

$(1 + \varepsilon)$ -ANN Data Structure for Curves via Subspaces of Bounded Doubling Dimension

Jacobus Conradi ✉ 

Department of Computer Science, University of Bonn, Germany

Anne Driemel ✉ 

Hausdorff Center for Mathematics, University of Bonn, Germany

Benedikt Kolbe ✉

Hausdorff Center for Mathematics, University of Bonn, Germany

Abstract

We consider the $(1 + \varepsilon)$ -Approximate Nearest Neighbour (ANN) Problem for polygonal curves in d -dimensional space consisting of at most k vertices under the Fréchet distance and ask to what extent known data structures for doubling spaces can be applied to this problem. Initially, this approach does not seem viable, since the doubling dimension of the target space is known to be unbounded — even for well-behaved polygonal curves of constant complexity in one dimension. In order to overcome this, we identify a subspace of curves which has bounded doubling dimension and small Gromov-Hausdorff distance to the target space. We then apply state-of-the-art techniques for doubling spaces and show how to obtain a data structure for the $(1 + \varepsilon)$ -ANN problem for any set of parametrized polygonal curves. The expected preprocessing time needed to construct the data-structure is $F(d, k, S, \varepsilon)n \log n$ and the space used is $F(d, k, S, \varepsilon)n$, with a query time of $F(d, k, S, \varepsilon) \log n + F(d, k, S, \varepsilon)^{-\log(\varepsilon)}$, where $F(d, k, S, \varepsilon) = O(2^{O(d)} k \Phi(S) \varepsilon^{-1})^k$ and $\Phi(S)$ denotes the spread of the set of vertices and edges of the curves in S . We extend these results to the realistic class of c -packed curves and show improved bounds for small values of c .

Keywords and phrases Fréchet Distance, Nearest Neighbour, Doubling Space

Digital Object Identifier 10.57717/cgt.v3i2.45

Funding *Jacobus Conradi*: Partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 313421352 (FOR 2535 Anticipating Human Behavior). Affiliated with Lamarr Institute for Machine Learning and Artificial Intelligence.

Anne Driemel: Affiliated with Lamarr Institute for Machine Learning and Artificial Intelligence.

1 Introduction

Given a set S of n points, the Nearest Neighbour Problem is the problem of finding the point in S that minimizes the distance to a given query point q . The Nearest Neighbour Problem is a fundamental problem whose variants have long been studied and applied in different areas, such as RNA sequencing [8], disease diagnosing [35], motion pattern detection [18], shape indexing [7] or handwritten digit recognition [27]. The problem has been studied as early as the 1960s [28], and classical results such as the one by Shamos via point location in a Voronoi diagram achieve a query time of $O(\log n)$ while using $O(n \log n)$ space in \mathbb{R}^2 [33], which was later improved upon by Kirkpatrick to only require linear space and preprocessing time with the logarithmic query time [25].

Various methods have been employed to design data structures for approximate solutions to the Nearest Neighbour Problem [6, 21, 26]. One complexity measure often used to generalize approaches for Euclidean spaces to more complicated metric spaces is the notion of doubling dimension [20, 24, 36]. The doubling dimension is the smallest number d such that any metric ball inside the metric space can be covered by 2^d many balls of half the radius. It is a well known fact that in Euclidean space the doubling dimension roughly



corresponds to the dimension, that is, the doubling dimension of \mathbb{R}^d is $\Theta(d)$. The Approximate Nearest Neighbour (ANN) Problem in spaces with low doubling dimension has been studied extensively [24] and results are known that roughly match the bounds known for \mathbb{R}^d [22].

The metric space we are interested in is the space of polygonal curves in \mathbb{R}^d under the Fréchet distance. Polygonal curves naturally arise from any sort of motion tracking, such as GPS data or motion capture data, and are therefore of much interest. The Fréchet distance is a natural distance measure on parametrized curves [4] that — unlike the Hausdorff distance — takes into account the parametrization of the curves and has received considerable attention [2, 5, 11, 12, 13, 14]. The metric space of curves under the Fréchet distance has been shown to have unbounded doubling dimension [15], which suggests that data structures designed for doubling spaces would perform poorly. Our primary motivation in this paper is to find a workaround to this problem, and to leverage the rich background of ANN results for doubling spaces to the Fréchet distance after all. We observe that the family showing unbounded doubling dimension has either arbitrarily long edges or arbitrarily short edges. Our approach circumvents these two scenarios and gives a small doubling dimension for all other cases by choosing an appropriate subspace.

1.1 Basic definitions

The metric space we study in this paper is the set of all polygonal curves with the Fréchet distance as its metric.

► **Definition 1** (polygonal curve). *An **edge** in \mathbb{R}^d is the continuous map obtained from the linear interpolation of two points a and b in \mathbb{R}^d . We may write \overline{ab} to denote this unique edge. A **polygonal curve** $T : [0, 1] \rightarrow \mathbb{R}^d$ of complexity n is defined by an ordered set of n points in \mathbb{R}^d and is the result of $(n - 2)$ concatenations of the $(n - 1)$ edges in \mathbb{R}^d defined by any two consecutive points. We call the underlying points of a polygonal curve of complexity n its **vertices**. For $0 \leq s \leq t \leq 1$ we denote the subcurve of T from $T(s)$ to $T(t)$ by $T[s, t]$.*

► **Definition 2** (Fréchet distance). *Given two curves X and Y in \mathbb{R}^d , their Fréchet distance is defined as*

$$d_{\mathcal{F}}(X, Y) = \inf_{f, g: [0, 1] \rightarrow [0, 1]} \max_{t \in [0, 1]} \|X(f(t)) - Y(g(t))\|$$

where f and g are non-decreasing and surjective.

In our definition of the Fréchet distance given above, we follow Alt and Godau [3]. The Fréchet distance is generally not a metric, but rather a pseudo-metric, as there are curves $X \neq Y$, such that $d_{\mathcal{F}}(X, Y) = 0$. However, this is easily remedied by considering the quotient space induced by the equivalence relation $X \sim Y \iff d_{\mathcal{F}}(X, Y) = 0$. If we assume that no vertex of an input curve lies inside the convex hull of its two neighbors along the curve, then the ordered sequence of vertices of two equivalent curves is the same. Therefore, we can eliminate all duplicates in near-linear time using lexicographical sorting. This results in a set of curves with pairwise non-zero distance. In the following, we simply assume that the Fréchet distance between any two curves in the input is non-zero.

► **Definition 3**. *Denote by $(\mathbb{X}^{d, k}, d_{\mathcal{F}})$ the metric space of polygonal curves in \mathbb{R}^d with complexity k under the continuous Fréchet distance. We further write $(\mathbb{X}_{\Lambda}^{d, k}, d_{\mathcal{F}})$ for the subspace of polygonal curves in $(\mathbb{X}^{d, k}, d_{\mathcal{F}})$ where the length of each edge is bounded by Λ .*

A well-studied variant of the continuous Fréchet distance is the discrete Fréchet distance where only the vertices—and not their connecting edges—are considered in the computation.

Similar to the continuous Fréchet distance, the discrete Fréchet distance is a pseudo-metric and can be regarded as a metric in a corresponding quotient space.

► **Definition 4** (discrete Fréchet distance). *Given two polygonal curves X and Y in \mathbb{R}^d defined by vertices x_1, \dots, x_n and y_1, \dots, y_m , their discrete Fréchet distance is defined as*

$$d_{d\mathcal{F}}(X, Y) = \inf_{f, g} \max_{t \in [0, 1]} \|x_{f(t)} - y_{g(t)}\|,$$

where $f : [0, 1] \rightarrow \{1, \dots, n\}$ and $g : [0, 1] \rightarrow \{1, \dots, m\}$ are non-decreasing and surjective functions.

The following well-known fact relates the discrete Fréchet distance to the continuous Fréchet distance.

► **Observation 5.** *Let P, Q be polygonal curves in \mathbb{R}^d . Then $d_{\mathcal{F}}(P, Q) \leq d_{d\mathcal{F}}(P, Q)$.*

Problem definition ((1 + ε)-ANN) Let $(\mathcal{M}, d_{\mathcal{M}})$ be a metric space. Let $P \subset \mathcal{M}$ be a set of points in \mathcal{M} and a parameter $\varepsilon > 0$ be given. For a given point $q \in \mathcal{M}$, the (1 + ε)-Approximate Nearest Neighbour Problem ((1 + ε)-ANN) is to find a point $\hat{x} \in P$ whose distance to q approximates the distance to the nearest neighbour in P . Specifically, $\hat{x} \in P$ is a valid solution iff for all $x \in P$ it holds that

$$d_{\mathcal{M}}(q, \hat{x}) \leq (1 + \varepsilon) d_{\mathcal{M}}(q, x)$$

1.2 State of the art

Most of the work on nearest-neighbour data structures for curves has focused on the discrete Fréchet distance. One such recent result is due to Filtser, Filtser and Katz [17] who presented a data structure for the (1 + ε, r)-Approximate Near Neighbour Problem of size $n \cdot O(1/\varepsilon)^{md}$, with a query time of $O(md)$, where n is the number of input curves, and m is the complexity of the input-curves. When the complexity k of the query curve is small compared to the complexity m of the input curves, the space can be improved to $n \cdot O(1/\varepsilon)^{kd}$ with query time $O(kd \log(nkd/\varepsilon))$. In the (1 + ε, r)-Approximate Near Neighbour Problem the goal is to construct a data structure on a set of input-curves which for a given query-curve outputs any of the input-curves that is at distance at most (1 + ε)r to the query curve, if there is an input-curve with distance at most r to the query-curve.

Results for the (1 + ε, r)-Approximate Near Neighbour problem readily extend to the (1 + ε)-Approximate Nearest Neighbour Problem by a result of Har-Peled, Indyk and Motwani [21]. This reduction incurs merely an additional logarithmic factor $O(\log^2(n))$ in the size, and $O(\log n)$ in the query time.

In contrast to the multitude of approaches to the discrete Fréchet distance for arbitrary dimension, results concerning the continuous Fréchet distance appear harder to come by. Consider the naive approach of approximating the continuous Fréchet distance via the discrete Fréchet distance. For this, let a set of n curves in $\mathbb{X}_{\Lambda}^{d, k}$ as well as the approximation parameter ε be given. For the discrete Fréchet distance to approximate the continuous Fréchet distance up to an additive term of $r\varepsilon$ for some $r > 0$, we require successive vertices to lie at most $\Theta(r\varepsilon)$ far apart. Thus we subdivide every edge into edges of length at most $r\varepsilon$, resulting in a set of n curves each of complexity $O(k\Lambda/(r\varepsilon))$. Building the (1 + ε, r)-ANN data structure from [17] results in a space requirement of $n \cdot O(1/\varepsilon)^{O(dk\Lambda/(r\varepsilon))}$. As the radii used in the reduction can be as small as r^*/n , where r^* is roughly the distance of some two input curves,

this extends to a data structure of size $O(n \log^2(n)) \cdot O(1/\varepsilon)^{O(ndk\Lambda/(r^*\varepsilon))}$ for the $(1+\varepsilon)$ -ANN problem.

Bringmann et al. [10] showed that there is a $(1+\varepsilon, r)$ -Approximate Near Neighbour data structure for the continuous Fréchet distance in one dimension, which uses $n \cdot O(\frac{m}{k\varepsilon})^k$ space, needs $O(nm) \cdot O(\frac{m}{k\varepsilon})^k$ expected preprocessing time and achieves a query time of $O(k \cdot 2^k)$. They also show tightness of their data structure bounds in several scenarios. More precisely, they show conditional lower bounds based on the Orthogonal Vectors Hypothesis that give reason to believe that one cannot achieve preprocessing time $\text{poly}(n)$ and query time $O(n^{1-\varepsilon'})$ at the same time, when k is $1 \ll k \ll \log n$ and $m > k \cdot n^{c/k}$, even if $d = 1$. Their arguments also apply to the c -ANN problem under the continuous Fréchet distance for any $c \in [1, 2)$.

