

Middle Curves Based on Discrete Fréchet Distance ^{*}

Hee-Kap Ahn[†] Helmut Alt[‡] Maike Buchin[§] Eunjin Oh[¶] Ludmila Scharf[‡]
Carola Wenk^{||}

July 29, 2018

Abstract

Given a set of polygonal curves, we present algorithms for computing a *middle curve* that serves as a representative for the entire set of curves. We require that the middle curve consists of vertices of the input curves and that it minimizes the maximum discrete Fréchet distance to all input curves. We consider three different variants of a middle curve depending on in which order vertices of the input curves may occur on the middle curve, and provide algorithms for computing each variant.

1 Introduction

Sequential point data, such as time series and trajectories, are ever increasing due to technological advances, and the analysis of these data calls for efficient algorithms. An important analysis task is to find a “representative” or “middle” curve for a set of similar curves. For instance, this could be the route of a group of people or animals traveling together. Or it could be a representation of a handwritten letter for a class of similar handwritten letters. Such a middle curve typically provides a concise representation of the data, which is useful for data analysis and for reducing the size of the data.

Since sampled locations are more reliable than positions interpolated in between those, we seek a middle curve consisting only of sampled point locations. The middle curve should then be as close as possible to the individual curves, hence we ask for it to minimize the maximum discrete Fréchet distance d_F to any of these. The Fréchet distance [1] and the discrete Fréchet distance [6] are well-known distance measures, which have been successfully used in analyzing handwritten characters [8] and trajectories [2, 10].

For simplicity, we restrict our definitions to sets of two polygonal curves P and Q , as the generalization to $k \geq 2$ such curves is straightforward. Given P and Q , we consider polygonal curves R with vertices from $P \cup Q$ where we assume that each vertex of R uniquely corresponds

^{*}This work was partially supported by research grant AL 253/8-1 from Deutsche Forschungsgemeinschaft (German Science Association), and by the National Science Foundation under grants CCF-1301911 and CCF-1618469. Work by Ahn and Oh was supported by the MSIT (Ministry of Science and ICT), Korea, under the SW Starlab support program (IITP-2017-0-00905) supervised by the IITP (Institute for Information & communications Technology Promotion) and the NRF grant 2011-0030044 (SRC-GAIA) funded by the government of Korea.

[†]Pohang University of Science and Technology, Korea. Email: heekap@postech.ac.kr

[‡]Free University of Berlin, Germany. Email: {alt, scharf}@mi.fu-berlin.de

[§]Technical University Dortmund, Germany. Email: maike.buchin@tu-dortmund.de

[¶]Max Planck Institute for Informatics, Germany. Email: eoh@mpi-inf.mpg.de

^{||}Tulane University, United States. Email: cwenk@tulane.edu

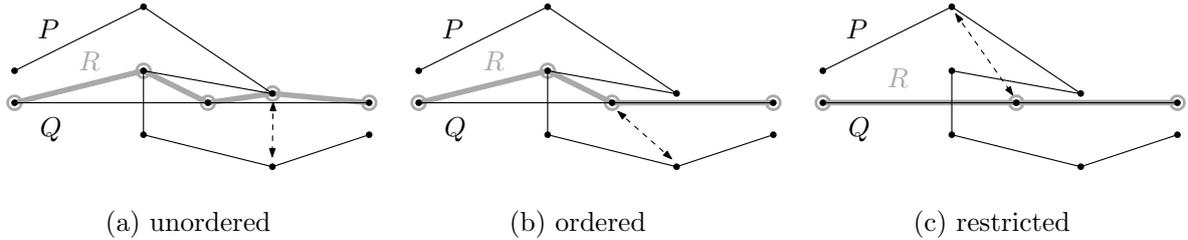


Figure 1: Illustration of the three different cases. The curve R is a middle curve for each case. The two-way arrow which points to a vertex in $P \cup Q$ and a vertex in R indicates a mapping between two vertices realizing the discrete Fréchet distance.

29 to a vertex of P or Q . That is, we identify any such vertex by its index in P or Q and not just
 30 by the geometric position of the corresponding point.

31 R is called *ordered* if any two elements of P occurring in R have the same order as in
 32 P , likewise with elements from Q . If R minimizes $\max\{d_F(R, P), d_F(R, Q)\}$, it is called a
 33 (unordered) *middle curve* of P and Q . If R is ordered and minimizes this expression for all
 34 ordered curves it is called an *ordered middle curve* of P and Q .

35 We obtain a third variant of the definition of middle curve, called *restricted middle curve*, by
 36 restricting the set of feasible matchings (see Section 2) in the definition of the discrete Fréchet
 37 distance. More precisely, R should be ordered and only matchings are allowed where elements
 38 of R are matched to their corresponding elements in P or Q . R then should be closest to P
 39 and Q among all ordered sequences with respect to this distance measure. Since vertices of R
 40 originate from P or Q , this seems a natural restriction.

41 Figure 1 illustrates the three cases we consider: unordered, ordered, and restricted middle
 42 curves. In this example, each case results in a different middle curve, and the associated distance
 43 increases as we add more restrictions. Note that from unordered to ordered we limit the middle
 44 curves we consider, whereas from ordered to restricted we limit the matchings we consider.
 45 Respecting the order of the input curves seems to be a natural requirement (ordered middle
 46 curve). Furthermore it seems intuitive that a vertex should be matched to itself on the middle
 47 curve (restricted middle curve). Among the different algorithms we present for the various cases,
 48 the most efficient algorithms are for computing an unordered middle curve in the Euclidean plane
 49 and (slightly less efficient) for computing a restricted middle curve.

50 **Related work.** The problem of finding a curve that represents a set of curves has been studied
 51 in the literature [4, 7, 9]. While there are different definitions of such a representative curve,
 52 none of them requires the representative curve to use vertices of the input curves. Buchin et
 53 al. [4] and van Kreveld et al. [9] both require the representative curve to use parts of the input
 54 edges. Buchin et al. aim for the curve to always “stay in the middle” in the sense of a median
 55 and give an $O(k^2n^2)$ -time algorithm, where k is the number of given curves and n is the number
 56 of vertices in each curve. van Kreveld et al. require the representative curve to be as close as
 57 possible to all trajectories at any time, allowing small jumps between different trajectories, and
 58 give an $O(k^3n)$ -time algorithm. Note that neither of these approaches makes use of the Fréchet
 59 distance or its variants. Using neither input vertices nor input edges, Har-Peled and Raichel [7]
 60 show that a curve minimizing the Fréchet distance to k input curves can be computed in $O(n^k)$
 61 time in the k -dimensional free space using the radius of the smallest enclosing disk as “distance”.

62 **Our Results.** We present algorithms for computing a middle curve that minimizes the discrete
 63 Fréchet distance to k input curves for $k \geq 2$ each of size at most n in three variants:

- 64 1. **Ordered case:** $O(n^{2k})$ -time algorithm for computing an ordered middle curve.
- 65 2. **Restricted case:** $O(n^k \log^k n)$ -time algorithm for computing a restricted middle curve.
- 66 3. **Unordered case:** $O(n^k \log n)$ -time algorithm for computing an unordered middle curve.

67 In the following sections, we present the algorithms for these three cases. For the ordered
 68 case and the restricted case, the algorithm works for any metric. For the unordered case, the
 69 algorithm works only for the Euclidean metric while we can modify our algorithm to work for
 70 any metric with running time of $O(n^{k+1})$.

71 We also distinguish two variants of the problem depending on whether multiple occurrences
 72 of vertices on R are allowed or not. The algorithms for the restricted and the unordered cases
 73 allow vertices to occur multiple times. In the ordered case, our algorithms can handle both
 74 variants.

75 Instead of minimizing the distance to the input curves, we may want to minimize the size of
 76 a middle curve for a fixed distance. We give algorithms for this variant as well. However in the
 77 end we discuss that middle curves may have high complexity.

78 2 Preliminaries

79 Let $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$ be two polygonal curves, represented as point se-
 80 quences in \mathbb{R}^d . And let $d(\cdot, \cdot)$ denote a metric used to measure point-wise distances. The discrete
 81 Fréchet distance, denoted by $d_F(P, Q)$, is defined as follows: A *matching*¹ is a sequence of pairs
 82 $(p, q) \in P \times Q$ such that (1) the sequence begins at (p_1, q_1) and ends at (p_n, q_m) , and (2) a pair
 83 (p_i, q_j) in the sequence is followed only by one of (p_{i+1}, q_j) , (p_i, q_{j+1}) , or (p_{i+1}, q_{j+1}) . The value
 84 of a matching is the maximum distance of p and q over all pairs (p, q) in the matching. Then
 85 $d_F(P, Q)$ is the minimum value over all possible matchings. If P and Q are empty, then we
 86 define $d_F(P, Q) = 0$, and if either P or Q is empty then $d_F(P, Q) = \infty$.

87 As mentioned in the introduction, given two polygonal curves represented as point sequences
 88 $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$, a (unordered) *middle curve* of P and Q is a sequence R
 89 of elements from $P \cup Q$ which minimizes $\max\{d_F(R, P), d_F(R, Q)\}$. Notice, that an element of
 90 R is identified by whether it originates from P or Q and by the index in that sequence. This
 91 is necessary when we consider the other variants of middle curve as defined in the introduction.
 92 We call R an *ordered middle curve* if elements of R respect the order given by the input curves
 93 and $\max\{d_F(R, P), d_F(R, Q)\}$ is the minimum among all such curves. We also consider the
 94 *restricted* variant of middle curves, where elements of R need to be matched to themselves in
 95 the curve they are taken from.

96 **2-Approximation.** A simple observation is that any of the input curves is a 2-approximate
 97 middle curve, i.e., the associated discrete Fréchet distance is at most double that of an optimal
 98 middle curve. Any input curve gives a restricted middle curve, and hence this holds for all three
 99 variants, i.e., unordered, ordered, and restricted.

100 The 2-approximation follows by the triangle inequality. In fact, let P and Q be two arbitrary
 101 input curves. For any middle curve R realizing the minimum distance, say d_{\min} , let (p_i, r_k) and
 102 (q_j, r_k) denote pairs in optimal matchings of R with P and Q , respectively. Then, we have
 103 $d(p_i, q_j) \leq d(p_i, r_k) + d(r_k, q_j) \leq 2d_{\min}$ by the triangle inequality, where $d(p, q)$ denotes the

¹Note that this is not a one-to-one matching, and for this reason has also been called a coupling.

