

Efficient Approximation of the Matching Distance for 2-Parameter Persistence

Michael Kerber 

Graz University of Technology, Austria
kerber@tugraz.at

Arnur Nigmatov 

Graz University of Technology, Austria
nigmatov@tugraz.at

Abstract

In topological data analysis, the matching distance is a computationally tractable metric on multi-filtered simplicial complexes. We design efficient algorithms for approximating the matching distance of two bi-filtered complexes to any desired precision $\varepsilon > 0$. Our approach is based on a quad-tree refinement strategy introduced by Biasotti et al., but we recast their approach entirely in geometric terms. This point of view leads to several novel observations resulting in a practically faster algorithm. We demonstrate this speed-up by experimental comparison and provide our code in a public repository which provides the first efficient publicly available implementation of the matching distance.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases multi-parameter persistence, matching distance, approximation algorithm

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.53

Related Version A full version of the paper is available at [14], <https://arxiv.org/abs/1912.05826>.

Supplementary Material Our code is available as part of the HERA library https://bitbucket.org/grey_narn/hera/src/master/matching/ and provides an efficient implementation for computing the matching distance for bi-filtrations.

Funding *Michael Kerber*: Supported by Austrian Science Fund (FWF) grant number P 29984-N35.

1 Introduction

Persistent homology [10, 4, 9, 18] is one of the major concepts in the quickly evolving field of topological data analysis. The concept is based on the idea that studying the topological properties of a data set across various scales yields valuable information that is more robust to noise than restricting to a fixed scale.

We distinguish the case of *single-parameter persistence*, where the scale is expressed by a single real parameter, and the case of *multi-parameter persistence*, in which the scale consists of two or more parameters that vary independently. The former case is the predominant one in the literature. The entire homological multi-scale evolution of the data set can be expressed by a multi-set of points in the plane, the so-called *persistence diagram*. Moreover, the *interleaving distance* yields a distance measure between two data sets by measuring the difference in their topological evolution. For a single parameter, this distance can be rephrased as a combinatorial matching problem of the corresponding persistence diagrams (known as the *bottleneck distance*) and computed efficiently [12]. These results are part of a rich theory of single-parameter persistence, with many algorithmic results and applications.


The case of *multi-parameter persistence* received significantly less attention until recently. One reason is the early result that a complete combinatorial structure such as the persistence diagram does not exist for two or more parameters [5]. Moreover, while the interleaving distance can be straight-forwardly generalized to several parameters, its computation becomes



© Michael Kerber and Arnur Nigmatov;
licensed under Creative Commons License CC-BY
36th International Symposium on Computational Geometry (SoCG 2020).
Editors: Sergio Cabello and Danny Z. Chen; Article No. 53; pp. 53:1–53:16



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



NP-hard already for two parameters, as well as any approximation to a factor less than 3 [3]. On the other hand, data sets with several scale parameters appear naturally in applications, and an efficiently computable distance measure is therefore highly important.

We focus on the *matching distance* [6, 2, 13] as a computationally tractable lower bound on the interleaving distance [15]. It is based on the observation that when restricting the multi-parameter space \mathbb{R}^d to a one-dimensional affine subspace (that is, a line in \mathbb{R}^d), we are back in the case of single-parameter persistence. We can compute the bottleneck distance between the two persistence diagrams restricted to the same subspace. The matching distance is then defined as the supremum of all bottleneck distances over all subspaces (see Section 3 for the precise definition). The matching distance has been used in shape analysis [6, 2] (where it is known as the matching distance between size functions), for virtual screening in computational chemistry [11], and a recent algorithm [13] computes the distance exactly in polynomial time (with a large exponent).

Our contribution is an improvement of an approximation algorithm by Biasotti et al. [2] for the 2-parameter case which we summarize next. We parameterize the space of all lines of interest as a bounded rectangle $R \subset \mathbb{R}^2$. To each point p in the rectangle, we assign $f(p)$ as the bottleneck distance between the two persistence diagrams when restricting the data sets to the line parameterized by p . The matching distance is then equal to $\sup_{p \in R} f(p)$. The major ingredient of the algorithm is a *variation bound* which tells how much $f(p)$ varies when p is perturbed by a fixed amount. For any subrectangle $S \subseteq R$ with center c , $f(c)$ and the variation bound yield an upper bound of f within S . We then obtain an ε -approximation with a simple branch-and-bound scheme, subdividing R with a quad-tree in BFS order and stopping the subdivision of a rectangle when its upper bound is sufficiently small.

Our contributions. We aim for a fast implementation useful for practical applications of multi-parameter persistent homology. Towards this goal, we make the following contributions:

1. We rephrase the approximation algorithm by Biasotti et al. entirely in elementary geometric terms. We think that the geometric point of view complements their formulation and makes the structure of the algorithm more accessible. Indeed, we are able to simplify several arguments from [2] (see Appendix E in [14]).
2. We provide a simple yet crucial algorithmic improvement: instead of using the global variation bound for all rectangles of the subdivision, we derive adaptive local variation bounds for each rectangle individually. This results in much smaller upper bounds and avoids many subdivisions in the approximation algorithm.
3. We experimentally compare our version of the global bound with the usage of the adaptive bounds. We show that the speed-up factor of the sharpest adaptive bound is typically between 3 and 8, depending on the input bi-filtrations (for some inputs the speed-up is 15).
4. Our code is available as part of HERA library¹ and provides an efficient implementation for computing the matching distance for bi-filtrations.

Outline. We start with a short introduction to filtrations and persistent homology in Section 2. We define the matching distance in Section 3. The algorithm to approximate it is described in Section 4, and our local variation bounds are derived in Section 5. We do experiments in Section 6 and conclude in Section 7.

