

Dimension Reduction for Curves: Simplified and Generalized

Matthijs Ebbens ✉

University of Cologne, Germany

Jie Lu ✉

University of Cologne, Germany

Alexander Munteanu ✉

TU Dortmund, Germany

Abstract

We revisit random projections for reducing the dimension of high-dimensional polygonal curves. Drawing from the toolbox of randomized linear algebra, we give a considerably simplified proof of the known $O(\varepsilon^{-2} \log(nm))$ bound on the target dimension of a random projection that preserves the continuous Fréchet distance of polygonal curves up to a factor $(1 \pm \varepsilon)$. Our proof is based on the concept of sparse oblivious subspace embeddings. While previous techniques were limited to the case of the Fréchet distance, our techniques are fairly general and extend to all possible distance measures that involve the maximum, a sum or an integral over Euclidean distances between pairs of points on both input curves. We define a generalized dissimilarity measure for curves that includes several popular measures such as Fréchet, q -DTW, Hausdorff, etc. as special cases and show that the same dimension reduction technique works for this generalized dissimilarity measure. Finally, we apply the same framework for dimension reduction to piecewise linear surfaces, after extending the distance measure suitably to such surfaces.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Random projections and metric embeddings

Keywords and phrases dimension reduction, Fréchet distance, dynamic time warping, polygonal curves, piecewise linear surfaces

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.116

Related Version *Full Version:* <https://arxiv.org/abs/2607.03112>

Funding *Matthijs Ebbens:* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project Number 459420781.

Alexander Munteanu: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project Number 535889065.

1 Introduction

In this paper, we study dimension reduction for polygonal curves and build up a fairly general framework, which is widely independent of the dissimilarity measure that is used to compare them. Thus, in contrast to metric embeddings, our framework aims to embed *the curves* into lower dimensional space, not a specific dissimilarity measure. We first motivate our work from a broader perspective before we return to polygonal curves and surfaces.

Dimension reduction is one of the most fundamental tasks in computational geometry and has many beneficial consequences. For instance, algorithms for a computational task may have a dependence on the dimension of their input objects, such as points, shapes, or polygonal curves. Reducing the dimension of these objects, say from d to some $t \ll d$, allows to apply the same algorithm on the reduced data, but with the dependence on d replaced by the significantly smaller value t . This comes at the cost of a small approximation error which



© Matthijs Ebbens, Jie Lu, and Alexander Munteanu;
licensed under Creative Commons License CC-BY 4.0

34th Annual European Symposium on Algorithms (ESA 2026).

Editors: Philip Bille, Seth Pettie, and Sabine Storandt; Article No. 116; pp. 116:1–116:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is typically bounded within a factor of $(1 \pm \varepsilon)$. Another area where dimension reduction plays a crucial role is the construction of so-called *coresets* [7, 2]. Coresets aim to reduce the number of data objects rather than their dimension. Classic coreset constructions typically had a small, though polynomial dependence on d , depending on the problem for which they were developed. Sometimes exponential dependencies on the dimension were present or even unavoidable [1, 14, 31]. More recently, for some problems such as k -means/median clustering, the dependence of the coreset size on the dimension could be completely eliminated [37, 23], which would not have been imaginable without (random) projection techniques such as the celebrated Johnson-Lindenstrauss (JL) Lemma [25], PCA [26], and their descendants. Extending these techniques to polygonal curves and surfaces is an intriguing open direction, as even the best coresets [16, 17] in this regime suffered from linear dimension dependence.

JL-type embeddings are key to dimension-reduced results in other areas as well: for instance in compressed sensing, one receives the promise that a d -dimensional dense and noisy signal is inherently k -sparse and the task is to reconstruct that signal with as few (random) measurements as possible. Here, the number of measurements corresponds to a reduced target dimension which is only $\Theta(k \log(d/k))$ [18, 13, 8].

Dimension reduction tools have recently brought up dimension-free bounds for learning problems, where we would like to approximate some problem-specific quantity defined as the expectation taken over an unknown distribution of points or polygonal curves [10, 27]. The task is to sample a small (dimension-independent) number of data objects from the distribution such that the sample mean over the quantity is within a small error from its expectation. In contrast to coresets, the ground set is not necessarily finite and can only be accessed via i.i.d. samples from the unknown distribution, which complicates the problem and requires some natural assumptions to succeed [4, 3].

As a first step towards opening these avenues for problems on polygonal curves and surfaces, we propose a fairly general framework for embedding these objects in lower dimensional space, while preserving their “distance” as measured by popular (continuous) metrics, dissimilarity measures and divergences such as the Fréchet metric, dynamic time warping (DTW), Hausdorff distance, and their variants. In the following, we present our main results, techniques and related work, then prove our technical claims before concluding with a summary, certain limitations and possible extensions of our work.

1.1 Results

We present our main result, which is a general framework for dimension reduction of parameterized piecewise γ -dimensional linear surfaces. Notably, this includes polygonal curves as special case for $\gamma = 1$. For the notation and definitions of terms such as traversals, we refer to Section 2. For a more detailed explanation of the Theorem, we refer to Section 4.

► **Theorem 1.** *Let $\sigma_1, \dots, \sigma_n : [0, 1]^\gamma \rightarrow \mathbb{R}^d$ be n piecewise linear surfaces with m pieces. Let $\mathcal{T}$ be a set of admissible traversals between surfaces and let μ be a measure on $[0, 1]^\gamma$. There exists a random linear mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2}(\gamma + \log(mn/\delta)))$ such that, with probability at least $1 - \delta$, it holds for all $1 \leq i < j \leq n$ that*

$$(1 - \varepsilon) d(\sigma_i, \sigma_j) \leq d(f(\sigma_i), f(\sigma_j)) \leq (1 + \varepsilon) d(\sigma_i, \sigma_j),$$

where the distance d between surfaces σ_i and σ_j is defined as

$$d(\sigma_i, \sigma_j) = \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0, 1]^\gamma} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2^q d\mu(r) \right)^{1/q}.$$

The distance is highly general and includes the most popular continuous dissimilarity measures as special cases: Fréchet (for $q \rightarrow \infty$), DTW (for $q = 1$), and its variants such as weak Fréchet, q -DTW for $q \in [0, \infty)$, etc. While Hausdorff distance is not directly subsumed by the generalized dissimilarity measure, an extension is immediate from our results.

Discrete variants are also included by choosing the measure μ of integration to be a counting measure. Discrete traversals for polygonal curves are well-known, but no such notion is known for surfaces; in particular there is no natural ordering of the vertices. We thus propose a set-valued generalization of discrete traversals to piecewise linear surfaces and prove that our generalization is consistent with the standard definition in the case of polygonal curves.

1.2 Related work

Random projections for curves under Fréchet distance. Early work [19] studied probabilistic embeddings of the Fréchet distance between pairs of polygonal curves of complexity bounded by at most m onto random lines, resulting in large multiplicative distortions of order $\Theta(m)$. Allowing larger embedding dimensions, we note that preserving the *discrete* Fréchet and DTW distances of n curves is straightforward from the JL lemma: it suffices to preserve all m^2 pairwise distances of vertices between $\binom{n}{2}$ pairs of curves, because every discrete traversal that we consider in a maximization or sum, consists of a subset of these distances. The embedding dimension is thus $O(\varepsilon^{-2} \log(nm))$.

Maybe surprisingly, the same embedding dimension suffices to preserve the continuous case as well. A first step was made in [29] who achieved mixed relative and additive errors using the same embedding of all vertices. A relative $(1 \pm \varepsilon)$ error guarantee was proven more recently in [34] using a complicated argument based on logical predicates to decide whether the Fréchet distance is within some value δ . These predicates check the aforementioned pairwise distances between vertices and additionally include distances to a couple of anchor points along edges to avoid certain complications with very long edges that caused additive errors in the previous analysis of [29]. The set of all vertices and anchor points builds a discrete “skeleton” for the Fréchet distance. The entire set of skeleton points is added to the embedding and since their number remains polynomially bounded in both n and m , this does not increase the asymptotical $O(\varepsilon^{-2} \log(nm))$ target dimension of the embedding. The journal version [35, Section 1.2.1] included an informal note on oblivious subspace embeddings similar to our work. They did not provide a proof and only considered the Fréchet distance, whereas our techniques extend and generalize to other distance measures and to surfaces.

