Probabilistic Matching of Solid Shapes in Arbitrary Dimension

Dissertation zur Erlangung des Doktorgrades

vorgelegt am

Fachbereich Mathematik und Informatik
der Freien Universität Berlin

2011

von

Daria Schymura

Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
daria.schymura@gmx.net

gefördert durch die EU im Rahmen des Forschungsprojektes *PROFI*
unter der Vertragsnummer IST-511572-2,
durch die DFG im Rahmen des Graduiertenkollegs 1408
Methods for Discrete Structures
und des Schwerpunktprogramms 1307 *Algorithm Engineering*

Betreuer: Prof. Dr. Helmut Alt
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
alt@mi.fu-berlin.de

Gutachter: Prof. Dr. Helmut Alt
Institut für Informatik
Freie Universität Berlin
Takustraße 9
14195 Berlin
Germany
alt@mi.fu-berlin.de

Prof. Dr. Otfried Cheong
Department of Computer Science
KAIST
291~Daehak-ro, Yuseong-gu
Daejeon 305-701
Korea
otfried@kaist.edu

Vorlage zur Begutachtung: 1. September 2011
Termin der Disputation: 18. November 2011
Fassung vom: 21. November 2011

To understand is to perceive patterns.

Isaiah Berlin

Abstract

Determining the similarity between objects is a fundamental problem in computer vision and pattern recognition, but also in other fields of computer science. This thesis concentrates on the matching problem, which has received a lot of attention in Computational Geometry.

Given a class of shapes S, a set of transformations T, mapping shapes onto shapes, and a distance measure d on S, the *matching problem* with respect to S, T, and d is defined as follows: Given two shapes A, $B \in S$, compute a transformation $t^* \in T$ that minimizes $d(t^*(A), B)$.

We consider solid shapes, i.e., full-dimensional shapes, in arbitrary dimension and assume that they are given by an oracle that generates uniformly distributed random points from the shapes. This is a very rich class of shapes that contains the class of finite unions of simplices as a subclass. We study matching under translations and rigid motions (translation and rotation). Throughout this work, the volume of the symmetric difference is used as distance measure for the matching problem. Maximizing the volume of the overlap is equivalent to minimizing the volume of the symmetric difference under translations and rigid motions.

We study a probabilistic approach to the shape matching problem. The main idea is quite simple. Given two shapes A and B, repeat the following random experiment very often: Select random point samples of appropriate size from each shape and compute a transformation that maps the point sample of one shape to the sample of the other shape. Store this transformation. In each step, we extend the collection of random transformations by one. Clusters in the transformation space indicate transformations that map large parts of the shapes onto each other. We determine a densest cluster and output its center.

This thesis describes probabilistic algorithms for matching solid shapes in arbitrary dimension under translations and rigid motions. The algorithms are a priori heuristics. The main focus is on analyzing them and on proving that they maximize the volume of overlap approximately by solving the following instance of the matching problem. Given two solid shapes A and B, an error tolerance $\varepsilon \in (0,1)$, and an allowed probability of failure $p \in (0,1)$, the problem is to compute a transformation t^* such that with probability at least $1-p$, we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all transformations t, in particular for transformations maximizing the volume of overlap. Therein $|\cdot|$ denotes the volume.

The approach is mainly of theoretical interest. Still, the algorithms are so simple that they can easily be implemented, which we show by giving experimental results of a test implementation for 2-dimensional shapes.

Zusammenfassung

Die zentrale Frage in der Musteranpassung ist, ob zwei gegebene Objekte A und B sich ähneln. Methoden, die diese Frage lösen, haben zahlreiche Anwendungen in verschiedenen Gebieten der Informatik. Muster können aufgrund verschiedener Merkmale verglichen werden, zum Beispiel aufgrund der Farbe, der Textur oder aufgrund von ausgezeichneten Punkten. In dieser Arbeit werden Muster aufgrund geometrischer Merkmale verglichen, wie es in der Algorithmischen Geometrie üblich ist.

Seien eine Menge von Mustern S, eine Menge von Transformationen T, die Muster auf Muster abbilden, und ein Abstandsmaß d gegeben. Das *Musteranpassungsproblem* bezüglich S, T und d besteht darin, für gegebene Muster $A, B \in S$ eine Transformation $t^* \in T$ zu berechnen, die den Abstand $d(t^*(A), B)$ minimiert.

Wir betrachten volldimensionale Muster in beliebiger Dimension und nehmen an, dass die Muster durch ein Orakel gegeben sind, das gleichverteilte Zufallspunkte aus ihnen erzeugt. Diese Klasse von Mustern ist sehr allgemein, da sie die Klasse der endlichen Vereinigungen volldimensionaler Simplizes als Teilklasse enthält. Als Transformationsklassen betrachten wir Translationen und starre Bewegungen. Als Abstandsmaß verwenden wir das Volumen der symmetrischen Differenz. Für Translationen und starre Bewegungen ist das Maximieren des Volumen des Durchschnitts äquivalent zum Minimieren des Volumens der symmetrischen Differenz.

In dieser Arbeit wird ein probabilistischer Ansatz für das Musteranpassungsproblem verfolgt. Die zentrale Idee ist relativ einfach. Für zwei gegebene Muster A und B wird das folgende Zufallsexperiment sehr oft wiederholt: Aus beiden Mustern wird eine gewisse Anzahl an Zufallspunkten erzeugt. Dann wird eine Transformation berechnet, die die Zufallspunkte aus A auf die Zufallspunkte aus B abbildet. Diese Transformation wird gespeichert. In jedem Schritt wird eine Zufallstransformation zur Menge der gespeicherten Transformationen hinzugefügt. Häufungen im Transformationsraum zeigen Transformationen an, die große Teile der Muster aufeinander abbilden. Es wird die dichteste Häufung im Transformationsraum gesucht und ihr Mittelpunkt als Ergebnis ausgegeben.

Wir beschreiben probabilistische Algorithmen zur Anpassung volldimensionaler Muster in beliebiger Dimension unter Translationen und starren Bewegungen. A priori sind diese Algorithmen Heuristiken. Der Schwerpunkt der Arbeit liegt darauf, die Algorithmen zu analysieren und zu beweisen, dass sie das Volumen des Durchschnitts approximativ maximieren. Genauer gesagt lösen die Algorithmen folgendes Problem: Gegeben zwei Muster A und B, eine Fehlerschranke $\varepsilon \in (0, 1)$ und eine erlaubte Fehlerwahrscheinlichkeit $p \in (0, 1)$, berechne eine Transformation t^*, so dass mit Wahrscheinlichkeit mindestens $1 - p$ gilt, dass $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ für alle Transformationen t gilt. Das gilt insbesondere für Transformationen, die das Volumen des Durchschnitts maximieren.

Unser Ansatz ist zwar hauptsächlich von theoretischem Interesse, dennoch sind die vorgestellten Algorithmen so einfach, dass sie leicht implementiert werden können, was wir durch experimentelle Ergebnisse einer Testimplementation für zweidimensionale Muster belegen.

Acknowledgments

First of all, I like to thank my advisor Helmut Alt for introducing me to the interesting topic of geometric shape matching and for giving me the opportunity to do research and write this thesis. I learned a lot from Helmut, and I am grateful to him for supporting me and discussing research problems.

I thank Otfried Cheong for agreeing to co-referee this thesis and for hosting me during a memorable stay at KAIST in Daejeon, Korea, in August 2008.

The Graduiertenkolleg *Methods for Discrete Structures* in Berlin provided a lively research environment. I am thankful for this and also, of course, for the financial support.

I have very much enjoyed being a member of the workgroup *Theoretical Computer Science* at Freie Universität Berlin. Many thanks go to my colleagues for informative, yet hilarious discussions in the legendary coffee rounds and for generally being great company. The frogday will be remembered! Particularly, I like to thank Frank Hoffmann for his encouragement.

The winter 2009 at the Centre de Recerca Matemàtica in Bellaterra, Spain, was filled with plenty of exciting mathematics and abundant fun. I am grateful to the organizers and participants of the DocCourse Combinatorics and Geometry for this wonderful time.

I thank my coauthors Otfried Cheong, Erik Demaine, Sándor Fekete, Hyo-Sil Kim, Fabian Stehn, Joachim Gudmundsson, Günter Rote, Nils Schweer, and Mariano Zelke. Although our joint research is not included in this thesis, working together with you has contributed to making my time as a doctoral student interesting and enjoyable. I especially thank Helmut Alt and Ludmila Scharf for doing research together on the topic of my thesis.

I am deeply grateful to my friends Kristina Alex, Thomas Berndt, Felix Breuer, Mareike Massow, Fabian Stehn, Julia Westendorf and to my sister Tamara Schymura for supporting me and believing in me. I very much thank Mareike for carefully proofreading this thesis. I am incredibly thankful to Matthias Henze for his valuable comments on this thesis, and for everything else.

Contents

Abstract **vii**

Zusammenfassung **ix**

1. Introduction **1**
1.1. General idea of our approach 3
1.2. Related work 5
1.3. Overview 6

2. Shape models **9**
2.1. The Hausdorff measure and related results 9
2.2. The oracle model 11
2.3. Unions of simplices 12
2.4. Fatness and the isoperimetric quotient 14

3. The objective function **21**
3.1. Known results 21
3.2. Upper bounds on the volume of the symmetric difference 24
3.3. Proof of the upper bounds 26
3.4. Lower bounds 31

4. The probabilistic toolbox **35**
4.1. Transformation of density functions 36
4.2. Uniform approximation of the density function by the empirical measure 36

5. Probabilistic matching under translations in arbitrary dimension **43**
5.1. The algorithm 44
5.2. Correctness proof 47
5.3. Adaptive Sampling 50

6. Probabilistic matching of planar regions under rigid motions **59**
6.1. Using the algorithm for translations 60
6.2. Vote generation with random angle 62
6.3. Vote generation with 3+1 points 66

7. Evaluating the probabilistic matching approach in 2D **73**
7.1. Comparison with a simple, deterministic approach 73
7.2. Similarities 81

7.3. Experimental results 83

8. Probabilistic matching under rigid motions in arbitrary dimension 99

8.1. The rotation group 99
8.2. The space of rigid motions 104
8.3. Using the algorithm for translations 105
8.4. Sampling of the transformation space 107
8.5. Improving the results in 3D 110

A. The volume of small balls in the rotation group 117

Bibliography 121

1. Introduction

The main question in shape matching is: Given two objects A and B, how similar are they? Methods for solving this question have numerous applications in various fields of computer science. In computer vision, the problem of *object recognition* is essential. A system knows a model of an object and has to find the object in a scene. In *image retrieval*, a large database of images is given, and the problem is to find images that are similar to a query image. Drug design, or more specific *molecular docking*, is a shape matching problem, which occurs in computational biology. Another example is the *registration of medical images* in computer-aided surgical navigation systems.

These different applications give rise to many different problem types. Examples of fundamental problems are the *matching*, *simplification*, *morphing*, and *clustering* of shapes [4]. This thesis concentrates on the shape matching problem, which is an important theoretical problem. Shapes can be compared with respect to different properties, for example color, texture, or feature points; we are interested in the problem of matching shapes based on their geometry. For a survey on shape matching, as dealt with in computational geometry, see [4]. For surveys describing a broader set of approaches that are still close to computational geometry, we refer the reader to [38, 53].

Given a class of shapes S, a set of transformations T, mapping shapes onto shapes, and a distance measure d on S, the *matching problem* with respect to S, T, and d is the following: Given two shapes A, $B \in S$, compute a transformation $t^* \in T$ that minimizes $d(t^*(A), B)$.

Many different notions of shapes, also called patterns, are studied in the literature. Examples in two dimensions are finite point sets, polygonal curves, polygons, sets of line segments and geometric graphs. In the case of polygons, shapes can be restricted to be convex, to be simple, or polygonal regions can be allowed, which can have holes and are not necessarily connected. Studied classes of shapes in dimension $d \geq 3$ are finite point sets, convex polyhedra, and unions of simplices. We consider solid shapes, i.e. full-dimensional shapes, and assume that they are given by an oracle that generates uniformly distributed random points from the shapes. This is a very rich class of shapes that contains the class of finite unions of simplices as a subclass.

As transformations, usually a subset of the affine transformations is considered. The most basic, yet important, example is the set of translations. If instead rotation is allowed, the resulting transformations are called rigid motions. Examples that involve scaling are homotheties (translation and scaling) and similarities (translation, rotation and scaling). In this thesis, we study matching under translations and rigid motions.

There is a variety of distance measures, some are restricted to certain classes of shapes. The simplest example is the discrete metric that equals 0 or 1, depending on the fact if the shapes are equal or not. The matching problem with respect to the discrete metric yields the exact congruence problem, that is the question if there is a transformation such that the transformed shape A equals B. If instead the transformation group is the trivial group containing only the identity transformation, then the matching problem becomes the problem of computing the distance of A and B.

For solid polygons, the area of the symmetric difference is a natural distance measure. The *symmetric difference* of two sets is the set of points that belong to exactly one of the sets. Throughout this work, the volume of the symmetric difference is used as the distance measure for the matching problem. On the class of compact sets in $\mathbb{R}^d$ that equal the closure of their interior, the volume of the symmetric difference is a metric. Moreover, the volume of the symmetric difference is a deformation, crack, blur and noise robust distance measure [31]. Chapter 3 says more about this.

Since we consider the volume of the symmetric difference as the distance measure for the matching problem, only the full-dimensional part of the shapes is important. We indicate this by speaking about *solid shapes*, in contrast to lower-dimensional shapes, as for example curves or finite point sets. In the literature, sets are sometimes defined to be solid if they equal the closure of their interior.

Under volume-preserving transformations, minimizing the volume of the symmetric difference is the same as maximizing the volume of overlap. When speaking about approximations with respect to a relative error, one has to be careful because they are not the same anymore. We present algorithms that approximately maximize the volume of overlap. We denote the volume (Lebesgue measure) by $|\cdot|$ and the symmetric difference by $\triangle$.

Since no solutions of low time complexity are known for maximizing the volume of overlap of sets of simplices under translations and rigid motions, approximation algorithms are interesting. We consider the following approximate versions of the matching problem. For an *absolute error approximation* of the maximal volume of overlap, two shapes A and B and an error tolerance $\varepsilon \in (0, 1)$ are given as input. The problem is to compute a transformation t^* such that $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all transformations t, in particular for transformations maximizing the volume of overlap.

We use the error term $\varepsilon|A|$, and not ε, for two reasons. First we want ε to lie in the range $(0, 1)$ for all shapes A and B in order to make it comparable. Second the inequality should be invariant under scaling of A and B since scaling A and B with the same factor does not change the quality of the result. Of course, there is nothing special about $|A|$ in comparison to $|B|$. One could also regard $\varepsilon|B|$ as the error term. Since the problem is symmetric in A and B, it is no restriction to use the error term $\varepsilon|A|$. We do not use an error term involving both volumes, as

$\varepsilon\frac{|A|+|B|}{2}$ or $\varepsilon \min\{|A|,|B|\}$, since this would later complicate the interpretation of certain bounds and eventually conceal their meaning.

If additionally an allowed probability of failure $p \in (0,1)$ is given as input, the problem is to compute a transformation t^* such that, with probability at least $1-p$, we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all transformations t. We call this a *probabilistic absolute error approximation.*

For shapes A, B and $\varepsilon \in (0,1)$, a transformation $\hat{t}$ is a $(1-\varepsilon)$-approximation, or *relative error approximation*, if $|\hat{t}(A) \cap B| \geq (1-\varepsilon)|t(A) \cap B|$, or equivalently $|t(A) \cap B| - |\hat{t}(A) \cap B| \leq \varepsilon|t(A) \cap B|$, for all $t \in \mathbb{R}^d$. This is a stronger requirement than the condition for the absolute error approximation.

A *probabilistic relative error approximation* also obtains two shapes A, B, an error tolerance $\varepsilon \in (0,1)$ and an allowed probability of failure $p \in (0,1)$ as input and computes a transformation $\hat{t}$ that is a $(1-\varepsilon)$-approximation of the maximal volume of overlap with probability at least $1-p$.

This thesis describes algorithms computing probabilistic absolute error approximations for translations and rigid motions in arbitrary dimension. If the maximal volume of overlap is at least as large as a constant fraction of the volume of A, say $\kappa|A|$, then a probabilistic absolute error approximation with error bound $\varepsilon\kappa$ yields a probabilistic relative error approximation with error bound ε.

Maximizing the volume of overlap of two sets is an interesting mathematical problem in its own, see [23]. To successfully describe and match shapes in practice, also the boundaries of the shapes have to be considered. As Mumford [44] says "a plane shape $S \subset \mathbb{R}^2$ has a 1-dimensional side given by features of its boundary $C = \partial S$; and a 2-dimensional side given by its interior. No successful theory of shape description can ignore one or the other".

We study very simple shape matching algorithms that a priori are heuristics. The main focus of this thesis is to analyze these algorithms and prove that they maximize the volume of overlap approximately. The approach is mainly of theoretical interest. Still, the algorithms are so easy that they can easily be implemented and the runtimes are reasonable. In contrast to other approaches, the volume of overlap is not computed by the algorithms, but the solution is found by a voting scheme. No complicated data structures are needed. We show that our approach is implementable by giving experimental results of our test implementation for 2-dimensional shapes.

For the analysis, we use the real RAM model, including the square root operation, and assume that random reals that are uniformly distributed on a fixed interval can be generated in constant time.

1.1. General idea of our approach

We study algorithms for matching solid shapes in $\mathbb{R}^d$ under translations and rigid motions. The algorithms are instances of the following algorithmic scheme, or at least are related to it.

The main idea is quite simple. Given two shapes A and B, repeat the following random experiment very often, say N times: Select random point samples of appropriate size from each shape and compute a transformation that maps the point sample of one shape to the sample of the other shape. Store this transformation, called a *vote*. In each step, we extend the collection of votes by one. Clusters of votes in the transformation space indicate transformations that map large parts of the shapes onto each other. We determine a densest cluster and output its center. We refer to this general algorithmic scheme as the *probabilistic matching approach*.

Per se, this algorithmic scheme is a heuristic approach; it captures the intuitive notion of matching. Transformations whose δ-neighborhoods contain many votes should be "good" transformations since they map many points from A onto points from B.

To define clusters, we choose a metric d on the transformation space. A *δ-neighborhood* of a transformation t is the ball of radius δ around t with respect to the metric d. A *densest cluster* of votes is defined to be a δ-neighborhood of a vote that contains the largest number of votes for some fixed parameter δ. Thus along with the shapes A and B we have two additional parameters: the number of random experiments N and the clustering size δ.

Let us look at the case where in each random experiment one point in each shape is generated, say $a \in A$ and $b \in B$. This uniquely determines a translation that maps a onto b. It turns out that in this case the algorithm provably approximates the maximal volume of overlap if δ is small and N is large enough. Figure 1.1 illustrates this idea for the case of translations in the plane.

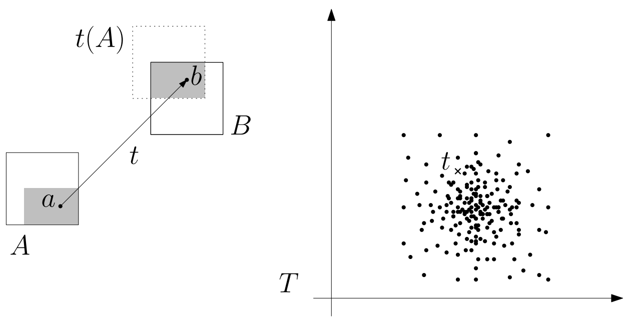

Figure 1.1.: We compare two copies of a square under translations. The area of overlap of $t(A)$ and B corresponds to the "chance" of choosing a point pair $(x, y) \in A \times B$ such that $y - x = t$. The point t is marked with a cross in the translation space.

Hence we choose more meaningful parameters as input for our algorithms. Along with two shapes A, B, the algorithms obtain an error tolerance ε and an allowed probability of failure p, both in $(0, 1)$, as input. In the algorithm, we compute N and δ such that, with probability at least $1 - p$, the output has an absolute error of at most $\varepsilon|A|$.

1.2. Related work

The probabilistic matching approach has been studied before for matching curves in the plane. Ludmila Scharf studied it in her thesis [45] for matching curves in the plane under translations, rigid motions, similarities and further subsets of affine transformations. Sven Scholz studied in his thesis [46] heuristics based on the algorithmic scheme for the retrieval of trademark images. See also [6, 7, 8] for instances of the algorithmic scheme applied to polygonal curves in the plane.

Our approach is related to the Hough transform and other voting schemes that are applied to shape matching. We refer the reader to [45] for a discussion of the relation of the probabilistic matching approach and these methods.

Dimension 2. In two dimensions, the problem of maximizing the area of overlap of two polygons with n and m vertices, say $m \leq n$, is studied for different settings. For convex polygons, de Berg et al. [23] give a solution in $O(n \log n)$ time for translations. Alt et al. [3] give a linear time constant factor approximation for convex polygons under translations and homotheties. Ahn et al. [2] give a $(1-\varepsilon)$-approximation in time $O((1/\varepsilon)\log n + (1/\varepsilon)\log(1/\varepsilon))$ for translations and $O((1/\varepsilon)\log n + (1/\varepsilon^2)\log(1/\varepsilon))$ for rigid motions, again for the case of convex polygons.

For simple polygons, the function that maps a translation t to the area of overlap of $t(A)$ and B has combinatorial complexity $O(n^2m^2)$ and can be maximized within the same time [43]. See Section 3.1 for the definition of combinatorial complexity.

For two polygonal regions, Cheong et al. [20] give an absolute error approximation that works with high probability for matching under translations in $O(m + (n^2/\varepsilon^4)\log^2 n)$ and matching under rigid motions in $O(m + (n^3/\varepsilon^8)\log^5 n)$. The runtime for rigid motions given in the paper is smaller because of a calculation error in the final derivation of the time bound, as was noted by Vigneron [54].

For maximizing the area of overlap of two simple polygons under rigid motions, no exact polynomial time algorithm is known. Vigneron [54] gives an FPTAS with relative error ε that runs in time $O((n^6/\varepsilon^3)\log^4(n/\varepsilon)\beta(n/\varepsilon))$ where β is a very slowly growing function related to the inverse Ackermann function. He also gives an FPTAS for minimizing the area of the symmetric difference of two polygons in the plane under rigid motions that has time complexity $O(n^{17+\varepsilon'}+(n^{15}/\varepsilon^3)\log^4(n/\varepsilon)\beta(n/\varepsilon)\beta(n))$ for any $\varepsilon' > 0$.

Arbitrary dimension. Hagedoorn determines the complexity of the function that maps a transformation t to the volume of $t(A)\triangle B$ for translations, rigid motions and other subgroups of affine transformations. He shows that the combinatorial complexity of the function is $\Theta((nm)^r)$ in the worst case when A is the union of n simplices, B is the union of m simplices, and r is the degree of freedom of the group. For example, $r = d$ for translations in $\mathbb{R}^d$ and $r = \frac{d^2+d}{2}$ for rigid motions in d dimensions.

Ahn et al. [1] minimize the volume of the symmetric difference of two convex polyhedra under translations in expected running time $O(n^{d+1-\frac{3}{d}} \log^{d+1} n)$ where n is the total number of simplices necessary in a triangulation of the input polyhedra.

In $d \geq 3$ dimensions, Vigneron [54] gives the only result for the matching of non-convex polyhedra under rigid motions so far. He describes FPTASs for maximizing the volume of overlap, as well as for minimizing the volume of the symmetric difference. 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 $O((\frac{nm}{\varepsilon})^{\frac{d^2}{2}+\frac{d}{2}+1}(\log \frac{nm}{\varepsilon})^{\frac{d^2}{2}+\frac{d}{2}+1})$. The time bound for approximating the minimal volume of the symmetric difference in dimension $d \geq 3$ is not given in the paper.

1.3. Overview

Chapter 2: Shape models

We introduce an oracle shape model, which defines the most general notion of solid shape for which our probabilistic matching approach works. We show that unions of simplices are a special case. We define the isoperimetric quotient and discuss a new notion of fatness.

Chapter 3: The objective function

We collect known facts about the function that maps a translation or rigid motion to the volume of overlap of two fixed shapes A and B, when A is transformed by the translation or rigid motion.

We bound the volume of the symmetric difference of $t(A)$ and B from above, when t is a translation or rigid motion. We use this bound to prove an inequality that implies that the objective function is Lipschitz continuous for many metrics on the transformation space. We determine a lower bound on the volume of the symmetric difference of $t(A)$ and B for the case of translations.

Chapter 4: The probabilistic toolbox

We review notions and theorems from probability theory, in particular we cite a theorem about the density function of transformed random vectors. We cite facts about the empirical measure. Devroye and Lugosi's book [25] is of great help. We prove the main theorem that we use for analyzing the convergence rate of the probabilistic algorithms.

Chapter 5: Probabilistic matching under translations in arbitrary dimension

We present the simplest variant of the probabilistic matching approach for solid shapes, which performs matching under translations in arbitrary dimension. We determine the density function on the translation space and prove bounds on the clustering size and the number of random experiments, guaranteeing that the algorithm computes a probabilistic absolute error approximation of the maximal volume of overlap. We analyze the runtime of the algorithm. We improve the bounds and the runtime by describing an adaptive algorithm that computes a matching by first computing a candidate region in the translation space and then searching in this region.

Chapter 6: Probabilistic matching of planar regions under rigid motions

We explore three possibilities to apply the probabilistic matching approach to rigid motions in the plane. First we show how to use the algorithm for matching under translations for matching under rigid motions by discretizing the angle space. Second we apply the probabilistic matching approach in two variants, which differ in the generation of votes. We prove that all three algorithms compute a probabilistic absolute error approximation of the area of overlap. For this, we give bounds on the clustering size and the sample size that guarantee approximation.

Chapter 7: Evaluating the probabilistic matching approach in 2D

We evaluate the probabilistic matching in the plane. We introduce a simple deterministic algorithm, related to the probabilistic approach, for matching under translations and rigid motions in the plane. We compare the probabilistic with the deterministic approach. We study whether the probabilistic matching approach can be used for matching solid shapes under similarities. We show experimental results of a preliminary implementation for matching 2-dimensional shapes under translations, rigid motions and similarities.

Chapter 8: Probabilistic matching under rigid motions in arbitrary dimension

We study probabilistic matching under rigid motions in arbitrary dimension by generalizing two of the algorithms from Chapter 6 to arbitrary dimension. We prove that both algorithms compute a probabilistic absolute error approximation of the maximal volume of overlap. We have to have a closer look at the rotation group $SO(d)$ to be able to do so. We improve the results in three dimensions by using a different representation of rotations.

Appendix A: The volume of small balls in the rotation group

We supply the proof of a theorem that we use in Chapter 8. The theorem describes the volume of balls in the rotation group asymptotically, as the radius tends to 0. For the proof, we introduce a number of definitions and results from measure theory.

2. Shape models

In the literature, a variety of shape models are studied, including finite (weighted) point sets, polygonal chains, (convex) polyhedra, unions of simplices, unions of (pairwise disjoint) balls, and plane geometric graphs. In this chapter, we describe the shape models for which we present matching algorithms in the next chapters. In Section 2.2, we introduce the oracle model, which is a very general shape model. In Section 2.3, we describe an important subclass, the class of finite unions of pairwise disjoint simplices. We review existing definitions of fatness and introduce yet another definition of fatness and compare it to existing definitions in Section 2.4.

Before presenting the shape models, we introduce some notions and results from measure theory, most importantly the definition of the Hausdorff measure.

2.1. The Hausdorff measure and related results

We assign a volume to measurable subsets A of $\mathbb{R}^d$ by the usual Lebesgue measure and denote it by $\mathcal{L}^d(A)$. Denote by ω_d the volume of the Euclidean unit ball in $\mathbb{R}^d$. For a set $A \subset \mathbb{R}^d$, the boundary ∂A is defined as the set of points that are in its closure $\mathrm{cl}(A)$, but not in its interior $\mathrm{int}(A)$. We want to measure not only the volume of sets, but also the surface area, or $(d-1)$-dimensional volume, of their boundaries.

We define the Hausdorff measure, which generalizes the Lebesgue measure. Moreover, with the Hausdorff measure lower-dimensional sets can be measured. We use it mainly because it allows us to measure boundaries smartly. The following definitions of the Hausdorff measure, the spherical measure, their properties, and rectifiability can be found in [28].

For $A \subset \mathbb{R}^d$, $0 \le k \le d$ and $\delta > 0$, let $\mathcal{H}^k_\delta(A)$ be the *size δ approximating k-dimensional Hausdorff measure* of A that is defined as follows

$$\mathcal{H}^k_\delta(A) = \omega_k\, 2^{-k} \inf \left\{ \sum_{j=0}^{\infty} (\operatorname{diam} B_j)^k \;\middle|\; \forall j \in \mathbb{N} : B_j \subset \mathbb{R}^d, \operatorname{diam} B_j \le \delta, A \subseteq \bigcup_{j=0}^{\infty} B_j \right\}.$$

We abbreviate $\delta \to 0$ and $\delta > 0$ by $\delta \to +0$. The k-dimensional Hausdorff measure of A is then defined as

$$\mathcal{H}^k(A) = \lim_{\delta \to +0} \mathcal{H}^k_\delta(A).$$

The limit exists and equals the supremum because $\mathcal{H}^k_\delta(A) \le \mathcal{H}^k_\eta(A)$ for $\eta \le \delta$, but it might equal $+\infty$.

A set $A \subset \mathbb{R}^d$ is called *$\mathcal{H}^k$-measurable* if it satisfies the Carathéodory property, meaning that for all sets $B \subset \mathbb{R}^d$, we have $\mathcal{H}^k(A) = \mathcal{H}^k(A \cap B) + \mathcal{H}^k(A \setminus B)$. The

Hausdorff measure $\mathcal{H}^k$ is defined for all $A \subset \mathbb{R}^d$, but, of course, it is more meaningful for $\mathcal{H}^k$-measurable sets.

The 0-dimensional Hausdorff measure equals the counting measure that gives the number of elements for finite sets and $+\infty$ for infinite sets. The $(d-1)$-dimensional Hausdorff measure of sufficiently nice $(d-1)$-dimensional sets, for example smooth manifolds, equals the surface area. In this context, "sufficiently nice" means $(\mathcal{H}^{d-1}, d-1)$-rectifiable, a notion which we define below. We measure the boundaries of sets in $\mathbb{R}^d$ by the $(d-1)$-dimensional Hausdorff measure. For Lebesgue measurable sets in $\mathbb{R}^d$, the d-dimensional Hausdorff measure coincides with the Lebesgue measure. Recall that the Lebesgue measure is invariant under rotation and translation, and note that the Hausdorff measure is invariant under rotation and translation, too.

The sets B_j in the definition of the approximating Hausdorff measure can be assumed to be convex because taking the convex hull does not increase the diameter of a set. Since convex sets are Lebesgue measurable [37], the sets B_j can also be assumed to be Lebesgue measurable.

If we restrict the coverings $\{B_j\}_{j\geq 0}$ in the definition of the approximating Hausdorff measure to be families of balls, then the resulting measure is called *spherical measure*. For $A \subset \mathbb{R}^d$, we denote the k-dimensional spherical measure of A by $\mathcal{S}^k(A)$. Since the choice of coverings is restricted, we have $\mathcal{H}^k(A) \leq \mathcal{S}^k(A)$.

Jung's theorem, which we cite from [28], gives a sharp bound on the radius of the smallest enclosing ball of a set of a fixed diameter. For regular, full-dimensional simplices, equality holds.

Theorem 2.1 (Jung's theorem)**.** *If $S \subset \mathbb{R}^d$ and $\operatorname{diam}(S) \in (0, +\infty)$, then S is contained in a unique closed ball with minimal radius, which does not exceed $\sqrt{\frac{d}{2d+2}}\operatorname{diam}(S)$.*

From this, it follows that we have $\mathcal{S}^k(A) \leq \left(\frac{2d}{d+1}\right)^{k/2}\mathcal{H}^k(A)$ for all $A \subset \mathbb{R}^d$. In general, the Hausdorff measure and the spherical measure are not equal, but for $(\mathcal{H}^k, k)$-rectifiable subsets of $\mathbb{R}^d$ they agree [28, Theorem 3.2.26]. We define the notion of rectifiability now.

A subset E of a metric space X is called *k-rectifiable* if there exists a Lipschitz continuous function that maps some bounded subset of $\mathbb{R}^k$ onto E. A union of countably many k-rectifiable sets is called *countably k-rectifiable*. E is called *countably (μ, k)-rectifiable* if μ is a measure defined on E and there is a countably k-rectifiable set that contains μ-almost all of E. If, additionally, $\mu(E) < +\infty$, then E is called *(μ, k)-rectifiable.*

To close this paragraph, we cite two inequalities [28]. The isodiametric inequality says that, among the Lebesgue measurable sets of a fixed diameter, Euclidean balls have the largest volume.

Theorem 2.2 (Isodiametric Inequality)**.** *If $\emptyset \neq S \subset \mathbb{R}^d$ is Lebesgue measurable, then $\mathcal{L}^d(S) \leq \omega_d 2^{-d}\operatorname{diam}(S)^d$.*

It can be also read as: The quotient $\frac{\mathrm{diam}(S)^d}{\mathcal{L}^d(S)}$ is at least as large as $2^d/\omega_d$, and equality holds for Euclidean balls. Therefore the quotient $\frac{\mathrm{diam}(S)^d}{\mathcal{L}^d(S)}$ is a way to measure how similar S is to ball. A different way to do so is provided by the isoperimetric inequality, which we will use in Section 2.4 for a new definition of fatness.

Theorem 2.3 (Isoperimetric Inequality [28, Theorem 3.2.43 and 4.5.9(31)]). *Let $A \subset \mathbb{R}^d$ be a Lebesgue measurable set such that $\mathcal{L}^d(\mathrm{cl}(A)) < +\infty$. Then*

$$d^d\omega_d \leq \frac{(\mathcal{H}^{d-1}(\partial A))^d}{(\mathcal{L}^d(A))^{d-1}},$$

and equality holds if and only if A is a d-dimensional Euclidean ball.

For giving definitions precisely and for being able to cite theorems from measure theory correctly, it has been necessary to explicitly indicate the dimensions of the measures, as in $\mathcal{H}^k$ and $\mathcal{L}^d$. From now on, we always have sets $A \subset \mathbb{R}^d$ and want to denote their volume (Lebesgue measure) and their surface area. For the shortness and clearness of the presentation, we return to denoting the volume by the slightly sloppy notation $|A|$. The surface area of A, which we measure by the $(d-1)$-dimensional Hausdorff measure of the boundary of A, will be denoted by $|\partial A|$. We will only use the notations $\mathcal{H}^k$ and $\mathcal{L}^d$ if the dimension of the measure is important.

The *isoperimetric quotient* $\frac{(\mathcal{H}^{d-1}(\partial A))^d}{(\mathcal{L}^d(A))^{d-1}}$ from the isoperimetric inequality thus is written as $\frac{|\partial A|^d}{|A|^{d-1}}$. It plays an important role in the analysis of our algorithms.

2.2. The oracle model

We always assume shapes to be Lebesgue measurable subsets of $\mathbb{R}^d$ that have positive finite volume, and whose boundary is $\mathcal{H}^{d-1}$-measurable and has positive finite $(d-1)$-dimensional volume. Furthermore, we assume the boundary to be $(\mathcal{H}^{d-1}, d-1)$-rectifiable. W.l.o.g. we assume that the shapes contain the origin. The dimension d is regarded to be constant.

Since sets of Lebesgue measure 0 do not make any difference for the volume of overlap, for some proofs, we assume the shapes to be Borel sets. Being a Borel set is a stronger requirement than being Lebesgue measurable, but every Lebesgue measurable set differs from some Borel set only by a measure 0 set.

The most general shapes we consider additionally satisfy the following assumptions:

- First and most importantly, we assume that A is *given by an oracle.* By this we mean that we can generate N points uniformly distributed from shape A in time $T(N, A)$. For the shortness of presentation, we will write $T(N)$ although the time may depend on the shape.

- For one of the two input shapes of our matching problem, say A, we know an upper bound K_A on the isoperimetric quotient $\frac{|\partial A|^d}{|A|^{d-1}}$. We will study the meaning of this parameter and its properties in Section 2.4.
- We are given lower bounds m_A and m_B on $|A|$ and $|B|$. Furthermore we are given an upper bound M_B on $|B|$.
- For matching under rigid motions, we assume to have an upper bound Δ_A on the diameter of A. Because of the lower bound on $|A|$, this gives us also an upper bound D_A on the quotient $\frac{\mathrm{diam}(A)^d}{|A|}$. In Section 5.3, we also use an upper bound Δ_B on the diameter of B.

We design and analyze our algorithms under these assumptions. A common representation of shapes is as finite unions of simplices, which we introduce in the next section. We show that finite unions of simplices are given by an oracle and fulfill the above assumptions.

For the algorithms in Sections 5.3 and 6.3, we not only assume that we can generate N uniformly distributed random points from B in time $T_q(N)$, but also can answer N membership queries of the type "$b \in B$?".

2.3. Unions of simplices

A d-dimensional *simplex* is defined as the convex hull of $d+1$ affinely independent points in $\mathbb{R}^k$ for some $k \geq d$. A common representation of shapes are finite unions of d-dimensional simplices in $\mathbb{R}^d$. The simplices of a shape are assumed to have pairwise disjoint interiors. Hence, in two dimensions, a shape is a polygonal region, which can have holes and is not necessarily connected. An example is depicted in Figure 2.1. A shape is given as a set of n simplices, but it is understood as the union of these simplices. This class of shapes is very rich; it contains the homogeneous geometric simplicial complexes as a subclass. Keep in mind that shapes do not need to be connected.

Next we check that this shape model satisfies our requirements, described in the previous section. So, let A be a finite set of d-dimensional simplices in $\mathbb{R}^d$ that have pairwise disjoint interiors. Let n be the number of simplices of A. First, we show how to compute bounds Δ_A and K_A in time $O(n)$.

The volume of a simplex Δ with vertices $v_0, v_1, \ldots, v_d$ in $\mathbb{R}^d$ equals

$$|\Delta| = \tfrac{1}{d!}|\det(v_1 - v_0, v_2 - v_0, \ldots, v_d - v_0)|,$$

so computing the volume is essentially computing the determinant of a d-dimensional matrix. This can be done in $O(d^3)$ time, or, by more intricate methods, in the time needed for matrix multiplication. Since we regard d as constant, we do not go into details here. As the simplices of A have pairwise disjoint interiors, we simply sum up the volumes of the simplices to compute $|A|$ in time $O(n)$.

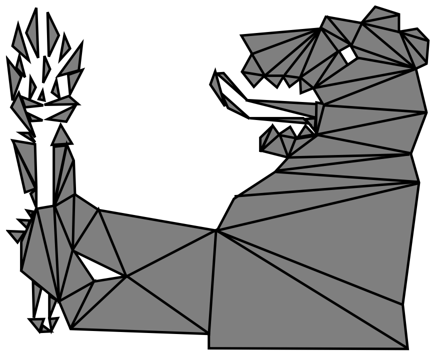

Figure 2.1.: An example shape in the plane given as the union of triangles that have pairwise disjoint interiors.

In linear time, the maximal distance D of a vertex from the origin can be computed. W.l.o.g. the shape contains the origin. Then $\frac{1}{2}\operatorname{diam}(A) \leq D \leq \operatorname{diam}(A)$, and we take $\Delta_A = 2D$ as upper bound on the diameter of A.

We do not know how to compute $|\partial A|$ in time linear in n since A is given as a set of simplices. But we can compute an upper bound on the size of the boundary in linear time by adding all $(d-1)$-dimensional volumes of the boundaries of the simplices. The boundary of a simplex is a collection of lower-dimensional simplices, its facets. Obviously, this bound can be bad as many of the facets may lie in the interior of the shape, but it can be sharp as it is matched for collections of pairwise disjoint simplices.

In Section 2.4, we will show that the isoperimetric quotient of a shape A is at most a constant multiple of n if A is the union of perimeter-fat simplices $\Delta_1, \dots, \Delta_n$ and $\max_{1\leq i,j\leq n} \frac{|\Delta_j|}{|\Delta_i|}$ is bounded by a constant. Definition 2.5 introduces the notion of perimeter-fatness.

To generate uniformly distributed random points from A, first draw a simplex from A with probability proportional to the volume of the simplex. Then we draw a point from this simplex uniformly at random (u.a.r.).

For drawing a simplex, we use the alias method [55]. Let A be the union of the simplices $\Delta_1, \dots, \Delta_n$ that have pairwise disjoint interiors. Let $p_i := \frac{|\Delta_i|}{|A|}$. The vector $(p_1, \dots, p_n)$ describes the discrete probability distribution according to which we have to pick a random point from simplex Δ_i. The alias method can be implemented such that after $O(n)$ preprocessing time, we can sample random points from $(p_1, \dots, p_n)$ in time $O(1)$ [48].

Generating a point from a d-dimensional simplex u.a.r. can be done in $O(d \log d)$ time by generating d numbers in $[0,1]$ u.a.r., sorting them $0 \leq x_1 < \cdots < x_d \leq 1$ and taking the spacings $x_1, x_2 - x_1, \dots, x_d - x_{d-1}, 1 - x_d$ as barycentric coordinates of the point in the simplex. This method generates a point in the simplex u.a.r. [24]. Thus generating N random points in A and B takes $T(N) = O(n+N)$ time, and we can generate the points one at a time.

When B is the finite union of at most n triangles in the plane that have pairwise disjoint interiors, we can answer N membership queries in time $T_q(N) = O\big((n+N)\log n\big)$ by preprocessing B in time $O(n \log n)$ to obtain a standard point location data structure and then answering membership queries in time $O(\log n)$. In dimension $d \geq 3$, we do not know how to answer membership queries for unions of at most n simplices faster than linear, so we have $T_q(N) = O(nN)$.

We summarize these properties in a lemma.

Lemma 2.4. *Let A be the union of n d-dimensional simplices in $\mathbb{R}^d$ that have pairwise disjoint interiors. After $O(n)$ preprocessing time, we can generate a point from A u.a.r. in time $O(1)$.*

We can answer a membership query in time $O(n)$. If $d = 2$, after $O(n \log n)$ preprocessing time, we can answer membership queries in time $O(\log n)$.

Higher-dimensional boxes behave similarly, and can be used instead of simplices; the computations are even easier for rectangles.

Unions of unit balls

A related shape model is given by finite unions of Euclidean unit balls that have pairwise disjoint interiors. This model is more restricted since the isoperimetric quotient of a Euclidean ball equals $d^d\omega_d$ by Theorem 2.3, implying that balls are fat. Fatness will be discussed in the next section. The isoperimetric quotient of the union of n balls of a fixed radius equals $nd^d\omega_d$ by Proposition 2.7.

A bound on the diameter can be computed in linear time. A random point can be computed in constant time if the centers of the balls are stored in an array, which gives $T(N, A) = N$.

Instead of the union of unit balls, we can take the union of congruent copies of any fixed convex set for which we can generate a random point in constant time and compute the volume and an upper bound on the $(d-1)$-dimensional volume of the boundary, for example a cube or a fixed simplex.

2.4. Fatness and the isoperimetric quotient

Many algorithms for geometric problems have better bounds on the time complexity if the input objects are not too “skinny”. In many applications, “realistic” input objects are not “skinny”, so it makes sense to analyze algorithms for these problems under this assumption. In order to capture not being too “skinny”, the notion of fatness has been proposed.

In our case, the bounds on the time complexity of our algorithms depend on the isoperimetric quotient $\frac{|\partial A|^d}{|A|^{d-1}}$. Demanding that this parameter is bounded from above by a constant is a fatness condition, which we take as definition of fatness.

Definition 2.5 (perimeter-fatness)**.** *A shape $A \subset \mathbb{R}^d$ is called* perimeter-fat with parameter κ *if* $\frac{|\partial A|^d}{|A|^{d-1}} \leq \kappa$*.*

This definition of fatness restricts the ratio of "skin" to volume. To make this ratio invariant under scaling, it is "homogenized" by taking the $(d-1)$-dimensional volume of the boundary to the d-th power and the volume to the $(d-1)$-th power. The isoperimetric inequality (Theorem 2.3) gives a lower bound on this ratio. Before discussing further properties of this notion, we review some important notions of fatness, which have been studied in the literature.

A convex object is defined to be *ball-fat* with parameter C if the ratio of the radius of the smallest enclosing and the radius of the largest inscribed ball is at most C. Most of the many different definitions of fatness that are present in the literature are equivalent to this definition for convex shapes, but differ for non-convex objects.

The notions of fatness depend on one (or two) parameters. By saying that one notion of fatness with parameter p_1 *is stronger* than another notion of fatness or *implies* another notion of fatness, we mean that there is a parameter p_2 for the latter notion that only depends on p_1 such that each object that is fat with respect to the first definition with parameter p_1 is fat with respect to the other definition with parameter p_2.

Fat shapes in the plane. Efrat [27] introduced *(α, β)-covered* objects, which are a generalization of convex ball-fat objects to the class of simply connected sets in the plane. He showed that the combinatorial complexity of the union of (α, β)-covered objects is better than quadratic. An object $A \subseteq \mathbb{R}^2$ is (α, β)-covered if it is simply connected and, for every point p on the boundary of A, there is a triangle that (1) is contained in A, (2) has p as a vertex, (3) each of its angles is at least as large as α, and (4) each of its side lengths is at least as large as $\beta \operatorname{diam}(A)$.

The most common definition of fatness is probably by van der Stappen, Halperin and Overmars [51]. A shape A is called *classically fat with parameter γ* if for all disks D whose center is contained in A and that do not fully contain A, the area $|A \cap D| \geq \gamma |D|$. In [51], this notion of fatness is called γ-fatness.

De Berg [21] invented locally fat objects as a relaxation of being (α, β)-covered while still having better than quadratic union complexity. A shape is *locally fat with parameter γ* if the intersection $A \cap D$ in the previous definition is replaced by the connected component of $A \cap D$ that contains the center of D. Obviously, being locally fat with parameter γ implies being classically fat with parameter γ, and the converse is not true.

A weak notion of fatness for sets in the plane is used by Mitchell [42] and many others. A set $A \subset \mathbb{R}^2$ is called diameter-fat if its area is at least as large as some constant fraction of the squared diameter. This generalizes immediately to higher dimensions: a set $A \subset \mathbb{R}^d$ is called *diameter-fat with parameter C* if $|A| \geq C \operatorname{diam}(A)^d$.

According to our definition, a shape $A \subset \mathbb{R}^2$ is perimeter-fat if the area of A is at least as large as a constant fraction of the squared length of the boundary. We compare this notion of fatness with the other mentioned notions of fatness.

Figure 2.2 summarizes the relations between the different definitions for the case of simply connected sets in the plane.

For connected shapes in the plane, perimeter-fatness is stronger than diameter-fatness since the length of ∂A is at least as large as $2\,\mathrm{diam}(A)$. Bose et al. [18] show that, for (α, β)-covered objects, the length of the boundary is linear in the diameter, which means that perimeter-fatness is weaker than being (α, β)-covered.

They also prove that locally fat shapes can have an arbitrarily long perimeter, which implies that perimeter-fatness is not weaker than local-fatness. The example of a square with a short spike attached shows that perimeter-fatness is also not stronger than local-fatness. For classically fat shapes, the ratio of the squared diameter and the area can be bounded by a constant by Theorem 2.1.

