Approximating Nearest Neighbor Distances[*]

Michael B. Cohen[†]   Brittany Terese Fasy[‡]   Gary L. Miller[§]

Amir Nayyeri[¶]   Donald R. Sheehy[‖]   Ameya Velingker[**]

March 2, 2015

Abstract

Several researchers proposed using non-Euclidean metrics on point sets in Euclidean space for clustering noisy data. Almost always, a distance function is desired that recognizes the closeness of the points in the same cluster, even if the Euclidean cluster diameter is large. Therefore, it is preferred to assign smaller costs to the paths that stay close to the input points.

In this paper, we consider the most natural metric with this property, which we call the nearest neighbor metric. Given a point set P and a path γ, our metric charges each point of γ with its distance to P. The total charge along γ determines its nearest neighbor length, which is formally defined as the integral of the distance to the input points along the curve. We describe a $(3+\varepsilon)$-approximation algorithm and a $(1+\varepsilon)$-approximation algorithm to compute the nearest neighbor metric. Both approximation algorithms work in near-linear time. The former uses shortest paths on a sparse graph using only the input points. The latter uses a sparse sample of the ambient space, to find good approximate geodesic paths.

[†]Massachusetts Institute of Technology, `micohen@mit.edu`

[‡]Department of Computer Science,Tulane University, `brittany.fasy@alumni.duke.edu`

[§]Department of Computer Science, Carnegie Mellon University, `glmiller@cs.cmu.edu`

[¶]School of electrical engineering and computer science, Oregon State University, `nayyeria@eecs.oregonstate.edu`

[‖]Computer Science and Engineering Department, University of Connecticut, `don.r.sheehy@gmail.com`

[**]Department of Computer Science, Carnegie Mellon University, `avelingk@cs.cmu.edu`

1 Introduction

Many problems lie at the interface of computational geometry, machine learning, and data analysis, including–but not limited to: clustering, manifold learning, geometric inference, and nonlinear dimensionality reduction. Although the input to these problems is often a Euclidean point cloud, a different distance measure may be more *intrinsic* to the data. In particular, we are interested in a distance that recognizes the closeness of two points in the same cluster, even if their Euclidean distance is large, and conversely, recognizes a large distance between points in different clusters, even if the Euclidean distance is small. For example, in Figure 1, the distance between a and b must be larger than the distance between b and c.
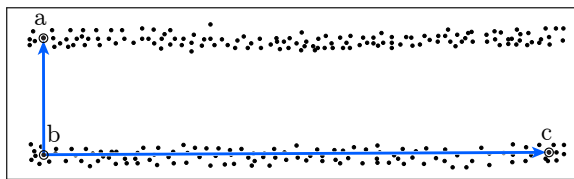

Figure 1: The intrinsic density-based distance can differ from the Euclidean distance.

There are at least two seemingly different approaches to define a non-Euclidean metric on a finite set of point on $\mathbb{R}^d$. The first approach is to form a graph metric on the point set. An example of such a graph is the kth nearest neighbor graph where a point is only connected to another point if one is a kth nearest neighbor of the other. The edge weights may be a constant or the Euclidean distances. In this paper we consider the complete graph but the edge lengths are a power of their Euclidean lengths. We are particularly interested in the squared length, which we will refer to as the edge-squared metric.

The second approach is to endow all of $\mathbb{R}^d$ with a new metric. We start with a cost function $c : \mathbb{R}^d \to \mathbb{R}$, which takes the point cloud into account. Then, the length of a path γ is the integral of the cost function along the path.

$$\ell_c(\gamma) = \int_\gamma c(s)ds = \int_0^1 c(\gamma(t)) \left| \frac{d\gamma}{dt}(t) \right| dt. \tag{1}$$

The distance between two points $x, y \in \mathbb{R}^d$ is then the length of the shortest path between them:

$$\mathbf{d}_c(x, y) = \inf_\gamma \ell_c(\gamma). \tag{2}$$

2

Note that the constant function, $c(x) = 1$ for all $x \in \mathbb{R}^d$, gives the Euclidean metric; whereas, other functions allow space to be stretched in various ways.

In almost all applications mentioned above for cost-based metrics, in order to reinforce paths within clusters, one would like to assign smaller lengths to paths that stay close to the point cloud. Therefore, the simplest natural cost function on $\mathbb{R}^d$ is the distance to the point cloud. More precisely, given a finite point set P the cost $c(x)$ for $x \in \mathbb{R}^d$ is chosen to be $N(x)$, the Euclidean distance from x to the nearest point in P. The *nearest neighbor length* (N-length) $\ell_N(\gamma)$ of a curve is given by (1), where we set $c(x) = N(x)$ for all points $x \in \mathcal{C}$. We refer to the corresponding metric given by (2) as the *nearest neighbor metric* or simply the N-distance.

In this paper, we investigate approximation algorithms for N-distance computation. We describe a $(3 + \varepsilon)$-approximation algorithm and a $(1 + \varepsilon)$-approximation algorithm. The former comes from comparing the nearest neighbor metric with the edge-squared metric. The latter is a tighter approximation that samples the ambient space to find good approximate geodesics.

1.1   Overview

In Section 4, we describe a constant factor approximation algorithm obtained via an elegant reduction into the edge-squared metric introduced by [BRS11] and [VB03]. This metric is defined between pairs of points in P by considering the graph distance on a complete weighted graph, where the weight of each edge is the square of its Euclidean length. We show that the N-distance and edge-squared metric are equivalent up to a factor of three (after a scaling by a factor of four). As a result, because spanners for the edge-squared metric can be computed in nearly linear time [LSV06], we obtain a $(3 + \varepsilon)$-approximation algorithm for computing N-distance.

Theorem 1.1. *Let P be a set of points in $\mathbb{R}^d$, and let $x, y \in P$. The nearest neighbor distance between x and y can be approximated within a $(3 + \varepsilon)$ factor in $O(n \log n + n\varepsilon^{-d})$ time, for any $0 < \varepsilon \leq 1$.*

In Section 5, we describe a $(1 + \varepsilon)$-approximation algorithm for the N-distance that works in time $\varepsilon^{-O(d)} n \log n$. Our algorithm computes a discretization of the space for points that are sufficiently far from P. Nevertheless, the sub-paths that are close to P are computed exactly. We can adapt our algorithm to work for any Lipschitz cost function that is bounded away from zero; thus, the algorithm can be applied to many of the scenarios describe in Appendix 2.

3

Theorem 1.2. *For any finite set of points $P \subset \mathbb{R}^d$ and any fixed number $0 < \varepsilon < 1$, the shortest N-distance between any pair of points of the space can be $(1 + \varepsilon)$-approximated in time $O(\varepsilon^{-O(d)} n \log n)$.*

2 Related Work

Computing the distance between a pair of points with respect to a cost function encompasses several significant problems that have been considered by different research communities for at least a few centuries. As early as 1696, Johann Bernoulli introduced the *brachistochrone* curve, the shortest path in the presence of gravity, as "an honest, challenging problem, whose possible solution will bestow fame and remain as a lasting monument" [Ber96]. With six solutions to his problem published just one year after it was posed, this event marked the birth of the field of *calculus of variations*. In this section, we review some work related to computing shortest paths in a weighted domain.

Models for Geometries. The cost function $c(\cdot)$ in (1) can be deliberately picked so that the metric $\mathbf{d}_c(\cdot)$ coincides with well-known metrics. For example, $c(x, y) = 1$ gives the Euclidean metric in the plane and $c(x, y) = 1/y$ gives the Poincaré metric in the half-plane model of hyperbolic (Lobachevsky) geometry [Kat10].

Motion Planning. Rowe and Ross [RR90] as well as Kime and Hespanha [KH03] consider the problem of computing anisotropic shortest paths on a terrain. An anisotropic path cost takes into account the (possibly weighted) length of the path and the direction of travel. Note that this problem can be translated into the problem of computing a shortest path between two compact subspaces of $\mathbb{R}^6$ under a certain cost function.

