Probabilistic matching of solids in arbitrary dimension

Daria Schymura*

Abstract

We give simple probabilistic algorithms that approximately maximize the volume of overlap of two solid shapes under translations and rigid motions. The shapes are subsets of $\mathbb{R}^d$, $d \geq 2$. The algorithms approximate w.r.t. an prespecified additive error and succeed with high probability. Apart from measurability assumptions, we only require from the shapes that uniformly distributed random points can be generated. An important example are finite unions of simplices that have pairwise disjoint interiors.

1 Introduction

We design and analyze simple probabilistic algorithms for matching solid shapes in $\mathbb{R}^d$ under translations and rigid motions. An important example of solid shapes are finite unions of simplices that have pairwise disjoint interiors. On the input of two shapes A and B, the algorithms compute a transformation t^* such that the volume of overlap of $t^*(A)$ and B is approximately maximal.

First, we explain the algorithm and its analysis for the case of translations. In Section 4, we show how to generalize the results to rigid motions.

For translations, the idea of the algorithm is as follows. Given two shapes A and B, a point $a \in A$ and a point $b \in B$ are picked uniformly at random. This tells us that the translation t that is given by the vector $b - a$ maps some part of A onto some part of B. We record this as a "vote" for t and repeat this procedure very often. Then we determine the densest cluster of the resulting point cloud of translation vectors, and output the center of this cluster. This translation maps a large part of A onto B. The details of the algorithm are explained in Section 2.

We show that this algorithm approximates the maximal volume of overlap under translations. More precisely, let t^{opt} be a translation that maximizes the volume of overlap of A and B, and let t^* be a translation that is computed by the algorithm. Given an error bound ε and an allowable probability of failure p, both between 0 and 1, we show bounds on the required number of random experiments, guaranteeing that the difference between approximation and optimum $|t^{\text{opt}}(A) \cap B| - |t^*(A) \cap B| \leq \varepsilon|A|$ with probability at least $1 - p$. Here $|\cdot|$ denotes the volume (Lebesgue measure) of a set. We use $\varepsilon|A|$ and not just ε as error bound because the inequality should be invariant under scaling of both shapes with the same factor.

In a previous publication [1] we considered the 2-dimensional case. Here we not only generalize the results to higher dimensions, but we also give new proofs that improve the bounds on the number of random samples.

Furthermore we considerably improve the time complexity of the algorithm by showing that a simpler definition of a cluster suffices to guarantee approximation. In [1], a translation whose neighborhood contained the maximal number of "votes" was computed, which boiled down to computing a deepest cell of an arrangement of boxes. For N boxes, the best known time bound is $O(N^{d/2}/(\log N)^{d/2-1} \operatorname{polyloglog} N)$ [2]. Here we show that it is sufficient to output the "vote" whose neighborhood contains the maximal number of "votes", which can be computed by brute force in time $O(N^2)$ in any dimension. The time bound can be further improved to $O(N(\log N)^{d-1})$ by using orthogonal range counting queries [4].

Cheong et al. [3] introduce a general probabilistic framework, which they use for approximating the maximal area of overlap of two unions of n and m triangles in the plane, with prespecified absolute error ε, in time $O(m + (n^2/\varepsilon^4)(\log n)^2)$ for translations and in time $O(m + (n^3/\varepsilon^8)(\log n)^5)$ for rigid motions. The latter time bound is smaller in their paper, due to a calculation error in the final derivation of the time bound, as was noted in [11]. Their algorithm works with high probability.

For two simple polygons with n and m vertices in the plane, Mount et al. [8] show that a translation that maximizes the area of overlap can be computed in time $O(n^2m^2)$.

For maximizing the volume of overlap of two unions of simplices under rigid motions, no exact algorithm that runs in polynomial time is known, not even in the plane. Vigneron gives an FPTAS with relative error ε for dimension $d \geq 2$ [11]. For two polyhedra P and Q in $\mathbb{R}^d$, given as the union of m and n simplices, respectively, the algorithm for approximating the maximal volume of overlap has time complexity $O\big((\frac{nm}{\varepsilon})^{d^2/2+d/2+1}(\log \frac{nm}{\varepsilon})^{d^2/2+d/2+1}\big)$, which can be improved to $O\big((n^6/\varepsilon^3)\log^4(n/\varepsilon)\beta(n/\varepsilon)\big)$ in the plane where β is a very slowly growing function related to the inverse Ackermann function.

