pathologies, we always assume the filtrations considered in this paper are tame and one-critical.

Now think of the poset $\mathbb{R}_{\preceq}^n = (\mathbb{R}^n, \preceq)$ as a category where the objects are elements of \mathbb{R}^n , and the morphisms are given by the order relation. Let SC be the category of finite simplicial complexes with inclusions as the morphisms. A filtered simplicial complex can be thought of as a functor $\mathcal{K} : \mathbb{R}_{\preceq}^n \rightarrow \text{SC}$. Further, for each $i \in \mathbb{Z}$, we can take the i -th homology of a simplicial complex to obtain a vector space. The two functors above can be composed to obtain a functor from \mathbb{R}_{\preceq}^n to $\text{Vect}_{\mathbb{K}}$, the category of \mathbb{K} -vector spaces. By Definition 1.1, a persistence module can be viewed as a functor from \mathbb{R}_{\preceq}^n to $\text{Vect}_{\mathbb{K}}$, so the construction above is a special case of a persistence module. This motivates the following definition.

Definition 1.10 (Persistent homology). For an integer i , the i -th *multi-parameter persistent homology group* is the persistence module

$$H_i \mathcal{K} : \mathbb{R}_{\leq}^n \longrightarrow \text{Vect}_{\mathbb{K}}$$

defined as the composition of the filtration functor \mathcal{K} with the i -th homology functor for simplicial complexes:

$$\mathbb{R}_{\leq}^n \xrightarrow{\mathcal{K}} \text{SC} \xrightarrow{H_i} \text{Vect}_{\mathbb{K}} .$$

$$\searrow \xrightarrow{H_i \mathcal{K}} \nearrow$$

We now introduce the basic constructions of discrete Morse theory, which we will heavily use in the remainder of the paper. More details may be found in [18].

Definition 1.11 (Discrete vector field). A *discrete vector field* \mathcal{V} on K is a collection of pairs of simplices (α, β) of K where α is a facet of β , such that each simplex is in at most one pair of \mathcal{V} . Given a discrete vector field \mathcal{V} on a simplicial complex K , a *V-path* is a sequence of simplices of dimensions p and $p + 1$,

$$\alpha_0, \beta_0, \alpha_1, \beta_1, \alpha_2, \dots, \beta_r, \alpha_{r+1}$$

such that for each $i = 0, \dots, r$, we have $(\alpha_i, \beta_i) \in \mathcal{V}$ and $\beta_i > \alpha_{i+1} \neq \alpha_i$. A path is called a *non-trivial closed path* if $r > 0$ and $\alpha_0 = \alpha_{r+1}$.

Definition 1.12. A discrete vector field \mathcal{V} on a simplicial complex K is called a *gradient vector field* if it contains no non-trivial closed V -paths. A simplex $\sigma \in K$ is *critical* if it is not paired in \mathcal{V} .

An example of such a discrete gradient vector field can be found in Figure 1.




Figure 1: An example of discrete gradient vector field. The pairs in the vector field are denoted by arrows and the critical cells are marked by a circle.

Definition 1.13 (Elementary collapse). Let K_1, K_2 be simplicial complexes such that $K_2 \subset K_1$, and $K_1 \setminus K_2 = \{\alpha, \beta\}$ where α is a free facet of β . Then the combinatorial deformation retract of K_1 to K_2 given by removing α and β is called an *elementary (simplicial) collapse*. The pair (α, β) is called a *collapsing pair*.

Remark. The definition of elementary collapse is also valid for CW-complexes.

Example 1.14. In Figure 2 we can see how to collapse the complex in Figure 1 using the given discrete gradient vector field. The first collapsing pair is formed by the edge α and the triangle β . The second collapse uses the pair (γ, δ) , given by the edge δ and the vertex γ .




Figure 2: The complex on the right is obtained by the one on the left by first collapsing the pair (α, β) , then the pair (γ, δ) .

An elementary simplicial collapse is a homotopy equivalence between two simplicial complexes, which in turn induces an isomorphism on the level of homology. Therefore, if K_1 and K_2 are related through a series of elementary collapses, then $H_i(K_2) \cong H_i(K_1)$ for all i .

In the case when a simplicial complex has no free faces available, it is still possible to simulate an elementary collapse, which in this case is called *internal*, by first removing a critical cell in order to obtain a free face, and then reinserting it after updating the incidence relations. The cell complex obtained in this case may no longer be simplicial, but internal collapses still induce isomorphisms in homology.

To turn simplicial collapses from just homology preserving into persistent homology preserving transformations, it is convenient to confine ourselves to considering gradient vector fields compatible with filtrations.

Definition 1.15 (Consistency). A discrete gradient vector field \mathcal{V} on a simplicial complex K is said to be *consistent* (or *compatible*) with a filtration $\mathcal{K} = \{K^u\}_{u \in \mathbb{R}^n}$ of K if the following condition is satisfied:

$$\forall (\sigma, \tau) \in \mathcal{V}, \sigma \in K^u \iff \tau \in K^u.$$

In this case, we say that a value $u \in \mathbb{R}^n$ is a *critical value of K* if it is the entrance value of a critical cell of \mathcal{V} .

Example 1.16. In Figure 3 we see a bifiltration, i.e. a filtration with 2 parameters, of a finite simplicial complex. This filtration is one-critical and the gradient vector field is consistent with the given filtration.




Figure 3: Each box corresponds to a multi-parameter in \mathbb{R}^2 , showing the simplices that are present at the corresponding step of the filtration. As before, the pairs of the discrete gradient vector field are indicated by arrows and the critical cells by circles. As we will see later on, we are interested in the entrance values of the critical cells. When the critical cells enter the filtration, they are denoted in a darker red, whereas they are denoted in a lighter orange for larger values, so as to not lose track of them throughout the filtration.

2 Computing the Rank Invariant

Let K be a finite simplicial complex, \mathcal{V} be a discrete gradient vector field on K , and $\mathcal{K} = \{K^u\}_{u \in \mathbb{R}^n}$ be a one-critical n -parameter filtration on K which is consistent with \mathcal{V} . Then, for each integer i , there is an n -parameter persistence module \mathbb{V}_i where the \mathbb{K} -vector space associated to $u \in \mathbb{R}^n$ is $H_i(K^u)$, the i -th homology group of K^u . Furthermore, for $u \preceq v \in \mathbb{R}^n$, $i_{u,v}$ is the induced map, on the level of homology, of the inclusion of K^u into K^v . The goal of this section is to identify a finite subset of values in \mathbb{R}^n from which the rank invariant of \mathbb{V}_i can be computed. In other words, we want some finite $U \subset \mathbb{R}^n$ such that, for all $u \preceq v \in \mathbb{R}^n$, there exists $\bar{u} \preceq \bar{v} \in U$ such that $\rho_{\mathbb{V}_i}(u, v) = \rho_{\mathbb{V}_i}(\bar{u}, \bar{v})$.

Define C to be the set of critical values of \mathcal{V} :

$$C = \{u \in \mathbb{R}^n \mid \sigma \text{ is critical in } \mathcal{V}, \text{ the entrance value of } \sigma \text{ is } u\}.$$

Let \overline{C} be the closure of C under least upper bound. Theorem 2.4 states that our candidate set U as described above is exactly given by \overline{C} . In order to identify each $u \in \mathbb{R}^n$ with an element of \overline{C} , define $\bar{u} = \max\{u' \in \overline{C} \mid u' \preceq u\}$.

Example 2.1. Figure 4 shows the set of critical values and its closure in the parameter space for the bifiltration in Example 1.16. We can see there that for the value u , its corresponding \bar{u} is the critical value c_2 .




Figure 4: Taking the set of critical values C to consist of $\{c_1, c_2, c_3\}$, $\overline{C} = C \cup \{c_4\}$. The grey area contains all values $x \preceq u$, and in this case $\bar{u} = c_2$.

Lemma 2.2. *Let U be a non-empty finite set closed under least upper bound. For all u in \mathbb{R}^n , the set $U' = \{u' \in U \mid u' \preceq u\}$, if non-empty, admits a (unique) maximum.*

Proof. From the fact that U is non-empty and finite, it admits maximal elements. Let a and b be maximal elements in U' . Because U is closed under least upper bound, there is c in U such that $a, b \preceq c \preceq u$. Hence, $c \in U'$. Thus, $c = a = b$ by maximality of a and b in U' . \square

The idea behind the definition of \overline{C} is as follows. Since the filtration \mathcal{K} has a consistent discrete gradient vector field \mathcal{V} , the tools of discrete Morse theory may be used to identify which elements of \mathcal{K} are guaranteed to be homotopy equivalent, and therefore have isomorphic homology groups. Lemma 2.3 can be used to show the existence of a simplicial collapse, induced by \mathcal{V} , between certain elements of \mathcal{K} ; this simplicial collapse is a homotopy equivalence.

Lemma 2.3. *Let $\mathcal{K} = \{K^u\}_{u \in \mathbb{R}^n}$ be a one-critical filtration and \mathcal{V} a discrete gradient vector field on a finite simplicial complex K consistent with \mathcal{K} . For $u \in \mathbb{R}^n$, let $\bar{u} = \max\{u' \in \overline{C} \mid u' \preceq u\}$, with C the set of critical values of \mathcal{V} and \overline{C} its closure under least upper bound. If $K^u - K^{\bar{u}}$ is non-empty, then it contains two simplices σ, τ such that $(\sigma, \tau) \in \mathcal{V}$ and σ is a free facet of τ .*

Proof. Because $K^u - K^{\bar{u}}$ is non-empty and finite, we can take $\tau \in K^u - K^{\bar{u}}$ of maximal dimension. Denote by $\text{l.u.b.}(u, v)$ the least upper bound of two

elements $u, v \in \mathbb{R}^n$. If τ is critical, there is a critical value \hat{u} such that $\bar{u} \neq \text{l.u.b.}(\hat{u}, \bar{u}) \preceq u$. Since $\text{l.u.b.}(\hat{u}, \bar{u}) \in \bar{\mathcal{C}}$, this contradicts the definition of \bar{u} .

Thus, it must be that τ is non-critical, and belongs to a vector $(\sigma, \tau) \in \mathcal{V}$. Because \mathcal{V} is consistent with \mathcal{K} , σ must belong to $K^u - K^{\bar{u}}$ as well, and since τ is of maximal dimension, it must be that σ is a face of τ . Hence, each of the simplices of maximal dimension in $K^u - K^{\bar{u}}$ belongs to a vector $(\sigma, \tau) \in \mathcal{V}$ with σ also in $K^u - K^{\bar{u}}$, and there are finitely many of them because K is finite.

We claim that one of the pairs of simplices $(\sigma, \tau) \in \mathcal{V}$, with τ of maximal dimension in $K^u - K^{\bar{u}}$, must be a collapsing pair, meaning σ is a free facet of τ . Supposing that this is not the case, we create a non-trivial cyclic path in \mathcal{V} , contradicting the assumption that \mathcal{V} is a discrete gradient vector field.