In two dimensions, Afshani and Driemel presented a data structure based on semi-algebraic range searching that solves the (exact) Near Neighbour Problem under the Fréchet distance [1]. The data structure needs $O(n(\log \log n)^{O(m^2)})$ space and has query time $O(\sqrt{n} \log^{O(m^2)} n)$.

In higher dimensions, Mirzanezhad recently presented a data structure result for the $(1+\varepsilon, r)$ -Approximate Near Neighbour Problem under the continuous Fréchet distance, using space in $n \cdot O((\max(1, D)\sqrt{d}/\varepsilon^2)^{kd})$ and query time in $O(kd)$, where D denotes the diameter of the underlying vertex set of the input curves [30]. However, as presented this data structure works only if $\varepsilon < r$ (refer to Theorem 8 in the arXiv version [29]). The data structure covers the entirety of the input-curves with a grid of edge-length roughly εr , to precompute an answer for every sequence of k gridpoints. A query-curve is then snapped to the closest grid points and the precomputed answer is given as the output. For smaller values of r one would have to scale the input, increasing D accordingly. As a result, combining this data structure with the standard reduction by Har-Peled, Indyk and Motwani [21] does not lead to an efficient data structure for the ANN-problem.

Independent to our work, Cheng and Huang very recently presented a $(1+\varepsilon)$ -ANN data structure for polygonal curves in arbitrary dimension under the continuous Fréchet distance. The data structure uses space in $\tilde{O}\left(k(mnd^d/\varepsilon^d)^{O(k+1/\varepsilon^2)}\right)$ and achieves query time in $\tilde{O}\left(k(mn)^{0.5+\varepsilon/\varepsilon^d} + k(d/\varepsilon)^{O(dk)}\right)$. Here, $\tilde{O}(\cdot)$ hides polylogarithmic factors with exponents in $O(d)$.

1.3 Our contribution

In this paper, we provide a $(1+\varepsilon)$ -ANN data structure for a set of curves in arbitrary dimensions under the continuous Fréchet distance. One of the main ingredients is the construction of a suitable subspace of the space of all polygonal curves. This subspace has small distance to the original space as metric spaces and additionally, unlike the space of all polygonal curves, bounded doubling dimension. Our approach can thus be seen as a particular instance of the more general idea of using results on the ANN problem for spaces of bounded doubling dimension in the context of spaces that are close to a space of bounded doubling dimension, even if they do not themselves have this property. This approach was inspired by the work of Sheehy and Sheth on the bottleneck distance for persistence diagrams [34].

Throughout this paper, for simplicity of the exposition, we will not distinguish between small or large k , that is, we assume $k = m$. A key result is the following (Section 5).

► **Theorem 6.** *Given a set S of n polygonal curves in $\mathbb{X}_\Lambda^{d,k}$ and parameters $0 < \varepsilon < 1$ and $\varepsilon' > 0$, one can construct a data structure that for given $q \in \mathbb{X}^{d,k}$ outputs an element $s^* \in S$ such that for all $s \in S$ it holds that $d_{\mathcal{F}}(s^*, q) \leq (1+\varepsilon)d_{\mathcal{F}}(s, q) + \varepsilon'$. The query time is*

$O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k \log n + O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^{-k \log(\varepsilon)}$, the expected preprocessing time is $O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k n \log n$ and the space used is $O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k n$.

To turn this into a data structure with a purely multiplicative error we choose ε' as the smallest Fréchet distance of any two distinct input curves over ε . Thus, the running time and space complexity of our data structure polynomially depends on a numerical value, which we call the bundledness of the set of input curves.

► **Definition 7 (bundledness).** Given a set of curves $S \in \mathbb{X}^{d,k}$, the bundledness $\mathcal{G}(S)$ of S is defined as

$$\mathcal{G}(S) = \frac{\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s')}{\max_{e \in E(S)} \|e\|},$$

where $E(S)$ denotes the set of edges of curves in S and $\|e\|$ denotes the length of the edge e .

The bundledness is reminiscent of the global stretch — a measure of complexity on geometric graphs — which was introduced by Erickson [16] and further analyzed by Bose et al. [9]. The bundledness is closely related to the spread of the set of vertices and edges of the input curves, in that it is lower bounded by the reciprocal of the spread (see Lemma 39).

► **Definition 8 (spread).** For a point set P in some metric space $(\mathcal{M}, d_{\mathcal{M}})$, we define the spread $\Phi(P)$ as the ratio between the maximal and minimal pairwise distance of points in P . Similarly, define the spread $\Phi(S)$ of a collection of sets as the ratio between the maximal and minimal non-zero pairwise distances of sets in S , where the distance between two sets $A, B \subset \mathcal{M}$ is defined as $d_{\mathcal{M}}(A, B) = \min_{a \in A} \min_{b \in B} d_{\mathcal{M}}(a, b)$.

With these definitions, our principal results can be summarized as follows (Section 5).

► **Theorem 9.** Given a set S of n polygonal curves in $\mathbb{X}^{d,k}$ and $0 < \varepsilon \leq 1$ one can construct a data structure answering $(1 + \varepsilon)$ -approximate nearest neighbour queries. The query time is $F(d, k, S, \varepsilon) \log n + F(d, k, S, \varepsilon)^{-\log(\varepsilon)}$, the expected preprocessing time is $F(d, k, S, \varepsilon)n \log n$ and the space used is $F(d, k, S, \varepsilon)n$, where $F(d, k, S, \varepsilon) = O(2^{O(d)}k(1 + \mathcal{G}(S)^{-1}\varepsilon^{-1}))^k$.

Replacing the bundledness with the more pessimistic spread of the vertices and edges of the input curves we get the following result.

► **Corollary 10.** Given a set S of n polygonal curves in $\mathbb{X}^{d,k}$ and $0 < \varepsilon \leq 1$ one can construct a data structure answering $(1 + \varepsilon)$ -approximate nearest neighbour queries. The query time is $F(d, k, S, \varepsilon) \log n + F(d, k, S, \varepsilon)^{-\log(\varepsilon)}$, the expected preprocessing time is $F(d, k, S, \varepsilon)n \log n$ and the space used is $F(d, k, S, \varepsilon)n$, where $F(d, k, S, \varepsilon) = O(2^{O(d)}k\Phi(S)\varepsilon^{-1})^k$, where $\Phi(S)$ denotes the spread of the set of vertices and edges of the curves in S .

In the special case that all curves in S are c -packed for some constant $c > 0$, that is the length of the intersection of the curve with any ball is at most c times the balls radius, the parameter $F(d, k, S, \varepsilon)$ is instead in $O(2^{O(d)}(1 + \mathcal{G}(S)^{-1}\varepsilon^{-1}))^k$, or, more pessimistically, in $O(2^{O(d)}\Phi(S)\varepsilon^{-1})^k$.

In Section 6 we also discuss sufficient conditions for our approach to generalize to other settings: when a space of unbounded doubling dimension is approximated by a well-suited family of subspaces of bounded doubling dimension.

1.4 Technical overview

In Section 2, we introduce a subspace \mathbb{X}^* of $(\mathbb{X}^{d,k}, d_{\mathcal{F}})$ and show how the subspace relates to $(\mathbb{X}^{d,k}, d_{\mathcal{F}})$. A significant part of our work is concerned with analyzing properties of this subspace \mathbb{X}^* . In Section 3, we present the analysis of the upper bound on the doubling constant and dimension of \mathbb{X}^* , which constitutes our main technical result, Theorem 21. Roughly, we show that for a curve Q to have a small Fréchet distance ($\leq \Delta$) to another curve C there is only a small subset of \mathbb{R}^d in which any vertex of Q can lie in. In fact, it turns out that this subset can be covered by $\approx 2^d k\mu$ (\mathbb{R}^d -)balls of radius $\Delta/2$, depending on the doubling dimension $\approx 2^d$ of \mathbb{R}^d . As Q consists of k vertices, Q is then described by one of at most $(2^d k\mu)^k$ sequences of (\mathbb{R}^d -)balls, implying that the doubling dimension is at most $\log((2^d k\mu)^k) \approx k(d + \log(k\mu))$. We furthermore extend this analysis to the special case that the curves are c -packed, resulting in an improvement in the upper bound. In Section 4, we extend the lower bound construction from [15] to argue that our bounds on the doubling dimension of \mathbb{X}^* are almost tight. In Section 5, we then use the bound on the doubling dimension of \mathbb{X}^* to construct an $(1 + \varepsilon)$ -ANN data structure based on the work of Har-Peled and Mendel [22]. In Section 6, we then generalize these results to present a general framework, which specifies how one can extend $(1 + \varepsilon)$ -ANN data structures for spaces of bounded doubling dimension to spaces of unbounded doubling dimension.

2 Curve spaces

► **Definition 11** (doubling constant and dimension). *Let $(\mathcal{M}, d_{\mathcal{M}})$ be a metric space. Define the r -ball in \mathcal{M} by $B_r^{\mathcal{M}}(x) = \{y \in \mathcal{M} \mid d_{\mathcal{M}}(x, y) \leq r\}$, as the closed ball of radius $r > 0$ centered at x . The **doubling constant** of $(\mathcal{M}, d_{\mathcal{M}})$ is defined as the minimal number ν , such that for any $x \in \mathcal{M}$ and any $r > 0$ the ball $B_r^{\mathcal{M}}(x)$ of radius r is contained in the union of at most ν balls with radius $r/2$. The **doubling dimension** of $(\mathcal{M}, d_{\mathcal{M}})$ is defined as $\log_2(\nu)$.*

We may omit the metric space in the notation of a ball whenever the metric space is clear, writing $B_r(x)$ instead of $B_r^{\mathcal{M}}(x)$ for $x \in \mathcal{M}$.

It turns out that for $k \geq 3$ the doubling dimension of $(\mathbb{X}^{d,k}, d_{\mathcal{F}})$ is unbounded [15]. A straightforward modification to their construction yields the following theorem.

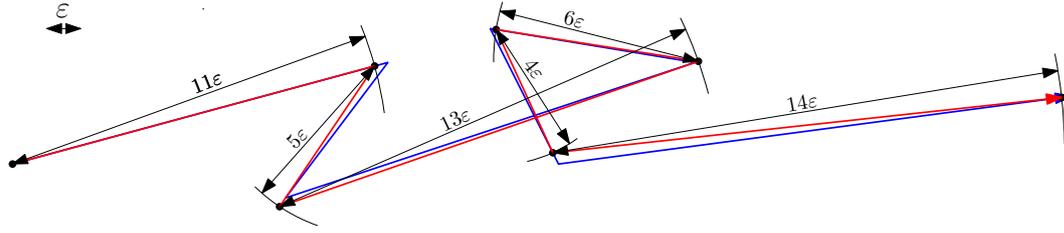
► **Theorem 12.** *The doubling constant of $(\mathbb{X}_{\Lambda}^{d,k}, d_{\mathcal{F}})$ is unbounded for any $k \geq 3$ and $\Lambda > 0$.*

Theorem 12 motivates the search for a subspace $(\mathbb{X}^*, d_{\mathcal{F}})$ of $(\mathbb{X}_{\Lambda}^{d,k}, d_{\mathcal{F}})$ of bounded doubling dimension. To answer $(1 + \varepsilon)$ -ANN queries in $\mathbb{X}_{\Lambda}^{d,k}$, the ansatz is to map the input curves to $(\mathbb{X}^*, d_{\mathcal{F}})$ and answer queries in this subspace. With this approach, the quality guarantee of the $(1 + \varepsilon)$ -ANN queries decreases by an additional additive factor, roughly depending on the distortion of the map between $(\mathbb{X}_{\Lambda}^{d,k}, d_{\mathcal{F}})$ and $(\mathbb{X}^*, d_{\mathcal{F}})$.

