Probabilistic Analysis of Multiparameter Persistence Decompositions into Intervals

Ángel Javier Alonso
Institute of Geometry, Graz University of Technology, Austria

Michael Kerber
Institute of Geometry, Graz University of Technology, Austria

Primoz Skraba
Queen Mary University of London, United Kingdom

Abstract

Multiparameter persistence modules can be uniquely decomposed into indecomposable summands. Among these indecomposables, intervals stand out for their simplicity, making them preferable for their ease of interpretation in practical applications and their computational efficiency. Empirical observations indicate that modules that decompose into only intervals are rare. To support this observation, we show that for numerous common multiparameter constructions, such as density- or degree-Rips bifiltrations, and across a general category of point samples, the probability of the homology-induced persistence module decomposing into intervals goes to zero as the sample size goes to infinity.

2012 ACM Subject Classification Mathematics of computing → Topology

Keywords and phrases Topological Data Analysis, Multi-Parameter Persistence, Decomposition of persistence modules, Poisson point processes

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.6

Related Version *Full Version*: `https://arxiv.org/abs/2403.11939`

Funding *Ángel Javier Alonso*: Austrian Science Fund (FWF) grant P 33765-N.
Michael Kerber: Austrian Science Fund (FWF) grant P 33765-N.

Acknowledgements The authors thank Jan Jendrysiak for helpful discussions. M.K. and P.S. acknowledge the Dagstuhl Seminar 23192 "Topological Data Analysis and Applications" that initiated this collaboration.

1 Introduction

Motivation. Persistence modules capture the topological evolution of data across a range of parameters and are a major object of study in topological data analysis (TDA). To understand the structure of persistence modules, they are often split into a direct sum of indecomposable elements. By the Krull-Remak-Schmidt theorem, such a decomposition is unique up to isomorphism. If the parameter space is one-dimensional, that is, the persistence module is filtered over the real line, every indecomposable has the structure of an interval, meaning that it represents a certain topological feature in the data that is active in the range of scales given by the interval boundaries. The collection of these intervals forms the famous barcode of persistent homology.

In this work, we study persistence modules over two real parameters. The concept of intervals generalizes into this setting, but it is not true that every 2-parameter persistence module decomposes into intervals. In fact, the 2-parameter case is already of *wild representation type*, meaning that indecomposables can become arbitrarily complicated. Those modules that do admit a decomposition into intervals are called *interval decomposable*.

Restricting attention to interval-decomposable modules seems attractive: an interval in the decomposition can be interpreted, as in the one-parameter case, as a topological feature that persists over a range of scales, providing a simple interpretation of these summands. Moreover, this subclass of modules allows for faster algorithms for certain problems, such as the computation of the bottleneck distance [19].

The advantages of interval-decomposable modules raise the question of how commonly we encounter this case in applications. Unfortunately, experiments suggests that, typically, persistence modules seem to contain non-interval summands; see [1, 24] for recent case studies. This leads to the motivating question for this paper: can we capture this empirical evidence through a probabilistic mathematical statement?

Contribution. We show that many commonly used 2-parameter persistence modules are unlikely to be interval decomposable when constructed over a large point sample. More precisely, let $\mathcal{P}_n$ denote a Poisson point sample in the d-dimensional unit cube with an expected number of n points and fix $p \leq d-1$. Let $F(\mathcal{P}_n)$ denote a bifiltration constructed over $\mathcal{P}_n$, and consider $H_p(F(\mathcal{P}_n))$, the induced persistence module in p-homology. Then, we prove that, for common choices of the bifiltration F, the probability that $H_p(F(\mathcal{P}_n))$ is interval-decomposable approaches zero as $n \to \infty$. In particular, we take F to be the sublevel offset (or Čech or Vietoris-Rips) bifiltration, with random function values, or given by a kernel density estimate or a fixed function, the degree offset (or Čech or Vietoris-Rips) bifiltration, or the multicover bifiltration (with $p = 0$); see the text for the detailed statements.

The proof consists of two major parts: we first show that for any fixed point set S with a constant number of points in the unit cube, the Poisson point sample $\mathcal{P}_n$ contains an approximate scaled copy of S in some subcube with probability going to 1 as $n \to \infty$. This result follows from basic properties of Poisson point processes and are straightforward for those familiar with stochastic geometry. However, we believe that our exposition is of benefit for the TDA community.

The second part is to identify point patterns that lead to non-intervals in the induced persistence module. The definition of this pattern depends on the bifiltration F chosen, so a separate proof is required for all aforementioned choices of F. Importantly, these patterns have to be stable under slight perturbation of the point coordinates to use the probabilistic result from the first part. We aim for small point patterns and achieve the provably minimal cardinality of points in many examples. Small point patterns are more likely to be realized for smaller samples, so our examples facilitate a quantified analysis of interval-decomposability for $\mathcal{P}_n$ with concrete values of n (which is not subject of this paper). On a technical level, we show the presence of non-intervals by restricting attention to a finite subposet of $\mathbb{R}^2$ and show that an indecomposable non-interval summand is present. This technique together with the small size of the point patters leads to short and pictorial proofs of our main theorems.

Related work. Perhaps closest to our result is the paper by Buchet and Escolar [16]. They provide a simple construction of a simplicial complex that yields indecomposables of arbitrary large degree, and the construction is stable under small perturbations. This is partially stronger than what we provide, as our construction only guarantees indecomposables of degree 2. However, their complexes are only shown to be realized via a dynamic point process in Euclidean space and does not directly apply to the bifiltrations considered in our work.

Also closely related is the work by Alonso and Kerber [1]. Adapted to our notation, they show that for F the sublevel offset bifiltration and $p = 0$, at least a quarter of the indecomposables of $H_p(F(X_n))$ will be intervals in expectation when n goes to infinity, with

X_n being a point process with rather weak assumptions. While the lower bound of 1/4 might not be tight (as suggested by experimental evidence in their paper), our result complements their result, saying that we cannot expect all indecomposables to be intervals.

Recently, Hiraoka et al. [24] do a case study on the indecomposables of a commutative ladder, a special case of bifiltrations. Their experiments confirm the observation that while intervals are frequent, non-intervals typically appear in the decomposition (Figure 17 in [24]).

In the purely algebraic setting, Bauer and Scoccola [6] have recently shown that being nearly indecomposable (under the interleaving distance) is a generic property of multiparameter persistence modules.

The bifiltrations considered in this work are the standard models for distance- and density-based bifiltrations in the literature; they are discussed, for instance, in the recent survey by Botnan and Lesnick [14] and investigated regarding their stability properties by Blumberg and Lesnick [8]. Efficient algorithms to compute such bifiltrations are an active area of research, with recent results for sublevel offset bifiltrations [2], for bifiltrations of clique complexes (including Vietoris-Rips complexes) [3] and for multicover [15, 18]. The meataxe algorithm [31] provides a method to compute the decomposition of the induced persistence module. That algorithm works for more generalized setups; a specialized algorithm for persistence modules has been proposed by Dey and Xin [20]. An important intermediate step from bifiltrations to their decomposition is the computation of a minimal presentation, for which efficient algorithms also exist [19, 23, 29].

The study of homology in random settings is also an active topic of research; see [10] for a survey of the geometric setting. The homology of random geometric complexes at a single scale has been extensively studied, e.g. [9, 26, 34]. Much less is known about the *persistent* homology of random geometric complexes, but the last few years have seen significant progress, including limit theorems [25] and the expected maximal persistence [11], with recent extensions to multiparameter persistence [13]. The technique of exhibiting a point set that appears with positive probability and has a given property is a standard approach, and is often used to show bounds in random settings; the work of Kahle [26], which proves lower bounds on the radius at which the k-dimensional Betti number is non-zero in Vietoris-Rips complexes, is an example of this.

2 Ball configurations and Poisson processes

We refer the reader to [27] for a reference on the basic facts below. The *Poisson distribution* $\mathrm{Poisson}(\lambda)$ with rate λ is a discrete random variable X with

$$\mathbb{P}(X = k) = \frac{\lambda^k}{k!} e^{-\lambda}.$$

For $n \in \mathbb{N}$, a *Poisson point sample* $\mathcal{P}_n$ over a measurable[1] set X is obtained by sampling a natural number N from $\mathrm{Poisson}(n)$ and then sampling N points $(x_1, \ldots, x_N)$ independently and uniformly at random from X. In this work, we will restrict to the case that X is the d-dimensional unit cube $[0,1]^d$.

A Poisson point sample has two important two properties: for every set $A \subset [0,1]^d$, the random variable $|\mathcal{P}_n \cap A|$ is a Poisson distribution with rate $n \operatorname{Vol}(A)$, and the sample satisfies *spatial independence*: for disjoint sets $A, B \subset [0,1]^d$, $|A \cap \mathcal{P}_n|$ and $|B \cap \mathcal{P}_n|$ are independent random variables. See Definition 3.1 in [27].

[1] More precisely, the set must be measurable with respect to the Lebesgue measure, however in our case the sets are nice so we generally omit the qualifier.

