Karl-Franzens-Universität Graz & Technische Universität Graz

SPEZIALFORSCHUNGSBEREICH F 003

OPTIMIERUNG
und
KONTROLLE

Projektbereich
DISKRETE OPTIMIERUNG

Helmut Alt   Ulrich Fuchs   Günter Rote   Gerald Weber

Matching Convex Shapes with Respect to the Symmetric Difference

Bericht Nr. 87 — März 1997 — Version 1c

to appear in Algorithmica (1998)

Matching Convex Shapes with Respect to the Symmetric Difference

Helmut Alt* Ulrich Fuchs* Günter Rote† Gerald Weber*

December 23, 1997

Abstract

This paper deals with questions from convex geometry related to shape matching. In particular, we consider the problem of moving one convex figure over another, minimizing the area of their symmetric difference. We show that if we just let the two centers of gravity coincide, the resulting symmetric difference is within a factor of 11/3 of the optimum. This leads to efficient approximate matching algorithms for convex figures.

1 Introduction

A very common problem arising in application areas like computer vision or pattern recognition is that two "figures" F_1 and F_2 are given and the question is how much these figures "look alike". F_1 might be an image of an unknown object, and F_2 might be one of several possible templates for this object. In other words, we want to *match* F_1 and F_2 as good as possible. The quality of a match is measured by some "distance" function $\delta(A, B)$ which assigns a number to any two sets A and B.

More precisely, assume that we consider a certain set $\mathcal{T}$ of feasible transformations that may be used for matching. Then we define the *shape matching problem* as follows:

> *Given* two figures F_1 and F_2, *find* a transformation $t^{\mathrm{opt}} \in \mathcal{T}$ minimizing $\delta(F_1, t(F_2))$ over all $t \in \mathcal{T}$.

Reasonable sets of matching transformations are for example translations, rigid motions (i.e., compositions of translations and rotations), similarities, arbitrary affine mappings, or projective transformations. (Note that for transformations that allow changes of size, such as homotheties, it makes a difference if we exchange F_1 and F_2.)

Most of the previous work has concentrated on the *Hausdorff distance* as a distance measure [ABB91, AST94, CGH⁺93, HKS93, AAR97]. Since solving the optimization problem exactly turns out to be rather difficult, more efficient *approximation algorithms* have been developed. These algorithms do not necessarily find the optimum but a solution whose quality is within a constant factor of the optimal one. The simplest approach for getting approximation algorithms uses *reference points*. Roughly speaking, a reference point is a characteristic point with the property that if two figures are matched optimally, their reference points lie close together. Conversely, if we restrict to matching transformations that map the reference point of F_2 onto the the reference point of F_1, the best solution in this restricted set cannot be much worse than the optimal solution. The restricted set of transformations has fewer degrees of freedom, and thus the restricted optimization problem is easier to solve.

For example, if $\mathcal{T}$ is the set of translations, the restricted optimal translation is directly available, namely the vector between the two reference points. In the case of rigid motions, the two reference

* *name*@inf.fu-berlin.de, Institut für Informatik, Freie Universität Berlin, Takustr. 9, D-14195 Berlin, Germany. This research was supported by grant No. Al 253/4-1 from Deutsche Forschungsgemeinschaft DFG (German Research Association).

† rote@opt.math.tu-graz.ac.at, Institut für Mathematik, Technische Universität Graz, Steyrergasse 30, A-8010 Graz, Austria. This research was supported by the Spezialforschungsbereich F 003 "Optimierung und Kontrolle", Projektbereich Diskrete Optimierung.

points are matched and then the optimal position of F_2 is sought among rotations around this point.

Formally, we call a map that assigns to each figure in a certain class of figures $\mathcal{F}$ a point in the plane a *reference point* for $\mathcal{T}$ (with respect to a distance δ and with respect to $\mathcal{F}$) if there is a constant $c \geq 1$, called the *approximation factor*, such that for any two figures $F_1, F_2 \in \mathcal{F}$ there exists $r \in \mathcal{T}$ mapping the reference point of F_2 onto the reference point of F_1 and fulfilling for all $t \in \mathcal{T}$ the inequality

$$\delta(F_1, r(F_2)) \leq c \cdot \delta(F_1, t(F_2)).$$

With respect to the Hausdorff distance and rigid motions (and also more general classes of transformations) the *center of gravity* or *centroid* of the convex hull is not a reference point: This is easily seen by considering a very long rectangle on the one hand and one of the triangles obtained from the rectangle by cutting along the diagonal on the other hand.

However, two reference points with respect to the Hausdorff distance and rigid motions have been found: the centroid of the boundary of the convex hull [ABB91] and the so-called *Steiner point* [AAR97]. The Steiner point of a convex polygon is obtained as the center of gravity when a mass proportional to the exterior angle is placed at each vertex. For a smooth convex body, the mass has to be distributed on the boundary proportional to the curvature.

Here we consider a different distance measure between figures, namely the *area of the symmetric difference* which we denote by

$$\delta(F_1, F_2) := \operatorname{area}(F_1 \triangle F_2) = \operatorname{area}(F_1 \setminus F_2) + \operatorname{area}(F_2 \setminus F_1).$$

For a planar region F, we will often just write F instead of area(F) when no confusion arises.

The symmetric difference δ is one of the standard error measures considered in the theory of convex approximation, see the surveys of Gruber [G83, G93]. In the area of computational geometry, δ has been investigated only in a few papers so far, including [ABGW90], where simplification problems are addressed, and a recent paper of de Berg et al. [BDK+96], which is also concerned with matching problems under translations.

