

Asymptotic Improvements on the Exact Matching Distance for 2-parameter Persistence

Håvard Bakke Bjerkevik*

Michael Kerber†

November 22, 2021

Abstract

In the field of topological data analysis, persistence modules are used to express geometrical features of data sets. The matching distance $d_{\mathcal{M}}$ measures the difference between 2-parameter persistence modules by taking the maximum bottleneck distance between 1-parameter slices of the modules. The previous fastest algorithm to compute $d_{\mathcal{M}}$ exactly runs in $O(n^{8+\omega})$, where ω is the matrix multiplication constant. We improve significantly on this by describing an algorithm with expected running time $O(n^5 \log^3 n)$. We first solve the decision problem $d_{\mathcal{M}} \leq \lambda$ for a constant λ in $O(n^5 \log n)$ by traversing a line arrangement in the dual plane, where each point represents a slice. Then we lift the line arrangement to a plane arrangement in \mathbb{R}^3 whose vertices represent possible values for $d_{\mathcal{M}}$, and use a randomized incremental method to search through the vertices and find $d_{\mathcal{M}}$. The expected running time of this algorithm is $O((n^4 + T(n)) \log^2 n)$, where $T(n)$ is an upper bound for the complexity of deciding if $d_{\mathcal{M}} \leq \lambda$.

1 Introduction

Persistent Homology [14, 15, 28] is a method to summarize the topological properties of a data set across different scales. The last 20 years have witnessed numerous applications of persistent homology in various scientific fields, including neuroscience [2, 30, 31], material science [18, 25], cosmology [29, 33], and many others, shaping the field of *topological data analysis* [5]. Part of the success of this field is the presence of fast algorithms for the most important sub-tasks, namely computing a discrete summary of the topology of the data (the so-called *barcode*) via Gaussian elimination of a sparse matrix, and computing the distance of two such barcodes (the so-called *bottleneck distance*) as a minimal cost matching via the Hopcroft-Karp algorithm.

The standard theory of persistent homology assumes a single scalar (real) parameter to determine the scale. In many applications, however, it is natural to alter more than one parameter: As a simple example, consider the situation of a point cloud in Euclidean space, where one parameter controls until what distance two points are considered as close, and a second parameter determines until which local density a point is labeled as an outlier. Every combination of these two parameters results in a shape approximating the data, and studying the evolution of topological properties across the 2-parameter space of shapes leads to the theory of *multi-parameter persistence*. Unfortunately, already for two parameters, there is no direct generalization of the barcode [6], and the most natural notion of distance between two 2-parameter data sets, the *interleaving distance*, is NP-hard to approximate for any factor smaller than 3 [4].

*Institute of Geometry, TU Graz; bjerkevik@tugraz.at

†Institute of Geometry, TU Graz; kerber@tugraz.at

Despite these negative results, there are alternative notions of distance which allow for a polynomial-time computation. Among them is the *matching distance* ($d_{\mathcal{M}}$), the topic of this paper. The main idea is that when imposing a linear condition on the parameters, we effectively restrict the data to one-parameter and can compute barcodes and bottleneck distance for this restriction. The matching distance is then defined as the supremum of the bottleneck distance over all possible linear restrictions.

Kerber et al. showed recently [21] that the matching distance can be computed (exactly) in polynomial time, more precisely, with an asymptotic complexity of $O(n^{8+\omega})$, where n is the size of the input and ω is the matrix multiplication constant.

Contribution. We describe a randomized algorithm that computes the exact matching distance in $\tilde{O}(n^5)$, where \tilde{O} means that we ignore logarithmic factors in n . We achieve this substantial improvement in complexity through a combination of several techniques from computational geometry, computational topology, and graph matchings. We only list the major ideas and refine them in Section 2 where we outline the approach in more detail.

- We focus on the decision question of whether $d_{\mathcal{M}} \leq \lambda$ or $d_{\mathcal{M}} > \lambda$ for a fixed $\lambda \in \mathbb{R}$. We show that this question can be addressed using a line arrangement with $O(n^2)$ lines, substantially reducing the size from the $O(n^4)$ lines needed in the previous approach by Kerber et al. [21].
- We traverse the line arrangement and maintain the barcodes and the bottleneck matching between them during this traversal. Using the vineyard algorithm by Cohen-Steiner et al. [10] and a result about computing bottleneck matchings from [22], this allows us to spend only linear time per visited arrangement cell (up to logarithmic factors), resulting in an $\tilde{O}(n^5)$ algorithm for the decision problem.
- We identify a set of $O(n^6)$ candidate values for λ that is guaranteed to contain the exact value of the matching distance. These values are obtained from triple intersections in a plane arrangement with $O(n^2)$ planes.
- We show how to avoid considering all candidate values through randomized incremental construction (e.g. [12, Ch 4]). More precisely, we maintain upper and lower bounds on the matching distance and cluster the candidates in groups. We can quickly rule out clusters if they do not improve the bounds, and by randomization, we expect only a logarithmic number of improvements of the bounds.

Motivation and Related work. A common theme in topological data analysis is to consider topological properties of a data set as a proxy and analyze a collection of data sets in terms of their topology. A distance measure allows for standard data analysis tasks on such proxies, such as clustering or visualization via multi-dimensional scaling.

For such a distance measure to be meaningful, it is required that it is *stable*, that is, a slightly perturbed version of a data set is in small distance to the unperturbed data set. At the same time, the distance should also be *discriminative*, that is, distinguish between data sets with fundamental differences. (In the extreme case, the distance measure that assigns 0 to all pairs is maximally stable, but also maximally non-discriminative.) Finally, for practical purposes, the distance, or at least an approximation of it, should also be efficiently computable.

For the case of one parameter, the bottleneck and Wasserstein distances satisfy all these properties. They are both stable under certain tameness conditions [8, 9, 32] and efficient software is

available for computation [22]. Moreover, the bottleneck distance is *universal*, meaning that it is the most discriminative distance among all stable distances [1].

In the multi-parameter setup, the situation is more complicated: the interleaving distance is universal [26], but NP-hard to compute [4]. There is also a multi-parameter extension of the bottleneck distance based on matching indecomposable elements, which is unstable and only efficiently computable in special cases [13]. The matching distance is an attractive alternative, because it is stable [24] and computable in polynomial time. Moreover, since it compares two multi-filtered data sets based on the *rank invariant* [6], we can expect good discriminative properties in practice, even if no theoretical statement can be made at this point. At least, first practical applications of the matching distance have been identified in shape analysis [3, 7] and computational chemistry [20]; see also [17] for a recent case study.

The matching distance allows for a rather straightforward approximation algorithm [3] through adaptive geometric subdivision with a quad-tree data structure. Recently, this approach has been refined [23] and the code has been released in the HERA library¹. These developments raised the question of practical relevance for our approach, given that a good approximation of the distance is usually sufficient in practical data. We believe that, besides the general interest in the theoretical complexity of the problem, an exact algorithm might turn out to be faster in practice than the proposed approximation scheme. The reason is that to achieve a satisfying approximation factor, the subdivision approach has to descend quite deep in the quad-tree, and while the adaptiveness of quad-trees avoids a full subdivision in some instances, there are others where large areas of the domain have to be fully subdivided to certify the result – see for instance Figure 11 in [23]. An exact approach might be able to treat such areas as one component and save time. At the same time, we emphasize that the approach presented in this paper is optimized for asymptotic complexity; we do not claim that it will be efficient in practice as described.

Acknowledgements. This work is supported by the Austrian Science Fund (FWF) grant number P 33765-N.

2 Extended overview of our approach

We give an overview describing all the main ideas of the paper, but skipping formal proofs of correctness and only introducing terms and technical constructions through simplified examples. Everything needed to prove our results will be introduced independently in later sections.

Bottleneck distance. The input data in topological data analysis often takes the form of a filtration $(X_s)_{s \geq 0}$ of simplicial complexes. Applying homology in some degree to the complexes and inclusion maps, one gets a 1-parameter *persistence module*, to which one can associate a barcode, which is a multiset of intervals (i.e., *bars*) of the form $[b, d)$. Intuitively, each bar represents a homological feature that is born at b and dies at d . Given two such persistence modules, one would like to say something about how similar they are, which is usually done by comparing their barcodes.

We consider two barcodes B_1 and B_2 as close if B_1 can be transformed into B_2 by changing all endpoints of bars in B_1 by a small value. More concretely, consider a real value $\lambda \geq 0$ and consider all barcodes that arise from B_1 by shifting both endpoints of every bar by at most λ to the left or right. In particular, a bar of length $\leq 2\lambda$ can be transformed to an empty interval, and

¹<https://github.com/grey-narn/hera>

thus removed. Likewise, any bar of length $\leq 2\lambda$ can be created by conceptually inserting an empty interval $[c, c)$ in B_1 , and shifting the left endpoint to the left and the right endpoint to the right. The *bottleneck distance* of B_1 and B_2 is the smallest λ such that B_2 can be obtained from B_1 by such a λ -perturbation.

The computation of the bottleneck distance can be reduced to computing maximum-cardinality matchings in bipartite graphs. The idea is to set up a graph for a fixed $\lambda \geq 0$ with the bars of B_1 and B_2 as vertex set and connect two bars if they can be transformed into each other via a λ -perturbation. Moreover, the graph has to be extended to account for the possibility of removing and creating bars in the perturbation – see [14, Ch. VIII.4] for details. Setting up this graph G^λ properly, we get that the bottleneck distance is at most λ if and only if G^λ has a perfect matching.

While the complexity for finding a maximum-cardinality matching is $O(n^{2.5})$ using the Hopcroft-Karp algorithm [19], it can be reduced to $O(n^{1.5} \log n)$ here because of the metric structure of the graph [16]. The idea is that a bar $[a, b)$ can be represented as a point (a, b) in the plane (resulting in the so-called *persistence diagram*), and two points are connected in G^λ if their ℓ_∞ -distance is at most λ . Then, in order to find an augmenting path, we can avoid traversing all edges of G^λ (whose size can be quadratic), but instead use a range tree data structure [12, Ch. 5] to find neighbors in the graph.

2-parameter filtrations In some important cases, our filtration is equipped with two parameters instead of one; in this case one has a family $(X_{a,b})_{a \geq 0, b \geq 0}$ of simplicial complexes with inclusion maps $X_{a,b} \rightarrow X_{a',b'}$ when $a \leq a'$ and $b \leq b'$. As in the 1-parameter setup, we are interested in the evolution of the homology as the parameters range over their domain. There is no summary for that evolution in the same way as the barcode for a single parameter. However, we can explore the 2-parameter filtration by fixing a line s of positive slope, called a *slice*, and consider the 1-parameter family of simplicial complexes $(X_{a,b})$, where (a, b) ranges over all points on s . One example is the family $(X_{a,a})_{a \geq 0}$, obtained from the slice $x = y$. Because slices have positive slope, the induced 1-parameter family is a filtration in the above case, so there is a well-defined barcode for the 2-parameter filtration with respect to every slice s . Moreover, having two 2-parameter filtrations and a slice s , there is a well-defined bottleneck distance between the slice barcodes.² The *matching distance* between two 2-parameter filtrations is the supremum of the bottleneck distance over all slices.

These definitions carry over to an arbitrary number of parameters, but the algorithm that follows will not, so we restrict to two parameters throughout.

