Signed Barcodes for Multi-Parameter Persistence via Rank Decompositions

Magnus Bakke Botnan ✉
Vrije Universiteit Amsterdam, The Netherlands

Steffen Oppermann ✉
Norwegian University of Science and Technology, Trondheim, Norway

Steve Oudot ✉
Inria, Palaiseau, France

──── **Abstract** ────

In this paper we introduce the signed barcode, a new visual representation of the global structure of the rank invariant of a multi-parameter persistence module or, more generally, of a poset representation. Like its unsigned counterpart in one-parameter persistence, the signed barcode encodes the rank invariant as a $\mathbb{Z}$-linear combination of rank invariants of indicator modules supported on segments in the poset. It can also be enriched to encode the generalized rank invariant as a $\mathbb{Z}$-linear combination of generalized rank invariants in fixed classes of interval modules. In the paper we develop the theory behind these rank decompositions, showing under what conditions they exist and are unique – so the signed barcode is canonically defined. We also illustrate the contribution of the signed barcode to the exploration of multi-parameter persistence modules through a practical example.

1 Introduction

1.1 Context and motivation

One of the central questions in the development of multi-parameter persistence theory is to find a proper generalization of the concept of a persistence barcode, which plays a key part in the one-parameter instance of the theory. Given a *one-parameter persistence module*, i.e. a functor M from some subposet $P \subseteq \mathbb{R}$ to the vector spaces over a fixed field $\mathbf{k}$, the *(persistence) barcode* $\mathrm{Dgm}\, M$ is a multi-set of intervals in P that fully characterizes the module M. Its role in applications is motivated by the fact that $\mathrm{Dgm}\, M$ provides a compact encoding of the so-called *rank invariant* $\mathrm{Rk}\, M$, a complete invariant that captures the ranks of the internal morphisms of M, more precisely:

$$\mathrm{Rk}\, M(s, t) = \mathrm{rank}\,[M(s) \to M(t)] \quad \text{for every } s \leq t \in P. \tag{1.1}$$
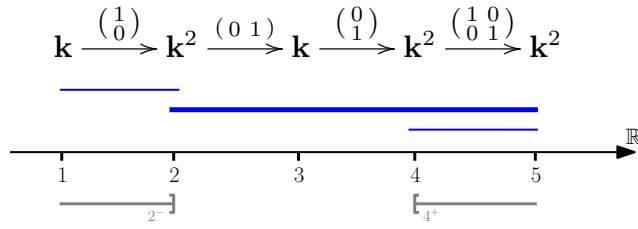
The encoding decomposes $\mathrm{Rk}\, M$ as a $\mathbb{Z}$-linear combination of rank invariants of *interval modules*, i.e. indicator modules supported on intervals:

$$\mathrm{Rk}\, M = \sum_{I \in \mathrm{Dgm}\, M} \mathrm{Rk}\, \mathbf{k}_I = \mathrm{Rk}\left(\bigoplus_{I \in \mathrm{Dgm}\, M} \mathbf{k}_I\right), \tag{1.2}$$

where each interval $I \in \mathrm{Dgm}\, M$ is considered with multiplicity, and where $\mathbf{k}_I$ denotes the interval module supported on I. Coefficients in the $\mathbb{Z}$-linear combination are all positive.

$$\mathbf{k} \xrightarrow{\binom{1}{0}} \mathbf{k}^2 \xrightarrow{(0\ 1)} \mathbf{k} \xrightarrow{\binom{0}{1}} \mathbf{k}^2 \xrightarrow{\binom{1\ 0}{0\ 1}} \mathbf{k}^2$$

🟨 **Figure 1** A one-parameter persistence module M (top) indexed over $\{1,2,3,4,5\}$, and its barcode (in blue). The corresponding rank decomposition is $\mathrm{Rk}\,M = \mathrm{Rk}\,\mathbf{k}_{[\![1,2]\!]} + \mathrm{Rk}\,\mathbf{k}_{[\![2,5]\!]} + \mathrm{Rk}\,\mathbf{k}_{[\![4,5]\!]}$. The rank $\mathrm{Rk}\,M(2,4) = 1$ is given by the one bar (thickened) connecting the down-set 2^- to the up-set 4^+.


🟨 **Figure 2** The indecomposable module M on the left does not have the same rank invariant as any direct sum of interval modules on the 3×3 grid. However, $\mathrm{Rk}\,M$ is equal to the difference between the rank invariants of the two direct sums of interval modules shown on the right. Blue is for intervals counted positively in the decomposition, while red is for intervals counted negatively.

The encoding in (1.2) is unique, i.e. there is no other way to decompose $\mathrm{Rk}\,M$ as a $\mathbb{Z}$-linear combination, with positive coefficients, of rank invariants of interval modules. This is because M itself decomposes essentially uniquely as a direct sum of interval modules [6]:

$$M \simeq \bigoplus_{I \in \mathrm{Dgm}\,M} \mathbf{k}_I. \tag{1.3}$$

Since the intervals in $\mathrm{Dgm}\,M$ are line segments – possibly closed, open, or half-open, $\mathrm{Dgm}\,M$ can be represented graphically as an actual barcode (see Figure 1) that reveals the global structure of the rank invariant $\mathrm{Rk}\,M$, as well as of the module M itself.

Major difficulties arise when trying to generalize the concept of barcode to *multi-parameter persistence modules* – i.e. functors M from $\mathbb{R}^d$ (equipped with the product order) to the vector spaces over $\mathbf{k}$. Foremost, while a direct-sum decomposition of M into indecomposables still exists and is essentially unique [3], the summands may no longer be interval modules as in (1.3), where intervals in $\mathbb{R}^d$ are defined to be connected convex subsets in the product order. For instance, the module on the left-hand side of Figure 2 is indecomposable yet not an interval module nor even an indicator module – its pointwise dimension is not everywhere ≤ 1. One may then ask whether rank decompositions such as (1.2) exist nonetheless. The answer is unfortunately negative: still in Figure 2, the module M on the left-hand side does not have the same rank invariant as any direct sum of interval modules, therefore it cannot decompose as in (1.2). Nevertheless, $\mathrm{Rk}\,M$ can be expressed as the difference between the
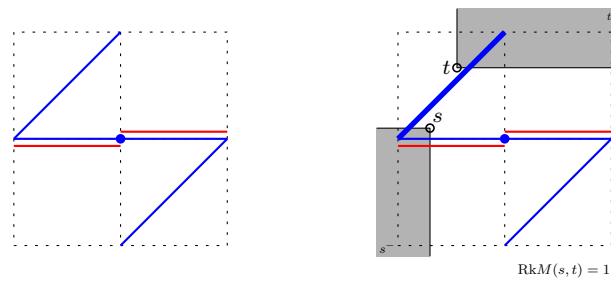
Figure 3 Left: the signed barcode corresponding to the rank decomposition of Figure 2. Each bar is the diagonal with positive slope of one of the rectangles involved in the decomposition, with the same color code (blue for positive sign, red for negative sign). Right: computing $\operatorname{Rk} M(s,t)$ for a pair of indices $s \leq t$ – the thick bar is the only one connecting the down-set s^- to the up-set t^+.

rank invariants of two direct sums of interval modules, as illustrated in the same figure. In other words, $\operatorname{Rk} M$ decomposes as a $\mathbb{Z}$-linear combination of rank invariants of interval modules, with possibly negative coefficients.

The fact that the signed rank decomposition of the module M in Figure 2 involves only rectangles is not mere chance: the point of our work is to show that such decompositions of the rank invariant (1.1) exist, and furthermore that they are essentially unique – which may not be the case for decompositions of this invariant using larger classes of intervals. Uniqueness in the case of rectangles comes from the known fact that the rank invariant is complete on direct sums of *rectangle modules*, i.e. interval modules supported on rectangles [4, 8].

Rectangles are also interesting because they are entirely determined by their upper bound and lower bound. They therefore allow for an alternative representation of the signed rank decomposition as a *signed barcode*, where each bar is the diagonal (with positive slope) of a particular rectangle in the decomposition, with the same sign. As illustrated in Figure 3, the signed barcode encodes visually the global structure of the rank invariant (1.1), and it gives access to the same information as the signed rank decomposition. For instance, the rank $\operatorname{Rk} M(s,t)$ between a pair of indices $s \leq t$ is given by the number of positive bars that connect the down-set $s^- = \{u \in P \mid u \leq s\}$ to the up-set $t^+ = \{u \in P \mid u \geq t\}$, minus the number of negative bars that connect s^- to t^+.

1.2   Our setting

We work more generally over a partially ordered set P, considered as a category in the obvious way, and we let $\mathbf{k}$ be an arbitrary but fixed field. Closed rectangles in $\mathbb{R}^d$ now become *closed segments* in P, defined by $\langle s,t \rangle = \{u \in P \mid s \leq u \leq t\}$. *Intervals* in P are defined as non-empty subsets I that are both convex and connected in the partial order. Denote by $\operatorname{Rep} P$ the functor category consisting of all functors $M \colon P \to \operatorname{Vec}_{\mathbf{k}}$ where $\operatorname{Vec}_{\mathbf{k}}$ is the category of vector spaces over $\mathbf{k}$. We shall refer to such a functor M either as a *representation* of P or as a *persistence module* over P, without distinction. Let $\operatorname{rep} P$ be the subcategory of *pointwise finite-dimensional* (pfd) representations, i.e. functors taking their values in the finite-dimensional vector spaces over $\mathbf{k}$. Denoting by $\{\leq_P\} = \{(a,b) \in P \times P \mid a \leq b\}$ the set of pairs defining the partial order in P, we see the rank invariant (1.1) as a map $\{\leq_P\} \to \mathbb{N}$ (in the literature, the rank invariant is sometimes defined as a map on $P \times P$ that vanishes outside $\{\leq_P\}$; such a map clearly holds the same information as our rank invariant). For $I \subseteq P$, $M|_I$ denotes $M \circ \iota$ where $\iota \colon I \hookrightarrow P$ is the canonical inclusion.

