

Transforming Dogs on the Line: On the Fréchet Distance Under Translation or Scaling in 1D

Lotte Blank  

University of Bonn, Germany

Jacobus Conradi  

University of Bonn, Germany

Anne Driemel  

University of Bonn, Germany

Benedikt Kolbe  

University of Bonn, Germany

André Nusser  

Université Côte d’Azur, CNRS, Inria, Sophia Antipolis, France

Marena Richter  

University of Bonn, Germany

Abstract

The Fréchet distance is a computational mainstay for comparing polygonal curves. The Fréchet distance under translation, which is a translation invariant version, considers the similarity of two curves independent of their location in space. It is defined as the minimum Fréchet distance that arises from allowing arbitrary translations of the input curves. This problem and numerous variants of the Fréchet distance under some transformations have been studied, with more work concentrating on the discrete Fréchet distance, leaving a significant gap between the discrete and continuous versions of the Fréchet distance under transformations. Our contribution is twofold: First, we present an algorithm for the Fréchet distance under translation on 1-dimensional curves of complexity n with a running time of $\mathcal{O}(n^{8/3} \log^3 n)$. To achieve this, we develop a novel framework for the problem for 1-dimensional curves, which also applies to other scenarios and leads to our second contribution. We present an algorithm with the same running time of $\mathcal{O}(n^{8/3} \log^3 n)$ for the Fréchet distance under scaling for 1-dimensional curves. For both algorithms we match the running times of the discrete case and improve the previously best known bounds of $\tilde{\mathcal{O}}(n^4)$. Our algorithms rely on technical insights but are conceptually simple, essentially reducing the continuous problem to the discrete case across different length scales.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Fréchet distance under translation, Fréchet distance under scaling, time series, shape matching

Digital Object Identifier 10.4230/LIPIcs.SoCG.2025.22

Related Version *Full Version:* <https://arxiv.org/abs/2501.12821>

Funding *Lotte Blank:* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 459420781 (FOR AlgoForGe).

Jacobus Conradi: Funded by the iBehave Network: Sponsored by the Ministry of Culture and Science of the State of North Rhine-Westphalia.

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

Benedikt Kolbe: This work was partially supported by the Lamarr Institute for Machine Learning and Artificial Intelligence.

André Nusser: This work was supported by the French government through the France 2030 investment plan managed by the National Research Agency (ANR), as part of the Initiative of Excellence of Université Côte d’Azur under reference number ANR-15-IDEX-01.




© Lotte Blank, Jacobus Conradi, Anne Driemel, Benedikt Kolbe, André Nusser, and Marena Richter; licensed under Creative Commons License CC-BY 4.0

41st International Symposium on Computational Geometry (SoCG 2025).
Editors: Oswin Aichholzer and Haitao Wang; Article No. 22; pp. 22:1–22:16



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



1 Introduction

The Fréchet distance [26] is one of the most well-studied distance measures for polygonal curves, with a long history in both applications [7, 17, 13] and algorithmic research [1, 21, 27]. For many areas of interest, including the analysis of stock market trends, electrocardiograms or electroencephalograms, the underlying behavior to be analyzed crucially involves the behavior of 1-dimensional curves, or *time series* [5, 29]. One of the advantages that the Fréchet distance offers over competing distance measures such as the Hausdorff distance is that the Fréchet distance considers a joint traversal of both curves, which naturally takes the orientation and ordering of the vertices of both curves into account. For 1-dimensional curves, the difference is particularly jarring, as representing a *continuous* curve by the points lying on it essentially discards all information concerning the underlying dynamics when measuring the Hausdorff distance. Intuitively, the Fréchet distance between two curves can be thought of the shortest possible length of a rope that can connect two climbers that cooperatively travel along their respective routes from start to end. Another well-known analogy stems from the metaphor of a person walking a dog. In this setting, the connecting rope corresponds to a dog leash and we ask for the shortest leash that can be used for the walk.

The algorithmic complexity of the computation of the Fréchet distance has been thoroughly investigated and is well understood. For curves of complexity n , the Fréchet distance can be computed in time $\tilde{O}(n^2)$ [1, 12, 16], where \tilde{O} hides polylogarithmic factors in n . Furthermore, this is hard in the sense that there cannot exist an $\mathcal{O}(n^{2-\epsilon})$ time algorithm in the plane for any $\epsilon > 0$, unless the Strong Exponential Time Hypothesis (SETH) fails even in one dimension [8, 15].

The Fréchet distance has been studied in many different contexts, leading to a plethora of variants, inspired both by applications as well as theoretical considerations, see [14, 19, 20, 24]. Each of the available variants fits into one of two classes, according to whether only the distances of vertices of the curves are considered (the discrete version), or whether the computed distance depends on all points along the straight line segments connecting them (the continuous variant).

While the continuous Fréchet distance naturally considers all points between two vertices and tends to give better practical results (e.g., no discretization artifacts), the discrete Fréchet distance is often easier to handle algorithmically as well as implementation-wise. One overarching goal of Fréchet related research is to match the algorithmic results that are attained in the discrete version with those in the continuous case.

Fréchet distance under translation or scaling. An important aspect of detecting movement patterns in different application areas is to consider the movement independent of its absolute location and scale, i.e., we want to know for which location and scale do the patterns look most similar. Consider, for example, a cardiogram of a patient with heart arrhythmia [5]. Intuitively, some of the shared characteristics of the observed heartbeats do not depend on their absolute location/scaling but instead on the relative movement of the wave. Similarly, identifying a trend in the stock market does not necessarily depend on the exact values of the stocks involved. Instead, a group of stocks may display similar relative behavior although their absolute values are far apart, in which case a natural approach is to allow arbitrary translation and rescaling of the curves to distill their concurrent behavior. These motivations directly lead to the idea of the Fréchet distance under translation or scaling. Concretely, the Fréchet distance under translation (resp. scaling) is the minimum Fréchet distance that we obtain by allowing an arbitrary translation (resp. scaling) of one curve.

Related work. As mentioned previously, due to the simpler setting, there are often more results in the discrete setting, as is also the case for the *discrete* Fréchet distance under translation. This problem was previously mainly considered in the Euclidean plane. The first algorithm due to Jiang, Xu, and Zhu [28] builds an arrangement of size $\mathcal{O}(n^4)$ over the translation space, and then queries an arbitrary representative of each cell of this arrangement with the quadratic-time Fréchet distance algorithm, leading to an $\mathcal{O}(n^6 \log n)$ running time (where the logarithmic factor stems from the usage of parametric search). Avraham, Kaplan, and Sharir [4] crucially observed that instead of re-computing the Fréchet distance for each cell of the arrangement, one can instead formulate the problem as an online dynamic reachability problem on a grid graph and provide a data structure with update time $\mathcal{O}(n)$, leading to an $\mathcal{O}(n^5 \log n)$ algorithm. The current best-known algorithm observes that these updates are offline, i.e., they can be known in advance, and design a tailored data structure for this problem reducing the update time to $\mathcal{O}(n^{2/3} \log^2 n)$, leading to a running time of $\mathcal{O}(n^{4+\frac{2}{3}} \log^3 n)$.¹ This result is complemented by a lower bound of $\mathcal{O}(n^{4-o(1)})$, conditional on the Strong Exponential Time Hypothesis (SETH), for the discrete Fréchet distance under translation in the plane [11].