Duality and bottleneck function We rephrase the definition of the matching distance through the well-known point-line duality of computational geometry [12, Ch. 8.2]. We assign for each line $y = ax + b$ with slope $a > 0$ a *dual point* $(a, b) \in \mathbb{R}_{>0} \times \mathbb{R}$. A point $p = (p_x, p_y) \in \mathbb{R}^2$ dualizes to a line $y = -p_x x + p_y$, and the points on that dual line are the duals of all slices that pass through p . We refer to the parameter plane as the *primal plane* and the space in which the dual points and lines are contained as the *dual plane* throughout. For every dual point (a, b) (with $a > 0$), we can assign a real value which is defined by the bottleneck distance of the two slice barcodes with respect to the slice $y = ax + b$. This induces a function $B : \mathbb{R}_{>0} \times \mathbb{R} \rightarrow \mathbb{R}$, which we call the *bottleneck function*. The matching distance is then just the supremum of B over its domain.

²Technically, the barcode and bottleneck distance depend on the parameterization of the line, but we postpone the discussion to the technical part.

Barcode templates The advantage of the described duality is that the space of slices can be decomposed into parts where the bottleneck function is simpler to analyze. As a first step, we define the *barcode pairing* of a one-parameter filtration of simplicial complexes as the collection of pairs (σ, τ) where σ is a simplex giving rise to a homology class and τ is the simplex eliminating it. In the same way we obtained slice barcodes, we also have a slice barcode pairing, and by duality, we can talk about the barcode pairing of a point in the dual plane.

The barcode can be obtained by just replacing σ and τ by the points at which they appear in the filtration, which we will refer to as their *weights*. Importantly, while the barcode depends on the weights, the barcode pairing only depends on the *order* of the weights. Hence, when moving in the dual plane, the corresponding barcode pairings remain the same as long as no two simplices switch order.

As observed by Lesnick and Wright [27], there is a line arrangement decomposing the dual plane into (2-dimensional) regions, such that the barcode pairing is identical for all dual points in a region. The line arrangement is not difficult to define: let us consider two simplices σ with weight (a, b) and σ' with weight (a', b') and assume $a < a'$. If also $b < b'$, σ will precede σ' in the filtration of any slice. If $b > b'$, the simplices σ and σ' will have the same critical value exactly for those slices that pass through the *join* of the two weights, which is the point (a', b) . Hence, the dual line $y = -a'x + b$ will split the dual plane into two half-planes where the order of the two simplices is fixed. Doing this for all pairs of simplices yield the line arrangement which we call \mathcal{L} . It consists of $O(n^2)$ lines and therefore has $O(n^4)$ regions.

Fast computation of barcode templates. Naively, computing the barcode template for each region can be done in $O(n^{4+\omega})$ time by computing the arrangement \mathcal{L} ($O(n^4)$ [12, Thm. 8.6]), picking one dual point per region ($O(n^4)$), generating the filtration along the primal slice ($O(n)$ per region), computing the slice barcodes ($O(n^\omega)$ per region), and computing their bottleneck distance ($O(n^{1.5} \log n)$ per region).

In [27], it is shown how to reduce the complexity to $O(n^5)$ using fast updates of persistence diagrams under transpositions as explained in [10]. Specifically, knowing the decomposition $RU = \partial$ for a fixed order of simplices of a simplicial complex and a transposition of two neighboring simplices in that order, we can update the decomposition (R, U) in linear time in the size of the matrix. Crossing a line of the arrangement \mathcal{L} corresponds exactly to such a transposition. This suggests the following approach: Construct a spanning tree for the dual graph of the arrangement (“dual” in a different sense than above: the vertices of the graph are the regions, and the edges correspond to pairs of regions that share an edge). The spanning tree defines a walk through \mathcal{L} with $O(n^4)$ steps that visits every region at least once (by doubling every edge and using an Eulerian tour). We compute the barcode pairing for the starting region and keep applying the transposition algorithm of [10] to get the barcode pairing for the next region in linear time. This yields an $O(n^5)$ time algorithm.

Though this method of computing the barcode at every slice is not new, using it to efficiently compute the matching distance is. Since our setting is rather different than that of [27], which allows us some simplifications, we give the details of the computation in Section 4.

Computing the matching distance – previous version. We outline the algorithm of [21]. They refine the arrangement \mathcal{L} to an arrangement \mathcal{L}' by adding $O(n^4)$ additional lines such that within each region, the bottleneck distance is determined by the distance of a fixed pair of simplices. Therefore, the bottleneck function takes a simple form in each region of \mathcal{L}' , and they show that in each region, the bottleneck function is maximized at a boundary vertex (or potentially as a limit

for unbounded cells – we ignore this case for simplicity). Hence, the algorithm simply evaluates the bottleneck function at all vertices of the arrangement and takes the maximum. This yields a running time of $O(n^{8+\omega})$ because the arrangement \mathcal{L}' has $O(n^8)$ vertices. Using the more efficient barcode computation from above, the complexity drops to $O(n^{9.5} \log n)$, because we obtain the barcodes in $O(n^9)$ time and the cost of computing the bottleneck distance at a fixed dual point is now dominated by the (geometric) Hopcroft-Karp algorithm with complexity $O(n^{1.5} \log n)$. Still, the large size of the arrangement \mathcal{L}' prevents us from further improvements with this approach.

Decision version of the matching distance. We first study the decision version of the problem: Given $\lambda \geq 0$, decide whether the matching distance is strictly greater than λ . We show how to answer this question in $\tilde{O}(n^5)$: for that, we refine the arrangement \mathcal{L} to an arrangement \mathcal{T}_λ by adding $O(n^2)$ further lines. The crucial property is that within each region of \mathcal{T}_λ , the bottleneck function B is either $\leq \lambda$ or $> \lambda$ throughout. The construction is based on the observation that if B changes from $\leq \lambda$ to $> \lambda$ along some path in the dual plane, this must happen at a dual point where the distance between two simplices along the primal slice changes from $\leq \lambda$ to $> \lambda$. However, we show that such a change can only happen at dual points that lie on constantly many lines in dual space per pair of simplices – note that some of these dual lines are vertical and therefore do not have a corresponding point in primal space.

With the arrangement \mathcal{T}_λ , the decision version is directly solvable in $O(n^{5.5} \log n)$ time: Pick one point in the interior of each region, compute the bottleneck function, and return if the function is $(> \lambda)$ for at least one point.

We further improve the complexity by the same idea of dynamic updates as for the computation of barcode templates. Recall that the question whether the bottleneck distance between two barcodes is at most λ reduces to the existence of a perfect matching in a graph G^λ . Every slice/dual point gives rise to such a graph, and by our construction, G^λ is fixed for every region of \mathcal{T}_λ . Moreover, G^λ changes in a very controlled way when crossing an edge of \mathcal{T}_λ . The most important example is that G^λ might gain or lose exactly one edge if we cross a line that changes the distance of two simplices from $\leq \lambda$ to $> \lambda$.

This suggests the following strategy: We fix a walk through \mathcal{T}_λ and maintain a perfect matching for G^λ . In many cases, we can either keep the matching or update it in constant time when walking to the next region. The exception is that an edge of the matching is removed from G^λ , turning two vertices of G^λ unmatched. Still, the remaining edges still form an “almost perfect” matching. We simply look for an augmenting path in the new G^λ to make the matching perfect again (if no augmenting path exists, the matching distance is larger than λ). In our geometric setting, computing one augmenting path can be done in $O(n \log n)$ time. So, by starting with an almost-perfect matching, we avoid the $O(\sqrt{n})$ rounds of the Hopcroft-Karp algorithm which leads to the improvement in the complexity.

Candidate values. To turn the decision version into an algorithm to compute the matching distance, we identify a candidate set $\lambda_1, \dots, \lambda_m$ with $m = O(n^6)$ with the guarantee that the exact matching distance δ is one of the candidates.

The cleanest way to describe the candidate set is by lifting all arrangements \mathcal{T}_λ in \mathbb{R}^3 parameterized by (a, b, λ) such that the intersection with $\lambda = \lambda_0$ yields \mathcal{T}_{λ_0} . It turns out that this lifting is a plane arrangement with $O(n^2)$ planes. We extend this arrangement by adding some other planes and call the resulting arrangement \mathcal{T} (which still has $O(n^2)$ planes).

For every open 3-dimensional region in \mathcal{T} , we have the property that either $B(a, b) < \lambda$ for all (a, b, λ) in the region, or $B(a, b) > \lambda$ for all (a, b, λ) in the region. Call a region *green* in the

former and *red* in the latter case. The matching distance is $\leq \lambda_0$ if and only if the plane $\lambda = \lambda_0$ intersects no red region, and the exact matching distance δ is determined by the lowest plane with that property. It follows that δ is the λ -value of a vertex of \mathcal{T} . Since such a vertex is the intersection of 3 planes, there are at most $O(n^6)$ such vertices.

Sorting the candidate set and using binary search, we can then find the exact matching distance using $\log m = O(\log n)$ queries to the decision algorithm. This results in a $O(n^6 \log n)$ algorithm for the matching distance.

Randomized construction. Our final ingredient is to avoid computing all candidates using randomization. The idea is to iterate through the $O(n^2)$ planes of \mathcal{T} in random order and maintain an interval containing the matching distance. After each iteration, the interval contains the λ -value of just one vertex that lies on the planes we have considered so far. After having run through all the planes of \mathcal{T} , we are left with just one possible value for the matching distance.

We prove that we can check if the interval needs to be updated in $\tilde{O}(n^2)$, and if it does, we can update it in $\tilde{O}(n^4 + T(n))$, where $T(n)$ is an upper bound on the complexity of solving the decision problem $d_{\mathcal{M}} \leq \lambda$. Because of the random order, we only expect $O(\log n)$ updates, leading to an expected runtime of $\tilde{O}(n^4 + T(n))$ in total. As mentioned previously, we will show that $T(n)$ can be taken to be $\tilde{O}(n^5)$, so we get a total expected runtime of $\tilde{O}(n^5)$.

We hope that our work can be used to prove an even better bound on the complexity of computing the matching distance. For instance, it follows from our results that to find an algorithm computing $d_{\mathcal{M}}$ in expected time $\tilde{O}(n^4)$, it is enough to describe an algorithm deciding if $d_{\mathcal{M}} \leq \lambda$ in expected time $\tilde{O}(n^4)$.

3 Persistence modules and the bottleneck distance

Persistence modules We consider \mathbb{R} and \mathbb{R}^2 as posets, \mathbb{R} with the usual poset structure \leq , and \mathbb{R}^2 with the poset structure given by $(a, b) \leq (a', b')$ if $a \leq a'$ and $b \leq b'$. In what follows, we fix a finite field F and let $P \in \{\mathbb{R}, \mathbb{R}^2\}$.

Definition 3.1. A persistence module M is a collection of vector spaces M_p over F for each $p \in P$ and linear transformations $M_{p \rightarrow q}: M_p \rightarrow M_q$ for all $p \leq q$. We require that $M_{p \rightarrow p}$ is the identity and that

$$M_{q \rightarrow r} \circ M_{p \rightarrow q} = M_{p \rightarrow r}$$

whenever these are defined.

If $P = \mathbb{R}$, we call M a 1-parameter (persistence) module, and if $P = \mathbb{R}^2$, we call M a 2-parameter (persistence) module.

We say that two modules M and N are *isomorphic* and write $M \simeq N$ if there is a collection $\{f_p\}_{p \in P}$ of isomorphisms $M_p \rightarrow N_p$ such that for all $p \leq q$, $f_q \circ M_{p \rightarrow q} = N_{p \rightarrow q} \circ f_p$. One can check that the isomorphism relation is an equivalence relation.

The direct sum $M \oplus N$ of modules M and N is defined straightforwardly: $(M \oplus N)_p = M_p \oplus N_p$, and $(M \oplus N)_{p \rightarrow q} = M_{p \rightarrow q} \oplus N_{p \rightarrow q}$.

Definition 3.2. Let $I \subset \mathbb{R}$ be an interval. The interval module \mathbb{I}^I is the 1-parameter module defined by $\mathbb{I}_p^I = F$ if $p \in I$, and $\mathbb{I}_p^I = 0$ if $p \notin I$. In addition, $\mathbb{I}_{p \rightarrow q}^I$ is the identity whenever $p, q \in I$.