In some applications, δ is more appropriate than the Hausdorff distance. Consider the case when F_1 is an image disturbed by noise: noise may add thin features to the boundary, but it is unlikely to change large areas. The Hausdorff distance may change dramatically, even if only a single point is added to F_1, whereas the optimal matching for the symmetric difference will hardly change, even if the noise adds some areas that are disconnected from F_1.

For measurable sets A, B, C with finite areas, the distance function δ satisfies the triangle inequality:

$$\delta(A, C) \leq \delta(A, B) + \delta(A, C) \tag{1}$$

This follows from the set-theoretic relation

$$A \triangle B \subseteq (A \triangle C) \cup (B \triangle C).$$

(To obtain a metric, some regularity conditions must be imposed on the sets A, B, C. For example, for bounded sets which are equal to the closure of their interior, or for compact convex sets of positive area, δ is a metric.)

In this paper we restrict our attention to *convex* figures only. For convex plane figures, we show that the centroid is a reference point for translations, rigid motions, and some other sets of transformations. In particular, if we translate a convex figure F_2 so that its centroid matches the centroid of another convex figure F_1, the resulting symmetric difference is at most $11/3$ times as large as the optimal one under translations. We give an example showing that this constant is optimal. This theorem is the main geometric result of the paper, and it is proved in Section 2.

A related theorem has been obtained by de Berg, Devillers, van Kreveld, Schwarzkopf, and Teillaud [BDK+96]. If instead of minimizing the area of the symmetric difference $F_1 \triangle t(F_2)$, we *maximize the area of the intersection* $I = F_1 \cap t(F_2)$, we get of course the same result since $\delta(F_1, t(F_2)) = F_1 + F_2 - 2I$. However, when it comes to the relative performance of approximation algorithms, there is a difference. De Berg et al. considered the very same heuristic as in our case, namely letting the centroids coincide. They showed that the area of the intersection that is obtained in this way is at least 9/25 of the maximum area that can be obtained by translating F_2. We will discuss the relation between this result and our result in the concluding section.

In Section 3, we extend our result from translations to more general sets of transformations. In Section 4, we apply this result to obtain approximate matching algorithms for various sets of transformations.

This paper was presented at the Fourth Annual European Symposium on Algorithms in Barcelona (ESA '96) [AFRW96].

2 A reference point for translations

The following is the key lemma for our main result.

Lemma 1 *Let $F \subset \mathbb{R}^2$ be a bounded convex set, let $f \subset F$ be a measurable subset of F with positive area. and let s_F and s_f denote the centroids of these sets. Let w be the length of the projection of F onto a line perpendicular to the vector $s_F - s_f$. Then*

$$w \cdot |s_F - s_f| \leq \frac{4}{3}(F - f).$$

The inequality is strict if $F \setminus f$ has positive area.

Proof. We assume w.l.o.g. that s_F and s_f have the same x-coordinate so that w is the horizontal width of F. We also assume that s_f lies below s_F. We will transform the sets f and F in five steps into more special sets (see Figure 1). Their area and the horizontal width w of F will not change. The centroids move, but in each step the distance of their y-coordinates will not decrease. After the final step we are able to prove the inequality directly. For simplicity, the sets are called f and F throughout the process.

Step 1.
Let L and R be the leftmost and rightmost point of F. (If L or R is not unique, we choose arbitrarily.) We use a *shearing* that preserves x-coordinates and transforms F and f in such a way that L and R have the same y-coordinate.

Step 2.
For a horizontal line g we denote by H^+ and H^- the upper and lower halfplane bounded by g, respectively. Choose g in such a way that $F \cap H^-$ has the same area as f and replace f by $F \cap H^-$. Since some part of f's area has been transferred from above g to below g, the y-coordinate of s_f is smaller than before.

Step 3.
We choose a vertical line s and apply *Steiner symmetrization* (see [BF33, §9]) to make F and f symmetric about the axis s. This operation can be imagined as cutting F into infinitesimally flat horizontal slices and arranging these slices symmetrically about s. The centroids s_f and s_F may move in this step, but their y-coordinates do not change. The equality $f = F \cap H^-$ is still valid. Because of the transformation carried out in Step 1, w is unchanged.

Now we would like to assume that f lies below the line segment LR. If this is not the case, we can exchange the roles of f and its complementary set $\overline{f} := F \setminus f$: the formula $s_F = \frac{f}{F} s_f + \frac{\overline{f}}{F} s_{\overline{f}}$ implies that

$$\frac{s_F - s_f}{F - f} = -\frac{s_F - s_{\overline{f}}}{F - \overline{f}}.$$

Thus, showing the upper bound of the lemma for $\overline{f}$ is as good as showing it for f.