Choose a pair $(\sigma_1, \tau_1) \in \mathcal{V}$ of simplices in $K^u - K^{\bar{u}}$, with τ_1 of maximal dimension. If σ_1 is not a free facet of τ_1 , there exists another cofacet $\tau_2 \in K^u$ such that $\sigma_1 < \tau_2 \neq \tau_1$. As \mathcal{V} is consistent with \mathcal{K} , σ_1 cannot be added to the filtration after τ_2 , so it must be that $\tau_2 \in K^u - K^{\bar{u}}$ as well. Moreover, since the dimension of τ_1 is maximal and σ_1 is a facet of τ_2 , τ_1 and τ_2 must have the same dimension. Thus also τ_2 is of maximal dimension. Thus, τ_2 is in some non-critical pair $(\sigma_2, \tau_2) \in \mathcal{V}$ of simplices in $K^u - K^{\bar{u}}$. Finally, since \mathcal{V} is a discrete gradient vector field, σ_1 can only exist in one pair of \mathcal{V} , so that $\sigma_2 \neq \sigma_1$.

We may iterate the above argument. After an appropriate finite number of iterations, we obtain the following V -path:

$$\sigma_1 < \tau_1 > \sigma_2 < \tau_2 > \cdots < \tau_{n-1} > \sigma_n < \tau_n > \sigma_{n+1}.$$

Since there are only finitely many possible choices for σ_i , it must be that $\sigma_i = \sigma_j$ for some $i \neq j \in \{1, \dots, n+1\}$. The portion of the V -path between σ_i and σ_j is non-trivial and cyclic. This is a contradiction, as \mathcal{V} is a discrete gradient vector field. Thus, for some pair $(\sigma, \tau) \in \mathcal{V}$ of simplices in $K^u - K^{\bar{u}}$ with τ of maximal dimension, σ must be a free facet of τ , proving our claim. \square

We may now prove the main result of the section, which gives a formula for computing the rank invariant for any pair (u, v) , using elements of $\bar{\mathcal{C}}$.

Theorem 2.4. *Let K be a finite simplicial complex, and let $\mathcal{K} = \{K^u\}_{u \in \mathbb{R}^n}$ be a filtration on K . Suppose \mathcal{V} is a discrete gradient vector field on K consistent with \mathcal{K} . Let $\mathbb{V}_i = H_i(\mathcal{K})$. Then, for all $u \preceq v$,*

$$\rho_{\mathbb{V}_i}(u, v) = \rho_{\mathbb{V}_i}(\bar{u}, \bar{v})$$

with

$$\bar{u} = \max\{u' \in \bar{\mathcal{C}} \mid u' \preceq u\},$$

$$\bar{v} = \max\{v' \in \bar{\mathcal{C}} \mid v' \preceq v\}.$$

if $\{u' \in \bar{\mathcal{C}} \mid u' \preceq u\}$ is non-empty, and $\rho_{\mathbb{V}_i}(u, v) = 0$ otherwise.

Proof. Let $u, v \in \mathbb{R}^n$. Because C is non-empty and \overline{C} contains all least upper bounds, if $\{u' \in \overline{C} \mid u' \preceq u\} = \emptyset$, then $K^u = \emptyset$, yielding $\rho_{\mathbb{V}_i}(u, v) = 0$. Otherwise, if $\{u' \in \overline{C} \mid u' \preceq u\}$ is non-empty, then also $\{v' \in \overline{C} \mid v' \preceq v\}$ is so, and $\overline{u}, \overline{v}$ exist and are unique by Lemma 2.2.

Set $K_0 := K^u$. If $K_0 = K^{\overline{u}}$, then obviously $H_i(K^u) = H_i(K^{\overline{u}})$ for all i . Otherwise, $K_0 - K^{\overline{u}}$ is non-empty and by Lemma 2.3 there is a pair $(\sigma_1, \tau_1) \in \mathcal{V}$ of simplices of $K_0 - K^{\overline{u}}$ such that σ_1 is a free facet of τ_1 . Let K_1 be the simplicial complex obtained by performing the elementary collapse of σ_1 onto τ_1 ; note that this means that the map induced on homology by the inclusion of K_1 into K_0 is an isomorphism: $H_i(K_1) \cong H_i(K_0)$. Now, either $K_1 = K^{\overline{u}}$, in which case obviously $H_i(K_1) = H_i(K^{\overline{u}})$ for all i , or, we may restrict \mathcal{V} to K^1 and consider the filtration of K_1 induced from that of K_0 . By applying Lemma 2.3 to K_1 , there is a vector in \mathcal{V} whose elementary collapse gives K_2 such that the map induced by the inclusion of K_2 into K_1 is an isomorphism: $H_i(K_2) \cong H_i(K_1)$. By induction, for any $r \geq 1$, either $K_r = K^{\overline{u}}$, in which case obviously $H_i(K_r) = H_i(K^{\overline{u}})$ for all i , or, we may restrict \mathcal{V} to K_r and consider the filtration of K_r induced from that of K_{r-1} . By applying Lemma 2.3 to K_r , there is a vector in \mathcal{V} whose elementary collapse gives K_{r+1} such that the map induced by the inclusion of K_{r+1} into K_r is an isomorphism: $H_i(K_{r+1}) \cong H_i(K_r)$. By the finiteness of K , there must be a value of r such that $K_r = K^{\overline{u}}$, yielding that the map induced by the inclusion of $K^{\overline{u}}$ into K^u is an isomorphism for all i :

$$H_i(K^u) \cong H_i(K^{\overline{u}})$$

Analogous argument works for v :

$$H_i(K^v) \cong H_i(K^{\overline{v}}).$$

Because the above isomorphisms are induced by inclusions, the following diagram commutes and gives equality of the i^{th} rank invariants of the pairs (u, v) and $(\overline{u}, \overline{v})$:

$$\begin{array}{ccc} H_i(K^u) & \xrightarrow{i_*^{u,v}} & H_i(K^v) \\ \downarrow \cong & & \downarrow \cong \\ H_i(K^{\overline{u}}) & \xrightarrow{i_*^{\overline{u},\overline{v}}} & H_i(K^{\overline{v}}). \end{array}$$

□

3 Computing the persistence space

For the sake of visualization, the rank invariant of an n -parameter persistence module \mathbb{V} can be completely encoded as a multiset of points known as a *persistence diagram* when $n = 1$ [15], and as a *persistence space* when $n \geq 1$ [14]. By completeness of the encoding we mean that the rank invariant can be exactly

reconstructed from the persistence space (cf. the k -Triangle Lemma in [15] and the Representation Theorem in [14]).

While the persistence space is easier to visualize than the rank invariant, as it is a set of points rather than a function, still for an n -parameter persistence module it lives in a $2n$ -dimensional space. So, it is convenient to visualize it along fibers [12]. For example, PHOG [9] and RIVET [24] visualize the persistence space of a 2-parameter persistence module by fibering it through lines.

The goal of this section is to propose a computational procedure to recover such fibration along lines for persistence modules with any number of parameters, by using critical values of gradient vector fields.

We start by reviewing the necessary definitions and properties.

A point $(u, v) \in \mathcal{H}^n$ belongs to the persistence space $\text{spc}(\mathbb{V})$ if and only if its *multiplicity*

$$\mu_{\mathbb{V}}(u, v) := \min_{\substack{\vec{e} \succ 0 \\ u + \vec{e} \prec v - \vec{e}}} \rho_{\mathbb{V}}(u + \vec{e}, v - \vec{e}) - \rho_{\mathbb{V}}(u - \vec{e}, v - \vec{e}) + \quad (1)$$

$$- \rho_{\mathbb{V}}(u + \vec{e}, v + \vec{e}) + \rho_{\mathbb{V}}(u - \vec{e}, v + \vec{e}) \quad (2)$$

is positive. This corresponds to the number of independent cycles that, along a positive direction in the parameter space, appear at u and become boundaries at v .

Similarly, a point (u, ∞) belongs to the persistence space of \mathbb{V} if and only if its multiplicity

$$\mu_{\mathbb{V}}(u, \infty) := \min_{\substack{\vec{e} \succ 0 \\ v \succeq u}} \rho_{\mathbb{V}}(u + \vec{e}, v) - \rho_{\mathbb{V}}(u - \vec{e}, v) \quad (3)$$

is positive. This corresponds to the number of independent cycles that, along a positive direction in the parameter space, appear exactly at u and persist for every larger value of the parameter.

In both cases, the multiplicity can be computed by fixing a direction for \vec{e} and only varying its length (with alternate sums of the ranks decreasing as the length decreases). Two convenient directions for \vec{e} are the diagonal direction and the $v - u$ direction. Moreover, for points at infinity, the multiplicity is reached for increasing values of v .

In particular, for $n = 1$, the persistence space is the persistence diagram of a 1-parameter persistence module. In terms of intervals in a persistence module bar decomposition, points in \mathcal{H}^n of positive multiplicity correspond to finite intervals, points at infinity of positive multiplicity correspond to infinite intervals.

3.1 Restriction of a persistence module to lines

Given a line L contained in the parameter space \mathbb{R}^n , each point $u \in L$ can be written as $u = \vec{m}t + u_0$, with u_0 a fixed starting point on L , $\vec{m} \in \mathbb{R}^n$ a fixed

velocity vector, and t a real parameter. If $\vec{0} \prec \vec{m}$, we say that L has *positive slope*.

For an n -parameter persistence module \mathbb{V} and a line $L \subseteq \mathbb{R}^n$ with positive slope, the *restriction* of \mathbb{V} to L is the persistence module \mathbb{V}_L that assigns V_u to each $u \in L$, and whose transition maps $i_L^{u,v} : V_u \rightarrow V_v$ for $u \preceq v \in L$ are the same as in \mathbb{V} . Once a parametrization $u = \vec{m}t + u_0$ of L is fixed, the persistence module \mathbb{V}_L is isomorphic to the 1-parameter persistence module, by abuse of notation denoted by \mathbb{V}_L , that assigns to each $t \in \mathbb{R}$ the vector space $(\mathbb{V}_L)_t = \mathbb{V}_u$, and to $s < t \in \mathbb{R}$, the transitions maps $i^{s,t} = i_L^{u,v} = i^{u,v}$.

By construction, for $u = \vec{m}s + u_0$ and $v = \vec{m}t + u_0$, it holds that

$$\rho_{\mathbb{V}}(u, v) = \rho_{\mathbb{V}_L}(s, t).$$

Hence, the multiplicity of a point $(u, v) \in \mathcal{H}^n$ in $\text{spc}(\mathbb{V})$ coincides with that of $(s, t) \in \mathcal{H}^1$ in $\text{dgm}(\mathbb{V}_L)$:

$$\mu_{\mathbb{V}}(u, v) = \mu_{\mathbb{V}_L}(s, t).$$

In conclusion, the persistence space $\text{spc}(\mathbb{V})$ can be viewed as the fibered union of infinitely many persistence diagrams $\text{dgm}(\mathbb{V}_L)$, each associated with a line L with positive slope.

