Maximizing the Area of Overlap of two Unions of Disks under Rigid Motion

Mark de Berg [a], Sergio Cabello [b,1], Panos Giannopoulos [*,b,1,2], Christian Knauer [c], René van Oostrum [b,1] and Remco C. Veltkamp [b]

[a] *Faculty of Mathematics and Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, the Netherlands.*
[b] *Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, the Netherlands.*
[c] *Institut für Informatik, Freie Universität Berlin, Takusstraße 9, D-14195 Berlin, Germany.*

Abstract

Let A and B be two sets of n resp. m $(m \geq n)$ disjoint unit disks in the plane. We consider the problem of finding a rigid motion of A that maximizes the total area of its overlap with B. The function describing the area of overlap is quite complex, even for combinatorially equivalent translations, and hence, we turn our attention to approximation algorithms. We give deterministic $(1-\epsilon)$-approximation algorithms for the maximum area of overlap under translation and rigid motion that run in $O((nm/\epsilon^2)\log(m/\epsilon))$ and $O((n^2m^2/\epsilon^3)\log m))$ time respectively. For the later, if Δ is the diameter of set A, we get an $(1-\epsilon)$-approximation in $O(\frac{m^2n^{4/3}\Delta^{1/3}\log n \log m}{\epsilon^3})$ time which is improvement when $\Delta = o(n^2)$. Under the condition that the maximum is at least a constant fraction of the area of B, we give a probabilistic $(1-\epsilon)$-approximation algorithm for rigid motions that runs in $O((cn^2/\epsilon^5)\log^2(n/\epsilon))$ time and succeeds with probability at least $1 - e^{-c} - n^{-6}$.

Key words: Geometric Optimization, Approximation Algorithms, Shape Matching, Area of Overlap, Unions of Disks, Rigid Motions

1. Introduction

Shape matching is a fundamental problem in computer vision: given two shapes A and B and a distance measure, one wants to determine a transformation of A - a translation, for instance, or a rigid motion - that minimizes its distance to B. Typical problems include: matching point sets with respect to the Hausdorff distance and matching polygons with respect to the Hausdorff or Fréchet distance between their boundaries; see Alt et al. [1] for a survey. The *area of overlap* (or the *area the symmetric difference*) of two polygons is less sensitive to noise and therefore more appropriate for certain applications. The function of the area of overlap of two translated simple polygons was first studied by Mount et al. [5]. They showed that this function is a continuous, $O((nm)^2)$ piecewise polynomial surface of at most degree two with a representation that can be computed in $O((nm)^2)$ time for two polygons with n and m vertices; efficient algorithms for determining the optimal translation for the case of *convex* polygons and constant-factor approximation of the minimum area of symmetric difference of two convex shapes under translation can be also found in the survey by Alt et al. [1].

Recently, de Berg et al. [3] examined the problem of matching unions of convex homothets or unions of fat objects and gave a nearly quadratic deterministic $(1-\epsilon)$-approximation algorithm for the maximum area of overlap under translations;

* Panos Giannopoulos
Email addresses: `m.t.d.berg@TUE.nl` (Mark de Berg), `sergio@cs.uu.nl` (Sergio Cabello), `panos@cs.uu.nl` (Panos Giannopoulos), `Christian.Knauer@inf.fu-berlin.de` (Christian Knauer), `rene@cs.uu.nl` (René van Oostrum), `Remco.Veltkamp@cs.uu.nl` (Remco C. Veltkamp).

[1] This research was partialy supported by the STW, project UIF 5055 and the Cornelis Lely Stichting
[2] The author was partially supported by a Marie Curie Fellowship of the EC programme "Combinatorics, Geometry, and Computation, FU Berlin, HPMT-CT-2001-00282.

we mention part of these results in this abstract as well. Their motivation comes from applications in matching shapes that are 'expressed' as unions of convex objects and weighted point set matching. Cheong et al. [2] gave a nearly linear time probabilistic approximation algorithm for the same problem.

Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$, $(n \leq m)$ be two sets of disjoint unit disks in the plane. We consider B to be fixed, while A can be translated and rotated relative to B. Let $\mathcal{I}$ be the infinite set of all possible isometries in the plane; we call $\mathcal{I}$ the *configuration space.* We denote by R_θ a rotation about the origin by some angle $\theta \in [0, 2\pi)$ and by $T_{\vec{t}}$ a translation by some $\vec{t} \in \mathbb{R}^2$. Rotated only versions of A are denoted by $A(\theta) = \{A_i(\theta), \ldots, A_n(\theta)\}$; translated only versions of A are denoted by $A(\vec{t}) = \{A_i(\vec{t}), \ldots, A_n(\vec{t})\}$. Any isometry $I \in \mathcal{I}$ can be uniquely defined as $I = I_{\vec{t},\theta} = R_\theta \circ T_{\vec{t}}$ or $I = I_{\vec{t}',\theta'} = T_{\vec{t}'} \circ R_{\theta'}$, for some $\theta, \theta' \in [0, 2\pi)$ and $\vec{t}, \vec{t}' \in \mathbb{R}^2$; transformed versions of A are denoted by $A(\vec{t}, \theta) = \{A_i(\vec{t}, \theta), \ldots, A_n(\vec{t}, \theta)\}$ for some $I_{\vec{t},\theta} \in \mathcal{I}$. Let $r_i(\vec{t})$ be the distance of $A_i(\vec{t}, \theta)$'s center to the origin.

Let $V(\mathcal{C})$ be the area of compact set $\mathcal{C} \in \mathbb{R}^2$ and $\mathcal{V}_{ij}(\vec{t}, \theta) := V(A_i(\vec{t}, \theta) \cap B_j)$. The *area of overlap* of $A(\vec{t}, \theta)$ and B, as $\vec{t}, \theta$ vary, is a function $\mathcal{V} : \mathcal{I} \rightarrow \mathbb{R}$ with $\mathcal{V}(\vec{t}, \theta) = V(A(\vec{t}, \theta) \cap B) = \sum_{A_i \in A, B_j \in B} V(A_i(\vec{t}, \theta) \cap B_j)$. We investigate the following problem: *For two sets A, B, defined as above, compute the optimal isometry $I_{\vec{t}_{opt}, \theta_{opt}}$ that maximizes $\mathcal{V}(\vec{t}, \theta)$.*

One can prove that, the maximum number of combinatorially distinct translations of A with respect to B can be as high as $\Theta(n^2 m)$ and the function describing the area of overlap is quite complex, even for combinatorially equivalent translations. If rotations are allowed as well the complexity of the configuration space can be as high as $\Theta(m^2 n^3)$.

Our major contributions are the following: (i) a lower bound on the area of overlap which is vital for almost all our algorithms, (ii) a nearly quadratic deterministic algorithm for translations, (iii) deterministic and nearly quadratic probabilistic algorithms for rigid motions. The algorithms are based on a simple two step framework in which an approximation of the best translation if followed by an approximation of the best rotation, whenever the latter applies. This way we first achieve an absolute error on $\mathcal{V}(\vec{t}_{opt}, \theta_{opt})$ which we then turn into a relative one using the lower bound theorem. In the deterministic one we employ a clever sampling of configuration space directed by some special properties of the function $\mathcal{V}_{ij}$. The probabilistic one is a combination of sampling of translation space, random sampling of both input sets and the technique by Cheong et al. [2].

2. The translational case

First, we present a lower bound on the maximum area of overlap under translation.

Theorem 1 *Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$, be two sets of disjoint unit disks in the plane. Let $\vec{t}_{opt}$ be the translation that maximizes the area of overlap $\mathcal{V}(\vec{t})$ of $A(\vec{t})$ and B over all possible translations $\vec{t}$ of set A. If k_{opt} is the number of overlapping pairs $A_i(\vec{t}_{opt})$ and B_j, then $\mathcal{V}(\vec{t}_{opt})$ is $\Theta(k_{opt})$.*

We proceed with the $(1-\epsilon)$-approximation algorithm. The algorithm is based on sampling of transformation space by using a uniform grid. This is possible due to the following lemma that states that $\mathcal{V}(\vec{t})$ is a sum of functions with some Lipschitz behaviour.