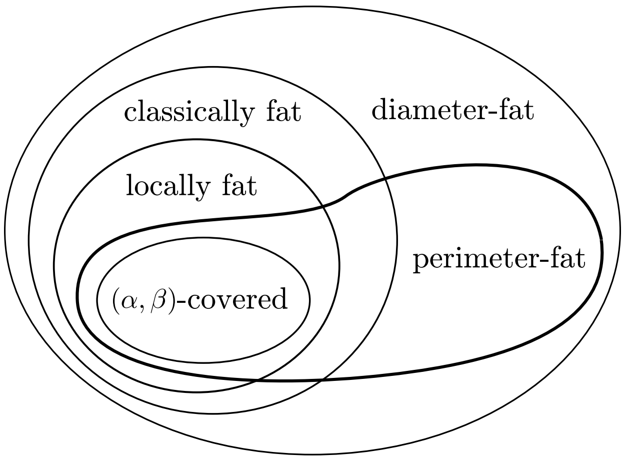

Figure 2.2.: Each region represents the class of simply connected sets in the plane that are fat with respect to one of the definitions. For example, the thick curve represents the class of perimeter-fat objects.

Observe that the condition $\frac{\mathrm{diam}(A)^2}{|A|} \leq C$ is weaker than $\frac{|\partial A|^2}{|A|} \leq \kappa$ only for connected shapes. If A is allowed to have several components as in our case, the two conditions are incomparable.

Fat shapes in arbitrary dimension. From dimension 3 on, the conditions $\frac{\mathrm{diam}(A)^d}{|A|} \leq C$ and $\frac{|\partial A|^d}{|A|^{d-1}} \leq \kappa$ are incomparable, even for connected shapes, which we show by giving examples. The volume and the surface area of a cube with an attached long, skinny spike can be bounded from above by a constant, but the diameter can be arbitrarily large. Hence such a shape is perimeter-fat, but not diameter-fat. On the other hand, the comb-like shape A that consists of the union of the rectangles $[0,1]^{d-1} \times [\frac{2i}{2n}, \frac{2i+1}{2n}]$ for $i = 0, \ldots, n-1$ and $\{0\} \times [0,1]^{d-1}$ has volume $\frac{1}{2}$ and diameter roughly $\sqrt{d}$. Therefore $\frac{\mathrm{diam}(A)^d}{|A|}$ is constant, while $\frac{|\partial A|^d}{|A|^{d-1}} \geq 2n^d$ since $|\partial A| \geq 2n + 2$.

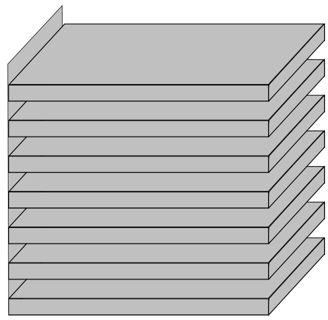

Figure 2.3.: A connected shape in dimension 3, which is diameter-fat, but not perimeter-fat.

Figure 2.3 shows a sketch of such a shape in dimension 3.

The isoperimetric quotient measures how similar a shape is to a ball; it is large for “skinny” shapes and shapes that are composed of many connected components. Next we show that, for convex sets, perimeter-fatness is equivalent to the standard definition of ball-fatness and also to diameter-fatness.

Lemma 2.6. *Let $K \subset \mathbb{R}^d$ be a convex set. Then the following three conditions are equivalent.*

1. *K is perimeter-fat.*
2. *K is ball-fat.*
3. *K is diameter-fat.*

Proof. 1. $\Longrightarrow$ 2.
We first show that it is sufficient to prove the claim for ellipsoids whose centroid is the origin and whose symmetry axes are parallel to the coordinate axes.

Every ellipsoid is the affine image $AB^d + t$ of the d-dimensional Euclidean unit ball B^d where A is a $(d \times d)$-matrix with $\det(A) \neq 0$, and $t \in \mathbb{R}^d$. For each convex set $K \subset \mathbb{R}^d$, there is an ellipsoid E such that $E \subseteq K \subseteq d\,E$ [35]. It is called the John ellipsoid. Therefore the volume, surface area of the boundary, and the radii of the smallest enclosed and largest inscribed ball of K and E only differ by constant factors. To see this for the surface area of the boundary, note that, for convex sets K_1 and K_2, $K_1 \subseteq K_2$ implies $|\partial K_1| \leq |\partial K_2|$; see [17, Section 31, Property (5)].

Let $\text{ins}(K)$ be the largest inscribed ball of K, and let $\text{enc}(K)$ be the smallest enclosing ball of K. Since the isoperimetric quotient, as well as the ratio of the radii of $\text{enc}(K)$ and $\text{ins}(K)$, are invariant under rigid motions, we can assume that E has the origin as centroid and its symmetry axes are the coordinate axes. Then $E = AB^d$, and A is a diagonal matrix with positive entries $\lambda_1, \ldots, \lambda_d$ on the diagonal. W.l.o.g. the value λ_1 is the maximal λ_i, the value λ_d is the minimal λ_i, and $\lambda_d = 1$. The volume $|E| = |AB^d| = |\det(A)||B^d| = \lambda_1 \cdot \ldots \cdot \lambda_d \cdot \omega_d$.

Next we estimate $|\partial E|$ from below by the boundary of the cross polytope $Q = \operatorname{conv}(\pm\lambda_1 e_1, \dots, \pm\lambda_d e_d)$, which is contained in E. The surface area $|\partial Q|$ is the sum of 2^d equal summands. Each summand is the $(d-1)$-dimensional volume of one of the congruent simplices $\operatorname{conv}(\pm\lambda_1 e_1, \dots, \pm\lambda_d e_d)$. By the Pythagorean theorem we have

$$|\operatorname{conv}(\lambda_1 e_1, \dots, \lambda_d e_d)|^2 = \sum_{j=1}^{d} |\operatorname{conv}(0, \lambda_1 e_1, \dots, \widehat{\lambda_j e_j}, \dots, \lambda_d e_d)|^2.$$

The hat $\widehat{\ }$ means that the labeled element is left out in the list. The $(d-1)$-dimensional volume $|\operatorname{conv}(0, \lambda_1 e_1, \dots, \widehat{\lambda_j e_j}, \dots, \lambda_d e_d)| = \frac{1}{(d-1)!}\lambda_1 \cdot \ldots \widehat{\lambda_j} \cdot \ldots \lambda_d$.

Observe that by assumption the radius of $\operatorname{enc}(E)$ equals λ_1 and the radius of $\operatorname{ins}(E)$ equals 1. Putting things together, we get the following inequalities. For the second of the inequalities note that, because of the Cauchy-Schwarz inequality, we have $|x| \geq \frac{1}{\sqrt{d}}||x||_1$ for all $x \in \mathbb{R}^d$.

$$\begin{aligned}
\frac{|\partial E|^d}{|E|^{d-1}} &\geq \frac{2^d\Big(\sum_{j=1}^d (\lambda_1 \cdot \ldots \cdot \widehat{\lambda_j} \cdot \ldots \cdot \lambda_d)^2\Big)^{d/2}}{(d-1)!(\omega_d \lambda_1 \cdot \ldots \cdot \lambda_d)^{d-1}} \\
&\geq \frac{2^d\Big(\sum_{j=1}^d \lambda_1 \cdot \ldots \cdot \widehat{\lambda_j} \cdot \ldots \cdot \lambda_d\Big)^{d}}{\sqrt{d}(d-1)!(\omega_d \lambda_1 \cdot \ldots \cdot \lambda_d)^{d-1}} \\
&\geq \frac{2^d}{\sqrt{d}\,(d-1)!\,\omega_d^{d-1}} \cdot \frac{(\lambda_1 \cdot \ldots \cdot \lambda_{d-1})^d}{(\lambda_1 \cdot \ldots \cdot \lambda_d)^{d-1}} \\
&= C \cdot \lambda_1 \cdot \ldots \cdot \lambda_{d-1} \geq C \cdot \lambda_1.
\end{aligned}$$

We have shown that $\frac{|\partial E|^d}{|E|^{d-1}} \geq C \cdot \frac{R}{r}$ where R is the radius of $\operatorname{enc}(E)$ and r is the radius of $\operatorname{ins}(E)$.

2. $\Longrightarrow$ 3.
Let r be the radius of $\operatorname{ins}(K)$, and let R be the radius of $\operatorname{enc}(K)$. The claim follows from

$$\frac{\operatorname{diam}(K)^d}{|K|} 2^{-d}\omega_d \leq \frac{|\operatorname{enc}(K)|}{|\operatorname{ins}(K)|} = \Big(\frac{R}{r}\Big)^d.$$

3. $\Longrightarrow$ 1.
By Jung's Theorem (Theorem 2.1), A is contained in a ball of radius $\sqrt{\frac{d}{2d+2}}\operatorname{diam}(A)$. This implies $|\partial A| \leq |\partial \operatorname{enc}(A)| \leq d\omega_d\big(\frac{d}{2d+2}\big)^{(d-1)/2} \operatorname{diam}(A)^{d-1}$, which implies the claim. $\square$

The next proposition says that the isoperimetric quotient of the union of perimeter-fat sets is linear in the number of the sets, if the sets have roughly the same volume.

Proposition 2.7. *1. Let $A \subset \mathbb{R}^d$ be the union of n interior-disjoint congruent copies of a set $B \subset \mathbb{R}^d$. Then $\frac{|\partial A|^d}{|A|^{d-1}} \leq n\frac{|\partial B|^d}{|B|^{d-1}}$, and equality holds if the copies of B are pairwise disjoint.*

2. Let A be the union of sets $B_1, B_2, \ldots, B_n$ of $\mathbb{R}^d$ that have pairwise disjoint interiors and that are perimeter-fat with parameter κ such that $\frac{\max_{1\leq i\leq n}|B_i|}{\min_{1\leq i\leq n}|B_i|} \leq \eta$ for some positive constant η. Then $\frac{|\partial A|^d}{|A|^{d-1}} \leq n\kappa\eta^{d-1}$.

Proof. The first claim follows from the fact that $|A| = n|B|$ and $|\partial A| \leq n|\partial B|$.

We prove the second claim. Let $j \in \{1, \ldots, n\}$ be such that $|\partial B_j| = \max_i |\partial B_i|$.

$$\begin{aligned}
\frac{|\partial A|^d}{|A|^{d-1}} &\leq \frac{(|\partial B_1| + |\partial B_2| + \cdots + |\partial B_n|)^d}{(|B_1| + |B_2| + \cdots + |B_n|)^{d-1}} \\
&\leq \frac{(n \max_i |\partial B_i|)^d}{(n \min_i |B_i|)^{d-1}} \\
&= \frac{(n|\partial B_j|)^d}{(n|B_j|)^{d-1}} \frac{(n|B_j|)^{d-1}}{(n \min_i |B_i|)^{d-1}} \\
&\leq n\kappa\eta^{d-1}
\end{aligned}$$

□

3. The objective function

For two shapes $A, B \subset \mathbb{R}^d$ and a class of transformations, the matching problem asks for a transformation t such that $t(A)$ and B match optimally. The quality of the match is measured by some distance measure, for instance the volume of the symmetric difference. Our matching problem can be formulated as finding a value that (approximately) minimizes the function $\bar{F}$ that maps a transformation t to the volume of the symmetric difference of $t(A)$ and B. Recall that we consider translations and rigid motions as transformations.

It is slightly sloppy to always denote the objective function with the same symbol since sometimes we consider rigid motions as transformations and sometimes we restrict our attention to translations, but surely no confusion will arise from this notation. The main goal of this chapter is to prove that the objective function is Lipschitz continuous for many metrics on the transformation space because we need this for the analysis of our algorithms in Chapters 5 to 8.

Let us mention an equivalent formulation of the matching problem. Under volume-preserving transformations, like translations and rigid motions, minimizing the volume of the symmetric difference is the same as maximizing the volume of overlap. When it comes to relative error approximation, this is no longer the case. For perfectly matching shapes there is a transformation t such that $|t(A) \triangle B| = 0$ and therefore a $(1+\varepsilon)$-approximation algorithm for minimizing the symmetric difference has to compute the exact optimum for any $\varepsilon > 0$, while a $(1-\varepsilon)$-approximation of the maximal volume of overlap does not have to be an exact match.

We denote the function that maps a translation or rigid motion t to the volume of overlap of $t(A)$ and B by F. The functions F and $\bar{F}$ are related through the following simple equation $\bar{F}(t) = |A| + |B| - 2F(t)$. This implies that both problem formulations are essentially the same for absolute error approximation, which we make precise in the next proposition. We allow ourselves to switch between both formulations and to speak of $\bar{F}$, as well as of F, as the objective function.

Proposition 3.1. *If $A, B \subset \mathbb{R}^d$ and r^*, r are rigid motions, then $|r^*(A) \cap B| \geq |r(A) \cap B| - \varepsilon|A|$ if and only if $|r^*(A) \triangle B| \leq |r(A) \triangle B| + 2\varepsilon|A|$.*

3.1. Known results

We discuss definitions and properties of the objective function. We start by summarizing some results about the volume of the symmetric difference from Hagedoorn's thesis [30]. The volume of the symmetric difference is obviously symmetric, and it fulfills the triangle inequality. Measured by the volume of the symmetric difference, each set clearly has distance 0 from itself. A distance measure with these three properties is usually called a *pseudometric*.

On the set of all subsets of $\mathbb{R}^d$, the volume of the symmetric difference is not a metric because sets that only differ by a measure-zero set have distance 0 although they are not equal. For the same reason, it is not a metric on the set of compact sets. It is a metric on the class of sets that equal the closure of their interior.

Denote the volume of the symmetric difference by $\sigma(A, B) = |A \triangle B|$. A pseudometric pattern space is a triple of a base space, a collection of subsets of the base space, called the patterns, and a pseudometric on the collection. Let $\mathcal{K}(\mathbb{R}^d)$ be the set of compact subsets of $\mathbb{R}^d$. Then $(\mathbb{R}^d, \mathcal{K}(\mathbb{R}^d), \sigma)$ is a pseudometric pattern space.

In applications, patterns are often obtained by the following process. Real world data is measured by some device, for example a camera, an MRT or CT scanner. Then some features are extracted and a pattern is computed by geometric modeling. In this process, errors can occur such that the result is not exact.

It is important that the distance measure that is used to compare the patterns is not too sensitive to inaccurate pattern modeling. For example, it is likely to happen that a pattern that is extracted from a pixel image contains outliers that are isolated black pixels that do not belong to the "real" pattern. A naive approach that compares this pattern without further processing by the Hausdorff distance would clearly fail.

How sensitive the matching to inexact pattern modeling is, depends on the used distance measure. Hagedoorn defines four types of robustness for pseudometric pattern spaces, and shows that $(\mathbb{R}^d, \mathcal{K}(\mathbb{R}^d), \sigma)$ satisfies these four axioms of robustness, namely deformation, blur, crack and noise robustness.

All four properties have a form that is analogous to the definition of continuity. If the "difference" between two patterns is sufficiently small, then also the distance measured by the pseudometric is arbitrarily small. Each of the four definitions captures a distinct type of "difference". Here, we only explain the intuitive meaning; we refer the reader to [30] for the precise definitions.

Deformation robustness means that the distance does not change much if the pattern is transformed by a homeomorphism that is close enough to the identity. A pseudometric pattern space is *blur robust* if adding boundary to a pattern in a sufficiently small neighborhood of its boundary does not increase the distance much. *Noise robustness* means that patterns that are equal outside a sufficiently small neighborhood of some point have a small distance.

A *crack* C is a subset of the boundary B of a pattern such that the closure of $B \setminus C$ equals B, and C is homeomorphic to a closed ball in some finite-dimensional Euclidean space. A line segment contained in a plane disk in $\mathbb{R}^3$ is an example of a crack. In a *crack robust* pseudometric pattern space, two patterns have a small distance if they are equal outside a small neighborhood of a crack.

After having discussed some properties of the distance measure, we summarize results about the combinatorial complexity of the objective function $\bar{F}$ from [30]. Let $A, B \subset \mathbb{R}^d$ be finite unions of d-dimensional simplices such that the intersection of two simplices is empty or a lower-dimensional simplex that is a face of both simplices.

In other words, A and B are homogeneous geometric simplicial complexes.

The *combinatorial complexity* of $\bar{F}$ is the number of "pieces" that $\bar{F}$ consists of. More precisely, two transformations s and t are combinatorially equivalent if the incidence relations between the simplices of $s(A)$ and $t(A)$ with B are the same. The number of connected components in the arrangement of combinatorially equivalent cells is the combinatorial complexity of $\bar{F}$.

In the worst case, the combinatorial complexity of $\bar{F}$ is $\Theta((nm)^k)$ where n and m are the numbers of simplices of A and B, respectively, and k is the dimension of the transformation group. Intuitively, the dimension of a transformation group is the number of degrees of freedom. The dimension of the group of translations in $\mathbb{R}^d$ is d, and the dimension of the group of rigid motions is $\binom{d+1}{2}$, as we will explain in Chapter 8. In two dimensions, the worst case combinatorial complexity is obtained already for connected polygons.

Mount, Silverman, and Wu [43] proved that in the case of translations in the plane, the function F that maps a translation t to the area of overlap of $t(A)$ and B is a piecewise polynomial of degree 2. The two variables represent the translation. Together with the bound $O((nm)^2)$ on the complexity of F, this implies an algorithm that maximizes the area of overlap of two simple polygons with n and m vertices. Ahn et al. [1] claim that this generalizes to an $O((nm)^d)$ algorithm for sets of simplices in $\mathbb{R}^d$. For further references on maximizing the volume of overlap of two shapes, we refer the reader to Chapter 5 for translations and Chapters 6 and 8 for rigid motions.

In two dimensions, Hagedoorn describes the function that maps an affine transformation t to the volume of overlap of $t(A)$ and B. Here, A and B are finite unions of n and m triangles in the plane that have pairwise disjoint interiors. An affine map is represented by the vector that contains the entries of the matrix, which is the linear part, and the entries of the translation vector. This function from $\mathbb{R}^6$ to $\mathbb{R}$ is piecewise rational, and the rationals are quotients of a polynomial and a monomial, both having degree $O(nm)$.

A restriction of F arises in other contexts. For $A \subset \mathbb{R}^d$, the function g_A that maps a translation vector $t \in \mathbb{R}^d$ to the volume of $(A + t) \cap A$ is called the *covariogram of A*, sometimes also the *set covariance*, and was introduced by Matheron [39] for compact sets.

We bound $|(A + t) \setminus A|$ in Section 3.2. Since $g_A(0) - g_A(t) = |(A + t) \setminus A|$, this volume is related to estimating the directional derivatives of g_A at 0. For convex, compact sets A, these are determined in [39]. For $u \in \mathcal{S}^{d-1}$, consider the function $\lambda \mapsto g_A(\lambda u)$ for $\lambda \in \mathbb{R}$; it has a continuous derivative that equals $-\mathcal{L}^{d-1}(A|u^\perp)$ where $A|u^\perp$ denotes the orthogonal projection of A to the orthogonal space of u.

Galerne [29] studies g_A for measurable sets A of finite Lebesgue measure. He computes the directional derivatives at the origin and proves that the perimeter of A can be computed from these derivatives. In this context, the perimeter of a set A is at most as large as $\mathcal{H}^{d-1}(\partial A)$. See [29] for the definition of perimeter. He also computes the Lipschitz constant of g_A in terms of the directional variation. For further details and definitions, we refer the reader to the paper and the references cited therein.

The inverse question whether the covariogram determines a convex body in $\mathbb{R}^d$, up to translations and reflections in the origin, is answered affirmatively for the planar case in [11]. For three dimensions, convex polytopes are determined by their covariogram [14]. In dimension ≥ 4, the question has a negative answer [16]. In the planar case, the class of sets among which the covariogram of a convex body is unique is extended in [13].

The function $g_{B,A}(t) := |(A + t) \cap B|$ is called *cross covariogram* for convex sets A and B. It equals F for the case of convex shapes and translations. Bianchi [15] proves that, for convex polygons A and B in the plane, $g_{B,A}$ determines the pair (A, B), except for a few exceptional pairs of parallelograms. The family of exceptions is completely described. For further references and other occurrences of the covariogram problem, see [11].

3.2. Upper bounds on the volume of the symmetric difference of a body and a congruent copy

A preliminary version of this and the next section can be found in [47]. Let A be a bounded subset of $\mathbb{R}^d$. We give an upper bound on the volume of the symmetric difference of A and $f(A)$ where f is a translation, a rotation, or the composition of both, a rigid motion.

We bound the volume of the symmetric difference of A and $f(A)$ in terms of the $(d-1)$-dimensional volume of the boundary of A and the maximal distance of a boundary point to its image under f. In the case of translations, our bound is sharp. In the case of rotations, we get a sharp bound under the assumption that the boundary is sufficiently nice. Recall that the volume is measured by the d-dimensional Hausdorff measure, which coincides with the Lebesgue measure for Lebesgue measurable sets, and the boundary is measured by the $(d-1)$-dimensional Hausdorff measure, which matches the surface area for sufficiently nice sets.

Let A and B be bounded subsets of $\mathbb{R}^d$, and let F be the function that maps a rigid motion r to the volume of overlap of $r(A)$ and B. Knowing that F is Lipschitz continuous helps to maximize it. We apply our results to bound the difference $|F(r) - F(s)|$ for rigid motions r, s that are close, implying that F is indeed Lipschitz continuous for many metrics on the space of rigid motions. Depending on the metric, also a Lipschitz constant can be deduced from the bound. We will use this repeatedly in later chapters.

In Section 3.3, we prove the following theorem, which gives an upper bound on the volume of $A \triangle (A + t)$ in terms of the $(d-1)$-dimensional volume of ∂A and the length of t.

Theorem 3.2. *Let $A \subset \mathbb{R}^d$ be a bounded set. Let $t \in \mathbb{R}^d$ be a translation vector. Then*

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

This inequality is best possible, in the sense that it becomes false when the right hand side is multiplied with any constant that is smaller than one. Let us assume for contradiction that the upper bound could be multiplied with $(1-\varepsilon)$ for some small, positive ε. Translate a rectangle R with side lengths 1 and ε in direction of the side of length ε. If the length of that translation t is $\varepsilon/2$, then the volume of $R \triangle (R+t)$ equals ε, but the modified bound gives $\varepsilon - \varepsilon^3$.

We also ask how the volume of the symmetric difference behaves when we rotate the set, instead of translating it. A *rotation matrix* is defined to be an orthogonal matrix that has determinant +1. For a rotation matrix $M \in \mathbb{R}^{d\times d}$, we give an upper bound on the volume of $A \triangle MA$, in terms of the $(d-1)$-dimensional volume of the boundary of A and a parameter w that measures the closeness of M and the identity matrix with respect to A. The parameter w is the maximal distance between a and Ma among all points $a \in \partial A$. We prove the following theorem in Section 3.3.

Theorem 3.3. *Let $A \subset \mathbb{R}^d$ be a bounded set. Let $M \in \mathbb{R}^{d\times d}$ be a rotation matrix and let $w = \max_{a\in\partial A} |a - Ma|$. Then*

$$\mathcal{H}^d(A \triangle MA) \leq \left(\tfrac{2d}{d+1}\right)^{\frac{d-1}{2}} w\, \mathcal{H}^{d-1}(\partial A).$$

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

Again, the constant 1 is the best possible constant because a rotation is arbitrarily close to a translation if the rotation center is sufficiently far away from the rotated set.

A rigid motion is the composition of a rotation and a translation. Let $SO(d) \subset \mathbb{R}^{d\times d}$ be the special orthogonal group, that is, the group of rotation matrices. Parameterize the space of rigid motions as $\mathcal{R} = SO(d) \times \mathbb{R}^d$ where $(M,t) \in \mathcal{R}$ denotes the rigid motion $x \mapsto Mx + t$.

Since the symmetric difference fulfills the triangle inequality, we have the following corollary for rigid motions. We assume that ∂A is $(\mathcal{H}^{d-1}, d-1)$-rectifiable such that we can use Theorem 3.3 with constant 1.

Corollary 3.4. *Let $A \subset \mathbb{R}^d$ be a bounded set whose boundary ∂A is $(\mathcal{H}^{d-1}, d-1)$-rectifiable. Let $r = (M,t) \in \mathcal{R}$ be a rigid motion, and let $w = \max_{a\in\partial A} |a - Ma|$. Then*

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

□

We are interested in studying the smoothness of the function F. In particular, we want to bound $|F(r) - F(s)|$ if the rigid motions r and s are close, meaning that they do not move points from A too far apart. Equip the space of rigid motions $\mathcal{R}$ with any metric that is induced by a norm on $\mathbb{R}^{d\times d} \times \mathbb{R}^d$. Then the following bound on $|F(r) - F(s)|$ implies that F is Lipschitz continuous if ∂A has a finite $(d-1)$-dimensional volume. Also a Lipschitz constant can be deduced.

Corollary 3.5. *Let $A, B \subset \mathbb{R}^d$ be bounded and $\mathcal{H}^d$-measurable sets such that ∂A is $(\mathcal{H}^{d-1}, d-1)$-rectifiable. Let $r = (M, p)$ and $s = (N, q)$ be rigid motions. Then*

$$|F(r) - F(s)| \leq \tfrac{1}{2}\big(|p-q| + w\big)\, \mathcal{H}^{d-1}(\partial A)$$

where $w = \max_{a \in \partial A} |Ma - Na|$.

We prove an easy proposition, which we then use to prove Corollary 3.5.

Proposition 3.6. *Let $(\mathcal{M}, \mu)$ be a measure space and let D, E, G be μ-measurable sets in $\mathcal{M}$. If $\mu(D) = \mu(G)$, then $|\mu(D \cap E) - \mu(G \cap E)| \leq \mu(D \setminus G)$.*

Proof. Because of $\mu(D) = \mu(G)$, we have $\mu(G \setminus D) = \mu(D \setminus G)$.

$$\begin{aligned} |\mu(D \cap E) - \mu(G \cap E)| &= |\mu((D \setminus G) \cap E) - \mu((G \setminus D) \cap E)| \\ &\leq \max\{\mu((D \setminus G) \cap E), \mu((G \setminus D) \cap E)\} \\ &\leq \max\{\mu(D \setminus G), \mu(G \setminus D)\} \\ &= \mu(D \setminus G). \end{aligned}$$ □

Proof of Corollary 3.5 from Corollary 3.4. Let $r = (M, p)$ and $s = (N, q)$ be rigid motions, and let $w = \max_{a \in \partial A} |Ma - Na|$. By the invariance of $\mathcal{H}^d$ under rigid motions we can apply Proposition 3.6 to get

$$|\mathcal{H}^d(r(A) \cap B) - \mathcal{H}^d(s(A) \cap B)| \leq \tfrac{1}{2}\, \mathcal{H}^d\big(r(A) \triangle s(A)\big) = \tfrac{1}{2}\, \mathcal{H}^d\big((s^{-1} \circ r)(A) \triangle A\big).$$

The map $s^{-1} \circ r$ is a rigid motion with rotation matrix $N^{-1}M$ and translation vector $N^{-1}(p-q)$. Therefore, $|\mathcal{H}^d(r(A) \cap B) - \mathcal{H}^d(s(A) \cap B)| \leq \frac{1}{2}\big(|p-q| + w\big)\, \mathcal{H}^{d-1}(\partial A)$ by Corollary 3.4. □

From Corollary 3.5, Lipschitz constants for various metrics can be deduced.

3.3. Proof of the upper bounds

We give the proofs of Theorems 3.2 and 3.3. More precisely, as Proof Step 1, we show that $A \triangle (A+t)$ and $A \triangle MA$ are contained in certain unions of line segments. For translations t, we show that $A \triangle (A+t) \subseteq \partial A \oplus [0,1]t$ where $\oplus$ denotes the Minkowski sum. For a rotation matrix M, we prove that $A \triangle MA$ is contained in the set of all line segments from a to Ma for $a \in \partial A$; see Figure 3.1. As Proof Step 2, we bound the volume of the unions of these line segments.

Proof Step 1 – Covering the symmetric difference of a body and a copy by line segments

For $x, y \in \mathbb{R}^d$, the line segment from x to y is the set $\{(1-\lambda)x + \lambda y : \lambda \in [0,1]\}$, and is denoted by $\ell(x, y)$. The Minkowski sum of two sets $A \oplus B$ equals the set of all sums $a + b$ for $a \in A$ and $b \in B$.

We show that for any bounded set $A \subset \mathbb{R}^d$ and a translation t, the set $A \triangle (A+t)$ is covered by the union of the line segments $\ell(a, a+t)$ where a is in the boundary of A. For a rotation matrix M, the set $A \triangle MA$ is covered by the union of the line segments $\ell(a, Ma)$ where a again is in the boundary of A. See Figure 3.1.

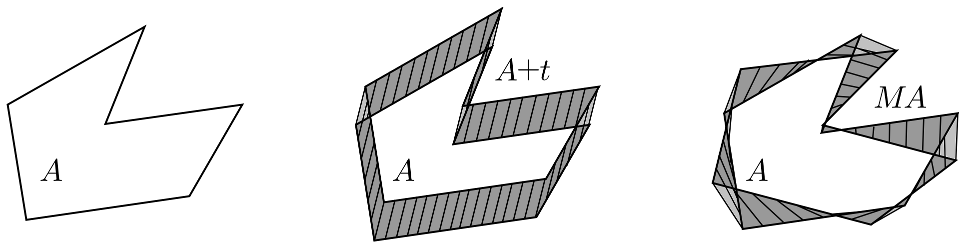

Figure 3.1.: On the left, the figure shows a body A. In the middle, the figure shows A and a translated copy of A. On the right, the figure shows A and a rotated copy. The symmetric differences are drawn in dark-gray. Examples of the line segments are drawn and the union of the line segments is drawn in light-gray.

Lemma 3.7. *Let $A \subset \mathbb{R}^d$ be a bounded set, and let $t \in \mathbb{R}^d$ be a translation vector. Then*

$$A \triangle (A+t) \subseteq \bigcup \{\ell(a, a+t) : a \in \partial A\}.$$

Proof. It suffices to show $A \setminus (A+t) \subseteq \bigcup\{\ell(a, a+t) : a \in \partial A\}$, because applying this to $A' = A + t$ and $t' = -t$ gives $(A+t) \setminus A \subseteq \bigcup\{\ell(a, a+t) : a \in \partial A\}$. Let $a \in A \setminus (A+t)$ and let l be the line $\{a + \lambda t : \lambda \in \mathbb{R}\}$. If $a \in \partial A$, we are done. Otherwise $a \in \operatorname{int}(A)$ and therefore $a + t \in \operatorname{int}(A+t)$. Since l intersects $\operatorname{int}(A+t)$ and $A+t$ is bounded, we have $\partial(A+t) \cap l \neq \emptyset$. Let $\lambda \in (0, 1]$ such that $a + \lambda t \in \partial(A+t)$. Then $a' = a + (\lambda - 1)t$ is a point in ∂A such that $a' + (1-\lambda)t = a$. □

Lemma 3.8. *Let $A \subset \mathbb{R}^d$ be a bounded set, and let $M \in \mathbb{R}^{d \times d}$ be a rotation matrix. Then*

$$A \triangle MA \subseteq \bigcup\{\ell(a, Ma) : a \in \partial A\}.$$

Proof. Consider the continuous function $\varphi : [0,1] \times \mathbb{R}^d \to \mathbb{R}^d$ that is defined by $\varphi(\lambda, x) = \varphi_\lambda(x) = (1-\lambda)x + \lambda M x$. We show that $MA \setminus A \subseteq \varphi([0,1] \times \partial A)$. Applying this to $A' = MA$ and $M' = M^{-1}$ gives the claim. We first prove that φ_λ is injective for all $\lambda \in [0,1] \setminus \{\frac{1}{2}\}$. This implies that for each $\lambda \in [0,1] \setminus \{\frac{1}{2}\}$ and each bounded set $S \subset \mathbb{R}^d$, the function $\varphi_\lambda : \operatorname{cl}(S) \to \varphi_\lambda(\operatorname{cl}(S))$ is a homeomorphism because it is bijective and linear.

Assume that there exist $x, y \in \mathbb{R}^d$, $x \neq y$, such that $\varphi_\lambda(x) = \varphi_\lambda(y)$. Since $x \neq y$, we have $\lambda \neq 0$. Then $M(x-y) = ((\lambda-1)/\lambda)(x-y)$, so $(\lambda-1)/\lambda$ is an eigenvalue of the rotation M. Since only 1 or -1 can occur as eigenvalues for a rotation, we get $\lambda = 1/2$.

Let $y \in MA \setminus A$. We show that there exist $\lambda \in [0,1]$ and $a \in \partial A$ such that $y = \varphi(\lambda, a)$. We distinguish two cases.

- **Case 1.** $\varphi_{1/2} : \mathrm{cl}(A) \to \varphi_{1/2}(\mathrm{cl}(A))$ is not bijective and $y \in \varphi_{1/2}(\mathrm{cl}(A))$.

  Let $a, b \in \mathrm{cl}(A)$ such that $a \neq b$ and $\varphi_{1/2}(a) = \varphi_{1/2}(b)$, and let $x \in \varphi_{1/2}^{-1}(y) \cap \mathrm{cl}(A)$. For each point v on the line $l = \{x + \mu(b-a) : \mu \in \mathbb{R}\}$, we have $\varphi_{1/2}(v) = y$, due to the linearity of φ_λ. Since A is bounded and $x \in \mathrm{cl}(A)$, we have $\partial A \cap l \neq \emptyset$ and for every point x' in this set $y = \varphi(1/2, x') \in \varphi([0,1] \times \partial A)$.

- **Case 2.** $\varphi_{1/2} : \mathrm{cl}(A) \to \varphi_{1/2}(\mathrm{cl}(A))$ is bijective or $y \notin \varphi_{1/2}(\mathrm{cl}(A))$.

  Since $y \in \varphi_1(A)$, we can define $t = \inf\{\lambda \in [0,1] : \forall \mu \in [\lambda, 1]\, y \in \varphi_\mu(\mathrm{cl}(A))\}$. We now show that $y \in \varphi_t(\mathrm{cl}(A))$. Assume that $y \notin \varphi_t(\mathrm{cl}(A))$. Then for all $a \in \mathrm{cl}(A)$ the distance $|\varphi_t(a) - y| > 0$. Since $\mathrm{cl}(A)$ is compact and φ_t is continuous, $\varphi_t(\mathrm{cl}(A)$ is compact. Since the distance is continuous, we would have $\min\{|\varphi_t(a) - y| : a \in \mathrm{cl}(A)\} = \eta > 0$. Let $w = \max\{|a - Ma| : a \in \mathrm{cl}(A)\} < \infty$. For all $a \in \mathrm{cl}(A)$, $\lambda \in [0, 1-\nu]$, $\nu \in [0, 1-\lambda]$, we have

  $$|\varphi_\lambda(a) - \varphi_{\lambda+\nu}(a)| \leq \nu w, \tag{3.1}$$

  so we have $|y - \varphi_{t+\frac{\eta}{2w}}(a)| \geq \Big||y - \varphi_t(a)| - |\varphi_t(a) - \varphi_{t+\frac{\eta}{2w}}(a)|\Big| \geq \frac{\eta}{2}$ by the triangle inequality. Therefore $t < t + \frac{\eta}{2w} \leq \inf\{\lambda \in [0,1] : \forall \mu \in [\lambda, 1]\, y \in \varphi_\mu(\mathrm{cl}(A))\}$, which is a contradiction to the definition of t. Therefore $y \in \varphi_t(\mathrm{cl}(A))$. By the case distinction, $\varphi_t : \mathrm{cl}(A) \to \varphi_t(\mathrm{cl}(A))$ is bijective: If $\varphi_{1/2}$ is not bijective, then $y \notin \varphi_t(A)$ as we are in case 2. That implies $t > \frac{1}{2}$, and so φ_t is bijective. Otherwise, $\varphi_{1/2}$ is bijective, and so φ_s is bijective for every choice of s.

  If $y \in \partial A$, we are done. Assume otherwise that $y \notin \partial A$. Next, we show that $y \in \partial \varphi_t(\mathrm{cl}(A))$. Assume on the contrary that $y \in \mathrm{int}(\varphi_t(\mathrm{cl}(A)))$. Then $y \notin \mathrm{cl}(A) = \varphi_0(\mathrm{cl}(A))$, so we have $t > 0$. We show that there exists $\varepsilon' > 0$ such that for all $\delta \in (0, \varepsilon']$ we have $y \in \varphi_{t-\delta}(\mathrm{cl}(A))$, which is a contradiction to the definition of t.

  Since $y \in \mathrm{int}(\varphi_t(\mathrm{cl}(A)))$, there exists $\varepsilon > 0$ such that $B(y, \varepsilon) \subseteq \varphi_t(\mathrm{cl}(A))$. Let $\varepsilon' \in (0, \frac{\varepsilon}{3w}]$ such that for all $\delta \in (0, \varepsilon']$ we have $0 < t - \delta \neq \frac{1}{2}$. Let $U = \varphi_t^{-1}(B(y, \varepsilon))$. For all $\delta \in (0, \varepsilon']$, the function $f = \varphi_{t-\delta} \circ \varphi_t^{-1} : B(y, \varepsilon) \to \varphi_{t-\delta}(U)$ is a homeomorphism. Let $a \in \mathrm{cl}(A)$ be such that $y = \varphi_t(a)$. For $a' = \varphi_{t-\delta}^{-1}(y)$, we have $\varphi_{t-\delta}(a') = y$. We prove $a' \in \mathrm{cl}(A)$.

  Let $v \in B(y, \varepsilon)$ and let $v' = \varphi_t^{-1}(v)$. We have $|f(v) - v| = |\varphi_{t-\delta}(v') - \varphi_t(v')| \leq \delta w$ by Inequality (3.1). Because of $\delta \leq \frac{\varepsilon}{3w}$, we have

  $$|f(v) - v| \leq \tfrac{\varepsilon}{3}. \tag{3.2}$$

  The set $\partial B(y, \frac{2\varepsilon}{3})$ is the set of all points that have distance $\frac{2\varepsilon}{3}$ from y. The set $f(B(y, \varepsilon))$ is homeomorphic to a d-dimensional ball and $f(\partial B(y, \frac{2\varepsilon}{3}))$ is homeomorphic to a $(d-1)$-dimensional sphere. Because of Inequality (3.2), $f(\partial B(y, \frac{2\varepsilon}{3})) \subset B(y, \varepsilon) \setminus B(y, \frac{\varepsilon}{3})$ and $f(y) \in B(y, \frac{\varepsilon}{3})$. By the Jordan-Brouwer separation theorem, which is the generalization of the Jordan curve theorem to higher dimensions, the topological sphere $f(\partial B(y, \frac{2\varepsilon}{3}))$ cuts $\mathbb{R}^d$ in two components, one of which containing $B(y, \frac{\varepsilon}{3})$ and the other containing $\mathbb{R}^d \setminus B(y, \varepsilon)$. Since $f(B(y, \varepsilon))$ is simply connected, it follows that $B(y, \frac{\varepsilon}{3}) \subset f(B(y, \varepsilon))$ and therefore $y \in f(B(y, \varepsilon)) \subseteq \varphi_{t-\delta}(\mathrm{cl}(A))$, which is a contradiction to the definition of t.

Since φ_t is a homeomorphism, we have that $\partial \varphi_t(\mathrm{cl}(A)) = \varphi_t(\partial A)$, which ends the proof of Lemma 3.8. □

Proof Step 2 – Bounding the volume of certain unions of line segments

Next, we prove that the volume of the union of line segments from a to $a+t$ for $a \in \partial A$ is bounded by the length of t times the $(d-1)$-dimensional volume of ∂A. Together with the results of Proof Step 1, this gives a bound on the volume of $A \triangle (A+t)$, proving Theorem 3.2.

Lemma 3.9. *Let $A \subset \mathbb{R}^d$ be a bounded set. Let $t \in \mathbb{R}^d$ be a translation vector. Then*

$$\mathcal{H}^d\left(\bigcup\{\ell(a, a+t) : a \in \partial A\}\right) \leq |t|\, \mathcal{H}^{d-1}(\partial A).$$

Proof. We abbreviate $L = \bigcup\{\ell(a, a+t) : a \in \partial A\}$. Let $\delta > 0$ and let $\{B_j : j \in \mathbb{N}\}$ be a covering of ∂A with $\operatorname{diam} B_j \leq \delta$ for all $j \in \mathbb{N}$.

For each $j \in \mathbb{N}$, we define a cylinder Z_j such that $L \subseteq \bigcup\{Z_j : j \in \mathbb{N}\}$. We denote the orthogonal space of t by $t^\perp$. The top and bottom of the cylinder are formed by copies of B_j projected to $t^\perp$. See Figure 3.2 as an illustration. The bottom Z_j^b of the cylinder Z_j sits in the hyperplane that contains a point of $\mathrm{cl}(B_j)$, but does not contain any point of $\mathrm{cl}(B_j)$, when translated in direction $-t$ by any small amount. The top Z_j^t of the cylinder Z_j is formed by $Z_j^b + \left(1 + \frac{\mathrm{diam}(B_j)}{|t|}\right)t$. By construction, the cylinder Z_j contains $\bigcup\{\ell(b, b+t) : b \in B_j\}$.

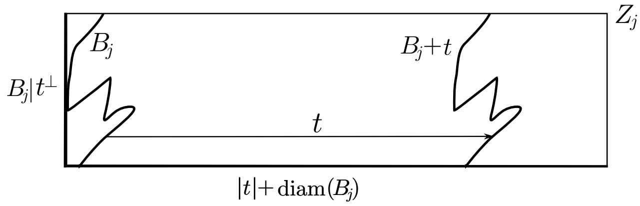

Figure 3.2.: The definition of the cylinder Z_j, which contains all line segments $\ell(b, b+t)$ for $b \in B_j$.

We have $\mathcal{H}^d(Z_j) = \mathcal{H}^{d-1}(B_j|t^\perp)\left(1 + \frac{\mathrm{diam}(B_j)}{|t|}\right)|t|$ due to [28, Theorem 3.2.23]. Note that $\mathrm{diam}(B_j|t^\perp) \leq \mathrm{diam}(B_j)$. We have $\mathcal{H}^{d-1}(B_j|t^\perp) = \mathcal{H}^{d-1}(\mathrm{conv}(B_j|t^\perp))$ and by Theorem 2.2, we have $\mathcal{H}^{d-1}(\mathrm{conv}(B_j|t^\perp)) \leq \omega_{d-1}(\mathrm{diam}(B_j)/2)^{d-1}$ and hence

$$\mathcal{H}^d(L) \;\leq\; \sum_{j\in\mathbb{N}} \mathcal{H}^d(Z_j) \;\leq\; (|t| + \delta) \sum_{j\in\mathbb{N}} \omega_{d-1}(\mathrm{diam}(B_j)/2)^{d-1}.$$

This implies that $\mathcal{H}^d(L) \leq (|t| + \delta)\, \mathcal{H}^{d-1}_\delta(\partial A)$ for all $\delta > 0$. Therefore $\mathcal{H}^d(L) \leq |t|\, \mathcal{H}^{d-1}(\partial A)$. $\square$

Next, we will bound the volume of the union of line segments from a to Ma for $a \in \partial A$ for a rotation matrix M. Together with the results of the previous section, this gives a bound on the volume of $A \triangle MA$, proving Theorem 3.3.

Lemma 3.10. *Let $A \subset \mathbb{R}^d$ be a bounded set. Let $M \in \mathbb{R}^{d \times d}$ be a rotation matrix and let $w = \max_{a \in \partial A} |a - Ma|$. Then*

$$\mathcal{H}^d\left(\bigcup\{\ell(a, Ma) : a \in \partial A\}\right) \leq \left(\tfrac{2d}{d+1}\right)^{\frac{d-1}{2}} w\, \mathcal{H}^{d-1}(\partial A).$$

Proof. The proof works similarly as the one of Lemma 3.9. We abbreviate $L = \bigcup\{\ell(a, Ma) : a \in \partial A\}$. Let $\delta > 0$ and let $\{B_j : j \in \mathbb{N}\}$ be a covering of ∂A with $\operatorname{diam} B_j \leq \delta$ for all $j \in \mathbb{N}$. We again cover L by a set of cylinders $\{Z_j : j \in \mathbb{N}\}$, which are defined using the sets B_j. Since the line segments in L are not parallel, top and bottom of Z_j have a volume that is larger than a ball of diameter $\operatorname{diam}(B_j)$, and therefore we get a constant in the inequality, which is larger than one.

Let $\operatorname{enc}(B_j)$ and $\operatorname{enc}(MB_j)$ be the smallest enclosing balls of B_j and MB_j. Clearly, both have the same radius. By Theorem 2.1, their radius r is at most $\sqrt{\frac{d}{2d+2}}\operatorname{diam}(B_j)$. Let t be the vector from the center of $\operatorname{enc}(B_j)$ to the center of $\operatorname{enc}(MB_j)$. The length of $|t|$ is bounded from above by $w + 2\sqrt{\frac{d}{2d+2}}\operatorname{diam}(B_j)$. See Figure 3.3 for illustration.

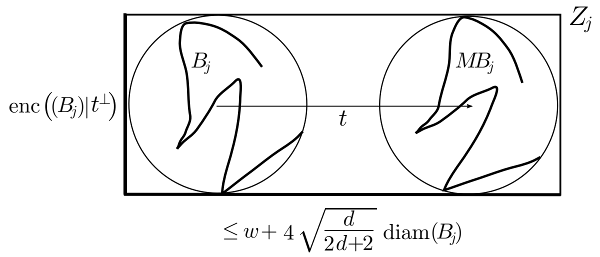

Figure 3.3.: The definition of the cylinder Z_j, which contains all line segments $\ell(b, Mb)$ for $b \in B_j$.

Already the convex hull of $\operatorname{enc}(B_j)$ and $\operatorname{enc}(MB_j)$ contains the union of all line segments $\ell(b, Mb)$ for $b \in B_j$. We enlarge this set by considering the cylinder Z_j having copies of the $(d-1)$-dimensional ball $\operatorname{enc}(B_j)|t^\perp$ as top and bottom. Top and bottom are touching $\operatorname{enc}(B_j)$ and $\operatorname{enc}(MB_j)$, respectively, so that the cylinder contains the convex hull of $\operatorname{enc}(B_j)$ and $\operatorname{enc}(MB_j)$. The volume of top and bottom equals $\mathcal{H}^{d-1}(\operatorname{enc}(B_j)|t^\perp) = \omega_{d-1} r^{d-1} \leq \omega_{d-1}\left(\frac{d}{2d+2}\right)^{\frac{d-1}{2}} \operatorname{diam}(B_j)^{d-1}$. The distance of top and bottom is at most $|t| + 2r \leq w + 4\sqrt{\frac{d}{2d+2}}\operatorname{diam}(B_j)$. By [28, Theorem

3.2.23] the volume of Z_j can be computed as the product of the area of the bottom and the height of the cylinder.