3.2 Critical values determine the persistence space

Our next goal is to demonstrate that, for a persistence module \mathbb{V} obtained from a tame and one-critical filtration \mathcal{K} of a simplicial complex, points of the persistence space $\text{spc}(\mathbb{V})$ are completely determined by the critical values of a discrete gradient vector field \mathcal{V} compatible with \mathcal{K} . This claim is proven in Proposition 3.7. The underlying idea to prove it is as follows.

As $\text{spc}(\mathbb{V})$ can be viewed as the fibered union of infinitely many persistence diagrams $\text{dgm}(\mathbb{V}_L)$, with each \mathbb{V}_L obtained by restricting \mathbb{V} to a line L with positive slope, the filtration \mathcal{K} may also be restricted to L . This way we obtain a 1-parameter filtration \mathcal{K}_L , and \mathbb{V}_L turns out to be the persistence module of \mathcal{K}_L . Moreover, if \mathcal{K} has a compatible discrete gradient vector field \mathcal{V} , then this discrete gradient vector field is inherited by \mathcal{K}_L . Each critical cell of \mathcal{V} has an entrance value in \mathcal{K}_L (as L has positive slope). As is the case with \mathcal{K} , the entrance values in \mathcal{K}_L of the critical cells in \overline{C} identify elements in \mathcal{K}_L where the filtration may undergo a change in homotopy type, and therefore a change in homology. Therefore, to determine $\text{spc}(\mathbb{V})$, it is enough to identify the entrance values of critical cells of \mathcal{K} in the restricted filtration \mathcal{K}_L . To this end, we introduce the following notation.

For every point u in \mathbb{R}^n , let $S_+(u)$ be the positive cone with vertex u :

$$S_+(u) = \{v \in \mathbb{R}^n : u \preceq v\}$$

The boundary of the positive cone, $\partial S_+(u)$, decomposes into open faces. In particular, $\partial S_+(u)$ can be partitioned by non-empty subsets A of $[n] = \{1, 2, \dots, n\}$ in the following way. For $\emptyset \neq A \subseteq [n]$, define

$$S_A(u) = \{(x_1, \dots, x_n) \in \mathbb{R}^n | x_i = u_i \text{ for } i \in A, x_j > u_j \text{ for } j \notin A\}.$$

Then, for $A \neq B \subseteq [n]$, $S_A \cap S_B = \emptyset$, and $\partial S_+(u) = \bigcup_{\emptyset \neq A \subseteq [n]} S_A$.

Example 3.1. If $n = 2$ and $u = (u_1, u_2)$, the open faces of $\partial S_+(u)$ consist of the vertex u and the two half-lines exiting from u rightwards and upwards, respectively as shown in Figure 5.

$$\begin{aligned} \partial S_+(u) &= S_{\{1,2\}}(u) \cup S_{\{1\}}(u) \cup S_{\{2\}}(u) \\ &= \{(u_1, u_2)\} \cup \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 = u_1, x_2 > u_2\} \\ &\quad \cup \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 > u_1, x_2 = u_2\}. \end{aligned}$$




Figure 5: The positive cone $S_+(u)$ of $u \in \mathbb{R}^2$ and the decomposition of its boundary into $S_{\{1\}}$, $S_{\{2\}}$ and $S_{\{1,2\}}$, which correspond respectively to the vertical boundary, horizontal boundary, and u .

It will be useful to consider the projection of points in the parameter space onto lines with positive slope (cf. [22]).

Definition 3.2. Given a line $L \subseteq \mathbb{R}^n$ with positive slope, for every $u \in \mathbb{R}^n$ define

$$\text{push}_L(u) := L \cap \partial S_+(u).$$

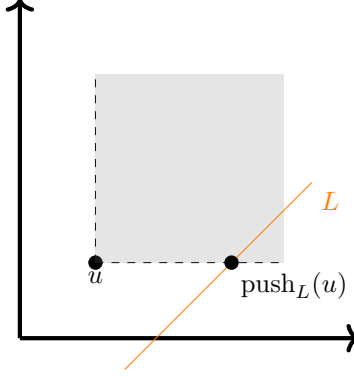


Figure 6: The push of u along the line L .

Proposition 3.3 (Properties of $\text{push}_L(u)$). *Some properties of $\text{push}_L(u)$ are (see also Figure 6):*

1. $\text{push}_L(u)$ consists of exactly one point because L has positive slope.
2. There is a unique non-empty subset A_u^L of $[n]$ such that

$$\text{push}_L(u) = L \cap S_{A_u^L}(u)$$

For ease of notation, we concisely write $S_L(u)$ meaning $S_{A_u^L}(u)$.

3. $u \preceq \text{push}_L(u)$ with equality only when $u \in L$.
4. $\text{push}_L(u)$ is the smallest point on L which is greater than or equal to u ; smaller points on L are either incomparable or less than u .
5. If $u \preceq v$, then $\text{push}_L(u) \preceq \text{push}_L(v)$.
6. Let $u \preceq v$. Let $\emptyset \neq A, B \subseteq [n]$ such that $\text{push}_L(u) \in S_A(u)$ and $\text{push}_L(v) \in S_B(v)$. We have:

- (a) $S_A(u) \cap S_B(v) \neq \emptyset$ implies that $A \subseteq B$.
- (b) $\text{push}_L(u) = \text{push}_L(v)$ if and only if $S_A(u) \cap S_B(v) \neq \emptyset$.

Proof. Properties 1, 2, 3 and 4 are immediate.

Proof of property 5: Suppose not; then $u \preceq v$ and $\text{push}_L(v) \prec \text{push}_L(u)$. Note that, since both $\text{push}_L(u)$ and $\text{push}_L(v)$ are points on L with positive slope, $\text{push}_L(v) \prec \text{push}_L(u)$ if and only if each coordinate is strictly less than, i.e., $(\text{push}_L(v))_i < (\text{push}_L(u))_i$ for all $i \in [n]$. Also, since L has positive slope, there must exist at least one $j \in [n]$ such that $u_j = (\text{push}_L(u))_j$. Combining these, we have

$$v_j \leq (\text{push}_L(v))_j < (\text{push}_L(u))_j = u_j$$

which contradicts $u \preceq v$. So, the claim holds.

Proof of property 6(a): Suppose there exists $y \in S_A(u) \cap S_B(v)$. By definition, $y_i > v_i \geq u_i$ for all $i \notin B$; $y_i > u_i$ implies that $i \notin A$. Thus, $i \notin B$ implies $i \notin A$, and the contrapositive must also be true, $j \in A$ implies $j \in B$.

Note that the converse is not necessarily true; one could have $A = B = \{1\}$ but $u_1 < v_1$, so that $S_A(u) \cap S_B(v) = \emptyset$.

Proof of property 6(b): Indeed, $\text{push}_L(u) = \text{push}_L(v)$ implies that $S_A(u) \cap S_B(v) \neq \emptyset$. Now, suppose $S_A(u) \cap S_B(v) \neq \emptyset$. By Property 6(a), this implies that $A \subseteq B$, and $u_i = v_i$ for $i \in A$. Now, since $\text{push}_L(u)$ and $\text{push}_L(v)$ both belong to L , a line with positive slope, either $\text{push}_L(u) = \text{push}_L(v)$ or $\text{push}_L(u)_j < \text{push}_L(v)_j$ for all $j \in [n]$. And, since $\text{push}_L(u) \in S_A(u)$ and $\text{push}_L(v) \in S_A(v)$, $\text{push}_L(u)_j = u_j = v_j = \text{push}_L(v)_j$ for all $j \in A$. Therefore, $\text{push}_L(u) = \text{push}_L(v)$. □

Recall the notations of Theorem 2.4 where a single bar on some value $u \in \mathbb{R}^n$, for which $\{u' \in \overline{C} \mid u' \preceq u\}$ is non-empty, denotes the greatest value in \overline{C} less than or equal to that value:

$$\bar{u} := \max\{u' \in \overline{C} \mid u' \preceq u\}.$$

We also introduce a double bar notation that depends on a given line L with positive slope (see also Figure 7):

$$\bar{u}^L := \max\{u' \in \overline{C} \mid u \preceq u' \text{ and } S_L(u) \cap S_L(u') \neq \emptyset\}.$$




Figure 7: The double bar of u with respect to the line L is the value $c_4 \in \overline{C}$. In this case we also have $\bar{c}_3^L = \bar{u}^L = c_4$.

Lemma 3.4. For every $u \in \overline{C}$, it holds that

$$\overline{\text{push}_L(u)} = \bar{u}^L.$$

Proof. First we note that for all $u \in \overline{C}$, using notation from Theorem 2.4,

$$\begin{aligned}\overline{\text{push}_L(u)} &= \max\{u' \in \overline{C} \mid u' \preceq \text{push}_L(u)\} \\ &= \max\{u' \in \overline{C} \mid u \preceq u' \preceq \text{push}_L(u)\}\end{aligned}$$

as $u \in \overline{C}$, and $u \preceq \text{push}_L(u)$. Note that $u \preceq u'$ implies $\text{push}_L(u) \preceq \text{push}_L(u')$, and $u' \preceq \text{push}_L(u)$ implies that $\text{push}_L(u') \preceq \text{push}_L(\text{push}_L(u)) = \text{push}_L(u)$. So, we may write

$$\overline{\text{push}_L(u)} = \max\{u' \in \overline{C} \mid u \preceq u' \text{ and } \text{push}_L(u) = \text{push}_L(u')\}.$$

By definition of $\emptyset \neq A_u^L, A_{u'}^L \subseteq [n]$ as the unique subsets $\emptyset \neq A_u^L, A_{u'}^L \subseteq [n]$ such that $\text{push}_L(u) \in S_{A_u^L}(u)$ and $\text{push}_L(u) \in S_{A_{u'}^L}(u')$. Hence, $u \preceq u'$ and $\text{push}_L(u) = \text{push}_L(u')$ if and only if $S_{A_u^L}(u) \cap S_{A_{u'}^L}(u') \neq \emptyset$. So, finally we obtain,

$$\begin{aligned}\overline{\text{push}_L(u)} &= \max\{u' \in \overline{C} \mid u \preceq u' \text{ and } S_{A_u^L}(u) \cap S_{A_{u'}^L}(u') \neq \emptyset\} \\ &= \bar{u}^L.\end{aligned}$$

that yields the claim recalling that the notation $S_L(u)$ is a shorthand for $S_{A_u^L}(u)$. \square

Lemma 3.5. *For all lines L with positive slope, and for all $u \prec v \in \overline{C}$, we have*

$$\rho_{\mathbb{V}}(\text{push}_L(u), \text{push}_L(v)) = \rho_{\mathbb{V}}(\bar{u}^L, \bar{v}^L).$$

Proof. Using Theorem 2.4 and Lemma 3.4, we obtain

$$\begin{aligned}\rho_{\mathbb{V}}(\text{push}_L(u), \text{push}_L(v)) &= \rho_{\mathbb{V}}(\overline{\text{push}_L(u)}, \overline{\text{push}_L(v)}) \\ &= \rho_{\mathbb{V}}(\bar{u}^L, \bar{v}^L)\end{aligned}$$