► **Definition 13** ((μ, ε) -curves). *For any $\varepsilon > 0$ and $\mu \in \mathbb{N}$ define the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$ as the subspace of $(\mathbb{X}^{d,k}, d_{\mathcal{F}})$ induced by the set of polygonal curves in $\mathbb{X}^{d,k}$ whose edge lengths are all exact multiples of ε . We further require the edge lengths to be bounded by $\mu\varepsilon$. The space of (μ, ε) -curves in $\mathbb{X}^{d,k}$ is a natural subspace of $(\mathbb{X}_{\mu\varepsilon}^{d,k}, d_{\mathcal{F}})$.*

We may abuse notation slightly, and not specify the ambient space $\mathbb{X}^{d,k}$ of the space of (μ, ε) -curves, if the ambient space is clear. We may likewise write \mathcal{M} when talking about the metric space $(\mathcal{M}, d_{\mathcal{M}})$, if the metric is clear.

► **Lemma 14.** *Let $P \in \mathbb{X}_{\Lambda}^{d,k}$ be a polygonal curve and $\varepsilon > 0$. We can construct a $(\lceil \Lambda/\varepsilon \rceil + 1, \varepsilon)$ -curve P' in $\mathbb{X}^{d,k}$ such that $d_{\mathcal{F}}(P, P') \leq \varepsilon/2$ in $O(k \log(\Lambda/\varepsilon))$ time.*



■ **Figure 1** Example of a curve $P \in \mathbb{X}_{\Lambda}^{d,k}$ in blue, and an ε -curve close to P resulting from Lemma 14 in red.

Proof. Refer to Figure 1. Let p_1, \dots, p_k be the vertices of P and $P' = \emptyset$. We begin by adding $p'_1 = p_1$ to P' . Now assume p'_{i-1} is the last vertex of P' . Compute the value μ_i , such that the magnitude $|\mu_i \varepsilon - \|p_i - p'_{i-1}\||$ is minimal. Then add $p'_i = p'_{i-1} + (\mu_i \varepsilon) \frac{p_i - p'_{i-1}}{\|p_i - p'_{i-1}\|}$ to P' . Note that by construction $\|p_i - p'_i\| \leq \varepsilon/2$. Hence $d_{\mathcal{F}}(P, P') \leq \varepsilon/2$. The length of the edges of P' are bounded by $(\lceil \Lambda/\varepsilon \rceil + 1)\varepsilon$. Indeed, $\|p_i - p'_i\| \leq \varepsilon/2$ and $\|p_i - p_{i-1}\| \leq \Lambda$ imply that $\|p'_i - p'_{i-1}\| \leq \Lambda + \varepsilon$. For the running time, for every $i \leq n$ the value μ_i such that the magnitude $|\mu_i \varepsilon - \|p_i - p'_{i-1}\||$ is minimal can be identified in $O(\log(\Lambda/\varepsilon))$ time. This is done by first identifying the smallest power of 2 larger than $\|p_i - p'_{i-1}\|/\varepsilon$ and then binary searching over the integer multiples of ε up to this power of 2. As we do this once for every edge, the claimed running time follows. ◀

The additive error incurred by mapping $\mathbb{X}_{\Lambda}^{d,k}$ to the space of (M, ε) -curves in $\mathbb{X}^{d,k}$ for $M \geq \lceil \Lambda/\varepsilon \rceil + 1$ depends only on the distortion of the map. The Gromov-Hausdorff distance is a related measure of the distance between metric spaces. For $\mu \rightarrow \infty$ the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$ will have (depending on ε) small Gromov-Hausdorff distance to the ambient space $\mathbb{X}^{d,k}$ of curves. We map any element from $\mathbb{X}^{d,k}$ into this subspace by rounding the length of every edge to a multiple of ε , thereby moving the vertices by a small amount each time (refer to Figure 1).

► **Definition 15** (Gromov-Hausdorff distance). *The Gromov-Hausdorff distance is a distance measure on metric spaces. Let \mathcal{M} and \mathcal{N} be two metric spaces. Then the Gromov-Hausdorff distance is defined as*

$$d_{GH}(\mathcal{M}, \mathcal{N}) = \inf_{\mathcal{Z}} \{d_H^{\mathcal{Z}}(f(\mathcal{M}), g(\mathcal{N})) \mid f : \mathcal{M} \rightarrow \mathcal{Z}, g : \mathcal{N} \rightarrow \mathcal{Z} \text{ isometric embeddings}\},$$

where \mathcal{Z} ranges over metric spaces and $d_H^{\mathcal{Z}}(X, Y)$ denotes the Hausdorff distance of two sets in the metric space \mathcal{Z} which is defined as $\max\{\sup_{x \in X} \inf_{y \in Y} d_{\mathcal{Z}}(x, y), \sup_{y \in Y} \inf_{x \in X} d_{\mathcal{Z}}(x, y)\}$.

► **Corollary 16.** *Let d, k and $\varepsilon > 0$ be given, $\mu \rightarrow \infty$, and denote by C_{μ} the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$. Then $d_{GH}(\mathbb{X}^{d,k}, C_{\mu}) \rightarrow \varepsilon/2$.*

Proof. For every $P \in \mathbb{X}^{d,k}$ and $\varepsilon > 0$ there is a finite integer $M < \infty$, such that $P \in \mathbb{X}_{M\varepsilon}^{d,k}$. Then by Lemma 14 there is a $P' \in C_{M+1}$ with $d_{\mathcal{F}}(P, P') \leq \varepsilon/2$. As C_{μ} is a natural subspace of $\mathbb{X}^{d,k}$ for all $\mu \in \mathbb{N}$, the claim holds. ◀

3 Bounding the doubling dimension of the space of (μ, ε) -curves

In this section, we study the doubling dimension of the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$. Unfortunately, our bound is non-constructive. As such, it does not provide a doubling oracle that, for a given ball of radius r in the metric space, outputs a set of balls of radius $r/2$ which cover the ball of radius r .

3.1 Properties of the Euclidean space

Before diving into the analysis of the doubling dimension of the space of (μ, ε) -curves, we begin by analyzing properties of the ambient space \mathbb{R}^d . For this we often inspect so called Δ -neighbourhoods of subsets of \mathbb{R}^d . For any subset A the Δ -neighbourhood of A is defined by $N_\Delta(A) = \{x \in \mathbb{R}^d \mid \exists a \in A : d(x, a) \leq \Delta\}$. Note that the Δ -neighbourhood of a single point x coincides with a ball of radius Δ centered at x .

It is well-known that the doubling dimension of \mathbb{R}^d under the Euclidean norm is in $\Theta(d)$. We prove a more general statement which relates the volume of some arbitrary set to the amount of balls needed to cover it. For this, we denote by λ^d the Lebesgue measure in \mathbb{R}^d and by $\lambda^d(A)$ the **volume** of a (measurable) set $A \subset \mathbb{R}^d$.

► **Lemma 17.** *Let $A \subset \mathbb{R}^d$ be a bounded set, and let $r > 0$ be fixed. Then there is a set of points $C \subset A$ of cardinality $\lceil \lambda^d(N_{r/2}(A)) / V_{r/2}^d \rceil$ such that $A \subset \bigcup_{c \in C} B_r(c)$, where $V_{r/2}^d$ denotes the volume of a d -dimensional ball of radius $r/2$.*

Proof. We construct the set C greedily. For this we start with $C = \emptyset$, and then iteratively add any point from $A \setminus \bigcup_{c \in C} B_r(c)$, until $A \subset \bigcup_{c \in C} B_r(c)$. As any two points in C have a distance of at least r , balls centered at points of C with radius $r/2$ are disjoint and clearly contained in $N_{r/2}(A)$. Thus the number of points in C is bound by $\lceil \lambda^d(N_{r/2}(A)) / V_{r/2}^d \rceil$. ◀

► **Lemma 18.** *For any $r > 0$ and $c > 1$, any ball $B_r(p) \subset \mathbb{R}^d$ can be covered by $O((2c+1)^d)$ balls of radius r/c .*

Proof. This follows from the classical result that $V_r^d = \frac{\pi^{d/2}}{\Gamma(d/2+1)} r^d$, where Γ denotes the Gamma-function [31]. As $N_{r/2c}(B_r(p)) = B_{r(2c+1)/2c}(p)$, Lemma 17 implies that any $B_r(p)$ can be covered with $\lceil (r(2c+1)/2c)^d / (r/2c)^d \rceil = O((2c+1)^d)$ balls of radius r/c . ◀

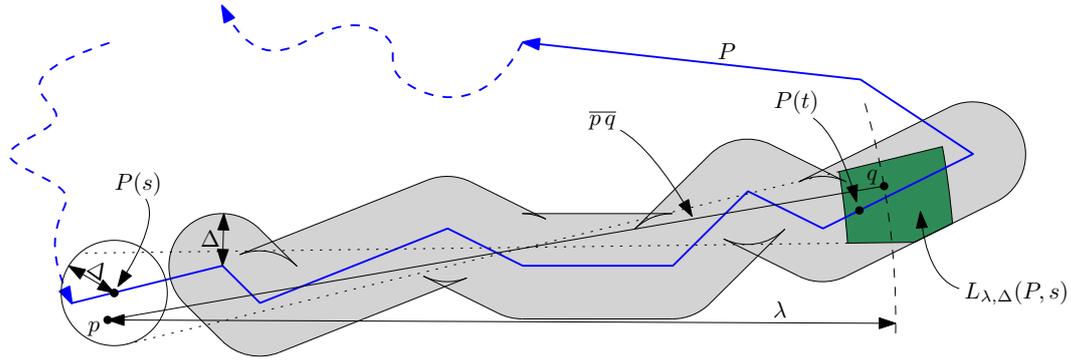
Later on, we will use this lemma to retrieve a set of points, namely the centers of the balls used in such a covering, as candidates for vertices of curves that will be centers of balls in $\mathbb{X}^{d,k}$.

When analyzing the doubling dimension, we often consider what an edge might look like that has Fréchet distance at most $\Delta > 0$ to a subcurve of an input curve. The basic tool we use for this analysis is the observation that the edge under consideration then is a Δ -stabber of the vertices (and indeed any ordered set of points along the subcurve) of the subcurve.

► **Definition 19 (Δ -stabber).** *Let an ordered set of points (p_1, \dots, p_n) in \mathbb{R}^d be given. A polygonal curve $l : [0, 1] \rightarrow \mathbb{R}^d$ between two points $a, b \in \mathbb{R}^d$ is called a Δ -stabber of (p_1, \dots, p_n) if there are values $0 \leq t_1 \leq \dots \leq t_n \leq 1$ such that $\|l(t_i) - p_i\| \leq \Delta$ for all $1 \leq i \leq n$.*

The notion of Δ -stabbers has been introduced by Guibas et al. [19], and is closely related to the Fréchet distance. Any edge that has Fréchet distance at most Δ to a polygonal curve defined by vertices p_1, \dots, p_n is a Δ -stabber of the ordered point set (p_1, \dots, p_n) . Similarly, a Δ -stabber of the ordered point set (p_1, \dots, p_n) contains an edge that has Fréchet distance at most Δ to the polygonal curve defined by the vertices p_1, \dots, p_n .