The currently best algorithm for discrete time series that is explicitly described in the literature has running time $\mathcal{O}(n^3)$ [25], based on the solution to the decision version of the problem in [4]. We note that the approach of Bringmann, Künnemann, and Nusser [11] is also directly applicable to the 1-dimensional problem, which leads to an $\mathcal{O}(n^{8/3} \log^3 n)$ algorithm.²

The *continuous* Fréchet distance under translations was first studied in 2001 [2, 3, 23]. Shortly after, Wenk [30] considered the Fréchet distance under scaling and developed the currently best published algorithms for both variants (translation or scaling) in the continuous setting in d -dimensional Euclidean space. For time series, these algorithms achieve a running time of $\mathcal{O}(n^5 \log n)$. We note that for time series there is an unpublished $\tilde{\mathcal{O}}(n^4)$ algorithm for both problems that can be considered folklore. We briefly sketch the idea here. First, we build the arrangement of intervals separated by events at which the decision problem for the Fréchet distance under translation/scaling is subject to change. The crux is that in one dimension, there are only $\mathcal{O}(n^2)$ such events, as the edge-vertex-vertex events that appear in higher dimensions can be omitted as they are captured by the vertex-vertex events (see Corollary 23 and the full version for more details). Thus, computing the Fréchet distance for an arbitrary representative translation/scaling in each cell together with parametric search [18] yields an $\tilde{\mathcal{O}}(n^4)$ algorithm. We want to stress that, while indeed the events in one dimension only depend on vertex-vertex alignments, this does not directly enable the usage of the grid reachability data structures of [4] and [11], as we still have to compute the *continuous* Fréchet distance for each cell.

1.1 Our Contributions

We present the first new result on the continuous Fréchet distance under translation and the continuous Fréchet distance under scaling since the results by Wenk [30] in 2002. Concretely, we obtain the following two theorems:

¹ Meanwhile, this problem was also tackled from the algorithm engineering side, making use of Lipschitz optimization instead of dynamic graph algorithms [10].

² This is a result of the arrangement size in \mathbb{R} being n^2 , while in \mathbb{R}^2 it is n^4 . Hence, the running time is reduced by a factor of n^2 .

► **Theorem 1.** *There exists an algorithm to compute the continuous Fréchet distance under translation between two time series of complexity n in time $\mathcal{O}(n^{8/3} \log^3 n)$.*

► **Theorem 2.** *There exists an algorithm to compute the continuous Fréchet distance under scaling between two time series of complexity n in time $\mathcal{O}(n^{8/3} \log^3 n)$.*

These results are made possible by the introduction of a novel framework for studying (continuous) time series [6] and the offline dynamic grid reachability data structure of [11]. Our approach essentially reduces (in an algorithmic sense) the continuous problems to their discrete counterparts and hence surprisingly matches the results for the two distance measures in the discrete setting. We note that in the Euclidean plane, there still is a large gap between the discrete and continuous setting and, while there was significant progress on the discrete side, there was no progress for the continuous Fréchet distance under translation or scaling. We believe that our approach is also of independent interest and can potentially be applied to other continuous 1-dimensional Fréchet distance problems, to match the running times of the discrete setting.

The asymmetric case. We remark that the offline dynamic grid reachability data structure of [11] is only described for the symmetric case in which both curves have the same complexity. While we believe that this restriction can be eliminated, it would drastically complicate the exposition of our results, which is why we also make this assumption in the statements of our results. Otherwise this restriction would not be necessary.

1.2 Technical Overview

On a high level, our approach for the decision problem for a given $\delta > 0$ for both the translation and scaling invariant Fréchet distance can be summarized in the following three steps.

1. The curves are represented by their (slightly adapted) δ -signatures [22], which are simplifications at the given length-scale δ and encode the large scale behavior of the curve in terms of a subset of selected vertices that essentially form a discrete curve. Curves naturally decompose into three parts, the beginning, middle, and end, with the middle part given by the δ -signature. The task is to design and choose data structures for all parts in such a way that they can be efficiently updated and combined when the transformation changes.
2. To process the beginning and end of the curves, we introduce another simplification, and subsequently identify the first point in the beginning of one curve that matches the entire beginning of the other, and likewise for the end parts, where we instead identify the last point that allows a matching. The crux is that we need to be able to do this efficiently for different transformations. We achieve this by introducing the structural notion of deadlocks.
3. We identify and precompute a set of $\mathcal{O}(n^2)$ representative transformations at which the answer to the decision problem is subject to change, and for which the answer can be updated efficiently as we sweep over them. Since each part of the curves essentially behaves like a discrete curve, the decision problem ultimately reduces to a reachability question for an object very similar to the well-known free space matrix used in the decision problem for the discrete Fréchet distance. Similar to the best algorithm for the discrete Fréchet distance under translation, we then make use of the offline dynamic grid reachability data structure [11] to implement updates and queries efficiently.

A note on parametric search. As is common, in this work we mostly describe how we solve the decision version, to then subsequently apply parametric search [18]. We note that the algorithm for the discrete Fréchet distance under translation on time series of [25] cleverly avoids parametric search, and hence shaves a logarithmic factor that would otherwise appear when directly applying [4]. One reason why we cannot take a similar approach – ignoring the fact that this technique is used for the discrete setting – is that the avoidance of parametric search in [25] leads to updates that crucially depend on the results of the decider calls; the updates are therefore *online*. Hence, a combination of [25] with the *offline dynamic* data structure of [11] seems infeasible.

2 Preliminaries

We denote with $[n]$ the set $\{1, \dots, n\}$. For any two points $p_1, p_2 \in \mathbb{R}$, we denote with $\overline{p_1 p_2}$ the directed line segment from p_1 to p_2 . A *time series* of complexity n is a 1-dimensional curve formed by n points $P(1), \dots, P(n) \in \mathbb{R}$ and the ordered line segments $\overline{P(i)P(i+1)}$ for $i = 1, \dots, n-1$. Such a time series can be viewed as a function $P : [1, n] \rightarrow \mathbb{R}$, where $P(i + \alpha) = (1 - \alpha)P(i) + \alpha P(i + 1)$ for $i = 1, \dots, n-1$ and $\alpha \in [0, 1]$. Sometimes, we denote P as $\langle P(1), P(2), \dots, P(n) \rangle$. The points $P(i)$ are called *vertices* of P and $\overline{P(i)P(i+1)}$ are called *edges* of P for $i = 1, \dots, n-1$. For $1 \leq s \leq t \leq n$, we denote by $P[s, t]$ the subcurve of P obtained from restricting the domain to the interval $[s, t]$. Further, we denote with $B(P, \delta)$ the set of points having distance at most δ to a point of P , i.e., $B(P, \delta) := \{x \mid \exists a \in [1, n] \text{ s.t. } |x - P(a)| \leq \delta\}$. We define $\text{im}(P)$ to be the image of the time series P , meaning that $\text{im}(P) := \{P(a) \mid a \in [1, n]\}$.

Degenerate point configurations. Our results do not depend on the assumption that the time series are in general position, as all such occurrences can be easily dealt with using arbitrary tie-breaks, or *symbolic perturbation*, where necessary. Every such instance is a result of degenerate vertex-vertex distances, and we discuss more concretely how they are dealt with when they arise.

