



Approximate Nearest Neighbor for Curves: Simple, Efficient, and Deterministic

Arnold Filtser¹ · Omrit Filtser²  · Matthew J. Katz³

Received: 20 June 2022 / Accepted: 28 November 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

In the $(1 + \varepsilon, r)$ -approximate near-neighbor problem for curves (ANNC) under some similarity measure δ , the goal is to construct a data structure for a given set \mathcal{C} of curves that supports approximate near-neighbor queries: Given a query curve Q , if there exists a curve $C \in \mathcal{C}$ such that $\delta(Q, C) \leq r$, then return a curve $C' \in \mathcal{C}$ with $\delta(Q, C') \leq (1 + \varepsilon)r$. There exists an efficient reduction from the $(1 + \varepsilon)$ -approximate nearest-neighbor problem to ANNC, where in the former problem the answer to a query is a curve $C \in \mathcal{C}$ with $\delta(Q, C) \leq (1 + \varepsilon) \cdot \delta(Q, C^*)$, where C^* is the curve of \mathcal{C} most similar to Q . Given a set \mathcal{C} of n curves, each consisting of m points in d dimensions, we construct a data structure for ANNC that uses $n \cdot O(\frac{1}{\varepsilon})^{md}$ storage space and has $O(md)$ query time (for a query curve of length m), where the similarity measure between two curves is their discrete Fréchet or dynamic time warping distance. Our method is simple to implement, deterministic, and results in an exponential improvement in both query time and storage space compared to all previous bounds. Further, we also consider the *asymmetric* version of ANNC, where the length of the query curves is $k \ll m$, and obtain essentially the same storage and query bounds as above, except that m is replaced by k . Finally, we apply our method to a version of approximate range counting for curves and achieve similar bounds.

A preliminary version of this paper excluding Sect. 6 and Appendices B and C has appeared in ICALP'20 [15].

✉ Omrit Filtser
omrit.filtser@gmail.com

Arnold Filtser
arnold273@gmail.com

Matthew J. Katz
matya@cs.bgu.ac.il

¹ Bar-Ilan University, Ramat Gan, Israel

² The Open University of Israel, Ra'anana, Israel

³ Ben-Gurion University of the Negev, Be'er Sheva, Israel

Keywords Nearest neighbor search · Polygonal curves · Fréchet distance · Dynamic time warping

Contents

1	Introduction
1.1	Related Work
1.2	Technical Ideas
2	Preliminaries
3	Discrete Fréchet Distance (DFD)
4	The Asymmetric Setting Under DFD
5	$\ell_{p,2}$ -Distance of Polygonal Curves
6	The Asymmetric Setting Under $\ell_{p,2}$ -distance
6.1	Simplification Under $\ell_{p,2}$ -Distance
6.2	The Data Structure
7	Approximate Range Counting
8	Simplification in d -dimensions
	Appendix: A Deterministic Construction Using a Prefix Tree
	Appendix: A.1 Discrete Fréchet Distance
	Appendix: A.2 $\ell_{p,2}$ -Distance
	Appendix: B Dealing with Query Curves and Input Curves of Varying Size
	Appendix: C One-Way Alignments
	References

1 Introduction

Nearest neighbor search is a fundamental and well-studied problem that has various applications in machine learning, data analysis, and classification. This important task also arises in applications where the recorded instances are trajectories or polygonal curves modeling, for example, epigenetic and surgical processes, market value fluctuations, population growth, the number of the requests per hour received at some web-page, and even the response of a football player in a given situation.

Let \mathcal{C} be a set of n curves, each consisting of at most m points in d dimensions, and let δ be some distance measure for curves. In the *nearest-neighbor* problem for curves, the goal is to construct a data structure for \mathcal{C} that supports nearest-neighbor queries, that is, given a query curve Q of length at most m , return the curve $C^* \in \mathcal{C}$ closest to Q (according to δ). The approximation version of this problem is the $(1 + \varepsilon)$ -*approximate nearest-neighbor* problem, where the answer to a query Q is a curve $C \in \mathcal{C}$ with $\delta(Q, C) \leq (1 + \varepsilon)\delta(Q, C^*)$. We study a decision version of this approximation problem, which is called the $(1 + \varepsilon, r)$ -*approximate near-neighbor* problem for curves (ANNC). Here, if there exists a curve in \mathcal{C} that lies within distance r of the query curve Q , one has to return a curve in \mathcal{C} that lies within distance $(1 + \varepsilon)r$ of Q . Note that there exists a reduction from the $(1 + \varepsilon)$ -approximate nearest-neighbor problem to the $(1 + \varepsilon, r)$ -approximate near-neighbor problem [16, 18, 24], at the cost of an additional logarithmic factor in the query time and an $O(\log^2 n)$ factor in the storage space.

In practice, it is often the case that the query curves are significantly shorter than the input curves (e.g., Google-search queries). Thus, we also study the *asymmetric*

setting of $(1 + \varepsilon, r)$ -ANNC, where each of the input curves has complexity at most m , while each query curve has complexity at most $k \ll m$.

There are many methods that are used in real-world applications for comparing curves, and one of the most prevalent is the (discrete) *Fréchet* distance (DFD for short), which is often described by the following analogy. Two frogs are hopping from vertex to vertex along two polygonal curves. At each step, one of the frogs or both frogs may advance to the next vertex on its curve. The discrete Fréchet distance is defined as the smallest maximum distance between the frogs that can be achieved in such a joint sequence of hops. Another useful distance measure for curves or time series is the *dynamic time warping* distance (DTW for short), in which instead of taking the smallest maximum distance we take the smallest sum of distances.

In the last several years, a series of papers have been written investigating the approximate near-neighbor problem for curves (ANNC) and its variants under the Fréchet distance [1, 2, 8, 9, 13, 14, 19] (see Table 1), and several different approaches and sophisticated methods were utilized in order to provide efficient data structures. Up to now, all data structures for ANNC under DFD have either a query time exponential in m , or an infeasible storage space bound. In this paper, for the first time, we manage to remove the exponential factor from the query time, while also significantly reducing the space consumption. Our approach consists of a discretization of space based on the input curves, which allows us to prepare a small set of curves that captures all possible queries approximately.

Indyk [19] was the first to give a deterministic near-neighbor data structure for curves under DFD. The data structure achieves an approximation factor of $O((\log m + \log \log n)^{t-1})$ given some trade-off parameter $t > 1$. Its space consumption is very high, $O(m^2 |X|)^{tm^{1/t}} \cdot n^{2t}$, where $|X|$ is the size of the domain on which the curves are defined, and the query time is $(m \log n)^{O(t)}$. In Table 1 we set $t = 1 + o(1)$ to obtain a constant approximation factor.

Later, Driemel and Silvestri [13] presented a locality-sensitive-hashing scheme for curves under DFD, improving the result of Indyk for short curves. Their data structure uses $O(2^{4md} n \log n)$ space and answers queries in $O(2^{4md} \log n)$ time with an approximation factor of $O(d^{3/2})$. They also provide a trade-off between approximation quality and computational performance: for $d = O(1)$, and given a parameter $k \in [m]$, a data structure of size $O(2^{2k} m^{k-1} n \log n + mn)$ is constructed that answers queries in $O(2^{2k} m^k \log n)$ time with an approximation factor of $O(m/k)$. For constant k , this data structure uses only $\text{poly}(m) \cdot n \log n$ space and has $\text{poly}(m) \cdot \log n$ query time, however, the approximation factor is $O(m)$. They also show that this result can be applied to DTW, but only for the extreme of the trade-off which gives an $O(m)$ approximation.

Recently, Emiris and Psarros [14] presented near-neighbor data structures for curves under both DFD and DTW. Their algorithm provides an approximation factor of $(1 + \varepsilon)$, at the expense of increased space usage and preprocessing time. They use the idea that for a fixed alignment between two curves (i.e., a given sequence of hops of the two frogs), the problem can be reduced to the near-neighbor problem for points in ℓ_∞ -product of ℓ_2 spaces. Their basic idea is to construct a data structure for every possible alignment. Once a query is given, they query all these data structures and return the

closest curve found. This approach is responsible for the 2^m factor in their query time. Furthermore, they generalize this approach using randomized projections of ℓ_p -products of Euclidean metrics (for any $p \geq 1$), and define the $\ell_{p,2}$ -distance for curves (for $p \geq 1$), which is exactly DFD when $p = \infty$, and DTW when $p = 1$ (see Sect. 2). The space used by their data structure is $\tilde{O}(n) \cdot (2 + \frac{d}{\log m})^{O(m^{1+1/\varepsilon} \cdot d \log(1/\varepsilon))}$ with query $\tilde{O}(dm^{1+1/\varepsilon} \cdot 2^{4m} \log n)$ for DFD and $\tilde{O}(n) \cdot \frac{1}{\varepsilon}^{O(md)}$ space and $\tilde{O}(d \cdot 2^{4m} \log n)$ query for DTW.

Har-Peled and Kumar [17] considered approximate nearest-neighbor for general metric spaces where the query points are constrained to lie on a subspace of low doubling dimension. It is easy to show that the doubling dimension of the metric space for curves of length m in d dimensions under DFD, is bounded by $O(md)$. Therefore, their result implies that one can construct a $(1 + \varepsilon)$ -approximate nearest-neighbor data-structure requiring space in $n\varepsilon^{-O(md)}$ and with query time in $2^{O(md)} \log(n/\varepsilon)$. Note that in [13, 14] the query time is exponential in md and m , respectively, while the space complexity is suboptimal (and in [13], the approximation factor is $O(d^{3/2})$ and not $(1 + \varepsilon)$). In this paper, we get the exact same space as in [17] (up to a $\log n$ factor), which seems to be optimal, while the query time is linear in md (using the more structured input).

Subsequent work In a recent work, Bringmann et al. [4] study the asymmetric setting of $(1 + \varepsilon, r)$ -ANNC under continuous Fréchet distance, for one-dimensional polygonal curves (time series). Improving upon the work of Driemel and Psarros [11], they use the idea of signatures to obtain some sort of a discrete version of the problem, and then follow our approach of preparing in advance the answers to all relevant queries on a discretization of the space (which was also used in [11]), to construct a data structure with space in $n \cdot O(\frac{m}{k\varepsilon})^k$ and query time in $O(1)^k$. They also show that an approximation factor of $(2 + \varepsilon)$ can be obtained with the same space bound and $O(k)$ query time. This provides additional evidence that our approach to ANNC, although quite simple and easy to implement, seems to produce more efficient data structures than those obtained using tools such as LSH and randomized projections. In addition, [4] present conditional lower bounds for several different settings of the problem.

Our results We present a data structure for the $(1 + \varepsilon, r)$ -approximate near-neighbor problem using a bucketing method. We construct a relatively small set of curves \mathcal{I} , such that, given a query curve Q , if there exists some curve in \mathcal{C} within distance r of Q , then one of the curves in \mathcal{I} must be very close to Q . The points of the curves in \mathcal{I} are chosen from a simple discretization of space, thus, while it is not surprising that we get the best query time, it is surprising that we achieve a better space bound. Moreover, while the analysis of the space bounds is rather involved, the implementation of our data structures remain simple in practice.

See Table 1 for a summary of our results. In the table, we do not state our result for the general $\ell_{p,2}$ -distance. Instead, we state our results for the two most important cases, i.e. DFD and DTW, and compare them with previous work. Note that our results substantially improve the current state of the art for any $p \geq 1$. In particular, we remove the exponential dependence on m in the query bounds and significantly improve the space bounds.

Our results for the asymmetric setting, where the query curve Q has complexity $k \ll m$, are summarized in Table 2. We show that in the asymmetric setting for DFD, our data structure can be slightly modified in order to achieve query time and storage space independent of m . Moreover, the storage space and query time matches those of the symmetric setting, by replacing m with k .

We also apply our methods to an approximation version of range counting for curves (for the general $\ell_{p,2}$ distance) and achieve bounds similar to those of our ANNC data structure. Moreover, at the cost of an additional $O(n)$ -factor in the space bound, we can also answer the corresponding approximation version of range searching, thus answering a question of Afshani and Driemel [1], with respect to DFD.

We note that our approach with obvious modifications works also in a dynamic setting, that is, we can construct an efficient dynamic data structure for ANNC as well as for other related problems such as range counting and range reporting for curves.

Another significant advantage of our approach is that, unlike some of the previous solutions, our data structure always returns an answer, and never returns a curve at distance greater than $(1 + \varepsilon)r$ from the query curve, i.e., there are no false positives. This is an important property of our solution, due to the fact that verifying the validity of the answer (i.e., computing the distance between two curves) cannot be done in strongly subquadratic time (assuming SETH, see [7]), which is already more than our query time (for $d < m$).

Summary of previous and current results for the asymmetric approximate near-neighbor data structure for curves. All the results in the table are w.r.t. DFD. Setting: the input consists of n curves with m points each, the query curve is of length k , the approximation ratio is $1 + \varepsilon$ for $\varepsilon \in (0, 1)$, and our data structures always succeed. Historic note: [12] is a subsequent work to the first version of this paper. In this version we also apply our counting techniques to the asymmetric cases

1.1 Related Work

De Berg, Gudmundsson, and Mehrabi [9] described a dynamic data structure for approximate nearest neighbor for curves (which can also be used for other types of queries such as range reporting), under the (continuous) Fréchet distance. Their data structure uses $n \cdot O\left(\frac{1}{\varepsilon}\right)^{2m}$ space and has $O(m)$ query time, but with an additive error of $\varepsilon \cdot reach(Q)$, where $reach(Q)$ is the maximum distance between the start vertex of the query curve Q and any other vertex of Q . Furthermore, their query procedure might fail when the distance to the nearest neighbor is relatively large.