\square

Lemma 3.6. *For all lines L with positive slope, and all $u \in L$, we have*

$$\bar{u} = \overline{\overline{\bar{u}}^L}.$$

Proof. By definition of double bar, $\bar{u} \preceq \overline{\bar{u}}^L$ and $S_L(\bar{u}) \cap S_L(\overline{\overline{\bar{u}}^L}) \neq \emptyset$. So, by Proposition 3.3.6(b),

$$\text{push}_L(\bar{u}) = \text{push}_L\left(\overline{\overline{\bar{u}}^L}\right).$$

Additionally, Proposition 3.3.3 implies that $\overline{\overline{u}}^L \preceq \text{push}_L\left(\overline{\overline{u}}^L\right)$, and, by Proposition 3.3.5, we have $\text{push}_L(\bar{u}) \preceq \text{push}_L(u)$ because $\bar{u} \preceq u$. Moreover, as $u \in L$, $\text{push}_L(u) = u$. So finally we have

$$\bar{u} \preceq \overline{\overline{u}}^L \preceq \text{push}_L\left(\overline{\overline{u}}^L\right) = \text{push}_L(\bar{u}) \preceq \text{push}_L(u) = u.$$

As $\overline{\overline{u}}^L \in \overline{C}$, the above inequalities imply that $\bar{u} = \overline{\overline{u}}^L$ by definition of \bar{u} . \square

Our next goal is to prove that critical values of the discrete vector field on K determine points of the persistence diagram of the restriction along a line through the parameter space. Given such a line L , we may define

$$\text{push}_L(\overline{C}) = \{\text{push}_L(c) \mid c \in \overline{C}\}.$$

Note that, since \overline{C} is finite, so is $\text{push}_L(\overline{C})$. We can order the elements of $\text{push}_L(\overline{C})$ as c^1, c^2, \dots, c^m , with $c^i \prec c^{i+1}$.

Proposition 3.7. *Let L be a line with positive slope. Let $\text{push}_L(\overline{C}) = \{c^1, c^2, \dots, c^m\}$ be increasingly ordered. For all points $u \prec v$ on L , it holds that:*

(i) *If $u = c^i$ and $v = c^j$, then*

$$\mu_{\mathbb{V}}(u, v) = \rho_{\mathbb{V}}(c^i, c^{j-1}) - \rho_{\mathbb{V}}(c^{i-1}, c^{j-1}) - \rho_{\mathbb{V}}(c^i, c^j) + \rho_{\mathbb{V}}(c^{i-1}, c^j),$$

and $\mu_{\mathbb{V}}(u, v) = 0$ if u or v not in \overline{C} .

(ii) *If $u = c^i$, then*

$$\mu_{\mathbb{V}}(u, \infty) = \rho_{\mathbb{V}}(c^i, c^m) - \rho_{\mathbb{V}}(c^{i-1}, c^m),$$

and $\mu_{\mathbb{V}}(u, \infty) = 0$ if u not in \overline{C} .

Proof. From $\overline{C} \neq \emptyset$ we get $\text{push}_L(\overline{C}) \neq \emptyset$. Note that we may partition the line L by points of $\text{push}_L(\overline{C})$. For each $c^i \in \text{push}_L(\overline{C})$, we have

$$c^i = \text{push}_L(\overline{c^i}).$$

Indeed, by the bar notation, $d \preceq \overline{\overline{c^i}} \preceq c^i$ for all $d \in \text{push}_L^{-1}(c^i)$. Thus,

$$c^i = \text{push}_L(d) \preceq \text{push}_L(\overline{\overline{c^i}}) \preceq \text{push}_L(c^i) = c^i.$$

We first consider the case when $u \in L$ and $u \prec c^1$. In this case, $\mu_{\mathbb{V}}(u, v) = 0$ for all $v \succeq u$, and $\mu_{\mathbb{V}}(u, \infty) = 0$. Indeed, we can take $\vec{e} \succ 0$ small enough so that, for all $0 \prec \vec{e}' \preceq \vec{e}$, we have $u - \vec{e}' \prec u \prec u + \vec{e}' \prec c^1$. Hence, Theorem 3.17 $\rho_{\mathbb{V}}(u - \vec{e}', v) = \rho_{\mathbb{V}}(u + \vec{e}', v) = 0$ for all $v \succeq u$.

We now consider the case of a point $u \in L$ such that $c^1 \preceq u$. Let c^i be the maximal element in $\text{push}_L(\overline{C})$ such that $c^i \preceq u$. In this case we claim that

$$\overline{u} = \overline{c^i}$$

Indeed, suppose not. Since $c^i \preceq u$, we must have $\overline{c^i} \preceq \overline{u}$. Therefore,

$$c^i = \text{push}_L(\overline{c^i}) \preceq \text{push}_L(\overline{u}) \preceq \text{push}_L(u) = u.$$

Since $\text{push}_L(\overline{u})$ is in $\text{push}_L(\overline{C})$ and c^i is maximal in $\text{push}_L(\overline{C})$ such that $c^i \preceq u$ (by assumption), we must have that $c^i = \text{push}_L(\overline{u})$. But then $\overline{u} \underline{\preceq} c^i$, with $\overline{u} \in \overline{C}$, implying $\overline{u} \preceq \overline{c^i}$ by the bar notation. So, it must be that $\overline{u} = \overline{c^i}$.

Now, we are ready to prove the first statement in the case $c^1 \preceq u$. First, suppose that u is not in $\text{push}_L(\overline{C})$. Then there is a maximal element in $\text{push}_L(\overline{C})$ such that $c^i \prec u$. If $i \neq m$, there exists $\vec{e} \succ 0$ such that

$$c^i \prec u - \vec{e} \prec u \prec u + \vec{e} \prec c^{i+1};$$

if $i = m$, then there exists $\vec{e} \succ 0$ such that

$$c^m \prec u - \vec{e} \prec u \prec u + \vec{e},$$

and the above inequalities hold for all $0 \prec \vec{e}' \preceq \vec{e}$. Moreover, by the above claim, we have that for all such \vec{e}' , $\overline{u - \vec{e}'} = \overline{u + \vec{e}'} = \overline{c^i}$.

Therefore,

$$\begin{aligned} \mu_{\mathbb{V}}(u, v) &= \min_{\substack{\vec{e} \succ 0 \\ u + \vec{e} \prec v - \vec{e}}} \rho_{\mathbb{V}}(u + \vec{e}, v - \vec{e}) - \rho_{\mathbb{V}}(u - \vec{e}, v - \vec{e}) + \\ &\quad - \rho_{\mathbb{V}}(u + \vec{e}, v + \vec{e}) + \rho_{\mathbb{V}}(u - \vec{e}, v + \vec{e}) \\ &= \min_{\substack{\vec{e} \succ 0 \\ u + \vec{e} \prec v - \vec{e}}} \rho_{\mathbb{V}}(\overline{u + \vec{e}}, v - \vec{e}) - \rho_{\mathbb{V}}(\overline{u - \vec{e}}, v - \vec{e}) + \\ &\quad - \rho_{\mathbb{V}}(\overline{u + \vec{e}}, v + \vec{e}) + \rho_{\mathbb{V}}(\overline{u - \vec{e}}, v + \vec{e}) \\ &= \min_{\substack{\vec{e} \succ 0 \\ u + \vec{e} \prec v - \vec{e}}} \rho_{\mathbb{V}}(\overline{c^i}, v - \vec{e}) - \rho_{\mathbb{V}}(\overline{c^i}, v - \vec{e}) + \\ &\quad - \rho_{\mathbb{V}}(\overline{c^i}, v + \vec{e}) + \rho_{\mathbb{V}}(\overline{c^i}, v + \vec{e}) \\ &= 0. \end{aligned}$$

Similarly, if v is not in $\text{push}_L(\overline{C})$, we obtain $\mu_{\mathbb{V}}(u, v) = 0$.

Now, if $c^j = v$ and $c^i = u$ for $c^i, c^j \in \text{push}_L(\overline{C})$, then we can find $\vec{e} \succ 0$ small enough such that both

$$c^{i-1} \prec c^i - \vec{e} \prec c^i \prec c^i + \vec{e} \prec c^{i+1}$$

and

$$c^{j-1} \prec c^j - \vec{e} \prec c^j \prec c^j + \vec{e} \prec c^{j+1};$$

(note that if $j = m$, then the second set of equalities does not have the final “ $\prec c^{j+1}$ ” term.)

Additionally, for all $0 \prec e^{\vec{j}} \prec \bar{e}$:

- $\overline{c^j - e^{\vec{j}}} = \overline{c^{j-1}}$,
- $\overline{c^j + e^{\vec{j}}} = \overline{c^j}$,
- $\overline{c^i - e^{\vec{j}}} = \overline{c^{i-1}}$, and
- $\overline{c^i + e^{\vec{j}}} = \overline{c^i}$.

Using these, we find that:

$$\begin{aligned}
\mu_{\mathbb{V}}(u, v) &= \min_{\substack{\bar{e} \succ 0 \\ c^i + \bar{e} \prec c^j - \bar{e}}} \rho_{\mathbb{V}}(c^i + \bar{e}, c^j - \bar{e}) - \rho_{\mathbb{V}}(c^i - \bar{e}, c^j - \bar{e}) + \\
&\quad - \rho_{\mathbb{V}}(c^i + \bar{e}, c^j + \bar{e}) + \rho_{\mathbb{V}}(c^i - \bar{e}, c^j + \bar{e}) \\
&= \min_{\substack{\bar{e} \succ 0 \\ c^i + \bar{e} \prec c^j - \bar{e}}} \rho_{\mathbb{V}}(\overline{c^i + \bar{e}}, \overline{c^j - \bar{e}}) - \rho_{\mathbb{V}}(\overline{c^i - \bar{e}}, \overline{c^j - \bar{e}}) + \\
&\quad - \rho_{\mathbb{V}}(\overline{c^i + \bar{e}}, \overline{c^j + \bar{e}}) + \rho_{\mathbb{V}}(\overline{c^i - \bar{e}}, \overline{c^j + \bar{e}}) \\
&= \min_{\substack{\bar{e} \succ 0 \\ c^i + \bar{e} \prec c^j - \bar{e}}} \rho_{\mathbb{V}}(\overline{c^i}, \overline{c^{j-1}}) - \rho_{\mathbb{V}}(\overline{c^{i-1}}, \overline{c^{j-1}}) - \rho_{\mathbb{V}}(\overline{c^i}, \overline{c^j}) + \rho_{\mathbb{V}}(\overline{c^{i-1}}, \overline{c^j}) \\
&= \rho_{\mathbb{V}}(\overline{c^i}, \overline{c^{j-1}}) - \rho_{\mathbb{V}}(\overline{c^{i-1}}, \overline{c^{j-1}}) - \rho_{\mathbb{V}}(\overline{c^i}, \overline{c^j}) + \rho_{\mathbb{V}}(\overline{c^{i-1}}, \overline{c^j}) \\
&= \rho_{\mathbb{V}}(c^i, c^{j-1}) - \rho_{\mathbb{V}}(c^{i-1}, c^{j-1}) - \rho_{\mathbb{V}}(c^i, c^j) + \rho_{\mathbb{V}}(c^{i-1}, c^j).
\end{aligned}$$