Step 4.
Steps 4 and 5 transform F into a union of two isosceles triangles with base LR. First we find a point O on the vertical axis s such that the area of the triangle LRO below g is equal to the area of f. There is a horizontal line a such that $F \setminus LRO$ lies completely above a, whereas $LRO \setminus F$ lies completely below a. (This line goes through the intersection points of the segments LO and RO with the boundary of F.) Now we set $f := LRO \cap H^-$ and $F := (F \cap H^+) \cup (LRO \cap H^-)$. F may now be temporarily non-convex, but f and F have the same area as before. As in Step 2, we may regard this transformation of f as a movement of some of its mass from above a ($f_1 = f \setminus LRO$) to

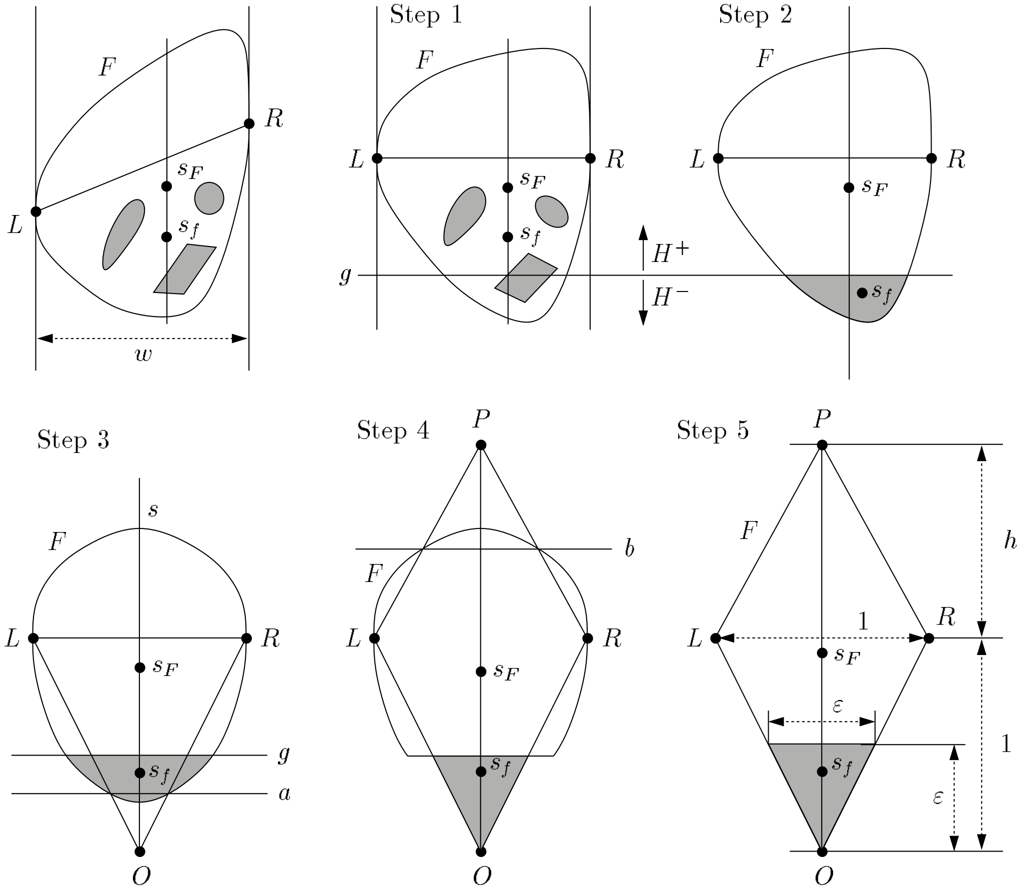

Figure 1: Transformations of F and f. The set f is the shaded area.

below a ($f_2 = LRO \setminus f$). Hence the y-coordinate of s_f has decreased by some amount $\varepsilon \geq 0$. The set F has undergone the same change (deletion of f_1 and addition of f_2), but as the set F is bigger than f, the y-coordinate of s_F decreases only by $\frac{f}{F} \cdot \varepsilon$. Hence the distance $|s_F - s_f|$ is at least as large as before.

Step 5.

We find a point P above LR on the vertical axis s such that the total area of the quadrilateral $LPRO$ is equal to the area of F. As in Step 4, we find a horizontal line b that separates $F \setminus LPRO$ from $LPRO \setminus F$. (Possibly this line goes through L and R.) We finally set $F := LPRO$, leaving f intact, and reasoning as above, we conclude that the y-coordinate of s_F increases by some nonnegative amount. Hence the distance $|s_F - s_f|$ is at least as large as before.

For the figure remaining after Step 5 we show the claim of the lemma directly. We assume without loss of generality that O is the origin and the height of the triangle LRO and its width $w = LR$ are both 1. Let the height of the triangle f be ε, and let the height of the triangle LRP be $h \geq 0$. Then area$(f) = \varepsilon^2/2$ and area$(F) = (1+h)/2$. The y-coordinate of s_f is $\frac{2}{3}\varepsilon$, and the y-coordinate of s_F is the same as for the centroid of the triangle LOP, which is $\frac{2}{3} \cdot (1 + h/2)$. Using the inequalities $0 < \varepsilon \leq 1$, we get

$$w \cdot |s_F - s_f| = \frac{2}{3}(1 + h/2 - \varepsilon) \leq \frac{2}{3}(1 + h - \varepsilon^2) = \frac{4}{3}(F - f). \tag{2}$$

If $F - f > 0$, we must have either $\varepsilon < 1$ or $h > 0$, and the inequality becomes strict. $\square$

Now we come to the proof of the main theorem. Consider convex bodies F_1 and F_2 in the plane. Let $\delta^{\text{opt}}(F_1, F_2)$ denote the minimal area of the symmetric difference between translates of F_1 and F_2, and let $\delta^{\text{C}}(F_1, F_2)$ denote the area of the symmetric difference between translates of F_1 and F_2 whose centroids coincide.

Theorem 1 *For convex plane bodies F_1 and F_2, we have*

$$\delta^{\mathrm{C}}(F_1, F_2) \leq \frac{11}{3} \cdot \delta^{\mathrm{opt}}(F_1, F_2).$$

The inequality is strict unless both sides are zero. The constant 11/3 in the inequality cannot be improved.