Afshani and Driemel [1] studied (exact) range searching under both the discrete and continuous Fréchet distance. In this problem, the goal is to preprocess \mathcal{C} such that given a query curve Q of length m_q and a radius r , all the curves in \mathcal{C} that are within distance r from Q can be found efficiently. For DFD, their data structure uses $O(n(\log \log n)^{m-1})$ space and has $O(n^{1-\frac{1}{d}} \cdot \log^{O(m)} n \cdot m_q^{O(d)})$ query time, where m_q is limited to $\log^{O(1)} n$. Additionally, they provide a lower bound in the pointer model, stating that every data structure with $Q(n) + O(k)$ query time, where k is the output size, has to use roughly $\Omega\left(\frac{n}{Q(n)}\right)^2$ space in the worst case (even for

Table 1 Our approximate near-neighbor data structure under DFD and DTW compared to the previous results

	Space	Query	Approx.	Comments
DFD	$O(m^2 X n^{1-\alpha(1)} \cdot n^{2+\alpha(1)})$	$(m \log n)^{O(1)}$	$O(1)$	deterministic [19]
	$O(2^{4md} n \log n)$	$O(2^{4md} \log n)$	$O(1)$	randomized, using LSH, $d = O(1)$ [13]
	$O(n \log n + mn)$	$O(m \log n)$	$O(m)$	
	$\tilde{O}(n \cdot (2 + \frac{d}{\log m})^{O(\frac{1}{\varepsilon})} d \log(\frac{1}{\varepsilon}))$	$\tilde{O}(dm^{O(\frac{1}{\varepsilon})} \cdot 2^{4m} \log n)$	$1 + \varepsilon$	randomized [14]
DTW	$n \cdot O(\frac{1}{\varepsilon})^{md}$	$O(md)$	$1 + \varepsilon$	deterministic (randomized construction), Theorem 8
	$O(n \log n + mn)$	$O(m \log n)$	$O(m)$	randomized, using LSH, $d = O(1)$ [13]
	$\tilde{O}(n) \cdot \frac{1}{\varepsilon}^{O(md)}$	$\tilde{O}(d \cdot 2^{4m} \log n)$	$1 + \varepsilon$	randomized [14]
	$n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$	$O(md)$	$1 + \varepsilon$	deterministic (randomized construction), Theorem 14

Table 2 Summary of previous and current results for the asymmetric approximate near-neighbor data structure for curves. All the results in the table are w.r.t. DFD. Setting: the input consists of n curves with m points each, the query curve is of length k , the approximation ratio is $1 + \varepsilon$ for $\varepsilon \in (0, 1)$, and our data structures always succeed

Space	Query	Deterministic construction?	References
$n \cdot \left(O\left(\frac{kd^{3/2}}{\varepsilon}\right)kd \right)$	$O(kd)$	No	[12]
$n \cdot \left(O\left(\frac{kd^{3/2}}{\varepsilon}\right)kd+1 \right)$	$O\left(\frac{k^2d^{5/2}}{\varepsilon}(\log n + kd \log\left(\frac{kd}{\varepsilon}\right))\right)$	Yes	[12]
$n \cdot O\left(\frac{1}{\varepsilon}\right)kd$	$O(kd)$	No	Theorem 10
$n \cdot O\left(\frac{1}{\varepsilon}\right)kd$	$O(kd \log\left(\frac{nkd}{\varepsilon}\right))$	Yes	Theorem 24

Historic note: [12] is a subsequent work to the first version of this paper. In this version we also apply our counting techniques to the asymmetric cases.

$m_q = 1$). Afshani and Driemel conclude their paper by asking whether more efficient data structures might be constructed if one allows approximation.

De Berg, Cook IV, and Gudmundsson [8], considered the following range counting problem under the continuous Fréchet distance. Given a polygonal curve C with m vertices, they show how to preprocess it into a data structure of size $O(k \cdot \text{polylog}(m))$, so that, given a query segment s , one can return a constant approximation of the number of subcurves of C that lie within distance r of s in $O(\frac{m}{\sqrt{k}} \cdot \text{polylog}(m))$ time, where k is a parameter between m and m^2 .

Aronov et al. [2] managed to obtain practical bounds for two cases of the asymmetric $(1 + \varepsilon, r)$ -ANNC under DFD: (i) when Q is a line segment (i.e., $k = 2$), or (ii) when C consists of line segments (i.e., $m = 2$). The bounds on the size of the data structure and query time are nearly linear in the size of the input and query curve, respectively. Specifically, for the case where $k = 2$, they achieve query time $O(\log^4(\frac{n}{\varepsilon}))$ and storage space $O(n \frac{1}{\varepsilon^4} \log^4(\frac{n}{\varepsilon}))$. They also provide efficient data structures for several other variants of the problem: the (exact) NNC where ℓ_∞ is used for interpoint distances, and the case where the location of the input curves is only fixed up to translation.

1.2 Technical Ideas

We use a discretization of the space, by considering a d -dimensional uniform grid with edge length $\frac{\varepsilon r}{\sqrt{d}}$. The main ingredient in our data structure is then a relatively small set \mathcal{I} of curves defined by grid points, which represents all possible queries. For each curve in \mathcal{I} we store an index of a close enough curve from the input set \mathcal{C} . Given a query Q sufficiently close to some curve in \mathcal{C} , we find a representative Q' in \mathcal{I} by simply rounding Q 's vertices and return the index of the curve stored for Q' .

Given a point $x \in \mathbb{R}^d$, the number of grid points that are within distance $(1 + \varepsilon)r$ from x is bounded by $O(\frac{1}{\varepsilon})^d$ (Corollary 6). Thus, given a curve C of length m , the total number of grid points that are within distance $(1 + \varepsilon)r$ from one of its vertices is $m \cdot O(\frac{1}{\varepsilon})^d$. Naively, the number of curves needed to represent all possible queries of length m within distance r of C is bounded by the number of ways to choose m points with repetitions from a set of grid points of size $m \cdot O(\frac{1}{\varepsilon})^d$, which is bounded by $m^m \cdot O(\frac{1}{\varepsilon})^{md}$. This infeasible bound on the storage space might be the reason why more sophisticated solutions for ANNC have been suggested throughout the years.

One of the main technical contributions of this paper is an analysis leading to a significantly better bound, if we store only candidate curves that are within distance $(1 + \varepsilon)r$ from C . Actually, in Sect. 3 we show that for the case of DFD, it is sufficient to store a set of representative curves of size only $O(\frac{1}{\varepsilon})^{md}$ for each input curve. The basic idea is to bound the number of representatives that can be obtained by some fixed alignment between C and the candidate curve (see Claim 7).

For the general case of $\ell_{p,2}$ -distance (including DTW), we are minimizing the sum of distances instead of the maximum distance (as in DFD). Thus, we have to use a denser grid (with edge length $\frac{\varepsilon r}{(2m)^{1/p} \sqrt{d}}$), and the situation becomes more complicated. First, unlike DFD, the triangle inequality does not hold for $\ell_{p,2}$ -distance in general (including DTW). Second, since DFD is a min-max measure, the choice of different

vertices for a representative curve is “independent” in a sense, whereas for $\ell_{p,2}$ -distance in general, the choice of different vertices depends on their sum of distances from the input curve. Using more careful counting arguments and analysis of the alignment between two curves, we are able to show that in this case the number of representative curves that our data structure has to store per input curve is bounded by $O(\frac{1}{\epsilon})^{m(d+1)}$ (see Claim 12).

To store the set \mathcal{I} we simply use a dictionary, which can be implemented using a hash table and guarantees a query time linear in the size of the query. To obtain a fully deterministic solution, one can use a search tree instead. However, a naive implementation using a binary search tree results in an additional factor of $O(\log |\mathcal{I}|) = O(md \log(\frac{n}{\epsilon}))$ to the query time, i.e., in a query time of $O(m^2 d^2 \log(\frac{n}{\epsilon}))$. We show how to implement the dictionary using a prefix tree, exploiting the fact that the vertices of the curves in \mathcal{I} are from a relatively small set of grid points, which improves the query time to $O(md \log(\frac{nm}{\epsilon}))$.

For the asymmetric setting (where the length of a query is $k \ll m$), we use simplifications of the input curves in order to obtain bounds that are independent of m . Given a curve C of length m , a simplification Π of C is a curve of length $k \ll m$ that is relatively close to C . Simplifications were used in order to provide approximate solutions in several asymmetric versions of problems on curves, such as clustering [3], and distance oracles [10, 12].

By the triangle inequality for DFD, every query curve Q within distance r from an input curve C is at distance at most $2r$ from the simplification Π (where Π is within distance r from C). Thus, it is enough to prepare for query curves at distance at most $2r$ from Π , which follows from previous arguments. Note that the query time and storage space are independent of m . For the asymmetric setting under general $\ell_{p,2}$ -distance, the situation again becomes more complicated. First, we present an algorithm that computes the closest (vertex-restricted) simplification of length k to a given curve of length m under the $\ell_{p,2}$ -distance. In order to adapt our data structure to the asymmetric setting, we need to increase the allowed distance between a simplification and a representative curve by a factor of $k^{1/p}$, for the triangle inequality to work. The counting arguments that we use for the symmetric case yield a bound of $O(\frac{k^{1/p}}{\epsilon})^{k(d+1)}$ on the storage space. In Sect. 6.2, we provide stronger counting arguments, which enable us to remove the $k^{1/p}$ factor from the base of the exponent. The main idea is to use the simplification in order to divide the input curve into $O(k)$ compact subsequences (see Claim 18).

2 Preliminaries

To simplify the presentation, we assume throughout the paper that all the input curves have exactly the same size, m , and all the query curves have exactly the same size, either m or k , depending on whether we are considering the standard or the asymmetric version. This assumption can be easily removed (see Appendix B).

Let \mathcal{C} be a set of n curves, each consisting of m points in \mathbb{R}^d , and let δ be some distance measure for curves.

Problem 1 ($(1 + \varepsilon)$ -approximate nearest-neighbor for curves) Given a parameter $0 < \varepsilon \leq 1$, preprocess \mathcal{C} into a data structure that given a query curve Q , returns a curve $C' \in \mathcal{C}$, such that $\delta(Q, C') \leq (1 + \varepsilon) \cdot \delta(Q, C)$, where C is the curve in \mathcal{C} closest to Q .

Problem 2 ($(1 + \varepsilon, r)$ -approximate near-neighbor for curves) Given a parameter r and $0 < \varepsilon \leq 1$, preprocess \mathcal{C} into a data structure that given a query curve Q , if there exists a curve $C_i \in \mathcal{C}$ such that $\delta(Q, C_i) \leq r$, returns a curve $C_j \in \mathcal{C}$ such that $\delta(Q, C_j) \leq (1 + \varepsilon)r$.

Problem 3 (Asymmetric $(1 + \varepsilon, r)$ -approximate near-neighbor for curves) Given parameters r, k , and $0 < \varepsilon \leq 1$, preprocess \mathcal{C} into a data structure that given a query curve Q of length k , if there exists a curve $C_i \in \mathcal{C}$ such that $\delta(Q, C_i) \leq r$, returns a curve $C_j \in \mathcal{C}$ such that $\delta(Q, C_j) \leq (1 + \varepsilon)r$.

Curve alignment Given two integers m_1, m_2 , let $\tau := \langle (i_1, j_1), \dots, (i_t, j_t) \rangle$ be a sequence of pairs where $i_1 = j_1 = 1, i_t = m_1, j_t = m_2$, and for each $1 \leq k < t$ we have $(i_{k+1}, j_{k+1}) \in \{(i_k + 1, j_k), (i_k, j_k + 1), (i_k + 1, j_k + 1)\}$. We call such a sequence τ an *alignment* w.r.t. two curves of lengths m_1 and m_2 , respectively.

Let $P = (p_1, \dots, p_{m_1})$ and $Q = (q_1, \dots, q_{m_2})$ be two curves of lengths m_1 and m_2 , respectively, in \mathbb{R}^d . We say that an alignment τ w.r.t. P and Q matches p_i and q_j if $(i, j) \in \tau$.

Discrete Fréchet distance (DFD) The *Fréchet cost* of an alignment τ w.r.t. P and Q is $\sigma_{dF}(\tau(P, Q)) := \max_{(i,j) \in \tau} \|p_i - q_j\|_2$. The discrete Fréchet distance is defined over the set \mathcal{T} of all alignments as

$$d_{dF}(P, Q) = \min_{\tau \in \mathcal{T}} \sigma_{dF}(\tau(P, Q)).$$

Dynamic time warping (DTW) The *time warping cost* of an alignment τ w.r.t. P and Q is $\sigma_{DTW}(\tau(P, Q)) := \sum_{(i,j) \in \tau} \|p_i - q_j\|_2$. The DTW distance is defined over the set \mathcal{T} of all alignments as

$$d_{DTW}(P, Q) = \min_{\tau \in \mathcal{T}} \sigma_{DTW}(\tau(P, Q)).$$

$\ell_{p,2}$ -distance for curves The $\ell_{p,2}$ -cost of an alignment τ w.r.t. P and Q is $\sigma_{p,2}(\tau(P, Q)) := \left(\sum_{(i,j) \in \tau} \|p_i - q_j\|_2^p \right)^{1/p}$. The $\ell_{p,2}$ -distance between P and Q is defined over the set \mathcal{T} of all alignments as

$$d_{p,2}(P, Q) = \min_{\tau \in \mathcal{T}} \sigma_{p,2}(\tau(P, Q)).$$

Notice that $\ell_{p,2}$ -distance is a generalization of DFD and DTW, in the sense that $\sigma_{dF} = \sigma_{\infty,2}$ and $d_{dF} = d_{\infty,2}$, $\sigma_{DTW} = \sigma_{1,2}$ and $d_{DTW} = d_{1,2}$. Also note that DFD satisfies the triangle inequality, but DTW and $\ell_{p,2}$ -distance (for $p \neq \infty$) do not (see Sect. 5 for details).

