A fast implementation of near neighbors queries for Fréchet distance (GIS Cup)

Julian Baldus[*]   Karl Bringmann[†]

Abstract

This paper describes an implementation of fast near-neighbours queries (also known as range searching) with respect to the Fréchet distance. The algorithm is designed to be efficient on practical data such as GPS trajectories. Our approach is to use a quadtree data structure to enumerate all curves in the database that have similar start and endpoints as the query curve. On these curves we run positive and negative filters to narrow the set of potential results. Only for those trajectories where these heuristics fail, we compute the Fréchet distance exactly, by running a novel recursive variant of the classic free-space diagram algorithm.

Our implementation won the ACM SIGSPATIAL GIS Cup 2017[1].

---

[*]Saarland University, Saarland Informatics Campus, Saarbrücken, Germany, `julianbaldus@gmail.com`

[†]Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany, `kbringma@mpi-inf.mpg.de`

[1]`http://sigspatial2017.sigspatial.org/giscup2017/`

1 Introduction

The Fréchet distance is a very popular measure of similarity of two given curves π, σ, which roughly speaking measures the minimal length of a leash connecting a dog to its owner as they walk without backtracking along π and σ, respectively. As a natural measure of curve similarity, it has broad applications in geographic information systems. Alt and Godau introduced the Fréchet distance to computer science in seminal papers in the early 90's, where they presented an $O(nm \log(nm))$ algorithm to compute the Fréchet distance of polygonal curves π and σ with n and m vertices [2, 11]. Over the decades, Fréchet distance developed into a rich field of research, in which many generalizations and variants are studied (see, e.g., [1, 8, 9, 10]). However, the quadratic worst-case complexity of Alt and Godau's algorithm is still state-of-the-art, apart from log-factor improvements [7]. The second author recently presented strong evidence that the Fréchet distance has no strongly subquadratic algorithms, by proving that any such algorithm would yield a breakthrough for the Satisfiability problem (specifically it would break the Strong Exponential Time Hypothesis) [5].

In this paper we study the problem of range searching with respect to the Fréchet distance, which comes up in many applications. Here, in the preprocessing phase we are given a set D of curves, which we call the *database*. As a query we are then given a *query curve* π and a *distance threshold* δ, and the task is to output all curves in D that have Fréchet distance at most δ to π.

The naive solution is to compute the Fréchet distance between π and σ for each $\sigma \in D$. If π has n vertices and D consists of k curves having m vertices each, then this takes time $O(knm \log(nm))$. We claim (without proof) that the construction of [5] can be adapted to show that any significant improvement over this running time would yield a breakthrough for the Satisfiability problem, specifically for any $\varepsilon > 0$ an $O((knm)^{1-\varepsilon})$-time algorithm would falsify the Strong Exponential Time Hypothesis. However, this only excludes better running time *guarantees in the worse case*, but it does not rule out algorithms that are efficient on realistic instances.

What are realistic curves? In the area of computational geometry, researchers tried to formalize what it means for a curve to be *realistic*. Several notions of curves have been developed, such as backbone curves [4], κ-bounded and κ-straight curves [3], ϕ-low density curves [10], and c-packed curves [10]. While these notions improved our understanding of characteristics of curves that make Fréchet distance computation hard, they are still very far from modeling real-world trajectories such as the example data set given in the ACM SIGSPATIAL GIS Cup 2017, which is supposedly derived from GPS traffic data in the San Francisco Bay Area. These curves are (close to) shortest paths in an underlying road network, and thus have the following features.

1. There are few sharp turns (narrow angles) in the input data, or at least every input curve has small Fréchet distance to some curve with few sharp turns.

2. Two curves with similar start and endpoints have small Fréchet distance.

Note that this is neither a complete set of characteristics of realistic curves, nor a formal definition of a model of such curves. However, it does show some features that are not incorporated in the standard models of realistic curves, or only to a small extend.

Our Contribution In this paper we describe an implementation of Fréchet distance range searching, with the goal of having a practically fast method. Our implementation is the winning submission to the ACM SIGSPATIAL GIS Cup 2017[2]. We also report on preliminary experiments.

Our algorithm consists of three phases. First, we determine all curves in the database whose start and endpoints are sufficiently close to the query curve's start and endpoints. Feature (2) ensures that most of these curves have small Fréchet distance to the query curve, and thus the number of listed curves is not much more than the output size. To efficiently enumerate these curves we use a multidimensional analogue of a quadtree. In the second phase, we use a greedy algorithm to filter out some of the curves that have sufficiently small Fréchet distance. We also have a negative filter that determines some of the curves with too large Fréchet distance. Only if both heuristics fail, we have to run an exact Fréchet distance decision algorithm, for which we use an novel recursive version of the standard free-space diagram algorithm.