We will see below that the 1-parameter modules we will consider are isomorphic to direct sums of interval modules. The associated intervals give rise to the barcodes mentioned in the introduction.




Figure 1: A 2-parameter module presented by $([1; -1], \{g_1, g_2\}, \{r\}, \text{gr})$ for certain choices of grades. The shade of red shows the dimension, and some of the vector spaces M_p and linear transformations $M_{p \rightarrow q}$ are shown.

Presentations We will only consider persistence modules that allow a finite description. To be precise, we restrict ourselves to *finitely presented* modules.

Let $P \in \{\mathbb{R}, \mathbb{R}^2\}$, and let $Q = ([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ be the collection of the following data:

- a finite $(m \times m')$ -matrix $[Q]$ with entries in F ,
- tuples $\mathcal{G} = (g_1, \dots, g_m)$ and $\mathcal{R} = (r_1, \dots, r_{m'})$ indexing the rows and columns of Q , respectively,
- a *grade* $\text{gr}(g) \in P$ for each $g \in \mathcal{G}$ and $\text{gr}(r) \in P$ for each $r \in \mathcal{R}$,

such that for every i and j with the entry in the i^{th} row and j^{th} column of $[Q]$ nonzero, we have $\text{gr}(g_i) \leq \text{gr}(r_j)$. We call the elements of \mathcal{G} *generators* and those of \mathcal{R} *relations*. We often treat \mathcal{G} and \mathcal{R} as sets instead of tuples when the order of their elements is understood or irrelevant.

Let e_1, \dots, e_m be the standard basis of F^m , c_j the j th column of $[Q]$, and for $p \in P$, define

$$\begin{aligned}\mathcal{G}_p &= \{e_i \mid \text{gr}(g_i) \leq p\}, \\ \mathcal{R}_p &= \{c_j \mid \text{gr}(r_j) \leq p\}.\end{aligned}$$

Then Q induces a module M^Q with

$$M_p^Q = \text{span}(\mathcal{G}_p) / \text{span}(\mathcal{R}_p)$$

and $M_{p \rightarrow q}^Q$ induced by the inclusion $\mathcal{G}_p \rightarrow \mathcal{G}_q$. We say that Q is a *presentation* of a module M if M is isomorphic to M^Q , and in this case, we say that M is finitely presented. See Fig. 1.

We say that a presentation $\partial = ([\partial], \mathcal{G}_\partial, \mathcal{R}_\partial, \text{gr})$ is a *permutation* of Q if there are permutations σ and τ such that $[\partial]_{i,j} = [Q]_{\sigma(i), \tau(j)}$ for all i and j , and $\mathcal{G}_\partial = (g_{\sigma(1)}, \dots, g_{\sigma(m)})$ and $\mathcal{R}_\partial = (r_{\tau(1)}, \dots, r_{\tau(m')})$. That is, we permute the rows and columns of $[Q]$ and permute the indexing elements correspondingly. If Q is a presentation of a module M , then so is ∂ .

A presentation $([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ of a 1-parameter module is *ordered* if the rows of $[Q]$ are nondecreasing with respect to $\text{gr}(g)$ for $g \in \mathcal{G}$ and the same holds for the columns and $r \in \mathcal{R}$. Observe

that for any presentation of a 1-parameter module, there always exists a permutation of it that is ordered.

From presentations to barcodes It is a classical result in persistence theory that 1-parameter modules admit an invariant up to isomorphism called a barcode:

Theorem 3.3 ([11, Theorem 1.1]). *Let M be a finitely presented 1-parameter module. Then there is a unique multiset B of intervals such that*

$$M \simeq \bigoplus_{I \in B} \mathbb{I}^I.$$

We call B the barcode of M . Isomorphic modules have the same barcode. The original statement of the theorem is for a larger set of modules, but we will not need the stronger version. In our case, where the modules are finitely presented, all the intervals are of the form $[a, b)$ with $a \in \mathbb{R}$ and $b \in \mathbb{R} \cup \{\infty\}$. (Equivalently, one can represent the barcode as a multiset of points in $\mathbb{R} \times (\mathbb{R} \cup \{\infty\})$, with a point (a, b) for each interval $[a, b)$ in the barcode. This alternative description is called a *persistence diagram*, and is used instead of barcodes in some of our references.)

For reasons that will become clear, we need to understand how these barcodes are computed. We give a quick summary following [10, Sec. 2; Computation]. First, some terminology: A *column operation* on a matrix is the addition of a multiple of the i^{th} column to the j^{th} column for some $i < j$, a *pivot* is a bottommost nonzero element of a (nonzero) column, and *reduced form* means that the pivots are all in different rows. In the standard algorithm to compute the barcode [34], one takes as input an ordered presentation $Q = ([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ of a 1-parameter module M and performs column operations on $[Q]$ until the matrix is in reduced form.

One can summarize this column reduction by $R = [Q]V$, or equivalently $RV^{-1} = [Q]$, where R is the reduced matrix and V is the upper-triangular matrix expressing the column operations we performed on $[Q]$ to get R . We call (R, V^{-1}) an *RU-decomposition* of Q , also in the case when Q is unordered and/or a presentation of a 2-parameter module. We consider the rows and columns of R to be indexed by \mathcal{G} and \mathcal{R} and view this information as part of the RU-decomposition.

We define the *barcode pairing* associated to Q as follows: It contains each pair (g, r) of a generator g and a relation r such that $\text{gr}(g) \neq \text{gr}(r)$ and the element in the row of g and the column of r in R is a pivot, as well as a pair (g, ∞) for each g whose row in R does not contain a pivot. We view ∞ as a dummy relation at infinity, and define $\text{gr}(\infty) = \infty$. The barcode of M turns out to be the multiset of intervals $[\text{gr}(g), \text{gr}(r))$ for the pairs (g, r) in the barcode pairing.

The RU-decomposition and barcode pairing are not invariants of M (as any module has several presentations), but, as stated by Theorem 3.3, the barcode is.

The bottleneck distance We will define the bottleneck distance on 1-parameter modules in terms of perfect matchings in graphs built on barcode pairings of the modules. The bottleneck distance is usually defined in terms of barcodes, not barcode pairings, but there are technical reasons why we want to keep track of the pairs of generators and relations instead of throwing that information away and just consider the intervals in the barcode. The end result is the same, as our definition is equivalent to the usual one.

Let M and N be 1-parameter modules with barcode pairings B and B' , respectively. Given $\lambda \in [0, \infty]$, define an undirected bipartite graph $G = G_{B, B'}^\lambda$ as follows: The vertex set of G is $(B \cup \overline{B'}) \sqcup (B' \cup \overline{B})$, where \overline{B} has an element $\overline{(g, r)}$ for each $(g, r) \in B$, and $\overline{B'}$ is defined similarly. There is an edge between $(g, r) \in B$ and $(g', r') \in B'$ if

$$|\text{gr}(g) - \text{gr}(g')| \leq \lambda \quad \text{and} \quad |\text{gr}(r) - \text{gr}(r')| \leq \lambda.$$

Let $U \subset B \sqcup B'$ be the set of (g, r) such that

$$|\text{gr}(g) - \text{gr}(r)| \leq 2\lambda.$$

For each $(g, r) \in U$, there is an edge between (g, r) and $\overline{(g, r)}$. Lastly, there is an edge between each $\overline{(g, r)} \in \overline{B}$ and $\overline{(g', r')} \in \overline{B'}$.

Observe that the isomorphism class of G only depends on the barcodes of M and N , so the existence of a perfect matching is independent of which barcode pairings we use. This justifies the following definition.

Definition 3.4. Let B and B' be barcode pairings of 1-parameter modules M and N . Define the bottleneck distance between M and N as

$$d_{\mathcal{B}}(M, N) = \min\{\lambda \mid G_{B, B'}^{\lambda} \text{ allows a perfect matching}\}.$$

The set we take the minimum of is never empty, as $G_{B, B'}^{\infty}$ is the complete bipartite graph. That the minimum is well-defined follows from the observation that if an edge is present in $G_{B, B'}^{\lambda}$ for all $\lambda \in (a, \infty]$, then it is also present in $G_{B, B'}^a$. This, in addition to the fact that as λ increases, the vertex set of G stays the same and no edges are removed, implies the following lemma:

Lemma 3.5. $d_{\mathcal{B}}(M, N) \leq \lambda$ holds if and only if $G_{B, B'}^{\lambda}$ allows a perfect matching.

We remark, leaving the proof to the reader, that the existence of a perfect matching in $G_{B, B'}^{\lambda}$ is equivalent to the existence of a matching in the subgraph induced by $B \sqcup B'$ covering (a superset of) $(B \sqcup B') \setminus U$. Intuitively, we are allowed to match intervals in the barcodes that are similar, and we do not have to match short (i.e. “insignificant”) intervals.

The matching distance We define a slice as a line $s \subset \mathbb{R}^2$ with positive slope. For $p \in \mathbb{R}^2$, let p_x and p_y be the x - and y -coordinates of p , respectively; that is, $p = (p_x, p_y)$.

Definition 3.6. Let M be a 2-parameter module and s a slice with slope at most one. We define M^s as the 1-parameter module with $M_r^s = M_p$, where p is the unique point on s such that $p_y = r$, and $M_{p_y \rightarrow q_y}^s = M_{p \rightarrow q}$ for $p \leq q \in s$.

If s has slope more than one, we define M^s as the 1-parameter module with $M_{p_x}^s = M_p$ for all $p \in s$, and morphisms defined by $M_{p_x \rightarrow q_x}^s = M_{p \rightarrow q}$ for $p \leq q \in s$.

Thus, if the slope of s is at most one, we can visualize M^s as being the restriction of M to s projected to the y -axis. If the slope is more than one, we can visualize M^s as being the restriction of M to s projected to the x -axis.

Definition 3.7. We define the matching distance as

$$d_{\mathcal{M}}(M, N) = \sup_{\text{slice } s} d_{\mathcal{B}}(M^s, N^s)$$

for finitely presented 2-parameter modules M and N .

In view of Lemma 4.1 below, M^s and N^s are finitely presented, so $d_{\mathcal{B}}(M^s, N^s)$ is indeed well-defined.

We will find it convenient to avoid the case distinction between slope less than and greater than one. Let

$$d_{\mathcal{M}}^{\leq 1}(M, N) = \sup_{\text{slice } s \text{ with slope } \leq 1} d_{\mathcal{B}}(M^s, N^s);$$




Figure 2: Points in red above, on and below a slice s and the pushes of the points on the real line in blue.

$d_{\mathcal{M}}^{\geq 1}(M, N)$ is defined similarly. Clearly,

$$d_{\mathcal{M}}(M, N) = \max \left\{ d_{\mathcal{M}}^{\leq 1}(M, N), d_{\mathcal{M}}^{\geq 1}(M, N) \right\},$$

so to compute $d_{\mathcal{M}}(M, N)$, it is sufficient to compute $d_{\mathcal{M}}^{\leq 1}(M, N)$ and $d_{\mathcal{M}}^{\geq 1}(M, N)$. Since computing $d_{\mathcal{M}}^{\leq 1}(M, N)$ and $d_{\mathcal{M}}^{\geq 1}(M, N)$ are completely symmetric problems (simply switch x - and y -coordinates), we can assume without loss of generality that $d_{\mathcal{M}}^{\leq 1}(M, N) \geq d_{\mathcal{M}}^{\geq 1}(M, N)$ for the purposes of asymptotic complexity, so that $d_{\mathcal{M}}^{\leq 1}(M, N) = d_{\mathcal{M}}(M, N)$. We will use this assumption throughout the rest of the paper.