Emiris and Psarros [14] showed that the number of all possible alignments of two curves is in $O(m \cdot 2^{2m})$. We reduce this bound by counting only alignments that can determine the $\ell_{p,2}$ -distance between two curves.¹ More formally, let τ be an alignment. If there exists an alignment τ' such that $\tau' \subset \tau$, then clearly $\sigma_{p,2}(\tau'(P, Q)) \leq \sigma_{p,2}(\tau(P, Q))$, for any $1 \leq p \leq \infty$ and for any two curves P and Q . In this case, we say that τ cannot determine the $\ell_{p,2}$ -distance between two curves.

Lemma 4 *The number of different alignments that can determine the $\ell_{p,2}$ -distance between two m -curves (for any $1 \leq p \leq \infty$) is at most $O(\frac{2^{2m}}{\sqrt{m}})$.*

Proof Let $\tau = \langle (i_1, j_1), \dots, (i_t, j_t) \rangle$ be an alignment. Notice that $m \leq t \leq 2m - 1$. By definition, τ has 3 types of (consecutive) subsequences of length two:

- (i) $\langle (i_k, j_k), (i_k + 1, j_k) \rangle$,
- (ii) $\langle (i_k, j_k), (i_k, j_k + 1) \rangle$, and
- (iii) $\langle (i_k, j_k), (i_k + 1, j_k + 1) \rangle$.

Denote by \mathcal{T}_1 the set of all alignments that do not contain any subsequence of type (iii). Then, any $\tau_1 \in \mathcal{T}_1$ is of length exactly $2m - 1$. Moreover, τ_1 contains exactly $2m - 2$ subsequences of length two, of which $m - 1$ are of type (i) and $m - 1$ are of type (ii). Therefore, $|\mathcal{T}_1| = \binom{2m-2}{m-1} = O(\frac{2^{2m}}{\sqrt{m}})$.

Assume that an alignment τ contains a subsequence of the form $(i_k, j_k - 1), (i_k, j_k), (i_k + 1, j_k)$, for some $1 < k \leq t - 1$. Notice that removing the pair (i_k, j_k) from τ results in a legal alignment τ' , such that $\sigma_{p,2}(\tau'(P, Q)) \leq \sigma_{p,2}(\tau(P, Q))$, for any $1 \leq p \leq \infty$ and two curves P, Q . We call the pair (i_k, j_k) a *redundant pair*. Similarly, if τ contains a subsequence of the form $(i_k - 1, j_k), (i_k, j_k), (i_k, j_k + 1)$, for some $1 < k \leq t - 1$, then the pair (i_k, j_k) is also a redundant pair. Therefore we only care about alignments that do not contain any redundant pairs. Denote by \mathcal{T}_2 the set of all alignments that do not contain a redundant pair, then any $\tau_2 \in \mathcal{T}_2$ contains at least one subsequence of type (iii).

We claim that for any alignment $\tau_2 \in \mathcal{T}_2$, there exists a unique alignment $\tau_1 \in \mathcal{T}_1$. Indeed, if we add the redundant pair $(i_l, j_l + 1)$ between (i_l, j_l) and $(i_l + 1, j_l + 1)$ for each subsequence of type (iii) in τ_2 , we obtain an alignment $\tau_1 \in \mathcal{T}_1$. Moreover, since τ_2 does not contain any redundant pairs, the reverse operation on τ_1 results in τ_2 . Thus we obtain $|\mathcal{T}_2| \leq |\mathcal{T}_1| = O(\frac{2^{2m}}{\sqrt{m}})$. □

Points and balls Given a point $x \in \mathbb{R}^d$ and a real number $R > 0$, we denote by $B_p^d(x, R)$ the d -dimensional ball under the ℓ_p norm with center x and radius R , i.e., a point $y \in \mathbb{R}^d$ is in $B_p^d(x, R)$ if and only if $\|x - y\|_p \leq R$, where $\|x - y\|_p = (\sum_{i=1}^d |x_i - y_i|^p)^{1/p}$. Let $B_p^d(R) = B_p^d(\mathbf{0}, R)$, and let $V_p^d(R)$ be the volume (w.r.t. Lebesgue measure) of $B_p^d(R)$, then

$$V_p^d(R) = \frac{2^d \Gamma(1 + 1/p)^d}{\Gamma(1 + d/p)} R^d,$$

¹ Since our storage space is already in $O(\frac{1}{\epsilon})^{md}$, and $m \cdot 2^{2m} \leq 3^{2m}$ is in $O(1)^{md}$, we could have used this larger upper bound. However, in Lemma 4 we show a tight upper bound on the number of relevant alignments, which may be useful for other applications.

where $\Gamma(\cdot)$ is Euler’s Gamma function (an extension of the factorial function). For $p = 2$ and $p = 1$, we get

$$V_2^d(R) = \frac{\pi^{d/2}}{\Gamma(1 + d/2)} R^d \quad \text{and} \quad V_1^d(R) = \frac{2^d}{d!} R^d.$$

Our approach consists of a discretization of the space using lattice points, i.e., points from \mathbb{Z}^d .

Lemma 5 *The number of lattice points in the d -dimensional ball of radius R under the ℓ_p norm (i.e., in $B_p^d(R)$) is bounded by $V_p^d(R + d^{1/p})$.*

Proof With each lattice point $z = (z_1, z_2, \dots, z_d)$, $z_i \in \mathbb{Z}$, we match the d -dimensional lattice cube $C(z) = [z_1, z_1 + 1] \times [z_2, z_2 + 1] \times \dots \times [z_d, z_d + 1]$. Notice that $z \in C(z)$, and the ℓ_p -diameter of a lattice cube is $d^{1/p}$. Therefore, the number of lattice points in the ℓ_p^d -ball of radius R is bounded by the number of lattice cubes that are contained in a ℓ_p^d -ball with radius $R + d^{1/p}$. This number is bounded by $V_p^d(R + d^{1/p})$ divided by the volume of a lattice cube, which is $1^d = 1$. \square

3 Discrete Fréchet Distance (DFD)

Consider the infinite d -dimensional grid with edge length $\frac{\epsilon r}{\sqrt{d}}$. Given a point x in \mathbb{R}^d , by rounding one can find in $O(d)$ time the grid point x' closest to x , and $\|x - x'\|_2 \leq \frac{\epsilon r}{2}$. Let $G(x, R)$ denote the set of grid points that are contained in $B_2^d(x, R)$.

Corollary 6 $|G(x, (1 + \epsilon)r)| = O\left(\frac{1}{\epsilon}\right)^d$.

Proof We scale our grid so that the edge length is 1, hence we are looking for the number of lattice points in $B_2^d(x, \frac{1+\epsilon}{\epsilon}\sqrt{d})$. By Lemma 5 we get that this number is bounded by the volume of the d -dimensional ball of radius $\frac{1+\epsilon}{\epsilon}\sqrt{d} + \sqrt{d} \leq \frac{3\sqrt{d}}{\epsilon}$. Using Stirling’s formula we conclude that

$$V_2^d\left(\frac{3\sqrt{d}}{\epsilon}\right) = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)} \cdot \left(\frac{3\sqrt{d}}{\epsilon}\right)^d \leq \left(\frac{\alpha}{\epsilon}\right)^d,$$

where α is a constant. For example, if d is even, then

$$\begin{aligned} V_2^d\left(\frac{3\sqrt{d}}{\epsilon}\right) &= \frac{\pi^{\frac{d}{2}}}{\left(\frac{d}{2}\right)!} \cdot \left(\frac{3\sqrt{d}}{\epsilon}\right)^d \leq \frac{\pi^{\frac{d}{2}}}{\sqrt{2\pi}(d/2)^{d/2+1/2}e^{-d/2}} \cdot \left(\frac{3\sqrt{d}}{\epsilon}\right)^d \\ &\leq \left(\frac{12.4}{\epsilon}\right)^d = O\left(\frac{1}{\epsilon}\right)^d. \end{aligned}$$

\square

Denote by p_j^i the j 'th point of C_i , and let $G_i = \bigcup_{1 \leq j \leq m} G(p_j^i, (1 + \varepsilon)r)$ and $\mathcal{G} = \bigcup_{1 \leq i \leq n} G_i$, then by the above corollary we have $|G_i| = m \cdot O(\frac{1}{\varepsilon})^d$ and $|\mathcal{G}| = mn \cdot O(\frac{1}{\varepsilon})^d$. Let \mathcal{I}_i be the set of all curves $\overline{Q} = (x_1, x_2, \dots, x_m)$ with points from G_i , such that $d_{dF}(C_i, \overline{Q}) \leq (1 + \frac{\varepsilon}{2})r$.

Claim 7 $|\mathcal{I}_i| = O(\frac{1}{\varepsilon})^{md}$ and it can be computed in $O(\frac{1}{\varepsilon})^{md}$ time.

Proof Let $\overline{Q} \in \mathcal{I}_i$ and let τ be an alignment with $\sigma_{dF}(\tau(C_i, \overline{Q})) \leq (1 + \frac{\varepsilon}{2})r$. For each $1 \leq k \leq m$ let j_k be the smallest index such that $(j_k, k) \in \tau$. In other words, j_k is the smallest index that is matched to k by the alignment τ . Since $d_{dF}(C_i, \overline{Q}) \leq (1 + \frac{\varepsilon}{2})r$, we have $x_k \in B_2^d(p_{j_k}^i, (1 + \frac{\varepsilon}{2})r)$, for $k = 1, \dots, m$. This means that for any curve $\overline{Q} \in \mathcal{I}_i$ such that $\sigma_{dF}(\tau(C_i, \overline{Q})) \leq (1 + \frac{\varepsilon}{2})r$, we have $x_k \in G(p_{j_k}^i, (1 + \frac{\varepsilon}{2})r)$, for $k = 1, \dots, m$. By Corollary 6, the number of ways to choose a grid point x_k from $G(p_{j_k}^i, (1 + \frac{\varepsilon}{2})r)$ is bounded by $O(\frac{1}{\varepsilon})^d$.

We conclude that given an alignment τ , the number of curves \overline{Q} with m points from G_i such that $\sigma_{dF}(\tau(C_i, \overline{Q})) \leq (1 + \frac{\varepsilon}{2})r$ is bounded by $O(\frac{1}{\varepsilon})^{md}$. Finally, by Lemma 4, the total number of curves in \mathcal{I}_i is bounded by $2^{2m} \cdot O(\frac{1}{\varepsilon})^{md} = O(\frac{1}{\varepsilon})^{md}$.

To construct \mathcal{I}_i we compute, for each of the $O(\frac{1}{\varepsilon})^{md}$ candidates, its discrete Fréchet distance to C_i . Thus, we construct \mathcal{I}_i in total time $O(\frac{1}{\varepsilon})^{md} \cdot O(m^2) = O(\frac{1}{\varepsilon})^{md}$. (The latter equality is true, since clearly $(\frac{\alpha}{\varepsilon})^{md} \cdot O(m^2) \leq (\frac{c\alpha}{\varepsilon})^{md}$, i.e., $O(m^2) \leq c^{md}$, where α is the constant from Corollary 6 and $c > 1$ is a sufficiently large constant.) \square

The data structure Denote $\mathcal{I} = \bigcup_{1 \leq i \leq n} \mathcal{I}_i$, so $|\mathcal{I}| \leq n \cdot O(\frac{1}{\varepsilon})^{md}$ and we construct \mathcal{I} in total time $n \cdot O(\frac{1}{\varepsilon})^{md}$. Next, we would like to store the set \mathcal{I} in a dictionary (a hash table or a lookup table) \mathcal{D} , such that given a query curve Q , one can find Q in \mathcal{D} (if it exists) in $O(md)$ time. We use Cuckoo Hashing [23] to construct a (dynamic) dictionary of linear space, constant worst-case query and deletion time, and constant expected amortized insertion time. We insert the curves of \mathcal{I} into the dictionary \mathcal{D} as follows. For each $1 \leq i \leq n$ and curve $\overline{Q} \in \mathcal{I}_i$, if $\overline{Q} \notin \mathcal{D}$, insert \overline{Q} into \mathcal{D} , and set $C(\overline{Q}) \leftarrow C_i$ ($C(\overline{Q})$ holds the “answer” for \overline{Q}). The storage space required for \mathcal{D} is $O(|\mathcal{I}|)$, and to construct it we perform $|\mathcal{I}|$ insertions and look-up operations which take in total $O(|\mathcal{I}| \cdot md) = O(|\mathcal{I}|)$ expected time.

A deterministic construction using a prefix tree Another way to implement the dictionary, which is also dynamic, simple, and does not require randomization at all, is using a binary search tree. Assuming that comparing two curves (given their binary representations) requires $O(md)$ time, the query time will be $O(md \log |\mathcal{I}|) = O((md)^2 \log \frac{n}{\varepsilon})$.