For our GIS Cup submission we also added a simple parallelization over multiple queries, but in this paper we focus on answering a single query.

The first phase (quadtrees) is described in Section 3, the second phase (filters) follows in Section 4, and the third phase (exact algorithm) is discussed in Section 5. We conclude with preliminary experiments in Section 6.

2 Preliminaries

All curves considered in this paper are *polygonal*. We define a polygonal curve π by its vertices $(\pi_1, \ldots, \pi_n)$ with $\pi_i \in \mathbb{R}^2$ for all $i \in \{1, \ldots, n\}$. We denote by $|\pi| = n$ be the number of vertices of π. Moreover, we denote by $\|\pi\|$ the total length $\sum_{i=1}^{n-1} \|\pi_i - \pi_{i+1}\|$, where $\|.\|$ denotes the Euclidean norm. We write $\pi_{p..b}$ for the subcurve $(\pi_p, \pi_{p+1}, \ldots, \pi_b)$. By interpolating between vertices, we can also view π as a continuous function $\pi \colon [1, n] \to \mathbb{R}^2$ with $\pi_{i+\lambda} = (1 - \lambda)\pi_i + \lambda\pi_{i+1}$ for $i \in \{1, \ldots, n-1\}$ and $\lambda \in [0, 1]$.

Formally, the Fréchet distance is defined as follows. Let Φ_n be the set of all continuous and non-decreasing functions ϕ from $[0, 1]$ onto $[1, n]$. Then two curves π, σ have Fréchet distance

$$d_{\mathrm{F}}(\pi, \sigma) := \inf_{\substack{\phi_1 \in \Phi_{|\pi|} \\ \phi_2 \in \Phi_{|\sigma|}}} \max_{t \in [0,1]} \|\pi_{\phi_1(t)} - \sigma_{\phi_2(t)}\|.$$

We call $\phi := (\phi_1, \phi_2)$ a *traversal* of (π, σ).

3 A Quadtree Finds Candidates

Consider two curves π, σ. Since any traversal starts at both starting points, we have $\|\pi_1 - \sigma_1\| \leq d_{\mathrm{F}}(\pi, \sigma)$. Similarly, since any traversal ends at both endpoints, we have $\|\pi_{|\pi|} - \sigma_{|\sigma|}\| \leq d_{\mathrm{F}}(\pi, \sigma)$.

Denote by min-$x(\pi)$ the minimal x-coordinate of any vertex of curve π. Then $|\text{min-}x(\pi) - \text{min-}x(\sigma)| \leq d_{\mathrm{F}}(\pi, \sigma)$. Indeed, if min-$x(\sigma) \geq$ min-$x(\pi)$ then the vertex with minimal x-coordinate on π has distance at least min-$x(\sigma) -$ min-$x(\pi)$ to any vertex of σ, and the remaining case symmetrically yields the lower bound min-$x(\pi) -$ min-$x(\sigma)$. We obtain

---

three symmetric lower bounds by considering maximal x-coordinates (max-$x(\pi)$), minimal y-coordinates (min-$y(\pi)$), and maximal y-coordinates (max-$y(\pi)$). In total, we obtain:

$$d_{\mathrm{F}}(\pi, \sigma) \geq LB_{\mathrm{F}}(\pi, \sigma), \text{ with}$$
$$LB_{\mathrm{F}}(\pi, \sigma) := \max \big\{ \|\pi_1 - \sigma_1\|, \|\pi_{|\pi|} - \sigma_{|\sigma|}\|,$$
$$|\text{min-}x(\pi) - \text{min-}x(\sigma)|, |\text{max-}x(\pi) - \text{max-}x(\sigma)|,$$
$$|\text{min-}y(\pi) - \text{min-}y(\sigma)|, |\text{max-}y(\pi) - \text{max-}y(\sigma)| \big\}.$$

In the first phase of our algorithm we determine the candidate set $D' := \{\sigma \in D \mid LB_{\mathrm{F}}(\pi, \sigma) \leq \delta\}$. Note that for any remaining curve $\sigma \in D \setminus D'$ we have $d_{\mathrm{F}}(\pi, \sigma) \geq LB_{\mathrm{F}}(\pi, \sigma) > \delta$, so there are no false negatives.