We have

$$\begin{aligned}
\mathcal{H}^d(L) &\leq \sum_{j\in\mathbb{N}} \mathcal{H}^d(Z_j) \\
&\leq \sum_{j\in\mathbb{N}} \omega_{d-1}\left(\tfrac{d}{2d+2}\right)^{\frac{d-1}{2}} \operatorname{diam}(B_j)^{d-1}\left(w + 4\sqrt{\tfrac{d}{2d+2}}\operatorname{diam}(B_j)\right) \\
&\leq \left(w + 4\sqrt{\tfrac{d}{2d+2}}\delta\right) \sum_{j\in\mathbb{N}} \omega_{d-1}\left(\tfrac{d}{2d+2}\right)^{\frac{d-1}{2}} \operatorname{diam}(B_j)^{d-1} \\
&\leq 2^{\frac{d-1}{2}}\left(\tfrac{d}{d+1}\right)^{\frac{d-1}{2}} \left(w + 4\sqrt{\tfrac{d}{2d+2}}\delta\right) \sum_{j\in\mathbb{N}} \omega_{d-1} 2^{-(d-1)} \operatorname{diam}(B_j)^{d-1}
\end{aligned}$$

This implies $\mathcal{H}^d(L) \leq \left(\frac{2d}{d+1}\right)^{\frac{d-1}{2}} \left(w + 4\sqrt{\frac{d}{2d+2}}\delta\right) \mathcal{H}^{d-1}_\delta(\partial A)$ for all $\delta > 0$. Therefore $\mathcal{H}^d(L) \leq \left(\frac{2d}{d+1}\right)^{\frac{d-1}{2}} w\, \mathcal{H}^{d-1}(\partial A)$. □

A close look into the proof of the above shows that, if we could assume in the proof that the covering $\{B_j\}_{j\geq 0}$ contains only balls, then the constant $\left(\frac{2d}{d+1}\right)^{\frac{d-1}{2}}$ in Lemma 3.10 could be replaced by 1. We can do this for sets for which the Hausdorff and the spherical measure coincide, as was discussed in Section 2.1.

Corollary 3.11. *Let $A \subset \mathbb{R}^d$ be a bounded set such that ∂A is $(\mathcal{H}^{d-1}, d-1)$-rectifiable. Let $M \in \mathbb{R}^{d\times d}$ be a rotation matrix and let $w = \max_{a\in\partial A} |a - Ma|$. Then*

$$\mathcal{H}^d\left(\bigcup\{\ell(a, Ma) : a \in \partial A\}\right) \leq w\, \mathcal{H}^{d-1}(\partial A).$$

For example, if $A \subset \mathbb{R}^d$ is a finite union of simplices, then ∂A is $(\mathcal{H}^{d-1}, d-1)$-rectifiable. Finite unions of simplices are a common representation of shapes, so this corollary is interesting for the shape matching application of this chapter.

3.4. Lower bounds

We prove a lower bound on the volume of $(A+t) \triangle A$ for a translation vector $t \in \mathbb{R}^d$ in terms of the length of t and the diameter and the volume of A.

Lemma 3.12. *Let $t \in \mathbb{R}^d$ be a vector of length at most* $\operatorname{diam}(A)$. *Then*

$$|(A+t) \triangle A| \geq \tfrac{|A|}{\operatorname{diam}(A)}\, |t|.$$

Proof. Let $k = \lceil \frac{\operatorname{diam}(A)}{|t|} \rceil$. Let H be a hyperplane orthogonal to t such that H touches A, and one closed halfspace defined by H, say H^+, contains A and $A+t$. For an illustration of the following definitions, see Figure 3.4. Define parallel hyperplanes $H_0, H_1, \ldots, H_{k+2}$ such that the H_i are orthogonal to t, the distance of H_i to H_{i+1} equals $|t|$ for all $i = 0, \ldots, k+1$ and $H_1 = H$. Denote the slices $H_i^+ \cap H_{i+1}^-$ by S_i and the volumes $|S_i \cap A|$ by m_i for $i = 0, \ldots, k$. We have $m_0 = m_{k+1} = 0$ and

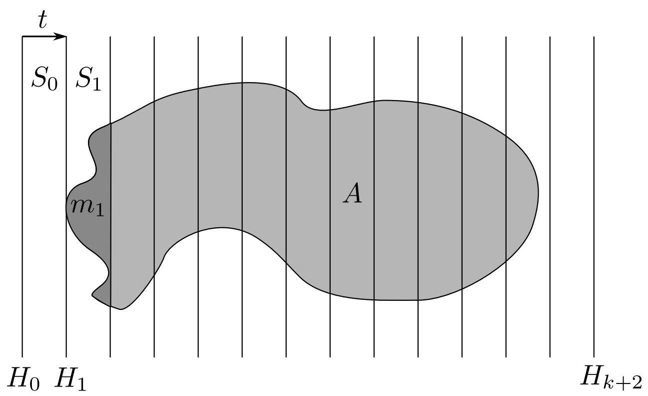

Figure 3.4.: The definition of the H_i, S_i and m_i is sketched.

$\sum_{i=1}^{k} m_i = |A|$. Hence there exists a j such that $m_j \geq \frac{|A|}{k}$. For the first step of the following computation, note that slice S_i is mapped onto S_{i+1} by the translation t.

$$
\begin{aligned}
|(A+t) \triangle A| &= \sum_{i=0}^{k} |t(A \cap S_i) \triangle (A \cap S_{i+1})| \\
&\geq \sum_{i=0}^{k} ||t(A \cap S_i)| - |A \cap S_{i+1}|| \\
&= \sum_{i=0}^{k} ||A \cap S_i| - |A \cap S_{i+1}|| \\
&= \sum_{i=0}^{k} |m_i - m_{i+1}| \\
&= \sum_{i=0}^{j-1} |m_i - m_{i+1}| + \sum_{i=j}^{k} |m_i - m_{i+1}| \\
&\geq \left| \sum_{i=0}^{j-1} m_i - m_{i+1} \right| + \left| \sum_{i=j}^{k} m_i - m_{i+1} \right| \\
&= 2m_j \\
&\geq \frac{2|A|}{k}.
\end{aligned}
$$

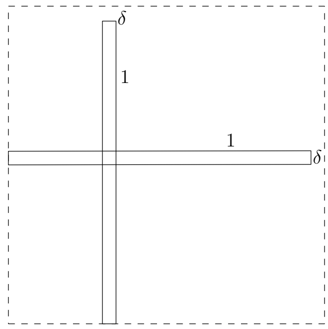

Figure 3.5.: Two rectangles of side lengths 1 and δ with small maximal overlap. The square has the same area as the Minkowski sum of the rectangles.

Because of

$$\frac{2|A|}{k} = \frac{2|A|}{\lceil \frac{\operatorname{diam}(A)}{|t|} \rceil} \geq \frac{2|A||t|}{\operatorname{diam}(A) + |t|} \geq \frac{|A||t|}{\operatorname{diam}(A)},$$

the claim of the lemma follows. □

We are interested in the maximal volume of overlap of A and B; still, it is an interesting question how small the maximal volume of overlap can be. We determine a simple lower bound in terms of the volumes and the diameters of A and B.

The maximal volume of overlap can be very small as the following example of skinny rectangles shows. The rectangles $[0,1]^k \times [0,\delta]^{d-k} \subset \mathbb{R}^d$ and $[0,\delta]^k \times [0,1]^{d-k} \subset \mathbb{R}^d$ have a maximal volume of overlap of δ^d, which equals the product of their volumes. The volume of their Minkowski sum is $(1+\delta)^d$. Figure 3.5 shows an example in dimension 2.

We consider the case of translations here. Of course, the bound is then also a lower bound in the case of rigid motions.

Proposition 3.13. *There exists $t \in \mathbb{R}^d$ such that $|t(A) \cap B| \geq \frac{|A|\,|B|}{|B \oplus (-A)|}$.*

Proof. As we will see soon, the function $f : \mathbb{R}^d \to \mathbb{R}^d$ with $f(t) = \frac{|t(A) \cap B|}{|A|\,|B|}$ is a density function that equals 0 outside the set $B \oplus (-A)$. If we had $f(t) < \frac{1}{|B \oplus (-A)|}$ for all $t \in \mathbb{R}^d$, then $1 = \int_{B \oplus (-A)} f(t)\, dt < \int_{B \oplus (-A)} \frac{1}{|B \oplus (-A)|}\, dt = 1$, which would be a contradiction. Hence there exists $t \in \mathbb{R}^d$ such that $|t(A) \cap B| \geq \frac{|A|\,|B|}{|B \oplus (-A)|}$. □

We now find an upper bound on $|B \oplus (-A)|$ in terms of the diameters of A and B. Let $\operatorname{enc}(-A)$ be the smallest enclosing ball of $-A$. We have $B \oplus (-A) \subseteq \operatorname{enc}(B) \oplus \operatorname{enc}(-A)$.

The Minkowski sum of two balls is again a ball. Its radius is the sum of the radii of the summands. Hence the volume of $\operatorname{enc}(B) \oplus \operatorname{enc}(-A)$ equals $(r_A + r_B)^d \omega_d$ where r_A is the radius of $\operatorname{enc}(-A)$ and r_B is the radius of $\operatorname{enc}(B)$.

By Theorem 2.1, we have $r_A \leq \sqrt{\frac{d}{2d+2}}\,\mathrm{diam}(A)$ and analogously for B. Hence there exist numbers C_1 and C_2 that depend on d and fulfill

$$|B \oplus (-A)| = C_1\big((\mathrm{diam}(A) + \mathrm{diam}(B))^d\big) = C_2\big(\mathrm{diam}(A)^d + \mathrm{diam}(B)^d\big).$$

Thus as a lower bound on $\frac{|A||B|}{|B\oplus(-A)|}$ we get the following:

$$\frac{|A||B|}{|B \oplus (-A)|} \geq \omega_d \left(\frac{2d+2}{d}\right)^{\frac{d}{2}} \frac{|A||B|}{(\mathrm{diam}(A) + \mathrm{diam}(B))^d}.$$

The factor $\omega_d \left(\frac{2d+2}{d}\right)^{\frac{d}{2}}$ tends to 0 as d tends to ∞, but we assume that the dimension is constant. If the shapes A and B fulfill the assumptions described in Section 2.2, we know an upper bound on $\mathrm{diam}(A)$ and a lower bound on $|A|$. If we are given such bounds also for the shape B, we can compute a lower bound on $\frac{|A||B|}{|B\oplus(-A)|}$ with the above inequality.

4. The probabilistic toolbox

In this short chapter, we introduce the tools from probability theory that we use for analyzing the probabilistic algorithms in Chapters 5, 6 and 8. We apply results from [25] to deduce a theorem that we call our *probabilistic toolbox*. It describes conditions under which the empirical measure approximates the density function uniformly and is central for the analysis of our algorithms.

We assume that the reader is familiar with the basic notions of probability theory, such as the concepts of random vectors, density functions and expected values, as they are explained, for example, in [36, 34]. We start by briefly discussing how to determine the density function of a random vector. We cite a transformation formula for density functions that allows us to compute the density function of a transformed random vector from the density function of the original random vector. Then we explain how to uniformly approximate the density function by the empirical measure, resulting in Theorem 4.6, the probabilistic toolbox. We also define the notion of Vapnik-Chervonenkis dimension, or short VC dimension, and give some examples and properties that we use in later chapters.

Let $\Omega \subseteq \mathbb{R}^k$ be a space with a metric d, a measure vol and a probability measure μ. Recall that μ has a *density function* $f : \Omega \to \mathbb{R}$ if the probability $\mu(E)$ of all events $E \subseteq \Omega$ can be computed as the integral over the density function $\int_E f(x)dx$. The function f is non-negative and we have $\int_\Omega f(x)dx = 1$. An important example is the *uniform distribution* on a measurable set $S \subset \mathbb{R}^k$ of finite positive volume. Its density function equals $\chi_S(x)/\operatorname{vol}(S)$ where χ_S denotes the characteristic function that equals 1 at points from S and 0 at points from $\mathbb{R}^k \setminus S$.

At this point, a remark about measurability seems appropriate. There are two approaches in the literature to define measures. The first is to define a measure μ on all subsets of a space Ω. A measure is then defined as a non-negative, subadditive function to the extended real line. A set $E \subseteq \Omega$ is said to be measurable if for all subsets $A \subseteq \Omega$, we have $\mu(E) = \mu(E \cap A) + \mu(E \setminus A)$. The set of all measurable sets forms a σ-algebra. We follow this approach, which is taken in the classic book of Federer [28], for example.

The other approach is to define a measure only on the sets of some σ-algebra such that all sets in this σ-algebra are measurable. A measure is then defined to be a non-negative, countably additive function to the extended real line. With the second approach, one often has to be careful to formulate results for measurable sets only. This caution is not necessary in the same way when taking the first approach. Both approaches lead to the same results, the difference is a difference of language only.

4.1. Transformation of density functions

The following transformation theorem for density functions can be found in its basic version in many probability theory books; we cite it from Hoffmann-Jørgensen [34, Volume II], adapting the notation to our usage.

Let X be an absolutely continuous random vector on $\mathbb{R}^k$ with density function f_X and probability measure P such that $P(X \in D) = 1$ for some open set $D \subset \mathbb{R}^k$. Let $\varphi : D \to \mathbb{R}^k$ be a continuously differentiable, injective map such that $J_\varphi(u) > 0$ for all $u \in D$ where the *Jacobian* $J_\varphi(u)$ is defined to be $\left|\det\left(\left(\frac{\partial \varphi_i}{\partial x_j}(u)\right)_{i,j=1,\dots,k}\right)\right|$. Then $Y = \varphi(X)$ is an absolutely continuous random vector on $\mathbb{R}^k$ with density function

$$f_Y(v) = \begin{cases} \frac{f_X(\varphi^{-1}(v))}{J_\varphi(\varphi^{-1}(v))} & \text{for all } v \in \varphi(D) \\ 0 & \text{for all } v \notin \varphi(D) \end{cases} .$$

This theorem follows from the change of variable formula for multiple integrals. We need a slightly more general version of the theorem for Borel sets D, which follows from the following special case of [34, Section 8.7].

Theorem 4.1. *Let $T : D \to \mathbb{R}^d$ be an almost everywhere smooth and injective Borel function, where D is a Borel subset of $\mathbb{R}^d$. Let $B \subseteq D$ be a Borel set and let $g : \mathbb{R}^d \to [0, \infty]$ be a Borel function. Then*

$$\int_B g(T(x)) J_T(x) dx = \int_{T(B)} g(y) dy.$$

4.2. Uniform approximation of the density function by the empirical measure

In a first step, we would like to approximate the values of the density function at a point $x \in \Omega$ by the probability of a small set containing x. For sufficiently nice functions f and small sets B containing x, the probability $\mu(B)$ is close to $f(x) \cdot \mathrm{vol}(B)$, implying that $f(x)$ is close to $\mu(B)/\mathrm{vol}(B)$. We make this precise in Proposition 4.2. As small sets containing x we use balls of a small radius δ, centered at x. In a metric space, we denote the closed ball of radius δ that is centered at x by $B(x, \delta)$.

Recall that a function g from a metric space (Ω, d) to $\mathbb{R}$ is called Lipschitz continuous if there is a constant L such that for all $x, y \in \Omega$ it holds $|g(x) - g(y)| \leq L\, d(x, y)$. The constant L is called a Lipschitz constant of f.

The following proposition implies that, if f is Lipschitz continuous, then the ratio $\frac{\mu(B(x,\delta))}{\mathrm{vol}(B(x,\delta))}$ converges uniformly to the value of the density function $f(x)$ as δ tends to zero.

Proposition 4.2. *Let Ω be a space with a metric d and a probability measure μ. Let $f : \Omega \to \mathbb{R}$ be a Lipschitz continuous density function of μ with Lipschitz constant L.*

Then, for all $x \in \Omega$ and $\delta > 0$,

$$\left| \frac{\mu(B(x,\delta))}{\mathrm{vol}(B(x,\delta))} - f(x) \right| \leq L\delta.$$

Proof. Let $x \in \Omega$. For all $y \in B(x,\delta)$, we have $|f(x)-f(y)| \leq L\delta$. Since $\mu(B(x,\delta)) = \int_{B(x,\delta)} f(y)dy$, it holds that

$$(f(x) - L\delta)\,\mathrm{vol}(B(x,\delta)) \leq \mu(B(x,\delta)) \leq (f(x) + L\delta)\,\mathrm{vol}(B(x,\delta)). \qquad \square$$

In a second step, we approximate the probability measure μ by the empirical measure μ_N, which we now define. Let $X_1, \dots, X_N$ be a sequence of independent and identically distributed (i.i.d.) random vectors with common distribution μ. The empirical measure μ_N of a measurable set C is defined as

$$\mu_N(C) = \frac{|\{i : X_i \in C\}|}{N} = \frac{1}{N}\sum_{i=1}^{N} \chi_C(X_i).$$

On a basic level, the deviation from the expected value can be controlled by the well-known Chernoff bound.

Theorem 4.3 (Chernoff bound)**.** *Let $X_1, \dots, X_N$ be independent binary random variables, let $X = \sum_{i=1}^{N} X_i$ and let $\varepsilon \in (0,1)$. Then*

$$P(|X - \mathbf{E}X| > \varepsilon N) < 2e^{-\varepsilon^2 N/2}$$

where $\mathbf{E}X$ is the expected value of X.

We need a tool that allows us to control not only the error at one point, but over a whole space, that is, the supremum of the error. We cite two theorems [25] that allow us to do so.

Together they tell us in which sense μ_N tends to μ as N tends to infinity. The first theorem states that, with high probability, the supremum of the error over a set system does not deviate much from its expected value.

Theorem 4.4 (Chapter 3, [25])**.** *Let $N \in \mathbb{N}$ and let $X_1, \dots, X_N$ be i.i.d. random variables taking values in $\mathbb{R}^k$ with common distribution μ and empirical distribution μ_N. Let $\mathcal{C}$ be a class of subsets of $\mathbb{R}^k$. For all $\tau > 0$,*

$$P\left(\left| \sup_{C \in \mathcal{C}} |\mu_N(C) - \mu(C)| - \mathbf{E}\left(\sup_{C \in \mathcal{C}} |\mu_N(C) - \mu(C)| \right) \right| > \tau \right) \leq 2e^{-2N\tau^2}.$$

For the formulation of the second theorem, we need to know what the Vapnik-Chervonenkis dimension is [52]. For a class $\mathcal{C}$ of subsets of $\mathbb{R}^k$, define the *Vapnik-Chervonenkis shatter coefficient* to be

$$S_{\mathcal{C}}(m) = \max_{x_1,\dots,x_m \in \mathbb{R}^k} |\{\{x_1, \dots, x_m\} \cap C : C \in \mathcal{C}\}|.$$

The *Vapnik-Chervonenskis dimension*, or shortly *VC dimension*, of a class $\mathcal{C}$ is defined as the largest integer m such that $S_{\mathcal{C}}(m) = 2^m$. We denote it by $\dim(\mathcal{C}) = m$. We say that $\mathcal{C}$ *shatters* $\{x_1, \ldots, x_m\}$ if $|\{\{x_1, \ldots, x_m\} \cap C : C \in \mathcal{C}\}| = 2^m$.

The next theorem says that, if the VC dimension of the set system is finite, then the expected value of the supremum of the approximation error over a set system converges to 0 as the number of random samples tends to infinity.

Theorem 4.5 (Chapter 4, [25])**.** *Let $X_1, \ldots, X_N$ be i.i.d. random variables taking values in $\mathbb{R}^k$ with common distribution μ and empirical distribution μ_N. Let $\mathcal{C}$ be a class of subsets of $\mathbb{R}^k$ with VC dimension V. Then*

$$\mathbf{E}\left(\sup_{C \in \mathcal{C}} |\mu_N(C) - \mu(C)|\right) \leq c\sqrt{V/N}$$

where c is a universal constant.

The constant c could be bounded explicitly.

Together, Proposition 4.2, Theorems 4.4 and 4.5 quantify the following intuitive statement:

$$f(x) \underset{\text{small } \delta}{\approx} \frac{\mu(B(x,\delta))}{\mathrm{vol}(B(x,\delta))} \underset{\text{large } N}{\approx} \frac{\mu_N(B(x,\delta))}{\mathrm{vol}(B(x,\delta))}.$$

Theorem 4.6 makes this precise. For optimization problems whose objective function can be expressed as the Lipschitz continuous density function of some probability distribution, this yields a method for solving these problems approximately by random sampling.

For the shape matching application, given two shapes A and B, we are interested in finding a translation or rigid motion f for which the volume of overlap of $f(A)$ and B is close to maximal. We will express this function that maps a transformation to the volume of overlap as a density function that can easily be sampled.

We summarize the result in the following longish theorem that specifies under which assumptions and in which way the maximal value of the density function can be approximated by a sample point of maximal empirical measure.

Theorem 4.6 (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 there is a $v_\delta > 0$ with $\mathrm{vol}(B(x,\delta)) = v_\delta$ for all $x \in \Omega$. 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 $X_1, \ldots, X_N$ be i.i.d. random variables taking values in Ω with common distribution μ and empirical measure μ_N.

Let $B^* \in \mathcal{B}$ a random set such that $\mu_N(B^*) = \max\{\mu_N(B(X_i,\delta)) : 1 \leq i \leq N\}$, and let the random vector X^* be the center of B^*. Then for all $\tau > 0$ with probability $\geq 1 - 2e^{-2N\tau^2}$ the following inequality is true:

$$f(X^*) \geq f(x) - 2(c\sqrt{V/N} + \tau)/v_\delta - 3L\delta \quad \textit{for all } x \in \Omega. \tag{4.1}$$

In this inequality, c is the universal constant from Theorem 4.5.

Proof. Because of the Lipschitz continuity of f and Proposition 4.2, for all $B \in \mathcal{B}$ with center x holds

$$|f(x) \; - \; \mu(B)/v_\delta| \quad \leq \quad L\,\delta. \tag{4.2}$$

By the triangle inequality and Theorems 4.4 and 4.5, we know that with probability at least $1 - 2e^{-2N\tau^2}$ for all $B \in \mathcal{B}$

$$|\mu(B) \; - \; \mu_N(B)| \quad \leq \quad c\sqrt{V/N} + \tau. \tag{4.3}$$

Let $\eta = (c\sqrt{V/N}+\tau)/v_\delta$. If $f(x) \leq \eta + L\delta$ for all $x \in \Omega$, then Equation (4.1) is true because f is non-negative. Assume that there is an $x \in \Omega$ such that $f(x) > \eta + L\delta$. Let $\bar{x} \in \Omega$ be a point for which f is maximal. If f does not attain its supremum, then let $\varepsilon' > 0$ and let $\bar{x}$ be a point such that $f(\bar{x}) > \eta + L\delta$ and $f(\bar{x}) + \varepsilon' > f(x)$ for all $x \in \Omega$.

Let $\bar{B} = B(\bar{x}, \delta)$. Because of $f(\bar{x}) \leq \mu_N(\bar{B})/v_\delta + \eta + L\delta$ by the triangle inequality and Equations (4.2) and (4.3), we have $\mu_N(\bar{B}) > 0$. Therefore there is an $i \in \{1, \ldots, N\}$ such that $X_i \in \bar{B}$. Since f is Lipschitz continuous, we have

$$|f(\bar{x}) - f(X_i)| \leq L\delta. \tag{4.4}$$

We put things together now. First we use Inequality (4.4). Second we use Inequalities (4.2), (4.3) and the triangle inequality. Third we use the assumption on B^*. Lastly we use Inequalities (4.2), (4.3) and the triangle inequality again. We get that with probability at least $1 - 2e^{-2N\tau^2}$ the following holds:

$$\begin{aligned} f(\bar{x}) &\leq f(X_i) + L\delta \\ &\leq \mu_N(B(X_i, \delta))/v_\delta + \eta + 2L\delta \\ &\leq \mu_N(B^*)/v_\delta + \eta + 2L\delta \\ &\leq f(X^*) + 2\eta + 3L\delta. \end{aligned}$$

Since this holds for all small $\varepsilon' > 0$, it implies $f(x) \leq f(X^*) + 2\eta + 3L\delta$ for all $x \in \Omega$ and thus Equation (4.1). □

Since Equations (4.2), (4.3) and (4.4) bound the absolute difference, choosing a sample point whose neighborhood has minimal empirical measure approximates the minimal value of the density function in a way completely analogous to Equation (4.1). More generally, the fraction of random samples in a neighborhood divided by the volume of the neighborhood is an approximation of the value of the density function at the center of the neighborhood for any neighborhood because of the uniform convergence.

We generalize Theorem 4.6 to cases where the objective function has the form $\frac{f}{g}$ for a density function f as above and a Lipschitz continuous function g that is positive and bounded away from zero.

Furthermore, we require no longer that all balls of radius δ have the same volume, but only that the volume is bounded away from zero.

Theorem 4.7 (Refined probabilistic toolbox)**.** *Let $\Omega \subseteq \mathbb{R}^k$ be a metric space, and let $\mathcal{B}$ be a set of balls of some fixed radius $\delta > 0$ in Ω. Let* vol *be a measure on Ω such that there is a $v_\delta > 0$ with $\mathrm{vol}(B(x,\delta)) \geq v_\delta$ for all $x \in \Omega$. 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_f, and let f be bounded from above by M_f. 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 $C > 0$ and let $g : \Omega \to [C, +\infty)$ be a Lipschitz continuous function with Lipschitz constant L_g.

Let $B^* \in \mathcal{B}$ a random set such that $\frac{\mu_N(B^*)}{g(X^*)\,\mathrm{vol}(B(X^*,\delta))} = \max\Big\{ \frac{\mu_N(B(X_i,\delta))}{g(X_i)\,\mathrm{vol}(B(X_i,\delta))} : i = 1, \dots, N \Big\}$ where the random vector X^* is the center of B^*. Then, for all $\tau > 0$, with probability $\geq 1 - 2e^{-2N\tau^2}$, for all $x \in \Omega$

$$\frac{f(X^*)}{g(X^*)} \geq \frac{f(x)}{g(x)} - \Big(\frac{2(c\sqrt{V/N} + \tau)}{C\, v_\delta} + \frac{3L_f\delta}{C} + \frac{M_f L_g \delta}{C^2} \Big). \tag{4.5}$$

In this inequality, c is the universal constant from Theorem 4.5.

Proof. The proof follows the same lines as the proof of Theorem 4.6. Abbreviate $\eta = (c\sqrt{V/N} + \tau)/v_\delta$. For all $i = 1, \dots, N$ and for all $\tau > 0$, with probability $\geq 1 - 2e^{-2N\tau^2}$, the following inequality is true, similar to Inequalities (4.2) and (4.3),

$$\Big| \frac{f(X_i)}{g(X_i)} - \frac{\mu_N(B(X_i,\delta))}{g(X_i)\,\mathrm{vol}(B(X_i,\delta))} \Big| \leq \frac{\eta + L_f\delta}{C}. \tag{4.6}$$

If $\frac{f(x)}{g(x)} \leq \frac{\eta + L_f\delta}{C}$ for all $x \in \Omega$, then Equation (4.5) is trivially true. Otherwise, let $\varepsilon' > 0$ and let $\bar{x} \in \Omega$ be such that $\frac{f(\bar{x})}{g(\bar{x})} > \frac{\eta + L_f\delta}{C}$ and $\frac{f(\bar{x})}{g(\bar{x})} + \varepsilon' > \frac{f(x)}{g(x)}$ for all $x \in \Omega$. We show that Equation (4.5) is true for $\bar{x}$. This implies that it is true for all $x \in \Omega$. Let $\bar{B} = B(\bar{x}, \delta)$. As in the proof of Theorem 4.6, we have $\mu_N(\bar{B}) > 0$ and therefore there is an $i \in \{1, \dots, N\}$ such that $X_i \in \bar{B}$. We estimate the difference between $\frac{f(\bar{x})}{g(\bar{x})}$ and $\frac{f(X_i)}{g(X_i)}$.

$$\begin{aligned}
\Big| \frac{f(\bar{x})}{g(\bar{x})} - \frac{f(X_i)}{g(X_i)} \Big| &\leq \Big| \frac{f(\bar{x})}{g(\bar{x})} - \frac{f(X_i)}{g(\bar{x})} \Big| + \Big| \frac{f(X_i)}{g(\bar{x})} - \frac{f(X_i)}{g(X_i)} \Big| \\
&\leq \frac{L_f\delta}{C} + \frac{f(X_i)\,|g(X_i) - g(\bar{x})|}{g(\bar{x})\,g(X_i)} \\
&\leq \frac{L_f\delta}{C} + \frac{M_f L_g \delta}{C^2}.
\end{aligned}$$

Using this estimate and Equation (4.6), we prove Equation (4.5):

$$\begin{aligned}\frac{f(X^*)}{g(X^*)} &\geq \frac{\mu_N(B^*)}{g(X^*)\ \mathrm{vol}(B(X^*,\delta))} - \frac{\eta + L_f\delta}{C}\\ &\geq \frac{\mu_N(B(X_i,\delta))}{g(X_i)\ \mathrm{vol}(B(X_i,\delta))} - \frac{\eta + L_f\delta}{C}\\ &\geq \frac{f(X_i)}{g(X_i)} - \frac{2(\eta + L_f\delta)}{C}\\ &\geq \frac{f(\bar{x})}{g(\bar{x})} - \frac{2\eta}{C} - \frac{3L_f\delta}{C} - \frac{M_f L_g\delta}{C^2}.\end{aligned}$$

□

To apply these theorems, we have to bound the VC dimension of classes of balls of fixed radius δ in a metric space. In preparation of this, we mention two easy results about the VC dimension.

Note that the VC dimension of a subclass of $\mathcal{C}$ can be only smaller than $\dim(\mathcal{C})$. Therefore, the following lemma bounds the VC dimension whenever our space at hand is some $\mathbb{R}^k$ with the metric that is induced by the maximum norm.

Lemma 4.8. *[25, Lemma 4.1] The VC dimension of the class of all rectangles in* $\mathbb{R}^k$ *equals* $2k$*.*

Lemma 4.9. *[26] The VC dimension of the class of all closed Euclidean balls in* $\mathbb{R}^k$ *equals* $k+1$*.*

5. Probabilistic matching under translations in arbitrary dimension

In this chapter, we design and analyze a simple probabilistic algorithm for matching shapes in $\mathbb{R}^d$ under translations. Given two shapes A and B, the algorithm computes a translation t^* such that the volume of overlap of $t^*(A)$ and B is approximately maximal.

The algorithm works as follows. Given two shapes A and B, we pick a point $a \in A$ and a point $b \in B$ 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 5.1.

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 tolerance ε and an allowable probability of failure p, both between 0 and 1, we derive bounds on the required number of random experiments, guaranteeing that the difference between approximation and optimum $|t^{\text{opt}}(A) \cap B| - |t^*(A) \cap B|$ is not larger than $\varepsilon|A|$ with probability at least $1 - p$. As before, $|\cdot|$ denotes the volume (Lebesgue measure) of a set. We use $\varepsilon|A|$ and not just ε as an error bound because the inequality should be invariant under scaling of both shapes with the same factor.

This algorithm is an instance of a probabilistic algorithmic scheme for approximating an optimal match of planar sets under a subgroup of affine transformations. Alt and Scharf [5] analyzed another instance of this algorithmic scheme that compares polygonal curves under translations, rigid motions, and similarities.

A preliminary version of the results in this chapter appeared in [9], together with Helmut Alt and Ludmila Scharf. In that paper, the algorithm was described and analyzed only for shapes in the plane. Here we not only generalize the results to higher dimensions, but also give new and simpler proofs, using the method that we presented in Chapter 4. We also slightly improve the bounds on the number of random samples for translations. Furthermore we improve the time complexity of the algorithm by showing that a simpler definition of a cluster suffices to guarantee approximation. In [9], a translation (not necessarily a vote) whose neighborhood contained the maximal number of votes was computed, which boils down to computing a deepest cell of an arrangement of boxes. For N boxes in $\mathbb{R}^d$, the best known time bound for solving this problem is $O(N^{d/2}/(\log N)^{d/2-1} \operatorname{polyloglog} N)$ [19]. 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 $O\big(N(\log N)^{d-1}\big)$ by using orthogonal range counting queries [22].

For two simple polygons with at most n vertices in the plane, Mount et al. [43] show that a translation that maximizes the area of overlap can be computed in time $O(n^4)$. Ahn et al. [1] state that the maximal volume of overlap of two sets of n simplices in $\mathbb{R}^d$ under translations can be computed in time $O(n^{2d})$ by a generalization of Mount et al.'s algorithm to higher dimensions.

Cheong et al. [20] 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. Their algorithm works with high probability. The method of Cheong et al. bears some similarity to our approach. For shapes A and B, they sample a set of points S uniformly from shape A. For each sample point $s \in S$, they compute the set of translations $W(s)$ that map s into B. For every $s \in S$, the set $W(s)$ is a translate of B. They return a point from a deepest cell in the arrangement of all regions $W(s)$ as a translation that approximates the maximal area of overlap.

In contrast to their approach, we do not need to compute the deepest cell of an arrangement. Furthermore our method does not require the shapes to be polygonal regions but works for a larger class of shapes. Apart from measurability assumptions, we only require that uniformly distributed random points can be generated from the shapes; see Chapter 2 for details on the shape model that we use. For unions of simplices, the runtime of our method depends only linearly on the number of vertices in any dimension, but more significantly on the error ε and a fatness parameter, namely the isoperimetric quotient.

5.1. The algorithm

We start by introducing some notation. A translation $f_t : \mathbb{R}^d \to \mathbb{R}^d,\ f_t : x \mapsto x + t$ is identified with its translation vector t, so the translation space equals $\mathbb{R}^d$. Observe that an ordered pair of points in $\mathbb{R}^d$ uniquely determines a translation mapping the first point onto the second. The set of translations that map some point of $A \subset \mathbb{R}^d$ onto some point of $B \subset \mathbb{R}^d$ equals the Minkowski sum $B \oplus (-A) = \{b - a : a \in A, b \in B\}$. We denote the Euclidean distance by $|\cdot|$, and the maximum norm by $||\cdot||_\infty$. The closed ball around $t \in \mathbb{R}^d$ of radius δ with respect to $||\cdot||_\infty$ is denoted by $B(t, \delta) = \{s \in \mathbb{R}^d : ||t - s||_\infty \leq \delta\}$.

Given an integer N and a positive number δ, the algorithm draws N times points $a \in A$ and $b \in B$ uniformly at random (u.a.r.) and stores the N random translations $b - a$ in the translation space. Then it determines a translation t^* among the N random translations such that the cube of side length 2δ, centered at t^*, contains the maximum number of the N random translations. The translation t^* is the result of the algorithm. We also output the number of votes that it received, which is an approximation of the maximal volume of overlap when suitably normalized. Note that, in contrast to [9], the output translation t^* is one of the votes, which simplifies its computation. First we state a heuristic version of the algorithm. On the input of a shape A, the function RandomPoint(A) returns a point from A uniformly at random.