Since there is a relatively small number of possible vertices (all the vertices are points of the grid \mathcal{G}) we can improve the query time to $O(md \log(\frac{nm}{\varepsilon}))$ by using a prefix tree instead of a search tree. For details, see Appendix A.

The query algorithm Let $Q = (q_1, \dots, q_m)$ be the query curve. The query algorithm is as follows: For each $1 \leq k \leq m$ find the grid point q'_k (not necessarily from \mathcal{G}) closest to q_k . This can be done in $O(md)$ time by rounding. Then, search for the curve $Q' = (q'_1, \dots, q'_m)$ in the dictionary \mathcal{D} . If Q' is in \mathcal{D} , return $C(Q')$, otherwise, return NO. The total query time is then $O(md)$.

Table 3 Comparing our ANN data structure to previous structures, for a fixed ε (say $\varepsilon = 1/2$)

m	References	Space	Query	Approx
$\log n$	[13]	$O(n^{4d+1} \log n)$	$\tilde{O}(n^{4d})$	$d\sqrt{d}$
	[14]	$n^{\Omega(d \log n)}$	$\tilde{O}(dn^4)$	$1 + \varepsilon$
	Theorem 8	$n^{O(d)}$	$O(d \log n)$	$1 + \varepsilon$
$O(1)$	[13]	$2^{O(d)} n \log n$	$2^{O(d)} \cdot \log n$	$d\sqrt{d}$
	[14]	$d^{O(d)} \tilde{O}(n)$	$O(d \log n)$	$1 + \varepsilon$
	Theorem 8	$2^{O(d)} n$	$O(d)$	$1 + \varepsilon$

Correctness Consider a query curve $Q = (q_1, \dots, q_m)$. Assume that there exists a curve $C_i \in \mathcal{C}$ such that $d_{dF}(C_i, Q) \leq r$. We show that the query algorithm returns a curve C^* with $d_{dF}(C^*, Q) \leq (1 + \varepsilon)r$.

Consider a point $q_k \in Q$. Denote by $q'_k \in \mathcal{G}$ the grid point closest to q_k , and let $Q' = (q'_1, \dots, q'_m)$. We have $\|q_k - q'_k\|_2 \leq \frac{\varepsilon r}{2}$, so $d_{dF}(Q, Q') \leq \frac{\varepsilon r}{2}$. By the triangle inequality,

$$d_{dF}(C_i, Q') \leq d_{dF}(C_i, Q) + d_{dF}(Q, Q') \leq r + \frac{\varepsilon r}{2} = (1 + \frac{\varepsilon}{2})r,$$

so Q' is in $\mathcal{I}_i \subseteq \mathcal{I}$. This means that \mathcal{D} contains Q' with a curve $C(Q') \in \mathcal{C}$ such that $d_{dF}(C(Q'), Q') \leq (1 + \frac{\varepsilon}{2})r$, and the query algorithm returns $C(Q')$. Now, again by the triangle inequality,

$$d_{dF}(C(Q'), Q) \leq d_{dF}(C(Q'), Q') + d_{dF}(Q', Q) \leq (1 + \frac{\varepsilon}{2})r + \frac{\varepsilon r}{2} = (1 + \varepsilon)r.$$

We obtain the following theorem (see Table 3).

Theorem 8 *There exists a data structure for the $(1 + \varepsilon, r)$ -ANNC under DFD, with $n \cdot O(\frac{1}{\varepsilon})^{md}$ space, $n \cdot O(\frac{1}{\varepsilon})^{md}$ expected preprocessing time, and $O(md)$ query time.*

4 The Asymmetric Setting Under DFD

In this section, we show how to easily adapt our data structure to the asymmetric setting, by using simplifications of length at most k instead of the original input curves.

Bereg et al. [6] showed that given a curve C consisting of m points in 3D, and a parameter $r > 0$, there is an algorithm that runs in $O(m \log m)$ time and returns a simplification Π with minimum number of vertices such that $d_{dF}(C, \Pi) \leq r$. Their algorithm generalizes to higher dimensions, using an approximation algorithm for the minimum enclosing ball problem (see Kumar et al. [20]). In this section, we use the

following generalization of their original approach ([6], Theorem 1). More details are given in Sect. 8.

Lemma 9 *Let C be a curve consisting of m points in \mathbb{R}^d . Given parameters $k \leq m$, $r > 0$, and $\varepsilon \in (0, 1]$, there is an algorithm that runs in $O\left(\frac{d \cdot m \log m}{\varepsilon} + m \cdot \text{poly}\frac{1}{\varepsilon}\right)$ time that either returns a simplification Π consisting of k points such that $d_{dF}(C, \Pi) \leq (1 + \varepsilon)r$, or declares that for every simplification Π with k points, it holds that $d_{dF}(C, \Pi) > r$.*

For each $C_i \in \mathcal{C}$, using Lemma 9 with parameter $\varepsilon = 1$, we find a curve Π_i of length k such that $d_{dF}(C_i, \Pi_i) \leq 2r$. If we fail to find such a curve, then we can ignore C_i , because it means that $d_{dF}(Q, C_i) > r$ for any curve Q of length k .

To reduce the space consumption of our data structure, we only store candidate curves of length k that are close enough to the simplifications Π_i . However, since the distance between the simplification Π_i and the input curve C_i could be up to $2r$, storing the answers for the set of candidate curves that are within distance $(1 + \frac{\varepsilon}{2})r$ from Π_i is not enough, because a query Q that is within distance $(1 + \varepsilon)r$ from C_i might be as far as $(3 + \varepsilon)r$ from Π_i . Thus, instead, we insert into our data structure all the curves that are within distance $4r$ from Π_i . This allows us to capture all query curves that are within distance r from C_i .

The data structure We construct our data structure for the original (symmetric) version, with the following modifications. The set of input curves is $\mathcal{P} = \{\Pi_1, \dots, \Pi_n\}$ (instead of \mathcal{C}), and the radius parameter is $4r$ (instead of r), but the grid edge length remains $\frac{\varepsilon r}{\sqrt{d}}$. In addition, we let \mathcal{I}'_i be the set of all curves \overline{Q} with k points from G_i , such that $d_{dF}(\overline{Q}, \Pi_i) \leq 4r$, and \mathcal{I}_i will be the set of all curves $\overline{Q} \in \mathcal{I}'_i$ such that $d_{dF}(\overline{Q}, C_i) \leq (1 + \frac{\varepsilon}{2})r$. We insert the curves in \mathcal{I}_i into the database \mathcal{D} as before: For each $\overline{Q} \in \mathcal{I}_i$, if $\overline{Q} \notin \mathcal{D}$, insert \overline{Q} into \mathcal{D} and set $C(\overline{Q}) \leftarrow C_i$.

Notice that using $4r$ instead of r , increases the ratio between the radius and the grid edge length by only a factor of 4, and therefore the bound on $|\mathcal{I}'_i|$ does not change, except that m is replaced by k . Therefore, the bounds on the storage space and query time are similar to those of the original data structure, where m is replaced by k . Thus, the storage space is in $n \cdot O(\frac{1}{\varepsilon})^{kd}$ and the query time is in $O(kd)$. As for the preprocessing time, we get an additional term of $O(nmd \log m)$ for computing the simplifications Π_1, \dots, Π_n . We also need to compute the distances $d_{dF}(C_i, \overline{Q})$ in the construction of \mathcal{I}_i , for $1 \leq i \leq n$, which takes $n \cdot O(\frac{1}{\varepsilon})^{kd} \cdot O(mkd) = nm \cdot O(\frac{1}{\varepsilon})^{kd}$ time in total (as $kd \leq 2^{kd}$). Thus the total expected preprocessing time is $O(nmd \log m) + nm \cdot O(\frac{1}{\varepsilon})^{kd} = nm \cdot (O(d \log m) + O(\frac{1}{\varepsilon})^{kd})$.

Correctness Consider a query curve Q , and assume that there exists a curve $C_i \in \mathcal{C}$ such that $d_{dF}(C_i, Q) \leq r$. Then, Π_i is a curve of length k and $d_{dF}(C_i, \Pi_i) \leq 2r$. As in the previous section, let Q' be the curve computed by the query algorithm, then $d_{dF}(Q', Q) \leq \frac{\varepsilon r}{2}$. By the triangle inequality, we have $d_{dF}(Q', C_i) \leq d_{dF}(Q', Q) + d_{dF}(Q, C_i) \leq (1 + \frac{\varepsilon}{2})r$, and

$$d_{dF}(Q', \Pi_i) \leq d_{dF}(Q', C_i) + d_{dF}(C_i, \Pi_i) \leq (1 + \frac{\varepsilon}{2})r + 2r \leq 4r.$$

Therefore our data structure contains Q' , and the query algorithm returns $C(Q')$, where $d_{dF}(C(Q'), Q') \leq (1 + \frac{\epsilon}{2})r$. Finally, again by the triangle inequality, we have

$$d_{dF}(C(Q'), Q) \leq d_{dF}(C(Q'), Q') + d_{dF}(Q', Q) \leq (1 + \frac{\epsilon}{2})r + \frac{\epsilon r}{2} = (1 + \epsilon)r.$$

We obtain the following theorem.

Theorem 10 *There exists a data structure for the asymmetric $(1 + \epsilon, r)$ -ANNC under DFD, with $n \cdot O(\frac{1}{\epsilon})^{dk}$ space, $nm \cdot (O(d \log m) + O(\frac{1}{\epsilon})^{kd})$ expected preprocessing time, and $O(kd)$ query time.*

5 $\ell_{p,2}$ -Distance of Polygonal Curves

For the near-neighbor problem under the $\ell_{p,2}$ -distance, we use the same basic approach as in Sect. 3, but with two small modifications. The first is that we set the grid's edge length to $\frac{\epsilon r}{(2m)^{1/p} \sqrt{d}}$, and redefine $G(x, R)$, G_i , and \mathcal{G} , as in Sect. 3 but with respect to the new edge length of our grid. The second modification is that we redefine \mathcal{I}_i to be the set of all curves $\bar{Q} = (x_1, x_2, \dots, x_m)$ with points from \mathcal{G} , such that $d_{p,2}(C_i, \bar{Q}) \leq (1 + \frac{\epsilon}{2})r$.

We assume without loss of generality from now and to the end of this section that $r = 1$ (we can simply scale the entire space by $1/r$), so the grid's edge length is $\frac{\epsilon}{(2m)^{1/p} \sqrt{d}}$. The following corollary is analogous to Corollary 6.

Corollary 11 $|G(x, R)| = O\left(1 + \frac{m^{1/p}}{\epsilon} R\right)^d$.

Proof We scale our grid so that the edge length is 1, hence we are looking for the number of lattice points in $B_2^d(x, \frac{(2m)^{1/p} \sqrt{d}}{\epsilon} R)$. By Lemma 5 we get that this number is bounded by the volume of the d -dimensional ball of radius $(1 + \frac{(2m)^{1/p}}{\epsilon} R) \sqrt{d}$. Using Stirling's formula we conclude,

$$\begin{aligned} V_2^d \left(\left(1 + \frac{(2m)^{1/p}}{\epsilon} R \right) \sqrt{d} \right) &= \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)} \cdot \left(\left(1 + \frac{(2m)^{1/p}}{\epsilon} R \right) \sqrt{d} \right)^d \\ &= \alpha^d \cdot \left(1 + \frac{m^{1/p}}{\epsilon} R \right)^d \end{aligned}$$

where α is a constant (approximately $4.13 \cdot 2^{1/p}$). □

In the following claim we bound the size of \mathcal{I}_i , which, surprisingly, is independent of p .

Claim 12 $|\mathcal{I}_i| = O(\frac{1}{\epsilon})^{m(d+1)}$ and it can be computed in $O(\frac{1}{\epsilon})^{m(d+1)}$ time.

Proof Let $\bar{Q} = (x_1, x_2, \dots, x_m) \in \mathcal{I}_i$, and let τ be an alignment with $\sigma_{p,2}(\tau(C_i, \bar{Q})) \leq (1 + \frac{\epsilon}{2})$. For each $1 \leq k \leq m$ let j_k be the smallest index such that $(j_k, k) \in \tau$. In other words, j_k is the smallest index that is matched to k by the alignment τ .

Set $R_k = \|x_k - p_{j_k}^i\|_2$, then we have $\|(R_1, \dots, R_m)\|_p \leq \sigma_{p,2}(\tau(C_i, \overline{Q})) \leq (1 + \frac{\varepsilon}{2})$.
 Let $\alpha_k = \lceil \frac{m^{1/p}}{\varepsilon} R_k \rceil$. By triangle inequality,

$$\begin{aligned} \|(\alpha_1, \alpha_2, \dots, \alpha_m)\|_p &\leq \frac{m^{1/p}}{\varepsilon} \|(R_1, R_2, \dots, R_m)\|_p + m^{1/p} \\ &\leq \frac{m^{1/p}}{\varepsilon} \left(1 + \frac{\varepsilon}{2}\right) + m^{1/p} < \left(2 + \frac{1}{\varepsilon}\right) m^{1/p}. \end{aligned}$$

Clearly, $x_k \in B_2^d(p_{j_k}^i, \alpha_k \frac{\varepsilon}{m^{1/p}})$.