*Institute of Computer Science, Freie Universität Berlin, schymura@inf.fu-berlin.de

2 A probabilistic algorithm for matching under translations

Before stating the main result for translations, we introduce some definitions. The boundary ∂A of a set $A \subseteq \mathbb{R}^d$ is the set of points that are contained in its closure, but not in its interior. We measure the boundary of d-dimensional sets by the $(d-1)$-dimensional Hausdorff measure, and denote it slightly sloppy by $|\partial A|$. For a definition of the Hausdorff measure and related definitions, we refer the reader to [7]. The *isoperimetric quotient* of $A \subset \mathbb{R}^d$ is defined to be $|\partial A|^d / |A|^{d-1}$. The isoperimetric quotient can be considered as a certain measure of the *fatness* of a figure A.

We always assume shapes to be Lebesgue measurable subsets of $\mathbb{R}^d$ that have positive, finite volume and whose boundary is measurable by the $(d-1)$-dimensional Hausdorff measure and has positive, finite $(d-1)$-dimensional volume.

Our algorithm can be applied to all shapes from which we can generate random sample points. For unions of simplices, the runtime of our method depends only linearly on the number of vertices, but it depends more significantly on fatness parameters as the isoperimetric quotient of one of the shapes and, in the case of rigid motions, on the ratio $\mathrm{diam}(A)^d / |A|$.

Theorem 1 (Runtime for translations) *Let A and B be shapes in constant dimension d, and let $\varepsilon, p \in (0,1)$ be parameters. Assume that we are given a lower bound on $|A|$ and an upper bound on $|B|$, which differ only by a constant, and an upper bound K_A on the isoperimetric quotient of A. Assume further that N uniformly distributed random points can be generated from a shape in time $T(N)$.*

Then a translation that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability $\geq 1-p$ can be computed in time $O\big(T(N) + N(\log N)^{d-1}\big)$ where $N = O\big(\varepsilon^{-(2d+2)} K_A^2 \log \frac{2}{p}\big)$.

If A and B are finite unions of at most n simplices that have pairwise disjoint interiors, then $T(N) = O(n + N \log n)$.

To prove this theorem, we describe an algorithm, Algorithm 1, and then prove that it is correct and has the runtime claimed in the theorem.

The following theorem gives bounds on the output of ClusteringSize(ε, A) and SampleSize$(B, \varepsilon, \delta, p)$ in Algorithm 1 that guarantee that the output of Algorithm 1 approximates the maximal volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$.

Of course, the roles of A and B can be swapped, which may result in better bounds. For the shortness of presentation, we do not reflect this fact in the following.

Algorithm 1: ProbMatchT

Input: shapes $A, B \subset \mathbb{R}^d$, error bound $\varepsilon \in (0,1)$, allowed probability of failure $p \in (0,1)$

```
real δ ← ClusteringSize(ε, A);
integer N ← SampleSize(B, ε, δ, p);
collection Q ← ∅;
for i = 1 . . . N do
    point a ← randomPoint(A);
    point b ← randomPoint(B);
    add(Q, b − a);
end
return FindDensestClusterT(Q, δ);
```

Function `FindDensestClusterT`(Q, δ)

Input: collection Q of points in $\mathbb{R}^d$, positive number δ

Output: point t in Q such that the cube of side length 2δ that is centered at t contains a maximal number of points from Q

Theorem 2 (Correctness of Algorithm 1) *Let A and B be shapes in constant dimension d, and let $\varepsilon, p \in (0,1)$ be parameters. If ClusteringSize(ε, A) returns a positive number $\delta \leq \frac{d}{d-1} \cdot \frac{\sqrt{2}}{3\sqrt{d}} \cdot \varepsilon \cdot \frac{|A|}{|\partial A|}$ and SampleSize$(B, \varepsilon, \delta, p)$ returns an integer $N \geq C\varepsilon^{-2}\,\delta^{-2d}\,|B|^2 \log \frac{2}{p}$ for some universal constant $C > 0$, then Algorithm 1 computes on the input (A, B, ε, p) a translation that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$.*

The universal constant C can be deduced from the proofs. Note that, when both shapes are scaled with the same factor, the sample size N does not change. In other words, it is homogeneous. The clustering size scales with the shapes. This is reasonable because blowing up the shapes coarsens fine features.