Algorithm 1: ProbMatchT(A, B, δ, N)

```
Input: shapes A, B ⊂ ℝ^d, clustering size δ > 0, sample size N ∈ ℕ
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: a pair of a point t in Q and an integer V such that the cube of side length 2δ that is centered at t contains a maximal number V of points from Q

To turn Algorithm 1 from a heuristic to an approximation algorithm, we modify the input. Algorithm 2 gets more meaningful parameters as input: along with the shapes A and B, an error tolerance ε and an allowable probability of failure p are given as input. The algorithm computes a clustering size δ and a sample size N that guarantee that the additive error of the output translation is less than $\varepsilon|A|$ with probability at least $1 - p$.

It turns out that such a clustering size δ can be computed depending only on A and ε, and such a sample size can be computed depending on B, δ, ε, and p. A smaller clustering size enforces a larger sample size. Of course, the roles of A and B can be interchanged. This may result in a smaller N, which would clearly improve the runtime of the algorithm. For the shortness of presentation, we do not reflect this fact in the following.

The algorithm also outputs a certain multiple of the number of votes that the output translation receives by the run of Algorithm 1. This number approximates the volume of overlap when A is translated by the output translation. It is used for matching under rigid motions in the next chapters.

The clustering size δ has to be sufficiently small to guarantee approximation; it depends on the shapes how small δ has to be. Skinny shapes require a smaller

Algorithm 2: MaxOverlapT

Input: shapes $A, B \subset \mathbb{R}^d$, error tolerance $\varepsilon \in (0,1)$, allowable probability of failure $p \in (0,1)$
real $\delta \leftarrow$ ClusteringSizeT(A, ε);
integer $N \leftarrow$ SampleSizeT$(B, \varepsilon, \delta, p)$;
(point,integer) $(t^*, V) \leftarrow$ ProbMatchT(A, B, δ, N);
`// Compute an estimate of` $|t^*(A) \cap B|/(|A|\,|B|)$
real $V^* \leftarrow V(2\delta)^{-d}N^{-1}$;
return (t^*, V^*);

Function ClusteringSizeT(A, ε)

Input: shape A, parameter $\varepsilon \in (0,1)$
Output: positive number $\delta \leq \frac{2\sqrt{d}}{3(d+1)} \frac{\varepsilon|A|}{|\partial A|}$

clustering size than fat shapes. The relevant concept of fatness is perimeter-fatness, introduced in Chapter 2. The clustering size also depends on the absolute size of the shape since scaling a shape with a constant greater than 1 enlarges the shape's features and allows a larger clustering size. Given shapes A, B and an error tolerance ε, we prove that for all positive clustering sizes below some constant multiple of $\varepsilon \frac{|A|}{|\partial A|}$ there is a sample size N that guarantees approximation with high probability.

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. The clustering size scales with the shapes. This is reasonable because blowing up the shapes coarsens fine features.

The following theorem states that Algorithm 2 computes a probabilistic absolute error approximation of the maximal volume of overlap under translations.

Theorem 5.1 (Approximation property of Algorithm 2)**.** *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters.*

Algorithm 2 with input (A, B, ε, p) computes a pair of a translation t^* and a number V^* such that, with probability at least $1-p$, we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all translations $t \in \mathbb{R}^d$, and $\big|V^*|A|\,|B| - |t^*(A) \cap B|\big| \leq \frac{\varepsilon}{2}|A|$.

Before proving this theorem, we determine the runtime of the algorithm. Recall that the parameters K_A, m_A, M_B are given for the shapes by the assumptions described in Section 2.2.

Theorem 5.2 (Runtime of Algorithm 2)**.** *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters.*

Algorithm 2 with input (A, B, ε, p) runs in time $O\big(T(N) + N(\log N)^{d-1}\big)$ where $N = O\big(\varepsilon^{-(2d+2)}(K_A M_B/m_A)^2 \log \frac{2}{p}\big)$.

Function SampleSizeT$(B, \varepsilon, \delta, p)$

Input: shape B, parameters $\varepsilon, p, \delta > 0$
Output: integer $N \geq C\varepsilon^{-2}\,\delta^{-2d}\,|B|^2 \log\frac{2}{p}$ for some universal constant $C > 0$

Proof of Theorem 5.2.

- **$2N$ calls of Function RandomPoint:**
  By assumption, N uniformly distributed random points can be generated from a shape in time $T(N)$. For the case of unions of simplices, see Chapter 2.

- **Functions ClusteringSizeT(A, ε) and SampleSizeT$(B, \varepsilon, \delta, p)$:**
  By assumption, we are given an upper bound K_A on the isoperimetric quotient of A, a lower bound m_A on $|A|$, and an upper bound M_B on $|B|$. In the case of unions of simplices, we can compute $|A|$ and $|B|$ in linear time as a preprocessing step and assume that $|B|$ is at most a constant multiple of $|A|$. Hence for all shapes that fulfill the assumption described in Section 2.2, we can compute in constant time a clustering size

$$\delta = \tfrac{2\sqrt{d}}{3(d+1)}\,\varepsilon\,(m_A/K_A)^{1/d} \leq \tfrac{2\sqrt{d}}{3(d+1)}\,\varepsilon\,\tfrac{|A|}{|\partial A|}$$

  and a sample size

$$\begin{aligned} N &= C(d+1)^2\left(\tfrac{3(d+1)}{2\sqrt{d}}\right)^{2d} K_A^2\,(M_B/m_A)^2\,\varepsilon^{-2d-2}\log\tfrac{2}{p} \\ &\geq C(d+1)^2\,\varepsilon^{-2}\,\delta^{-2d}\,|B|^2\,\log\tfrac{2}{p}. \end{aligned}$$

- **Function FindDensestClusterT(Q, δ):**
  Let $Q = \{t_1, \ldots, t_N\}$. A translation t_i maximizing $|B(t_i, \delta) \cap Q|$ can be computed in time $O\big(N(\log N)^{d-1}\big)$ for all $t_1, \ldots, t_N \in \mathbb{R}^d$ and $\delta > 0$. For this, we build a standard orthogonal range counting query data structure for $t_1, \ldots, t_N$ in time $O\big(N(\log N)^{d-1}\big)$. Then we query the data structure N times with cubes of side length 2δ, centered at the t_i. Each query takes time $O\big((\log N)^{d-1}\big)$. For details and references on range counting queries, see for example the book [22].

□

5.2. Correctness proof

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. 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. Then the goal is to find a translation at which the density function is approximately maximal.

Conceptually, the density function f is approximated in a two step process. 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 Theorem 4.6, which we proved in Chapter 4. We first prove that the density function of μ is proportional to the objective function.

In the following, for a function φ on $\mathbb{R}^d \times \mathbb{R}^d$, we may write $\varphi(a,b)$ for a function value, omitting the double brackets.

Lemma 5.3. *Let Z be the random vector on $\mathbb{R}^d$ that draws translations $b-a$ where $(a,b) \in A \times B$ is drawn from the uniform distribution. The density function of Z is given by $f(t) = \frac{|t(A)\cap B|}{|A|\,|B|}$.*

Proof. Let X be a random vector on $\mathbb{R}^d \times \mathbb{R}^d$ that is uniformly distributed on $A \times B$; its density function f_X is given by $f_X(x) = \chi_{A\times B}(x)/(|A|\,|B|)$. Let $\varphi : A \times B \to A \times (-A \oplus B)$ be the function defined by $\varphi(a,b) = (a, b-a)$. Denote by I_d the $(d \times d)$-identity matrix. The function φ is linear with matrix representation $J = \begin{pmatrix} I_d & 0 \\ -I_d & I_d \end{pmatrix}$, and its determinant equals 1. By Section 4.1, the density function of $Y = \varphi(X)$ equals $f_Y(a,t) = \chi_{A\times B}(a, a+t)/(|A|\,|B|)$. Let π be the projection defined by $\pi : A \times (-A \oplus B) \to (-A \oplus B)$, $\pi : (a,t) \mapsto t$, and define $Z = \pi(Y)$. The density function f_Z of Z equals $\int_{\mathbb{R}^d} f_Y(a,t)da$.

Since $\chi_{A\times B}(a, a+t) = \chi_A(a)\chi_{B-t}(a) = \chi_{A\cap(B-t)}(a)$, we have

$$f_Z(t) = \int_{\mathbb{R}^d} \frac{\chi_{A\cap(B-t)}(a)}{|A|\,|B|} da = \frac{|A \cap (B-t)|}{|A|\,|B|} = \frac{|(A+t)\cap B|}{|A|\,|B|}.$$

□

Next we deduce from Corollary 3.5 that f is Lipschitz continuous with respect to $||\cdot||_\infty$. Applying the Cauchy-Schwarz inequality yields $|t-s| \le \sqrt{d}\,||t-s||_\infty$, which is best possible.

Corollary 5.4. *The function f on $\mathbb{R}^d$, given by $f(t) = \frac{|t(A)\cap B|}{|A|\,|B|}$, is Lipschitz continuous with constant $L = \frac{\sqrt{d}|\partial A|}{2|A|\,|B|}$ with respect to the metric that is induced by the maximum norm.*

Now we are ready to prove Theorems 5.1 and 5.2.

Proof of Theorem 5.1. We apply our probabilistic toolbox Theorem 4.6 to the translation space $\mathbb{R}^d$, equipped with the metric that is induced by the maximum norm. Let $\mathcal{B}_\delta$ be the set of closed balls that have radius δ. For every $B \in \mathcal{B}_\delta$, we have $|B| = 2^d\delta^d$. By Lemma 4.8, the VC dimension of $\mathcal{B}_\delta$ is at most $2d$.

Let $X_1, X_2, \dots, X_N$ be the random translations in the algorithm. Denote their common distribution by μ. By Lemma 5.3, μ has a density function f that is given by $f(t) = \frac{|t(A)\cap B|}{|A|\,|B|}$, and by Lemma 5.4, f is Lipschitz continuous with constant $L = \frac{\sqrt{d}|\partial A|}{2|A|\,|B|}$.

Let μ_N be the empirical measure of μ. Algorithm 1 computes a $B^* \in \mathcal{B}_\delta$ such that $\mu_N(B^*) = \max\{\mu_N(B(X_i, \delta)) : i = 1, \dots, N\}$. Let t^* be the center of B^*. By Theorem 4.6, for all $\tau > 0$, with probability at least $1 - 2e^{-2N\tau^2}$, for all $t \in \mathbb{R}^d$, we have

$$f(t^*) \geq f(t) - \left(\frac{2c\sqrt{2d}}{\sqrt{N}(2\delta)^d} + \frac{2\tau}{(2\delta)^d} + 3L\delta \right) \tag{5.1}$$

for some constant $c > 0$.

Let $\varepsilon, p \in (0,1)$. We want to keep N as small as possible as the runtime of our algorithm grows with N. We now determine the minimal value of N and a compatible value of δ for that we can guarantee that with probability at least $1-p$ the error is at most $\varepsilon|A|$. For the sake of clarity, we compute N only up to a constant.

Since $p \in (0,1)$ is the allowable probability of failure, we should have $2e^{-2N\tau^2} \leq p$, which is equivalent to $\tau \geq \sqrt{\frac{1}{2N}\log\frac{2}{p}}$. The larger τ is, the larger is the error term in Inequality (5.1), so we set $\tau = \sqrt{\frac{1}{2N}\log\frac{2}{p}}$.

Substituting τ, L and the definition of f shows that

$$|t^*(A) \cap B| \geq |t(A) \cap B| - \left(\frac{\sqrt{2}\,|A|\,|B|}{\sqrt{N}(2\delta)^d} \left(2c\sqrt{d} + \sqrt{\log\frac{2}{p}} \right) + \frac{3}{2}\sqrt{d}\,\delta|\partial A| \right). \tag{5.2}$$

We want the additive error in Inequality (5.2) to be at most $\varepsilon|A|$. Therefore it is necessary that $\delta < \frac{2\varepsilon|A|}{3\sqrt{d}|\partial A|}$. Let $\eta \in (0,1)$ be such that $\delta = \eta\frac{2\varepsilon|A|}{3\sqrt{d}|\partial A|}$. Then as condition for N we get the following:

$$\frac{\sqrt{2}\,|A|\,|B|}{\sqrt{N}(2\delta)^d} \left(2c\sqrt{d} + \sqrt{\log\frac{2}{p}} \right) < (1-\eta)\varepsilon|A|.$$

There is a universal constant $C > 0$, such that this condition is fulfilled if

$$N \geq C(1-\eta)^{-2}\varepsilon^{-2}\,\delta^{-2d}\,|B|^2\,\log\tfrac{2}{p}.$$

Differentiating with respect to η shows that $\eta = \frac{d}{d+1}$ minimizes the bound on N, so we use this value, which gives $\delta \leq \frac{2\sqrt{d}}{3(d+1)}\frac{\varepsilon|A|}{|\partial A|}$.

A look at Inequalities (4.2) and (4.3) in the proof of Theorem 4.6 tells us that $|f(t^*) - \mu_N(B(t^*, \delta))/(2\delta)^d|$ is at most as large as half of the error term in Inequality (5.1). This means that $\big||t^*(A) \cap B| - |A|\,|B|V/((2\delta)^d N)\big|$ is as most as large as half of the error term in Inequality (5.2) where V is the number of votes that t^* received. □

5.3. Adaptive Sampling

We improve the probabilistic matching algorithm for translations by computing the output translation in two steps. First we compute a candidate region of translations that we prove to contain all optimal translations with high probability. We do so by matching the two input shapes with an error bound of ε^k for some $k \in (0,1)$ and keeping all sufficiently good translations. Second we search for a translation in the candidate region with error bound ε. We choose k later such that the runtime for both steps is balanced; it will turn out that it is best to choose $k = \frac{2d+2}{3d+1}$.

For the search in the candidate region, we show how to restrict the algorithm to certain regions in the translation space. For this, we need to be able to answer membership queries of the type "Is $b \in B$?" Therefore, in this section, we assume that in time $T_q(N)$ we cannot only generate N random points from B, but also answer N membership queries. For details, see Chapter 2.

The improvement in the runtime depends on the size of the candidate region: the more translations exist that match A and B with small error, the larger is the candidate region, and the larger the candidate region, the smaller is the improvement. More precisely, for a lower bound m_B on $|B|$ and an error bound ε, we show a time bound on the runtime of the algorithm in which the volume of the set

$$\mathcal{W} = \{t \in \mathbb{R}^d : |t^{opt}(A) \cap B| - |t(A) \cap B| \leq (2\frac{|B|}{m_B} + 1)\varepsilon^k |A|\} \tag{5.3}$$

appears as a multiplicative factor.

Despite the fact that the candidate region may be large, the adaptive algorithm has a better time bound in ε: the dependence on ε is $\varepsilon^{-(\frac{4}{3}d+\frac{8}{3})}$, ignoring log-factors, in contrast to $\varepsilon^{-(2d+2)}$, which we proved for Algorithm 2 in the previous section.

Computation of a candidate region

For the first step, the computation of a candidate region, we modify Algorithm 2. Instead of Function FindDensestClusterT, we call the following Function FindAllDenseClustersT. It returns a collection of translations. The error bound ε and a lower bound m_B on the volume of shape B are given as additional input to the function.

Function FindAllDenseClustersT($Q, \delta, \varepsilon, m_B$)

Input: collection Q of N points in $\mathbb{R}^d$, positive numbers δ, ε, m_B
Output: Subcollection Q' of Q, containing exactly those $q \in Q$ for which $V^* - V_q \leq \frac{2^{d+1}\delta^d \varepsilon^k N}{m_B}$ where $V_q = |B(q,\delta) \cap Q|$ and $V^* = \max_{q \in Q} V_q$.

Next we give a pseudocode description of the computation of a candidate region, Algorithm 3.

Let $A, B \subset \mathbb{R}^d$ be shapes in constant dimension d, and let $\varepsilon, p \in (0,1)$ be as usual.

Algorithm 3: MaxOverlapTCandidates(A, B, ε, p)

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

real $\delta \leftarrow$ ClusteringSizeT(A, ε^k);
integer $N \leftarrow$ SampleSizeT($B, \varepsilon^k, \delta, p$);
collection $Q \leftarrow \emptyset$;
for $i = 1 \ldots N$ **do**
    point $a \leftarrow$ RandomPoint(A);
    point $b \leftarrow$ RandomPoint(B);
    add($Q, b - a$);
end
collection $Q' \leftarrow$ FindAllDenseClustersT($Q, \delta, \varepsilon, m_B$); // $m_B \leq |B|$
return (Q', δ);

Lemma 5.5. *Assume that* $|t^{opt}(A) \cap B| \geq \frac{\varepsilon^k}{2}|A|$ *for all translations* t^{opt} *that maximize the volume of overlap of* A *and* B*. Then Algorithm 3 with input* (A, B, ε, p) *computes a collection of translations* Q' *such that* $\bigcup_{q \in Q'} B(q, \delta)$ *contains, with probability at least* $1 - p$*, all translations* t^{opt} *that match* A *and* B *optimally. For all* $t \in \bigcup_{q \in Q'} B(q, \delta)$*, we have* $|t^{opt}(A) \cap B| - |t(A) \cap B| \leq \big(2\frac{|B|}{m_B} + 1\big)\varepsilon^k |A|$.

Proof. We use the method developed in Chapter 4. For all $t \in \mathbb{R}^d$, define

$$V_t = |B(t, \delta) \cap Q| \quad \text{and} \quad e_t = \frac{V_t |A|\,|B|}{N(2\delta)^d}.$$

The number e_t is the empirical measure of $B(t, \delta)$, multiplied by $|A|\,|B|$, which makes it an estimate of $|t(A) \cap B|$. We use the notation of the proof of Theorem 5.1, but replace ε by ε^k. Let $\tau > 0$ and abbreviate $\eta = \frac{2c\sqrt{2d}}{\sqrt{N}(2\delta)^d} + \frac{2\tau}{(2\delta)^d}$. By Inequalities (4.2) and (4.3), we have that, with probability at least $1 - 2e^{-\tau^2 N/2}$, for all $t \in \mathbb{R}^d$,

$$\big|e_t - |t(A) \cap B|\big| \leq |A|\,|B|(\eta + L\delta). \tag{5.4}$$

Choosing $\tau = \sqrt{\frac{1}{2N} \log \frac{2}{p}}$, Inequality (5.4) is true with probability at least $1 - p$. There is a constant $C > 0$ such that, if $\delta \leq \frac{2\sqrt{d}}{3(d+1)} \frac{\varepsilon^k |A|}{|\partial A|}$ and $N \geq C\varepsilon^{-1}\,\delta^{-2d}\,|B|^2 \log \frac{2}{p}$, then it follows with the same argument as in the proof of Theorem 5.1 that

$$|A|\,|B|(2\eta + 3L\delta) \leq \varepsilon^k |A|. \tag{5.5}$$

We first prove that for all optimal solutions t^{opt} there exists a $q_0 \in Q'$ such that $t^{opt} \in B(q_0, \delta)$. Because of Inequalities (5.4) and (5.5), assuming that $|t^{opt}(A) \cap B| \geq \frac{\varepsilon^k}{2}|A|$ implies that $e_{t^{opt}} > 0$. This implies that there exists $q_0 \in B(t^{opt}, \delta) \cap Q$. Hence we have found $q_0 \in Q$ such that $t^{opt} \in B(q_0, \delta)$.

Next we show
$$V^* - V_{q_0} \leq \frac{2^{d+1}\delta^d\varepsilon^k N}{|B|} \leq \frac{2^{d+1}\delta^d\varepsilon^k N}{m_B},$$
which implies $q_0 \in Q'$. We have $\big||q_0(A) \cap B| - |t^{opt}(A) \cap B|\big| \leq |A|\,|B|L\delta$ by the Lipschitz continuity of f. Applying Inequality (5.4) to q_0 and using the Lipschitz continuity gives

$$e_{q_0} \geq |q_0(A) \cap B| - |A|\,|B|(\eta + L\delta) \geq |t^{opt}(A) \cap B| - |A|\,|B|(\eta + 2L\delta). \tag{5.6}$$

Observe that $e_{t^*} - e_{q_0} \leq 2\varepsilon^k|A|$ is equivalent to $V^* - V_{q_0} \leq \frac{2^{d+1}\delta^d\varepsilon^k N}{|B|}$ where $t^* \in Q$ such that $V^* = V_{t^*}$. We prove that $e_{t^*} - e_{q_0} \leq 2\varepsilon^k|A|$. Using Theorem 5.1, we have

$$|t^{opt}(A) \cap B| \geq |t^*(A) \cap B| - |A|\,|B|(2\eta + 3L\delta) \geq e_{t^*} - |A|\,|B|(3\eta + 4L\delta). \tag{5.7}$$

Inequalities (5.6) and (5.7) give

$$e_{q_0} \geq e_{t^*} - |A|\,|B|(4\eta + 6L\delta) \geq e_{t^*} - 2\varepsilon^k|A|.$$

Lastly we show

$$\big||t(A) \cap B| - |t^{opt}(A) \cap B|\big| \leq \Big(2\frac{|B|}{m_B} + 1\Big)\varepsilon^k|A|$$

for all $t \in \bigcup_{q \in Q'} B(q, \delta)$. Because of $|V_t - V_{q_0}| \leq \frac{2^{d+1}\delta^d\varepsilon^k N}{m_B}$, we have $|e_t - e_{q_0}| \leq 2\frac{|B|}{m_B}\varepsilon^k|A|$. Hence

$$\begin{aligned}
|t(A) \cap B| &\geq e_t - |A|\,|B|(\eta + L\delta) \\
&\geq e_{q_0} - |A|\,|B|(\eta + L\delta) - 2\tfrac{|B|}{m_B}\varepsilon^k|A| \\
&\geq |t^{opt}(A) \cap B| - |A|\,|B|(2\eta + 3L\delta) - 2\tfrac{|B|}{m_B}\varepsilon^k|A| \\
&\geq |t^{opt}(A) \cap B| - \big(2\tfrac{|B|}{m_B} + 1\big)\varepsilon^k|A|.
\end{aligned}$$

□

As a direct corollary of Proposition 3.13, we get a sufficient condition on ε such that the assumption $|t^{opt}(A) \cap B| \geq c\,\varepsilon^k|A|$ for some positive constant c is satisfied.

Corollary 5.6. *If $\varepsilon^k \leq \frac{|B|}{c|B \oplus (-A)|}$ for any $c > 0$, then there exists a translation $t \in \mathbb{R}^d$ such that $|t(A) \cap B| \geq c\,\varepsilon^k|A|$.*

Searching in the candidate cubes

Let Q' be as in Lemma 5.5. Denote the points in Q' by $t_1, \ldots, t_I$. We have $I \leq N$. For $i = 1, \ldots, I$, we search for good translations in $B(t_i, \delta)$, where $\delta \leq \frac{2\sqrt{d}}{3(d+1)}\frac{\varepsilon^k|A|}{|\partial A|}$. In fact, later we thin out the set Q' and search in fewer than I cubes $B(t_i, \delta)$.

We draw $(a, b) \in A \times B$ u.a.r., but accept only samples such that $||(b-a)-t_i||_\infty \leq \delta$. We generate the samples (a, b) by drawing $a \in A$ and $t \in B(0, \delta)$ u.a.r. and checking if $b = a + t_i + t \in B$. If $b \notin B$, we discard the sample. In this way, we generate samples u.a.r. from $\{(a, b) \in A \times B : b - a \in B(t_i, \delta)\}$.

We describe the modified version of Algorithm 2 for the second step of the approximation, the search for a solution in a candidate cube $B(q, \delta)$.

Algorithm 4: MaxOverlapTCube($A, B, q, \delta, \varepsilon, p$)

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

real $\delta_\varepsilon \leftarrow \delta\varepsilon^{1-k}$;
`// Compute the necessary number of samples`
integer $N \leftarrow C\varepsilon^{2kd-k-2d-2} \log \frac{2}{p}$; `//` C `is a positive constant`
collection $Q \leftarrow \emptyset$;
for $i = 1 \ldots N$ **do**
  point $a \leftarrow$ RandomPoint(A);
  point $t \leftarrow$ RandomPoint($B(q, \delta)$);
  point $b \leftarrow q + t + a$;
  if $b \in B$ **then**
    add($Q, b - a$);
  end
end
return *FindDensestClusterT(Q, δ_ε)*;

The set $\mathcal{W}$ is defined in Equation (5.3).

Lemma 5.7. *Assume that $|t^{opt}(A) \cap B| \geq \big(4\frac{|B|}{m_B} + 2\big)\varepsilon^k|A|$ for any optimal translation t^{opt}. Then, for all $q \in \mathcal{W}$ and for all $\delta \in \big(0, \frac{2\sqrt{d}}{3(d+1)} \frac{\varepsilon^k|A|}{|\partial A|}\big]$, Algorithm 4 with input $(A, B, q, \delta, \varepsilon, p)$ computes a pair of a translation $t^* \in B(q, \delta)$ and a number V^* such that with probability at least $1 - p$ we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all $t \in B(q, \delta)$ and $\big|V^*|A||B| - |t^*(A) \cap B|\big| \leq \frac{\varepsilon}{2}|A|$.*

Proof. Not every sample generated in the algorithm is valid, that is, not every sample t satisfies $a + q + t \in B$. First we show that there is a constant C such that, if the algorithm generates $M \geq C\varepsilon^{2kd-2d-2} \log \frac{2}{p}$ valid samples, then the claim of the lemma is true.

By Lemma 5.5, for all $t \in B(q, \delta)$, we have

$$|t^{opt}(A) \cap B| - |t(A) \cap B| \leq \big(2\frac{|B|}{m_B} + 1\big)\varepsilon^k|A|.$$

By this and the assumption, we have

$$|t^{opt}(A) \cap B| \geq \big(4\frac{|B|}{m_B} + 2\big)\varepsilon^k|A| \geq 2|t^{opt}(A) \cap B| - 2|t(A) \cap B|,$$

implying $|t(A) \cap B| \geq \frac{1}{2}|t^{opt}(A) \cap B|$.

We restricted the random experiment to $B(q,\delta)$, so the density function is a function on $B(q,\delta)$, given by $f(t) = c\,|t(A) \cap B|$ where $c = \left(\int_{B(q,\delta)} |t(A) \cap B| dt\right)^{-1}$. Hence

$$2^d \delta^d |t^{opt}(A) \cap B| \;\geq\; \int_{B(q,\delta)} |t(A) \cap B|\, dt \;\geq\; 2^{d-1}\delta^d |t^{opt}(A) \cap B|.$$

This implies that the density function is given by $f(t) = c'\delta^{-d}\frac{|t(A)\cap B|}{|t^{opt}(A)\cap B|}$ for some $c' \in [\frac{1}{2^d}, \frac{1}{2^{d-1}}]$ and f is Lipschitz continuous with constant $c'\delta^{-d}\frac{|\partial A|}{|t^{opt}(A)\cap B|}$.

Completely analogously to Inequality (5.2) in the proof of Theorem 5.1, we have, with probability at least $1-p$,

$$|t^*(A) \cap B| \geq |t(A) \cap B| - \left(\frac{|t^{opt}(A) \cap B|\delta^d}{c'\sqrt{M}(2\delta_\varepsilon)^d}\left(2c\sqrt{2d} + 2\sqrt{\log\frac{2}{p}}\right) + \frac{3}{2}\sqrt{d}\delta_\varepsilon|\partial A|\right). \tag{5.8}$$

As before, there is a constant $C > 0$ such that for every positive $\delta_\varepsilon \leq \frac{2\sqrt{d}}{3(d+1)}\varepsilon\frac{|A|}{|\partial A|}$ and every integer $M \geq C\varepsilon^{-2}\delta^{2d}\delta_\varepsilon^{-2d}\log\frac{2}{p}$, we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$. We can choose $\delta_\varepsilon = \delta\varepsilon^{1-k}$, which simplifies the bound on M to $M \geq C\varepsilon^{2kd-2d-2}\log\frac{2}{p}$.

Since $\log\frac{4}{p} = \log 2 + \log\frac{2}{p}$ and $\log\frac{2}{p} \geq \log 2$ for $p \in (0,1)$, we have $\log\frac{4}{p} = \Theta(\log\frac{2}{p})$. By multiplying the constant C with 2, we can safely assume that Inequality (5.8) holds with probability at least $1 - \frac{p}{2}$.

Next we have to estimate how many samples we have to generate such that, with probability at least $1 - \frac{p}{2}$ there are at least M valid samples. Let N be the total number of generated samples.

We estimate the probability that for $(a,t) \in A \times B(0,\delta)$ u.a.r., we have $b \in B$.

$$\begin{aligned}
P(a+q+t \in B) &= \frac{1}{|A|\delta^d}\int_{B(0,\delta)}\int_A \chi_{(B-q-t)}(a)\,da\,dt \\
&= \frac{1}{|A|\delta^d}\int_{B(0,\delta)} |(q+t)(A) \cap B|\,dt \\
&\geq \frac{1}{|A|\delta^d}\int_{B(0,\delta)} (|t^{opt}(A) \cap B| - (2\frac{|B|}{m_B} + 1)\varepsilon^k|A|)dt \\
&\geq \frac{|t^{opt}(A) \cap B|}{|A|} - \left(2\frac{|B|}{m_B} + 1\right)\varepsilon^k \\
&\geq \left(2\frac{|B|}{m_B} + 1\right)\varepsilon^k \\
&\geq 3\varepsilon^k.
\end{aligned}$$

Let X_j be a binary random variable indicating if the jth sample t fulfills $a+q+t \in B$, and let $X = \sum_{j=1}^N X_j$. The expected value $\mathbf{E}X$ equals $P_1 N$ where $P_1 = P(X_j = 1) = P(a+q+t \in B)$, which is as least as large as ε^k. By the Chernoff bound (Theorem 4.3), $P(X < \frac{P_1}{2}N) < 2e^{-P_1^2 N/8}$. For $N \geq \frac{8}{P_1^2}\log\frac{4}{p}$, we have

$2e^{-P_1^2 N/8} \le \frac{p}{2}$. If additionally $N \ge \frac{2M}{P_1}$, then with probability at least $1-\frac{p}{2}$, we have at least M valid samples. For $N \ge \max\{\frac{8}{9\varepsilon^{2k}} \log\frac{4}{p}, \frac{2M}{3\varepsilon^k}\}$, we have with probability at least $1-\frac{p}{2}$ at least M valid samples. We have $M \ge C\varepsilon^{2kd-2d-2}\log\frac{2}{p}$ and because of $2kd-2d-2 < -k$, which is equivalent to the correct inequality $k(2d+1) < 2d+2$, we can choose C such that $\frac{2M}{3\varepsilon^k} \ge \frac{8}{9\varepsilon^{2k}}\log\frac{4}{p}$. To summarize, the condition on N is $N \ge \frac{2M}{3\varepsilon^k}$. □

Reducing the number of candidate cubes

Next we thin out the set Q' in the following simple way. As long as we have a point $q_0 \in Q'$ left, we put q_0 in the output collection Q'' and delete all points in $B(q_0,\delta)$ from Q'.

Function thinOut(Q', δ)

Input: collection Q' of points from $\mathbb{R}^d$, positive number δ
collection $Q'' \leftarrow \emptyset$;
while $Q' \ne \emptyset$ **do**
  $q \leftarrow$ element(Q');
  $Q' \leftarrow$ removeAll(q, δ, Q'); `// removes all points in` $B(q,\delta)$ `from` Q'
  $Q'' \leftarrow$ add(q, Q'');
end
return Q'';

Define

$$\mathcal{Q} = \bigcup_{q\in Q''} B(q,\delta).$$

Then by construction $\bigcup_{q\in Q'} B(q,\delta) = \mathcal{Q}$ and for all $q_0 \in \mathcal{Q}$, we have

$$|t^{opt}(A)\cap B| - |q_0(A)\cap B| \le \Big(2\frac{|B|}{m_B}+1\Big)\varepsilon^k|A|.$$

For $c > 0$, define

$$\mathcal{W}_c = \{t \in \mathbb{R}^d : |t^{opt}(A)\cap B| - |t(A)\cap B| \le c\varepsilon^k|A|\}.$$

Let $c_0 = 2\frac{|B|}{m_B}+1$, then $\mathcal{Q} \subseteq \mathcal{W}_{c_0} = \mathcal{W}$. Next we bound the number of points in Q'' from above. To a point $q \in Q''$, let us associate the cube $B(q,\frac{\delta}{2}) \subset \mathcal{Q}$. Observe that for $q_1, q_2 \in Q''$, the distance $||q_1-q_2||_\infty > \delta$ and therefore the associated cubes are disjoint. Hence $|Q''|(\frac{\delta}{2})^d \le |\mathcal{Q}|$, and therefore

$$|Q''| \le (\tfrac{\delta}{2})^{-d}|\mathcal{W}_{c_0}| \le (\tfrac{\delta}{2})^{-d}|B\oplus(-A)|. \tag{5.9}$$

We do not know how to find a better upper bound than the volume of the translation space on $|\mathcal{W}_{c_0}|$ in general. For the special case of perfectly matching shapes A and B, we prove that, if $\varepsilon < 1/c_0^2$, then $|\mathcal{W}_{c_0}| \leq (2c_0\varepsilon^k \operatorname{diam}(A))^d\omega_d$. We know there is $\eta \in (0,1)$ such that $\delta = \eta\frac{2\sqrt{d}}{3(d+1)}\frac{\varepsilon^k|A|}{|\partial A|}$. If we can assume η to be constant, then for perfectly matching shapes we have $|Q''| = O\big(\frac{\operatorname{diam}(A)^d|\partial A|^d}{|A|^d}\big) = O(D_A K_A)$.

Lemma 5.8. *Let $A, B \subset \mathbb{R}^d$ be shapes such that there is $t^{opt} \in \mathbb{R}^d$ with $t^{opt}(A) = B$. Let $\varepsilon \in (0,1)$ and $c > 0$ such that $c\varepsilon^k < 1$. Then $|\mathcal{W}_c| \leq (2c\varepsilon^k \operatorname{diam}(A))^d\,\omega_d$.*

Proof. If there is $t^{opt} \in \mathbb{R}^d$ with $t^{opt}(A) = B$, then this translation t^{opt} is unique, and $|t^{opt}(A) \cap B| = |A|$. Let $t \in \mathcal{W}_c$, and denote $t^* = t^{opt} - t$. We have

$$|t(A) \cap B| = |t(A) \cap t^{opt}(A)| = |t^*(A) \cap A| = |A| - \tfrac{1}{2}|t^*(A) \triangle A|.$$

If $|t^*| \geq \operatorname{diam}(A)$, then $|t^*(A) \triangle A| = 2|A|$ and hence $|t(A) \cap B| = 0$. Therefore $|A| = |t^{opt}(A) \cap B| \leq c\varepsilon^k|A| < |A|$, which is a contradiction.

By Lemma 3.12, it holds that $|t^*(A) \triangle A| \geq \frac{|A|}{\operatorname{diam}(A)}|t^*|$. Using that $t \in \mathcal{W}_c$, we have

$$|A| - c\varepsilon^k|A| \leq |t(A) \cap B| \leq |A| - \tfrac{|A|}{2\operatorname{diam}(A)}|t^*|.$$

Hence $|t^*| \leq 2c\varepsilon^k \operatorname{diam}(A)$, and $\mathcal{W}_c$ is contained in the Euclidean ball of radius $2c\varepsilon^k \operatorname{diam}(A)$ around t^{opt}. □

Putting things together

First we describe Algorithm 5. Recall that k is not a variable, although the algorithm works for every $k \in (0,1)$. We will set k to some optimal constant, as soon as we analyzed the runtime of the algorithm. Therefore it is not part of the input.

Let E_1 be the event that MaxOverlapTCandidates$(A, B, \varepsilon, \frac{p}{2})$ fails and let E_2 be the event that one of the $|Q''|$ calls of MaxOverlapTCube$(A, B, q, \delta, \varepsilon, \frac{p}{4|Q''|})$ fails. We have $P(E_1) \leq \frac{p}{2}$ and $P(E_2) \leq \frac{p}{4}$.

Denote the event that MaxOverlapTCandidates$(A, B, \varepsilon, \frac{p}{2})$ succeeds by $\overline{E_1}$. Then $P(\overline{E_1}) \geq 1 - \frac{p}{2}$. The probability that Algorithm 5 fails is less or equal to $P(E_1) + P(E_2 \,|\, \overline{E_1})$. We have

$$P(E_2 \,|\, \overline{E_1}) = \frac{P(E_2 \cap \overline{E_1})}{P(\overline{E_1})} \leq \frac{\frac{p}{4}}{2(1-\frac{p}{2})} \leq \frac{p}{2}.$$

With this observation, Lemmas 5.5 and 5.7 imply the following corollary.

Corollary 5.9 (Correctness of Algorithm 5)**.** *Let A and B be shapes in constant dimension d, and let $\varepsilon, k, p \in (0,1)$ be parameters. Assume that there exists $t \in \mathbb{R}^d$ such that $|t(A) \cap B| \geq (4\frac{|B|}{m_B} + 2)\varepsilon^k|A|$. Then Algorithm 5 with input (A, B, ε, p) computes a pair of a translation t^* and a number V^* such that, with probability at least $1 - p$, we have $|t^*(A) \cap B| \geq |t(A) \cap B| - \varepsilon|A|$ for all translations $t \in \mathbb{R}^d$, and $|V^*|A|\,|B| - |t^*(A) \cap B|| \leq \frac{\varepsilon}{2}|A|$.*

Algorithm 5: AdaptiveMaxOverlapT(A, B, ε, p)

```
Input: shapes A, B ⊂ ℝ^d, error bound ε ∈ (0,1), allowed probability of
       failure p ∈ (0,1)
// Step 1:  Computation of a candidate region Q'
(collection, real) (Q', δ) ← MaxOverlapTCandidates(A, B, ε, p/2);
collection Q'' ← thinOut(Q');
// Step 2:  Search in cubes B(q, δ) for all q ∈ Q''
integer V* ← 0;                                   // current maximum
point t* ← (0, 0, ..., 0);                        // corresponding translation
for q ∈ Q'' do
    (point, integer) (t, V) ← MaxOverlapTCube(A, B, q, δ, ε, p/(4|Q''|));
    if V > V* then
        V* ← V;
        t* ← t;
    end
end
return (t*, V*);
```

Theorem 5.10 (Runtime of Algorithm 5)**.** *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters. Assume that there exists $t \in \mathbb{R}^d$ such that $|t(A) \cap B| \geq (4\frac{|B|}{m_B} + 2)\varepsilon^k |A|$.*

Then Algorithm 5 with input (A, B, ε, p) runs in time $O\big(T_q(N') + N'(\log N')^{d-1}\big)$ where $N' = O\Big(\varepsilon^{-\frac{4d^2+8d+4}{3d+1}} \log \frac{2}{p}\big((K_A M_B/m_A)^2 + K_A \frac{\Delta_A^d + \Delta_B^d}{m_A}\big)\Big)$.

Proof. We determine the time complexity of the different steps of the algorithm.

- **Algorithm MaxOverlapTCandidates($A, B, \varepsilon, \frac{p}{2}$):**
  For Function FindAllDenseClustersT, we proceed as in the case of Function FindDensestClusterT: We build an orthogonal range counting query data-structure for the N random translations in Q. Then we query every $B(q, \delta)$ for $q \in Q$ twice. In the first round, we determine the translation with the maximum number of points in its neighborhood. During the second round we keep all translations $q \in Q$ that are sufficiently close to the maximum. We have the same time bound for FindAllDenseClustersT($Q, \delta, \varepsilon, m_B$) as for FindDensestClusterT(Q, δ), that is $O(N(\log N)^{d-1})$ where $N = |Q|$.

  Hence the runtime of Algorithm MaxOverlapTCandidates($A, B, \varepsilon, \frac{p}{2}$) equals the runtime of Algorithm MaxOverlapT($A, B, \varepsilon^k, \frac{p}{2}$), which by Theorem 5.2 is $O\big(T(N) + N(\log N)^{d-1}\big)$ where $N = O\big(\varepsilon^{-2kd-2k}(K_A M_B/m_A)^2 \log \frac{2}{p}\big)$.

- **Function thinOut(Q'):**
  The while-loop is entered $|Q''|$ times since in every pass one element is added

to Q''. We build a orthogonal range reporting data structure for Q'. We have $|Q'| \leq N$. For a $q \in Q'$, we query $B(q, \delta)$. Let $b_q = B(q, \delta) \cap Q'$. Then the query time can be bounded by $O\big((\log N)^{d-1} + b_q\big)$. We can also remove the elements from Q' within the same time bound since we do not care about rebalancing the tree. We have $\sum_{q \in Q'} b_q = |Q'| \leq N$, and in total we need time $O\big(|Q''|(\log N)^{d-1} + N\big) = O(N(\log N)^{d-1})$.

- **Algorithm MaxOverlapTCube($A, B, q, \delta, \varepsilon, \frac{p}{4|Q''|}$):**
  Recall that we assume in this section that in time $T_q(N)$ we can generate N random samples from a shape and answer N membership queries. Algorithm MaxOverlapTCube$(A, B, q, \delta, \varepsilon, \frac{p}{4|Q''|})$ runs in time $O(T_q(M) + M(\log M)^{d-1})$ where $M = O(\varepsilon^{2kd-k-2d-2} \log \frac{4|Q''|}{p})$.

- **Upper bound on $|Q''|$:**
  The value δ that is computed in MaxOverlapTCandidates$(A, B, \varepsilon^k, \frac{p}{2})$ fulfills $\delta \leq \frac{2\sqrt{d}}{3(d+1)} \frac{\varepsilon^k |A|}{|\partial A|}$ and we can assume that $\delta^{-1} = O(\varepsilon^{-k} K_A^{\frac{1}{d}} m_A^{-\frac{1}{d}})$. By Inequality (5.9) and Section 3.4, we have $|Q''| = O(\varepsilon^{-kd} K_A m_A^{-1}(\Delta_A^d + \Delta_B^d))$.

In total, the runtime is $O\Big(T(N) + N(\log N)^{d-1} + |Q''|\big(T_q(M) + M(\log M)^{d-1}\big)\Big)$.

The bound $O\big(N(\log N)^{d-1}\big)$ increases with increasing k, while the corresponding bound for the second step $O\big(\varepsilon^{-kd} K_A m_A^{-1}(\Delta_A^d + \Delta_B^d) M(\log M)^{d-1}\big)$ decreases with increasing k. Hence we choose k such that both bounds are equal in the growth rate of ε. If we have reason to believe that $|Q''|$ is small for our instance, than a different choice of k might be appropriate.

Denote by O_ε^* the growth rate in ε, ignoring log-factors. Then $O(N(\log N)^{d-1}) = O_\varepsilon^*(\varepsilon^{-2kd-2k})$ and $O\big(\varepsilon^{-kd} K_A m_A^{-1}(\Delta_A^d + \Delta_B^d) M(\log M)^{d-1}\big) = O_\varepsilon^*(\varepsilon^{kd-k-2d-2})$. For $-2kd - 2k = kd - k - 2d - 2$, which is equivalent to $k = \frac{2d+2}{3d+1}$, both growth rates are equal to $O_\varepsilon^*(\varepsilon^{-\frac{4d^2+8d+4}{3d+1}})$.

We can choose δ such that there is

$$N' = N + |Q''|M = O\Big(\varepsilon^{-\frac{4d^2+8d+4}{3d+1}} \log \tfrac{2}{p}\big((K_A M_B/m_A)^2 + K_A \tfrac{\Delta_A^d + \Delta_B^d}{m_A}\big)\Big).$$

Assuming that $T_q(N)$ grows at least linearly in N, the runtime is $O\big(T_q(N') + N'(\log N')^{d-1}\big)$. □

6. Probabilistic matching of planar regions under rigid motions

We present three simple probabilistic algorithms that approximately maximize the area of overlap of two given 2-dimensional shapes under rigid motions. More precisely, given shapes $A, B \subset \mathbb{R}^2$ and parameters $p, \varepsilon \in (0,1)$, the presented algorithms compute a rigid motion r^* such that with probability at least $1-p$ we have $|r^*(A) \cap B| \geq |r(A) \cap B| - \varepsilon|A|$ for all rigid motions r.

These algorithms are based on the ideas explained in Chapter 5. The algorithm described in Section 6.1 uses the algorithm for translations. The algorithms described in Sections 6.2 and 6.3 are generalizations of the idea of the probabilistic matching algorithm for translations. They differ in the generation of random rigid motions. We provide an experimental comparison of these algorithms in Chapter 7.

Variants of the two algorithms in Sections 6.2 and 6.3 were analyzed in [9]. Here we improve the results by showing that a simpler definition of cluster suffices to guarantee approximation, as in the case of translations. We also slightly improve the bounds on the number of random samples, but more importantly, the clustering size now scales with the shapes and the sample size is homogeneous in all variants of the probabilistic matching algorithm.

For two polygonal regions with n and m vertices, Cheong et al. [20] give an absolute error approximation of the maximal area of overlap under rigid motions in time $O(m + (n^3/\varepsilon^8)\log^5 n)$, which works with high probability. The runtime given in the paper is smaller because of a calculation error in the final derivation of the time bound, as was noted by Vigneron [54]. We compared Cheong et al.'s and our approach in Chapter 5 for the case of translations; basically the same remarks apply to the case of rigid motions.

For maximizing the area of overlap of two simple polygons with at most n vertices under rigid motions, no exact polynomial time algorithm is known. Vigneron [54] gives an FPTAS with relative error ε that runs in time $O((n^6/\varepsilon^3)\log^4(n/\varepsilon)\beta(n/\varepsilon))$ where β is a very slowly growing function related to the inverse Ackermann function. He also gives an FPTAS for minimizing the area of the symmetric difference of two polygons in the plane under rigid motions that has time complexity $O\big(n^{17+\varepsilon'} + (n^{15}/\varepsilon^3)\log^4(n/\varepsilon)\beta(n/\varepsilon)\beta(n)\big)$ for any $\varepsilon' > 0$.

The space of rigid motions in the plane R is given as $\frac{1}{2\pi}\mathcal{S}^1 \times \mathbb{R}^2$. Therein, the 1-dimensional sphere $\mathcal{S}^1$ is understood as the interval $[0, 2\pi]/\sim$, where the equivalence relation $\sim$ identifies 0 and 2π. A point $(\alpha, t) \in R$ denotes the rigid motion

$$x \mapsto M_\alpha x + t, \quad M_\alpha = \begin{pmatrix} \cos 2\pi\alpha & -\sin 2\pi\alpha \\ \sin 2\pi\alpha & \cos 2\pi\alpha \end{pmatrix}.$$

6.1. Using the algorithm for translations

The idea of the first algorithm for rigid motions is very simple. We discretize the angle space $[0,1]$ in m angles $(\alpha_1,\dots,\alpha_m)$ and apply Algorithm 2 with input $(M_{\alpha_k}A, B, \varepsilon/3, p/m)$ for all $k=1,\dots,m$. Let $(t_1,V_1),\dots,(t_m,V_m)$ be the outputs of Algorithm 2. The result is a rigid motion (α_i,t_i) such that the area estimate $V_i|A||B|$ is maximal among $V_1,\dots,V_m$, which means that t_i received the overall most votes.

Algorithm 6: 2DMaxOverlapRMT

```
Input: shapes A, B ⊂ ℝ², error tolerance ε ∈ (0,1), allowed probability of
       failure p ∈ (0,1)
// discretization size 1/m
integer m ← 2DDiscretizationSizeRMT(A, ε);
real V* ← 0 ;                            // current maximal area estimate
point t* ← (0,0) ;                  // corresponding translation vector
real α* ← 0 ;                           // corresponding rotation angle
for k = 1...m do
    A_k ← M_{k/m} A;                                // compute rotated A
    // use algorithm for translations to match A_k and B
    (point,real) (t, V) ← MaxOverlapT(A_k, B, ε/3, p/m);
    if V > V* then
        V* ← V ;                                // update area estimate
        t* ← t;                                   // update translation
        α* ← k/m;                               // update rotation angle
    end
end
return (α*, t*);
```

Observe that we do not need to compute rotated shapes $M_\alpha A$. Instead we can pass the angle α to the algorithm for translation and rotate every generated random point from A in constant time. In this way, we generate uniformly distributed points from $M_\alpha A$ without actually computing $M_\alpha A$. We can also use Algorithm 5 instead of Algorithm 2 for the matching under translations.

A bound on the discretization size that guarantees approximation is given in the following theorem.

Theorem 6.1 (Approximation property of Algorithm 6)**.** *Let $A,B\subset\mathbb{R}^2$ be shapes and let $\varepsilon,p\in(0,1)$ be parameters. If 2DDiscretizationSizeRMT(A,ε) returns an integer $m\geq 3\pi\varepsilon^{-1}|A|^{-1}\operatorname{diam}(A)|\partial A|$, then Algorithm 6 with input (A,B,ε,p) computes a rigid motion r^* that maximizes the area of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$.*

Proof. Let (a^*,t^*) be the output of Algorithm 6 and V^* the area estimate that Algorithm 2 computed for t^*. Each of the m calls of Algorithm 2 fails with probability at most p/m. Hence the probability that one of the m calls fails, which is the probability of the union of these events, is at most p, and the probability that all m calls succeed is at least $1-p$.

Let us assume that this is the case, meaning that for all $k = 1, \dots, m$, the output (t_k, V_k) of MaxOverlapT$(A_k, B, \varepsilon/3, p/m)$ fulfills for all $t \in \mathbb{R}^2$ that

$$|t_k(A_k) \cap B| \geq |t(A_k) \cap B| - \varepsilon/3|A| \quad \text{and} \quad \big||t_k(A_k) \cap B| - V_k\big| \leq \varepsilon/6|A|.$$

Let (α^{opt}, t^{opt}) be an optimal solution. There exists k such that $|\alpha_k - \alpha^{opt}| < 1/m$. We show that $|t_k(A_k) \cap B|$ is close to the maximal area of overlap, as well as to the area of overlap with respect to the output, implying that also the maximal area of overlap and the area of overlap with respect to to (α^*, t^*) are close.

First, we show that

$$\big||t_k(A_k) \cap B| - |t^{opt}(M_\alpha^{opt} A) \cap B|\big| \leq 2/3\varepsilon|A|.$$

By the optimality of (α^{opt}, t^{opt}), we have $|t^{opt}(M_\alpha^{opt} A) \cap B| \geq |t_k(A_k) \cap B|$. Using Proposition 3.6 and Theorem 3.3, we have for all translations t, in particular for t^{opt},

$$\begin{aligned} \big||(A_k + t) \cap B| - |(M_{\alpha^{opt}} A + t) \cap B|\big| &= \big||A_k \cap (B - t)| - |M_{\alpha^{opt}} A \cap (B - t)|\big| \\ &\leq \tfrac{1}{2}\, |A_k \triangle M_{\alpha^{opt}} A| \\ &\leq \tfrac{w}{2}\, |\partial A| \end{aligned}$$

where $w = \sup_{a \in \partial A} |M_{k/m} a - M_{\alpha^{opt}} a|$. We have $w \leq 2\pi/m \,\mathrm{diam}(A)$. If m is as least as large as $3\pi\,\varepsilon^{-1}|A|^{-1}\,\mathrm{diam}(A)|\partial A|$, then $\frac{w}{2}|\partial A| \leq \varepsilon/3|A|$. By this and the assumption on the outputs of Algorithm 6, we have

$$|t_k(A_k) \cap B| \geq |t^{opt}(A_k) \cap B| - \varepsilon/3|A| \geq |t^{opt}(M_{\alpha^{opt}} A) \cap B| - 2/3\varepsilon|A|.$$

If the area of overlap with respect to the output is as least as large as $|t_k(A_k) \cap B|$, then $\big||t^*(M_{\alpha^*} A) \cap B| - |t^{opt}(M_\alpha^{opt} A) \cap B|\big| \leq 2/3\varepsilon|A|$ and we are done. Assume on the contrary that $|t^*(M_{\alpha^*} A) \cap B| < |t_k(A_k) \cap B|$.

Since $V^* \geq V_k$ and $\big||t_k(A_k) \cap B| - V_k\big| \leq \varepsilon/6|A|$ as well as $\big||t^*(M_{\alpha^*} A) \cap B| - V^*\big| \leq \varepsilon/6|A|$, we have $\big||t^*(M_{\alpha^*} A) \cap B| - |t_k(A_k) \cap B|\big| \leq \varepsilon/3|A|$.

In total, we get $\big||t^*(M_{\alpha^*} A) \cap B| - |t^{opt}(M_{\alpha^{opt}} A) \cap B|\big| \leq \varepsilon|A|$. □

The runtime of Algorithm 6 equals m times the runtime of Algorithm 2 in two dimensions plus the time to compute the discretization size m since we do not actually rotate shape A, but rotate the generated random points from A. For a better time bound, we use Algorithm 5 instead of Algorithm 2 for the matching under translations. We can compute a valid m, using the given upper bound D_A on $\frac{\mathrm{diam}(A)^2}{|A|}$ (see Chapter 2). For unions of simplices, the diameter and the volume can be computed in linear time, so only an upper bound on $|\partial A|$ that is better than $\sum_{\Delta \in A} |\partial \Delta|$ is useful.

Corollary 6.2 (Runtime of Algorithm 6). *Let $A, B \subset \mathbb{R}^2$ be shapes, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters. Then Algorithm 6 with input (A, B, ε, p) computes a rigid motion r^* that maximizes the area of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$ in time $O\big(\varepsilon^{-1}\sqrt{D_A K_A}\,(T_q(N) + N\log N)\big)$ where $N = O\Big(\varepsilon^{-(5+\frac{1}{7})}\big((K_A M_B/m_A)^2 + K_A \frac{\Delta_A^2+\Delta_B^2}{m_A}\big)\log\frac{2}{p}\Big)$.*

If we use Algorithm 2, the runtime is $O\big(\varepsilon^{-1}\sqrt{D_A K_A}\,\big(T(N) + N\log N\big)\big)$ where $N = O(\varepsilon^{-6}(K_A M_B/m_A)^2 \log\frac{2}{p})$.

6.2. Vote generation with random angle

The second algorithm for rigid motions, Algorithm 7, is a generalization of the algorithm for translations. In each step, we select uniformly distributed an angle α and random points $a \in A$ and $b \in B$. We give one vote to the unique rigid motion with counterclockwise rotation angle α that maps a onto b, which is the rigid motion $(\alpha, b - M_\alpha a)$.

For the definition of a best cluster, we first define a metric on the space of rigid motions in the plane R. We use a scaled version of the metric that is induced by the maximum norm. The scaling factor depends on the input shape A. More precisely, we scale the axis that corresponds to the rotation angle with $\operatorname{diam}(A)$. Instead of $\operatorname{diam}(A)$, an upper bound on the diameter D_A can be used. Let $r = (\alpha, p)$ and $s = (\beta, q)$ be rigid motions. Define

$$d(r,s) = \max\{\operatorname{diam}(A)\,|\alpha - \beta|\,,\ ||p - q||_\infty\,\}.$$

The angle distance is scaled because otherwise it is not comparable to the distance of the translational part. When scaling the shape with a factor, the clustering size for the rotational part should not change, but the clustering size for the translational part should change linearly.

Any usual metric, which does not depend on A, for example the Euclidean distance or the L_∞ metric, behaves differently for the rotational part and for the translational part. If two translation vectors p and q have distance δ, then for each point in $\mathbb{R}^2$, the images under the translations p and q have distance exactly δ, too. If two rotation angles α and β have distance δ, it depends on the distance from a point $x \in \mathbb{R}^2$ to the origin how large the distance of $M_\alpha x$ and $M_\beta x$ is. This distance depends linearly on δ and the length of the vector x.

W.l.o.g. the shape A contains the origin. Then the maximal distance Δ from the origin to any point of A has the same order of magnitude as $\operatorname{diam}(A)$ since $\frac{1}{2}\operatorname{diam}(A) \le \Delta \le \operatorname{diam}(A)$.

The *δ-neighborhood* of r is the closed δ-ball around r with respect to the metric d. It is a 3-dimensional box with side lengths $2\delta/\operatorname{diam}(A)$ and 2δ and 2δ that is centered at r.

Algorithm 7: 2DMaxOverlapRMRA

Input: shapes $A, B \subset \mathbb{R}^2$, error tolerance $\varepsilon \in (0,1)$, allowed probability of failure $p \in (0,1)$

real $\delta \leftarrow$ 2DClusteringSizeRMRA(A, ε);
integer $N \leftarrow$ 2DSampleSizeRMRA$(B, \varepsilon, \delta, p)$;
collection $Q \leftarrow \emptyset$;
for $i = 1 \ldots N$ **do**
  real $\alpha \leftarrow$ RandomAngle(); `// returns a real in` $[0,1]$ `u.a.r.`
  point $a \leftarrow$ RandomPoint(A);
  point $b \leftarrow$ RandomPoint(B);
  add$(Q, (\alpha, b - M_\alpha a))$;
end
return *2DFindDensestClusterRM(*$Q, \delta, \mathrm{diam}(A)$*)*;

Function 2DFindDensestClusterRM computes a rigid motion whose neighborhood contains the most rigid motions from the input collection Q. The radius of the neighborhood δ and the scaling factor of the metric Δ are given as input parameters.

Function 2DFindDensestClusterRM(Q, δ, Δ)

Input: collection of points in the plane Q, positive numbers δ and Δ
Output: point $r \in [0,1] \times \mathbb{R}^2$ such that the rectangle of side lengths $2\delta/\Delta$, 2δ and 2δ that is centered at r contains a maximal number of points from Q

Theorem 6.3 (Approximation property of Algorithm 7)**.** *Let $A, B \subset \mathbb{R}^2$ be shapes and let $\varepsilon, p \in (0,1)$ be parameters. If 2DClusteringSizeRMRA(A, ε) returns a positive number $\delta \leq \frac{2}{3\sqrt{2}+6\pi} \frac{\varepsilon |A|}{|\partial A|}$ and 2DSampleSizeRMRA$(B, \varepsilon, \delta, p)$ returns an integer $N \geq C \varepsilon^{-2} \delta^{-6} \,\mathrm{diam}(A)^2 |B|^2 \log \frac{2}{p}$ for some universal constant $C > 0$, then Algorithm 7 with input (A, B, ε, p) computes a rigid motion r^* such that with probability at least $1 - p$ we have $|r^*(A) \cap B| \geq |r(A) \cap B| - \varepsilon |A|$ for all rigid motions r.*

As one would expect, the clustering size scales with the shapes. In contrast to the earlier version of this result [9], the sample size N is invariant under scaling of both shapes with the same factor. The constant C can be computed from the proofs.

We analyze the density functions of the probability distribution in the transformation space that is induced by the random experiments in the algorithms. We show the surprising fact that, also in the case of rigid motions, the value of the density function for a rigid motion r is proportional to the area of overlap $|r(A) \cap B|$.

Lemma 6.4. *Let X be a random vector that draws rigid motions $(\alpha, b - M_\alpha a) \in [0,1] \times \mathbb{R}^2$ where $(\alpha, a, b) \in [0,1] \times A \times B$ is drawn u.a.r. The density function of X is given by $g_{RA}(r) = \frac{|r(A) \cap B|}{|A|\,|B|}$.*

Proof. Our random experiment consists in selecting uniformly distributed points from $\Omega = \frac{1}{2\pi}\mathcal{S}^1 \times A \times B$. Let $I = \frac{1}{2\pi}\mathcal{S}^1$. We are interested in the density function f_Y of the random variable

$$Y : \Omega \to R, \quad Y : (\alpha, a, b) \mapsto (\alpha, b - M_\alpha a).$$

We will express the density function of Y in terms of the conditional probability densities of the following two random variables Y_I and Y_T defined as

$$Y_I : \Omega \to I, \qquad Y_I : (\alpha, a, b) \mapsto \alpha,$$
$$Y_T : \Omega \to \mathbb{R}^2, \qquad Y_T : (\alpha, a, b) \mapsto b - M_\alpha a.$$

The density function of Y is the joint density of the random variables Y_I and Y_T. Recall that the counterclockwise rotation angle is selected uniformly distributed in I independently from the points a and b. So the marginal probability density of Y_I, i.e., probability density of $Y_I = \alpha$ allowing all possible values of Y_T, is

$$f_I(\alpha) = \tfrac{1}{|I|} = 1.$$

The value of Y_T depends on the selected points a and b and on the value of Y_I. The conditional probability density of $Y_T = t$ given $Y_I = \alpha$ is exactly the probability density in the space of translations for shapes $M_\alpha A$ and B:

$$f_T(t \mid Y_I = \alpha) = \frac{|(M_\alpha A + t) \cap B|}{|A||B|}.$$

The conditional probability density can also be expressed in terms of the joint probability density $f_T(t \mid Y_I = \alpha) = f_Y(\alpha, t)/f_I(\alpha)$. Thus we get for any rigid motion $r = (\alpha, t)$

$$g_{RA}(r) = f_Y(r) = \frac{|r(A) \cap B|}{|A|\,|B|}.$$

□

Lemma 6.5. *The density function g_{RA} is Lipschitz continuous with the constant $L = L_{RA} = (\frac{1}{\sqrt{2}} + \pi)\frac{|\partial A|}{|A|\,|B|}$.*

Proof. For two angles α and β, we have $|M_\beta a - M_\alpha a| \leq 2\pi|\beta - \alpha|\,|a|$ because $2\pi|\beta - \alpha|\,|a|$ is the length of the circular arc from $M_\alpha a$ to $M_\beta a$ that is a segment of the circle of radius $|a|$ around the origin. W.l.o.g. let A contain the origin, then we have $\sup_{a \in A} |M_\beta a - M_\alpha a| \leq 2\pi|\beta - \alpha| \operatorname{diam}(A)$. By Corollary 3.5,

$$\big||r(A) \cap B| - |s(A) \cap B|\big| \leq \tfrac{|\partial A|}{2|A|\,|B|}\big(|t - s| + 2\pi|\beta - \alpha| \operatorname{diam}(A)\big).$$

Because of $|t - s| \leq \sqrt{2}||t - s||_\infty$, the claim follows. □

Proof of Theorem 6.3. The proof works similarly as in the case of translations. Again, we start by checking that all assumptions of Theorem 4.6 are satisfied. The space of rigid motions $R = \frac{1}{2\pi}\mathcal{S}^1 \times \mathbb{R}^2$ is a metric space. Recall the definition of the metric $d((\alpha,p),(\beta,q)) = \max\{\,\mathrm{diam}(A)|\alpha-\beta|\,,\,||p-q||_\infty\,\}$. Let $\mathcal{B}_\delta$ be the set of closed balls in the space of rigid motions that have radius δ. For every $B \in \mathcal{B}_\delta$, we have $|B| = 8\delta^3/\,\mathrm{diam}(A)$. The VC dimension of $\mathcal{B}_\delta$ is at most 6 by Lemma 4.8.

Let μ be the probability distribution that is implicitly given by the random experiment in Algorithm 7. The density function of μ equals $g_{RA}(r) = \frac{|r(A)\cap B|}{|A|\,|B|}$ by Lemma 6.4. It is Lipschitz continuous with constant $L = (\frac{1}{\sqrt{2}}+\pi)\frac{|\partial A|}{|A|\,|B|}$ by Lemma 6.5.

Let $Y_1, Y_2, \ldots, Y_N$ be the random rigid motions in the algorithm. They have the common distribution μ. Let μ_N be the empirical measure. Algorithm 7 computes a Y_j such that $\mu_N(B(Y_j,\delta)) = \max\{\mu_N(B(Y_i,\delta)) : 1 \le i \le N\}$. By Theorem 4.6, there is a constant $c > 0$ such that for all $\tau > 0$, we have with probability at least $1-2e^{-2N\tau^2}$ that for all $r \in R$,

$$g_{RA}(r^*) \ge g_{RA}(r) - \left(\frac{c\sqrt{6}\,\mathrm{diam}(A)}{4\sqrt{N}\delta^3} + \frac{\tau\,\mathrm{diam}(A)}{4\delta^3} + 3L\delta\right). \tag{6.1}$$

As in the case of translations, we set $\tau = \sqrt{\frac{1}{2N}\log\frac{2}{p}}$ because this minimizes the error term while guaranteeing a probability of success $\ge 1-p$.

We now determine a minimal N and a compatible δ such that with probability at least p, for all $r \in R$, we have $|r^*(A)\cap B| \ge |r(A)\cap B| - \varepsilon|A|$. Therefore by Inequality (6.1) we need to have

$$\frac{1}{4}\,\delta^{-3}\,\mathrm{diam}(A)|A|\,|B|\left(\frac{c\sqrt{6}}{\sqrt{N}}+\tau\right) + \left(\frac{3}{\sqrt{2}}+3\pi\right)|\partial A|\,\delta \;\le\; \varepsilon|A|.$$

Clearly, we need $\delta < \frac{2}{3\sqrt{2}+6\pi}\frac{\varepsilon|A|}{|\partial A|}$. Let $\eta \in (0,1)$ such that $\delta = \eta\frac{2}{3\sqrt{2}+6\pi}\frac{\varepsilon|A|}{|\partial A|}$. There is a $C > 0$ such that for $N \ge C(1-\eta)^{-2}\varepsilon^{-2}\delta^{-6}\,\mathrm{diam}(A)^2|B|^2\log\frac{2}{p}$, we have that the error term is at most $\varepsilon|A|$. Differentiating shows that $\eta = 3/4$ gives the optimal value for δ, that is the value minimizing N. □

The runtime of Algorithm 7 is bounded by the time to compute δ and N, to generate N random points from a shape, and to find a best δ-neighborhood. For the computation of δ, we use an upper bound on the isoperimetric quotient and a lower bound on the volume of A.

The angle space $\frac{1}{2\pi}\mathcal{S}^1$ can be represented as $[0, 1+\frac{\delta}{\mathrm{diam}(A)}]$. Angles $\alpha \in [0, \frac{\delta}{\mathrm{diam}(A)}]$ are then inserted twice as α and $1+\alpha$. This ensures that the number of random samples in their neighborhood will be computed correctly for one of the copies.

Corollary 6.6 (Runtime of Algorithm 7)**.** *Let $A, B \subset \mathbb{R}^2$ be shapes, which are given by an oracle and fulfill the assumption described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters.*

Then Algorithm 7 with input (A, B, ε, p) computes a rigid motion r that maximizes the volume of overlap of $r(A)$ and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$ in time $O\big(T(N)+N(\log N)^2\big)$ where $N = O\big(\varepsilon^{-8}K_A^3 D_A(M_B/m_A)^2\log\frac{2}{p}\big)$.

6.3. Vote generation with 3+1 points

Algorithms 6 and 7 do not use any information about the shape for the generation of the rotation angles. We consider a variant of the probabilistic algorithm for rigid motions that uses the directions present in the shapes to determine the rotation angle in the random experiments. For this variant, we need to be able to test whether a given point is contained in a shape. Therefore, in this section, we assume that in time $T_q(N)$ we can not only generate N random points from a shape B, but also execute N membership queries "$b \in B$?".

A rigid motion is determined by selecting two points $a_1, a_2 \in A$ and one point $b_1 \in B$ uniformly at random. Then we select another point $b_2 \in \mathbb{R}^2$ such that the distances between the points a_1 and a_2 and b_1 and b_2 are the same, i.e., $b_2 = b_1 + |a_2 - a_1| M_\beta \binom{1}{0}$, where $\beta \in [0, 1]$ is selected u.a.r. If b_2 happens to be in B, then (a_1, a_2, b_1, b_2) is a valid random sample. Otherwise, we discard the sample and select new points. In this way, we select uniformly distributed tuples from

$$\mathcal{S} = \{(a_1, a_2, b_1, \beta) \in A^2 \times B \times [0, 1] : a_1 \neq a_2, b_2 = b_1 + |a_2 - a_1| M_\beta \tbinom{1}{0} \in B\}.$$

Generating one random sample for this algorithm involves testing whether a point b_2 lies in the shape B, which takes more time than simply generating a random rotation angle, as in Algorithm 7. Also, we need to analyze how many random samples are discarded by the algorithm. This increases the bound on the required number of random samples. However, experiments with our implementation of the algorithms show that the quality of random samples is significantly better than in the random rotation case. For the same number of random samples, Algorithm 8 computes a much better matching transformation. See Chapter 7 for experimental results.

How often we have to discard samples depends on the shapes in a way that is not immediately clear. For skinny shapes, we cannot control the probability of discarding a sample. We bound this probability for fat shapes. We discussed a number of fatness notions in Chapter 2. Here we use a very simple, home-made definition of fatness that is designed for the purpose. We say that a shape A is κ-fat if it contains a disk of area at least $\kappa|A|$. This simple notion of fatness has the advantage that it is a very weak assumption. Whenever fat shapes are used in this chapter, we refer to this definition.

Algorithm 8 does not only get two shapes, an error tolerance and a desired probability of success as input, but also a fatness constant κ. The shapes are assumed to be κ-fat.

Theorem 6.7 (Approximation property of Algorithm 8)**.** *Let $\kappa, \varepsilon, p \in (0, 1)$ be parameters, and let $A, B \subset \mathbb{R}^2$ shapes such that A and B contain disks of area $\mu_A \geq \kappa|A|$ and $\mu_B \geq \kappa|B|$, respectively, and $\mu_A \leq \mu_B$.*

There is a constant $C > 0$ such that, if 2DClusteringSizeRM3+1(A, ε, κ) returns a positive number $\delta < \frac{\varepsilon\kappa}{3(\sqrt{2}+2\pi)} \frac{|A|}{|\partial A|}$ and 2DSampleSizeRM3+1($B, \varepsilon, \delta, p, \kappa$)

Algorithm 8: 2DMaxOverlap3+1

Input: shapes $A, B \subset \mathbb{R}^2$, fatness constant $\kappa \in (0,1)$, error tolerance $\varepsilon \in (0,1)$, allowed probability of failure $p \in (0,1)$