Proof. We first assume that $F_2 \subseteq F_1$, so $\delta^{\mathrm{opt}}(F_1, F_2) = F_1 - F_2 = F_1 \triangle F_2$. Let s_1, s_2 be the centroid of F_1 and F_2, respectively. Now suppose that F_2 is translated by the vector $s_1 - s_2$ from F_2—the optimal position—into the position F_2' where s_1 and s_2 are matched (see Figure 2).

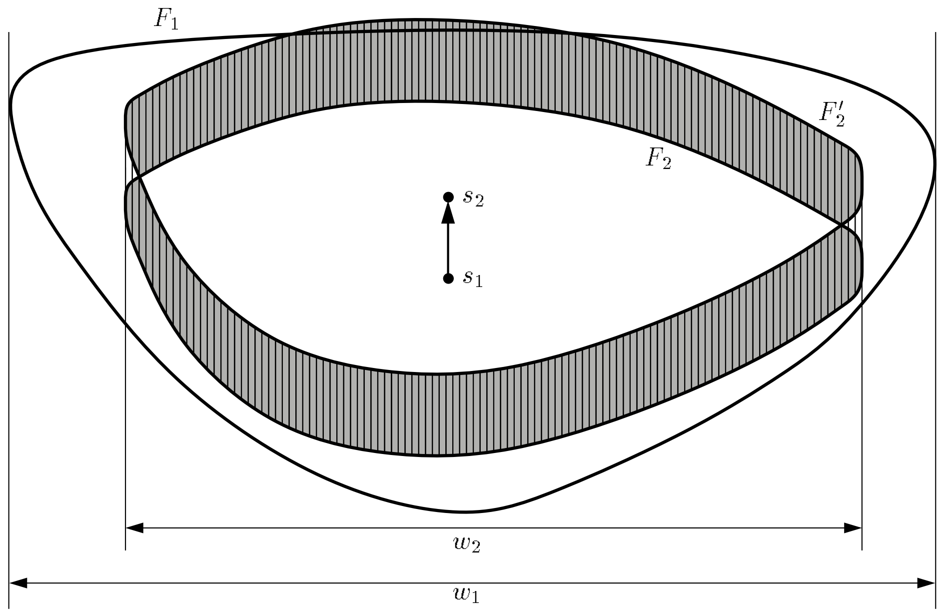

Figure 2: Difference between optimal and heuristic position

The symmetric difference $F_2' \triangle F_2$ is contained in the area that is “swept” by the boundary of F_2 during the translation. By Cavalieri's principle, this area is bounded by twice the length of the translation vector $s_1 - s_2$ times the width w_2 of the projection of F_2 onto a line normal to this vector. So we have

$$\begin{aligned}
\delta^{\mathrm{C}}(F_1, F_2) &= F_1 \triangle F_2' \\
&\leq (F_1 \triangle F_2) + (F_2 \triangle F_2') && \text{by (1)} \\
&\leq (F_1 \triangle F_2) + 2 \cdot |s_1 - s_2| \cdot w_2 \\
&\leq (F_1 \triangle F_2) + 2 \cdot |s_1 - s_2| \cdot w_1 \\
&\leq (F_1 \triangle F_2) + 2 \cdot \frac{4}{3} \cdot (F_1 - F_2) && \text{by Lemma 1} \\
&= \frac{11}{3} \cdot \delta^{\mathrm{opt}}(F_1, F_2),
\end{aligned}$$

where w_1 denotes the width of the projection of F_1 onto a line normal to $s_1 - s_2$. If $F_2 - F_1 > 0$, we even get strict inequality from Lemma 1.

Let's now consider the general case. Assume that F_1, F_2 are in optimal position and let $I = F_1 \cap F_2$. Applying (1) to translates of F_1, F_2 and I with coinciding centroids we get

$$\begin{aligned}
\delta^{\mathrm{C}}(F_1, F_2) &\leq \delta^{\mathrm{C}}(F_1, I) + \delta^{\mathrm{C}}(F_2, I) \\
&\leq \frac{11}{3} \cdot \delta^{\mathrm{opt}}(F_1, I) + \frac{11}{3} \cdot \delta^{\mathrm{opt}}(F_2, I) && \text{by the first case} \\
&= \frac{11}{3} \cdot \delta^{\mathrm{opt}}(F_1, F_2),
\end{aligned}$$

which proves the inequality of the theorem. Again, we get strict inequality whenever $F_1 \triangle F_2 > 0$.

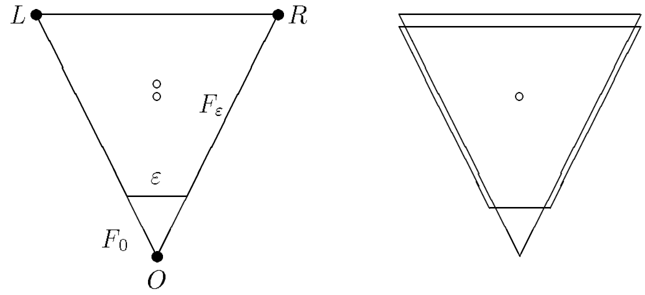

Figure 3: $\delta^{\mathrm{opt}}(F_0, F_\varepsilon)$ and $\delta^{\mathrm{C}}(F_0, F_\varepsilon)$