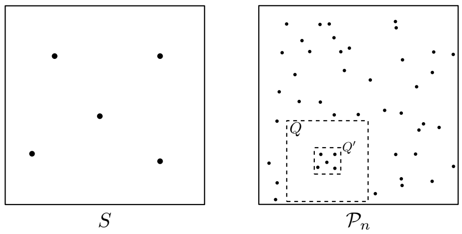

Figure 1 Left: A point set S of 5 points. Right: A Poisson point sample with a subcube that contains an isolated scaled copy of S.

Fix now an arbitrary point set S of constant size in $(0,1)^d$ and $\varepsilon > 0$. We want to show that for n large enough, it is likely that $\mathcal{P}_n$ contains a scaled version of the configuration S up to a perturbation of every point by at most ε; see Figure 1 for an illustration.

We define the property of containing a scaled version of S formally. Fix $S = (a_1, \dots, a_m) \in [0,1]^d$, and let $B_1, \dots, B_m$ denote the balls centered at $a_1, \dots, a_m$ with radius ε. We assume for simplicity that ε is sufficiently small, so that $B_1, \dots, B_m$ are pairwise disjoint and every B_i is contained in the unit cube.

Next, fix a realization of the Poisson point process $\mathcal{P}_n$. For a fixed subcube $Q \subset [0,1]^d$, let α denote the side length of Q and x its center. Furthermore, let Q' denote the subcube of Q with same center and side length $\frac{\alpha}{4\sqrt{d}}$.

There is a canonical axis-preserving bijection that maps $[0,1]^d$ to Q'. This bijection maps $a_1, \dots, a_m$ to points in Q', and the ε-balls $B_1, \dots, B_m$ to pairwise disjoint balls $B'_1, \dots, B'_m$ of radius $\frac{\alpha\varepsilon}{4\sqrt{d}}$. We say that $\mathcal{P}_n$ *contains a scaled ε-copy of S in Q*, if $\mathcal{P}_n$ contains exactly one point in B'_i, for every $i = 1, \dots, m$, and $\mathcal{P}_n$ contains no further point in Q. We say that $\mathcal{P}_n$ *contains a scaled ε-copy of S* if there exists a subcube Q such that $\mathcal{P}_n$ contains a scaled ε-copy of S in Q. We note that the size of the subcube Q' and the condition that Q contains no other points ensures that the points in Q' are "well-separated" from the remaining points in $\mathcal{P}_n$. Specifically, the distance between any point in Q' and any point outside Q is at least the diameter of Q'. See Figure 1 for an example.

▶ **Theorem 1.** *For every finite point set S and every $\varepsilon > 0$ small enough, there exists a constant α independent of n such that*

$$\mathbb{P}(\mathcal{P}_n \textit{ contains a scaled } \varepsilon\textit{-copy of } S) \geq 1 - e^{-\alpha n}.$$

Proof. We consider a subcube Q of side length $n^{-1/d}$. The scaled balls $B'_1, \dots, B'_m$ then have radius $\frac{\varepsilon n^{-1/d}}{4\sqrt{d}}$ and, hence, volume $\frac{\varepsilon^d n^{-1}}{(4\sqrt{d})^d}\omega_d$, where ω_d is the volume of the d-dimensional unit ball. By the first characteristic property of Poisson point samples, the number of points in $\mathcal{P}_n \cap B'_i$ is a Poisson distribution with rate $\frac{\varepsilon^d}{(4\sqrt{d})^d}\omega_d$, and therefore is independent of n. Thus, the probability that B'_i contains exactly one point of P_n is some constant λ_1 that is also independent of n.

Set $T \coloneqq Q \setminus (B'_1 \cup \ldots \cup B'_m)$ as the complement of the scaled ε-balls in Q. For $\mathcal{P}_n$ to contain a scaled ε-copy of S in Q, T must not contain a point. The probability for this to occur is lower bounded by the probability that Q does not contain a point. The number of points in $\mathcal{P}_n \cap Q$ is a Poisson distribution with rate 1 (since the volume of Q is $1/n$). So, the probability for no points in $\mathcal{P}_n \cap T$ is lower-bounded by a constant λ_0 independent of n.

$\mathcal{P}_n$ contains a scaled copy of S in Q if and only if each B'_i contains exactly one point of $\mathcal{P}_n$, and T contains no point of $\mathcal{P}_n$. By spatial independence, we have that these events are independent, and, therefore, the probability that $\mathcal{P}_n$ contains a scaled copy of S in Q is at least $\lambda = \lambda_0 \cdot (\lambda_1)^m$, which is a constant independent of n.

To complete the proof, note that we can pack at least $\lfloor n^{1/d} \rfloor^d$ disjoint d-dimensional subcubes of side length $n^{-1/d}$ in the unit cube without overlap, which is larger than $\frac{1}{2}n$ for n large enough. In each of these cubes, the probability of not containing a scaled copy of S is at most $1-\lambda$. Again by spatial independence, the probability that none of the subcubes contain a scaled copy of S is at most $(1-\lambda)^{\frac{1}{2}n}$ which may be upper bounded by $e^{-\alpha n}$ for some constant α. ◀

The above is for *homogenous* Poisson processes, but below we show that the result holds for non-homogeneous Poisson processes as well. A non-homogeneous Poisson point process is determined by an (integrable) intensity function f. In practice, f is a probability density function. The number of points in any set $A \subset \mathbb{R}^d$ is a Poisson random variable with rate $n \cdot \int_A f(x)dx$. We note that this type of process retains the property that the number of points (and point configurations) in any countable number of fixed disjoint sets are independent.

The proof, analogous to the one of Theorem 1, is included in the full version.

▶ **Corollary 2.** *Let $\mathcal{P}_n^f$ be a non-homogeneous Poisson point process with a continuous intensity $n \cdot f(x)$ for $x \in \mathbb{R}^d$. If for some $p \in \mathbb{R}^d$ and constant $\delta > 0$, there exists a cube $p + [0, \delta]^d$ where f is strictly positive, then for every finite finite point set S and every $\varepsilon > 0$ small enough, there exists a constant α independent of n such that*

$$\mathbb{P}(\mathcal{P}_n^f \textit{ contains a scaled } \varepsilon\textit{-copy of } S) \geq 1 - e^{-\alpha n}.$$

3 Decompositions of persistence modules

Persistence modules. For a partially ordered set (poset) P, a *persistence module* M over P is a functor from P to Vec, the category of (for us, finite-dimensional) vector spaces over a fixed field K. This means that a persistence module assigns to each $p \in P$ a vector space M_p and to any two $p \leq q$ a linear map $M_{p \to q} : M_p \to M_q$, such that $M_{p \to s} = M_{q \to s} \circ M_{p \to q}$ for all $p \leq q \leq s$, and $M_{p \to p}$ is the identity. Morphisms and isomorphisms are defined as usual for functor categories.

In this work, we only consider product posets of $\mathbb{R}_+$ and $\mathbb{R}_+^{\mathrm{op}}$, where $\mathbb{R}_+ := [0, \infty)$ and $\mathbb{R}_+^{\mathrm{op}}$ is the opposite poset, and finite subposets of these products.

Decomposition. For two persistence modules M, N over a common poset P, its *direct sum* $M \oplus N$ is the persistence module defined by taking direct sums pointwise: for $p \in P$, we take $M_p \oplus N_p$, and, for every $p \leq q$, the map $M_p \oplus N_p \to M_q \oplus N_q$, $(x, y) \mapsto (M_{p \to q}(x), N_{p \to q}(y))$.

A module M is called *indecomposable* if $M \cong A \oplus B$ implies that $A = 0$ or $B = 0$, where 0 is the persistence module for which all vector spaces are 0. All the persistence modules we consider, which come from finite point sets, can be decomposed uniquely, up to isomorphism, as a direct sum of indecomposable summands by the Krull-Remak-Schmidt theorem.

Restrictions. Given a subposet $Q \subset P$, a persistence module M over P induces a persistence module over Q in the natural way: taking the vector spaces and linear maps of M at the places of Q. We call this persistence module $M\big|_Q$ over Q the *restriction* of M. Moreover, if $M \cong A \oplus B$ is a decomposition, we have $M\big|_Q \cong A\big|_Q \oplus B\big|_Q$.

A persistence module M over a poset P is *thin* if $\dim M_p \leq 1$ for all $p \in P$. Clearly, if M is a thin persistence module over $\mathbb{R}_+^2$, the restricted module $M\big|_Q$ is thin as well for any subposet $Q \subset \mathbb{R}_+^2$. A persistence module M is called *interval decomposable*, if it admits a decomposition into thin modules (this is not the usual definition, see below). We obtain:

$$\begin{tikzcd} K \arrow[r, "\binom{1}{1}"] & K^2 \arrow[r, "(0\ 1)"] & K \\ & K \arrow[u, "\binom{0}{1}"] & \end{tikzcd} \qquad \begin{tikzcd} K & \\ K^2 \arrow[u, "(0\ 1)"] \arrow[r, "(1\ 0)"] & K \\ K \arrow[u, "\binom{1}{1}"] & \end{tikzcd}$$

Figure 2 Two non-thin indecomposable persistence modules over finite posets. Both posets are subposets of $\mathbb{R}^2_+$.

