

Hausdorff Distance under Translation for Points and Balls*

Pankaj K. Agarwal[†] Sariel Har-Peled[‡] Micha Sharir[§] Yusu Wang[¶]

February 6, 2006

Abstract

We study the shape matching problem under the Hausdorff distance and its variants. In the first part of the paper, we consider two sets \mathcal{A}, \mathcal{B} of balls in \mathbb{R}^d , $d = 2, 3$, and wish to find a translation t that minimizes the Hausdorff distance between $\mathcal{A} + t$, the set of all balls in \mathcal{A} shifted by t , and \mathcal{B} . We consider several variants of this problem. First, we extend the notion of Hausdorff distance from sets of points to sets of balls, so that each ball has to be matched with the nearest ball in the other set. We also consider the problem in the standard setting, by computing the Hausdorff distance between the unions of the two sets (as point sets). Second, we consider either all possible translations t (as is the standard approach), or consider only translations that keep the balls of $\mathcal{A} + t$ disjoint from those of \mathcal{B} . We propose several exact and approximation algorithms for these problems. In the second part of the paper, we note that the Hausdorff distance is sensitive to outliers, and thus consider two more robust variants—the root-mean-square (rms) and the summed Hausdorff distance. We propose efficient approximation algorithms for computing the minimum rms and the minimum summed Hausdorff distances under translation, between two point sets in \mathbb{R}^d . In order to obtain a fast algorithm for the summed Hausdorff distance, we propose a deterministic efficient dynamic data structure for maintaining an ε -approximation of the 1-median of a set of points in \mathbb{R}^d , under insertions and deletions.

*P.A. is supported by by NSF grants EIA-98-70724, EIA-99-72879, ITR-333-1050, CCR-97-32787, and CCR-00-86013, and by a grant from the U.S.-Israeli Binational Science Foundation (jointly with M.S.). S.H. is supported by NSF CAREER award CCR-01-32901. M.S. is also supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University. Y.W. is supported by NSF grants ITR-333-1050 and CCR-02-04118.

Part of this work was done while the last author was at Duke University.

A preliminary version of this paper appeared in *Proc. 19th Annual Sympos. Comput. Geom.*, 2003, pp. 282–291.

[†]Department of Computer Science, Duke University, Durham, NC 27708-0129, U.S.A. Email: pankaj@cs.duke.edu.

[‡]Department of Computer Science, University of Illinois, Urbana, IL 61801, U.S.A. Email: sariel@cs.uiuc.edu.

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. Email: michas@post.tau.ac.il.

[¶]Department of Computer Science and Engineering, Ohio State University, 2015 Neil Avenue, Columbus, OH 43210, USA. Email: yusu@cse.ohio-state.edu

1 Introduction

The problem of shape matching in two and three dimensions arises in a variety of applications, including computer graphics, computer vision, pattern recognition, computer aided design, and molecular biology [7, 18, 28]. For example, proteins with similar shapes are likely to have similar functionalities, therefore classifying proteins (or their fragments) based on their shapes is an important problem in computational biology. Similarly, the proclivity of two proteins binding with each other also depends on their shapes, so shape matching is central to the so-called *docking* problem in molecular biology [18].

Informally, the shape-matching problem can be described as follows: Given a distance measure between two sets of objects in \mathbb{R}^2 or \mathbb{R}^3 , determine a transformation, from an allowed set, that minimizes the distance between the sets. In many applications, the allowed transformations are all possible rigid motions. However, in certain applications there are constraints on the allowed transformations. For example, in matching the pieces of a jigsaw puzzle, it is important that no two pieces overlap each other in their matched positions. Another example is the aforementioned docking problem, where two molecules bind together to form a compound, and, clearly, at this docking position the molecules should occupy disjoint portions of space [18]. Moreover, because of efficiency considerations, one sometimes restricts further the set of allowed transformations, most typically to translations only.

Several distance measures between objects have been proposed, varying with the kind of input objects and the application. One common distance measure is the *Hausdorff distance* [7], originally proposed for point sets. In this paper we adopt this measure, extend it to sets of non-point objects (mainly, disks and balls), and apply it to several variants of the shape matching problem, with and without constraints on the allowed transformations. In many applications (e.g., molecular biology), shapes can be approximated by a finite union of balls [8], which is therefore the main type of input assumed in the first part of this paper.

1.1 Problem statement

Let \mathcal{A} and \mathcal{B} be two (possibly infinite) sets of geometric objects (e.g., points, balls, simplices) in \mathbb{R}^d , and let $d : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ be a distance function between objects in \mathcal{A} and in \mathcal{B} . For $a \in \mathcal{A}$, we define $d(a, \mathcal{B}) = \inf_{b \in \mathcal{B}} d(a, b)$. Similarly, we define $d(\mathcal{B}, a) = \inf_{b \in \mathcal{B}} d(b, a)$, for $b \in \mathcal{B}$. The *directional Hausdorff distance* between \mathcal{A} and \mathcal{B} is defined as

$$h(\mathcal{A}, \mathcal{B}) = \sup_{a \in \mathcal{A}} d(a, \mathcal{B}),$$

and the *Hausdorff distance* between \mathcal{A} and \mathcal{B} is defined as

$$H(\mathcal{A}, \mathcal{B}) = \max\{h(\mathcal{A}, \mathcal{B}), h(\mathcal{B}, \mathcal{A})\}.$$

(It is important to note that in this definition each object in \mathcal{A} or in \mathcal{B} is considered as a single entity, and not as the set of its points.) In order to measure similarity between \mathcal{A} and \mathcal{B} , we compute the minimum value of the Hausdorff distance over all translates of \mathcal{A} within a given set $T \subseteq \mathbb{R}^d$ of

allowed translation vectors. Namely, we define

$$\sigma(\mathcal{A}, \mathcal{B}; T) = \inf_{t \in T} H(\mathcal{A} + t, \mathcal{B}),$$

where $\mathcal{A} + t = \{a + t \mid a \in \mathcal{A}\}$. In our applications, T will either be the entire \mathbb{R}^d or the set of *collision-free* translates of \mathcal{A} at which none of its objects intersects any object of \mathcal{B} . The collision-free matching between objects is useful for applications (such as those mentioned above) in which the goal is to locate a transformation where the collective shape of one set of objects best complements that of the other set. We will use $\sigma(\mathcal{A}, \mathcal{B})$ to denote $\sigma(\mathcal{A}, \mathcal{B}; \mathbb{R}^d)$.

As already mentioned, our definition of (directional) Hausdorff distance is slightly different from the one typically used in the literature [7], in which one considers the two unions $\cup \mathcal{A}, \cup \mathcal{B}$ as two (possibly infinite) point sets, and computes the standard Hausdorff distance

$$H(\cup \mathcal{A}, \cup \mathcal{B}) = \max\{h(\cup \mathcal{A}, \cup \mathcal{B}), h(\cup \mathcal{B}, \cup \mathcal{A})\},$$

where

$$h(\cup \mathcal{A}, \cup \mathcal{B}) = \sup_{p \in \cup \mathcal{A}} d(p, \cup \mathcal{B}) = \sup_{p \in \cup \mathcal{A}} \inf_{q \in \cup \mathcal{B}} d(p, q).$$

We will denote $\cup \mathcal{A}$ (resp., $\cup \mathcal{B}$) as $U_{\mathcal{A}}$ (resp., $U_{\mathcal{B}}$), and use the notation $h_U(\mathcal{A}, \mathcal{B})$ to denote $h(U_{\mathcal{A}}, U_{\mathcal{B}})$. Analogous meanings hold for the notations $H_U(\mathcal{A}, \mathcal{B})$ and $\sigma_U(\mathcal{A}, \mathcal{B}; T)$.

A drawback of the directional Hausdorff distance (and thus of the Hausdorff distance) is its sensitivity to outliers in the given data. One possible approach to circumvent this problem is to use “partial matching” [15], but then one has to determine how many (and which) of the objects in \mathcal{A} should be matched to \mathcal{B} . Another possible approach is to use the *root-mean-square* (rms, for brevity) Hausdorff distance between \mathcal{A} and \mathcal{B} , defined by

$$\begin{aligned} h_R(\mathcal{A}, \mathcal{B}) &= \left(\frac{\int_{\mathcal{A}} d^2(a, \mathcal{B}) da}{\int_{\mathcal{A}} da} \right)^{1/2} \quad \text{and} \\ H_R(\mathcal{A}, \mathcal{B}) &= \max\{h_R(\mathcal{A}, \mathcal{B}), h_R(\mathcal{B}, \mathcal{A})\}, \end{aligned}$$

with an appropriate definition of integration (usually, summation over a finite set or the Lebesgue integration over infinite point sets). Define $\sigma_R(\mathcal{A}, \mathcal{B}; T) = \inf_{t \in T} H_R(\mathcal{A} + t, \mathcal{B})$. Finally, as in [24], we define the *summed Hausdorff distance* to be

$$h_S(\mathcal{A}, \mathcal{B}) = \frac{\int_{\mathcal{A}} d(a, \mathcal{B}) da}{\int_{\mathcal{A}} da},$$

and similarly define H_S and σ_S . Informally, $h(\mathcal{A}, \mathcal{B})$ can be regarded as an L_{∞} -distance over the sets of objects \mathcal{A} and \mathcal{B} . The two new definitions replace L_{∞} by L_2 and L_1 , respectively.

1.2 Previous results

It is beyond the scope of this paper to discuss all the results on shape matching. We refer the reader to [7, 18, 28] and references therein for a sample of known results. Here we summarize known results on shape matching using the Hausdorff distance measure.

Most of the early work on computing Hausdorff distance focused on finite point sets. Let \mathcal{A} and \mathcal{B} be two families of m and n points, respectively, in \mathbb{R}^d . In the plane, $H(\mathcal{A}, \mathcal{B})$ can be computed in $O((m+n) \log mn)$ time using Voronoi diagrams [5]. In \mathbb{R}^3 , it can be computed in time $O((m+n)^{4/3+\varepsilon})$, where $\varepsilon > 0$ is an arbitrarily small constant, using the data structure of Agarwal and Matoušek [2]. Huttenlocher *et al.* [23] showed that $\sigma(\mathcal{A}, \mathcal{B})$ can be computed in time $O(mn(m+n)\alpha(mn) \log mn)$ in \mathbb{R}^2 , and in time $O((mn)^2(m+n)^{1+\varepsilon})$ in \mathbb{R}^3 , for any $\varepsilon > 0$. Chew *et al.* [15] presented an $O((m+n)^{\lceil 3d/2 \rceil + 1} \log^3 mn)$ -time algorithm to compute $\sigma(\mathcal{A}, \mathcal{B})$ in \mathbb{R}^d for any $d \geq 2$. The minimum Hausdorff distance between \mathcal{A} and \mathcal{B} under rigid motion in \mathbb{R}^2 can be computed in $O((m+n)^6 \log mn)$ time [22].

Faster approximation algorithms to compute $\sigma(\mathcal{A}, \mathcal{B})$ were first proposed by Goodrich *et al.* [17]. Aichholzer *et al.* proposed a framework of approximation algorithms using *reference points* [4]. In \mathbb{R}^2 , their algorithm approximates the optimal Hausdorff distance within a constant factor, in $O((m+n) \log mn)$ time over all translations, and in $O(mn \log(mn) \log^* mn)$ time over rigid motions. The reference point approach can be extended to higher dimensions. However, it neither approximates the directional Hausdorff distance over a set of transformations, nor can it cope with the partial-matching problem.

Indyk *et al.* [24] study the partial matching problem, i.e., given a query $r > 0$, compute a rigid transform τ so that the number of points $p \in \mathcal{A}$ for which $d(\tau(p), \mathcal{B}) \leq r$ is maximized. They present algorithms for ε -approximating the maximum-size partial matching over the set of rigid motions in $O(mn\Delta/(r\varepsilon^2) \text{polylog}(\frac{n\Delta}{\varepsilon r}))$ time in \mathbb{R}^2 , and in $O(mn\Delta^3/(r^3\varepsilon^3) \text{polylog}(\frac{n\Delta}{\varepsilon r}))$ time in \mathbb{R}^3 , where Δ is the maximum of the spreads of the two point sets.¹ Their algorithm can be extended to approximate the minimum summed Hausdorff distance over rigid motions. Similar results were independently achieved in [14] using a different technique.

Algorithms for computing $H_U(\mathcal{A}, \mathcal{B})$ and $\sigma_U(\mathcal{A}, \mathcal{B})$, where \mathcal{A} and \mathcal{B} are sets of segments in the plane, or sets of simplices in higher dimensions are presented in [3, 5, 6]. Atallah [12] presents an algorithm for computing $H_U(\mathcal{A}, \mathcal{B})$ for two convex polygons in \mathbb{R}^2 . Agarwal *et al.* [3] provide an algorithm for computing $\sigma_U(\mathcal{A}, \mathcal{B})$, where \mathcal{A} and \mathcal{B} are two sets of m and n segments in \mathbb{R}^2 , respectively, in time $O((mn)^2 \log^3 mn)$. If rigid motions are allowed, the minimum Hausdorff distance between two sets of points in the plane can be computed in time $O((mn)^3 \log^2(mn))$ (Chew *et al.* [16]). Aichholzer *et al.* [4] present algorithms for approximating the minimum Hausdorff distance under different families of transformations for sets of points or of segments in \mathbb{R}^2 , and for sets of triangles in \mathbb{R}^3 , using reference points. Other than that, little is known about computing $\sigma_U(\mathcal{A}, \mathcal{B})$ or $\sigma(\mathcal{A}, \mathcal{B})$ where \mathcal{A} and \mathcal{B} are sets of simplices or other geometric shapes in higher dimensions.