► **Observation 20.** *Let a polygonal curve P and $\Delta > 0$ be given. Let $e = \overline{pq}$ be an edge such that for given $0 \leq s \leq t \leq 1$ the Fréchet distance $d_{\mathcal{F}}(P[s, t], e)$ is at most Δ . Then for any $s \leq m \leq t$ the edge e is a Δ -stabber of $(P(s), P(m), P(t))$.*



■ **Figure 2** Illustration of the set $L_{\lambda, \Delta}(P, s)$ in dark green, together with the points p and $P(t)$ realizing a point q in $L_{\lambda, \Delta}(P, s)$, that is $\|p - q\| = \lambda$ and $d_{\mathcal{F}}(P[s, t], \overline{pq}) \leq \Delta$.

3.2 Packing the metric ball

The result we now want to prove is the following theorem. Proving it constitutes the main bulk of work in this section.

► **Theorem 21.** *Let $k, \mu, d \in \mathbb{N}$ and $\varepsilon > 0$. The doubling constant of the space of (μ, ε) -curves in $\mathbb{X}^{d, k}$ is bounded by $O(43^d k \mu)^k$ and thus the doubling dimension of the space of (μ, ε) -curves is bounded by $O(k(d + \log(k\mu)))$.*

Let P be a (μ, ε) -curve in $\mathbb{X}^{d, k}$. Our objective is to cover the Δ -neighbourhood of P with respect to $d_{\mathcal{F}}$ with balls of radius $\Delta/2$. We encounter the question of where, given $p \in \mathbb{R}^d$, we may place a second point q such that there is a subcurve of P that is close to \overline{pq} . Indeed, for any curve Q with $d_{\mathcal{F}}(P, Q) \leq \Delta$, the endpoint q of any edge \overline{pq} of Q is a potential such point for the start point p . Thus if we answer the above stated question, we can iteratively build up any curve Q with $d_{\mathcal{F}}(P, Q) \leq \Delta$.

► **Definition 22.** *Let P be a polygonal curve and $\lambda \geq 0$ and $\Delta \geq 0$. For $s \in [0, 1]$ define the locus of edge endpoints of edges close to subcurves of P starting at the parameter s as the set*

$$L_{\lambda, \Delta}(P, s) = \{ q \in \mathbb{R}^d \mid \exists p \in \mathbb{R}^d \text{ and } \exists t \in [s, 1] \text{ with } \|p - q\| = \lambda, d_{\mathcal{F}}(P[s, t], \overline{pq}) \leq \Delta \}.$$

The following lemma motivates discretising the lengths of edges, as for a fixed length λ the set $L_{\lambda, \Delta}(P, s)$ is contained in a single ball of constant size.

► **Lemma 23.** *Let P be a polygonal curve. Let $\lambda \geq 0$ and $\Delta \geq 0$ be given. Then for every $s \in [0, 1]$ there is a point $p^* \in \mathbb{R}^d$, such that*

$$L_{\lambda, \Delta}(P, s) \subset B_{5\Delta}(p^*).$$

Proof. Assume $L_{\lambda, \Delta}(P, s)$ is nonempty, as otherwise we are done. Similarly assume $\lambda \geq 4\Delta$ as otherwise $L_{\lambda, \Delta}(P, s)$ is trivially contained in $B_{\lambda+\Delta}(P(s)) \subset B_{5\Delta}(P(s))$. Now let $t^* \geq s$ be the smallest value such that $\|P(s) - P(t^*)\| \geq \lambda - 2\Delta$. Then we claim that the sought-after point is $p^* = P(t^*)$.

For the remainder of this proof refer to Figure 3. Let $q \in L_{\lambda, \Delta}(P, s)$ be given. Then by definition there is a point p and a value t , such that \overline{pq} has length λ and $d_{\mathcal{F}}(P[s, t], \overline{pq}) \leq \Delta$. Thus $\|P(s) - p\| \leq \Delta$ and similarly $\|P(t) - q\| \leq \Delta$. But then by the triangle inequality $\|P(t) - P(s)\| \geq \lambda - 2\Delta$. Thus $t \geq t^*$. Hence, \overline{pq} is a Δ -stabber of $(P(s), P(t^*), P(t))$. This implies that there is a point m along \overline{pq} with $\|m - P(t^*)\| \leq \Delta$. As m lies on \overline{pq} , we have

Proof. Let $q \in \mathcal{L}_{\lambda, \Delta}(P, p)$. This implies that there are values $0 \leq s \leq t \leq 1$, such that $d_{\mathcal{F}}(P[s, t], \overline{pq}) \leq \Delta$ and $\|p - q\| = \lambda$. Let $q' = q + (p' - p)$. Then, as $\|p - p'\| \leq \Delta/c$ and thus $\|q - q'\| \leq \Delta/c$, we get that $d_{\mathcal{F}}(\overline{pq}, \overline{p'q'}) \leq \Delta/c$ by Observation 20, and thus $d_{\mathcal{F}}(P[s, t], \overline{p'q'}) \leq (1 + c^{-1})\Delta$. Finally, $\|p' - q'\| = \|p - q\| = \lambda$ and thus the point q' lies in $\mathcal{L}_{\lambda, (1+c^{-1})\Delta}(P, p')$, implying the claim. \blacktriangleleft

We now prove a stronger version of Theorem 21, which allows us to analyze the doubling constant of (μ, ε) -curves in $\mathbb{X}^{d, k}$.

► **Lemma 28.** *Let $k, \mu, d \in \mathbb{N}$ and $\varepsilon > 0$. Let a (μ, ε) -curve P in $\mathbb{X}^{d, k}$ be given, as well as $\Delta > 0$ and $c \geq 1$. There is a family of curves $\mathcal{C}_P \subset \mathbb{X}_{\mu\varepsilon + \Delta/c}^{d, k}$ of size $O(k\mu(10c + 3)^d)^k$, such that for any (μ, ε) -curve Q with $d_{\mathcal{F}}(P, Q) \leq \Delta$ there is a $Q^* \in \mathcal{C}_P$ with $d_{\mathcal{F}}(Q, Q^*) \leq \Delta/c$.*

Proof. We construct the set \mathcal{C}_P as follows. First, choose an element $(m_1, \dots, m_{k-1}) \in \{1, \dots, \mu\}^{k-1}$. Next, choose one circle center of a cover of $B_{\Delta}(P(0))$ consisting of $O((2c+1)^d)$ many balls of radius r/c , which exists by Lemma 18. Iteratively choose one point among the circle centers of a cover of $N_{\Delta/c}(\mathcal{L}_{m_{i-1}\varepsilon, (1+c^{-1})\Delta}(P, q_{i-1}^*))$ of Corollary 26, consisting of $O(k(10c+3)^d)$ many balls of radius r/c as the vertex q_i^* of Q^* for $i \leq k$. Then $Q^* \in \mathbb{X}_{\mu\varepsilon + \Delta/c}^{d, k}$, as for any i the fact that q_i^* lies in $N_{\Delta/c}(\mathcal{L}_{m_{i-1}\varepsilon, (1+c^{-1})\Delta}(P, q_{i-1}^*))$ implies that there is a point $q \in \mathcal{L}_{m_{i-1}\varepsilon, (1+c^{-1})\Delta}(P, q_{i-1}^*)$, with $\|q_{i-1}^* - q\| = m_{i-1}\varepsilon$ and $\|q - q_i^*\| \leq \Delta/c$. Hence, $\|q_i^* - q_{i-1}^*\| \leq m_1\varepsilon + \Delta/c \leq \mu\varepsilon + \Delta/c$. To account for all the choices, we have that $|\mathcal{C}_P| = O(k\mu(10c+3)^d)^k$.

Let Q be a given (μ, ε) -curve, with $d_{\mathcal{F}}(P, Q) \leq \Delta$. The curve Q consists of $k-1$ edges and induces an ordered set $(m_1, \dots, m_{k-1}) \in \{0, \dots, \mu\}^{k-1}$ representing the lengths of the edges in order. Let q_1, \dots, q_k be the vertices of Q . For all $1 \leq i \leq k$ it holds that $q_i \in \mathcal{L}_{m_{i-1}\varepsilon, \Delta}(P, q_{i-1})$, by construction.

As $d_{\mathcal{F}}(P, Q) \leq \Delta$, the first vertex q_1 lies in $B_{\Delta}(P(0))$, and thus there is a point q_1^* of the cover of $B_{\Delta}(P(0))$ consisting of balls of radius r/c , that lies at distance at most Δ/c to q_1 . For every subsequent q_i , by Lemma 27 and because $q_i \in \mathcal{L}_{m_{i-1}\varepsilon, \Delta}(P, q_{i-1})$, $q_i \in N_{\Delta/c}(\mathcal{L}_{m_{i-1}\varepsilon, (1+c^{-1})\Delta}(P, q_{i-1}^*))$ and thus there is a point q_i^* of the Δ/c -cover of $N_{\Delta/c}(\mathcal{L}_{m_{i-1}\varepsilon, (1+c^{-1})\Delta}(P, q_{i-1}^*))$ that is at distance at most Δ/c to q_i . This implies that there is an element Q^* (defined by exactly this choice of points) in \mathcal{C}_P that has distance $d_{\mathcal{F}}(Q, Q^*) \leq \Delta/c$ and thus, by Observation 5, it holds that $d_{\mathcal{F}}(Q, Q^*) \leq \Delta/c$. \blacktriangleleft

We are now ready to prove Theorem 21 as a straightforward application of Lemma 28.

Proof of Theorem 21. Let P be a (μ, ε) -curve in $\mathbb{X}^{d, k}$ and a value Δ be given. By Lemma 28, there is a family \mathcal{C}_P of curves of size $O(k\mu(43)^d)^k$ in $\mathbb{X}_{\mu\varepsilon + \Delta/4}^{d, k} \subset \mathbb{X}^{d, k}$, such that for any (μ, ε) -curve Q with $d_{\mathcal{F}}(P, Q) \leq \Delta$ there is curve $Q^* \in \mathcal{C}_P$ with $d_{\mathcal{F}}(Q, Q^*) \leq \Delta/4$. For any $Q^* \in \mathcal{C}_P$ identify some (μ, ε) -curve \widehat{Q}^* such that $d_{\mathcal{F}}(Q^*, \widehat{Q}^*) \leq \Delta/4$. If no such element exists, ignore Q^* . Otherwise for any (μ, ε) -curve Q with $d_{\mathcal{F}}(P, Q) \leq \Delta$ there is a curve \widehat{Q}^* in \mathcal{C}_P with $d_{\mathcal{F}}(Q, \widehat{Q}^*) \leq \Delta/4$, and thus by the triangle inequality there is a (μ, ε) -curve \widehat{Q}^* with $d_{\mathcal{F}}(Q, \widehat{Q}^*) \leq \Delta/2$, proving the bounded doubling dimension. \blacktriangleleft

► **Corollary 29.** *Let $k, \mu, d \in \mathbb{N}$ and $\varepsilon > 0$. The doubling dimension of the space of (μ, ε) -curves in $\mathbb{X}^{d, k}$ under the discrete Fréchet distance is bounded by $O(k(d + \log(k\mu)))$.*

Proof. This is a consequence of the proof of Lemma 28 and Theorem 21. By Observation 5, for any two curves P and Q in $\mathbb{X}^{d, k}$ it holds that $d_{\mathcal{F}}(P, Q) \leq d_{\text{d}\mathcal{F}}(P, Q)$, and thus any Δ -ball centered at a curve P under the discrete Fréchet distance is contained in the Δ -ball under the continuous Fréchet distance. In the proof of Lemma 28 the Δ -ball is covered by $\Delta/2$ -balls

under the discrete Fréchet distance, thus a proof similar to that of Theorem 21 implies the claim. ◀

3.3 Improvements for c -packed curves

In this section, we add the assumption that the curves in the space of (μ, ε) -curves are c -packed, which leads to an improvement of the above bounds on the doubling dimension.