```
real δ ← 2DClusteringSizeRM3+1(A, ε, κ);
integer N ← 2DSampleSizeRM3+1(B, ε, δ, p, κ);
collection Q ← ∅;
for i = 1 ... N do
    point a1 ← RandomPoint(A);
    point a2 ← RandomPoint(A);
    point b1 ← RandomPoint(B);
    real β ← RandomAngle();
    point b2 ← b1 + |a2 − a1| Mβ (1 0)ᵀ;
    if b2 ∈ B then
        α ← ∠(a2 − a1, b2 − b1)/(2π);
        add(Q, (α, b1 − Mα a1));
    end
end
return 2DFindDensestClusterRM(Q, δ, diam(A));
```

returns an integer $N \geq C\,\varepsilon^{-2}\delta^{-6}\kappa^{-5}\operatorname{diam}(A)^2|B|^2\log\frac{2}{p}$*, then Algorithm 8 with input* $(A, B, \kappa, \varepsilon, p)$ *computes a rigid motion* r^* *such that* $|r^*(A) \cap B|$ *is maximal up to an additive error of* $\varepsilon|A|$ *with probability at least* $1-p$*.*

Note that also in this version of the probabilistic matching δ scales with the shapes and N is invariant under scaling of both shapes with the same factor.

Lemma 6.8. *The density function on the space of rigid motions R that is induced by the random experiment in Algorithm 8 is given by*

$$g_{3+1}(r) = |r(A) \cap B|^2/c,$$

where c is a positive real depending on A and B that satisfies $c \leq |A|^2|B|$*.*

Proof. Let $I = \frac{1}{2\pi}\mathcal{S}^1$. In one random experiment, we select uniformly distributed random elements from the set

$$\mathcal{S} = \{(a_1, a_2, b_1, \beta) \in A^2 \times B \times I : a_1 \neq a_2, b_2 = b_1 + |a_2 - a_1| M_\beta \tbinom{1}{0} \in B\}.$$

The density function of the random variable $X = id_{\mathcal{S}}$ is then

$$f_X(a_1, a_2, b_1, \beta) = \frac{\chi_{\mathcal{S}}(a_1, a_2, b_1, \beta)}{|\mathcal{S}|} = \frac{\chi_A(a_1)\chi_A(a_2)\chi_B(b_1)\chi_B(b_2)}{|\mathcal{S}|},$$

where b_2 is as in the definition of $\mathcal{S}$.

Let Y denote the random variable that corresponds to the rigid motion resulting from one random experiment:

$$Y : \mathcal{S} \to R, \quad Y : (a_1, a_2, b_1, \beta) \mapsto (\alpha, b_1 - M_\alpha a_1),$$

where $\alpha = \angle(a_2 - a_1, b_2 - b_1)/2\pi$. We represent Y as a composite function of the random variable X. Define functions φ_1, φ_2 as follows:

$$\begin{aligned} &\varphi_1 : \mathcal{S} \to \mathbb{R}^6 \times I && \varphi_1 : (a_1, a_2, b_1, \beta) \mapsto (a_1, a_2, b_1 - M_\alpha a_1, \alpha) \\ &\varphi_2 : \mathbb{R}^6 \times I \to R && \varphi_2 : (a_1, a_2, t, \alpha) \mapsto (\alpha, t). \end{aligned}$$

Then $Y = \varphi_2 \circ \varphi_1 \circ X$.

Observe that $\mathcal{S}$ is a Borel set and the function φ_1 and its inverse are bijective and differentiable, so we can apply the transformation formula from Section 4.1 to compute the density function of $\varphi_1(X)$. We first compute the absolute value $J_{\varphi_1}(x)$ of the determinant of the Jacobian matrix of φ_1.

Denote $x = (x_1, \ldots, x_7) = (a_1, a_2, b_1, \beta)$. Because of $\varphi_{1,i}(x) = x_i$ for $i = 1, \ldots, 4$, we have $\frac{\partial \varphi_{1,i}}{\partial x_j} = \delta_{ij}$ for $(i,j) \in \{1,2,3,4\} \times \{1, \ldots, 7\}$. Hence $\det(\frac{\partial \varphi_{1,i}}{\partial x_j})_{i,j=1,\ldots,7} = \det(\frac{\partial \varphi_{1,i}}{\partial x_j})_{i,j=5,\ldots,7}$.

Note that the angle α does not depend on b_1 and it depends linearly on β: We have $\alpha = \beta + \gamma$ where $\gamma = \angle(a_2 - a_1, \binom{1}{0})/2\pi$. See Figure 6.1.

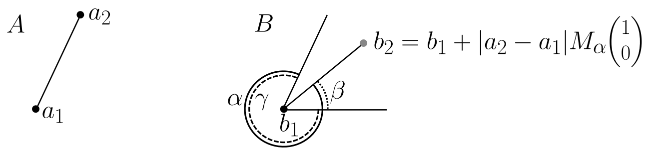

Figure 6.1.:

Therefore $\frac{\partial (b_1 - M_\alpha a_1)_i}{\partial b_{1j}} = \delta_{ij}$, $\frac{\partial \alpha}{\partial b_{1j}} = 0$ for $j = 1, 2$, and $\frac{\partial \alpha}{\partial \beta} = 1$. Now we have

$$\left(\frac{\partial \varphi_{1,i}}{\partial x_j} \right)_{i,j=5,\ldots,7} = \begin{pmatrix} 1 & 0 & \lambda_1 \\ 0 & 1 & \lambda_2 \\ 0 & 0 & 1 \end{pmatrix},$$

for some $\lambda_1, \lambda_2 \in \mathbb{R}$. Thus $J_{\varphi_1}(x) = 1$. The inverse function of φ_1 maps a tuple (a_1, a_2, t, α) to $(a_1, a_2, t + M_\alpha a_1, \alpha + \gamma)$. By Section 4.1, the density function of $\varphi_1(X)$ is given by

$$\begin{aligned} f_{\varphi \circ id_{\mathcal{S}}}(a_1, a_2, t, \alpha) &= \chi_A(a_1)\chi_A(a_2)\chi_B(t + M_\alpha a_1)\chi_B(t + M_\alpha a_2)/|\mathcal{S}| \\ &= \chi_A(a_1)\chi_A(a_2)\chi_{M_{-\alpha}(B-t)}(a_1)\chi_{M_{-\alpha}(B-t)}(a_2)/|\mathcal{S}|. \\ &= \chi_{(A \cap r^{-1}(B))^2}(a_1, a_2)/|\mathcal{S}| \end{aligned}$$

for $r = (\alpha, t) \in R$. The density function of random variable Y on R is then

$$\begin{aligned} g_{3+1}(r) &= \frac{1}{|\mathcal{S}|} \int_{A^2} \chi_{(A \cap r^{-1}(B))^2}(a_1, a_2) d(a_1, a_2) \\ &= |A \cap r^{-1}(B)|^2 / |\mathcal{S}| \\ &= |r(A) \cap B|^2 / |\mathcal{S}|. \end{aligned}$$

□

Lemma 6.9. *The density function g_{3+1} is Lipschitz continuous with constant $L = (\sqrt{2} + 2\pi)\, |\partial A|\, \min\{|A|, |B|\}/c$, where c is as in Lemma 6.8.*

Proof. The density function g_{3+1} induced by Algorithm 8 has the form $f^2(r)/c$, where c is as in Lemma 6.8. Observe that if f is bounded and Lipschitz continuous with Lipschitz constant L_f, then f^2 is Lipschitz continuous with constant $L_{f^2} = 2L_f \sup_x f(x)$ due to the following consideration:

$$\begin{aligned} |f^2(x) - f^2(y)| &= |f(x) - f(y)| \cdot |f(x) + f(y)| \\ &< L_f \cdot |x - y| \cdot |f(x) + f(y)| \\ &\leq 2L_f |x - y| \sup_z f(z). \end{aligned}$$

Because of Lemma 6.5, the Lipschitz constant of the function g_{3+1} is

$$L_{3+1} = (\sqrt{2} + 2\pi)|\partial A| \sup_z f(z)/c. \tag{6.2}$$

The maximal possible area of overlap of two shapes under rigid motions is clearly bounded by the area of the smaller shape. Therefore, the function g_{3+1} is Lipschitz continuous with constant $L_{3+1} = (\sqrt{2} + 2\pi)|\partial A|\, \min\{|A|, |B|\}/c$. □

Proof of Theorem 6.7. Since in Algorithm 8 some random samples are rejected and do not induce a vote in the transformation space, we first determine the necessary number of not rejected experiments, in the following denoted by M, in order to guarantee the required error bound with probability at least $1 - \frac{p}{2}$. Afterwards we determine the total number N of random samples that the algorithm needs to generate in order to record at least M votes in the transformation space with probability at least $1 - \frac{p}{2}$. Then the probability that less than M votes are recorded or the error is larger than $\varepsilon|A|$ is at most p.

Let $f(r)$ denote the area of overlap $|r(A) \cap B|$ and let g_{3+1} be the density function of the probability distribution μ in the space of rigid motions induced by Algorithm 8. Let r^* be a random rigid motion with the maximum number of random rigid motions generated by the algorithm in its δ-neighborhood.

For Algorithm 8, the density function equals $g_{3+1}(r) = f^2(r)/c$, where c is as in Lemma 6.8. Its Lipschitz constant is $L = (\sqrt{2} + 2\pi)|\partial A|\, \min\{|A|, |B|\}/c$ by Lemma 6.9.

Let r be a rigid motion that maximizes the area of overlap of A and B. By Theorem 4.6, we have that

$$\begin{aligned} \tfrac{1}{c}\left(f(r) - f(r^*)\right)\left(f(r) + f(r^*)\right) &= g_{3+1}(r) - g_{3+1}(r^*) \\ &\leq \tfrac{1}{4}\operatorname{diam}(A)\delta^{-3}\Big(C\sqrt{\tfrac{6}{M}} + \tau\Big) + 3L\delta. \end{aligned}$$

The universal constant c from Theorem 4.6 is here called C to avoid a clash of notation.

We want to have at least $1 - \frac{p}{2}$ as probability of success, that is at most $\frac{p}{2}$ as probability of failure, which lets us set $\tau = \sqrt{\frac{1}{2M}\log\frac{4}{p}}$. Using $f(r) \geq \kappa|A|$, which holds by assumption, and substituting τ shows

$$f(r) - f(r^*) \leq \frac{c}{\kappa|A|}\left(\left(C\sqrt{6} + \sqrt{\frac{1}{2}\log\frac{4}{p}}\right)\frac{\operatorname{diam}(A)}{4\delta^3\sqrt{M}} + 3L\delta\right).$$

We need to have $\frac{3c}{\kappa|A|}L\delta < \varepsilon|A|$ to get the error small enough, which is fulfilled for all clustering sizes

$$\delta < \frac{\varepsilon\,\kappa}{3(\sqrt{2} + 2\pi)}\frac{|A|}{|\partial A|}.$$

For each such δ there is an

$$M = O\left(\varepsilon^{-2}\kappa^{-2}\delta^{-6}|B|^2\operatorname{diam}(A)^2\log\tfrac{2}{p}\right)$$

such that, if M is the number of not rejected experiments in Algorithm 8, the additive error is at most $\varepsilon|A|$ with probability at least $1 - \frac{p}{2}$.

Next we determine the total number N of random samples that the algorithm needs to generate in order to record at least M votes with probability at least $1 - \frac{p}{2}$. For that purpose we first determine the probability that one randomly generated sample is not rejected. In the following this probability is denoted by q.

Let $\mathcal{C}_A$ and $\mathcal{C}_B$ denote the largest inscribed circles in A and B, respectively, and let r_A, r_B denote the radii of $\mathcal{C}_A$ and $\mathcal{C}_B$. By the assumptions $|\mathcal{C}_A| \geq \kappa|A|$ and $|\mathcal{C}_B| \geq \kappa|B|$, and by a precondition of the theorem $r_A \leq r_B$.

Consider a circle $\mathcal{C}'_A$ of radius $r_A/4$ contained in A and a circle $\mathcal{C}'_B$ concentric with $\mathcal{C}_B$ of radius $r_B - r_A/2 \geq r_B/2$ in B. The area of $\mathcal{C}'_A$ is at least $\kappa|A|/16$ and the area of $\mathcal{C}'_B$ is at least $\kappa|B|/4$. Then the probability that two randomly selected points a_1 and a_2 from A are both contained in $\mathcal{C}'_A$ is at least $(\kappa/16)^2$. The distance between a_1 and a_2 is at most $r_A/2$. The probability that a randomly selected point b_1 from B is contained in $\mathcal{C}'_B$ is at least $\kappa/4$, and by construction, the complete circle centered at b_1 with radius equal to the distance between a_1 and a_2 is completely contained in $\mathcal{C}_B \subset B$. Therefore, for every choice of two points in $\mathcal{C}'_A$, a point in $\mathcal{C}'_B$ and for every randomly chosen direction, the random sample induces a vote in the transformation space. The probability q that one random sample is not rejected is then at least as large as $\kappa^3/4^5$.

Our algorithm generates random samples independently in every step. For each sample the probability not to be rejected is at least q. Then the expected number of valid samples after N steps is at least qN. Let X denote the number of valid samples after N steps. Using the Chernoff bound (Theorem 4.3) we can determine the number of steps N for which X is not much smaller than qN with high probability:

$$P\left(X < qN - \xi N\right) < 2e^{-\xi^2 N/2}$$

for all $\xi \in (0,1)$. For $\xi = q/2$ we have $P\left(X < \frac{q}{2}N\right) < 2e^{-q^2N/8}$. Restricting this failure probability to be at most $\frac{p}{2}$, we get that for $N \geq 8q^{-2}\log\frac{4}{p}$ the number of votes X is at least $qN/2$ with probability at least $1-\frac{p}{2}$. Then with $N \geq 2M/q$ random samples the algorithm generates at least M votes with probability at least $1-\frac{p}{2}$.

Finally, choosing

$$N \geq \max\left\{\tfrac{2M}{q},\ \tfrac{8}{q^2}\log\tfrac{4}{p}\right\} = O\Big(\max\big\{\varepsilon^{-2}\delta^{-6}\kappa^{-5}\,\mathrm{diam}(A)^2|B|^2\log\tfrac{2}{p}\,,\ \kappa^{-6}\log\tfrac{2}{p}\big\}\Big),$$

we have that with probability at least $1-\frac{p}{2}$ the number of recorded votes is sufficiently large, and therefore with probability at least $1-p$ the approximation of the maximum area of overlap differs from the optimum by at most $\varepsilon|A|$. Since $\kappa^{-1} = O(\delta^{-1})$, the term $\kappa^{-6}\log\frac{2}{p}$ can be dropped. □

Corollary 6.10 (Runtime of Algorithm 8)**.** *Let $\kappa, \varepsilon, p \in (0,1)$ be parameters, and let $A, B \subset \mathbb{R}^2$ shapes, which are given by an oracle and fulfill the assumptions described in Section 2.2, such that A and B contain disks of area $\mu_A \geq \kappa|A|$ and $\mu_B \geq \kappa|B|$, respectively, and $\mu_A \leq \mu_B$. Then Algorithm 8 with input $(A, B, \kappa, \varepsilon, p)$ computes a rigid motion that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$ in time $O\big(T_q(N) + N(\log N)^2\big)$ where $N = O(\varepsilon^{-8}\kappa^{-11}K_A^3 D_A(M_B/m_A)^2\log\frac{2}{p})$.*

7. Evaluating the probabilistic matching approach in 2D

The purpose of this chapter is to further investigate the probabilistic matching approach in the plane, as described in the two previous chapters. Matching under translations for arbitrary dimension is described in Chapter 5 in Algorithms 2 and 5. In the previous chapter, three algorithms for matching solid shapes in the plane under rigid motions were studied.

In the first section of this chapter, we present simple, deterministic, grid-based matching algorithms for translations and rigid motions in the plane, which are in some sense the deterministic versions of the probabilistic algorithms. We compare them to the probabilistic algorithms.

In Section 7.2, we study whether the probabilistic matching approach can be applied to matching under similarities. We describe a random experiment that results in matching under similarities and determine the density function for two parameterizations of the space of similarities.

In the last section of this chapter, we present experimental results for translations in dimension 2 (Algorithm 1) and rigid motions (Algorithms 7 and 8), which indicate that our theoretical bounds are overly pessimistic and that all but one of the presented algorithms work well. The experimental results are not meant as rigorous study, but to prove that the algorithms are easily implementable and give reasonable results.

7.1. Comparison with a simple, deterministic approach

One could ask whether it is necessary to make such heavy use of randomness in our algorithms. There might be a simple deterministic algorithm that serves the same purpose and is equally simple as the probabilistic matching approach. We study a deterministic analogue of the probabilistic algorithm to check to what extent randomness helps. We only do so for the plane, but we immediately take care of the case of rigid motions.

We present a deterministic algorithm that approximately maximizes the area of overlap of two shapes under translations and rigid motions in the plane. Here we model a shape A as a region in the plane whose boundary is a set of simple closed curves. Additionally we assume that we can generate all points from $A \cap \delta\mathbb{Z}^2$ for each grid width $\delta > 0$. We formulate the results for the example of finite unions of triangles that have non-intersecting interiors.

The deterministic algorithm for translations

The idea of the algorithm for translations is as follows. The shapes A and B are approximated by orthogonal polygons, namely by unions of squares from the grid $\delta\mathbb{Z}^2$, for some positive grid width δ, whose centers are contained in the shapes. For these orthogonal polygons, only translations that map grid points onto each other are considered, which is a finite number of translations.

The grid width is chosen so small that the area of the union of the squares that intersects the boundary of the shapes is only a small fraction of the area of A and B.

The algorithm counts how many pairs of grid points from $A \times B$ are mapped onto each other by each translation and outputs the most frequent translation. All pairs of grid points (a, b) with $t = b - a$ for which the square centered at a is contained in A and the square centered at b is contained in B contribute δ^2 to the area of overlap $|t(A) \cap B|$. Therefore, for a translation t, the number of pairs (a, b) such that a is the center of a square in A and b is the center of a square in B and $t = b - a$ gives an estimate of the area of overlap $|t(A) \cap B|$. From the definition of the Lebesgue measure follows $\lim_{\delta \to 0} |t(A) \cap B \cap \delta\mathbb{Z}^2| \cdot \delta^2 = |t(A) \cap B|$. Figure 7.1 illustrates the idea of the algorithm.

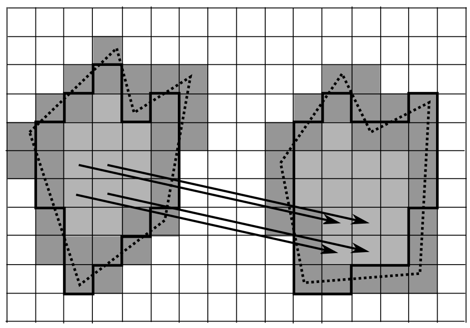

Figure 7.1.: The figure shows two polygonal shapes, drawn with dotted lines, and a grid of some example width. In the algorithm, the grid width is chosen so small that the union of the squares that intersect the boundary, which is drawn in dark gray, has a small area. The boldly drawn orthogonal polygons are the unions of squares whose centers are contained in the shapes. The figure also shows four example pairs of centers of squares that give the same translation vector.

Our algorithm, which is given a pseudocode description in Algorithm 9, obtains as input shapes A, B and an error tolerance ε. It computes some grid width δ that depends on A, B and ε. The next theorem gives an upper bound on the grid width that guarantees that the output translation t^* is optimal up to an additive error of at most $\varepsilon|A|$. It also outputs the number of pairs of grid points $m(t^*)$ that have t^* as difference vector. This information will be used for matching under rigid motions.

Algorithm 9: 2DGridMaxOverlapT

```
Input: finite sets of triangles A, B ⊂ ℝ² that have pairwise disjoint
       interiors, error tolerance ε ∈ (0,1)
δ ← 2DGridWidthT(A, B, ε);          // Compute positive grid with δ
Initialize hashtable H;
for (a,b) ∈ (A ∩ δℤ²) × (B ∩ δℤ²) do
    t ← b − a;
    if t ∈ H then
        increase value m(t) in H by 1;
    else
        insert key t with value m(t) = 1 into H;
end
Compute t* in H with maximal value m(t*);
return (t*, m(t*));
```

Theorem 7.1 (Correctness of Algorithm 9)**.** *Let $A, B \subset \mathbb{R}^2$ be unions of finitely many triangles that have pairwise disjoint interiors, and let $\varepsilon \in (0,1)$ be a parameter. Let $M(A,B)$ be the minimal length of the boundaries of the triangles of A and B, and let $L(A,B)$ be the total length of the boundaries of A and B. If 2DGridWidthT(A,B,ε) returns a positive number $\delta \le \min\left\{\frac{1}{12}\frac{\varepsilon|A|}{L(A,B)}, \frac{1}{4}M(A,B)\right\}$, then Algorithm 9 with input (A,B,ε) computes a translation t^* such that $|t^*(A) \cap B| \ge |t(A) \cap B| - \varepsilon|A|$ for all translations $t \in \mathbb{R}^2$, and an integer $m(t^*)$ such that $\big||t^*(A) \cap B| - m(t^*)\delta^2\big| \le \frac{5}{12}\varepsilon|A|$.*

We prepare the correctness proof of Algorithm 9 with a series of lemmas. A *(half open) square of the grid $\delta\mathbb{Z}^2$* is a set of the form $[x\delta,(x+1)\delta) \times [y\delta,(y+1)\delta)$ for two integers x, y. We show that a closed curve cannot intersect too many squares in a grid.

Lemma 7.2. *Let C be a simple closed curve of length ℓ in the plane. Let $\delta > 0$ and let k be the number of squares of $\delta\mathbb{Z}^2$ that C intersects. Then, $k \le 4(\ell/\delta + 1)$.*

Proof. If C intersects at most four squares, the claim is clearly true. If C intersects more than four squares, let $c_1 \in C$ be a point at which C crosses the boundary of a square. Following C clockwise from c_1, let S_1 be the first square that C passes through, i.e. $C \cap S_1 \neq \emptyset$. Let c_2 be the next point on the curve C that is on the boundary of a square such that from c_2 on C passes through a previously not seen square. Proceeding in this way, $c_1, \ldots, c_k$ are defined, and as a consequence k equals the number of squares that C intersects. Compare Figure 7.2 for the definition of $c_1, \ldots, c_k$. Each piece of the curve from c_i to c_{i+4} passes through at least 4 squares that were previously not visited. By inspection of the connected arrangements of 4 squares, each such piece of the curve has length at least δ. Therefore $\ell \ge \delta\lfloor k/4 \rfloor \ge \delta(k/4 - 1)$, which implies the claim. □

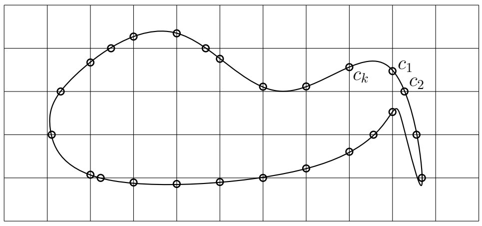

Figure 7.2.: The figure shows an example curve C and points $c_1, \ldots c_k$.

In our model, the boundary of a plane shape is a set of simple closed curves. We bound the number of grid squares that the boundary of a shape can intersect.

Corollary 7.3. *Let D be a finite set of simple closed curves, each of which has length at least $c > 0$. Let L be the total length of the curves in D. Let $0 < \delta \leq c/4$ and let K be the number of squares of $\delta\mathbb{Z}^2$ that are intersected by at least one of the curves in D. Then, $K \leq 5L/\delta$.*

Proof. Let C be a simple closed curve of length $\ell \geq c > 0$. Let $\delta \leq c/4$ and let k be the number of squares of $\delta\mathbb{Z}^2$ that C intersects. By Lemma 7.2, we have $k \leq 4\ell/\delta + 4 \leq 4\ell/\delta + c/\delta \leq 5\ell/\delta$.

Let $C_1, \ldots, C_n$ be simple closed curves of lengths $\ell_1, \ldots, \ell_n$ such that C_i intersects k_i squares of $\delta\mathbb{Z}^2$. Let K be the number of squares that are intersected by some C_i. Then, $K \leq k_1 + \cdots + k_n \leq 5\ell_1/\delta + \cdots + 5\ell_n/\delta = 5L/\delta$. □

Let $G(A) = A \cap \delta\mathbb{Z}^2$ for some fixed $\delta > 0$. Let m be the number of pairs of grid points in $G(A) \times G(B)$ that have some fixed vector $t \in \mathbb{R}^2$ as their difference vector. Next we prove the key lemma for the algorithm. It states that $m\delta^2$ is indeed an approximation of $|t(A) \cap B|$.

Lemma 7.4. *Let t be a translation in $G(B) \oplus (-G(A))$, and let m be the number of pairs $(a, b) \in G(A) \times G(B)$ such that $t = b - a$. Let k be the number of squares of $\delta\mathbb{Z}^2$ that $\partial A \cup \partial B$ intersects. Then $(m - k)\delta^2 \leq |t(A) \cap B| \leq (m + k)\delta^2$.*

Proof. For two distinct pairs (a, b) and (a', b') that yield the same translation t, $a \neq a'$ and $b \neq b'$ hold. Each pair (a, b) such that the square with center a is contained in A, the square with center b is contained in B, and $t = b - a$ contributes the area of one square, which is δ^2, to the area of overlap. Consider the shape S_A that consists of all squares of $\delta\mathbb{Z}^2$ that are intersected by A. Also, consider the shape s_A that consists of all squares that are contained in A. Clearly $s_A \subseteq A \subseteq S_A$. Similarly, let S_B be the union of all squares of $\delta\mathbb{Z}^2$ that are intersected by B and s_B the union of all squares that are contained in B. We have

$$(m - k)\delta^2 \leq |t(s_A) \cap s_B| \leq |t(A) \cap B| \leq |t(S_A) \cap S_B| \leq (m + k)\delta^2.$$

□

Finally we prove that for every translation t for which the area of overlap of $t(A)$ and B is not too small, there is a translation close to t that is the difference of a grid point from A and a grid point from B.

Lemma 7.5. *Let k be the number of squares of $\delta\mathbb{Z}^2$ that $\partial A \cup \partial B$ intersects. For every translation $\bar{t} \in \mathbb{R}^2$ such that $|\bar{t}(A) \cap B| > k\delta^2$, there is a translation $t \in G(B) \oplus (-G(A))$ such that $|t - \bar{t}| \leq \sqrt{2}\delta$.*

Proof. Let $\bar{t} \in \mathbb{R}^2$ be a translation vector. If for all $(a, b) \in A \times B$ such that $\bar{t} = b - a$ hold that ∂A intersects the square of $\delta\mathbb{Z}^2$ that contains a or ∂B intersects the square that contains b, then $|\bar{t}(A) \cap B| \leq k\delta^2$. Therefore, if $|\bar{t}(A) \cap B| > k\delta^2$, then there is $(\bar{a}, \bar{b}) \in A \times B$ such that $\bar{t} = \bar{b} - \bar{a}$ and the squares containing $\bar{a}$ and $\bar{b}$ are contained in A and B, respectively. Let a be the center of the square containing $\bar{a}$ and let b be the center of the square containing $\bar{b}$. Then $t = b - a \in G(B) \oplus (-G(A))$ and $|t - \bar{t}| = |b - a - (\bar{b} - \bar{a})| \leq |a - \bar{a}| + |b - \bar{b}| \leq \sqrt{2}\delta$. □

Proof of Theorem 7.1. Let A, B, ε, t^* and $m(t^*)$ be defined as in the theorem, and assume that δ is at most as large as the given upper bound. Let k be the number of squares of $\delta\mathbb{Z}^2$ intersected by $\partial A \cup \partial B$. The set $\partial A \cup \partial B$ is a set of simple closed curves. Let $L(A, B)$ the sum of the lengths of the boundaries of A and B. Using $\delta \leq \frac{1}{4}M(A, B)$, we have $k \leq 5L(A, B)/\delta$ by Corollary 7.3. We have $k\delta^2 \leq 5L(A, B)\delta \leq \frac{5}{12}\varepsilon|A|$. By Lemma 7.4, it holds that $m(t^*)$ is the claimed approximation of $|t^*(A) \cap B|$.

If the maximal area of overlap is less than or equal to $\varepsilon|A|$, then every translation t^* is a valid output. Assume that there is a translation $t \in \mathbb{R}^2$ such that $|t(A) \cap B| > \varepsilon|A|$.

By Lemma 7.4 and the above inequality, for every $t \in G(B) \oplus (-G(A))$ holds that $m(t)\delta^2 + \frac{5}{12}\varepsilon|A| \geq |t(A) \cap B| \geq m(t)\delta^2 - \frac{5}{12}\varepsilon|A|$ where $m(t)$ is the number of pairs $(a, b) \in G(A) \times G(B)$ such that $t = b - a$.

Let $\bar{t} \in \mathbb{R}^2$ be a translation that maximizes the area of overlap, and let t^* be the output of Algorithm 9. By Lemma 7.5, there is a $t \in G(B) \oplus (-G(A))$ such that $|\bar{t} - t| \leq \sqrt{2}\delta$. By Corollary 3.5, we have

$$|t(A) \cap B| \geq |\bar{t}(A) \cap B| - \tfrac{\sqrt{2}}{2}\delta|\partial A| > |\bar{t}(A) \cap B| - \tfrac{1}{12}\varepsilon|A|.$$

Since $m(t^*) \geq m(t)$, we have

$$|t^*(A) \cap B| + \tfrac{5}{12}\varepsilon|A| \geq |t(A) \cap B| - \tfrac{5}{12}\varepsilon|A|.$$

Because $\frac{5}{12} + \frac{5}{12} + \frac{1}{12} < 1$, we have $|t^*(A) \cap B| > |\bar{t}(A) \cap B| - \varepsilon|A|$. □

The deterministic algorithm for rigid motions

A rigid motion in the plane is defined by a counterclockwise rotation angle $\alpha \in [0, 2\pi)$ and a translation vector $t \in \mathbb{R}^2$. We denote the rotation matrix of the angle α by M_α.

We use the algorithm for translations to compute a matching under rigid motions. For $k = 1, \dots, q$, where q is a large enough integer, we apply the algorithm for translations to the shapes $M_{2\pi k/q}(A)$ and B with error tolerance $\varepsilon/4$. We output the translation t for which $m(t)$ is the largest and the corresponding angle $2\pi k/q$. See the pseudocode description of Algorithm 10.

Algorithm 10: 2DGridMaxOverlapRMT