1.3 Our results

In this paper, we develop efficient algorithms for computing $\sigma(\mathcal{A}, \mathcal{B}; T)$ and $\sigma_U(\mathcal{A}, \mathcal{B}; T)$ for sets of balls, and for approximating $\sigma_R(\mathcal{A}, \mathcal{B}), \sigma_S(\mathcal{A}, \mathcal{B})$ for sets of points in \mathbb{R}^d . Consequently, the paper consists of three parts, where the first two deal with the two variants of Hausdorff distances for balls, and the third part studies the rms and summed Hausdorff-distance problems for point sets.

Let $D(c, r)$ denote the ball in \mathbb{R}^d of radius r centered at c . Let $\mathcal{A} = \{A_1, \dots, A_m\}$ and $\mathcal{B} =$

¹The spread of a set of points is the ratio of its diameter to the closest-pair distance.

$\{B_1, \dots, B_n\}$ be two families of balls in \mathbb{R}^d , where $A_i = D(a_i, \rho_i)$ and $B_j = D(b_j, r_j)$, for each i and j . Let \mathcal{F} be the set of all translation vectors $t \in \mathbb{R}^d$ so that no ball of $\mathcal{A} + t$ intersects any ball of \mathcal{B} . We note, though, that balls of the same family can intersect each other, as is typically the case, e.g., in modeling molecules as collections of balls.

Section 2 considers the problem of computing the Hausdorff distance between two sets \mathcal{A} and \mathcal{B} of balls under the collision-free constraint, where the distance between two disjoint balls $A_i \in \mathcal{A}$ and $B_j \in \mathcal{B}$ is defined as $d(A_i, B_j) = d(a_i, b_j) - \rho_i - r_j$. We can regard this distance as an additively weighted Euclidean distance between the centers of A_i and B_j , and it is a common way of measuring distance between atoms in molecular biology [18]. In Section 2 we describe algorithms for computing $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in two and three dimensions. The running time is $O(mn(m+n) \log^4 mn)$ in \mathbb{R}^2 , and $O(m^2 n^2 (m+n) \log^4 mn)$ in \mathbb{R}^3 . The approach can be extended to solve the (collision-free) partial-matching problem under this variant of Hausdorff distance in the same asymptotic time complexity.

Section 3 considers the problem of computing $\sigma_U(\mathcal{A}, \mathcal{B})$ and $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$, i.e., of computing the Hausdorff distance between the union of \mathcal{A} and the union of \mathcal{B} , minimized over all translates of \mathcal{A} in \mathbb{R}^d or in \mathcal{F} . We first describe an $O(mn(m+n) \log^4 mn)$ -time algorithm for computing $\sigma_U(\mathcal{A}, \mathcal{B})$ and $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in \mathbb{R}^2 , which relies on several geometric properties of the *medial axis* of the union of disks. A straightforward extension of our algorithm to \mathbb{R}^3 is harder to analyze, and does not yield efficient bounds on its running time, mainly because little is known about the complexity of the medial axis of the union of balls in \mathbb{R}^3 [8]. We therefore consider approximation algorithms. In particular, given a parameter $\varepsilon > 0$, we compute a translation t , in time $O(((m+n)/\varepsilon^2) \log^3 mn)$ in \mathbb{R}^2 and in time $O(((m^2 + n^2)/\varepsilon^3) \log^2 mn)$ in \mathbb{R}^3 , such that $H_U(\mathcal{A} + t, \mathcal{B}) \leq (1 + \varepsilon)\sigma_U(\mathcal{A}, \mathcal{B})$. We also present a “pseudo-approximation” algorithm for computing $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$: Given an $\varepsilon > 0$, the algorithm computes a region $X \subseteq \mathbb{R}^d$ that serves as an ε -approximation of \mathcal{F} (in a sense defined formally in Section 3). It then returns a placement $t \in X$ such that

$$H_U(\mathcal{A} + t, \mathcal{B}; X) \leq (1 + \varepsilon)\sigma_U(\mathcal{A}, \mathcal{B}; X),$$

in time $O(((m^2 + n^2)/\varepsilon^3) \log^2 mn)$ in \mathbb{R}^3 . This variant of approximation makes sense in applications where the data is noisy and shallow penetrations between objects are allowed, as is the case in the docking problem [18].

Finally, we turn, in Section 4, to the two variants of rms and summed Hausdorff distances. Given two sets of points \mathcal{A} and \mathcal{B} in \mathbb{R}^d of size m and n , respectively, Section 4 describes an $O((mn/\varepsilon^d) \log(1/\varepsilon) \log(mn/\varepsilon))$ -time algorithm for computing an ε -approximation of $\sigma_R(\mathcal{A}, \mathcal{B})$. It also provides a data structure of size $O((mn/\varepsilon^d) \log(1/\varepsilon))$ so that, for a query vector $t \in \mathbb{R}^d$, an ε -approximation of $H_R(\mathcal{A} + t, \mathcal{B})$ can be computed in $O(\log(mn/\varepsilon))$ time. In fact, we solve a more general problem, which is interesting in its own right. Given a family P_1, \dots, P_ℓ of point sets in \mathbb{R}^d , with a total of N points, we construct a decomposition of \mathbb{R}^d into $O((N/\varepsilon^d) \log(1/\varepsilon))$ cells, which is an ε -approximation of each of the Voronoi diagrams of P_1, \dots, P_ℓ , in the sense defined in [10, 19]. We can preprocess this decomposition in $O((N/\varepsilon^d) \log(1/\varepsilon) \log(N/\varepsilon))$ time into a data structure of size $O((N/\varepsilon^d) \log(1/\varepsilon))$, so that for a query point q , an ε -approximate nearest neighbor of q in every P_i can be computed in a total time of $O(\log(N/\varepsilon) + \ell)$. Moreover, given a semigroup operation $+$, an ε -approximation of $\sum_{i=1}^{\ell} d^2(q, P_i)$ can be computed

in $O(\log(N/\varepsilon))$ time. We also extend the approach to obtain an algorithm that ε -approximates $\sigma_S(\mathcal{A}, \mathcal{B})$ in $O((mn/\varepsilon^{2d}) \text{polylog}(mn, 1/\varepsilon))$ time. This result relies on an efficient dynamic data structure for maintaining an ε -coreset of P for 1-median. That is, we show that we can maintain a weighted subset $Q \subseteq P$ of size $O((1/\varepsilon^d) \log(1/\varepsilon))$ so that the weighted sum of distances from any point $x \in \mathbb{R}^d$ to the points of Q is an ε -approximation of the sum of distances from x to the points of P . Q can be updated efficiently as the points in P are inserted or deleted. Using similar ideas, an algorithm for computing an ε -coreset of P for k -medians has been proposed in [21].

Indyk *et al.* [24] outline a near-linear-time approximation algorithm for computing $\sigma_R(\mathcal{A}, \mathcal{B})$ providing only sketchy details. Furthermore, the running time of their algorithm depends on the spread of the point set. We believe our technique have other applications and is of independent interest, and it has the advantage of supporting fast queries.

2 Collision-Free Hausdorff Distance between Sets of Balls

Let $\mathcal{A} = \{A_1, \dots, A_m\}$ and $\mathcal{B} = \{B_1, \dots, B_n\}$ be two sets of balls in \mathbb{R}^d , $d = 2, 3$. For two disjoint balls $A_i = D(a_i, \rho_i) \in \mathcal{A}$ and $B_j = D(b_j, r_j) \in \mathcal{B}$, we define

$$d(A_i, B_j) = d(a_i, b_j) - \rho_i - r_j,$$

namely, the (minimum) distance between A_i and B_j as point sets. Let \mathcal{F} be the set of placements t of \mathcal{A} such that no ball in $\mathcal{A} + t$ intersects any ball of \mathcal{B} . In this section, we describe an exact algorithm for computing $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$, and show that it can be extended to partial matching.

2.1 Computing $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in \mathbb{R}^2 and \mathbb{R}^3

As is common in geometric optimization, we first present an algorithm for the *decision problem*, namely, given a parameter $\delta > 0$, we wish to determine whether $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F}) \leq \delta$. We then use the parametric-searching technique [3, 26] to compute $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$. Given $\delta > 0$, for $1 \leq i \leq m$, let $V_i \subseteq \mathbb{R}^d$ be the set of vectors $t \in \mathbb{R}^d$ such that

$$(V) \quad 0 < \min_{1 \leq j \leq n} d(A_i + t, B_j) \leq \delta.$$

(In particular, $A_i + t$ does not intersect the interior of any $B_j \in \mathcal{B}$.)

Let $D_{ij}^- = D(b_j - a_i, \rho_i + r_j)$ and $D_{ij}^+ = D(b_j - a_i, \rho_i + r_j + \delta)$. Then $U_i^+ = \bigcup_{j \leq n} D_{ij}^+$ is the set of vectors that satisfy $\min_{1 \leq j \leq n} d(A_i + t, B_j) \leq \delta$, and the interior of $U_i^- = \bigcup_{j \leq n} D_{ij}^-$ violates $0 < \min_{1 \leq j \leq n} d(A_i + t, B_j)$. Hence, $V_i = \text{cl}(U_i^+ \setminus U_i^-)$. Let

$$V(\mathcal{A}, \mathcal{B}) = \bigcap_{1 \leq i \leq m} V_i = \text{cl} \left(\left(\bigcap_i U_i^+ \right) \setminus \left(\bigcup_i U_i^- \right) \right).$$

See Figure 1 for an illustration. By definition, $V(\mathcal{A}, \mathcal{B}) \subseteq \mathcal{F}$ is the set of vectors $t \in \mathcal{F}$ such that $h(\mathcal{A} + t, \mathcal{B}) \leq \delta$. Similarly, we define

$$V(\mathcal{B}, \mathcal{A}) \subseteq \mathcal{F} = \{t \in \mathcal{F} \mid h(\mathcal{B}, \mathcal{A} + t) \leq \delta\}.$$

Thus $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F}) \leq \delta$ if and only if $V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A}) \neq \emptyset$.

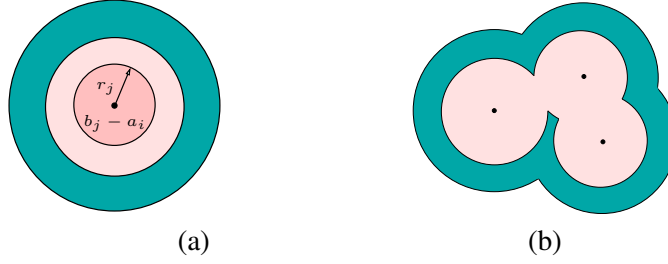


Figure 1: (a) Inner, middle and outer disks are $B_j - a_i$, D_{ij}^- , and D_{ij}^+ , respectively; (b) an example of V_i (dark region), which is the difference between U_i^+ (the whole union) and U_i^- (inner light region).

Lemma 2.1 *The combinatorial complexity of $V(\mathcal{A}, \mathcal{B})$ in \mathbb{R}^2 is $O(m^2n)$.*

Proof: If an edge of $\partial V(\mathcal{A}, \mathcal{B})$ is not adjacent to any vertex, then it is the entire circle bounding a disk of D_{ij}^+ or D_{ij}^- . There are $O(mn)$ such disks, so it suffices to bound the number of vertices in $V(\mathcal{A}, \mathcal{B})$.