¹ https://bitbucket.org/grey_narn/hera/src/master/matching/




Figure 1 Left: Mono-filtration of a simplicial complex K of dimension 2. The critical value of each simplex is displayed. Right: Examples of the complexes K_v for various values of v .

2 Background

Mono-Filtrations. Fixing a base set V , a k -simplex σ is a non-empty subset of V of cardinality $k + 1$. A *face* τ of σ is a non-empty subset of σ . A *simplicial complex* K is a collection of simplices such that whenever $\sigma \in K$, every face of σ is in K as well. The dimension of a simplicial complex K is the maximal k such that K contains a k -simplex. As an example, a graph is merely a simplicial complex of dimension 1. Following graph-theoretic notations, we call 0-simplices of K *vertices* and 1-simplices of K *edges*. A *subcomplex* of K is a subset $L \subseteq K$ such that L is again a simplicial complex. We will henceforth assume that the base set V is finite, which implies also that the simplicial complex K is finite.


A *mono-filtration* is a simplicial complex K equipped with a function $\varphi : K \mapsto \mathbb{R}$ such that for any simplex σ and any face τ of σ , it holds that $\varphi(\tau) \leq \varphi(\sigma)$. We call $\varphi(\sigma)$ the *critical value of σ* and define for any $v \in \mathbb{R}$

$$K_v := \{\sigma \in K \mid \varphi(\sigma) \leq v\}.$$

By the condition on φ from above, K_v is a subcomplex of K for each v . Moreover, whenever $v \leq w$, we have that $K_v \subseteq K_w$. Hence, the collection $(K_v)_{v \in \mathbb{R}}$ yields a nested sequence of simplicial complexes, which is entirely determined by the critical values of each simplex in K . See Figure 1 for an illustration.

Persistence diagrams. We are interested in the topological changes of $(K_v)_{v \in \mathbb{R}}$ when v increases continuously. A *persistence diagram* is a multi-set of points in $\mathbb{R} \times (\mathbb{R} \cup \{\infty\})$ with all points strictly above the diagonal $x = y$. The general definition requires a digression into representation theory and homological algebra (e.g., see [18]). Instead, we explain the idea on the problem of tracking connected components of K_v within the filtration, which is a special case of the general theory. Assume for simplicity that no two simplices have the same critical value. Whenever we reach the critical value b of a vertex, a new connected component comes into existence. We call this a *birth*. We say that the component born at b *dies* at value d if there is an edge with critical value d that merges the connected component with another connected component which was born before b . In that case, (b, d) is a point in the persistence diagram, denoting that the corresponding connected component persisted from scale b to scale d . Assuming that K is connected, each component gets assigned a unique death value except the component born at the minimal critical value. We assign the death value ∞ to this component, adding an infinite point to the diagram. The resulting diagram is called the *persistence diagram in (homological) dimension 0*. See Figure 2 (left). Similar diagrams can be defined for detecting tunnels, voids, and higher-dimensional holes in the simplicial complex.

We define a distance on two persistence diagrams D_1 and D_2 next. Fixing a partial matching between D_1 and D_2 , we assign to each match of $p \in D_1$ and $q \in D_2$ the cost $\|p - q\|_\infty = \max\{|p_x - q_x|, |p_y - q_y|\}$, with the understanding that $\infty - \infty = 0$. In particular, the cost of matching a finite to an infinite point is ∞ . Every unmatched point p gets assigned



■ **Figure 2** Left: The persistence diagram in dimension 0 of the example from Figure 1 (plus the point $(0, \infty)$ that is not drawn). Note that indeed, connected components are born at 0, 0.1, 0.4 and 0.5, and the latter three components die at 0.2, 0.6 and 0.8, respectively. Right: A partial matching of two diagrams (depicted by circles and x-shapes). The cost of each match and of each unmatched vertex is displayed. The cost of this matching is 0.125 which is in fact the optimal cost in this example, so the bottleneck distance between the diagrams is 0.125.

the cost $\frac{p_y - p_x}{2}$ which corresponds to the L_∞ -distance from p to the diagonal. Taking the maximum over all matched and unmatched points in D_1 and D_2 results in the cost of the chosen partial matching. The *bottleneck distance* between D_1 and D_2 is then the minimum cost over all possible partial matchings between D_1 and D_2 . See Figure 2 (right) for an example. Since filtrations give rise to persistence diagrams, we also talk about the bottleneck distance between two filtrations and denote it by $d_B(\cdot, \cdot)$ from now on.

We will need the following properties of the bottleneck distance. The proofs of the first three of them follow directly from the definition.

- d_B satisfies the triangle inequality: $d_B(F, H) \leq d_B(F, G) + d_B(G, h)$ for three filtrations F, G, H .
- d_B is shift-invariant: let $F = (K, \varphi)$ be a filtration, define F_r be the filtration $(K, \varphi + r)$, where the critical value of each simplex is shifted by r . Then $d_B(F, G) = d_B(F_r, G_r)$.
- d_B is homogeneous: let $F = (K, \varphi)$ be a filtration, λ be a positive number, define λF be the filtration $(K, \lambda\varphi)$. Then $d_B(\lambda F, \lambda G) = \lambda d_B(F, G)$.
- d_B is stable [7]: let $F_1 = (K, \varphi_1)$ and $F_2 = (K, \varphi_2)$ be two filtrations of the same complex such that for each $\sigma \in K$, $|\varphi_1(\sigma) - \varphi_2(\sigma)| \leq \varepsilon$. Then, $d_B(F_1, F_2) \leq \varepsilon$.