3 Proof idea of Theorem 2

Let μ be the probability distribution on the translation space that is induced by the random experiment in the algorithm, and let μ_N be the empirical measure. Recall that μ has a *density function* f if the probability $\mu(E)$ of any event $E \subseteq \Omega$ can be computed as the integral over the density function $\int_E f(x)dx$.

The main idea is to prove that the density function of μ is proportional to the objective function, that is the function that maps a translation vector $t \in \mathbb{R}^d$ to the volume of the intersection of $t(A)$ and B. Thus the goal is to find a translation at which the density function is approximately maximal.

Conceptually, the density function is approximated in a two step process. Let $B(t, \delta)$ be a ball of radius δ, centered at t, w.r.t. the metric that is induced by the maximum norm. First, $f(t)\cdot|B(t,\delta)|$ is close to $\mu(B(t,\delta))$ if f is nice enough and δ is sufficiently small. Second, the probability of a small cube $\mu(B(t,\delta))$ is close to $\mu_N(B(t,\delta))$ if N is sufficiently large.

The analysis of the algorithm is based on these simple ideas whose details are not that easy. They are hidden in the following longish theorem that follows from theorems in Chapters 3 and 4 of [5].

Theorem 3 (Probabilistic toolbox) *Let $\Omega \subseteq \mathbb{R}^k$ be a metric space, and let $\mathcal{B}$ be the set of balls of some fixed radius $\delta > 0$ in Ω. Let* vol *be a measure on Ω such that, for all $x \in \Omega$, the volume $\mathrm{vol}(B(x,\delta)) = v_\delta$ for some $v_\delta > 0$. Assume further that $\mathcal{B}$ has finite VC dimension V.*

Let μ be a probability measure on Ω that has a Lipschitz continuous density function f with Lipschitz constant L. Let $X_1,\dots,X_N$ be i.i.d. random variables taking values in Ω with common distribution μ and empirical measure μ_N.

Let $j \in \{1,\dots,N\}$ be such that $\mu_N(B(X_j,\delta)) = \max_{1\le i\le N}\mu_N(B(X_i,\delta))$. Then, for all $\tau > 0$ and for all $x \in \Omega$, with probability $\ge 1 - 2e^{-2N\tau^2}$, we have $f(X_j) \ge f(x) - 2(c\sqrt{V/N}+\tau)/v_\delta - 3L\delta$.

In this inequality, c is a universal constant.

The key lemma to apply Theorem 3 states that the density function of μ is proportional to the objective function. It follows from a transformation rule for density functions.

Lemma 4 (Key lemma) *Let X be the random vector on $\mathbb{R}^d$ that draws translations $t = b - a$ where $(a,b) \in A \times B \subset \mathbb{R}^{2d}$ is drawn uniformly at random. The density function of X is given by $f(t) = \frac{|t(A)\cap B|}{|A|\cdot|B|}$.*

Furthermore we have to show that the density function f is Lipschitz continuous. In the following theorem, let $\mathcal{H}^k$ denote the k-dimensional Hausdorff measure. For Lebesgue measurable sets in $\mathbb{R}^d$, the d-dimensional Hausdorff measure and the Lebesgue measure coincide. The symmetric difference is denoted by $\triangle$.

Theorem 5 *[9] Let $A \subset \mathbb{R}^d$ be bounded. Let $t \in \mathbb{R}^d$ be a translation vector.*

Then $\mathcal{H}^d(A \triangle (A+t)) \le |t|\,\mathcal{H}^{d-1}(\partial A)$.

This implies that the density function f is Lipschitz continuous. Applying the Cauchy-Schwarz inequality yields $||t-s||_2 \le \sqrt{d}||t-s||_\infty$, which is best possible.