► **Definition 30** (c -packed). *A curve P is said to be c -packed for some $c > 0$, if for any ball of radius r , the length of P inside the ball is at most cr . The set of c -packed (μ, ε) -curves in $\mathbb{X}^{d,k}$ is defined as the intersection of the set of (μ, ε) -curves with the set of c -packed curves.*

► **Lemma 31.** *Let $P \in \mathbb{X}^{d,k}$ be a polygonal c -packed curve with complexity at most k . Let $\lambda \geq 0$ and $\Delta \geq 0$ be given. Then for every $p \in \mathbb{R}^d$ there are $2c = O(c)$ points $p_1^*, \dots, p_{2c}^* \in \mathbb{R}^d$ such that*

$$\mathcal{L}_{\lambda, \Delta}(P, p) \subset \bigcup_{i=1}^{2c} B_{5\Delta}(p_i^*).$$

Proof. Assume $\lambda \geq 5\Delta$ as otherwise $\mathcal{L}_{\lambda, \Delta}(P, p) \subset B_{5\Delta}(p)$ clearly holds, implying the claim. For the sake of contradiction, assume that $\mathcal{L}_{\lambda, \Delta}(P, p)$ cannot be covered by $2c$ balls of radius 5Δ . This implies, that there are at least $2c + 1$ points $\{p_1, \dots\} =: \mathcal{P}$ in $\mathcal{L}_{\lambda, \Delta}(P, p)$ with a pairwise distance of at least 10Δ . For any $p_i \in \mathcal{P}$ we know that $\|p_i - p\| = \lambda$ and there are values $s_i < t_i$ such that $d_{\mathcal{F}}(P[s_i, t_i], \overline{pp_i}) \leq \Delta$. For any two distinct $p_i, p_j \in \mathcal{P}$ the points p_i, p_j and p form an isosceles triangle with side lengths λ, λ and $\|p_i - p_j\| \geq 10\Delta$. This implies that the distance to p_j from any point q along $\overline{pp_i}$ is at least 2Δ . This means that t_j cannot lie in the interval $[s_i, t_i]$ as $\|p_j - P(t_j)\| \leq \Delta$. This in turn implies that all intervals $[s_1, t_1], \dots$ are pairwise disjoint. We have thus identified $2c + 1$ disjoint (in the domain) subcurves of P with a total length of at least $(2c + 1)(\lambda - 2\Delta)$, contained in the ball $B_{\lambda}(p)$. However, since $\lambda \geq 5\Delta$, we have that $(2c + 1)(\lambda - 2\Delta) > c\lambda$, contradicting the fact that P is c -packed. This in turn implies the claim. ◀

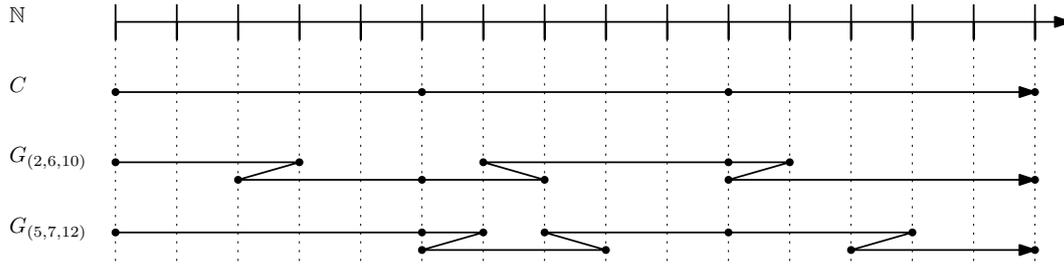
► **Corollary 32.** *Let $k, \mu, d \in \mathbb{N}$ and $c, \varepsilon > 0$. The doubling constant of the space of c -packed (μ, ε) -curves in $\mathbb{X}^{d,k}$ is bounded by $O(43^d c \mu)^k$ and thus its doubling dimension is bounded by $O(k(d + \log(c\mu)))$.*

Proof. This follows from a straightforward modification of Theorem 21 via Lemma 31. ◀

4 Lower bounds for the doubling constant of (μ, ε) -curves

In this section we want to show that the bound on the doubling dimension of $O(k(d + \log(k\mu)))$ is not too pessimistic. We begin with a straight-forward argument which implies a lower bound of $\Omega(d)$, before we discuss the lower bound of $\Omega(k \log \mu)$, resulting in a lower bound of $\Omega(d + k \log \mu)$.

The lower bound of $\Omega(d)$ follows trivially as the space of (μ, ε) -curves in $\mathbb{X}^{d,1}$ consists of every singleton in \mathbb{R}^d and thus the doubling dimension of the (μ, ε) -curves in $\mathbb{X}^{d,1}$ must be at least that of \mathbb{R}^d . For spaces with curves of higher complexity any cover of a ball (with respect to $d_{\mathcal{F}}$) the bound follows similarly, by inspecting the starting points of the curves.



■ **Figure 4** Illustration of the construction from Lemma 33 of two $(5, 1)$ -curves $G_{(2,6,10)}$ and $G_{(5,7,12)}$ in $\mathbb{X}^{1,9}$ that have Fréchet distance $1/2$ to the center curve C at the top.

4.1 Lower bound of $\Omega(k \log \mu)$

In this section, we give a lower bound for the doubling dimension of both c -packed and non- c -packed (μ, ε) -curves in $\mathbb{X}^{d,k}$ that shows the necessity of the factor μ^k in the bound of the doubling constant. The construction we give is an adaptation of a construction by Driemel et al. [15] that shows an unbounded doubling dimension for the space $(\mathbb{X}^{1,3}, d_{\mathcal{F}})$.

► **Lemma 33.** *Let $d = 1$. Given $\mu > 1$, $k \in \mathbb{N}$ and $m \leq k/2$, we can construct a $(\mu, 1)$ -curve C , of complexity $k - 2m$, and $\binom{(k-2m)\mu}{m}$ $(\mu, 1)$ -curves G_i , such that $d_{\mathcal{F}}(C, G_i) \leq 1/2$ for all i . Additionally, for all $i \neq j$ there is no $(\mu, 1)$ -curve X in $\mathbb{X}^{d,k}$ such that $d_{\mathcal{F}}(X, G_i) \leq 1/4$ and $d_{\mathcal{F}}(X, G_j) \leq 1/4$.*

The main intuition behind this proof is a variant of the main ingredient of the proof of Theorem 12. Namely, we imagine a center curve C in \mathbb{R}^1 (refer to Figure 4) that goes strictly to the right (in positive direction) for some length L . Modifying C by introducing small zig-zags of length 1 results in a new curve G that has Fréchet distance $1/2$ to C . Importantly, any curve that has Fréchet distance $1/4$ to G must also imitate the zig-zag course whenever G has a zig-zag. Now imagine there are two curves G and G' resulting from C by introducing m zig-zags each. The Fréchet distance of G and G' may be $1/2$, but as long as the introduced zig-zags are sufficiently far from one another, any $(\mu, 1)$ -curve, that has Fréchet distance $\leq 1/4$ to G has Fréchet distance $> 1/4$ to G' . Conversely, we also show that there is also no $(\mu, 1)$ -curve M such that $\{G, G'\} \subset B_{1/4}(M)$.

Proof of Lemma 35. Define C via the points (c_1, \dots, c_{k-2m+1}) , where $c_i = ((i - 1) \cdot \mu)$ for $1 \leq i \leq k - 2m + 1$. We now identify a large set of curves such that no two distinct elements of this set are at a distance at most $1/4$ to any $(\mu, 1)$ -curve in $\mathbb{X}^{1,k}$. For this, choose an ordered subset $(n_1, \dots, n_m) \subset \{0, \dots, (k - 2m)\mu - 1\}$. Clearly, there are $\binom{(k-2m)\mu}{m}$ such choices. Based on the choice, we construct a curve $G_{(n_1, \dots, n_m)}$ from C , first cutting $m + 1$ pieces C_0, \dots, C_m from C , where C_0 goes from 0 to $n_1 + 1$, C_i from n_i to $n_{i+1} + 1$ for $1 \leq i \leq m$, and C_m from n_m to $((k - 2m)\mu)$. For $1 \leq i < m$, construct curves T_i defined by two vertices $s_i = (n_i + 1)$, $t_i = (n_i)$. Then, retrieve $G_{(n_1, \dots, n_m)}$ via the concatenation $C_0 * T_1 * C_1 * \dots * T_m * C_m$. The set of curves constructed in this way forms the sought-after large set of curves. We show that $G_{(n_1, \dots, n_m)}$ consists of k edges. For this we show that every cut introduces 2 new edges. If neither n_i nor $n_i + 1$ are vertices of C , then we are done, so assume that n_i is a vertex of C . Then we introduce exactly two new edges: one from (n_i) to $(n_i + 1)$ and one from $(n_i + 1)$ to (n_i) . Similarly, if $n_i + 1$ is a vertex of C , we also introduce two new edges, one from $(n_i + 1)$ to (n_i) and one from (n_i) to $(n_i + 1)$. Observe that $d_{\mathcal{F}}(C, G_{(n_1, \dots, n_m)}) = 1/2$, as C goes to the right, whereas $G_{(n_1, \dots, n_m)}$ follows C except in the introduced pieces T_i , where it goes left for a distance of 1.

Let X be a curve that has Fréchet distance at most $1/4$ to some curve $G = G_{(n_1, \dots, n_m)}$ constructed above. Let $n_i \in (n_1, \dots, n_m)$. Then the vertex $s_i = (n_i + 1)$ and $t_i = (n_i)$ define the connecting piece T_i between C_{i-1} and C_i of G . Now, X has to first enter the interval $[n_i + 3/4, n_i + 5/4]$ and then $[n_i - 1/4, n_i + 1/4]$. As these two intervals are disjoint and $n_i < n_i + 1$, the first interval lies to the right of the second interval. By construction, the vertex of G before s_i also lies to the left of t_i . Similarly, the vertex after t_i lies to the right of s_i . Hence, X has a vertex to the left of t_i , then a vertex in $[n_i + 3/4, n_i + 5/4]$, then a vertex inside $[n_i - 1/4, n_i + 1/4]$, and subsequently a vertex to the right of s_i . Thus, a curve X also has to go to the left near $n_i + 1$. Note that as every edge of X has to have a length which is a multiple of 1 , any curve that does not go left near $n_i + 1$ has to have a Fréchet distance of at least $1/2$ to X . This has to hold for all $n_i \in (n_1, \dots, n_m)$. Since for any two distinct such constructed curves, there is a point where one travels to the left while the other does not, this then implies the claim as there are $\binom{(k-2m)\mu}{m}$ such constructed curves. ◀

► **Corollary 34.** *For $d = 1$ and given μ and k , the doubling dimension of the space of $(\mu, 1)$ -curves in $\mathbb{X}^{d,k}$ is in $\Omega(k \log \mu)$.*

Proof. Apply $m = k/3$ to Lemma 33. This, together with the fact that

$$\binom{(k/3)\mu}{k/3} \geq \frac{((k/3)\mu)^{(k/3)}}{(k/3)^{(k/3)}} = \Omega\left(\mu^{(k/3)}\right)$$

holds implies the claim. ◀

► **Theorem 35.** *For $d = 1$ and given μ , k and $\varepsilon > 0$, the doubling dimension of the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$ is in $\Omega(d + k \log \mu)$.*

Proof. This is a straight forward consequence of Corollary 34 and the trivial $\Omega(d)$ lower bound. ◀

► **Corollary 36.** *For $d = 1$, $c \geq 6$ and given μ , k and $\varepsilon > 0$, the doubling dimension of the space of c -packed (μ, ε) -curves in $\mathbb{X}^{d,k}$ is in $\Omega(d + k \log \mu)$.*

Proof. The constructed curves are clearly 6 -packed. Thus, Theorem 35 implies the claim. ◀

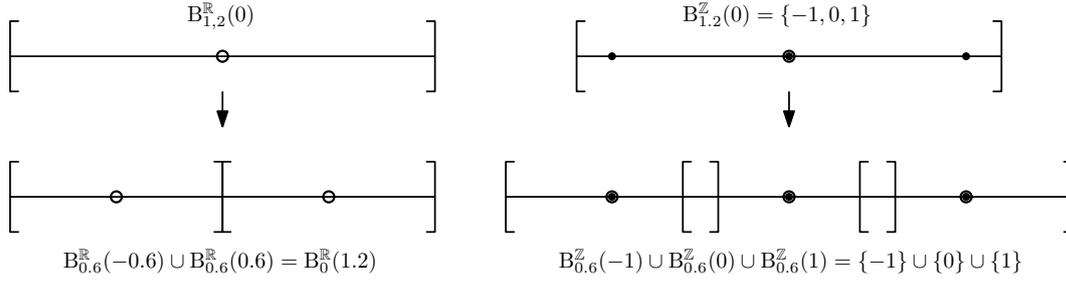
5 Approximate Nearest Neighbour

Har-Peled et al. [22] showed that $(1 + \varepsilon)$ -ANN can be solved in metric spaces of bounded doubling dimension.