Let v be a vertex of $V(\mathcal{A}, \mathcal{B})$; v is either a vertex of V_i , for some $1 \leq i \leq m$, or an intersection point of an edge in V_i and an edge in V_k , for some $1 \leq i \neq k \leq m$. In the latter case,

$$v \in V_i \cap V_k = (U_i^+ \cap U_k^+) \setminus (U_i^- \cup U_k^-).$$

In other words, a vertex of $V(\mathcal{A}, \mathcal{B})$ is a vertex of $U_i^+ \cap U_k^+$, $U_i^+ \setminus U_k^-$, $U_k^+ \setminus U_i^-$, or $U_i^- \cup U_k^-$, for $1 \leq i, k \leq m$. Observe that a vertex of $U_i^+ \cap U_k^+$ (resp., of $U_i^+ \setminus U_k^-$) that lies on both ∂U_i^+ and ∂U_k^+ (resp., ∂U_k^-) is also a vertex of $U_i^+ \cup U_k^+$ (resp., $U_i^+ \cup U_k^-$). Therefore, every vertex in $V(\mathcal{A}, \mathcal{B})$ is a vertex of $U_i^+ \cup U_k^+$, $U_i^+ \cup U_k^-$, $U_k^+ \cup U_i^-$, or $U_i^- \cup U_k^-$, for some $1 \leq i, k \leq m$. Since each U_i^+ , U_i^- is the union of a set of n disks, each of $U_i^- \cup U_k^+$, $U_i^+ \cup U_k^-$, $U_k^+ \cup U_i^-$, $U_i^- \cup U_k^-$ is the union of a set of $2n$ disks and thus has $O(n)$ vertices [25]. Hence, $V(\mathcal{A}, \mathcal{B})$ has $O(m^2n)$ vertices. \blacksquare

Lemma 2.2 *The combinatorial complexity of $V(\mathcal{A}, \mathcal{B})$ in \mathbb{R}^3 is $O(m^3n^2)$.*

Proof: The number of faces or edges of $V(\mathcal{A}, \mathcal{B})$ that do not contain any vertex is $O(n^2m^2)$ since they are defined by at most two balls in a family of $2mn$ balls. We therefore focus on the number of vertices in $V(\mathcal{A}, \mathcal{B})$. As in the proof of Lemma 2.1, any vertex $V(\mathcal{A}, \mathcal{B})$ satisfies:

$$v \in V_i \cap V_j \cap V_k = (U_i^+ \cap U_j^+ \cap U_k^+) \setminus (U_i^- \cup U_j^- \cup U_k^-),$$

for some $1 \leq i \leq j \leq k \leq m$. Again, such a vertex is also a vertex of $X_i \cup X_j \cup X_k$, where X_i is U_i^+ or U_i^- , and similarly for X_j , X_k . Since the union of r balls in \mathbb{R}^3 has $O(r^2)$ vertices, $X_i \cup X_j \cup X_k$ has $O(n^2)$ vertices, thereby implying that $V(\mathcal{A}, \mathcal{B})$ has $O(m^3n^2)$ vertices. \blacksquare

Similarly, we can prove that the complexity of $V(\mathcal{B}, \mathcal{A})$ is $O(n^2m)$ in \mathbb{R}^2 and $O(n^3m^2)$ in \mathbb{R}^3 . Extending the preceding arguments a little, we obtain the following.

Lemma 2.3 $V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A})$ has a combinatorial complexity of $O(mn(m+n))$ in \mathbb{R}^2 , and $O(m^2n^2(m+n))$ in \mathbb{R}^3 .

Remark. The above argument in fact bounds the complexity of the arrangement of $\mathcal{V} = \{V_1, V_2, \dots, V_m\}$. For example, in \mathbb{R}^2 , any intersection point of ∂V_i and ∂V_k lies on the boundary of $\partial(V_i \cap V_k)$, and we have argued that $V_i \cap V_k$ has $O(n)$ vertices. Hence, the entire arrangement has $O(m^2n)$ vertices in \mathbb{R}^2 .

We exploit a divide-and-conquer approach, combined with a plane-sweep, to compute $V(\mathcal{A}, \mathcal{B})$, $V(\mathcal{B}, \mathcal{A})$, and their intersections in \mathbb{R}^2 . For example, to compute $V(\mathcal{A}, \mathcal{B})$, we compute $V' = \bigcap_{i=1}^{n/2} V_i$ and $V'' = \bigcap_{i=n/2+1}^n V_i$ recursively, and merge $V(\mathcal{A}, \mathcal{B}) = V' \cap V''$ by a plane-sweep method. The overall running time is $O((m+n)mn \log mn)$.

To decide whether $V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A}) = \emptyset$ in \mathbb{R}^3 , it suffices to check whether

$$\Gamma_D = V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A}) \cap \partial D$$

is empty for all balls $D \in \{D_{ij}^-, D_{ij}^+ \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. Using the fact that the various D_{ij}^-, D_{ij}^+ meet any ∂D in a collection of spherical caps, we can compute Γ_D in time $O(mn(m+n) \log mn)$, by the same divide-and-conquer approach as computing $V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A})$ in \mathbb{R}^2 . Therefore we can determine in $O(m^2n^2(m+n) \log mn)$ time whether $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F}) \leq \delta$ in \mathbb{R}^3 .

Finally, the optimization problem can be solved by the parametric search technique [3]. In order to apply the parametric search technique, we need a parallel version of the above procedure. However, this divide-and-conquer paradigm uses plane-sweep during the conquer stage, which is not easy to parallelize. Instead, we use the algorithm of [3] to compute the union/intersection of two planar or spherical regions. It yields an overall parallel algorithm for determining whether $V(\mathcal{A}, \mathcal{B}) \cap V(\mathcal{B}, \mathcal{A})$ is empty in $O(\log^2 mn)$ time using $O(mn(m+n) \log mn)$ processors in \mathbb{R}^2 , and $O(m^2n^2(m+n) \log mn)$ processors in \mathbb{R}^3 . The standard technique of parametric searching then implies the following result.

Theorem 2.4 Given two sets \mathcal{A} and \mathcal{B} of m and n disks (or balls), we can compute $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in time $O(mn(m+n) \log^4 mn)$ in \mathbb{R}^2 , and in time $O(m^2n^2(m+n) \log^4 mn)$ in \mathbb{R}^3 .

2.2 Partial matching

Extending the definition of partial matching in [24], we define the partial collision-free Hausdorff distance problem as follows.

Given an integer k , let $h_k(\mathcal{A}, \mathcal{B})$ denote the k^{th} largest value in the set $\{d(a, \mathcal{B}) \mid a \in \mathcal{A}\}$; note that $h(\mathcal{A}, \mathcal{B}) = h_1(\mathcal{A}, \mathcal{B})$. We define $h_k(\mathcal{B}, \mathcal{A})$ in a fully symmetric manner, and then define $H_k(\mathcal{A}, \mathcal{B})$, $\sigma_k(\mathcal{A}, \mathcal{B}; T)$ as above. The preceding algorithm can be extended to compute $\sigma_k(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in the same asymptotic time complexity. We briefly illustrate the two-dimensional case. Let $\mathcal{V} = \{V_1, V_2, \dots, V_m\}$ be as defined above, and let $\Xi(\mathcal{V})$ be the arrangement of \mathcal{V} . For each cell $\Delta \in \Xi(\mathcal{V})$, let $\chi(\Delta)$ be the number of V_i 's that fully contain Δ . Note that for any point t in a cell Δ with $\chi(\Delta) > (m-k)$, $h_k(\mathcal{A} + t, \mathcal{B}) \leq \delta$, and vice versa. Hence, we compute $\Xi(\mathcal{V})$ and

$\chi(\Delta)$ for each cell $\Delta \in \Xi(\mathcal{V})$, and then discard all the cells Δ for which $\chi(\Delta) \leq (m - k)$. The remaining cells form the set $T_1 = \{t \mid h_k(\mathcal{A} + t, \mathcal{B}) \leq \delta\}$. By the Remark following Lemma 2.2, Ξ has $O(m^2n)$ vertices, and it can be computed in $O(m^2n \log mn)$ time. Therefore, T_1 can be computed in $O(m^2n \log mn)$ time. Similarly, we can compute $T_2 = \{t \mid h_k(\mathcal{B}, \mathcal{A} + t) \leq \delta\}$ in $O(mn^2 \log mn)$ time, and we can determine in $O(mn(m+n) \log mn)$ time whether $T_1 \cap T_2 \neq \emptyset$. Similar arguments can solve the partial matching problem in \mathbb{R}^3 , by computing the sets T_1, T_2 , and by checking for their intersection along the boundary of each of the balls D_{ij}^+, D_{ij}^- . Putting everything together, we obtain the following.

Theorem 2.5 *Let \mathcal{A} and \mathcal{B} be two families of m and n balls, respectively, and let $k \geq 0$ be an integer; we can compute $\sigma_k(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in $O(mn(m+n) \log^4 mn)$ time in \mathbb{R}^2 , and in $O(m^2n^2(m+n) \log^4 mn)$ time in \mathbb{R}^3 .*

3 Hausdorff Distance between Unions of Balls

In Section 3.1 we describe an algorithm for computing $\sigma_U(\mathcal{A}, \mathcal{B})$ in \mathbb{R}^2 . The same approach can be extended to compute $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$ within the same asymptotic time complexity. In Section 3.2, we present approximation algorithms for the same problem in \mathbb{R}^2 and \mathbb{R}^3 .

3.1 The exact 2D algorithm

Let $\mathcal{A} = \{A_1, \dots, A_m\}$ and $\mathcal{B} = \{B_1, \dots, B_n\}$ be two sets of disks in the plane. Write, as above, $A_i = D(a_i, \rho_i)$, for $i = 1, \dots, m$, and $B_j = D(b_j, r_j)$, for $j = 1, \dots, n$. Let $U_{\mathcal{A}}$ (resp., $U_{\mathcal{B}}$) be the union of the disks in \mathcal{A} (resp., \mathcal{B}). As in Section 2, we focus on the decision problem for a given distance parameter $\delta > 0$.

For any point p , we have

$$\begin{aligned} d(p, U_{\mathcal{B}}) &= \min_{q \in U_{\mathcal{B}}} d(p, q) = \min_{1 \leq j \leq n} d(p, B_j) \\ &= \min_{1 \leq j \leq n} \max\{d(p, b_j) - r_j, 0\}. \end{aligned}$$

This value is greater than δ if and only if

$$\min_{1 \leq j \leq n} \left(d(p, b_j) - (r_j + \delta) \right) > 0.$$

In other words, $h_U(\mathcal{A} + t, \mathcal{B}) > \delta$ if and only if there exists a point $p \in U_{\mathcal{A}}$ such that $p + t \notin U_{\mathcal{B}}(\delta) = \bigcup_{j=1}^n B_j(\delta)$, where $B_j(\delta) = D(b_j, r_j + \delta)$ is the disk B_j expanded by δ .

Let

$$T_1 = \{t \mid h_U(\mathcal{A} + t, \mathcal{B}) \leq \delta\};$$

T_1 is the set of all translations t such that $U_{\mathcal{A}} + t \subseteq U_{\mathcal{B}}(\delta)$. Our decision procedure computes the set T_1 and the analogously defined set

$$T_2 = \{t \mid h_U(\mathcal{B}, \mathcal{A} + t) \leq \delta\},$$

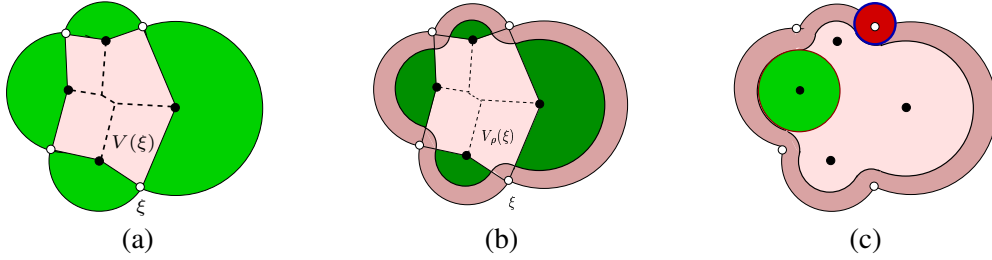


Figure 2: (a) The medial axis (dotted segments) of the union of four disks centered at the solid points: The Voronoi diagram of the boundary decomposes the union into 8 cells; (b) shrinking by ρ the Voronoi cell $V(\xi)$ of each boundary element ξ of the union; (c) The boundary of the lighter-colored disk contains a convex arc, and the boundary of the darker-colored disk contains a concave arc.

and then tests whether $T_1 \cap T_2 \neq \emptyset$. To understand the structure of T_1 , we first study the case in which \mathcal{A} consists of just one disk A , with center a and radius ρ . For simplicity of notation, we denote $U_{\mathcal{B}}(\delta)$ temporarily by U . Let Q denote the set of vertices of ∂U , and Γ the set of (relatively open) edges of ∂U ; we have $|Q| \leq |\Gamma| \leq 6n - 12$ [25].

Consider the Voronoi diagram $\text{Vor}(Q \cup \Gamma)$ of the boundary features of U , clipped to within U . This is a decomposition of U into cells, so that, for each $\xi \in Q \cup \Gamma$, the cell $V(\xi)$ of ξ is the set of points $x \in U$ such that $d(x, \xi) \leq d(x, \xi')$, for all $\xi' \in Q \cup \Gamma$. The diagram is closely related to the medial axis of ∂U . See Figure 2(a). For each $\gamma \in \Gamma$, let $W(\gamma)$ denote the circular sector spanned by γ within the disk $B_j(\delta)$ whose boundary is γ , and let $U' = U \setminus \bigcup_{\gamma \in \Gamma} W(\gamma)$. The diagram has the following structure. (A slight variant of the following lemma was observed in [8].)

Lemma 3.1 (a) For each $\gamma \in \Gamma$, we have $V(\gamma) = W(\gamma)$.
 (b) For each $\xi \in Q$, we have $V(\xi) = U' \cap V'(\xi)$, where $V'(\xi)$ is the Voronoi cell of ξ in the Voronoi diagram $\text{Vor}(Q)$ of Q . Moreover, $V(\xi)$ is a convex polygon.

Lemma 3.1 implies that $\text{Vor}(Q \cup \Gamma)$ yields a convex decomposition of U of linear size. The medial axis of ∂U consists of all the edges of those cells $V'(\xi)$, for $\xi \in Q$, that are also edges of $\text{Vor}(Q \cup \Gamma)$ (the dashed cells of Figure 2(a)).

