

# Sampling-based bottleneck pathfinding with applications to Fréchet matching \*

Kiril Solovey      Dan Halperin

Blavatnik School of Computer Science, Tel Aviv University, Israel

## Abstract

We describe a general probabilistic framework to address a variety of Fréchet-distance optimization problems. Specifically, we are interested in finding minimal *bottleneck*-paths in  $d$ -dimensional Euclidean space between given start and goal points, namely paths that minimize the maximal value over a continuous cost map. We present an efficient and simple sampling-based framework for this problem, which is inspired by, and draws ideas from, techniques for robot motion planning. We extend the framework to handle not only standard bottleneck pathfinding, but also the more demanding case, where the path needs to be monotone in all dimensions. Finally, we provide experimental results of the framework on several types of problems.

## 1 Introduction

This paper studies the problem of finding near-optimal paths in  $d$ -dimensional Euclidean space. Specifically, we are interested in *bottleneck* paths which minimize the maximal value the path obtains over a generally-defined continuous cost map. As an example, suppose that one wishes to plan a hiking route in a mountainous region between two camping grounds, such that the highest altitude along the path is minimized [17]. In this case, the map assigns to each two-dimensional point its altitude. A similar setting, albeit much more complex, requires to find a pathway of low energy for a given protein molecule (see, e.g., [37]).

Our main motivation for studying bottleneck optimization over cost maps is its tight relation to the Fréchet distance (or matching), which is a popular and widely studied similarity measure in computational geometry. The problem has applications to various domains such as path simplification [19], protein alignment [27], handwritten-text search [48], and signature verification [53]. The Fréchet distance, which was initially defined for curves, is often considered to be a more informative measure than the popular *Hausdorff* distance as it takes into consideration not only each curve as a whole but also the location and the ordering of points along it. Usually one is interested not only in the Fréchet distance between two given curves, but also in the parametrization which attains the optimal alignment.

---

\*This work has been supported in part by the Israel Science Foundation (grant no. 1102/11), by the Blavatnik Computer Science Research Fund, and by the Hermann Minkowski–Minerva Center for Geometry at Tel Aviv University. Kiril Solovey is also supported by the Clore Israel Foundation.

Since its introduction by Alt and Godau [2] in 1995, a vast number of works has been devoted to the subject, and many algorithms have been developed to tackle various settings of the problem. However, from a practical standpoint the problem is far from being solved: for many natural extensions of the Fréchet problem only prohibitively-costly algorithms are known. Furthermore, in some cases it was shown, via hardness proofs, that efforts for finding polynomial-time algorithms are doomed to fail. For some variants of the problem efficient algorithms are known to exist, however their implementation requires complex geometric machinery that relies on geometric kernels with infinite precision [30].

**Contribution.** We describe a generic, efficient and simple algorithmic framework for solving pathfinding optimization problems over cost maps. The framework is inspired by, and draws ideas from, sampling-based methods for robot motion planning. We provide experimental results of the framework on various scenarios. Furthermore, we theoretically analyze the framework and show that the cost of the obtained solution converges to the optimum, as the number of samples increases. We also consider the more demanding case, where paths need to be monotone in all dimensions.

**Organization.** In Section 2 we review related work. In Section 3 we provide a formal definition of the bottleneck pathfinding problem. In Section 4 we describe an algorithmic framework for solving this problem. In Section 5 we provide an analysis of the method. Finally, in Section 6 we report on experimental results.

## 2 Related work

This section is devoted to related work on Fréchet distance and robot motion planning.

### 2.1 Fréchet distance

The *Fréchet distance* between two curves is often described by an analogy to a person walking her dog: each of the two creatures is required to walk along a predefined path and the person wishes to know the length of the shortest *leash* which will make this walk possible. In many cases one also likes to know how to advance along the path given the short leash.

Formally, let  $\sigma_1, \sigma_2 : [0, 1] \rightarrow \mathbb{R}^d$  be two continuous curves. We wish to find a traversal along the two curves which minimizes the distance between the two traversal points. The traversal is defined by two continuous parametrizations  $\alpha_1, \alpha_2 : [0, 1] \rightarrow [0, 1]$  of  $\sigma_1, \sigma_2$  respectively, where for a given point in time  $\tau \in [0, 1]$ , the positions of the person and her dog are specified by  $\sigma_1(\alpha_1(\tau))$  and  $\sigma_2(\alpha_2(\tau))$ , respectively. The *Fréchet distance* between  $\sigma_1, \sigma_2$  is defined by the expression

$$\min_{\alpha_1, \alpha_2: [0,1] \rightarrow [0,1]} \max_{\tau \in [0,1]} \|\sigma_1(\alpha_1(\tau)) - \sigma_2(\alpha_2(\tau))\|_2.$$

Alt and Godau [2] described an  $O(n^2 \log n)$ -time algorithm for the setting of two polygonal curves, where  $n$  is the number of vertices in each of the two curve. Buchin et al. [12] described a different method for solving this problem for the same running time. Recently, Buchin et al. [10] developed an algorithm with a slightly improved running time  $O(n^2 \log^2 \log n)$ . Har-Peled and Raichel [22] introduced a simpler randomized algorithm with running time of  $O(n^2 \log n)$ . Bringmann [6] showed that an algorithm

with running time of  $O(n^{2-\delta})$ , for some constant  $\delta > 0$ , does not exist, unless a widely accepted conjecture, termed SETH [25], is wrong. In a following work [7] this conditional lower bound was extended to  $(1 + \varepsilon)$ -approximation algorithms of the Fréchet problem, where  $\varepsilon \leq 0.399$ .

The notion of Fréchet distance can be extended to  $k$  curves in various ways. One natural extension can be described figuratively as having a pack of  $k$  dogs, where each of the dogs has to walk along a predefined path, and every pair of dogs is connected with a leash. The goal now is to find a parametrization which minimizes the length of the longest leash. Dumitrescu and Rote [18] introduced a generalization of the Alt-Godau algorithm to this case, which runs in  $O(kn^k \log n)$  time, i.e., exponential in the number of input curves. They also describe a 2-approximation algorithm with a much lower running time of  $O(k^2 n^2 \log n)$ . In the work of Har-Peled and Raichel [22] mentioned above they also consider the case of  $k$  input curves and devise an  $O(n^k)$  algorithm. Notably, their technique is flexible enough to cope with different Fréchet-type goal functions over the  $k$  curves. Furthermore, their algorithm is also applicable when the  $k$  curves are replaced with  $k$  *simplicial complexes*, and the problem is to find  $k$  curves—one in each complex—which minimize the given goal function. A recent work [9], which extends the conditional lower bound mentioned earlier for the setting of multiple curves, suggests that a running time that is exponential in the number of curves is unavoidable.