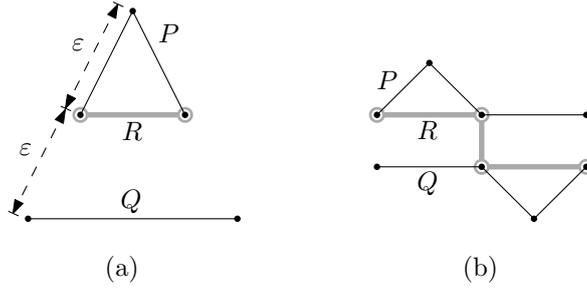


Figure 2: (a) The 2-approximation is tight. (b) The middle curve may need to consist of vertices from both curves.

104 distance between p and q of the underlying metric. As can easily be seen, this implies that there
 105 is a matching of P and Q whose distance is at most $2d_{\min}$. Thus, we have a 2-approximation in
 106 constant time (not counting the time to output the vertices of the approximate middle curve).
 107 Note that the 2-approximation is tight, as the example in Figure 2(a) shows. We observe also
 108 that for a middle curve we may need to choose a subset of vertices from both curves, as in
 109 Figure 2(b).

110 3 Algorithms for the Ordered Case

111 In this section we present a dynamic programming algorithm for computing an ordered middle
 112 curve R . We first give a decision algorithm for two input curves P and Q , and we allow the
 113 same vertex to occur at most once on R . Later we show how to generalize the algorithm to more
 114 than two input curves, to allow vertices to occur multiple times in R , and how to solve the two
 115 optimization variants.

116 3.1 Decision Problem for Two Curves

117 Let $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$ be two input curves, and let $\varepsilon \geq 0$ be given. In
 118 this section we allow the same vertex to occur at most once on R . Let P_i denote the *prefix*
 119 (p_1, \dots, p_i) of P for $1 \leq i \leq n$, and let Q_j denote the prefix (q_1, \dots, q_j) of Q for $1 \leq j \leq m$. We
 120 use P_0 and Q_0 to denote an empty subsequence of P and Q , respectively.

The dynamic programming algorithm operates with four-dimensional Boolean arrays of the form $X[i, j, k, l]$ for $0 \leq i \leq k \leq n$ and $0 \leq j \leq l \leq m$, where $X[i, j, k, l]$ is **true** if and only if there exists an ordered sequence R from points in $P_i \cup Q_j$ for a fixed $\varepsilon \geq 0$ such that

$$\max\{d_F(R, P_k), d_F(R, Q_l)\} \leq \varepsilon.$$

121 We say in this case that R *covers* P_k and Q_l . Clearly, the decision problem has a positive answer
 122 if and only if $X[i, j, n, m]$ is true for some i and j .

In order to determine the values of $X[i, j, k, l]$ with $0 \leq i \leq k \leq n$ and $0 \leq j \leq l \leq m$, we need more information, particularly, whether there is a covering sequence R in which the points p_i and q_j occur, and if they do, whether they occur in the interior or at the end of the sequence. To this end, we can represent the array X as the component-wise disjunction of seven Boolean arrays

$$X = A \vee B \vee C \vee D \vee E \vee F \vee G.$$

123 The entries of the Boolean arrays defined below, indicate whether an ordered sequence R from
 124 points in $P_i \cup Q_j, 0 \leq i \leq n, 0 \leq j \leq m$ covering P_k and Q_l exists with the following properties,
 125 respectively:

- 126 $A[i, j, k, l]$: R contains neither p_i nor q_j .
 127 $B[i, j, k, l]$: R contains p_i in its interior but does not contain q_j .
 128 $C[i, j, k, l]$: R ends in p_i but does not contain q_j .
 129 $D[i, j, k, l]$: R contains q_j in its interior but does not contain p_i .
 130 $E[i, j, k, l]$: R ends in q_j but does not contain p_i .
 131 $F[i, j, k, l]$: R contains q_j in its interior and ends in p_i .
 132 $G[i, j, k, l]$: R contains p_i in its interior and ends in q_j .

133 If $i = 0$ or $j = 0$, the described properties involve the “nonexisting points” p_0 or q_0 which should
 134 be interpreted as not being contained in any curve. Also, it might be the case that i, j, k , or l are
 135 0 which means that the corresponding sequences are empty. In general, observe that R cannot
 136 contain both p_i and q_j in its interior. See Figure 3 for an illustration of the seven different cases
 137 that can occur.

Our dynamic programming algorithm is based on the recursive identities for the Boolean arrays given in the following paragraphs. Each identity holds only if all index ranges of the arrays in the formulas are nonnegative.

$$\begin{aligned}
 A[0, 0, 0, 0] &= \text{true} \\
 A[0, 0, k, l] &= \text{false} && , \text{ for } k \geq 1 \text{ or } l \geq 1 \\
 A[i, 0, k, l] &= X[i - 1, 0, k, l] \\
 A[0, j, k, l] &= X[0, j - 1, k, l] \\
 A[i, j, k, l] &= X[i - 1, j - 1, k, l] \\
 B[i, 0, k, l] &= B[0, j, k, l] = \text{false} \\
 B[i, j, k, l] &= G[i, j - 1, k, l] \vee B[i, j - 1, k, l]
 \end{aligned}$$

138 As easily can be verified, each of these identities holds with the meaning given to the Boolean
 139 arrays previously. For example, the first two equalities for A hold because an empty curve covers
 140 an empty sequence, but not any other sequence, and the following equalities for A are straight-
 141 forward. The first line of equalities for B holds because p_i must be at the end of R if no points
 142 from Q are available, and there is no point p_0 which a middle curve could contain. In the last
 143 equality for B , the entry $G[i, j - 1, k, l]$ accounts for the case that R contains q_{j-1} (which then
 144 must be at the end), and $B[i, j - 1, k, l]$ for the case that it does not.

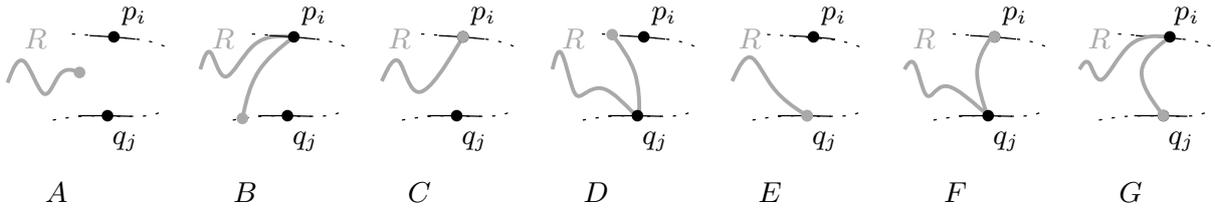


Figure 3: Illustration of cases in the dynamic programming.

In the following, let $cl(p, q) = \mathbf{true}$ if and only if $d(p, q) \leq \varepsilon$, for two points p and q . The following identities hold for C :

$$\begin{aligned} C[i, j, 0, l] &= C[i, j, k, 0] = C[0, j, k, l] = \mathbf{false} && \text{and otherwise,} \\ C[i, j, k, l] &= cl(p_i, p_k) \wedge cl(p_i, q_l) \wedge \\ & \quad (A[i, j, k-1, l-1] \vee A[i, j, k-1, l] \vee A[i, j, k, l-1] \vee \\ & \quad C[i, j, k-1, l-1] \vee C[i, j, k-1, l] \vee C[i, j, k, l-1]) \end{aligned}$$

145 The equalities in the first line hold because only an empty curve can cover an empty sequence,
 146 and because a middle curve cannot end in the nonexistent p_0 . The equality in the second line
 147 models that the final point p_i in R can cover p_k and q_l only, or it can also cover additional points
 148 that occur earlier in the sequences P_k and Q_l .

The entries of D and E can be determined analogously to the ones of B and C with the roles of p_i and q_j exchanged. The identities for F have similar explanations as the ones for C :

$$\begin{aligned} F[0, j, k, l] &= F[i, 0, k, l] = F[i, j, 0, l] = F[i, j, k, 0] = \mathbf{false} \\ F[i, j, k, l] &= cl(p_i, p_k) \wedge cl(p_i, q_l) \wedge \\ & \quad (D[i, j, k-1, l-1] \vee D[i, j, k-1, l] \vee D[i, j, k, l-1] \vee \\ & \quad E[i, j, k-1, l-1] \vee E[i, j, k-1, l] \vee E[i, j, k, l-1] \vee \\ & \quad F[i, j, k-1, l-1] \vee F[i, j, k-1, l] \vee F[i, j, k, l-1]) \end{aligned}$$

149 The entries of G can be determined analogously to the ones of F with the roles of p_i and q_j
 150 exchanged.

151 The dynamic programming algorithm takes $O(n^2m^2)$ time and space to decide if an ordered
 152 middle curve of two curves of size n and m exists for a fixed distance. While filling the dynamic
 153 programming matrices we can compute an additional pointer array $Y[i, j, k, l]$ that, for each
 154 \mathbf{true} assignment in one of the equalities, stores a pointer to one of the 4-tuples of indices on the
 155 right hand side of the equality that made the assignment \mathbf{true} . A covering sequence R can then
 156 be computed by backtracking these pointers. Note that there can be an exponential number of
 157 valid middle curves (e.g., if all points are within distance ε of each other).

158 3.2 Optimization Problems

159 We can solve the optimization problem of minimizing ε by adapting the dynamic programming
 160 approach to compute the smallest value for which a covering middle curve exists.

161 Instead of storing truth values, X stores the minimum value of the seven arrays, A to G .
 162 Array entries are initialized to $0|\infty$ instead of $\mathbf{true}|\mathbf{false}$, any $cl(p, q)$ in the formulas is replaced
 163 by the distance $d(p, q)$, \vee becomes \min , and \wedge becomes \max . The running time for computing
 164 the minimum ε for which a middle curve exists is the same as for the decision problem described
 165 in Section 3.1.