Bi-filtrations. Define the partial order \leq on \mathbb{R}^2 as $p \leq q$ if and only if $p_x \leq q_x$ and $p_y \leq q_y$. Geometrically, $p \leq q$ if and only if q lies in the upper-right quadrant with corner p . A *(1-critical) bi-filtration* is a simplicial complex K together with a function $\varphi : K \rightarrow \mathbb{R}^2$ such that for every simplex σ and every face τ of σ , we have that $\varphi(\tau) \leq \varphi(\sigma)$. As before, $\varphi(\sigma)$ is called the *critical value* of σ . Our assumption that every simplex has a unique critical value is just for the sake of simpler exposition; our ideas extend to the *k-critical* case where each σ has up to k incomparable critical values (Appendix H, [14]). Fixing $p \in \mathbb{R}^2$, we define

$$K_p := \{\sigma \in K \mid \varphi(\sigma) \leq p\}.$$

Similar to the mono-filtration case, K_p is a subcomplex and $K_p \subseteq K_q$, whenever $p \leq q$.

It is worth visualizing the construction of K_p geometrically. We can represent the bi-filtration as a multi-set of points in \mathbb{R}^2 , where each point corresponds to a simplex and is placed at the critical value of the simplex. The complex K_p then consists of all simplices that are placed in the lower-left quadrant with p at its corner. See Figure 3 for an example.




Figure 3 Left: Bi-filtration of a simplicial complex K of dimension 2. Middle: Every point in the plane denotes the critical value of a simplex. The shaded rectangle yields the simplices that belong to K_p . Right: Illustration of K_p as a subcomplex of K .




Figure 4 Left: The slice parameterized by $(\frac{\pi}{6}, 0.1)$. For two critical values of the bi-filtration from above, we illustrate the construction of the point q (displayed by a cross shape). The push of the critical value is simply the Euclidean distance to the point $(0, 0.1)$, which is the origin of the slice. Right: The non-weighted restriction on the slice. Each simplex gets its push as critical value.


3 The matching distance

Slices. Bi-filtrations are too wild to admit a simple combinatorial description such as a persistence diagram. But we can obtain a persistence diagram when restricting to a one-dimensional affine subspace. For all concepts in this subparagraph, see Figure 4 for an illustration. We consider a non-vertical line L with positive slope, which we call a *slice*. For every slice, we distinguish a point \mathcal{O} , called the *origin* of the slice. We let \mathcal{L} denote the set of all slices. Since the slope is positive, for any two distinct points p, q on L either $p \leq q$ or $q \leq p$ holds. Hence, \leq becomes a total order along L .

Given $p \in \mathbb{R}^2$, let q be the minimal point on L (with respect to \leq) such that $p \leq q$. Geometrically, q is the intersection of L with the boundary of the upper-right quadrant of p , or equivalently, the horizontally-rightwards projection of p to L if p lies above L , or the vertically-upwards projection of p to L if p lies below L . Since q lies on L , q can be written as

$$\mathcal{O} + \lambda_p \begin{pmatrix} \cos \gamma \\ \sin \gamma \end{pmatrix}$$

where γ is the angle between L and x -axis, and $\lambda_p \in \mathbb{R}$. We define λ_p as the *push* of p to L , which can be formally written as a function $\text{push} : \mathbb{R}^2 \times \mathcal{L} \rightarrow \mathbb{R}$. Geometrically, the push is simply the (signed) distance of the point q to the origin of the slice. Fixing a bi-filtration $F = (K, \varphi)$, the composition $\text{push}(\cdot, L) \circ \varphi$ yields a function $K \rightarrow \mathbb{R}$, and it can be readily checked that this function yields a mono-filtration, which we call the *non-weighted restriction* of F onto L . See Figure 4 (right) for an example.



■ **Figure 5** Two points that are close to each other might have pushes far from each other. Note that by making the slice more flat, the distance between the pushes can be made arbitrarily large.

Matching distance. Given two bi-filtrations F^1, F^2 , we could try to define a distance between them by taking the supremum of bottleneck distances between their non-weighted restrictions on all slices. However, this does not yield a meaningful result. The reason is that for almost horizontal and almost vertical slices, the pushes of two close-by points can move very far away from each other – see Figure 5 for an example. As a result, the bottleneck distance along such slices becomes arbitrarily large.

Instead, we introduce a weight for each slice. Let γ denote the angle between the slice L and x -axis. We call L *flat* if $\gamma \leq \frac{\pi}{4}$ (i.e., if its slope is ≤ 1) and *steep* if $\gamma \geq \frac{\pi}{4}$. Then we set

$$w(L) := \begin{cases} \sin \gamma & \text{if } L \text{ is flat} \\ \cos \gamma & \text{if } L \text{ is steep.} \end{cases}$$

We define the matching distance between the bi-filtrations $F^1 = (K^1, \varphi^1)$ and $F^2 = (K^2, \varphi^2)$ as

$$d_M(F^1, F^2) := \sup_{L \in \mathcal{L}} w(L) \cdot d_B(\text{restr}(F^1, L), \text{restr}(F^2, L)),$$

where $\text{restr}(F^i, L)$ denotes the non-weighted restriction of F^i onto L . Note that while the non-weighted restrictions depend on the choice of the origin, a different choice of origin for a slice only results in a uniform translation of the critical values of both mono-filtrations. Hence, the bottleneck distance does not change because of shift-invariance. This means that the matching distance is independent of the choice of the origins.

Moreover, the shift-invariance of d_B implies that if we alter φ_1 and φ_2 such that every value is translated by the same vector $v \in \mathbb{R}^2$, the matching distance does not change. Recalling that we can visualize bi-filtrations as finite multi-sets of points in \mathbb{R}^2 , we can hence assume without loss of generality that all these points are in the upper-right quadrant of the plane, that is, $\varphi^i(\sigma) \in [0, \infty) \times [0, \infty)$.