Corollary 6 *The function f on $\mathbb{R}^d$ that is given by $f(t) = \frac{|t(A)\cap B|}{|A|\cdot|B|}$ is Lipschitz continuous with constant $L = \frac{\sqrt{d}\,|\partial A|}{2|A|\cdot|B|}$ w.r.t. the metric that is induced by the maximum norm.*

With these results, the proof of Theorem 2 is an application of Theorem 3. Note that the VC dimension of the class of rectangles in $\mathbb{R}^d$ equals $2d$ [5].

4 Rigid motions

A rigid motion r on $\mathbb{R}^d$ is given by $r(x) = Mx + t$ where $M \in \mathbb{R}^{d\times d}$ is a rotation matrix and $t \in \mathbb{R}^d$ is a translation vector. A matrix M is contained in the group of rotation matrices $SO(d) \subset \mathbb{R}^{d\times d}$ if it is orthogonal, meaning $M^T = M^{-1}$, and $\det M = 1$. We identify each rigid motion with the pair of its rotation matrix and translation vector. Thus the space of rigid motions equals $SO(d)\times\mathbb{R}^d$.

The algorithm for rigid motions works similarly as the algorithm for translations. We draw a rotation matrix $M \in SO(d)$, a point $a \in A$ and a point $b \in B$ uniformly at random. Then we register the unique rigid motion that has M as rotation matrix and maps a onto b as a "vote" in the transformation space. After many rounds, say N, we determine the best cluster in the space of rigid motions.

There are many different methods to compute a random rotation matrix described in the literature; see for example [6, 10]. To define the uniform distribution formally, a volume has to be defined. The group $SO(d)$ is $\binom{d}{2}$-dimensional. The volume $|\cdot|$ in $SO(d)$ is measured by the $\binom{d}{2}$-dimensional Haar measure.

For a matrix $M = (m_{ij})_{1\le i,j\le d}$, let $||M||_2 = \sqrt{\sum_{1\le i,j\le d}(m_{ij})^2}$ be the Frobenius norm. Denote the Euclidean norm on $\mathbb{R}^d$ also by $||\cdot||_2$. Define $B_2(M,\delta)$ and $B_2(t,\delta)$ to be the closed balls of radius δ w.r.t. the metrics induced by the Frobenius and the Euclidean norm.

We define a *δ-neighborhood* of a rigid motion (M,t) by $B((M,t),\delta) = B_2(M,\delta/\operatorname{diam}(A)) \times B_2(t,\delta)$. The radius of the rotational part of the neighborhood depends on the diameter of A because it should not change if A is scaled. The "rotational distance" of shapes does not depend on their absolute size. A best cluster is the random rigid motion whose neighborhood contains the maximal number of other random rigid motions. We give a pseudocode description of the algorithm for rigid motions, Algorithm 2.

Theorem 7 (Correctness of Algorithm 2) *Let A and B be shapes in constant dimension d, and let $\varepsilon, p \in (0,1)$ be parameters.*

There are constants $C, C' > 0$ such that, if ClusteringSize(ε, A) returns a positive $\delta \le C\varepsilon\frac{|A|}{|\partial A|}$ and SampleSize(B,ε,δ,p) returns an integer $N \ge C'\varepsilon^{-2}\,\delta^{-d^2-d}\,\operatorname{diam}(A)^{-d^2+d}\,|B|^2\log\frac{2}{p}$, then Algorithm 2 computes on the input (A,B,ε,p) a rigid motion that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability $\ge 1-p$.

Algorithm 2: ProbMatchRM

```
Input: shapes A, B ⊂ ℝ^d, error bound ε ∈ (0,1),
       allowed probability of failure p ∈ (0,1)
real δ ← ClusteringSize(ε, A);
integer N ← SampleSize(B, ε, δ, p);
collection Q ← ∅;
for i = 1 ... N do
    rotation matrix M ← randomRotation();
    point a ← randomPoint(A);
    point b ← randomPoint(B);
    add(Q, (M, b − Ma));
end
return FindDensestClusterRM(Q, δ, diam(A));
```

Function `FindDensestClusterRM`(Q, δ, Δ)

```
Input: collection Q of points in ℝ^{d×d} × ℝ^d,
       positive numbers δ and Δ
Output: point (M, t) in Q such that the
        neighborhood B_2(M, δ/Δ) × B_2(t, δ) of
        (M, t) contains a maximal number of
        points from Q
```