To prove the second statement, we again first suppose that u is not in $\text{push}_L(\bar{C})$, and that c^i is the maximal element in $\text{push}_L(\bar{C})$ such that $c^i \prec u$. Then, as in the proof of the first statement, we can find $0 \prec \bar{e}$ such that, for all $0 \prec e^{\vec{j}} \preceq \bar{e}$, $u - e^{\vec{j}} = u + e^{\vec{j}} = \bar{c}^i$. Thus,

$$\begin{aligned}
\mu_{\mathbb{V}}(u, \infty) &= \min_{\substack{\bar{e} \succ 0 \\ v \succeq u}} \rho_{\mathbb{V}}(u + \bar{e}, v) - \rho_{\mathbb{V}}(u - \bar{e}, v) \\
&= \min_{\substack{\bar{e} \succ 0 \\ v \succeq u}} \rho_{\mathbb{V}}(\overline{u + \bar{e}}, v) - \rho_{\mathbb{V}}(\overline{u - \bar{e}}, v) \\
&= \min_{\substack{\bar{e} \succ 0 \\ v \succeq u}} \rho_{\mathbb{V}}(\bar{c}^i, v) - \rho_{\mathbb{V}}(\bar{c}^i, v) \\
&= 0.
\end{aligned}$$

If $u = c^i$ for some $c^i \in \text{push}_L(\bar{C})$, then we can find $0 \prec \bar{e}$ such that, for all $0 \prec e^{\vec{j}} \preceq \bar{e}$, $u - e^{\vec{j}} = \overline{c^{i-1}}$ and $u + e^{\vec{j}} = \bar{c}^i$. We also note that for all $v \in L$ such that $c^m \prec v$, $\bar{v} = \overline{c^m}$.

Thus,

$$\begin{aligned}
\mu_{\mathbb{V}}(u, \infty) &= \min_{\substack{\vec{e} \succ 0 \\ v \succ u}} \rho_{\mathbb{V}}(u + \vec{e}, v) - \rho_{\mathbb{V}}(u - \vec{e}, v) \\
&= \min_{\substack{\vec{e} \succ 0 \\ v \succ u}} \rho_{\mathbb{V}}(\overline{u + \vec{e}}, v) - \rho_{\mathbb{V}}(\overline{u - \vec{e}}, v) \\
&= \min_{v \succ u} \rho_{\mathbb{V}}(\overline{c^{i-1}}, v) - \rho_{\mathbb{V}}(\overline{c^i}, v) \\
&= \min_{v \succ u} \rho_{\mathbb{V}}(c^{i-1}, v) - \rho_{\mathbb{V}}(c^i, v) \\
&= \min_{v \succ u} \rho_{\mathbb{V}}(c^{i-1}, \bar{v}) - \rho_{\mathbb{V}}(c^i, \bar{v}). \\
&= \min_{v \succ u} \rho_{\mathbb{V}}(c^{i-1}, \overline{c^m}) - \rho_{\mathbb{V}}(c^i, \overline{c^m}) \\
&= \rho_{\mathbb{V}}(c^{i-1}, \overline{c^m}) - \rho_{\mathbb{V}}(c^i, \overline{c^m}) \\
&= \rho_{\mathbb{V}}(c^{i-1}, c^m) - \rho_{\mathbb{V}}(c^i, c^m).
\end{aligned}$$

□

3.3 Grouping fibers of persistence spaces by equivalence

We now use critical values to partition the set of all lines of \mathbb{R}^n into equivalence classes, as illustrated in Figure 8, such that the persistence diagrams of the restriction along lines in the same class are easily obtainable from each other by a bijective correspondence.

Definition 3.8. Two lines $L, L' \subseteq \mathbb{R}^n$ with positive slope are said to have the *same reciprocal position* with respect to u if and only if $\text{push}_L(u)$ and $\text{push}_{L'}(u)$ belong to the same open face of $\partial S_+(u)$. Given a non-empty subset U of \mathbb{R}^n , we write $L \sim_U L'$, if L and L' have the same reciprocal position with respect to u for all $u \in U$.

Example 3.9. Figure 8 shows two examples of the equivalence classes of lines yielded by the set \overline{C} of Example 1.16.

Lines with the same reciprocal position with respect to \overline{C} are characterized by the property of hitting the same face of the positive cone of u for each $u \in \overline{C}$:

Lemma 3.10. $L \sim_{\overline{C}} L'$ if and only if $S_L(u) = S_{L'}(u)$ for all values $u \in \overline{C}$. Hence, $A_u^L = A_u^{L'}$ for all values $u \in \overline{C}$.

Proof. Recall that $S_L(u) = S_{A_u^L}(u)$ and A_u^L is the unique non-empty subset of $[n]$ such that $\text{push}_L(u) = L \cap S_{A_u^L}(u)$. Therefore, $\text{push}_L(u) \in S_{A_u^L}(u)$ and $\text{push}_{L'}(u) \in S_{A_u^{L'}}(u)$.

By definition, if $L \sim_{\overline{C}} L'$, then for all $u \in \overline{C}$, $\text{push}_L(u)$ and $\text{push}_{L'}(u)$ belong to the same open face of $\partial S_+(u)$, i.e.




Figure 8: The dashed lines represent the boundaries of the positive cones of the values in \overline{C} . Here L and L' have the same reciprocal position with respect to \overline{C} but L'' does not.

$$S_{A_u^L} = S_{A_u^{L'}}$$

for all $u \in \overline{C}$. This can only happen if $A_u^L = A_u^{L'}$ for all $u \in \overline{C}$. \square

Proposition 3.11. *Given a non-empty subset U in \mathbb{R}^2 , $L \sim_U L'$ defines an equivalence relation on the set of lines with positive slope.*

Proof. We can define $\mathcal{A}_L(U) = \{A_u^L\}_{u \in U}$. By Lemma 3.10, $L \sim_U L'$ if and only if $\mathcal{A}_L(U) = \mathcal{A}_{L'}(U)$. Using this equivalent definition of $L \sim_U L'$, it is clear that \sim_U is reflexive, transitive, and symmetric, and therefore an equivalence relation. \square

The rank invariant on equivalent lines satisfies the following condition.

Proposition 3.12. *If $u \prec v \in L$, with $\{u' \in \overline{C} \mid u' \preceq u\}$ non-empty, and $L \sim_{\overline{C}} L'$, then it holds that*

$$\rho_{\mathbb{V}}(u, v) = \rho_{\mathbb{V}}(\text{push}_{L'}(\bar{u}), \text{push}_{L'}(\bar{v})).$$

Proof. Since $u \preceq v$, it follows from Theorem 2.4 that

$$\rho_{\mathbb{V}}(u, v) = \rho_{\mathbb{V}}(\bar{u}, \bar{v}).$$

As $u, v \in L$, $\bar{u} = \overline{\overline{u}}^L$ and $\bar{v} = \overline{\overline{v}}^L$ by Lemma 3.6, implying

$$\rho_{\mathbb{V}}(\bar{u}, \bar{v}) = \rho_{\mathbb{V}}\left(\overline{\overline{u}}^L, \overline{\overline{v}}^L\right).$$

As $\bar{u}, \bar{v} \in \bar{C}$, and $u \preceq v$ implies $\bar{u} \preceq \bar{v}$, from Lemma 3.5 we get

$$\rho_{\mathbb{V}}(\overline{\bar{u}}^L, \overline{\bar{v}}^L) = \rho_{\mathbb{V}}(\text{push}_L(\bar{u}), \text{push}_L(\bar{v})).$$

□

Lemma 3.13. *If $L \sim_{\bar{C}} L'$, then $\overline{\text{push}_L(u)} = \overline{\text{push}_{L'}(u)}$ for all values $u \in \bar{C}$.*

Proof. By Lemma 3.4,

$$\begin{aligned} \overline{\text{push}_L(u)} &= \max\{u' \in \bar{C} \mid u \preceq u' \text{ and } S_L(u) \cap S_L(u') \neq \emptyset\}, \\ \overline{\text{push}_{L'}(u)} &= \max\{u' \in \bar{C} \mid u \preceq u' \text{ and } S_{L'}(u) \cap S_{L'}(u') \neq \emptyset\}. \end{aligned}$$

So the claim follows because $S_L(u) = S_{L'}(u)$ and $S_L(u') = S_{L'}(u')$ by Lemma 3.10. □

Lemma 3.14. *If $L \sim_{\bar{C}} L'$, then $\overline{\text{push}_{L'}(\bar{u})} = \bar{u}$ for all values $u \in L$.*

Proof. It follows by successively applying Lemmas 3.13, and 3.4, 3.6. □

Lemma 3.15. *If $L \sim_{\bar{C}} L'$, then the correspondence $\sigma : \text{push}_L(\bar{C}) \rightarrow \text{push}_{L'}(\bar{C})$ defined by $\sigma(d) = \text{push}_{L'}(\text{push}_L^{-1}(d))$ for all $d \in \text{push}_L(\bar{C})$, is an order preserving bijective function. In particular, $\sigma(d) = \text{push}_{L'}(\bar{d})$ with $\bar{d} = \max\{u \in \bar{C} : u \preceq d\}$ as usual, for all $d \in \text{push}_L(\bar{C})$.*

Proof. Let $d \in \text{push}_L(\bar{C})$. Then there exists at least one $c \in \bar{C}$ such that $d = \text{push}_L(c)$. We first show that, for all $c \in \text{push}_L^{-1}(d)$,

$$\text{push}_{L'}(c) = \text{push}_{L'}(\bar{d}). \quad (4)$$

By definition of double bar, we have that $S_L(\bar{c}^L) \cap S_L(c) \neq \emptyset$. So, as $L \sim_{\bar{C}} L'$ implies that $S_L(\bar{c}^L) = S_{L'}(\bar{c}^L)$ and $S_L(c) = S_{L'}(c)$ by Lemma 3.10. This means that

$$S_{L'}(\bar{c}^L) \cap S_{L'}(c) \neq \emptyset.$$

Therefore, by Proposition 3.3.6(b), we have that

$$\text{push}_{L'}(\bar{c}^L) = \text{push}_{L'}(c).$$

Now note that $\bar{d} = \overline{\text{push}_L(c)} = \bar{c}^L$ by Lemma 3.4, and so

$$\text{push}_{L'}(\bar{d}) = \text{push}_{L'}(\bar{c}^L) = \text{push}_{L'}(c).$$

Equality (4) implies that σ is a well defined function because \bar{d} is unique by Lemma 2.2 and $\text{push}_{L'}(\bar{d})$ is also unique by Property 1 of Proposition 3.3.

Now for any $d' \in \text{push}_{L'}(\bar{C})$, there exists (at least one) $c' \in \bar{C}$ such that $d' = \text{push}_{L'}(c')$ and $\sigma(\text{push}_L(c')) = \text{push}_{L'}(c')$, showing that $\sigma : \text{push}_L(\bar{C}) \rightarrow \text{push}_{L'}(\bar{C})$ is surjective.