Oblivious subspace embeddings and sparsity. The finite set embedding guarantee of the celebrated JL lemma [25] was first extended to preserve k -dimensional linear subspaces in the compressed sensing literature [18, 13, 8], where we get the promise that some signal is k -sparse. We note that this only requires an embedding of axis aligned subspaces, whose bases are subsets comprising k elements of the standard basis in $\mathbb{R}^d$. This is often referred to as the restricted isometry property (RIP). The next step towards arbitrary k -dimensional subspaces represented in arbitrary bases¹ was made by Sarlós [36], who paved the way for the fastest approximation algorithms for numerical linear algebra [15, 38] with a plethora of applications in Machine Learning, cf. [20, 30]. Finally, the step back to k -dimensional

¹ While an explicit basis is useful for conducting formal proofs, the embeddings are *oblivious* and do not require an explicit representation to work; the existence of a subspace is sufficient to embed it.

subspaces with a k -sparse support in an arbitrary, not necessarily axis-aligned basis, was recently studied in the context of sparse regression [28]. Our contributions show how to exploit these works for dimension reduction for polygonal curves and piecewise linear surfaces.

1.3 Technical overview

Warm up: dimension reduction for the Fréchet distance. Our investigations start with the continuous Fréchet distance as a warm up. We follow a different approach than recent work [34, 35] detailed in Section 1.2 on related work. Instead of discretizing the polygonal curves directly to preserve the Fréchet distance and applying the mapping f asserted by the JL lemma to the discrete set, we interpret the Euclidean distances between *any* points p, q on either curve as norms of their connecting vectors $v = p - q$. Our problem reduces to preserving the norm of all possible vectors v . Now, the multiplicative approximation guarantee that we get from the JL lemma is

$$(1 - \varepsilon)\|v\|_2 \leq \|f(v)\|_2 \leq (1 + \varepsilon)\|v\|_2$$

which is equivalent to a (1-dimensional) so-called subspace embedding, which guarantees that

$$\forall x \in \mathbb{R}: (1 - \varepsilon)\|vx\|_2 \leq \|f(vx)\|_2 \leq (1 + \varepsilon)\|vx\|_2.$$

The latter guarantee can be generalized to preserve arbitrary k -dimensional subspaces based on the JL lemma with embedding dimension $t \in O(\varepsilon^{-2}(k + \log(1/\delta)))$ where δ denotes the failure probability [36, 38]. Let $V \in \mathbb{R}^{d \times k}$ be a matrix whose columns build a basis for a k dimensional linear subspace of $\mathbb{R}^d$. Then choosing a properly rescaled i.i.d. Gaussian matrix $S \in \mathbb{R}^{t \times d}$ yields the guarantee that with probability at least $1 - \delta$ it satisfies [36, 38]

$$\forall x \in \mathbb{R}^k: (1 - \varepsilon)\|Vx\|_2 \leq \|SVx\|_2 \leq (1 + \varepsilon)\|Vx\|_2.$$

Now, we reduce the task of preserving the Fréchet distance via Gaussian matrices to preserving a *collection* of k -dimensional subspace embeddings for some $k \in O(1)$ as follows. We consider two polygonal curves π and τ in $\mathbb{R}^d$. We observe that any pair of points $p \in \pi$ and $q \in \tau$, not necessarily vertices, can be expressed as a convex combination of their two neighboring vertices. The vector $v = p - q$ may thus be expressed as a linear combination of 4 vertices. In fact *any* vector determining the Fréchet distance lies in some 4-dimensional linear subspace. Preserving the collection of all such “4-sparse” subspaces will thus suffice.

Now, consider a set of n curves with vertices in $\mathbb{R}^d$ of complexity at most m . We let $A \in \mathbb{R}^{d \times nm}$ be the matrix whose columns comprise all up to nm vertices of all input curves. Then any norm we are interested in, takes the form $\|Ax\|_2$, where we get the promise that at most $k = 4$ many coordinates of x are non-zero at a time. We can now construct a k -sparse subspace embedding [28] that satisfies

$$\forall x \in \Psi_k: (1 - \varepsilon)\|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon)\|Ax\|_2,$$

where $\Psi_k = \{x \in \mathbb{R}^{nm} \mid \|x\|_0 \leq k\}$. Recall that one single k -dimensional subspace embedding can be achieved with target embedding dimension $O(\varepsilon^{-2}(k + \log(1/\delta)))$. A simple combinatorial calculation yields that there exist at most $\binom{nm}{k} \leq (\frac{enm}{k})^k$ choices. Using the union bound it thus suffices to choose a target dimension of $t \in O(\varepsilon^{-2}k \log(nm/(\delta k))) \subseteq O(\varepsilon^{-2} \log(nm/\delta))$, since $k = 4$ is an absolute constant.

Finally, by an exchange argument between the optimal traversals w.r.t. the original and w.r.t the embedded curves, and using the subspace embedding to switch between original and low-dimensional curves, we show that the Fréchet distance is preserved up to $(1 \pm \varepsilon)$ factors.

Extension beyond Fréchet. We observe that the dimension reduction and the final exchange argument are much more general. Indeed, preserving an infimum over integrals for the continuous DTW in the same way seems even more direct than preserving the inf/sup combination for the Fréchet distance. Regardless of the difficulty of the exchange argument, by embedding the collection of 4-sparse subspaces, we have actually embedded all Euclidean distances that can ever appear in a plethora of dissimilarity measures that are variants of Fréchet and DTW. Even others are included, such as Hausdorff distance, which has a different formal definition but is again based on a vector space equipped with the Euclidean norm.

Piecewise linear surfaces. Our techniques can be extended to handle piecewise linear surfaces of dimension γ , where polygonal curves constitute the special case $\gamma = 1$. Our argument works as follows. Consider two piecewise linear surfaces. Let p, q be points on some linear piece of either of the two surfaces. By Carathéodory's theorem there exists a convex combination of vertices v_i such that $p = \sum_{i=1}^{\gamma+1} \lambda_i v_i$. The same holds for q , so the difference vector $p - q$ can be expressed as a linear combination of $k = 2(\gamma + 1)$ vertices. To preserve the norms of all such linear combinations, we can again apply the k -sparse subspace embedding of [28] with target embedding dimension $O(\varepsilon^{-2}(\gamma \log(nm/(\delta\gamma))))$. In our technical proofs in Section 4, we exploit more geometric structure of our problem to reduce this to $O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$.

Our generalized dissimilarity measure again serves as a general framework for the surface variants of continuous q -DTW, including the limiting cases $q = 1$ (continuous DTW), and $q \rightarrow \infty$ (Fréchet). Hausdorff distance can also be handled in a similar way as for curves. Continuous reparameterizations between two surfaces can be handled using homeomorphisms.

Discrete dissimilarity measures. Dimension reduction for discrete dissimilarity measures for curves is folklore, and can be handled by a straightforward application of the JL lemma to all pairwise distances between vertices of two curves. It is subsumed in our framework, and can be implemented by using a counting measure instead of a continuous measure μ for integration. For surfaces, however, it is unclear how a discrete traversal should be defined. We thus handle this issue by taking a detour over a continuous homeomorphism h , mapping vertices to one another whenever h or h^{-1} maps a vertex to the Voronoi cell of the other. For a fixed choice $h: A \rightarrow B$, we consider the Voronoi cells $V(a_i), V(b_i)$ of the vertices $a_1, \dots, a_m$ respectively $b_1, \dots, b_m$ on each surface. Now, whenever $h(a_i) \in V(b_j)$, we add the pair $(a_i, b_j) \in T$ without duplicates. Similarly, we also consider the inverse homeomorphism and add $(a_j, b_i) \in T$ without duplicates whenever $h^{-1}(b_i) \in V(a_j)$. We prove that a discrete traversal constructed in this way is consistent with the standard definition of discrete traversals for polygonal curves. Ranging over all homeomorphisms, we thus enumerate the set of all possible discrete traversals $T \in \mathcal{T}$ between surfaces or curves.