166 Similarly, if we want to minimize the size of an ordered middle curve for a fixed distance ε ,
 167 we can adapt the dynamic program to store the smallest size of a middle curve in the following
 168 way:

Again, $\mathbf{true}|\mathbf{false}$ are replaced by $0|\infty$, \vee becomes \min , and the recursion for C is modified as follows:

$$\begin{aligned} C[i, j, k, l] &= \infty && \text{if } \neg(cl(p_i, p_k) \wedge cl(p_i, q_l)) \text{ and, otherwise,} \\ C[i, j, k, l] &= \min(A[i, j, k-1, l-1] + 1, A[i, j, k-1, l] + 1, A[i, j, k, l-1] + 1, \\ & \quad C[i, j, k-1, l-1], C[i, j, k-1, l], C[i, j, k, l-1]) \end{aligned}$$

and the one for F becomes:

$$F[i, j, k, l] = \infty \quad \text{if } \neg(\text{cl}(p_i, p_k) \wedge \text{cl}(p_i, q_l)) \text{ and, otherwise,}$$

$$F[i, j, k, l] = \min(D[i, j, k-1, l-1] + 1, D[i, j, k-1, l] + 1, D[i, j, k, l-1] + 1, \\ E[i, j, k-1, l-1] + 1, E[i, j, k-1, l] + 1, E[i, j, k, l-1] + 1, \\ F[i, j, k-1, l-1], F[i, j, k-1, l], F[i, j, k, l-1])$$

169 Again, the entries of E and G can be computed analogously to the ones of C and F with the
170 roles of p_i and q_j exchanged.

171 3.3 Generalization to Multiple Curves and Multiple Vertex Occurrences

172 The decision and optimization algorithms can be generalized to k sequences P^1, \dots, P^k of
173 sizes n_1, \dots, n_k , respectively. The corresponding arrays then have $2k$ indices $[i_1, \dots, i_k, j_1, \dots, j_k]$,
174 $1 \leq i_1, j_1 \leq n_1, \dots, 1 \leq i_k, j_k \leq n_k$, describing the covering of the partial sequences (prefixes)
175 $P_{j_1}^1, \dots, P_{j_k}^k$ by a middle curve R using points from $P_{i_1}^1 \cup \dots \cup P_{i_k}^k$.

176 As in the case $k = 2$, we obtain different arrays depending on which points $p_{i_1}^1, \dots, p_{i_k}^k$ are
177 at the end, in the interior, or not contained in R . Since there are $k2^{k-1}$ possibilities to put
178 one of these points at the end and others in the interior, and 2^k possibilities to put a subset
179 of $\{p_{i_1}^1, \dots, p_{i_k}^k\}$ in the interior and not using the remaining ones, and the possibility of having
180 all of them in the interior is excluded, there must be $k2^{k-1} + 2^k - 1$ arrays reflecting all these
181 possibilities. (Which is, in fact, 7 for $k = 2$).

182 For any constant k , there is a constant number of arrays each of which has $n_1^2 \cdots n_k^2$ entries.
183 The recursive formulas to compute the entries of the arrays which, for simplicity, we do not
184 explain for the general case, can be derived in a way analogous to the ones in the case $k = 2$.
185 They have constant size for constant k , and therefore the runtime of the dynamic programming
186 algorithm is $O(n_1^2 \cdots n_k^2)$.

187 The dynamic programming algorithm can also be modified to allow multiple occurrences of
188 points on R , which requires distinguishing slightly more cases than before: Whether a point
189 appears at the end only, both at the end and in the interior, in the interior only, or not at all in
190 R . This results in $2^k(k+1) - 1$ arrays: now all k points may appear in the interior or not, any
191 of the k or no point appears at the end, but not all k points can appear in the interior only.

192 We summarize the results of this section:

193 **Theorem 1.** *For $k \geq 2$ curves of size at most n each, we can compute an ordered middle curve
194 in $O(n^{2k})$ time using $O(n^{2k})$ space. For fixed $\varepsilon > 0$, an ordered middle curve with minimum
195 complexity and distance at most ε to the input curves can be computed in the same time.*

196 4 Algorithms for the Restricted Case

197 Now we consider the case where the matchings realizing the discrete Fréchet distance are re-
198 stricted to map every vertex of R to itself in the input curve it originated from. This case allows
199 for a more efficient dynamic programming algorithm.

200 4.1 Decision Problem for Two Curves

Dynamic Programming Formulation. Let $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$ be two
input curves with $m \leq n$, and let $\varepsilon \geq 0$. We define arrays similar to the ones in Section 3. Let

$X[i, j]$ be **true** for $0 \leq i \leq n, 0 \leq j \leq m$ if and only if there exists an ordered sequence R from points in $P_i \cup Q_j$ with

$$\max\{d_F(R, P_i), d_F(R, Q_j)\} \leq \varepsilon,$$

with the restriction that there exist matchings that map every vertex of R to itself in the input curve it originated from. We say in this case that R *restrictively covers* P_i and Q_j . Clearly, the decision problem has a positive answer if and only if $X[n, m]$ is **true**.

Similar to the one in Section 3 we can write X as a disjunction of three Boolean arrays

$$X = A' \vee C' \vee E'.$$

For each array defined below, a sequence R from points in $P_i \cup Q_j$ restrictively covering P_i and Q_j exists with the following properties, respectively:

$A'[i, j]$: R ends in neither p_i nor q_j (but may contain one of them in its interior).

$C'[i, j]$: R ends in p_i (and may contain q_j in its interior).

$E'[i, j]$: R ends in q_j (and may contain p_i in its interior).

In contrast to the one in Section 3, we now only distinguish the cases by the last point of R . Hence, we only distinguish three cases. (In comparison to the ordered case, A' combines A, B, D , and C' combines C, F , and E' combines E, G). However, we will only explicitly compute X and not A', C', E' .

We compute X incrementally using dynamic programming, iterating over all (i, j) in increasing order. For fixed (i, j) we consider adding p_i or q_j to a middle curve R from $P_i \cup Q_j$, where p_i and q_j are matched to each other. That is, we consider adding say p_i to R , where it is matched to itself on P and to the point q_j on Q . We iterate over all (i, j) maintaining in the table X the coverage of P and Q by a restricted middle curve from $P_i \cup Q_j$.

For this, we initialize $X[0, 0]$ to **true** and $X[i, 0]$ and $X[0, j]$ to **false** for $i, j > 0$. Then we incrementally process all entries (i, j) for $i, j > 0$. We use the following order (which we will need later when we introduce \bar{X}): For $j = 1$, we process all (i, j) incrementally from $i = 1$ to n ; Then we increase the index j by 1 and process all (i, j) incrementally from $i = 1$ to n , and we repeat this up to $j = m$.

When processing (i, j) , we need to first check, whether p_i and q_j can be matched to each other, i.e., they have distance $\leq \varepsilon$. If they do, we then need to check if there exists a middle curve from $P_i \cup Q_j$ that covers up to here, or whose coverage can be extended up to here after adding p_i or q_j . If this is the case, we consider how the coverage extends by adding p_i or q_j or both. We maintain this information in the table X as we iterate over (i, j) . To make this more explicit, we introduce the notions of upper and lower wedges in the table X , which reflect the coverage on P and Q when adding a point p_i or q_j , which are matched to each other.

Assume p_i is matched to q_j . We use the *upper wedge* $U_P(i, j)$ to describe the resulting coverage of P and Q when adding p_i to R . Specifically, $U_P(i, j)$ denotes the set of index pairs (i', j') such that $d(p_{i''}, p_i) \leq \varepsilon$ and $d(q_{j''}, p_i) \leq \varepsilon$ for all $i \leq i'' \leq i'$ and $j \leq j'' \leq j'$. That is, $U_P(i, j)$ consists of consecutive index pairs $(i', j') \geq (i, j)$ that are covered by p_i . The *lower wedge* $L_P(i, j)$ denotes the set of index pairs (i', j') such that $d(p_{i''}, p_i) \leq \varepsilon$ and $d(q_{j''}, p_i) \leq \varepsilon$ for all $i' \leq i'' \leq i$ and $j' \leq j'' \leq j$. Furthermore, we define the *extended lower wedge* $\hat{L}_P(i, j)$ which, in addition to all index pairs in the lower wedge $L_P(i, j)$ also contains (i', j') immediately to the left or below, i.e., for which $(i' + 1, j')$, $(i', j' + 1)$, or $(i' + 1, j' + 1)$ is contained in $L_P(i, j)$. The wedges $U_Q[i, j]$, $L_Q[i, j]$, and $\hat{L}_Q[i, j]$ are defined analogously, consisting of index pairs (i', j') for which $p_{i'}$ and $q_{j'}$ are both close to q_j , that is, $cl(p_{i'}, q_j) = cl(q_{j'}, q_j) = \mathbf{true}$. Figure 4 illustrates these wedges for a pair (i, j) .

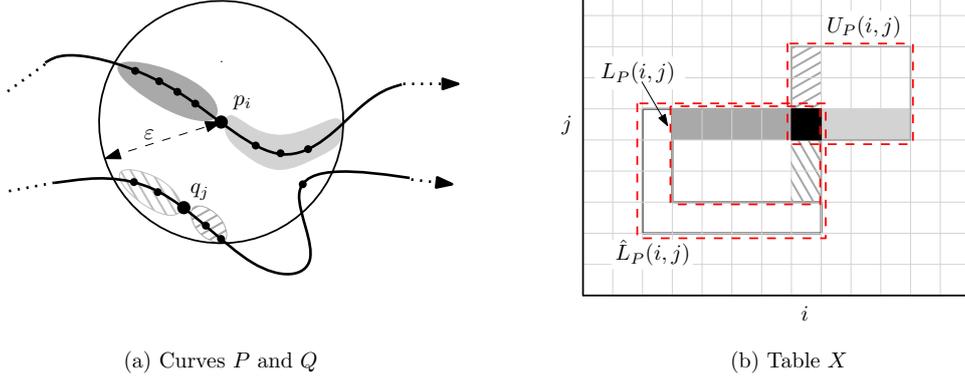
(a) Curves P and Q (b) Table X

Figure 4: (a) Points of P and Q that are at distance at most ε from p_i . (b) The upper wedge $U_P(i, j)$, the lower wedge $L_P(i, j)$, and the expanded lower edge $\hat{L}_P(i, j)$ at (i, j) on P .