To see that the approximation factor 11/3 is best possible, we construct an example where this factor can be approached arbitrarily closely. In fact, such an example can be found by examining the proof of Lemma 1. In order to get the ratio between the sides of the inequality in (2) as small as possible. we must choose h and ε very small. For example, we may take $h = 0$, but ε must be positive. So we take an isosceles triangle $F_0 = LRO$ whose base and the corresponding height have unit length. For $\varepsilon > 0$ denote by F_ε the trapezoid obtained from F_0 by cutting off a tip of height ε (Figure 3). Clearly,

$$\delta^{\mathrm{opt}}(F_0, F_\varepsilon) = \varepsilon^2/2. \tag{3}$$

The cut moves the centroid by $\frac{2}{3}\frac{\varepsilon^2}{1+\varepsilon}$ towards the base. The area of the symmetric difference of translates of F_0 and F_ε with the same centroid is shown in Fig. 3. A straightforward computation yields

$$\delta^{\mathrm{C}}(F_0, F_\varepsilon) = \varepsilon^2 \cdot \frac{33 + 18\varepsilon - 11\varepsilon^2}{18(1+\varepsilon)^2}. \tag{4}$$

It follows from (3) and (4) that

$$\lim_{\varepsilon \to 0} \frac{\delta^{\mathrm{C}}(F_0, F_\varepsilon)}{\delta^{\mathrm{opt}}(F_0, F_\varepsilon)} = \frac{11}{3}. \qquad \square$$

3 Transformations other than translations

In many applications, more general matching transformations than just translations are considered. These include, for example,

- *rigid motions*, i.e., combinations of translations and rotations;
- rigid motions where only a restricted set of rotations is allowed;
- (positive) *homotheties*, i.e., mappings of the form $x \mapsto a + \lambda(x - a)$, for some fixed scaling factor $\lambda \geq 0$ and some fixed center $a \in \mathbb{R}^2$; this allows scaling and translation but no rotation;
- *similarities*, i.e., combinations of homotheties and rigid motions;
- arbitrary *affine mappings*.

We propose the *centroid heuristic* for finding approximate solutions to the shape matching problem for convex sets. This approach considers only those transformations in $\mathcal{T}$ that map the centroid of F_2 onto the centroid of F_1. Let t^{C} be an optimal transformation of this kind, and denote $\delta^{\mathrm{C}}(F_1, F_2) = \delta(F_1, t^{\mathrm{C}}(F_2))$ and $\delta^{\mathrm{opt}}(F_1, F_2) = \delta(F_1, t^{\mathrm{opt}}(F_2))$. Showing that the centroid is a reference point for $\mathcal{T}$ amounts to proving that there is a constant $c \geq 1$ such that

$$\delta^{\mathrm{C}}(F_1, F_2) \leq c \cdot \delta^{\mathrm{opt}}(F_1, F_2).$$

We denote the centroid of a set F by s_F.

Theorem 2 *If a set of transformations $\mathcal{T}$ has the following properties:*

(i) Equivariance *with respect to the centroid:* $t(s_F) = s_{t(F)}$ *for all convex figures* F *and all* $t \in \mathcal{T}$.

(ii) $\mathcal{T}$ *is closed under compositions with translations.*

Then the centroid is a reference point for $\mathcal{T}$ *with respect to the class of convex figures, with approximation factor* $c = 11/3$.

Proof. Let F_1, F_2 be two figures and $F_2' := t^{\mathrm{opt}}(F_2)$. Translate F_2' so that the resulting figure F_2'' has the same centroid as F_1. Then, by Theorem 1,

$$\delta(F_1, F_2'') \leq c \cdot \delta(F_1, F_2') = c \cdot \delta^{\mathrm{opt}}(F_1, F_2).$$

We have $F_2'' = t'(F_2)$ for a transformation t' which is a composition of t^{opt} and a translation. By condition (ii), $t' \in \mathcal{T}$, and by condition (i), $t'(s_{F_2}) = s_{F_2''} = s_{F_1}$. Since t^{C} is the optimal match of F_1 and F_2 under transformations satisfying this condition, we have

$$\delta^{\mathrm{C}}(F_1, F_2) \leq \delta(F_1, t'(F_2)) = \delta(F_1, F_2'') \leq c \cdot \delta^{\mathrm{opt}}(F_1, F_2). \qquad \square$$

A class of transformations for which condition (i) of Theorem 2 does not hold is the set of projective transformations. However, all sets of transformations mentioned at the beginning of this section satisfy conditions (i) and (ii) of Theorem 2. So for all these sets of transformations we obtain a simplified matching problem, whose optimal solution is an approximate solution for the original problem. Since the number of degrees of freedom in the simplified matching problem is reduced by 2, this problem is hopefully easier to solve. We will consider algorithms based on this idea in the next section.

4 Algorithms

The results of the previous sections can be used to design efficient approximate matching algorithms for convex polygons under various sets of transformations. These algorithms will produce a solution which is at most by a factor of 11/3 worse than the optimal one. Throughout this section we assume that we are given two convex polygons F_1 and F_2 (by a sorted list of their vertices) which are to be matched. Let n be the total number of vertices of F_1 and F_2.

4.1 Translations

In the case of *translations* we just have to compute the centroids s_1 and s_2 and then to translate F_2 by the vector $s_1 - s_2$. The centroids can easily be computed in linear time, for example by triangulating each figure, determining the centroids and areas of all triangles, and then determining the total centroid as the weighted sum of the triangle centroids.