2 Preliminaries

Notation. We set $[n] := \{1, \dots, n\}$. For a vector $v \in \mathbb{R}^d$, and $p \in (0, \infty)$, we denote its ℓ_p -norm by $\|v\|_p = \left(\sum_{i=1}^d |v_i|^p\right)^{1/p}$. We remark that the cases $0 < p < 1$ only define quasi-norms. Other special cases are $\|v\|_\infty := \lim_{p \rightarrow \infty} \|v\|_p = \max_{i \in [d]} |v_i|$, and $\|v\|_0 = |\{i \in [d] \mid v_i \neq 0\}|$. We denote by $\Psi_k := \{x \in \mathbb{R}^d \mid \|x\|_0 \leq k\}$ the set of k -sparse vectors in $\mathbb{R}^d$.

Polygonal curves and surfaces. A *curve* in Euclidean space $\mathbb{R}^d$, is a continuous function $\sigma : [0, 1] \rightarrow \mathbb{R}^d$. A *polygonal curve* is a curve such that there exist a finite number $m \in \mathbb{N}$ of values $0 = s_1 \leq s_2 \leq \dots \leq s_m = 1$, with $v_i = \sigma(s_i)$ which we call *vertices*. We say that σ has *complexity* m . Further, for each $i \in \{1, \dots, m-1\}$ the two consecutive vertices $\sigma(s_i)$ and $\sigma(s_{i+1})$ are connected by an affine straight-line segment, i.e., it holds that

$$\sigma(s_i + \varsigma) = \left(1 - \frac{\varsigma}{s_{i+1} - s_i}\right) \cdot \sigma(s_i) + \frac{\varsigma}{s_{i+1} - s_i} \cdot \sigma(s_{i+1}),$$

for all $\varsigma \in [0, s_{i+1} - s_i]$. For simplicity of notation, we can fully characterize a polygonal curve σ by defining the sequence of its vertices $\sigma = (v_1, \dots, v_m)$.

Analogously, a piecewise linear *surface* in $\mathbb{R}^d$ is a continuous function $\sigma : [0, 1]^\gamma \rightarrow \mathbb{R}^d$ whose image consist of m linear pieces which are polygons/polyhedrons bounded by a total of N vertices $v_i \in \mathbb{R}^d$, such that every piece lies in a γ -dimensional affine subspace of $\mathbb{R}^d$.

Traversals and dissimilarity measures. We work with discrete and continuous “transformations”. One way to define them is the notion of *traversals*. Let $\sigma = (u_1, \dots, u_{m'})$ and $\tau = (v_1, \dots, v_{m''})$ be two curves of complexities m' and m'' respectively. Traversals are usually defined by inductively constructed sequences of point pairs (p, q) such that $p \in \sigma$ and $q \in \tau$. It will be helpful for the sake of generalization to surfaces, to work with equivalent (unordered) sets of point pairs instead. We prove their equivalence in Appendix A.

Then, a discrete *traversal* of σ and τ is a set T that consists of pairs of vertices (p, q) with $p \in \sigma$ and $q \in \tau$, such that

1. $(u_1, v_1) \in T \wedge (u_{m'}, v_{m''}) \in T$,
2. $\forall i < m', j < m'': (u_i, v_j) \in T_h \Rightarrow (u_i, v_{j+1}) \in T_h \vee (u_{i+1}, v_j) \in T_h \vee (u_{i+1}, v_{j+1}) \in T_h$,
3. $(u_i, v_{j+1}) \in T_h \Rightarrow (u_{i+1}, v_j) \notin T_h$.

We remark that the property of discrete traversals of not being allowed to backtrack, as common in the literature, is encapsulated by the third criterion. However, the equivalence between defining traversals as sequences of point pairs or as sets of point pairs holds even if traversals are allowed to backtrack: in this case, the set definition can be easily adapted. The notion of discrete traversals can be extended to the continuous case by specifying pairs of “reparameterizations”, which are functions $\alpha, \beta : [0, 1]^\gamma \rightarrow [0, 1]^\gamma$. Then we let the induced traversal be the set $T = \{(\sigma(\alpha(r)), \tau(\beta(r))) \mid r \in [0, 1]^\gamma\}$. Sometimes it is appropriate to impose further restrictions on α, β such as injectivity, surjectivity, continuity, monotonicity, as required, and denote $\mathcal{T} \subseteq \{(\alpha, \beta) \mid \alpha, \beta : [0, 1]^\gamma \rightarrow [0, 1]^\gamma\}$ the corresponding subset of *admissible* traversals.

Given two curves σ, τ , we now define a general dissimilarity measure $d(\sigma, \tau)$ between surfaces including curves as the special case where $\gamma = 1$. We let

$$d(\sigma, \tau) := \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0, 1]^\gamma} \|\sigma(\alpha(r)) - \tau(\beta(r))\|_2^q d\mu(r) \right)^{1/q}.$$

This dissimilarity measure subsumes a plethora of popular distances. For instance, the Fréchet distance, $d_{\text{F}}(\sigma, \tau) = \inf_{(\alpha, \beta)} \max_{r \in [0, 1]} \|\sigma(\alpha(r)) - \tau(\beta(r))\|_2$, is implemented by the following choices: we take the limit of $q \rightarrow \infty$, in which case the $\mathcal{L}^q$ integration becomes the supremum, and in case of compact shapes, even simply the maximum. A uniform measure μ over $[0, 1]^\gamma$ is chosen. For curves, the set of admissible traversals are restricted to be pairs of continuous, non-decreasing surjections (in which case the infimum is attained). An alternative that generalizes to surfaces is to take homeomorphisms. The infimum is then attained in the limit of a sequence of homeomorphisms.

For continuous DTW, we set $q = 1$. The other items are similar to the Fréchet distance except that for technical reasons, we would need to extend the interval $[0, 1]$ to $[0, p + q]$, where p, q equal the arc-lengths of σ, τ respectively, see [12] for details. We stress that these technical details do not matter for our dimension reduction. Discrete variants can be realized by choosing a counting measure for μ , which transforms the integral to a finite sum.

Subspace embeddings. The notion of ε -subspace embeddings [36] is central to our work. The following version is due to [38] and is slightly modified to fit our notation and setting.

► **Theorem 2** (ε -Subspace Embedding, Theorem 2.3 of [38]). *Let $A \in \mathbb{R}^{d \times k}$. There exists a distribution over random matrices $S \in \mathbb{R}^{t \times d}$ with $t \in O(\varepsilon^{-2}(k + \log(1/\delta)))$ such that it holds with probability $1 - \delta$ that*

$$\forall x \in \mathbb{R}^d : (1 - \varepsilon)\|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon)\|Ax\|_2.$$

We note that the matrix S can be chosen such that every $S_{i,j} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1/t)$. The same result can be achieved by rescaled random sign (Rademacher) matrices, which is more convenient in streaming and other memory constrained settings [15]. A crucial property is that the construction of S is *oblivious* to the matrix A . That means we can apply it to preserve subspaces of dimension k for which we do not have an explicit representation or whose basis column vectors are hidden in a superset of columns. The existence of a subspace suffices for the guarantee to apply, cf. [38]. The idea behind the subspace embedding construction of [36] is that the unit ball in the k -dimensional subspace of $\mathbb{R}^d$ spanned by the columns of A can be covered by an ε -net of size $(1 + 2/\varepsilon)^k$ [33]. By a simple vector chaining argument covered in [38, Theorem 2.1], the net can be constructed with an absolute constant $\varepsilon_0 := 1/2$ instead of ε to still hold for every $\varepsilon \leq \varepsilon_0$. This yields a net $\mathcal{N}$ of size $|\mathcal{N}| = C^k$ for an absolute constant C . Applying the JL lemma with a union bound over this net yields a $(1 \pm \varepsilon)$ approximation to the norm of every point in the net within target dimension $O(\varepsilon^{-2} \log(|\mathcal{N}|/\delta)) \subseteq O(\varepsilon^{-2}(k \log(C) + \log(1/\delta))) \subseteq O(\varepsilon^{-2}(k + \log(1/\delta)))$. To complete the argument, every point in the ball has a point in the net at distance at most ε to which it can be related to achieve $(1 \pm O(\varepsilon))$ error. By linearity of the embedding and the multiplicative error guarantee, the argument extends from the unit ball to arbitrary radius by scaling.

We will exploit ε -subspace embeddings for *collections* of k -dimensional subspaces. To this end we use the following theorem of [28] slightly modified to fit our notation and setting.