▶ **Lemma 3.** *If a persistence module M over $\mathbb{R}^2_+$ is interval decomposable, then for any subposet $Q \subset \mathbb{R}^2_+$ the restriction $M\big|_Q$ is interval decomposable.*

One usually defines interval decomposability as decomposition into interval modules, a subcase of thin indecomposables whose support is convex and connected (e.g., [14, Def 2.1]). However, for all persistence modules we consider in this paper, the two notions coincide: every thin module can be decomposed into (finitely many) interval modules [4, Thm 24].

Two non-thin indecomposables. As an interface to the examples in the next section, we consider the two modules depicted in Figure 2. To check that they are indecomposable is elementary: For the right hand side module M, suppose that $M \cong A \oplus B$ and note that if A is non-zero at the source vertex of the underlying graph, then it is also not zero at the sinks of the graph, and thus necessarily $B = 0$. The left example can be similarly checked.

We will use the following frequently. It follows immediately from Lemma 3.

▶ **Lemma 4.** *Let M be a persistence module over $\mathbb{R}^2_+$ and P be a subposet of $\mathbb{R}^2_+$ such that $M\big|_P \cong A \oplus B$, where A is one of modules in Figure 2. Then, M is not interval decomposable.*

4 Non-interval decomposability with high probability

A *bifiltration* F is a functor from the poset $\mathbb{R}^2_+$ to the category Top of topological spaces or the category Simp of simplicial complexes. Writing $H_k \colon \mathsf{Top} \to \mathsf{Vec}$ for the k-homology functor with coefficient in the base field K, $H_k(F) : \mathbb{R}^2_+ \to \mathsf{Vec}$ is a persistence module.

We are interested in properties of point sets that are preserved under perturbation. We say that a finite point set $S' \subset \mathbb{R}^d$ is an *ε-perturbation* of another point set $S \subset \mathbb{R}^d$ if there exists a bijection $f \colon S' \to S$ such that $\|x - f(x)\| \leq \varepsilon$.

Below, we go over typical bifiltrations based on point sets. We argue the same way for each type of bifiltration: we first show that there exists a point set S whose bifiltration induces a non-interval-decomposable persistence module, and that this is preserved by ε-perturbation. We then conclude that the Poisson process $\mathcal{P}_n$ of Section 2 has the same property with high probability precisely because S exists. Note that with high probability means with probability greater than $1 - \frac{1}{n^c}$ with $c > 0$.

4.1 Offset bifiltrations

The *offset (or union-of-balls) filtration* $\mathcal{O}_\bullet(S) \colon \mathbb{R}_+ \to \mathsf{Top}$ of a point set $S \subset \mathbb{R}^d$, is given by setting $\mathcal{O}_r(S) := \bigcup_{x \in S} B_r(x)$, where $B_r(x)$ is the closed ball with center x and radius r. The *Čech filtration* $\mathcal{C}_\bullet(S) \colon \mathbb{R}_+ \to \mathsf{Simp}$ is the *nerve* of $\mathcal{O}_\bullet(S)$, that is, $\mathcal{C}_r(S) = \{\sigma \subset S \mid$

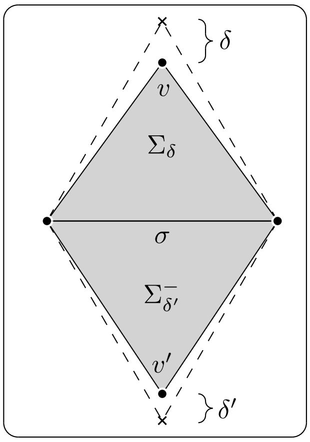

Figure 3 The construction of Lemma 6 for $\mathbb{R}^2$: two perturbed equilateral triangles glued along σ.

$\bigcap_{x\in\sigma} B_r(x) \neq \varnothing\}$, or, equivalently, a simplex $\sigma \subset S$ is in $\mathcal{C}_r(S)$ if the minimum enclosing ball of σ has radius less than or equal to r. Given a function $\gamma\colon S \to \mathbb{R}_+$, we define the *sublevel offset bifiltration* $\mathcal{O}_\bullet(\gamma)$ of S with respect to γ by taking

$$\mathcal{O}_{r,s}(\gamma) := \mathcal{O}_r(\gamma^{-1}([0,s])).$$

And analogously for the *sublevel Čech bifiltration* $\mathcal{C}_\bullet(\gamma)$. As a consequence of the nerve theorem [5], $H_k(\mathcal{O}_\bullet(\gamma))$ and $H_k(\mathcal{C}_\bullet(\gamma))$ are isomorphic persistence modules for any k.

For the rest of the subsection, we fix $d > 1$ to be a fixed constant and $1 \leq k \leq d-1$ as the homology dimension; the case $k = 0$ will be handled separately in Section 4.3.

▶ **Theorem 5.** *Let $\mathcal{P}_n$ be a Poisson point process in $\mathbb{R}^d$. For each point x in $\mathcal{P}_n$, assign $\gamma(x)$ uniformly at random in $[0,1]$. Then, the persistence module $H_k(\mathcal{O}_\bullet(\gamma))$ is not interval decomposable with high probability.*

We follow the strategy hinted at the start of the section. The first step is the following:

▶ **Lemma 6.** *There is a finite point set $S \subset \mathbb{R}^d$ of $d+2$ points and a function $\gamma\colon S \to \mathbb{R}$ such that $H_{d-1}(\mathcal{O}_\bullet(\gamma))$ (or, equivalently, $H_{d-1}(\mathcal{C}_\bullet(\gamma))$) is not interval decomposable.*

Moreover, the same holds for any ε-perturbation $S' \subset \mathbb{R}^d$ of S, for a small enough $\varepsilon > 0$.

Proof. Let Σ be a regular d-simplex with unit edge length in $\mathbb{R}^d$. Fix a facet σ of Σ. For $\delta > 0$, we denote by Σ_δ a perturbed version of Σ given by moving the vertex opposite to σ perpendicularly towards σ by a distance of δ. Define Σ^- as the regular d-simplex obtained by reflecting Σ along the hyperplane of σ, and define $\Sigma^-_{\delta'}$ as its perturbed version in the analogous way, using $\delta' < \delta$. See Figure 3.

We take S to be the vertex set of $\Sigma_\delta \cup \Sigma^-_{\delta'}$, for small enough $\delta' < \delta$ such that the circumcenter of Σ_δ is still in the interior of Σ_δ. Denote by v the vertex in Σ_δ opposite to σ, and let $\gamma\colon S \to \mathbb{R}$ be any function such that $\gamma(v) > \gamma(x)$ for any other $x \in S$. We argue that $H_{d-1}(\mathcal{C}_\bullet(\gamma))$ is not interval decomposable by looking at a restriction of it. The minimum enclosing ball (given by its circumscribed hypersphere) of Σ has radius $R := \sqrt{\frac{d}{2(d+1)}}$, and the minimum enclosing ball of its facets has radius $r := \sqrt{\frac{d-1}{2d}}$ (see, for instance, [22, Theorem 4.5.1]). Let r_δ be the radius of the minimum enclosing ball of the facets that are not σ of Σ_δ.

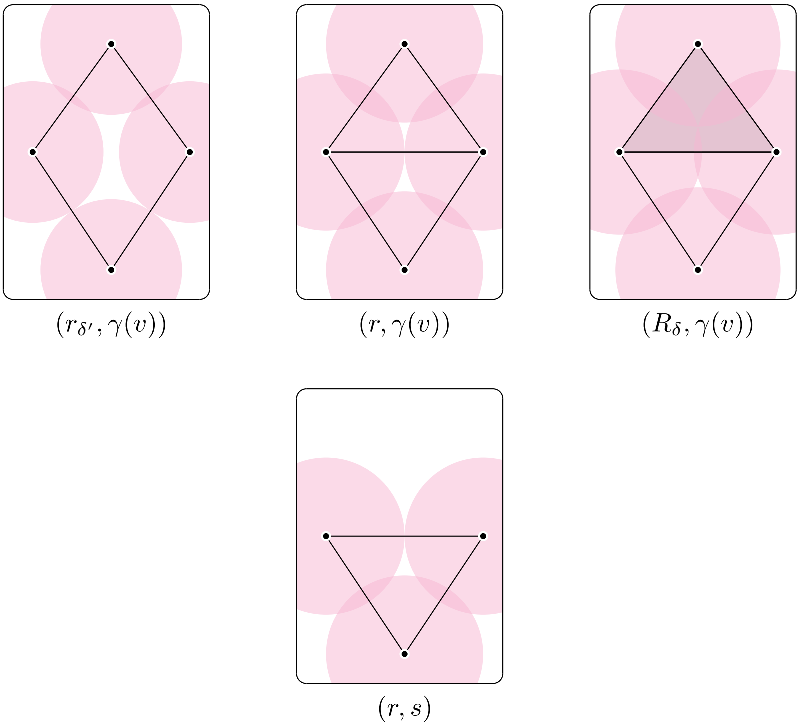

Figure 4 The Čech and offset bifiltration of Lemma 6 restricted to the subposet $P \subset \mathbb{R}^2_+$.