The notion of Fréchet distance can be generalized to more complex objects. Buchin et al. [13] considered the problem of finding a mapping between two simple polygons, which minimizes the maximal distance between a point and its image in the other polygon. More formally, given two simple polygons  $P, Q \subset \mathbb{R}^2$  the problem consists of finding a mapping  $\delta : P \rightarrow Q$  which minimizes the expression  $\max_{p \in P} \|p - \delta(p)\|_2$ , subject to various constraints on  $\delta$ . They introduced a polynomial-time algorithm for this case. In a different paper, Buchin et al. [11] showed that the decision problem is NP-hard for more complex geometric objects, e.g., pairs of polygons with holes in the plane or pairs of two-dimensional terrains. Another interesting NP-hard problem that was studied by Sherette and Wenk [41] is *curve embedding* in which one wishes to find an embedding of a curve in  $\mathbb{R}^3$  to a given plane, which minimizes the Fréchet distance with the curve. In a similar setting Meulemans [34] showed that it is NP-hard to decide whether there exists a simple cycle in a plane-embedded graph that has at most a given Fréchet distance to a simple closed curve.

The Fréchet distance between curves in the presence of obstacles have earned some attention. Cook and Wenk [16] studied the *geodesic* variant, which consists of a simple polygon and two polygonal curves inside it. As in the standard formulation, the main goal is to minimize the length of the leash, but now the leash may wrap or bend around obstacles. Their algorithm has running time of  $O(m + n^2 \log mn \log n)$ , where  $m$  is the complexity of the polygon and  $n$  is defined as the total complexity of the two curves, as before. The more complex *homotopic* setting is a special case of the aforementioned geodesic setting, with the additional constraint that the leash must continuously deform. Chambers et al. [14] considered this problem for the specific setting of two curves in planar environment with polygonal obstacles. They developed an algorithm whose running time is  $O(N^9 \log N)$ , where  $N = n + m$  for  $n$  and  $m$  as defined above.

## 2.2 Motion planning

Motion planning is a fundamental problem in robotics. In its most basic form, the problem consists of finding a collision-free path for a robot  $\mathcal{R}$  in a workspace environment

$\mathcal{W}$  cluttered with obstacles. Typically, the problem is approached from the configuration space  $\mathcal{C}$ —the set of all robot configurations. The problem can be reformulated as finding a continuous curve in  $\mathcal{C}$ , which entirely consists of collision-free configurations and represents a path for the robot from a given start configuration to another, target, configuration. An important attribute of the problem is the number of *degrees of freedom* of  $\mathcal{R}$ , using which one can specify every configuration in  $\mathcal{C}$ . Typically the dimension of  $\mathcal{C}$  equals the number of degrees of freedom.

For some cases of the problem, which involve a small number of degrees of freedom, efficient and exact analytical techniques exist (see, e.g., [4, 21, 40]), which are guaranteed to find a solution if one exists, or report that none exists otherwise. Recently, it was shown [1, 46, 50] that efficient and complete techniques can be developed for the *multi-robot* motion-planning problem, which entails many degrees of freedom, by making several simplifying assumptions on the separation of the start and target positions. However, it is known that the general setting of the motion-planning problem is computationally intractable (see, e.g., [23, 38, 43, 47]) with respect to the number of degrees of freedom.

Sampling-based algorithms for motion planning, which were first described about two decades ago, have revolutionized the field of robotics by providing simple yet effective tools to cope with challenging problems involving many degrees of freedom. Such algorithms (see, e.g., PRM by Kavraki et al. [29], RRT by Kuffner and LaValle [32], and EST by Hsu et al. [24]) explore the high-dimensional configuration space by random sampling and connecting nearby samples, which result in a graph data structure that can be viewed as an approximation of the *free space*—a subspace of  $\mathcal{C}$ , which consists entirely of collision-free configurations. While such techniques have weaker theoretical guarantees than analytical methods, many of them are *probabilistically complete*, i.e., guaranteed to find a solution if one exists, given sufficient processing time. More recently, *asymptotically optimal* sampling-based algorithms, whose solution converges to the optimum, for various criteria, have started to emerge: Karaman and Frazzoli introduced the RRT\* and PRM\* [28] algorithms, which are asymptotically optimal variants of RRT and PRM. Following their footsteps Arslan and Tsiotras introduced RRT# [3]. A different approach was taken by Janson and Pavone who introduced the FMT\* algorithm [26], which was later refined by Salzman and Halperin [39].

### 3 Problem statement

In this section we describe the general problem of bottleneck pathfinding over a given cost map, to which we describe an algorithmic framework in Section 4. We conclude this section we several concrete examples of the problems that will be used for experiments in Section 6.

We start with several basic definitions. Given  $x, y \in \mathbb{R}^d$ , for some fixed dimension  $d \geq 2$ , let  $\|x - y\|_2$  denote the Euclidean distance between two points. Denote by  $\mathcal{B}_r(x)$  the  $d$ -dimensional Euclidean ball of radius  $r > 0$  centered at  $x \in \mathbb{R}^d$  and  $\mathcal{B}_r(\Gamma) = \bigcup_{x \in \Gamma} \mathcal{B}_r(x)$  for any  $\Gamma \subseteq \mathbb{R}^d$ . We will use the terms “path” and “curve” interchangeably, to refer to a continuous curve in  $\mathbb{R}^d$  parametrized over  $[0, 1]$ . Given a curve  $\sigma : [0, 1] \rightarrow \mathbb{R}^d$  define  $\mathcal{B}_r(\sigma) = \bigcup_{\tau \in [0, 1]} \mathcal{B}_r(\sigma(\tau))$ . Additionally, denote the image of a curve  $\sigma$  by  $\text{Im}(\sigma) = \bigcup_{\tau \in [0, 1]} \{\sigma(\tau)\}$ . Let  $A_1, A_2, \dots$  be random variables in some probability space and let  $B$  be an event depending on  $A_n$ . We say that  $B$  occurs *almost surely* (a.s., in short) if  $\lim_{n \rightarrow \infty} \Pr[B(A_n)] = 1$ .

Let  $\mathcal{M} : [0, 1]^d \rightarrow \mathbb{R}$  be a *cost map* that assigns to each point in  $[0, 1]^d$  a real value. For simplicity, we assume that the domain of  $\mathcal{M}$  is a  $d$ -dimensional unit hypercube. Let  $S, T \in [0, 1]^d$  denote the start and target points. Denote by  $\Sigma(S, T)$  the collection of paths that start in  $S$  and end in  $T$ . Formally, every  $\sigma \in \Sigma(S, T)$  is a continuous path  $\sigma : [0, 1] \rightarrow [0, 1]^d$ , where  $\sigma(0) = S, \sigma(1) = T$ . Given a path  $\sigma$  we use the notation  $\mathcal{M}(\sigma) = \max_{\tau \in [0, 1]} \mathcal{M}(\sigma(\tau))$  to represent its bottleneck cost.

In some applications, monotone paths are desired. For instance, in the classical problem of Fréchet matching between two curves it is often the case that backward motion along the curves is forbidden. Here we consider monotonicity in all  $d$  coordinates of points along the path. Formally, given two points  $p, p' \in \mathbb{R}^d$ , where  $p = (p_1, \dots, p_d), p' = (p'_1, \dots, p'_d)$ , we use the notation  $p \preceq p'$  to indicate that  $p_i \leq p'_i$ , for every  $1 \leq i \leq d$ . A path  $\sigma \in \Sigma(S, T)$  is said to be *monotone* if for every  $0 \leq \tau \leq \tau' \leq 1$  it holds that  $\sigma(\tau) \preceq \sigma(\tau')$ .