Lemma 2 *Let k_{opt} be the number of overlapping pairs $A_i(\vec{t}_{opt})$ and B_j. For any given $\delta > 0$ and any $\vec{t} \in \mathbb{R}^2$ for which $|\vec{t}_{opt} - \vec{t}| = O(\delta)$, we have $\mathcal{V}(\vec{t}_{opt}) - \mathcal{V}(\vec{t}) = O(k_{opt}\delta)$.*

Figure 1 shows algorithm TRANSLATION(A, B, ϵ).

TRANSLATION(A, B, ϵ):

(i) Initialize an empty binary search tree $\mathcal{S}$ with entries of the form $(\vec{t}, \mathcal{V}(\vec{t}))$ where $\vec{t}$ is the key.

(ii) For each pair of disks $A_i \in A$ and $B_j \in B$ do:
  - (a) Impose a uniform grid of spacing $\Theta(\epsilon)$ on $T_{ij} = B_j \ominus A_i$.
  - (b) For each grid point $\vec{t}_g \in T_{ij}$ do:
    - If $\vec{t}_g$ is in $\mathcal{S}$, then $\mathcal{V}(\vec{t}_g) := \mathcal{V}(\vec{t}_g) + \mathcal{V}_{ij}(\vec{t}_g)$ otherwise, insert $\vec{t}_g$ in $\mathcal{S}$ with $\mathcal{V}(\vec{t}_g) := \mathcal{V}_{ij}(\vec{t}_g)$.

(iii) Report the grid point $\vec{t}_{apx}$ that maximizes $\mathcal{V}(\vec{t}_g)$, and $\mathcal{V}(\vec{t}_{apx})$.

Fig. 1. Algorithm TRANSLATION(A, B, ϵ).

Theorem 3 *Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$ be two sets of disjoint unit disks in the plane. Let $\vec{t}_{opt}$ be the translation that maximizes $\mathcal{V}(\vec{t})$. Then, for any given $\epsilon > 0$,* TRANSLATION(A, B, ϵ) *computes a translation*

$\vec{t}_{apx}$, *for which* $\mathcal{V}(\vec{t}_{apx}) \geq (1-\epsilon)\mathcal{V}(\vec{t}_{opt})$, *in* $O((mn/\epsilon^2)\log(m/\epsilon))$ *time.*
All these results hold for sets of convex homothets where the ratio of the areas of any two objects is bounded and any object intersects only a constant number of other objects in its set [3].

3. The rotational case

We consider the following restricted scenario: set B is fixed, and set A can be rotated around the origin. This problem has a one-dimensional configuration space: the angle of rotation. Consider the function $\mathcal{V} : [0, 2\pi) \to \mathbb{R}$ defined by $\mathcal{V}(\theta) := V(A(\theta) \cap B)$. We start with a result that states that the term $\mathcal{V}_{ij}(\theta)$ has some Lipschitz behaviour.
Lemma 4 *Let* A_i, B_j *be any fixed pair of disks. For any given* $\delta > 0$ *and any* θ_1, θ_2 *for which* $|\theta_1 - \theta_2| \leq \delta/(2r_i)$, *we have* $|\mathcal{V}_{ij}(\theta_1) - \mathcal{V}_{ij}(\theta_2)| \leq 2\delta$.