Let R_δ be the minimum enclosing ball radius of Σ_δ. By the way we construct S, we have

$r_\delta < r_{\delta'} < r$, and $R_\delta < R_{\delta'} < R$.

Finally, let s be the maximum value of γ in $S \setminus \{v\}$.

We look at the persistence module $H_{d-1}(\mathcal{C}_\bullet(\gamma))$ restricted to the finite subposet $P \subset \mathbb{R}^2_+$ given by $(r_{\delta'}, \gamma(v)) < (r, \gamma(v)) < (R_\delta, \gamma(v))$ and (r, s), see Figure 4. Note that $\mathcal{C}(\gamma)_{(r_{\delta'},\gamma(v))}$ consists of all $(d-1)$-simplices of Σ_δ and $\Sigma^-_{\delta'}$ except σ, $\mathcal{C}(\gamma)_{(r,\gamma(v))}$ consists of all $(d-1)$-simplices, while the only d-simplex in $\mathcal{C}(\gamma)_{(R_\delta,\gamma(v))}$ is Σ_δ. All in all, the restricted persistence module can be seen to be isomorphic to the left example of Figure 2, and Lemma 4 yields the desired statement. Finally, note that there exists a sufficiently small $\varepsilon > 0$ such that any ε-perturbation of S induces a non-interval decomposable persistence module. ◀

▶ Remark 7. The example above is minimal: if the $(d-1)$-homology of an offset filtration is not a thin persistence module, then the underlying point set has at least $d+2$ points.

Proof of Theorem 5. Given homological degree k and dimension d, set $\tilde{d} = k + 1 \leq d$. Let $S \subset \mathbb{R}^{\tilde{d}}$ be the construction of Lemma 6. We adopt the notation used in the proof of that lemma. We scale down S and embed it in $[0,1]^d$. By the triangle inequality, if S is stable under any ε-perturbation in $\mathbb{R}^{\tilde{d}}$, it is stable under any $\frac{\varepsilon}{2}$-perturbation in $\mathbb{R}^d$. Theorem 1 applied to S yields that $\mathcal{P}_n$ contains an (isolated) scaled $\frac{\varepsilon}{2}$-copy of S with high probability. The only requirement on the function γ of Lemma 6 is that it has a unique maximum at $v \in S$, as defined in the proof of the lemma. Since $\gamma(x)$ is sampled uniformly, this occurs with probability $\frac{1}{\tilde{d}+2}$ (since $|S| = \tilde{d} + 2$) which is a constant independent of n. It follows that with high probability, $\mathcal{P}_n$ contains a scaled $\frac{\varepsilon}{2}$-copy of S where v has maximal γ-value.

Let S' be the scaled $\frac{\varepsilon}{2}$-copy of S contained in $\mathcal{P}_n$. Consider the subposet $P \subset \mathbb{R}^2_+$ given in the proof of Lemma 6, and change its values so that the construction of the lemma applies to the scaled-down copy S', obtaining $P' \subset \mathbb{R}^2_+$. Theorem 1 guarantees that the distance between a point in S' and a point of $\mathcal{P}_n$ not in S' is at least the diameter of S'. It follows that, at every $p \in P'$, the offset of S', $\mathcal{O}_p(\gamma\big|_{S'})$, is an isolated connected component in $\mathcal{O}_p(\gamma)$. It follows that $H_{\tilde{d}-1}(\mathcal{O}_\bullet(\gamma\big|_{S'}))$, restricted to P', is a summand in the decomposition of $H_{\tilde{d}-1}(\mathcal{O}_\bullet(\gamma))$, restricted to P'. Since this summand is not interval decomposable and since $\tilde{d} = k+1$, we conclude that $H_k(\mathcal{O}_\bullet(\gamma))$ is not interval decomposable, as required. ◀

Probability density functions and density estimation. The function γ of the sublevel offset bifiltration $\mathcal{O}_\bullet(\gamma)$ is usually given by the application itself or taken so that γ captures the *density* of the points: lower values of γ mean points of lower density, and thus more likely to be noise [17]. Therefore, we are usually more interested in regions of higher density. Note that when working with sublevel sets, as in $\mathcal{O}_\bullet(\gamma)$, we would need to invert the order of the densities – our methods work either way.

If γ captures the density, we consider a Poisson process with intensity γ, $\mathcal{P}^\gamma_n$, and then we would ideally use γ itself in $\mathcal{O}_\bullet(\gamma)$. However, γ must usually be estimated. We first treat the case of estimating the density, in Theorem 8 below, and then, later in Theorem 10, we treat the more technical case of a fixed γ, under generic smooth conditions.

A common approach to density estimation is to use a *kernel* [32]. Choosing a kernel K_h, where h is the *bandwidth*, the estimated density at a point p is

$$\hat{\gamma}(p) = \frac{1}{|S|}\sum_{q\in S} K_h(d(p,q)),$$

where S is the point set. We note that the kernel must satisfy certain conditions but, since we do not use them, we refer the reader to any text on statistics e.g. [33]. The most basic kernel is a *ball kernel*, $K_h(x) = 1$ if $x \leq h$; the estimated density is

$$\hat{\gamma}(p) = \frac{|S \cap B_h(p)|}{|S|},$$

where $B_h(p)$ is the closed ball of radius h. This is the kernel we will now work with. We believe our methods work for more general choices of kernel, but the probabilistic elements of the proof become much more delicate; we comment on this in the conclusion.

▶ **Theorem 8.** *Let $\mathcal{P}^\gamma_n \in \mathbb{R}^d$ be a Poisson point process with intensity γ, where γ is a Morse function. If $\hat{\gamma}$ is the ball estimator with $h = (\log n/n)^{1/d}$, the persistence module $H_k(\mathcal{O}_\bullet(\hat{\gamma}))$ is not interval decomposable with high probability.*

Proof. The proof is similar to Theorem 5. The main difference is we require the kernel estimates to be independent for each subcube Q_i. That is for any $p \in Q_i$ and $q \in Q_j$, the distance between p and q must be at least twice the bandwidth. Therefore in the proof of Corollary 2 (resp. the proof of Theorem 1), we can only pack $(\lfloor n/4^d \log n)^{1/d}\rfloor)^d = O(n/\log n)$ subcubes into the unit cube such that the estimator supports are disjoint.

It remains to analyze the case for one subcube. We again fix a point set S and ε which leads to a non-interval decomposition. While the estimates for each sample are dependent, we show that the probability of any permutation of ordering is strictly positive. For any two points in $s_1, s_2 \in S$, consider any points $p \in B_\varepsilon(s_1)$ and $q \in B_\varepsilon(s_2)$. Define the sets $A_p = B_h(p) - B_h(q) \cap B_h(p)$ and $A_q = B_h(q) - B_h(p) \cap B_h(q)$. By construction, A_p and A_q of not intersect each other, nor any of the subcubes we consider (see Figure 5). As A_p

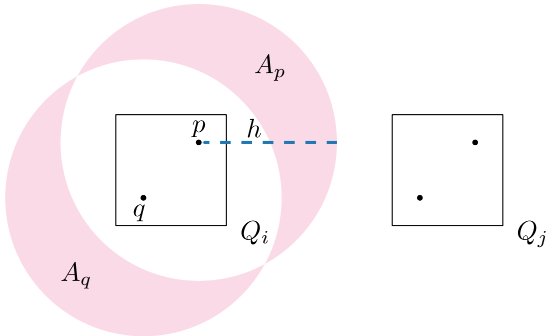

Figure 5 The setup of the proof of Theorem 8.

and A_q are disjoint, the number of points in in A_p and A_q are independent Poisson random variables. Further, as they have volume $\Omega(1/n)$. Hence, $\mathbb{P}(|A_p \cap \mathcal{P}_n| > |A_q \cap \mathcal{P}_n|)$ and $\mathbb{P}(|A_p \cap \mathcal{P}_n| < |A_q \cap \mathcal{P}_n|)$ are both strictly positive, which implies that any ordering occurs with positive probability. ◀

▶ Remark 9. The argument above requires that $h \to 0$, as for a fixed h the ordering of the points in each Q_i are no longer independent, requiring a more delicate analysis. A similar problem occurs for kernels with non-compact support.