**Definition 1.** Given the triplet  $\langle \mathcal{M}, S, T \rangle$ , the *bottleneck-pathfinding problem* (BPP, for short) consists of finding a path  $\sigma \in \Sigma(S, T)$  which minimizes the expression  $\max_{\tau \in [0, 1]} \mathcal{M}(\sigma(\tau))$ . A special case of the bottleneck pathfinding problem, termed *strong-BPP*, requires that the path will be *monotone*.

### 3.1 Examples

We provide three examples of BPPs, which will be used for experiments in Section 6. Each example is paired with the  $d$ -dimensional configuration space  $\mathcal{C} := [0, 1]^d$ , start and target points  $S, T \in \mathcal{C}$ , and a cost map  $\mathcal{M} : [0, 1]^d \rightarrow \mathbb{R}$ . The examples below are defined for two-dimensional input objects, but can be generalized to higher dimensions.

**Problem 1:** We start with the classical *Fréchet distance among  $k$  curves* (see, e.g., [22]). Let  $\sigma_1, \dots, \sigma_k : [0, 1] \rightarrow [0, 1]^2$  be  $k$  continuous curves embedded in Euclidean plane. Here  $\mathcal{C} = [0, 1]^k$  is defined as the Cartesian product of the various positions along the  $k$  curves. Namely, a point  $P = (p_1, \dots, p_k) \in \mathcal{C}$  describes the location  $\sigma_i(p_i)$  along  $\sigma_i$ , for each  $1 \leq i \leq k$ . To every such  $P$  we assign the cost  $\mathcal{M}(P) = \max_{1 \leq i < j \leq k} \|\sigma_i(p_i) - \sigma_j(p_j)\|_2$ . We note that more complex formulations of  $\mathcal{M}$  can be used, depending on the exact application. The start and target positions are defined to be  $S = (\sigma_1(0), \dots, \sigma_k(0)), T = (\sigma_1(1), \dots, \sigma_k(1))$ .

**Problem 2:** We introduce the problem of *Fréchet distance with visibility*, whose basis is similar to **P1** with  $k = 3$ . In addition to the curves, we are given a subspace  $\mathcal{F} \subseteq [0, 1]^2$ . The goal is to find a traversal of the curves which minimizes  $\mathcal{M}$  as defined in **P1**, with the additional constraint that the traversal point along  $\sigma_1$  must be “seen” by one of the traversal points of  $\sigma_2, \sigma_3$ . Formally, for every  $P = (p_1, p_2, p_3) \in \mathcal{C}$  it must hold that  $p_1 p_2 \subset \mathcal{F}$  or  $p_1 p_3 \subset \mathcal{F}$  (but not necessarily both), where  $p_i p_j$  is the straight-line path from  $p_i$  to  $p_j$ .

**Problem 3:** In *curve embedding* (see, [34, 41]), the input consists of a curve  $\sigma : [0, 1] \rightarrow [0, 1]^2$ , a subspace  $\mathcal{F} \subseteq [0, 1]^2$  and a pair of two-dimensional points  $s, t \in \mathcal{F}$ . A point  $P = (p_1, p_2, p_3) \in \mathcal{C} = [0, 1]^3$  describes the location  $\sigma(p_1)$  along  $\sigma$  and the point  $(p_2, p_3) \in \mathcal{F}$ . The BPP is defined for the start and target points  $S = (0, s), T = (1, t) \in \mathcal{C}$  and the cost map  $\mathcal{M}(P) = \|\sigma(p_1) - (p_2, p_3)\|_2$ .

## 4 Algorithmic framework

In this section we describe an algorithmic framework that will be used for solving standard and strong regimes of BPP (Definition 1). The framework can be viewed as a variant of the PRM algorithm [28], and we chose to describe it here in full detail for completeness. However, the analysis provided in Section 5 is brand new.

The framework consists of three conceptually simple steps: In the first step, we construct a random graph embedded in  $[0, 1]^d$ , whose vertices consist of the start and target points  $S, T$ , and of a collection of randomly sampled points; the edges connect points that are separated by a distance of at most a given connection threshold  $r_n$ . In the second step the edges of the graph are assigned with weights corresponding to their bottleneck cost over  $\mathcal{M}$ . In the third and final step, the discrete graph is searched for a path connecting  $S$  to  $T$  which minimizes the bottleneck cost.

Before proceeding to a more elaborate description of the framework we provide a formal definition of the random graphs that are at the heart of the technique. Let  $\mathcal{X}_n = \{X_1, \dots, X_n\}$  be  $n$  points chosen independently and uniformly at random from the Euclidean  $d$ -dimensional cube  $[0, 1]^d$ . The following definition corresponds to the standard and well-studied model of random geometric graphs (see, e.g., [5, 36, 51] and the literature review in [45]).

**Definition 2.** The *random geometric graph* (RGG)  $\mathcal{G}_n = \mathcal{G}(\mathcal{X}_n; r_n)$  is a *directed* graph with vertex set  $\mathcal{X}_n$  and edge set  $\{(x, y) : x \neq y, x, y \in \mathcal{X}_n, \|x - y\|_2 \leq r_n\}$ .

We are ready to describe the framework, which has two parameters:  $n$  represents the number of samples generated and  $r_n$  defines the Euclidean connection radius used in the construction of the graphs. In the next section we show that for a range of values of  $r_n$ , which is a function of the number of samples  $n$ , the cost of the returned solution converges to the optimum, as  $n$  tends to infinity. The framework consists of the following steps:

**Step I:** We construct the RGG  $\mathcal{G}_n = (\mathcal{X}_n \cup \{S, T\}; r_n)$ . For the purpose of generating  $\mathcal{G}_n$  a collection of  $n$  samples  $\mathcal{X}_n$  is generated and a nearest-neighbor structure is employed to find for every  $x \in \mathcal{X}_n \cup \{S, T\}$  the set of samples that located within a Euclidean distance of  $r_n$  from it.

**Step II:** We assign to each edge of the graph the bottleneck cost of the straight-line path connecting its endpoints under  $\mathcal{M}$ . In particular, for the standard BPP, for every edge  $(x, y)$  the cost  $\max_{\tau \in [0, 1]} \mathcal{M}(x + \tau(y - x))$  is assigned. The same applies for strong-BPP, unless  $x \not\leq y$ , in which case the value  $+\infty$  is assigned.

**Step III:** For the final step we find a path over  $\mathcal{G}_n$  from  $S$  to  $T$  which minimizes the bottleneck cost. Several efficient algorithms solving this problem exist (see, e.g., [15, 52]).

## 5 Theoretical foundations

We study the behavior of the framework for the standard and the strong case of BPP (Definition 1). Recall the framework uses the two parameters  $n$  and  $r_n$ , which specify the number of samples and the connection radius. We establish a range of connection

radii,  $r_n$ , for which the cost of the returned solution is guaranteed to converge to a relaxed notion of the optimum.

