On Flipping the Fréchet distance

Omrit Filtser[1], Mayank Goswami[*2], Joseph S.B.Mitchell[†3], and Valentin Polishchuk[4]

[1]The Open University of Israel, omrit.filtser@gmail.com
[2]Queens College CUNY, mayank.goswami@qc.cuny.edu
[3]Stony Brook University, joseph.mitchell@stonybrook.edu
[4]Linköping University, valentin.polishchuk@liu.se

Abstract

The classical and extensively-studied *Fréchet distance* between two curves is defined as an *inf max*, where the infimum is over all traversals of the curves, and the maximum is over all concurrent positions of the two agents. In this article we investigate a "flipped" Fréchet measure defined by a *sup min* – the supremum is over all traversals of the curves, and the minimum is over all concurrent positions of the two agents. This measure produces a notion of "social distance" between two curves (or general domains), where agents traverse curves while trying to stay as far apart as possible.

We draw connections between our proposed flipped Fréchet measure and existing related work in computational geometry, hoping that our new measure may spawn investigations akin to those performed for the Fréchet distance, and into further interesting problems that arise.

1 Introduction

The classical Fréchet distance between two curves P and Q is defined as the minimum length of a leash required for a person to walk their dog, with the person and the dog traversing P and Q from start to finish, respectively. Inspired by the challenge of maintaining social distancing among groups and individuals, we consider the question of developing a notion opposite to the Fréchet distance, where instead of keeping the agents close (short leash), we keep them as far apart as possible.

In this paper we propose a new measure, called the *Flipped Fréchet* measure, to capture the amount of social distancing possible while traversing two curves. While Fréchet distance is defined as an *inf max*, where the infimum is over all traversals of the curves, and the maximum is over all concurrent positions of the two agents, the flipped Fréchet measure[1] is defined as a *sup min* – the supremum is over all traversals of the curves, and the minimum is over all concurrent positions of the two agents. How efficiently can this measure be computed, for curves in one or two dimensions? What if the two agents are walking on edges of a graph, which may or may not be embedded in the plane? Such questions have been considered for Fréchet distance, and in this paper we initiate their study for the flipped Fréchet measure.

We refer to the two agents as "Red" and "Blue" henceforth. Considering the social distancing problem further, what if Blue is not restricted to move along some given curve; rather, it can choose its own path? We now start arriving at a class of problems that have no analogues in the Fréchet version. Of course, if Blue had no restrictions at all, it could just go to infinity and thus be far from Red (on any path). It therefore makes sense to restrict the domain for Blue, e.g. to a simple polygon P in which Red is traveling, and measure separation using geodesic distance in P. We consider questions regarding the complexity of

*This work is supported by US National Science Foundation (NSF) awards CRII-1755791 and CCF-1910873.

†This work is partially supported by the National Science Foundation (CCF-2007275), the US-Israel Binational Science Foundation (BSF project 2016116), Sandia National Labs, and DARPA (Lagrange).

[1]One observes that this measure is not a metric/distance as it does not satisfy the triangle inequality.

calculating a strategy for Blue to stay away from Red, when Red is traveling on a given path, which may or not be a geodesic in P.

An architect designing spaces within a building is faced with choices about the shapes of these spaces, where one must choose, say, between two polygons P and Q where agents will move, in hopes to maximize the potential for social distancing. In order to do so, it would be useful to have a notion of a "social distancing width" of a polygon, that captures the difficulty or ease with which two agents can move around in a polygon while maintaining separation. Consider a simple polygon P where the Red agent is on a mission to follow a path, e.g. to traverse the boundary of P, while the Blue agent moves within P (with a starting point of Blue's choice), in order to maximize the minimum Red-Blue distance. We define the *social distance width* (SDW for short) of a polygon P to be the minimum Red-Blue distance that can be maintained throughout the movement, maximized over all possible movement strategies, and study algorithms to compute the SDW of a polygon.

For all of the above problems, in addition to developing algorithms for the general versions, we also consider special scenarios which facilitate faster algorithms; for example, while our algorithm for computing the SDW for general polygons runs in quadratic time, we show that for skinny polygons (or a tree), one can compute the SDW in linear time.

Although this article mostly considers the above problems in the case of $k = 2$ agents, in general one may be given k agents and k associated domains. Each agent is restricted to move only within its respective domain, and at least one of the agents has some *mission*, e.g., to move from a given start point to a given end point, or to traverse a given path inside the domain. In addition, the domains may be shared or distinct, and different agents may have different speeds. The goal is to find a movement strategy for all the agents, such that the *minimum pairwise* distance between the agents at any time is maximized. Additionally, one may seek to minimize the time necessary to complete one or more missions.