In order to determine D', for each curve $\sigma \in D$ we store an 8-dimensional vector consisting of the coordinates for the start and endpoints as well as the largest and smallest coordinates in both dimensions. We store all these vectors in an 8-dimensional analogue of a quadtree. Given a query (π, δ), for the query curve π we also compute this 8-dimensional vector. Then using the quadtree we can enumerate all curves $\sigma \in D'$.

Because of feature (2) from the introduction, curves $\sigma \in D'$ typically also have $d_{\mathrm{F}}(\pi, \sigma)$ not much larger than δ. This means that this enumeration has not too many false positives in practical situations, i.e., there are not too many curves $\sigma \in D'$ with $d_{\mathrm{F}}(\pi, \sigma) > \delta$, see also the last line of Table 1.

On all enumerated curves $\sigma \in D'$, we first run the filters from the next section, and only if these heuristics fail we run the exact decision algorithm from Section 5. This determines for each curve $\sigma \in D'$ whether $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$ and thus solves the range searching problem.

4 Positive and Negative Filters

Let π, σ be curves with n, m vertices, respectively, and let $\delta > 0$. Our task is to decide whether $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$. We make use of the following heuristics.

Greedy Algorithm We first run a simple greedy algorithm, which was previously discussed in [6]: We start at $i = 1, j = 1$. When we are at (i, j), we either increment i or j or both, specifically we go to the pair $(i', j') \in \{(i+1, j), (i, j+1), (i+1, j+1)\}$ minimizing the distance $\|\pi_{i'} - \sigma_{j'}\|$. From these three possible steps we ignore the ones that would increase i beyond n or j beyond m. We end once we reach (n, m). If during the whole process we stayed within distance δ, then it follows that $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$. This easy and very efficient ($O(n + m)$ time) algorithm filters out some curves $\sigma \in D'$ with $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$.

Negative Filter Set $q_1 := 1$. For any $i \in \{2, \ldots, n\}$, let q_i be the first point on σ after q_{i-1} that is within distance δ of π_i, i.e., set $q_i := \min\{q \in [q_{i-1}, m] \mid \|\sigma_q - \pi_i\| \leq \delta\}$. Note that here q may be a real number, i.e., σ_q may be an interpolated point. An easy inductive proof shows that in any traversal that stays in distance δ, any point σ_z that is visited while being at π_i satisfies $z \geq q_i$, i.e., q_i is a lower bound on the position in σ while being at π_i. Hence, if it happens that some value q_i is undefined (since the set that we minimize over is empty), then we can conclude $d_{\mathrm{F}}(\pi, \sigma) > \delta$. This filters out some curves $\sigma \in D'$ that have too large Fréchet distance.

The above procedure can be implemented in time $O(n+m)$. However, the constant factors are relatively large, since in each step we compute the intersection of a circle ($\{x \in \mathbb{R}^2 \mid \|x - \pi_i\| = \delta\}$) and a line segment ($\{\sigma_q \mid q \in [j, j+1]\}$). Hence, we instead implemented a discrete version of this filter, where intuitively we aim to round down each q_i to an integer, so that we only consider vertices in both π and σ. Formally, we replace the definition of q_i by $\tilde{q}_i := \min\{j \in \{\tilde{q}_{i-1}, \ldots, m\} \mid \|\sigma_j - \pi_i\| - \|\sigma_j - \sigma_{j+1}\| \leq \delta\}$. Observe that we relax the distance constraint by the length of the current segment in σ, which indeed implies $\tilde{q}_i \leq q_i$. Thus, when it happens that $\tilde{q}_i$ is undefined, we can conclude that $d_{\mathrm{F}}(\pi, \sigma) > \delta$.

As this procedure is not symmetric in π and σ, we can run it a second time with their roles swapped.

5 Fréchet Distance Decider

It remains to handle the curves $\sigma \in D'$ reported by the first phase (quadtree) that are not filtered out in the second phase (filters). Preliminary experiments show that this third step dominates the running time compared to the first two phases, see Table 1, which is why we spent most of our development time on an optimized implementation of this phase.

We start by reviewing the standard free-space diagram algorithm for the Fréchet distance. Then we describe our intuition for an optimized implementation, and finally we present the details of our new algorithm. In the whole section, let π, σ be curves with n, m vertices, respectively, and let $\delta > 0$.