Returning to the study of the structure of T_1 , we have, by definition, $A + t \subseteq U$ if and only if $d(a + t, \xi) \geq \rho$, where ξ is the feature of $Q \cup \Gamma$ whose cell contains $a + t$. This implies that the set $T_1(A)$ of all translations t of A for which $A + t \subseteq U$ is given by

$$T_1(A) = \left(\bigcup_{\xi \in Q \cup \Gamma} V_\rho(\xi) \right) - a,$$

where

$$V_\rho(\xi) = \{x \in V(\xi) \mid d(x, \xi) \geq \rho\}.$$

For $\gamma \in \Gamma$, $V_\rho(\gamma)$ is the sector obtained from $W(\gamma)$ by shrinking it by distance ρ towards its center. For $\xi \in Q$, $V_\rho(\xi) = V(\xi) \setminus D(\xi, \rho)$. See Figure 2(b) for an illustration.

Now return to the original case in which \mathcal{A} consists of m disks; we obtain

$$T_1 = \bigcap_{i=1}^m T_1(A_i) = \bigcap_{i=1}^m \bigcup_{\xi \in Q \cup \Gamma} (V_{\rho_i}(\xi) - a_i).$$

Note that each $T_1(A_i)$ is bounded by $O(n)$ circular arcs, some of which are *convex* (those bounding shrunk sectors), and some are *concave* (those bounding shrunk Voronoi cells of vertices). Convex arcs are bounded by disks $D(b_k - a_i, r_k + \delta - \rho_i)$, for some $1 \leq k \leq n$, while concave arcs are bounded by disks $D(\xi - a_i, \rho_i)$ for $\xi \in Q$. Furthermore, since $T_1(A_i)$ is obtained by removing all points $x \in U$ such that the nearest distance from x to ∂U is smaller than ρ_i , we have that: (i) $D(b_k - a_i, r_k + \delta - \rho_i) \subseteq T_1(A_i)$; and (ii) $D(\xi - a_i, \rho_i) \cap (T_1(A_i) \setminus \partial(T_1(A_i))) = \emptyset$. See Figure 2 (c) for an illustration.

Lemma 3.2 *For any pair of disks $A_i, A_j \in \mathcal{A}$, the complexity of $T_1(A_i) \cap T_1(A_j)$ is $O(n)$.*

Proof: Clearly, $T_1(A_i) \cap T_1(A_j)$ is bounded by circular arcs, whose endpoints are either vertices of $T_1(A_i)$ or $T_1(A_j)$, or intersection points between an arc of $\partial T_1(A_i)$ and an arc of $\partial T_1(A_j)$. It suffices to estimate the number of vertices of the latter kind.

Consider the set \mathcal{B}'_{ij} of the $2n + 2|Q|$ disks

$$\{D(b_k - a_i, r_k + \delta - \rho_i), D(b_k - a_j, r_k + \delta - \rho_j)\}_{1 \leq k \leq n} \cup \{D(\xi - a_i, \rho_i), D(\xi - a_j, \rho_j)\}_{\xi \in Q}.$$

We claim that any intersection point between two arcs, one from $\partial T_1(A_i)$ and one from $\partial T_1(A_j)$, lies on $\partial(\cup \mathcal{B}'_{ij})$. Indeed, assume that x is such an intersection point that does not lie on $\partial(\cup \mathcal{B}'_{ij})$. Then it has to lie in the interior of $\cup \mathcal{B}'_{ij}$. That is, there is a disk $D \in \mathcal{B}'_{ij}$ that contains x . There are two possibilities for the choice of D .

- (i) $D = D(b_k - a_i, r_k + \delta - \rho_i)$ (resp., $D = D(b_k - a_j, r_k + \delta - \rho_j)$), for some $1 \leq k \leq n$. The boundary of such a disk contains some convex arc on $\partial T_1(A_i)$ (resp., $\partial T_1(A_j)$), and $D \subseteq T_1(A_i)$ (resp., $D \subseteq T_1(A_j)$). As such, x cannot appear on the boundary of $\partial T_1(A_i)$ (resp., $\partial T_1(A_j)$), contrary to assumption.
- (ii) $D = D(\xi - a_i, \rho_i)$ (resp., $D = D(\xi - a_j, \rho_j)$), for some $\xi \in Q$. Recall that Q is the set of vertices on the boundary of ∂U . Therefore, by definition, $A_i + x$ (resp., $A_j + x$) contains ξ in its interior, so it cannot be fully contained in U , implying that $x \notin T_1(A_i)$ (resp., $x \notin T_1(A_j)$), again a contradiction.

These contradictions imply the claim. It then follows, using the bound of [25], that the number of intersections under consideration is at most $6 \cdot (2n + 2|Q|) - 12 = O(n)$. \blacksquare

Each vertex of T_1 is also a vertex of some $T_1(A_i) \cap T_1(A_j)$. Applying the preceding lemma to all the $O(m^2)$ pairs A_i, A_j , we obtain the following.

Lemma 3.3 *The complexity of T_1 is $O(m^2n)$, and it can be computed in $O(m^2n \log mn)$ time.*

Similarly, the set T_2 has complexity $O(mn^2)$ and can be computed in time $O(mn^2 \log mn)$. Finally, we can determine whether $T_1 \cap T_2 \neq \emptyset$, by plane sweep, in time $O(mn(m+n) \log mn)$. Using parametric search, as in [3], $\sigma_U(\mathcal{A}, \mathcal{B})$ can be computed in $O(mn(m+n) \log^4 mn)$ time.

To compute $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$, we follow the same approach as computing $\sigma(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in the preceding section. Specifically, we need to modify the definition of T_1 and of T_2 , to require also that no disk of $\mathcal{A} + t$ intersect any disk of \mathcal{B} . This amounts, in the case of T_1 , to redefine each $T_1(A)$ to consist of all $t \in \mathbb{R}^2$ such that $A + t \subseteq U$ and $(A + t) \cap U_{\mathcal{B}} = \emptyset$. The latter is equivalent to requiring that $t \notin U_{\mathcal{B}}(\rho) - a$. Hence

$$T_1 = \bigcap_{i=1}^m T_1(A_i) = \bigcap_{i=1}^m \bigcup_{\xi \in Q \cup \Gamma} (V_{\rho_i}(\xi) - a_i) \setminus \bigcup_{i=1}^m (U_{\mathcal{B}}(\rho_i) - a_i).$$

It is now easy to modify the arguments in the proof of Lemma 3.2, to conclude that the complexity of T_1 (and, symmetrically of T_2) remains asymptotically the same, which then implies the following result.

Theorem 3.4 *Given two families \mathcal{A} and \mathcal{B} of m and n disks in \mathbb{R}^2 , we can compute both $\sigma_U(\mathcal{A}, \mathcal{B})$ and $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in time $O(mn(m+n) \log^4 mn)$.*

3.2 Approximation algorithms

No good bounds are known for the complexity of the Voronoi diagram of the boundary of the union of n balls in \mathbb{R}^3 , or, more precisely, for the complexity of the portion of the diagram inside the union [8]. The best known bound is $O(n^4)$. Hence, a naïve extension of the preceding exact algorithm to \mathbb{R}^3 yields an algorithm whose running time is hard to calibrate, and only rather weak upper bounds can be derived. We therefore resort to approximation algorithms.

Approximating $\sigma_U(\mathcal{A}, \mathcal{B})$ in \mathbb{R}^2 and \mathbb{R}^3 . Given a parameter $\varepsilon > 0$, we wish to compute a translation t of \mathcal{A} such that $H_U(\mathcal{A} + t, \mathcal{B}) \leq (1 + \varepsilon)\sigma_U(\mathcal{A}, \mathcal{B})$, i.e., $H_U(\mathcal{A} + t, \mathcal{B})$ is an ε -approximation of $\sigma_U(\mathcal{A}, \mathcal{B})$. Our approximation algorithm for $\sigma_U(\mathcal{A}, \mathcal{B})$ follows the same approach as the one used in [4, 5]. That is, let $r(\mathcal{A})$ (resp., $r(\mathcal{B})$) denote the point with smallest coordinates, called the *reference point*, of the axis-parallel bounding box of $U_{\mathcal{A}}$ (resp., $U_{\mathcal{B}}$). Set $\tau = r(\mathcal{B}) - r(\mathcal{A})$. It is shown in [5] that in \mathbb{R}^d ,

$$\sigma_U(\mathcal{A}, \mathcal{B}) \leq H_U(\mathcal{A} + \tau, \mathcal{B}) \leq (1 + \sqrt{d})\sigma_U(\mathcal{A}, \mathcal{B}),$$

and that the optimal translation lies in a disk of radius $H_U(\mathcal{A} + \tau, \mathcal{B})$ centered at τ . Computing τ takes $O(m+n)$ time. We compute $H_U(\mathcal{A} + \tau, \mathcal{B})$ using the parametric search technique [3], which is based on the following simple implementation of the decision procedure:

Fix a parameter $\delta > 0$, and put $U_{\mathcal{A}}(\delta) = \bigcup_i D(a_i, \rho_i + \delta)$ and $U_{\mathcal{B}}(\delta) = \bigcup_j D(b_j, r_j + \delta)$. We observe that $H_U(\mathcal{A} + t, \mathcal{B}) \leq \delta$ if and only if $U_{\mathcal{A}} + t \subseteq U_{\mathcal{B}}(\delta)$ and $U_{\mathcal{B}} \subseteq U_{\mathcal{A}}(\delta) + t$. To test whether $U_{\mathcal{A}} + t \subseteq U_{\mathcal{B}}(\delta)$, we compute $(U_{\mathcal{A}} + t) \cup U_{\mathcal{B}}(\delta)$, the union of the balls in $\mathcal{A} + t$ and of the δ -expanded balls in \mathcal{B} , and check whether any ball of \mathcal{A} appears on its boundary. If not, then

$U_{\mathcal{A}} + t \subseteq U_{\mathcal{B}}(\delta)$. Similarly, we test whether $U_{\mathcal{B}} \subseteq U_{\mathcal{A}}(\delta) + t$. The total time spent is proportional to the time needed to compute the union of $m + n$ balls, which is $O((m + n) \log(m + n))$ in \mathbb{R}^2 , and $O((m + n)^2)$ in \mathbb{R}^3 , as it can be reduced to computing a convex hull of $m + n$ points in \mathbb{R}^3 and \mathbb{R}^4 , respectively [29]. Plugging this into the parametric searching technique, the time spent in computing $H_U(\mathcal{A} + t, \mathcal{B})$ in \mathbb{R}^2 and \mathbb{R}^3 is $O((m + n) \log^3(mn))$ and $O((m^2 + n^2) \log^2(mn))$, respectively.

In order to compute an ε -approximation of $\sigma_U(\mathcal{A}, \mathcal{B})$ from this constant-factor approximation, we use the standard trick [4] of placing a grid of cell size $\frac{\varepsilon}{1+\sqrt{d}} \cdot H_U(\mathcal{A} + \tau, \mathcal{B})$ in the disk of radius $H_U(\mathcal{A} + t, \mathcal{B})$ centered at τ , and returning the smallest $H_U(\mathcal{A} + t, \mathcal{B})$, where t ranges over the grid points. We thus obtain the following result.

Theorem 3.5 *Given two sets of balls, \mathcal{A} and \mathcal{B} , of size m and n , respectively, and $\varepsilon > 0$, an ε -approximation of $\sigma_U(\mathcal{A}, \mathcal{B})$ can be computed in $O(((m + n)/\varepsilon^2) \log^3 mn)$ time in \mathbb{R}^2 , and in $O(((m^2 + n^2)/\varepsilon^3) \log^2 mn)$ time in \mathbb{R}^3 .*

Pseudo-approximation for $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$. Currently, we do not have an efficient algorithm to ε -approximate $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$ in \mathbb{R}^3 . Instead, we present a ‘‘pseudo-approximation’’ algorithm, in the following sense.

The set $\mathcal{K} = U_{\mathcal{B}} \oplus (-U_{\mathcal{A}})$, where \oplus denotes the Minkowski sum, is the set of all placements of \mathcal{A} at which $U_{\mathcal{A}}$ intersects $U_{\mathcal{B}}$; we have $\mathcal{K} = \bigcup_{i,j} D(b_j - a_i, \rho_i + r_j)$, and $\mathcal{F} = cl(\mathbb{R}^3 \setminus \mathcal{K})$. For a parameter $\varepsilon \geq 0$, let

$$\mathcal{K}(\varepsilon) = \bigcup_{i,j} D(b_j - a_i, (1 - \varepsilon)(\rho_i + r_j)),$$

and $\mathcal{F}(\varepsilon) = cl(\mathbb{R}^3 \setminus \mathcal{K}(\varepsilon))$. We call a region $X \subseteq \mathbb{R}^3$ ε -free if $\mathcal{F} \subseteq X \subseteq \mathcal{F}(\varepsilon)$.

This notion of approximating \mathcal{F} is motivated by some applications in which the data is noisy, and/or shallow penetration is allowed. For example, each atom in a protein is best modeled as a ‘‘fuzzy’’ ball rather than a hard ball [18]. We can model this fuzziness by allowing any atom $D(b, r)$ to be intersected by other atoms, but only within the shell $D(b, r) \setminus D(b, (1 - \varepsilon)r)$ for some $\varepsilon > 0$. In this way, the atoms of two docking molecules may penetrate a little at the desired placement. Although \mathcal{F} can have large complexity, namely, up to $O(m^2 n^2)$ in \mathbb{R}^3 , we present a technique for constructing an ε -free region X of considerably smaller complexity. We thus compute X and a placement $t^* \in X$ such that $H_U(\mathcal{A} + t^*, \mathcal{B}) \leq (1 + \varepsilon)\sigma_U(\mathcal{A}, \mathcal{B}; X)$. We refer to such an approximation $H_U(\mathcal{A} + t^*, \mathcal{B})$ as a *pseudo- ε -approximation* for $\sigma_U(\mathcal{A}, \mathcal{B}; \mathcal{F})$.