Computational Geometry. Indeed, computational geometers are interested in different versions of this problem. In the simplest case, c takes values from $\{1, \infty\}$, i.e., the space is divided into free space and obstacles. This problem can be solved in polynomial time using visibility graph in two dimensions, and it can be ε-approximated in three dimensions. For example, the computation of the Fréchet distance can be posed in this way [AG95]. A slightly more complicated case occurs when we let c be a piecewise-constant function. Mitchell and Papadimitriou [MP91] formulated this problem in two dimensions and designed a linear-time algorithm to find the solution

within ε-accuracy. They list the problem for more general cost functions as an open problem (See Section 10, problem number (3)). A series of works [ALMS98, AMS00, RS00, AMS05] has resulted in an ε-approximation algorithm that computes the shortest paths in a weighted polyhedral surface in $O((n/\sqrt{\varepsilon})\log(n/\varepsilon)\log(1/\varepsilon)$ time.

Machine Learning. Sajama and Orlitsky [SO05] first applied density-based distance (DBDs) to semi-supervised learning. Assuming that the sample points are taken from an underlying distribution with density f, a density-based distance can be defined by setting $c(x) = f(x)^{-1/d}$ in (1) and (2).[1] The goal here is to place points that can be connected through dense regions into the same cluster. Vincent and Bengio [VB03] and Bousquet et al. [BCH04] suggest estimating f using a KDE and then approximating the metric by discretizing the space in a similar fashion to Tsitsiklis [Tsi95]. However, they do not provide any analysis on the complexity of the discretized space. Bijral et al. [BRS11] bypass estimating f by building a complete graph over a set of points $\{x_1, x_2, \ldots, x_n\}$ sampled with respect to f, in which the length of an edge (x_i, x_j) is $||x_i - x_j||_p^q$ for fixed p and q, and computing pairwise shortest paths in this graph. Hwang et al. [HDI14] prove that for certain values of p and q, the latter metric and the density based metric are equivalent up to a linear factor for sufficiently large values of n.

The nearest neighbor metric can be viewed as a special case of density-based distance when the underlying density is the nearest neighbor density estimator.

3 Preliminaries

In this section, we define some basic concepts that are used in the paper.

3.1 Metrics

In this paper, we consider three metrics. Each metric is defined by a length function on a set of paths between two points of the space. The distance between two points is the length of the shortest path between them.

---

[1]The definition of DBDs in [SO05] is more general in that it allows for a choice of discount functions (we take the inverse raised to the power $1/d$ as this seems to be the most natural way of turning a measure of volumetric density to a measure of length).

Euclidean metric. This is the most natural metric defined by the Euclidean length. We use $\ell(\gamma)$ to denote the Euclidean length of a curve γ; $\ell(\gamma)$ can also be defined by setting $c(x) = 1$ for all $x \in \mathbb{R}^d$ in (1). We use $\mathbf{d}(x, y)$ to denote the distance between two points $x, y \in \mathbb{R}^d$ based on the Euclidean metric.

Nearest neighbor metric. As mentioned above, the nearest neighbor length of a curve with respect to a set of points P, is defined by setting $c(\cdot)$ to be $\mathrm{N}(\cdot)$ in (1). The nearest neighbor length of a curve γ is denoted by $\ell_{\mathrm{N}}(\gamma)$, and the distance between two points $x, y \in \mathbb{R}^d$ with respect to the nearest neighbor metric is denoted by $\mathbf{d}_{\mathrm{N}}(x, y)$.

Edge-squared metric. Finally, the edge-squared metric is defined as the shortest path metric on a complete graph on a point set P, where the length of each edge is its Euclidean length squared. The length of a path γ in this graph is naturally the total length of its edges and it is denoted by $\ell_{\mathrm{sq}}(\gamma)$. The edge-squared distance between two points $x, y \in P$ is the length of the shortest path and is denoted by $\mathbf{d}_{\mathrm{sq}}(x, y)$.

3.2 Voronoi Diagrams and Delaunay Triangulations

Let P be a finite set of points, called *sites*, in $\mathbb{R}^n$, for some $n \geq 1$. The Delaunay triangulation $\mathrm{Del}(P)$ is a decomposition of the plane into simplices such that for each simplex $\sigma \in \mathrm{Del}(P)$, the Delaunay empty circle property is satisfied; that is, there exists a circle C such that the vertices of σ are on the boundary of C and $\mathrm{int}(C) \cap V$ is empty. The Voronoi diagram, denoted $\mathrm{Vor}(P)$, is the planar dual to $\mathrm{Del}(P)$. We define the in-ball of a Voronoi cell with site p to be the maximal ball centered at p that is contained in the cell. The inradius of a Voronoi cell is the radius of its in-ball. We refer the reader to [DBVKOS00] for more details.

4 Nearest Neighbor Distance Versus Edge-Squared Distance

In this section, we show that the nearest neighbor distance of two points $x, y \in P$ can be approximated within a factor of three by looking at their edge-squared distance. More precisely, $\mathbf{d}_{\mathrm{sq}}(x, y)/4 \geq \mathbf{d}_{\mathrm{N}}(x, y) \geq \mathbf{d}_{\mathrm{sq}}(x, y)/12$ (see Lemma 4.1 and Lemma 4.4).

As a consequence, a constant factor approximation of the N-distance can be obtained via computing shortest paths on a weighted graph, in nearly-quadratic time. This approximation algorithm becomes more efficient, if the shortest paths are computed on a Euclidean spanner of the points, which is computable in nearly linear time [Hp11]. A result of Lukovszki et al. (Theorem 16(ii) of [LSV06]) confirms that a $(1+\varepsilon)$-Euclidean spanner is a $(1+\varepsilon)^2$-spanner for the edge squared metric. Therefore, we obtain Theorem 1.1.

Before, starting the technical part of this section, we remark that both the nearest neighbor and edge-squared metric can have $\Omega(\log n)$ doubling dimension. An illustrative example is a star with dense points sets on its edges.

4.1 The Upper Bound

We show that the edge-squared distance between any pair of points $x, y \in P$ (with respect to the point set P) is always larger than four times the N-distance between x and y (with respect to P). To this end, we consider any shortest path with respect to the edge-squared measure and observe that its N-length is an upper bound on the N-distance between its endpoints.

Lemma 4.1. *Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $\mathbf{d}_N$ and $\mathbf{d}_{sq}$ be the associated nearest neighbor and edge-squared distances, respectively. Then, for any distinct points $x, y \in P$, we have that $\mathbf{d}_N(x, y) \leq \frac{1}{4}\mathbf{d}_{sq}(x, y)$.*

Proof. Consider the shortest path $x = q_1 \to q_2 \to \cdots \to q_k = y$ with respect to the edge-squared metric. Let γ be the same path in $\mathbb{R}^d$ parameterized by arc length that uses straight line segments between each pair q_i, q_{i+1}. By the definition of the edge-squared distance, we have

$$\mathbf{d}_{sq}(x, y) = \sum_{i=1}^{k-1} \mathbf{d}(q_i, q_{i+1})^2.$$

On the other hand, by the definition of N-distance we have

$$\mathbf{d}_N(x, y) \leq \ell_N(\gamma) = \int_\gamma N(s)\, ds.$$

The following sequence of equalities follow by basic rules of integration.

$$
\begin{aligned}
\int_\gamma \mathrm{N}(s)\, ds \;&=\; \sum_{i=1}^{k-1} \int_{q_i \to q_{i+1}} \mathrm{N}(s)\, ds \\
&=\; \sum_{i=1}^{k-1} \left(\int_{q_i \to \frac{q_i+q_{i+1}}{2}} \mathrm{N}(s)\, ds + \int_{\frac{q_i+q_{i+1}}{2} \to q_{i+1}} \mathrm{N}(s)\, ds \right) \\
&\leq\; \sum_{i=1}^{k-1} \left(\int_{q_i \to \frac{q_i+q_{i+1}}{2}} \mathbf{d}(s, q_i)\, ds + \int_{\frac{q_i+q_{i+1}}{2} \to q_{i+1}} \mathbf{d}(s, q_{i+1})\, ds \right) \\
&=\; \sum_{i=1}^{k-1} \left(\int_0^{\mathbf{d}(q_i, q_{i+1})/2} t\, dt + \int_0^{\mathbf{d}(q_i, q_{i+1})/2} t\, dt \right) \\
&=\; \sum_{i=1}^{k-1} \frac{\mathbf{d}(q_i, q_{i+1})^2}{4} \\
&=\; \frac{1}{4} \mathbf{d}_{\mathrm{sq}}(x, y).
\end{aligned}
$$

Thus, the proof is complete. $\qquad\square$

4.2 The Lower Bound

Next, we show that the edge-squared distance between any pair of points from P cannot be larger than twelve times their N-distance. To this end, we break a shortest path of the N-distance into segments in a certain manner, and shadow the endpoints of each segment into their closest point of P to obtain a short edge-squared path. The following definition formalizes our method of discretizing paths.