Since the usual rank invariant is incomplete, even on the subcategory of *interval-decomposable* modules (i.e. modules isomorphic to direct sums of interval modules), we will consider a generalization of the rank invariant that probes the existence of "features" in the module across arbitrary intervals $I \subseteq P$, not just across closed segments. This generalization is known to be complete on interval-decomposable modules – see [8] or our Proposition 2.8:

▶ **Definition 1.1.** *Let $M \in \operatorname{Rep} P$. Given an interval $I \subseteq P$, the* generalized rank *of M over I, denoted by $\operatorname{Rk}_I M$, is defined by:*

$$\operatorname{Rk}_I M = \operatorname{rank} \left[\varprojlim M|_I \to \varinjlim M|_I \right].$$

Given a collection $\mathcal{I}$ of intervals, the generalized rank invariant *of M over $\mathcal{I}$ is the map $\operatorname{Rk}_\mathcal{I} M \colon \mathcal{I} \to \mathbb{N} \cup \{\infty\}$ defined by $\operatorname{Rk}_\mathcal{I} M(I) = \operatorname{Rk}_I M$.*

▶ **Remark 1.2.** To see that Definition 1.1 generalizes the standard rank invariant, observe that when I is a closed segment $\langle i, j \rangle = \{u \in P \mid i \leq u \leq j\}$, we have $\operatorname{Rk}_I M = \operatorname{rank} [M(i) \to M(j)]$. Hence, taking $\mathcal{I} = \{\langle i, j \rangle \mid i \leq j \in P\} \simeq \{\leq_P\}$ in the above definition gives $\operatorname{Rk}_\mathcal{I} M = \operatorname{Rk} M$, the usual rank invariant of M.

In this work we focus on the subcategory $\operatorname{rep}_\mathcal{I} P$ of representations M that have a *finite* generalized rank invariant over a fixed collection $\mathcal{I}$ of intervals, i.e. such that $\operatorname{Rk}_I M \in \mathbb{N}$ for all $I \in \mathcal{I}$. Our setting considers in fact arbitrary functions $\mathcal{I} \to \mathbb{Z}$. Note that $\operatorname{rep}_\mathcal{I} P \supseteq \operatorname{rep} P$, since the morphism $\varprojlim M|_I \to \varinjlim M|_I$ factors through the internal spaces of $M|_I$.

▶ **Definition 1.3.** *Given a collection $\mathcal{I}$ of intervals in P, and a function $r \colon \mathcal{I} \to \mathbb{Z}$, a* (signed) rank decomposition *of r over $\mathcal{I}$ is given by the following kind of identity:*

$$r = \operatorname{Rk}_\mathcal{I} \mathbf{k}_\mathcal{R} - \operatorname{Rk}_\mathcal{I} \mathbf{k}_\mathcal{S},$$

where $\mathcal{R}$ and $\mathcal{S}$ are multi-sets of elements taken from $\mathcal{I}$ such that $\mathbf{k}_\mathcal{R}$ and $\mathbf{k}_\mathcal{S}$ lie in $\operatorname{rep}_\mathcal{I} P$, and where by definition $\mathbf{k}_\mathcal{R} = \bigoplus_{R \in \mathcal{R}} \mathbf{k}_R$ and $\mathbf{k}_\mathcal{S} = \bigoplus_{S \in \mathcal{S}} \mathbf{k}_S$ (note that elements $R \in \mathcal{R}$ and $S \in \mathcal{S}$ are considered with multiplicity). By extension, we call the pair $(\mathcal{R}, \mathcal{S})$ itself a rank decomposition of r over $\mathcal{I}$. It is minimal *if $\mathcal{R}$ and $\mathcal{S}$ are disjoint as multi-sets.*

Note that $\operatorname{Rk}_\mathcal{I} \mathbf{k}_R(I) = \mathbb{1}_{R \supseteq I}$ for any $I \in \mathcal{I}$ and $R \in \mathcal{R}$ (Proposition 2.1), so $\operatorname{Rk}_\mathcal{I} \mathbf{k}_\mathcal{R}(I)$ counts the number of elements in $\mathcal{R}$ that contain I. This number is requested to be finite in the definition ($\mathbf{k}_\mathcal{R} \in \operatorname{rep}_\mathcal{I} P$): a sufficient condition for this is that $\mathcal{R}$ is *pointwise finite*, i.e. that every index in P belongs to only finitely many elements of $\mathcal{R}$, for then $\mathbf{k}_\mathcal{R} \in \operatorname{rep} P$.

An important consequence of having $\operatorname{Rk}_\mathcal{I} \mathbf{k}_R(I) = \mathbb{1}_{R \supseteq I}$ is that adding the same interval $I \in \mathcal{I}$ to both $\mathcal{R}$ and $\mathcal{S}$ does not change the difference $\operatorname{Rk}_\mathcal{I} \mathbf{k}_\mathcal{R} - \operatorname{Rk}_\mathcal{I} \mathbf{k}_\mathcal{S}$, so rank decompositions cannot be unique. This motivates the notion of minimal rank decomposition.

1.3   Contributions and structure of the paper

In Section 2 we study the existence and uniqueness of minimal rank decompositions. We show in Theorem 2.9 and Corollary 2.10 that a minimal rank decomposition $(\mathcal{R}, \mathcal{S})$ of a given map $r : \mathcal{I} \to \mathbb{Z}$ exists as soon as at least one rank decomposition of r exists, that it is always unique, and that it satisfies a universality property justifying its name. To complete the picture, in Corollary 2.5 we provide mild sufficient conditions for the existence of rank decompositions in the first place. Our proofs emphasize the role played by the family of generalized rank invariants $(\operatorname{Rk}_\mathcal{I} \mathbf{k}_I)_{I \in \mathcal{I}}$, which acts as a generalized basis (Theorem 2.4).

In Section 3 we reformulate our results in the specific context of multi-parameter persistence. We thus obtain existence and uniqueness results for minimal rank decompositions of finitely presented persistence modules over $\mathbb{R}^d$ (Theorem 3.3), and of pfd persistence modules

over finite grids (Corollary 3.2). In the latter case, we derive an explicit inclusion-exclusion formula to compute the coefficients in the minimal rank decompositions, which generalizes the known formula for counting multiplicities in persistence diagrams in the one-parameter case. We also discuss the stability of the minimal rank decompositions, and propose a metric in which to compare them, based on the matching (pseudo-)distance from [9]. In this metric we show that the minimal rank decompositions are the ones maximizing the distance (Proposition 3.8), and that replacing the modules by their rank decompositions does not expand their pairwise distances (Theorem 3.7).

In Section 4 we introduce the signed barcode as a visual representation of the minimal rank decomposition of the usual rank invariant. We explain how the signed barcode reflects the global structure of the usual rank invariant, and how its role in multi-parameter persistence is similar to the one played by the unsigned barcode in one-parameter persistence. We also discuss its extension to generalized rank invariants, for which it takes the form of a "decorated" signed barcode with similar properties and extra information. The use of these barcodes is illustrated on a practical example coming from 2-parameter clustering.

1.4   Related work

Rank decompositions have strong ties with the concept of *generalized persistence diagram*, introduced by Patel [14] and further studied in [2, 8, 11]. This diagram is defined from the rank invariant via a Möbius inversion, from which our inclusion-exclusion formula for computing the coefficients in the minimal rank decomposition derives. Indeed, in the full version of this paper [5] we show that, whenever it is defined, the generalized persistence diagram does correspond to the minimal rank decomposition. However, our framework allows us to prove the existence and uniqueness of the minimal rank decomposition using direct arguments that: (1.) emphasize the role played by the family of rank invariants of interval modules as a generalized basis for the space of maps $\mathcal{I} \to \mathbb{Z}$, and (2.) hold in more general settings where the Möbius inversion is not defined. It also allows us to derive stability results for rank decompositions in general (not just minimal ones), in terms of the matching distance $\mathrm{d}_{\mathrm{match}}$ [9], and to introduce the signed barcodes as a practical graphical representation of minimal rank decompositions – hence of generalized persistence diagrams as well.

2   Rank Decompositions: Existence and Uniqueness

Let P be an arbitrary poset. The following result will be instrumental throughout our analysis. It generalizes [8, Proposition 3.17] by dropping the assumption of local finiteness of the poset P and allowing for generalized ranks, moreover it is given a more direct proof – see the full version of this paper [5]. Note that the result is immediate when working with segments.