► **Theorem 3** (k -Sparse ε -Subspace Embedding, Theorem 7 of [28]). *Let $A \in \mathbb{R}^{d \times N}$. There exists a distribution over random matrices $S \in \mathbb{R}^{t \times d}$ with $t \in O(\varepsilon^{-2}k \log(N/(\delta k)))$ such that it holds with probability $1 - \delta$ that*

$$\forall x \in \Psi_k : (1 - \varepsilon)\|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon)\|Ax\|_2.$$

The idea is that there are $\binom{N}{k} \leq (eN/k)^k$ choices of k non-zero coordinates out of N . Every choice corresponds to a linear subspace of dimension at most k that is spanned solely by the columns associated with non-zero entries in x . Using a union bound over all these choices, Theorem 2 applies with target dimension $t \in O(\varepsilon^{-2}k \log(N/(\delta k)))$ and yields $(1 \pm \varepsilon)$ distortion for the Euclidean norm of all vectors in the collection of k -sparse subspaces.

3 Fréchet distance between curves

3.1 Dimension reduction

In this section, we will prove the following proposition, which can be seen as the restriction of Theorem 1 to the continuous Fréchet distance. This result is already known [34, 35] as discussed in Section 1.2. However, our approach is different, in particular considerably simpler and shorter and it serves as a warm-up for proving Theorem 1 in full generality.

► **Proposition 4.** *Let $\sigma_1, \dots, \sigma_n : [0, 1] \rightarrow \mathbb{R}^d$ be n piecewise linear curves with m vertices. There exists a random linear mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2} \log(mn/\delta))$ such that, with probability at least $1 - \delta$ it holds for all $1 \leq i < j \leq n$ that*

$$(1 - \varepsilon) d_{\text{cF}}(\sigma_i, \sigma_j) \leq d_{\text{cF}}(f(\sigma_i), f(\sigma_j)) \leq (1 + \varepsilon) d_{\text{cF}}(\sigma_i, \sigma_j),$$

where $d_{\text{cF}}(\sigma_i, \sigma_j) = \inf_{(\alpha, \beta)} \max_{r \in [0, 1]} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2$ denotes the continuous Fréchet distance between curves.

Proof. The proof consists of two parts: first, we show that there exists a random linear mapping $S : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2} \log(mn/\delta))$ such that

$$(1 - \varepsilon) \|p - q\|_2 \leq \|S(p) - S(q)\|_2 \leq (1 + \varepsilon) \|p - q\|_2$$

for every pair of points $p \in \sigma_i, q \in \sigma_j$ on every pair of curves σ_i, σ_j . We emphasize that p and q do not have to be vertices of σ_i and σ_j . Then, we show that this leads to the desired approximation for the continuous Fréchet distance between σ_i and σ_j .

Consider any pair of curves σ_i, σ_j and any pair of points $p \in \sigma_i$ and $q \in \sigma_j$. Denote the vertices of σ_i and σ_j by $p_1, \dots, p_m$ and $q_1, \dots, q_m$, respectively. Observe that p and q can be expressed as convex combinations of their neighboring vertices, i.e., there exist $\lambda_p, \lambda_q \in [0, 1]$ and indices $h, \ell \in [m - 1]$ such that $p = (1 - \lambda_p)p_h + \lambda_p p_{h+1}$ and $q = (1 - \lambda_q)q_\ell + \lambda_q q_{\ell+1}$. Then, $p - q = (1 - \lambda_p)p_h + \lambda_p p_{h+1} - (1 - \lambda_q)q_\ell - \lambda_q q_{\ell+1}$ is a linear combination of 4 vertices.

Hence, to preserve $\|p - q\|_2$ up to a factor $(1 \pm \varepsilon)$ for all $p \in \sigma_i, q \in \sigma_j$, it suffices to preserve the norm of all linear combinations of 4 vertices of σ_i and σ_j . Let $A \in \mathbb{R}^{d \times mn}$ be the matrix whose columns comprise all vertices of all curves $\sigma_1, \dots, \sigma_n$. Then, any linear combination of 4 vertices of σ_i and σ_j can be written as Ax for some 4-sparse vector $x \in \mathbb{R}^{mn}$. By setting $k = 4$ and applying Theorem 3 with $N = nm$, there exists a random linear mapping $S : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2} \log(mn/\delta))$ such that $(1 - \varepsilon) \|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon) \|Ax\|_2$ holds for all 4-sparse x . It follows for every pair of points $p \in \sigma_i, q \in \sigma_j$ that

$$(1 - \varepsilon) \|p - q\|_2 \leq \|S(p) - S(q)\|_2 \leq (1 + \varepsilon) \|p - q\|_2$$

holds for every pair of curves σ_i, σ_j , which is what we wanted to show.

Now, to show that S preserves the continuous Fréchet distance between any pair of curves σ_i and σ_j , let (α^*, β^*) and (α^S, β^S) be reparameterizations realizing $d_{\text{cF}}(\sigma_i, \sigma_j)$ and $d_{\text{cF}}(S(\sigma_i), S(\sigma_j))$ respectively. Let $r' \in \arg \max_{r \in [0, 1]} \|S(\sigma_i(\alpha^*(r))) - S(\sigma_j(\beta^*(r)))\|_2$. Then,

$$\begin{aligned} d_{\text{cF}}(S(\sigma_i), S(\sigma_j)) &= \inf_{(\alpha, \beta)} \max_{r \in [0, 1]} \|S(\sigma_i(\alpha(r))) - S(\sigma_j(\beta(r)))\|_2 \\ &\leq \max_{r \in [0, 1]} \|S(\sigma_i(\alpha^*(r))) - S(\sigma_j(\beta^*(r)))\|_2 \\ &= \|S(\sigma_i(\alpha^*(r'))) - S(\sigma_j(\beta^*(r')))\|_2 \\ &\leq (1 + \varepsilon) \|\sigma_i(\alpha^*(r')) - \sigma_j(\beta^*(r'))\|_2 \end{aligned}$$

$$\begin{aligned}
&\leq (1 + \varepsilon) \max_{r \in [0,1]} \|\sigma_i(\alpha^*(r)) - \sigma_j(\beta^*(r))\|_2 \\
&= (1 + \varepsilon) \inf_{(\alpha, \beta)} \max_{r \in [0,1]} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2 \\
&= (1 + \varepsilon) d_{\text{cF}}(\sigma_i, \sigma_j),
\end{aligned}$$

which proves the dilation bound. Next, let $r'' \in \arg \max_{r \in [0,1]} \|\sigma_i(\alpha^S(r)) - \sigma_j(\beta^S(r))\|_2$. Then it holds that

$$\begin{aligned}
d_{\text{cF}}(S(\sigma_i), S(\sigma_j)) &= \inf_{(\alpha, \beta)} \max_{r \in [0,1]} \|S(\sigma_i(\alpha(r))) - S(\sigma_j(\beta(r)))\|_2 \\
&= \max_{r \in [0,1]} \|S(\sigma_i(\alpha^S(r))) - S(\sigma_j(\beta^S(r)))\|_2 \\
&\geq \|S(\sigma_i(\alpha^S(r''))) - S(\sigma_j(\beta^S(r''))) \|_2 \\
&\geq (1 - \varepsilon) \|\sigma_i(\alpha^S(r'')) - \sigma_j(\beta^S(r''))\|_2 \\
&= (1 - \varepsilon) \max_{r \in [0,1]} \|\sigma_i(\alpha^S(r)) - \sigma_j(\beta^S(r))\|_2 \\
&\geq (1 - \varepsilon) \inf_{(\alpha, \beta)} \max_{r \in [0,1]} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2 \\
&= (1 - \varepsilon) d_{\text{cF}}(\sigma_i, \sigma_j),
\end{aligned}$$

which proves the contraction bound. This finishes the proof. $\blacktriangleleft$

4 Distance measures between piecewise linear surfaces

In this section, we first extend our framework for dimension reduction to distances between surfaces. Then, we move to discrete distance measures by giving a consistent definition of discrete traversals for surfaces.

4.1 Dimension reduction

In this subsection, we prove Theorem 1. The outline of our proof is very similar to the special case of the Fréchet distance treated in Section 3. However, to achieve the target dimension bound of $O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$, it is not sufficient to apply the sparse subspace embedding given in Theorem 3. Instead, we exploit the geometric structure of the problem to avoid considering the same γ -dimensional subspaces multiple times and other overcounting issues that did not matter in the case of polygonal curves, where $\gamma \in O(1)$. Moreover, standard reparameterizations are insufficient to handle the case of surfaces. We thus work with homeomorphisms, which require to handle the infimum via sequences of homeomorphisms, because, unlike reparameterizations, there are cases where the infimum is not attained.