Definition 4.1. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $x, y \in P$. Let $\gamma : [0, 1] \to \mathbb{R}^d$ be an (x, y)-path that is internally disjoint from P. A sequence $0 < t_0 \leq t_1 \leq \cdots \leq t_k < 1$ is a proper breaking sequence of γ if it has the following properties:*

1. *The nearest neighbors of $\gamma(t_0)$ and $\gamma(t_k)$ in P are x and y, respectively.*

2. *For all $1 \leq i \leq k$, we have $\ell(\gamma[t_{i-1}, t_i]) = \frac{1}{2}(\mathrm{N}(\gamma(t_{i-1})) + \mathrm{N}(\gamma(t_i)))$*

The following lemma guarantees the existence of breaking sequences.

Lemma 4.2. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $x, y \in P$. Let γ be a path from x to y that is internally disjoint from P. There exists a proper breaking sequence of γ.*

Proof. Pick t_0 such that the closest neighbor to $\gamma(t_0)$ in P is x. Inductively, pick $t_i > t_{i-1}$ so that property (2) of a proper breaking sequence holds until $\gamma(t_i)$'s closest neighbor is y.

We need to prove two properties: **(I)** a t_i with property (2) always exists, **(II)** this process ends after a finite number of steps (i.e., t_i falls in the Voronoi cell of y for some i).

Property I. Let t_{i-1} be the last selected point in the process. We prove that a $t_i \in (t_{i-1}, 1)$ exists such that

$$\ell(\gamma[t_{i-1}, t_i]) = \frac{1}{2}(\mathrm{N}(\gamma(t_{i-1})) + \mathrm{N}(\gamma(t_i))).$$

Let the functions $f : [t_{i-1}, 1] \to \mathbb{R}$ and $g : [t_{i-1}, 1] \to \mathbb{R}$ be defined as follows:

$$\begin{aligned}
f(t) &= \ell(\gamma[t_{i-1}, t]) \\
g(t) &= \frac{1}{2}(\mathrm{N}(\gamma(t_{i-1})) + \mathrm{N}(\gamma(t))).
\end{aligned}$$

In particular, $f(t_{i-1}) = 0$ and $g(t_{i-1}) = \mathrm{N}(\gamma(t_{i-1})) > 0$; so $g(t_{i-1}) > f(t_{i-1})$. On the other hand,

$$f(1) = \ell(\gamma(t_{i-1}, 1)) \geq \mathrm{N}(\gamma(t_{i-1})) > \mathrm{N}(\gamma(t_{i-1}))/2 = g(1).$$

The first inequality is a result of the Lipschitz property and the fact that $\mathrm{N}(\gamma(1)) = \mathrm{N}(y) = 0$.

Since both f and g are continuous functions, the intermediate value theorem implies the existence of a $t_i \in [t_{i-1}, 1]$ such that $f(t_i) = g(t_i)$, which in turn implies property (I).

Property II. For $t \in [t_0, 1]$, let $p(t)$ be the Euclidean distance from $\gamma(t)$ to $P \setminus \{y\}$. By definition, $p(t)$ is positive everywhere. Since p is continuous and defined on a closed interval, by the extreme value theorem, it attains a minimum value $p_{\min} > 0$. In any inductive step, if the nearest neighbor of neither t_{i-1} nor t_i is y then $\ell(\gamma[t_{i-1}, t_i]) \geq p_{\min}$, by the second property of a breaking sequence. This implies the existence of a finite k such that $\gamma(t_k)$'s nearest neighbor in P is y. $\square$

Lemma 4.3. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$. Furthermore, let γ be any path in $\mathbb{R}^d$ and x be an endpoint of γ. If $\ell(\gamma) = s$, then $\ell_{\mathrm{N}}(\gamma) \geq s\mathrm{N}(x) - s^2/2$.*

Proof. Let γ_u be a unit speed reparameterization of γ (i.e., $|\gamma_u(t)| = 1$ for all t). Suppose, without loss of generality, that $x = \gamma(0) = \gamma_u(0)$. Then, by definition of the nearest neighbor metric,

$$\ell_N(\gamma) = \ell_N(\gamma_u) = \int_0^s N(\gamma_u(t))dt.$$

Then, the Lipschitz property of the N function implies:

$$\begin{aligned} \ell_N(\gamma) \quad &\geq \quad \int_0^s (N(\gamma_u(0)) - t)dt \\ &= \quad \int_0^s (N(x) - t)dt \\ &= \quad sN(x) - s^2/2. \end{aligned}$$

$\square$

Given a path γ that realizes the nearest neighbor distance between two points x and y, in the proof of the following lemma we show how to obtain another (x, y)-path with bounded edge-squared length. The proof heavily relies on the idea of breaking sequences.

Lemma 4.4. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $\mathbf{d}_N$ and $\mathbf{d}_{sq}$ be the associated nearest neighbor and edge-squared distances, respectively. Then, for any distinct points $x, y \in P$, $\mathbf{d}_N(x, y) \geq \frac{1}{12}\mathbf{d}_{sq}(x, y)$.*

Proof. Let γ be a path from x to y that realizes the nearest neighbor distance between x and y, i.e., $\ell_N(\gamma) = \mathbf{d}_N(x, y)$.

Suppose, without loss of generality, that γ intersects P only at $\{x, y\}$. Otherwise, we break γ into pieces that are internally disjoint from P and prove the bound for each piece separately.

Let $\{t_0, t_1, \ldots, t_k\}$ be a proper breaking sequence of γ. For $0 \leq i \leq k$, let $n_i \in P$ be the nearest neighbor of $\gamma(t_i)$; in particular $n_0 = x$ and $n_k = y$.
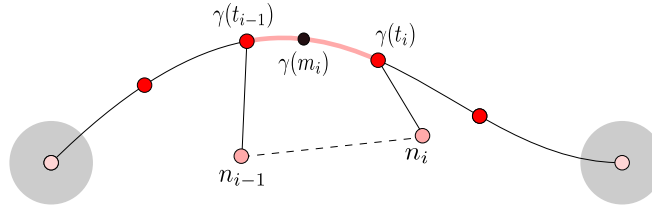

Figure 2: Proof of Lemma 4.4

10

We show that $\ell_N(\gamma[t_{i-1}, t_i]) \geq \frac{1}{12}(\mathbf{d}(n_{i-1}, n_i))^2$ for any $1 \leq i \leq k$, which in turn implies

$$\ell_N(\gamma) \geq \frac{1}{12} \sum_{i=1}^{k} (\mathbf{d}(n_{i-1}, n_i))^2 \geq \frac{1}{12} \mathbf{d}_{sq}(x, y).$$

Property 2 of a breaking sequence implies that $\ell(\gamma[t_{i-1}, t_i]) = 2s$, where $s = \frac{1}{4}(N(\gamma(t_{i-1})) + N(\gamma(t_i)))$. We pick $m_i \in (t_{i-1}, t_i)$ so that $\ell(\gamma[t_{i-1}, m_i]) = \ell(\gamma[m_i, t_i]) = s$. Lemma 4.3 implies

$$\begin{aligned}
\ell_N(\gamma[t_{i-1}, t_i]) &= \ell_N(\gamma[t_{i-1}, m_i]) + \ell_N(\gamma[m_i, t_i]) \\
&\geq (sN(\gamma(t_{i-1})) - s^2/2) + (sN(\gamma(t_i)) - s^2/2) \\
&= 3s^2.
\end{aligned} \tag{3}$$