The analysis below does not restrict itself to a specific type of cost maps  $\mathcal{M}$ , e.g., continuous or smooth. Thus, due to the stochastic nature of the framework, and the general definition of  $\mathcal{M}$ , we cannot guarantee that the returned solution will tend to the absolute optimum. As an example consider the cost map  $\mathcal{M}$  such that for a given  $x = (x_1, x_2)$ ,  $\mathcal{M}(x) = 0$  if  $x_1 = x_2$ , and  $\mathcal{M}(x) = 1$  otherwise. For the start and target points  $S = (0.1, 0.1), T = (0.9, 0.9)$  the optimal solution is a subset of the diagonal. Obviously, the probability of having a single point of  $\mathcal{X}_n$ , let alone a whole path in  $\mathcal{G}_n$ , that lie on the diagonal is equal to 0.

We can however guarantee convergence to a *robustly-optimal* path, which is defined below. Informally, such paths have “well-behaved” neighborhoods, in terms of the value of  $\mathcal{M}$ . We provide below a formal definition of this notion for the bottleneck cost function. Recall that given a path  $\sigma$  the notation  $\mathcal{M}(\sigma)$  represents its bottleneck cost.

**Definition 3.** Given the triplet  $\langle \mathcal{M}, S, T \rangle$ , a path  $\sigma \in \Sigma(S, T)$  is called *robust* if for every  $\varepsilon > 0$  there exists  $\delta > 0$  such that  $\mathcal{M}(\sigma') \leq (1 + \varepsilon)\mathcal{M}(\sigma)$ , for any  $\sigma' \in \Sigma(S, T)$  such that  $\text{Im}(\sigma') \subset \mathcal{B}_\delta(\sigma)$ . A path that attains the infimum cost, over all robust paths, is termed *robustly optimal*.

## 5.1 (Standard) Bottleneck cost

For a given triplet  $\langle \mathcal{M}, S, T \rangle$  representing an instance of BPP, denote by  $\sigma^*$  a robustly-optimal solution. Note that we do not require here that  $\sigma^*$  or the returned solution will be monotone. We obtain the following result. All logarithms stated henceforth are to base  $e$ .

**Theorem 1.** Let  $\mathcal{G}_n = \mathcal{G}(\mathcal{X}_n \cup \{S, T\}; r_n)$  be an RGG with

$$r_n = \gamma \left( \frac{\log n}{n} \right)^{1/d}, \quad \gamma > 2(2d\theta_d)^{-1/d},$$

where  $\theta_d$  denotes the Lebesgue measure of a unit ball in  $\mathbb{R}^d$ . Then  $\mathcal{G}_n$  contains a path  $\sigma_n \in \Sigma(S, T)$  such that  $\mathcal{M}(\sigma_n) = (1 + o(1))\mathcal{M}(\sigma^*)$ , a.s.

We mention that this connection radius is also essential for connectivity of RGGs, i.e., a smaller radius results in a graph that is disconnected with high probability (see, e.g., [8]). This fact is instrumental to our proof. We also mention that a result similar to Theorem 1 can be obtained through a different proof technique [28], albeit with a larger value of the constant  $\gamma$ .

For simplicity, we assume for the purpose of the proof that exists a finite constant  $\delta' > 0$  such that  $\mathcal{B}_{\delta'}(\sigma^*) \subset [0, 1]^d$ , namely the robustly-optimal solution is at least  $\delta'$  away from the boundary of the domain  $[0, 1]^d$ . This constraint can be easily relaxed by transforming  $\langle \mathcal{M}, S, T \rangle$  into an equivalent instance  $\langle \mathcal{M}', S', T' \rangle$  where this condition is met. In particular the original input can be embedded to a cube of side length  $1 - \varepsilon$  for some constant  $\varepsilon > 0$ , which is centered in the middle of  $[0, 1]^d$ . The cost along the boundaries of the smaller cube should be extended to the remaining parts of the  $[0, 1]^d$  cube.

Given an RGG  $\mathcal{G}_n$  and a subset  $\Gamma \subset [0, 1]^d$  denote by  $\mathcal{G}_n(\Gamma)$  the graph obtained from the intersection of  $\mathcal{G}_n$  and  $\Gamma$ : it consists of the vertices of  $\mathcal{G}_n$  that are contained in

$\Gamma$  and the subset of edges of  $\mathcal{G}_n$  that are fully contained in  $\Gamma$ . Each edge is considered as a straight-line segment connecting its two end points.

A main ingredient in the proof of Theorem 1 is the following Lemma. We employ the *localization-tessellation* framework [45], which was developed by the authors. The framework allows to extend properties of RGGs to domains with complex geometry and topology.

**Lemma 1.** *Let  $\mathcal{G}_n$  be the RGG defined in Theorem 1. Additionally let  $\Gamma \subset [0, 1]^d$  be a fixed subset, where  $S, T \in \Gamma$ , and let  $\rho > 0$  be some fixed constant, such that  $\mathcal{B}_\rho(\Gamma) \subset [0, 1]^d$ . Then  $S, T$  are connected in  $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$  a.s.*

*Proof.* We rely on the well-known result that  $\mathcal{G}_n$  is connected a.s. in the domain  $[0, 1]^d$  for the given connection radius  $r_n$  (see, e.g., [45, Theorem 1]). We then use Lemma 1 and Theorem 6 in [45] which state that if  $\mathcal{G}_n$  is connected a.s., and is *localizable* (see, Definition 6 therein), then  $S, T$  are connected a.s. over  $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$ .  $\square$

*Proof of Theorem 1.* We first show that for any  $\varepsilon > 0$  it follows that  $\mathcal{M}(\sigma_n) \leq (1 + \varepsilon)\mathcal{M}(\sigma^*)$  a.s. Fix some  $\varepsilon > 0$ . Due to the fact that  $\sigma^*$  is robustly optimal, there exists  $\delta_\varepsilon > 0$  independent of  $n$  such that for every  $\sigma \in \Sigma(S, T)$  such that  $\text{Im}(\sigma) \subset \mathcal{B}_{\delta_\varepsilon}(\sigma^*)$  we have that  $\mathcal{M} \leq (1 + \varepsilon)\mathcal{M}(\sigma^*)$  a.s. Additionally, recall that there exists some  $\delta' > 0$  such  $\mathcal{B}_{\delta'}(\sigma^*) \subset [0, 1]^d$ .

Set  $\delta = \min\{\delta_\varepsilon, \delta'\}$  and define the sets  $\Gamma_{\delta/2} = \mathcal{B}_{\delta/2}(\sigma^*)$ ,  $\Gamma_\delta = \mathcal{B}_\delta(\sigma^*)$  and notice that  $S, T \in \Gamma_{\delta/2}$ . By Lemma 1 we have that  $S, T$  are connected in  $\mathcal{G}_n(\Gamma_\delta)$ . Moreover, a path connecting  $S, T$  in  $\mathcal{G}_n(\Gamma_\delta)$  must have a bottleneck cost of at most  $(1 + \varepsilon)\mathcal{M}(\sigma^*)$ .