▶ **Proposition 2.1.** *Let $\mathcal{R}$ be a multi-set of intervals of P. Then, for any interval $I \subseteq P$:*

$$\mathrm{Rk}_I(\mathbf{k}_{\mathcal{R}}) = \#\{R \in \mathcal{R} \mid I \subseteq R\}.$$

▶ **Corollary 2.2.** *Let $\mathcal{I}$ be a collection of intervals in P. For a multi-set $\mathcal{R}$ of intervals, we have that $\mathbf{k}_{\mathcal{R}} \in \mathrm{rep}_{\mathcal{I}} P$ if and only if $\#\{R \in \mathcal{R} \mid I \subseteq R\} < \infty$ for all $I \in \mathcal{I}$.*

2.1   The locally finite case

Let $\mathcal{I}$ be a locally finite collection of intervals in P. That is, for any two comparable intervals in $\mathcal{I}$, there are only finitely many intervals in $\mathcal{I}$ between the two. We say a map $\mathcal{I} \to \mathbb{Z}$ has locally finite support if its restriction to the up-set of any element of $\mathcal{I}$ has finite support.

▶ **Remark 2.3.** For any fixed $I \in \mathcal{I}$, the map $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I : J \mapsto \mathrm{Rk}_J \, \mathbf{k}_I$ has locally finite support, by the description in Proposition 2.1. More generally, for any multi-set $\mathcal{R}$ of elements in $\mathcal{I}$, if $\mathbf{k}_{\mathcal{R}} \in \mathrm{rep}_{\mathcal{I}} P$ then the map $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}}$ has locally finite support: for any fixed $I \in \mathcal{I}$, by Corollary 2.2 $\mathcal{R}$ only contains finitely many elements containing I, and these are the only ones relevant when considering the restriction of $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}}$ to the up-set of I.

▶ **Theorem 2.4.** *Let $\mathcal{I}$ be a locally finite collection of intervals in P. Then any function $r : \mathcal{I} \to \mathbb{Z}$ with locally finite support can uniquely be written as a (possibly infinite, but pointwise finite) $\mathbb{Z}$-linear combination of the functions $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I$ with $I \in \mathcal{I}$.*

Proof. Existence: Given $I \in \mathcal{I}$ we let $S_I = \{ J \supseteq I \mid \exists K \supseteq J \text{ with } r(K) \neq 0 \}$. Since r is locally finite, its support restricted to the up-set of I is finite, and so is S_I since $\mathcal{I}$ is locally finite. Now we can define a collection of scalars $\alpha_I \in \mathbb{Z}$ for $I \in \mathcal{I}$, inductively on the size of S_I. If $S_I = \varnothing$ we set $\alpha_I = 0$. Otherwise we set

$$\alpha_I = r(I) - \sum_{J \in S_I \setminus \{I\}} \alpha_J .$$

Note that for $J \in S_I \setminus \{I\}$ we have $S_J \subsetneq S_I$, so the terms in the sum are already defined.

Now, using the description of the map $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I$ in Proposition 2.1, one immediately verifies that $r = \sum_{I \in \mathcal{I}} \alpha_I \, \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I$. Note in particular that this infinite sum is pointwise finite – on a given interval J the only possibly non-zero terms are the ones in S_J – hence well-defined.

Uniqueness: subtracting two different $\mathbb{Z}$-linear combinations realizing r from each other, we get a single linear combination $\sum_{I \in \mathcal{I}} \alpha_I \, \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I$ with non-zero coefficients which sums up to zero. Note that there is at least one maximal $I \in \mathcal{I}$ such that $\alpha_I \neq 0$, for otherwise the sum would not be defined. It follows, again using Proposition 2.1, that $(\sum_{J \in \mathcal{I}} \alpha_J \, \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_J)(I) = \alpha_I \neq 0$, contradicting our assumption. ◀

▶ **Corollary 2.5.** *Let $\mathcal{I}$ be a locally finite collection of intervals in P. Then, for any map $r : \mathcal{I} \to \mathbb{Z}$ with locally finite support, there is a unique pair $\mathcal{R}, \mathcal{S}$ of disjoint multi-sets of elements of $\mathcal{I}$ such that $r = \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}} - \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{S}}$ and $\mathbf{k}_{\mathcal{R}}, \mathbf{k}_{\mathcal{S}}$ both lie in $\mathrm{rep}_{\mathcal{I}} P$.*

Proof. By Theorem 2.4, there is a unique (possibly infinite, but pointwise finite) $\mathbb{Z}$-linear combination of functions $r = \sum_{I \in \mathcal{I}} \alpha_I \, \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_I$. Let then $\mathcal{R} = \{ I \in \mathcal{I} \mid \alpha_I > 0 \}$ with multiplicities $I \mapsto \alpha_I$, and $\mathcal{S} = \{ I \in \mathcal{I} \mid \alpha_I < 0 \}$ with multiplicities $I \mapsto |\alpha_I|$. It follows from the pointwise-finiteness of the linear combination that $\mathcal{R}$ and $\mathcal{S}$ satisfy the condition in Corollary 2.2, so in particular $\mathbf{k}_{\mathcal{R}}$ and $\mathbf{k}_{\mathcal{S}}$ lie in $\mathrm{rep}_{\mathcal{I}} P$. ◀

Specializing Theorem 2.4 and Corollary 2.5 to the case where P is finite and $\mathcal{I} = \{ \langle i, j \rangle \mid i \leq j \in P \} \simeq \{ \leq_P \}$ yields the following results – where $\mathrm{Rk}_{\mathcal{I}}$ becomes the usual rank invariant Rk according to Remark 1.2:

▶ **Corollary 2.6.** *If P is finite, then the maps $\mathrm{Rk} \, \mathbf{k}_{\langle a, b \rangle}$ for all $a \leq b \in P$ is a basis of $\mathbb{Z}^{\{\leq_P\}}$.*

▶ **Corollary 2.7.** *Given a finite poset P, for any map $r : \{ \leq_P \} \to \mathbb{Z}$ there is a unique pair $\mathcal{R}, \mathcal{S}$ of disjoint finite multi-sets of closed segments such that $r = \mathrm{Rk} \, \mathbf{k}_{\mathcal{R}} - \mathrm{Rk} \, \mathbf{k}_{\mathcal{S}}$.*

2.2   The general case

We now drop our previous finiteness assumptions and consider arbitrary maps $r : \mathcal{I} \to \mathbb{Z}$ over an arbitrary collection $\mathcal{I}$ of intervals in an arbitrary poset P. Our first result shows that $\mathrm{Rk}_{\mathcal{I}}$ is a complete invariant when restricted to interval-decomposable representations supported on intervals in $\mathcal{I}$. In fact, we show that the rank invariant is complete on a slightly larger collection of intervals. This generalizes [8, Theorem 3.14].

▶ **Proposition 2.8.** *Given a collection $\mathcal{I}$ of intervals in P, let $\widehat{\mathcal{I}} \supseteq \mathcal{I}$ be the collection of limit intervals (which by construction are also intervals):*

$$\widehat{\mathcal{I}} := \left\{ \bigcup_{x \in X} I_x \mid X \text{ totally ordered}, I_x \in \mathcal{I} \text{ and } I_x \subseteq I_y \ \forall x \leq y \in X \right\}.$$

If $\mathcal{R}$ and $\mathcal{R}'$ are two multi-sets of elements in $\widehat{\mathcal{I}}$, such that $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}} = \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}'}$ and this common rank invariant is finite, then $\mathcal{R} = \mathcal{R}'$.

Proof. Since the rank of a direct sum is the sum of the ranks we may remove the common elements from $\mathcal{R}$ and $\mathcal{R}'$, and thus assume that the two multi-sets are disjoint. It follows from the description of multi-sets giving rise to finite ranks in Corollary 2.2 that $\mathcal{R} \cup \mathcal{R}'$ contains at least one maximal element, say J. Without loss of generality we assume $J \in \mathcal{R}$. By definition of $\widehat{\mathcal{I}}$ we have $J = \bigcup_{x \in X} J_x$ with $J_x \in \mathcal{I}$ and $J_x \subseteq J_y$ for all $x \leq y \in X$. Now, by assumption, for every $x \in X$ we have

$$\mathrm{Rk}_{J_x} \, \mathbf{k}_{\mathcal{R}'} = \mathrm{Rk}_{J_x} \, \mathbf{k}_{\mathcal{R}} \geq \mathrm{Rk}_{J_x} \, \mathbf{k}_J,$$

which is at least 1 by Proposition 2.1. It also follows from Proposition 2.1 that, for each $x \in X$, there is some interval $I_x \in \mathcal{R}'$ such that $J_x \subseteq I_x$. Since $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}'}$ is finite, Corollary 2.2 says that there are actually only finitely many choices for I_x. Hence, there is an $I \in \mathcal{R}'$ independent of x such that $J_x \subseteq I$ for all $x \in X$. Thus $J \subseteq I$. If this is a proper inclusion then it contradicts the maximality of J, otherwise it contradicts the disjointness of $\mathcal{R}$ and $\mathcal{R}'$. ◄

We can now show minimal rank decompositions satisfy a universality property when they exist.