4 Deciding if $d_{\mathcal{M}} \leq \lambda$

Our goal is to compute $d_{\mathcal{M}}(M, N)$, but first we consider an easier problem: given $\lambda \geq 0$, decide if $d_{\mathcal{M}}(M, N) \leq \lambda$. We explain how to do this in $O(n^5 \log n)$. In the next section, we show how to use this to compute $d_{\mathcal{M}}(M, N)$ exactly.

Fix $\lambda \in [0, \infty)$. Since we assume $d_{\mathcal{M}}^{\leq 1}(M, N) = d_{\mathcal{M}}(M, N)$, every slice in this section will be assumed to have slope at most one. To decide if $d_{\mathcal{M}}(M, N) \leq \lambda$ for modules M and N , we need to decide if $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ for every slice s . Since there are infinitely many slices, it is not obvious how to do this in finite time, let alone efficiently. To decide if $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ for a given slice s , the most obvious approach is to find presentations of M^s and N^s , then barcode pairings B and B' , construct the graph $G_{B, B'}^{\lambda}$, decide if it has a perfect matching and apply Lemma 3.5. Our strategy will be to group the slices into a finite number of sets such that we only need to go through this procedure for one slice in every set. We begin by explaining how a presentation Q of a 2-parameter module M induces a canonical presentation of M^s for every slice s .

Induced presentations Let $p \in \mathbb{R}^2$, let s be a slice, and let q be the least point on s that is greater than or equal to p . We define $\text{push}_s p$ to be q_y . If p is on or above s , then $\text{push}_s p = p_y$; in particular, if $p \in s$ and M is a 2-parameter module, then $M_{\text{push}_s p}^s = M_p$. If p is below s and $(p_x, r) \in s$, then $\text{push}_s p = r$. See Fig. 2. Let Q be a presentation of a 2-parameter module

with indexing sets $\mathcal{G} = (g_1, \dots, g_m)$ and $\mathcal{R} = (r_1, \dots, r_{m'})$. Let $\text{push}_s(Q)$ be the 1-parameter presentation with $[\text{push}_s(Q)] = [Q]$, indexing sets

$$\text{push}_s(\mathcal{G}) = (\text{push}_s(g_1), \dots, \text{push}_s(g_m)), \quad \text{push}_s(\mathcal{R}) = (\text{push}_s(r_1), \dots, \text{push}_s(r_{m'}))$$

(here $\text{push}_s(g_i)$ and $\text{push}_s(r_j)$ are to be understood simply as (labels for) the i^{th} generator and j^{th} relation of $\text{push}_s(Q)$) and

$$\text{gr}(\text{push}_s(g_i)) = \text{push}_s(\text{gr}(g_i)), \quad \text{gr}(\text{push}_s(r_j)) = \text{push}_s(\text{gr}(r_j))$$

for all i and j . This is indeed a valid presentation, as

$$\text{gr}(p) \leq \text{gr}(q) \implies \text{gr}(\text{push}_s(p)) \leq \text{gr}(\text{push}_s(q))$$

for all $p, q \in \mathbb{R}^2$.

Lemma 4.1. *Let s be a slice. If $Q = ([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ is a presentation of a 2-parameter module M , then $\text{push}_s(Q)$ is a presentation of M^s .*

Proof. For $h \in \mathcal{G} \cup \mathcal{R}$ and $p \in s$, we have $\text{gr}(h) \leq p$ if and only if $\text{push}_s(\text{gr}(h)) \leq p_y$. We get

$$\begin{aligned} M_{p_y}^{\text{push}_s(Q)} &= \text{span}\{e_i \mid \text{push}_s(\text{gr}(g_i)) \leq p_y\} / \text{span}\{c_j \mid \text{push}_s(\text{gr}(r_j)) \leq p_y\} \\ &= \text{span}\{e_i \mid \text{gr}(g_i) \leq p\} / \text{span}\{c_j \mid \text{gr}(r_j) \leq p\} \\ &= \text{span}(\mathcal{G}_p) / \text{span}(\mathcal{R}_p) \\ &= M_p^Q. \end{aligned}$$

The morphisms $M_{p_y \rightarrow q_y}^{\text{push}_s(Q)}$ and $M_{p \rightarrow q}^Q$ are induced in the same canonical way, so they are equal. Since $M_{p_y}^s = M_p$, $M_{p_y \rightarrow q_y}^s = M_{p \rightarrow q}$ and $M \simeq M^Q$ all hold, we get $M^{\text{push}_s(Q)} \simeq M^s$. Thus, $\text{push}_s(Q)$ is a presentation of M^s . \square

We say that $\text{push}_s(Q)$ is the presentation of M^s induced by Q .

RU-decompositions and barcode pairings at slices Let $Q = ([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ be a presentation of M , and let $\partial = ([\partial], \mathcal{G}_\partial, \mathcal{R}_\partial, \text{gr})$ be a permutation of Q such that $\text{push}_s(\partial)$ is an ordered presentation of M^s . Such a permutation always exists, as it is always possible to sort $\text{push}_s(\mathcal{G})$ and $\text{push}_s(\mathcal{R})$ in nondecreasing order. If (R, U) is an RU-decomposition of ∂ , we say that it is an *RU-decomposition of Q at s* . We say that the set B of $(g, r) \in \mathcal{G}_\partial \times \mathcal{R}_\partial = \mathcal{G} \times \mathcal{R}$ obtained from the pivots of R is a *barcode pairing of Q at s* , and let $\text{push}_s(B)$ be the induced barcode pairing of M^s .

Computing the barcode at every slice The idea of defining RU-decompositions and barcode pairings of a fixed presentation at different slices is that we want to group many slices together and use the same RU-decompositions, barcode pairings and (as we will explain later) bipartite graph and perfect matching for all of them. We will use the following geometric representation exploiting point-line duality to make it easier to work with the set of slices.

Definition 4.2. *Let $\omega = (0, 1] \times \mathbb{R}$. We say that $(a, b) \in \omega$ represents the slice given by $y = ax + b$.*

We say that a point in ω and the slice it represents are *dual* to each other. One can see that the set of slices through a point p is represented by the line $y = -p_x x + p_y$. We call this the line *dual to p* . We use the notation $-^*$ for duality – for instance, p^* is the line given by $y = -p_x x + p_y$, and $(p^*)^* = p$. We say that p lies in the *primal plane*, and p^* in the *dual plane*. We usually abuse notation and identify a slice with the point in ω that represents it, but we sometimes write s^* if we want to emphasize that we treat it as a point in the dual plane and not as a line in the primal plane. The vertical line $x = r$ contains exactly the (points representing) slices with slope r . We remind the reader of some standard facts: If p is a point and s a slice, then $p \in s$ if and only if $s^* \in p^*$, and p is above (below) s if and only if p^* is above (below) s^* .

For the remainder of the section, we fix two 2-parameter modules M and N along with presentations $Q = ([Q], \mathcal{G}, \mathcal{R}, \text{gr})$ and $Q' = ([Q'], \mathcal{G}', \mathcal{R}', \text{gr})$, respectively. Let $\mathfrak{G} = \mathcal{G} \cup \mathcal{R} \cup \mathcal{G}' \cup \mathcal{R}'$, and let $n = |\mathfrak{G}|$. Our next step is to define a line arrangement in ω and show that in a certain precise sense, in each region of the arrangement, RU-decompositions and thus barcode decompositions stay constant. We also show that we can update the RU-decompositions efficiently as we cross a line in the arrangement. Later, as we traverse the arrangement and decide if $d_{\mathcal{B}}(M^s, N^s)$ at every slice s , this will come in handy.

For $p, p' \in \mathbb{R}^2$, let $p \vee p'$ denote the join of p and p' , i.e., the least element that is greater than or equal to p and p' .

Definition 4.3. Define the following set of lines in ω :

$$\mathcal{L} = \{(\text{gr}(h) \vee \text{gr}(h'))^* \mid h, h' \in \mathfrak{G}\}.$$

This is the barcode template described in Section 2. There are $O(n^2)$ lines in \mathcal{L} , so the line arrangement has $O(n^4)$ open regions [12, Theorem 8.4 (iii)].

For $h \in \mathfrak{G}$, let $h_x = \text{gr}(h)_x$ and $h_y = \text{gr}(h)_y$. When working with \mathcal{L} and the line arrangement \mathcal{T}_λ defined below, we run into the technical issue of coinciding lines. In \mathcal{L} , this happens for instance if there are $h \neq h' \in \mathfrak{G}$ with $h_x = h'_x$, as in that case, we get $\text{gr}(h) \vee \text{gr}(h') = \text{gr}(h') \vee \text{gr}(h')$ for any $h'' \in \mathfrak{G}$ with $h''_y \geq h_y, h'_y$. The only reason why this might be a problem for us is that we will traverse these arrangements, and the complexity of our algorithm depends on how many times we have to cross a line. When lines coincide, we have to count the number of times we cross them with multiplicity, as we need to perform certain operations for each coinciding line. To avoid distracting technical details, we will gloss over this issue and assume that the lines we consider have multiplicity one in what follows. Then, in the proof of Theorem 4.10, we will explain how to handle lines with multiplicity greater than one.

Lemma 4.4. Let S be an open region in \mathcal{L} , let $s \in S$, and let (R, U) be an RU-decomposition of Q at s .

- (i) For any $s' \in S$, (R, U) is an RU-decomposition of Q at s' .
- (ii) If S' is another open region in \mathcal{L} that is separated from S by a single line, then for any $s' \in S'$, an RU-decomposition of Q at s' can be obtained from (R, U) in $O(n)$.

The analogous statements hold if Q is replaced by Q' .

Proof. (i): Any slice s induces a total preorder \leq_s on \mathfrak{G} by letting $h \leq_s h'$ if $\text{push}_s(\text{gr}(h)) \leq \text{push}_s(\text{gr}(h'))$. If (the points represented by) s and s' are both in the same open half-plane defined by $(h \vee h')^*$, then $h \leq_s h'$ if and only if $h \leq_{s'} h'$: If $h_x < h'_x$ and $h_y > h'_y$, then $h <_s h'$ if and only if s is above $h \vee h'$, and $h >_s h'$ if and only if s is below $h \vee h'$. These correspond to s^* lying above $(h \vee h')^*$ and below $(h \vee h')^*$, respectively. We leave the other cases to the reader.

This means that \leq_s only depends on where s is in relation to each line in \mathcal{L} , i.e., whether it is on, above or below it. Thus, the preorder is constant in S . But if ∂ is a permutation of Q , whether $\text{push}_{s'}(\partial)$ is an ordered presentation of $M^{s'}$ depends only on $\leq_{s'}$. Therefore, (R, U) is an RU-decomposition of Q at s if and only if it is an RU-decomposition of Q at s' .

(ii): Suppose the only line in \mathcal{L} separating S and S' is $(\text{gr}(h) \vee \text{gr}(h'))^*$, and let $s \in S$ and $s' \in S'$. We have just shown that $j \leq_s j'$ and $j \leq_{s'} j'$ are equivalent for all $\{j, j'\} \neq \{h, h'\}$. Thus, only the relations between h and h' can change between the two regions. Suppose ∂ is a permutation of Q and $\text{push}_s(\partial)$ is an ordered presentation of M^s . In this case, $\text{push}_{s'}(\partial)$ is not necessarily an ordered presentation of $M^{s'}$. But if it is not, we know that we can make the rows and columns of $[\text{push}_{s'}(\partial)]$ nondecreasing by either switching two rows or two columns, namely those indexed by h and h' . Suppose ∂' is the permutation of ∂ (and thus of Q) we get by switching h and h' , and their corresponding rows/columns. By [10, Sec. 3], we can find an RU-decomposition of ∂' in $O(n)$, since we already have an RU-decomposition of ∂ and we only need to make a single transposition. Because $[\text{push}_{s'}(\partial')]$ is nondecreasing, an RU-decomposition of ∂' is by definition an RU-decomposition of Q at s' .