► **Theorem 1.** *Let $\sigma_1, \dots, \sigma_n : [0, 1]^\gamma \rightarrow \mathbb{R}^d$ be n piecewise linear surfaces with m pieces. Let $\mathcal{T}$ be a set of admissible traversals between surfaces and let μ be a measure on $[0, 1]^\gamma$. There exists a random linear mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2}(\gamma + \log(mn/\delta)))$ such that, with probability at least $1 - \delta$, it holds for all $1 \leq i < j \leq n$ that*

$$(1 - \varepsilon) d(\sigma_i, \sigma_j) \leq d(f(\sigma_i), f(\sigma_j)) \leq (1 + \varepsilon) d(\sigma_i, \sigma_j),$$

where the distance d between surfaces σ_i and σ_j is defined as

$$d(\sigma_i, \sigma_j) = \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0,1]^\gamma} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2^q d\mu(r) \right)^{1/q}.$$

116:10 Dimension Reduction for Curves: Simplified and Generalized

Proof. The proof again consists of two parts: first, we show that there exists a random linear mapping $S : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$ such that

$$(1 - \varepsilon)\|p - q\|_2 \leq \|S(p) - S(q)\|_2 \leq (1 + \varepsilon)\|p - q\|_2$$

for every pair of points $p \in \sigma_i, q \in \sigma_j$ on every pair of distinct piecewise linear surfaces σ_i, σ_j . Note that p and q do not have to be vertices of σ_i and σ_j . Then, we show that this leads to the desired approximation for the generalized distance measure between σ_i and σ_j .

Consider any pair of piecewise linear surfaces σ_i, σ_j and any pair of points $p \in \sigma_i$ and $q \in \sigma_j$. Observe that p lies in some linear piece P of σ_i that may be non-convex or contain holes. Consider the convex hull of P and a triangulation thereof into simplices whose vertices are subsets of $\gamma + 1$ vertices of σ_i .

Now, p lies in one of these simplices and by Carathéodory's theorem, p can be expressed as a convex combination of the vertices of this simplex. This means that there exist $\gamma + 1$ vertices $v_1, \dots, v_{\gamma+1}$ of σ_i and coefficients $\lambda \in \Delta_{\gamma+1}$ such that $p = \sum_{i=1}^{\gamma+1} \lambda_i v_i$, where $\Delta_{\gamma+1} := \{\lambda \in \mathbb{R}^{\gamma+1} \mid \sum_{i=1}^{\gamma+1} \lambda_i = 1, \forall i \in [\gamma + 1]: \lambda_i \geq 0\}$ is the probability simplex in $\gamma + 1$ dimensions. The same argumentation holds for q and thus we have $\lambda, \eta \in \Delta_{\gamma+1}$, and vertices $v_1, \dots, v_{\gamma+1}$ of σ_i and vertices $w_1, \dots, w_{\gamma+1}$ of σ_j such that

$$\|p - q\|_2 = \left\| \sum_{i=1}^{\gamma+1} \lambda_i v_i - \sum_{i=1}^{\gamma+1} \eta_i w_i \right\|_2.$$

Hence, the difference vector can be expressed as a linear combination of $2(\gamma + 1)$ vertices.

At this point, we could again apply the sparse subspace embedding, given in Theorem 3, to the matrix $A \in \mathbb{R}^{d \times N}$ carrying the collection of all $N \geq nm(\gamma + 1)$ vertices as its columns, where the sparsity bound parameter is $k = 2(\gamma + 1)$. This would result in an embedding dimension $t \in O(\varepsilon^{-2}(\gamma \log(N/(\delta\gamma))))$. In particular, N , the number of vertices could be much larger than $nm(\gamma + 1)$, and additionally we wish to reduce the factor γ outside the logarithm to an additive γ . Our ultimate goal is to reduce from d to $t \in O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$ dimensions.

To this end, we need to eliminate some overcounting in our previous argument. Observe that every linear piece lies in an affine subspace which requires a basis of only $\gamma + 1$ columns to be preserved. In particular *any* non-degenerate set of vertices of one single simplex will suffice as a representative. If we consider any other simplex within the same piece, it will span the same subspace. The only change is that instead of convex combinations $\lambda, \eta \in \Delta_{\gamma+1}$ of these vertices, we need to consider linear combinations $\lambda, \eta \in \mathbb{R}^{\gamma+1}$ to express p resp. q . The resulting distance vector $p - q$ is again a linear combination of at most $k = 2(\gamma + 1)$ vertices, half of which are located on a linear piece on either surface σ_i resp. σ_j .

By our refined argument, the number of k -dimensional linear subspaces we need to consider is thus bounded by $\binom{n}{2}$ pairs of surfaces, times m^2 pairings of linear pieces. Thus, we can apply Theorem 2 with $k = 2(\gamma + 1)$ and union bound over $M \in O(n^2 m^2)$ choices by adjusting the failure probability to $\delta' = \delta/M$. Then, we obtain an embedding dimension of $t \in O(\varepsilon^{-2}(k + \log(M/\delta))) \subseteq O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$.

Let $A \in \mathbb{R}^{d \times N}$ be the matrix whose N columns comprise all vertices of all piecewise linear surfaces $\sigma_1, \dots, \sigma_n$. The above argumentation yields the existence of a suitable collection of $2(\gamma + 1)$ -sparse vectors $\Xi_{2(\gamma+1)} \subset \Psi_{2(\gamma+1)}$ and a linear mapping $S : \mathbb{R}^d \rightarrow \mathbb{R}^t$ with $t \in O(\varepsilon^{-2}(\gamma + \log(nm/\delta)))$ such that $\forall x \in \Xi_{2(\gamma+1)}: (1 - \varepsilon)\|Ax\|_2 \leq \|SAx\|_2 \leq (1 + \varepsilon)\|Ax\|_2$. This finally implies that

$$(1 - \varepsilon)\|p - q\|_2 \leq \|S(p) - S(q)\|_2 \leq (1 + \varepsilon)\|p - q\|_2$$

holds for every pair of points $p \in \sigma_i, q \in \sigma_j$ on every pair of piecewise linear surfaces σ_i, σ_j , for all $1 \leq i < j \leq n$, which is what we wanted to show.

Now, to show that S preserves the generalized distance measure between any given pair of piecewise linear surfaces σ_i and σ_j , let $\{(\alpha_l^*, \beta_l^*)\}_{l \in \mathbb{N}}$ and $\{(\alpha_l^S, \beta_l^S)\}_{l \in \mathbb{N}}$ be sequences of homeomorphism pairs whose limits determine the values of $d(\sigma_i, \sigma_j)$ and $d(S(\sigma_i), S(\sigma_j))$ respectively. Then,

$$\begin{aligned}
 d(S(\sigma_i), S(\sigma_j)) &= \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0,1]^\gamma} \|S(\sigma_i(\alpha(r))) - S(\sigma_j(\beta(r)))\|_2^q d\mu(r) \right)^{1/q} \\
 &\leq \lim_{l \rightarrow \infty} \left(\int_{[0,1]^\gamma} \|S(\sigma_i(\alpha_l^*(r))) - S(\sigma_j(\beta_l^*(r)))\|_2^q d\mu(r) \right)^{1/q} \\
 &\leq \lim_{l \rightarrow \infty} \left(\int_{[0,1]^\gamma} (1 + \varepsilon)^q \|\sigma_i(\alpha_l^*(r)) - \sigma_j(\beta_l^*(r))\|_2^q d\mu(r) \right)^{1/q} \\
 &= (1 + \varepsilon) \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0,1]^\gamma} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2^q d\mu(r) \right)^{1/q} \\
 &= (1 + \varepsilon) d(\sigma_i, \sigma_j),
 \end{aligned}$$

which proves the dilation bound. Similarly,

$$\begin{aligned}
 d(S(\sigma_i), S(\sigma_j)) &= \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0,1]^\gamma} \|S(\sigma_i(\alpha(r))) - S(\sigma_j(\beta(r)))\|_2^q d\mu(r) \right)^{1/q} \\
 &= \lim_{l \rightarrow \infty} \left(\int_{[0,1]^\gamma} \|S(\sigma_i(\alpha_l^S(r))) - S(\sigma_j(\beta_l^S(r)))\|_2^q d\mu(r) \right)^{1/q} \\
 &\geq \lim_{l \rightarrow \infty} \left(\int_{[0,1]^\gamma} (1 - \varepsilon)^q \|\sigma_i(\alpha_l^S(r)) - \sigma_j(\beta_l^S(r))\|_2^q d\mu(r) \right)^{1/q} \\
 &\geq (1 - \varepsilon) \inf_{(\alpha, \beta) \in \mathcal{T}} \left(\int_{[0,1]^\gamma} \|\sigma_i(\alpha(r)) - \sigma_j(\beta(r))\|_2^q d\mu(r) \right)^{1/q} \\
 &= (1 - \varepsilon) d(\sigma_i, \sigma_j),
 \end{aligned}$$

which proves the contraction bound. This finishes the proof. $\blacktriangleleft$