▶ **Theorem 2.9.** *Let $\mathcal{R}, \mathcal{S}, \mathcal{R}^*, \mathcal{S}^*$ be multi-sets of elements of $\widehat{\mathcal{I}}$, whose corresponding representations lie in $\mathrm{rep}_{\mathcal{I}} P$, and such that $\mathcal{R}^* \cap \mathcal{S}^* = \emptyset$. If*

$$\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}} - \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{S}} = \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}^*} - \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{S}^*}$$

then $\mathcal{R} \supseteq \mathcal{R}^*$, $\mathcal{S} \supseteq \mathcal{S}^*$, and $\mathcal{R} \setminus \mathcal{R}^* = \mathcal{S} \setminus \mathcal{S}^*$.

Proof. Rewriting the equation yields $\mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}} + \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{S}^*} = \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{R}^*} + \mathrm{Rk}_{\mathcal{I}} \, \mathbf{k}_{\mathcal{S}}$, and by additivity of the rank invariant, $\mathrm{Rk}_{\mathcal{I}}(\mathbf{k}_{\mathcal{R}} \oplus \mathbf{k}_{\mathcal{S}^*}) = \mathrm{Rk}_{\mathcal{I}}(\mathbf{k}_{\mathcal{R}^*} \oplus \mathbf{k}_{\mathcal{S}})$. It follows then by Proposition 2.8 that $\mathcal{R} \cup \mathcal{S}^* = \mathcal{R}^* \cup \mathcal{S}$. As $\mathcal{R}^* \cap \mathcal{S}^* = \emptyset$, we conclude that $\mathcal{R} \supseteq \mathcal{R}^*$, $\mathcal{S} \supseteq \mathcal{S}^*$, and $\mathcal{R} \setminus \mathcal{R}^* = \mathcal{S} \setminus \mathcal{S}^*$. ◄

As an immediate consequence of Theorem 2.9, we obtain uniqueness and conditional existence of minimal rank decompositions:

▶ **Corollary 2.10.** *The minimal rank decomposition $(\mathcal{R}^*, \mathcal{S}^*)$ of any map $r : \mathcal{I} \to \mathbb{Z}$ is unique if it exists. Furthermore, it exists as soon as any rank decomposition $(\mathcal{R}, \mathcal{S})$ of r does, being obtained from it by removing common intervals, that is: $(\mathcal{R}^*, \mathcal{S}^*) = (\mathcal{R} \setminus \mathcal{R} \cap \mathcal{S}, \ \mathcal{S} \setminus \mathcal{R} \cap \mathcal{S})$.*

We also get a connection between the various rank decompositions of a map $\mathcal{I} \to \mathbb{Z}$:

▶ **Corollary 2.11.** *$\mathcal{R} \cup \mathcal{S}' = \mathcal{R}' \cup \mathcal{S}$ for any rank decompositions $(\mathcal{R}, \mathcal{S})$, $(\mathcal{R}', \mathcal{S}')$ of $r : \mathcal{I} \to \mathbb{Z}$.*

Proof. Let $(\mathcal{R}^*, \mathcal{S}^*)$ be the minimal rank decomposition of r. By Theorem 2.9, we have $\mathcal{R} = \mathcal{R}^* \cup \mathcal{T}$ and $\mathcal{S} = \mathcal{S}^* \cup \mathcal{T}$ for some finite multi-set $\mathcal{T}$ of elements of $\mathcal{I}$, while $\mathcal{R}' = \mathcal{R}^* \cup \mathcal{T}'$ and $\mathcal{S}' = \mathcal{S}^* \cup \mathcal{T}'$ for some multi-set $\mathcal{T}'$. Then, $\mathcal{R} \cup \mathcal{S}' = \mathcal{R}^* \cup \mathcal{S}^* \cup \mathcal{T} \cup \mathcal{T}' = \mathcal{R}' \cup \mathcal{S}$. ◄

$$\mathrm{Rk}_{\mathcal{I}}\left(\begin{array}{c} \mathbf{k} \xrightarrow{\mathrm{id}} \mathbf{k} \to 0 \\ \mathrm{id}\uparrow \quad \uparrow [1\;0] \\ \mathbf{k} \xrightarrow[{[\begin{smallmatrix}1\\0\end{smallmatrix}]}]{} \mathbf{k}^2 \xrightarrow{[1\;1]} \mathbf{k} \\ \uparrow [\begin{smallmatrix}0\\1\end{smallmatrix}] \quad \uparrow \mathrm{id} \\ 0 \to \mathbf{k} \xrightarrow{\mathrm{id}} \mathbf{k} \end{array}\right) = \mathrm{Rk}_{\mathcal{I}}\left(\text{[blue grid]} \oplus \text{[blue grid]} \oplus \text{[blue grid]} \oplus \text{[blue grid]} \right)$$

$$- \; \mathrm{Rk}_{\mathcal{I}}\left(\text{[red grid]} \oplus \text{[red grid]} \right)$$

■ **Figure 4** Minimal rank decomposition of the generalized rank invariant of the module M from Figure 2 over the full collection $\mathcal{I}$ of intervals in the 3×3 grid. Blue is for intervals in $\mathcal{R}$, red is for intervals in $\mathcal{S}$.

$$\mathrm{Rk}_{\mathcal{I}}\left(\mathbf{k} \right) = \mathrm{Rk}_{\mathcal{I}}\left(\mathbf{k} \oplus \mathbf{k} \right) - \mathrm{Rk}_{\mathcal{I}}\left(\mathbf{k} \right)$$

■ **Figure 5** Taking $\mathcal{I}$ to be the collection of all intervals with one generator and at most two cogenerators (which includes in particular all rectangles), the generalized rank invariant of the interval 2-parameter persistence module $\mathbf{k}_I$ on the left-hand side decomposes minimally as the difference between the generalized rank invariants of the two modules on the right-hand side. Blue is for intervals in $\mathcal{R}$, red is for intervals in $\mathcal{S}$.

3 Application to multi-parameter persistence

Here the poset P under consideration is either $\mathbb{R}^d$, viewed as a product of d copies of the totally ordered real line, or a subposet of $\mathbb{R}^d$ – usually $\mathbb{Z}^d$ or some finite grid $\prod_{i=1}^{d} [\![1, n_i]\!]$. The role of segments is played by *rectangles*, i.e. products of 1-d intervals.

3.1 The finite grid case

In this case, Corollary 2.5 reformulates as follows:

▶ **Corollary 3.1.** *Given an arbitrary collection $\mathcal{I}$ of intervals in a finite grid $G = \prod_{i=1}^{d} [\![1, n_i]\!] \subset \mathbb{R}^d$, the generalized rank invariant $\mathrm{Rk}_{\mathcal{I}} M$ of any pfd persistence module M indexed over G admits a unique minimal rank decomposition $(\mathcal{R}, \mathcal{S})$ over $\mathcal{I}$.*

Taking $\mathcal{I}$ to be the collection of all closed rectangles in the grid G yields the following reformulation of Corollary 2.7:

▶ **Corollary 3.2.** *The usual rank invariant of any pfd persistence module M indexed over a finite grid $G = \prod_{i=1}^{d} [\![1, n_i]\!] \subset \mathbb{R}^d$ admits a unique minimal rank decomposition $(\mathcal{R}, \mathcal{S})$, where $\mathcal{R}$ and $\mathcal{S}$ are finite multi-sets of (closed) rectangles in G.*

Figures 4 and 5 illustrate Corollary 3.1, while Figures 2 and 6 illustrate Corollary 3.2.

Figure 6 The usual rank invariant of the interval module $\mathbf{k}_I$ on the left-hand side decomposes minimally as the difference between the usual rank invariants of the two rectangle-decomposable modules on the right-hand side. Blue is for rectangles in $\mathcal{R}$, red is for rectangles in $\mathcal{S}$.

Computation

Given $\mathrm{Rk}_{\mathcal{I}} M$, computing its minimal rank decomposition can be done by applying the inclusion-exclusion formula defining the so-called *generalized persistence diagram* of M – see e.g. Definition 3.13 in [8]. Indeed, whenever it exists, the generalized persistence diagram of M coincides with the minimal rank decomposition of $\mathrm{Rk}_{\mathcal{I}}(M)$. This connection happens as both objects derive from the Möbius inversion of the generalized rank invariant – see Section 3 in the full version of our paper [5] for the details. While simple, this approach does not scale up well because the inclusion-exclusion formula must be applied for every interval in the collection $\mathcal{I}$, whose size can be up to exponential in the size of the indexing grid G, even in two dimensions [1].