Let us now define the *weighted push* of a point p to a slice L as $\text{wpush}(p, L) = w(L) \text{push}(p, L)$, and let F_L denote the mono-filtration induced by $\sigma \mapsto \text{wpush}(\varphi(\sigma), L)$. We call F_L a *weighted restriction* of F onto L . Note that F_L equals $\text{restr}(F, L)$ except that all critical values are scaled by the factor $w(L)$. Using homogeneity of d_B , we see that

$$d_M(F^1, F^2) = \sup_{L \in \mathcal{L}} d_B(F_L^1, F_L^2). \tag{1}$$

We will use this equivalent definition of the matching distance in the remaining part of the paper, and “restriction” will always mean “weighted restriction”.

4 The approximation algorithm

The idea of the approximation algorithm for d_M is to sample the set of slices through a finite sample, and chose the maximal bottleneck distance between the (weighted) restriction encountered as the approximation value. In order to execute this plan, we need to parameterize the space of slices and need to compute the restriction of a parameterized slice efficiently.



■ **Figure 6** A steep y -slice (I), a flat y -slice (II), a steep x -slice(III) and a flat x -slice. The slopes are 2 for the steep and $\frac{1}{2}$ for the flat slices, and the origin is at $(0, 2)$ for the y -slices and at $(2, 0)$ for the x -slices. Consequently, all four slices are parameterized by $(\frac{1}{2}, 2)$.

Slice Parameterization. Every slice has a unique point where the line enters the positive quadrant of \mathbb{R}^2 , which is either its intersection with the positive x -axis, the positive y -axis, or the point $(0, 0)$. From now on, we always use this point as the origin of the slice.

We call a slice an x -slice if its origin lies on the positive x -axis, and call it a y -slice if its origin lies on the positive y -axis (slices through the origin are both x - and y -slices). Recall also that a slice is flat if its slope is less than 1, and steep if it is larger than 1. Thus, a slice belongs to one of the four types: flat x -slices, flat y -slices, steep x -slices and steep y -slices. Every slice is represented as a point $(\lambda, \mu) \in (0, 1] \times [0, \infty)$ where the interpretation of the parameters depends on the type of the slice as follows: Let $\mathcal{O} = (\mathcal{O}_x, \mathcal{O}_y)$ be the origin of L , and recall that γ is the angle of the slice with the x -axis. Then

$$\lambda = \begin{cases} \tan(\gamma), & \text{if } L \text{ is flat} \\ \cot(\gamma), & \text{if } L \text{ is steep} \end{cases}, \quad \mu = \begin{cases} \mathcal{O}_x & \text{if } L \text{ is } x\text{-slice} \\ \mathcal{O}_y & \text{if } L \text{ is } y\text{-slice} \end{cases}.$$

In other words, λ is the slope of the line in the flat case, and the inverse of the slope in the steep case, and μ contains the non-trivial coordinate of the origin. Note that the same pair of parameters can parameterize different slices depending on the type. Figure 6 illustrates this.

Weighted pushes. We next show a simple formula for the value of $\text{wpush}(p, L)$ depending on the type of the slice.

► **Lemma 1.** *With the chosen parameterization and choice of origin on L , $\text{wpush}(p, L)$ is computed according to the formulas given in Table 1.*

Proof. The proof of the lemma is a series of elementary calculations. Let us consider, for example, the case of a flat y -slice. In this case the slice $L = (\lambda, \mu)$ is given by

$$\left\{ \left(\begin{array}{c} 0 \\ \mu \end{array} \right) + \rho \left(\begin{array}{c} \cos \gamma \\ \sin \gamma \end{array} \right) \mid \rho \in \mathbb{R} \right\},$$


If $p = (p_x, p_y)$ is above L , we consider the point $q = (q_x, q_y)$ which is the intersection of L and the line $y = p_y$. Obviously, $q_y = p_y$, and, since q lies on L , the second coordinate yields

$$\rho = \frac{q_y - \mu}{\sin \gamma} = \frac{p_y - \mu}{\sin \gamma}.$$

By definition, ρ is the push of p to L and since L is flat, we have that $w(L) = \sin \gamma$ which cancels with the denominator. The other 7 cases are proved analogously. ◀

■ **Table 1** Formulas for weighted push of (p_x, p_y) onto a slice $L = (\lambda, \mu)$.

	y-slices		x-slices	
	flat	steep	flat	steep
p above L	$p_y - \mu$	$\lambda(p_y - \mu)$	p_y	λp_y
p below L	λp_x	p_x	$\lambda(p_x - \mu)$	$p_x - \mu$



■ **Figure 7** Illustration for the fact that slices with larger value of μ can be ignored.

All 8 expressions in Table 1 involve only addition and multiplication without trigonometric functions. Hence we can extend them continuously to $\lambda = 0$, which corresponds to horizontal lines (in the flat case) or vertical lines (in the steep case). With this interpretation, we can extend \mathcal{L} to a set $\bar{\mathcal{L}}$ in (1), containing these limit cases, without changing the supremum.

Next, we observe that we can restrict our attention to a bounded range of μ -parameters. For that, let X denote the maximal x -coordinate and Y be the maximal y -coordinate among all critical values of F^1 and of F^2 . For a y -slice (steep or flat) $L = (\lambda, \mu)$ with $\mu > Y$, let $L' = (\lambda, Y)$ be the parallel slice with origin at $(0, Y)$. All critical values of $F^{1,2}$ are below L and L' by construction (recall that all critical points are assumed in the upper-right quadrant), hence we obtain the push by projecting vertically upwards. Looking at the second row in Table 1, we see that the weighted pushes are independent of μ , and therefore equal for L and L' . Hence, the weighted bottleneck distances along L and L' are equal: $d_B(F_L^1, F_L^2) = d_B(F_{L'}^1, F_{L'}^2)$. See Figure 7 for an illustration. We conclude that for y -slices it suffices to consider $0 \leq \mu \leq Y$ in (1) without changing the matching distance. An analogous argument shows that for x -slices, it is only necessary to consider $0 \leq \mu \leq X$.