We have shown that for any fixed  $\varepsilon > 0$ ,  $\mathcal{M}(\sigma_n) \leq (1 + \varepsilon)\mathcal{M}(\sigma^*)$  a.s. By defining the sequence  $\varepsilon_i = 1/i$  one can extend the previous result and show that  $\mathcal{M}(\sigma_n) \leq (1 + o(1))\mathcal{M}(\sigma^*)$ . This part is technical and its details are omitted (see a similar proof in [44, Theorem 6]). This concludes the proof.  $\square$

## 5.2 Strong bottleneck cost

We now focus on the strong case of the problem, where the solution is restricted to paths that are monotone in each of the  $d$  coordinates. Denote by  $\vec{\sigma}^*$  the robustly-optimal monotone solution for a given instance  $\langle \mathcal{M}, S, T \rangle$ .

**Theorem 2.** *Let  $\mathcal{G}_n = \mathcal{G}(\mathcal{X}_n \cup \{S, T\}; r_n)$  be an RGG with  $r_n = \omega(1) \left(\frac{\log n}{n}\right)^{1/d}$ . Then  $\mathcal{G}_n$  contains a monotone path  $\vec{\sigma}_n \in \Sigma(S, T)$  such that  $\mathcal{M}(\vec{\sigma}_n) = (1 + o(1))\mathcal{M}(\vec{\sigma}^*)$ , a.s.*

Let  $x, x' \in [0, 1]^d$  be two points such that  $x \preceq x'$ . For a given  $\delta > 0$  the notation  $x \preceq_\delta x'$  indicates that  $\delta = \min\{x'_i - x_i\}_{i=1}^d$ , where  $x = (x_1, \dots, x_d)$ ,  $x' = (x'_1, \dots, x'_d)$ . Given two points  $x, x' \in [0, 1]^d$ , such that  $x \preceq x'$ , denote by  $\mathcal{H}(x, x')$  the  $d$ -dimensional box  $[x_1, x'_1] \times \dots \times [x_d, x'_d]$ . In addition to the assumption that the robustly-optimal solution  $\vec{\sigma}^*$  is separated from the boundary of  $[0, 1]^d$  that we have taken in the previous analysis, we also assume that there exists a constant  $0 < \delta'' \leq 1$  such that  $S \preceq_{\delta''} T$ .

In preparation for the main proof we prove the following lemma.

**Lemma 2.** *Choose any<sup>1</sup>  $f_n \in \omega(1)$  and set  $r_n = \omega(1) \left(\frac{\log n}{n}\right)^{1/d}$ . Let  $q, q' \in [0, 1]^d$  be two points such that  $q \preceq_\delta q'$ , where  $\delta$  is independent of  $n$ . Then a.s. there exist*

<sup>1</sup>For instance,  $f_n$  can be either one of the following functions:  $\log n$ ,  $\log^* n$ , or the inverse Ackerman function  $\alpha(n)$ .



$X, X' \in \mathcal{X}_n$  with the following properties: (i)  $\|X - q\|_2 \leq r_n/2, \|X' - q'\| \leq r_n/2$ ; (ii)  $q \preceq X, X' \preceq q'$ ; (iii)  $X, X'$  are connected in  $\mathcal{G}_n$  with a monotone path.

*Proof.* We apply a tessellation argument similar to the one used to show that the standard (and undirected) RGG is connected (see, e.g., [51, Section 2.4]). Set  $\ell = \left\lceil \frac{2\|q'-q\|_2}{r_n} \right\rceil$  and observe that  $\ell \leq 2\sqrt{d}/r_n$ . Define the normalized vector  $\vec{v} = \frac{q'-q}{\|q'-q\|_2}$  and let  $H_1, \dots, H_\ell$  be a sequence of  $\ell$  hyperboxes, where

$$H_j = \mathcal{H}\left(q + (j-1) \cdot \frac{r_n}{2} \cdot \vec{v}, q + j \cdot \frac{r_n}{2} \cdot \vec{v}\right),$$

for every  $1 \leq j \leq \ell$  (see Figure 1). Observe that for every  $1 \leq j < \ell$  and every  $X_j \in H_j, X_{j+1} \in H_{j+1}$ , we have

$$X_j \preceq X_{j+1}, \|X_{j+1} - X_j\|_2 \leq r_n. \quad (1)$$

We show that for every  $1 \leq j \leq \ell$  it follows that  $\mathcal{X}_n \cap H_j \neq \emptyset$ , a.s. We start by bounding the volume of  $H_j$ . Denote by  $c_1, \dots, c_d$  the side lengths of  $H_j$ , and denote by  $\delta_1, \dots, \delta_d$  the side lengths of  $\mathcal{H}(q, q')$ . Note that  $\delta_i$  is independent of  $n$  and  $c_i = \delta_i/\ell$ . Consequently, we can represent  $c_i = \alpha_i r_n$ , where  $\alpha_i > 0$  is constant, for every  $1 \leq i \leq d$ . Thus,  $|H_j| = cr_n^d$  for some constant  $c > 0$ . Now,

$$\Pr[\mathcal{X}_n \cap H_j = \emptyset] = (1 - |H_j|)^n \leq \exp\{-n|H_j|\} = \exp\{-\omega(1) \cdot c \log n\} \leq n^{-1}.$$

In the last transition we used the fact that the function  $f_n \in \omega(1)$  can “absorb” any constant  $c$ . We are ready to show that every  $H_i$  contains a point from  $\mathcal{X}_n$  a.s.:

$$\begin{aligned} \Pr[\exists H_j : \mathcal{X}_n \cap H_j = \emptyset] &\leq \sum_{j=1}^{\ell} \Pr[\mathcal{X}_n \cap H_j = \emptyset] \\ &\leq \ell \cdot n^{-1} \leq \frac{2\sqrt{d}}{r_n} \cdot n^{-1} = \frac{2\sqrt{d}}{\omega(1) \cdot n^{1-1/d} \log^{1/d} n}. \end{aligned}$$

Thus, a.s. there exists for every  $1 \leq j \leq \ell$  a point  $x_j \in H_j$ . Observe that  $X := X_1, X' := X_\ell$  satisfy (i),(ii). Condition (iii) follows from Equation 1.  $\square$

*Proof of Theorem 2.* Similarly to the proof of Theorem 1, we fix  $\varepsilon > 0$  and select  $\delta \leq \min\{\delta', \delta''\}$  such that  $\mathcal{M}(\vec{\sigma}) \leq (1 + \varepsilon)\mathcal{M}(\vec{\sigma}^*)$  for every  $\vec{\sigma} \in \vec{\Sigma}(S, T)$  with  $\text{Im}(\vec{\sigma}) \subset \mathcal{B}_\delta(\vec{\sigma}^*) \subset [0, 1]^d$ .