We conclude that for each curve $\overline{Q} = (x_1, x_2, \dots, x_m) \in \mathcal{I}_i$ there exists an alignment τ such that $\sigma_{p,2}(\tau(C_i, \overline{Q})) \leq 1 + \frac{\varepsilon}{2}$, and a sequence of integers $(\alpha_1, \dots, \alpha_m)$ such that $\|(\alpha_1, \alpha_2, \dots, \alpha_m)\|_p \leq (2 + \frac{1}{\varepsilon})m^{1/p}$ and $x_k \in B_2^d(p_{j_k}^i, \alpha_k \frac{\varepsilon}{m^{1/p}})$, for $k = 1, \dots, m$. Therefore, the number of curves in \mathcal{I}_i is bounded by the multiplication of three numbers:

1. The number of alignments that can determine the distance, which is at most 2^{2m} by Lemma 4.
2. The number of ways to choose a sequence of m positive integers $\alpha_1, \dots, \alpha_m$ such that $\|(\alpha_1, \alpha_2, \dots, \alpha_m)\|_p \leq (2 + \frac{1}{\varepsilon})m^{1/p}$, which is bounded by the number of lattice points in $B_p^m((2 + \frac{1}{\varepsilon})m^{1/p})$ (the m -dimensional ℓ_p -ball of radius $(2 + \frac{1}{\varepsilon})m^{1/p}$). By Lemma 5, this number is bounded by

$$\begin{aligned} V_p^m((2 + \frac{1}{\varepsilon})m^{1/p} + m^{1/p}) &\leq V_p^m(\frac{4m^{1/p}}{\varepsilon}) \\ &= \frac{2^m \Gamma(1 + 1/p)^m}{\Gamma(1 + m/p)} \left(\frac{4m^{1/p}}{\varepsilon}\right)^m = O(\frac{1}{\varepsilon})^m, \end{aligned}$$

where the last equality follows as by Stirling's formula $\frac{m^{m/p}}{\Gamma(1+m/p)} = O(1)^m$.

3. The number of ways to choose a curve (x_1, x_2, \dots, x_m) , such that $x_k \in G(p_{j_k}^i, \alpha_k \frac{\varepsilon}{m^{1/p}})$, for $k = 1, \dots, m$. By Corollary 11, the number of grid points in $G(p_{j_k}^i, \alpha_k \frac{\varepsilon}{m^{1/p}})$ is $O(1 + \alpha_k)^d$, so the number of ways to choose (x_1, x_2, \dots, x_m) is at most $\prod_{k=1}^m O(1 + \alpha_k)^d = O(1)^{md} (\prod_{k=1}^m (1 + \alpha_k))^d$. By the inequality of arithmetic and geometric means we have

$$\begin{aligned} (\prod_{k=1}^m (1 + \alpha_k)^p)^{1/p} &\leq \left(\frac{\sum_{k=1}^m (1 + \alpha_k)^p}{m}\right)^{m/p} \\ &= \left(\frac{\|(1 + \alpha_1, \dots, 1 + \alpha_m)\|_p}{m^{1/p}}\right)^m \\ &\leq \left(\frac{\|1\|_p + \|(\alpha_1, \dots, \alpha_m)\|_p}{m^{1/p}}\right)^m \\ &\leq \left(\frac{m^{1/p} + (2 + \frac{1}{\varepsilon})m^{1/p}}{m^{1/p}}\right)^m = O(\frac{1}{\varepsilon})^m, \end{aligned}$$

$$\text{so } \prod_{k=1}^m O(1 + \alpha_k)^d = O(1)^{md} O(\frac{1}{\varepsilon})^{md} = O(\frac{1}{\varepsilon})^{md}.$$

Finally, $|\mathcal{I}_i| \leq 2^{2m} \cdot O(\frac{1}{\varepsilon})^m \cdot O(\frac{1}{\varepsilon})^{md} \leq O(\frac{1}{\varepsilon})^{m(d+1)}$. □

The data structure and query algorithm are similar to those we described for DFD, and the size of \mathcal{I}_i and \mathcal{I} is roughly the same (here there is an additional $O(\frac{1}{\varepsilon})^m$ factor in the space bound). Therefore, the query time, storage space, and preprocessing time are roughly similar, but we still need to show that the algorithm is correct.

Correctness Consider a query curve $Q = (q_1, \dots, q_m)$. Assume that there exists a curve $C_i \in \mathcal{C}$ such that $d_{p,2}(C_i, Q) \leq 1$. We will show that the query algorithm returns a curve C^* with $d_{p,2}(C^*, Q) \leq 1 + \varepsilon$.

Consider a point $q_k \in Q$. Denote by $q'_k \in \mathcal{G}$ the grid point closest to q_k , and let $Q' = (q'_1, \dots, q'_m)$. We have $\|q_k - q'_k\|_2 \leq \frac{\varepsilon}{2(2m)^{1/p}}$. Let τ be an alignment such that the $\ell_{p,2}$ -cost of τ w.r.t. C_i and Q is at most 1. Unlike the Fréchet distance, $\ell_{p,2}$ -distance for curves does not satisfy the triangle inequality. However, by the triangle inequality under ℓ_2 and ℓ_p , we get that the $\ell_{p,2}$ -cost of τ w.r.t. C_i and Q' is

$$\begin{aligned} \sigma_{p,2}(\tau(C_i, Q')) &= \left(\sum_{(j,t) \in \tau} \|p_j^i - q'_t\|_2^p \right)^{1/p} \leq \left(\sum_{(j,t) \in \tau} (\|p_j^i - q_t\|_2 + \|q_t - q'_t\|_2)^p \right)^{1/p} \\ &\leq \left(\sum_{(j,t) \in \tau} \|p_j^i - q_t\|_2^p \right)^{1/p} + \left(\sum_{(j,t) \in \tau} \|q_t - q'_t\|_2^p \right)^{1/p} \\ &\leq 1 + \left(2m \left(\frac{\varepsilon}{2(2m)^{1/p}} \right)^p \right)^{1/p} = 1 + \frac{\varepsilon}{2}. \end{aligned}$$

So $d_{p,2}(C_i, Q') \leq 1 + \frac{\varepsilon}{2}$, and thus Q' is in $I_i \subseteq \mathcal{I}$. This means that \mathcal{D} contains Q' with a curve $C(Q') \in \mathcal{C}$ such that $d_{p,2}(C(Q'), Q') \leq 1 + \frac{\varepsilon}{2}$, and the query algorithm returns $C(Q')$. Now, again by the same argument (using an alignment with $\ell_{p,2}$ -cost at most $1 + \frac{\varepsilon}{2}$ w.r.t. $C(Q')$ and Q'), we get that $d_{p,2}(C(Q'), Q) \leq 1 + \frac{\varepsilon}{2} + \left(2m \left(\frac{\varepsilon}{2(2m)^{1/p}} \right)^p \right)^{1/p} = 1 + \varepsilon$.

We obtain the following theorem.

Theorem 13 *There exists a data structure for the $(1 + \varepsilon, r)$ -ANNC under $\ell_{p,2}$ -distance, with $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ space, $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ expected preprocessing time, and $O(md)$ query time.*

As mentioned in the preliminaries section, the DTW distance between two curves equals to their $\ell_{1,2}$ -distance, and therefore we obtain the following theorem.

Theorem 14 *There exists a data structure for the $(1 + \varepsilon, r)$ -ANNC under DTW, with $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ space, $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ expected preprocessing time, and $O(md)$ query time.*

6 The Asymmetric Setting Under $\ell_{p,2}$ -distance

In Sect. 4, we strongly rely on the fact that DFD satisfies the triangle inequality, in order to provide a data structure with storage space independent of m . However, the

general $\ell_{p,2}$ distance does not satisfy the triangle inequality, not even up to a constant factor. Lemire [21] proved the following weak version of the triangle inequality for $\ell_{p,2}$ distance, and showed it to be tight: For any three curves A, B, C of length k , it holds that $d_{p,2}(A, B) \leq k^{1/p} (d_{p,2}(A, C) + d_{p,2}(C, B))$.

Nonetheless, in Sect. 5 we showed how to apply our approach to ANNC under the general $\ell_{p,2}$ -distance. This was possible due to the fact that we always match Q' to Q in a “one-to-one” alignment, a special case where the triangle inequality does hold. In our solution to the asymmetric case, we use the triangle inequality between Π_i, C_i , and Q , in which case we do not have a one-to-one matching. Therefore, our analysis for the asymmetric Fréchet distance does not trivially apply to the $\ell_{p,2}$ -distance in general. In this section we show how we can still adapt our algorithm to the asymmetric ANNC under $\ell_{p,2}$ distance. The storage space and query time is almost the same as in Theorem 13, where we replace m by k .

We begin by describing some important characterizations of curve alignments. First, notice that a curve alignment $\tau = \langle (i_1, j_1), \dots, (i_t, j_t) \rangle$ can be viewed as a bipartite graph on the vertices of the two curves. Moreover, we can assume that each connected component in this graph is a star graph, i.e., a single vertex from the first curve is connected to one or more vertices from the second curve, and vice versa. Indeed, as we claimed in the proof of Lemma 4, if there exist pairs (or edges) $(i_k, j_k - 1), (i_k, j_k), (i_k + 1, j_k)$, then removing (i_k, j_k) from τ results in a legal curve alignment with a smaller (or equal) cost. The following special type of curve alignment is crucial in the construction and proof of our algorithm.

One-way alignment An alignment $\tau = \langle (i_1, j_1), \dots, (i_t, j_t) \rangle$ is a *one-way alignment* if for any $1 \leq s \leq t$, we have $i_s = s$. In other words, in the view of τ as a bipartite graph, we get a set of stars such that all the centers are in the second curve, and thus each index of the first curve appears in exactly one pair of τ .

Remark 15 A nice property of one-way alignments is that, while the triangle inequality does not apply for $\ell_{p,2}$ distance between curves in general, it does hold when the $\ell_{p,2}$ distances between the curves are obtained by one-way alignments. See Claim 26 in Appendix C for details.

6.1 Simplification Under $\ell_{p,2}$ -Distance

As for DFD, in order to construct a data structure with storage space independent of m , we will need to compute simplifications of length k for the input curves.² The points in a simplification can be arbitrary, but for the purpose of this section it is enough to use vertex-restricted simplifications. A simplification Π of a curve C is *vertex-restricted* if the points of Π are from the vertices of C , and follow the same ordering.

Lemma 16 *Let $C = (p_1, \dots, p_m)$ be a curve consisting of m points in \mathbb{R}^d . Denote by $\bar{\Pi}$ the closest simplification to C with k points under the $\ell_{p,2}$ -distance, i.e. $\bar{\Pi} = \arg \min_{\Pi, |\Pi|=k} d_{p,2}(C, \Pi)$. Then there exist a vertex-restricted simplification Π of C with k points such that $d_{p,2}(C, \Pi) \leq 2 \cdot d_{p,2}(C, \bar{\Pi})$.*

² See [5] for a closely related more recent result on simplifications with bounded length.

Proof Set $\bar{\Pi} = (x_1, \dots, x_k)$, and let τ be an alignment such that $d_{p,2}(C, \bar{\Pi}) = \sigma_{p,2}(\tau(C, \bar{\Pi}))$. We can assume that τ is a one-way alignment, as otherwise we can remove points from $\bar{\Pi}$ without increasing the distance to C . For $1 \leq j \leq k$, let $A_j = \{i \in [m] \mid (i, j) \in \tau\}$ be the set of indices matched to j by τ . Let $i_j = \arg \min_{t \in A_j} \|p_t - x_j\|_2$. Note that for every $i \in A_j$ it holds that $\|p_i - p_{i_j}\|_2 \leq \|p_i - x_j\|_2 + \|x_j - p_{i_j}\|_2 \leq 2\|p_i - x_j\|_2$. Set $\Pi = (p_{i_1}, \dots, p_{i_k})$, so Π is a vertex-restricted simplification of C consisting of k points from C . It holds that

$$\begin{aligned} \sigma_{p,2}(\tau(C, \bar{\Pi})) &= \left(\sum_{j=1}^k \sum_{i \in A_j} \|p_i - x_j\|_2^p \right)^{1/p} \geq \left(\sum_{j=1}^k \sum_{i \in A_j} \frac{1}{2^p} \|p_i - p_{i_j}\|_2^p \right)^{1/p} \\ &= \frac{1}{2} \sigma_{p,2}(\tau(C, \Pi)). \end{aligned}$$

The lemma follows as $d_{p,2}(C, \Pi) \leq \sigma_{p,2}(\tau(C, \Pi)) \leq 2\sigma_{p,2}(\tau(C, \bar{\Pi})) = 2d_{p,2}(C, \bar{\Pi})$. □

Lemma 17 *Given a curve $C = (p_1, \dots, p_m)$ consisting of m points in \mathbb{R}^d , and parameters $k < m, p \geq 1$, there exists an algorithm that runs in $O(m^3k + m^2d)$ time and computes a vertex-restricted simplification Π of C with k points, such that $d_{p,2}(C, \Pi)$ is minimized.*

Proof We show how to compute Π using a dynamic programming technique. We begin by precomputing all the pairwise distances in C , such that for every i, j , we will have a constant time access to $\|p_i - p_j\|_2^p$. This takes $O(m^2d)$ time (ignoring the time it takes to compute the p -power of a number).

We define $OPT[i, j, x]$ as follows. Let $C[1 : i] = (p_1, \dots, p_i)$, and let $\Pi_j^x = (p_{i_1}, \dots, p_{i_x})$ be a vertex-restricted simplification with x points from $C[1 : j]$, such that $i_x = j$, and $d_{p,2}(C[1 : i], \Pi_j^x)$ is minimized. Then $OPT[i, j, x] = \left(d_{p,2}(C[1 : i], \Pi_j^x) \right)^p$. We compute $OPT[i, j, x]$ for any $1 \leq x \leq k, x \leq j \leq m$, and $x \leq i \leq m$ as follows.