On the other hand, by the triangle inequality,

$$\mathbf{d}(n_{i-1}, n_i) \leq N(\gamma(t_{i-1})) + 2s + N(\gamma(t_i)) = 6s. \tag{4}$$

The last equality follows by property I. Finally, Inequalities (3) and (4) imply

$$\ell_N(\gamma_i) \geq \frac{1}{12}(\mathbf{d}(n_{i-1}, n_i))^2,$$

and the proof is complete. $\qquad\square$

5 A $(1+\varepsilon)$-Approximation for the Nearest Neighbor Metric

In this section, we describe a polynomial time approximation scheme to compute the N-distance between a pair of points from a finite set $P \subset \mathbb{R}^d$. The running time of our algorithm is $\varepsilon^{-O(d)} n \log n$ for n points in d-dimensional space. We start with Section 5.1, which describes an exact algorithm for the simple case in which P consists of just one site. Section 5.3 describes how to obtain a piecewise linear path using infinitely many Steiner points. Section 5.4 combines ideas from 5.3 and 5.1 to cut down the required Steiner points to a finite number. Finally, Section 5.5 describes how to generate the necessary Steiner points.

5.1 Nearest Neighbor Distance with One Site

We describe a method for computing $\mathbf{d}_N$ for the special case that P is a single point using complex analysis. This case will be important since distances will go to zero at an input point and thus we must be more careful at input points. Far away from input points, we will use a piecewise constant approximation for the nearest neighbor function but near input points we will us exact distances. More than likely this case has been solved by others since the solution is so elegant. We refer the interested reader to [Str] for more general methods to solve similar problems in the field of calculus of variations.
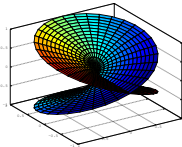

Figure 3: To make the complex function one-to-one, one needs to extend the complex plane to the two-fold cover called the two-fold Riemann Surface.

Suppose we want to compute $\mathbf{d}_N(x, y)$ where $P = \{(0,0)\}$. Writing $(x, y) \in \mathbb{C}$ in polar coordinates as $z = re^{i\theta}$, we define the *quadratic transformation* $f \colon \mathbb{C} \to \mathcal{R}$ by

$$f(z) = z^2/2 = (r^2/2)e^{i2\theta},$$

where $\mathcal{R}$ is the two-fold Riemann surface; see Figure 3. The important point here is that the image is a double covering of $\mathbb{C}$. For example, the points 1 and -1 are mapped to different copies of $1/2$. Therefore, on the Riemann surface, the distance between 1 and -1 is one and the shortest path goes through the origin. More generally, given any two nonzero points p and q on the surface, the minimum angle between them (measured with respect to the origin) will be between 0 and 2π. Moreover, if this angle is $\geq \pi$, then the shortest path between them will consist of the two straight lines $[p, 0]$ and $[0, q]$. Otherwise, the line $[p, q]$ will be a line on the surface and, thus, the geodesic from p to q.

Let $\mathbf{d}_{\mathcal{R}}$ denote the distance on the Riemann surface. We next show that for a single point, the nearest neighbor geodesic is identical to the geodesic on the Riemann surface.

Lemma 5.1. *Let* $\gamma \colon [0, 1] \to \mathbb{C}$ *be a curve. Then, the image of γ under f, denoted by $f \circ \gamma$ satisfies the following property:*

$$\mathbf{d}_{\mathcal{R}}(f \circ \gamma) = \ell_N(\gamma).$$

Proof. Suppose $\gamma \colon [0,1] \to \mathbb{C}$ is any piecewise differentiable curve, and let $\alpha := f \circ \gamma$. The N-length $\ell_{\mathrm{N}}(\gamma)$ of γ is the finite sum of the N-length of all differentiable pieces of γ. If the path γ goes through the origin, we further break the path at the origin so that α is also differentiable. Thus, it suffices to consider $(a, b) \subset [0,1]$ so that $\gamma[a,b]$ is a differentiable piece of γ. Then, we have

$$\begin{aligned}
\ell_{\mathrm{N}}(\gamma[a,b]) &= \int_a^b |\gamma(t)||\gamma'(t)|\, dt && |\cdot| \text{ is modulus.}\\
&= \int_a^b |\gamma(t)\gamma'(t)|\, dt && \text{Modulus commutes with product.}\\
&= \int_a^b |\alpha'(t)|\, dt && \text{Chain rule.}\\
&= \ell_{\mathcal{R}}(\alpha[f(a), f(b)]).
\end{aligned}$$

$\square$

Corollary 1 (Reduction to Euclidean Distances on a Riemann Surface). *Given three points x, y, and p in $\mathbb{R}^d$ such that $p = \mathrm{NN}(x) = \mathrm{NN}(y)$, the nearest neighbor geodesic G from x to y satisfies the following properties:*

1. *G is in the plane determined by x, y, p.*

2. (a) *If the angle formed by x, p, y is $\pi/2$ or more, then G consists of the two straight segments $\overline{xp}$ and $\overline{py}$.*

   (b) *Otherwise, G is the preimage of the straight line from $f(x)$ to $f(y)$, where f is the quadratic map in the plane given by x, y, p to the Riemann Surface.*

5.2 Discretizing a Path

First we approximate any (x, y)-path γ with a piecewise linear (x, y)-path α that is composed of segments that are sufficiently long. The properties of α, formalized in the following definition, are used in the rest of this section to obtain a piecewise linear approximation using Steiner points.

Definition 5.1. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $x, y \in P$. Assume γ is a path from x to y that is internally disjoint from P. Also, let $0 < \varepsilon_\alpha < 1$ be a real number. A sequence $\{0 = t_0, t_1, \ldots, t_k, t_{k+1} = 1\}$ is a ε_α-discretizing sequence of γ if it has the following properties:*

1. *The nearest neighbor of both $\gamma(t_1)$ and $\gamma(t_2)$ is x, while the nearest neighbor of both $\gamma(t_{k-1})$ and $\gamma(t_k)$ is y.*

2. *For all $1 \le i \le k$, we have $\mathbf{d}(\gamma(t_i), \gamma(t_{i+1})) = \varepsilon_\alpha \cdot \mathrm{N}(\gamma(t_i))$.*

In this case, the piecewise linear path $\alpha = \{x, \gamma(t_1), \ldots, \gamma(t_k), y\}$ is a ε_α-discretized path of γ.

The following lemma guarantees the existence of ε_α-discretizing sequences and, hence, ε_α-discretized paths. Its proof is essentially the same as the proof of Lemma 4.2.

Lemma 5.2. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $x, y \in P$. Let γ be a path from x to y that is internally disjoint from P and $0 < \varepsilon_\alpha < 1$ a real number. There exists a ε_α-discretizing sequence of γ.*

Next, we show that the N-length of a ε_α-discretized path of γ is not much longer than the N-length of γ.

Lemma 5.3. *If α is a ε_α-discretized path of γ for some $\varepsilon_\alpha \le 1/2$, then $\ell_\mathrm{N}(\alpha) \le (1 + 4\varepsilon_\alpha)\ell_\mathrm{N}(\gamma)$.*

Proof. For $1 \le i \le k-1$, let α_i be the line segment $(\gamma(t_i), \gamma(t_{i+1}))$. Furthermore, let α_0 and α_k be the line segments $(x, \gamma(t_1))$ and $(\gamma(t_k), y)$, respectively. Also, let $\gamma_i = \gamma[t_i, t_{i+1}]$, for $1 \le i \le k-1$, and $\gamma_0 = \gamma[0, t_1]$ and $\gamma_k = \gamma[t_k, 1]$.