To define the Fréchet distance, let P and Q be two time series of complexity n and m . Further, let \mathcal{H}_n be the set of all continuous non-decreasing functions $h : [0, 1] \rightarrow [1, n]$ with $h(0) = 1$ and $h(1) = n$. The *continuous Fréchet distance* is defined as

$$d_F(P, Q) = \min_{h_P \in \mathcal{H}_n, h_Q \in \mathcal{H}_m} \max_{\alpha \in [0, 1]} |P(h_P(\alpha)) - Q(h_Q(\alpha))|.$$


In this paper, we aim to determine the optimal translation or scaling of a time series that minimizes the Fréchet distance. For a time series $P = \langle P(1), \dots, P(n) \rangle$ and a value $t \in \mathbb{R}$, the *translated time series* P_t is $\langle P(1) + t, \dots, P(n) + t \rangle$ and the *Fréchet distance under translation* is defined as

$$d_F^T(P, Q) = \min_{t \in \mathbb{R}} d_F(P, Q_t).$$

For a value $s \in \mathbb{R}$, we define the *scaled time series* sP of P to be $sP = \langle sP(1), \dots, sP(n) \rangle$ and the *Fréchet distance under scaling* is defined as

$$d_F^S(P, Q) = \min_{s \in \mathbb{R}_{\geq 0}} d_F(P, sQ).$$

In contrast to the Fréchet distance under translation, the Fréchet distance under scaling is not symmetric. The above is the directed version, and one obtains the undirected version by minimizing over both directed variants. Whenever we just refer to the Fréchet distance, we mean the continuous Fréchet distance.



■ **Figure 1** Throughout this paper, vertices of time series are drawn as vertical segments for clarity. The red vertices of the time series P are its δ -signature vertices. After linearly interpolating those red vertices, we get the extended δ -signature of P .



■ **Figure 2** The left (resp. right) time series is δ -monotone increasing (resp. decreasing).

2.1 Signatures and Coupled Visiting Orders

A δ -signature of a time series P is a time series that consists of important maxima and minima of P . The concept of signatures was introduced by Driemel, Krivošija and Sohler [22]. Later, Blank and Driemel [6] extended the δ -signature by adding one vertex in the beginning and one in the end to obtain an extended δ -signature. In this paper, we only use extended δ -signatures. See Figure 1 for an example. Hence, whenever we say signature we mean extended signature. To define extended δ -signatures, we use the notion of δ -monotone time series. Examples of δ -monotone time series can be found in Figure 2.

► **Definition 3.** A time series $P : [1, n] \rightarrow \mathbb{R}$ is δ -monotone increasing (resp. decreasing) if it holds that $P(s) \leq P(s') + \delta$ (resp. $P(s) \geq P(s') - \delta$) for all $s < s' \in [1, n]$.


► **Definition 4 (extended δ -signature).** Let $P = \langle P(1), \dots, P(n) \rangle$ be a time series and $\delta \geq 0$. Then, an extended δ -signature $P' = \langle P(i_1), \dots, P(i_t) \rangle$ with $1 = i_1 \leq i_2 < \dots < i_{t-1} \leq i_t = n$ of P is a time series with the following properties:

- (a) (non-degenerate) For $k = 2, \dots, t - 1$, it holds that $P(i_k) \notin \overline{P(i_{k-1})P(i_{k+1})}$.
- (b) (2δ -monotone) For $k = 1, \dots, t - 1$, $P[i_k, i_{k+1}]$ is 2δ -monotone increasing or decreasing.
- (c) (minimum edge length) If $t > 4$, then for $k = 2, \dots, t - 2$, $|P(i_k) - P(i_{k+1})| > 2\delta$.
- (d) (range)
 - For $k = 2, \dots, t - 2$, it holds that $\text{im}(P[i_k, i_{k+1}]) = \overline{P(i_k)P(i_{k+1})}$, and
 - $\text{im}(P[1, i_2]) \subset [P(i_2), P(i_2) + 2\delta]$ or $\text{im}(P[1, i_2]) \subset [P(i_2) - 2\delta, P(i_2)]$, and
 - $\text{im}(P[i_{t-1}, n]) \subset [P(i_{t-1}), P(i_{t-1}) + 2\delta]$ or $\text{im}(P[i_{t-1}, n]) \subset [P(i_{t-1}) - 2\delta, P(i_{t-1})]$.

The vertices of P' are called δ -signature vertices of P .

Driemel, Krivošija and Sohler [22] showed that there always exists a δ -signature and that it can be computed in $\mathcal{O}(n)$ time.³ Bringmann, Driemel, Nusser and Psarros [9] defined δ -visiting orders and Blank and Driemel [6] used this concept to define coupled δ -visiting orders (see Figure 3), which can be used to decide whether the Fréchet distance between two time series is at most δ .

³ They proved this statement for δ -signatures, but it can easily be shown for extended δ -signatures as well.



■ **Figure 3** Visualization of a coupled δ -visiting order. Edges are drawn between vertices $P(i)$ and $Q(j)$, when one is a δ -signature vertex and $|P(i) - Q(j)| \leq \delta$. Then, a coupled δ -visiting order consists of a subset of the drawn edges, shown as an orange bipartite graph, where no two edges cross and it contains one incident edge for every δ -signature vertex.

► **Definition 5** (coupled δ -visiting order). Consider two time series $P = \langle P(1), \dots, P(n) \rangle$ and $Q = \langle Q(1), \dots, Q(m) \rangle$. Then, a δ -visiting order of Q on P for the δ -signature vertices $Q(j_1), \dots, Q(j_{s_Q})$ is a sequence of indices $x_1 \leq \dots \leq x_{s_Q}$ such that $|P(x_k) - Q(j_k)| \leq \delta$ for all k . We similarly define a δ -visiting order $y_1 \leq \dots \leq y_{s_P}$ of P on Q for the δ -signature vertices $P(i_1), \dots, P(i_{s_P})$ of P . These two δ -visiting orders are said to be crossing-free if there exists no k, l such that $i_k < x_l$ and $j_l < y_k$, or $x_k < i_l$ and $y_l < j_k$. In this case, the ordered sequence containing all tuples (x_k, j_k) and (i_l, y_l) , where $k = 1, \dots, s_Q$ and $l = 1, \dots, s_P$, is called coupled δ -visiting order.

► **Lemma 6** (Lemma 9 of [6]). Let P, Q be two time series and $\langle P(i_1), \dots, P(i_{s_P}) \rangle$ and $\langle Q(j_1), \dots, Q(j_{s_Q}) \rangle$ their δ -signatures. Then, $d_F(P, Q) \leq \delta$ if and only if there exist a coupled δ -visiting order $((v_1, w_1), (v_2, w_2), \dots, (v_t, w_t))$ of P and Q such that

- (i) if $v_2 = i_2$ (resp. $w_2 = j_2$), then $\exists w^* \in [w_2, \min\{w_2 + 1, j_2\}]$ (resp. $v^* \in [v_2, \min\{v_2 + 1, i_2\}]$) such that $d_F(P[1, v_2], Q[1, w^*]) \leq \delta$ (resp. $d_F(P[1, v^*], Q[1, w_2]) \leq \delta$), and
- (ii) if $v_{t-1} = i_{s_P-1}$ (resp. $w_{t-1} = j_{s_Q-1}$), then $\exists w^{**} \in [\max\{j_{s_Q-1}, w_{t-1} - 1\}, w_{t-1}]$ (resp. $v^{**} \in [\max\{i_{s_P-1}, v_{t-1} - 1\}, v_{t-1}]$) such that $d_F(P[v_{t-1}, v_t], Q[w^{**}, w_t]) \leq \delta$ (resp. $d_F(P[v^{**}, v_t], Q[w_{t-1}, w_t]) \leq \delta$).