To summarize the last two observations, we arrive at the following statement. There are sets \mathcal{L}_1 of flat x -slices, \mathcal{L}_2 of steep x -slices, \mathcal{L}_3 of flat y -slices, and \mathcal{L}_4 of steep y -slices (with each set containing some vertical/horizontal lines as limit case) such that

$$d_M(F^1, F^2) = \sup_{L \in \mathcal{L}_1 \cup \dots \cup \mathcal{L}_4} d_B(F_L^1, F_L^2) \tag{2}$$

and such that \mathcal{L}_1 and \mathcal{L}_2 are parameterized by $[0, 1] \times [0, X]$ and \mathcal{L}_3 and \mathcal{L}_4 by $[0, 1] \times [0, Y]$.

Approximation. We present an approximation algorithm that, given two bi-filtrations F^1 and F^2 and some $\varepsilon > 0$ returns a number δ such that

$$d_M(F^1, F^2) - \varepsilon \leq \delta \leq d_M(F^1, F^2).$$

We assume that the two bi-filtrations are given as simplicial complexes, i.e., a list of simplices, where each simplex is annotated with two real values denoting the critical value of the simplex. In the description, we set $T := \{x\text{-flat}, x\text{-steep}, y\text{-flat}, y\text{-steep}\}$ for the type of a slice. The algorithm is based on the following two primitives:

Eval(F^1, F^2, L) Computes $d_B(F_L^1, F_L^2)$, where L is specified by the triple (λ, μ, t) where (λ, μ) are the parameterization of L and $t \in T$ denotes its type.

Bound(F^1, F^2, B, t) If B is an axis-parallel rectangle and $t \in T$, the pair (B, t) specifies a set of slices \mathcal{L}_0 . The primitive computes a number $\mu \in \mathbb{R}$ such that, for every $L \in \mathcal{L}_0$,

$$d_B(F_L^1, F_L^2) \leq \mu.$$

With these two primitives, we can state our approximation algorithm: from now on, we refer to axis-parallel rectangles as *boxes* for brevity. We start by computing maximal coordinates X and Y of critical values of F^1 and F^2 and enqueueing the four initial items $([0, 1] \times [0, Y], y\text{-steep})$, $([0, 1] \times [0, Y], y\text{-flat})$, $([0, 1] \times [0, X], x\text{-steep})$, and $([0, 1] \times [0, X], x\text{-flat})$ into a FIFO-queue. We also maintain a variable ρ storing the largest bottleneck distance encountered so far, initialized to 0.

Now, we pop items from the queue and repeat the following steps: for an item (B, t) , let L denote the slice that corresponds to the center point of B . We call **Eval**(F^1, F^2, L) and update ρ if the computed value is bigger than the current maximum. Then, we compute $\mu \leftarrow \text{Bound}(F^1, F^2, B, t)$. If $\mu > \rho + \varepsilon$, we split B into 4 sub-boxes B_1, \dots, B_4 of equal dimensions (using the center as splitpoint) and enqueue $(B_1, t), \dots, (B_4, t)$. When the queue is empty, we return $\delta \leftarrow \rho$. This ends the description of the algorithm.

Assuming that the above algorithm terminates (which is unclear at this point because it depends on the implementation of the **Bound** primitive), the output is indeed an ε -approximation. This can be derived directly from the termination condition of the subdivision and the fact that ρ is non-decreasing during the algorithm. See Appendix A, [14] for details.

A variant of the above algorithm computes a relative approximation of the matching distance, that is, a number δ such that

$$d_M(F^1, F^2) \leq \delta \leq (1 + \varepsilon) d_M(F^1, F^2).$$

The algorithm is analogous to the above, with the difference that a box is subdivided if $\mu > (1 + \varepsilon)\rho$, and at the end of the algorithm $(1 + \varepsilon)\rho$ is returned as δ . The correctness of this variant follows similarly. However, the algorithm terminates only if $d_M(F^1, F^2) > 0$, and its complexity depends on the value of the matching distance.

What is needed to realize the **Eval** primitive? First, we compute the weighted pushes of each critical value of F^1 and of F^2 in time proportional to the number of critical values using Lemma 1. Then, we compute the persistence diagrams of F^1 and of F^2 , and their bottleneck distance. Both steps are well-studied standard tasks in persistent homology, and several practically efficient algorithms have been studied. We use PHAT [1] for computing persistence diagrams and HERA [12] for the bottleneck computation.

5 The Bound primitive

Recall that the input of **Bound** is (F^1, F^2, B, t) , where (B, t) specifies a collection of slices of type t . In what follows, we will identify points in B with the parameterized slice, writing $L \in B$ to denote that L is obtained from a pair of parameters $(\lambda, \mu) \in B$ with respect to type t (which we skip for notational convenience).