4.2 Discrete traversals for piecewise linear surfaces

In this subsection, we provide a notion of discrete traversals for surfaces and show that this definition is consistent with discrete traversals for curves. It thus suits the requirements for generalizing the discrete Fréchet distance, and variants of discrete DTW to surfaces.

In contrast with curves, for which there are definitions and algorithms for both the continuous and the discrete Fréchet distance, for surfaces only a continuous Fréchet distance has been defined in the literature [5]. Given two piecewise linear surfaces $f, g : \mathbb{R}^\gamma \rightarrow \mathbb{R}^d$, the continuous Fréchet distance between f and g is defined as

$$\inf_{h: \mathbb{R}^\gamma \rightarrow \mathbb{R}^\gamma} \sup_{x \in \mathbb{R}^\gamma} \|f(x) - g(h(x))\|_2,$$

where the infimum is taken over all orientable homeomorphisms h from $\mathbb{R}^\gamma$ to itself. This is consistent with the previous definition using reparameterizations by setting h such that $h \circ \alpha = \beta$. For simplicity of notation, we will treat $h: \mathbb{R}^d \rightarrow \mathbb{R}^d$ as a mapping between the ambient spaces of f and g instead of their parameter spaces.

Recall that in the definition of discrete Fréchet and DTW distances for curves, a minimum is taken over the set of discrete traversals. Hence, a definition for discrete Fréchet and DTW distance between surfaces could be obtained by defining a suitable notion of discrete traversals between surfaces. However, unlike curves, the vertices on a piecewise linear surface do not have a natural ordering that can be used to define a discrete traversal as a sequence. Instead, we define discrete traversals starting from an orientable homeomorphism.

We assume that each surface comes with a fixed “starting” point, say u^* for f and v^* for g and restrict ourselves to homeomorphisms h satisfying $h(u^*) = v^*$. Let $u_1, \dots, u_{m'} \in \mathbb{R}^d$ and $v_1, \dots, v_{m''} \in \mathbb{R}^d$ be the sets of vertices of the surfaces f and g , respectively. A discrete traversal T_h of the surfaces is a set that consists of pairs of vertices, induced by h in the following way. For each vertex u_i of f , consider the Voronoi cell $V(u_i)$ of u_i within f . Similarly, consider the Voronoi cells $V(v_j)$ for the vertices v_j of g . We add the pair (u_i, v_j) to the traversal T_h (without duplicates) if and only if $h(u_i) \in V(v_j)$ or $h^{-1}(v_j) \in V(u_i)$, where ties lying on common boundaries of two or more distinct Voronoi cells are broken in lexicographic order. We prove that a traversal constructed in this way is consistent with the standard definition of a discrete traversal in the special case of polygonal curves.

► **Proposition 5.** *Let $f, g : [0, 1] \rightarrow \mathbb{R}^d$ be polygonal curves, let h be an orientable homeomorphism between f and g such that $h(f(0)) = g(0)$ and let T_h be the discrete traversal induced by h as defined above. Then, T_h is a discrete traversal between curves f and g . Additionally, for every discrete traversal T , there exists an orientable homeomorphism h such that $T = T_h$.*

Proof.

Consistency. Let $f(s_1) = f(0) = u_1, \dots, f(s_{m'}) = f(1) = u_{m'}$ and $g(t_1) = g(0) = v_1, \dots, g(t_{m''}) = g(1) = v_{m''}$ be the vertices of f and g , respectively. We need to show that

1. $(u_1, v_1) \in T_h \wedge (u_{m'}, v_{m''}) \in T_h$,
2. $\forall i < m', j < m'': (u_i, v_j) \in T_h \Rightarrow (u_i, v_{j+1}) \in T_h \vee (u_{i+1}, v_j) \in T_h \vee (u_{i+1}, v_{j+1}) \in T_h$,
3. $(u_i, v_{j+1}) \in T_h \Rightarrow (u_{i+1}, v_j) \notin T_h$.

For the first property, we have by assumption, $h(u_1) = v_1$, so in particular, $h(u_1) \in V(v_1)$, which means that $(u_1, v_1) \in T_h$ holds by definition. Now, because h is a homeomorphism, which maps the boundary ∂f of f bijectively to the boundary ∂g of g , and since $\partial f = \{u_1, u_{m'}\}$ and $\partial g = \{v_1, v_{m''}\}$, it follows that $h(u_{m'}) = v_{m''}$, so $(u_{m'}, v_{m''}) \in T_h$.

To prove the second property, assume that $(u_i, v_j) \in T_h$. This means that $h(u_i) \in V(v_j)$ or $h^{-1}(v_j) \in V(u_i)$. Without loss of generality, assume that $h(u_i) \in V(v_j)$. Suppose, for the sake of a contradiction, that $(u_i, v_{j+1}) \notin T_h$ and $(u_{i+1}, v_j) \notin T_h$ and $(u_{i+1}, v_{j+1}) \notin T_h$. This means that $h(u_{i+1}) \notin V(v_j) \cup V(v_{j+1})$ and $h^{-1}(v_{j+1}) \notin V(u_i) \cup V(u_{i+1})$. From $h(u_i) \in V(v_j)$ and the fact that h is a homeomorphism, it follows that $h(u_{i+1}) \in V(v_k)$ for some $k \geq j+2$. This implies that $h^{-1}(v_\ell) \in V(u_i) \cup V(u_{i+1})$ for all $\ell \in \{j+1, \dots, k-1\}$. In particular, $h^{-1}(v_{j+1}) \in V(u_i) \cup V(u_{i+1})$, which contradicts the statement before.

For the third property, let $(u_i, v_{j+1}) \in T_h$. This implies that either $h(u_i) \in V(v_{j+1})$ or $h^{-1}(v_{j+1}) \in V(u_i)$. Then there exist $r, s \in [0, 1]$ parameterizing $p = f(r), q = g(s)$ such that $h(p) = q$, and moreover $r \leq \bar{r} \in [0, 1], s \geq \bar{s} \in [0, 1]$, where $\bar{r}$ parameterizes the boundary between $V(u_i)$ and $V(u_{i+1})$, and $\bar{s}$ parameterizes the boundary between $V(v_j)$ and $V(v_{j+1})$, respectively. Assume for the sake of contradiction that $(u_{i+1}, v_j) \in T_h$. Then we have analogously $h(p') = q'$ with $p' = f(r'), q' = g(s')$ parameterized by $r' \geq \bar{r}$, and $s' \leq \bar{s}$. This, however, violates the continuity of either h or h^{-1} , or either $p = p'$ or $q = q'$ hold (if they lie exactly on the boundary). Any case contradicts the fact that h is a homeomorphism.