Using this terminology we observe for $i, j > 0$:

$$\begin{aligned}
 A'[i, j] &= (\exists i' < i, j' \leq j : (C'[i', j'] \wedge (i, j) \in U_P(i', j'))) \vee (\exists i' \leq i, j' < j : (E'[i', j'] \wedge (i, j) \in U_Q(i', j'))) \\
 C'[i, j] &= cl(p_i, q_j) \wedge (\exists i' \leq i, j' < j : (X[i', j'] \wedge (i', j') \in \hat{L}_P(i, j))) \\
 E'[i, j] &= cl(p_i, q_j) \wedge (\exists i' < i, j' \leq j : (X[i', j'] \wedge (i', j') \in \hat{L}_Q(i, j)))
 \end{aligned}$$

241 During the dynamic programming, in order to efficiently answer queries of the form $(i, j) \in$
 242 $U_P(i', j')$ we maintain the upper envelope \bar{X} of all true elements in X . More specifically, we
 243 define $\bar{X}[i] = \max\{j \mid X[i, j] = \text{true}\}$. Using \bar{X} , the query whether a rectangle is nonempty (i.e.,
 244 contains a true point) reduces to querying whether the upper envelope intersects the rectangle.
 245 Note that X as well as \bar{X} change during the dynamic programming for increasing i and j in the
 246 specified order.

247 **Querying and Updating the Upper Envelope \bar{X} .** We store \bar{X} in an augmented balanced
 248 binary search tree sorted on i . Each leaf corresponds to an index i and stores $\bar{X}[i]$. We sometimes
 249 use an index i to denote the leaf corresponding to i . For an internal node v of the search tree,
 250 let $I(v)$ denote the set of all indices corresponding to the leaves of the subtree rooted at v . Each
 251 internal node v stores two key values $m[v]$ and $M[v]$, where $m[v]$ is the minimum of $\bar{X}[i]$ over all
 252 indices i in $I(v)$ and $M[v]$ is the maximum. Note that the node set, and hence the structure, of
 253 the tree is static. The only changes we make, is increasing the values $\bar{X}[i]$ at leaves, and hence
 254 also the augmented values at inner nodes.

255 We need the following two operations.

- 256 1. *Querying whether a rectangle intersects \bar{X} .* Given an extended lower wedge with bottom-
 257 left corner (i_B, j_B) and top-right corner (i_T, j_T) , we need to check if there is an index i
 258 such that $j_B \leq \bar{X}[i]$ and $i_B \leq i \leq i_T$.

259 This can be done as follows. Consider the search paths from the root to i_B and i_T . Let u_c
 260 be the lowest common ancestor of i_B and i_T . Whenever we descend into the right child at
 261 a node v on the path from u_c to i_T , we check the maximum key value of the left child v_L
 262 of v . The set $I(v_L)$ is contained in the interval $[i_B, i_T]$. Thus, if $M[v_L] \geq j_B$, the correct
 263 answer for the query is “yes”. Otherwise, we do not need to consider the subtree rooted
 264 at v_L further. Whenever we descend into the left child at a node v on the path to i_B , we
 265 check the answer for the query on the right child of v analogously. Hence we can answer
 266 the query while we traverse the two paths, which takes logarithmic time.

267 2. *Updating \bar{X} by adding a rectangle.* Given an upper wedge whose bottom-left corner is
 268 (i_B, j_B) and top-right corner is (i_T, j_T) , we want to add this to \bar{X} , i.e., all values in the
 269 wedge are set to **true** in X . For this, we need to update $\bar{X}[i]$ to j_T for all $i_B \leq i \leq i_T$
 270 with $\bar{X}[i] < j_T$.

271 We traverse the balanced binary search tree from the root as follows. Assume that we
 272 reach a node v . If j_T is at most $m[v]$ or $I(v)$ has no element lying in $[i_B, i_T]$, then we do
 273 not need to update the values stored in the leaves of the subtree rooted at v . Hence we
 274 do not traverse this subtree. If $m[v]$ is smaller than j_T and $I(v)$ has an element lying in
 275 $[i_B, i_T]$, then we need to search further in the subtree rooted at v . So, we move to both
 276 children of v .

277 Finally we reach some leaf, which is updated if $j_T > \bar{X}[i]$. Then we go back to the root
 278 from those leaves and update the key values for internal nodes lying on the paths. It is
 279 easy to see that the running time of the update is $O(c \log n)$, where c is the number of
 280 indices which are updated.

281 We can now formulate the decision algorithm, which first precomputes all wedges and then
 282 iteratively computes the table X .

283 **Computing all wedges.** We compute the upper wedge $U_P(i, j)$ as follows: For fixed p_i , we
 284 first find the largest $k \geq i$ such that all p_i, \dots, p_k are at distance at most ε from p_i . Then we
 285 find the largest $l \geq j$ such that all q_j, \dots, q_l are at distance at most ε from p_i . This determines
 286 the upper right corner (k, l) of $U_P(i, j)$. Note that (k, l) is also the upper right corner for all
 287 $U_P(i, j')$ for $j \leq j' \leq l$. Hence, all upper wedges $U_P(i, j)$ for a fixed i can be computed in $O(n)$
 288 time using two linear scans, one over P and one over Q . The wedges $U_Q(i, j)$, $L_P(i, j)$, $L_Q(i, j)$
 289 are computed in a similar manner.

290 **Computing the table X .** First, we initialize all $X[i, j]$ to **false**, except for $X[0, 0]$ which is
 291 set to **true**. Then we compute $X[i, j]$ incrementally in the specified order. In each iteration, we
 292 process (p_i, q_j) only if they can be matched to each other, i.e., if $cl(p_i, q_j) = \mathbf{true}$.

293 When we compute an entry $X[i, j]$, we have two cases. If $X[i, j]$ is still set to **false**, i.e.,
 294 we do not know of a middle curve covering P_i and Q_j yet, we first check whether adding p_i or
 295 q_j to a covering sequence would extend the coverage to here. For this, we check if $\hat{L}_P(i, j)$ or
 296 $\hat{L}_Q(i, j)$ intersects \bar{X} . If $\hat{L}_P(i, j)$ intersects \bar{X} , then p_i can be added to a covering sequence, and
 297 we set $X[i, j] = \mathbf{true}$. Conversely, if $\hat{L}_Q(i, j)$ intersects \bar{X} , then q_j can be added to a covering
 298 sequence, and we do the same.

299 If $X[i, j]$ is **true**, then both p_i and q_j can be added to a covering sequence, and hence we
 300 add the points covered by p_i or q_j , i.e., $U_P(i, j)$ and $U_Q(i, j)$, to X and \bar{X} . The wedge $U_P(i, j)$
 301 is *added to X and \bar{X}* as follows: We update \bar{X} with $U_P(i, j)$. During the update step we can
 302 identify all pairs $(i', j') \in U_P(i, j)$ with $\neg X[i', j']$; these are all (i', j') such that i' is a leaf in \bar{X}
 303 that gets updated and $\max\{j_B, \bar{X}[i']\} \leq j' \leq j_T$, where (i_B, j_B) is the lower left and (i_T, j_T)
 304 the upper right corner of $U_P(i, j)$. We set all $X[i', j'] = \mathbf{true}$ and store a pointer from (i', j') to
 305 (i, j) that is labeled with P . Adding $U_Q(i, j)$ to X and \bar{X} is done in a similar manner, but the
 306 pointers are labeled with Q . Note that the upper wedges are added to X in such a way that
 307 each $X[i, j]$ is set to **true** only once.

308 The algorithm can be summarized as follows.

Set $X[i, j] = \mathbf{false}$ for all index pairs (i, j) , except $X[0, 0]$ which is set to \mathbf{true} .

Set $\bar{X}[i] = -1$ for all indices $i > 0$, except $\bar{X}[0]$ which is set to 0.

for $j = 1$ **to** m **do**

for $i = 1$ **to** n **do**

if $cl(p_i, q_j) = \mathbf{true}$:

if $\neg X[i, j]$: If $\hat{L}_P(i, j)$ or $\hat{L}_Q(i, j)$ intersects \bar{X} , set $X[i, j]$ to \mathbf{true}

if $X[i, j]$: Add $U_P(i, j)$ and $U_Q(i, j)$ to X and \bar{X} , i.e. set all entries to \mathbf{true}

For the correctness of the algorithm, observe that if $X[i, j]$ holds because of $A'[i, j]$, then it is set to \mathbf{true} when the last point of a covering is processed. If $X[i, j]$ holds by $C'[i, j]$ or $E'[i, j]$, then this is handled in the $\neg X[i, j]$ case of the algorithm.

Recall that we assume $m \leq n$. The running time for computing all wedges is $O(n^2)$ since for each point $p_i \in P$ or $q_j \in Q$, we perform a constant number of linear scans. For the main part of the dynamic programming algorithm, when we consider an index pair (i, j) , we perform a query on \bar{X} which takes $O(\log n)$ time, and we add one or two upper wedges to X . The update operation that is part of adding a wedge takes $O(c \log n)$ time, where c is the number of indices that are updated. However note that $\bar{X}[i]$ is updated at most m times for each index i in total, and $X[i, j]$ is updated at most once for each index pair (i, j) . Thus the running time for the decision algorithm is $O(n^2 + nm \log n)$.

To extract a middle curve from the table X we additionally store labeled pointers in the table as follows. Each entry (i, j) in X that is set to \mathbf{true} (except for $(0, 0)$) gets a pointer to a \mathbf{true} entry (i', j') with $i' < i$ and $j' < j$. Furthermore, the pointer receives two labels: $p|q$ and $f|b$ where the first indicates the curve P or Q , and the second indicates *forward* or *backward*. The two labels together exactly specify one of the four points $p_{i'}, q_{j'}, p_i, q_j$. A middle curve can be reconstructed from these pointers by following them from the entry (n, m) to $(0, 0)$ and choosing the points according to the labels. The pointers can be set in the algorithm as follows: when $X[i, j]$ is set to \mathbf{true} in the $\neg X[i, j]$ case, a pointer is set to a true entry in $\hat{L}_P(i, j)$ or $\hat{L}_Q(i, j)$ labeled $(p|q, b)$. When $X[i', j']$ is set to true when adding an upper wedge $U_P(i, j)$ or $U_Q(i, j)$ in the $X[i, j]$ case, a pointer is set to (i, j) labeled $(p|q, f)$. Thus in both cases the pointer refers to p_i or q_j depending on whether the P or Q wedge is involved.