For a pair A_i, B_j, we define the rotational interval $R_{ij} = \{\theta \in [0, 2\pi) : A_i(\theta) \cap B_j \neq \emptyset\}$. Note that R_{ij} is connected under the considered topology; we denote its length by $|R_{ij}|$. Instead of computing $V_{ij}(\theta)$ at each $\theta \in R_{ij}$, we would like to sample it at regular intervals whose length is at most $\delta/(2r_i)$. At first, it looks as if we would have to take an infinite number of sample points as $r_i \to \infty$. However, as the following lemma shows, the number of samples we need to consider is bounded.
Lemma 5 *For any* A_i, B_j *with* $r_i > 0$, *and any given given* $\delta > 0$, *we have* $|R_{ij}|/(\delta/(2r_i)) = O(1/\delta)$.
Algorithm ROTATION(A, B, δ), see Figure 2, maximizes $\mathcal{V}(\theta)$ up to an absolute error. By computing a value $\tilde{\mathcal{V}}(\theta)$ that approximates $\mathcal{V}(\theta)$, we save a linear factor in the running time.
Lemma 6 *Let* θ_{opt} *be the rotation that maximizes* $\mathcal{V}(\theta)$ *and let* k_{opt} *be the number of overlapping pairs* $A_i(\theta_{opt})$ *and* B_j. *For any given* $\delta > 0$, *the rotation* θ_{apx} *reported by* ROTATION(A, B, δ) *satisfies* $\mathcal{V}(\theta_{opt}) - \mathcal{V}(\theta_{apx}) = O(k_{opt}\delta)$, *and can be computed in* $O((mn/\delta)\log m)$ *time.*

ROTATION(A, B, δ):

(i) For each pair of disks $A_i \in A$ and $B_j \in B$, choose points $\Theta_{ij} := \{\theta_{ij}^1, \ldots, \theta_{ij}^{s_{ij}}\}$ with an uniform spacing of $\delta/(2r_i)$ and such that $R_{ij} \subset [\theta_{ij}^1, \theta_{ij}^{s_{ij}}]$. Make sure that the midpoint of R_{ij} is in Θ_{ij}.
(ii) Sort the values $\Theta := \bigcup_{i,j} \Theta_{ij} = \{\theta_{ij}^s\}$, keeping repetitions and solving ties arbitrarily, if applicable. Let $\theta_0, \theta_1, \ldots$ be the ordering of Θ.
(iii) Let $\tilde{\mathcal{V}} : \Theta \to \mathbb{R}$ with $\tilde{\mathcal{V}}(\theta_0) := \mathcal{V}(\theta_0)$.
(iv) For each $\theta_l = \theta_{ij}^s$ in increasing order of l do:
- If $\mathcal{V}_{ij}$ is decreasing at θ_{ij}^s, or θ_{ij}^s is the midpoint of R_{ij}, then $\tilde{\mathcal{V}}(\theta_l) := \tilde{\mathcal{V}}(\theta_{l-1}) - \mathcal{V}_{ij}(\theta_{ij}^{s-1}) + \mathcal{V}_{ij}(\theta_{ij}^{s+1})$
- If $\mathcal{V}_{ij}$ is increasing at θ_{ij}^s, then $\tilde{\mathcal{V}}(\theta_l) := \tilde{\mathcal{V}}(\theta_{l-1}) - \mathcal{V}_{ij}(\theta_{ij}^{s-1}) + \mathcal{V}_{ij}(\theta_{ij}^s)$
(v) Report the $\theta_{apx} \in \Theta$ that maximizes $\tilde{\mathcal{V}}(\theta)$.

Fig. 2. Algorithm ROTATION(A, B, δ).

4. A $(1-\epsilon)$-approximation algorithm for rigid motions

A simple algorihtm RIGIDMOTION(A, B, ϵ) which gives a $(1-\epsilon)$-approximation of the optimum is described in Figure 3. It uses the fact that any isometry can be expressed as a translation followed by a rotation around the origin.