The proofs for Q' are exactly the same. \square

Since we can keep track of how the barcode pairing changes as the RU-decomposition changes, Lemma 4.4 allows us to efficiently update the barcode pairing as we move from a region to a neighboring region.

Deciding if $d_{\mathcal{B}} \leq \lambda$ at every slice Now that we have dealt with the RU-decompositions, we want to refine the line arrangement \mathcal{L} so that deciding if $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ holds for all slices s in a given region boils down to checking if a single graph has a perfect matching.

Let s be a slice, and let $B \subset \mathcal{G} \times \mathcal{R}$ and $B' \subset \mathcal{G}' \times \mathcal{R}'$ be barcode pairings at s (so $\text{push}_s(B)$ and $\text{push}_s(B')$ are barcode pairings of M^s and N^s). Recall from Lemma 3.5 that $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ if and only if there is a perfect matching in the graph $G_{C, C'}^\lambda$, where C and C' are arbitrary barcode pairings of M^s and N^s ; in particular we can choose $C = \text{push}_s(B)$ and $C' = \text{push}_s(B')$. Define $G_{B, B'}^s$ (we write $G_{B, B'}^{s, \lambda}$ if we want to emphasize that it depends on λ) as being equal to $G_{\text{push}_s(B), \text{push}_s(B')}^\lambda$, except that each vertex $(\text{push}_s(g), \text{push}_s(r)) \in \text{push}_s(B) \cup \text{push}_s(B')$ is replaced by (g, r) , and $(\text{push}_s(g), \text{push}_s(r)) \in \text{push}_s(B) \cup \text{push}_s(B')$ by (g, r) . The graphs $G_{B, B'}^s$ and $G_{\text{push}_s(B), \text{push}_s(B')}^\lambda$ are clearly isomorphic, so one has a perfect matching if and only if the other has. Thus, $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ if and only if there is a perfect matching in $G_{B, B'}^s$.

For $h, h' \in \mathfrak{G}$, define

$$d^s(h, h') = |\text{gr}(\text{push}_s(h)) - \text{gr}(\text{push}_s(h'))|.$$

Observe that in $G_{B, B'}^s$, there is an edge between $(g, r) \in B$ and $(g', r') \in B'$ if and only if $d^s(g, g'), d^s(r, r') \leq \lambda$. Similarly, there is an edge between (g, r) and $\overline{(g, r)}$ if and only if $d^s(g, r) \leq 2\lambda$.

We now construct a line arrangement \mathcal{T}_λ and show that for each open region S in \mathcal{T}_λ , there are B and B' such that $G_{B, B'}^s$ is well-defined and $G_{B, B'}^s = G_{B, B'}^{s'}$ for all $s, s' \in S$.

Definition 4.5. Define the following sets of lines in ω :

$$\begin{aligned} \mathcal{P}_\lambda &= \{(h_x, h'_y + i\lambda)^* \mid h, h' \in \mathfrak{G}, i \in \{-2, -1, 1, 2\}\}, \\ \mathcal{S}_\lambda &= \left\{ \left\{ \frac{i\lambda}{|h_x - h'_x|} \right\} \times \mathbb{R} \mid h, h' \in \mathfrak{G}, i \in \{1, 2\} \right\}, \\ \mathcal{T}_\lambda &= \mathcal{P}_\lambda \cup \mathcal{S}_\lambda \cup \mathcal{L}. \end{aligned}$$

For $h, h' \in \mathfrak{G}$, let $\mathcal{T}_\lambda^{h, h'}$ be the set of all lines in \mathcal{T}_λ involving h and h' , i.e.,

$$\begin{aligned} \mathcal{T}_\lambda^{h, h'} = & \{(h_x, h'_y + i\lambda)^* \mid i \in \{-2, -1, 1, 2\}\} \\ & \cup \{(h'_x, h_y + i\lambda)^* \mid i \in \{-2, -1, 1, 2\}\} \\ & \cup \left\{ (\text{gr}(h) \vee \text{gr}(h'))^*, \left\{ \frac{\lambda}{|h_x - h'_x|} \right\} \times \mathbb{R}, \left\{ \frac{2\lambda}{|h_x - h'_x|} \right\} \times \mathbb{R} \right\} \end{aligned}$$

if $h_x \neq h'_x$, and the same except the last two elements if $h_x = h'_x$.

Observe that $\mathcal{T}_\lambda = \bigcup_{h, h' \in \mathfrak{G}} \mathcal{T}_\lambda^{h, h'}$.

Definition 4.6. For $h, h' \in \mathfrak{G}$, let $d^-(h, h') : \omega \rightarrow [0, \infty)$ be the function defined by $s \mapsto d^s(h, h')$. Define an equivalence relation $\sim_{h, h'}$ on ω by letting the equivalence classes be $d^-(h, h')^{-1}(I)$ for $I \in \{[0, \lambda), \{\lambda\}, (\lambda, 2\lambda), \{2\lambda\}, (2\lambda, \infty)\}$.

Thus, $s \sim_{h, h'} s'$ if and only if $d^s(h, h')$ and $d^{s'}(h, h')$ are contained in the same of the sets $[0, \lambda)$, $\{\lambda\}$, $(\lambda, 2\lambda)$, $\{2\lambda\}$, $(2\lambda, \infty)$.

Lemma 4.7. Let $h, h' \in \mathfrak{G}$. If $s, s' \in \omega$ lie in the same open region of $\mathcal{T}_\lambda^{h, h'}$, then $s \sim_{h, h'} s'$.

Proof. Fix h and h' . It is straightforward to check that the function $\omega \rightarrow \mathbb{R}$ given by $r \mapsto \text{gr}(\text{push}_r(h))$ is continuous. Thus, also $d^-(h, h')$ is continuous.

Assume that s and s' lie in the same open region S of $\mathcal{T}_\lambda^{h, h'}$. We will show that if $\lambda \in d^-(h, h')(S)$, then $d^-(h, h')(S) = \{\lambda\}$, and if $2\lambda \in d^-(h, h')(S)$, then $d^-(h, h')(S) = \{2\lambda\}$. By connectivity of S and continuity of $d^-(h, h')$, $d^-(h, h')(S)$ is connected, so this implies that $d^-(h, h')(S)$ is contained in one of the sets $[0, \lambda)$, $\{\lambda\}$, $(\lambda, 2\lambda)$, $\{2\lambda\}$, $(2\lambda, \infty)$, which proves the lemma.

Suppose $d^r(h, h') = \lambda$ for some $r \in S$. We will only consider this case; the argument for $d^r(h, h') = 2\lambda$ is practically the same. First assume that $\text{gr}(h)$ and $\text{gr}(h')$ are above or on r^* . (As the arguments that follow are heavy on geometry and point-line duality, we will be careful about differentiating between the point r in ω and the slice r^* in the primal plane.) Then $d^r(h, h') = |h_y - h'_y|$, so $|h_y - h'_y| = \lambda$. In this case, $\text{gr}(h) = (h_x, h'_y \pm \lambda)$ and $\text{gr}(h') = (h'_x, h_y \pm \lambda)$, so $\text{gr}(h)^*, \text{gr}(h')^* \in \mathcal{P}_\lambda$. Since $\text{gr}(h)$ and $\text{gr}(h')$ are above or on r^* in the primal plane, r is below or on $\text{gr}(h)^*$ and $\text{gr}(h')^*$ in ω . Since $\text{gr}(h)^*, \text{gr}(h')^* \in \mathcal{P}_\lambda$, this implies that all elements of S are below $\text{gr}(h)^*$ and $\text{gr}(h')^*$ in ω , and thus the slices they represent lie below $\text{gr}(h)$ and $\text{gr}(h')$ in the primal plane. Thus, $d^{r'}(h, h') = |h_y - h'_y| = \lambda$ holds for all $r' \in S$.

Next, assume that $\text{gr}(h)$ and $\text{gr}(h')$ are below or on r^* . Then there are two points

$$(h_x, a), (h'_x, a \pm \lambda) \in r^*,$$

since $d^r(h, h') = \lambda$. In this case, r^* has slope $\frac{\lambda}{|h_x - h'_x|}$ and thus r lies on the line $x = \frac{\lambda}{|h_x - h'_x|}$ in ω , which is one of the lines in $\mathcal{T}_\lambda^{h, h'}$. (Consider $r = s_3$ in Fig. 3.) This is a contradiction, as we assumed that r is in an open region S , which does not intersect any line in $\mathcal{T}_\lambda^{h, h'}$.

Lastly, assume that $\text{gr}(h)$ and $\text{gr}(h')$ are on opposite sides of r^* ; say $\text{gr}(h)$ is above and $\text{gr}(h')$ below. Then $\text{gr}(\text{push}_r(h)) = h_y$. This means that $\text{gr}(\text{push}_r(h')) = h_y \pm \lambda$, so there is a point $(h'_x, h_y \pm \lambda) \in r^*$. (Consider $r = s_1$ or $r = s_2$ in Fig. 3.) Thus, $r \in (h'_x, h_y \pm \lambda)^* \in \mathcal{T}_\lambda^{h, h'}$, which is again a contradiction. \square

Lemma 4.8. Let S be an open region in \mathcal{T}_λ and $s \in S$. Let $B \subset \mathcal{G} \times \mathcal{R}$ and $B' \subset \mathcal{G}' \times \mathcal{R}'$ be barcode pairings at $s \in S$. Then for any $s' \in S$, $G_{B, B'}^{s'}$ is well-defined, and $G_{B, B'}^s = G_{B, B'}^{s'}$.




Figure 3: On the left: the grades of h and h' in the primal plane, and three slices dual to $s_1, s_2, s_3 \in \omega$ with $d^{s_i}(h, h') = \lambda$. s_1^* and s_2^* pass through $(h'_x, h_y \pm \lambda)$, and s_3^* has slope $\frac{\lambda}{h'_x - h_x}$. On the right: s_1, s_2, s_3 all lie on lines in \mathcal{T}_λ . The lines $(h'_x, h_y - \lambda)^*$ and $(h'_x, h_y + \lambda)^*$ are given by the equations $y = -h'_x x + (h_y - \lambda)$ and $y = -h'_x x + (h_y + \lambda)$, respectively.

Proof. Since $\mathcal{L} \subset \mathcal{T}_\lambda$, S is contained in an open region of \mathcal{L} . By Lemma 4.4 (i), B and B' are then barcode pairings at s' , so $G_{B, B'}^{s'}$ is well-defined.

Let $h, h' \in \mathfrak{G}$. Since s, s' lie in the same open region of \mathcal{T}_λ , they lie in the same open region of $\mathcal{T}_\lambda^{h, h'}$. By Lemma 4.7, we get $s \sim_{h, h'} s'$. But whether there is an edge between two vertices of $G_{B, B'}^r$ depends only on whether inequalities of the form $d^r(h, h') \leq \lambda$ and $d^r(h, h') \leq 2\lambda$ hold, and since $s \sim_{h, h'} s'$, these inequalities hold for $r = s$ if and only if they hold for $r = s'$. \square

With this in mind, we define $G_{B, B'}^S$ as $G_{B, B'}^{s'}$ for any $s \in S$ and say that B and B' are barcode pairings in S . Similarly, we say that an RU-decomposition at $s \in S$ is an RU-decomposition in S , as justified by Lemma 4.4 (i).