Note that the algorithm allows multiple occurrences of vertices. However, the restriction enforces that if a vertex occurs multiple times, then all vertices of the other curve that occur in between are matched to that vertex in the discrete Fréchet matching. Figure 5 shows an example of this.

4.2 Optimization Problems for Two Curves

The optimal distance is attained between a pair of points from $P \cup Q$. We therefore sort all distances between such pairs of points in $O(n^2 \log n)$ time and then find the smallest distance by binary search using the decision algorithm.

If we wish to minimize the size of an ordered middle curve for a fixed distance ε , we can, similar to the case in Section 3.2, adapt the dynamic program by storing the smallest size of a middle curve covering up to p_i and q_j in the arrays, and substituting \vee by a \min . In this case, however, we cannot use the upper envelope \bar{X} anymore to facilitate fast queries and updates. So, in each iteration we instead update each entry in X individually, replacing the

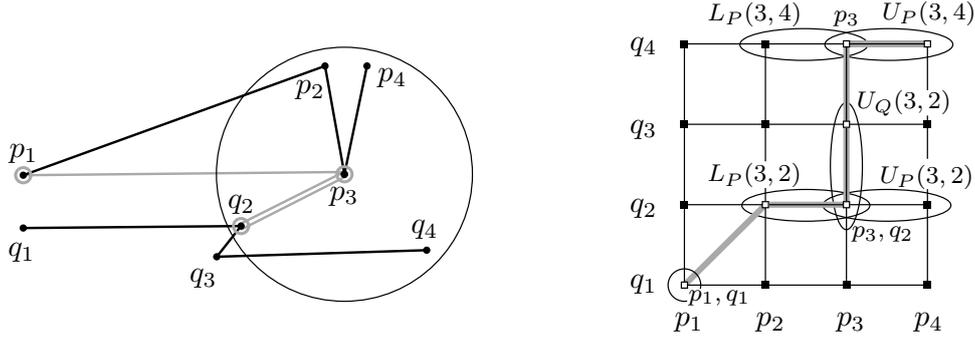


Figure 5: (a) Two curves $P = (p_1, p_2, p_3, p_4)$ and $Q = (q_1, q_2, q_3, q_4)$, and a restricted middle curve $R = (p_1, p_3, q_2, p_3)$ that uses p_3 twice. (b) A diagram illustrating array X and wedges.

345 “if $cl(p_i, q_j) = \text{true}$ ” statement with two nested for-loops. Hence, the total running time is
 346 $O((mn)^4)$. If we wish to optimize ε as well, we can do so using binary search, resulting in a
 347 total running time of $O((mn)^4 \log n)$.

348 4.3 Generalization to Multiple Curves

349 For $k > 2$ curves of size at most n each, the decision algorithm can be generalized to work with
 350 a $(k - 1)$ -dimensional range tree for \bar{X} and running time $O(n^k \log^{k-1} n)$. We search over all
 351 distances between two points from any curves, so a middle curve can be decided in $O(n^k \log^k n)$
 352 time. Using the pointers set by the algorithm, the algorithm can also output a middle curve.

353 To compute a minimal restricted middle curve for a fixed distance, we can modify the simpler
 354 $O(n^{2k})$ algorithm, where we iterate over all entries $X[i, j]$ and if reachable add the k upper
 355 wedges. We now store in each entry the minimum complexity needed to cover the curves up
 356 to here (and possibly a pointer to the last point added), and take the minimum when adding
 357 wedges. This simpler algorithm has complexity $O(n^{2k})$ and our technique of speeding up the
 358 algorithm using the upper envelope does not apply.

359 We summarize the results of this section in the following theorem:

360 **Theorem 2.** *For two polygonal curves with n and m vertices for $m \leq n$, and $\varepsilon > 0$, it can be*
 361 *decided whether there exists a restricted middle curve with distance at most ε in $O(n^2 + mn \log n)$*
 362 *time. A restricted middle curve with minimum distance can be computed in $O(n^2 \log n +$*
 363 *$mn \log^2 n)$ time. A restricted middle curve with minimum complexity and distance at most*
 364 *$\varepsilon > 0$ to the input curves can be computed in $O((mn)^4)$ time, and ε can be optimized in total*
 365 *$O((mn)^4 \log n)$ time. For $k \geq 2$ curves of size at most n each, a restricted middle curve with*
 366 *minimal distance can be computed in $O(n^k \log^k n)$ time.*

367 5 Algorithms for the Unordered Case

368 In this section we consider the problem of computing an unordered middle curve. We present a
 369 straight-forward approach for the problem in Section 5.1. Then in Sections 5.2–5.4 we present
 370 faster algorithms for the case of curves in the Euclidean plane.

371 In the unordered case we can formulate the problem more generally, that we are given two
 372 curves P, Q and a set S of points from which we build the middle curve R . The algorithms we
 373 present in this section work for any set S of points. When we build R from P and Q we use
 374 $S = P \cup Q$.

375 5.1 Straight-Forward Decision and Optimization

376 Let again $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$ be two input curves with $m \leq n$. Let S be a set
 377 of ℓ points (possibly $S = P \cup Q$) from which we build the middle curve R . To solve the decision
 378 problem for the unordered case, we modify the dynamic programming algorithm for computing
 379 the discrete Fréchet distance of two curves [6] as follows. Let $\varepsilon > 0$ be an input of the decision
 380 problem. We consider an $n \times m$ matrix X , which we call the *free space matrix*. Each entry
 381 $X[i, j]$ corresponds to the pair (p_i, q_j) of points. In contrast to the original algorithm, we mark
 382 an entry $X[i, j]$ **free** if and only if there exists a point v from S such that v has distance at
 383 most ε to both p_i and q_j . Then we search for a monotone path from $X[1, 1]$ to $X[n, m]$ within
 384 the **free** entries in X .

385 One way to determine whether $X[i, j]$ is **free** for an index pair (i, j) is to test all possibilities
 386 for a point $v \in S$, each of which can be tested in $O(1)$ time. When $|S| = \ell$ the free space matrix
 387 of size $m \times n$ can be computed in $O(\ell mn)$ time. When $S = P \cup Q$, i.e., $\ell = m + n$ this results
 388 in $O(mn^2)$ time for $m \leq n$. The running time for searching for a monotone path in the matrix
 389 is $O(mn)$.

390 Similarly, we can compute an unordered middle curve with minimal distance in the same
 391 time as follows. We let $X[i, j]$ be the minimum of $\max\{d(v, p_i), d(v, q_j)\}$ over all points $v \in S$.
 392 Then we search for a monotone path from $X[1, 1]$ to $X[n, m]$ such that the maximum entry
 393 $X[i, j]$ in the path is minimized.

394 For $k \geq 2$ input curves of size at most n each, where k is constant, and $|S| = \ell$, we can
 395 compute an unordered middle curve of minimal distance in $O(\ell n^k)$ time in total using a k -
 396 dimensional free space matrix of size n^k .

397 To compute an unordered middle curve of minimal size for a given distance $\varepsilon > 0$, we consider
 398 the following graph: vertices are all tuples of points from the k curves, and there is a directed
 399 edge from (i_1, \dots, i_k) to (j_1, \dots, j_k) if there is a vertex in S that covers the subcurves from
 400 i_h to j_h on each curve $1 \leq h \leq k$. This graph has n^k vertices and up to n^{2k} edges. It can
 401 be constructed in $O(\ell n^{2k})$ time, by computing for each of n^k vertices all of its outgoing edges
 402 in $O(\ell n^k)$ time. In the graph we search for a shortest path from vertex $(1, \dots, 1)$ to vertex
 403 (n, \dots, n) which corresponds to a minimum middle curve.

404 **Theorem 3.** *For $k \geq 2$ curves of size at most n each and a set S of ℓ points, an unordered*
 405 *middle curve built from S can be computed in $O(\ell n^k)$ time. For fixed $\varepsilon > 0$, an unordered middle*
 406 *curve with minimum complexity and distance at most ε to the input curves can be computed in*
 407 *$O(\ell n^{2k})$ time.*

408 5.2 Decision Algorithm for Two Curves in the Euclidean Plane

409 In this section, we assume that the two input curves lie in the Euclidean plane and use d_E to
 410 denote the Euclidean metric. A further assumption is that now $S = P \cup Q$. We describe how to
 411 determine whether $X[i, j]$ is **free** more efficiently for the decision problem in this setting.

412 We will use a circular sweep to determine, for each point q_j in Q , all points p_i in P such
 413 that $X[i, j]$ is **free**, i.e., there is some point v of $P \cup Q$ which has distance at most ε to both p_i
 414 and q_j . Let $U_j(\varepsilon)$ be the union of disks of radius ε centered at points in $P \cup Q$ and containing
 415 $q_j \in Q$, and $\partial U_j(\varepsilon)$ be the boundary of $U_j(\varepsilon)$. Then, for a point $p_i \in P$ contained in $U_j(\varepsilon)$,
 416 $X[i, j]$ is **free**. To compute $X[i, j]$ for all $p_i \in P$, we construct $U_j(\varepsilon)$ and perform a circular
 417 sweep around q_j for the points in P . Once the endpoints of the circular arcs of $\partial U_j(\varepsilon)$ and all
 418 points $p_i \in P$ are sorted around q_j in clockwise order, the circular sweep takes $O(m + n)$ time.

419 We design an algorithm that computes $U_j(\varepsilon)$ efficiently by constructing two data structures,
 420 the *history list* \mathcal{H}_j and the *deletion list* \mathcal{D}_j . In the *preprocessing* phase, we increase ε gradually

421 and consider the combinatorial structure of $U_j(\varepsilon)$. In doing so, we maintain the changes of the
 422 combinatorial structure of $U_j(\cdot)$ using the two data structures. In the *construction* phase, we
 423 compute, for a given ε , the union $U_j(\varepsilon)$ of disks using the two data structures. This will allow
 424 us to solve the decision problem efficiently. The construction phase takes $O(m+n) = O(n)$ time
 425 while the preprocessing phase takes $O(mn \log n)$ time. The space we use for the data structures
 426 is $O(mn)$.