There is a small difference to the cited lemma, which we discuss in the full version.

2.2 Grid Reachability

► **Definition 7** (Offline Dynamic Grid Reachability). Let G be the directed $n \times n$ -grid graph in which every node is either activated or deactivated. We are given updates u_1, \dots, u_U , which are of the form “activate node (i, j) ” or “deactivate node (i, j) ” in an offline manner. The task of offline dynamic grid reachability is to compute for all $1 \leq \ell \leq U$ if (n, n) can be reached by $(1, 1)$ after updates u_1, \dots, u_ℓ are performed.

Our main results depend on the following theorem.

► **Theorem 8** (Theorem 3.1 of [11]). Offline Dynamic Grid Reachability can be solved in time $\mathcal{O}(n^2 + Un^{2/3} \log^2 n)$.



■ **Figure 4** Minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$, w_{pre} , and minimal $\text{suf}(P)$ -matcher on $\text{suf}(Q)$, w_{suf} .

3 The Static Algorithm

First, we discuss how the *modified free-space matrix* can be used to decide whether the Fréchet distance between two time series is at most a given threshold δ . Lemma 6 implies that, except for the beginning and the end, it is enough to look at pairs $(P(i), Q(j))$ of vertices, where one of the two vertices is a δ -signature vertex. Missing proofs for this section can be found in the full version.

► **Definition 9** (prefix and suffix). *Let P be a time series of complexity n and let $P(i_2)$ denote the second δ -signature vertex of P , and $P(i_{s_P-1})$ denote the penultimate δ -signature vertex of P . Then $P[1, i_2]$ is the prefix $\text{pre}(P)$ and $P[i_{s_P-1}, n]$ in reverse order defines the suffix $\text{suf}(P)$ of P . That is, the time series $P[i_{s_P-1}, n]$ is $\widehat{\text{suf}(P)}$, where \widehat{P} denotes the time series defined by the vertices of P in reverse order.*

Observe that by the definition of the δ -signature vertices the time series $\text{pre}(P)$ and $\text{suf}(P)$ are each contained in some ball of radius δ .

► **Definition 10** (minimal matcher). *Let P and Q be two time series of complexity n and m respectively. The minimal P -matcher on Q is the smallest $w \in [m]$ such that*

1. $|P(n) - Q(w)| \leq \delta$, and
 2. *there is a $w^* \in [w, \min(w + 1, m)]$ such that $d_F(P, Q[1, w^*]) \leq \delta$.*
- If no such value exists, we set it to ∞ .*

We denote by w_{pre} the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ and by v_{pre} the minimal $\text{pre}(Q)$ -matcher on $\text{pre}(P)$. For the suffix, we denote by w_{suf} the index of the vertex in Q corresponding to the minimal $\text{suf}(P)$ -matcher on $\text{suf}(Q)$ (and similarly v_{suf} when the roles of P and Q are swapped). See Figure 4 for an example.

Due to the next observation, we will only discuss how to process the prefix in Section 3.1, Section 4.2 and Section 5.

► **Observation 11.** *Let \widehat{w}_{pre} be the minimal $\text{pre}(\widehat{P})$ -matcher on $\text{pre}(\widehat{Q})$. If m is the complexity of Q , then $w_{\text{suf}} = m - (\widehat{w}_{\text{pre}} - 1)$.*

With these notions, we define the modified free-space matrix. See Figure 5 for an example.

► **Definition 12** (Modified Free-Space Matrix). *Let P and Q be two time series of complexity n and m . Further, let $\langle P(i_1), \dots, P(i_{s_P}) \rangle$ and $\langle Q(j_1), \dots, Q(j_{s_Q}) \rangle$ be their δ -signatures. We construct a matrix $M_\delta \in \{0, 1\}^{n \times m}$, where the entry $M_\delta(i, j) = 1$ if*

- a) $P(i)$ and $Q(j)$ are both not δ -signature vertices, or
- b) $|P(i) - Q(j)| \leq \delta$ and $(i, j) \notin [1, i_2] \times [1, j_2] \cup [i_{s_P-1}, n] \times [j_{s_Q-1}, m]$, or
- c) $|P(i) - Q(j)| \leq \delta$ and $(i, j) \in \{(1, 1), (n, m)\}$, or
- d) $|P(i) - Q(j)| \leq \delta$ and $(i, j) = (i_2, j_2)$ and $d_F(\text{pre}(P), \text{pre}(Q)) \leq \delta$, or




Figure 5 Example of a Modified Free-Space Matrix M_δ . The colored columns (resp. rows) correspond to the δ -signature vertices of P (resp. Q). The white entries are all 1 by Definition 12 a), the red entries are defined by b), the purple entries by c) and d) and e), and the yellow entries by f). The traversal drawn in blue uses only 1-entries of M_δ . Hence, $M_\delta(n, m)$ is reachable.

- e) $|P(i) - Q(j)| \leq \delta$ and $(i, j) = (i_{s_P-1}, j_{s_Q-1})$ and $d_F(\text{suf}(P), \text{suf}(Q)) \leq \delta$, or
 - f) $(i, j) \in \{(i_2, w_{\text{pre}}), (v_{\text{pre}}, j_2), (i_{s_P-1}, w_{\text{suf}}), (v_{\text{suf}}, j_{s_Q-1})\}$.
- Otherwise, the entry $M_\delta(i, j) = 0$. Further, we say that an entry $M_\delta(i, j)$ is reachable if there exists a traversal $(1, 1) = (a_1, b_1), (a_2, b_2), \dots, (a_k, b_k) = (i, j)$ such that $M_\delta(a_l, b_l) = 1$ and $(a_l, b_l) \in \{(a_{l-1} + 1, b_{l-1}), (a_{l-1}, b_{l-1} + 1), (a_{l-1} + 1, b_{l-1} + 1)\}$ for all $l = 2, \dots, k$.

The importance of the modified free-space matrix is that it can be used to answer the decision problem of whether or not $d_F(P, Q) \leq \delta$ essentially in the same way as it is done for the discrete Fréchet distance using the free-space matrix. The next lemma follows mainly from a result of [6]. For completeness, a proof is contained in the full version.

► **Lemma 13.** *It holds that $d_F(P, Q) \leq \delta$ if and only if $M_\delta(n, m)$ is reachable. This can be checked in $\mathcal{O}(nm)$ time.*

3.1 Prefix and Suffix


This section is dedicated to identifying the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$. We often use the fact that $\text{pre}(P)$ and $\text{pre}(Q)$ are each contained in some δ -ball. We will show that the vital behavior of such constricted time series can be captured by their extreme points.