► **Theorem 37** ([22]). *Given a set S of n points in a metric space \mathcal{M} of bounded doubling dimension ν , one can construct a data structure for answering $(1 + \varepsilon)$ -approximate nearest neighbour queries. The query time is $2^{O(\nu)} \log n + \varepsilon^{-O(\nu)}$, the expected preprocessing time is $2^{O(\nu)} n \log n$ and the space used is $2^{O(\nu)} n$.*

A careful reading reveals an important specification for our purposes, namely, that the doubling dimension is that of the n -point metric space defined by S induced by the metric space \mathcal{M} and not of the ambient metric space \mathcal{M} . Note that by Lemma 38 the doubling dimension of the metric space induced on the subset is at most twice the doubling dimension of the ambient space. For an example where the doubling dimension increases, refer to Figure 5.

► **Lemma 38.** *Let $(\mathcal{M}, d_{\mathcal{M}})$ be a metric space, and let S be some subset of \mathcal{M} . Then the doubling dimension of $(S, d_{\mathcal{M}})$ is at most twice the doubling dimension of $(\mathcal{M}, d_{\mathcal{M}})$.*



■ **Figure 5** Example of a subset \mathbb{Z} of the metric space \mathbb{R} whose doubling dimension is larger than that of its ambient space. The disk centers are marked by circles.

Proof. Let ν be the doubling dimension of \mathcal{M} . Let $s \in S$ and $r > 0$ be given and let $B_r^S(s) \subset S$ be the ball in S , of radius r and centered at s , i.e., $B_r^S(s) = B_r^{\mathcal{M}}(s) \cap S$. But $B_r^{\mathcal{M}}(s)$ can be covered by $(2^\nu)^2$ balls of radius $r/4$. For any such ball check if the intersection with S is nonempty. If this is the case, pick some element from this intersection, and center a ball in S of radius $r/2$ around it. Clearly, any such larger ball contains the intersection of the smaller ball with S . Therefore, $B_r^S(s)$ is contained in the union of at most $(2^\nu)^2$ balls of radius $r/2$ in S , which implies the claim, as $\log_2((2^\nu)^2) = 2\nu$. ◀

By Theorem 21, we know that the doubling dimension of the space of (μ, ε) -curves in $\mathbb{X}^{d,k}$ is bounded. We further know that for any $\varepsilon > 0$ we can map any curve of $\mathbb{X}_\Lambda^{d,k}$ into the space of $(\lceil \Lambda/\varepsilon \rceil + 1, \varepsilon)$ -curves in $\mathbb{X}^{d,k}$ with a distortion of at most $\varepsilon/2$, by Lemma 14. Hence, Theorem 37 together with Lemma 38 imply Theorem 6, a central piece to constructing a data structure solving the $(1 + \varepsilon)$ -ANN problem for polygonal curves under the Fréchet distance.

► **Theorem 6.** *Given a set S of n polygonal curves in $\mathbb{X}_\Lambda^{d,k}$ and parameters $0 < \varepsilon < 1$ and $\varepsilon' > 0$, one can construct a data structure that for given $q \in \mathbb{X}^{d,k}$ outputs an element $s^* \in S$ such that for all $s \in S$ it holds that $d_{\mathcal{F}}(s^*, q) \leq (1 + \varepsilon)d_{\mathcal{F}}(s, q) + \varepsilon'$. The query time is $O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k \log n + O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^{-k \log(\varepsilon)}$, the expected preprocessing time is $O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k n \log n$ and the space used is $O(2^{O(d)}k(1 + \Lambda/\varepsilon'))^k n$.*

Proof. Define $\hat{\varepsilon} = \varepsilon'/2$. Let $\mu = \lceil \Lambda/\hat{\varepsilon} \rceil + 1 = \Theta(1 + \Lambda/\varepsilon')$. We begin by simplifying every polygonal curve $s \in S$ via Lemma 14, resulting in a set S' of $(\mu, \hat{\varepsilon})$ -curves. This takes $O(\log(\mu)nk)$ time, which is in $O(2^\nu n)$. As S' lies in the space of $(\mu, \hat{\varepsilon})$ -curves, the doubling dimension of the set S' with the Fréchet distance is bounded by $\nu = O(k(d + \log(k(1 + \Lambda/\varepsilon'))))$ via Theorem 21 and Lemma 38. Note that for every $s \in S$ and its simplification $s' \in S'$ it holds that $d_{\mathcal{F}}(s, s') \leq \hat{\varepsilon}/2$. We apply Theorem 37 to the set S' and ε . Note that Theorem 37 assumes that the distance between any two points in the metric space of (μ, ε) -curves can be computed in $O(1)$ time. However, the computation of the continuous Fréchet distance takes polynomial time in k . On the other hand, both $2^{k \log k}$ and $\varepsilon^{-k \log k}$ dominate $\text{poly}(k)$ for $\varepsilon < 1$. Thus the running time is indeed as claimed. We then query the data structure with q , returning an element \hat{s}' such that for every $s' \in S'$ it holds that $d_{\mathcal{F}}(q, \hat{s}') \leq (1 + \varepsilon)d_{\mathcal{F}}(q, s')$. Lastly, the element of S returned by the data structure will be the element $\hat{s} \in S$ which corresponds to \hat{s}' . We then get for every $s \in S$ that

$$\begin{aligned}
 d_{\mathcal{F}}(q, \hat{s}) &\leq d_{\mathcal{F}}(q, \hat{s}') + \hat{\varepsilon}/2 \leq (1 + \varepsilon)d_{\mathcal{F}}(q, s') + \hat{\varepsilon}/2 \leq (1 + \varepsilon)(d_{\mathcal{F}}(q, s) + \hat{\varepsilon}/2) + \hat{\varepsilon}/2 \\
 &\leq (1 + \varepsilon)d_{\mathcal{F}}(q, s) + \hat{\varepsilon} + \varepsilon\hat{\varepsilon}/2 = (1 + \varepsilon)d_{\mathcal{F}}(q, s) + \hat{\varepsilon}(1 + \varepsilon/2) \\
 &\leq (1 + \varepsilon)d_{\mathcal{F}}(q, s) + \varepsilon'.
 \end{aligned}$$

To get rid of the additive error, we want to set ε' to $O(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s'))$. As ε' impacts the running time, we reformulate the running time in terms of the spread.

Note that this measure is scale- and translation-independent, as translating all elements in S by the same offset changes neither their pairwise Fréchet distances nor the length of the edges of the curves. Similarly scaling the elements of S scales both the pairwise Fréchet distances as well as the length of the edges of the curves, thus not changing the bundledness.

► **Lemma 39.** *Given a set of curves $S \in \mathbb{X}^{d,k}$, then*

$$\mathcal{G}(S)^{-1} = O(\Phi(S))$$

where $\Phi(S)$ denotes the spread of the set of vertices and edges of curves in S .

Proof. Observe that no two curves in S have a Fréchet distance of 0. Note that the Fréchet distance between two curves P and Q is approximated up to a constant by the Euclidean distance of either one vertex of P and one vertex of Q or the Euclidean distance of a vertex and an edge (one of P and one of Q) as sets. Thus

$$\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') = \Omega \left(\min_{\substack{o, o' \in V(S) \cup E(S) \\ d(o, o') > 0}} d(o, o') \right)$$

where $V(S)$ denotes the set of vertices and $E(S)$ denotes the set of edges defining the curves in S . Further note that as the length of any edge of a curve s in S is defined as the distance of two vertices defining s . As such, observe that

$$\max_{s \in S, e \in E(s)} \|e\| \leq \max_{p, q \in V(S)} d(p, q) \leq \max_{o, o' \in V(S) \cup E(S)} d(o, o').$$

Thus the claim follows. ◀

► **Theorem 9.** *Given a set S of n polygonal curves in $\mathbb{X}^{d,k}$ and $0 < \varepsilon \leq 1$ one can construct a data structure answering $(1 + \varepsilon)$ -approximate nearest neighbour queries. The query time is $F(d, k, S, \varepsilon) \log n + F(d, k, S, \varepsilon)^{-\log(\varepsilon)}$, the expected preprocessing time is $F(d, k, S, \varepsilon)n \log n$ and the space used is $F(d, k, S, \varepsilon)n$, where $F(d, k, S, \varepsilon) = O(2^{O(d)}k(1 + \mathcal{G}(S)^{-1}\varepsilon^{-1}))^k$.*

Proof. Let $\varepsilon' = \varepsilon/4$ and $\varepsilon'' = \varepsilon'(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s'))$. Let $E(S)$ be the set of edges of curves in S and let further $\Lambda = \max_{e \in E(S)} \|e\|$, thus clearly $S \subset \mathbb{X}_{\Lambda}^{d,k}$. We then apply Theorem 6 with ε' and ε'' resulting in the described data structure. Let $q \in \mathbb{X}^{d,k}$ be given. Let s^* be the element in S minimizing the distance to q . Querying the data structure with q results in an element \hat{s} with the property that

$$d_{\mathcal{F}}(q, \hat{s}) \leq (1 + \varepsilon')d_{\mathcal{F}}(q, s^*) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') \right).$$

We now look at two different cases. First assume $(2 + \varepsilon')d_{\mathcal{F}}(q, s^*) < (1 - \varepsilon') \min_{s \neq s' \in S} d_{\mathcal{F}}(s, s')$. But then for every $s' \neq s^*$ we know that

$$\begin{aligned} d_{\mathcal{F}}(q, s') &\geq d_{\mathcal{F}}(s^*, s') - d_{\mathcal{F}}(q, s^*) \geq \min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') - d_{\mathcal{F}}(q, s^*) \\ &> (1 - \varepsilon') \left(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') \right) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') \right) - d_{\mathcal{F}}(q, s^*) \\ &= (1 + \varepsilon')d_{\mathcal{F}}(q, s^*) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') \right) \end{aligned}$$

and thus $\widehat{s} = s^*$, implying $d_{\mathcal{F}}(q, \widehat{s}) \leq (1 + \varepsilon)d_{\mathcal{F}}(q, s^*)$.