This gives a matching algorithm of runtime $O(n)$. As far as asymptotic runtime is concerned, this is not too big an improvement over the algorithm of de Berg et al. [BDK+96] which computes the *optimal* match under translations in $O(n \log n)$ time. But our algorithm may be a viable alternative in practice since it is much simpler.

Usually, after the two figures have been matched, one also wants to compute the resulting area of the symmetric difference. This can be done in linear time in a straightforward way. In fact, the problem is equivalent to computing the area of the intersection $I = F_1 \cap F_2$, since $\delta(F_1, F_2) = F_1 + F_2 - 2I$. The sets F_1, F_2, and I are convex polygons, and I can be computed from the sorted lists of vertices of F_1 and F_2 in linear time.

4.2 Homotheties

According to Theorem 2 we get an 11/3-approximate solution by first computing the two centroids s_1 and s_2, then translating F_2 by $s_1 - s_2$ obtaining F_2', and finally stretching F_2' about s_1 by a factor λ minimizing the symmetric difference.

It remains to explain the last step. Suppose w.l.o.g. that s_1 is the origin, and denote by

$$q(\lambda) := \delta(F_1, \lambda F_2')$$

the function that we want to minimize.

Lemma 2 *Suppose that F_1 and F_2 are convex polygons with a total number of n vertices. Then the function $q\colon \mathbb{R}^+ \to \mathbb{R}^+$ is a piecewise quadratic function with at most $2n+1$ quadratic pieces. When viewed as a function of $A = \lambda^2$, the function $\bar{q}(A) := q(\sqrt{A})$ is convex. The function $\bar{q}$, and hence also q, has a unique local minimum, which is the global minimum.*

Proof. If we draw a ray from the origin through each vertex of F_1 and F_2', we partition the plane into at most n wedges. Within a wedge W, the boundary of each of the two figures consists of a line segment. Now, suppose that F_2' is stretched by the factor λ which is 0 in the beginning and is then continuously increased. We obtain the successive configurations a), b), and c) of the edges e_1 of F_1 and e_2 of $\lambda F_2'$ shown in Figure 4. In each case the symmetric difference within W is

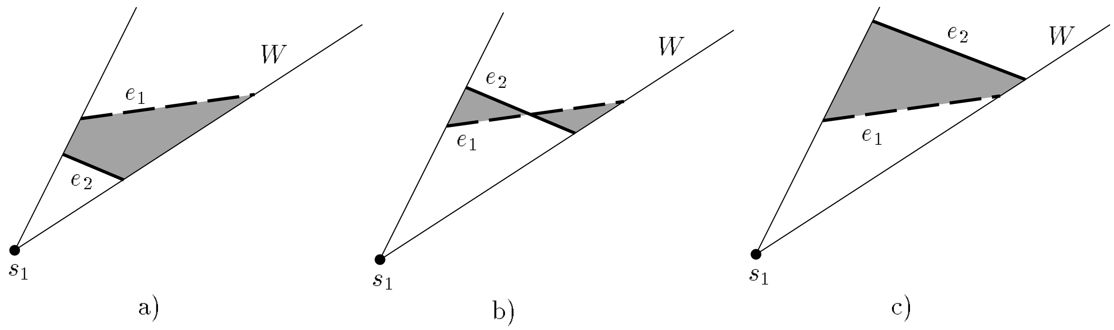

Figure 4: Configurations of two edges inside a wedge

a quadratic function. The symmetric difference within the i-th wedge, which we denote by $q_i(\lambda)$, is thus a piecewise quadratic function with three quadratic pieces. The total symmetric difference $q(\lambda)$ is the sum of the n functions $q_i(\lambda)$. It is piecewise quadratic with $2n$ breakpoints.

We want to show that each function q_i is convex when considered as a function of $A = \lambda^2$. The parameter A is proportional to the area of F_2' and to the area of F_2' inside the wedge W. Therefore, the first and third piece of $\bar{q}_i(A) := q_i(\sqrt{A})$ are *linear* functions of the form $|\text{area}(F_1 \cap W) - A \cdot \text{area}(F_2' \cap W)|$. The first piece is a strictly decreasing function, and the third piece is a strictly increasing function. These two properties hold also for the first and third piece of q_i. If the two edges e_1 and e_2 are parallel, the second piece is missing. Otherwise, the second piece of q_i is a quadratic function which smoothly joins the first piece, which is strictly decreasing, to the second piece, which is strictly increasing. If follows that the second piece is strictly convex, positive, with a unique local minimum C_3 inside its domain of validity. Hence, the second piece of q_i can be written in the form $q_i(\lambda) = C_1 + C_2(\lambda - C_3)^2$, with positive constants C_1, C_2, C_3. The second piece of $\bar{q}_i$ takes the form $\bar{q}_i(A) = q_i(\sqrt{A}) = C_1 + C_2(\sqrt{A} - C_3)^2 = D_1 + D_2 A - D_3\sqrt{A}$ with positive constants D_1, D_2, D_3. This function is strictly convex for $A \geq 0$.

Summarizing, $\bar{q}$, as a sum of the convex functions $\bar{q}_i$, is a convex function of A. To prove that the minimum A^* is unique, we must exclude the possibility that A^* occurs at a point where $\bar{q}$ is linear. But then all functions $\bar{q}_i$ must be linear at A^*. It is impossible that A^* lies in the first part of all functions $\bar{q}_i$, because then $\bar{q}$ would be strictly decreasing at A^*. Similarly, A^* cannot lie in the third part of all functions $\bar{q}_i$. However, if A^* lies in the first part of some functions $\bar{q}_i$ and in the third part of some other functions $\bar{q}_i$, this means that in some wedges, λF_2 lies completely outside F_1, whereas in other wedges λF_2 lies completely inside F_1. But then there must be some wedges where the boundaries of F_1 and F_2 cross, and therefore $\bar{q}$ is strictly convex. □

We remark that, as the proof shows, the functions q and $\bar{q}$ are even continuously differentiable unless some edge of F_1 is parallel to an edge of F_2.