- **Definition 14** (extreme point sequence). *Let P be a time series of complexity n , such that $P(n)$ is a global extremum of P and P is contained in some δ -ball. We define the extreme point sequence $1 = a_1 < a_2 < \dots < a_p = n$ to be a sequence of indices, such that*
 - (range-preserving) $\text{im}(P[1, a_{k+1}]) = \overline{P(a_k)P(a_{k+1})}$ for all k , and
 - (extreme point) $P(a_k) = \min(P[1, a_k])$ or $P(a_k) = \max(P[1, a_k])$.

See Figure 6 and 7 for an example of both an extreme point sequence and their associated preliminary assignments, defined next.

- **Definition 15** (preliminary assignment). *Let P and Q be two time series such that the image of each time series is contained in some ball of radius δ . Let $a_1 < \dots < a_p$ and $b_1 < \dots < b_q$ be extreme point sequences of time series P and Q respectively. Define the*

22:10 Transforming Dogs on the Line



■ **Figure 6** An extreme point sequence of the depicted time series P is $1 = a_1 < a_2 < \dots < a_p$ and the preliminary assignment $Y(7)$ of Q on P is 7, since $Q(a_7)$ is the first vertex on P in $B(Q(b_7), \delta)$.

preliminary assignment $X(k)$ of P on Q for every $k \leq p$ to be the smallest index such that $|P(a_k) - Q(b_{X(k)})| \leq \delta$. If no such index exists, we set $X(k) = \infty$. Let similarly $Y(l)$ be the preliminary assignment of Q on P . We say that $X(k)$ and $Y(l)$ form a deadlock, if $l < X(k)$ and $k < Y(l)$.

Before we show how the preliminary assignment can be used to find the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$, we state some of its structural properties.

► **Lemma 16.** Let P and Q be two time series such that the image of each time series is contained in some ball of radius δ . Let $a_1 < \dots < a_p$ and $b_1 < \dots < b_q$ be extreme point sequences of time series P and Q . If $X(1) = 1$, then the following holds.

- i) For any $3 \leq k \leq p$ it holds that $X(k) \geq X(k - 2)$.
- ii) It holds that $X(2k) = 1$ for all k or $X(2k - 1) = 1$ for all k .
- iii) If $k < p$ and $X(k) \neq 1$, then for any index $l \geq X(k)$ there is a $t \in [b_{l-1}, b_l]$ such that $\text{im}(P[1, a_{k+1}]) \subset B(Q(t), \delta)$.
- iv) If $X(k) > 1$ and $Y(l)$ form a deadlock, then $X(k)$ and $Y(X(k) - 1)$ form a deadlock.

The importance of deadlocks is summarized in the following pivotal lemma.

► **Lemma 17.** Let P and Q be time series such that the image of each time series is contained in some ball of radius δ . Let $a_1 < \dots < a_p$ and $b_1 < \dots < b_q$ be extreme point sequences of P and Q . For any $w^* \leq b_q$, it holds that $d_F(P[1, a_p], Q[1, w^*]) \leq \delta$ if and only if

- (i) $|P(1) - Q(1)| \leq \delta$ and $|P(a_p) - Q(w^*)| \leq \delta$,
- (ii) $\text{im}(P[1, a_p]) \subset B(Q[1, w^*], \delta)$,
- (iii) $\text{im}(Q[1, w^*]) \subset B(P[1, a_p], \delta)$, and
- (iv) $X(k)$ and $Y(l)$ do not form a deadlock for all $k = 1, \dots, p$ and $l = 1, \dots, q$.

Observe that if $X(1)$ and $Y(1)$ do not form a deadlock then already $|P(1) - Q(1)| \leq \delta$.

Proof Sketch. First assume that $d_F(P[1, a_p], Q[1, w^*]) \leq \delta$. Then Condition (i)-(iii) hold by definition. Observe that if $X(k)$ and $Y(l)$ form a deadlock, then $|P(a_k) - Q(t)| > \delta$ for all $t \leq b_l$. However, in this case any traversal of P and Q would have to match $Q(b_l)$ with some $P(s)$ with $s < a_k$, since the part of Q before $Q(b_l)$ has to be traversed before reaching $P(a_k)$ as $d_F(P[1, a_p], Q[1, w^*]) \leq \delta$. As $X(k)$ and $Y(l)$ form a deadlock, also $|Q(b_l) - P(s)| > \delta$ for all $s \leq a_k$. Therefore, any traversal induces a cost greater than δ . Thus $X(k)$ and $Y(l)$ cannot form a deadlock, implying Condition (iv).

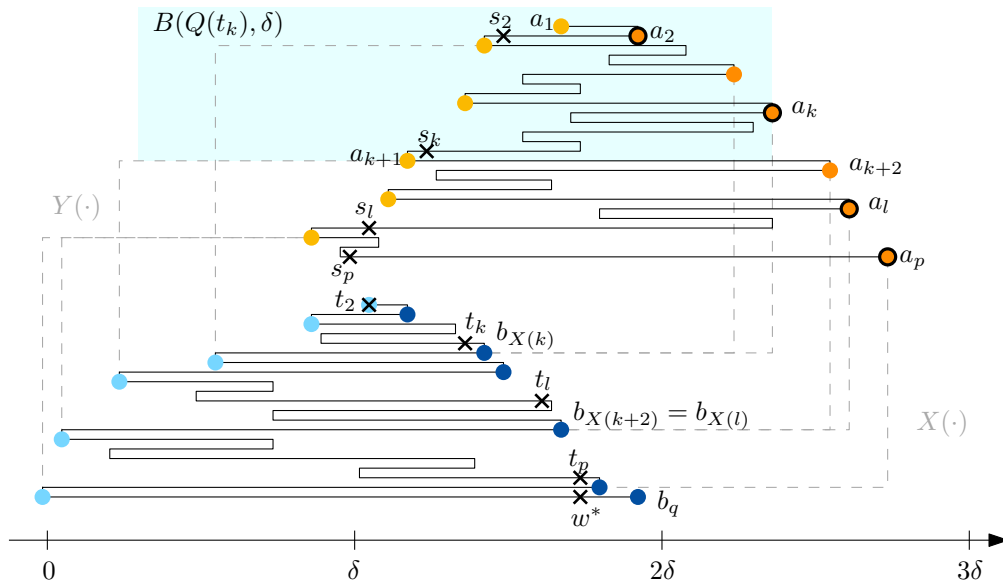


Figure 7 The gray dashed lines mark the preliminary assignments $X(\cdot)$ and $Y(\cdot)$. The points defined in the proof of Lemma 17 are marked. The circled points are the transition indices.

Now assume that Condition (i)-(iv) hold. Then, we explicitly construct a traversal of the two time series P and Q with cost $\leq \delta$. If $X \equiv 1$, then by Condition (iii) and (iv) there exists a point $s \in [a_{p-1}, a_p]$, such that $\text{im}(Q[1, w^*]) \subset B(P(s), \delta)$ and $\text{im}(P[s, a_p]) \subset B(Q(w^*), \delta)$. Then, matching $P[1, s]$ to $Q(1)$, and $Q[1, w^*]$ to $P(s)$, and $P[s, a_p]$ to $Q(w^*)$ results in a matching with cost $\leq \delta$. Otherwise, $X \not\equiv 1$. Let $\mathcal{K} = \{k \leq p-2 \mid X(k) \neq X(k+2)\} \cup \{p-1 \leq k \mid X(k) \neq 1\}$ be the set of transition indices, at which $X(k)$ is about to change when ignoring all 1-entries. To construct a matching between $P[1, a_p]$ and $Q[1, w^*]$, we define values $s_k \in [1, a_p]$ and $t_k \in [1, w^*]$ for all transition indices $k \in \mathcal{K}$ with the following properties. See Figure 7 and 8. The existence of those values is shown in the full version.