We prove the stronger statement that $\ell_\mathrm{N}(\alpha_i) \le (1 + 4\varepsilon_\alpha)\ell_\mathrm{N}(\gamma_i)$, for any $0 \le i \le k$. For $i \in \{0, k\}$, the straight line is indeed the shortest path in the nearest neighbor metric (since $\mathrm{N}(\gamma(t_1)) = \gamma(t_0) = x$ and $\mathrm{N}(\gamma(t_k)) = \gamma(t_{k+1}) = y$), and so, $\ell_\mathrm{N}(\alpha_i) \le \ell_\mathrm{N}(\gamma_i)$. For $1 \le i \le k-1$, by Lemma 5.4,

$$\ell_\mathrm{N}(\alpha_i) \le \ell_\mathrm{N}(\gamma_i) \cdot \frac{\mathrm{N}(\gamma(t_i)) + \ell(\alpha_i)}{\mathrm{N}(\gamma(t_i)) - \ell(\alpha_i)}.$$

Then, by the second property of ε_α-discretizing sequences,

$$\ell_\mathrm{N}(\alpha_i) \le \ell_\mathrm{N}(\gamma_i) \cdot \frac{\mathrm{N}(\gamma(t_i)) + \varepsilon_\alpha \mathrm{N}(\gamma(t_i))}{\mathrm{N}(\gamma(t_i)) - \varepsilon_\alpha \mathrm{N}(\gamma(t_i))}.$$

After simplifying and using the fact that $\varepsilon_\alpha \le 1/2$, we obtain

$$\ell_\mathrm{N}(\alpha_i) \le \ell_\mathrm{N}(\gamma_i) \cdot \frac{1 + \varepsilon_\alpha}{1 - \varepsilon_\alpha} \le \ell_\mathrm{N}(\gamma_i) \cdot (1 + 4\varepsilon_\alpha).$$

$\square$

5.3 Approximating with Steiner Points

Assume $P \subset \mathbb{R}^d$, $x, y \in P$, and let γ be an arbitrary (x,y)-path. We show how to approximate γ with a piecewise linear path through a collection of Steiner points in $\mathbb{R}^d$. To obtain an accurate estimation of γ, we require the Steiner points to be sufficiently dense. The following definition formalizes this density with a parameter δ.

Definition 5.2 (δ-sample). *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $D \subseteq \mathbb{R}^d$. For a real number $0 < \delta < 1$, a δ-sample is a (possibly infinite) set of points $T \subseteq D$ such that if $z \in D \setminus P$, then $\mathbf{d}(z, T) \leq \delta \cdot \mathrm{N}(z)$.*

The following lemmas guarantees that an accurate estimation of γ can be computed using a δ-sample.

Lemma 5.4. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let $x, y \in \mathbb{R}^d$. Let γ be any path from x to y, and l be the straight (x,y)-segment. Assume further that $\mathrm{N}(x) > \ell(l)$. Then,*

$$\ell_{\mathrm{N}}(l) \leq \ell_{\mathrm{N}}(\gamma) \cdot \frac{\mathrm{N}(x) + \ell(l)}{\mathrm{N}(x) - \ell(l)}$$

Proof. Let $\gamma' = \gamma \cap B(x, \ell(l))$. Observe that γ' is a finite collection paths, and let $\ell(\gamma')$ be the total length of these paths. For any point $z \in B(x, \ell(l))$
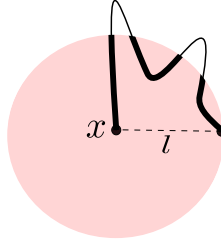

Figure 4: Proof of Lemma 5.4; γ' is bold.

we have $\mathrm{N}(x) - \ell(l) \leq \mathrm{N}(z) \leq \mathrm{N}(x) + \ell(l)$ because $\mathrm{N}(\cdot)$ is Lipschitz. Since γ' and l are both contained in $B(x, \ell(l))$, in particular, we have:

$$\ell_{\mathrm{N}}(\gamma') \geq \ell(\gamma')(\mathrm{N}(x) - \ell(l))$$

and

$$\ell_{\mathrm{N}}(l) \leq \ell(l)(\mathrm{N}(x) + \ell(l)) \leq \ell(\gamma')(\mathrm{N}(x) + \ell(l)).$$

Thus, since $\gamma' \subseteq \gamma$, we have $\ell_N(\gamma') \leq \ell_N(\gamma)$. Finally, because both sides of the inequalities are positive numbers, we have:

$$\frac{\ell_N(l)}{\ell_N(\gamma)} \leq \frac{\ell_N(l)}{\ell_N(\gamma')} \leq \frac{N(x) + \ell(l)}{N(x) - \ell(l)}.$$

$\square$

Lemma 5.5. *Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let S be a δ-sample, and let $0 < \delta < 1/10$. Then, for any pair of points $x, y \in P$, there is a piecewise linear path $\eta = (x, s_1, \ldots, s_k, y)$, where $s_1, \ldots, s_k \in S$, such that:*

$$\ell_N(\eta) \leq (1 + C_1 \delta^{2/3}) \mathbf{d}_N(x, y),$$

and, for all $1 \leq i \leq k - 1$,

$$\ell_N((s_i, s_{i+1})) \leq C_2 \cdot \delta^{2/3} \cdot N(s_i).$$

C_1 and C_2 are universal constants.

Proof. Let γ be an (x, y)-shortest path under the nearest neighbor metric. We assume, without loss of generality, that γ is internally disjoint from P. Of course, we can break any path to such pieces and use induction to prove the lemma for the general case.

Let $\alpha = (x = \gamma(t_0), \gamma(t_1), \ldots, \gamma(t_{k+1}) = y)$ be a $(2\delta^{2/3})$-discretization of γ. For all $0 \leq i \leq k$, let $\alpha_i = (\gamma(t_i), \gamma(t_{i+1}))$ for $0 \leq i \leq k$, let m_i be the midpoint of α_i, and let $s_i \in S$ be the closest Steiner point to m_i. Then, let $\beta_i = (\gamma(t_i), s_i, \gamma(t_{i+1}))$ be a two-segment path, and let $\beta = \beta_0 \circ \beta_1 \circ \cdots \circ \beta_k$. Finally, let $\eta = (x, s_1, s_2, \ldots, s_{k-1}, y)$, and let $\eta_i = (s_i, s_{i+1})$ where $0 \leq i < k$ and $s_0 = x$ and $s_k = y$ to simplify notation.

We need to prove that $\ell_N(\eta) \leq (1 + C_1 \delta^{2/3}) \ell_N(\gamma)$. Lemma 5.3 implies that

$$\ell_N(\alpha) \leq (1 + 8\delta^{2/3}) \ell_N(\gamma)$$

Showing the following two facts imply the statement of the lemma.

$$\ell_N(\beta) \leq (1 + 20\delta^{2/3}) \ell_N(\alpha) \tag{5}$$
$$\ell_N(\eta) \leq (1 + 8\delta^{2/3}) \ell_N(\beta) \tag{6}$$

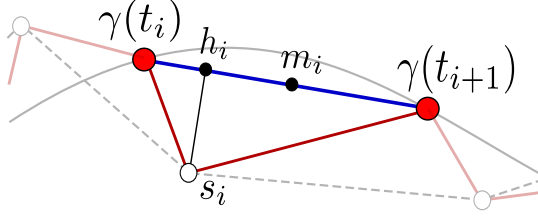**Figure 5:** The lines segment α_i is blue, and the two-segment path β_i is red.

A proof of Inequality (5): We prove the stronger statement that for all $0 \leq i \leq k$, $\ell_N(\beta_i) \leq (1 + 8\delta^{2/3})\ell_N(\alpha_i)$. Recall that m_i is the midpoint of α_i and $s_i \in S$ is the closest Steiner point to m_i. Let h_i be the closes point of α_i to s_i; see Figure 5.