RIGIDMOTION(A, B, ϵ):

(i) For each pair of disks $A_i \in A$ and $B_j \in B$ do:
  (a) Set the center of rotation, i.e the origin, to be B_j's center by translating B appropriately. Impose a uniform grid of spacing $\Theta(\epsilon)$ on $T_{ij} = B_j \ominus A_i$.
  (b) For each grid point $\vec{t}_g \in T_{ij}$ do:
  - run ROTATION$(A(\vec{t}_g), B, c\epsilon)$, where c is an appropriate constant; let θ_{apx}^g be the rotation returned. Compute $\mathcal{V}(\vec{t}_g, \theta_{\text{apx}}^g)$.
(ii) Report the pair $(t_{apx}, \theta_{\text{apx}})$ that maximizes $\mathcal{V}(\vec{t}_g, \theta_{\text{apx}}^g)$, and $\mathcal{V}(t_{apx}, \theta_{\text{apx}})$.

Fig. 3. Algorithm RIGIDMOTION(A, B, ϵ).

Theorem 7 *Let* $A = \{A_1, \ldots, A_n\}$ *and* $B = \{B_1, \ldots, B_m\}$ *be two sets of disjoint unit disks in the plane. Let* $I_{\vec{t}_{opt}, \theta_{opt}}$ *be the isometry that maximizes* $\mathcal{V}(\vec{t}, \theta)$. *Then, for any given* $\epsilon > 0$, RIGIDMOTION(A, B, ϵ) *computes an isometry* $I_{\vec{t}_{apx}, \theta_{apx}}$, *for which* $\mathcal{V}(\vec{t}_{apx}, \theta_{apx}) \geq (1-\epsilon)\mathcal{V}(\vec{t}_{opt}, \theta_{opt})$, *in* $O((n^2m^2/\epsilon^3)\log m)$ *time.*
We can modify the algorithm such that its running time depends on the diameter Δ of the set A. The main idea is to convert our algorithm into one that is sensitive to the number of pairs of disks in A and B that have approximately the same distance, and then use the combinatorial bounds by Gavrilov et al. [4, Theorem 4.1]. In many applications it

is reasonable to assume bounds of the type $\Delta = O(n)$ [4], and therefore the result below is relevant.

Theorem 8 *Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$ be two sets of disjoint unit disks in the plane. Let Δ be the diameter of A, and let $I_{\vec{t}_{opt}, \theta_{opt}}$ be the isometry that maximizes $\mathcal{V}(\vec{t}, \theta)$. Then, for any given $\epsilon > 0$, an isometry $I_{\vec{t}_{apx}, \theta_{apx}}$ such that $\mathcal{V}(\vec{t}_{apx}, \theta_{apx}) \geq (1-\epsilon)\mathcal{V}(\vec{t}_{opt}, \theta_{opt})$ can be found in $O(\frac{m^2 n^{4/3} \Delta^{1/3} \log n \log m}{\epsilon^3})$ time.*

5. A Monte Carlo Algorithm

In this section we present a nearly quadratic Monte Carlo randomized algorithm that computes a $(1-\epsilon)$-approximation of the maximum area of overlap with high probability. The algorithm works under the assumption that $\mathcal{V}(\vec{t}_{opt}, \theta_{opt}) \geq \alpha V(B) = \alpha m \pi$, for some $0 < \alpha \leq 1$, which implies that $n = \Theta(m)$ and $k_{opt} = \Theta(n)$. As noted by Cheong et al. [2], this is a reasonable assumption to make in many shape matching applications.

The first step is a combination of random sampling of set A and the sampling of the space of translations. This is based on the observation that the deterministic algorithm of Section 4 will compute a $(1-\epsilon)$-approximation k_{opt} times. Intuitively, the larger the k_{opt}, the less the neccessary number of pairs that have to be tried out in this step. The second step is based on a direct application of a technique by Cheong et al. [2]. The technique allows us to maximize, up to an absolute error, the area of overlap under rotation in linear time, by computing a point of maximum depth in a one dimensional arrangement. We sumarize it below.

We choose a uniform random sample S of points in B, and for a rotation θ we compute the estimate $e_S(\theta) = \frac{|A(\theta) \cap S|}{S}$. For a point $s \in S$ we define $W(s) = \{\theta | s \in A(\theta)\}$. Let $\Theta_A(S)$ be the arrangement of all regions $W(s)$; it is a one-dimensional arragement of rotational intervals.