Lemma 4.9. *Let R and S be regions in \mathcal{T}_λ that are separated by a single line, and suppose we are given RU-decompositions with associated barcode pairings $B \subset \mathcal{G} \times \mathcal{R}$ and $B' \subset \mathcal{G}' \times \mathcal{R}'$ in R and a perfect matching in $G_{B, B'}^R$. Then we can find RU-decompositions with associated barcode pairings $B_S \subset \mathcal{G} \times \mathcal{R}$ and $B'_S \subset \mathcal{G}' \times \mathcal{R}'$ in S and either find a perfect matching in G_{B_S, B'_S}^S or determine that a perfect matching does not exist in $O(n \log n)$.*

Proof. We first assume that the line ℓ separating R and S is in \mathcal{L} . Then we can update the RU-decompositions and thus B and B' in $O(n)$ using Lemma 4.4 (ii). By [10, Fig. 4, left]³, either B and B' do not change, or one of the following happens:

- a pair (g, r) appears or disappears in B or B' (in this case, the RU-decomposition does not necessarily change, but (g, r) gives rise to an empty interval in R and a nonempty interval in S , or vice versa),

³The setup in [10] is a little different than ours. In their framework, an element can act both as a generator and a relation. Since this is impossible for us, we can ignore the right part of Fig. 4.

- two pairs (g, r) and (g', r') are replaced by (g', r) and (g, r') in B or B' .

In both cases, the graph $G_{B, B'}^R$ only has to be modified “locally” to get a graph G_{B_S, B'_S}^S : At most two vertices are added, removed or replaced, and only the edges adjacent to the modified vertices are changed. (Actually, in the second bullet point, the graphs are isomorphic. We will not need this fact, so we do not prove it.) The matching in $G_{B, B'}^R$ therefore gives us a matching in G_{B_S, B'_S}^S covering all but at most four vertices.

Now assume that ℓ is in $\mathcal{T}_\lambda \setminus \mathcal{L} = \mathcal{P}_\lambda \cup \mathcal{S}_\lambda$. Then by Lemma 4.4(i), the RU-decompositions and barcode pairings are valid also in S , as R and S are contained in the same region of \mathcal{L} . The only relevant change crossing ℓ can make is that there might be two elements $h, h' \in \mathfrak{G}$ such that $d^r(h, h') \leq \lambda$ and $d^s(h, h') > \lambda$ or vice versa for $r \in R$ and $s \in S$, or the same statement holds for 2λ instead of λ . The only change this can make to the graph is that an edge between the elements in $B \cup B'$ involving h and h' has to be added or removed, or, in the case $(h, h') \in B \cup B'$, an edge between (h, h') and (\overline{h}, h') has to be added or removed. In either case, the matching in $G_{B, B'}^R$ induces a matching in $G_{B, B'}^S$ (which is well-defined, as B and B' are valid in S) covering all but at most two vertices.

Thus, we have reduced the problem to finding a perfect matching, or decide that a perfect matching does not exist, in a graph G_{B_S, B'_S}^S where we already have a matching covering all but at most four vertices. By [22, Theorem 3.1], which builds on [16], the bottleneck distance can be computed in $O(n^{1.5} \log n)$. But a factor of $O(n^{0.5})$ comes from the number of times one has to augment the matching, and after each augmentation one has a matching with at least one more edge. Since we only need to add a constant number of edges, we can replace this factor by a constant, so the complexity of constructing a perfect matching or determining that it does not exist is only $O(n \log n)$. \square

Theorem 4.10. *For any $\lambda \in [0, \infty)$, we can decide if $d_{\mathcal{M}}(M, N) \leq \lambda$ in $O(n^5 \log n)$.*

Proof. By [21, Lemma 3], the function f sending a slice s to $d_{\mathcal{B}}(M^s, N^s)$ is continuous. Since the union U of open regions of \mathcal{T}_λ is dense in ω , we have $f(U) \subset [0, \lambda]$ if and only if $f(\omega) \subset [0, \lambda]$. Thus, to find out if $d_{\mathcal{M}} \leq \lambda$, it is enough to check if $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ for all s in open regions. By Lemma 3.5 and Lemma 4.8, for any open region S , $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ is equivalent to $d_{\mathcal{B}}(M^{s'}, N^{s'}) \leq \lambda$ for all $s, s' \in S$. Thus, it is enough to check if $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ for one slice s in each open region.

We construct a spanning tree T of the dual graph of \mathcal{T}_λ ; that is, the graph with a vertex for each open region and an edge between two open regions if they are separated by a line. We begin by computing RU-decompositions of Q and Q' by column operations in an open region S of \mathcal{T}_λ , finding the associated barcode pairings B and B' , constructing $G_{B, B'}^S$ and either finding a perfect matching or determining that it does not exist with the Hopcroft-Karp algorithm [19]. All of this can be done in $O(n^3)$. If there is no perfect matching, then $d_{\mathcal{B}}(M^s, N^s) > \lambda$ for all $s \in S$, so $d_{\mathcal{M}}(M, N) > \lambda$, and we are done.

Assuming that we found a perfect matching, we walk through T in such a way that we visit all the vertices and do not cross any edge more than twice. When we move along an edge in T , which corresponds to crossing an edge in \mathcal{T}_λ going from an open region S to a neighboring region S' , we update the RU-decompositions, barcode pairings and the perfect matchings. If a perfect matching does not exist, we conclude that $d_{\mathcal{M}}(M, N) > \lambda$ and stop; otherwise, we continue. If we have visited every vertex of T (i.e., every open region of \mathcal{T}_λ) and have never failed to find a perfect matching, we conclude that $d_{\mathcal{M}}(M, N) \leq \lambda$.

If the line between S and S' has multiplicity one, then by Lemma 4.9, the complexity of the updates we make as we go from S to S' is $O(n \log n)$. However, some of the lines we cross might




Figure 4: \mathcal{T}_λ on the left with the multiplicities of the lines shown, and \mathcal{T}'_λ on the right. The arrows show part of a walk visiting the open regions of \mathcal{T}_λ . Crossing a line with multiplicity k in \mathcal{T}_λ corresponds to crossing k lines in \mathcal{T}'_λ .

have multiplicity greater than one. To handle this, we consider a line arrangement \mathcal{T}'_λ where for all k , each line in \mathcal{T}_λ with multiplicity k has been replaced with k different parallel lines, all within some small ϵ of the original line. (To be precise, we pick ϵ such that all open regions of \mathcal{T}_λ contain a closed disk with radius ϵ .) Since \mathcal{T}'_λ has $O(n^2)$ lines, it has $O(n^4)$ open regions and edges between open regions. See Fig. 4. Moving from one region across a line with multiplicity k to a neighboring region in \mathcal{T}_λ then corresponds to moving across k lines with multiplicity one in \mathcal{T}'_λ . If we do not cross any edge in \mathcal{T}_λ more than twice, we do not cross any edge in \mathcal{T}'_λ more than twice, either. Thus, we only perform $O(n^4)$ edge crossings in \mathcal{T}'_λ as we walk through \mathcal{T}_λ . Each edge crossing in \mathcal{T}'_λ costs $O(n \log n)$ by Lemma 4.9, as we only need to perform updates associated to a single line. Thus, the total running time is $O(n^5 \log n)$. \square

5 Computing $d_{\mathcal{M}}$

From now on, we use the shorthand $d_{\mathcal{M}}$ for $d_{\mathcal{M}}(M, N)$. Before we explain the algorithm to compute $d_{\mathcal{M}}$, we deal with the special cases $d_{\mathcal{M}} = 0$ and $d_{\mathcal{M}} = \infty$. Let $T(n)$ be an upper bound for the complexity of deciding if $d_{\mathcal{M}} \leq \lambda$ for an arbitrary λ . (In the last section we proved that we can take $T(n)$ to be $O(n^5 \log n)$.) First, we check if $d_{\mathcal{M}} \leq 0$ in time $T(n)$. If the answer is yes, we know that $d_{\mathcal{M}} = 0$. To find out if $d_{\mathcal{M}} = \infty$, we check if $d_{\mathcal{B}}(M^s, N^s) = \infty$ for one slice s in $O(n^3)$. These are equivalent: One can check that $d^s(h, h') - d^{s'}(h, h')$ is bounded above by $\max\{|h_x - h'_x|, |h_y - h'_y|\}$ for any s, s' in the same region of \mathcal{L} . One can then use Lemma 4.4 (i) to prove that $|d_{\mathcal{B}}(M^s, N^s) - d_{\mathcal{B}}(M^{s'}, N^{s'})|$ is bounded above by

$$L := \max_{h, h' \in \mathfrak{G}} \max\{|h_x - h'_x|, |h_y - h'_y|\}.$$

Since $d_{\mathcal{B}}(M^s, N^s)$ varies continuously with s in all of ω and there are only finitely many regions of \mathcal{L} , $d_{\mathcal{B}}(M^s, N^s)$ has a global (finite) upper bound for all $s \in \omega$ if it is finite for one such s . Considering this, we will assume $d_{\mathcal{M}} \in (0, \infty)$.

The plane arrangement \mathcal{T} We will now move up a dimension from ω and work with $\Omega := \omega \times [0, \infty)$. Elements in Ω are of the form (s, λ) , where s should be viewed as a slice and λ as a

potential value for $d_{\mathcal{M}}$. We now define a plane arrangement \mathcal{T} that can be viewed as the union of the line arrangements \mathcal{T}_λ for all λ with some extra planes added in. We define these planes as subsets of \mathbb{R}^3 . Ω should be understood as a subset of $\mathbb{R}^2 \times \mathbb{R} \simeq \mathbb{R}^3$.

Definition 5.1. *We define the following sets of planes in Ω . Lines of the form p^* should be understood to be the extension of p^* in ω to \mathbb{R}^2 .*

$$\mathcal{P} = \{P_{h,h',i} \mid h, h' \in \mathfrak{G}, i \in \{-2, -1, 0, 1, 2\}\},$$

where

$$P_{h,h',i} = \bigcup_{\lambda \in \mathbb{R}} (h_x, h'_y + i\lambda)^* \times \{\lambda\};$$

$$\mathcal{S} = \{S_{h,h',i} \mid h, h' \in \mathfrak{G}, |h_x - h'_x| \neq 0, i \in \{1, 2\}\} \cup \{S_0, S_1\},$$

where

$$S_{h,h',i} = \left\{ ((a, b), \lambda) \mid a = \frac{i\lambda}{|h_x - h'_x|} \right\},$$

$$S_i = \{i\} \times \mathbb{R}^2;$$

$$\mathcal{H} = \{\mathbb{R}^2 \times \{\lambda\} \mid \exists (h, h') \in \mathfrak{G} \text{ s.t. } |h_y - h'_y| \in \{\lambda, 2\lambda\}\}.$$

Finally, define $\mathcal{T} = \mathcal{P} \cup \mathcal{S} \cup \mathcal{H}$.

Note that $|\mathcal{T}| = O(n^2)$. The reader can verify that the intersection of the planes in $\mathcal{P} \cup \mathcal{S}$ with $\mathbb{R}^2 \times \{\lambda\}$ gives a set of lines that contains all the lines in \mathcal{T}_λ , assuming that we identify $\mathbb{R}^2 \times \{\lambda\}$ with \mathbb{R}^2 in the obvious way. In particular, $(h_1 \vee h_2)^* \in \mathcal{L}$ is of the form $(h_x, h'_y + i\lambda)^*$ for $i = 0$ and some choice of $h, h' \in \{h_1, h_2\}$.

It is clear that the sets in \mathcal{S} and \mathcal{H} are indeed planes, as they are given by linear equations. The same holds for the sets in \mathcal{P} , as $((a, b), \lambda) \in P_{h,h',i}$ is equivalent to $b = -h_x a + h'_y + i\lambda$.