Because $\ell(\alpha_i) \approx \delta^{2/3} N(m_i)$, $\ell((m_i, s_i)) \approx \delta N(m_i)$, and $\delta \leq \delta^{2/3}$, we have (s_i, h_i) is orthogonal to α_i. Then, the following bound on the length of the piecewise linear path $\beta_i = (\gamma(t_i), s_i, \gamma(t_{i+1}))$ is implied by the Pythagorean theorem using derivative to maximize a function.

$$\ell(\beta_i) \leq 2 \cdot \sqrt{\left(\frac{\ell(\alpha_i)}{2}\right)^2 + (\ell(s_i, h_i))^2} \leq 2 \cdot \sqrt{\left(\delta^{2/3}N(\gamma(t_i))\right)^2 + (\delta N(m_i))^2} \quad (7)$$

Then, the Lipschitz property of the N-function implies:

$$N(m_i) \leq N(\gamma(t_i)) + \frac{\ell(\alpha_i)}{2} \leq N(\gamma(t_i)) + \delta^{2/3} \cdot N(\gamma(t_i)) \leq (1 + \delta^{2/3})N(\gamma(t_i)).$$

By substituting $(1 + \delta^{2/3})N(\gamma(t_i))$ for $N(m_i)$ in (7), we obtain

$$\begin{aligned}
\ell(\beta_i) &\leq 2 \cdot \sqrt{\left(\delta^{2/3}N(\gamma(t_i))\right)^2 + \left(\delta(1 + \delta^{2/3})N(\gamma(t_i))\right)^2} \\
&= \delta^{2/3}N(\gamma(t_i)) \cdot \sqrt{1 + \left(\delta^{1/3}(1 + \delta^{2/3})\right)^2} \\
&= \ell(\alpha_i) \cdot \left(1 + 2\delta^{2/3}\right)
\end{aligned}$$

On the other hand, because both α_i and β_i are contained in $B(\gamma(t_i), \varepsilon_\alpha N(\gamma(t_i)))$, for any point $z \in \alpha_i \cup \beta_i$ we have $N(\gamma(t_i))(1 - 2\delta^{2/3}) \leq N(z) \leq N(\gamma(t_i))(1 + 2\delta^{2/3})$ and so

$$\ell_N(\beta_i) \leq \ell(\beta_i) \cdot N(\gamma(t_i))(1 + 2\delta^{2/3}).$$

and

$$\ell_N(\alpha_i) \geq \ell(\alpha_i) \cdot N(\gamma(t_i))(1 - 2\delta^{2/3}).$$

Last three inequalities together imply

$$
\begin{aligned}
\frac{\ell_{\mathrm{N}}(\beta_i)}{\ell_{\mathrm{N}}(\alpha_i)} &\leq \frac{1 + 2\delta^{2/3}}{1 - 2\delta^{2/3}} \cdot \left(1 + 2\delta^{2/3}\right) \\
&\leq (1 + 20\delta^{2/3}).
\end{aligned}
$$

Thus, we obtain Inequality (5) by summing over all α_i's and β_i's.

A proof of Inequality (6):   First, by the triangle inequality, we obtain:

$$
\begin{aligned}
\ell(\eta_i) &\leq \ell(s_i, t_{i+1}) + \ell(t_{i+1}, s_{i+1}) \\
&= 2\delta^{2/3} \cdot \left(\frac{\mathrm{N}(\gamma(t_i)) + \mathrm{N}(\gamma(t_{i+1}))}{2}\right) \\
&\leq 2\delta^{2/3} \cdot \max\left(\mathrm{N}(\gamma(t_i)), \mathrm{N}(\gamma(t_{i+1}))\right)
\end{aligned}
$$

Then, Lemma 5.4, letting x be the endpoint with larger N value, implies

$$
\ell_{\mathrm{N}}(\eta_i) \leq \ell_{\mathrm{N}}(s_i, t_{i+1}, s_{i+1}) \cdot \frac{1 + 2\delta^{2/3}}{1 - 2\delta^{2/3}} \leq \ell_{\mathrm{N}}(s_i, t_{i+1}, s_{i+1}) \cdot (1 + 8\delta^{2/3}).
$$

Finally, by summing over all η_i's we obtain Inequality 6, and the proof is complete. $\qquad\square$

5.4   The Approximation Graph

So far we have shown that any shortest path can be approximated using a δ-sample that is composed of infinitely many points. In addition, we know how to compute the exact N-distance between any pair of points if they reside in the same Voronoi cell of $\mathrm{Vor}(P)$. Here, we combine these two ideas to be able to approximate any shortest path using only a finite number of Steiner points. The high-level idea is to use the Steiner point approximation while γ passes through regions that are far from P and switch to the exact distance computation as soon as γ is sufficiently close to one of the points in P.

Let $P = \{p_1, p_2, \cdots, p_n\}$ be a set of points in $\mathbb{R}^d$, and let B be any convex body that contains P. Fix $\delta \in (0, 1)$, and for any $1 \leq i \leq n$, let $r_i = r_P(p_i)$ be the inradius of the Voronoi cell with site p_i. Also, let $u_i = (1 - \delta^{2/3})r_i$. Finally, let S be a δ-sample on the domain $B \setminus \bigcup_{1 \leq i \leq n} B(p_i, u_i)$.

Definition 5.3 (Approximation Graph)**.** *The approximation graph $\mathcal{A} = \mathcal{A}(P, \{u_1, \ldots, u_n\}, S, \delta) = (V_{\mathcal{A}}, E_{\mathcal{A}})$ is a weighted undirected graph, with*

weight function $w : E_{\mathcal{A}} \to \mathbb{R}^+$. *The vertices in* $V_{\mathcal{A}}$ *are in one to one correspondence with the points in* $S \cup P$; *for simplicity we use the same notation to refer to corresponding elements in* $S \cup P$ *and* $V_{\mathcal{A}}$. *The set* $E_{\mathcal{A}}$ *is composed of three types of edges:*

1. *If* $s_1, s_2 \in S$ *and* $s_1, s_2 \in B(p_i, r_i)$ *for any* p_i, *then* $(s_1, s_2) \in E_{\mathcal{A}}$ *and* $w(s_1, s_2) = \mathbf{d}_N(s_1, s_2)$. *We compute this distance using Corollary* 1.

2. *Otherwise, if* $s_1, s_2 \in S$ *and* $\ell(s_1, s_2) \leq C_2 \delta^{2/3} \max(N(s_1), N(s_2))$, *where* C_2 *is the constant of Lemma* 5.5, *then* $(s_1, s_2) \in E_{\mathcal{A}}$ *and* $w(s_1, s_2) = \max(N(s_1), N(s_2)) \cdot \ell(s_1, s_2)$.

3. *If* $s_1 \in S$ *and* $s_1 \in B(p_i, r_i)$ *then* $(p_i, s_1) \in E_{\mathcal{A}}$ *and* $w(p_i, s_1) = \mathbf{d}_N(p_i, s_1) = (\mathbf{d}(p_i, s_1))^2/2$; *see Corollary* 1.

For $x, y \in V_{\mathcal{A}}$ *let* $\mathbf{d}_{\mathcal{A}}(x, y)$ *denote the length of the shortest path from* x *to* y *in the graph* $\mathcal{A}$.

The following lemma guarantees that the shortest paths in the approximation graph are sufficiently accurate estimations.

Lemma 5.6. *Let* $\{u_1, \ldots, u_n\}$, S *and* δ *be defined as above. Let* $\mathcal{A}(P, \{u_1, \ldots, u_n\}, S, \delta)$ *be the approximation graph for* P. *For any pair of points* $x, y \in P$ *we have:*

$$(1 - C_2 \delta^{2/3}) \cdot \mathbf{d}_N(x, y) \leq \mathbf{d}_{\mathcal{A}}(x, y) \leq (1 + C_4 \delta^{2/3}) \cdot \mathbf{d}_N(x, y),$$

where C_2 *and* C_4 *are constants computable in* $O(1)$ *time.*

Proof. First, we prove the upper bound. Let $D = B \setminus \bigcup_{1 \leq i \leq n} B(p_i, u_i)$, and recall that S is a δ-sample on D. Let S^+ be a δ-sample of B such that $S^+ \cap D = S$. According to Lemma 5.5, there is a piecewise linear path $\eta = \{x, s_1, \ldots, s_k, y\}$, for $s_1, \ldots, s_k \in S^+$, such that $\ell_N(\eta) \leq (1 + C_1 \delta^{2/3}) \mathbf{d}_N(x, y)$. We use η to obtain a (x, y)-path σ in $\mathcal{A}$ with length at most $(1 + 2C_2 \delta^{2/3}) \ell_N(\eta)$; C_2 is the constant in Lemma 5.5.