The crux of this proof is that there exists a sequence of  $k$  points  $q_1, \dots, q_k \in \text{Im}(\vec{\sigma}^*)$ , where  $S = q_1, T = q_k$ , such that  $q_j \prec_{\delta/2} q_{j+1}$  for every  $1 \leq j < k$  (see Figure 2). Moreover, due to fact that  $\vec{\sigma}^*$  is monotone we can determine that such  $k$  is finite and independent of  $n$ . Thus, by Lemma 2, for every  $1 \leq j < k$  there exist  $X_j, X'_j \in \mathcal{X}_n$  which satisfy the following conditions a.s.: (i)  $\|X_j - q_j\|_2 \leq r_n/2, \|X'_j - q_{j+1}\|_2 \leq r_n/2$ ; (ii)  $q_j \preceq X_j, X'_j \preceq q_{j+1}$ ; (iii)  $X_j, X'_j$  are connected in  $\mathcal{G}_n$ . By conditions (i),(ii), for every  $1 \leq j < k$  the graph  $\mathcal{G}_n$  contains the edge  $(X'_j, X_{j+1})$ . Combined with condition (iii) this implies that  $S$  is connected to  $T$  in  $\mathcal{G}_n$  a.s.

It remains to show that the path constructed above has a cost of at most  $(1 + \varepsilon)\mathcal{M}(\vec{\sigma}^*)$ . For every  $1 \leq j \leq k$  denote by  $\vec{\sigma}_j$  the path induced by Lemma 2 from  $X_j$  to  $X'_j$ , i.e.,  $\vec{\sigma}_j(0) = X_j, \vec{\sigma}_j(1) = X'_j$  and  $\text{Im}(\vec{\sigma}_j) \subset H_j$ . Additionally, for




Figure 1: Visualization of the proof of Lemma 2 for  $d = 2$ . The blue rectangle represents  $\mathcal{H}(q, q')$  and the three red rectangles represent  $H_1, \dots, H_\ell$  for  $\ell = 3$  (the small value of  $\ell$  was selected for the clarity of visualization and in reality  $r_n \ll \delta_1$ ). The length of the largest diagonal in each of the small rectangles is  $r_n/2$ , which implies that a distance between  $X_j \in H_j, X_{j+1} \in H_{j+1}$  is at most  $r_n$ . The blue dashed arrows represent the directed graph edges  $(X_1, X_2), (X_2, X_3)$  which correspond to a monotone path connecting  $X_1$  to  $X_3$ .

every  $1 \leq j < k$  denote by  $\vec{\sigma}'_j$  the straight-line segment (sub-path) from  $X'_j$  to  $X_{j+1}$ . Now, define  $\vec{\sigma}$  to be a concatenation of  $\vec{\sigma}'_1, \vec{\sigma}'_2, \dots, \vec{\sigma}'_{k-1}, \vec{\sigma}'_k$ . We showed in the previous paragraph that such a path exists in  $\mathcal{G}_n$  a.s. Observe that for every  $1 \leq j \leq k$  it holds that  $\vec{\sigma}_j \subset \mathcal{H}(q_j, q_{j+1})$ , where  $\mathcal{H}(q_j, q_{j+1}) \subset \mathcal{B}_{\delta/2}(\vec{\sigma}^*)$ . This implies that  $\mathcal{M}(\vec{\sigma}_i) \leq (1 + \varepsilon)\mathcal{M}(\vec{\sigma}^*)$ . Additionally, recall that for every  $1 \leq j < k$  it holds that  $\|X'_j - q_{j+1}\|_2 \leq r_n/2, \|X_{j+1} - q_{j+1}\|_2 \leq r_n/2$ , which implies that  $\text{Im}(\vec{\sigma}'_j) \subset \mathcal{B}_{r_n}(q_{j+1}) \subset \mathcal{B}_\delta(\vec{\sigma}^*)$ , and consequently  $\mathcal{M}(\vec{\sigma}'_j) \leq (1 + \varepsilon)\mathcal{M}(\vec{\sigma}^*)$ . Finally,  $\mathcal{M}(\vec{\sigma}_n) \leq \mathcal{M}(\vec{\sigma}) \leq (1 + \varepsilon)\mathcal{M}(\vec{\sigma}^*)$ . This concludes the proof.  $\square$

## 6 Experimental results

In this section we validate the theoretical results that were described in the previous section. We observe that the framework can cope with complex scenarios involving two or three degrees of freedom ( $d \in \{2, 3\}$ ), and converges quickly to the optimum.

Before proceeding to the results we provide details regarding the implementation. We implemented the framework in C++, and tested it on scenarios involving two-dimensional objects. Nearest-neighbor search, which is used for the construction of RGGs, was implemented using FLANN [35]. We note that other nearest-neighbor search data structures that are tailored for the implementation of RGGs exist (see, e.g., [31]). Geometric objects, such as points, curves, and polygons were represented with CGAL [49]. For the representation of graphs and related algorithms we used BOOST [42]. Experiments were conducted on a PC with Intel i7-2600 3.4GHz processor with 8GB of memory, running a 64-bit Windows 7 OS.

We proceed to describe the implementation involving the computation of non-trivial cost maps. For curve embedding we used PQP [20] for collision detection, i.e., determining whether a given point lies in the forbidden region  $[0, 1]^2 \setminus \mathcal{F}$ . Finally, the cost of an edge with respect to a given cost map was approximated by dense sampling along the edge, as is customary in motion planning (see, e.g., [33]).




Figure 2: Visualization of the proof of Theorem 2 for  $d = 2$  and  $k = 3$ . The red curve represents  $\vec{\sigma}^*$ , on which lie the points  $q_1, q_2, q_3$  such that  $q_1 \prec_{\delta/2} q_2 \prec_{\delta/2} q_3$ . The dashed blue curves represent  $\vec{\sigma}_1, \vec{\sigma}_2$ . The gray area represents  $\mathcal{B}_\delta(\vec{\sigma}^*)$ .

The majority of running time (over %90) in the experiments below is devoted to the computation of  $\mathcal{M}$  for given point samples or edges. Thus, we report only the overall running time in the following experiments. We mention that we also implemented a simple grid-based method for the purpose of comparison with the framework. However, it performed poorly in easy scenarios and did not terminate in hard cases. Thus, we chose to omit these results here.

Unless stated otherwise, we use in the experiments the connection radius which is described in Theorem 1, and denote it by  $r_n^*$ . This applies both to the standard and strong regimes of the problem. A discussion regarding the connection radius in the strong regime appears below in Section 6.3.

## 6.1 Various scenarios

In this set of experiments we demonstrate the flexibility of the framework and test it on the three different scenarios. We emphasize that we employ a shared code framework to solve these three problems and the ones described later. The only difference in the implementation lies in the type of cost function used. The following problems are solved using a planner for the *strong* case of BPP.

Figure 3 (left) depicts an instance of **P1** (see Section 3.1), which consists of two geometrically-identical curves (red and blue). The curves are bounded in  $[0, 1]^2$  and the red curve is translated by  $(0.05, 0.05)$  from the blue curve. The optimal solution has a cost of 0.07, in which the curves are traversed identically. Our program was able to produce a solution of cost 0.126 in 27 seconds and  $n = 100,000$  samples. Results reported throughout this section are the averaged over 10 trials.