We can analogously define a function $\tau : \text{push}_{L'}(\bar{C}) \rightarrow \text{push}_L(\bar{C})$ by setting $\tau(d') = \text{push}_L(\text{push}_{L'}^{-1}(d'))$ for all $d' \in \text{push}_{L'}(\bar{C})$.

Now we prove that σ and τ are bijective by showing that σ is the inverse of τ : for all $d' \in \text{push}_{L'}(\overline{C})$,

$$\sigma(\tau(d')) = \sigma(\text{push}_L(\text{push}_{L'}^{-1}(d'))) = \text{push}_{L'}(\text{push}_{L'}^{-1}(d')) = d'$$

and, similarly, for all $d \in \text{push}_L(\overline{C})$,

$$\tau(\sigma(d)) = \tau(\text{push}_{L'}(\text{push}_L^{-1}(d))) = \text{push}_L(\text{push}_L^{-1}(d)) = d .$$

Finally, we show that σ is order-preserving: Assume that $d, e \in \text{push}_L(\overline{C})$ with $d \preceq e$. Then $\bar{d} \in \overline{C}$ and $\bar{d} \preceq d \preceq e$ by definition of bar, and therefore

$$\bar{d} \preceq \bar{e} = \max\{u' \in \overline{C} \mid u' \preceq e\} .$$

Hence, by Proposition 3.3.5, we have that

$$\sigma(d) = \text{push}_{L'}(\bar{d}) \preceq \text{push}_{L'}(\bar{e}) = \sigma(e) ,$$

as required. □

Lemma 3.16. *Let $L \sim_{\overline{C}} L'$ be equivalent lines with positive slope. For any $c^i \in \text{push}_L(\overline{C}) = \{c^1, c^2, \dots, c^m\}$ increasingly ordered, let $d^i = \sigma(c^i)$ with σ as in Lemma 3.15. Then,*

$$\mu_{\mathbb{V}}(d^i, d^j) = \mu_{\mathbb{V}}(c^i, c^j)$$

for $c^i \prec c^j$, and

$$\mu_{\mathbb{V}}(d^i, \infty) = \mu_{\mathbb{V}}(c^i, \infty).$$

Proof. Assuming $c^i \prec c^j \in \text{push}_L(\overline{C})$, by Proposition 3.7(i),

$$\mu_{\mathbb{V}}(c^i, c^j) = \rho_{\mathbb{V}}(c^i, c^{j-1}) - \rho_{\mathbb{V}}(c^{i-1}, c^{j-1}) - \rho_{\mathbb{V}}(c^i, c^j) + \rho_{\mathbb{V}}(c^{i-1}, c^j)$$

On the other hand, setting $d^h = \sigma(c^h)$ for $h \in \{i, i-1\}$, and $d^k = \sigma(c^k)$ for $k \in \{j, j-1\}$ yield

$$\begin{aligned} \rho_{\mathbb{V}}(c^h, c^k) &= \rho_{\mathbb{V}}(\overline{c^h}, \overline{c^k}) \quad (\text{by Theorem 2.4}) \\ &= \rho_{\mathbb{V}}(\text{push}_{L'}(\overline{c^h}), \text{push}_{L'}(\overline{c^k})) \quad (\text{by Lemma 3.14}) \\ &= \rho_{\mathbb{V}}(\text{push}_L(\overline{c^h}), \text{push}_L(\overline{c^k})) \quad (\text{by Theorem 2.4}) \\ &= \rho_{\mathbb{V}}(d^h, d^k) \quad (\text{by definition}). \end{aligned}$$

Therefore,

$$\mu_{\mathbb{V}}(c^i, c^j) = \rho_{\mathbb{V}}(d^i, d^{j-1}) - \rho_{\mathbb{V}}(d^{i-1}, d^{j-1}) - \rho_{\mathbb{V}}(d^i, d^j) + \rho_{\mathbb{V}}(d^{i-1}, d^j) = \mu_{\mathbb{V}}(d^i, d^j)$$

with the second equality holding by Proposition 3.7(i) applied to L' .

Analogously, by Proposition 3.7(ii), we can see that

$$\mu_{\mathbb{V}}(c^i, \infty) = \rho_{\mathbb{V}}(c^i, c^m) - \rho_{\mathbb{V}}(c^{i-1}, c^m) = \rho_{\mathbb{V}}(d^i, d^m) - \rho_{\mathbb{V}}(d^{i-1}, d^m) = \mu_{\mathbb{V}}(d^i, \infty).$$

□

Theorem 3.17. *Let $L \sim_{\overline{C}} L'$ be equivalent lines with positive slope, parametrized by $L : u = \vec{m}s + u_0$ and $L' : u = \vec{m}'s' + u'_0$, respectively. Let $\text{dgm}(\mathbb{V}_L)$ and $\text{dgm}(\mathbb{V}_{L'})$ be the persistence diagrams of the restrictions of \mathbb{V} to L and L' , respectively. Then, there exists a multi-bijection (that is, a bijection between sets of points with multiplicities), $\gamma : \text{dgm}(\mathbb{V}_L) \rightarrow \text{dgm}(\mathbb{V}_{L'})$ such that, for all $(s, t) \in \text{dgm}(\mathbb{V}_L)$,*

$$\gamma(s, t) = (s', t') \in \text{dgm}(\mathbb{V}_{L'})$$

with

- $s' \in \mathbb{R}$ such that $\vec{m}'s' + u'_0 = \text{push}_{L'}(\overline{\vec{m}s + u_0})$, and
- $t' \in \mathbb{R}$ such that $\vec{m}'t' + u'_0 = \text{push}_{L'}(\overline{\vec{m}t + u_0})$ if $t \in \mathbb{R}$, while $t' = \infty$ if $t = \infty$.

Proof. For $(s, t) \in \text{dgm}(\mathbb{V}_L)$ with $s < t < \infty$, we have $\mu_{\mathbb{V}_L}(s, t) = \mu_{\mathbb{V}}(u, v)$ with $u = \vec{m}s + u_0$ and $v = \vec{m}t + u_0$ in L . By Proposition 3.7, we can see that (u, v) will be of the form (c^i, c^j) where $c^i \prec c^j \in \text{push}_L(\overline{C})$. By Lemma 3.16, the bijection $\sigma : \text{push}_L(\overline{C}) \rightarrow \text{push}_{L'}(\overline{C})$ such that $\sigma(c^i) = d^i$ with $d^i = \text{push}_{L'}(\overline{c^i})$ for $1 \leq i \leq m$, satisfies $\mu_{\mathbb{V}}(c^i, c^j) = \mu_{\mathbb{V}}(d^i, d^j)$ and $\mu_{\mathbb{V}}(c^i, \infty) = \mu_{\mathbb{V}}(d^i, \infty)$. The parametrization of L' uniquely determines values $s' < t' \in \mathbb{R}$ such that $d^i = \vec{m}'s' + u'_0$ and $d^j = \vec{m}'t' + u'_0$. Since

$$\mu_{\mathbb{V}_{L'}}(s', t') = \mu_{\mathbb{V}}(d^i, d^j) = \mu_{\mathbb{V}}(c^i, c^j) = \mu_{\mathbb{V}_L}(s, t)$$

and

$$\mu_{\mathbb{V}_{L'}}(s', \infty) = \mu_{\mathbb{V}}(d^i, \infty) = \mu_{\mathbb{V}}(c^i, \infty) = \mu_{\mathbb{V}_L}(s, \infty)$$

we see that σ induces the desired multi-bijection

$$\begin{aligned} \gamma : \text{dgm}(\mathbb{V}_L) &\rightarrow \text{dgm}(\mathbb{V}_{L'}) \\ (s, t) &\rightarrow (s', t') \end{aligned}$$

□

4 Conclusions and discussion

Based on the results of this paper, we can derive a method to fiber the rank invariant of multi-parameter persistence along lines of positive slope chosen by a user. This way, the persistence space of an n -parameter persistence module is sliced into persistence diagrams.

The method consists of an offline preprocessing step, where we compute the representative lines and their persistence diagrams, and an interactive step where we can compute the rank invariant for any chosen line in real time. Starting from a discrete gradient vector field \mathcal{V} consistent with the multi-filtration at hand as input data, the offline step requires:

- computing the set C of entrance values of the critical cells of \mathcal{V} ,

- taking the closure \overline{C} of C with respect to their least upper bound,
- partitioning the set of lines with positive slope by the equivalence relation $\sim_{\overline{C}}$ and picking a representative line from each equivalence class (for example, following the procedure shown in Appendix A).
- storing the persistence diagrams of the restriction of filtration to each representative line.

Having pre-computed these data, the interactive part

- takes as input from the user a line L with positive slope,
- detects the representative line L_0 of its equivalence class with respect to $\sim_{\overline{C}}$,
- computes the persistence diagram relative to L by pushing onto L the bar of the persistence diagram relative to L_0 .

The correctness of the method is guaranteed by Theorem 3.17.

The method requires additional routines from computational geometry in order to efficiently detect representatives of equivalence classes of lines, computing the bars of points, and pushing points onto lines. In Appendix A, we propose a method applied to 2-parameter persistence modules to find representative lines for the equivalence classes defined in Section 3. The method is based on a bijection between segments linking points of \overline{C} and lines cutting \overline{C} in two non-empty subsets.

Moreover, the method requires routines for the persistence diagram computation such as those implemented in [1], [2], or [3]. It is worth noticing that, based on the tests using [2] presented in [27], it is more efficient to compute persistence diagrams of the persistence module restricted to lines starting with the Morse complex obtained from \mathcal{V} rather than directly from the original cell complex. In other words, the gradient vector field \mathcal{V} is used twice: its critical cells allow us both to determine representative lines, and to reduce persistence computation

In conclusion, the presented method allows for: computational efficiency, by requiring only linear asymptotic time complexity to obtain the input gradient field, e.g. with the algorithm of [27]; theoretical improvements, by permitting any number of parameters; data analysis and understanding advantages, by making explicit the correspondence between persistence features and critical cells. For future work we plan to extend the algorithm to a larger number of parameters.

The rank invariant fibering along lines is central also in the definition of the *matching distance*, a metric on rank invariants of multi-parameter persistence modules [13, 12] ensuring their stability. In [21], the exact computation of the matching distance is achieved for at most two parameters in polynomial runtime in the number of simplices requiring $O(m^{11})$ runtime and $O(m^4)$ memory, with m denoting the number of simplices. Motivated by the practical need of

decreasing the number of operations, and increasing the number of allowed parameters, our next project will be to extend equivalence classes of lines to pairs of persistence modules and to apply the method presented here to the matching distance exact computation in any number of parameters.

Acknowledgements

This research began at the 2019 Women in Computational Topology (Win-CompTop) workshop in Canberra. We thank Ashleigh Thomas and Elizabeth Stephenson for joining us in the initial discussions during that week. The results of this paper have been presented within the Summer 2020 AATRN Seminars. We thank Anand Deopurkar, Anthony Licata, and Nicholas Proudfoot for helpful conversations related to Appendix A.