We now tackle the case that we construct $\mathcal{O}_\bullet(\varphi)$ for some fixed function φ, under generic smoothness conditions:

▶ **Theorem 10.** *Let $\mathcal{P}_n$ be a Poisson point process in $\mathbb{R}^d$ and φ a real-valued (C^2)Morse function on $\mathbb{R}^d$ with finitely many critical points. The persistence module $H_k(\mathcal{O}_\bullet(\varphi))$ is not interval decomposable with high probability.*

The main technical tool is a form of linearization. Essentially, we show that there exists an isometric transformation (with scaling) that the function will induce the required ordering on the point set S. We first define a technical condition on the point sets:

▶ **Definition 11.** *An ordering on a point set $S \subset \mathbb{R}^d$ is called f-linear if there exists a linear function $f : \mathbb{R}^d \to \mathbb{R}$ which realizes the ordering on S, i.e. $\forall x, y \in S$, $x \preceq y \Leftrightarrow f(x) \leq f(y)$.*

It is straightforward to check that all of the obstructions to being interval indecomposable we present are f-linear. We further make the assumption that f induces a total order on the point set (which is the case in all our constructions). The main technical result follows:

▶ **Lemma 12.** *Given a f-linear ordering on a point set S and a real-valued (C^2) Morse function g, at any regular point of g, for all sufficiently small scalings there exists an isometry of (scaled) S such that f and g induce the same order on S.*

Proof. We assume the values of f at S are distinct. Let a be a regular point of g and index the points in S such that $f(p_i) < f(p_j)$ if $i < j$. Without loss of generality, we can assume p_0 is at a which we can take to be the origin and $f(p_0) = 0$.

We require the notion of a *cone ordering* [30]. This is a partial ordering on a convex cone D where for $x, y \in D$, $x \preccurlyeq y$ if and only if $y - x \in D$. Consider the cone induced by the span of the vectors $U = \{p - q \mid p, q \in S, f(p) \geq f(q)\}$. By construction, the cone ordering agrees with the order induced by f on S. As the values of f are distinct, this cone is acute, i.e. the angle between any two points in the cone is less than π.

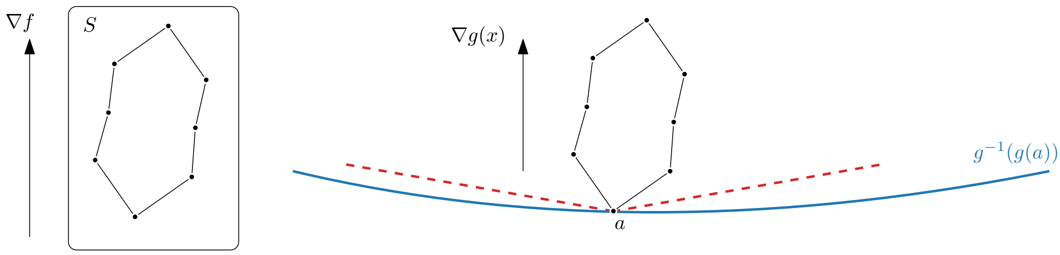

Figure 6 The construction from Lemma 12. We tilt S to ensure the values at the points are unique. We then place the minimum point at the desired point. If we scale S sufficiently, the cone given by the red lines which contains the set U and will be contained in the superlevel set locally.

Consider this cone at p_0, i.e. the origin. If for all $u \in U$, $\nabla g(p) \cdot u \geq 0$ for all $p \in \text{conv}(S)$, where conv(S) denotes the convex hull of S. Then [30, Corollary 4] states that the order induced by g is equivalent to the ordering induced by the cone ordering and hence f.

As a is a regular a point of g, $g^{-1}(g(a))$ is a $(d-1)$-dimensional surface with positive reach and hence bounded curvature [21]. Placing p_0 at a, rotate S such that $\nabla f(p_0) = \nabla g(a)$. Sufficiently scaling down S, we can ensure that the cone spanned by U lies on one side of the tangent plane of $g^{-1}(g(a))$, so $\nabla g(p_0) \cdot u \geq 0$. As the gradient is continuous for C^2 functions, for all sufficiently scaled S, $\nabla g(p) \cdot u \geq 0$ for all $p \in \text{conv}(S)$, completing the proof. ◀

An illustration of the construction in the lemma for the S in Figure 8 can be seen in Figure 6. Note that we tilt it slightly to ensure the function values at the points are unique and the red lines indicate an outer bound for the cone U. We are now ready to finish the proof:

Proof of Theorem 10. Since we assume that φ only has a finite number of critical points, we can place a cube of constant side length such that all points in the cube are regular points of φ. As it is Morse and C^2, this implies that the gradient of $\nabla\varphi(x)$ is non-zero at all points x in the cube. As in Theorem 1, we can pack $O(n)$ subcubes Q_i of side-length $O(1/n^{1/d})$ within the constant sized cube. By Lemma 12, for each Q_i, we can orient S such that the ordering of the vertices induced by φ is as in Lemma 6. The result follows. ◀

We note that by Corollary 2 this result holds in the case of a non-homogeneous Poisson process, e.g. if γ is the intensity of the Poisson process.

Degree offset bifiltration. A problem with using density estimation is that we need to choose a bandwidth parameter. The degree bifiltrations [28] improve on this by being parameter-free. For a point set S and $r \in \mathbb{R}$, we define the *degree* of a point $x \in S$ at scale r as the number of points $y \in S$ with $y \neq x$ such that $\|x - y\| \leq 2r$. Intuitively, the degree of a point estimates its density: higher degree implies higher density. We denote by $\mathcal{D}_{r,k}(S)$ the subset of points of S that have at least degree k at scale r. Then, the *degree offset bifiltration* $\mathcal{DO}_\bullet(S)\colon \mathbb{R}_+ \times \mathbb{R}_+^{\text{op}} \to \mathsf{Top}$ is given by $\mathcal{DO}_{r,k}(S) = \mathcal{O}_r(\mathcal{D}_{r,k}(S))$. Similarly for the *degree Čech bifiltration* $\mathcal{DC}_\bullet(S)$.

▶ **Lemma 13.** *There is a finite point configuration S in $\mathbb{R}^d$ of $d+3$ points and an $\varepsilon > 0$ such that for any ε-perturbation S' of S, $H_{d-1}(\mathcal{DO}_\bullet(S'))$ is not interval decomposable.*

Proof. We consider the point set S and the notation of the proof of Lemma 6. Recall that v is the perturbed point of Σ_δ, and v' is that of $\Sigma^-_{\delta'}$.

We construct $S' := S \cup \{w\}$ by placing a new point w close to v' in the direction perpendicular to σ. This distance between w and v' is chosen small enough such that, at each scale $s \in \{r_{\delta'}, r, R_\delta\}$, the offsets $\mathcal{O}_s(S)$ and $\mathcal{O}_s(S')$ are essentially the same: that there is a strong deformation retraction from $\mathcal{O}_s(S')$ to $\mathcal{O}_s(S)$. If follows that each inclusion $\mathcal{O}_s(S) \hookrightarrow \mathcal{O}_s(S')$ induces an isomorphism in $(d-1)$-homology. The same applies to the inclusion $\mathcal{O}_s(S \setminus \{v\}) \hookrightarrow \mathcal{O}_s(S' \setminus \{v\})$.

Let $P' \subset \mathbb{R}_+ \times \mathbb{R}_+^{\mathrm{op}}$ be the following subposet, meant to be compared to the one in Figure 4,

$$\begin{array}{ccccc} (r_{\delta'}, d) & < & (r, d) & < & (R_\delta, d) \\ & & \vee & & \\ & & (r, d+1). & & \end{array}$$

Note that at scale $r_{\delta'}$, all points have degree at least d: this is clear if $d = 2$ and, otherwise, $r_{\delta'} > \frac{1}{2}$, so every edge of Σ_δ and $\Sigma^-_{\delta'}$ is present. At scale r, v has degree d (it is connected only to the d vertices of σ) and all other points have degree at least $d+1$, precisely because we have added w. By the isomorphisms in $(d-1)$-homology noted above, and noting that $\mathcal{D}_{r_{\delta'},d}(S') = S'$ and that $\mathcal{D}_{r,d+1}(S') = S' \setminus v$, in parallel to the proof of Lemma 6, we have that the persistence module $H_{d-1}(\mathcal{DO}_\bullet(S'))$ restricted to P' is isomorphic to the left example of Figure 2. The desired statement follows by Lemma 3.

It is clear that there exists a sufficiently small $\varepsilon > 0$ such that the same holds for any ε-perturbation of S. ◀

Using the same argument as before, we obtain:

▶ **Theorem 14.** *Let $\mathcal{P}_n$ be a Poisson point process in $\mathbb{R}^d$. The persistence module $H_k(\mathcal{DO}_\bullet(\mathcal{P}_n))$ is not interval decomposable with high probability.*

4.2 Rips bifiltrations

A variant of the Čech complex $\mathcal{C}_r(S)$ for a finite point set $S \subset \mathbb{R}^d$ is the *(Vietoris-)Rips complex* $\mathcal{R}_r(S)$, whose simplexes are those subsets of S of diameter at most $2r$:

$$\mathcal{R}_r(S) \coloneqq \{\sigma \subset S \mid \operatorname{diam} \sigma \leq 2r\}, \tag{1}$$

where $\operatorname{diam} \sigma$ is the maximum distance between two points in σ. The Rips complex assembles into the *Rips filtration* $\mathcal{R}_\bullet(S)$ over $\mathbb{R}_+$, and, given a function $\gamma \colon S \to \mathbb{R}_+$, we define the *sublevel Rips bifiltration* $\mathcal{R}_\bullet(\gamma) \colon \mathbb{R}_+ \to \mathsf{Simp}$ by $\mathcal{R}_{r,s}(\gamma) \coloneqq \mathcal{R}_r(\gamma^{-1}([0,s]))$.