We have thus established that q is a very well-behaved function for optimization. The minimizer λ^* of the function q can be determined in $O(n)$ time by the prune-and-search technique: We search for the quadratic piece in which λ^* lies by performing a binary search among the $2n$ breakpoints, successively narrowing down the interval $[\lambda_0, \lambda_1]$ in which λ^* is known to lie. The decision whether λ^* is bigger or smaller than the current decision point λ depends just on the sign of the derivative $q'(\lambda)$ at this point. When the interval $[\lambda_0, \lambda_1]$ contains only k breakpoints, there are at most k functions q_i for which the definition changes inside the interval; the remaining functions are plain quadratic functions and their sum can be accumulated in one quadratic function. This means that $q(\lambda)$ and the derivative $q'(\lambda)$ can be evaluated in $O(k)$ time. The next trial value for the binary search is the median of the k remaining breakpoints and it can also be computed in $O(k)$ steps. This reduces k by a factor of 2. Thus, in $O(n)$ time the interval in which λ^* must lie is narrowed down to one quadratic piece of the function q. The optimum λ^* is then found by solving the linear equation $q'(\lambda^*) = 0$.

Summarizing the results of the last two subsections, we have:

Theorem 3 *For the shape matching problem for two convex polygons with a total number of n vertices, with respect to*

1. *the set of translations, or*

2. *the set of homotheties,*

an 11/3-approximate solution can be found in $O(n)$ time.

We remark that the technique of this subsection can also be used when the allowable transformations $\mathcal{T}$ consist only of translations in a given fixed direction d. This is an optimization problem with one degree of freedom, just like the problem for homotheties after the scaling center is fixed. Instead of the wedges, we consider strips formed by the lines parallel to d through every vertex of F_1 and F_2. The translation minimizing δ can be found in linear time, since, in the range where the intersection I is nonempty, $\sqrt{I}$ is a concave function of the translation vector. Another linear-time algorithm for this task is described in Avis et al. [ABS$^+$96], where the problem is approached in a more indirect way.

4.3 Rigid Motions

As in the previous case we will first perform a translation such that the centroids of F_1 and F_2 coincide. We will assume for simplicity that F_1 and F_2 already have their common centroid at the origin O. Now we have to rotate F_2 around O by some angle φ in order to minimize the symmetric difference. We denote by $F_2' = t_\varphi(F_2)$ the rotated copy of F_2, and by $q(\varphi)$ the symmetric difference $\delta(F_1, t_\varphi(F_2))$ as a function of φ. It is possible to work out an expression for $q(\varphi)$ in terms of φ, but since we are interested in the minimum, we will only describe how the derivative $q'(\varphi)$ can be computed.

Lemma 3 *The function $q\colon \mathbb{R} \to \mathbb{R}^+$ is continuous; it is continuously differentiable except when a vertex of $F_2' := t_\varphi(F_2)$ lies on an edge of F_1 or vice versa. The derivative $q'(\varphi)$ can be computed as follows. Let $I = F_1 \cap F_2'$, and let $P_1, \dots, P_{2m}$ $(m \geq 0)$ be a sequence of* crossing points *between F_1 and F_2' in the following sense: between P_{2i-1} and P_{2i}, the boundary of I is formed by the boundary of F_2'; and between P_{2i} and P_{2i+1} (and between P_{2m} and P_1), the boundary of I is formed by the boundary of F_1. (If $m = 0$, then one of the sets F_1, F_2' is contained in the other.)*

Then $q'(\varphi)$ is the alternating sum of squared distances of the points $P_1, \dots, P_{2m}$ from O:

$$q'(\varphi) = \sum_{i=1}^{2m} (-1)^i \cdot \overline{OP_i}^2 \tag{5}$$

Proof. Instead of the symmetric difference we may equally well consider the area of the intersection I, since $q(\varphi) = F_1 + F_2 - 2I$, and so $q'(\varphi) = -2 \cdot \frac{\partial I}{\partial \varphi}$. Let us consider the change of I in the vicinity of a point P_i when φ changes to $\varphi + \varepsilon > \varphi$ (Figure 5). Assume that the part of I's boundary that is formed by F_2' lies to the left of P_i and thus moves away from the part that is formed by F_1. We see that, between these two parts, a small wedge-like quadrilateral is inserted, which has area

$$\varepsilon \cdot \overline{OP_i}^2/2 + O(\varepsilon^2).$$

(The area of a circular sector of radius $\overline{OP_i}$ and angle ε would be $\varepsilon \cdot \overline{OP_i}^2/2$. But in fact, the distance from O to the boundary of I inside the sector of interest differs from the "ideal" radius $\overline{OP_i}$ by $O(\varepsilon)$. This accounts for the error term $O(\varepsilon^2)$.) If i has different parity, then an analogous area of the same size is *subtracted* from I. Summing up the different contributions, dividing by ε, and taking the limit $\varepsilon \to 0$ gives (5). □