References

- [1] GUDHI library, geometry understanding in higher dimensions. <http://gudhi.gforge.inria.fr/>.
- [2] Phat (persistent homology algorithm toolbox). <https://github.com/blazs/phant>.
- [3] TTK the topological toolkit. <https://topology-tool-kit.github.io/>.
- [4] Madjid Allili, Tomasz Kaczynski, and Claudia Landi. Reducing complexes in multidimensional persistent homology theory. *J. Symbolic Comput.*, 78:61–75, 2017.
- [5] Madjid Allili, Tomasz Kaczynski, Claudia Landi, and Filippo Mazoni. Acyclic partial matchings for multidimensional persistence: algorithm and combinatorial interpretation. *J. Math. Imaging Vision*, 61(2):174–192, 2019.
- [6] S. Barannikov. Framed Morse complexes and its invariants. *Advances in Soviet Mathematics*, 21:93–116, 1994.
- [7] Paul Bendich, J. S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *Ann. Appl. Stat.*, 10(1):198–218, 2016.
- [8] Subhrajit Bhattacharya, Robert Ghrist, and Vijay Kumar. Persistent homology for path planning in uncertain environments. *IEEE Transactions on Robotics*, 31(3):578–590, Jun 2015.
- [9] S. Biasotti, A. Cerri, D. Giorgi, and M. Spagnuolo. PHOG: Photometric and geometric functions for textured shape retrieval. *Computer Graphics Forum*, 32(5):13–22, Aug 2013.

- [10] Gunnar Carlsson, Gurjeet Singh, and Afra Zomorodian. Computing multidimensional persistence. In *Algorithms and computation*, volume 5878 of *Lecture Notes in Comput. Sci.*, pages 730–739. Springer, Berlin, 2009.
- [11] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete Comput. Geom.*, 42(1):71–93, 2009.
- [12] Andrea Cerri, Barbara Di Fabio, Massimo Ferri, Patrizio Frosini, and Claudia Landi. Betti numbers in multidimensional persistent homology are stable functions. *Math. Methods Appl. Sci.*, 36(12):1543–1557, 2013.
- [13] Andrea Cerri, Marc Ethier, and Patrizio Frosini. The coherent matching distance in 2D persistent homology. In *Computational topology in image context*, volume 9667 of *Lecture Notes in Comput. Sci.*, pages 216–227. Springer, 2016.
- [14] Andrea Cerri and Claudia Landi. Hausdorff stability of persistence spaces. *Found. Comput. Math.*, 16(2):343–367, 2016.
- [15] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete Comput. Geom.*, 37(1):103–120, 2007.
- [16] Tamal K. Dey and Cheng Xin. Generalized persistence algorithm for decomposing multi-parameter persistence modules, 2020.
- [17] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, 28(4):511–533, November 2002.
- [18] Robin Forman. A user’s guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48:Art. B48c, 35, 2002.
- [19] Patrizio Frosini. Discrete computation of size functions. *Journal of Combinatorics, Information & System Sciences*, 17(3-4):232–250, 1992.
- [20] Sheridan B. Green, Abby Mintz, Xin Xu, and Jessi Cisewski-Kehe. Topology of our cosmology with persistent homology. *CHANCE*, 32(3):6–13, Jul 2019.
- [21] Michael Kerber, Michael Lesnick, and Steve Oudot. Exact computation of the matching distance on 2-parameter persistence modules. In Gill Barequet and Yusu Wang, editors, *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA*, volume 129 of *LIPICs*, pages 46:1–46:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [22] C. Landi. *The Rank Invariant Stability via Interleavings*, volume 13 of *Association for Women in Mathematics Series*, pages 1–10. 2018.
- [23] Yongjin Lee, Senja D. Barthel, Paweł Dłotko, S. Mohamad Moosavi, Kathryn Hess, and Berend Smit. Quantifying similarity of pore-geometry in nanoporous materials. *Nature Communications*, 8(1), May 2017.

- [24] Michael Lesnick and Matthew Wright. Interactive Visualization of 2-D Persistence Modules. 2015. arXiv:1512.00180 [math.AT].
- [25] Michael Lesnick and Matthew Wright. Computing minimal presentations and bigraded betti numbers of 2-parameter persistent homology, 2020.
- [26] Vanessa Robins. *Computational topology at multiple resolutions: Foundations and applications to fractals and dynamics*. PhD thesis, University of Colorado at Boulder, 2000.
- [27] Sara Scaramuccia, Federico Iuricich, Leila De Floriani, and Claudia Landi. Computing multiparameter persistent homology through a discrete Morse-based approach. *Comput. Geom.*, 89:101623, 19, 2020.
- [28] Michael Sinhuber and Nicholas T. Ouellette. Phase coexistence in insect swarms. *Phys. Rev. Lett.*, 119:178003, Oct 2017.
- [29] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, 33(2):249–274, Nov 2004.

Appendix

A Enumerating equivalence classes of lines

In this appendix we describe an algorithm to enumerate the equivalence classes of lines with positive slope with respect to a set \overline{C} in \mathbb{R}^2 . The results of this section will be initially presented in a slightly more general form.

Consider a set P of n points in the real plane, not necessarily in general position. We say that any division of P into two non-empty subsets via a line not passing through any point of P is a *cut* in the plane. Note that a cut determines an equivalence class of lines if and only if the cut can be realised by a line of positive slope.

A.1 Cuts are determined by primitive pairs

A key observation is that we can associate a certain pair of points in P to any given cut. We say that a pair of distinct points in P is *primitive* if no other point of P lies in the interior of the segment joining the two points.

Proposition A.1. *There is a bijection between the primitive pairs of points in P and the cuts of P .*

Proof. We give an explicit bijection, as follows: given a primitive pair of points in P , we can rotate the line joining the two points clockwise by a small amount, around the midpoint of the segment. It is possible to choose a small enough angle so that this rotated line does not pass through any point of P . Moreover, the two original points lie on opposite sides of this line, and so this line defines a cut.




Figure 9: Choosing the initial center of rotation a

Let us now prove that to every cut we can associate a pair of points as above. Let $\{A, B\}$ be a cut, where A and B are non-empty disjoint subsets of P with $P = A \cup B$. We can represent this cut by a dividing line L .

Now fix a direction \vec{v} orthogonal to L . Translate L along \vec{v} until it hits at least one point of P for the first time. Let H be the set of points that L hits. If H has only one element, set a to be the unique element of H . If H has more than one element, then set a to be the one with the following property: for any other point $a_1 \in H$, the dot product of the vector \vec{v} with the vector $\overrightarrow{aa_1}$ is positive. In other words, if we consider \vec{v} to be the positive direction along the x -axis, then a is the point with the minimum y -coordinate (see Figure 9).

Note that until the line hits a , it still defines the same cut. Without loss of generality, assume that $a \in A$. Taking a to be the centre of rotation, rotate the line L counter-clockwise until it hits another point of P . Note that the line defines a cut of all points of P other than those on this line, and it is the same cut except for these points.

If the line simultaneously hits multiple points, then exactly one of the following is true of these points: they all points lie in A , they all lie in B , or they include both points in A and in B . In the third case, a lies between the rest of the points in A and those in B on the line. Now we have an algorithm to generate the primitive pair a, b for the cut L , as follows.

1. If at least one point hit lies in B , set b to be the point of B hit by the line that is closest to a , as shown in Figure 10 (Left).

In this case, rotating the line ab clockwise around the midpoint of the segment ab gets us back the original cut. This is because the line yields the original cut excluding the points of P that lie on it, and therefore rotating it slightly clockwise restores these points to the correct subsets (either A or B). We are done.

2. If all points hit lie in A , reset the center of rotation to be the point of A hit by the line that is farthest away from a as in Figure 10 (Right). Rename




Figure 10: Left: Rotating L around a will hit at least one point in B . Right: Rotating L around a hits only points in A .

this point as a . Continue to rotate the line counter-clockwise around the new centre of rotation until it hits another point of P , and repeat the above steps until the algorithm terminates.

The algorithm terminates when the rotating line hits a point of B . This always happens, for the following reason. As the line makes a full rotation around the convex hull of A , it sweeps through the entire plane except for the convex hull of A , while B is a non-empty set outside the convex hull of A . \square

Remark A.2. There is a similar, counter-clockwise, bijection between cuts and primitive pairs: rotate a line segment joining a primitive pair by a small amount counter-clockwise. Thus each cut corresponds to a “clockwise primitive pair”, and another “counter-clockwise primitive pair”. These pairs are distinct unless all points of P lie on a single line. Correspondingly, there is a counter-clockwise version of the algorithm explained in the previous proposition (see Figure 11). We use this fact in the next section.

A.2 Achieving positive slope

We have determined that every primitive pair of points in P determines a cut by rotating the line segment through this pair slightly clockwise about the midpoint of the segment. This is the *clockwise primitive pair* associated to this cut. Similarly, every primitive pair of points determines another cut by rotating the line segment through this pair slightly counter-clockwise around the midpoint of the segment. This is the *counter-clockwise primitive pair* associated to this cut. These are shown in Figure 11. The clockwise and counter-clockwise primitive pairs associated to any cut can be found by the algorithm in the previous section, and its variant explained in Remark A.2 respectively. We now tackle the problem




Figure 11: The clockwise and counter-clockwise primitive pairs associated to the cut $\{A, B\}$. The pair (a_1, b_1) is the clockwise primitive pair and (a_2, b_2) is the counter-clockwise primitive pair.

of determining whether such a cut can be defined by a line of positive slope. The answer is given by Algorithm 1.

We will need some preparatory lemmas to prove the correctness of the algorithm. We use the same setup for all of these lemmas, as follows. Let (a_1, b_1) and (a_2, b_2) be the clockwise and anti-clockwise primitive pairs respectively for a cut $\{A, B\}$. It may be the case that either $a_1 = a_2$ or $b_1 = b_2$, but we suppose that not both are true. This supposition is true unless all points in P lie on a line.

Let p be the intersection point of the segments a_1b_1 and a_2b_2 . Note that p must exist, for the following reason. If all four points are distinct, then a_2 and b_2 lie on opposite sides of the line a_1b_1 , and so the segments intersect somewhere in their interiors. Otherwise, if $a_1 = a_2$ (resp. $b_1 = b_2$), then $p = a_1 = a_2$ (resp. $p = b_1 = b_2$).

Lemma A.3. *Suppose that the four points a_1, a_2, b_1, b_2 are distinct. If we start at the line a_1b_1 and rotate clockwise around p until we hit a_2b_2 , every intermediate line defines the cut $\{A, B\}$.*

Proof. Let L_1 be the line a_1b_1 and L_2 be the line a_2b_2 . Note that each of the two lines has a well-defined “A” side and a well-defined “B” side: the side of the line on which the remaining points of A lie is the “A” side, and the side on which the remaining points of B lie is the “B” side.

The lines L_1 and L_2 cut up the plane into four open cones. We can label these cones as $C_{AA}, C_{AB}, C_{BB}, C_{BA}$, where for example C_{BA} is the intersection of the “B” side of L_1 with the “A” side of L_2 . The cones C_{AB} and C_{BA} contain no points of P . This is precisely because these cones lie on the “A” side of one of the lines and on the “B” side of the other. Moreover, their closures only intersect at the point p .