- a) There is a $t_k \in [b_{X(k)-1}, \min\{w^*, b_{X(k)}\}]$ such that $\text{im}(P[1, a_{k+1}]) \subset B(Q(t_k), \delta)$, with $a_{p+1} = a_p$.
- b) If $k < p-1$, then there is an $s_k \in [a_k, a_{k+1}]$ such that $\text{im}(Q[t_k, t_l]) \subset B(P(s_k), \delta)$, where l is the transition index after k .
- c) If $k \geq p-1$, then there is an $s_k \in [a_{p-1}, a_p]$ such that $\text{im}(Q[t_k, w^*]) \subset B(P(s_k), \delta)$ and $\text{im}(P[s_k, a_p]) \subset B(Q(w^*), \delta)$.


Utilizing a)-d), we construct a traversal of P and Q as follows: Let k be the smallest transition index. Match $P[1, s_k]$ to $Q(t_k)$ with $t_k = 1$ via a). Now let k be the last handled transition index and let l be the next transition index after k . First match $Q[t_k, t_l]$ to $P(s_k)$ via b) (lightblue in Figure 8) and then match $P[s_k, s_l]$ to $Q(t_l)$ via a) (yellow in Figure 8). Lastly, let k be the last transition index. So far, the constructed traversal matches $Q[1, t_k]$ to $P[1, s_k]$. We now match $Q(t_k, w^*)$ to $P(s_k)$ and lastly $P[s_k, a_p]$ to $Q(w^*)$ via d) (darkblue and orange in Figure 8). This yields a traversal with cost $\leq \delta$, which concludes the proof. ◀

We use the next lemma to decide whether a minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ exists.

► **Lemma 18.** Let P and Q be two time series and $P(i_2)$ be the second δ -signature vertex of P and $Q(j_2)$ of Q . There exists a minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ if and only if $X(k) < \infty$, and $X(k)$ and $Y(l)$ do not form a deadlock for all $k \in [p]$ and $l \in [q]$. If it exists, it is

$$\min\{w = 1, \dots, j_2 \mid |P(i_2) - Q(w)| \leq \delta, \text{im}(\text{pre}(P)) \subset B(Q([1, w + 1], \delta))\}.$$

22:12 Transforming Dogs on the Line



■ **Figure 8** The transition indices are circled in black. The arrows link a_k and $X(k)$ as well as b_k and $Y(k)$ and for all a_k (resp. b_k) without an outgoing arrow it holds $X(k) = 1$ (resp. $Y(k) = 1$). The traversal described in the proof of Lemma 17 is visualized by the triangles.

Finally, we are ready to show how the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ can be computed.

► **Lemma 19.** *Let P and Q be time series and let $\text{im}(Q[1, j])$ be given for every $j = 1, \dots, q$. Further, assume the preliminary assignments X and Y of $\text{pre}(P)$ and $\text{pre}(Q)$ do not form a deadlock. Then we can compute the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ in $\mathcal{O}(\log n)$ time.*

► **Lemma 20.** *Let P and Q be two time series and let extreme point sequences of $\text{pre}(P)$ and $\text{pre}(Q)$ be given. If the preliminary assignments of $\text{pre}(P)$ and $\text{pre}(Q)$ do not form a deadlock, then we can decide whether $d_F(\text{pre}(P), \text{pre}(Q)) \leq \delta$ in $\mathcal{O}(1)$ time. Otherwise, it holds that $d_F(\text{pre}(P), \text{pre}(Q)) > \delta$.*

The results of this section show how, using deadlocks, we can fill in the modified free-space matrix. In the following, we show that deadlocks allow efficient processing of updates of the transformations, which we crucially exploit for our algorithms.

4 1D Continuous Fréchet Distance Under Translation

Missing proofs for the statements in this section can be found in the full version. We first describe the reduction of the 1D continuous Fréchet distance under translation to an offline dynamic grid graph reachability problem. Subsequently we describe how we solve this problem for two time series P and Q both of complexity n in time $\mathcal{O}(n^{8/3} \log^3 n)$.

4.1 The $\mathcal{O}(n^2)$ Events

In this section, we show how to compute a set of $\mathcal{O}(n^2)$ translation values such that the Fréchet distance under translation between P and Q is at most δ if and only if there exists a translation value t in this set such that the Fréchet distance between P and Q_t is at most δ . We start by observing that it is sufficient to consider vertex-vertex events.

► **Lemma 21.** *Let $M(P, Q) = (m_{i,j})_{(i,j) \in [n] \times [n]}$ be the matrix where $m_{i,j} = \mathbb{1}_{|P(i) - Q(j)| \leq \delta}$. Let $t, t' \in \mathbb{R}$ be given. If $M(P, Q_t) = M(P, Q_{t'})$, then $d_F(P, Q_t) \leq \delta \iff d_F(P, Q_{t'}) \leq \delta$.*

► **Observation 22.** *Let P and Q be two time series of complexity n . There are at most $\mathcal{O}(n^2)$ many values t such that there are values $i, j \in [n]$ such that $|P(i) - Q_t(j)| = \delta$. These values are precisely the interval boundaries of intervals of the form $[P(i) - Q(j) - \delta, P(i) - Q(j) + \delta]$.*

We briefly discuss how to deal with degenerate cases, that is, cases where it is not true that for every $t \in \mathbb{R}$ exactly two values $i \neq j \in [n]$ exist such that $|P(i) - Q_t(j)| = \delta$, in the full version.

► **Corollary 23** (translation representatives). *There exist a sorted set $\mathfrak{T} \subset \mathbb{R}$ containing $\mathcal{O}(n^2)$ points and computable in $\mathcal{O}(n^2 \log n)$ time with the following properties.*

- *It holds that $d_F^T(P, Q) \leq \delta$ if and only if $\exists t \in \mathfrak{T}$ such that $d_F(P, Q_t) \leq \delta$.*
- *For two consecutive t, t' in \mathfrak{T} , there exist only one pair of indices $k, l \in [n]$ such that $|P(k) - Q_t(l)| \leq \delta$ and $|P(k) - Q_{t'}(l)| > \delta$, or the other way round. Further, for the set of all pairs of consecutive $t, t' \in \mathfrak{T}$, those indices can be computed in $\mathcal{O}(n^2 \log n)$ time.*

We call the set \mathfrak{T} of Corollary 23 $\mathcal{O}(n^2)$ points the *translation representatives*.

4.2 Prefix and Suffix under Translation

In Section 3.1, we showed that the preliminary assignment can be used to determine the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$. In particular, the minimal $\text{pre}(P)$ -matcher on $\text{pre}(Q)$ can be computed in $\mathcal{O}(\log n)$ time, if we know whether there exists a deadlock or not. Therefore, we define the translated preliminary assignment and show how to keep track of whether a deadlock exists while translating the time series Q with the values in \mathfrak{T} .