Lemma 9 *Let θ_{opt} be the rotation that maximizes $\mathcal{V}(\theta)$. For any given $\delta > 0$, let S be a uniform random sample of points in B with $|S| \geq c_1 \frac{\log n}{\delta^3}$ where c_1 is an appropriate constant. The vertex θ_{apx} of $\Theta_A(S)$ that maximizes $e_S(\theta)$ satisfies $\mathcal{V}(\theta_{opt}) - \mathcal{V}(\theta_{apx}) \leq \delta V(B)$ with probability at least $1 - 1/n^6$.*

Algorithm RandomRigidMotion$(A, B, \alpha, c, \epsilon)$ is given in Figure 4.

RandomRigidMotion$(A, B, \alpha, c, \epsilon)$:

(i) Choose a uniform random sample $R_A \subset A$, with $|R_A| = \Omega(c/\alpha)$, and a uniform random sample S of points in B, with $|S| = \Omega(\log|A|/\epsilon^3)$.

(ii) For each pair of disks $A_i \in R_A$ and $B_j \in B$ do:
  - (a) Impose a uniform grid of spacing $\Theta(\epsilon)$ on $T_{ij} = B_j \ominus A_i$.
  - (b) For each grid point $\vec{t}_g \in T_{ij}$ do:
    - Compute a vertex θ^g_{apx} of maximum depth in $\Theta_{A(\vec{t}_g)}(S)$, and $\mathcal{V}(\vec{t}_g, \theta^g_{apx})$.

(iii) Report the pair $(\vec{t}_{apx}, \theta_{apx})$ that maximizes $\mathcal{V}(\vec{t}_g, \theta^g_{apx})$, and $\mathcal{V}(\vec{t}_{apx}, \theta_{apx})$.

Fig. 4. Algorithm RandomRigidMotion$(A, B, \alpha, c, \epsilon)$.

Theorem 10 *Let $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_m\}$, be two sets of disjoint unit disks in the plane. Let $I_{\vec{t}_{opt}, \theta_{opt}}$ be the isometry that maximizes $\mathcal{V}(\vec{t}, \theta)$ with the assumption that $\mathcal{V}(\vec{t}_{opt}, thopt) \geq \alpha V(B)$, for some $0 < \alpha \leq 1$. Then, for any given $\epsilon > 0$, $c \geq \alpha/6$,* RandomRigidMotion$(A, B, \alpha, c, \epsilon)$ *computes an isometry $I_{\vec{t}_{apx}, \theta_{apx}}$, for which $\mathcal{V}(\vec{t}_{apx}, \theta_{apx}) \geq (1-\epsilon)\mathcal{V}(\vec{t}_{opt}, \theta_{opt})$, in $O((cn^2/\epsilon^5)\log^2(n/\epsilon))$ time. The algorithm succeeds with probability at least $1 - \epsilon^{-c} - 1/n^6$.*

References

[1] H. Alt and L. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In J.R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 121–153. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 1999.

[2] O. Cheong, A. Efrat, and S. Har-Peled. On finding a guard that sees most and a shop that sells most. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, page to appear, 2004.

[3] M. de Berg, P. Giannopoulos, C. Knauer, R. van Oostrum, and R. C. Veltkamp. The area of overlap of two unions of convex objects under translations. Technical Report UU-CS-2003-025, Institute of Information and Computing Sciences, Utrecht University, The Netherlands, 2003.

[4] M. Gavrilov, P. Indyk, R. Motwani, and S. Venkatasubramanian. Combinatorial and experimental methods for approximate point pattern matching. *Algorithmica*, 38(2):59–90, 2003.

[5] D. M. Mount, R. Silverman, and A. Y. Wu. On the area of overlap of translated polygons. *Computer Vision and Image Understanding*, 64:53–61, 1996.