This new class of problems is different from the usual motion planning problems between robots, or disjoint disks, in some fundamental aspects. Most, if not all, literature on robot motion planning assumes robots are cooperating on some task. One then considers optimizing objectives like makespan, or total distance travelled, etc. However, the kind of movement we consider is far from cooperative – in fact, some agents may not care about social distancing, while others do. Some may be "on a mission" while others are just trying to maintain a safe distance. In addition, different agents may have different starting times and deadlines. Furthermore, as mentioned above, one also encounters design problems, where one may want to configure a layout of a building, a floor plan, or designate rules for traffic flow, in order to facilitate social distancing.

2 Polygonal curves

We begin by considering the scenario in which the two domains are polygonal curves P and Q. The agents' missions are to traverse their respective curves, from the start point to the end point, in order to maximize the minimum distance between the agents. The Flipped Fréchet measure between the two curves is the maximum separation that can be maintained. Formally, let $P : [1, n] \to \mathbb{R}^d$ and $Q : [1, m] \to \mathbb{R}^d$ be two polygonal curves. A *traversal* of P and Q is a pair of continuous, non-decreasing, surjective functions $f : [0, 1] \to [1, n]$ and $g : [0, 1] \to [1, m]$.

Definition 1 (Flipped Fréchet Measure)**.** *The* Flipped Fréchet measure *(FF) of P and Q is* $FF(P, Q) = \sup_{\tau=(f,g)} \min_{t\in[0,1]} \|P(f(t)) - Q(g(t))\|$.

Note that the well-studied *Fréchet distance* between P and Q is $\inf_{\tau=(f,g)} \max_{t\in[0,1]} \|P(f(t))-Q(g(t))\|$, where f, g are a traversal of P and Q. We consider both the continuous case (agents move continuously along the edges of their curves), and the discrete case (agents "jump" between consecutive vertices of their curves).

We first show that the Flipped Fréchet measure between two n-vertex curves in one dimension (1D) can be computed in near-linear time, demonstrating that "flipping" the objective function makes this setting easier: for continuous Fréchet there exist conditional (SETH-based) quadratic lower bounds [BM16].

Theorem 2. *Given two polygonal curves P, Q of length n in 1D, their social distance width, $FF(P, Q)$, can be computed in $O(n \log^2 n)$ time.*

Next, we give quadratic algorithms and then conditional lower bounds for computing or approximating other variants (1D discrete, 2D continuous and discrete) of FF measure, specifically:

Theorem 3. *There exists*

- *a quadratic lower bound, conditioned on the Strong Exponential Time Hypothesis (SETH), on approximating FF measure for curves in 2D, with approximation factor $\frac{\sqrt{5}}{2\sqrt{2}}$, and*
- *a quadratic lower bound, conditioned on the Strong Exponential Time Hypothesis (SETH), on approximating dFF measure for 1D curves, with approximation factor $\frac{2}{3}$.*

Our algorithm for continuous FF in 1D generalizes to any number $k \geq 2$ of agents, with a running time of $O(kn \log n)$. The question of whether or not there exists an algorithm in 2D with running time fully polynomial in k remains open (for the Fréchet distance of a set of curves, the best known running time is roughly $O(n^k)$; see [DR04]).

3 SDW of polygons

We then consider distancing problems in which the given domain (for both Red and Blue) is a simple polygon. Since the two agents are moving inside the same polygon, it is natural to consider geodesic distance (i.e., the shortest path inside the polygon) instead of Euclidean distance to measure separation.

Consider a scenario in which Red and Blue have to traverse two polygonal paths R and B, both inside a given polygon P, and their goal is to find a movement strategy (a traversal) that maintains geodesic distance of at least δ between them. For the analogous Fréchet problem (Red and Blue have to maintain geodesic distance of at most δ), Cook and Wenk [CW10] presented an algorithm that runs in $O(n^2 \log N)$ time, where N is the complexity of P and n is the complexity of R and B. A "flipped" version of their algorithm can be applied for computing the $FF(R, B)$ under geodesic distance in nearly quadratic time.

Closed curves and polygons. Consider a scenario in which the polygonal curves R and B are closed curves. Here, the starting points of Red and Blue are not given as an input, and the goal is to decide whether they can traverse their respective curves while maintaining distance at least δ. The analogous Fréchet problem has been investigated ([AG95],[SVY14]), and again similar "flipped" versions of these near quadratic time algorithms can be applied to FF.