Let L_c be the slice corresponding to the center of B . The *variation* of a point $p \in \mathbb{R}^2$ for B denotes how much the weighted push of p changes when the slice is changed within B :

$$v(p, B) := \max_{L \in B} |\text{wpush}_p(L) - \text{wpush}_p(L_c)|.$$

53:10 Matching Distance Approximations

For a bi-filtration F , we define

$$v(F, B) := \max_{p \text{ critical value of } F} v(p, B).$$

The variation yields an upper bound for the bottleneck distance within a box:

► **Lemma 2.** *With the notation as before, we have that for two filtrations F^1, F^2 that*

$$\sup_{L \in B} d_B(F_L^1, F_L^2) \leq v(F^1, B) + d_B(F_{L_c}^1, F_{L_c}^2) + v(F^2, B)$$

Proof. By triangle inequality of the bottleneck distance,

$$d_B(F_L^1, F_L^2) \leq d_B(F_L^1, F_{L_c}^1) + d_B(F_{L_c}^1, F_{L_c}^2) + d_B(F_{L_c}^2, F_L^2).$$

Looking at the first term on the right, we have two filtrations of the same simplicial complex, and every critical values changes by at most $v(F^1, B)$ by definition of the variation. Hence, by stability of the bottleneck distance, $d_B(F_L^1, F_{L_c}^1) \leq v(F^1, B)$. The same argument applies to the third term which proves the theorem. ◀

Note that the second term in the bound of Lemma 2 is the value at the center slice, which is already computed in the algorithm. It remains to compute the variation of a bi-filtration within B . This, in turn, we do by analyzing the variation of a point p within B . We show

► **Theorem 3.** *For a box B , let L_1, \dots, L_4 be the four slices on the corners of B . Then*

$$v(p, B) = \max_{i=1, \dots, 4} |\text{wpush}_p(L_i) - \text{wpush}_p(L_c)|$$

The theorem gives a direct algorithm to compute $v(p, B)$, just by computing the weighted pushes at the four corners (in constant time) and return the maximal difference to the weighted push in the center. Doing so for every critical point of a bi-filtration F yields $v(F, B)$, and with Lemma 2 an algorithm for the **Bound** primitive that runs in time proportional to the number of critical points of F^1 and F^2 . We refer to this bound as *local linear bound* (where the term “linear” refers to the computational complexity), or as *L-bound*.

The proof of the theorem is presented in Appendix B, [14]. The main idea is that the expression $|\text{wpush}_p(L_{\lambda, \mu}) - \text{wpush}_p(L_c)|$ (with $L_{\lambda, \mu}$ the slice given by (λ, μ)) has no isolated local maxima, even for a fixed λ or a fixed μ . That implies that from any (λ, μ) in the box, there is rectilinear path to a corner on which the expression above is non-decreasing.

A coarser bound. We have derived a method to compute $v(p, B)$ exactly which takes linear time. Alternatively, we can derive an upper bound as follows:

► **Theorem 4.** *Let B be a box $[\lambda_{\min}, \lambda_{\max}] \times [\mu_{\min}, \mu_{\max}]$ with center (λ_c, μ_c) , width $\Delta\lambda = \lambda_{\max} - \lambda_{\min}$ and height $\Delta\mu = \mu_{\max} - \mu_{\min}$. Then, for any point $p \in [0, X] \times [0, Y]$, $v(p, B)$ is at most*

$$\begin{cases} \frac{1}{2}(\Delta\mu + X\Delta\lambda) & \text{for flat } y\text{-slices} \\ \frac{1}{2}(\lambda_c\Delta\mu + (Y - \mu_{\min})\Delta\lambda) \} & \text{for steep } y\text{-slices} \\ \frac{1}{2}(\lambda_c\Delta\mu + (X - \mu_{\min})\Delta\lambda) & \text{for flat } x\text{-slices} \\ \frac{1}{2}(\Delta\mu + Y\Delta\lambda) & \text{for steep } x\text{-slices.} \end{cases}$$

Importantly, the bound is independent of p , and hence also an upper bound for $v(F, B)$ that can be computed in constant time; we refer to it as *local constant bound* or *C-bound*.

The proof of Theorem 4 is based on deriving a bound of how much $\text{wpush}_p(L)$ and $\text{wpush}_p(L')$ can differ for two slices $L = (\lambda, \mu)$ and $L' = (\lambda', \mu')$ in dependence of $|\lambda - \lambda'|$ and $|\mu - \mu'|$. This bound, in turn, is derived separately for all four types of boxes and involves an inner case distinction depending on whether p lies above both slices, below both slices, or in-between. In either case, the claim of the statement follows from the bound by plugging in the center slice of a box for either L or L' . See [14] for the detailed proof.

Termination and complexity. We show next that our absolute approximation algorithm terminates when realized with either the local linear bound or the local constant bound. In what follows, set $C := \max\{X, Y\}$. In the subdivision process, each box B considered is assigned a *level*, where the level of the four initial boxes is 0, and the four sub-boxes obtained from a level- k -box have level $k + 1$. Since every box is subdivided by the center, we have immediately that for a level- k -box, $\Delta\lambda = 2^{-k}$, and $\Delta\mu \leq C2^{-k}$. Using these estimates in Theorem 4, we obtain

$$v(F^i, B) \leq \frac{1}{2}(C2^{-k} + C2^{-k}) = C2^{-k}. \quad (3)$$

for $i = 1, 2$ and every level- k -box B considered by the algorithm. Note that the local constant bound yields a bound on $v(F^i, B)$ that is not worse, and so does the local linear bound (which computes $v(F^i, B)$ exactly). Hence we have

► **Lemma 5.** *Let B be a level- k -box considered in the algorithm. Then, μ , the result of the Bound primitive in the algorithm, satisfies*

$$\mu \leq d_B(F_{L_c}^1, F_{L_c}^2) + 2C2^{-k}$$

both for the local linear and local constant bound.

It follows easily that if $2C2^{-k} \leq \varepsilon$, or equivalently $2^k \geq \frac{2C}{\varepsilon}$, a box is not further subdivided in the algorithm (see Lemma 9 in [14]). Moreover, if the maximal subdivision depth is k , the algorithm visits $O(4^k)$ boxes, and requires $O(n^3)$ time per box because of the computation of two persistence diagrams and their bottleneck distance. Combining these results leads to the complexity bound.

► **Theorem 6.** *Our algorithm to compute an absolute ε -approximation terminates in*

$$O\left(n^3 \left(\frac{C}{\varepsilon}\right)^2\right)$$

steps in the worst case (both for the linear and constant bound).