Existence of a consistent homeomorphism. Let us now turn to showing that every traversal T can be realized as $T = T_h$ for some homeomorphism h through our transformation. Consider a fixed traversal T and for each vertex w of either of the two curves, let $\Gamma(w) = \{p \mid (w, p) \in T \vee (p, w) \in T\}$ be its neighborhood in T . The idea of our construction is to map the entire neighborhood of a vertex w to its Voronoi cell $V(w)$ in the order induced by T . First, let $h(u_1) = v_1$, and $h(u_{m'}) = v_{m''}$. Note that since there are no crossing edges by Item 3 above, only one of $\Gamma(u_1), \Gamma(v_1)$ can contain more than one element and the two cases are symmetric. W.l.o.g. let $\Gamma(u_1)$ contain $n_1 > 1$ elements $\{v_1, v_2, \dots, v_{n_1}\}$, then we identify points $p_2, \dots, p_{n_1} \in V(u_1)$ that appear after u_1 in the order given by the indices and let $h^{-1}(v_i) = p_i$ for $i \in \{2, \dots, n_1\}$. In case the next element in T is (u_2, v_{n_1+1}) , we can “remove” all preceding vertices and start the process over. Otherwise, the next element is (u_2, v_{n_1}) and we “remove” preceding vertices. We distinguish two sub-cases: if $|\Gamma(v_{n_1})| \geq 1$, then we map all vertices to points in $V(v_{n_1})$ after v_{n_1} using h in the same way described above. Otherwise $|\Gamma(u_2)| \geq 1$, and we map $h(u_2) = p$ for a point $p \in V(v_{n_1})$ after v_{n_1} . All remaining vertices in $\Gamma(u_2)$ are mapped using h^{-1} to points in $V(u_2)$ after u_2 in the same way described above. We repeat this algorithm until we reach the pair $(u_{m'}, v_{m''})$. Finally, we complete the homeomorphism between the points mapped during this process in a continuous bijective way preserving the ordering of points in the image and pre-image. By construction, the traversal T_h induced by h equals T again. ◀

5 Further Considerations

Extension to ℓ_p . An interesting generalization of our framework is offered by the possibility of using ℓ_p normed Banach spaces as underlying vector spaces instead of the Euclidean space. This, however requires changing the i.i.d. Gaussian distribution of the embedding matrix to an i.i.d. p -stable distribution, which exists for any $p \in (0, 2]$, see [24, 21]. Even more crucial, while in the case of the ℓ_2 norm, the estimator is again the ℓ_2 norm, this does not continue to hold for other choices of $p \in (0, 2)$. Instead, we get the following guarantee from [28]

$$\forall x \in \Psi_k: (1 - \varepsilon)\|Ax\|_p \leq \|SAx\|_{\text{med}} \leq (1 + \varepsilon)\|Ax\|_p,$$

where for target dimension t , $\|y\|_{\text{med}} := \text{median}\{|y_i| \mid i \in [t]\}$.

Any algorithm that can be implemented to query only the ℓ_p distances of point pairs p, q located on different curves (or surfaces) will work unchanged up to $(1 \pm \varepsilon)$ errors using median queries. Thus an ℓ_p extension can benefit from our framework. However, algorithms that perform more complicated operations in the original dimension may need sophisticated modifications to work with the median estimator in the reduced dimension.

Algorithmic considerations. The continuous Fréchet and Hausdorff distances have efficient algorithms, see [6] and descendants thereof. Also discrete variants of Fréchet, DTW, Hausdorff, etc. can be computed efficiently and with linear dependencies on the dimension [22]. These algorithms directly benefit from our framework by reducing the dimension of the input curves and running the algorithms unchanged on the low dimensional embedded curves.

As we have outlined in the previous paragraph, ℓ_p extensions may require changes to the algorithm. For continuous DTW the situation is even worse, since algorithms for these distance are only known in the $O(1)$ -dimensional case [9, 12] (where dimension reduction makes no sense). We face similar limitations in the context of dissimilarity measures between surfaces: efficient algorithms are only known for simple polygons in $O(1)$ -dimensions [11]. The Fréchet distance for general surfaces is only known to be semi-computable [5]. However, for some of these issues, relaxations such as a *weak* Fréchet distance or practical heuristics are

available and benefit from dimension reduction while preserving their original performance up to $(1 \pm \varepsilon)$ errors. Even in cases where no algorithms are known, our framework provides dimension reduction guarantees that will continue to hold for future developments. For instance, one intriguing open direction would be to extend the algorithm for the Fréchet distance between simple polygons [11] (or more general classes of surfaces) to arbitrary d -dimensional space. Our framework then applies directly to remove any dependence on d .

6 Conclusion

We have introduced a widely general framework for dimension reduction for polygonal curves and piecewise linear surfaces, and discussed popular special cases as well as further extensions. We hope that our framework and methods will motivate and contribute to future developments and ultimately lead to dimension-reduced or dimension-independent coresets, dimension-free learning theoretical bounds, and finally improved algorithms for important proximity, classification and clustering problems for curves and surfaces. A first step in exploiting random projections with collections of low-dimensional subspaces has been made recently in [32]. The authors designed the first *terminal embeddings* for preserving the Fréchet distance of input curves to arbitrary curves in the ambient space. This also enabled the first dimension-free and near-optimal coresets for k -clustering under the Fréchet distance. Extensions to other distance measures and higher dimensional linear objects are intriguing directions for further research.

References

- 1 Pankaj K. Agarwal, Sarel Har-Peled, and Kasturi R. Varadarajan. Approximating extent measures of points. *J. ACM*, 51(4):606–635, 2004. doi:10.1145/1008731.1008736.
- 2 Pankaj K. Agarwal, Sarel Har-Peled, and Kasturi R. Varadarajan. Geometric approximation via coresets. *Combinatorial and computational geometry*, 52(1):1–30, 2005.
- 3 Meysam Alishahi, Alexander Munteanu, Simon Omlor, and Jeff M. Phillips. Optimal dimension-free sampling for regularized classification. *CoRR*, abs/2605.23726, 2026. doi:10.48550/arXiv.2605.23726.
- 4 Meysam Alishahi and Jeff M. Phillips. No dimensional sampling coresets for classification. In *41st International Conference on Machine Learning (ICML)*, pages 1008–1049, 2024. URL: <https://proceedings.mlr.press/v235/alishahi24a.html>.
- 5 Helmut Alt and Maike Buchin. Can we compute the similarity between surfaces? *Discrete & Computational Geometry*, 43(1):78–99, 2010. doi:10.1007/S00454-009-9152-8.
- 6 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.*, 5:75–91, 1995. doi:10.1142/S0218195995000064.
- 7 Mihai Bădoiu, Sarel Har-Peled, and Piotr Indyk. Approximate clustering via core-sets. In *Proc. of the 34th ACM Symposium on Theory of Computing (STOC)*, pages 250–257, 2002.
- 8 Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- 9 Milutin Brankovic, Kevin Buchin, Koen Klaren, André Nusser, Aleksandr Popov, and Sampson Wong. (k,l)-medians clustering of trajectories using continuous dynamic time warping. In *28th International Conference on Advances in Geographic Information Systems (SIGSPATIAL)*, pages 99–110, 2020. doi:10.1145/3397536.3422245.
- 10 Maria Sofia Bucarelli, Matilde Fjeldsø Larsen, Chris Schwiegelshohn, and Mads Toftrup. On generalization bounds for projective clustering. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, 2023. URL: http://papers.nips.cc/paper_files/paper/2023/hash/e30bf4765ae6b16a87fb4d7b0b3b3dec-Abstract-Conference.html.