5.1 Standard Free-space Algorithm

The set $F := \{(p, q) \in [1, n] \times [1, m] \mid \|\pi_p - \sigma_q\| \leq \delta\}$, which describes the pairs (p, q) such that π_p and σ_q are in distance δ, is called the *free-space diagram* of π and σ. We also define the *reachable free-space* $R \subseteq F$ as the pairs (p, q) such that there exists an x- and y-monotone path through the free-space diagram from the lower left corner $(1, 1)$ to (p, q). These concepts are used in all known algorithms for the Fréchet distance, because of the fact that $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$ holds if and only if $(n, m) \in R$.

A subset $C_{i,j} := F \cap [i, i+1] \times [j, j+1]$ for $i \in \{1, \ldots, n-1\}, j \in \{1, \ldots, m-1\}$ is called a *free-space cell*. It is a well-known fact that each $C_{i,j}$ is the intersection of an ellipse with the box $[i, i+1] \times [j, j+1]$, in particular $C_{i,j}$ is convex. We denote by $F_{i,j}^l$ the left boundary of $C_{i,j}$, i.e., the interval $F \cap [i, i+1] \times \{j\}$, and similarly by $F_{i,j}^b$ the bottom interval $F \cap \{i\} \times [j, j+1]$. Similarly, $F_{i,j}^r := F_{i,j+1}^l$ and $F_{i,j}^t := F_{i+1,j}^b$. Any such interval can be computed in constant time by one intersection of a circle and a line segment.

Furthermore, we define the reachable subset $R_{i,j}^l := F_{i,j}^l \cap R$ and similarly $R_{i,j}^b, R_{i,j}^r, R_{i,j}^t$. Then $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$ holds if and only if $(n, m) \in R_{n-1,m-1}^t$. A simple procedure computes the output intervals $R_{i,j}^r$ and $R_{i,j}^t$ from the input intervals $R_{i,j}^l$ and $R_{i,j}^b$ in constant time. Since also the leftmost boundaries $R_{i,1}^l$ and bottommost boundaries $R_{1,j}^b$ can each be computed in constant time, we can determine all reachability intervals in total time $O(nm)$. This is the standard algorithm for deciding whether $d_{\mathrm{F}}(\pi, \sigma) \leq \delta$.

5.2 Intuition for our Algorithm

For intuition, think of π and σ being equal curves, specifically a line segment of length L subdivided n times. In this very simple case, the reachable free-space is a tube of the form

$R = \{(p, q) \in [1, n] \times [1, m] \mid |p - q| \leq \Delta\}$ with $\Delta = \delta n / L$. Our intuition is that up to a "smooth" transformation and some "noise" at the boundary of R, this case models practical input curves quite well.

Observe that in this specific case R consists of $\Theta(\Delta n)$ non-empty free-space cells. An algorithm enumerating all non-empty free-space cells thus takes time $\Omega(\Delta n)$. However, there are large blocks of the reachable free-space that are either completely filled (i.e., reachable) or completely empty (i.e., unreachable). Let B denote the number of "boundary" free-space cells, which are neither completely filled nor completely empty. In our exemplary situation we have $B = O(n)$. By an analogy to (standard 2-dimensional) quadtrees, it can be shown that there is a partitioning of $[1, n] \times [1, m]$ into $O(B)$ blocks of the form $[p, b] \times [q, d]$ such that each block is either completely filled or completely empty or a boundary free-space cell. The goal of our recursive algorithm is to compute such a partitioning. In the motivating example of π and σ being the same subdivided straight line, this complexity is $O(B) = O(n)$, which yields a speedup compared to the time $\Omega(\Delta n)$ of the standard algorithm for large Δ.

However, since checking whether a block is completely filled or completely empty is costly, we replace this step by a heuristic involving the triangle inequality. For this heuristic variant, we cannot prove a bound of $O(B)$ on size of the generated partitioning, but on practical data it behaves well.

In a nutshell, the idea is to identify large blocks of completely filled or completely empty reachable free-space, and handle them efficiently (in constant time).

5.3   Details of our Algorithm

We construct a partitioning of the reachable free-space recursively as follows. Given subcurves $\pi_{p..b}, \sigma_{q..d}$ and their input intervals $R_{i,q}^l$ for $i \in \{p, \ldots, b\}$ and $R_{p,j}^b$ for $j \in \{q, \ldots, d\}$, the recursive subproblem is to determine all output intervals $R_{i,d}^r$ for $i \in \{p, \ldots, b\}$ and $R_{b,j}^t$ for $j \in \{q, \ldots, d\}$. At the root of the recursion we have $[p, b] = [1, n]$ and $[q, d] = [1, m]$.