Lemma 3.6 *In \mathbb{R}^3 , an ε -free region X of size $O(mn/\varepsilon^3)$ can be computed in time $O((mn/\varepsilon^3) \log(mn/\varepsilon))$.*

Proof: Let $\mathcal{D} = \{D_{ij} = D(b_j - a_i, \rho_i + r_j) \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. We insert each ball $D_{ij} \in \mathcal{D}$ into an oct-tree T . Let C_v denote the cube associated with a node v of T . In order to insert D_{ij} , we visit T in a top-down manner. Suppose we are at a node v . If $C_v \subseteq D_{ij}$, we mark v as black and stop. If $C_v \cap D_{ij} \neq \emptyset$ and the size of C_v is at least $\varepsilon(\rho_i + r_j)/2$, then we recursively visit the children of v . Otherwise, we stop, leaving v unmarked. After we insert all balls from \mathcal{D} , if all eight

children of a node v are marked black, we mark v as black too. Let $V = \{v_1, v_2, \dots, v_k\}$ be the set of highest marked nodes, i.e., each v_i is marked black but none of its ancestors is black. It is easy to verify that each D_{ij} marks at most $O(1/\varepsilon^3)$ nodes as black, because the nodes a fixed D_{ij} marks are disjoint and of size at least $\varepsilon(\rho_i + r_j)/2$; thus $|V| = O(mn/\varepsilon^3)$. The whole construction takes $O((mn/\varepsilon^3) \log(mn/\varepsilon))$ time, and obviously $\mathcal{K}(\varepsilon) \subseteq \bigcup_{v \in V} C_v \subseteq \mathcal{K}$. Set $X = \text{cl}(\mathbb{R}^3 \setminus \bigcup_{v \in V} C_v)$; it is an ε -free region, as claimed. \blacksquare

Furthermore, let $r(\mathcal{B}), r(\mathcal{A})$, and $\tau = r(\mathcal{B}) - r(\mathcal{A})$ be as defined earlier in this section. We prove the following result.

Lemma 3.7 *Let $t^* \in X$ be the closest point of τ in X . Then*

$$H_U(\mathcal{A} + t^*, \mathcal{B}) \leq (1 + 2\sqrt{3})\sigma_U(\mathcal{A}, \mathcal{B}; X).$$

Proof: Let $\hat{\delta} = \sigma_U(\mathcal{A}, \mathcal{B}; X)$ and $\hat{t} \in X$ the placement so that $H_U(\mathcal{A} + \hat{t}, \mathcal{B}) = \hat{\delta}$. Then

$$\|\hat{t} - \tau\| = \|\hat{t} - r(\mathcal{B}) + r(\mathcal{A})\| = d(r(\mathcal{A}) + \hat{t}, r(\mathcal{B})).$$

A result in [4] implies that $d(r(\mathcal{A}) + \hat{t}, r(\mathcal{B})) \leq \sqrt{3}\hat{\delta}$. On the other hand,

$$\begin{aligned} H_u(\mathcal{A} + t^*, \mathcal{B}) &\leq \hat{\delta} + \|\hat{t} - t^*\| \leq \hat{\delta} + \|\hat{t} - \tau\| + \|\tau - t^*\| \\ &\leq \hat{\delta} + 2\|\tau - \hat{t}\| \leq \hat{\delta} + 2\sqrt{3}\hat{\delta} = (1 + 2\sqrt{3})\sigma_U(\mathcal{A}, \mathcal{B}; X). \end{aligned}$$

The point $t^* \in X$ closest to τ can be computed as follows. Recall that in Lemma 3.6, $X = \text{cl}(\mathbb{R}^3 \setminus \bigcup_{v \in V} C_v)$. Set $\bar{X} = \text{cl}(\mathbb{R}^3 \setminus X) = \bigcup_{v \in V} C_v$; \bar{X} consists of a set of openly disjoint cubes. We first check whether $\tau \in \bar{X}$ by a point-location operation. If the answer is no, then $\tau \in X$, and we return $t^* = \tau$. Otherwise, t^* is a point on $\partial X = \partial \bar{X}$ that is closest to τ . In that case, t^* is either a vertex of a cube in V , or lies in the interior of an edge or of a face of a cube in V . For each node $v \in V$ and for each boundary feature $\xi \subset C_v$, that is, a face, an edge, or a vertex of C_v , we compute the point in ξ closest to τ . Let Q_v be the resulting set of closest points. We then check, for each $q \in Q_v$, whether $q \in \partial \bar{X}$, by testing whether at least one neighboring cube is unmarked; there are $O(1)$ neighboring cubes of q . This can be achieved by performing point-location operations in T . Finally, from among those points of Q_v that lie on $\partial \bar{X}$ (thus on ∂X), we return the one that is closest to τ . There are $O(mn/\varepsilon^3)$ cubes, and each has constant number of boundary features. Furthermore, at most a constant number of nodes in V contain a given point, and each point-location operation takes $O(\log(mn/\varepsilon))$ time. Hence, t^* can be computed in $O((mn/\varepsilon^3) \log(mn/\varepsilon))$ time. \blacksquare

We can compute $H_U(\mathcal{A} + t^*, \mathcal{B})$ in $O((n^2 + m^2) \log^2 mn)$ time, as described in Section 3.2, so we can approximate $\sigma_U(\mathcal{A}, \mathcal{B}; X)$, up to a constant factor, in $O((n^2 + m^2) \log^2 mn)$ time. We then draw an appropriate grid around t^* and use it to compute an ε -approximation of $\sigma_U(\mathcal{A}, \mathcal{B}; X)$, as in Section 3.2, with the difference that we only test those grid points that lie in X . We thus obtain the following result.

Theorem 3.8 Given \mathcal{A}, \mathcal{B} in \mathbb{R}^3 and $\varepsilon > 0$, we can compute in $O(((n^2 + m^2)/\varepsilon^3) \log^2 mn)$ time, an ε -free region $X \subseteq \mathbb{R}^3$ and a placement $t \in X$ of \mathcal{A} , such that

$$H_U(\mathcal{A} + t, \mathcal{B}) \leq (1 + \varepsilon)\sigma_U(\mathcal{A}, \mathcal{B}; X).$$

4 RMS and Summed Hausdorff Distance between Points

We first establish a result on simultaneous approximation of the Voronoi diagrams of several point sets, which we believe to be of independent interest, and then we apply this result to approximate $\sigma_R(\mathcal{A}, \mathcal{B})$ and $\sigma_S(\mathcal{A}, \mathcal{B})$ for point sets $\mathcal{A} = \{a_1, \dots, a_m\}$ and $\mathcal{B} = \{b_1, \dots, b_n\}$ in any dimension.

4.1 Simultaneous approximation of Voronoi diagrams

Given a family $\{P_1, \dots, P_\ell\}$ of point sets in \mathbb{R}^d , with a total of N points, and a parameter $\varepsilon > 0$, we wish to construct a subdivision of \mathbb{R}^d , so that, for any $x \in \mathbb{R}^d$, we can quickly compute points $p_i \in P_i$, for all $1 \leq i \leq \ell$, with the property that $d(x, p_i) \leq (1 + \varepsilon)d(x, P_i)$, where $d(x, P_i) = \min_{q \in P_i} d(x, q)$. Arya and Malamatos [10] proposed a data structure that can answer an ε -approximate nearest-neighbor query in time $O(\log(n/\varepsilon))$ using $O((n/\varepsilon^d) \log(1/\varepsilon))$ space and $O((n/\varepsilon^d) \log(n/\varepsilon) \log(1/\varepsilon))$ preprocessing. Constructing this data structure for each P_i separately, one can answer the above query in time $O(\ell \log(n/\varepsilon))$ using $O((N/\varepsilon^d) \log(1/\varepsilon))$ space and $O((N/\varepsilon^d) \log(N/\varepsilon) \log(1/\varepsilon))$ preprocessing. We adapt this data structure so that the query time can be improved to $O(\log(N/\varepsilon) + \ell)$ at the cost of increasing the space by a $\log(N/\varepsilon)$ factor. This data structure also constructs a subdivision of \mathbb{R}^d of size $O((N/\varepsilon^d) \log(1/\varepsilon))$, which is an ε -approximate Voronoi diagram of each P_i . Besides being interesting in its own right, the modified data structure will be used in the subsequent subsections.

We begin by describing how we adapt the data structure by Arya and Malamatos [10]. Let S be a set of n points in \mathbb{R}^d , let $\varepsilon > 0$ be a parameter, and let $\mathcal{H} \supset S$ be a hypercube so that any point of S is at least $\text{diam}(S)/\varepsilon$ away from $\partial\mathcal{H}$. Any point of S is an ε -approximate nearest neighbor for a point outside \mathcal{H} . A *quad-tree box* of \mathcal{H} is a hypercube that can be obtained by recursively dividing each side of \mathcal{H} into two equal parts. A useful property of quad-tree boxes (of \mathcal{H}) is that any two of them are either disjoint or one of them is contained in the other. Arya and Malamatos choose a set $\mathcal{B}(S)$ of $O((n/\varepsilon^d) \log(1/\varepsilon))$ quad-tree boxes of \mathcal{H} that cover \mathcal{H} ; each box $\Delta \in \mathcal{B}(S)$ is associated with a point $p_\Delta \in S$. $\mathcal{B}(S)$ has the following crucial property: For a point $x \in \mathcal{H}$, let $\Delta \in \mathcal{B}(S)$ be the smallest box containing x ; then $d(x, p_\Delta) \leq (1 + \varepsilon)d(x, S)$. In order to find the smallest box containing x , they store $\mathcal{B}(S)$ in a BBD-tree, proposed by Arya *et al.* [11]. Instead, we store them in a *compressed quad tree* (see e.g. [20]), as follows. We first construct a quad tree \mathcal{Q} on $\mathcal{B}(S)$. Let H_v be the hypercube associated with the node v of \mathcal{Q} ; the root is associated with \mathcal{H} . A box $\Delta \in \mathcal{B}(S)$ is stored at a node $v \in \mathcal{Q}$ if $H_v = \Delta$. If an interior node $v \in \mathcal{Q}$ does not store a box of \mathcal{B} and if the degree of both v and $p(v)$, the parent of v , is one, we compress v , i.e., we delete v and the child of v becomes the child of $p(v)$. We repeat this step until there is no such node in \mathcal{Q} . The size of \mathcal{Q} is $O(|\mathcal{B}(S)|)$. For a node $v \in \mathcal{Q}$, if a box in $\mathcal{B}(S)$ contains H_v , then we set $\varphi(v) := p_\Delta$ where Δ is

the smallest such box; if there is no such box, then $\varphi(v)$ is undefined. The compressed quadtree \mathcal{Q} can be constructed in time $O(|\mathcal{B}(S)| \log |\mathcal{B}(S)|) = O((n/\varepsilon^d) \log(1/\varepsilon) \log(n/\varepsilon))$ [20].

We call a node $v \in \mathcal{Q}$ *exposed* if its degree is at most one. We associate a region R_v with each exposed node v . If v is a leaf, then $R_v = H_v$, and if v has one child w , then $R_v = H_v \setminus H_w$. For a point $x \in R_v$, v is the lowest node in \mathcal{Q} such that $x \in H_v$. The regions R_v form a partition of \mathcal{H} , and this partition is an ε -approximate Voronoi diagram of S . It can be checked that $\varphi(v)$ is defined for every exposed node v . By construction,

$$d(x, S) \leq d(x, \varphi(v)) \leq (1 + \varepsilon)d(x, S) \quad \forall x \in R_v.$$

The depth of \mathcal{Q} is linear in the worst case. In order to quickly find the node v for which R_v contains a query point x , we construct another tree \mathcal{T} , of depth $O(\log(n/\varepsilon))$, on the nodes of \mathcal{Q} , as follows. We identify in $O(|\mathcal{Q}|)$ time a node ξ in \mathcal{Q} so that the removal of ξ decomposes \mathcal{Q} into at most $u \leq 2^d + 1$ connected components, each of size at most $|\mathcal{Q}|/2$. Let $\mathcal{Q}_0(\xi)$ be the subtree that contains the ancestors of ξ , and let $\mathcal{Q}_1(\xi), \dots, \mathcal{Q}_{u-1}(\xi)$ be the subtrees of \mathcal{Q} rooted at the children of ξ . We recursively construct a tree $\mathcal{T}_i(\xi)$ on each $\mathcal{Q}_i(\xi)$, for $0 \leq i < u$, and attach it as a subtree of ξ . \mathcal{T} can be constructed in time $O(|\mathcal{Q}| \log |\mathcal{Q}|) = O((n/\varepsilon^d) \log(1/\varepsilon) \log(n/\varepsilon))$. Given a point $x \in \mathcal{H}$, we visit a path in \mathcal{T} , starting from its root. Suppose, we are at a node ξ . If ξ is exposed and $x \in R_\xi$, we return $\varphi(\xi)$ and stop. If $x \notin H_\xi$, we recursively visit the subtree $\mathcal{T}_0(\xi)$; otherwise, we visit the child η so that $x \in H_\eta$. The query time is proportional to the depth of \mathcal{T} , which is $O(\log(n/\varepsilon))$.