See again [14] for more details. In there, we also derive a similar bound for the variant of computing a relative approximation.

► **Theorem 7.** *Our algorithm to compute a relative $(1 + \varepsilon)$ -approximation terminates in*

$$O\left(n^3 \left(\frac{C(1 + \varepsilon)}{\varepsilon d_M(F^1, F^2)}\right)^2\right)$$

steps in the worst case if $d_M(F^1, F^2) > 0$.

6 Experiments

Experimental setup. Our experiments were performed on a workstation with an Intel(R) Xeon(R) CPU E5-1650 v3 CPU (6 cores, 3.5GHz) and 64 GB RAM, running GNU/Linux (Ubuntu 16.04.5). The code was written in C++ and compiled with gcc-8.1.0.

We generated two datasets, which we call GH and ED, following [2] (unfortunately, we were unable to get either the code or the data used by the authors). Each of the 70 files in the datasets is a lower-star bi-filtration of a triangular mesh (2-dimensional complex), representing a 3D shape. We also generated dataset RND of larger random bi-filtrations with up to 2,000 vertices. A more detailed description of the datasets can be found in [14]. In all our experiments we used persistence diagrams in dimension 0. In the experiments with the datasets GH and ED, we computed all pairwise distances; in the experiments with RND we computed distances only between bi-filtration with the same number of vertices. We used relative error threshold, which we call ε in this section throughout (i.e., we always compute $(1 + \varepsilon)$ -approximation).

Comparison of different bounds. First, we experimentally compare the performance of our algorithm with the L-bound from Theorem 3 and with the C-bound from Theorem 4. Obviously, the L-bound is sharper, so it allows us to subdivide fewer boxes and in this sense is more efficient. However, it is not a priori clear that the L-bound is preferable, since its computation takes $O(n)$ time per box, in contrast to the constant bound.

Secondly, we compare the local bounds L and C with the bound provided by Equation (3), which we call the *global* bound or G-bound because it only depends on the size of the box. A bound of that sort (worse than Equation (3)) is used in [2].

Recall that the dominating step in the complexity analysis (and in practice) is the computation of persistence, and we perform two such computations when we call the `Eval` primitive. Therefore we are interested in the number of calls of `Eval`; for brevity, we refer to this number simply as the number of *calls*.

In Table 2, we give the average number of calls and timings for different datasets and values of ε . Actually, the variance behind the average in these tables is large, so we additionally provide Table 3 and Table 4, where we report the average, maximal, and minimal ratios of the number of calls and time that the algorithm needs with different bounds. For instance, the third line of Table 3 shows that for all pairs from ED that we tested with relative error $\varepsilon = 0.5$, switching from the local constant to the local linear bound reduces the number of calls by a factor between 1.78 and 4.92.

Table 4 shows that, as expected, the C-bound always performs better than the G-bound, with the average speed-up around 2. The L-bound brings an additional speed-up by a factor of 1.5-2 in terms of the running time; the number of calls is reduced more significantly, by a factor of 3. However, the second from the right column of Table 4 shows that the running time can sometimes moderately increase, if we switch to the L-bound from the C-bound. If we compare the G-bound with the L-bound directly (these numbers are not present in Table 4), the best speed-up factor is 15.6, the worst one is 1.14, and the average is between 3 and 8, depending on the dataset.

Breadth-first search, depth-first search, and error decay. In the formulation of our algorithm we used a FIFO-queue. This means that we traverse the quad-tree in breadth-first order (i.e., level by level). Other traversal strategies are possible, for instance a depth-first order, or a greedy algorithm where boxes with large bottleneck distance at the center are

■ **Table 2** Average number of calls and average running time with the L-, C- and G-bounds for different datasets and relative error ε .


	#Calls			Time (min)		
	L	C	G	L	C	G
GH, $\varepsilon = 0.5$	938	2502	11082	2.08	3.67	18.78
ED, $\varepsilon = 0.1$	1455	3920	27529	2.58	3.26	25.96
ED, $\varepsilon = 0.5$	169	531	2112	0.28	0.42	1.67

■ **Table 3** Comparison of number of calls between the global, local constant, and local linear bounds. G / C denotes the ratio of the G-bound compared with the C-bound; C / L denotes the ratio of the C-bound compared with the L-bound.

Dataset, ε	#Calls: G / C			#Calls: C / L		
	Avg	Min	Max	Avg	Min	Max
GH, $\varepsilon = 0.5$	1.80	1.21	3.28	3.10	1.51	7.02
ED, $\varepsilon = 0.1$	2.93	1.43	5.07	3.00	1.88	6.82
ED, $\varepsilon = 0.5$	1.94	1.17	2.81	3.29	1.78	4.92
RND, $\varepsilon = 0.1$	6.06	2.76	10.58	2.08	1.91	2.47

■ **Table 4** Comparison of running time between the global, local constant, and local linear bounds.

Dataset, ε	Time: G / C			Time: C / L		
	Avg	Min	Max	Avg	Min	Max
GH, $\varepsilon = 0.5$	1.66	1.00	3.18	2.03	0.75	6.32
ED, $\varepsilon = 0.1$	3.12	1.44	5.21	1.64	0.81	3.40
ED, $\varepsilon = 0.5$	2.08	1.07	3.38	1.59	0.92	3.89
RND, $\varepsilon = 0.1$	5.73	2.83	10.67	1.93	1.66	2.20



■ **Figure 8** Error decay with time for the C- and G-bounds.

picked first. We experimented with these variants and found no significant difference. The explanation is that a good lower bound is achieved after a small number of iterations in every variant, and the remaining part of the computations is mostly to certify the answer.