- 11 Kevin Buchin, Maike Buchin, and Carola Wenk. Computing the Fréchet distance between simple polygons in polynomial time. In *Proceedings of the 22nd ACM Symposium on Computational Geometry, (SoCG)*, pages 80–87, 2006. doi:10.1145/1137856.1137870.
- 12 Kevin Buchin, André Nusser, and Sampson Wong. Computing continuous dynamic time warping of time series in polynomial time. In *38th International Symposium on Computational Geometry (SoCG)*, pages 22:1–22:16, 2022. doi:10.4230/LIPIcs.SOCG.2022.22.
- 13 Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- 14 Timothy M Chan. Faster core-set constructions and data stream algorithms in fixed dimensions. In *Proc. of the 20th Symposium on Computational Geometry, (SoCG)*, pages 152–159, 2004. doi:10.1145/997817.997843.
- 15 Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *Proc. of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 205–214, 2009. doi:10.1145/1536414.1536445.
- 16 Vincent Cohen-Addad, Andrew Draganov, Matteo Russo, David Saulpic, and Chris Schwiegelshohn. A tight VC-dimension analysis of clustering coresets with applications. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 4783–4808, 2025. doi:10.1137/1.9781611978322.162.
- 17 Jacobus Conradi, Benedikt Kolbe, Ioannis Psarros, and Dennis Rohde. Fast approximations and coresets for (k, ℓ) -median under dynamic time warping. In *40th Symposium on Computational Geometry, (SoCG)*, pages 42:1–42:17, 2024. doi:10.4230/LIPIcs.SOCG.2024.42.
- 18 David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. doi:10.1109/TIT.2006.871582.
- 19 Anne Driemel and Amer Krivosija. Probabilistic embeddings of the Fréchet distance. In *Approximation and Online Algorithms - 16th International Workshop, (WAOA)*, pages 218–237, 2018. doi:10.1007/978-3-030-04693-4_14.
- 20 Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. *J. Mach. Learn. Res.*, 13:3475–3506, 2012. doi:10.5555/2503308.2503352.
- 21 Alex Dytso, Ronit Bustin, H. Vincent Poor, and Shlomo Shamai. Analytical properties of generalized Gaussian distributions. *Journal of Statistical Distributions and Applications*, 5:1–40, December 2018.
- 22 Thomas Eiter and Heikki Mannila. Computing discrete Fréchet distance. Technical Report CD-TR 94/64, Christian Doppler Laboratory, 1994.
- 23 Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning Big Data into tiny data: Constant-size coresets for k-means, PCA, and projective clustering. *SIAM J. Comput.*, 49(3):601–657, 2020. doi:10.1137/18M1209854.
- 24 Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM*, 53(3):307–323, 2006. doi:10.1145/1147954.1147955.
- 25 William Johnson and Joram Lindenstrauss. Extensions of Lipschitz maps into a Hilbert space. *Contemporary Mathematics*, 26:189–206, January 1984.
- 26 Ravi Kannan and Santosh S. Vempala. Spectral algorithms. *Found. Trends Theor. Comput. Sci.*, 4(3-4):157–288, 2009. doi:10.1561/04000000025.
- 27 Amer Krivosija, Alexander Munteanu, André Nusser, and Chris Schwiegelshohn. Improved learning via k-DTW: A novel dissimilarity measure for curves. In *42nd International Conference on Machine Learning (ICML)*, pages 31696–31726, 2025. URL: <https://proceedings.mlr.press/v267/krivosija25a.html>.
- 28 Tung Mai, Alexander Munteanu, Cameron Musco, Anup Rao, Chris Schwiegelshohn, and David P. Woodruff. Optimal sketching bounds for sparse linear regression. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 11288–11316, 2023. URL: <https://proceedings.mlr.press/v206/mai23a.html>.

- 29 Stefan Meintrup, Alexander Munteanu, and Dennis Rohde. Random projections and sampling algorithms for clustering of high-dimensional polygonal curves. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, pages 12807–12817, 2019. URL: <https://proceedings.neurips.cc/paper/2019/hash/0394ea68951e3299bcdafa75a097d7c11-Abstract.html>.
- 30 Alexander Munteanu. Coresets and sketches for regression problems on data streams and distributed data. In *Machine Learning under Resource Constraints - Volume 1: Fundamentals*, pages 85–97. De Gruyter, 2022. doi:10.1515/9783110785944-003.
- 31 Alexander Munteanu, Simon Omlor, and Jeff M. Phillips. Hardness of high-dimensional linear classification. In *42nd International Symposium on Computational Geometry, (SoCG)*, pages 80:1–80:16, 2026. doi:10.4230/LIPIcs.SOCG.2026.80.
- 32 Alexander Munteanu, Matteo Russo, David Saulpic, and Chris Schwiegelshohn. Terminal dimension reduction for time series with applications. In *Forty-third International Conference on Machine Learning*, 2026. URL: <https://openreview.net/forum?id=gxbCZ1y8XS>.
- 33 Gilles Pisier. *The volume of convex bodies and Banach space geometry*. Cambridge Tracts in Mathematics. 94, 1999.
- 34 Ioannis Psarros and Dennis Rohde. Random projections for curves in high dimensions. In *39th International Symposium on Computational Geometry, (SoCG)*, pages 53:1–53:15, 2023. doi:10.4230/LIPIcs.SOCG.2023.53.
- 35 Ioannis Psarros and Dennis Rohde. Random projections for curves in high dimensions. *Discrete & Computational Geometry*, 74(2):374–398, 2025. doi:10.1007/S00454-024-00710-5.
- 36 Tamás Sarlós. Improved approximation algorithms for large matrices via random projections. In *47th Annual IEEE Symposium on Foundations of Computer Science, (FOCS)*, pages 143–152, 2006. doi:10.1109/FOCS.2006.37.
- 37 Christian Sohler and David P. Woodruff. Strong coresets for k-median and subspace approximation: Goodbye dimension. In *59th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 802–813, 2018. doi:10.1109/FOCS.2018.00081.
- 38 David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014. doi:10.1561/04000000060.

A Equivalence of traversal definitions

We prove that the usual definition of traversals as a sequence of vertex pairs is equivalent to the definition using unordered sets, that is used in our paper.

► **Lemma 6.** *Let $\sigma = (u_1, \dots, u_{m'})$, $\tau = (v_1, \dots, v_{m''})$ be polygonal curves. For every sequence of vertex pairs $T = ((u, v)_\ell)_{\ell \in [L]}$, where $u \in \sigma, v \in \tau$ and L is the length of the sequence, consider the corresponding set $S = \{(u, v)_\ell \mid \ell \in [L]\}$. Then T satisfies*

- A. T starts with (u_1, v_1) and ends with $(u_{m'}, v_{m''})$,
- B. $\forall i < m', j < m''$: the successor of (u_i, v_j) in T is exactly one element $(u_{i'}, v_{j'}) \in \{(u_{i+1}, v_j), (u_i, v_{j+1}), (u_{i+1}, v_{j+1})\}$,

if and only if S satisfies

1. $(u_1, v_1) \in S \wedge (u_{m'}, v_{m''}) \in S$,
2. $\forall i < m', j < m''$: $(u_i, v_j) \in S \Rightarrow (u_i, v_{j+1}) \in S \vee (u_{i+1}, v_j) \in S \vee (u_{i+1}, v_{j+1}) \in S$,
3. $(u_i, v_{j+1}) \in S \Rightarrow (u_{i+1}, v_j) \notin S$.

Proof. We split the proof into several items:

A. $\Leftrightarrow$ 1.

If T starts with (u_1, v_1) and ends with $(u_{m'}, v_{m''})$, then clearly $(u_1, v_1) \in S$ and $(u_{m'}, v_{m''}) \in S$. Reversely, if $(u_1, v_1) \in S$ and $(u_{m'}, v_{m''}) \in S$, we can put them to be the first, respectively the last element in T .

B. $\Rightarrow$ **2.**

Let $(u_i, v_j) \in S$, then it is also an element of T and has one element $(u_{i'}, v_{j'}) \in \{(u_{i+1}, v_j), (u_i, v_{j+1}), (u_{i+1}, v_{j+1})\}$ as its direct successor in T . Thus $(u_{i'}, v_{j'}) \in S$.

B. $\Rightarrow$ **3.**

Let $(u_i, v_{j+1}) \in S$. Then it is also in T and all its (not necessarily direct) successors $(u_{i'}, v_{j'})$ in T satisfy $i' \geq i$ and $j' \geq j + 1$. So (u_{i+1}, v_j) can never appear after (u_i, v_{j+1}) in T . Similarly, all predecessors $(u_{i'}, v_{j'})$ in T satisfy $i' \leq i$ and $j' \leq j + 1$. So (u_{i+1}, v_j) can never appear before (u_i, v_{j+1}) in T . It follows that $(u_{i+1}, v_j) \notin S$.

2. $\wedge$ **3.** $\Rightarrow$ **B.**

By Item 2 we have that at least one of the admissible successors of $(u_i, v_j) \in S$ (w.r.t. T) must be present in S . If it is exactly one, then we are done. Otherwise, by Item 3, we have that only one element $(u_{i'}, v_{j'}) \in \{(u_i, v_{j+1}), (u_{i+1}, v_j)\}$ can be in S in addition to $(u_{i+1}, v_{j+1}) \in S$. We can thus construct T such that $T = (\dots, (u_i, v_j), (u_{i'}, v_{j'}), (u_{i+1}, v_{j+1}), \dots)$, which obeys the rules of B.

This finishes our proof. ◀