Returning to the problem of computing ε -approximate nearest neighbors of P_1, \dots, P_ℓ , let C be a hypercube containing $P = \bigcup_{i=1}^\ell P_i$ so that any point of P is at least $\text{diam}(P)/\varepsilon$ away from ∂C . For each $1 \leq i \leq \ell$, we first construct the family $\mathcal{B}(P_i)$ of quad-tree boxes of C , using the algorithm by Arya and Malamatos [10], and then construct the trees \mathcal{Q} and \mathcal{T} on $\bigcup_i \mathcal{B}(P_i)$, as described above. For a node $v \in \mathcal{Q}$, let $\varphi_i(v) = p_{\Delta_i}$, where Δ_i is the smallest box of $\mathcal{B}(P_i)$ that contains H_v ; if there is no such box, then $\varphi_i(v)$ is undefined. By definition, for any point $x \in R_v$ and for $1 \leq i \leq \ell$,

$$d(x, P_i) \leq d(x, \varphi_i(v)) \leq (1 + \varepsilon)d(x, P_i).$$

The regions R_v form a subdivision of C , which is an ε -approximate Voronoi diagram of each P_i .

For a node v , let $\Phi(v)$ be a list of ℓ items whose i th item is $\varphi_i(v)$ if it is defined and NULL otherwise. We store in an implicit manner the set of $\Phi(v)$ for all nodes $v \in \mathcal{Q}$, using a persistent data structure [27] of size $O((N/\varepsilon^d) \log(1/\varepsilon))$ so that $\Phi(v)$ for a node v can be reported in $O(\ell)$ time.

Note that $\varphi_i(v) \neq \varphi_i(p(v))$ if a box $\Delta \in \mathcal{B}(P_i)$ is stored at v , i.e., $H_v \in \mathcal{B}(P_i)$, and $\varphi_i(v) = p_\Delta$ in this case. Let $J(v) = \{i \mid H_v \in \mathcal{B}(P_i)\}$. Then $\Phi(v)$ can be constructed from $\Phi(p(v))$ by updating $\varphi_i(v)$ for $i \in J(v)$. We therefore perform an in-order traversal of \mathcal{Q} and maintain the lists $\Phi(v)$ in a linked list Ψ , using persistence, so that $\Phi(u)$ for the nodes u that have been visited so far can be retrieved in $O(\ell)$ time. If we are currently visiting a node w , then the “current version” of Ψ contains $\Phi(w)$. Moreover, we maintain an array that stores pointers to each item in the current version of Ψ so that the i th item, for any $i \leq \ell$, can be accessed in $O(1)$ time. When we arrive at a node $v \in \mathcal{Q}$ for the first time, for each $i \in J(v)$, we do the following. Let $\Delta \in \mathcal{B}(P_i)$ be a box stored at v . Then we set $\varphi_i(v) = p_\Delta$ and update the i th item of Ψ . We store at v a pointer $\pi(v)$ to the

root of the current version of Ψ . We also store the values of $\varphi_i(p(v))$, for $i \in J(v)$, in a stack at v . After we have processed the subtree rooted at v , the current version of Ψ contains $\Phi(v)$. We delete the values stored in the stack at v and restore them in Ψ so that its current version contains $\Phi(p(v))$. Hence, we perform $2|J(v)|$ updates on Ψ while processing v . Sarnak and Tarjan [27] showed that the amortized time and space of performing an update in a persistent linked list is $O(1)$, therefore the total time spent in this step is $\sum_{v \in \mathcal{Q}} O(|J(v)|) = O((n/\varepsilon^d) \log(1/\varepsilon))$.

Let q be a query point. Following the procedure described above, we first find in $O(\log(N/\varepsilon))$ time the node $\xi \in \mathcal{T}$ such that $q \in R_\xi$. Next, following the pointer $\pi(\xi)$, we access the version of Ψ that stores $\Phi(\xi)$ and report $\varphi_i(\xi)$ for all $1 \leq i \leq \ell$ in $O(\ell)$ time. Hence, we conclude the following.

Theorem 4.1 *Given a family $\{P_1, \dots, P_\ell\}$ of point sets in \mathbb{R}^d , with a total of N points, and a parameter $\varepsilon > 0$, we can compute in $O((N/\varepsilon^d) \log(N/\varepsilon) \log(1/\varepsilon))$ time a subdivision of \mathbb{R}^d and a data structure of size $O((N/\varepsilon^d) \log(1/\varepsilon))$ so that, for any point $q \in \mathbb{R}^d$, one can ε -approximate $d(q, P_i)$, for all $1 \leq i \leq \ell$, in $O(\log(N/\varepsilon) + \ell)$ time.*

4.2 Approximating $\sigma_R(\mathcal{A}, \mathcal{B})$

For $1 \leq i \leq m$, let $P_i = \mathcal{B} - a_i = \{b_j - a_i \mid 1 \leq j \leq n\}$, and for $m < i \leq m+n$, let $P_i = b_{i-m} - \mathcal{A} = \{b_{i-m} - a_j \mid 1 \leq j \leq m\}$. We construct the compressed quad-tree \mathcal{Q} for P_1, \dots, P_{m+n} , with the given parameter ε ; $|\mathcal{Q}| = O((mn/\varepsilon^d) \log(1/\varepsilon))$. Define

$$f_i(t) = d^2(t, P_i) = \begin{cases} \min_{1 \leq j \leq n} d^2(t, b_j - a_i), & 1 \leq i \leq m, \\ \min_{1 \leq j \leq m} d^2(t, b_{i-m} - a_j), & m < i \leq m+n. \end{cases}$$

Let

$$\begin{aligned} F_{\mathcal{A}}(t) &= h_R^2(\mathcal{A} + t, \mathcal{B}) = \frac{1}{m} \sum_{i=1}^m d^2(a_i + t, \mathcal{B}) = \frac{1}{m} \sum_{i=1}^m f_i(t) \\ F_{\mathcal{B}}(t) &= h_R^2(\mathcal{B}, \mathcal{A} + t) = \frac{1}{n} \sum_{i=1}^n d^2(\mathcal{A} + t, b_i) = \frac{1}{n} \sum_{i=m+1}^{m+n} f_i(t). \end{aligned}$$

For each exposed node $v \in \mathcal{Q}$, $\varphi_i(v)$ is defined for all $1 \leq i \leq m+n$. Let

$$\widehat{F}_{\mathcal{A},v}(t) = \frac{1}{m} \sum_{i=1}^m d^2(t, \varphi_i(v)) \quad \text{and} \quad \widehat{F}_{\mathcal{B},v}(t) = \frac{1}{n} \sum_{i=1}^n d^2(t, \varphi_{m+i}(v)).$$

By construction, for any $t \in R_v$,

$$F_{\mathcal{A}}(t) \leq \widehat{F}_{\mathcal{A},v}(t) = \sum_{i=1}^m d^2(t, \varphi_i(v)) \leq \sum_{i=1}^m (1 + \varepsilon)^2 \cdot d^2(t, P_i) \leq (1 + \varepsilon)^2 F_{\mathcal{A}}(t),$$

implying that

$$\sqrt{\widehat{F}_{\mathcal{A},v}(t)} \leq (1 + \varepsilon) h_R(\mathcal{A} + t, \mathcal{B}).$$

Similarly, $\sqrt{\widehat{F}_{\mathcal{B},v}(t)} \leq (1 + \varepsilon)h_R(\mathcal{B}, \mathcal{A} + t)$. Hence, it suffices to store $\widehat{F}_{\mathcal{A},v}(t), \widehat{F}_{\mathcal{B},v}(t)$ at each exposed node $v \in \Omega$. Since they are quadratic functions in $t \in \mathbb{R}^d$, they can be stored using $O(1)$ space (where the constant depends on d) and updated in $O(1)$ time for each change in $\varphi_i(v)$.

If we compute $\widehat{F}_{\mathcal{A},v}$ for each exposed node $v \in \Omega$ independently, the total time spent is $O((m^2n/\varepsilon^d) \log(1/\varepsilon))$. We therefore proceed as in the previous section. We perform an in-order traversal of Ω . For each node $v \in \Omega$, let $J^*(v) = \{i \mid \varphi_i(v) \text{ is defined}\}$, and let

$$\widehat{F}_{\mathcal{A},v}(t) = \frac{1}{m} \sum_{i \in J^*(v), i \leq m} d^2(t, \varphi_i(v)) \quad \text{and} \quad \widehat{F}_{\mathcal{B},v}(t) = \frac{1}{n} \sum_{i \in J^*(v), i > m} d^2(t, \varphi_i(v)).$$

If v is an exposed node then $J^*(v) = [1..m + n]$, therefore the above definition of $\widehat{F}_{\mathcal{A},v}, \widehat{F}_{\mathcal{B},v}$ is consistent with that for exposed nodes defined earlier. Following the same idea as in the previous subsection,

$$\begin{aligned} \widehat{F}_{\mathcal{A},v}(t) &= \widehat{F}_{\mathcal{A},p(v)}(t) + \frac{1}{m} \sum_{i \in J^*(v), i \leq m} [d^2(t, \varphi_i(v)) - d^2(t, \varphi_i(p(v)))], \\ \widehat{F}_{\mathcal{B},v}(t) &= \widehat{F}_{\mathcal{B},p(v)}(t) + \frac{1}{n} \sum_{i \in J^*(v), i > m} [d^2(t, \varphi_i(v)) - d^2(t, \varphi_i(p(v)))]. \end{aligned}$$

Hence, $\widehat{F}_{\mathcal{A},v}, \widehat{F}_{\mathcal{B},v}$ can be computed in $O(|J^*(v)|)$ time and the total time spent in the traversal is $O((n/\varepsilon^d) \log(1/\varepsilon))$. Finally, for each exposed node $v \in \Omega$, we compute

$$t_v = \arg \min_{t \in \mathbb{R}^d} \max \left\{ \sqrt{\widehat{F}_{\mathcal{A},v}(t)}, \sqrt{\widehat{F}_{\mathcal{B},v}(t)} \right\}$$

and return

$$\min_v H_R(\mathcal{A} + t_v, \mathcal{B}) \leq (1 + \varepsilon)\sigma_R(\mathcal{A}, \mathcal{B})$$

where the minimum is taken over all exposed nodes v of Ω .

If we wish to compute an ε -approximate value of $H_R(\mathcal{A} + t, \mathcal{B})$, for a given $t \in \mathbb{R}^d$, we construct the tree \mathcal{T} as described in the previous section. Instead of storing the lists $\Phi(\cdot)$ at each node, we now store $\widehat{F}_{\mathcal{A},v}, \widehat{F}_{\mathcal{B},v}$ at each node $v \in \mathcal{T}$. The total storage needed is $O((mn/\varepsilon^d) \log(1/\varepsilon))$. Hence, we obtain the following.

Theorem 4.2 *Given two sets \mathcal{A} and \mathcal{B} of m and n points in \mathbb{R}^d and a parameter $\varepsilon > 0$, we can:*

- i. *compute a vector $t^* \in \mathbb{R}^d$ in $O((mn/\varepsilon^d) \log(mn/\varepsilon) \log(1/\varepsilon))$ time, so that*

$$H_R(\mathcal{A} + t^*, \mathcal{B}) \leq (1 + \varepsilon)\sigma_R(\mathcal{A}, \mathcal{B});$$

- ii. *construct a data structure of size $O((mn/\varepsilon^d) \log(1/\varepsilon))$, in time $O((mn/\varepsilon^d) \log(mn/\varepsilon) \log(1/\varepsilon))$, so that for any query vector $t \in \mathbb{R}^d$, we can compute an ε -approximate value of $H_R(\mathcal{A} + t, \mathcal{B})$ in $O(\log(mn/\varepsilon))$ time.*

4.3 Approximating $\sigma_S(\mathcal{A}, \mathcal{B})$

Modifying the above scheme, we approximate $\sigma_S(\mathcal{A}, \mathcal{B})$ as follows. Let P_i and \mathcal{Q} be the same as defined above. We define

$$\begin{aligned} G_{\mathcal{A}}(t) &= \frac{1}{m} \sum_{i=1}^m d(t, P_i) = h_S(\mathcal{A} + t, \mathcal{B}), \\ G_{\mathcal{B}}(t) &= \frac{1}{n} \sum_{j=1}^n d(t, P_{m+j}) = h_S(\mathcal{B}, \mathcal{A} + t). \end{aligned}$$

For each exposed node $v \in \mathcal{Q}$, let

$$\begin{aligned} \widehat{G}_{\mathcal{A},v}(t) &= \frac{1}{m} \sum_{i=1}^m d(t, \varphi_i(v)) \leq (1 + \varepsilon) h_S(\mathcal{A} + t, \mathcal{B}) \\ \widehat{G}_{\mathcal{B},v}(t) &= \frac{1}{n} \sum_{j=1}^n d(t, \varphi_{m+j}(v)) \leq (1 + \varepsilon) h_S(\mathcal{B}, \mathcal{A} + t) \\ t_v &= \arg \min_{t \in R_v} \max \left\{ \widehat{G}_{\mathcal{A},v}(t), \widehat{G}_{\mathcal{B},v}(t) \right\}. \end{aligned}$$

Since $\widehat{G}_{\mathcal{A},v}$ and $\widehat{G}_{\mathcal{B},v}$ are not simple algebraic functions, we do not know how to compute, store, and update them efficiently. Nevertheless, we can compute an ε -approximation for $\widehat{G}_{\mathcal{A},v}$ (resp., $\widehat{G}_{\mathcal{B},v}$) that is easier to handle. More precisely, for a given set P of points in \mathbb{R}^d , define the 1-median function

$$\text{med}_P(t) = \frac{1}{|P|} \sum_{p \in P} d(p, t).$$

For any $v \in \mathcal{Q}$, let $\Phi_{\mathcal{A}}(v) = \{\varphi_i(v) \mid i \leq m, i \in J^*(v)\}$, $\Phi_{\mathcal{B}}(v) = \{\varphi_i(v) \mid i > m, i \in J^*(v)\}$, $\widehat{G}_{\mathcal{A},v}(t) = \text{med}_{\Phi_{\mathcal{A}}(v)}(t)$, and $\widehat{G}_{\mathcal{B},v}(t) = \text{med}_{\Phi_{\mathcal{B}}(v)}(t)$. In Section 4.4, we describe a dynamic data structure that, given a point set P of size n , maintains an ε -approximation of the function $\text{med}_P(\cdot)$ as a function $\widehat{\text{med}}_P(\cdot)$ defined by set of $O((1/\varepsilon^d) \log(1/\varepsilon))$ weighted points. A point can be inserted into or deleted from P in $O((1/\varepsilon^d) \log((\log n)/\varepsilon) \log^{d+1} n)$ time. Furthermore, given two point sets P and Q in \mathbb{R}^d , this can be used to maintain an ε -approximation of $\arg \min_t \max\{\text{med}_P(t), \text{med}_Q(t)\}$ within the same time bound.