A variant of our algorithm is that instead of ε , we are given a time budget and want to compute the best possible (relative) approximation in this time limit. In such a case, we propose to traverse the quad-tree by always subdividing the box with the largest upper bound (i.e., the output of the `Bound` primitive). When the time is over, it suffices to peek at the top of the priority queue to get the current upper bound, and we can output the lower bound ρ and the relative error that we can guarantee at this moment. It is instructive to plot how the relative error decreases as the algorithm runs; see Figure 8. For instance, we can see that it takes approximately 3.5 times longer to bring the relative error below 0.1 than below 0.2, if we use the constant bound. This agrees with the complexity estimate in Theorem 7.

One detail in this plot is relevant for the experiments of the previous subparagraph. If we choose a relative error ε_0 and draw a horizontal line $\varepsilon = \varepsilon_0$ in Figure 8 until it intersects the plotted curves, then the x -coordinate of the intersection is the time that our algorithm needs to guarantee a $1 + \varepsilon_0$ approximation with the corresponding bound. We can see that the difference between the time needed with the global bound and the time needed with the constant bound is not large for some values of ε_0 , but a small change of ε_0 can rapidly increase it. Clearly, this is highly input-specific, and this partially explains the large variation in the improvement ratios that we observed above, when we ran experiments with fixed ε .

We provide additional experimental results in [14]. It contains the reduction rate (the measure used in [2]) for the experiments of this section, scaling results on the dataset RND, and heatmap visualization of $d_B(F_L^1, F_L^2)$.

7 Conclusion

We presented an algorithm for the matching distance that keeps subdividing boxes until a sufficiently close approximation of the matching distance can be guaranteed. This high-level description also applies to the previous approach by Biasotti et al., which raises the question of how the approaches compare in the details. Instead of pointing out similarities and differences in the technical part, we give a detailed discussion on this topic in [14].

We have restricted to the case of bi-filtrations in this work. Generalizations in several directions are possible. First of all, instead of bi-filtrations, our algorithm works the same when the input is a pair of presentations of persistence modules [16, 13, 3]. Since a minimal presentation of a bi-filtration can be of much smaller size than the bi-filtration itself, and its computation is feasible [16, 8], switching to a minimal presentation will most likely increase the performance further. We plan to investigate this in further work. Moreover, the case of k -critical bi-filtrations can be handled with our methodology, just by defining the push of a simplex as the minimal push over all its critical values (Appendix H in [14]). Our approach can also be combined with *barcode templates* as introduced in the RIVET library [17]. Finally, an extension of our approach to 3 and more parameters should be possible in principle, but we point out that the space of affine lines through \mathbb{R}^d is $2(d-1)$ dimensional. Hence, already the next case of tri-filtrations requires a subdivision in \mathbb{R}^4 , and it is questionable whether reasonably-sized instances could be handled by an extended algorithm.

As the experiments show, in some cases the cost of computing the L-bound makes the variant with the C-bound faster. However, parallelization of the L-bound is trivial, and we believe that on larger instances it will make the L-bound the best choice.

Finally, since the matching distance can be computed exactly in polynomial time [13], the question is whether there is a practical algorithm for this exact computation. Our current implementation can serve as a base-line for a comparison between exact and approximate version of matching distance computations that hopefully lead to further improvements for the computation of the matching distance.

References

- 1 Ulrich Bauer, Michael Kerber, Jan Reininghaus, and Hubert Wagner. Phat - persistent homology algorithms toolbox. *J. Symb. Comput.*, 78:76–90, 2017. doi:10.1016/j.jsc.2016.03.008.
- 2 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.
- 3 Håvard Bjerkevik, Magnus Botnan, and Michael Kerber. Computing the interleaving distance is NP-hard. [arXiv:1811.09165](https://arxiv.org/abs/1811.09165).
- 4 G. Carlsson. Topology and data. *Bulletin of the AMS*, 46:255–308, 2009.
- 5 G. Carlsson and A. Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, 2009. doi:10.1007/s00454-009-9176-0.
- 6 A. Cerri, B. Di Fabio, M. Ferri, P. Frosini, and C. Landi. Betti numbers in multidimensional persistent homology are stable functions. *Mathematical Methods in the Applied Sciences*, 36(12):1543–1557, 2013.
- 7 D. Cohen-Steiner, H. Edelsbrunner, and J. Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37:103–120, 2007.
- 8 Tamal Dey and Cheng Xin. Generalized persistence algorithm for decomposing multi-parameter persistence modules. [arXiv:1904.03766](https://arxiv.org/abs/1904.03766).
- 9 H. Edelsbrunner and J. Harer. *Computational Topology. An Introduction*. American Mathematical Society, 2010.
- 10 H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002. doi:10.1007/s00454-002-2885-2.
- 11 Bryn Keller, Michael Lesnick, and Theodore L Willke. Persistent homology for virtual screening, 2018.
- 12 M. Kerber, D. Morozov, and A. Nigmatov. Geometry helps to compare persistence diagrams. *Journal of Experimental Algorithms*, 22:1.4:1–1.4:20, September 2017.

53:16 Matching Distance Approximations

- 13 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.
- 14 Michael Kerber and Arnur Nigmatov. Efficient approximation of the matching distance for 2-parameter persistence. *arXiv preprint*, 2019. [arXiv:1912.05826](#).
- 15 Claudia Landi. The rank invariant stability via interleavings. In *Research in Computational Topology*, pages 1–10. Springer, 2018.
- 16 Michael Lesnick and Matthew Wright. Computing minimal presentations and bigraded betti numbers of 2-parameter persistent homology. [arXiv:1902.05708](#).
- 17 Michael Lesnick and Matthew Wright. Interactive visualization of 2-D persistence modules persistence modules. *arXiv*, 2015. [arXiv:1512.00180](#).
- 18 S. Oudot. *Persistence theory: From Quiver Representation to Data Analysis*, volume 209 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2015.