Figure 4 (left) depicts an instance of **P2**. The goal is to find a traversal of the three curves such that the traversal point along the purple curve is visible from either the blue or red curve, while of course minimizing the lengths of the leashes between the three curves. Note that the view can be obstructed by the gray rectangular obstacles. A trivial, albeit poor, solution is to move the point along the purple curve from start to end, while the traversal point of, say, the red curve stays put in the start position. A much

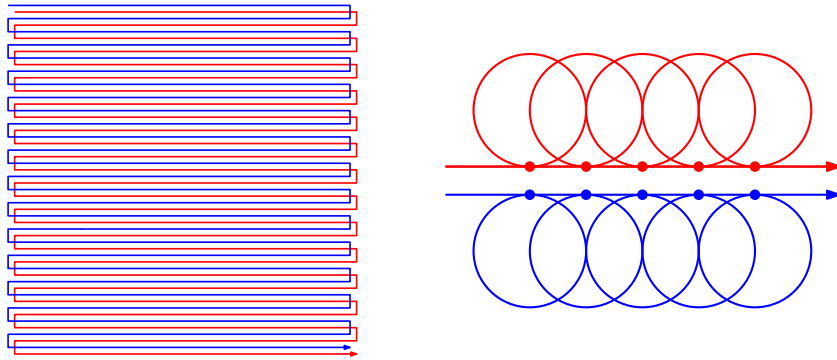


Figure 3: Scenarios involving two curve.




Figure 4: Scenarios involving curves and obstacles.

better solution, which maintains short leashes, is described as follows: we move along the purple curve until reaching the first resting point, indicated by the leftmost black disc. Then we move along the red curve until we reach to the position directly below the black circle. Only then we move along the blue curve from start until reaching the point directly below the first black disc. We use a similar parametrization with respect to the second “pit stop”, and so on. Such a solution was obtained by our program in 11 seconds using  $n = 20,000$  samples.

Figure 4 (right) depicts an instance of **P3**. The input consists of a curve (depicted in red), and polygonal obstacles (depicted in gray). The solution obtained by our program after 600 seconds with  $n = 100,000$ , is drawn in blue.

## 6.2 Increasing difficulty

Here we focus on **P1** for two curves in the standard regime. We study how the difficulty of the problem affects the running time and the rate of convergence of the returned cost. We start with a base scenario, depicted in Figure 3 (right), and gradually increase its difficulty. In the depicted scenario the bottom (blue) curve consists of five circular loops of radius 0.15, where the entrance and exit point to each circle is indicated by a bullet. The top curve is similarly defined, and the two curves are separated by a vertical distance of 0.04. The optimal matching of cost 0.34 is obtained in the following manner: when a given circle of the red curve is traversed, the position along the blue




Figure 5: Results for scenarios of increasing difficulty, as described in Section 6.2.

curve is fixed to the entrance point of the circle directly below the traversed circle, and vice versa. In a similar fashion we construct scenarios with 10,20,40 and 80 loops in each curve.

In Figure 5 we report for each of the scenarios the cost of the obtained solution as a function of the number of samples  $n$ . We set  $n = 2^i$  for the integer value  $i$  between 12 and 18. For  $i = 12$  and  $i = 18$  the running times were roughly 2 and 66 seconds, respectively. In between, the values were linearly proportional to the number of samples (results omitted). Observe that as the difficulty of the problem increases the convergence rate of the cost slightly decreases, but overall a value near the optimum is reached fairly quickly.

### 6.3 Connection radius in the strong regime

Here we consider the *strong* regime and study the behavior of the framework for varying connection radii. For this purpose, we use the two-curves scenario with 20 loops that was described in Section 6.2. We set the connection radius to  $r_n := g_n \cdot r_n^*$ , where  $r_n^*$  is the radius of the standard regime (see Theorem 1). We set  $g_n \in \{1, 1.1, \log \log n + 1, \sqrt{\log n}\}$ . Results are depicted in Figure 6.

Not surprisingly, larger values of  $r_n$  lead to quicker convergence, in terms of the number of samples required, to the optimum. However, this comes at the price of a denser RGG, which results in poor running times. Note that the program terminated due to lack of space for the two largest functions of  $g_n$  for  $n = 128,000$ . Interestingly, the connection radius  $r_n^*$  of the standard regime seems to converge to the optimum, albeit slowly. This leads to the question whether such a function also results in connectivity in the strong regime. Note that our proof of the convergence in the strong regime requires a large value of  $r_n$  (see Theorem 2).

### 6.4 Increasing dimensionality

We test how the dimension of the configuration space  $d$  affects the performance. For this purpose we study the behavior of the framework on *weak*  $k$ -curve Fréchet distance with  $k$  ranging from 2 to 5. For  $k = 2$  we use the scenario described in Section 6.2 with 10 loops. For  $k = 3$  we add another copy of the blue curve, for  $k = 4$  an additional copy of the red curve, and another blue curve for  $k = 5$ . We report running time and cost in Figure 7 for various values of  $n$ , as described earlier.

Note that that for  $k = 4$  the program ran out of memory for  $n = 64,000$ , and for  $k = 5$  around  $n = 32,000$ . This phenomena occurs since the connection radius




Figure 6: Results for varying connection radii in the strong regime, as described in Section 6.3.




Figure 7: Figures for the first set of experiments, as described in Section 6.4

obtained in Theorem 1 grows exponentially in  $d$ . In particular, for  $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$ , where  $\gamma = 2(2d\theta_d)^{-1/d}$ , each sample has in expectancy  $\Theta(2^d \log n)$  neighbors in the obtained RGG.

## References

- [1] A. Adler, M. de Berg, D. Halperin, and K. Solovey. Efficient multi-robot motion planning for unlabeled discs in simple polygons. *IEEE Trans. Automation Science and Engineering*, 12(4):1309–1317, 2015.
- [2] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geometry Appl.*, 5:75–91, 1995.
- [3] O. Arslan and P. Tsotras. Use of relaxation methods in sampling-based algorithms for optimal motion planning. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 2421–2428. IEEE, 2013.
- [4] F. Avnaim, J.-D. Boissonnat, and B. Faverjon. A practical exact motion planning algorithm for polygonal objects amidst polygonal obstacles. In *Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on*, pages 1656–1661. IEEE, 1988.
- [5] P. Balister, A. Sarkar, and B. Bollobás. Percolation, connectivity, coverage and colouring of random geometric graphs. In B. Bollobás, R. Kozma, and D. Miklós, editors, *Handbook of Large-Scale Random Networks*, pages 117–142. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