```
Input: finite sets of triangles A, B ⊂ ℝ² that have pairwise disjoint
       interiors, error tolerance ε ∈ (0,1)
q ← 2DAngleWidthRM(A, B, ε);
// Discretization size for the angle is 1/q
V ← 0;   t* ← (0,0);   α* ← 0;
for  k = 1, ..., q do
    A_k ← M_{2πk/q} A;                          // Compute rotated A
    // Use algorithm for translations to match A_k and B
    (point, integer) (t, v) ← 2DGridMaxOverlapT(A_k, B, ε/4);
    if V < v then
        t* ← t;                                 // Update translation
        α* ← 2πk/q;                             // Update angle
        V ← v;                                  // Update estimated overlap
end
return (α*, t*);
```

We prove a bound on the discretization size of the angle that ensures approximation up to the given error.

Theorem 7.6 (Correctness of Algorithm 10)**.** *Let $A, B \subset \mathbb{R}^2$ be unions of finitely many triangles that have pairwise disjoint interiors, and let $\varepsilon \in (0,1)$ be a parameter. If 2DAngleWidthRM(A, B, ε) returns a positive integer $q \geq \frac{\operatorname{diam}(A)|\partial A|}{2\varepsilon|A|}$, then Algorithm 10 computes with input (A, B, ε) a rigid motion r^* such that $|r^*(A) \cap B| \geq |r(A) \cap B| - \varepsilon|A|$ for all rigid motions r.*

Proof. Denote by t_β the output of Algorithm 9 on input of $(M_\beta A, B, \varepsilon)$. Let (α^{opt}, t^{opt}) be a rigid motion that maximizes the area of overlap of A and B. For all $x \in \mathbb{R}^2$ and all angles $\beta \in [0, 2\pi)$, we have that $|M_{\alpha^{opt}}x - M_\beta x| \leq |\alpha^{opt} - \beta| \cdot |x|$. Let β be an angle such that $|\alpha^{opt} - \beta| < 1/q$. Since by assumption $0 \in A$, we have $|M_{\alpha^{opt}}a - M_\beta a| \leq \operatorname{diam}(A)/q$ for all $a \in A$. By Theorem 3.5, we have

$$\big||(M_{\alpha^{opt}}A + t^{opt}) \cap B| - |(M_\beta A + t^{opt}) \cap B|\big| \leq \operatorname{diam}(A)|\partial A|/(2q). \tag{7.1}$$

For $q \geq \frac{\operatorname{diam}(A)|\partial A|}{2\varepsilon|A|}$, the right hand side is at most $\varepsilon|A|$.

Let $m_\beta(t)$ be the number of pairs of grid points that have t as their difference, as counted by Algorithm 9 with input $(M_\beta A, B, \varepsilon)$. Let (α^*, t^*) be the output of

Algorithm 10. We know

$$|(M_{\alpha^*}A + t^*) \cap B| \geq m_{\alpha^*}(t^*)\delta^2 - \varepsilon|A|$$

by Theorem 7.1. By the definition of Algorithm 10, $m_{\alpha^*}(t^*) \geq m_\beta(t_\beta)$. Applying Theorem 7.1 two more times and using Inequality (7.1), we obtain

$$\begin{aligned} |(M_{\alpha^*}A + t^*) \cap B| &\geq m_{\alpha^*}(t^*)\delta^2 - \varepsilon|A| \\ &\geq m_\beta(t_\beta)\delta^2 - \varepsilon|A| \\ &\geq |(M_\beta A + t_\beta) \cap B| - 2\varepsilon|A| \\ &\geq |(M_\beta A + t^{opt}) \cap B| - 3\varepsilon|A| \\ &\geq |(M_{\alpha^{opt}}A + t^{opt}) \cap B| - 4\varepsilon|A|. \end{aligned}$$

□

Runtimes and comparison with the probabilistic approach

We briefly discuss the runtimes of Algorithm 9 and Algorithm 10 and compare them with the time bounds on the probabilistic matching algorithms. We assume that the input shapes A and B both are sets of at most n triangles that have pairwise disjoint interiors.

Let $L(A, B)$ the total length of the boundary. We can compute $L(A, B)$ in quadratic time as a preprocessing step. We will see in a moment that the time bound involves a factor of n^2 anyway. After also having computed $|A|$ and $M(A, B)$, which we can do in linear time, we compute $\delta = \min\left\{ \frac{\varepsilon|A|}{12\,L(A,B)}, \frac{1}{4}M(A, B) \right\}$.

For each triangle Δ, the set $\Delta \cap \delta\mathbb{Z}^2$ can be computed in time $O(|\Delta \cap \delta\mathbb{Z}^2|) = O(|\Delta|/\delta^2 + |\partial\Delta|/\delta + 1)$ in a straightforward way. The set $G(A) = A \cap \delta\mathbb{Z}^2$ has size $O\big(|A|/\delta^2 + |\partial A|/\delta + n\big)$ and can be computed within the same time bound. Note that we do not have to store $G(A)$ at once, but that we can also output one grid point after the other while storing only a constant amount of information. Using universal hashing, we can compute a translation t with maximal value $m(t)$ in expected time $O(|G(A)| \cdot |G(B)|)$ and space $O(|G(B) \oplus (-G(A))|)$. This time bound involves the sum of n^2 and ε^{-4}, as well as terms in $|A|$, $|\partial A|$, $|B|$, $|\partial B|$, $L(A, B)$ and $M(A, B)$.

Table 7.1 summarizes the time bounds for the different algorithms for the case of matching unions of triangles in the plane under translations and rigid motions. For unions of at most n triangles, we have $T(N) = O(n + N \log n)$, if we generate one random point at a time. Furthermore $T_q(N) = O((n + N) \log n)$. See Chapter 2 for details.

The runtime of the probabilistic algorithm for translations in the plane is linear in n, so the dependence on n is significantly better for the probabilistic approach. However, the growth rate of the approximation error is ε^{-6} in the time bound on the probabilistic approach, compared to ε^{-4} for the deterministic approach. For the probabilistic adaptive algorithm, described in Section 5.3, the growth rate of the time bound in n is $n \log n$ and the growth rate in ε is $\varepsilon^{-\frac{36}{7}}$.

Translations			
Algo. 2	prob. incr.	$O^g(n + N \log N)$	where $N = O^g(\varepsilon^{-6} \log \frac{2}{p})$
Algo. 5	prob. incr.	$O^g(n \log n + N \log N)$	where $N = O^g(\varepsilon^{-\frac{36}{7}} \log \frac{2}{p})$
Algo. 9	determ.	$O^g(n^2 + \varepsilon^{-4})$	

Rigid motions			
Algo. 6	prob.	$O^g(n\varepsilon^{-1} \log n + N \log N)$	where $N = O^g(\varepsilon^{-\frac{43}{7}} \log \frac{2}{p})$
Algo. 7	prob. incr.	$O^g(n + N(\log N)^2)$	where $N = O^g(\varepsilon^{-8} \log \frac{2}{p})$
Algo. 8	prob. incr.	$O^g(n \log n + N(\log N)^2)$	where $N = O^g(\varepsilon^{-8} \log \frac{2}{p})$
Algo. 10	determ.	$O^g(n^2\varepsilon^{-1} + \varepsilon^{-5})$	

Table 7.1.: Comparison of the time bounds for different algorithms for the matching of unions of at most n triangles in the plane with error bound ε and allowed probability of failure p under translations and rigid motions. The second column indicates if the algorithm works incremental (incr.), probabilistic (prob.) or deterministic (determ.). The notation O^g indicates that there is a hidden dependence on geometric parameters as the area, length of the boundary and the diameter of the input shapes and in the case of Algorithm 8 on the fatness parameter κ (Section 6.3).

Let us determine the runtime of the deterministic algorithm for rigid motions. For each angle α, it takes time linear in n to compute $M_\alpha(A)$. We can do the same trick as in Algorithm 11 and not actually rotate A, but rotate each sample point $a \in A$. So the runtime for rigid motions is

$$O(\varepsilon^{-1}|A|^{-1} \operatorname{diam}(A)|\partial A| \cdot |G(A)| \cdot |G(B)|).$$

This bound involves the summands $\varepsilon^{-1}n^2$ and ε^{-5} and terms depending on the areas, diameters and lengths of the boundaries of the shapes.

The probabilistic approach has the advantage of working incrementally. The experiments indicate that the bound on the number of random points that we proved is overly pessimistic, as we will see in Section 7.3.

The discretization in the deterministic algorithm introduces a new type of error since for a fixed grid width the result cannot be arbitrarily good. The optimal transformation has a certain distance to the grid and the quality of the output depends on this distance. When using the probabilistic approach with bucket clustering, this type of error can also occur. By doing the latter, we will see in Section 7.3 that this error type can worsen the result considerably.

When used with the clustering method that is described in the algorithms, the probabilistic approach does not have this problem. If the clustering size is below a certain threshold for the probabilistic matching approach, as described in Algorithms 2, 5, 7 and 8, then the error of the output tends to 0 with high probability, when the number of random samples tends to $+\infty$, even if the parameter δ stays fixed.

7.2. Similarities

A *similarity* σ in the plane is a map that consists of rotation, scaling and translation. It has the form $\sigma(x) = \lambda M_\alpha x + t$ where $\lambda \in [0, +\infty)$ is a scaling factor, $\alpha \in [0, 2\pi)$ a counterclockwise rotation angle and $t \in \mathbb{R}^2$ a translation vector. We use two different parameterizations of similarities, namely $(\lambda, \alpha, t_1, t_2)$ and $(\lambda \cos\alpha, -\lambda \sin\alpha, t_1, t_2)$.

It is natural to generalize the probabilistic matching approach to similarities since, for two of points $a, a' \in A$, $a \neq a'$, and two points $b, b' \in B$, $b \neq b'$, there is exactly one similarity that maps a onto b and a' onto b'. We have $\sigma(a) = b$ if and only if the translation vector t equals $-\lambda M_\alpha a + b$. With this, $\sigma(a') = b'$ is equivalent to $\sigma(a) - \sigma(a') = b - b'$. Because of $\sigma(a) - \sigma(a') = \lambda M_\alpha(a - a')$, this is true if and only if the rotation angle α is the counterclockwise angle between the vectors $a' - a$ and $b' - b$ and the scaling factor λ equals $\frac{|b-b'|}{|a-a'|}$.

With this observation, we can apply the probabilistic matching approach to matching under similarities in the plane. In each random experiment, we draw points $a, a' \in A$ and points $b, b' \in B$ uniformly at random and store a vote for the unique similarity that maps a onto a' and b onto b'.

In the following, we determine the density function on the space of similarities for the two parameterizations of the space of similarities mentioned above.

Consider the parameterization of similarities as 4-tuples $(s_1, s_2, t_1, t_2) \in \mathbb{R}^4$ denoting the map $x \mapsto \begin{pmatrix} s_1 & s_2 \\ -s_2 & s_1 \end{pmatrix} x + \begin{pmatrix} t_1 \\ t_2 \end{pmatrix}$.

By the above observations, we have the following formulas:

$$\begin{aligned}
\lambda &= \frac{|b' - b|}{|a' - a|}, \\
s_1 &= \lambda \cos\angle(a' - a, b' - b) = \frac{(a'_1 - a_1)(b'_1 - b_1) + (a'_2 - a_2)(b'_2 - b_2)}{|a' - a|^2}, \\
s_2 &= -\lambda \sin\angle(a' - a, b' - b) = \frac{-(a'_1 - a_1)(b'_2 - b_2) + (a'_2 - a_2)(b'_1 - b_1)}{|a' - a|^2}, \\
t_1 &= \lambda(s_1 a_1 + s_2 a_2) + b_1, \\
t_2 &= \lambda(-s_2 a_1 + s_1 a_2) + b_2.
\end{aligned}$$

Let X be a random vector that draws tuples $(a, a', b, b') \in \mathbb{R}^8$ u.a.r. from $A^2 \times B^2$, i.e. with density function $f_X(a, a', b, b') = \frac{\chi_{A^2 \times B}(a, a', b, b')}{|A|^2 \cdot |B|^2}$. Let $\sigma = (s_1, s_2, t_1, t_2)$ be the unique similarity that maps a onto b and a' onto b'. Define

$$\varphi : (a, a', b, b') \mapsto (a, a', s_1, s_2, t_1, t_2).$$

We are interested in the density function of the projection of the random vector $\varphi(X)$ to the last four coordinates. Denote $x = (x_1, \ldots, x_8) = (a_1, a_2, a'_1, a'_2, b_1, b_2, b'_1, b'_2)$. Applying elementary rules to compute the partial derivatives shows that the Jacobian

$$\det\Big(\frac{\partial \varphi_i}{\partial x_j}\Big)_{i,j=1,\ldots,8} = -\frac{1}{|a' - a|^2}.$$

The inverse function is given by $\varphi^{-1}(a, a', s_1, s_2, t_1, t_2) = (a, a', \sigma(a), \sigma(a'))$ where σ denotes the similarity map (s_1, s_2, t_1, t_2) since σ is the unique similarity that maps a onto b and a' onto b'. By Section 4.1, $(a, a', s_1, s_2, t_1, t_2)$ is distributed with density function $f_{\varphi(X)}(a, a', s_1, s_2, t_1, t_2) = f_X(a, a', \sigma(a), \sigma(a'))|a-a'|^2$. Therefore σ is distributed with density function

$$h(\sigma) = \frac{1}{|A|^2|B|^2} \int_{A\times A} |a - a'|^2 \chi_{A^2\times B^2}(a, a', \sigma(a), \sigma(a'))da\, da'.$$

We have $\chi_B(\sigma(a)) = \chi_{\sigma^{-1}(B)}(a)$, and hence $h(\sigma) = \frac{1}{|A|^2|B|^2} \int_{(A\cap\sigma^{-1}(B))^2} |a-a'|^2 da\, da'$. Furthermore

$$\int_{(A\cap\sigma^{-1}(B))^2} |a - a'|^2 da\, da' = \frac{1}{\lambda^2} \int_{(\sigma(A)\cap B)^2} |a - a'|^2 da\, da'$$

where $\lambda = \sqrt{s_1^2 + s_2^2}$ is the scaling factor of σ, as $\int_C f(x)dx = \int_{\lambda C} f(\frac{x}{\lambda})dx$ for $C \subset \mathbb{R}^k$ and $\lambda > 0$.

We apply the parallel axis theorem from physics to compute the integral expression $\int_{(\sigma(A)\cap B)^2} |a - a'|^2 da\, da'$. For a measurable set $C \subset \mathbb{R}^2$ of finite area and a point $z \in C$, let $I_z(C) = \int_C |x - z|^2 dx$. In physics, $I_z(C)$ is called the *moment of inertia* of C about the point z. If c is the center of mass of C, $I_c(C)$ is of special interest. The parallel axis theorem states that $I_z(C) = I_c(C) + |C| \cdot |z - c|^2$. Let c be the center of mass of $\sigma(A) \cap B$. We obtain

$$\begin{aligned} \int_{(\sigma(A)\cap B)^2} |a - a'|^2 da\, da' &= \int_{\sigma(A)\cap B} I_{a'}(\sigma(A) \cap B)da' \\ &= \int_{\sigma(A)\cap B} \big(I_c(\sigma(A) \cap B) + |\sigma(A) \cap B| \cdot |c - a'|^2\big)da' \\ &= 2\, |\sigma(A) \cap B|\, I_c(\sigma(A) \cap B). \end{aligned}$$

In total,

$$h(\sigma) = \frac{2\, |\sigma(A) \cap B|\, I_c(\sigma(A) \cap B)}{\lambda^2\, |A|^2\, |B|^2},$$

where c is the center of mass of $\sigma(A) \cap B$.

Next we consider the parameterization of planar similarities as 4-tuples $(\lambda, \alpha, t_1, t_2) \in \mathbb{R}^+ \times [0, 2\pi) \times \mathbb{R}^2$ denoting the map $x \mapsto \lambda M_\alpha x + (t_1\ t_2)^T$. To determine the density function on this space, we study the function ψ that maps (s_1, s_2, t_1, t_2) to $(\lambda, \alpha, t_1, t_2)$. We have

$$\psi(\sigma) = \begin{cases} \left(\sqrt{s_1^2 + s_2^2}, \arccos\Big(\frac{s_1}{\sqrt{s_1^2+s_2^2}}\Big), t_1, t_2\right) & \text{for } s_2 \le 0 \\ \left(\sqrt{s_1^2 + s_2^2}, \arccos\Big(-\frac{s_1}{\sqrt{s_1^2+s_2^2}}\Big) - \pi, t_1, t_2\right) & \text{for } s_2 \ge 0 \end{cases}$$

where arccos : $[-1,1] \to [0,\pi]$. An easy computation shows

$$\left|\det\left(\frac{\partial \psi_i}{\partial \sigma_j}\right)_{i,j=1,\dots,4}\right| = \frac{1}{\sqrt{s_1^2+s_2^2}}.$$

By the results of Section 4.1, the density function on $\mathbb{R}^+ \times [0,2\pi) \times \mathbb{R}^2$ is given by

$$h^*(\lambda,\alpha,t_1,t_2) = \sqrt{s_1^2+s_2^2}\, h(s_1,s_2,t_1,t_2) = \frac{2\,|\sigma(A)\cap B|\, I_c(\sigma(A)\cap B)}{\lambda\,|A|^2\,|B|^2},$$

where σ is the map represented by (λ,α,t_1,t_2) and (s_1,s_2,t_1,t_2) and c is the center of mass of $\sigma(A)\cap B$.

Thus the probabilistic scheme maximizes the function h or h^*, depending on the parameterization of the space. Observe that the maxima of these functions change if A and B are swapped. The functions h and h^* do not seem to be useful objective functions. With the help of Theorem 4.7, we could use the probabilistic scheme for matching shapes under similarities with respect to the area of overlap, if we could compute moments of inertia of the form $I_c(\sigma(A)\cap B)$ fast.

When weighting the votes with the scaling factor λ, the product of the area of overlap and the moment of inertia is maximized. The problem is that it is then a good solution to blow up shape A such that it contains B. When the votes are not weighted with the scaling factor, shrinking the shape with a very small scaling factor obtains a high value.

7.3. Experimental results

We show and discuss experimental results of a test software, which was written by Philipp Preis in Java.

Shapes. The shapes are modeled as unions of polygons in the plane that have pairwise disjoint interiors. The shape files are text files that contain a list of polygonal chains, which may be closed or open. For our purposes, they are always closed. Each polygonal chain is represented by the sequence of the Euclidean coordinates of its vertices. Figure 7.3 shows the shapes with which the experiments in this section are performed. The shapes were drawn with the shape editor provided by the software.

Shape FUBear.seg depicts the bear from the logo of Freie Universität Berlin. The shape consists of 10 polygons with 344 vertices in total. It has an area of 241,757 square unit length, a boundary of 7414 unit length, a diameter of 868, and a bounding box of 725×572. We will match FUBear.seg to itself under translations and rigid motions. Whenever matching results for only one shape are mentioned, this shape is taken as source and target shape.

Shape BTor.seg depicts the “Brandenburger Tor”. The shapes sightsT.seg and sightsR.seg show three Berlin landmarks each. We will match BTor.seg to sightsT.seg under translations and to sightsR.seg under rigid motions. This is a partial matching task since the “Brandenburger Tor” is contained in shape sightsT.seg and sightsR.seg.

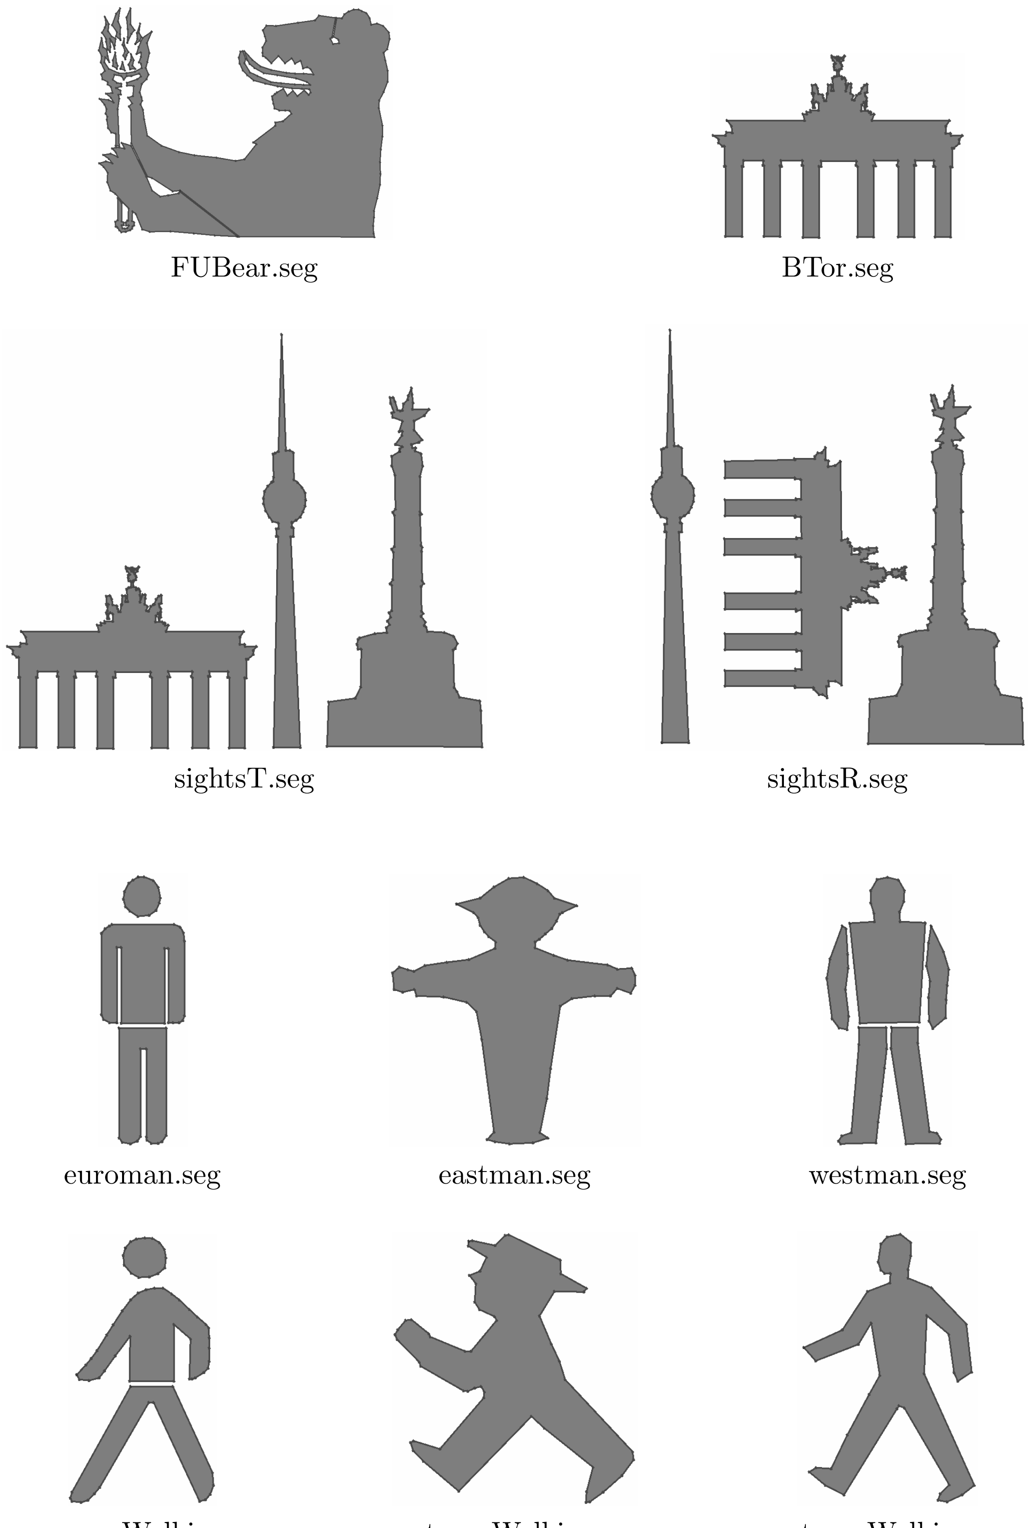

Figure 7.3.: The input shapes for the experiments.

Shapes euroman.seg and euromanWalking.seg show the standardized European traffic light symbols, eastman.seg and eastmanWalking.seg show the former GDR traffic light symbols, and westman.seg and westmanWalking.seg show the former FRG traffic light symbols.

We use the shapes depicting traffic light symbols for giving examples of matching results for shapes that differ considerably. Figure 7.4 summarizes relevant parameters of some of the shapes, for example the area and number of vertices.

shape A	# vertices	$\lvert A\rvert$	$\lvert\partial A\rvert$	diam(A)	$\frac{\lvert\partial A\rvert^2}{\lvert A\rvert}$	$\frac{\lvert A\rvert}{\lvert\partial A\rvert}$
FUBear.seg	344	241,757	7,414	868	227	33
BTor.seg	161	64,820	3,339	172	172	19
sightsT.seg	294	157,940	7,041	995	314	22
sightsR.seg	294	158,368	7,048	1,019	314	22
euromanWalking.seg	67	48,933	2,280	515	106	21
westmanWalking.seg	43	58,928	2,251	509	86	26

Figure 7.4.: Relevant parameters of some of the input shapes. The values are rounded to integers.

The algorithms. For matching solid shapes under translations, the software implements Algorithm 1 in dimension 2. We refer to the implemented algorithm for translations as Algorithm 'translations'. Recall that Algorithm 1 obtains as input two shapes, a clustering size, and a sample size.

For matching solid shapes under rigid motions, the software implements versions of Algorithms 7 and 8 that obtain as input two shapes, a clustering size for the translational part, a clustering size for the rotation angle and a sample size. In the following, we refer to the implemented versions of Algorithm 7 and 8 as Algorithms 'random angle' and Algorithm '3+1 points'. Unless otherwise stated, Algorithm '3+1 points' is used for matching under rigid motions.

The program also provides probabilistic matching of solid shapes under similarities for test purposes. Different parameterization of the space of similarities can be chosen. Furthermore the software is able to perform probabilistic matching of polylines. Figure 7.5 shows a screenshot of the main window of the program's graphical user interface.

The program offers a discretized clustering algorithm, additionally to implementations of the Functions FindDensestClusterT (in dimension 2) and 2DFindDensestClusterRM. In this section, we call the clustering algorithms as described in Functions FindDensestClusterT and 2DFindDensestClusterRM the *exact clustering*.

Recall that neighborhoods (or clusters) are rectangles. The *discretized clustering algorithm* rounds every sample to a grid whose widths are given by the translational and rotational clustering sizes. The clustering size in the program corresponds to the side length of the rectangle and not half of the side length as in the analysis in previous chapters. In the descriptions, we stay with our notation and refer to the

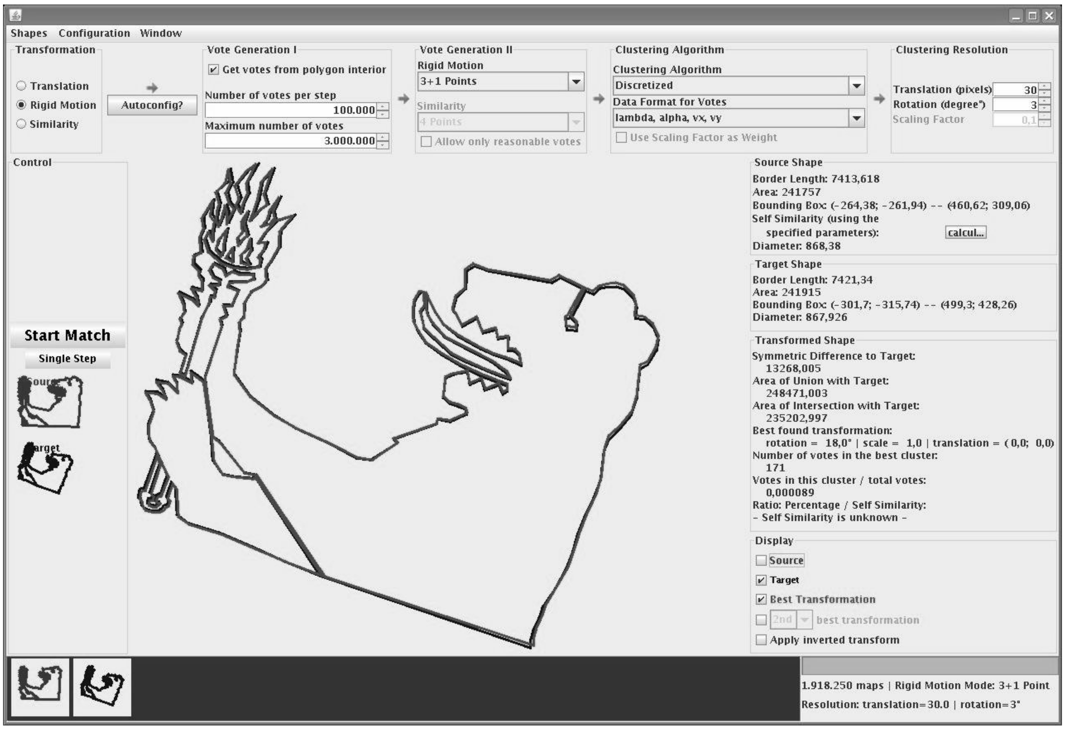

Figure 7.5.: The GUI of the test software. The transformation class, the type of algorithm and two clustering methods can be chosen in the upper part of the window. Also, the clustering sizes and the number of random samples can be set. On the right side, a number of parameters, as the area of the shapes and the area of the symmetric difference of the result are displayed. On the left side, the source and target shape are shown as small pictures. In the middle, the matching result is displayed. Here the target shape is shown in black, and the transformed source shape is shown in gray. In the screenshot matching under rigid motions of FUBear.seg and a rotated copy, FUBearRot.seg, is performed.

clustering size as half of the side length of the rectangle. The discretized clustering algorithm, which is often called bucket clustering, is much faster than the exact clustering algorithm. The non-empty grid cells are stored in a hash table, together with the number of points in the grid cell.

But the discretized clustering introduces an additional error. The quality of the result depends on the distance of the optimal transformation to the closest grid point. In general, the result cannot be arbitrarily close to the optimum. Furthermore an optimal translation might lie on the boundary of four grid cells, and an optimal rigid motion might even lie on the boundary of eight grid cells. In this case, the best cluster is cut and distributed on several grid cells, which can corrupt the result. On the other hand, if the optimal transformation happens to be a grid point, a large clustering size and a small number of random samples might result reliably in the optimal transformation.

We use the exact clustering whenever it is feasible with respect to the runtime. The experimental results for translations are performed using the exact clustering algorithm. In the experiments for rigid motions, the clustering is done by the discretized algorithm.

The absolute error. Denote the input shapes by A and B. Let t^* be an output transformation and t^{opt} be the optimal transformation. In the theoretical analysis, we considered $\frac{|t^{opt}(A)\cap B|}{|A|} - \frac{|t^*(A)\cap B|}{|A|}$ to be the absolute error, where $|A|$ could be replaced by $|B|$. By Proposition 3.1, we have $|t^{opt}(A)\cap B| - |t^*(A)\cap B| \le \varepsilon|A|$ if and only if $|t^*(A)\triangle B| - |t^{opt}(A)\triangle B| \le 2\varepsilon|A|$. Hence the absolute error expressed in the area of the symmetric difference would be $\frac{|t^*(A)\triangle B|}{2|A|} - \frac{|t^{opt}(A)\triangle B|}{2|A|}$. However it seems to be more reasonable to consider $\frac{|t^*(A)\triangle B|}{|A|+|B|} - \frac{|t^{opt}(A)\triangle B|}{|A|+|B|}$ to be the absolute error, since for any transformation t, we have $\frac{|t(A)\triangle B|}{|A|+|B|} \in [0,1]$, and every transformation t such that $t(A)$ and B are disjoint yields $\frac{|t(A)\triangle B|}{|A|+|B|} = 1$. Hence in this section the *absolute error* is defined to be

$$\frac{|t^*(A)\triangle B|}{|A|+|B|} - \frac{|t^{opt}(A)\triangle B|}{|A|+|B|}.$$

Figure 7.6 shows example values of the absolute error for matches of FUBear.seg. The reader is invited to take these values as a reference to interpret the plots in Figures 7.9 to 7.14.

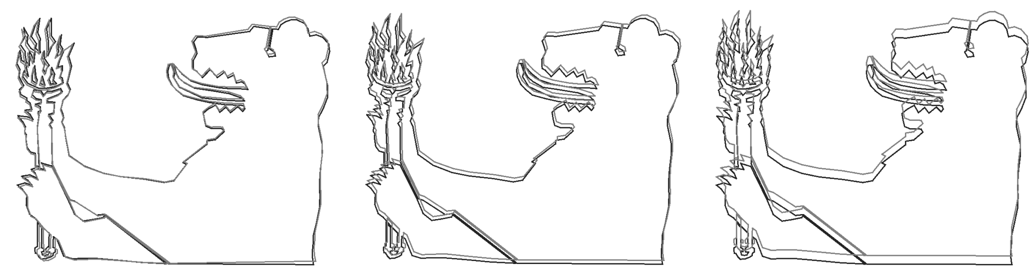

0.037 0.081 0.12

Figure 7.6.: Three different matches of FUBear.seg with the respective values of the absolute error below the pictures.

Matching results. Figures 7.7 and 7.8 show matching results for the shapes depicting traffic light symbols. The results were obtained with the described software.

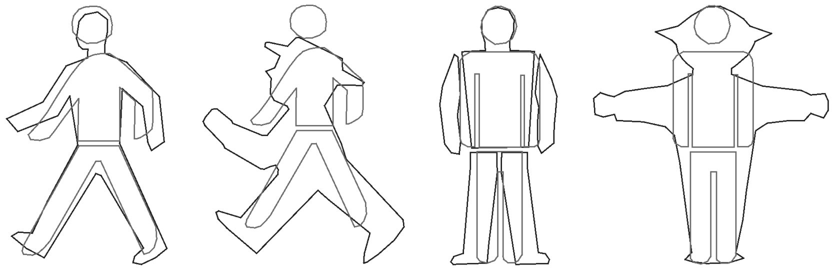

Figure 7.7.: Matching results for matching the standardized European to the former GDR and former FRG symbols under translations.

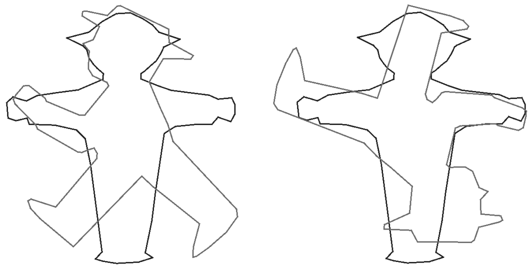

Figure 7.8.: Matching results for matching the former GDR symbols to one another under translations (left) and rigid motions (right).

Let us look at a few examples of how the error of the matching result behaves for a fixed clustering size, as the number of random samples grows. First, we look at matching FUBear.seg to itself under translations. Figure 7.9 shows the results of 50 runs of Algorithm 'translations' for four different clustering sizes. The exact clustering algorithm is used. For $k = 1, \ldots, 50$, the minimum, median and maximum absolute errors for $20,000k$ random samples are shown, which gives a total number of $1,000,000$ random samples. The median is defined to be the value at position $\lceil \frac{k}{2} \rceil$ in the increasingly ordered list of values. One can see that the convergence of the error is the slower the smaller the clustering size is, but that for a clustering size of 20 the results are slightly worse than for 5, 10, and 15. This behavior coincides with the results of the theoretical analysis (Theorem 5.1).

Figure 7.10 displays partial matching results. BTor.seg is matched to sightsT.seg under translations. The optimal match translates BTor.seg onto the congruent shape in sightsT.seg. The optimal area of the symmetric difference is the difference between the areas of the two shapes. The plot shows a maximum number of 200,000 votes. Note that the results are far better than for FUBear.seg. For a sample size of 50,000, the median error is about 0.05, while for FUBear.seg the median error is at least 0.1 for all tested clustering sizes. One reason for this is that for a partial matching problem the normalizing term of the absolute error, i.e. $|A| + |B|$, is larger than

$2|A|$, which is the normalizing term for matching a shape A to itself. Furthermore FUBear.seg has a more complex boundary structure than BTor.seg.

Figure 7.11 shows results of matching euromanWalking.seg to westmanWalking.seg under translations. Since we do not know the optimal area of the symmetric difference in this case, the y-axis shows $\frac{|t^*(A) \triangle B|}{|A|+|B|}$, which is the sum of the absolute error and the optimal area of the symmetric difference, normalized by $|A|+|B|$. The latter is called the *optimal distance* in the plots. It is believable that the minima are close to the optimum. Already for 8,000 votes, the median results are close to the minimum, which might be due to the fact that euromanWalking.seg and westmanWalking.seg are relatively simple shapes.

Next we turn to matching under rigid motions. Figure 7.12 depicts the results of 20 runs of Algorithm '3+1 points'. The clustering size for the translational part was set to 15, and the clustering size for the angle was set to 1.5 degree. The random experiments are performed in 10 steps of 200,000 votes, which gives a total number of 2,000,000 votes. As described above, the discretized clustering introduces an additional error. The results for Algorithm '3+1 points' are far better than for Algorithm 'random angle' for any given sample size.

Figure 7.13 shows results for matching euromanWalking.seg to westmanWalking.seg for Algorithm '3+1 points' and Algorithm 'random angle'. Again, Algorithm '3+1 points' outperforms Algorithm 'random angle' with respect to the quality of the matching results. Comparing the results to the results of matching euromanWalking.seg to westmanWalking.seg under translations, as depicted in Figure 7.11, shows the influence of the discretization error. The best found rigid motion has a larger absolute error than the best found translation.

Figure 7.14 depicts results for partial matching under rigid motions. BTor.seg is matched to sightsR.seg. Note that an absolute error of 0 should be possible, but the best match has an error of about 0.1. Probably this is due to the discretization error.

Summing up, the experiments show that Algorithm 'translations' and Algorithm '3+1 points' work well, whereas the convergence of Algorithm 'random angle' is quite slow. For the same sample size, the results of Algorithm '3+1 points' are far better than the results of Algorithm 'random angle'. The computation of one random rigid motion however is faster for Algorithm 'random angle' than for Algorithm '3+1 points'. Furthermore the experiments show that the speed of convergence of the error is very different for different shapes.

The empirical distribution on the transformation space. The program is able to display the distribution of votes on the transformation space, Figures 7.15 and 7.16 show examples for translations and rigid motions.

For the space of rigid motions, the translation space and the angle space are depicted seperately. The translation space shows the orthogonal projection of all votes onto the plane that is given by $z = 0$. The distribution of the angles is shown by a curve that represents the number of times an angle occured. Of course, for

Algorithm ‘random angle’ this curve is expected to be close to a constant function.

Comparing the first two pictures of Figure 7.16, one can see that the density function for Algorithm ‘3+1 points’, which is proportional to the square of the area of overlap, behaves better than the density function for Algorithm ‘random angle’, which is proportional to the area of overlap. This is the theoretical reason for the fact that the matching results of Algorithm ‘3+1 points’ are better than the results of Algorithm ‘random angle’ for any fixed clustering size and any fixed number of random samples.

Figure 7.9.: Results of matching shape FUBear.seg (Figure 7.3) under translations. The plots show the results of 50 runs of Algorithm 'translations' with different clustering sizes. A clustering size of 15 gives the best results.

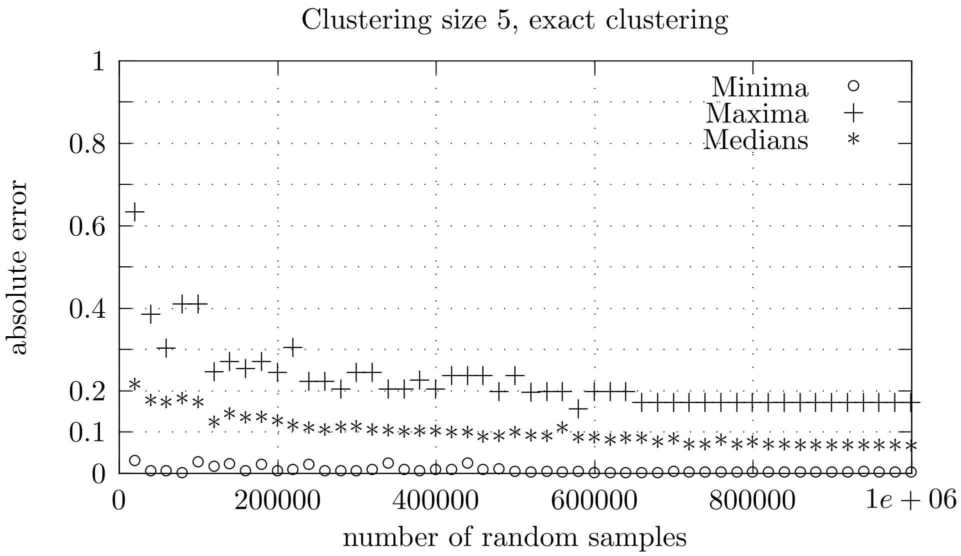

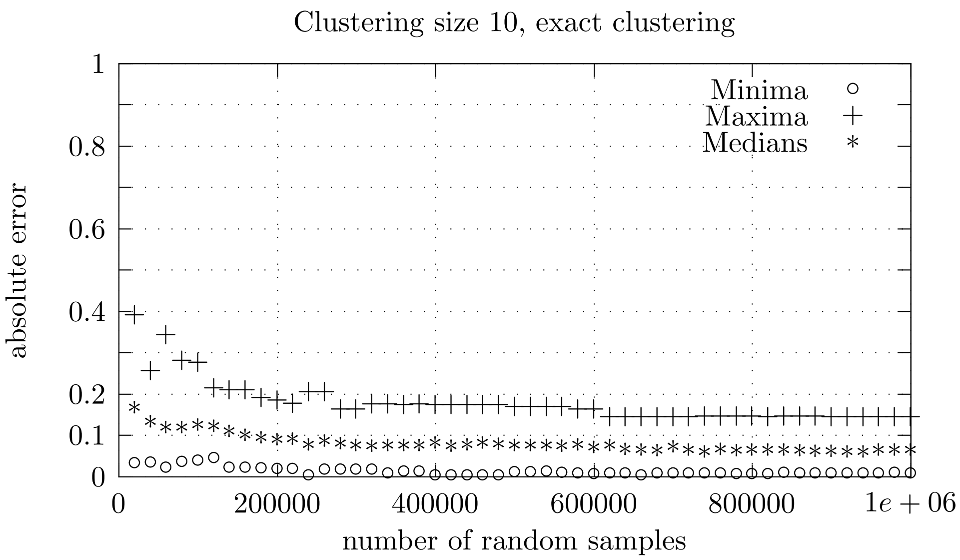

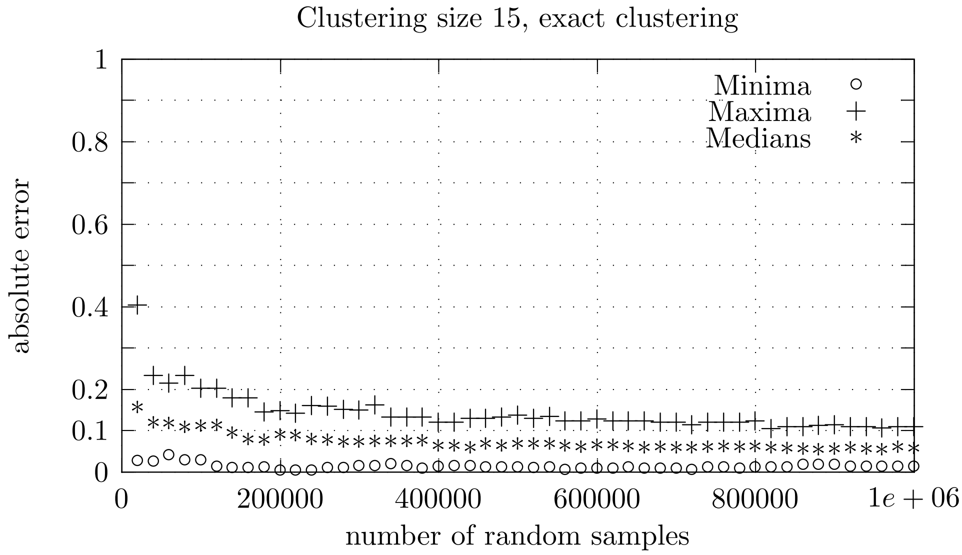
Clustering size 15, exact clustering
Minima
Maxima
Medians
absolute error
number of random samples
0
0.1
0.2
0.4
0.6
0.8
1
0
200000
400000
600000
800000
1e + 06

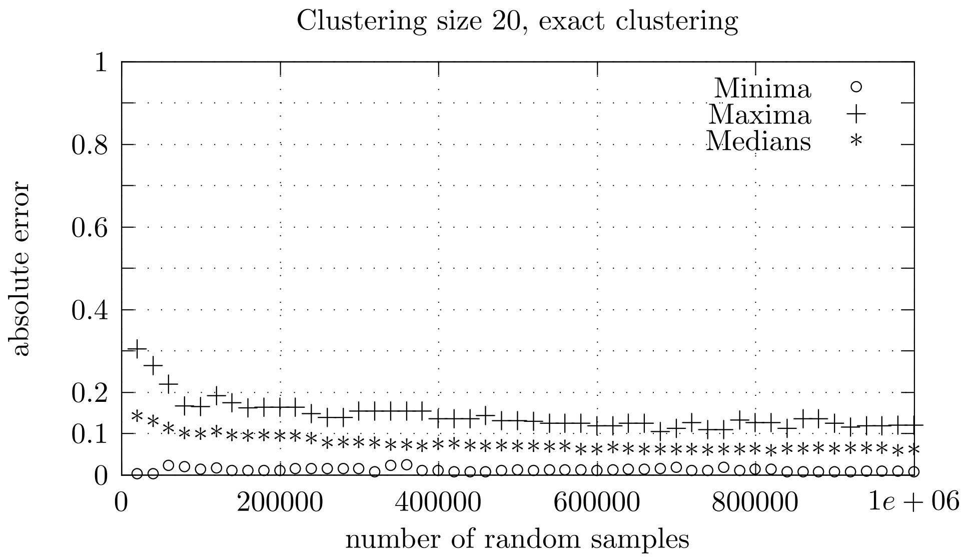
Clustering size 20, exact clustering
Minima
Maxima
Medians
absolute error
number of random samples
0
0.1
0.2
0.4
0.6
0.8
1
0
200000
400000
600000
800000
1e + 06

Figure 7.10.: Results of 50 runs of Algorithm 'translations' on the input of shapes BTor.seg (Figure 7.3) and sightsT.seg.

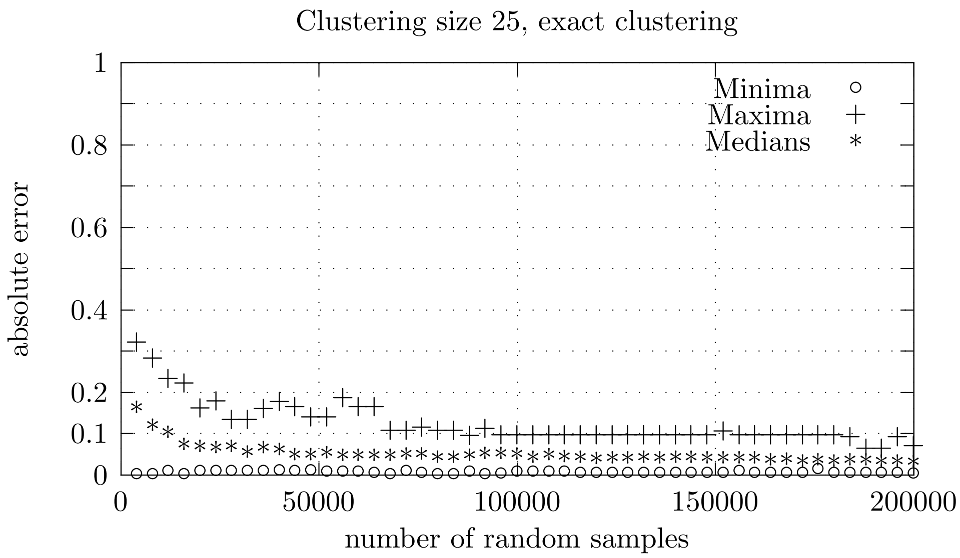

Figure 7.11.: Results of 50 runs of matching euromanWalking.seg (Figure 7.3) to westmanWalking.seg with Algorithm 'translations'.

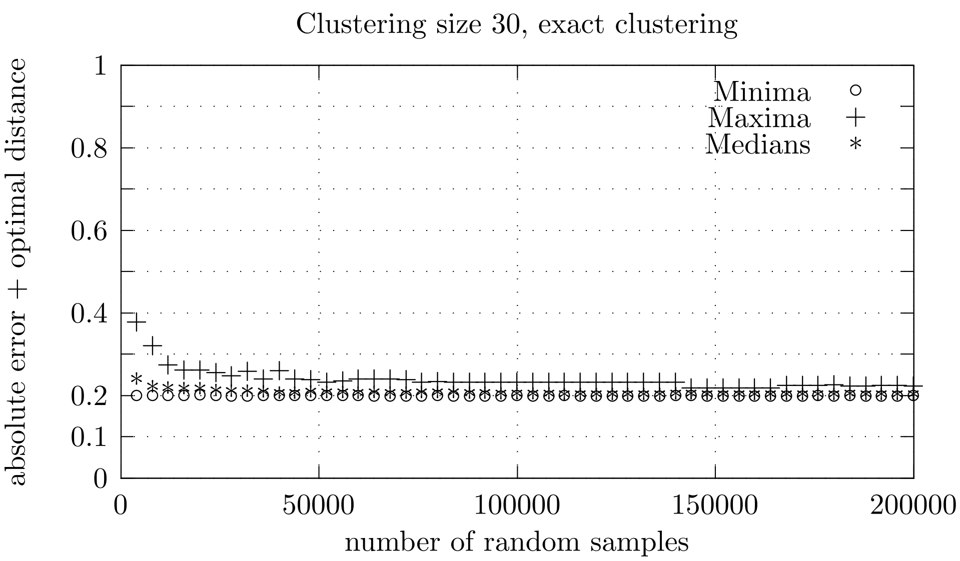

Figure 7.12.: Results of matching under rigid motions of FUBear.seg (Figure 7.3). The plots show the results of 20 runs of Algorithm '3+1 points' and 'random angle' with a clustering size of 15 for the translational part and 1.5 degree for the angle. The discretized clustering algorithm is used.

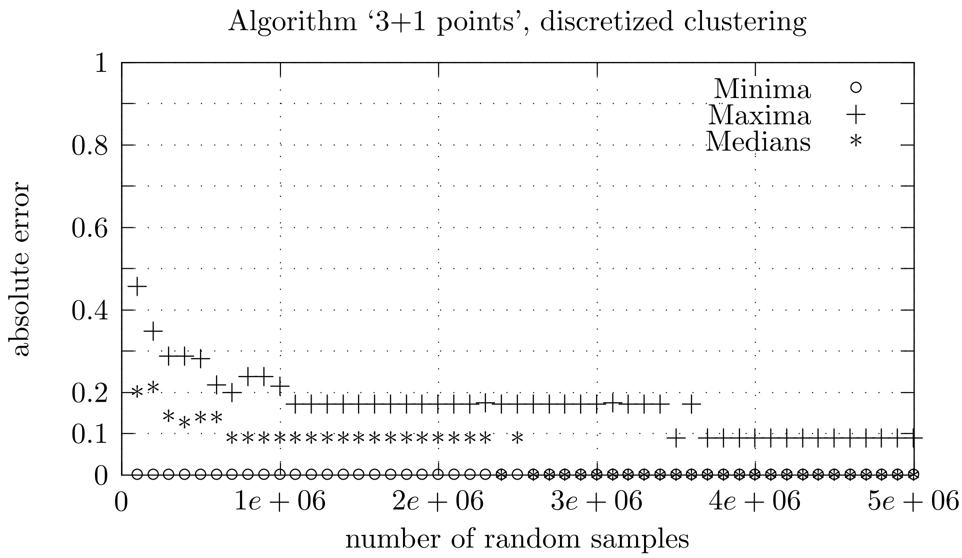

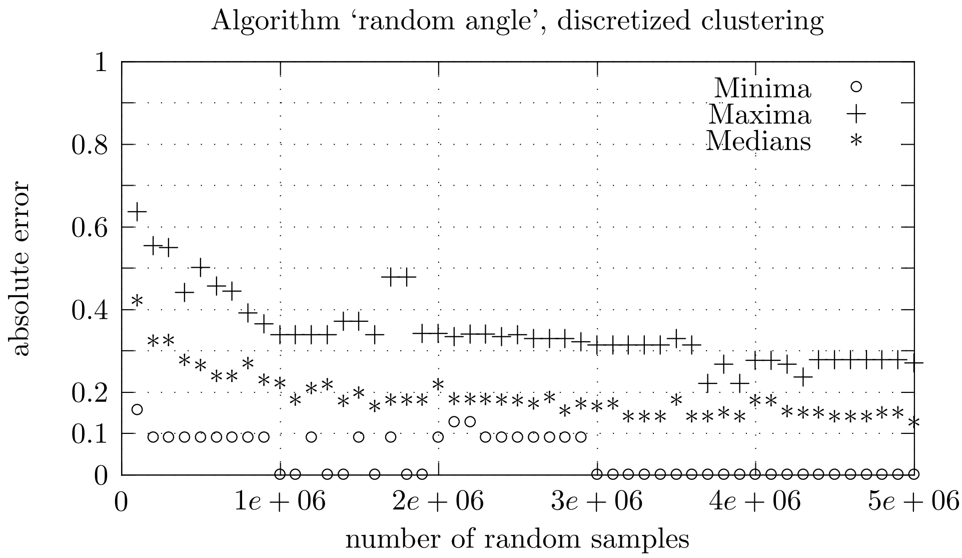

Figure 7.13.: Results of matching under rigid motions of euromanWalking.seg to westmanWalking.seg (Figure 7.3). Both plots show the results of 20 runs of the respective algorithm. The clustering size for the translational part equals 20 and the clustering size for the angle equals 1.5 degrees. Note that the area of the symmetric difference cannot vanish since the shapes are different.

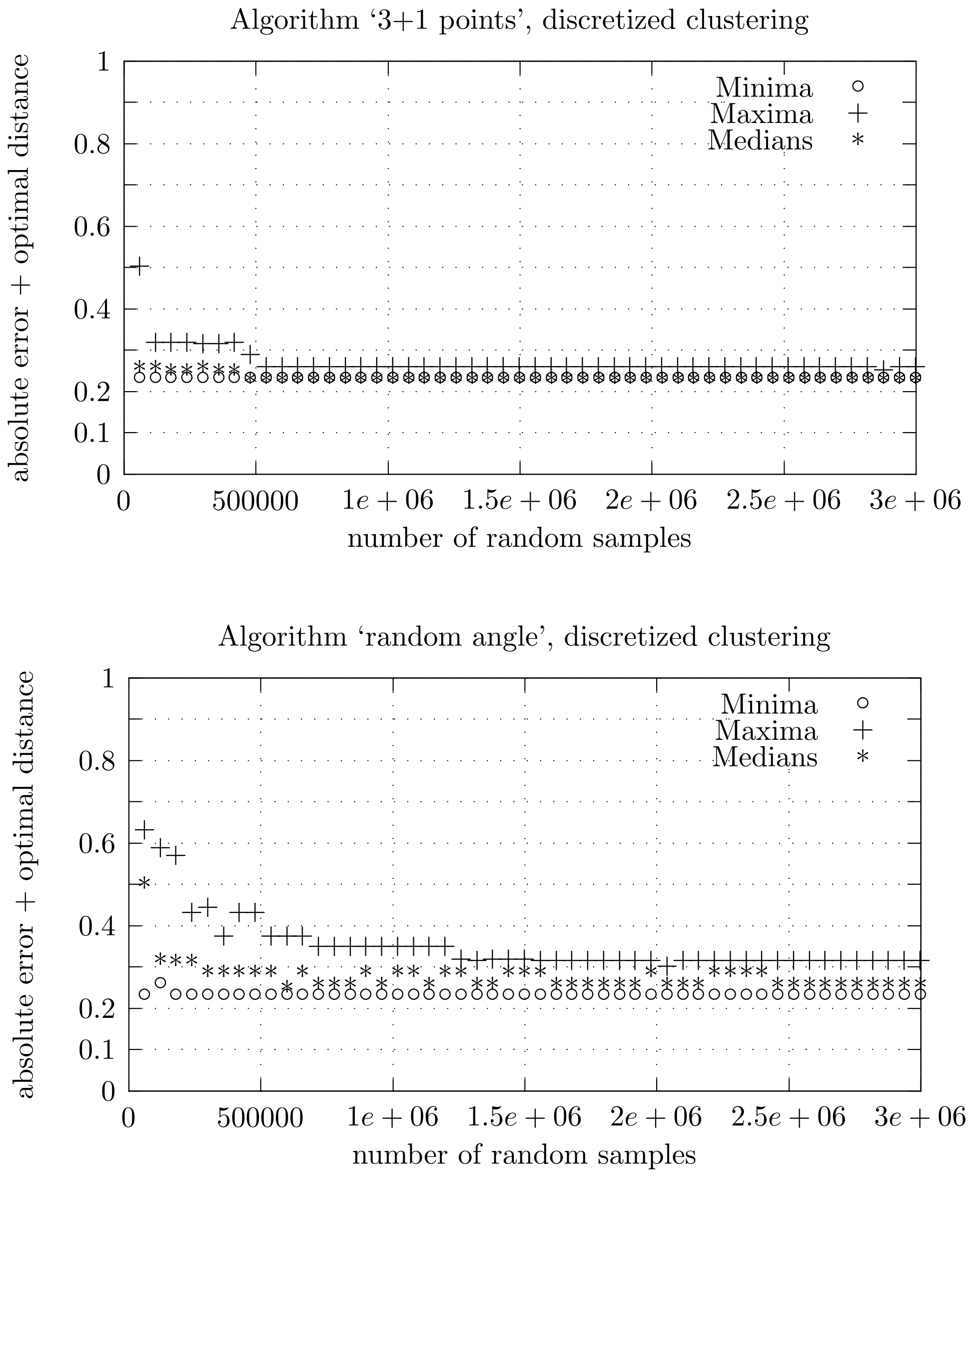

Figure 7.14.: Results of matching BTor.seg to sightsR.seg under rigid motions. The plots show the results of 20 runs of Algorithm '3+1 points' and Algorithm 'random angle'. The translational clustering size is set to 15 and the rotational clustering size is set to 1.5 degrees.

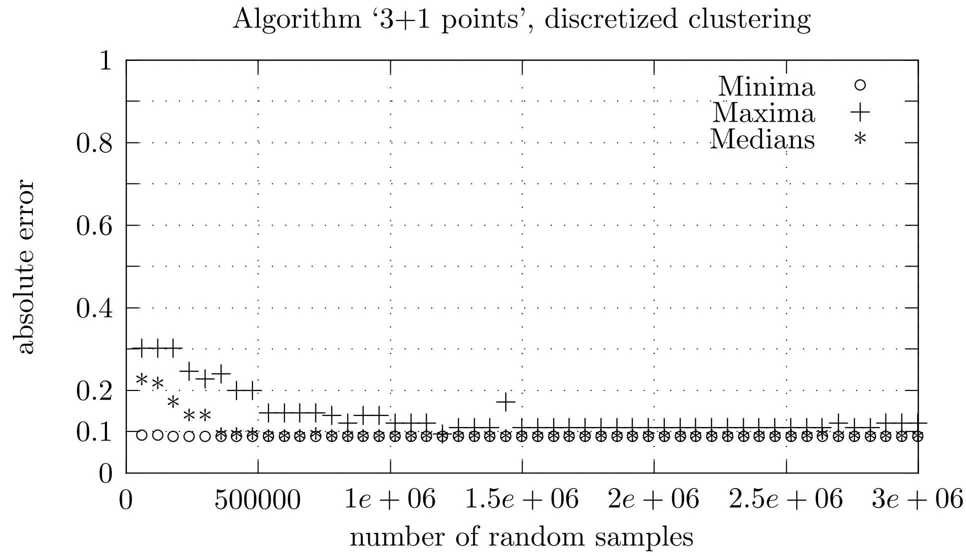

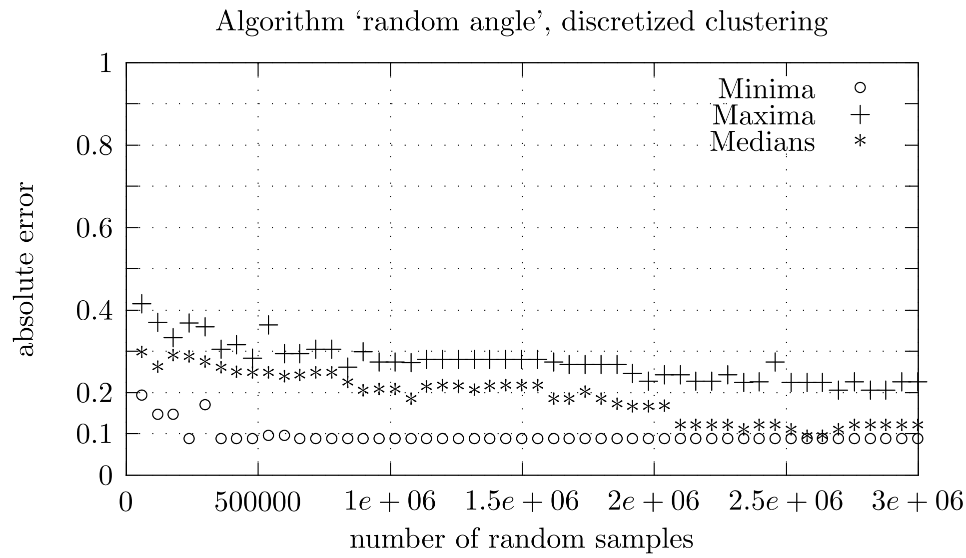

Figure 7.15.: The empirical distribution on the translation space for matching under translations. In all pictures, the number of votes is 3,000,000.

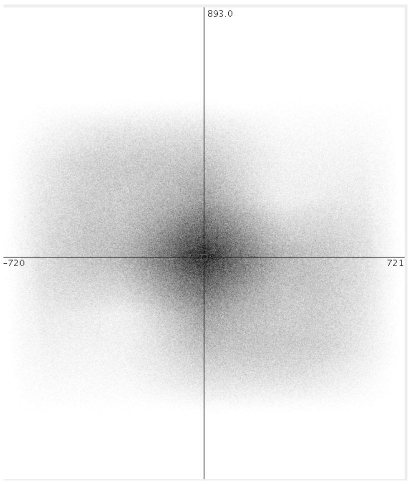

FUBear.seg matched to itself

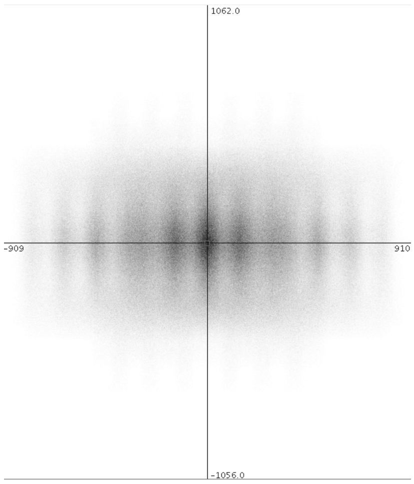

BTor.seg matched to itself

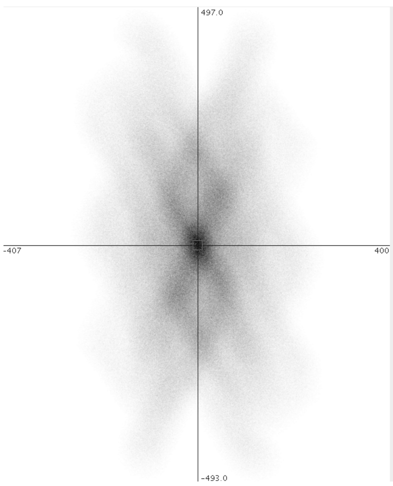

euromanWalking.seg matched to itself

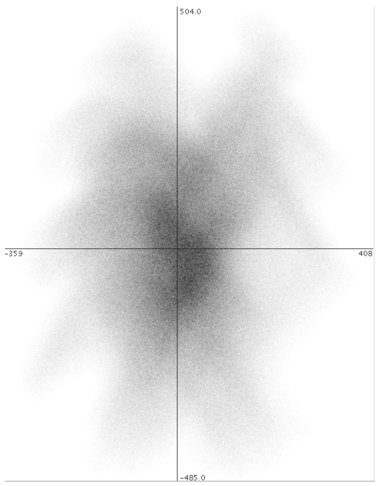

euromanWalking.seg matched to eastmanWalking.seg

Figure 7.16.: The empirical distribution on the space of rigid motions. The translation space is depicted on the left and the angle space on the right. The algorithms generated 3,000,000 random samples.

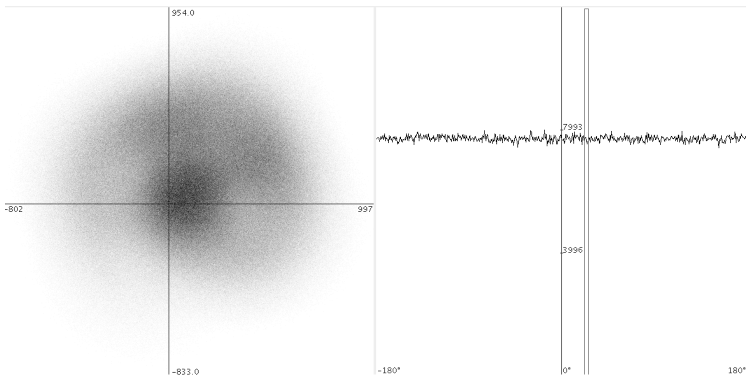

FUBear.seg matched to itself by Algorithm ‘random angle’

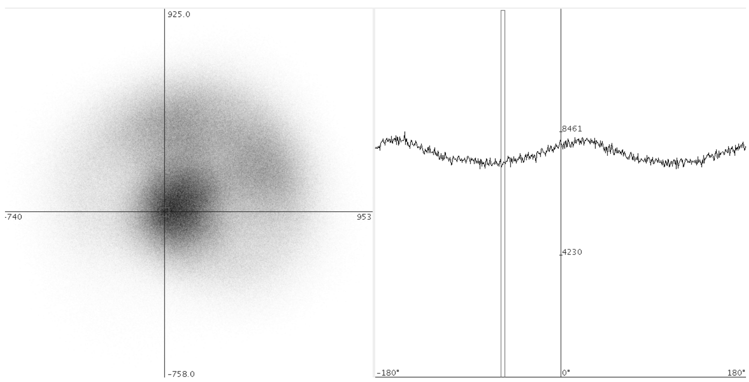

FUBear.seg matched to itself by Algorithm ‘3+1 points’

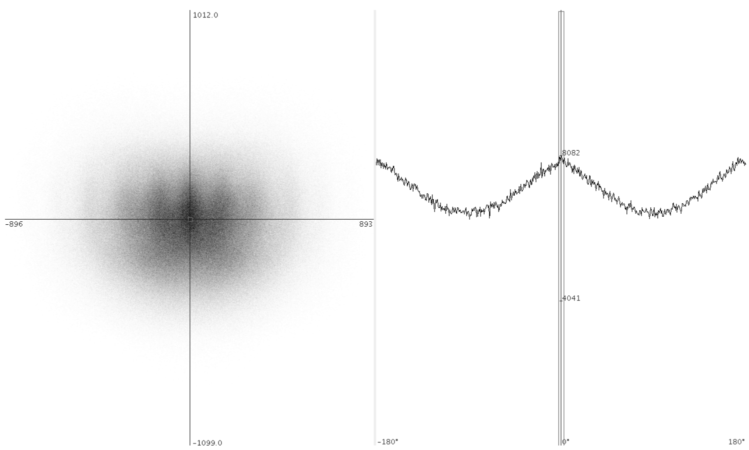

BTor.seg matched to itself by Algorithm ‘3+1 points’

8. Probabilistic matching under rigid motions in arbitrary dimension

We study the probabilistic matching of d-dimensional shapes under rigid motions. Given two shapes A and B, an error tolerance $\varepsilon \in (0,1)$ and an allowed probability of failure $p \in (0,1)$, we compute a rigid motion r^* such that with probability $\geq 1-p$ for all rigid motions r, we have $|r^*(A) \cap B| \geq |r(A) \cap B| - \varepsilon|A|$.

In $d \geq 3$ dimensions, Vigneron [54] gives the only result for the matching of non-convex polyhedra under rigid motions so far. He describes FPTASs for maximizing the volume of overlap, as well as for minimizing the volume of the symmetric difference. 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 $O((\frac{nm}{\varepsilon})^{\frac{d^2}{2}+\frac{d}{2}+1}(\log \frac{nm}{\varepsilon})^{\frac{d^2}{2}+\frac{d}{2}+1})$. The time bound for approximating the minimal volume of the symmetric difference is not given in the paper for dimension $d \geq 3$.

As shapes, we consider shapes that are given by an oracle and fulfill the assumptions described in Section 2.2. We present two approaches. First we use the matching algorithm for translations that we studied in Chapter 5 to develop an algorithm for matching under rigid motions. Second we generalize Algorithm 7, which we presented in Chapter 6 from two to arbitrary dimension.

Our algorithms are again very simple and could easily be implemented, but the time complexity increases quickly with the dimension. We begin with a section about rotations and a section about rigid motions. Among other things, we have to understand how to measure distances and volume in the space of rigid motions because balls of a fixed radius δ and their volume will again play a central role in the analysis.

8.1. The rotation group

A *rotation* is a linear map $r : \mathbb{R}^d \to \mathbb{R}^d$, $r(x) = Mx$ that preserves distances and orientations; the matrix $M \in \mathbb{R}^{d \times d}$ is called a *rotation matrix*. A matrix M is a rotation matrix if and only if its determinant equals 1 and it is *orthogonal*, the latter meaning that it is non-singular and $M^{-1} = M^T$. The set of real rotation matrices, together with the matrix multiplication, which corresponds to the composition of the maps, forms the well-known *special orthogonal group* over the reals $SO(d)$.

Measuring distances between rotations. To measure distances in $SO(d)$, we use the metric that is induced by the Frobenius norm $||M||_2 = \sqrt{\sum_{1\leq i,j\leq d} m_{ij}^2}$ where $M = (m_{ij})_{1\leq i,j\leq d}$. We write $d(M, N)$ for the distance $||M - N||_2$. For $M \in SO(d)$ and $\delta > 0$, denote the closed ball $B(M, \delta) = \{N \in SO(d) : d(M, N) \leq \delta\}$.

A pair of vector norm $|| \cdot ||_v$ and matrix norm $|| \cdot ||_m$ is called *compatible* if for all vectors $x \in \mathbb{R}^d$ and matrices $M \in R^{d\times d}$, submultiplicativity holds: $\|Mx\|_v \leq \|M\|_m \|x\|_v$. The Euclidean norm $| \cdot |$ is compatible with the Frobenius norm $|| \cdot ||_2$.

It is convenient that the metric d is invariant under the group operation.

Proposition 8.1. *For all $M, N \in \mathbb{R}^{d\times d}$ and $R \in SO(d)$, we have $d(RM, RN) = d(M, N)$.*

Proof. For a matrix M, let M_i be the ith column vector. We write $M = (M_1 \dots M_d)$. Then $RM = (RM_1 \dots RM_d)$ by the definition of matrix multiplication. For $R \in SO(d)$ and $x \in \mathbb{R}^d$, we have $|Rx| = |x|$. Therefore

$$\begin{aligned} ||M - N||_2^2 &= \sum_{j=1}^{d}\sum_{i=1}^{d}(m_{ij} - n_{ij})^2 = \sum_{j=1}^{d} |M_j - N_j|^2 \\ &= \sum_{j=1}^{d} |RM_j - RN_j|^2 = ||RM - RN||_2^2. \end{aligned}$$

□

Proposition 8.1 implies the following.

Corollary 8.2. *For rotation matrices $M, N \in SO(d)$ and $\delta > 0$, we have $B(M, \delta) = MN^T B(N, \delta)$.*

Proof. If $R \in B(N, \delta)$, then $d(N, R) = d(M, MN^T R) \leq \delta$ and therefore $MN^T R \in B(M, \delta)$.

If $R \in B(M, \delta)$, then $d(M, R) = d(N, NM^T R) \leq \delta$. Therefore $NM^T R \in B(N, \delta)$ and $R \in (MN^T)B(N, \delta)$. □

Measuring the volume in $SO(d)$. In the algorithm, we draw random rotation matrices from the uniform distribution. This means that a drawn matrix lies in a set $E \subseteq SO(d)$ with probability proportional to the volume of E. So to define the uniform distribution on $SO(d)$ properly, we have to know how to measure the volume in $SO(d)$.

The elements of $SO(d)$ are $(d \times d)$-matrices, so $SO(d)$ is a subset of $\mathbb{R}^{d^2}$. But $SO(d)$ is not a full-dimensional subset of $\mathbb{R}^{d^2}$. When determining a rotation matrix, there are less than d^2 degrees of freedom. For choosing the first column vector, there are $d - 1$ degrees of freedom since any point from S^{d-1} can be chosen. The second column vector has also length 1, but it has to be orthogonal to the first one. So it can be chosen from a $(d - 2)$-dimensional sphere that is orthogonal to the first column vector, as described above. Proceeding in this manner, it becomes clear that the number of degrees of freedom equals $(d-1)+(d-2)+\cdots+1=\binom{d}{2}$. In fact, $SO(d)$ is a $\binom{d}{2}$-dimensional manifold in $\mathbb{R}^{d^2}$. We measure the volume $|\cdot|$ in $SO(d)$ by the $\binom{d}{2}$-dimensional Hausdorff measure $\mathcal{H}^{\binom{d}{2}}$, as defined in Section 2.1.

The algorithm is based on the fact that, for a fixed $\delta > 0$, all δ-balls $B(M,\delta)$ have the same volume. Corollary 8.2 implies that the set of all δ-balls $\{B(M,\delta) : M \in SO(d)\}$ equals $\{MB(I,\delta) : M \in SO(d)\}$ where I is the identity matrix. Thus the fact that all δ-balls have the same volume boils down to $|B(I,\delta)| = |MB(I,\delta)|$ for all $M \in SO(d)$.

A natural requirement on a measure μ over a group G, which is slightly more general, is that for all $g \in G$ and $A \subset G$ we have $\mu(A) = \mu(gA)$. Such a measure is called a *Haar measure*. It turns out that there is up to a scaling factor a unique Haar measure on $SO(d)$ if we require the measure to be sufficiently "nice". In this context, "nice" means that we require the measure to be a Radon measure. These notions are defined in Appendix A. It follows from the definition of the Hausdorff measure that it is a Haar measure on $SO(d)$. Since $SO(d)$ is a $\binom{d}{2}$-dimensional set, we measure $SO(d)$ by $\mathcal{H}^{\binom{d}{2}}$ and we have $0 < \mathcal{H}^{\binom{d}{2}}(SO(d)) < +\infty$ [28].

A different approach is to observe that $SO(d)$ is a $\binom{d}{2}$-dimensional smooth manifold in $\mathbb{R}^{d^2}$ and to use the surface area that is given by integration over manifolds. Both approaches give, up to a constant, the same measure.

The analysis of the algorithm relies on the following theorem, which is proved in Appendix A. Recall that ω_k is the volume of the Euclidean unit ball in $\mathbb{R}^k$.

Theorem 8.3. *For all $M \in SO(d)$, we have*

$$\lim_{\delta\to+0} \frac{\mathcal{H}^{\binom{d}{2}}\big(B(M,\delta)\big)}{\delta^{\binom{d}{2}}} = \omega_{\binom{d}{2}}.$$

Computing rotation matrices uniformly at random. Within the algorithms, random rotation matrices have to be generated. There are many methods to generate uniformly distributed random orthogonal matrices described in the literature. The first algorithm [33], with a correction in [49], is based on the following observation. The first column vector v_1 from a random orthogonal matrix is uniformly distributed on $S^{d-1} \subset \mathbb{R}^d$. The second column vector is uniformly distributed on S^{d-1}, intersected with the hyperplane orthogonal to v_1, which is a $(d-2)$-dimensional sphere in $\mathbb{R}^d$. Proceeding inductively, the ith column vector v_i is picked in a $(d-i)$-dimensional sphere in $\mathbb{R}^d$ u.a.r., which is then intersected with $(v_i)^\perp$, the orthogonal space of v_i, resulting in a $(d-i-1)$-dimensional sphere. The last vector is chosen from S^0, which consists of two antipodal points. If we want to generate a matrix with positive determinant, we do not have any choice, but have to take the vector that gives determinant 1. Otherwise, if we allow all isometries, we choose one of the two vectors at random. A uniformly distributed random orthogonal matrix can be computed in $O(d^3)$ time [24]. A random orthogonal matrix can also be determined by computing $\binom{d}{2}$ random rotations of two axes, while all other axes are fixed.

Sampling $SO(d)$**.** In 2D, we described an algorithm for matching under rigid motions that uses the algorithm for matching under translations. Algorithm 6 runs the algorithm for translations for all shapes $M_{2\pi k/m}A$ and B where $k = 1 \dots m$ and outputs the rigid motion that got the overall largest number of votes. We will do something similar in arbitrary dimension. Instead of discretizing the angle space, as we did in 2D, we use a random "dense" set of angles. For this, we introduce the notion of ε-nets, which we quote from [40].

Let X be a set, let μ be a probability measure on X. Let $\mathcal{F}$ be a system of μ-measurable subsets of X, and let $\varepsilon \in [0, 1]$. A subset $N \subseteq X$ is called an *ε-net* for $(X, \mathcal{F})$ with respect to μ if $N \cap S \neq \emptyset$ for all $S \in \mathcal{F}$ with $\mu(S) \geq \varepsilon$.

Theorem 8.4 (Epsilon net theorem [40])**.** *If X is a set with probability measure μ, $\mathcal{F}$ is a system of μ-measurable subsets of X with VC dimension at most $d \geq 2$, and $r \geq 2$ is a parameter, then there exists a $\frac{1}{r}$-net for $(X, \mathcal{F})$ with respect to μ of size at most $Cdr \ln r$, where C is an absolute constant.*

The proof of this theorem shows that any random sample of X of size $Cdr \ln r$ where each sample is drawn independently according to μ is an $\frac{1}{r}$-net with positive probability, even with probability $\geq \frac{1}{2}$. A sample of size $kCdr \ln r$ is then a $\frac{1}{r}$-net of X with probability $\geq 1 - (\frac{1}{2})^k$. For $k \geq \log_2 \frac{2}{p}$, such a sample is an $\frac{1}{r}$-net of X with probability $\geq 1 - \frac{p}{2}$.

For shapes A, B and the random sample S of $SO(d)$, we want to have the property

$$\forall M \in SO(d)\ \exists M' \in S\ \forall t \in \mathbb{R}^d : \quad \big| |M(A+t) \cap B| - |M'(A+t) \cap B| \big| \leq \tfrac{\varepsilon}{2}|A| \quad (8.1)$$

Because of $\sup_{a \in A} |Ma - Ma'| \leq \operatorname{diam}(A) d(M, M')$, as we will see in Proposition 8.9, and Corollary 3.5, this is fulfilled if for all $M \in SO(d)$ there exists $M' \in S$ such that $d(M, M') \leq \frac{\varepsilon|A|}{\operatorname{diam}(A)|\partial A|}$.

Recall that $B(I, \delta)$ denotes the closed ball of radius δ with center I in $SO(d)$ with respect to the metric d. Define

$$\nu = \frac{|B(I,\delta)|}{|SO(d)|} \quad \text{where} \quad \delta = \frac{\varepsilon|A|}{\operatorname{diam}(A)|\partial A|}. \tag{8.2}$$

To see $|A| \leq \operatorname{diam}(A)|\partial A|$, take $a, a' \in \operatorname{cl}(A)$ such that $|a - a'| = \operatorname{diam}(A)$. A is contained in the cylinder of height $\operatorname{diam}(A)$ with bottom $a + A|(a-a')^\perp$ and top $a' + A|(a-a')^\perp$. The $(d-1)$-dimensional volume of $a + A|(a-a')^\perp$ and $a' + A|(a-a')^\perp$ is at most $|\partial A|$. For details on this type of argument, see Chapter 3. Since $0 < \delta \leq 1$ and $\mathcal{H}^{\binom{d}{2}}\big(B(M,\delta)\big)\delta^{-\binom{d}{2}}$ is continuous in δ, we have $\nu \geq C\delta^{\binom{d}{2}}$ for some constant C, depending on d, by Theorem 8.3.

Define

$$\mathcal{C} = \{B(M, \delta) : M \in SO(d)\}. \tag{8.3}$$

Then each ν-net S for $(SO(d), \mathcal{C})$ with respect to the uniform distribution satisfies Property (8.1). By Lemma 8.11, the VC dimension of $\mathcal{C}$ is bounded by a constant. Therefore there is a constant C such that any random sample S of $SO(d)$ of size $C\frac{1}{\nu} \log \frac{1}{\nu} \log \frac{2}{p}$ fulfills Property (8.1) with probability $\geq 1 - \frac{p}{2}$.

Reducing the number of parameters to represent a rotation. We would prefer not to represent a rotation by the d^2 entries of a rotation matrix, but only by $\binom{d}{2}$ parameters. A rotation matrix can be identified with $\binom{d}{2}$ parameters by using the representation of a rotation matrix as a product of rotations along the $\binom{d}{2}$ planes, each rotation keeping $d-2$ axes fixed.

Let $G_{ij} = G_{ij}(\alpha_{ij})$ be the following simple rotation matrix.

$$\begin{pmatrix} I & 0 & 0 & 0 & 0 \\ 0 & \cos\alpha_{ij} & 0 & -\sin\alpha_{ij} & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & \sin\alpha_{ij} & 0 & \cos\alpha_{ij} & 0 \\ 0 & 0 & 0 & 0 & I \end{pmatrix}$$

In this matrix, each I denotes an identity matrix and each 0 a zero matrix of the appropriate sizes. The four entries that define the rotation along a plane have the indices ii, ij, ji, and jj. The matrix G_{ij} is called the *Givens matrix* in [10]. We call the vector of angles $(\alpha_{12}, \ldots, \alpha_{1d}, \alpha_{23}, \ldots, \alpha_{2d}, \ldots, \alpha_{d-1,d})$ the *Givens angles* of the rotation.

Theorem 8.5. *[10] Each orthogonal matrix $R \in O(d)$ has a representation as a product $R = (G_{12}G_{13} \cdots G_{1d})(G_{23} \cdots G_{2d}) \cdots (G_{d-1,d})D$ of Givens matrices G_{ij} with angles $\alpha_{ij} \in (-\frac{\pi}{2}, \frac{\pi}{2}]$ for all $i, j \in \{1, \ldots, d\}$ such that $i < j$, and a diagonal matrix D with entries ± 1 on the diagonal.*

A random orthogonal matrix from the uniform distribution can be generated by picking the entries on the diagonal of D independently as 1 or -1 with probability $\frac{1}{2}$ and letting the $\binom{d}{2}$ angles α_{ij} be mutually independent with joint density function proportional to

$$\left(\prod_{j=2}^{d} \cos^{j-2} \alpha_{1j}\right)\left(\prod_{j=3}^{d} \cos^{j-3} \alpha_{2j}\right) \cdots \left(\prod_{j=d}^{d} \cos^{j-d} \alpha_{d-1,j}\right).$$

We use this result to generate random rotation matrices in 3D. Rotation matrices are orthogonal matrices that have a positive determinant. Hence we have the following corollary.

Corollary 8.6. *Each rotation matrix $R \in SO(3)$ has a representation as a product $R = G_1G_2G_3D$, where G_1, G_2, G_3 are the Givens matrices correspoding to some angles $\alpha, \beta, \gamma \in (-\frac{\pi}{2}, \frac{\pi}{2}]$ and D is one of the four 3×3 diagonal matrices with entries ± 1 on the diagonal and an even number of negative entries.*

A random 3×3 rotation matrix from the uniform distribution can be generated by picking one of the four possible matrices D with probability $\frac{1}{4}$ and by picking the Givens angles α, β, γ mutually independent with density functions $f_1(\alpha) \equiv \frac{1}{\pi}$, $f_2(\beta) \equiv \frac{1}{\pi}$ and $f_3(\gamma) = \frac{1}{2}\cos\gamma$.

8.2. The space of rigid motions

A *rigid motion* is an orientation preserving isometry. For each rigid motion r, there is a rotation matrix M and a translation vector t such that $r(x) = Mx + t$ for all $x \in \mathbb{R}^d$. We identify a rigid motion with the pair (M, t) of its rotation matrix and its translation vector. Thus the transformation space equals $SO(d) \times \mathbb{R}^d$.

Recall that we are given a parameter Δ_A that fulfills $\mathrm{diam}(A) \leq \Delta_A$ if A satisfies the assumptions described in Section 2.2. We define a metric on the space of rigid motions, which depends on the input shape A, by

$$d((M,p),(N,q)) = \max\{\ \Delta_A \cdot d(M,N),\ |p-q|\ \}.$$

Denote by $B(r, \delta)$ the closed ball of radius δ with respect to d, centered at r, in the space of rigid motions. From these definitions, the following proposition is obvious.

Proposition 8.7. *For all rigid motions (M, t) and all $\delta > 0$ holds that*

$$B((M,t),\delta) = B(M, \delta/\Delta_A) \times B_2(t,\delta)$$

where $B_2(t, \delta)$ is the closed Euclidean ball of radius δ, centered at t, in $\mathbb{R}^d$.

The set of rigid motions forms a locally compact topological group. The group operation is the composition of maps. The volume on the group of rigid motions is the Haar measure. It equals the product measure of $\mathcal{H}^{\binom{d}{2}}$ restricted to $SO(d)$ and the d-dimensional Lebesgue measure on $\mathbb{R}^d$. Therefore for all rigid motions (M, t) and all $\delta > 0$ holds that $|B((M, t), \delta)| = |B(M, \delta/\Delta_A)| \cdot |B_2(t, \delta)|$. Because of Theorem 8.3 the following theorem is true.

Theorem 8.8. *For all rigid motions r and all $\delta > 0$ holds*

$$\lim_{\delta \to +0} |B(r,\delta)|\, \Delta_A^{\binom{d}{2}} \delta^{-\binom{d+1}{2}} = \omega_{\binom{d}{2}}\, \omega_d.$$

If the distance of two translation vectors t_1 and t_2 is δ with respect to some norm, then, for each x, the translated points $x + t_1$ and $x + t_2$ have distance δ, too. If two rotation matrices M_1, M_2 have distance δ with respect to the matrix norm $\|\cdot\|_M$, then we only know that $\|M_1x - M_2x\|_V \leq \|M_1 - M_2\|_M \|x\|_V = \delta\|x\|_V$ if $\|\cdot\|_V$ is a to $\|\cdot\|_M$ compatible vector norm. The following proposition gives an upper bound on how much the images of two rigid motions differ in terms of the distance of the rigid motions. We assume w.l.o.g. that $0 \in A$.

Proposition 8.9. *Let $r = (M, p)$ and $s = (N, q)$ be rigid motions. Then, for all $x \in A$,*

$$|r(x) - s(x)| \leq 2d(r,s) \qquad \textit{and} \qquad |Mx - Nx| \leq \mathrm{diam}(A)d(M,N).$$

Proof. $|r(x) - s(x)| \leq ||M - N||_2|x| + |p - q| \leq \mathrm{diam}(A)||M - N||_2 + |p - q| = \mathrm{diam}(A)d(M, N) + d(p, q) \leq 2d(r, s)$. □

Note that the factor diam(A) in the proposition cannot be omitted. Let $r = (I, 0)$ be the identity, let e_{11} be the matrix that has a 1-entry in the upper left corner and 0-entries otherwise. Let $S = (I + \delta e_{11}, 0)$. Then $d(r, s) = \delta$ and for $x = \lambda e_1$ holds $|r(x) - s(x)| = |\lambda \delta e_1| = |\lambda| \delta = |x| d(r, s)$. From this proposition and Corollary 3.5 follows:

Corollary 8.10. *For all rigid motions $r, s \in SO(d) \times \mathbb{R}^d$*

$$\big| |r(A) \cap B| - |s(A) \cap B| \big| \leq |\partial A| \; d(r, s).$$

Shatter coefficients and the VC dimension were introduced in Chapter 4.

Lemma 8.11. *The class $\mathcal{C} = \{B(r, \delta) : \delta > 0, r \in SO(d) \times \mathbb{R}^d\}$ has VC dimension $\leq 3(d^2 + d + 2) \log(3(d^2 + d + 2))$.*

Proof. If $\mathcal{C}_1$ and $\mathcal{C}_2$ are set systems with $\mathcal{C}_1 \subset \mathcal{C}_2$, then $\dim(\mathcal{C}_1) \leq \dim(\mathcal{C}_2)$. Let $\mathcal{T} = \{B_2(t, \delta) : \delta > 0, t \in \mathbb{R}^d\}$. The family of all closed Euclidean balls in $\mathbb{R}^d$ has VC dimension $d + 1$ [26].

Let $\mathcal{M}' = \{B_2(M, \delta/\Delta_A) : \delta > 0, M \in \mathbb{R}^{d^2}\}$. Thus, $\dim(\mathcal{M}') = d^2 + 1$. Let $\mathcal{M} = \{B(M, \delta/\Delta_A) : \delta > 0, M \in SO(d)\} = \{B_2(M, \delta/\Delta_A) \cap SO(d) : \delta > 0, M \in SO(d)\}$. It follows that $\dim(\mathcal{M}) \leq d^2 + 1$ since in comparison to $\mathcal{M}'$ the point sets that possibly can be shattered are restricted to $SO(d)$.

The set system $\mathcal{C} = \mathcal{M} \times \mathcal{T}$. Let $x_1, \ldots, x_k$ be points in $\mathbb{R}^{d^2+d}$ that can be shattered by $\mathcal{M} \times \mathcal{T}$. For all $1 \leq i \leq k$, let $x_i = (y_i, z_i)$ where y_i are the first d^2 coordinates of x_i and z_i are the last d coordinates. For all $I \subseteq \{1, \ldots, k\}$, there exists $(M, T) \in \mathcal{M} \times \mathcal{T}$ such that $I = \{i : x_i \in (M, T)\} = \{i : y_i \in M\} \cap \{i : z_i \in T\}$.

This implies $|\{(J, L) : \exists M \in \mathcal{M} : J = \{i : y_i \in M\}, \exists T \in \mathcal{T} : L = \{i : z_i \in T\}\}| = \mathcal{S}_{\mathcal{M}}(k)\, \mathcal{S}_{\mathcal{T}}(k)$ is as least as large as 2^k. Using the fact that, for all set systems $\mathcal{N}$ with finite VC dimension, $\mathcal{S}_{\mathcal{N}}(k) \leq (k+1)^{\dim \mathcal{N}}$ (Corollary 4.1. in [25]), we get that $2^k \leq (k+1)^{d^2+d+2}$.

We prove that, if $k \geq 3(d^2 + d + 2) \log(3(d^2 + d + 2))$, then $(k + 1)^{d^2+d+2} < 2^k$, which shows that $\dim(\mathcal{C}) < 3(d^2 + d + 2) \log(3(d^2 + d + 2))$. Let $K = 3m \log(3m)$ for $m = d^2 + d + 2$. By the following computation we have $(K + 1)^{d^2+d+2} < 2^K$, which implies that, for all $k \geq K$, we have $(k + 1)^{d^2+d+2} < 2^k$.

$$\frac{2^{K/m-1}}{K} = \frac{2^{3 \log(3m)}}{6m \log(3m)} = \frac{(3m)^3}{6m \log(3m)} > \frac{9}{2} m \geq 1$$

Therefore, $2^{K/m} > 2K \geq (K + 1)$. This implies the claim. □

8.3. Using the algorithm for translations

Given d-dimensional shapes A and B, the idea of Algorithm 11 is to sample the set of rotations $SO(d)$ and to apply Algorithm 2 for matching under translations to the rotated shape MA and B for every matrix M in the sample. Instead of Algorithm 2,

Algorithm 11: MaxOverlapRMT

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