For a set \mathcal{X} of planes, we write $\bigcup \mathcal{X} = \bigcup_{H \in \mathcal{X}} H$. We call a connected component of the complement of $\bigcup \mathcal{T}$ an *open cell*. Letting $h = h'$, we see that $\mathbb{R}^2 \times \{0\} \in \mathcal{H}$. The union $\mathbb{R}^2 \times \{0\} \cup S_0 \cup S_1$ contains the boundary of Ω . Therefore, any open cell is either contained in Ω or contained in $\mathbb{R}^3 \setminus \Omega$.

$d_{\mathcal{M}}$ is attained at a vertex of \mathcal{T} Define a function $f: \Omega \rightarrow \mathbb{R}$ by

$$(s, \lambda) \mapsto \lambda - d_{\mathcal{B}}(M^s, N^s).$$

Thus, $d_{\mathcal{B}}(M^s, N^s) \leq \lambda$ if and only if $f(s, \lambda) \geq 0$.

Proposition 5.2. *f is continuous.*

Proof. As we have already noted, [21, Lemma 3] shows that the function sending a slice s to $d_{\mathcal{B}}(M^s, N^s)$ is continuous. It follows that f is continuous. \square

Lemma 5.3. *Let $(s, \lambda) \in \Omega$. If $f(s, \lambda) = 0$, then $(s, \lambda) \in \bigcup \mathcal{T}$.*

Proof. If $f(s, \lambda) = 0$, there is an edge in $G_{B,B'}^{s,\lambda}$ that is not in $G_{B,B'}^{s,\lambda'}$ for any $\lambda' < \lambda$ (where B and B' are any barcode pairings of M^s and N^s). By the definition of the graphs, this means that there are $h, h' \in \mathfrak{G}$ such that $d^s(h, h') \in \{\lambda, 2\lambda\}$. We showed in the proof of Lemma 4.7 that this implies that either $s \in \ell$ for some $\ell \in \mathcal{T}_\lambda$, or $|h_y - h'_y| \in \{\lambda, 2\lambda\}$. In the first case, $s \in \bigcup (\mathcal{P} \cup \mathcal{S})$, and in the second, $s \in \bigcup \mathcal{H}$. \square




Figure 5: For ease of visualization, we illustrate \mathcal{T}_λ as a line arrangement in \mathbb{R}^2 instead of as a plane arrangement in \mathbb{R}^3 . The function f is negative in the orange cells and nonnegative in the blue cells. The levels of the vertices are shown on the real line on the left. These are the potential values for $d_{\mathcal{M}}$, which is the supremum of the projection of the orange region to the real line. In \mathcal{T}_λ , the vertices are triple intersections of planes instead of (double) intersections of lines.

Corollary 5.4. *Suppose $C \subset \Omega$ is an open cell. Either $f(s, \lambda) > 0$ for all $(s, \lambda) \in C$, or $f(s, \lambda) < 0$ for all $(s, \lambda) \in C$.*

Proof. Since C is connected and f is continuous, $f(C)$ is connected. By Lemma 5.3, $0 \notin f(C)$, so either $f(C) \subset (0, \infty)$ or $f(C) \subset (-\infty, 0)$. \square

By definition,

$$\begin{aligned} d_{\mathcal{M}} &= \sup_{\text{slice } s} d_{\mathcal{B}}(M^s, N^s) \\ &= \sup\{\lambda \mid \exists s \text{ such that } f(s, \lambda) \leq 0\} \\ &= \sup\{\lambda \mid \exists s \text{ such that } f(s, \lambda) < 0\} \\ &= \sup\{\lambda \mid \exists s \text{ such that } (s, \lambda) \in f^{-1}(-\infty, 0)\}. \end{aligned}$$

Since $d_{\mathcal{M}} > 0$ by assumption, there exists an s such that $d_{\mathcal{B}}(M^s, N^s) > 0$, so none of the sets we are taking the supremum over is empty. Since $\Omega \setminus \bigcup \mathcal{T}$ (i.e., the union of all open cells in Ω) is dense in Ω and $f^{-1}(-\infty, 0)$ is open, this is equivalent to

$$d_{\mathcal{M}} = \sup\{\lambda \mid \exists s \text{ such that } (s, \lambda) \in f^{-1}(-\infty, 0) \cap (\Omega \setminus \bigcup \mathcal{T})\}.$$

By Corollary 5.4, it follows that

$$d_{\mathcal{M}} \in \left\{ \sup_{(s, \lambda) \in C} \lambda \mid C \text{ open cell in } \Omega \right\} \subset \left\{ \sup_{(s, \lambda) \in C} \lambda \mid C \text{ open cell} \right\}, \quad (1)$$

where we now assume that s is a point in \mathbb{R}^2 that is not necessarily in ω . Call a cell *small* if its projection onto \mathbb{R} is bounded above. Since we have assumed $d_{\mathcal{M}} \neq \infty$, we can restrict ourselves to small cells:

$$d_{\mathcal{M}} \in \left\{ \sup_{(s, \lambda) \in C} \lambda \mid C \text{ small open cell} \right\} = \left\{ \max_{(s, \lambda) \in \overline{C}} \lambda \mid C \text{ small open cell} \right\}. \quad (2)$$

(\overline{C} denotes the closure of C .) Each value $\max_{(s,\lambda) \in \overline{C}} \lambda$ must be attained at a point that is the unique intersection point between three planes.⁴ Let the *level* of a point in \mathbb{R}^3 be its third coordinate. That is, the level of (a, b, λ) is λ . By a *vertex* of \mathcal{T} , we mean a point p such that there exist $P, P', P'' \in \mathcal{T}$ with $P \cap P' \cap P'' = \{p\}$. See Fig. 5. We sometimes abuse terminology and call $\{p\}$ a vertex when we mean that p is a vertex.

Searching through the vertices to find $d_{\mathcal{M}}$ Recall that we assume that $T(n)$ is an upper bound on the time needed to decide $d_{\mathcal{M}} \leq \lambda$ for a given λ , and that we proved that we can take $T(n)$ to be $O(n^5 \log n)$ in the previous section. To find $d_{\mathcal{M}}$, we could do the following: Compute the intersections of all triples of planes in \mathcal{T} (for the triples whose intersection is a single point), sort the levels of these points, and find $d_{\mathcal{M}}$ by binary search. This would give us a runtime of $O(n^6 \log n + T(n) \log n) = O(n^6 \log n)$, as the most expensive operation is to sort the $O(n^6)$ triple intersections.

However, we want to do better and describe an algorithm that runs in expected time

$$O((n^4 + T(n)) \log^2 n).$$

Roughly, our idea is to start with an interval $(a, b]$ that contains $d_{\mathcal{M}}$ and run through the planes in \mathcal{T} in random order, narrowing the interval as we go. After having dealt with a plane $P \in \mathcal{T}$, we want the updated a and b to be such that the interval (a, b) does not contain the level of any vertex in P . In the end, we will be left with an interval $(a, b]$ such that b is the only point in the interval that is the level of a vertex, so we can conclude that $d_{\mathcal{M}} = b$. By a randomization argument, we will prove that the expected number of times we have to update a and b is $O(\log n)$. Thus, since $|\mathcal{T}| = O(n^2)$, it is sufficient to prove that we can check if we have to update a and b for a given plane in $O(n^2 \log n)$ (also $O(n^2 \log^2 n)$ would be good enough), and that we do not spend more time than $O((n^4 + T(n)) \log n)$ every time we update a and b .

Definition 5.5. For a plane $P \in \mathcal{T}$, let $c_1 < \dots < c_K$ be the levels of the vertices of \mathcal{T} in P . Let $c_{K+1} = \infty$, and if $c_1 > 0$, let $c_0 = 0$. Let $I_P = (\alpha_P, \beta_P]$ be the interval of the form $(c_i, c_{i+1}]$ that contains $d_{\mathcal{M}}$.

Lemma 5.6. For any $P \in \mathcal{T}$ and any interval $(\alpha, \beta]$ containing $d_{\mathcal{M}}$, we can decide if $(\alpha, \beta] \subseteq I_P$ in $O(n^2 \log n)$.

The idea of the proof is illustrated in Fig. 6: α and β induce two parallel lines on P , and the planes in \mathcal{T} not parallel to P induce a line arrangement on P . Now, $(\alpha, \beta] \subseteq I_P$ if and only if this line arrangement does not have an intersection point in the strip between α and β . This in turn is equivalent to the order of these arrangement lines along the α line and along the β -line being the same, which can be checked in $O(n^2 \log n)$ time by sorting. The subsequent proof contains more details:

Proof. If P is of the form $\mathbb{R}^2 \times \{\lambda\}$, then all the vertices in P are at level λ . Thus, I_P is either $(0, \lambda]$ or $(\lambda, \infty]$, so $(\alpha, \beta] \subseteq I_P$ holds if and only if $\lambda \notin (\alpha, \beta)$, which we can check in constant time. (The intersection $(\alpha, \beta] \cap I_P$ cannot be empty, as both intervals contain $d_{\mathcal{M}}$.)

Assume P is not of the form $\mathbb{R}^2 \times \{\lambda\}$. Then P intersects the planes $\mathbb{R}^2 \times \{\alpha\}$ and $\mathbb{R}^2 \times \{\beta\}$ in lines P_α and P_β , respectively. Let $P_{(\alpha, \beta)}$ be the subset of P that lies strictly between P_α and P_β . We have $(\alpha, \beta] \subseteq I_P$ if and only if there are no vertices in $P_{(\alpha, \beta)}$.

⁴This requires that not all the planes in \mathcal{T} are parallel to a single line. Since M and N are not both the zero module, \mathfrak{G} contains a generator g . Then there are three planes $P_{g,g,1}$, S_0 and $\mathbb{R}^2 \times \{0\}$ in \mathcal{T} that are not all parallel to one line.




Figure 6: Illustration of P with P_α and P_β as well as the intersection with slanted planes P^1, \dots, P^5 . $P_{(\alpha,\beta)}$ is the open region in pink. We have $P^1 <_\alpha P^2 =_\alpha P^3 <_\alpha P^4 <_\alpha P^5$ and $P^1 =_\beta P^2 <_\beta P^4 <_\beta P^3 =_\beta P^5$. The only i and j for which both $P^i <_\alpha P^j$ and $P^j <_\beta P^i$ hold are $i = 3$ and $j = 4$, and this corresponds to the only vertex in $P_{(\alpha,\beta)}$ arising from the planes we have drawn, namely $p \in P \cap P^3 \cap P^4$.

Call a plane $P' \in \mathcal{T}$ *straight* if $P \cap P'$ is a line contained in a plane of the form $\mathbb{R}^2 \times \{\lambda\}$. That is, all the points in $P \cap P'$ are at the same level. Call a plane in \mathcal{T} *slanted* if it is neither straight nor parallel to P .

We first find all the straight planes in \mathcal{T} and check if the level of the intersection of each plane with P (i.e., the level of any point in this intersection) is strictly between α and β . We can do this in $O(n^2)$. If we find one such plane P' , then all the vertices in $P \cap P'$ are at a level between α and β , and we conclude that $(\alpha, \beta] \not\subseteq I_P$. Otherwise, we know that no straight plane contains a vertex in $P_{(\alpha,\beta)}$, so we can ignore the straight planes.

Assuming that the straight planes can be safely disregarded, the only vertices we have left to check are those of the form $P \cap P' \cap P''$ where P' and P'' are both slanted. If P' is slanted, $P' \cap P_\alpha$ and $P' \cap P_\beta$ both contain exactly one point. This is because the intersection $P \cap P'$ is a line, as P and P' are not parallel, and the line contains exactly one point at every level, as the line does not lie at a fixed level. In $O(n^2)$, we can compute $P' \cap P_\alpha$ and $P' \cap P_\beta$ for all slanted P' . The set of slanted planes can be sorted by when they intersect P_α as we traverse P_α in one direction. This gives a total preorder \leq_α on the set of slanted planes with $P' \leq_\alpha P''$ if P' intersects P_α before or at the same point as P'' . This is not necessarily a total order, as several planes might intersect P_α in the same point. Define a second total preorder \leq_β by doing the same for P_β , traversing it in the same direction as P_α . See Fig. 6.