We can then define the Social Distance Width of two polygons P_1, P_2 as a special case in which R is the boundary of P_1 and B is the boundary of P_2; i.e., $SDW(P_1, P_2) = SDW(\partial P_1, \partial P_2)$. Similarly, the Social Distance Width of a (single) polygon P is $SDW(P) = FF(\partial P, \partial P)$. We show

Theorem 4. *The social distance width of a polygon P of n vertices can be computed in$O(n^2)$ time.*

The notion of SDW of a polygon is possibly related to other characteristics of polygons, such as fatness. Intuitively, if the polygon P is fat under standard definitions, then the SDW of P will be large. However, the exact connection is yet unclear, and we leave open the question of what exactly is the relation between the two definitions.

When both Red and Blue are restricted to traverse a given path, it seems that the Fréchet-like nature of the problem leads to near-quadratic time algorithms. Thus, we consider the scenario where Blue has more freedom, and it is not required to traverse a given path.

Red on an arbitrary path mission. We first consider the case when Red moves along an arbitrary path R in P, and Blue may wander around in P, starting from some given point b. The free-space diagram can be adapted to the case of a path and a polygon, by partitioning the polygon into a linear number of convex cells (for example, a triangulation). This is a three-dimensional structure, which contains $O(n^2)$ cells (assuming that the complexity of both R and P is $O(n)$). However, for maintaining geodesic separation, building the free-space may be cumbersome, because it would involve building the parametric shortest path map in a triangle as the source moves along a segment, and such SPM may have $\Omega(n)$ combinatorial changes. Instead, we show that Blue may stay on the boundary, thus reducing the problem to the standard free-space diagram between a path and a closed curve. We show:

Theorem 5. *Let P be a polygon with n vertices, b a point in P, and R a path between two points r and r' in P. There exists an $O(n^2)$-time algorithm to decide whether there exists a path B in P starting from b, such that $FF(R, B) > 1$ under geodesic distance.*

We now show that in special cases, the SDW can be computed in linear-time.

Red on a shortest path mission. Assume that Red moves along a geodesic path R in P (Red is on a mission and does not care about social distancing) while Blue may wander around anywhere within P starting from a given point b. We show that the decision problem, whether Blue can maintain (geodesic) social distance at least 1 from Red, can be solved in linear time.

Theorem 6. *Let P be a polygon with n vertices, b a point in P, and R a geodesic shortest path between two points r and r' in P. There exists an $O(n)$-time algorithm to decide whether there exists a path B in P starting from b, such that $FF(R, B) > 1$ under geodesic distance.*

Social distancing in a skinny polygon (or a tree). We then consider the case in which the shared domain of Red and Blue is a tree T, and the distance is the shortest-path distance in the tree (the distance between vertices u and v denoted $|uv|$). Red moves around T in a depth-first fashion: there is no start and end point, it keeps moving ad infinitum. In particular, if T is embedded in the plane, the motion is the limiting case of moving around the boundary of an infinitesimally thin simple polygon, and the distance is the geodesic distance inside the polygon.

Theorem 7. *Let T be a tree with n vertices, embedded in the plane, and R be a traversal of T in a depth-first fashion. There exists an $O(n)$-time algorithm to find a path B in T that maximizes $SDW(R, B)$ under the geodesic distance.*

The discussion on this special case of thin polygons, leads us to investigating the more general case of Social Distance Width of two graphs, which we discuss in the full version of the paper.

References

[AG95] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.*, 5:75–91, 1995, doi:10.1142/S0218195995000064. 3

[BM16] K. Bringmann and W. Mulzer. Approximability of the discrete Fréchet distance. *JoCG*, 7(2):46–76, 2016, doi:10.20382/jocg.v7i2a4. 3

[CW10] A. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms (TALG)*, 7(1):1–19, 2010. 3

[DR04] A. Dumitrescu and G. Rote. On the fréchet distance of a set of curves. In *Proceedings of the 16th Canadian Conference on Computational Geometry, CCCG'04, Concordia University, Montréal, Québec, Canada, August 9-11, 2004*, pages 162–165, 2004. 3

[SVY14] M. I. Schlesinger, E. Vodolazskiy, and V. Yakovenko. Similarity of closed polygonal curves in Fréchet metric. *arXiv preprint arXiv:1409.4613*, 2014. 3