```
integer m ← NumberOfRotations(A, ε, p);
real V ← 0 ;                                   // current maximal area estimate
vector t* ← (0,0) ;                       // corresponding translation vector
matrix M* ← I ;                             // corresponding rotation matrix
for k = 1...m do
    matrix M ← RandomRotationMatrix();
    shape A' = M(A);                                // compute rotated A
    (vector,real) (t, v) ← MaxOverlapT(A', B, ε/4, p/(2m));
    if v > V then
        V ← v;                                       // update area estimate
        t* ← t;                                        // update translation
        M* ← M;                                           // update rotation
    end
end
return (M*, t*);
```

Function NumberOfRotations(A, ε, p)

Input: shape A, parameters $\varepsilon, p \in (0,1)$

Output: $m \geq C\big(\varepsilon^{-1}|A|^{-1}\operatorname{diam}(A)|\partial A|\big)^{\binom{d}{2}} \log(\varepsilon^{-1}|A|^{-1}\operatorname{diam}(A)|\partial A|) \log \frac{2}{p}$, where $m \in \mathbb{N}$ and C is some positive universal constant

we also could use Algorithm 5. For each run of Algorithm 2, we get a triple (M, t, V) where M is from the sample and (t, V) is the output of Algorithm 2. The result is the rigid motion (M, t) with the maximal volume of overlap estimate $V\,|A|\,|B|$.

Theorem 8.12 (Correctness of Algorithm 11)**.** *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$. Algorithm 11 with input (A, B, ε, p) computes a rigid motion r^* such that $|r^*(A) \cap B|$ is maximal up to an additive error of $\varepsilon|A|$ with a probability $\geq 1 - p$.*

Proof. Let ν and $\mathcal{C}$ be defined as in Equations (8.2) and (8.3). Recall that there exists a constant $C > 0$ such that, if

$$m \geq C\big(\varepsilon^{-1}|A|^{-1}\operatorname{diam}(A)|\partial A|\big)^{\binom{d}{2}} \log(\varepsilon^{-1}|A|^{-1}\operatorname{diam}(A)|\partial A|) \log \frac{2}{p},$$

then any random sample S picked by m independent draws from $SO(d)$, where each matrix is drawn according to the uniform distribution, fulfills Property (8.1) with probability $\geq 1 - \frac{p}{2}$.

Let (M^*, t^*) be the output of Algorithm 11, and let (M^{opt}, t^{opt}) be an optimal solution. The probability that one of the m calls of Algorithm 2 fails is $\leq \frac{p}{2}$. With probability at least $1-p$, the following conditions are satisfied: the set of random rotation matrizes S drawn in the algorithm is a ν-net and thus fulfills Property (8.1) and for all $M \in S$ the output (t, V) of Algorithm 2 satisfies

$$|(MA+t) \cap B| \geq |(MA+t^{opt}) \cap B| - \tfrac{\varepsilon}{4}|A| \tag{8.4}$$

and

$$\big||(MA+t) \cap B| - V|A||B|\big| \leq \tfrac{\varepsilon}{8}|A|. \tag{8.5}$$

Let us assume that this is the case. Let $M \in S$ such that $d(M^{opt}, M) \leq \frac{\varepsilon|A|}{2|\partial A| \operatorname{diam}(A)}$, and let (t, V) be the output that is computed by Algorithm 2, when started with input $(MA, B, \frac{\varepsilon}{4}, \frac{p}{2m})$. We show that $|(MA+t) \cap B|$ is close to both, $|(M^*A+t^*) \cap B|$ and $|(M^{opt}A + t^{opt}) \cap B|$.

First we have $\big||(MA+t) \cap B| - |(M^{opt}A + t^{opt}) \cap B)|\big| \leq \frac{3}{4}\varepsilon|A|$ by Inequality (8.4) and Property (8.1). If $|(M^*A+t^*) \cap B| \geq |(MA+t) \cap B|$, then we are done.

Assume that this is not the case. Let V^* be the volume estimate corresponding to t^*. Since (M^*, t^*) is the output of the algorithm, we have $V^* \geq V$. Together with Inequality (8.5), this implies $\big||(M^*A+t^*) \cap B| - |(MA+t) \cap B)|\big| \leq \frac{\varepsilon}{4}|A|$ and therefore $\big||(M^*A+t^*) \cap B| - |(M^{opt}A+t^{opt}) \cap B)|\big| \leq \varepsilon|A|$. □

Recall that a random rotation matrix can be generated in constant time. The runtime of Algorithm 11 is dominated by m times the runtime of Algorithm 2, plus the time needed to compute m, which proves the next theorem.

Theorem 8.13 (Runtime of Algorithm 11)**.** *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters.*

Then Algorithm 11 with input (A, B, ε, p) computes a rigid motion that maximizes the volume of overlap of A and B up to an additive error of $\varepsilon|A|$ with probability at least $1-p$ in time $O\Big(\varepsilon^{-\binom{d}{2}}(D_A K_A)^{\frac{d-1}{2}} \log(\varepsilon^{-1} D_A K_A) \log \frac{2}{p}\big(T(N) + N(\log N)^{d-1}\big)\Big)$ where $N = O\big(\varepsilon^{-(2d+2)}(K_A M_B/m_A)^2 \log \frac{2}{p}\big)$.

8.4. Sampling of the transformation space

Next we generalize Algorithm 7 for matching under rigid motions in the plane, in which we chose the rotation angles uniformly at random, from two to higher dimensions.

Given two d-dimensional shapes A and B, we sample $SO(d) \times A \times B$ uniformly at random. To each triple (M, a, b) of a rotation matrix, a point from A and a point from B, we associate the unique rigid motion that has rotation matrix M and maps a onto b, which is the map $(M, b - Ma)$. By this process, we sample the transformation space with a certain probability distribution μ.

We prove in Lemma 8.15 that the density function of μ is proportional to $f(r) = |r(A) \cap B|$. Our goal is to find a rigid motion r at which the density function is close to maximal with high probability. We find such a transformation by computing that random rigid motion for which the empirical measure of a small neighborhood is maximal. We apply our probabilistic toolbox, Theorem 4.6, to work out how to choose the clustering size and the number of random experiments. Next we give a pseudocode description of the generalized algorithm, which is Algorithm 12.

Algorithm 12: MaxOverlapRMRA

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