We now comment on how to extend the results of the previous section to the Rips setting. The construction in Lemma 6 does not work immediately. Indeed, in Figure 4 at $(r, \gamma(v)) \in P$, the Rips complex of such a point set consists of two (filled) triangles: the 1-skeleton (the graph given by its vertices and edges) of $\mathcal{C}_r(S)$ and $\mathcal{R}_r(S)$ coincide, and $\mathcal{R}_r(S)$ is given by the cliques of this 1-skeleton. Focusing on the case of $\mathbb{R}^2$, and referring back to the proof of Lemma 6, Figure 7 proves by picture the following lemma:

▶ **Lemma 15.** *There is a finite point configuration S in $\mathbb{R}^2$, a function $\gamma \colon S \to \mathbb{R}$ such that $H_1(\mathcal{R}_\bullet(\gamma))$ is not interval decomposable. Moreover, the same holds for any ε-perturbation of S, for an $\varepsilon > 0$ small enough.*

The example S uses as building block the d-dimensional cross-polytope in $\mathbb{R}^d$: the convex hull of the $2d$ points $\{\pm e_1, \dots, \pm e_d\}$, where $e_1, \dots, e_d$ are the endpoints of the standard basis vectors. We denote the vertex set of the d-dimensional cross-polytope by O_d. Note that, for all $r \in [\sqrt{2}/2, 1)$, $\mathcal{R}_r(O_d)$ is the boundary of the d-cross-polytope, which implies that $H_{d-1}(\mathcal{R}_r(O_d)) \cong K$, and, for $r = 1$, we have that $\mathcal{R}_1(O_d)$ contains every subset of O_d.

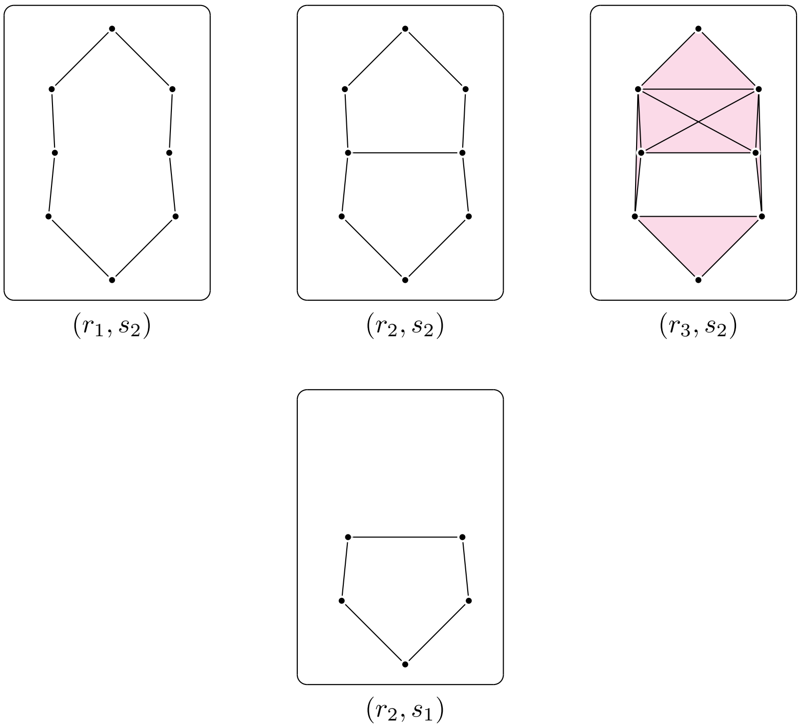

Figure 7 A point set S in the plane and its associated sublevel Rips bifiltration restricted to a finite subposet. Shaded triangles are the 2-simplices of the Rips complex.

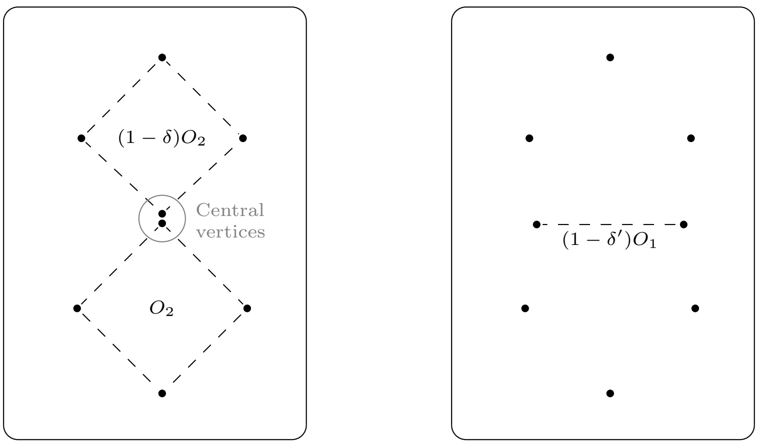

Figure 8 How to construct the point set of Figure 7. On the left, a copy of $(1-\delta)O_2$ and O_2 lying side by side. On the right, the marked central vertices are replaced by $(1-\delta')O_1$.

The point set S consists of a copy of O_2 and a scaled-down version $(1-\delta)O_2$, for a small enough $\delta > 0$, lying side by side, see Figure 8, where we replace the two central vertices by a copy $(1-\delta')O_1$, with a sufficiently small $\delta' > \delta$. An analogous constructions works in higher dimensions $\mathbb{R}^d$, $d > 2$, and $(d-1)$-homology; we omit the details.

The arguments used to prove the main theorems of Section 4.1 apply in the same way to the example of Lemma 15 and thus to the Rips setting.

4.3 Zero-dimensional homology and clustering

We now handle the case of zero-dimensional homology. Here, the Rips and Čech complexes coincide, $H_0(\mathcal{R}_r(S)) = H_0(\mathcal{C}_r(S))$, since both have the same 1-skeleton, and, by the nerve theorem [5], $H_0(\mathcal{C}_r(S)) \cong H_0(\mathcal{O}_r(S))$. We use the same strategy as before:

▶ **Lemma 16.** *There is a finite point configuration S in $\mathbb{R}^2$ such that $H_0(\mathcal{O}_\bullet(\gamma))$ is not interval decomposable. The same holds for any ε-perturbation of S, with $\varepsilon > 0$ small enough.*

Proof. Consider the point set S of Figure 9, and a function $\gamma\colon S \to \mathbb{R}_+$ such that $\gamma(A) < \gamma(B) < \gamma(C) < \gamma(D)$. Let $P \subset \mathbb{R}^2_+$ be the subposet given by $(0,\gamma(B))$, $(0,\gamma(C))$, $(3.2,\gamma(C))$, and $(2.7,\gamma(D))$, which is also described in Figure 9.

Recalling that $\mathcal{O}_{r,s}(\gamma)$ is generated by the set of connected components, and that, for any two $(r,s) \le (r',s')$, the map $\mathcal{O}_{(r,s)\to(r',s')}(S)$ is induced by the inclusion of connected components, one can see that $H_0(\mathcal{O}_\bullet(\gamma))$ restricted to P is isomorphic to and decomposes as:

$$\begin{tikzcd}
K^2 & & & K & & & K & \\
K^3 \arrow[u, "{\begin{pmatrix}1&0&1\\0&1&0\end{pmatrix}}"] \arrow[r, "{\begin{pmatrix}1&0&0\\0&1&1\end{pmatrix}}"] & K^2 & \cong & K \arrow[u, "(1)"] \arrow[r, "(1)"] & K & \oplus & K^2 \arrow[u, "{\begin{pmatrix}0&1\end{pmatrix}}"] \arrow[r, "{\begin{pmatrix}1&0\end{pmatrix}}"] & K \\
K^2 \arrow[u, "{\begin{pmatrix}1&0\\0&1\\0&0\end{pmatrix}}"] & & & K \arrow[u, "(1)"] & & & K. \arrow[u, "{\begin{pmatrix}1\\1\end{pmatrix}}"] &
\end{tikzcd}$$

Lemma 4 yields that $H_0(\mathcal{O}_\bullet(S))$ is not interval-decomposable, as required. Finally, it is clear that the indecomposables do not change by perturbing the points by a small enough ε. ◀

The same arguments used in the theorems of Section 4.1 apply to the example above.

Multicover. We now show that multicover bifiltrations have non-interval-decomposable 0-homology persistence modules with probability going to 1. The *multicover bifiltration* $\mathrm{Cov}_\bullet(S)\colon \mathbb{R}_+ \times \mathbb{R}_+^{\mathrm{op}} \to \mathsf{Top}$ of a finite point set $S \subset \mathbb{R}^d$ is given by

$$\mathrm{Cov}_{r,k}(S) := \{y \in \mathbb{R}^d \mid \|y - x\| \le r \text{ for at least } k \text{ points of } x \in S\},$$

that is, $\mathrm{Cov}_{r,k}$ is the region of $\mathbb{R}^d$ that is covered by at least k of the r-balls centered at the points in S. Note that $\mathrm{Cov}_{r,1}(S) = \mathcal{O}_r(S)$, and that the multicover bifiltration is sensitive to density. Figure 11 shows an example for various r and k.