Using this data structure, we can traverse all cells of \mathcal{Q} , as in Section 4.2, and compute an $(\varepsilon/3)$ -approximation of $\widehat{G}_{\mathcal{A},v}$ and $\widehat{G}_{\mathcal{B},v}$ (thus an ε -approximation of $G_{\mathcal{A}}$ and $G_{\mathcal{B}}$) for each node $v \in \mathcal{Q}$. However, we now spend

$$O(|J^*(v)| \cdot (1/\varepsilon^d) \text{polylog}(mn, 1/\varepsilon))$$

time to compute an $(\varepsilon/3)$ -approximation of $\widehat{G}_{\mathcal{A},v}$ from that of $\widehat{G}_{\mathcal{A},p(v)}$. Putting everything together, we conclude the following.

Theorem 4.3 *Given two sets \mathcal{A} and \mathcal{B} of m and n points in \mathbb{R}^d and a parameter $0 < \varepsilon \leq 1$, we can compute:*

i. a vector $t^* \in \mathbb{R}^d$, in $O((mn/\varepsilon^{2d}) \text{polylog}(mn, 1/\varepsilon))$ time, so that

$$H_S(\mathcal{A} + t^*, \mathcal{B}) \leq (1 + \varepsilon)\sigma_S(\mathcal{A}, \mathcal{B});$$

ii. a data structure of $O((mn/\varepsilon^{2d}) \text{polylog}(mn, 1/\varepsilon))$ size in time $O((mn/\varepsilon^{2d}) \text{polylog}(mn, 1/\varepsilon))$, so that for any query vector $t \in \mathbb{R}^d$, we can ε -approximate $H_S(\mathcal{A}+t, \mathcal{B})$ in time $O(\text{polylog}(mn, 1/\varepsilon))$.

4.4 Maintaining the 1-median function

Let P be a set of n points in \mathbb{R}^2 . In this subsection, we describe an algorithm for maintaining an ε -approximation of the 1-median of a point set P as points are inserted into or deleted from P . Using the ideas in [1, 21], we construct a *coreset*, a weighted subset of P , whose 1-median approximates that of P , and argue that it can be updated efficiently as the set P changes.

Let (P, w) be a *weighted* point set in \mathbb{R}^d with the weight function $w : P \rightarrow \mathbb{R}^+$. For a subset $A \subseteq P$, let $w(A) = \sum_{p \in A} w(p)$. If the weight function is not important or obvious from the context, we will use P to denote (P, w) . For a point $x \in \mathbb{R}^d$, we define $\nu((P, w), x) = \sum_{p \in P} w(p) \|px\|$ as the *price* of the 1-median placed at x . Furthermore, let $\nu_{\text{opt}}((P, w)) = \min_{x \in \mathbb{R}^d} \nu((P, w), x)$ denote the price of the *optimal* 1-median for (P, w) .

Definition 4.4 (Coreset) Let (P, w) be weighted point set in \mathbb{R}^d . A weighted set (\mathcal{S}, χ) with $\mathcal{S} \subseteq P$ is an ε -coreset of (P, w) for 1-median if

$$(1 - \varepsilon)\nu((P, w), x) \leq \nu((\mathcal{S}, \chi), x) \leq (1 + \varepsilon)\nu((P, w), x) \quad \forall x \in \mathbb{R}^d. \quad (1)$$

The following lemma describes an algorithm for computing a small ε -coreset for 1-median.

Lemma 4.5 Let (P, w) be a weighted set of n points in \mathbb{R}^d , and let $0 < \varepsilon \leq 1/2$ be a parameter. An ε -coreset (\mathcal{S}, χ) for (P, w) of size $O((1/\varepsilon^d) \log(1/\varepsilon))$ for 1-median can be computed in time $O(n \log(n/\varepsilon) + (1/\varepsilon^d) \log(1/\varepsilon))$.

Proof: Let $\mathbf{0} \in \mathbb{R}^d$ be the point realizing $\nu_{\text{opt}}(P)$. Set $r = \nu_{\text{opt}}(P)/w(P)$, $M = \lceil \log_2(c_1 d/\varepsilon^2) \rceil$, where $c_1 \geq 40$ is a constant, and $\ell = 2^M r = c_1 d r / \varepsilon^2$. Let $\mathcal{C}(x, \rho)$ be the axis-parallel hypercube of side length ρ centered at x . Let $B_0 = \mathcal{C}(\mathbf{0}, r)$, and let $B_i = \mathcal{C}(\mathbf{0}, 2^i r) \setminus \mathcal{C}(\mathbf{0}, 2^{i-1} r)$ for $1 \leq i \leq M$. Next, we partition each B_i into $O(1/\varepsilon^d)$ hypercubes of side length $\varepsilon 2^i r / (c_2 d)$ by drawing a uniform grid, where c_2 is a sufficiently large constant. This forms an exponential grid G that covers the hypercube $\mathcal{H} = \mathcal{C}(\mathbf{0}, \ell)$; see Figure 3. Set $P_I = P \cap \mathcal{H}$ and $P_O = P \setminus P_I$.

For every grid cell $\square \in G$, we pick a “representative” point $p \in P \cap \square$, add it to \mathcal{S} and set its weight $\chi(p) = w(P \cap \square)$. Finally, let $u \in P$ be the point farthest away from $\mathbf{0}$ in P . We add u to \mathcal{S} and set its weight $\chi(u) = \nu(P_O, \mathbf{0}) / \|u\mathbf{0}\|$. By construction, $|\mathcal{S}| = O((1/\varepsilon^d) \log(1/\varepsilon))$. We claim that (\mathcal{S}, χ) is a ε -coreset of (P, w) for 1-median, i.e., it satisfies (1).

For any $x \in \mathbb{R}^d$, we claim that

$$|\nu(P_I, x) - \nu(\mathcal{S}_I, x)| \leq (\varepsilon/4)\nu(P, \mathbf{0}) = (\varepsilon/4)\nu_{\text{opt}}(P). \quad (2)$$

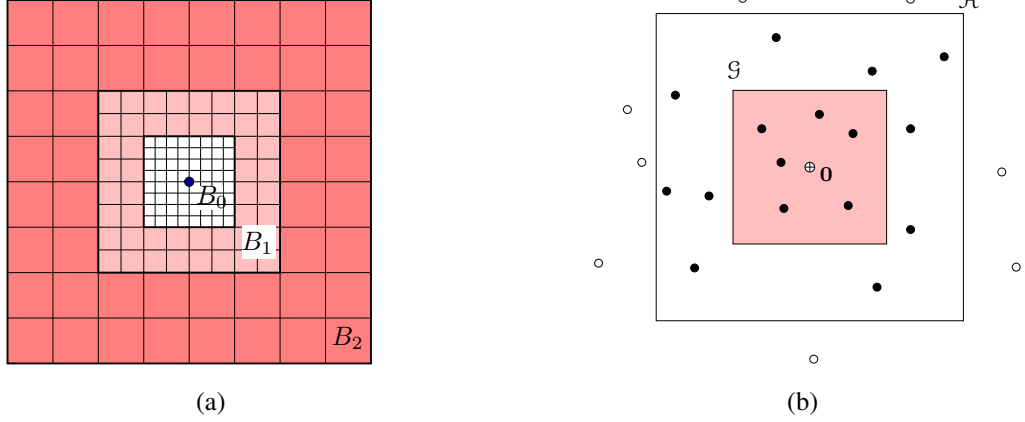


Figure 3: (a) An exponential grid with 3 layers. (b) The larger (resp., smaller) box is \mathcal{H} (resp., \mathcal{G}), and the set of hollow circles is P_O .

Indeed, every point in a cell of $P \cap \mathcal{H}$ can be interpreted as “traveling” to its representative point in that cell. The total (weighted) distance traveled by all the points inside \mathcal{H} to the coresets is smaller than $(\varepsilon/4)\nu_{\text{opt}}(P)$. As such, the error contributed by the points of P_I is smaller than $(\varepsilon/4)\nu_{\text{opt}}(P)$. This implies that we only need to consider the error contributed by points in P_O .

In particular, we need to bound the difference in the price of $\nu(P_O, x)$ and $\nu(\mathcal{S}_O, x)$. (Of course, if $P_O = \emptyset$, the lemma trivially holds and the following is unnecessary.) Let $\mathcal{S}_I = \mathcal{S} \cap \mathcal{H}$ and $\mathcal{S}_O = \{u\}$. Let $\mathcal{G} = \mathcal{C}(\mathbf{0}, c_1 r/4\varepsilon) \subseteq \mathcal{H}$; see Figure 3. We consider two cases.

- (i) $x \in \mathcal{G}$: By the triangle inequality, for any $x, y \in \mathbb{R}^d$,

$$\|y\mathbf{0}\| - \|x\mathbf{0}\| \leq \|yx\| \leq \|y\mathbf{0}\| + \|x\mathbf{0}\|.$$

Therefore,

$$\nu(P_O, \mathbf{0}) - \omega(P_O) \|x\mathbf{0}\| \leq \nu(P_O, x) \leq \nu(P_O, \mathbf{0}) + \omega(P_O) \|x\mathbf{0}\|. \quad (3)$$

Since $\ell \geq c_1 d r / \varepsilon^2$ and the side length of \mathcal{G} is at most $c_1 r / 4\varepsilon$, $\|px\| \geq 4 \|x\mathbf{0}\| / \varepsilon$, for all $p \in P_O$, and thus (3) becomes

$$(1 - \varepsilon/4)\nu(P_O, \mathbf{0}) \leq \nu(P_O, x) \leq (1 + \varepsilon/4)\nu(P_O, \mathbf{0}). \quad (4)$$

Similarly, we have

$$(1 - \varepsilon/4)\nu(P_O, \mathbf{0}) \leq \nu(\mathcal{S}_O, x) \leq (1 + \varepsilon/4)\nu(P_O, \mathbf{0}). \quad (5)$$

Putting (4) and (5) together, we have

$$|\nu(P_O, x) - \nu(\mathcal{S}_O, x)| \leq (\varepsilon/2)\nu(P_O, \mathbf{0}). \quad (6)$$

Using (2) and (6), we conclude

$$\begin{aligned}
|\nu(P, x) - \nu(\mathcal{S}, x)| &\leq |\nu(P_I, x) - \nu(\mathcal{S}_I, x)| + |\nu(P_O, x) - \nu(\mathcal{S}_O, x)| \\
&\leq (\varepsilon/4)\nu(P, x) + (\varepsilon/2)\nu(P_O, \mathbf{0}) \\
&\leq \varepsilon\nu(P, x).
\end{aligned}$$

(ii) $x \in \mathbb{R}^d \setminus \mathcal{G}$: We first claim that

$$\omega(P_O) \leq \varepsilon^2\omega(P)/c_1. \quad (7)$$

Indeed,

$$\nu_{\text{opt}}(P) = \sum_{p \in P} \omega(p) \|p\mathbf{0}\| \geq \sum_{p \in P_O} \omega(p)\ell \geq \omega(P_O) \cdot \frac{c_1 d \nu_{\text{opt}}(P)}{\varepsilon^2 \omega(P)},$$

which implies the claim. Next,

$$\omega(P) \|x\mathbf{0}\| \geq \omega(P) \frac{c_1 r}{4\varepsilon} \geq c_1 \nu_{\text{opt}}(P)/4\varepsilon \quad (8)$$