In the special case of the usual rank invariant however, the approach scales up reasonably as the number of rectangles is at most quadratic in the size of the grid G. To be more specific, the inclusion-exclusion formula writes as follows in this case, where $\alpha_{\langle s,t \rangle}$ is the coefficient assigned to $\mathrm{Rk}\,\mathbf{k}_{\langle s,t \rangle}$ in the minimal rank decomposition of $\mathrm{Rk}\,M$:

$$\forall s \le t \in G, \ \alpha_{\langle s,t \rangle} = \sum_{\substack{s' \le s \\ \|s'-s\|_\infty \le 1}} \sum_{\substack{t' \ge t \\ \|t'-t\|_\infty \le 1}} (-1)^{\|s'-s\|_1 + \|t'-t\|_1} \,\mathrm{Rk}\,M(s',t'). \tag{3.1}$$

By applying (3.1) successively to every pair of comparable indices $s \le t$ in the grid $G = \prod_{i=1}^{d} [\![1, n_i]\!] \subset \mathbb{R}^d$, one computes the minimal rank decomposition of $\mathrm{Rk}\,M$ in time $O\left(2^{2d} \#\{\le_G\}\right)$, assuming constant-time access to the ranks $\mathrm{Rk}\,M(s',t')$ and constant-time arithmetic operations[1]. This bound is in $O\left(2^{2d} \prod_{i=1}^{d} n_i^2\right)$, and when d is fixed, it is linear in the size of the encoding of the usual rank invariant as a map $\{\le_G\} \to \mathbb{Z}$. When the module M comes from a simplicial filtration over the grid G with $n = \max_i n_i$ simplices in total, the usual rank invariant itself can be pre-computed and stored, e.g. by naively computing the ranks $\mathrm{Rk}\,M(s,t)$ for each pair $s \le t \in G$ independently, which takes $O(n^{2d+\omega})$ time in total, where $2 \le \omega < 2.373$ is the exponent for matrix multiplication [12]. Adding in the

---

[1] We are considering an implementation that iterates over the indices s', t' such that $\|s'-s\|_\infty \le 1$ and $\|t'-t\|_\infty \le 1$ by increasing order of the 1-norms $\|s'-s\|_1$ and $\|t'-t\|_1$, so that the 1-norms do not have to be re-computed from scratch at each step. Such an implementation boils down to iterating over the vertices of the unit hypercube in $\mathbb{R}^d$ by increasing order of the number of 1's in their coordinates.

■ **Figure 7** Left: an indecomposable persistence module with 2 generators, at a and b, and a relation equating them at h (indices d, e, f, g, i lie at infinity). Right: minimal rank decomposition of the usual rank invariant of the module, over the right-open rectangles. Blue is for rectangles in $\mathcal{R}$, red is for rectangles in $\mathcal{S}$. Solid boundaries belong to the rectangles while dotted boundaries do not.

computation time for the minimal rank decomposition yields a bound in $O(n^{2d+\omega} + (2n)^{2d})$. While naive, this approach already compares favorably, in terms of running time, to the computation of other (stronger) invariants such as for instance the direct-sum decomposition of M – for which the best known algorithm runs in $O(n^{d(2\omega+1)})$ time [7]. Moreover, the running time of computing the minimal rank decomposition is dominated by the running time of pre-computing the usual rank invariant, for which there is room for improvement. In the special case where $d = 2$ for instance, assuming the filtration is 1-critical (i.e. each simplex has a unique minimal time of appearance in the filtration), there is an $O(n^4)$-time algorithm to compute the usual rank invariant [4, 13], and computing its minimal rank decomposition also takes $O(n^4)$ time. By comparison, the best known algorithm to compute the direct-sum decomposition of M in this setting takes $O(n^{2\omega+1})$ time [7], and computing the line arrangement data structure in RIVET takes $O(n^5)$ time in the worst case [10].

3.2   The $\mathbb{R}^d$ case

The rectangles in this context are *right-open rectangles*, i.e. products $\prod_{i=1}^d [a_i, b_i)$ of right-open intervals of the real line ($a_i < b_i \in \mathbb{R} \cup \{\infty\}$ for each $i = 1, \cdots, d$).

▶ **Theorem 3.3.** *The usual rank invariant of every finitely presented persistence module M over $\mathbb{R}^d$ admits a unique minimal rank decomposition over the right-open rectangles in $\mathbb{R}^d$.*

This result, illustrated in Figure 7, follows from Corollary 2.10 and from the fact that rank decompositions of finitely presented persistence modules over $\mathbb{R}^d$ exist in the first place – this fact itself is a consequence of the existence of so-called *rank-exact resolutions* of such modules, which are the subject of Section 4 in the full version of the paper [5].

3.3   Restrictions to lines

A line ℓ in $\mathbb{R}^d$ is called *monotone* if it can be parametrized by $\lambda \mapsto (1 - \lambda)s + \lambda t$ where $s \leq t \in \mathbb{R}^d$ are fixed. If $s_i < t_i$ in every dimension $i = 1, \cdots, d$, then ℓ is called *strictly monotone*. The restriction of M to ℓ is a one-parameter persistence module and thus has a well-defined barcode called the *fibered barcode*. We shall now see that the fibered barcode of $M|_\ell$ can be obtained by a rank decomposition of M. In the following we employ the notation $\mathcal{R}|_\ell = \{R \cap \ell \colon R \in \mathcal{R}\}$. Note that the elements of $\mathcal{R}|_\ell$ are intervals in ℓ.
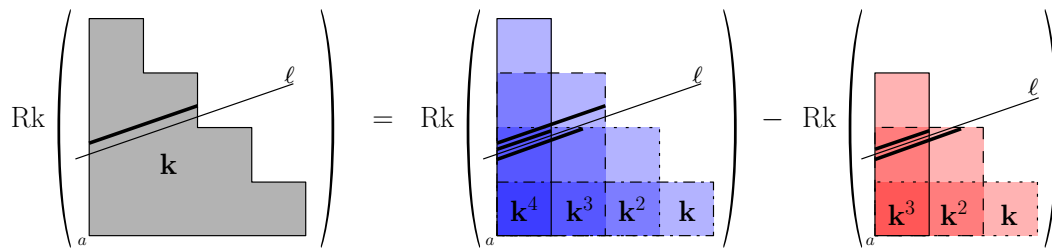
Figure 8 Restricting an interval module $\mathbf{k}_I$ to a monotone line ℓ (left) yields a restriction of the minimal rank decomposition of $\mathrm{Rk}\,\mathbf{k}_I$ to ℓ (right) – for clarity, the rectangles' boundaries are shown with different line styles. Here, the restricted rank decomposition is not minimal, as the two interval summands of $\mathbf{k}_{\mathcal{S}|\ell}$ cancel out with two of the three interval summands of $\mathbf{k}_{\mathcal{R}|\ell}$.

▶ **Proposition 3.4.** *Let M be a pfd persistence module over $\mathbb{R}^d$ such that the usual rank invariant $\mathrm{Rk}\,M$ admits a rank decomposition $(\mathcal{R}, \mathcal{S})$. Then, for any monotone line ℓ in $\mathbb{R}^d$, $(\mathcal{R}, \mathcal{S})$ restricts to a rank decomposition $(\mathcal{R}_{|\ell}, \mathcal{S}_{|\ell})$ of $\mathrm{Rk}\,M_{|\ell}$.*

Proof. Observe that $\mathrm{Rk}\,\mathbf{k}_{\mathcal{R}}|_\ell = \mathrm{Rk}\,\mathbf{k}_{\mathcal{R}\cap\ell}$ (and likewise for $\mathcal{S}$). Thus, $\mathrm{Rk}\,M|_\ell = \mathrm{Rk}\,\mathbf{k}_{\mathcal{R}}|_\ell - \mathrm{Rk}\,\mathbf{k}_{\mathcal{S}}|_\ell = \mathrm{Rk}\,\mathbf{k}_{\mathcal{R}\cap\ell} - \mathrm{Rk}\,\mathbf{k}_{\mathcal{S}\cap\ell}$. ◀

▶ Remark 3.5. For a general discussion on restrictions of rank decompositions to subposets, see Section 5 in the full version of the paper [5].

Note that the restriction of a minimal decomposition may not be minimal, as different rectangles in $\mathcal{R}$ and $\mathcal{S}$ may restrict to the same 1-d interval – see Figure 8 for an illustration. However, by Corollary 2.10, the minimal rank decomposition $(\mathcal{R}^*, \mathcal{S}^*)$ of $\mathrm{Rk}\,M_{|\ell}$ is easily obtained by removing all the common elements in $\mathcal{R}_{|\ell}$ and $\mathcal{S}_{|\ell}$. Furthermore, as illustrated in Figure 8 and formalized in the following result, $(\mathcal{R}^*, \mathcal{S}^*)$ actually coincides with the persistence barcode of the one-parameter module $M|_\ell$.

▶ **Corollary 3.6.** *Every pfd persistence module M over $\mathbb{R}$ admits a unique minimal rank decomposition $(\mathcal{R}, \mathcal{S})$, given by $\mathcal{R} = \mathrm{Dgm}\,M$, the persistence barcode of M, and $\mathcal{S} = \emptyset$.*

Proof. Follows from (1.3) and Corollary 2.10. ◀

3.4   Stability

We conclude this section by saying a few words about the stability of our rank decompositions. Recall from Corollary 2.11 that we have $\mathbf{k}_{\mathcal{R}} \oplus \mathbf{k}_{\mathcal{S}'} \simeq \mathbf{k}_{\mathcal{R}'} \oplus \mathbf{k}_{\mathcal{S}}$ for any two rank decompositions $(\mathcal{R}, \mathcal{S})$ and $(\mathcal{R}', \mathcal{S}')$ of the same persistence module M, or of two persistence modules M, M' sharing the same (usual) rank invariant. In effect, this is telling us that two rank decompositions are equivalent whenever their ground modules have the same rank invariant. Using the matching (pseudo-)distance $\mathrm{d}_{\mathrm{match}}$ from [9], we can derive a metric version of this statement (Theorem 3.7), which bounds the defect of equivalence between two rank decompositions in terms of the fibered distance between the rank invariants of their ground modules. Recall that the matching distance between two pfd persistence modules M, N in $\mathbb{R}^d$ is defined as follows:

$$\mathrm{d}_{\mathrm{match}}(M, N) = \sup_{\ell \text{ strictly monotone}} \omega(\ell)\, \mathrm{d}_{\mathrm{b}}(M|_\ell, N|_\ell), \qquad (3.2)$$

where d_{b} denotes the usual bottleneck distance between one-parameter persistence modules, and where the weight of ℓ (parametrized as in Section 3.3) is

$$\omega(\ell) = (\min_i t_i - s_i)/(\max_i t_i - s_i) > 0.$$

▶ **Theorem 3.7.** *Let M, M' be pfd persistence modules indexed over $\mathbb{R}^d$. Then, for any rank decompositions $(\mathcal{R}, \mathcal{S})$ and $(\mathcal{R}', \mathcal{S}')$ of M and M' respectively, we have:*

$$\mathrm{d}_{\mathrm{match}}(\mathbf{k}_{\mathcal{R}} \oplus \mathbf{k}_{\mathcal{S}'}, \mathbf{k}_{\mathcal{R}'} \oplus \mathbf{k}_{\mathcal{S}}) \leq \mathrm{d}_{\mathrm{match}}(M, M').$$