▶ **Lemma 17.** *There is a finite point configuration S in $\mathbb{R}^2$ such that $H_0(\mathrm{Cov}_\bullet(S))$ is not interval decomposable. The same holds for any ε-perturbation of S, with $\varepsilon > 0$ small enough.*

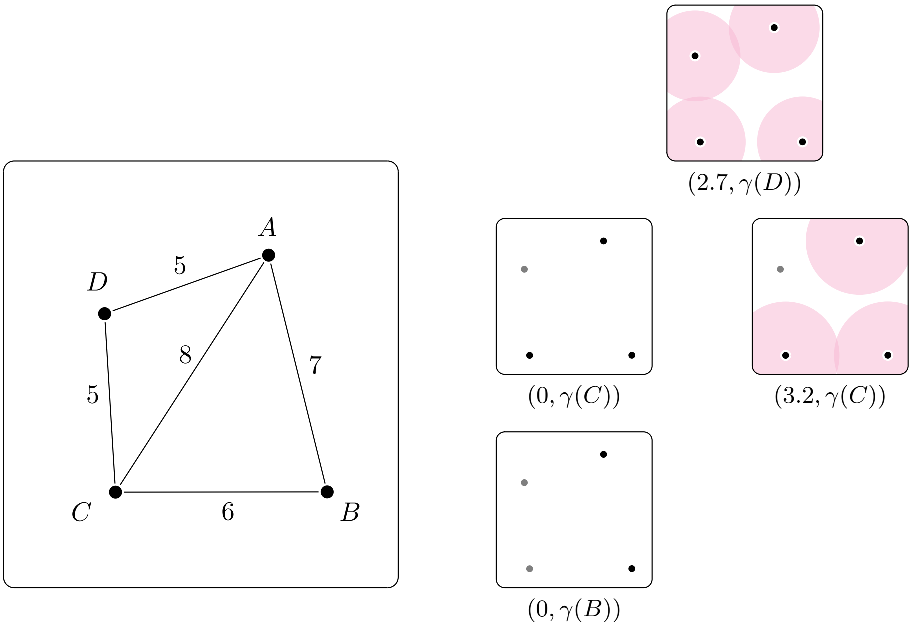

Figure 9 The point set $S \subset \mathbb{R}^2$ of the proof of Lemma 16. On the right we describe the offset bifiltration restricted to the subposet $P \subset \mathbb{R}^2_+$ of the same proof.

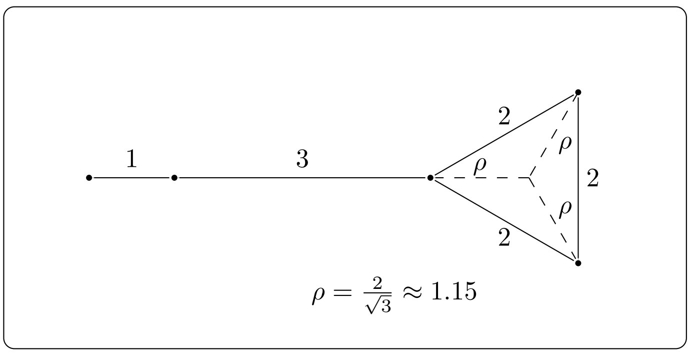

Figure 10 The point set $S \subset \mathbb{R}^2$ of Lemma 17. The edges mark the distance between the points.

Proof. The point set S is the one described in Figure 10. Let $P \subset \mathbb{R}_+ \times \mathbb{R}_+^{\mathrm{op}}$ be the subposet

$$\begin{array}{ccc} & & (r = 2, k = 1) \\ & & \vee \\ (r = 0.6, k = 2) & < & (r = 1.2, k = 2) \\ & & \vee \\ & & (r = 1.2, k = 3). \end{array}$$

We draw $\mathrm{Cov}_p(S)$ for each $p \in P$ in Figure 11. Tracking how the connected components merge, it can be seen that $H_0(\mathrm{Cov}_\bullet(S))\big|_P$ is isomorphic to

$$\begin{tikzcd} & K \\ K \arrow[r, "\binom{1}{0}"] & K^2 \arrow[u, "\begin{pmatrix}1 & 1\end{pmatrix}"'] \\ & K, \arrow[u, "\binom{0}{1}"] \end{tikzcd} \quad \cong \quad \begin{tikzcd} & K \\ K \arrow[r, "\binom{1}{1}"] & K^2 \arrow[u, "\begin{pmatrix}0 & 1\end{pmatrix}"'] \\ & K. \arrow[u, "\binom{0}{1}"] \end{tikzcd}$$

Comparing this module with the left example of Figure 2, Lemma 4 yields that $H_0(\mathrm{Cov}_\bullet(S))$ is not interval decomposable. This holds for any small enough ε-perturbation of the points. ◀

5 Conclusion

Building on the empirical observation that in most cases, bifiltrations do not admit a complete decomposition into intervals, we showed that this occurs with probability going to 1 whenever there is random sampling. We focused on the intervals (or rather thin decomposables), but our technique extends to prove the presence of arbitrarily complicated indecomposables: if one can give an ε-stable example for a point set for which the persistence module of its bifiltration contains that indecomposable when restricted to a subposet, then a Poisson point process will contain that indecomposable with probability going to 1. From these results, there are many interesting directions and open questions one can consider.

First, our examples occur due to noise, and so have short lifetimes. A natural question from the perspective of TDA is whether approximations of bifiltrations or persistence modules, through discretization or the erosion strategy by Bjerkevik [7], would yield a different expected decomposition structure. Furthermore, we have only scratched the surface of possible probabilistic questions. While we show existence, we would like to understand of the distribution of non-interval summands. Within the range of parameters corresponding to “noise”, what is the most persistent non-interval summand? Does the number of such summands obey a central limit theorem or other universality law, e.g. as it is conjectured in the single parameter case for geometric random complexes [12]?

Though we cover several settings, there are some remaining. For example, while we covered a simple kernel density estimator, we expect the same results to hold for more general kernels, even though the probabilistic elements of the proof become much more delicate; we leave this case for future work. Additionally, can one show similar results for more general processes, such as binomial or determinantal processes where there is dependence between the points, or random functions such as Gaussian random fields? In such cases, significantly more advanced probabilistic techniques will be needed.