$P \cap P' \cap P''$ is a vertex in $P_{(\alpha,\beta)}$ if and only if $P' <_\alpha P''$ and $P'' <_\beta P'$, or vice versa. Whether P' and P'' satisfying these inequalities exist can be checked in $O(n^2 \log n)$: Fix an arbitrary total order $<$ of the slanted planes. List the slanted planes in increasing order by \leq_α , using \leq_β as a tiebreaker if $P' =_\alpha P''$ and $<$ as a tiebreaker if in addition $P' =_\beta P''$ (in this case, $P \cap P' \cap P''$ is a line) in $O(n^2 \log n)$. This gives a total order on the set of slanted planes. Then list the slanted planes by the same method with \leq_α and \leq_β switching roles. There are P' and P'' such that $P' <_\alpha P''$ and $P'' <_\beta P'$ if and only if these two lists are not identical: If $P' <_\alpha P''$ and $P'' <_\beta P'$, then P' and P'' are in opposite orders in the two lists, and if $P' \leq_\alpha P''$ and $P' \leq_\beta P''$, then they are in the same order. We can check if the lists are identical by running through them simultaneously in $O(n^2)$. \square

Lemma 5.7. *If $T(n)$ is an upper bound on the runtime for deciding if $d_{\mathcal{M}} \leq \lambda$ for a given λ , then for any $P \in \mathcal{T}$, we can compute I_P in $O((n^4 + T(n)) \log n)$.*

Proof. For each pair P', P'' of planes in \mathcal{T} , we can compute $P \cap P' \cap P''$ in constant time. If this is a single point, it is a vertex. In $O(n^4 \log n)$, we can find all the vertices of \mathcal{T} in P by running through all pairs of planes in \mathcal{T} , list their levels and sort them in increasing order. Adding $c_{K+1} = \infty$ and possibly $c_0 = 0$, we get a list $(c_0 <) c_1 < \dots < c_{K+1}$ as in Definition 5.5. Using binary search, we can find c_i and c_{i+1} such that $c_i < d_{\mathcal{M}} \leq c_{i+1}$ in $O(T(n) \log n)$. By definition, $I_P = (c_i, c_{i+1}]$. \square

Theorem 5.8. *If $T(n)$ is an upper bound on the runtime for deciding if $d_{\mathcal{M}} \leq \lambda$ for a given λ , then we can compute $d_{\mathcal{M}}$ in expected time $O((n^4 + T(n)) \log^2 n)$.*

Proof. We want to compute $(\alpha, \beta] = \bigcap_{P \in \mathcal{T}} I_P$: Since all $I_P = (\alpha_P, \beta_P]$ contain $d_{\mathcal{M}}$, $(\alpha, \beta]$ contains $d_{\mathcal{M}}$. Since (α_P, β_P) does not contain the level of any vertex in P , (α, β) does not contain the level of any vertex and thus does not contain $d_{\mathcal{M}}$. We conclude that $d_{\mathcal{M}} = \beta$.

We iterate through all planes in \mathcal{T} in a random order $P_1, P_2, \dots, P_{|\mathcal{T}|}$. We first compute I_{P_1} and let $I_1 = I_{P_1}$. For all subsequent P_i , we assume $I_{i-1} = \bigcap_{j=1}^{i-1} I_{P_j}$ and check if $I_{i-1} \subseteq I_{P_i}$. If the answer is yes, we let $I_i = I_{i-1}$ and move on to the next iteration. If the answer is no, we compute I_{P_i} and let $I_i = I_{i-1} \cap I_{P_i}$. In both cases, $I_i = \bigcap_{j=1}^i I_{P_j}$, so it is clear that $I_{|\mathcal{T}|} = \bigcap_{P \in \mathcal{T}} I_P$.

Let N be the number of i for which $I_{i-1} \not\subseteq I_{P_i}$. By Lemma 5.6 and Lemma 5.7, the runtime of this algorithm is

$$O(n^2 \cdot n^2 \log n + N(n^4 + T(n)) \log n). \quad (3)$$

To find the expected runtime, we need to estimate N . For $I_{i-1} = \bigcap_{j=1}^{i-1} I_{P_j} \not\subseteq I_{P_i}$ to hold, we must either have $\alpha_{P_i} > \alpha_{P_j}$ for all $j < i$, or $\beta_{P_i} < \beta_{P_j}$ for all $j < i$. The probability for the first inequality to hold is at most $\frac{1}{i}$, as in the sequence $\alpha_{P_1}, \dots, \alpha_{P_i}$, there has to be a unique smallest number, this number has to be placed last, and it is equally likely to be in each of the i positions. A similar argument shows that also the probability of $\beta_{P_i} < \beta_{P_j}$ is at most $\frac{1}{i}$. Putting the two together, the probability of $I_{i-1} \not\subseteq I_{P_i}$ is bounded above by $\frac{2}{i}$. Summing over all i then shows that the expected value of N is $O(\log n)$. We get the expected runtime given in the lemma by putting this into Eq. (3). \square

Corollary 5.9. *We can compute $d_{\mathcal{M}}$ in expected time $O(n^5 \log^3 n)$.*

Proof. By Theorem 4.10, we can take $T(n)$ to be $O(n^5 \log n)$. Inserting this in the expected runtime $O((n^4 + T(n)) \log^2 n)$ in Theorem 5.8, we get $O(n^5 \log^3 n)$. \square

References

- [1] Ulrich Bauer and Michael Lesnick. Induced matchings and the algebraic stability of persistence barcodes. *J. Comput. Geom.*, 6(2):162–191, 2015.
- [2] Paul Bendich, J.S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. *Annals of Applied Statistics*, 10(1):198–216, 2016.
- [3] Silvia Biasotti, Andrea Cerri, Patrizio Frosini, and Daniela Giorgi. A new algorithm for computing the 2-dimensional matching distance between size functions. *Pattern Recognition Letters*, 32(14):1735–1746, 2011.

- [4] Håvard Bakke Bjerkevik, Magnus Bakke Botnan, and Michael Kerber. Computing the interleaving distance is NP-hard. *Foundations of Computational Mathematics*, 20(5):1237–1271, 2020.
- [5] Gunnar Carlsson. Topology and data. *Bulletin of the AMS*, 46:255–308, 2009.
- [6] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, 2009.
- [7] Andrea Cerri, Barbara Di Fabio, Massimo Ferri, Patrizio Frosini, and Claudia Landi. Betti numbers in multidimensional persistent homology are stable functions. *Mathematical Methods in the Applied Sciences*, 36(12):1543–1557, 2013.
- [8] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007.
- [9] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have L_p -stable persistence. *Foundations of Computational Mathematics*, 10(2):127–139, 2010.
- [10] David Cohen-Steiner, Herbert Edelsbrunner, and Dmitriy Morozov. Vines and vineyards by updating persistence in linear time. In *Proceedings of the 22nd ACM Symposium on Computational Geometry, Sedona, Arizona, USA, June 5-7, 2006*, pages 119–126. ACM, 2006.
- [11] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.
- [12] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational geometry: algorithms and applications, 3rd Edition*. Springer, 2008.
- [13] Tamal K. Dey and Cheng Xin. Computing bottleneck distance for 2-d interval decomposable modules. In *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, pages 32:1–32:15, 2018.
- [14] Herbert Edelsbrunner and John Harer. *Computational Topology. An Introduction*. American Mathematical Society, 2010.
- [15] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.
- [16] Alon Efrat, Alon Itai, and Matthew J. Katz. Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31(1):1–28, 2001.
- [17] So Mang Han, Taylor Okonek, Nikesh Yadav, and Xiaojun Zheng. Distributions of matching distances in topological data analysis. *SIURO*, 2020.
- [18] Yasuaki Hiraoka, Takenobu Nakamura, Akihiko Hirata, Emerson G. Escobar, Kaname Matsue, and Yasumasa Nishiura. Hierarchical structures of amorphous solids characterized by persistent homology. *Proceedings of the National Academy of Sciences*, 113(26):7035–7040, 2016.
- [19] John E Hopcroft and Richard M Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on computing*, 2(4):225–231, 1973.

- [20] Bryn Keller, Michael Lesnick, and Theodore L. Willke. Persistent homology for virtual screening. *ChemRxiv*, 2018.
- [21] Michael Kerber, Michael Lesnick, and Steve Oudot. Exact computation of the matching distance on 2-parameter persistence modules. In *35th International Symposium on Computational Geometry (SoCG 2019)*, pages 46:1–46:15, 2019.
- [22] Michael Kerber, Dmitriy Morozov, and Arnur Nigmatov. Geometry helps to compare persistence diagrams. *Journal of Experimental Algorithms*, 22:1.4:1–1.4:20, 2017.
- [23] Michael Kerber and Arnur Nigmatov. Efficient approximation of the matching distance for 2-parameter persistence. *36th International Symposium on Computational Geometry (SoCG 2020)*, pages 53:1–53:16, 2020.
- [24] Claudia Landi. The rank invariant stability via interleavings. In *Research in Computational Topology*, pages 1–10. Springer, 2018.
- [25] Yongjin Lee, Senja Barthel, Pawel Dlotko, Seyed Mohamad Moosavi, Kathryn Hess, and Berend Smit. High-throughput screening approach for nanoporous materials genome using topological data analysis: Application to zeolites. *Journal of Chemical Theory and Computation*, 14(8):4427–4437, 2018.
- [26] Michael Lesnick. The theory of the interleaving distance on multidimensional persistence modules. *Foundations of Computational Mathematics*, 15(3):613–650, 2015.
- [27] Michael Lesnick and Matthew Wright. Interactive visualization of 2-D persistence modules persistence modules. *arXiv:1512.00180*, 2015.
- [28] Steve Oudot. *Persistence theory: From Quiver Representation to Data Analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2015.
- [29] Pratyush Pranav, Herbert Edelsbrunner, Rien van de Weygaert, Gert Vegter, Michael Kerber, Bernard J. T. Jones, and Mathijs Wintraecken. The topology of the cosmic web in terms of persistent betti numbers. *Monthly Notices of the Royal Astronomical Society*, 465(4):4281–4310, 2017.
- [30] Michael W. Reimann, Max Nolte, Martina Scolamiero, Katharine Turner, Rodrigo Perin, Giuseppe Chindemi, Pawel Dlotko, Ran Levi, Kathryn Hess, and Henry Markram. Cliques of neurons bound into cavities provide a missing link between structure and function. *Frontiers Comput. Neurosci.*, 11:48, 2017.
- [31] Erik Rybakken, Nils A. Baas, and Benjamin Dunn. Decoding of neural data using cohomological feature extraction. *Neural Comput.*, 31(1), 2019.
- [32] Primož Skraba and Katharine Turner. Wasserstein stability for persistence diagrams. *arXiv:2006.16824*, 2021.
- [33] Rien van de Weygaert, Gert Vegter, Herbert Edelsbrunner, Bernard J. T. Jones, Pratyush Pranav, Changbom Park, Wojciech A. Hellwing, Bob Eldering, Nico Kruithof, E. G. P. (Patrick) Bos, Johan Hidding, Job Feldbrugge, Eline ten Have, Matti van Engelen, Manuel Caroli, and Monique Teillaud. Alpha, betti and the megaparsec universe: On the topology of the cosmic web. *Trans. Comput. Sci.*, 14:60–101, 2011.

- [34] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.