427 Each arc of $\partial U_j(\varepsilon)$ changes continuously as ε increases. We will show that there are at most
 428 two arcs in $\partial U_j(\varepsilon)$, which come from the same circle (Lemma 6). We treat them as distinct
 429 elements in $\partial U_j(\varepsilon)$.

430 Data Structures for a Point $q_j \in Q$

431 1. The data structures we maintain during the preprocessing phase:

- 432 (a) $S_j(\varepsilon)$ denotes the set of points of $P \cup Q$ that are within distance ε from the point q_j .
- 433 (b) $\Gamma_j(\varepsilon)$ is the balanced binary search tree representing the union of disks of radius ε
 434 centered at points of $S_j(\varepsilon)$ for a fixed ε . The union $U_j(\varepsilon)$ of disks is star-shaped
 435 and its boundary $\partial U_j(\varepsilon)$ is a sequence of circular arcs with vertices in between. We
 436 maintain the balanced binary search tree $\Gamma_j(\varepsilon)$ of these circular arcs in clockwise
 437 order around q_j . Each element in $\Gamma_j(\varepsilon)$ corresponds to an element in the history list
 438 (defined below) and they reference each other with a pointer.

439 2. The data structures used for constructing $U_j(\varepsilon)$ in the construction phase:

- 440 (a) The history list $\mathcal{H}_j = \{x_1, \dots, x_l\}$: Each element of the list corresponds to an arc
 441 of $\partial U_j(\varepsilon)$ for some ε . Each such arc is part of a disk boundary of radius ε centered
 442 at a point of $S_j(\varepsilon)$ that appears on $\partial U_j(\varepsilon)$. This list represents the order of circular
 443 arcs of $\partial U_j(\varepsilon)$ for *every* $\varepsilon > 0$. That is, for any three elements in \mathcal{H}_j , if all arcs
 444 corresponding to the elements appear on $\partial U_j(\varepsilon)$ for some fixed $\varepsilon > 0$, then the order
 445 of them on $\partial U_j(\varepsilon)$ is the same as the order of the three elements in \mathcal{H}_j .
- 446 (b) The deletion list $\mathcal{D}_j = \{(\varepsilon_1, \varepsilon'_1), \dots, (\varepsilon_t, \varepsilon'_t)\}$ with $\varepsilon_k \leq \varepsilon'_k$ for every $1 \leq k \leq t$: The
 447 k -th element of the list is defined by the k -th closest point in $P \cup Q$ from q_j . By
 448 Lemma 6, the disk centered at the k -th closest point of radius ε has at most two arcs
 449 appearing on $\partial U_j(\varepsilon)$ for any fixed $\varepsilon > 0$. An arc of the disk disappears from $\partial U_j(\varepsilon)$
 450 at $\varepsilon = \varepsilon_k$, and the other arc disappears from $\partial U_j(\varepsilon)$ at $\varepsilon = \varepsilon'_k$. (ε_k or ε'_k might be
 451 infinity or zero.) Since the list is an array of size $m+n$, we can access each element
 452 in $O(1)$ time.

453 **Preprocessing Phase: Constructing the Data Structures.** For each $q_j \in Q$, we construct
 454 the data structures mentioned above. To do this, we imagine that ε increases from zero to infinity.
 455 As ε changes, the combinatorial structure of $U_j(\varepsilon)$ changes. More specifically, a new arc appears
 456 on $U_j(\varepsilon)$ and an arc disappears from $U_j(\varepsilon)$. A value ε' is called an *event* if the combinatorial
 457 structure of $U_j(\varepsilon)$ changes at $\varepsilon = \varepsilon'$.

458 There are two types of events: *point events* and *radius events*. A new arc appears on $U_j(\varepsilon)$
 459 only if a new point is inserted to $S_j(\varepsilon)$. A new point $p \in P \cup Q$ is inserted to $S_j(\varepsilon)$ only when ε
 460 is the distance between q_j and p by definition. We call such a value ε a *point event*. We say p
 461 *defines* this point event. For a point event e , we let $p(e)$ be the point defining e . There are two
 462 cases how an arc disappears from $U_j(\varepsilon)$: the case that the arc is contained in the disk centered
 463 at $p(\varepsilon')$ with radius ε' at $\varepsilon = \varepsilon'$ for a point event ε' and the case that $\partial U_j(\varepsilon)$ passes through the
 464 point equidistant from the centers of the arc and its two neighboring arcs of $U_j(\varepsilon')$ at some point

465 ε' . The first case is handled by the point event defined by p , and the second case is handled by a
466 *radius event* which is defined as follows. For any three consecutive arcs of $\partial U_j(\cdot)$, we call $d(p, c)$
467 a *radius event*, where p is the center of any of the three arcs and c is the point equidistant from
468 the centers of the three arcs. We say the three arcs *define* this event.

469 In the following, we show how to handle each event as ε increases.

- 470 1. Sort all points of P around q_j in clockwise order. Let \mathcal{L}_j denote the sorted list.
- 471 2. Initialize $\mathcal{H}_j := \{q_j\}$, $\Gamma_j(0) := \{p_j\}$, $S_j(0) := \{p_j\}$ and $U_j(0) := \{p_j\}$. Initialize \mathcal{D}_j to the
472 array of size $m + n$ each of whose elements is initialized to a null value.
- 473 3. Sort all points of $P \cup Q$ in increasing order of distance from q_j and store the distances
474 together with their corresponding points as events in an event queue \mathcal{E} . Note that they
475 are point events.
- 476 4. While \mathcal{E} is not empty, handle the earliest event $e \in \mathcal{E}$ as follows. Let e' denote the event
477 we just handled.

478 (a) If e is a point event, the boundary of the disk centered at $p(e)$ with radius e appears
479 on $\partial U_j(e)$ in one connected circular arc γ if it appears on $\partial U_j(e)$ (see Lemma 4.) The
480 endpoints of γ can be computed in $O(\log n)$ time by Lemma 5.

481 Let γ_1 and γ_2 denote the neighboring arcs of γ , respectively, along $\partial U_j(e)$ so that γ_1, γ
482 and γ_2 appear on $\partial U_j(e)$ in clockwise order. See Figure 6(a) for an illustration. When
483 we compute γ , we can obtain γ_1 and γ_2 . We find the element in \mathcal{H}_j corresponding to
484 γ_1 and insert an element corresponding to γ to \mathcal{H}_j next to the element. We remove
485 all arcs of $\partial U_j(e')$ coming from γ_1 to γ_2 in clockwise order along $\partial U_j(e')$ and update
486 the corresponding elements in \mathcal{D}_j to e . We update $\Gamma_j(e')$ accordingly to obtain $\Gamma_j(e)$
487 in $O(c \log n)$ time, where c is the number of the deleted arcs.

488 Then we have new triples of consecutive arcs along $U_j(e)$, which induce radius events.
489 Note that such a triple contains γ , and thus there are at most three new radius events.
490 We insert all such events to \mathcal{E} .

491 (b) If e is a radius event, we first check if all three arcs defining e appear on $U_j(e')$. If
492 so, we remove the arc in the middle among the three arcs from $\Gamma_j(e')$ and update \mathcal{D}_j
493 accordingly by setting the value of the element corresponding to γ in \mathcal{D}_j to e . Due to
494 the deletion of γ , the two neighboring arcs of γ become adjacent in $U_j(e)$ for which
495 we insert a new radius event.

496 The number of point events is $O(n)$ and the number of radius events is bounded by the num-
497 ber of distinct arcs appearing on $U_j(\varepsilon)$ over all increasing ε values, which is $O(n)$ by Lemma 6.
498 Thus, the number of events in the preprocessing phase and the size of the data structures are
499 $O(n)$. The preprocessing phase takes $O(n \log n)$ time for each $q_j \in Q$.

500 **Construction Phase: Constructing the Free Space Matrix.** Given $\varepsilon > 0$, the construc-
501 tion phase works as follows. For each $q_j \in Q$:

- 502 1. Scan the list \mathcal{H}_j from the first element to the last element and check the list \mathcal{D}_j to determine
503 whether each arc appears on $\partial U_j(\varepsilon)$. In $O(n)$ time, we can obtain the sequence of the arcs
504 appearing on $\partial U_j(\varepsilon)$ in clockwise order, which represents $\partial U_j(\varepsilon)$ itself.
- 505 2. Perform a circular sweep by a ray from q_j around q_j with the points in \mathcal{L}_j and the vertices
506 of $\partial U_j(\varepsilon)$. During the sweep, the ray always intersects an arc of $\partial U_j(\varepsilon)$. We can determine

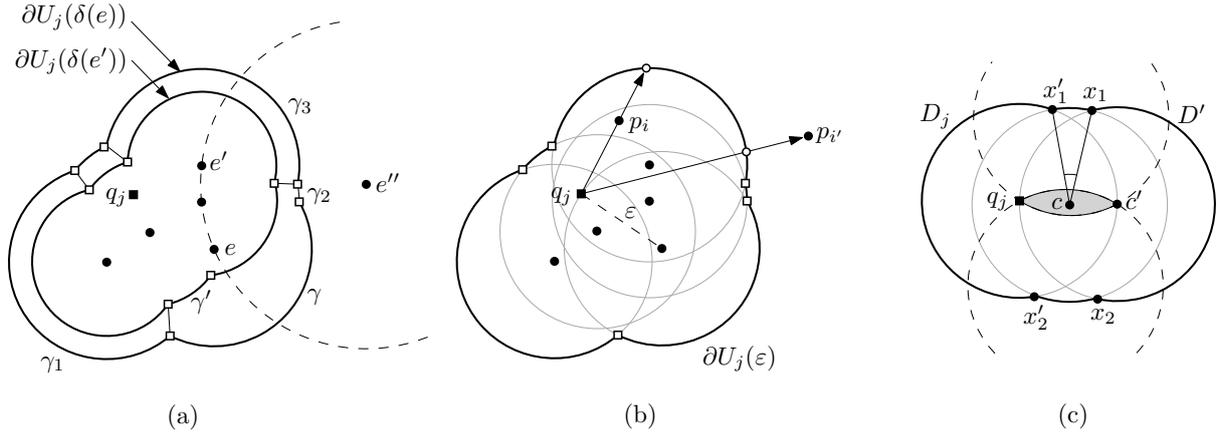


Figure 6: (a) Preprocessing. At the point event e , the arc γ' disappears and the new radius event e'' is created, where e'' is the point equidistant from the centers of γ_1, γ_2 , and γ . γ_2 will disappear from the boundary at radius event e'' unless it disappears before the event. (b) Construction of the free space matrix. During the circular sweep, we compare each point $p_i \in P$ with the intersection point of $\partial U_j(\varepsilon)$ with the ray from q_j through p_i to determine whether p_i is in $U_j(\varepsilon)$. Here, $X[i, j]$ is **free** and $X[i, j']$ is not **free**. (c) The arc centered at c is subdivided into two subarcs by the event e . Then $\angle x_1 c x'_1$ and $\angle x_2 c x'_2$ are at most $2\pi/3$.