Now assume $(2 + \varepsilon')d_{\mathcal{F}}(q, s^*) \geq (1 - \varepsilon') \min_{s \neq s' \in S} d_{\mathcal{F}}(s, s')$. Then we know that

$$\begin{aligned} d_{\mathcal{F}}(q, \widehat{s}) &\leq (1 + \varepsilon')d_{\mathcal{F}}(q, s^*) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') \right) \\ &\leq (1 + \varepsilon')d_{\mathcal{F}}(q, s^*) + \varepsilon' \left(\frac{2 + \varepsilon'}{1 - \varepsilon'} \right) d_{\mathcal{F}}(q, s^*). \end{aligned}$$

Now since $\varepsilon \leq 1$, we know that $\varepsilon' \leq 1/4$ and thus $\frac{2 + \varepsilon'}{1 - \varepsilon'} \leq 3$. This then concludes the case-distinction, as

$$d_{\mathcal{F}}(q, \widehat{s}) \leq (1 + \varepsilon')d_{\mathcal{F}}(q, s^*) + \varepsilon' \left(\frac{2 + \varepsilon'}{1 - \varepsilon'} \right) d_{\mathcal{F}}(q, s^*) \leq (1 + 4\varepsilon')d_{\mathcal{F}}(q, s^*) = (1 + \varepsilon)d_{\mathcal{F}}(q, s^*).$$

Regarding the running time, observe that $\Lambda / \min_{s \neq s' \in S} d_{\mathcal{F}}(s, s') = \mathcal{G}(S)^{-1}$. Hence, as $\varepsilon' = \Theta(\varepsilon)$ and $\varepsilon'' = \Theta(\varepsilon \min_{s \neq s' \in S} d_{\mathcal{F}}(s, s'))$, the preprocessing time, query time and space is as claimed. \blacktriangleleft

► **Corollary 10.** *Given a set S of n polygonal curves in $\mathbb{X}^{d,k}$ and $0 < \varepsilon \leq 1$ one can construct a data structure answering $(1 + \varepsilon)$ -approximate nearest neighbour queries. The query time is $F(d, k, S, \varepsilon) \log n + F(d, k, S, \varepsilon)^{-\log(\varepsilon)}$, the expected preprocessing time is $F(d, k, S, \varepsilon)n \log n$ and the space used is $F(d, k, S, \varepsilon)n$, where $F(d, k, S, \varepsilon) = O(2^{O(d)}k\Phi(S)\varepsilon^{-1})^k$, where $\Phi(S)$ denotes the spread of the set of vertices and edges of the curves in S .*

Proof. This is a direct consequence of Theorem 9 and Lemma 39 together with the fact that the spread of any collection of sets is at least 1, implying $(1 + \Phi(S)\varepsilon^{-1}) = O(\Phi(S)\varepsilon^{-1})$, since both $\Phi(S) \geq 1$ and $\varepsilon \leq 1$. \blacktriangleleft

6 Generalization to arbitrary metric spaces

In this section, we observe that our techniques directly generalize to other metric spaces. In particular, we can solve $(1 + \varepsilon)$ -ANN in any space with unbounded doubling dimension as long as a subspace with bounded doubling dimension that is close (under the Gromov-Hausdorff distance) exists. This approach is closely related to that of Sheehy and Seth [34] in that they extended Clarkson's algorithm for finding a λ -net in some metric space X if there is another space of bounded doubling dimension with Gromov-Hausdorff distance at most $\lambda/3$ to X .

► **Definition 40** (tractably nearly-doubling space). *Let \mathcal{M} be a metric space. We say \mathcal{M} is tractably nearly-doubling if there are functions $\nu_{\mathcal{M}} : \mathbb{R} \rightarrow \mathbb{R}$ and $\rho_{\mathcal{M}} : \mathcal{M} \times \mathbb{R} \rightarrow \mathbb{R}$ such that for every $\varepsilon > 0$ there is some subspace $\mathcal{M}_{\varepsilon} \subset \mathcal{M}$ with doubling dimension at most $\nu_{\mathcal{M}}(\varepsilon)$ and projection $\pi_{\varepsilon} : \mathcal{M} \rightarrow \mathcal{M}_{\varepsilon}$. Furthermore we require $d_{\mathcal{M}}(s, \pi_{\varepsilon}(s)) \leq \varepsilon$ for any $s \in \mathcal{M}$ and the element $\pi_{\varepsilon}(s) \in \mathcal{M}_{\varepsilon}$ can be computed in $\rho_{\mathcal{M}}(s, \varepsilon)$ time.*

Observe that for a tractably nearly-doubling space \mathcal{M} the Gromov-Hausdorff distance $d_{GH}(\mathcal{M}, \mathcal{M}_{\varepsilon})$ is at most ε . Note that by Lemma 14 and Theorem 21, the space $\mathbb{X}_{\Lambda}^{d,k}$ is tractably nearly-doubling for any $d, k \in \mathbb{N}$ and $\Lambda > 0$. The proof of the following lemma and theorem conceptually are the same proof as for Theorem 6 and Theorem 9.

► **Lemma 41.** *Let \mathcal{M} be a metric space, and let $S \subset \mathcal{M}$ be a set of n points and $\varepsilon > 0$. Let $\mathcal{M}' \subset \mathcal{M}$ be some subspace of doubling dimension ν , and let $S' \subset \mathcal{M}'$ be another set with a*

surjection $\pi : S \rightarrow S'$ with a running time of $O(1)$. Then one can compute a data structure which for given $q \in \mathcal{M}$ outputs $\hat{s} \in S$, such that for every $s \in S$

$$d_{\mathcal{M}}(q, \hat{s}) \leq (1 + \varepsilon)d_{\mathcal{M}} + (2 + \varepsilon) \left(\max_{s \in S} d_{\mathcal{M}}(s, \pi_{\varepsilon}(s)) \right).$$

The expected preprocessing time is $2^{O(\nu)} + \varepsilon^{-O(\nu)}$ and the space used is $2^{O(d)}n \log n$, with the query time being $2^{O(\nu)} \log n + \varepsilon^{-O(\nu)}$, where ν is the doubling dimension of \mathcal{M}' .

Proof. This proof is parallel to the proof of Theorem 6. We first store the inverse map $\pi^{-1} : S' \rightarrow S$ in $O(n)$ time. We then compute the data structure from Theorem 37 on the set S' . This takes $2^{O(\nu)} + \varepsilon^{-O(\nu)}$ expected preprocessing time and $2^{O(\nu)}n \log n$ space. Further the query time for this data structure is $2^{O(\nu)} \log n + \varepsilon^{-O(\nu)}$, as ν is the doubling dimension of \mathcal{M}' .

Then for any $q \in \mathcal{M}$ the data structure outputs a point $\hat{s}' \in S'$. We then output $\hat{s} = \pi^{-1}(\hat{s}')$. Then overall for every point $s' \in S'$ it holds that $d_{\mathcal{M}}(q, \hat{s}') \leq (1 + \varepsilon)d_{\mathcal{M}}(q, s')$. Then for any $s' \in S'$ it holds that

$$d_{\mathcal{M}}(q, \hat{s}) \leq d_{\mathcal{M}}(q, \hat{s}') + d_{\mathcal{M}}(\hat{s}, \hat{s}') \leq (1 + \varepsilon)d_{\mathcal{M}}(q, s') + \max_{s \in S} d_{\mathcal{M}}(s, \pi_{\varepsilon}(s))$$

And thus for any $s \in S$ it holds that

$$\begin{aligned} d_{\mathcal{M}}(q, \hat{s}) &\leq (1 + \varepsilon)(d_{\mathcal{M}}(q, s) + d_{\mathcal{M}}(s, s')) + \max_{s \in S} d_{\mathcal{M}}(s, \pi_{\varepsilon}(s)) \\ &\leq (1 + \varepsilon)d_{\mathcal{M}}(q, s) + (2 + \varepsilon) \left(\max_{s \in S} d_{\mathcal{M}}(s, \pi_{\varepsilon}(s)) \right), \end{aligned}$$

implying the claim. ◀

► **Theorem 42.** Let \mathcal{M} be a metric space that is tractably nearly-doubling. Let $S \subset \mathcal{M}$ be a set of n points. Then for every $\varepsilon > 0$ one can construct a data structure which for any given $q \in \mathcal{M}$ returns a point $\hat{s} \in S$ such that for every $s \in S$ it holds that

$$d_{\mathcal{M}}(q, \hat{s}) \leq (1 + \varepsilon)d_{\mathcal{M}}(q, s).$$

The expected preprocessing time is given by $\sum_{s \in S} \rho_{\mathcal{M}}(s, e) + 2^{O(\nu_{\mathcal{M}}(e))} + \varepsilon^{-O(\nu_{\mathcal{M}}(e))}$ and the space used is in $2^{O(\nu_{\mathcal{M}}(e))}n \log n$. The query time is given by $2^{O(\nu_{\mathcal{M}}(e))} \log n + \varepsilon^{-O(\nu_{\mathcal{M}}(e))}$, where $e = (\varepsilon/4) \min_{s \neq s' \in S} d_{\mathcal{M}}(s, s')$.

Proof. This proof follows the proof of Theorem 9. Just like before, it is a consequence from Lemma 41, by selecting specific parameters.

First let $\varepsilon' = \varepsilon/4$ and $\varepsilon'' = \varepsilon \min_{s \neq s' \in S} d_{\mathcal{M}}(s, s')$. We then apply Lemma 41 with the subspace $\mathcal{M}_{\varepsilon''}$, the set $\pi_{\varepsilon''}(S) \subset \mathcal{M}_{\varepsilon''} \subset \mathcal{M}$ and the map $\pi_{\varepsilon''}$, where the doubling dimension of $\mathcal{M}_{\varepsilon''}$ is given by $\nu_{\mathcal{M}}(\varepsilon'')$, as \mathcal{M} is tractably nearly-doubling. Thus $d_H(S, \pi_{\varepsilon''}(S)) \leq \varepsilon''$. The expected preprocessing time is given by $\sum_{s \in S} \rho_{\mathcal{M}}(s, \varepsilon) + 2^{O(\nu_{\mathcal{M}}(\varepsilon''))} + \varepsilon^{-O(\nu_{\mathcal{M}}(\varepsilon''))}$ and the space used is in $2^{O(\nu_{\mathcal{M}}(\varepsilon''))}n \log n$. The query time is given by $2^{O(\nu_{\mathcal{M}}(\varepsilon''))} \log n + \varepsilon^{-O(\nu_{\mathcal{M}}(\varepsilon''))}$.

Let $q \in \mathcal{M}$ be given. Let \hat{s} be the output of the data structure, and let s^* be the point in S minimizing the distance to q . We again start, by assuming $(2 + \varepsilon')d_{\mathcal{M}}(q, s^*) \leq (1 + \varepsilon') \min_{s \neq s' \in S} d_{\mathcal{M}}(s, s')$. But then for every $s' \neq s^*$ we know that

$$\begin{aligned} d_{\mathcal{M}} &\geq d_{\mathcal{M}}(s^*, s') - d_{\mathcal{M}}(q, s^*) \geq \min_{s \neq s' \in S} d_{\mathcal{M}}(s, s') - d_{\mathcal{M}}(q, s^*) \\ &> (1 - \varepsilon') \left(\min_{s \neq s' \in S} d_{\mathcal{M}}(s, s') \right) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{M}}(s, s') \right) - d_{\mathcal{M}}(q, s^*) \\ &= (1 + \varepsilon')d_{\mathcal{M}}(q, s^*) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{M}}(s, s') \right) \end{aligned}$$

and thus $\hat{s} = s^*$, implying $d_{\mathcal{M}}(q, \hat{s}) \leq (1 + \varepsilon)d_{\mathcal{M}}(q, s^*)$.