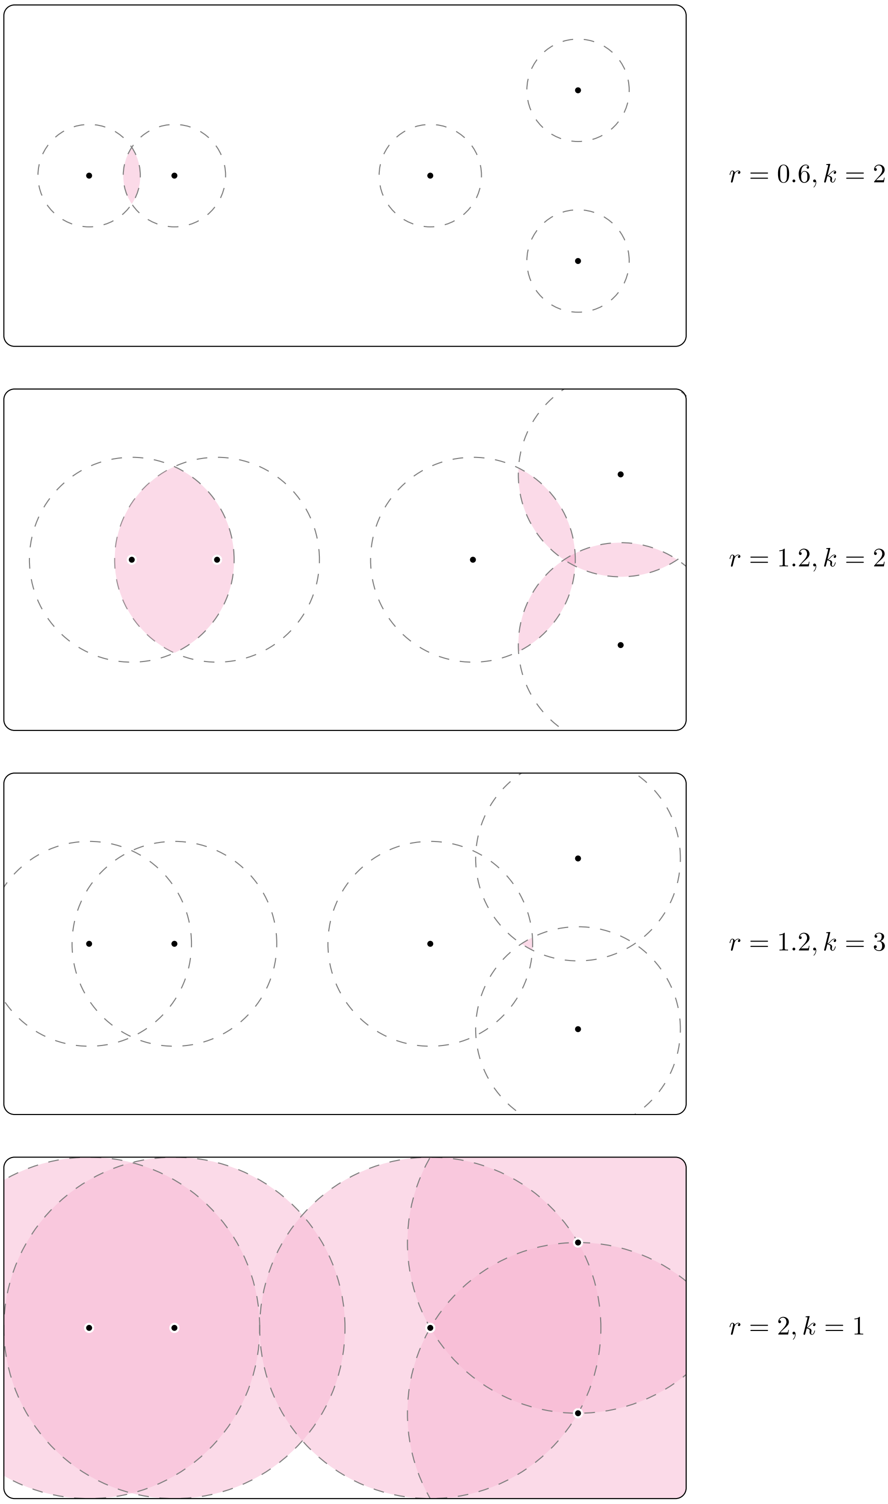

Figure 11 The multicover bifiltration $\mathrm{Cov}_\bullet(S)$ of the point set $S \subset \mathbb{R}^2$ of Lemma 17.

References

1 Ángel Javier Alonso and Michael Kerber. Decomposition of zero-dimensional persistence modules via rooted subsets. In *39th International Symposium on Computational Geometry, SoCG 2023*, 2023. doi:10.4230/LIPIcs.SoCG.2023.7.

2 Ángel Javier Alonso, Michael Kerber, Tung Lam, and Michael Lesnick. Delaunay bifiltrations of functions on point clouds. In *Symposium on Discrete Algorithms SODA 2024*, 2024. doi:10.1137/1.9781611977912.173.

3 Ángel Javier Alonso, Michael Kerber, and Siddharth Pritam. Filtration-domination in bifiltered graphs. In *Proceedings of the Symposium on Algorithm Engineering and Experiments, ALENEX 2023*, 2023. doi:10.1137/1.9781611977561.CH3.

4 Hideto Asashiba, Mickael Buchet, Emerson G. Escolar, Ken Nakashima, and Michio Yoshiwaki. On interval decomposability of 2D persistence modules. *Computational Geometry*, 105-106:101879, 2022. doi:10.1016/j.comgeo.2022.101879.

5 Ulrich Bauer, Michael Kerber, Fabian Roll, and Alexander Rolle. A unified view on the functorial nerve theorem and its variations. *Expo. Math.*, 41(4):125503, 2023. doi:10.1016/j.exmath.2023.04.005.

6 Ulrich Bauer and Luis Scoccola. Generic multi-parameter persistence modules are nearly indecomposable, 2023. arXiv:2211.15306.

7 Håvard Bakke Bjerkevik. Stabilizing decomposition of multiparameter persistence modules, 2023. arXiv:2305.15550.

8 Andrew J. Blumberg and Michael Lesnick. Stability of 2-parameter persistent homology. *Foundations of Computational Mathematics*, 2022. doi:10.1007/s10208-022-09576-6.

9 Omer Bobrowski and Robert J. Adler. Distance functions, critical points, and the topology of random Čech complexes. *Homology, Homotopy and Applications*, 16(2):311–344, 2014. doi:10.4310/hha.2014.v16.n2.a18.

10 Omer Bobrowski and Matthew Kahle. Topology of random geometric complexes: a survey. *Journal of Applied and Computational Topology*, 1(3-4):331–364, 2018. doi:10.1007/s41468-017-0010-0.

11 Omer Bobrowski, Matthew Kahle, and Primoz Skraba. Maximally persistent cycles in random geometric complexes. *The Annals of Applied Probability*, pages 2032–2060, 2017. doi:10.1214/16-AAP1232.

12 Omer Bobrowski and Primoz Skraba. A universal null-distribution for topological data analysis. *Scientific Reports*, 13(1), July 2023. doi:10.1038/s41598-023-37842-2.

13 Magnus Bakke Botnan and Christian Hirsch. On the consistency and asymptotic normality of multiparameter persistent Betti numbers. *Journal of Applied and Computational Topology*, December 2022. doi:10.1007/s41468-022-00110-9.

14 Magnus Bakke Botnan and Michael Lesnick. An introduction to multiparameter persistence. In *2020 International Conference on Representations of Algebras*, pages 77–150. EMS Press, November 2023. doi:10.4171/ecr/19/4.

15 Mickaël Buchet, Bianca B. Dornelas, and Michael Kerber. Sparse higher order čech filtrations. In *39th International Symposium on Computational Geometry, SoCG 2023, June 12-15, 2023, Dallas, Texas, USA*, pages 20:1–20:17, 2023. doi:10.4230/LIPICS.SOCG.2023.20.

16 Mickaël Buchet and Emerson G. Escolar. Realizations of indecomposable persistence modules of arbitrarily large dimensions. *Journal of Computational Geometry*, 13(1), 2022. doi:10.20382/jocg.v13i1a12.

17 Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete Comput. Geom.*, 42(1):71–93, 2009. doi:10.1007/s00454-009-9176-0.

18 René Corbet, Michael Kerber, Michael Lesnick, and Georg Osang. Computing the multicover bifiltration. In *37th International Symposium on Computational Geometry, SoCG 2021*, pages 27:1–27:17, 2021. doi:10.4230/LIPICS.SOCG.2021.27.

19 Tamal K. Dey and Cheng Xin. Computing bottleneck distance for 2-D interval decomposable modules. In *34th International Symposium on Computational Geometry, SoCG 2018*, pages 32:1–32:15, 2018. doi:10.4230/LIPIcs.SoCG.2018.32.

20 Tamal K. Dey and Cheng Xin. Generalized persistence algorithm for decomposing multiparameter persistence modules. *Journal of Applied and Computational Topology*, 6:271–322, 2022. doi:10.1007/s41468-022-00087-5.

21 Herbert Federer. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. doi:10.2307/1993504.

22 Miroslav Fiedler. *Matrices and graphs in geometry*, volume 139 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2011. doi:10.1017/CBO9780511973611.

23 Ulderico Fugacci, Michael Kerber, and Alexander Rolle. Compression for 2-parameter persistent homology. *Computational Geometry*, 109:101940, 2023. doi:10.1016/j.comgeo.2022.101940.

24 Yasuaki Hiraoka, Ken Nakashima, Ippei Obayashi, and Chenguang Xu. Refinement of interval approximations for fully commutative quivers, 2023. arXiv:2310.03649.

25 Yasuaki Hiraoka, Tomoyuki Shirai, and Khanh Duy Trinh. Limit theorems for persistence diagrams. *Ann. Appl. Probab.*, 28(5):2740–2780, 2018. doi:10.1214/17-AAP1371.

26 Matthew Kahle. Random geometric complexes. *Discrete Comput. Geom.*, 45(3):553–573, 2011. doi:10.1007/s00454-010-9319-3.

27 Günter Last and Mathew Penrose. *Lectures on the Poisson process*, volume 7. Cambridge University Press, 2017.

28 Michael Lesnick and Matthew Wright. Interactive Visualization of 2-D Persistence Modules. arXiv:1512.00180.

29 Michael Lesnick and Matthew Wright. Computing minimal presentations and bigraded betti numbers of 2-parameter persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 6(2):267–298, 2022. doi:10.1137/20M1388425.

30 A. W. Marshall, D. W. Walkup, and R. J.-B. Wets. Order-preserving functions: Applications to majorization and order statistics. *Pacific J. Math.*, 23:569–584, 1967.

31 R. A. Parker. The computer calculation of modular characters (the meat-axe). In Michael D. Atkinson, editor, *Computational Group Theory*, pages 267–274. New York: Academic Press, 1984.

32 Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Springer Netherlands, 1986. doi:10.1007/978-1-4899-3324-9.

33 Larry Wasserman. *All of statistics: a concise course in statistical inference*, volume 26. Springer, 2004.

34 D. Yogeshwaran and Robert J. Adler. On the topology of random complexes built over stationary point processes. *Ann. Appl. Probab.*, 25(6):3338–3380, 2015. doi:10.1214/14-AAP1075.