507 whether p_i is in $U_j(\varepsilon)$ by comparing each point p_i in \mathcal{L}_j encountered by the ray with the
 508 current circular arc of $\partial U_j(\varepsilon)$ intersected by the ray. If so, set $X[i, j]$ to **free**. Figure 6(b)
 509 illustrates the circular sweep. This again can be done in $O(n)$ time once \mathcal{L}_j has been
 510 computed in the preprocessing step.

511 For the correctness of the algorithm, we show the following lemma.

512 **Lemma 4.** *For a point event e , at most one arc of the disk of radius e centered at $p(e)$ appears*
 513 *on $\partial U_j(e)$.*

514 *Proof.* Assume to the contrary that there are two maximal circular arcs, γ and γ' , on $\partial U_j(e)$
 515 such that both arcs are on the boundary of the disk of radius e centered at $p(e)$. Let D_j be the
 516 disk of radius e centered at q_j . Since $d_E(p(e), q_j) = e$, the boundary of D_j contains $p(e)$. Then
 517 there must be a disk D' splitting the arc of $\partial D \setminus D_j$ into two, one containing γ and the other
 518 containing γ' .

519 Here, the center of D' is contained in D_j since the center of D' has already been handled.
 520 Thus $\partial D \cap D_j$ intersects D' at a point, say x_4 . Since D' splits $\partial D \setminus D_j$ into two, there are three
 521 points x_1, x_2 and x_3 appearing on $\partial D \setminus D_j$ in clockwise order such that x_1 and x_3 are contained
 522 in D' , and x_2 and x_4 are not contained in D' . This means that D and D' cross each other, which
 523 is a contradiction. \square

524
 525

To analyze the running time of the algorithm, we need the following lemma.

526 **Lemma 5.** *The endpoints of the circular arc to be inserted at step 4(a) of the preprocessing*
 527 *phase can be computed in $O(\log n)$ time.*

528 *Proof.* Let e and e' be the current event and the event previous to e , respectively. We maintain
 529 $\Gamma_j(\varepsilon)$, which is the balanced binary search tree of the arcs of $\partial U_j(\varepsilon)$ in clockwise order around
 530 q_j . Note that $U_j(\varepsilon)$ is star-shaped with respect to q_j and there is no structural change to $\partial U_j(\varepsilon)$
 531 for $e' \leq \varepsilon < e$. Since the disk of radius e centered at $p(e)$ is also star-shaped and contributes

532 only one connected circular arc γ to $\partial U_j(\varepsilon)$, the two endpoints of γ can be computed in $O(\log n)$
533 time by a binary search on the vertices of $\partial U_j(\varepsilon)$. \square

534

535 **Lemma 6.** *Once an arc is divided into two subarcs, the subarcs will never be divided again.*

536 *Proof.* Let γ_1 and γ_2 be the arcs appearing on $\partial U_j(\varepsilon)$ for some $\varepsilon > 0$, which come from the
537 same disk D centered at a point, say c . Let x_1, x'_1, x_2 and x'_2 be the endpoints of γ_1 and γ_2 in
538 clockwise order along the disk centered at c with radius ε .

539 We first claim that $\angle x_1 c x'_1$ and $\angle x_2 c x'_2$ are at most $2\pi/3$. See Figure 6(c). Let D_j be the
540 disk centered at q_j with radius ε . Since there are two arcs which come from ∂D , there is a disk,
541 say D' , splitting $\partial D \setminus D_j$. Note that the center of D' , say c' , is contained in D_j . Without loss of
542 generality, assume that $q_j c'$ is contained in the x -axis. Since D' splits $\partial D \setminus D_j$, D contains the
543 intersection points between $\partial D'$ and ∂D_j . Thus the center c of D is contained in the intersection
544 I of D' , D_j and the disks centered at the intersection points in $\partial D' \cap \partial D_j$ of radius ε . Note
545 that I consists of two circular arcs whose common endpoints are q_j and c' . Also, notice that
546 each of x_1, x'_1, x_2 and x'_2 lies on the bisector of c and q_j (or c'). By construction, the intersection
547 point between the two bisectors, one between c and q_j and one between c and c' , lies outside of
548 $D \cup D_j$. Therefore, $\angle x_1 c x'_1$ and $\angle x_2 c x'_2$ are at most $2\pi/3$.

549 Now we show that γ_1 is not divided further. The case of γ_2 can be shown analogously.
550 Assume to the contrary that γ_1 is divided at ε . Again, let D_j be the disk centered at q_j with
551 radius ε . This means that there is a disk, say D'' , centered at a point in D_j , say c'' , with radius
552 ε such that x_1 and x'_1 are not contained in D'' but a point, say x , other than its endpoints is
553 contained in D'' . Imagine the set of points whose distance to x_1 (and x'_1) is larger than ε and
554 whose distance to x is at most ε . The set lies outside of D because $\angle x_1 c x'_1$ is at most $2\pi/3$.
555 Since c'' lies in D_j and lies outside of D , the line segment cc'' intersects $\partial D \setminus \gamma_1$. This means
556 that there are four points on ∂D such that the first and third points are contained in D'' but
557 the second and fourth points are not contained in D'' in clockwise order around ∂D . This con-
558 tradicts that D and D'' are disks. Therefore, γ_1 is not divided further, and the lemma holds. \square

559

560 To obtain a covering sequence in addition to a yes-answer, for each entry $X[i, j]$ of the free
561 space matrix, we mark the center of the circular arc of $\partial U_j(\varepsilon)$ intersected by the ray starting
562 from q_j towards p_i . Then the sequence of labels of a monotone path gives a feasible unordered
563 sequence for the middle curve.

564 **Theorem 7.** *For two polygonal curves with n and m vertices for $m \leq n$ in the Euclidean plane,
565 the decision problem for the unordered case can be solved in $O(mn)$ time with $O(mn \log n)$
566 preprocessing time. A covering sequence can be computed in the same time.*

567 5.3 Optimization Algorithm for Two Curves in the Euclidean Plane

568 We apply binary search on the sorted list of distances of pairs of points from $P \cup Q$ involved in
569 each step. There are $O((m+n)^2) = O(n^2)$ distinct distances each defined by two points from
570 $P \cup Q$. We will show that we need only $O(mn)$ of them to compute the optimal distance ε^* .
571 The optimization algorithm we propose works as shown in the following four steps.

- 572 1. Compute the set \mathcal{D} of distances each defined by two points that are either both from Q ,
573 or one from P and one from Q .
- 574 2. Sort the $O(mn)$ distances of \mathcal{D} and apply binary search on the sorted list with the decision
575 algorithm in Section 5.2. Let ε_1 be the largest distance of \mathcal{D} that the decision algorithm

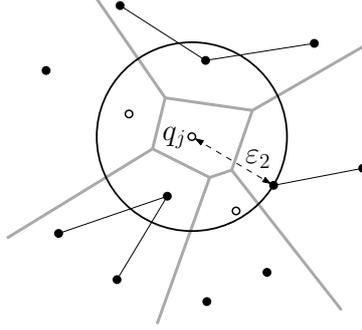


Figure 7: The white (circle) points are in Q and the black points are in P . The lengths of the line segments connecting two black points are candidates of ε^* .

576 returns “no” and ε_2 be the smallest distance of \mathcal{D} that the decision algorithm returns “yes”.
 577 Then we know that $\varepsilon_1 \leq \varepsilon^* \leq \varepsilon_2$. If $\varepsilon_1 \neq \varepsilon^*$ and $\varepsilon_2 \neq \varepsilon^*$, then ε^* is the distance defined
 578 by two points in P . See Figure 7.

- 579 3. To find ε^* , for each point $q_j \in Q$,
- 580 (a) compute the set S_j of points in $P \cup Q$ that are at distance at most ε_2 from q_j , and
 581 construct the Voronoi diagram $\text{VD}(S_j)$.
- 582 (b) For each point p_i in $P \setminus S_j$, locate the cell of $\text{VD}(S_j)$ that contains p_i . If the site
 583 x associated with the cell is from P and $\varepsilon_1 < d_E(p_i, x) < \varepsilon_2$, then $d_E(p_i, x)$ is a
 584 candidate for ε^* .
- 585 4. Sort the $O(mn)$ candidate distances and again apply binary search on the sorted list with
 586 the decision algorithm above.

587 **Analysis.** Let (p_i, q_j, x) be a tuple realizing ε^* . Then $\max\{d_E(p_i, x), d_E(q_j, x)\} = \varepsilon^*$. Clearly,
 588 x is the point in $P \cup Q$ that minimizes $\max\{d_E(p_i, x), d_E(q_j, x)\}$. If $x \in P$ and $\varepsilon_1 < \varepsilon^* < \varepsilon_2$,
 589 then $d_E(p_i, x) > d_E(q_j, x)$. Thus x is the point in S_j that is closest to p_i . Thus, x is the point
 590 site associated with the Voronoi cell in $\text{VD}(S_j)$ that contains p_i . This proves that ε^* is in the
 591 set of all candidates.

592 Let us analyze the running time of the optimization algorithm. The set \mathcal{D} can be constructed
 593 in $O(mn)$ time. It takes $O(mn \log n)$ time to sort the distances in \mathcal{D} . The binary search on
 594 the sorted list with the decision algorithm takes $O(mn \log n)$ time as the preprocessing phase
 595 is executed only once for each $q_j \in Q$ and the history and deletion lists are used for different
 596 radii. In Step 3, the Voronoi diagram $\text{VD}(S_j)$ can be constructed in $O(n \log n)$ time for each
 597 $q_j \in Q$, and the point location for n points can be performed in the same time. Step 3(b) takes
 598 $O(n \log n)$ time for each $q_j \in Q$.