We now describe how to handle a recursive subproblem. Suppose that the lower left corner (p, q) is in the input intervals. Also suppose that the distances $\|\pi_i - \sigma_j\|$ are at most δ for all $i \in \{p, \ldots, b\}, j \in \{q, \ldots, d\}$. Then all free-space cells $C_{i,j}$ for $i \in \{p, \ldots, b-1\}, j \in \{q, \ldots, d-1\}$ are completely filled, and thus we are done with the subcurves $\pi_{p..b}, \sigma_{q..d}$. However, testing the condition requires iterating over all i's and j's and thus is too costly. Hence, we replace it with a cheaper condition: We test whether $\|\pi_p - \sigma_q\| + \|\pi_{p..b}\| + \|\sigma_{q..d}\| \leq \delta$, i.e., whether the distance between the startpoints of the current subcurves plus the curvelengths of the current subcurves is at most δ. Note that by triangle inequality this implies that all distances $\|\pi_i - \sigma_j\|$ are at most δ for any $i \in \{p, \ldots, b\}, j \in \{q, \ldots, d\}$, and thus the free-space spanned by $\pi_{p..b}$ and $\sigma_{q..d}$ is completely filled. Also, this check can be performed in constant time after the following preprocessing: For any curve in the database as well as for the query curve we store for each prefix $\pi_{1..i}$ the curvelength $\|\pi_{1..i}\|$. This allows us to determine the curvelength of $\pi_{p..b}$ as $\|\pi_{1..b}\| - \|\pi_{1..p}\|$ in constant time. This yields a quick test handling large completely filled blocks in constant time.

We use a similar test to determine whether the free-space spanned by $\pi_{p..b}$ and $\sigma_{q..d}$ is completely empty.

In case these tests do not apply, we split one of the two curves into two halves, say we split $\pi_{p..b}$ into $\pi_{p..r}$ and $\pi_{r..b}$ for $r = \lfloor (p+b)/2 \rfloor$. We then recursively determine the reachable free-space in the subcurves $(\pi_{p..r}, \sigma_{q..d})$ and then in $(\pi_{r..b}, \sigma_{q..d})$.

In case we reach a cell, i.e., $b = p+1$ and $d = q+1$, we handle it as in the usual free-space diagram algorithm. This finishes the description of our recursive algorithm.

The worst-case complexity of this algorithm is $O(nm \log(nm))$, however, this does not reflect the behaviour on realistic inputs. We leave it as an open problem for the theory community to explain the efficiency of our algorithm on a reasonable model of curves.

6 Experimental Results

A set of $\approx 20\,000$ trajectories was provided by the GIS Cup organizers as test data. They each have between 11 and 769 vertices. We performed preliminary experiments on this database on a MacBook Air with an 1,8 GHz Intel Core i5 and 8GB RAM. The time to read the database from the hard disc is about 9700ms, which dominates the time of 310ms to build our data structures.

On this database D, we generate random queries as follows. Fix an outputsize $k \in \{0, 1, 10, 100, 1000\}$. We pick a random curve $\pi \in D$. Using a binary-search-like procedure, we determine a random distance threshold δ such that there are exactly k curves in D within distance δ to π. (For some curves $\pi \in D$ such a threshold does not exist since there are multiple curves with the same distance to π; in this case we ignore π and pick a new random curve.) We then measure the running time for the query (π, δ).

The results are shown in Table 1, where each entry gives the mean and standard deviation over 500 such random queries. The first line shows the time per query. Fitting this first line to the model $a + k^b \cdot c$ yields an approximation of $\approx 0.44\text{ms} + k^{0.79} \cdot 0.27\text{ms}$ for the average query time in terms of the outputsize k.

The next four lines show the time spent for our main components. Note that for $k \geq 10$ the dominant part is the exact decision procedure, while for smaller k our solution could possibly benefit from an improved quadtree implementation.

The next lines show the average query time if we (1) remove the greedy filter, (2) remove the negative filter, or (3) replace our novel decision procedure with the standard decision algorithm for the Fréchet distance. Note that all components are necessary as the running time increases by at least a factor ≈ 2 – except for the negative filter, but we believe that the latter is benefitial for other types of data.