First, for any $1 \leq j \leq m$ we have $OPT[1, j, 1] = \|p_1 - p_j\|_2^p$, and for any $2 \leq i \leq m$ we have $OPT[i, j, 1] = OPT[i - 1, j, 1] + \|p_i - p_j\|_2^p$. Thus $OPT[i, j, 1]$ can be computed in $O(m^2)$ time for all $1 \leq i \leq m, 1 \leq j \leq m$.

Next, we compute $OPT[i, j, x]$ for all $2 \leq x \leq k, x \leq j \leq m$, and $x \leq i \leq m$, using the following formula, in $O(m^3k)$ time:

$$OPT[i, j, x] = \|p_i - p_j\|_2^p + \min \left\{ \min_{x-1 \leq j' < j} OPT[i - 1, j', x - 1], OPT[i - 1, j, x] \right\}.$$

Any alignment w.r.t. $C[1 : i]$ and a simplification Π_j^x has to match p_i and p_j . Let τ be an alignment such that $\sigma_{p,2}(\tau(C[1 : i], \Pi_j^x))$ is minimized. If τ matches p_j to p_{i-1} then $OPT[i, j, x] = \|p_i - p_j\|_2^p + OPT[i - 1, j, x]$. Otherwise, there exists some $x - 1 \leq j' < j$ such that τ matches $p_{j'}$ to p_{i-1} , and $OPT[i, j, x] = \|p_i - p_j\|_2^p + OPT[i - 1, j', x - 1]$.

Finally, we return

$$\min_{\substack{1 \leq x \leq k \\ x \leq j \leq m}} OPT[m, j, x],$$

which can be computed in $O(mk)$ time.

Clearly, the vertex-restricted simplification Π that minimizes $d_{p,2}(C, \Pi)$, and the corresponding alignment, can be found by backtracking in $O(mk)$ time. \square

6.2 The Data Structure

Our data structure for the asymmetric case under $\ell_{p,2}$ -distance is very similar to the asymmetric case under DFD, but, without the triangle inequality, we have to use more involved counting arguments in order to achieve similar bounds.

Fix some $p > 1$ and assume w.l.o.g., as in Sect. 5, that $r = 1$. For every input curve C_i , compute the optimal vertex-restricted simplification Π_i of C_i , with at most k points, using Lemma 17 in $O(m^3k + m^2d)$ time. Note that in addition we obtain a one-way alignment τ_i for which $d_{p,2}(C_i, \Pi_i) = \sigma_{p,2}(\tau_i(C_i, \Pi_i))$. We can assume that $d_{p,2}(C_i, \Pi_i) \leq 2$, as otherwise we can just ignore C_i , since by Lemma 16 we get that for every curve Q of length $\leq k$, $d_{p,2}(C_i, Q) > 1$.

We again construct our data structure for the original problem (for $\ell_{p,2}$ -distance), but with the following modifications. The set of input curves is $\mathcal{P} = \{\Pi_1, \dots, \Pi_n\}$, and the radius parameter is $7k^{1/p}$ (the length of grid edges is still $\frac{\epsilon}{(2m)^{1/p}\sqrt{d}}$). Now let \mathcal{I}'_i be the set of all curves \overline{Q} with k points from G_i , such that $d_{p,2}(\overline{Q}, \Pi_i) \leq 7k^{1/p}$. In addition, let \mathcal{I}_i be the set of all curves $\overline{Q} \in \mathcal{I}'_i$ such that $d_{p,2}(\overline{Q}, C_i) \leq 1 + \frac{\epsilon}{2}$. *Correctness* Consider a query curve $Q = (q_1, \dots, q_k)$ such that $d_{p,2}(Q, C_i) \leq 1$, and let $Q' = (q'_1, \dots, q'_k)$ be the grid curve that was computed by the query algorithm. Following the same arguments as in Sect. 5, there exists an alignment τ such that $\sigma_p(\tau(Q', C_i)) \leq 1 + \frac{\epsilon}{2}$, and thus $d_{p,2}(Q', C_i) \leq 1 + \frac{\epsilon}{2}$. We have that

$$\begin{aligned} d_{p,2}(Q', \Pi_i) &\leq (2k - 1)^{1/p} \cdot d_{dF}(Q', \Pi_i) \\ &\leq (2k - 1)^{1/p} \cdot (d_{dF}(Q', C_i) + d_{dF}(C_i, \Pi_i)) \\ &\leq (2k - 1)^{1/p} \cdot (d_{p,2}(Q', C_i) + d_{p,2}(C_i, \Pi_i)) \\ &\leq (2k - 1)^{1/p} \cdot \left(1 + \frac{\epsilon}{2} + 2\right) \cdot d_{p,2}(Q, C_i) \leq 7k^{1/p}, \end{aligned}$$

where the first and third inequalities follow by the fact that for every $\mathbf{x} \in \mathbb{R}^d$, $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_p \leq d^{1/p} \cdot \|\mathbf{x}\|_\infty$, and the second inequality follows by the triangle inequality of discrete Fréchet distance. It follows that $Q' \in \mathcal{I}_i$. The query algorithm then returns a curve $C(Q')$ such that $d_{p,2}(Q', C(Q')) \leq 1 + \frac{\epsilon}{2}$, and $d_{p,2}(Q, C(Q')) \leq 1 + \epsilon$.

Storage space A calculation using arguments similar to those in Sect. 5 (where \mathcal{I}'_i is computed using the same algorithm, but with a radius multiplied by $7k^{1/p}$), bounds the storage by $O(\frac{1}{\epsilon})^{k(d+1)} \cdot (k \cdot m^d)^{\frac{k}{p}}$, and the preprocessing time by

$nm \cdot \left(O(m^2k + md) \right) + O\left(\frac{1}{\epsilon}\right)^{k(d+1)} \cdot \left(k \cdot m^d \right)^{\frac{k}{p}}$ only (we leave this as an exercise for the diligent reader). Next, we provide a better counting argument and remove the $k^{1/p}$ factor from the base of the exponent.

Fix an input curve $C_i = (p_1, \dots, p_m)$, a simplification Π_i of length k , and a one-way alignment τ_i such that $d_{p,2}(C_i, \Pi_i) = \sigma_{p,2}(\tau_i(C_i, \Pi_i)) \leq 2$. We say that a subset $A \subseteq C_i$ is *consecutive* if it is of the form $(p_s, p_{s+1}, \dots, p_t)$, for some $s, t \in [m]$.

Claim 18 *The points of C_i can be partitioned into $k' \leq 3k$ consecutive subsets $A_1, \dots, A_{k'}$ such that for every $1 \leq s \leq k'$ there exists a center point y_s for which $A_s \subset B_2^d(y_s, \frac{2}{k^{1/p}})$.*

Proof For any $p_j \in C_i$, denote by $\pi(p_j)$ the single point in Π_i that is matched to p_j by τ_i . Let $L = \{p_j \in C_i \mid \|p_j - \pi(p_j)\|_2 > \frac{2}{k^{1/p}}\}$, then $|L| \leq k$, as otherwise we have

$$\begin{aligned} \sigma_{p,2}(\tau_i(C_i, \Pi_i)) &= \left(\sum_{j=1}^m \|p_j - \pi(p_j)\|_2^p \right)^{1/p} \geq \left(\sum_{p_j \in L} \|p_j - \pi(p_j)\|_2^p \right)^{1/p} \\ &> \left(k \cdot \frac{2^p}{k} \right)^{1/p} = 2, \end{aligned}$$

which is a contradiction.

We construct a set \mathcal{A} of consecutive subsets $A_1, \dots, A_{k'}$ and the corresponding center points $y_1, \dots, y_{k'}$ as follows. First, for every $1 \leq s' \leq k$ let $\tilde{A}_{s'} = \{p_j \in C_i \mid (j, s') \in \tau_i\}$. Then $\tilde{A}_1, \dots, \tilde{A}_k$ is a set of consecutive subsets that partition C_i . For every set $\tilde{A}_{s'}$, let $\tilde{A}_{s'} \cap L = \{p_{i_1}, \dots, p_{i_t}\}$, and insert into \mathcal{A} the partition of $\tilde{A}_{s'}$ into $2t + 1$ consecutive subsets $A_{s'}^1, \{p_{i_1}\}, A_{s'}^2, \{p_{i_2}\}, \dots, A_{s'}^t, \{p_{i_t}\}, A_{s'}^{t+1}$. The number of subsets in \mathcal{A} is at most $2|L| + k \leq 3k$. For each subset $A_s \in \mathcal{A}$, if it is a singleton $\{p_j\}$ then we set $y_s = p_j$, otherwise, $A_s = A_{s'}^\ell$ and we set y_s to be the s' point of Π_i . Clearly, $\|p_j - y_s\|_2 \leq \frac{2}{k^{1/p}}$ for every $p_j \in A_s$. □

For each curve $C_i \in \mathcal{C}$ we can compute in $O(m)$ time the set of at most $3k$ consecutive subsets $A_1, \dots, A_{k'}$ and centers $y_1, \dots, y_{k'}$ from Claim 18.

Now, consider a curve $\overline{Q} = (x_1, \dots, x_k) \in \mathcal{I}_i$, and let τ be an alignment such that $\sigma_{p,2}(\tau(C_i, \overline{Q})) \leq 1 + \frac{\epsilon}{2}$. For every $1 \leq j \leq k$, let $1 \leq c(j) \leq m$ be the smallest index such that $(c(j), j) \in \tau$, and let $1 \leq f(j) \leq k'$ be the index such that $p_{c(j)} \in A_{f(j)}$. Since the sets $A_1, \dots, A_{k'}$ are consecutive and τ is an alignment, f is a monotonically non-decreasing function from $[k]$ to $[k']$. It follows that the number of possible functions f is equivalent to the number of ways to distribute k identical balls into k' distinct boxes, thus $\binom{k'-1+k}{k} < \binom{4k}{k} \leq O(2)^k$.

Set $R_j = \|x_j - p_{c(j)}\|_2$ and $\alpha_j = \left\lceil \frac{k^{1/p}}{\epsilon} R_j \right\rceil$. We have $\|(R_1, \dots, R_k)\|_p \leq \sigma_{p,2}(\tau(C_i, \overline{Q})) \leq 1 + \frac{\epsilon}{2}$. By the triangle inequality,

$$\|(\alpha_1, \alpha_2, \dots, \alpha_k)\|_p \leq \left\| \left(\frac{k^{1/p}}{\epsilon} R_1 + 1, \dots, \frac{k^{1/p}}{\epsilon} R_k + 1 \right) \right\|_p$$

$$\leq \frac{k^{1/p}}{\varepsilon} \|(R_1, \dots, R_k)\|_p + k^{1/p} \leq \frac{k^{1/p}}{\varepsilon} (1 + \frac{3}{2}\varepsilon).$$

The number of ways to choose such a sequence $\alpha_1, \dots, \alpha_k$ of k positive integers is bounded by the number of lattice points in $B_p^k(\frac{k^{1/p}}{\varepsilon}(1 + \frac{3}{2}\varepsilon))$ (the k -dimensional ℓ_p -ball of radius $\frac{k^{1/p}}{\varepsilon}(1 + \frac{3}{2}\varepsilon)$). By Lemma 5, this number is bounded by $V_p^k(\frac{k^{1/p}}{\varepsilon}(1 + \frac{3}{2}\varepsilon) + k^{1/p}) = O(\frac{1}{\varepsilon})^k$.

For every $1 \leq j \leq k$ we have $\|p_{c(j)} - y_{f(j)}\|_2 \leq \frac{2}{k^{1/p}}$, and thus

$$\|x_j - y_{f(j)}\|_2 \leq \|x_j - p_{c(j)}\|_2 + \|p_{c(j)} - y_{f(j)}\|_2 \leq \alpha_j \frac{\varepsilon}{k^{1/p}} + \frac{2}{k^{1/p}}.$$

Therefore, once f and $(\alpha_1, \dots, \alpha_k)$ are fixed, it is sufficient to count the number of ways to choose a curve (x_1, x_2, \dots, x_k) such that $x_j \in G(y_{f(j)}, \alpha_j \frac{\varepsilon}{k^{1/p}} + \frac{2}{k^{1/p}})$ for every $1 \leq j \leq k$.

By Corollary 11, as we use the $\frac{\varepsilon}{m^{1/p}}$ grid, the number of grid points in $G(y_{c(j)}, \alpha_j \frac{\varepsilon}{k^{1/p}} + \frac{2}{k^{1/p}})$ is $O(1 + \frac{m^{1/p}}{\varepsilon}(\alpha_j \frac{\varepsilon}{k^{1/p}} + \frac{2}{k^{1/p}}))^d \leq O(\frac{m}{k})^{d/p} \cdot (\frac{3}{\varepsilon} + \alpha_j)^d$. Thus the number of ways to choose (x_1, x_2, \dots, x_k) is at most $O(\frac{m}{k})^{dk/p} (\prod_{j=1}^k (\frac{3}{\varepsilon} + \alpha_j))^d$. By the inequality of arithmetic and geometric means we have

$$\begin{aligned} \prod_{j=1}^k (\frac{3}{\varepsilon} + \alpha_j) &= \left(\prod_{j=1}^k (\frac{3}{\varepsilon} + \alpha_j)^p \right)^{1/p} \\ &\leq \left(\frac{\sum_{j=1}^k (\frac{3}{\varepsilon} + \alpha_j)^p}{k} \right)^{k/p} \\ &\leq \left(\frac{\|(\frac{3}{\varepsilon}, \dots, \frac{3}{\varepsilon})\|_p + \|(\alpha_1, \dots, \alpha_k)\|_p}{k^{1/p}} \right)^k \\ &\leq \left(\frac{\frac{3}{\varepsilon} \cdot k^{1/p} + \frac{k^{1/p}}{\varepsilon}(1 + \frac{3}{2}\varepsilon)}{k^{1/p}} \right)^k = O\left(\frac{1}{\varepsilon}\right)^k. \end{aligned}$$

It follows that for some fixed f and $(\alpha_1, \dots, \alpha_k)$, the number of ways to choose (x_1, \dots, x_k) is bounded by $O(\frac{1}{\varepsilon})^{kd} \cdot (\frac{m}{k})^{kd/p}$.