Proof. Take any strictly monotone line ℓ in $\mathbb{R}^d$. By (3.2), we have:

$$\mathrm{d}_{\mathrm{b}}(M_{|\ell}, M'_{|\ell}) \leq \omega(\ell)^{-1}\, \mathrm{d}_{\mathrm{match}}(M, M').$$

Meanwhile, by Corollary 3.4, $(\mathcal{R}_{|\ell}, \mathcal{S}_{|\ell})$ is a rank decomposition of $M_{|\ell}$, and $(\mathcal{R}'_{|\ell}, \mathcal{S}'_{|\ell})$ is a rank decomposition of $M'_{|\ell}$. By Proposition 2.8, we then have $M_{|\ell} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}} \simeq \mathbf{k}_{\mathcal{R}_{|\ell}}$ and $M'_{|\ell} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}} \simeq \mathbf{k}_{\mathcal{R}'_{|\ell}}$, from which we deduce:

$$\mathrm{d}_{\mathrm{b}}(M_{|\ell}, M'_{|\ell}) \geq \mathrm{d}_{\mathrm{b}}(M_{|\ell} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}}, M'_{|\ell} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}}) = \mathrm{d}_{\mathrm{b}}(\mathbf{k}_{\mathcal{R}_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}}, \mathbf{k}_{\mathcal{R}'_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}}).$$

Combined with the previous equation, this gives:

$$\mathrm{d}_{\mathrm{b}}(\mathbf{k}_{\mathcal{R}_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}}, \mathbf{k}_{\mathcal{R}'_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}}) \leq \omega(\ell)^{-1}\, \mathrm{d}_{\mathrm{match}}(M, M').$$

The result follows then by taking the supremum on the left-hand side over all possible choices of strictly monotone lines ℓ.　　　　◀

Note that different choices of rank decompositions $(\mathcal{R}, \mathcal{S})$ and $(\mathcal{R}', \mathcal{S}')$ for M and M' may yield different values for the matching distance $\mathrm{d}_{\mathrm{match}}(\mathbf{k}_{\mathcal{R}} \oplus \mathbf{k}_{\mathcal{S}'}, \mathbf{k}_{\mathcal{R}'} \oplus \mathbf{k}_{\mathcal{S}})$. It turns out that the rank decompositions maximizing this distance are precisely the minimal rank decompositions, which therefore also satisfy a universal property in terms of the metric:

▶ **Proposition 3.8.** *Let M, M' be pfd persistence modules indexed over $\mathbb{R}^d$. Then, for any rank decompositions $(\mathcal{R}, \mathcal{S})$ and $(\mathcal{R}', \mathcal{S}')$ of M and M' respectively, we have:*

$$\mathrm{d}_{\mathrm{match}}(\mathbf{k}_{\mathcal{R}} \oplus \mathbf{k}_{\mathcal{S}'}, \mathbf{k}_{\mathcal{R}'} \oplus \mathbf{k}_{\mathcal{S}}) \leq \mathrm{d}_{\mathrm{match}}(\mathbf{k}_{\mathcal{R}^*} \oplus \mathbf{k}_{\mathcal{S}'^*}, \mathbf{k}_{\mathcal{R}'^*} \oplus \mathbf{k}_{\mathcal{S}^*}),$$

where $(\mathcal{R}^*, \mathcal{S}^*)$ and $(\mathcal{R}'^*, \mathcal{S}'^*)$ are the minimal rank decompositions of M and M' respectively – which exist as soon as $(\mathcal{R}, \mathcal{S})$ and $(\mathcal{R}', \mathcal{S}')$ do, by Corollary 2.10.

Proof. Let $\mathcal{T} := \mathcal{R} \setminus \mathcal{R}^* = \mathcal{S} \setminus \mathcal{S}^*$, and $\mathcal{T}' := \mathcal{R}' \setminus \mathcal{R}'^* = \mathcal{S}' \setminus \mathcal{S}'^*$. Note that $\mathcal{T}, \mathcal{T}'$ are well-defined by Theorem 2.9. Then, for any strictly monotone line ℓ, we have:

$$\mathrm{d}_{\mathrm{b}}(\mathbf{k}_{\mathcal{R}_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'_{|\ell}}, \mathbf{k}_{\mathcal{R}'_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}_{|\ell}}) = \mathrm{d}_{\mathrm{b}}(\mathbf{k}_{\mathcal{R}^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{T}_{|\ell}} \oplus \mathbf{k}_{\mathcal{T}'_{|\ell}}, \mathbf{k}_{\mathcal{R}'^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{T}_{|\ell}} \oplus \mathbf{k}_{\mathcal{T}'_{|\ell}})$$

$$\leq \mathrm{d}_{\mathrm{b}}(\mathbf{k}_{\mathcal{R}^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}'^*_{|\ell}}, \mathbf{k}_{\mathcal{R}'^*_{|\ell}} \oplus \mathbf{k}_{\mathcal{S}^*_{|\ell}}).$$

The result follows then after multiplying by $\omega(\ell)$ and taking the supremum on both sides over all possible choices of strictly monotone lines ℓ.　　　　◀

4　Signed barcodes and prominence diagrams

In the context of topological data analysis, the minimal rank decomposition $(\mathcal{R}, \mathcal{S})$ of $\mathrm{Rk}\, M$ encodes visually the structure of the rank invariant of $M \colon \mathbb{R}^d \to \mathrm{Vec}_{\mathbf{k}}$. However, representing rectangles as rectangles quickly leads to arrangements that are hard to read – see e.g. Figure 8.
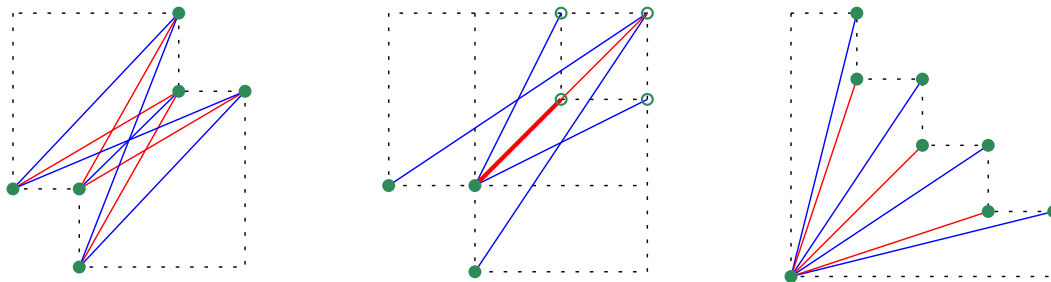
Figure 9 From left to right: signed barcodes corresponding to the usual rank decompositions of Figures 6, 7, and 8 respectively. Blue bars are diagonals of rectangles in $\mathcal{R}$ and therefore counted positively, while red bars are diagonals of rectangles in $\mathcal{S}$ and therefore counted negatively. The bars' endpoints are marked in green (as a solid dot when the endpoint lies in the rectangle, as a circled dot when it does not – e.g. when it lies at infinity), to discriminate them from intersections. The thick red line segment in the center picture shows the overlap between a shorter red bar and a longer red bar sharing the same lower endpoint and slope.