599 5.4 Generalization to Multiple Curves in the Euclidean Plane

600 The decision algorithm can be extended to k curves P^1, \dots, P^k of size at most n each for
 601 a constant k in the Euclidean plane. We construct a k -dimensional free space matrix whose
 602 entries correspond to k -tuples of points from distinct curves. An entry of the matrix is marked
 603 as **free** if there is an input point in the intersection of the disks centered at points in the k -tuple
 604 corresponding to the entry with radius ε , where ε is an input distance for the decision problem.

605 To construct the matrix, we use an approach similar to the one for $k = 2$. For every
606 $(k - 1)$ -tuple (p_1, \dots, p_{k-1}) with $p_i \in P^i$ for $i = 1, \dots, k - 1$, we do the following. Let D be
607 the intersection of the disks centered at p_i with radius ε for all $i = 1, \dots, k - 1$. We compute
608 the union U of the disks with radius ε centered at input points lying in D , and check for each
609 point in P^k whether it is contained in the union. Here, the boundary of U has the star-shaped
610 property for any point p_i in the $(k - 1)$ -tuple. We mark the entry in the matrix corresponding
611 to the k -tuple (p_1, \dots, p_k) as **free** if and only if p_k is contained in the union. Then we check if
612 there is a monotone path from $X[1, \dots, 1]$ to $X[n, \dots, n]$ within the **free** entries in X .

613 The construction of U takes $O(kn \log(kn))$ time for fixed $\varepsilon > 0$. However we can compute it
614 more efficiently by maintaining the history data structure as we did for $k = 2$. Imagine that ε
615 increases from zero to infinity, and consider the combinatorial changes of U . There are two types
616 of events: the point events and the radius events. The radius event is defined in the same way as
617 the case of $k = 2$. For the point events, observe that the centers of the circular arcs of ∂U are in
618 D . As ε increases, D changes as well. A new arc appears on ∂U when its center appears on the
619 boundary of D . Also, its center appears on the boundary of D when ε is the distance between
620 the center and the point in the $(k - 1)$ -tuple farthest from the center. Such distances are defined
621 as point events. Clearly, there are $O(n)$ point events. Also, Lemmas 4, 5 and 6 hold for a larger
622 k . Thus we can maintain the combinatorial structure of U in $O(kn \log(kn)) = O(n \log n)$ time.
623 We keep track of the combinatorial changes using the history and deletion data structures as we
624 did for $k = 2$. Then after the preprocessing, we can construct the union in $O(kn) = O(n)$ time.
625 We can check for each point in P^k whether it is contained in the union in $O(n)$ time. Thus,
626 the decision algorithm takes $O(n^k)$ time once the history data structures are constructed for all
627 $(k - 1)$ -tuples.

628 To compute a middle curve, we first construct history data structures for all $(k - 1)$ -tuples
629 in $O(n^k \log n)$ time. Then we sort all distances defined by point pairs from $P^1 \cup \dots \cup P^k$ and
630 search the optimal distance among them. Thus, we can compute an optimal covering sequence
631 in $O(n^k \log n)$ time.

632 **Theorem 8.** *For two polygonal curves with n and m vertices for $m \leq n$ in the Euclidean plane,*
633 *the optimization problem for the unordered case can be solved in $O(mn \log n)$ time. An optimal*
634 *covering sequence can be computed in the same time. For a fixed $k \geq 2$, the optimization of k*
635 *curves of size at most n each in the Euclidean plane can be solved in $O(n^k \log n)$ time.*

636 6 Discussion

637 We presented algorithms for computing a middle curve of minimal discrete Fréchet distance to
638 the input curves. All our algorithms run in time exponential in k , the number of input curves.
639 Hence these are practical only for small k . However, other algorithms that compute variants of
640 the Fréchet distance for k curves such as [5] and [7] also take time exponential in k due to the
641 use of a k -dimensional free space diagram. Assuming the Strong Exponential Time Hypothesis
642 it is known that essentially no faster algorithms are possible [3]. Hence we also do not expect
643 any substantially faster algorithms for finding a middle curve based on the (discrete) Fréchet
644 distance. An interesting open problem is to find more efficient approximation algorithms.

645 Also note that a middle curve (unordered, ordered, or restricted) computed by our algorithms
646 can have complexity at most $nk - k + 1$. This follows because each vertex of the middle curve
647 “advances” by at least one vertex on one of the input curves, of nk vertices in total. More
648 formally, let R be a middle curve of size r . Consider the discrete Fréchet mappings of R to
649 each of the input curves. Each of these gives a segmentation with duplicates of size r of the
650 input curve. Now consider the set of all segmentations. For two consecutive vertices of R , the

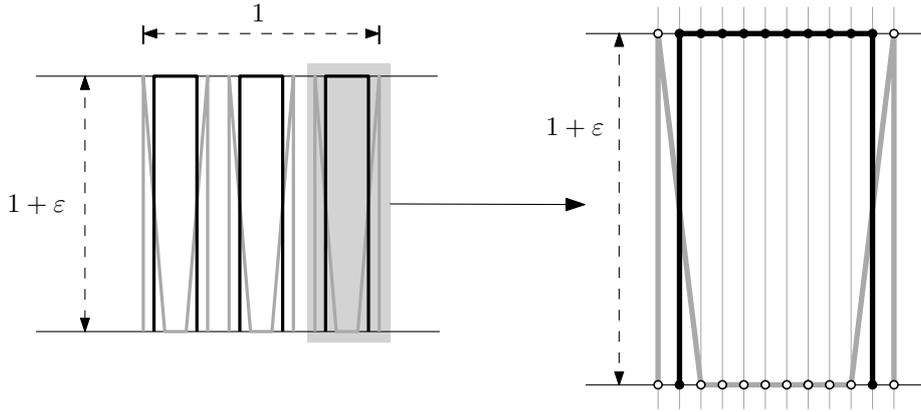


Figure 8: Example of a middle curve of complexity $nk/2$ (here $k = 6$). It consists of $k/2$ (here 3) gray curves and $k/2$ black curves. The curves are given in pairs of one gray and one black curve as shown on the right. Each gray curve has $n - 2$ vertices on the bottom line and 2 vertices on the top line whereas each black curve has 2 vertices on the bottom line and $n - 2$ vertices on the top line. The vertices are aligned along the two parallel lines.

651 corresponding sets of segmentations differ, i.e., one of the input curves advances by at least one
 652 vertex. Furthermore, the first vertex of the middle curve covers all first vertices of each of the
 653 input curves. Hence these do not increase the size of the output middle curve.

654 We observe that this bound is essentially tight by giving an example of k curves of complexity
 655 n each where a minimal middle curve has complexity $nk/2$. Consider the curves shown in
 656 Figure 8. In this example the only middle curve achieving the minimal Fréchet distance $1 + \varepsilon$ is
 657 the set of all points on the bottom line, either ordered from left to right or from right to left. For
 658 this, first observe that for a pair of corresponding gray and black curve, a middle curve consists
 659 of n points, one of each pair of points at distance $1 + \varepsilon$, arbitrarily from the top or bottom line.
 660 However a middle curve for all k curves may only contain points on the bottom line, which are
 661 (horizontally) close to the start and end point of all other curves.

662 However, in practice we expect middle curves to be much smaller. In the example in Figure 8
 663 we observe that if we increase the distance, the complexity decreases fairly quickly. On the other
 664 hand, if we decrease the distance, then a middle curve is no longer possible.

665 **Acknowledgments.** This work was initiated at the 17th Korean Workshop on Computational
 666 Geometry. We thank the organizers and all participants for the stimulating atmosphere. In
 667 particular we thank Fabian Stehn and Wolfgang Mulzer for discussing this paper.

668 References

- 669 [1] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves.
 670 *International Journal of Computational Geometry & Applications*, 5(1-2):75–91, 1995.
- 671 [2] K. Buchin, M. Buchin, J. Gudmundsson, M. Löffler, and J. Luo. Detecting commuting
 672 patterns by clustering subtrajectories. *International Journal of Computational Geometry
 673 & Applications*, 21(3):253–282, 2011.

- 674 [3] K. Buchin, M. Buchin, W. Mulzer, M. Konzack, and A. Schulz. Fine-grained analysis of
675 problems on curves. In G. Barequet and E. Papadopoulou, editors, *Proceedings of the 32nd*
676 *European Workshop on Computational Geometry (EuroCG'16)*, 2016.
- 677 [4] K. Buchin, M. Buchin, M. van Kreveld, M. Löffler, R. I. Silveira, C. Wenk, and L. Wiratma.
678 Median trajectories. *Algorithmica*, 66(3):595–614, 2013.
- 679 [5] A. Dumitrescu and G. Rote. On the Fréchet distance of a set of curves. In *Proceedings*
680 *of the 16th Canadian Conference on Computational Geometry (CCCG'04)*, pages 162–165,
681 2004.
- 682 [6] T. Eiter and H. Mannila. Computing discrete Fréchet distance. Technical report, Technische
683 Universität Wien, 1994.
- 684 [7] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. *ACM Transac-*
685 *tions on Algorithms*, 10(1):3:1–3:22, Jan. 2014.
- 686 [8] E. Sriraghavendra, K. Karthik, and C. Bhattacharyya. Fréchet distance based approach
687 for searching online handwritten documents. In *Proceedings of the Ninth International*
688 *Conference on Document Analysis and Recognition (ICDAR'07)*, volume 1, pages 461–465.
689 IEEE Computer Society, 2007.
- 690 [9] M. J. van Kreveld, M. Löffler, and F. Staals. Central trajectories. In *31st European Work-*
691 *shop on Computational Geometry (EuroCG'15)*, pages 129–132, 2015.
- 692 [10] H. Zhu, J. Luo, H. Yin, X. Zhou, J. Z. Huang, and F. B. Zhan. Mining trajectory corridors
693 using fréchet distance and meshing grids. In *Advances in Knowledge Discovery and Data*
694 *Mining*, pages 228–237. Springer Berlin Heidelberg, 2010.