Consider any maximal subpath $\eta' = (s_i, s_{i+1}, \ldots, s_{j-1}, s_j)$ of η that resides in $B\left(p_t, u_t/(1 - \delta^{2/3})\right)$ for a $p_t \in P$. The lengths of the type (2) edges and the second inequality of Lemma 5.5 imply that $s_i, s_j \notin B(p_t, u_t)$, which in turn implies that $s_i, s_j \in D$. Further, (s_i, s_j) is a type (1) edge of $\mathcal{A}$ and $w(s_i, s_j) = \ell_N(s_i, s_j) \leq \ell_N(\eta')$. So, η' in η corresponds to a single type (1) edge (s_i, s_j) in σ with no longer length. In fact, all type (1) edges of σ correspond to not shorter paths in η.

Similarly, we consider maximal paths $(p, s_i, s_{i+1}, \ldots, s_{j-1}, s_j)$, where p is an input point, and replacing them with type (3) edges of σ. Again, all type (3) edges of σ correspond to not shorter paths of η.

It remains to show that type (2) edges of σ are not too long. Observe, that each type (2) edge of σ corresponds to a line segment of η. Let $(s_1, s_2) \in E_{\mathcal{A}}$ be a type (2) edge and let $\max(N(s_1), N(s_2)) = N(s_1)$ without loss of generality. By the Lipschitz property and the length property of type (2) edges

$$\ell_N(s_1, s_2) \geq (N(s_1) - \ell(s_1, s_2))\ell(s_1, s_2) \geq (1 - C_2\delta^{2/3})N(s_1)\ell(s_1, s_2) = (1 - C_2\delta^{2/3})w(s_1, s_2),$$

which implies
$$w(s_1, s_2) \leq (1 + 2C_2\delta^{2/3})\ell_N(s_1, s_2),$$

and so, by Lemma 5.5,

$$\ell_{\mathcal{A}}(\sigma) \leq (1 + 2C_2\delta^{2/3})\ell_N(\eta) \leq (1 + C_4\delta^{2/3})\mathbf{d}_N(x, y).$$

To prove the lower bound, let $x, y \in V_{\mathcal{A}}$ and let $\sigma = (x, s_1, \ldots, s_k, y)$ be a path from x to y in $\mathcal{A}$. Also, let γ be the (x, y) path in $\mathbb{R}^d$ that hits $(x = s_0, s_1, \ldots, s_k, y = s_{k+1})$ in this order, and let the (s_i, s_j) subpath of γ be a line segment if (s_i, s_j) is a type (2) edge in $\mathcal{A}$ and be the shortest (s_i, s_j) nearest neighbor path otherwise. Then, let γ_i be the subpath of γ from s_i to s_{i+1}. We compare $w(s_i, s_{i+1})$ with $\ell_N(\gamma_i)$.

If (s_i, s_{i+1}) is type (1) or (3) then $\ell_N(\gamma_i) = w(s_1, s_2)$. Otherwise, (s_i, s_{i+1}) is type (2) and by Lipschitz property

$$\ell_N(s_1, s_2) \leq (N(s_1) + \ell(s_1, s_2))\ell(s_1, s_2) \leq (1 + C_2\delta^{2/3})N(s_1)\ell(s_1, s_2) = (1 + C_2\delta^{2/3})w(s_1, s_2).$$

Overall,
$$\mathbf{d}_N(x, y) \leq \ell_N(\gamma) \leq (1 + C_2\delta^{2/3})\ell_{\mathcal{A}}(\sigma),$$

and the proof is complete. □

5.5 Construction of Steiner Points

The only remaining piece that we need to obtain an approximation scheme is an algorithm for computing a δ-sample. For this section, given a point set T and $x \in T$, let $r_T(x)$ denote the inradius of the Voronoi cell of $\text{Vor}(T)$ that contains x. Also, given a set T and an arbitrary point x (not necessarily in T), let $\mathbf{f}_T(x)$ denote the distance from x to its *second* nearest neighbor in T.

We can apply existing algorithms for generating meshes and well-spaced points to compute a δ-sample on $\mathcal{D}\setminus\bigcup_i B(p_i, u_i)$, where $\mathcal{D} \subseteq \mathbb{R}^d$ is a domain, and $u_i = (1 - \delta^{2/3})r_P(p_i)$. The procedure consists of two steps:

1. Use the algorithm of [MSV13] to construct a well-spaced point set M (along with its associated approximate Delaunay graph) with aspect ratio τ in time $2^{O(d)}(n \log n + |M|)$.

2. Then over-refine M to S for the sizing function $g(x) = \frac{2\delta}{11\tau}\mathbf{f}_P(x)$ (while maintaining aspect ratio τ) in time $2^{O(d)}|S|$ by using the algorithm of Section 3.7 in [She11]. (see also [HOMS10] for an earlier use of this technique)

In the above algorithm, we will choose τ to be a fixed constant, say, $\tau = 6$. Both of the meshing algorithms listed above are chosen for their theoretical guarantees on running time. In practice, one could use any quality Delaunay meshing algorithm, popular choices include Triangle [She96] in $\mathbb{R}^2$ and Tetgen [Si11] or CGAL [ART$^+$12] in $\mathbb{R}^3$.

From the guarantees in ([She11]), we know that

$$|S| = O\left(\int_{\mathcal{D}} \frac{dx}{g(x)^d}\right) = \delta^{-O(d)} n \log \Delta, \tag{8}$$

where Δ is the *spread* of P, i.e., the ratio of the largest distance between two points in P to the smallest distance between two points in P.

Lemma 5.7. *If $x \in \mathcal{D} \setminus B(p_i, u_i)$, where $p_i = \mathrm{NN}(x)$, then $\mathrm{N}(x) \leq f_P(x) \leq 5\mathrm{N}(x)$.*

Proof. Note that $N(x) \leq f_P(x)$ is trivial. To prove the second half of the inequality, let q be the closest point to p in $P \setminus \{p_i\}$. Then, by the triangle inequality, we have that

$$
\begin{aligned}
f_P(x) &\leq & \mathbf{d}(x, q) \\
&\leq & \mathbf{d}(x, p_i) + \mathbf{d}(p_i, q) \\
&\leq & \mathbf{d}(x, p_i) + 2r_P(p_i) \\
&\leq & \mathbf{d}(x, p_i) + \frac{2}{1 - \delta^{2/3}}\mathbf{d}(x, p_i) \\
&= & 5\mathbf{d}(x, p_i) \\
&= & 5\mathrm{N}(x).
\end{aligned}
$$

$\square$

Now, it remains to show that the point set S is indeed a δ-sample on $\mathcal{D} \setminus \bigcup_i B(p_i, u_i)$. This is provided by the following lemma.

Lemma 5.8. *S is a δ-sample on $\mathcal{D} \setminus \bigcup_i B(p_i, u_i)$.*

Proof. Let $x \in \mathcal{D} \setminus \bigcup_i B(p_i, u_i)$ be an arbitrary point. Then, let q be the closest point to x that lies in S. Note that

$$
\begin{aligned}
\mathbf{d}(x, q) &\leq \tau \cdot r_S(q) \\
&\leq \tau\left(\frac{2\delta}{11\tau}\mathbf{f}_P(q)\right) \\
&\leq \frac{2\delta}{11}(\mathbf{f}_P(x) + \mathbf{d}(x, q)).
\end{aligned}
$$

This implies that

$$
\begin{aligned}
\mathbf{d}(x, q) &\leq \frac{2\delta}{11 - 2\delta}\mathbf{f}_P(x) \\
&\leq \frac{10\delta}{11 - 2\delta}\mathrm{N}(x) \\
&\leq \delta\mathrm{N}(x),
\end{aligned}
$$

as desired. $\qquad\square$

Now, we calculate the number of edges that will be present in the approximation graph defined in the previous section. For this, we require a few lemmas.