- [6] K. Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Foundations of Computer Science*, pages 661–670, 2014.
- [7] K. Bringmann and W. Mulzer. Approximability of the discrete Fréchet distance. *Journal of Computational Geometry*, 7(2):46–76, 2016.
- [8] N. Broutin, L. Devroye, N. Fraiman, and G. Lugosi. Connectivity threshold of bluetooth graphs. *Random Struct. Algorithms*, 44(1):45–66, 2014.
- [9] K. Buchin, M. Buchin, M. Konzack, W. Mulzer, and A. Schulz. Fine-grained analysis of problems on curves. In *EuroCG, Lugano, Switzerland*, 2016.
- [10] K. Buchin, M. Buchin, W. Meulemans, and W. Mulzer. Four soviets walk the dog - with an application to Alt’s conjecture. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 1399–1413, 2014.
- [11] K. Buchin, M. Buchin, and A. Schulz. Fréchet distance of surfaces: Some simple hard cases. In *European Symposium of Algorithms*, pages 63–74, 2010.
- [12] K. Buchin, M. Buchin, R. van Leusden, W. Meulemans, and W. Mulzer. Computing the Fréchet distance with a retractable leash. In *European Symposium of Algorithms*, pages 241–252, 2013.
- [13] K. Buchin, M. Buchin, and C. Wenk. Computing the Fréchet distance between simple polygons. *Comput. Geom.*, 41(1-2):2–20, 2008.
- [14] E. W. Chambers, É. C. de Verdière, J. Erickson, S. Lazard, F. Lazarus, and S. Thite. Homotopic Fréchet distance between curves or, walking your dog in the woods in polynomial time. *Comput. Geom.*, 43(3):295–311, 2010.
- [15] S. Chechik, H. Kaplan, M. Thorup, O. Zamir, and U. Zwick. Bottleneck paths and trees and deterministic graphical games. In *Symposium on Theoretical Aspects of Computer Science*, pages 27:1–27:13, 2016.
- [16] A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1):9, 2010.
- [17] M. de Berg and M. J. van Kreveld. Trekking in the alps without freezing or getting tired. *Algorithmica*, 18(3):306–323, 1997.
- [18] A. Dumitrescu and G. Rote. On the Fréchet distance of a set of curves. In *Canadian Conference on Computational Geometry*, pages 162–165, 2004.
- [19] C. Fan, O. Filtser, M. J. Katz, T. Wylie, and B. Zhu. On the chain pair simplification problem. In *Symposium on Algorithms and Data Structures*, pages 351–362, 2015.
- [20] GAMMA group. PQP - a proximity query package, 1999. University of North Carolina at Chapel Hill, USA.
- [21] D. Halperin and M. Sharir. A near-quadratic algorithm for planning the motion of a polygon in a polygonal environment. *Discrete & Computational Geometry*, 16(2):121–134, 1996.

- [22] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. *ACM Transactions on Algorithms*, 10(1):3, 2014.
- [23] J. E. Hopcroft, J. T. Schwartz, and M. Sharir. On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the “Warehouseman’s problem”. *International Journal of Robotics Research*, 3(4):76–88, 1984.
- [24] D. Hsu, J. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *Int. J. Comput. Geometry Appl.*, 9(4/5):495–512, 1999.
- [25] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [26] L. Janson, E. Schmerling, A. A. Clark, and M. Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *I. J. Robotic Res.*, 34(7):883–921, 2015.
- [27] M. Jiang, Y. Xu, and B. Zhu. Protein structure–structure alignment with discrete Fréchet distance. *Journal of bioinformatics and computational biology*, 6(01):51–64, 2008.
- [28] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011.
- [29] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [30] L. Kettner, K. Mehlhorn, S. Pion, S. Schirra, and C. Yap. Classroom examples of robustness problems in geometric computations. *Comput. Geom.*, 40(1):61–78, 2008.
- [31] M. Kleinbort, O. Salzman, and D. Halperin. Efficient high-quality motion planning by fast all-pairs r-nearest-neighbors. In *IEEE International Conference on Robotics and Automation*, pages 2985–2990, 2015.
- [32] J. J. Kuffner and S. M. LaValle. RRT-Connect: An efficient approach to single-query path planning. In *International Conference on Robotics and Automation (ICRA)*, pages 995–1001, 2000.
- [33] S. M. LaValle. *Planning algorithms*. Cambridge University Press, 2006.
- [34] W. Meulemans. Map matching with simplicity constraints. *CoRR*, abs/1306.2827, 2013.
- [35] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *VISSAPP*, pages 331–340. INSTICC Press, 2009.
- [36] M. Penrose. *Random geometric graphs*, volume 5. Oxford University Press, 2003.
- [37] B. Raveh, A. Enosh, O. Schueler-Furman, and D. Halperin. Rapid sampling of molecular motions with prior information constraints. *PLoS Computational Biology*, 5(2), 2009.



- [38] J. H. Reif. Complexity of the movers problem and generalizations: Extended abstract. In *Foundations of Computer Science*, pages 421–427, 1979.
- [39] O. Salzman and D. Halperin. Asymptotically-optimal motion planning using lower bounds on cost. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4167–4172, 2015.
- [40] M. Sharir. Algorithmic motion planning. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry, Second Edition.*, pages 1037–1064. Chapman and Hall/CRC, 2004.
- [41] J. Sherette and C. Wenk. Simple curve embedding. *CoRR*, abs/1303.0821, 2013.
- [42] J. G. Siek, L.-Q. Lee, and A. Lumsdaine. *The Boost Graph Library User Guide and Reference Manual*. Addison-Wesley, 2002.
- [43] K. Solovey and D. Halperin. On the hardness of unlabeled multi-robot motion planning. In *Robotics: Science and Systems (RSS)*, 2015.
- [44] K. Solovey, O. Salzman, and D. Halperin. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *International Journal of Robotics Research*, 35(5):501–513, 2016.
- [45] K. Solovey, O. Salzman, and D. Halperin. New perspective on sampling-based motion planning via random geometric graphs. In *Robotics: Science and Systems (RSS)*, Ann Arbor, Michigan, 2016.
- [46] K. Solovey, J. Yu, O. Zamir, and D. Halperin. Motion planning for unlabeled discs with optimality guarantees. In *Robotics: Science and Systems (RSS)*, 2015.
- [47] P. G. Spirakis and C.-K. Yap. Strong NP-hardness of moving many discs. *Information Processing Letters*, 19(1):55–59, 1984.
- [48] R. Sriraghavendra, K. Karthik, and C. Bhattacharyya. Fréchet distance based approach for searching online handwritten documents. In *Document Analysis and Recognition*, volume 1, pages 461–465. IEEE, 2007.
- [49] The CGAL Project. *CGAL user and reference manual*. CGAL editorial board, 4.8 edition, 2016.
- [50] M. Turpin, N. Michael, and V. Kumar. Concurrent assignment and planning of trajectories for large teams of interchangeable robots. In *International Conference on Robotics and Automation (ICRA)*, pages 842–848, 2013.
- [51] M. Walters. Random geometric graphs. In R. Chapman, editor, *Surveys in Combinatorics 2011*, chapter 8, pages 365–401. Cambridge University Press, 2011.
- [52] V. V. Williams. *Efficient Algorithms for Path Problems in Weighted Graphs*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, 2008.
- [53] J. Zheng, X. Gao, E. Zhan, and Z. Huang. Algorithm of on-line handwriting signature verification based on discrete Fréchet distance. In *Advances in Computation and Intelligence*, pages 461–469. Springer, 2008.