We conclude that, $|\mathcal{I}_i| = O(2)^k \cdot O(\frac{1}{\varepsilon})^k \cdot (\frac{1}{\varepsilon})^{kd} \cdot (\frac{m}{k})^{kd/p} = O(\frac{1}{\varepsilon})^{k(d+1)} \cdot (\frac{m}{k})^{kd/p}$, and thus $|\mathcal{I}| = n \cdot O(\frac{1}{\varepsilon})^{k(d+1)} \cdot (\frac{m}{k})^{kd/p}$.

Following the same lines as in the counting argument, an efficient implementation of the preprocessing time is self evident.

Theorem 19 *There exists a data structure for the Asymmetric $(1 + \varepsilon, r)$ -ANNC under $d_{p,2}$, with $n \cdot O(\frac{1}{\varepsilon})^{k(d+1)} \cdot (\frac{m}{k})^{kd/p}$ space, $nm \cdot (O(m^2k + md) + O(\frac{1}{\varepsilon})^{k(d+1)} \cdot (\frac{m}{k})^{kd/p})$ preprocessing time, and $O(kd)$ query time.*

7 Approximate Range Counting

In the range counting problem for curves, we are given a set \mathcal{C} of n curves, each consisting of m points in d dimensions, and a distance measure for curves δ . The goal is to preprocess \mathcal{C} into a data structure that given a query curve Q and a threshold value r , returns the number of curves that are within distance r from Q .

In this section we consider the following approximation version of range counting for curves, in which r is part of the input. Note that by storing pointers to curves instead of just counters, we can obtain a data structure for the approximate range searching problem (at the cost of an additional $O(n)$ -factor to the storage space).

Problem 20 $[(1 + \varepsilon, r)$ -approximate range-counting for curves] Given a parameter r and $0 < \varepsilon \leq 1$, preprocess \mathcal{C} into a data structure that given a query curve Q , returns the number of all the input curves whose distance to Q is at most r plus possibly additional input curves whose distance to Q is greater than r but at most $(1 + \varepsilon)r$.

We construct the dictionary \mathcal{D} (implemented as a dynamic hash table, or a prefix tree) for the curves in \mathcal{I} as in Sect. 5, as follows. For each $1 \leq i \leq n$ and curve $\overline{Q} \in \mathcal{I}_i$, if \overline{Q} is not in \mathcal{D} , insert it into \mathcal{D} and initialize $C(\overline{Q}) \leftarrow 1$. Otherwise, if \overline{Q} is in \mathcal{D} , update $C(\overline{Q}) \leftarrow C(\overline{Q}) + 1$. Notice that $C(\overline{Q})$ holds the number of curves from \mathcal{C} that are within distance $(1 + \frac{\varepsilon}{2})r$ to \overline{Q} . Given a query curve Q , we compute Q' as in Sect. 5. If Q' is in \mathcal{D} , we return $C(Q')$, otherwise, we return 0.

Clearly, the storage space, preprocessing time, and query time are similar to those in Sect. 5. We claim that the query algorithm returns the number of curves from \mathcal{C} that are within distance r to Q plus possibly additional input curves whose distance to Q is greater than r but at most $(1 + \varepsilon)r$. Indeed, let C_i be a curve such that $d_{p,2}(C_i, Q) \leq r$. As shown in Sect. 5 we get $d_{p,2}(C_i, Q') \leq (1 + \frac{\varepsilon}{2})r$, so Q' is in \mathcal{I}_i and C_i is counted in $C(Q')$. Now let C_i be a curve such that $d_{p,2}(C_i, Q) > (1 + \varepsilon)r$. If $d_{p,2}(C_i, Q') \leq (1 + \frac{\varepsilon}{2})r$, then by a similar argument (switching the rolls of Q and Q') we get that $d_{p,2}(C_i, Q') \leq (1 + \varepsilon)r$, a contradiction. So $d_{p,2}(C_i, Q') > (1 + \frac{\varepsilon}{2})r$, and thus C_i is not counted in $C(Q')$.

We obtain the following theorem.

Theorem 21 *There exists a data structure for the $(1 + \varepsilon, r)$ -approximate range-counting for curves under $\ell_{p,2}$ -distance, with $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ space, $n \log(\frac{n}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ preprocessing time, and $O(md \log(\frac{nm}{\varepsilon}))$ query time. (Under DFD, the exponent in the bounds for the space and preprocessing time is md rather than $m(d + 1)$.)*

8 Simplification in d -dimensions

The algorithm of Bereg et al. [6] receives as input a curve C consisting of m points in \mathbb{R}^3 , and a parameter $r > 0$. In $O(m \log m)$ time, it returns a curve Π such that $d_{dF}(C, \Pi) \leq r$, and Π has the minimum number of vertices among all curves within distance r from C . The algorithm operates in a greedy manner, by repeatedly executing Megiddo's [22] minimum enclosing ball (MEB) algorithm for points in \mathbb{R}^3 , which takes linear time.

We generalize the algorithm of Bereg et al. for curves in \mathbb{R}^d , by using an algorithm presented by Kumar et al. [20] for approximated minimum enclosing ball (AMEB) in \mathbb{R}^d . Formally, given a set A of n points in \mathbb{R}^d and a parameter $\varepsilon \in (0, 1]$, the goal is to find an enclosing ball of A with radius $r > 0$, where the minimum enclosing ball of A has radius at least $\frac{r}{1+\varepsilon}$. The algorithm of [20] can find an AMEB in $O(\frac{nd}{\varepsilon} + \varepsilon^{-4.5} \log \frac{1}{\varepsilon})$ time. In particular, given an additional parameter $r > 0$, this algorithm either returns an enclosing ball of A with radius $(1 + \varepsilon)r$, or declares that the minimum enclosing ball of A has radius larger than r .

Next, we describe our modified algorithm. Consider a curve $C = (x_1, \dots, x_m)$, and denote $C[i, j] = (x_i, \dots, x_j)$. The following sub-procedure takes as input a curve A and returns a point y and an index s , such that the ball with radius $(1 + \varepsilon)r$ centered at y covers the prefix $A[1, s]$, and (if $s < |A|$) the minimum enclosing ball of $A[1, s + 1]$ has radius larger than r .

1. By iterative probing, using an algorithm for AMEB, find some t such that $A[1, 2^t]$ can be covered by a ball of radius $(1 + \varepsilon)r$, while $A[1, 2^{t+1}]$ cannot be covered by a ball of radius r . If all the points in A can be enclosed by a single ball of radius $(1 + \varepsilon)r$ centered at y , simply return y and $|A|$.
2. By binary search, again using an algorithm for AMEB, find some $s \in [2^t, 2^{t+1})$ such that $A[1, s]$ can be covered by a ball of radius $(1 + \varepsilon)r$, and $A[1, s + 1]$ cannot be covered by a ball of radius r . Let $y \in \mathbb{R}^d$ be the center of this ball. Return y and s .

Starting from the input $A = C[1, m]$, repeat the above sub-procedure such that in each step the input is the suffix of C that was not yet covered by the previous steps (i.e. $A[s + 1, m]$). Let (y_1, \dots, y_q) be the sequence of output points.

Lemma 9 is an easy corollary of the following lemma.

Lemma 22 *Let C be a curve consisting of m points in \mathbb{R}^d . Given parameters $r > 0$, and $\varepsilon \in (0, 1]$, the algorithm above runs in $O\left(\frac{d \cdot m \log m}{\varepsilon} + m \cdot \varepsilon^{-4.5} \log \frac{1}{\varepsilon}\right)$ time and returns a curve $\Pi = (y_1, \dots, y_q)$ such that $d_{dF}(C, \Pi) \leq (1 + \varepsilon)r$. Furthermore, for every curve Π' with less than q points, it holds that $d_{dF}(C, \Pi') > r$.*

Proof We start by analyzing the running time for a single iteration of the sub-procedure, when using the algorithm of [20] to find an AMEB. The total time for the first step of the sub-procedure (finding t) is

$$\sum_{i=1}^{t+1} O\left(\frac{2^i \cdot d}{\varepsilon} + \varepsilon^{-4.5} \log \frac{1}{\varepsilon}\right) = O\left(\frac{2^t \cdot d}{\varepsilon} + t \cdot \varepsilon^{-4.5} \log \frac{1}{\varepsilon}\right).$$

In the second step, there are $O(t)$ executions of [20] on a set of size at most 2^{t+1} , so the total time for this step is $t \cdot O(\frac{2^t \cdot d}{\varepsilon} + \varepsilon^{-4.5} \log \frac{1}{\varepsilon})$.

Let m_i be the length of the subcurve covered by the point y_i that was found in the i 'th iteration of the sub-procedure. The total time spent for finding y_i is therefore $\log m_i \cdot O(\frac{m_i \cdot d}{\varepsilon} + \varepsilon^{-4.5} \log \frac{1}{\varepsilon})$, and the total running time of the algorithm is

$$\sum_{i=1}^q \log m_i \cdot O\left(\frac{m_i \cdot d}{\varepsilon} + \varepsilon^{-4.5} \log \frac{1}{\varepsilon}\right) = O\left(\frac{d \cdot m \log m}{\varepsilon} + m \cdot \varepsilon^{-4.5} \log \frac{1}{\varepsilon}\right),$$

where we used the concavity of the log function, and the fact $\sum_{i=1}^q m_i = m$.

Next we argue the correctness. Clearly, $d_{dF}(C, \Pi) \leq (1 + \varepsilon)r$. Let $s_0 = 0, s_1, \dots, s_q = m$ be the sequence of indices (of vertices in C) found during the execution of the algorithm, such that the ball of radius $(1 + \varepsilon)r$ around y_i covers $C[s_{i-1} + 1, s_i]$. It follows by a straightforward induction that every curve Π' with less than i points will be at distance greater than r from $C[1, s_{i-1} + 1]$. The lemma now follows. □

Acknowledgements We wish to thank Boris Aronov for helpful discussions on the problems studied in this paper.

Funding Arnold Filtser was partially supported by Grant 1042/22 from the Israel Science Foundation. Omrit Filtser was supported by the Eric and Wendy Schmidt Fund for Strategic Innovation, by the Council for Higher Education of Israel, and by Ben-Gurion University of the Negev. Matthew J. Katz was partially supported by Grant 1884/16 from the Israel Science Foundation.

Appendix: A Deterministic Construction Using a Prefix Tree

When implementing the dictionary \mathcal{D} as a hash table, the construction of the data structure is randomized and thus in the worst case we might get higher preprocessing time. To avoid this, we can implement \mathcal{D} as a prefix tree.

Appendix: A.1 Discrete Fréchet Distance

In this section we describe the implementation of \mathcal{D} as a prefix tree in the case of ANNC under DFD.

We can construct a prefix tree \mathcal{T} for the curves in \mathcal{I} , where any path in \mathcal{T} from the root to a leaf corresponds to a curve that is stored in it. For each $1 \leq i \leq n$ and curve $\overline{Q} \in \mathcal{I}_i$, if $\overline{Q} \notin \mathcal{T}$, insert \overline{Q} into \mathcal{T} , and set $C(\overline{Q}) \leftarrow C_i$.

Each node $v \in \mathcal{T}$ corresponds to a grid point from \mathcal{G} . Denote the set of v 's children by $N(v)$. We store with v a multilevel search tree on $N(v)$, with a level for each coordinate. The points in \mathcal{G} are the grid points contained in nm balls of radius $(1 + \varepsilon)r$. Thus when projecting these points to a single dimension, the number of 1-dimensional points is at most $nm \cdot \frac{\sqrt{d}(1+\varepsilon)2r}{\varepsilon} = O(\frac{nm\sqrt{d}}{\varepsilon})$. So in each level of the search tree on $N(v)$ we have $O(\frac{nm\sqrt{d}}{\varepsilon})$ 1-dimensional points, so the query time is $O(d \log(\frac{nm\sqrt{d}}{\varepsilon}))$.

Inserting a curve of length m to the tree \mathcal{T} takes $O(md \log(\frac{nm\sqrt{d}}{\varepsilon}))$ time. Since \mathcal{T} is a compact representation of $|\mathcal{I}| = n \cdot O(\frac{1}{\varepsilon})^{dm}$ curves of length m , the number of nodes in \mathcal{T} is $m \cdot |\mathcal{I}| = nm \cdot O(\frac{1}{\varepsilon})^{dm}$. Each node $v \in \mathcal{T}$ contains a search tree for its children of size $O(d \cdot |N(v)|)$, and $\sum_{v \in \mathcal{T}} |N(v)| = nm \cdot O(\frac{1}{\varepsilon})^{dm}$ so the total space complexity is $O(nmd) \cdot O(\frac{1}{\varepsilon})^{md} = n \cdot O(\frac{1}{\varepsilon})^{md}$. Constructing \mathcal{T} takes $O(|\mathcal{I}| \cdot md \log(\frac{nm\sqrt{d}}{\varepsilon})) = n \log(\frac{nm\sqrt{d}}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{md}$ time.