```
real δ ← ClusteringSizeRMRA(A, ε);
integer N ← SampleSizeRMRA(B, ε, δ, p);
collection Q ← ∅;
for i = 1 ... N do
    matrix M ← RandomRotationMatrix();
    point a ← RandomPoint(A);
    point b ← RandomPoint(B);
    add(Q, (M, b − Ma));
end
return FindDensestClusterRM(Q, δ, Δ_A);
```

Recall that for a rigid motion r a ball $B(r, \delta)$ equals the Cartesian product of a ball around the rotation matrix and a Euclidean ball around the translation vector. The ball around the matrix is with respect to the $|| \cdot ||_2$ norm, scaled by a factor Δ_A that is at least the diameter of A. In Function FindDensestClusterRM, the factor Δ_A is given as input.

Function FindDensestClusterRM(Q, δ, Δ)

Input: collection Q of points in $SO(d) \times R^d$, positive parameters δ, Δ

Output: point $r^* \in Q$ such that $|B(r^*, \delta) \cap Q| = \max_{r \in Q} |B(r, \delta) \cap Q|$, where Δ is the scaling factor of the matrix distance.

We chose the metric on the spaces of rigid motions in a way such that the clustering size scales with the shapes. The required number of random experiments N does not change if both shapes are scaled with the same constant.

Theorem 8.14. *Let A and B be shapes in constant dimension d, and let $\varepsilon, p \in (0,1)$ be parameters. There are $C_1, C_2 > 0$ such that, if ClusteringSizeRMRA(A, ε) returns a positive $\delta \leq C_1 \varepsilon \frac{|A|}{|\partial A|}$ and SampleSizeRMRA($B, \varepsilon, \delta, p$) returns an integer $N \geq C_2 \varepsilon^{-2} \delta^{-d^2-d} \Delta_A^{d^2-d} |B|^2 \log \frac{2}{p}$, then Algorithm 12 with input (A, B, ε, p) computes a rigid motion that maximizes the volume of overlap of A and B with probability at least $1 - p$ up to an additive error of $\varepsilon |A|$.*

First, we prove the remarkable fact that the density function at a rigid motion r is proportional to $|r(A) \cap B|$.

Lemma 8.15. *Let Y be the random variable on $SO(d) \times \mathbb{R}^d$ that draws pairs $(M, b - Ma)$ where $(M, a, b) \in SO(d) \times A \times B$ is drawn uniformly at random. The density function of Y is given by $f(r) = \frac{|r(A) \cap B|}{|A|\,|B|\,|SO(d)|}$.*

Proof. Our random experiment consists in drawing points from $\Omega = SO(d) \times A \times B$ uniformly at random. Our goal is to determine the density function f_Y of the random variable

$$Y : \Omega \to SO(d) \times \mathbb{R}^d,\ Y : (M, a, b) \mapsto (M, b - Ma).$$

The density function of Y can be expressed in terms of the conditional densities of the two random variables

$$\begin{aligned} &Y_{SO(d)} : \Omega \to SO(d), && Y_{SO(d)} : (M, a, b) \mapsto M, \\ &Y_{A \times B} : \Omega \to \mathbb{R}^d, && Y_{A \times B} : (M, a, b) \mapsto b - Ma. \end{aligned}$$

The density function of Y is the joint density of the random variables $Y_{SO(d)}$ and $Y_{A \times B}$. The conditional density function can also be expressed in terms of the joint density function

$$f_{A \times B}(t \mid Y_{SO(d)} = M) = f_Y(M, t)/f_{SO(d)}(M). \tag{8.6}$$

Clearly, $f_{SO(d)} \equiv 1/|SO(d)|$. The conditional density function $f_{A \times B}(t \mid Y_{SO(d)} = M)$ equals the density function of the translational case for the shapes MA and B and thus is $\frac{|(MA+t) \cap B|}{|A|\,|B|}$. Therefore, $f_Y(M, t) = \frac{|(MA+t) \cap B|}{|A|\,|B|\,|SO(d)|}$. □

Proof of Theorem 8.14. We apply Theorem 4.6 to $(SO(d) \times \mathbb{R}^d, d)$. Using Lemma 8.15, we get that with probability at least $1 - 2e^{-2N\tau^2}$ for all rigid motions r,

$$|r^*(A) \cap B| \geq |r(A) \cap B| - |A|\,|B| \Big(\frac{2c\sqrt{V/N} + 2\tau}{|B(r, \delta)|} + 3L\delta \Big)$$

where V is the VC dimension of $\mathcal{C}$ (Lemma 8.11). The number $L = \frac{|\partial A|}{|A|\,|B|\,|SO(d)|}$ is a Lipschitz constant of the density function by Corollary 3.5 and Proposition 8.9.

Let $\varepsilon, p \in (0, 1)$. We now determine bounds on N and δ such that this error term is at most $\varepsilon|A|$ with probability at least $1 - p$. More precisely, we determine the minimal value of N and a compatible value of δ for that we can guarantee an error of at most $\varepsilon|A|$ with probability at least $1 - p$. We compute δ and N only up to constants.

In order to get $2e^{-2N\tau^2} \leq p$, we set $\tau = \sqrt{\log \frac{2}{p}/(2N)}$. For sufficiently small δ, the volume of a δ-ball $B(r, \delta)$ is bounded from below by $C\Delta_A^{-\binom{d}{2}} \delta^{\binom{d+1}{2}}$ for some constant C, depending on d, by Theorem 8.3. Thus we get for the error term

$$|A|\,|B| \left(\frac{2c\sqrt{V/N} + 2\tau}{|B(r, \delta)|} + 3L\delta \right) =$$
$$O\left(|A|\,|B|\,\Delta_A^{\binom{d}{2}} \delta^{-\binom{d+1}{2}} N^{-1/2} \big(\log \tfrac{2}{p}\big)^{1/2} + 3|\partial A|\delta \right).$$

Therefore there are positive constants C_1 and C_2 such that, if $\delta \leq C_1 \varepsilon |A|/|\partial A|$ and $N \geq \varepsilon^{-2} \delta^{-d^2-d} \Delta_A^{d^2-d} |B|^2 \log \frac{2}{p}$, then the error is at most $\varepsilon|A|$ with probability at least $1-p$. □

The runtime of the algorithm equals the sum of the time to compute the clustering size δ and the sample size N, the time to generate N random rigid motions as described in the algorithm, and the time to determine the densest cluster.

Generating N random rigid motions costs $O(N + T(N))$ time. To compute N, we additionally need an explicit lower bound on $|B(M, \delta)|$ for a rotation matrix M. Theorem 8.3 does not quantify the rate of convergence of $|B(M, \delta)|/\delta^{\binom{d}{2}}$. Such a lower bound can be achieved by estimating $|B(M, \delta)|$ by integration over the manifold $SO(d)$, but this is not done here. If we did this and traced the constants that are hidden in the Big-Oh notation in the proof of Theorem 8.14, we could compute N and δ in constant time for shapes that fulfill the assumptions described in Section 2.2.

To determine the densest cluster, we simply check by brute force for every pair $\{r, r'\}$ of random rigid motions that is generated in the algorithm if $d(r, r') \leq \delta$. One test can be done in constant time. So we can determine a densest cluster in time $O(N^2)$.

Theorem 8.16. *Let A and B be shapes in constant dimension d, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters.*

Then Algorithm 12 with input (A, B, ε, p) computes a rigid motion that maximizes the volume of overlap of A and B with probability at least $1-p$ up to an additive error of $\varepsilon|A|$ in time $O(T(N) + N^2)$ where $N = O(\varepsilon^{-(d^2+d+2)} K_A^{d+1} D_A^{d-1} (M_B/m_A)^2 \log \frac{2}{p})$.

8.5. Improving the results in 3D

The main idea to improve the results in 3 dimensions is to represent rigid motions with fewer parameters. We then cope with the arising technical difficulties by using the refined probabilistic toolbox Theorem 4.7. The order of magnitude of the number of random samples N in the modified approach is the same as in Algorithm 12, but the modified representation of a rotation allows us to use neighborhoods that are rectangles in $\mathbb{R}^6$, compared to the more complex neighborhood in $\mathbb{R}^{12}$ that were used before. A best cluster can then be found in $O(N(\log N)^5)$ instead of $O(N^2)$. In principle this approach works in arbitrary dimension. To set limits to the technicalities, we concentrate on the most interesting case, namely on the case of 3 dimensions.

We use the representation from Corollary 8.6. For a fixed choice of D, we represent a rotation by the Givens angles $(\alpha, \beta, \gamma) \in (-\frac{\pi}{2}, \frac{\pi}{2}]^3$. Given the 3 angles, we can compute the corresponding rotation matrix in constant time. We apply the algorithm to the rotated shape DA for all four choices of diagonal matrices D.

We draw rotations $(\alpha,\beta,\gamma) \in (-\frac{\pi}{2},\frac{\pi}{2}]^3$ according to the distribution that has a density function proportional to $f(\alpha,\beta,\gamma) = \cos\gamma$ and corresponds to the uniform distribution on $SO(3)$. How to draw angles with respect to this distribution is described in [10], and it can be done in constant time.

The definition of a best cluster is modified as follows. We define a δ-neighborhood $B(r,\delta)$ of a rigid motion r to be a box in $\mathbb{R}^6$ with center r, side lengths $2\delta/\Delta_A$ for the rotational part and side lengths 2δ for the translational part. The definition of the function FindDensestClusterRM is modified to this notion of neighborhood.

It turns out that the density function g of the random experiment in the improved algorithm is proportional to $\cos\gamma\cdot|r(A)\cap B|$ where γ is the third Givens angle of r. To be able to apply Theorem 4.7 to $\frac{g(r)}{\cos\gamma}$, we have to avoid angles γ for which the cosine is close to 0. Observe that each rotation in 3D is of the form $G_1(\alpha)G_2(\beta)G_3(\gamma+\gamma')D$ where $\alpha,\beta \in (-\frac{\pi}{2},\frac{\pi}{2}]$, $\gamma \in (-\frac{\pi}{6},\frac{\pi}{6}]$ and $\gamma' \in \{-\frac{\pi}{3},0,\frac{\pi}{3}\}$.

Denote a diagonal matrix with entries $\varepsilon_1,\ldots,\varepsilon_d$ on the diagonal by $\mathrm{diag}(\varepsilon_1,\ldots,\varepsilon_d)$. We give a pseudocode description of the algorithm.

Function 3DFindDensentClusterRM(Q,δ,Δ)

Input: collection Q of points in $\Omega = (-\frac{\pi}{2},\frac{\pi}{2}]^2 \times (-\frac{\pi}{6},\frac{\pi}{6}] \times \mathbb{R}^3$, positive parameters δ,Δ

Output: point $r^* \in Q$ that maximizes the number $\frac{|B(r)\cap Q|}{\cos r_3 |B(r)\cap\Omega|}$ where $B(r) \subset \mathbb{R}^6$ is the box with center r and side lengths $2\delta/\Delta$ along the first three axes and side lengths 2δ along the second three axes.

Function RandomGivensAngles

Output: random triple $(\alpha,\beta,\gamma) \in (-\frac{\pi}{2},\frac{\pi}{2}]^2 \times (-\frac{\pi}{6},\frac{\pi}{6}]$ from the probability distribution that has the density function $f(\alpha,\beta,\gamma) = \frac{1}{\pi^2}\cos\gamma$.

Conceptually, the algorithm is applied to the rotated shape $(G_3(\gamma')D)A$ and B for each of the twelve choices of (γ',D) where $\gamma' \in \{-\frac{\pi}{3},0,\frac{\pi}{3}\}$ and D is one of the four diagonal matrices with entries ±1 on the diagonal and positive determinant. In each run, the approximate maximal volume of overlap over a partial set of rigid motions is computed such that the union of the sets covers the space of rigid motions. The sample size N in the twelve runs is chosen large enough such that with probability $\geq 1-p$ all twelve runs succeed, and therefore the rigid motion with the largest density estimate is an approximation of the overall maximum with probability $\geq 1-p$.

Let $(-\frac{\pi}{2},\frac{\pi}{2}]^3 \times \mathbb{R}^3$ be the (partial) space of rigid motions in 3 dimensions. Next we prove that the density function of this random experiment on $(-\frac{\pi}{2},\frac{\pi}{2}]^3 \times \mathbb{R}^3$ is proportional to $|r(A)\cap B|\cdot\cos\gamma$ for $r = (\alpha,\beta,\gamma,t_1,t_2,t_3)$.

Lemma 8.17. *Let Y be a random vector that draws $((\alpha,\beta,\gamma),b-Ma) \in (-\frac{\pi}{2},\frac{\pi}{2}]^3 \times \mathbb{R}^3$ where (α,β,γ), a and b are drawn independently, M is the rotation matrix with*

Given angles (α, β, γ), $a \in A$ and $b \in B$ are drawn u.a.r. and (α, β, γ) is drawn according to the density function given by $f(\alpha, \beta, \gamma) = \frac{1}{2\pi^2} \cos \gamma$, which corresponds to the uniform distribution on $SO(3)$. Then Y is distributed with density function $g(r) = \frac{|r(A) \cap B| \cdot \cos \gamma}{2\pi^2 |A|\,|B|}$ for $r = (\alpha, \beta, \gamma, t_1, t_2, t_3)$.

Proof. The proof proceeds along the same lines as the proof of Lemma 6.4. Consider the random experiment that consists in selecting triples $(R, a, b) \in (-\frac{\pi}{2}, \frac{\pi}{2}]^3 \times A \times B$ where R, a and b are selected independently, $a \in A$ and $b \in B$ are drawn from the uniform distribution and R is drawn from the distribution that has the density function $f(\alpha, \beta, \gamma) = \frac{1}{2\pi^2} \cos \gamma$ and corresponds to the uniform distribution on $SO(3)$ [10]. Let $I = (-\frac{\pi}{2}, \frac{\pi}{2}]^3$, let $\Omega = I \times A \times B$ and let M_R be the Givens matrix of R. We are interested in the density function f_Y of the random variable

$$Y : \Omega \to I \times \mathbb{R}^3, \quad Y : (R, a, b) \mapsto (R, b - M_R a).$$

We will express the density function of Y in terms of the conditional probability densities of the following two random variables Y_I and Y_T defined as

$$\begin{aligned} Y_I &: \Omega \to I, & Y_I &: (R, a, b) \mapsto R, \\ Y_T &: \Omega \to \mathbb{R}^3, & Y_T &: (R, a, b) \mapsto b - M_R a. \end{aligned}$$

The density function of Y is the joint density of the random variables Y_I and Y_T. By assumption, the marginal probability density of Y_I, i.e., probability density of $Y_I = \alpha$ allowing all possible values of Y_T, is

$$f_I(\alpha, \beta, \gamma) = \frac{\cos \gamma}{2\pi^2}.$$

The value of Y_T depends on the selected points a and b and on the value of Y_I. The conditional probability density of $Y_T = t$ given $Y_I = R$ is exactly the probability density in the space of translations for shapes $M_R A$ and B:

$$f_T(t \mid Y_I = R) = \frac{|(M_R A + t) \cap B|}{|A|\,|B|}.$$

The conditional probability density can also be expressed in terms of the joint probability density $f_T(t \mid Y_I = R) = f_Y(R, t)/f_I(R)$. Thus we get for any rigid motion $r = (R, t)$, $R = (\alpha, \beta, \gamma)$

$$g(r) = f_Y(r) = \frac{\cos \gamma \, |r(A) \cap B|}{2\pi^2 |A|\,|B|}. \qquad \square$$

Define the distance d of two rigid motions (R_1, t_1) and (R_2, t_2) as a scaled version of the distance that is induced by the maximum norm:

$$d\big((R_1, t_1), (R_2, t_2)\big) = \max\{\, \Delta_A ||R_1 - R_2||_\infty \,, \, ||t_1 - t_2||_\infty \,\}.$$

We prove that the density function g is Lipschitz continuous and determine a Lipschitz constant with respect to this distance.

Lemma 8.18. *The function g on the (partial) space of rigid motions $(-\frac{\pi}{2}, \frac{\pi}{2}]^3 \times \mathbb{R}^3$, given by $g(r) = \frac{\cos\gamma\, |r(A)\cap B|}{2\pi^2 |A|\,|B|}$ for $r = (\alpha, \beta, \gamma, t_1, t_2, t_3)$, is Lipschitz continuous with respect to the metric d with constant $L = \frac{42\,|\partial A|\,\Delta_A + |A|}{2\pi^2 |A|\,|B|\,\Delta_A}$.*

Proof. Let $r_1 = (\alpha_1, \beta_1, \gamma_1, t_1), r_2 = (\alpha_2, \beta_2, \gamma_2, t_2) \in (-\frac{\pi}{2}, \frac{\pi}{2}]^3 \times \mathbb{R}^3$ where $t_i \in \mathbb{R}^3$ denotes the translation vector. Let M_i be the Givens matrices of $(\alpha_i, \beta_i, \gamma_i)$ for $i = 1, 2$. Then

$$\begin{aligned} |g(r_1) - g(r_2)| &\leq \tfrac{1}{2\pi^2 |A|\,|B|} \Big(|\cos\gamma_1| \cdot \big| |r_1(A)\cap B| - |r_2(A) \cap B| \big| \\ &\qquad + |\cos\gamma_1 - \cos\gamma_2| \cdot |r_2(A)\cap B| \Big) \qquad (8.7) \\ &\leq \tfrac{1}{2\pi^2 |A|\,|B|} \Big(\big| |r_1(A)\cap B| - |r_2(A)\cap B| \big| + |\gamma_1 - \gamma_2| \cdot |A| \Big). \end{aligned}$$

To apply Corollary 3.5, we need to estimate $|M_1 a - M_2 a|$ for $a \in \mathrm{cl}(A)$. For a $(d \times d)$-matrix M, define $||M||_G = d \max_{ij} |m_{ij}|$. The matrix norm $||\cdot||_G$ is compatible with the Euclidean norm. Therefore $|M_1 a - M_2 a| \leq ||M_1 - M_2||_G \cdot |a| \leq ||M_1 - M_2||_G \Delta_A$. We have $M_i = G_{12}(\alpha_i) G_{13}(\beta_i) G_{23}(\gamma_i)$ for $i = 1, 2$. Abbreviate $G_{12}(\alpha_i) = M_\alpha^i$ etc. We have $||M_\alpha^i||_G = 3$. Next we upper bound $||M_1 - M_2||_G$. For this note that $||M_\alpha^1 - M_\alpha^2||_G = 3\max\{|\cos\alpha_1 - \cos\alpha_2|, |\sin\alpha_1 - \sin\alpha_2|\} \leq 3|\alpha_1 - \alpha_2|$ because for all $\alpha_1, \alpha_2 \in \mathbb{R}, \alpha_1 < \alpha_2$, there is $\xi \in [\alpha_1, \alpha_2]$ such that $\frac{\cos\alpha_2 - \cos\alpha_1}{\alpha_2 - \alpha_1} = (\cos\xi)' = -\sin\xi \in [-1, 1]$ by the mean value theorem, and similarly $\frac{\sin\alpha_2 - \sin\alpha_1}{\alpha_2 - \alpha_1} \in [-1, 1]$. Thus

$$\begin{aligned} ||M_1 - M_2||_G &= ||M_\alpha^1 M_\beta^1 M_\gamma^1 - M_\alpha^2 M_\beta^1 M_\gamma^1 + M_\alpha^2 M_\beta^1 M_\gamma^1 - M_\alpha^2 M_\beta^2 M_\gamma^2||_G \\ &\leq ||M_\alpha^1 - M_\alpha^2||_G ||M_\beta^1 M_\gamma^1||_G + ||M_\alpha^2||_G ||M_\beta^1 M_\gamma^1 - M_\beta^2 M_\gamma^2||_G \\ &\leq 27|\alpha_1 - \alpha_2| + 3||M_\beta^1 M_\gamma^1 - M_\beta^2 M_\gamma^2||_G \\ &\leq 27|\alpha_1 - \alpha_2| + 3(||M_\beta^1 - M_\beta^2||_G ||M_\gamma^1||_G + ||M_\beta^2||_G ||M_\gamma^1 - M_\gamma^2||_G) \\ &\leq 27(|\alpha_1 - \alpha_2| + |\beta_1 - \beta_2| + |\gamma_1 - \gamma_2|) \\ &\leq 81\, d(r_1, r_2)/\Delta_A. \end{aligned}$$

Thus we get $|M_1 a - M_2 a| \leq 81\, d(r_1, r_2)$. By Corollary 3.5, we have $\big| |r_1(A)\cap B| - |r_2(A)\cap B| \big| \leq \frac{1}{2}(81 d(r_1, r_2) + |t_1 - t_2|)|\partial A|$. We have $|t_1 - t_2| \leq \sqrt{3}||t_1 - t_2||_\infty$ and therefore $\frac{1}{2}(81 d(r_1, r_2) + |t_1 - t_2|) \leq \frac{81 + \sqrt{3}}{2} d(r_1, r_2) < 42\, d(r_1, r_2)$.

In total, we get by Inequality (8.7) $|g(r_1) - g(r_2)| \leq (\frac{21\,|\partial A|}{\pi^2 |A|\,|B|} + \frac{1}{2\pi^2 \Delta_A\,|B|})\, d(r_1, r_2) = \frac{42\,|\partial A|\,\Delta_A + |A|}{2\pi^2 |A|\,|B|\,\Delta_A}\, d(r_1, r_2)$. □

We now state the modified approximation result for rigid motions in 3 dimensions.

Theorem 8.19. *Let A and B be shapes in dimension 3, and let $\varepsilon, p \in (0, 1)$ be parameters. If 3DClusteringSizeRM(A, ε) returns a positive $\delta \leq \frac{\varepsilon\pi^2 |A| \Delta_A}{180|\partial A|\Delta_A + 20|A|}$ and 3DSampleSizeRM($B, \varepsilon, \delta, p$) returns an integer $N \geq 18c^2\varepsilon^{-2}\delta^{-12}|B|^2(\Delta_A^6 + \log\frac{24}{p})$,* *then Algorithm 13 with input (A, B, ε, p) computes a rigid motion 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 comes from Theorem 4.5.*

Proof.

1. Let $\delta \in \big(0, \frac{\varepsilon\pi^2|A|\Delta_A}{252|\partial A|\Delta_A+14|A|}\big]$.

2. Let $\Omega = (-\frac{\pi}{2}, \frac{\pi}{2}]^2 \times (-\frac{\pi}{6}, \frac{\pi}{6}] \times \mathbb{R}^3$ be a subspace of the space of rigid motions where the first three coordinates correspond to the Givens angles of a rotation, as explained above. Define twelve subsets of $SO(3)$ by

   $$\mathcal{R}_{\gamma',D} = \{M : M = G_1(\alpha)G_2(\beta)G_3(\gamma)G_3(\gamma')D,\ \alpha,\beta \in (-\tfrac{\pi}{2}, \tfrac{\pi}{2}],\ \gamma \in (-\tfrac{\pi}{6}, \tfrac{\pi}{6}]\}$$

   for each $D \in \{\text{diag}(1,1,1), \text{diag}(1,-1,-1), \text{diag}(-1,1,-1), \text{diag}(-1,-1,1)\}$ and $\gamma' \in \{-\frac{\pi}{3}, 0, \frac{\pi}{3}\}$. The union of the sets $\mathcal{R}_{\gamma',D}$ covers $SO(3)$ because of Corollary 8.6.

   Let $D \in \{\text{diag}(1,1,1), \text{diag}(1,-1,-1), \text{diag}(-1,1,-1), \text{diag}(-1,-1,1)\}$, $\gamma' \in \{-\frac{\pi}{3}, 0, \frac{\pi}{3}\}$ and let $R = G_3(\gamma')D$.

3. By Lemma 8.17, the density function of the random experiment in Line 9 to 15 of Algorithm 13 is given by $f(r) = \frac{|r(RA)\cap B|\cdot\cos\gamma}{2\pi^2|A|\,|B|}$ where $r = (\alpha, \beta, \gamma, t_1, t_2, t_3) \in \Omega$ and R is the rotation matrix described by (α, β, γ). Instead of sampling the shape RA, we sample A and rotate the sample points by R, which is a way of sampling RA uniformly distributed.

   By Lemma 8.18, the density function f is Lipschitz continuous with constant $L_f = \frac{42|\partial A|\Delta_A+|A|}{2\pi^2|A|\,|B|\Delta_A}$. The function f is bounded from above by $M_f = \frac{1}{2\pi^2|B|}$.

4. The cosine is Lipschitz continuous with constant 1, so the function $g : \Omega \to [\frac{1}{2}, 1]$, given by $g(\alpha, \beta, \gamma, t_1, t_2, t_3) = \cos\gamma$ is Lipschitz continuous with constant $L_g = \frac{1}{\Delta_A}$ with respect to the metric d.

5. For all $r \in \Omega$, the ball around r with radius δ with respect to d is a rectangle. If r has Givens angles α, β, γ such that α and β have distance $\geq \delta/\Delta_A$ to $-\frac{\pi}{2}$ and $\frac{\pi}{2}$ and γ has distance $\geq \delta/\Delta_A$ to all of the angles $-\frac{\pi}{2}, -\frac{\pi}{6}, \frac{\pi}{6}\frac{\pi}{2}$, then $|B(r,\delta)| = 64\,\delta^6/\Delta_A^3$. Otherwise, $|B(r,\delta)|$ is smaller, but it is at least $8\,\delta^6/\Delta_A^3$ since, for small δ, no angle can be closer than δ/Δ_A to $-\frac{\pi}{2}$ and $\frac{\pi}{2}$, for example.

   Let $\mathcal{B}$ be the set of balls $B(r,\delta)$ for all $r \in \Omega$. Because $\mathcal{B}$ is a set of rectangles in 6 dimensions, the VC dimension of $\mathcal{B}$ is ≤ 12 by Lemma 4.8.

6. From Line 7 to 16, Algorithm 13 computes a rigid motion r such that $|r(A)\cap B|$ is approximately maximal over the set of rigid motions $\mathcal{R}_{\gamma',D} \times \mathbb{R}^3$ with high probability. We make this statement precise by applying Theorem 4.7 to (Ω, d). For all $\tau > 0$, with probability $\geq 1 - 2e^{-2N\tau^2}$, we have

   $$|r^*(RA)\cap B| \geq |r(RA)\cap B| - |A|\,|B|\Big(\frac{c\Delta_A^3\sqrt{12/N}+\tau}{2\delta^6} + 6L_f\delta + 4M_fL_g\delta\Big).$$

   For $\tau = \sqrt{\log\frac{24}{p}/(2N)}$, we have this estimate with probability $\geq 1 - \frac{p}{12}$.

7. We show that under the assumptions in the theorem, the error term in Step 6 is at most $\varepsilon|A|$. We have $|A|\,|B|(6L_f + 4M_f L_g)\,\delta \le \frac{\varepsilon}{2}|A|$ because of
$$\frac{\varepsilon}{2(6L_f + 4M_f L_g)|B|} = \frac{\varepsilon\pi^2|A|\Delta_A}{252|\partial A|\Delta_A + 14|A|} \ge \delta.$$
The inequality $|A|\,|B|\Big(\frac{c\Delta_A^3\sqrt{12/N}+\tau}{2\delta^6}\Big) \le \frac{\varepsilon}{2}|A|$ is equivalent to
$$|B|^2\delta^{-12}\varepsilon^{-2}\big(\sqrt{12}c\,\Delta_A^3 + \sqrt{\tfrac{1}{2}\log\tfrac{24}{p}}\big)^2 \le N.$$
By an easy calculation, we have $\big(\sqrt{12}c\,\Delta_A^3 + \sqrt{\frac{1}{2}\log\frac{24}{p}}\big)^2 < 18c^2(\Delta_A^6 + \log\frac{24}{p})$, using that $c \ge 1$. In total, the error term is $\le \varepsilon|A|$ with probability $\ge 1 - \frac{p}{12}$.

8. The probability that one of the 12 loops fails is $\le p$. So with probability $\ge 1-p$, all runs succeed and the overall largest estimate $\frac{|B(X^*,\delta)\cap\{X_1,\dots,X_N\}|}{\cos X_3^*|B(X^*,\delta)|}$ gives an approximation of the maximal volume of overlap with the desired precision.

□

Theorem 8.20. *Let A and B be shapes in dimension 3, which are given by an oracle and fulfill the assumptions described in Section 2.2. Let $\varepsilon, p \in (0,1)$ be parameters. Then Algorithm 13 with input (A, B, ε, p) computes a rigid motion that maximizes the volume of overlap of A and B with probability at least $1-p$ up to an additive error of $\varepsilon|A|$ in time $O(T(N) + N(\log N)^5)$ where N is given as in Theorem 8.19.*

Algorithm 13: 3DMaxOverlapRM

Input: shapes $A, B \subset \mathbb{R}^3$, error tolerance $\varepsilon \in (0,1)$, allowed probability of failure $p \in (0,1)$

real $\delta \leftarrow$ 3DClusteringSizeRM(A, ε);
integer $N \leftarrow$ 3DSampleSizeRM$(B, \varepsilon, \delta, p)$;
collection Q; matrix D^*; real γ'^*;
real $V^* \leftarrow 0$; vector $r^* \leftarrow \mathbf{0}$;
for $D \in \{\mathrm{diag}(1,1,1), \mathrm{diag}(1,-1,-1), \mathrm{diag}(-1,1,-1), \mathrm{diag}(-1,-1,1)\}$ **do**
  for $\gamma' \in \{-\frac{\pi}{3}, 0, \frac{\pi}{3}\}$ **do**
    $Q \leftarrow \emptyset$;
    matrix $R \leftarrow G_3(\gamma')D$;   `// rotation of` A
    for $i = 1 \ldots N$ **do**
      point $(\alpha, \beta, \gamma) \leftarrow$ RandomGivensAngles();
      `// Compute corresponding rotation matrix`
      matrix $M \leftarrow G_1(\alpha)G_2(\beta)G_3(\gamma)$;
      point $a \leftarrow$ RandomPoint(A);
      point $b \leftarrow$ RandomPoint(B);
      add$(Q, (\alpha, \beta, \gamma, b - MRa))$;
    end
    (point,real) $(r, V) \leftarrow$ 3DFindDensentClusterRM(Q, δ, Δ_A);
    if $V > V^*$ **then**
      $V^* \leftarrow V$; $r^* \leftarrow r$;
      $D^* \leftarrow D$; $\gamma'^* \leftarrow \gamma'$;
    end
  end
end
`// Compute output rotation matrix`
matrix $M^* \leftarrow G_1(r_1^*)G_2(r_2^*)G_3(r_3^* + \gamma'^*)D^*$;
`// Output translation vector`
vector $t^* \leftarrow (r_4^*, r_5^*, r_6^*)$;
return (M^*, t^*);

A. The volume of small balls in the rotation group

We prove Theorem 8.3, which we used for the analysis of the algorithm for probabilistic matching of solid shapes under rigid motions in arbitrary dimension. For the proof, we have to introduce some notions from measure theory. For the convenience of the reader, we repeat the theorem here. Recall that we measure distances in the rotation group $SO(d)$ by the metric that is induced by the Frobenius norm. Further recall that $\mathcal{H}^k$ denotes the k-dimensional Hausdorff measure and that ω_k denotes the volume of the k-dimensional Euclidean unit ball.

Theorem 8.3. *For all $M \in SO(d)$, we have*

$$\lim_{\delta \to +0} \frac{\mathcal{H}^{\binom{d}{2}}\big(B(M,\delta)\big)}{\delta^{\binom{d}{2}}} = \omega_{\binom{d}{2}}.$$

Measure theoretic notions. We cite the following definitions and facts from [41]. A measure ϕ on a metric space Z is *locally finite* if for every $z \in Z$ there is a $\delta > 0$ such that $\phi(B(z,\delta)) < +\infty$. A measure ϕ on a topological space Z is a *Borel measure* if every Borel set in Z is ϕ-measurable. It is *Borel regular* if it is a Borel measure and each $X \subseteq Z$ is contained in a Borel set Y for which $\phi(X) = \phi(Y)$. A measure ϕ on Z is a *Radon measure* if it is a Borel measure, compact sets have finite measure ϕ, for open sets $U \subseteq Z$ holds $\phi(U) = \sup\{\phi(K) : K \subset U \text{ compact}\}$ and for all $X \subseteq Z$ holds $\phi(X) = \inf\{\phi(U) : X \subseteq U, U \text{ open}\}$. Every Radon measure is Borel regular. The converse is not true in general, but locally finite Borel regular measures on complete separable metric spaces are Radon measures.

The Haar measure: definition, existence and uniqueness. A *topological group* is a topological space with a group structure such that the group operation and the inverse function of the group are continuous with respect to the given topology. A topological group is *locally compact* if every point has a compact neighbourhood. The rotation group $SO(d)$ is a compact topological group.

We cite the following definition and theorems from [32], but we use the notions Borel measure, Borel regular and Radon measure as defined above and translate the theorems into this language.

A *Haar measure* is a Borel measure ϕ in a locally compact topological group G such that every non-empty open Borel set has positive measure ϕ, every compact set has finite measure ϕ and $\phi(E) = \phi(gE)$ for all $g \in G$ and all Borel sets E. The following theorems imply that there exists a unique normalized Haar measure on $SO(d)$ if we restrict our attention to Radon measures.

Theorem A.1. *In every locally compact topological group G exists at least one Radon measure that is a Haar measure.*

Theorem A.2. *If μ and ν are Haar measures on a locally compact topological group G that are Radon measures, then there exists a constant $c > 0$ such that $\mu = c\nu$.*

The Hausdorff measure on $SO(d)$. With any measure ϕ on Z, and any set $Y \subseteq Z$, one associates another measure $\phi \angle Y$ on Z, the *restriction of ϕ to Y*, by the formula

$$(\phi \angle Y)(X) = \phi(X \cap Y) \text{ for } X \subseteq Z.$$

The Hausdorff measure $\mathcal{H}^m$, which is defined in Section 2.1, is Borel regular. If $Y \subset \mathbb{R}^n$ is $\mathcal{H}^m$-measurable and $\mathcal{H}^m(Y) < +\infty$, then $\mathcal{H}^m \angle Y$ is a Radon measure (Theorem 1.9(2), Corollary 1.11 and Corollary 4.5 in [41]).

The set of orthogonal matrices $O(d) \subset \mathbb{R}^{d^2}$ is $\mathcal{H}^m$-measurable for every m since it is a closed set. The value $\mathcal{H}^{\binom{d}{2}}(O(d))$ is finite [28, 3.2.28]. Since $SO(d) \subset O(d)$, we also have $\mathcal{H}^{\binom{d}{2}}(SO(d)) < +\infty$ and $\mathcal{H}^{\binom{d}{2}} \angle SO(d)$ is a Radon measure. We will consider it as a measure on $SO(d)$.

In general, if m is an integer and S is a sufficiently regular m-dimensional surface in $\mathbb{R}^n$, for example a C^1-submanifold, then the restriction $\mathcal{H}^m \angle S$ is a constant multiple of the surface area on S. The surface area on S is the measure that is given by integration over the characteristic function of a set over S. The computation of such an integral can be reduced to an integral over $\mathbb{R}^m$ by transformation formulas; see for example [50].

Proposition A.3. *The measure $\mathcal{H}^{\binom{d}{2}} \angle SO(d)$ on $SO(d)$ is a Radon measure and a Haar measure.*

Proof. We have already seen that $\mathcal{H}^{\binom{d}{2}} \angle SO(d)$ is a Radon measure on $\mathbb{R}^{d^2}$. By the definition of relative topology, the open sets of $SO(d)$ are exactly the sets $U \cap SO(d)$ for U open in $\mathbb{R}^{d^2}$. The compact sets of $SO(d)$ are exactly the sets $K \cap SO(d)$ for K compact in $\mathbb{R}^{d^2}$ because $SO(d)$ is bounded. It follows directly from the definition of a Radon measure that $\mathcal{H}^{\binom{d}{2}} \angle SO(d)$ is a Radon measure on $SO(d)$.

Let M be a rotation matrix and $E \subset SO(d)$. For all $\delta > 0$ and all coverings $\{S_1, S_2, \dots\}$ of E with non-empty sets S_i of diameter $\leq \delta$, $\{MS_1, MS_2, \dots\}$ is a covering of ME with non-empty sets S_i of diameter $\leq \delta$. Therefore, $\mathcal{H}^m_\delta(E) = \mathcal{H}^m_\delta(ME)$ and $\mathcal{H}^m(E) = \mathcal{H}^m(ME)$ for all m. □

Proof of Theorem 8.3 We need the notion of the density of a measure and another theorem for the proof of Theorem 8.3, which we cite from [28].

Whenever ϕ measures a metric space Z, $z \in Z$, and m is a natural number, the *m-dimensional upper and lower densities of ϕ at z* are defined as $\Theta^{*m}(\phi, z) = \limsup_{\delta \to +0} \frac{\phi(B(z,\delta))}{\omega_m \delta^m}$ and $\Theta_*^m(\phi, z) = \liminf_{\delta \to +0} \frac{\phi(B(z,\delta))}{\omega_m \delta^m}$ where $B(z, \delta)$ is the ball of radius δ centered at z, as usual. If the upper and lower density are equal, their common value $\Theta^m(\phi, z) = \lim_{\delta \to +0} \frac{\phi(B(z,\delta))}{\omega_m \delta^m}$ is called the *m-dimensional density of ϕ at z*. Recall the definition of rectifiability from Section 2.1. We now know all definitions to understand the following theorem.

Theorem A.4 (Theorem 3.2.19 in [28])**.** *If W is an $(\mathcal{H}^m, m)$-rectifiable and $\mathcal{H}^m$-measurable subset of $\mathbb{R}^n$, then, for $\mathcal{H}^m$-almost all $w \in W$,*

$$\Theta^m(\mathcal{H}^m \angle W, w) = 1.$$

A common matrix norm is the operator norm, defined by $||M|| = \sup_{|x|=1} |Mx|$. For $M \in SO(d)$, this norm equals 1. For $\delta > 0$, denote the open ball $B_*(M, \delta) = \{N \in SO(d) : ||M - N|| < \delta\}$. The Euclidean norm $|\cdot|$ is compatible with the operator norm. The proof of Theorem 8.3 is an application of Theorem A.4 to the set $W = B_*(I, 1)$ for $m = \binom{d}{2}$. Let us check that W and m satisfy the assumptions of the theorem.

Lemma A.5. *For all $\varepsilon \in [0, 1)$, we have that the open ball $B_*(I, \varepsilon) \subset SO(d)$ is $\binom{d}{2}$-rectifiable.*

Proof. We have to show that $B_*(I, \varepsilon)$ is the image of a Lipschitz function of some bounded set in $\mathbb{R}^{\binom{d}{2}}$. Define the usual matrix exponential as

$$\exp : \mathbb{R}^{d \times d} \to \mathbb{R}^{d \times d}, \ \exp(M) = \sum_{k \geq 0} \frac{M^k}{k!}.$$

For the definition of the matrix exponential and its properties, we follow [12]. The function exp is invertible on the ball $B_*(I, 1)$. Its inverse log is defined as $\log(M) = \sum_{k \geq 1} \frac{(-1)^{k-1}}{k} (M - I)^k$ for all $||M - I|| < 1$, and it is continuous. Since the continuous image of a compact set is bounded, $\log(\mathrm{cl}(B_*(I, \varepsilon)))$ is bounded for every $\varepsilon \in [0, 1)$. Therefore, $\log(B_*(I, \varepsilon))$ is bounded for every $\varepsilon \in [0, 1]$.

The function exp maps exactly the skew-symmetric matrices to rotation matrices. A matrix $M = (m_{ij})_{1 \leq i,j \leq d}$ is skew-symmetric if $m_{ij} = -m_{ji}$ for all $1 \leq i, j \leq d$. The set of skew-symmetric matrices can be identified with $\mathbb{R}^{\binom{d}{2}}$ by forgetting the zero diagonal and the lower triangular matrix. So, we can identify $\log(B_*(I, \varepsilon))$ with a bounded subset of $\mathbb{R}^{\binom{d}{2}}$ for every $\varepsilon \in [0, 1)$.

It remains to show that exp is a Lipschitz map on every bounded set of matrices. This is true because of

$$||\exp(X) - \exp(Y)|| \leq ||X - Y|| \exp(||X - Y||) \exp(||X||)$$

for all matrices X, Y and each submultiplicative matrix norm $||\cdot||$. Equivalently, we show that $||\exp(X+Y)-\exp(X)|| \leq ||Y||\exp(||X||)\exp(||Y||)$:

$$\begin{aligned}
||\exp(X+Y)-\exp(X)|| &\leq \sum_{k=1}^{\infty}\frac{1}{k!}||(X+Y)^k - X^k|| \\
&\leq \sum_{k=1}^{\infty}\frac{1}{k!}\sum_{l=0}^{k-1}\binom{k}{l}||X||^l||Y||^{k-l} \\
&\leq ||Y||\sum_{k=0}^{\infty}\sum_{l=0}^{k}\frac{||X||^l}{l!}\frac{||Y||^{k-l}}{(k-l)!} \\
&= ||Y||\exp(||X||)\exp(||Y||).
\end{aligned}$$

□

The fact that $B_*(I,1) = \bigcup_{k=1}^{\infty} B_*(I, 1-\frac{1}{k})$ and Lemma A.5 imply the following.

Corollary A.6. *The ball $B_*(I,1) \subset SO(d)$ is $(\mathcal{H}^{\binom{d}{2}}, \binom{d}{2})$-rectifiable.*

Proof of Theorem 8.3. Let $W = B_*(I,1) \subset SO(d)$ and $m = \binom{d}{2}$. W is $\mathcal{H}^{\binom{d}{2}}$-measurable because it is a Borel set and it is $(\mathcal{H}^{\binom{d}{2}}, \binom{d}{2})$-rectifiable due to Corollary A.6. We now apply Theorem A.4, giving that

$$\lim_{\delta\to+0}\frac{\mathcal{H}^{\binom{d}{2}}(B(w,\delta))}{\delta^{\binom{d}{2}}} = \omega_{\binom{d}{2}}$$

for almost all $w \in B_*(I,1)$. Since all δ-balls in $SO(d)$ have the same measure for a fixed $\delta > 0$, the claim holds even for all $w \in SO(d)$. □

Bibliography

[1] H.-K. Ahn, P. Brass, and C.-S. Shin. Maximum overlap and minimum convex hull of two convex polyhedra under translations. *Computational Geometry, Theory and Applications (CGTA)*, 40:171–177, 2008.

[2] H.-K. Ahn, O. Cheong, C.-D. Park, C.-S. Shin, and A. Vigneron. Maximizing the overlap of two planar convex sets under rigid motions. In *Symposium on Computational Geometry*, pages 356–363, 2005.

[3] H. Alt, U. Fuchs, G. Rote, and G. Weber. Matching convex shapes with respect to the symmetric difference. *Algorithmica*, 21:89–103, 1998.

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

[5] H. Alt and L. Scharf. Shape matching by random sampling. In *3rd Annual Workshop on Algorithms and Computation (WALCOM 2009)*, volume 5431 of *Lecture Note in Computer Science*, pages 381–393, 2009.

[6] H. Alt and L. Scharf. Shape matching by random sampling. *Theoretical Computer Science*, In Press, Corrected Proof, 2010.

[7] H. Alt, L. Scharf, and S. Scholz. Probabilistic matching of sets of polygonal curves. In *Proceedings of the 22nd European Workshop on Computational Geometry (EWCG)*, pages 107–110, Delphi, Greece, March 2006.

[8] H. Alt, L. Scharf, and S. Scholz. Probabilistic matching and resemblance evaluation of shapes in trademark images. In *Proceedings of the ACM International Conference on Image and Video Retrieval (CIVR)*, pages 533–540, Amsterdam, The Netherlands, July 2007.

[9] 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).

[10] T.W. Anderson, I. Olkin, and L.G. Underhill. Generation of random orthogonal matrices. *SIAM Journal on Scientific and Statistical Computing*, 8:625–629, 1987.

[11] G. Averkov and G. Bianchi. Confirmation of Matheron's conjecture on the covariogram of a planar convex body. *Journal of the European Mathematical Society*, 11:1187–1202, 2009.

[12] A. Baker. *Matrix Groups: An Introduction to Lie Group Theory.* Springer, 2002.

[13] C. Benassi, G. Bianchi, and G. D'Ercole. Covariogram of non-convex sets. *Mathematika*, 56:267–284, 2010.

[14] G. Bianchi. The covariogram determines three-dimensional convex polytopes. *Advances in Mathematics*, 220:1771–1808, 2009.

[15] G. Bianchi. The cross covariogram of a pair of polygons determines both polygons, with a few exceptions. *Advances in Applied Mathematics*, 42:519–544, 2009.

[16] G. Bianchi. Matheron's conjecture for the covariogram problem. *Journal of the London Mathematical Society*, 11:1187–1202, 2009.

[17] T. Bonnesen and W. Fenchel. *Theory of convex bodies.* BSC Associates, 1987. Translated from the German and edited by L. Boron and C. Christenson and B. Smith with the collaboration of W. Fenchel. Originally published as *Theorie der konvexen Körper*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 3 (berichtigter reprint, 1974).

[18] P. Bose, O. Cheong, and V. Dujmović. A note on the perimeter of fat objects. *Computational Geometry*, 44(1):1–8, 2011.

[19] T.M. Chan. A (slightly) faster algorithm for Klee's measure problem. *Computational Geometry*, 43(3):243 – 250, 2010. Special Issue on 24th Annual Symposium on Computational Geometry (SoCG'08).

[20] 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.

[21] M. de Berg. Improved bounds on the union complexity of fat objects. *Discrete and Computational Geometry*, 40(1):127–140, 2008.

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

[23] M. de Berg, O. Devillers, M. van Kreveld, O. Schwarzkopf, and M. Teillaud. Computing the maximum overlap of two convex polygons under translations. *Theory of computing systems*, 31:613–628, 1998.

[24] L. Devroye. *Non-Uniform Random Variate Generation.* Springer, 1986.

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

[26] R.M. Dudley. Balls in $\mathbb{R}^k$ do not cut all subsets of $k+2$ points. *Advances in Mathematics*, 31:306 – 308, 1979.

[27] A. Efrat. The complexity of the union of (α, β)-covered objects. *SIAM J. Comput.*, 34(4):775–787, 2005.

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

[29] B. Galerne. Computation of the perimeter of measurable sets via their covariogram. applications to random sets. Preprint. `http://hal.archives-ouvertes.fr/hal-00480825/en/`, 16 June 2010, 17 pages.

[30] M. Hagedoorn. *Pattern matching using similarity measures.* PhD thesis, Utrecht University, 2000. `ftp://ftp.cs.uu.nl/pub/RUU/CS/phdtheses/Hagedoorn/thesis.pdf`.

[31] M. Hagedoorn and R.C. Veltkamp. Metric pattern spaces. Technical Report UU-CS-1999-03, Department of Information and Computing Sciences, Utrecht University, 1999. `http://www.cs.uu.nl/research/techreps/repo/CS-1999/1999-03.pdf`.

[32] P.R. Halmos. *Measure Theory.* Springer, 1950.

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

[34] J. Hoffmann-Jørgensen. *Probability with a view toward statistics, volume I & II.* Chapman & Hall, 1994.

[35] F. John. *Extremum problems with inequalities as subsidiary conditions*, pages 187–204. Interscience Publishers, New York, 1948.

[36] U. Krengel. *Einführung in die Wahrscheinlichkeitstheorie und Stochastik.* Vieweg, fifth edition, 2000.

[37] R. Lang. A note on the measurability of convex sets. *Archiv der Mathematik*, 47:90–92, 1986.

[38] L.J. Latecki and R.C. Veltkamp. Properties and performances of shape similarity measures. In *Proceedings of International Conference on Data Science and Classification (IFCS)*, 2006.

[39] G. Matheron. *Random Sets and Integral Geometry.* Wiley, 1975.

[40] J. Matoušek. *Lectures on Discrete Geometry.* Springer, 2002.

[41] P. Mattila. *Geometry of Sets and Measures in Euclidean Spaces.* Cambridge University Press, 1995.

[42] J.S.B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *SODA '07 Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, 2007.

[43] 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.

[44] D. Mumford. Mathematical theories of shape: Do they model perception? In B. Vemuri, editor, *Geometric Methods in Computer Vision, SPIE*, volume 1570, pages 2–10, 1991.

[45] L. Scharf. *Probabilistic Matching of Planar Shapes.* PhD thesis, Freie Universität Berlin, Institut für Informatik, June 2009. `http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000010520`.

[46] S. Scholz. *Visual shape similarity and retrieval of figurative images.* PhD thesis, Freie Universität Berlin, Institut für Informatik, June 2011. `http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000023035`.

[47] 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`.

[48] W.D. Smith. How to sample from a probability distribution, 2002. `http://introcs.cs.princeton.edu/java/98simulation/smith-sampling.ps`.

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

[50] M.E. Taylor. *Measure Theory and Integration.* AMS Graduate Studies in Mathematics, Volume 76, 2006.

[51] A.F. van der Stappen, D. Halperin, and M.H. Overmars. The complexity of the free space for a robot moving amidst fat obstacles. *Computational Geometry*, 3(6):353 – 373, 1993.

[52] V.N. Vapnik and A.Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theor. Probab. Appl.*, 16:264–280, 1971.

[53] R.C. Veltkamp. Shape matching: Similarity measures and algorithms. In *Shape Modeling International*, pages 188–197. IEEE Computer Society, 2001.

[54] 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.

[55] A.J. Walker. An efficient method for generating discrete random variables with general distributions. *ACM Trans. Math. Softw.*, 3:253–256, September 1977.

Erklärung

Hiermit versichere ich, Daria Schymura, dass ich die vorliegende Arbeit selbständig angefertigt und ohne fremde Hilfe verfasst habe, keine außer den von mir angegebenen Hilfsmitteln und Quellen dazu verwendet habe und die den benutzten Werken inhaltlich oder wörtlich entnommenen Stellen als solche kenntlich gemacht habe.

Berlin, den 1. September 2011

Daria Schymura