Now assume $(2 + \varepsilon')d_{\mathcal{M}}(q, s^*) \geq (1 - \varepsilon') \min_{s \neq s' \in S} d_{\mathcal{M}}(s, s')$. Then we know that

$$\begin{aligned} d_{\mathcal{M}}(q, \hat{s}) &\leq (1 + \varepsilon')d_{\mathcal{M}}(q, s^*) + \varepsilon' \left(\min_{s \neq s' \in S} d_{\mathcal{M}}(s, s') \right) \\ &\leq (1 + \varepsilon')d_{\mathcal{M}}(q, s^*) + \varepsilon' \left(\frac{2 + \varepsilon'}{1 - \varepsilon'} \right) d_{\mathcal{M}}(q, s^*). \end{aligned}$$

Now since $\varepsilon \leq 1$, we know that $\varepsilon' \leq 1/4$ and thus $\frac{2 + \varepsilon'}{1 - \varepsilon'} \leq 3$. This then concludes the proof, as

$$\begin{aligned} d_{\mathcal{M}}(q, \hat{s}) &\leq (1 + \varepsilon')d_{\mathcal{M}}(q, s^*) + \varepsilon' \left(\frac{2 + \varepsilon'}{1 - \varepsilon'} \right) d_{\mathcal{M}}(q, s^*) \\ &\leq (1 + 4\varepsilon')d_{\mathcal{M}}(q, s^*) = (1 + \varepsilon)d_{\mathcal{M}}(q, s^*). \end{aligned} \quad \blacktriangleleft$$

7 Conclusion

In this work we established that despite the unbounded doubling dimension of the metric space of the Fréchet distance of curves $\mathbb{X}^{d,k}$, there are spaces which are arbitrarily close to $\mathbb{X}^{d,k}$ which all have bounded doubling dimension, depending on the distance to $\mathbb{X}^{d,k}$. We then constructed an approximate nearest neighbour data structure for $\mathbb{X}^{d,k}$ by answering approximate nearest neighbour queries in these arbitrarily close spaces via well-established data structures which can be constructed in spaces of bounded doubling dimension. As a special case we considered the setting of c -packed curves to which our approach extends, resulting in an improved running time in this special setting.

We further gave a constructive argument as to why the doubling dimension of the space of (μ, ε) -curves in $\mathbb{X}_{\Lambda}^{d,k}$ is large. The gap between the upper bound on the doubling dimension and the construction is quite small — especially in the case of c -packed curves for constant c — but any further improvements to this gap would be interesting to see. We do not believe that either of the given bounds are necessarily tight.

Intuitively, it seems reasonable to assume that the ‘dimension’ of the problem should be at least kd for $d \in O(\log nm)$ and $k = m$. It is well-known [23] that the L_{∞} -metric in \mathbb{R}^d embeds isometrically into a Fréchet metric space of one-dimensional curves of complexity $k = 3d$. Recently, Rohde and Psarros showed that random projections can be used to obtain dimensionality reductions for the Fréchet metric [32] when $d \in \Omega(\log(nm))$.

An important future research direction is the reduction of the dependence on the bandedness (or spread of the underlying set of vertices and edges) in the running time. In our case the dependence on the spread is a result of turning the additive error of the Gromov-Hausdorff distance between $\mathbb{X}^{d,k}$ and the space of (μ, ε) -curves into a multiplicative error. With our approach this seems to be inevitable.

It would further be interesting to see what other results from spaces of bounded doubling dimension can be extended to spaces (not restricted to that of polygonal curves) of unbounded doubling dimension in this way.

References

- 1 Peyman Afshani and Anne Driemel. On the complexity of range searching among curves. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 898–917. Society for Industrial and Applied Mathematics, Philadelphia, PA, jan 2018. [arXiv:1707.04789](https://arxiv.org/abs/1707.04789), [doi:10.1137/1.9781611975031.58](https://doi.org/10.1137/1.9781611975031.58).

- 2 Pankaj K. Agarwal, Rinat Ben Avraham, Haim Kaplan, and Micha Sharir. Computing the Discrete Fréchet Distance in Subquadratic Time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, page 156–167. Society for Industrial and Applied Mathematics, 2013. doi:10.1137/1.9781611973105.12.
- 3 Helmut Alt, Bernd Behrends, and Johannes Blömer. Approximate matching of polygonal shapes. *Annals of Mathematics and Artificial Intelligence*, 13:251–265, 1995. doi:10.1007/BF01530830.
- 4 Helmut Alt and Michael Godau. Computing the Fréchet Distance between Two Polygonal Curves. *International Journal of Computational Geometry & Applications*, 5:75–91, 1995. doi:10.1142/S0218195995000064.
- 5 Boris Aronov, Sariel Har-Peled, Christian Knauer, Yusu Wang, and Carola Wenk. Fréchet Distance for Curves, Revisited. In *Proceedings of the 14th Conference on Annual European Symposium - Volume 14*, pages 52–63. Springer-Verlag, 2006. doi:10.1007/11841036_8.
- 6 Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, and Angela Y. Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM*, 45(6):891–923, nov 1998. doi:10.1145/293347.293348.
- 7 J.S. Beis and D.G. Lowe. Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1000–1006. IEEE Comput. Soc, 1997. doi:10.1109/CVPR.1997.609451.
- 8 Eckart Bindewald and Bruce A Shapiro. Rna secondary structure prediction from sequence alignments using a network of k-nearest neighbor classifiers. *RNA*, 12(3):342–352, 2006. doi:10.1261/rna.2164906.
- 9 Prosenjit Bose, Sébastien Collette, Stefan Langerman, Anil Maheshwari, Pat Morin, and Michiel Smid. Sigma-local graphs. *Journal of Discrete Algorithms*, 8(1):15–23, 2010. doi:https://doi.org/10.1016/j.jda.2008.10.002.
- 10 Karl Bringmann, Anne Driemel, André Nusser, and Ioannis Psarros. Tight Bounds for Approximate Near Neighbor Searching for Time Series under the Fréchet Distance. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 517–550. Society for Industrial and Applied Mathematics, Philadelphia, PA, jan 2022. arXiv:2107.07792v2, doi:10.1137/1.9781611977073.25.
- 11 Kevin Buchin, Maike Buchin, Wouter Meulemans, and Wolfgang Mulzer. Four Soviets Walk the Dog: Improved Bounds for Computing the Fréchet Distance. *Discrete & Computational Geometry*, 58(1):180–216, 2017. doi:10.1007/s00454-017-9878-7.
- 12 Timothy M. Chan and Zahed Rahmati. An improved approximation algorithm for the discrete Fréchet distance. *Information Processing Letters*, 138:72–74, 2018. doi:10.1016/j.ipl.2018.06.011.
- 13 Connor Colombe and Kyle Fox. Approximating the (Continuous) Fréchet Distance. In *37th International Symposium on Computational Geometry*, volume 189, pages 26:1–26:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.SoCG.2021.26.
- 14 Anne Driemel, Sariel Har-Peled, and Carola Wenk. Approximating the Fréchet Distance for Realistic Curves in near Linear Time. In *Proceedings of the Twenty-Sixth Annual Symposium on Computational Geometry*, pages 365–374. Association for Computing Machinery, 2010. doi:10.1145/1810959.1811019.
- 15 Anne Driemel, Amer Krivošija, and Christian Sohler. Clustering time series under the fréchet distance. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 766–785. SIAM, 2016. doi:1.9781611974331.ch55.
- 16 Jeff Erickson. Local polyhedra and geometric graphs. *Computational Geometry*, 31(1):101–125, 2005. Special Issue on the 19th Annual Symposium on Computational Geometry - SoCG 2003. doi:https://doi.org/10.1016/j.comgeo.2004.08.004.
- 17 Arnold Filtser, Omrit Filtser, and Matthew J. Katz. Approximate Nearest Neighbor for Curves: Simple, Efficient, and Deterministic. *Algorithmica*, 85(5):1490–1519, may 2023. arXiv:1902.07562, doi:10.1007/s00453-022-01080-1.

- 18 Joachim Gudmundsson, Marc van Kreveld, and Bettina Speckmann. Efficient detection of motion patterns in spatio-temporal data sets. In *Proceedings of the 12th Annual ACM International Workshop on Geographic Information Systems*, page 250–257. ACM, 2004. doi:10.1145/1032222.1032259.
- 19 Leonidas Guibas, John Hershberger, Joseph Mitchell, and Jack Snoeyink. Approximating Polygons and Subdivisions with Minimum-Link Paths. In *International Journal of Computational Geometry & Applications*, volume 3, pages 151–162. World Scientific, 1994. doi:10.1007/3-540-54945-5_59.
- 20 Anupam Gupta, Robert Krauthgamer, and James R Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 534–543. IEEE, 2003. doi:10.1109/SFCS.2003.1238226.
- 21 Sariel Har-Peled, Piotr Indyk, and Rajeev Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of Computing*, 8(14):321–350, 2012. doi:10.4086/toc.2012.v008a014.
- 22 Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006. doi:10.1137/S0097539704446281.
- 23 Piotr Indyk, Jiří Matoušek, and Anastasios Sidiropoulos. Low-distortion embeddings of finite metric spaces. In *Handbook of discrete and computational geometry*, pages 211–231. Chapman and Hall/CRC, 2017.
- 24 David R. Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, page 741–750. Association for Computing Machinery, 2002. doi:10.1145/509907.510013.
- 25 David Kirkpatrick. Optimal search in planar subdivisions. *SIAM Journal on Computing*, 12(1):28–35, 1983. doi:10.1137/0212002.
- 26 Eyal Kushilevitz, Rafail Ostrovsky, and Yuval Rabani. Efficient Search for Approximate Nearest Neighbor in High Dimensional Spaces. *SIAM Journal on Computing*, 30(2):457–474, jan 2000. doi:10.1137/S0097539798347177.
- 27 Yuchun Lee. Handwritten digit recognition using k nearest-neighbor, radial-basis function, and backpropagation neural networks. *Neural computation*, 3(3):440–449, 1991. doi:10.1162/neco.1991.3.3.440.
- 28 Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, 1969.
- 29 Majid Mirzanezhad. On the approximate nearest neighbor queries among curves under the fréchet distance. *arXiv preprint arXiv:2004.08444*, 2020. URL: <https://arxiv.org/abs/2004.08444>.
- 30 Majid Mirzanezhad. On Approximate Near-Neighbors Search under the (Continuous) Fréchet Distance in Higher Dimensions. *Information Processing Letters*, page 106405, 2023. doi: <https://doi.org/10.1016/j.ip1.2023.106405>.
- 31 Harold R. Parks. The Volume of the Unit n-Ball. *Mathematics Magazine*, 86(4):270–274, 2013. doi:10.4169/math.mag.86.4.270.
- 32 Ioannis Psarros and Dennis Rohde. Random projections for high-dimensional curves, 2022. doi:10.48550/ARXIV.2207.07442.
- 33 Michael Ian Shamos and Dan Hoey. Closest-point problems. In *16th Annual Symposium on Foundations of Computer Science (sfcs 1975)*, pages 151–162, 1975. doi:10.1109/SFCS.1975.8.
- 34 Donald R. Sheehy and Siddharth S. Sheth. Nearly-Doubling Spaces of Persistence Diagrams. In *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224, pages 60:1–60:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SocG.2022.60.

- 35 Mai Shouman, Tim Turner, and Rob Stocker. Applying k-nearest neighbour in diagnosing heart disease patients. *International Journal of Information and Education Technology*, 2(3):220–223, 2012. doi:10.7763/IJJET.2012.V2.114.
- 36 Kunal Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 281–290, 2004. doi:10.1145/1007352.1007399.