Finally, the last line of Table 1 shows the number p of false positives of our quadtree data structure, i.e., while the correct output size is k the quadtree returns $k + p$ candidates. Note that this number p is very small on average compared to k, showing that our lower bound $LB_\text{F}(\pi, \sigma)$ and the resulting quadtree data structure are very effective. (Recall that these false positives are caught by our filters or exact decider; the complete algorithm has no false positives or negatives.)

7 Conclusion

We described a novel implementation for near-neighbors queries with respect to the Fréchet distance, which won the ACM SIGSPATIAL GIS Cup 2017. It remains to incorporate this or a similarly fast implementation into one of the standard libraries for geographic information systems, in order to make it usable in a wide range of applications.

We leave it as an open problem to theoretically explain the running time of our algorithm on a suitable model of realistic input curves. The biggest issue for this is that it seems like our algorithm uses some features that are present in many realistic input curves, but which are not guaranteed by the standard models of realistic input curves used in computational

geometry. One such property is that curves with similar start- and endpoints have small Fréchet distance.

References

[1] Helmut Alt and Maike Buchin. Can we compute the similarity between surfaces? *Discrete & Computational Geometry*, 43(1):78–99, 2010.

[2] Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5(1–2):78–99, 1995.

[3] Helmut Alt, Christian Knauer, and Carola Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2004.

[4] Boris Aronov, Sariel Har-Peled, Christian Knauer, Yusu Wang, and Carola Wenk. Fréchet distance for curves, revisited. In *Proc. 14th Annu. European Symp. Algorithms (ESA'06)*, pages 52–63. 2006.

[5] 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.

[6] Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete Fréchet distance. *JoCG*, 7(2):46–76, 2016.

[7] 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.

[8] Kevin Buchin, Maike Buchin, and Yusu Wang. Exact algorithms for partial curve matching via the Fréchet distance. In *Proc. 20th Annu. ACM-SIAM Symp. Discrete Algorithms (SODA'09)*, pages 645–654, 2009.

[9] Erin Wolf Chambers, Éric Colin de Verdière, Jeff Erickson, Sylvain Lazard, Francis Lazarus, and Shripad Thite. Homotopic Fréchet distance between curves or, walking your dog in the woods in polynomial time. *Computational Geometry*, 43(3):295–311, 2010.

[10] Anne Driemel, Sariel Har-Peled, and Carola Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete & Comp. Geom.*, 48(1):94–127, 2012.

[11] Michael Godau. A natural metric for curves - computing the distance for polygonal chains and approximation algorithms. In *Proc. 8th Sympos. Theoret. Aspects Comput. Sci. (STACS'91)*, pages 127–136, 1991.

output size k:	0	1	10	100	1000
total time:	0.44 ± 0.12 ms	0.72 ± 0.47 ms	2.00 ± 2.11 ms	11.58 ± 15.97 ms	60.44 ± 92.25 ms
time for quadtree:	0.42 ± 0.10 ms	0.44 ± 0.10 ms	0.45 ± 0.09 ms	0.55 ± 0.12 ms	0.77 ± 0.14 ms
time for greedy filter:	0.01 ± 0.01 ms	0.02 ± 0.02 ms	0.07 ± 0.06 ms	0.53 ± 0.32 ms	5.39 ± 2.96 ms
time for negative filter:	0.01 ± 0.02 ms	0.04 ± 0.06 ms	0.16 ± 0.20 ms	0.92 ± 1.32 ms	4.44 ± 7.33 ms
time for exact decider:	0.01 ± 0.05 ms	0.23 ± 0.40 ms	1.31 ± 1.87 ms	9.58 ± 14.44 ms	49.77 ± 83.39 ms
total without greedy filter:	1.02 ± 0.48 ms	1.52 ± 0.86 ms	5.93 ± 4.25 ms	46.55 ± 37.76 ms	345.42 ± 290.81 ms
total without negative filter:	0.44 ± 0.14 ms	0.78 ± 0.66 ms	2.56 ± 3.14 ms	12.00 ± 15.71 ms	65.74 ± 94.74 ms
total without exact decider:	0.60 ± 1.42 ms	2.58 ± 4.90 ms	12.20 ± 21.72 ms	91.16 ± 163.71 ms	571.88 ± 1114.99 ms
# false positives of quadtree:	0.16 ± 0.53	1.06 ± 1.76	4.05 ± 5.92	9.63 ± 13.48	69.43 ± 171.58

Table 1: Experimental evaluation on the GIS Cup example dataset.