Algorithm 1 An algorithm to determine whether a cut can be obtained by a line of positive slope.

```

1: Let  $(a_1, b_1)$  and  $(a_2, b_2)$  be the clockwise and anti-clockwise primitive pairs
   respectively for a cut  $\{A, B\}$ .
2: Let  $m_1$  and  $m_2$  be the slopes of the lines  $a_1b_1$  and  $a_2b_2$  respectively.
3: if  $a_1 = b_1$  and  $a_2 = b_2$  then
4:   return true
5: end if
6: if  $0 < m_1 \leq \infty$  then
7:   return true
8: else if  $0 \leq m_2 < \infty$  then
9:   return true
10: else
11:   We have  $-\infty < m_1 \leq 0$  and  $-\infty \leq m_2 < 0$ .
12:   if  $m_1 < m_2$  then
13:     return true
14:   else
15:     return false
16:   end if
17: end if

```

It is clear that any line rotated clockwise around p starting from L_1 until we hit L_2 , excluding L_1 and L_2 itself, lies completely in the set $C_{AB} \cup C_{BA} \cup \{p\}$ and has a_1 on its “A” side and b_1 on its “B” side. So any such line continues to define the same cut $\{A, B\}$. \square

Lemma A.4. *Suppose that $a_1 = a_2$. If we start at the line a_1b_1 and rotate clockwise around $a_1 = a_2$ until we hit a_1b_2 keeping track of the trace of b_1 under this rotation, then every intermediate line, except for a_1b_2 itself, defines the cut $\{A, B\}$ after a sufficiently small clockwise rotation about the midpoint between a_1 and the trace of b_1 .*

Proof. As in the proof of Lemma A.3, let L_1 be the line a_1b_1 and L_2 be the line a_1b_2 , and notice that L_1 and L_2 cut up the plane into four open cones $C_{AA}, C_{AB}, C_{BB}, C_{BA}$. As in the proof of Lemma A.3, the cones C_{AB} and C_{BA} contain no points of P and their closures only intersect in the point of intersection of L_1 and L_2 , which is a_1 in this case.

It is clear that any line rotated clockwise around a_1 starting from L_1 until we hit L_2 , excluding L_1 and L_2 itself, lies completely in the set $C_{AB} \cup C_{BA} \cup \{a_1\}$ and has b_1 on its “B” side, so defines the cut $\{A, B\}$ excluding a_1 . A sufficiently small clockwise rotation of such a line about the midpoint between a_1 and the trace of b_1 moves a to its “A” side without crossing any other points in P , so the resulting line indeed defines the cut $\{A, B\}$. \square

Lemma A.5. *Let S be the set of possible slopes of lines obtained by starting at the line a_1b_1 and rotating clockwise through p until we hit the line a_2b_2 , excluding the slopes of the lines a_1b_1 and a_2b_2 themselves. Then the slope of any line L that defines the same cut $\{A, B\}$ lies in S .*

Proof. For the proof of this lemma, the four points a_1, a_2, b_1, b_2 need not all be distinct. First note that if L is any line defining the cut $\{A, B\}$, then it intersects the interiors of the segments a_1b_2 and a_2b_1 . This is precisely because $\{a_1, a_2\}$ and $\{b_1, b_2\}$ lie on opposite sides of L .

Now suppose L is *any* line that intersects the interiors of the segments a_1b_2 and a_2b_1 . These are two opposite sides of the (possibly degenerate) quadrilateral $a_1a_2b_1b_2$. Therefore L must also intersect the interiors of the diagonals of this (possibly degenerate) quadrilateral, namely the segments a_1b_1 and a_2b_2 . In particular, because L intersects both a_1b_1 and a_2b_2 , it cannot have slope equal to either m_1 or m_2 .

The set of possible slopes of lines in the plane can be identified with the real projective line, by noting that slopes can lie between $[-\infty, \infty]$ with $\infty = -\infty$. We now have a continuous map

$$a_1b_2 \times a_2b_1 \rightarrow \mathbb{RP}^1, \tag{5}$$

defined by mapping an ordered pair of points to the slope of the line joining the two points. By the previous argument, the image of this map lies in $\mathbb{RP}^1 \setminus \{m_1, m_2\}$, which has two connected components. Since the domain is connected, the image of the map must lie in exactly one of the connected components.

The set S is precisely one of the two connected components: we start at m_1 , rotate clockwise until we hit m_2 . The other connected component is obtained by rotating counter-clockwise starting at m_1 until we hit m_2 .

We already know by either Lemma A.3 or Lemma A.4 (depending on whether or not the four points a_1, a_2, b_1, b_2 are distinct) that there are points in the image of the map in Equation 5 that lie in S . By connectedness, all lines that intersect the interiors of a_1b_2 and a_2b_1 have slopes that lie in S . In particular, all lines that define the same cut $\{A, B\}$ have slopes that lie in S . \square

Now we can prove the correctness of the algorithm.

Proposition A.6. *Algorithm 1 correctly determines whether a cut can be obtained by a line of positive slope.*

Proof. Recall that (a_1, b_1) and (a_2, b_2) are the clockwise and anti-clockwise primitive pairs respectively for a cut $\{A, B\}$. Recall that m_1 and m_2 are the slopes of the lines a_1b_1 and a_2b_2 respectively. We treat each step of the algorithm in order.

First, $a_1 = a_2$ and $b_1 = b_2$ if and only if all points of P lie on a single line. In this case it is clearly always possible to achieve any cut by a line of positive slope. Now assume that not all points of P lie on a single line, which implies that either $a_1 \neq a_2$ or $b_1 \neq b_2$. This is the setting of the previous lemmas.

If $0 < m_1 \leq \infty$, then a small clockwise rotation of the line a_1b_1 has positive slope. Since the rotated line determines the desired cut $\{A, B\}$, we are done. Similarly, if $0 \leq m_2 < \infty$, then a small counter-clockwise rotation of the line a_2b_2 has positive slope. Since the rotated line determines the desired cut $\{A, B\}$, we are done.

Now suppose that $-\infty < m_1 \leq 0$ and $-\infty \leq m_2 < 0$. Let p be the intersection point of a_1b_1 and a_2b_2 . Let S be the set of possible slopes of lines obtained by starting at the line a_1b_1 and rotating clockwise through p until we hit the line a_2b_2 .

Suppose first that $-\infty < m_1 < m_2 < 0$. As we sweep clockwise from a_1b_1 , we begin at m_1 , decrease slope until we hit a vertical line with slope $-\infty = \infty$, and then decrease again from ∞ until we cross 0 down to m_2 . In particular, at least one of the intermediate lines has positive slope. If the points a_1, a_2, b_1, b_2 are all distinct, then by Lemma A.3 we have a line of positive slope that gives the cut $\{A, B\}$. If two of the four points are equal, then Lemma A.4 states that a sufficiently small clockwise rotation of one of the intermediate lines (which will also have positive slope) gives the cut $\{A, B\}$.

Now suppose that $m_1 \geq m_2$. In this case, the set S consists only of negative numbers: these are the slopes starting from m_1 and decreasing down to m_2 . By Lemma A.5, we see that there is no line of positive slope that defines this cut. \square

A.3 Cuts through a fixed point

In order to address equivalence classes of lines that pass through a given point c of \overline{C} , we now say that a division of a non-empty set of points P of the plane, with $c \notin P$, into two disjoint subsets A and B , of which at most one can be empty, via a line passing through the given point c and disjoint from P , is a c -cut of P . Reciprocally, we say that a line through c and a point of P is a c -primitive line of P .

Proposition A.7. *There is a bijection between the c -primitive lines of P and the c -cuts of P .*

Proof. We construct an explicit bijection, called the *clockwise bijection*, as follows. Given a c -primitive line of P , we can rotate this line clockwise by a small amount, around c , so that this rotated line does not pass through any point of P . This line defines a c -cut. Vice versa, with every c -cut of P we can associate a c -primitive line by rotating the line realing the c -cut counter-clockwise until it hits some point of P , which exists because P is non-empty. Note that by a completely symmetric argument we also have a *counter-clockwise bijection*. \square

We have determined that any c -cut L of P can be determined by rotating both a c -primitive line L_1 clockwise and a c -primitive line L_2 counter-clockwise. Let m_1 and m_2 be the slopes of L_1 and L_2 , respectively. In the case $m_1 = m_2$, because L_1 and L_2 both pass through c , we have $L_1 = L_2$. In this case all points

of P belong to L_1 and positive slope can be achieved by rotation around c for every value of m_1 .

In the case when $m_1 \neq m_2$, Algorithm 1 applied to c -primitive lines instead of primitive pairs of points achieves the goal. Indeed, again, L_1 and L_2 have a well defined A -side and B -side, and $L - \{c\}$ is contained in $C_{AB} \cup C_{BA}$. Then the argument follows as in Proposition A.6.

A.4 Retrieving representatives lines

With reference to the equivalence relation on lines defined by their reciprocal position with respect to the set \overline{C} of critical values and their least upper bounds as given in Definition 3.8, our goal is to retrieve a representative line with positive slope for each possible equivalence class. Recall from Lemma 3.10 that two lines belong to the same equivalence class with respect to \overline{C} if and only if they hit the positive cone of each point of \overline{C} at the same facet.

There are three possible situations for lines in the same class: (i) There is only one line in the equivalence class passing through two points c and c' of \overline{C} ; (ii) the lines in the considered equivalence class contain exactly one point of \overline{C} , say c ; (iii) the lines in the considered equivalence class do not contain any point of \overline{C} .

Case (i) can be easily solved by taking lines through all possible pairs of distinct points c and c' in \overline{C} , provided that $c \preceq c'$, paying attention to not taking the same line multiple times if there are more than two points on the same line.

Case (ii). In this case, each such line partitions $\overline{C} - \{c\}$ into two subsets A and B . For each equivalence class of lines for which A and B are both empty, c is the only point of \overline{C} , so there is only one such equivalence class and any line through c with positive slope is a representative of it. In the case when at least one between A and B are non-empty, we can obtain a representative line by applying the algorithm presented in Subsection A.2 as explained in Subsection A.3. Note that since \overline{C} is closed under least upper bound, the case where $-\infty < m_1 < m_2 < 0$ (line 12 in Algorithm 1) cannot occur.

Case (iii). In this case, each such line partitions \overline{C} into two subsets A and B . For each equivalence class of lines for which A and B are both non-empty, we can obtain a representative by applying the algorithm presented in Subsection A.2. For the case when either A or B is empty, the other one is necessarily equal to \overline{C} . There are exactly two such equivalence classes of lines depending on whether the lines hit all the positive cones of points of \overline{C} at their horizontal or vertical facets. As a representative of the first class, we can take a line parallel to the diagonal of \mathbb{R}^2 passing to a point with abscissa greater than the maximum abscissa of points of \overline{C} , and ordinate smaller than the minimum ordinate of points of \overline{C} . Symmetrically, as a representative of the second class, we can take a line parallel to the diagonal of \mathbb{R}^2 passing to a point with abscissa smaller than the minimum abscissa of points of \overline{C} , and ordinate greater than the maximum ordinate of points of \overline{C} .