$$\mathrm{Rk}\, M(s,t) = 2 - 1 = 1$$
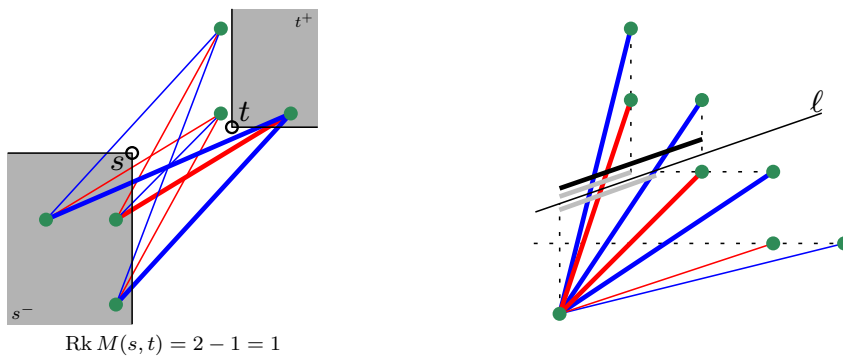
Figure 10 Left: computing $\mathrm{Rk}\, M(s,t)$ for a pair of indices $s \leq t$ – the thick bars are the ones connecting the down-set s^- to the up-set t^+. Right: restricting the minimal rank decomposition of $\mathrm{Rk}\, M$ to a strictly monotone line ℓ – the thick blue and red bars are the ones projecting to non-empty bars along ℓ, and among those projections, the thick gray bars get cancelled out during the simplification while the thick black bar remains in the barcode of $M|_\ell$.
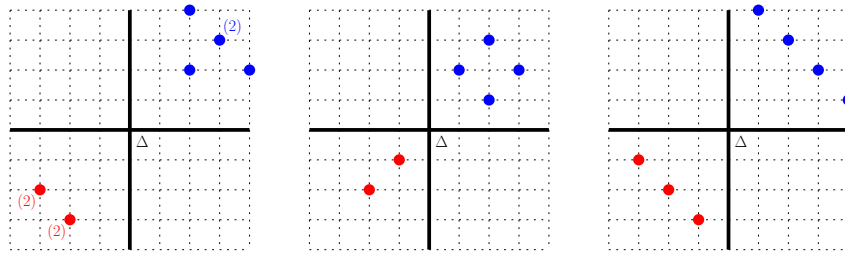
4.1 Signed barcodes

An alternate representation of the rectangles is by their diagonal with positive slope. We call this representation the *signed barcode* of $\mathrm{Rk}\, M$, where each bar is the diagonal (with positive slope) of a particular rectangle in $\mathcal{R}$ or $\mathcal{S}$, and where the sign is positive for bars coming from $\mathcal{R}$ and negative for bars coming from $\mathcal{S}$ – see Figure 9 for an illustration. Like the rectangles, the bars are considered with multiplicity. The signed barcode of $\mathrm{Rk}\, M$ gives direct access to the same pieces of information as the rectangular representation, as shown in Figure 10. Furthermore, the signed barcode makes it possible to visually grasp the global structure of the usual rank invariant $\mathrm{Rk}\, M$, and in particular, to infer the directions along which topological features have the best chances to persist.

When the collection $\mathcal{I}$ of intervals under consideration contains more than just the rectangles, the intervals involved in the corresponding minimal rank decomposition $(\mathcal{R}, \mathcal{S})$ of M are no longer described by a single diagonal. Nevertheless, each interval $I \in \mathcal{R} \sqcup \mathcal{S}$ is still uniquely described by the signed barcode of the corresponding interval module $\mathbf{k}_I$. We
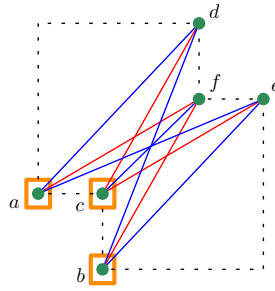
Figure 11 Decorated signed barcode corresponding to the generalized rank decomposition $(\mathcal{R}, \mathcal{S})$ of Figure 5. The orange squares indicate how the bars are grouped together according to their originating element $I \in \mathcal{R} \sqcup \mathcal{S}$.

Figure 12 The signed prominence diagrams corresponding to the signed barcodes of Figure 9, in the same order. Blue dots correspond to blue bars (hence to rectangles in $\mathcal{R}$), while red dots correspond to red bars (hence to rectangles in $\mathcal{S}$). Multiplicities differing from 1 are indicated explicitly. The union Δ of the two coordinate axes plays the role of the diagonal.

can then collate all these signed barcodes together, negating the ones coming from intervals in $\mathcal{S}$, which yields a global *decorated* signed barcode for $\mathrm{Rk}_{\mathcal{I}}\, M$, where the decoration groups the bars according to which element $I \in \mathcal{R} \sqcup \mathcal{S}$ they originate from – see Figure 11.
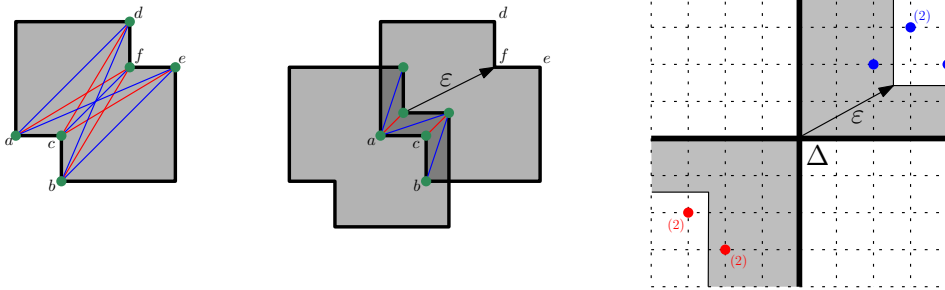
4.2   Signed prominence diagrams

To each bar with endpoints $s \leq t$ in the (undecorated) signed barcode of $\mathrm{Rk}\, M$, we can associate its *signed prominence*, which is the d-dimensional vector $t - s$ if the bar corresponds to a rectangle in $\mathcal{R}$, or $s - t$ if the bar corresponds to a rectangle in $\mathcal{S}$. We call *signed prominence diagram* of M the resulting collection of vectors in $\mathbb{R}^d$ – see Figure 12.

In a signed prominence diagram, the union Δ of the hyperplanes perpendicular to the coordinate axes and passing through the origin plays the role of the diagonal: a bar whose signed prominence lies close to Δ can be viewed as noise, whereas a bar whose signed prominence lies far away from Δ can be considered significant for the structure of $\mathrm{Rk}\, M$. The right way to formalize this intuition is via smoothings, as in the one-parameter case.

▶ **Definition 4.1.** *Given $\varepsilon \in \mathbb{R}^d_{\geq 0}$, the ε-shift $M[\varepsilon]$ is the persistence module defined pointwise by $M[\varepsilon](t) = M(t + \varepsilon)$ and $M[\varepsilon](s \leq t) = M(s + \varepsilon \leq t + \varepsilon)$. There is a canonical morphism of persistence modules $M \to M[\varepsilon]$, whose image M^ε is called the ε-smoothing of M.*

▶ **Example 4.2.** The ε-shift of a rectangle module $\mathbf{k}_R$ is $\mathbf{k}_{R-\varepsilon}$, where by definition $R - \varepsilon$ is the shifted rectangle $\{t - \varepsilon \mid t \in R\}$. The ε-smoothing of $\mathbf{k}_R$ is $\mathbf{k}_{R^\varepsilon}$, where by definition R^ε is the rectangle $R \cap (R - \varepsilon)$, obtained from R by shifting its upper-right corner by $-\varepsilon$.

Figure 13 Behavior of the signed barcode and prominence diagram under ε-smoothing. Left: the input module M from Figure 6, overlaid with its signed barcode. Center: the ε-smoothing M^ε of M (in dark gray), overlaid with its own signed barcode – obtained by shifting the right endpoints in the signed barcode of M by $-\varepsilon$. Right: effect of the ε-smoothing on the signed prominence diagram.

As it turns out, rank decompositions of the usual rank invariant commute with smoothings:

▶ **Lemma 4.3.** *If $(\mathcal{R}, \mathcal{S})$ is a rank decomposition of $\mathrm{Rk}\, M$, then, for any $\varepsilon \in \mathbb{R}^d_{\geq 0}$, the pair $(\mathcal{R}^\varepsilon, \mathcal{S}^\varepsilon)$ where $\mathcal{R}^\varepsilon = \{R^\varepsilon \mid R \in \mathcal{R}\}$ and $\mathcal{S}^\varepsilon = \{S^\varepsilon \mid S \in \mathcal{S}\}$ is a rank decomposition of $\mathrm{Rk}\, M^\varepsilon$. If $(\mathcal{R}, \mathcal{S})$ is minimal, then so is $(\mathcal{R}^\varepsilon, \mathcal{S}^\varepsilon)$ after removing the empty rectangles from $\mathcal{R}^\varepsilon$ and $\mathcal{S}^\varepsilon$.*

See our full version [5] for an elementary proof of this result, which says that the effect of ε-smoothing M on its signed barcode is to shift the right endpoints of the bars by $-\varepsilon$, removing those bars for which the shifted right endpoint is no longer greater than or equal to the left endpoint. The effect on its signed prominence diagram is to shift the positive vectors by $-\varepsilon$ and the negative vectors by ε, removing those vectors that cross Δ. Alternatively, one can inflate Δ by ε, and remove the vectors that lie in the inflated Δ, as illustrated in Figure 13.

4.3 A practical example: two-parameter clustering

We consider the point set P shown in Figure 14, which consists of $N = 90$ planar points sampled from three different Gaussian distributions. We build the *Vietoris–Rips bifiltration* from this dataset, given by $\mathrm{VR}(P)_{r,s} := \mathrm{VR}(f_\varepsilon^{-1}(-\infty, s])_r$, where $\mathrm{VR}(\cdot)_r$ denotes the usual Vietoris-Rips complex of parameter r, and where $f_\varepsilon \colon P \to \mathbb{R}$ is a local co-density estimator:

$$f_\varepsilon(p) = \#\{q \in P : d(p, q) > \epsilon\}, \qquad \text{for a fixed parameter } \epsilon \geq 0.$$

As the Vietoris-Rips complex $\mathrm{VR}(P)_{r,s}$ can be hard to visualize, we replace it in our plots by a proxy union of balls, $U_{r,s} = \left\{ z \in \mathbb{R}^2 : \min\limits_{p \in P, f_\varepsilon(p) \leq s} ||p - z|| \leq r/2 \right\}$, which is known to be interleaved multiplicatively with it.