Theorem 8 (Runtime for rigid motions) *Let A, B, ε, and p be given under the same assumptions as in Theorem 1. Additionally, let D_A be given such that $\frac{\operatorname{diam}(A)^d}{|A|} \leq D_A$. Then a rigid motion that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability $\geq 1-p$ can be computed in time $O\big(T(N) + N^2\big)$ where $N = O\big(\varepsilon^{-(d^2+d+2)} K_A^{d+1} D_A^{d-1} \log \frac{2}{p}\big)$.*

As in the case of translations, the density function on the transformation space is proportional to the objective function, and it is Lipschitz continuous:

Lemma 9 *Let Z be the random vector that draws rigid motions $(M, b - Ma) \in SO(d) \times \mathbb{R}^d$ where $(M, a, b) \in SO(d) \times A \times B$ is drawn u.a.r. The density function of Z is given by $g(r) = \frac{|r(A) \cap B|}{|SO(d)| \cdot |A| \cdot |B|}$.*

The function g is Lipschitz continuous with constant $L = \frac{|\partial A|}{|SO(d)| \cdot |A| \cdot |B|}$ w.r.t. the metric $d(r,s) = \max\{\operatorname{diam}(A) \cdot ||M - N||_2\,,\, ||p - q||_2\}$ for rigid motions $r = (M,p)$ and $s = (N,q)$.

Observe that a δ-neighborhood of a rigid motion, as defined above, equals the closed ball of radius δ w.r.t. the metric d. In order to apply Theorem 3, one has to prove that all neighborhoods have the same volume, which follows from the fact that the $\binom{d}{2}$-dimensional Hausdorff measure is a Haar measure on $SO(d)$. Additionally, one has to compute a lower bound on the volume of such a neighborhood.

The Lipschitz continuity of g can be deduced using Theorems 5, 10 and the fact that the volume of the symmetric difference fulfills the triangle inequality. We assume that A and B have $(\mathcal{H}^{d-1}, d-1)$-rectifiable boundaries.

Theorem 10 *[9] Let $A \subset \mathbb{R}^d$ be bounded. Let $M \in \mathbb{R}^{d \times d}$ be a rotation matrix and let $w = \max_{a \in \partial A} ||a - Ma||_2$.*

Then $\mathcal{H}^d(A \triangle MA) \leq (2d/d+1)^{\frac{d-1}{2}} \cdot w \cdot \mathcal{H}^{d-1}(\partial A)$.

The constant $(2d/d+1)^{\frac{d-1}{2}}$ can be replaced by 1 for sets that have an $(\mathcal{H}^{d-1}, d-1)$-rectifiable boundary.

In the 2-dimensional case, the results can be significantly improved by representing a rigid motion by the pair of its rotation angle (instead of the rotation matrix) and the translation vector, as it was done in [1].

References

[1] H. Alt, L. Scharf, and D. Schymura. Probabilistic matching of planar regions. *Computational Geometry, Theory and Applications (CGTA)*, 43:99–114, 2010. Special Issue on the 24th European Workshop on Computational Geometry (EuroCG'08).

[2] T.M. Chan. A (slightly) faster algorithm for Klee's measure problem. *Computational Geometry*, 43(3):243 – 250, 2010.

[3] O. Cheong, A. Efrat, and S. Har-Peled. Finding a guard that sees most and a shop that sells most. *Discrete and Computational Geometry*, 37:545–563, 2007.

[4] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry.* Springer, third edition, 2008.

[5] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation.* Springer, 2001.

[6] R.M. Heiberger. Generation of random orthogonal matrices. *Applied Statistics*, 27:199–206, 1978.

[7] H. Federer. *Geometric Measure Theory.* Springer, 1969.

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

[9] D. Schymura. An upper bound on the volume of the symmetric difference of a body and a congruent copy. *ArXiv e-prints*, 2010. `http://arxiv.org/abs/1010.2446`.

[10] M.A. Tanner and R.A. Thisted. A remark on AS127. Generation of random orthogonal matrices. *Applied Statistics*, 31:190–192, 1982.

[11] A. Vigneron. Geometric optimization and sums of algebraic functions. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, pages 906–917, 2010.