► **Definition 24** (translated preliminary assignment). *Let P and Q be time series and let $t \in \mathbb{R}$. Define the translated preliminary assignment X_t to be the preliminary assignment of $\text{pre}(P)$ on $\text{pre}(Q_t)$. Let similarly Y_t be the preliminary assignment of $\text{pre}(Q_t)$ on $\text{pre}(P)$. We say that $X_t(k)$ and $Y_t(l)$ form a deadlock, if $l < X_t(k)$ and $k < Y_t(l)$.*

Given two time series P and Q together with extreme point sequences a_1, \dots, a_p and b_1, \dots, b_q of $\text{pre}(P)$ and $\text{pre}(Q)$, we show that we can compute the set $\mathcal{T} \subset \mathfrak{T}$ of translations for which X_t and Y_t do not form a deadlock in total time $\mathcal{O}(n^2 \log n)$.

► **Observation 25.** *For every $t \in \mathbb{R}$, the extreme point sequences of Q_t and Q coincide.*

► **Observation 26.** $\mathcal{T} \subset \{t \in \mathbb{R} \mid X_t(1) = 1\} = \{t \in \mathbb{R} \mid Y_t(1) = 1\}$.

The next observation and next lemma show that the preliminary assignments of two consecutive translations of the translation representatives differ only slightly.

► **Observation 27.** *Let two consecutive translation representatives t, t' be given. Then, X_t and $X_{t'}$ (resp. Y_t and $Y_{t'}$) agree everywhere except at the index that participates in the point-point distance inducing the arrangement boundary separating t and t' .*

► **Lemma 28.** *Let two consecutive translation representatives t, t' be given. Let k be the index of P and l the index of Q participating in the point-point distance inducing the arrangement boundary separating t and t' . Then, it holds that*

1. $|Y_t(l) - Y_{t'}(l)| \leq 2$,
2. $Y_t(l) = \infty$ and $Y_{t'}(l) \in \{p-1, p\}$, or
3. $Y_{t'}(l) = \infty$ and $Y_t(l) \in \{p-1, p\}$.

Similarly, it holds that

1. $|X_t(k) - X_{t'}(k)| \leq 2$,
2. $X_t(k) = \infty$ and $X_{t'}(k) \in \{q-1, q\}$, or
3. $X_{t'}(k) = \infty$ and $X_t(k) \in \{q-1, q\}$.

► **Lemma 29.** *There is an algorithm that correctly computes the set \mathcal{T} of translations $t \in \mathfrak{T}$ for which X_t and Y_t do not form a deadlock in $\mathcal{O}(n^2 \log n)$. This is the set of translations in \mathfrak{T} for which the preliminary assignments of $\text{pre}(P)$ and $\text{pre}(Q_t)$ do not form a deadlock. Similarly, we can compute the set of translations in \mathfrak{T} for which $\text{suf}(P)$ and $\text{suf}(Q_t)$ do not form a deadlock in $\mathcal{O}(n^2 \log n)$ time.*

4.3 Computing the Fréchet Distance under Translation

We are now ready to prove our result for the decision problem of the Fréchet distance under translation of time series.

► **Lemma 30.** *We can decide the Fréchet distance under translation between two time series in time $\mathcal{O}(n^{8/3} \log^2 n)$.*

Proof. The goal is to apply Theorem 8 to maintain reachability in the modified free-space matrix defined in Definition 12 for all translations in \mathfrak{T} from Corollary 23.

In a pre-processing step, we first compute the signatures of both input time series and extreme point sequences of $\text{pre}(P), \text{pre}(Q), \text{suf}(P), \text{suf}(Q)$, all in time $\mathcal{O}(n)$. We use Corollary 23 to compute the set \mathfrak{T} of candidate translations that we have to check reachability for, and fix the order in which we check them (e.g., from left to right). Additionally, we invoke Lemma 29 to compute the set $\mathcal{T} \subset \mathfrak{T}$ for which there are no deadlocks of the prefixes and suffixes. The above steps take time $\mathcal{O}(n^2 \log n)$ each. It now suffices to show that if we consider the translations \mathfrak{T} in order, there are only $\mathcal{O}(1)$ many updates per translation to the modified free-space matrix of Definition 12, for each type of entry (a) to (f). Furthermore, we need to determine all of these updates in time $\mathcal{O}(n^2 \log n)$.

Combining Corollary 23 and the information from the pre-computed signatures, we can pre-compute the $\mathcal{O}(1)$ many updates per translation in \mathfrak{T} of type (a),(b), and (c) in time $\mathcal{O}(n^2 \log n)$. For entries of type (d) and (e), we first check whether the translation at hand incurs a deadlock or not (i.e., we check for containment in the pre-computed set \mathcal{T}). If it does not incur a deadlock, then we use Lemma 20 to determine in constant time whether $d_F(\text{pre}(P), \text{pre}(Q)) \leq \delta$ and $d_F(\text{suf}(P), \text{suf}(Q)) \leq \delta$. For type (f), we use the pre-computed signatures and Lemma 19 to compute $w_{\text{pre}}, v_{\text{pre}}, w_{\text{suf}}, v_{\text{suf}}$ in $\mathcal{O}(\log n)$ time.

Note that by the above, each entry type incurs at most $\mathcal{O}(1)$ updates per translation, and we therefore have $\mathcal{O}(n^2)$ updates in total. All of these updates can be pre-determined, i.e., they are offline. Consequently, we can apply Theorem 8 and obtain a running time of $\mathcal{O}(n^2 \cdot n^{2/3} \log^2 n) = \mathcal{O}(n^{8/3} \log^2 n)$. ◀

Using parametric search, the standard technique for solving the optimization problem of computing the Fréchet distance from the solution to the decision problem, we obtain the following theorem, proved in more detail in the full version.

► **Theorem 1.** *There exists an algorithm to compute the continuous Fréchet distance under translation between two time series of complexity n in time $\mathcal{O}(n^{8/3} \log^3 n)$.*

5 1D Continuous Fréchet Distance Under Scaling

Our algorithm for the Fréchet distance under scaling works in a similar way as our algorithm for translations. An additional obstacle to overcome is that scaling a time series can change the δ -signature. The next lemma shows that the inclusion of δ -signature vertices is monotonously increasing in the scaling. We can prove it using the observation that, if Q' is the δ -signature of Q , then sQ' is the δ -signature of sQ and previous results about signatures in [22].

► **Lemma 31.** *We can compute in $\mathcal{O}(n \log n)$ time all values s_j with $j = 1, \dots, n$ such that the following holds. The vertex $sQ(j)$ is a δ -signature vertex of sQ if and only if $s > s_j$.*

As a result, by adding the events where the δ -signature changes to the arrangement for the scaling parameters, the decision problem can be solved similarly to the translation case, yielding the following theorem. See the full version for details.