Using (7) and (8), we obtain

$$\begin{aligned}
\nu(P_I, x) &\geq \omega(P_I) \|x\mathbf{0}\| - \nu(P_I, \mathbf{0}) \\
&\geq (1 - \varepsilon^2/c_1) \omega(P) \|x\mathbf{0}\| - \nu_{\text{opt}}(P) \\
&\geq (1 - \varepsilon^2/c_1) \omega(P) \|x\mathbf{0}\| - (4\varepsilon/c_1) \omega(P) \|x\mathbf{0}\| \\
&\geq (1 - \varepsilon/8) \omega(P) \|x\mathbf{0}\|.
\end{aligned} \quad (9)$$

The last inequality follows since $c_1 \geq 40$. A similar argument shows that

$$\begin{aligned}
\nu(P_O, x) &\leq \omega(P_O) \|x\mathbf{0}\| + \nu(P_O, \mathbf{0}) \\
&\leq \frac{\varepsilon^2}{c_1} \omega(P) \|x\mathbf{0}\| + \nu_{\text{opt}}(P) \\
&\leq \frac{\varepsilon}{8} \omega(P) \|x\mathbf{0}\| \quad (\text{using (8)}).
\end{aligned} \quad (10)$$

Plugging (9) into (10), we obtain

$$\nu(P_O, x) \leq \frac{\varepsilon}{8} \frac{\nu(P_I, x)}{1 - \varepsilon/8} \leq \frac{\varepsilon}{4} \nu(P, x).$$

Similarly, we can argue that

$$\nu(\mathcal{S}_O, x) \leq (\varepsilon/4)\nu(P, x). \quad (11)$$

Thus, using (2), (10), (11), we obtain

$$\begin{aligned}
|\nu(P, x) - \nu(\mathcal{S}, x)| &\leq |\nu(P_I, x) - \nu(\mathcal{S}_I, x)| + |\nu(P_O, x) - \nu(\mathcal{S}_O, x)| \\
&\leq (\varepsilon/4)\nu(P, x) + (\varepsilon/2)\nu(P, x) \\
&\leq \varepsilon\nu(P, x).
\end{aligned}$$

This completes the proof that (\mathcal{S}, χ) is an ε -coreset of (P, ω) . The above argument works even if $\nu(P, \mathbf{0}) \leq c\nu_{\text{opt}}(P)$ for some constant $c \geq 1$, provided that c_1, c_2 are chosen appropriately.

In order to compute \mathcal{S} , we first compute in $O(n)$ time the centroid $\bar{c} \in \mathbb{R}^2$. It is well known that $\nu(P, \bar{c}) \leq 2\nu_{\text{opt}}(P)$. We then compute the B_i 's and the exponential grid in time $O((1/\varepsilon^d) \log(1/\varepsilon))$ and find the grid cell that contains each point of P in a total time of $O(n \log(1/\varepsilon))$. Hence, the total time spent in computing \mathcal{S} is $O(n \log(n/\varepsilon) + (1/\varepsilon^d) \log(1/\varepsilon))$. \blacksquare

Let (P_1, ω) and (P_2, ω) be two weighted point sets so that $P_1 \cap P_2 = \emptyset$, and let (\mathcal{S}_i, χ) be an ε -coreset of (P_i, ω) for 1-median. Then $(\mathcal{S}_1 \cup \mathcal{S}_2, \chi)$ is a ε -coreset of $(P_1 \cup P_2, \omega)$. Moreover, if (\mathcal{S}_2, χ_2) is an ε_2 -coreset of (\mathcal{S}_1, χ_1) , and (\mathcal{S}_1, χ_1) is an ε_1 -coreset of (P, ω) , then (\mathcal{S}_2, χ_2) is a $2(\varepsilon_1 + \varepsilon_2)$ -coreset of (P, ω) . Using these observations and plugging Lemma 4.5 into the dynamic data structure by Agarwal et al. [1], one can maintain an ε -coreset of (P, ω) of size $O((1/\varepsilon^d) \log(1/\varepsilon))$ for 1-median of (P, ω) efficiently. Omitting all the details, we conclude the following.

Theorem 4.6 *Let P be a set of n points in \mathbb{R}^d , and let $\varepsilon > 0$ be a parameter. One can maintain a ε -coreset of P of size $O((1/\varepsilon^d) \log(1/\varepsilon))$ under insertions and deletions of points in P so that each update takes $O((1/\varepsilon^d) \log(\log(n)/\varepsilon) \log^{d+1} n)$ time.*

4.5 A randomized algorithm

We briefly describe below a simple randomized algorithm to approximate $\sigma_R(\mathcal{A}, \mathcal{B})$. The algorithm for approximating $\sigma_S(\mathcal{A}, \mathcal{B})$ is similar. Let t^* be the optimal translation, i.e., $H_R(\mathcal{A} + t^*, \mathcal{B}) = \sigma_R(\mathcal{A}, \mathcal{B})$.

Lemma 4.7 *For a random point a_k from \mathcal{A} , $d(a_k + t^*, \mathcal{B}) \leq 2\sigma_R(\mathcal{A}, \mathcal{B})$, with probability greater than $1/2$. The same claim holds for $\sigma_S(\mathcal{A}, \mathcal{B})$.*

Proof: Let a_k be a random point from \mathcal{A} , where each point of \mathcal{A} is chosen with equal probability. Let Y be the random variable $Y = d(a_k + t^*, \mathcal{B})$. Then

$$E[Y] = \frac{1}{m} \sum_{i=1}^m d(a_i + t^*, \mathcal{B}) = H_R(\mathcal{A} + t^*, \mathcal{B}) = \sigma_R(\mathcal{A}, \mathcal{B}).$$

The lemma now follows immediately from Markov's inequality. \blacksquare

Choose a random point $a_k \in \mathcal{A}$. Let $t_j = b_j - a_k$ and $\delta_j = H_R(\mathcal{A} + t_j, \mathcal{B})$, for $1 \leq j \leq n$. It then follows from Lemma 4.7 and the same argument as in Lemma 3.7, that $\min_j \delta_j$ is a constant-factor approximation of $\sigma_R(\mathcal{A}, \mathcal{B})$, with probability greater than $1/2$. Computing δ_j exactly is expensive in \mathbb{R}^d , therefore we compute an approximate value of δ_j , for $1 \leq j \leq n$, in time $O((m+n) \log mn)$, by performing approximate nearest-neighbor queries [10]. We can improve this constant-factor approximation algorithm to compute a $(1+\varepsilon)$ -approximation of $\sigma_R(\mathcal{A}, \mathcal{B})$ using the same technique as in Section 3. We thus obtain the following result.

Theorem 4.8 *Given two sets \mathcal{A} and \mathcal{B} of m and n points, respectively, in \mathbb{R}^d , and a parameter $\varepsilon > 0$, we can compute, in randomized expected time $O((mn/\varepsilon^d) \log mn)$, two translation vectors*

t_1 and t_2 , such that, with probability greater than $1/2$,

$$H_R(\mathcal{A} + t_1, \mathcal{B}) \leq (1 + \varepsilon)\sigma_R(\mathcal{A}, \mathcal{B}) \text{ and } H_S(\mathcal{A} + t_2, \mathcal{B}) \leq (1 + \varepsilon)\sigma_S(\mathcal{A}, \mathcal{B}).$$

5 Conclusions

We provide in this paper some initial study of various problems related to minimizing Hausdorff distance between sets of points, disks, and balls. One natural question following our study is to compute exactly or approximately the smallest Hausdorff distance over all possible rigid motions in \mathbb{R}^2 and \mathbb{R}^3 . Given two sets of points \mathcal{A} and \mathcal{B} of size n and m , respectively, let Δ be the maximum of the diameters of \mathcal{A} and \mathcal{B} . We believe that there is a randomized algorithm with roughly $mn\sqrt{\Delta}$ expected time, that approximates the optimal summed-Hausdorff distance (or rms-Hausdorff distance) under rigid motions in the plane. The algorithm that we envisage combines our randomized approach from Section 4.5, a framework to convert the original problem to a pattern matching problem [24], and a result by Amir *et al.* on string matching [9]. However, this approach does not extend to families of balls. We leave the problem of computing the smallest Hausdorff distance between sets of points or balls under rigid motions as an open question for further research. Another question is to approximate efficiently the best Hausdorff distance under certain transformations when partial matching is allowed. The traditional approaches using reference points break down with partial matching.

References

- [1] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan, Approximating extent measures of points, *J. ACM*, 51 (2004), 606–635.
- [2] P. K. Agarwal and J. Matoušek, Ray shooting and parametric search, *SIAM J. Comput.*, 22 (1993), 540–570.
- [3] P. K. Agarwal, M. Sharir, and S. Toledo, Applications of parametric searching in geometric optimization, *J. Algorithms*, 17 (1994), 292–318.
- [4] O. Aichholzer, H. Alt, and G. Rote, Matching shapes with a reference point, *Intl. J. Comput. Geom. Appl.*, 7 (1997), 349–363.
- [5] H. Alt, B. Behrends, and J. Blömer, Approximate matching of polygonal shapes, *Ann. Math. Artif. Intell.*, 13 (1995), 251–266.
- [6] H. Alt, P. Brass, M. Godau, C. Knauer, and C. Wenk, Computing the Hausdorff distance of geometric patterns and shapes, in: *Discrete and Computational Geometry — The Goodman-Pollack Festschrift* (B. Aronov, S. Basu, J. Pach, and M. Sharir, eds.), Springer-Verlag, Heidelberg, 2003, pp. 65–76.
- [7] H. Alt and L. J. Guibas, Discrete geometric shapes: Matching, interpolation, and approximation, in: *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, eds.), Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000, pp. 121–153.

- [8] N. Amenta and R. Kolluri, The medial axis of a union of balls, *Comput. Geom. Theory Appl.*, 20 (2001), 25–37.
- [9] A. Amir, E. Porat, and M. Lewenstein, Approximate subset matching with Don't Cares, *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, 2001, pp. 305–306.
- [10] S. Arya and T. Malamatos, Linear-size approximate Voronoi diagrams, *Proc. 13th ACM-SIAM Sympos. on Discrete Algorithms*, 2002, pp. 147–155.
- [11] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu, An optimal algorithm for approximate nearest neighbor searching fixed dimensions, *J. ACM*, 45 (1998), 891–923.
- [12] M. J. Atallah, A linear time algorithm for the Hausdorff distance between convex polygons, *Inform. Process. Lett.*, 17 (1983), 207–209.
- [13] P. Bose, A. Maheshwari, and P. Morin, Fast approximations for sums of distances, clustering and the fermit–weber problem, *Comput. Geom. Theory Appl.*, 24 (2003), 135–146.
- [14] D. Cardoze and L. Schulman, Pattern matching for spatial point sets, *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.*, 1998, pp. 156–165.
- [15] L. P. Chew, D. Dor, A. Efrat, and K. Kedem, Geometric pattern matching in d -dimensional space, *Discrete Comput. Geom.*, 21 (1999), 257–274.
- [16] L. P. Chew, M. T. Goodrich, D. P. Huttenlocher, K. Kedem, J. M. Kleinberg, and D. Kravets, Geometric pattern matching under euclidean motion, *Comput. Geom. Theory Appl.*, 7 (1997), 113–124.
- [17] M. T. Goodrich, J. S. B. Mitchell, and M. W. Orletsky, Practical methods for approximate geometric pattern matching under rigid motion, *Proc. 10th Annu. Sympos. Comput. Geom.*, 1994, pp. 103–112.
- [18] I. Halperin, B. Ma, H. Wolfson, and R. Nussinov, Principles of docking: An overview of search algorithms and a guide to scoring functions, *Proteins: Structure, Function, and Genetics*, 47 (2002), 409–443.
- [19] S. Har-Peled, A replacement for Voronoi diagrams of near linear size, *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci.*, 2001, pp. 94–103.
- [20] S. Har-Peled, Class notes on compressed quadtrees,
http://valis.cs.uiuc.edu/~sariel/teach/2004/a_aprx/lec/03_quadtree.pdf, 2005.
- [21] S. Har-Peled and S. Mazumdar, Coresets for k -means and k -median clustering and their applications, *Proc. 36th ACM Sympos. Theory Comput.*, 2004, pp. 291–300.
- [22] D. P. Huttenlocher, K. Kedem, and J. M. Kleinberg, On dynamic Voronoi diagrams and the minimum Hausdorff distance for point sets under Euclidean motion in the plane, *Proc. 8th Annu. Sympos. Comput. Geom.*, 1992, pp. 110–120.
- [23] D. P. Huttenlocher, K. Kedem, and M. Sharir, The upper envelope of Voronoi surfaces and its applications, *Discrete Comput. Geom.*, 9 (1993), 267–291.
- [24] P. Indyk, R. Motwani, and S. Venkatasubramanian, Geometric matching under noise: Combinatorial bounds and algorithms, *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, 1999, pp. 457–465.

- [25] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.*, 1 (1986), 59–71.
- [26] N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, *J. ACM*, 30 (1983), 852–865.
- [27] N. Sarnak and R. E. Tarjan, Planar point location using persistent search trees, *Commun. ACM*, 29 (1986), 669–679.
- [28] S. Seeger and X. Labourey, Feature extraction and registration: An overview, *Principles of 3D Image Analysis and Synthesis*, Kluwer Academic Publishers, 2002, pp. 153–166.
- [29] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.