Applying simplicial 0-homology with coefficients in the field $\mathbb{Z}_2$ yields a bipersistence module $M \colon M(r, s) = H_0(\mathrm{VR}(P)_{r,s})$. In practice we discretize M over a 10×10 regular grid G, which we identify with the grid $\{0, 1, \ldots, 9\} \times \{0, 1, \ldots, 9\}$ in our plots. We know that[2], if $(\mathcal{R}, \mathcal{S})$ is a rank decomposition of M, then $(\mathcal{R}|_G, \mathcal{S}|_G)$ is a rank decomposition of $M|_G$. Note that the persistence module thus obtained is not interval-decomposable. Geometrically, this is due to three clusters A, B, C merging in three different ways at incomparable grades, as shown in the highlighted squares of Figure 14, so that we have the following diagrams:

[2] This comes from an extension of Proposition 3.4 to lattices, proven in the full version of the paper [5].
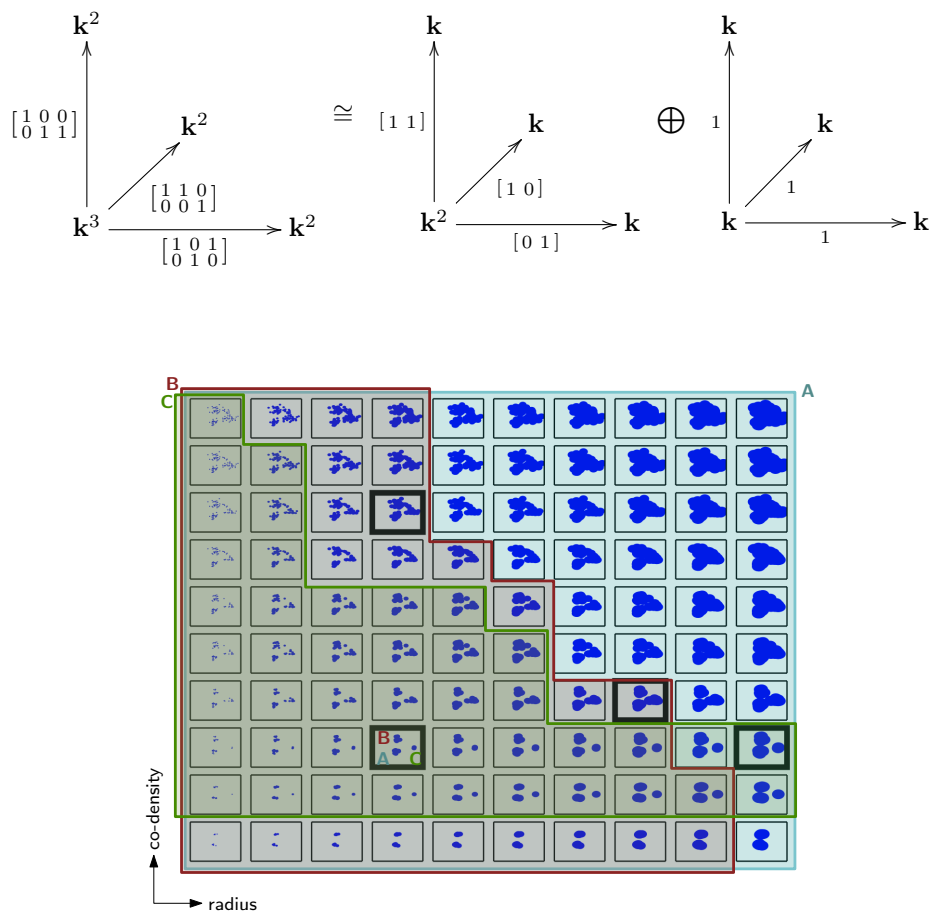
Figure 14 The bifiltration in our experiment. The highlighted black squares show that three clusters (named A, B, C) merge in three different ways at incomparable scales. The lifespan of each one of these three clusters is marked by an interval with matching color.
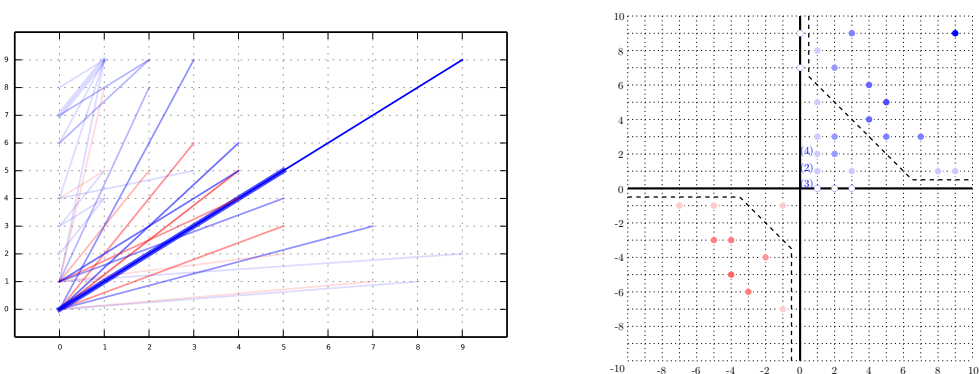
Figure 15 Left: signed barcode of our experiment over the 10×10 grid G, where thicker bars overlap with another bar. Right: corresponding prominence diagram, where the bars coming from the lifespans of A, B, C are separated from the rest of the bars by the dashed curves. Each bar with endpoints $s \leq t$ in the barcode (and diagram) has an intensity proportional to $\min\{t_x - s_x, \, t_y - s_y\}$.
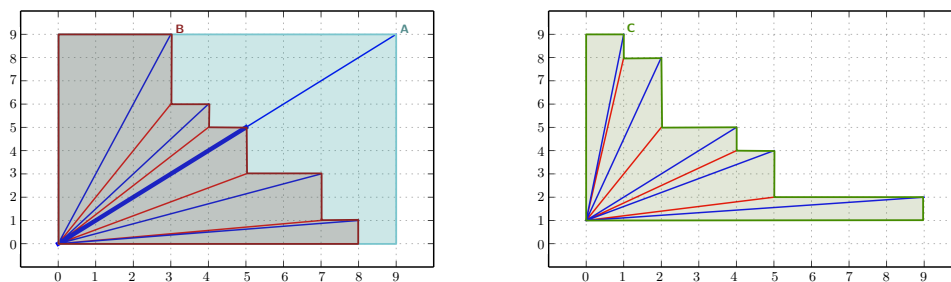
Figure 16 Lifespans of A, B (left) and C (right) in the signed barcode.

The resulting signed barcode is shown in Figure 15. As expected, the lifespans of the three clusters A, B, C appear as separate subsets of the bars, as shown in Figure 16. Checking whether one of these three subsets does correspond to the lifespan of some feature can be done by computing the coefficient assigned to the corresponding interval in the generalized rank decomposition of M. The decorated barcode would provide this information as well.

References

1 Hideto Asashiba, Mickaël Buchet, Emerson G. Escolar, Ken Nakashima, and Michio Yoshiwaki. On interval decomposability of 2d persistence modules, 2018. `arXiv:1812.05261`.

2 Hideto Asashiba, Emerson G Escolar, Ken Nakashima, and Michio Yoshiwaki. On approximation of 2d persistence modules by interval-decomposables. *arXiv preprint*, 2019. `arXiv:1911.01637`.

3 Magnus Botnan and William Crawley-Boevey. Decomposition of persistence modules. *Proceedings of the American Mathematical Society*, 148(11):4581–4596, 2020.

4 Magnus Bakke Botnan, Vadim Lebovici, and Steve Oudot. On Rectangle-Decomposable 2-Parameter Persistence Modules. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. `doi:10.4230/LIPIcs.SoCG.2020.22`.

5 Magnus Bakke Botnan, Steffen Oppermann, and Steve Oudot. Signed barcodes for multiparameter persistence via rank decompositions and rank-exact resolutions. *arXiv preprint*, 2021. `arXiv:2107.06800`.

6 William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.

7 Tamal K Dey and Cheng Xin. Generalized persistence algorithm for decomposing multiparameter persistence modules. *arXiv preprint*, 2019. `arXiv:1904.03766`.

8 Woojin Kim and Facundo Memoli. Generalized persistence diagrams for persistence modules over posets. *arXiv preprint*, 2018. `arXiv:1810.11517`.

9 Claudia Landi. The rank invariant stability via interleavings. In *Research in computational topology*, pages 1–10. Springer, 2018.

10 Michael Lesnick and Matthew Wright. Interactive visualization of 2-d persistence modules. *arXiv preprint*, 2015. `arXiv:1512.00180`.

11 Alexander McCleary and Amit Patel. Edit distance and persistence diagrams over lattices, 2021. `arXiv:2010.07337`.

12 Nikola Milosavljević, Dmitriy Morozov, and Primoz Skraba. Zigzag persistent homology in matrix multiplication time. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 216–225. ACM, 2011.

13 Dmitriy Morozov. *Homological illusions of persistence and stability*. PhD thesis, Duke University, 2008.

14 Amit Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, 1(3):397–419, 2018.