► **Theorem 2.** *There exists an algorithm to compute the continuous Fréchet distance under scaling between two time series of complexity n in time $\mathcal{O}(n^{8/3} \log^3 n)$.*

References

- 1 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5(1–2):75–91, 1995. doi:10.1142/S0218195995000064.
- 2 Helmut Alt, Christian Knauer, and Carola Wenk. Matching polygonal curves with respect to the Fréchet distance. In Afonso Ferreira and Horst Reichel, editors, *STACS 2001, 18th Annual Symposium on Theoretical Aspects of Computer Science, Dresden, Germany, February 15–17, 2001, Proceedings*, volume 2010 of *Lecture Notes in Computer Science*, pages 63–74. Springer, 2001. doi:10.1007/3-540-44693-1_6.
- 3 Helmut Alt, Christian Knauer, and Carola Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2004. doi:10.1007/S00453-003-1042-5.
- 4 Rinat Ben Avraham, Haim Kaplan, and Micha Sharir. A faster algorithm for the discrete Fréchet distance under translation. *CoRR*, abs/1501.03724, 2015. arXiv:1501.03724.
- 5 Abdullah Biran and Aleksandar Jeremic. Ecg bio-identification using Fréchet classifiers: A proposed methodology based on modeling the dynamic change of the ecg features. *Biomedical Signal Processing and Control*, 82:104575, 2023. doi:10.1016/j.bspc.2023.104575.
- 6 Lotte Blank and Anne Driemel. A faster algorithm for the Fréchet distance in 1d for the imbalanced case. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, September 2–4, 2024, Royal Holloway, London, United Kingdom*, volume 308 of *LIPICs*, pages 28:1–28:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.ESA.2024.28.
- 7 Sotiris Brakatsoulas, Dieter Pfoser, Randall Salas, and Carola Wenk. On map-matching vehicle tracking data. In *Proc. 31st International Conf. Very Large Data Bases (VLDB’05)*, pages 853–864, 2005. URL: <http://www.vldb.org/archives/website/2005/program/paper/fri/p853-brakatsoulas.pdf>.
- 8 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Proc. 55th Ann. IEEE Symposium on Foundations of Computer Science (FOCS’14)*, pages 661–670, 2014. doi:10.1109/FOCS.2014.76.
- 9 Karl Bringmann, Anne Driemel, André Nusser, and Ioannis Psarros. Tight bounds for approximate near neighbor searching for time series under the Fréchet distance. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022*, pages 517–550. SIAM, 2022. doi:10.1137/1.9781611977073.25.
- 10 Karl Bringmann, Marvin Künnemann, and André Nusser. When Lipschitz walks your dog: Algorithm engineering of the discrete Fréchet distance under translation. In *28th Annual European Symposium on Algorithms (ESA)*, volume 173 of *LIPICs*, pages 25:1–25:17, 2020. doi:10.4230/LIPICs.ESA.2020.25.
- 11 Karl Bringmann, Marvin Künnemann, and André Nusser. Discrete Fréchet distance under translation: Conditional hardness and an improved algorithm. *ACM Trans. Algorithms*, 17(3):25:1–25:42, 2021. doi:10.1145/3460656.
- 12 Kevin Buchin, Maike Buchin, Wouter Meulemans, and Wolfgang Mulzer. Four soviets walk the dog - with an application to Alt’s conjecture. In *Proc. 25th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA’14)*, pages 1399–1413, 2014. doi:10.1137/1.9781611973402.103.
- 13 Kevin Buchin, Anne Driemel, Natasja van de L’Isle, and André Nusser. klcluster: Center-based clustering of trajectories. In Farnoush Banaei Kashani, Goce Trajcevski, Ralf Hartmut Güting, Lars Kulik, and Shawn D. Newsam, editors, *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2019, Chicago, IL, USA, November 5–8, 2019*, pages 496–499. ACM, 2019. doi:10.1145/3347146.3359111.
- 14 Kevin Buchin, André Nusser, and Sampson Wong. Computing continuous dynamic time warping of time series in polynomial time. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022, June 7–10, 2022, Berlin, Germany*, volume 224 of *LIPICs*, pages 22:1–22:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.SOCG.2022.22.

- 15 Kevin Buchin, Tim Ophelders, and Bettina Speckmann. SETH says: Weak Fréchet distance is faster, but only if it is continuous and in one dimension. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2887–2901. SIAM, 2019. doi:10.1137/1.9781611975482.179.
- 16 Siu-Wing Cheng and Haoqiang Huang. Fréchet distance in subquadratic time. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5100–5113, 2025. doi:10.1137/1.9781611978322.173.
- 17 Ian R. Cleasby, Ewan D. Wakefield, Barbara J. Morrissey, Thomas W. Bodey, Steven C. Votier, Stuart Bearhop, and Keith C. Hamer. Using time-series similarity measures to compare animal movement trajectories in ecology. *Behavioral Ecology and Sociobiology*, 73(11):151, November 2019. doi:10.1007/s00265-019-2761-1.
- 18 Richard Cole. Slowing down sorting networks to obtain faster sorting algorithms. *Journal of the ACM*, 34(1):200–208, 1987. doi:10.1145/7531.7537.
- 19 Jacobus Conradi, Anne Driemel, and Benedikt Kolbe. $(1 + \epsilon)$ -ANN Data Structure for Curves via Subspaces of Bounded Doubling Dimension. *Computing in Geometry and Topology*, 3(2):1–22, 2024. doi:10.57717/cgt.v3i2.45.
- 20 Anne Driemel and Sariel Har-Peled. Jaywalking your dog: Computing the Fréchet distance with shortcuts. *SIAM J. Comput.*, 42(5):1830–1866, 2013. doi:10.1137/120865112.
- 21 Anne Driemel, Sariel Har-Peled, and Carola Wenk. Approximating the Fréchet distance for realistic curves in near linear time. In *Proceedings of the Twenty-Sixth Annual Symposium on Computational Geometry, SoCG '10*, pages 365–374, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1810959.1811019.
- 22 Anne Driemel, Amer Krivošija, and Christian Sohler. Clustering time series under the Fréchet distance. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 766–785. SIAM, 2016. doi:10.1137/1.9781611974331.CH55.
- 23 Alon Efrat, Piotr Indyk, and Suresh Venkatasubramanian. Pattern matching for sets of segments. In *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '01)*, pages 295–304, 2001. URL: <http://dl.acm.org/citation.cfm?id=365411.365463>.
- 24 Omrit Filtser and Matthew J. Katz. Algorithms for the discrete Fréchet distance under translation. *J. Comput. Geom.*, 11(1):156–175, 2020. doi:10.20382/JOCG.V11I1A7.
- 25 Omrit Filtser and Matthew J. Katz. Variants of the discrete Fréchet distance under translation. *Journal of Computational Geometry*, 11(1):156–175, 2020.
- 26 M Maurice Fréchet. Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 22(1):1–72, 1906.
- 27 Sariel Har-Peled and Benjamin Raichel. The Fréchet distance revisited and extended. *ACM Trans. Algorithms*, 10(1):3:1–3:22, 2014. doi:10.1145/2532646.
- 28 Minghui Jiang, Ying Xu, and Binhai Zhu. Protein structure–structure alignment with discrete Fréchet distance. *J. Bioinformatics and Computational Biology*, 6(01):51–64, 2008. doi:10.1142/S0219720008003278.
- 29 Hongli Niu and Jun Wang. Volatility clustering and long memory of financial time series and financial price model. *Digital Signal Processing*, 23(2):489–498, 2013. doi:10.1016/j.dsp.2012.11.004.
- 30 Carola Wenk. *Shape matching in higher dimensions*. PhD thesis, Freie Universität Berlin, 2002. PhD Thesis.