Theorem 23 *There exists a data structure for the $(1 + \varepsilon, r)$ -ANNC under DFD, with $n \cdot O(\frac{1}{\varepsilon})^{dm}$ space, $n \cdot \log(\frac{n}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{md}$ preprocessing time, and $O(md \log(\frac{nm}{\varepsilon}))$ query time.*

Similarly, for the asymmetric case we obtain the following theorem.

Theorem 24 *There exists a data structure for the asymmetric $(1 + \varepsilon, r)$ -ANNC under DFD, with $n \cdot O(\frac{1}{\varepsilon})^{dk}$ space, $nm \log(\frac{n}{\varepsilon}) \cdot (O(d \log m) + O(\frac{1}{\varepsilon})^{kd})$ preprocessing time, and $O(kd \log(\frac{nk}{\varepsilon}))$ query time.*

Appendix: A.2 $\ell_{p,2}$ -Distance

For the case of ANNC under $\ell_{p,2}$ -distance, the total number of curves stored in the tree \mathcal{T} is roughly the same as in the case of DFD. We only need to show that for a given node v of the tree \mathcal{T} , the upper bound on the size and query time of the search tree associated with it are similar.

The grid points corresponding to the nodes in $N(v)$ are from n sets of m balls with radius $(1 + \varepsilon)$. When projecting the grid points in one of the balls to a single dimension, the number of 1-dimensional points is at most $\frac{m^{1/p} \sqrt{d}}{\varepsilon} \cdot (1 + \varepsilon)$, so the total number of projected points is at most $\frac{nm^{1+\frac{1}{p}} \sqrt{d}}{\varepsilon} \cdot (1 + \varepsilon)$.

Thus in each level of the search tree of v we have $O(\frac{nm^2 \sqrt{d}}{\varepsilon})$ 1-dimensional points, so the query time is $O(d \log(\frac{nm}{\varepsilon}))$, and inserting a curve of length m into the tree \mathcal{T} takes $O(md \log(\frac{nm}{\varepsilon}))$ time. Note that the size of the search tree of v remains $O(d \cdot |N(v)|)$.

We conclude that the total space complexity is $O(\frac{nm^2 \sqrt{d}}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{m(d+1)} = n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$, constructing \mathcal{T} takes $O(|\mathcal{I}| \cdot md \log(nmd/\varepsilon)) = n \log(\frac{n}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ time, and the total query time is $O(md \log(\frac{nm}{\varepsilon}))$.

Theorem 25 *There exists a data structure for the $(1 + \varepsilon, r)$ -ANNC under $\ell_{p,2}$ -distance, with $n \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ space, $n \cdot \log(\frac{n}{\varepsilon}) \cdot O(\frac{1}{\varepsilon})^{m(d+1)}$ preprocessing time, and $O(md \log(\frac{nm}{\varepsilon}))$ query time.*

Appendix: B Dealing with Query Curves and Input Curves of Varying Size

Notice that if an input curve C_i has length $t < m$, then the size of the set of candidates \mathcal{I}_i (and \mathcal{I}'_i in the asymmetric case) can only decrease.

In addition, our assumption that all query curves are of length exactly k can be easily removed by constructing k data structures $\mathcal{D}_1, \dots, \mathcal{D}_k$, where \mathcal{D}_i is our data structure constructed for query curves of length i (instead of k), for $1 \leq i \leq k$. Clearly, the query time does not change. The storage space is multiplied by k , so for the case of DFD we have storage space $nk \cdot O(\frac{1}{\varepsilon})^{kd}$, but $k < 2^{kd}$, so the storage space remains $n \cdot O(\frac{1}{\varepsilon})^{kd}$. Similarly, for the case of $\ell_{p,2}$ -distance we obtain storage space of $n \cdot O(\frac{1}{\varepsilon})^{k(d+1)} \cdot (\frac{m}{k})^{kd/p}$.

Appendix: C One-Way Alignments

Claim 26 *Let A, B, C be three curves, and let τ_1, τ_2 be two one-way alignments such that τ_1 matches C to A and τ_2 matches C to B . Then $d_{p,2}(A, B) \leq \sigma_{p,2}(\tau_1(C, A)) + \sigma_{p,2}(\tau_2(C, B))$.*

Proof Denote by k_A, k_B, k_C the lengths of the curves A, B, C respectively. Consider the following algorithm that constructs an alignment τ . For every $1 \leq x \leq k_C$, denote by i_x, j_x the unique indexes such that $(x, i_x) \in \tau_1$ and $(x, j_x) \in \tau_2$. Add the pair (i_x, j_x) to τ if it is not already there.

First, we need to show that $\tau = \langle (i_1, j_1), \dots, (i_t, j_t) \rangle$ is a valid alignment. Clearly, $(i_1, j_1) = (1, 1)$ because $(1, 1) \in \tau_1$ and $(1, 1) \in \tau_2$. Similarly, $(i_t, j_t) = (k_A, k_B)$ because $(k_C, k_A) \in \tau_1$ and $(k_C, k_B) \in \tau_2$.

For any $1 \leq s < t$, consider the two consecutive pairs $(i_s, j_s), (i_{s+1}, j_{s+1}) \in \tau$. Let x_1 be an index such that $(x_1, i_s) \in \tau_1$ and $(x_1, j_s) \in \tau_2$, and x_2 an index such that $(x_2, i_{s+1}) \in \tau_1$ and $(x_2, j_{s+1}) \in \tau_2$. Since τ_1, τ_2 are one-way alignments, we have $x_1 \neq x_2$. Moreover, since the algorithm added (i_s, j_s) to τ before (i_{s+1}, j_{s+1}) , we have $x_1 < x_2$. This implies that $i_{s+1} \geq i_s$ and $j_{s+1} \geq j_s$. Assume by contradiction that $i_{s+1} > i_s + 1$, and let x be the index such that $(x, i_s + 1) \in \tau_1$, then $x_1 < x < x_2$ and thus the algorithm adds a pair $(i_s + 1, j)$ for some index j after (i_s, j_s) and before (i_{s+1}, j_{s+1}) , a contradiction. So we have $i_s \leq i_{s+1} \leq i_s + 1$, and by symmetric arguments, $j_s \leq j_{s+1} \leq j_s + 1$, and therefore τ is valid.

Using the triangle inequality for the ℓ_p norm, we get that

$$\begin{aligned} d_{p,2}(A, B) &\leq \sigma_{p,2}(\tau(A, B)) = \left(\sum_{(i,j) \in \tau} \|a_i - b_j\|_2^p \right)^{1/p} \\ &\leq \left(\sum_{x=1}^{k_C} \|a_{i_x} - b_{j_x}\|_2^p \right)^{1/p} \\ &\leq \left(\sum_{x=1}^{k_C} \|a_{i_x} - c_x\|_2^p \right)^{1/p} + \left(\sum_{x=1}^{k_C} \|c_x - b_{j_x}\|_2^p \right)^{1/p} \\ &= \sigma_{p,2}(\tau_1(C, A)) + \sigma_{p,2}(\tau_2(C, B)) . \end{aligned}$$

□

References

1. Afshani, P., Driemel, A.: On the complexity of range searching among curves. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018, pp 898-917, (2018), <https://doi.org/10.1137/1.9781611975031.58>
2. Aronov, B., Filtser, O., Horton, M., Katz, M.J., Sheikhan, K.: Efficient nearest-neighbor query and clustering of planar curves. In: Algorithms and Data Structures—16th International Symposium, WADS 2019, Edmonton, AB, Canada, August 5-7, 2019, Proceedings, pp 28-42 (2019), https://doi.org/10.1007/978-3-030-24766-9_3
3. Buchin, K., Driemel, A., Gudmundsson, J., Horton, M., Kostitsyna, I., Löffler, M., Struijs, M.: Approximating (k, l) -center clustering for curves. In: Proceedings of the Thirtieth Annual ACM-SIAM

- Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019, pp 2922–2938, (2019), <https://doi.org/10.1137/1.9781611975482.181>
4. Bringmann, K., Driemel, A., Nusser, A., Psarros, I.: Tight bounds for approximate near neighbor searching for time series under the Fréchet distance. In: Symposium on Discrete Algorithms, SODA (2022)
 5. Buchin, M., Driemel, A., van Greevenbroek, K., Psarros, I., Rohde, D.: Approximating length-restricted means under dynamic time warping. In: Approximation and Online Algorithms—20th International Workshop, WAOA, volume 13538, pp 225–253, (2022), https://doi.org/10.1007/978-3-031-18367-6_12
 6. Bereg, S., Jiang, M., Wang, W., Yang, B., Zhu, B.: Simplifying 3D polygonal chains under the discrete Fréchet distance. In LATIN 2008: Theoretical Informatics, 8th Latin American Symposium, Búzios, Brazil, April 7-11, 2008, Proceedings, pp 630–641, (2008), https://doi.org/10.1007/978-3-540-78773-0_54
 7. Bringmann, K.: Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014, pp 661–670, 2014, <https://doi.org/10.1109/FOCS.2014.76>
 8. de Berg, M., Cook, A.F., IV., Gudmundsson, J.: Fast Fréchet queries. *Comput. Geom.* **46**(6), 747–755 (2013). <https://doi.org/10.1016/j.comgeo.2012.11.006>
 9. de Berg, M., Gudmundsson, J., Mehrabi, A. D.: A dynamic data structure for approximate proximity queries in trajectory data. In: Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS 2017, Redondo Beach, CA, USA, November 7–10, 2017, pp 48:1–48:4, (2017), <https://doi.org/10.1145/3139958.3140023>
 10. Driemel, A., Har-Peled, S.: Jaywalking your dog: Computing the Fréchet distance with shortcuts. *SIAM J. Comput.* **42**(5), 1830–1866 (2013). <https://doi.org/10.1137/120865112>
 11. Driemel, A., Psarros, I.: ANN for time series under the Fréchet distance. In: A. Lubiw and M. R. Salavatipour, editors, Algorithms and Data Structures—17th International Symposium, WADS 2021, Virtual Event, August 9-11, 2021, Proceedings, volume 12808 of Lecture Notes in Computer Science, pp 315–328. Springer, (2021), https://doi.org/10.1007/978-3-030-83508-8_23
 12. Driemel, A., Psarros, I., Schmidt, M.: Sublinear data structures for short Fréchet queries. *CoRR*, abs/1907.04420, 2019, [arXiv:1907.04420](https://arxiv.org/abs/1907.04420)
 13. Driemel, A., Silvestri, F.: Locality-sensitive hashing of curves. In Proceedings of the 33rd International Symposium on Computational Geometry, volume 77, pp 37:1–37:16. Brisbane, Australia, July 2017. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, <https://doi.org/10.4230/LIPIcs.SoCG.2017.37>
 14. Emiris, I.Z., Psarros, I.: Products of Euclidean metrics, applied to proximity problems among curves: unified treatment of discrete Fréchet and dynamic time warping distances. *ACM Trans. Spatial Algorithms Syst.* **6**(4), 27:1–27:20 (2020). <https://doi.org/10.1145/3397518>
 15. Filtser, A., Filtser, O., Katz, M. J.: Approximate nearest neighbor for curves—simple, efficient, and deterministic. In: A. Czumaj, A. Dawar, and E. Merelli, editors, 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference), volume 168 of LIPIcs, pages 48:1–48:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, <https://doi.org/10.4230/LIPIcs.ICALP.2020.48>
 16. Har-Peled, S., Indyk, P., Motwani, R.: Approximate nearest neighbor: towards removing the curse of dimensionality. *Theory Comput.* **8**(1), 321–350 (2012). <https://doi.org/10.4086/toc.2012.v008a014>
 17. Har-Peled, S., Kumar, N.: Approximate nearest neighbor search for low dimensional queries. In: D. Randall, editor, Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23–25, 2011, pp 854–867. SIAM, 2011, <https://doi.org/10.1137/1.9781611973082.67>
 18. Indyk, P.: High-dimensional computational geometry. PhD thesis, Stanford University, 2000
 19. Indyk, P.: Approximate nearest neighbor algorithms for Fréchet distance via product metrics. In: Proceedings of the 8th Symposium on Computational Geometry, pp 102–106, Barcelona, Spain, June 2002. ACM Press, <https://doi.org/10.1145/513400.513414>
 20. Kumar, P., Mitchell, J. S. B., Yildirim, E. A.: Computing core-sets and approximate smallest enclosing hyperspheres in high dimensions. In: Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments, Baltimore, MD, USA, January 11, 2003, pp 45–55, (2003), <https://doi.org/10.1145/996546.996548>

21. Lemire, D.: Faster retrieval with a two-pass dynamic-time-warping lower bound. *Pattern Recogn.* **42**(9), 2169–2180 (2009). <https://doi.org/10.1016/j.patcog.2008.11.030>
22. Megiddo, N.: Linear programming in linear time when the dimension is fixed. *J. ACM* **31**(1), 114–127 (1984). <https://doi.org/10.1145/2422.322418>
23. Pagh, R., Rodler, F.F.: Cuckoo hashing. *J. Algorithms* **51**(2), 122–144 (2004). <https://doi.org/10.1016/j.jalgor.2003.12.002>
24. Shakhnarovich, G., Darrell, T., Indyk, P.: *Nearest-Neighbor Methods in Learning and Vision: Theory and Practice* (neural Information Processing). The MIT press, Cambridge (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.