Lemma 5.9. *Let $A = B(p_i, r_P(p_i)) \setminus B(p_i, u_i)$ be an annulus around p_i. Then, $|A \cap S| = \delta^{-O(d)}$.*

Proof. By the meshing guarantees of [She11], we know that for any point $s \in A \cap S$, $B(s, t)$ does not contain a point from $S \setminus \{s\}$ for $t = \Omega(r_S(s)) = \Omega(\delta \cdot r_P(p))$. Thus, the desired result follows using a simple sphere packing argument. $\qquad\square$

Lemma 5.10. *If $s \in S$, then $|B(s, C_2\delta^{2/3}\mathrm{N}(s)) \cap S| = \delta^{-O(d)}$, where C_2 is the constant in Lemma 5.5.*

Proof. As in the previous lemma, meshing guarantees tell us that for any $s' \in B(s, C_2\delta^{2/3}\mathrm{N}(s))$, we have that $B(s', t)$ does not contain a point from $S \setminus \{s'\}$ for $t = \Omega(\delta \cdot \mathrm{N}(s')) = \Omega(\delta \cdot \mathrm{N}(s))$. Thus, we again obtain the desired result from a sphere packing argument. $\qquad\square$

From the above lemmas, we see that $\mathcal{A}$ is composed of $|S| = \delta^{-O(d)}n \log \Delta$ vertices and $n\delta^{-O(d)} + |S| \cdot \delta^{-O(d)} = |S| \cdot \delta^{-O(d)}$ edges.

Remark. Note that the right hand side of (8) is in terms of the spread, a non-combinatorial quantity. Indeed, one can construct examples of P for

which the integral in (8) is not bounded from above by any function of n. However, for many classes of inputs, one can obtain a tighter analysis. In particular, if P satisfies a property known as *well-paced*, one can show that the resulting set S will satisfy $|S| = 2^{O(d)}n$ (see [MPS08, She12]).

In a more general setting (without requiring that P is well-paced), one can modify the algorithms to produce output in the form of a *hierarchical mesh* [MPS11]. This then produces an output of size $2^{O(d)}n$, and $(1 + \varepsilon)$-approximation algorithm for the nearest neighbor metric can be suitably modified so that the underlying approximation graph uses a hierarchical set of points instead of a full δ-sample. However, we ignore the details here for the sake of simplicity of exposition.

The above remark, along with the edge count of $\mathcal{A}$ and the running time guarantees from [MSV13], yields Theorem 1.2, the main theorem of this section.

6 Discussion

Motivated by estimating geodesic distances within subsets of $\mathbb{R}^n$, we consider two distance metrics in this paper: the N-distance and the edge-squared distance. The main focus of this paper is to find an approximation of the N-distance. One possible drawback of our $(1 + \varepsilon)$-approximation algorithm is its exponential dependency on d. To alleviate this dependency a natural approach is using a Johnson-Lindenstrauss type projection. Thereby, we would like to ask which properties are preserved under random projections such as those in Johnson-Lindenstrauss transforms.

We are currently working on implementing the approximation algorithm presented in Section 4. We hope to show that this approximation is fast in practice as well as in theory.

Acknowledgement

References

[AG95]   Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications*, 5(01n02):75–91, 1995.

[ALMS98]   Lyudmil Aleksandrov, Mark Lanthier, Anil Maheshwari, and Jörg-Rüdiger Sack. An epsilon-approximation for weighted shortest paths on polyhedral surfaces. In *6th WS on Algorithm Theory*, pages 11–22, London, UK, UK, 1998.

[AMS00]   Lyudmil Aleksandrov, Anil Maheshwari, and Jörg-Rüdiger Sack. Approximation algorithms for geometric shortest path problems. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, STOC '00, pages 286–295, New York, NY, USA, 2000. ACM.

[AMS05]   L. Aleksandrov, A. Maheshwari, and J.-R. Sack. Determining approximate shortest paths on weighted polyhedral surfaces. *J. ACM*, 52(1):25–53, January 2005.

[ART+12]   Pierre Alliez, Laurent Rineau, Stéphane Tayeb, Jane Tournois, and Mariette Yvinec. 3D mesh generation. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.1 edition, 2012.

[BCH04]   Olivier Bousquet, Olivier Chapelle, and Matthias Hein. Measure based regularization. In *16th NIPS*, 2004.

[Ber96]   Johann Bernoulli. Branchistochrone problem. *Acta Eruditorum*, June 1696.

[BRS11]   Avleen Singh Bijral, Nathan D. Ratliff, and Nathan Srebro. Semi-supervised learning with density based distances. In Fabio Gagliardi Cozman and Avi Pfeffer, editors, *UAI*, pages 43–50. AUAI Press, 2011.

[DBVKOS00] Mark De Berg, Marc Van Kreveld, Mark Overmars, and Otfried Cheong Schwarzkopf. *Computational Geometry*. Springer, 2000.

[HDI14]   Sung Jin Hwang, Steven B. Damelin, and Alfred O. Hero III. Shortest path through random points. 2014. arXiv/1202.0045v3.

[HOMS10]   Benoît Hudson, Steve Y. Oudot, Gary L. Miller, and Donald R. Sheehy. Topological inference via meshing. In *SOCG: Proceedings of the 26th ACM Symposium on Computational Geometry*, 2010.

[Hp11]   Sariel Har-peled. *Geometric Approximation Algorithms*. American Mathematical Society, Boston, MA, USA, 2011.

[Kat10]   Svetlana Katok. Fuchsian groups, geodesic flows on surfaces of constant negative curvature and symbolic coding of geodesics. http://www.personal.psu.edu/sxk37/cmi.pdf, 2010.

[KH03]   J. Kim and J.P. Hespanha. Discrete approximations to continuous shortest-path: Application to minimum-risk path planning for groups of uavs. In *42nd IEEE ICDC*, Jan 2003.

[LSV06]   Tams Lukovszki, Christian Schindelhauer, and Klaus Volbert. Resource efficient maintenance of wireless network topologies. *Journal of Universal Computer Science*, 12(9):1292–1311, 2006.

[MP91]   Joseph S. B. Mitchell and Christos H. Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *J. ACM*, 38(1):18–73, January 1991.

[MPS08]   Gary L. Miller, Todd Phillips, and Donald R. Sheehy. Linear-size meshes. In *CCCG: Canadian Conference in Computational Geometry*, 2008.

[MPS11]   Gary L. Miller, Todd Phillips, and Donald R. Sheehy. Beating the spread: Time-optimal point meshing. In *SOCG: Proceedings of the 27th ACM Symposium on Computational Geometry*, 2011.

[MSV13]   Gary L. Miller, Donald R. Sheehy, and Ameya Velingker. A fast algorithm for well-spaced points and approximate delaunay graphs. In *29th SOCG*, SoCG '13, pages 289–298, New York, NY, USA, 2013. ACM.

[RR90]   Neil Rowe and Ron Ross. Optimal grid-free path planning across arbitrarily contoured terrain with anisotropic friction and gravity effects. pages 540–553, 1990.

[RS00]   John Reif and Zheng Sun. An efficient approximation algorithm for weighted region shortest path problem. In *Proceedings of the 4th Workshop on Algorithmic Foundations of Robotics*, WAFR '00, pages 191–203, 2000.

[She96]   Jonathan Richard Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In *Applied Computational Geometry*, volume 1148 of *Lecture Notes in Computer Science*, pages 203–222, 1996.

[She11]   Donald Sheehy. *Mesh Generation and Geometric Persistent Homology*. PhD thesis, Carnegie Mellon University, Pittsburgh, July 2011. CMU CS Tech Report CMU-CS-11-121.

[She12]   Donald R. Sheehy. New Bounds on the Size of Optimal Meshes. *Computer Graphics Forum*, 31(5):1627–1635, 2012.

[Si11]   Hang Si. TetGen: A quality tetrahedral mesh generator and a 3D Delaunay triangulator. http://tetgen.org/, January 2011.

[SO05]   Sajama and Alon Orlitsky. Estimating and computing density based distance metrics. In *ICML '05*, pages 760–767, New York, NY, USA, 2005. ACM.

[Str]   John Strain. Calculus of variation. http://math.berkeley.edu/~strain/170.S13/cov.pdf.

[Tsi95]   John N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40:1528–1538, 1995.

[VB03]   Pascal Vincent and Yoshua Bengio. Density sensitive metrics and kernels. In *Snowbird Workshop*, 2003.