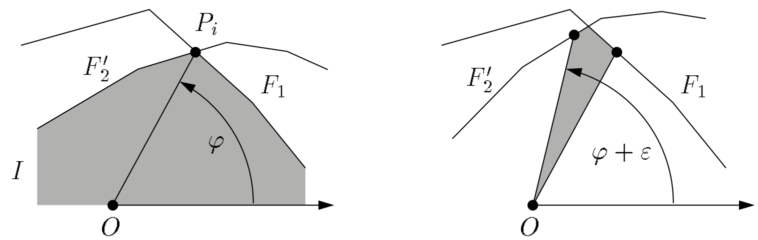

Figure 5: The derivative of the symmetric difference under rotations

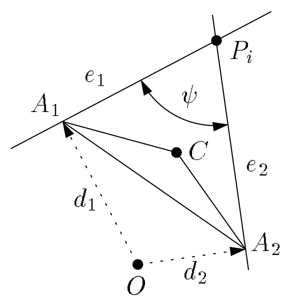

Figure 6: Computing the distance $\overline{OP_i}$

From the lemma it follows that the minimum of $q(\varphi)$ is either one of the points where q is not differentiable, or one of the stationary points where $q'(\varphi) = 0$. There are up to $O(n^2)$ critical points of nondifferentiability: each vertex of F_1 can lie on a particular edge of F_2' for at most two values of φ, and vice versa. These critical points are also the points where the points P_i may change: there are $O(n^2)$ intervals such that, inside an interval, every point P_i is given as the intersection of a fixed edge of F_1 with a fixed edge of F_2', and thus $q(\varphi)$ has a fixed analytic expression in terms of φ.

Summarizing, we have to check $O(n^2)$ single points plus $O(n^2)$ intervals. For finding the possible candidates for a minimum in each interval, we have to look for points where the derivative is zero, i.e., for solutions of $q'(\varphi) = 0$, where $q'(\varphi)$ is given by (5). Let us discuss how the equation $q'(\varphi)$ can be solved. Suppose that P_i is determined as the intersection of two edges e_1 and e_2. Then the distance $\overline{OP_i}$ can be expressed as follows. Let ψ denote the angle between e_1 and e_2, and let d_1, d_2 denote the distance from e_1, e_2, respectively, to the origin O (Figure 6). Then we have

$$\overline{OP_i}^2 = \frac{d_1^2 + d_2^2 + 2d_1 d_2 \cos\psi}{\sin^2\psi}.$$

(This can be checked by observing that the quadrilateral $OA_2P_iA_1$ is inscribed in a circle with diameter OP_i and center $C = (O + P_i)/2$. The distance $\overline{A_1A_2}$ can be expressed in terms of d_1, d_2, and the angle $A_2OA_1 = \pi - \psi$, and is can also be computed from the isosceles triangle A_1A_2C with sides $\overline{A_1C} = \overline{A_2C} = \overline{OP_i}/2$ and angle $A_1CA_2 = 2\psi$ or $A_1CA_2 = 2\pi - 2\psi$, as appropriate.) When φ varies and F_2' rotates, d_1 and d_2 remain fixed, but ψ must be substituted by $\psi_0 \pm \varphi$, for some fixed ψ_0.

It is convenient to use $t = \tan(\varphi/2)$ as a parameter instead of φ. Then $\sin\varphi = 2t/(1+t^2)$ and $\cos\varphi = (1-t^2)/(1+t^2)$, so $\overline{OP_i}^2$ can be written as a rational function with bounded degree. The expression $q'(\varphi)$ is the sum of at most n such functions. Thus, inside each interval, $q'(\varphi) = 0$ can be solved as the root of a polynomial of degree $O(n)$, and $q(\varphi)$ has at most $O(n)$ local minima.

Altogether, this gives $O(n^3)$ candidates for the minimum. In principle, these candidates can be found exactly, since only computations with algebraic numbers are required. However, this would be very expensive. A more practical approach would be to solve the polynomial equations numerically. This still involves highly nontrivial numerical problems, whoses detailed investigation goes beyond the scope of this paper. For an intensive treatment of this problem with respect to bit complexity see [S82].

Even just computing $q(\varphi)$ at all candidate points (assuming they are given to us) would take $O(n^4)$ arithmetic operations. It does not make sense to spend so much time for finding the optimal rotation in the present context, since the computed value is only a rough approximation to the overall problem. An algorithm that computes a good approximation to the optimal rotation within reasonable time bounds is called for.

5 Conclusion and open problems

As mentioned in the introduction, de Berg et al. [BDK+96] showed that, when the centroids of two convex sets F_1 and F_2 coincide, the area of their intersection is at least $9/25$ of the maximum area that can be obtained by translating F_2. As in Section 3, this result extends to more general sets of transformations.

Let us briefly relate the bound of Theorem 1 to this result. This bound is more powerful than our result when the area of the intersection is relatively small and the area of the symmetric difference is relatively large. For example, when $F_1 \leq \frac{4}{7}F_2$, our bound is worthless: $\delta(F_1, t(F_2)) \geq F_2 - F_1 \geq \frac{3}{7}F_2$, and if we multiply this by $11/3$, the bound on $\delta^{\mathrm{C}}(F_1, F_2)$ that we get is larger than the trivial bound of $F_1 + F_2$ that we get by placing F_2 anywhere. In contrast to this situation, the bound of de Berg et al. makes a nontrivial statement in any case. On the other hand, Theorem 1 gives the strongest statement when the sets F_1 and F_2 have a very similar shape and δ is small, as in the cases of practical interest for pattern matching. One may check that Theorem 1 gives a stronger bound precisely if $\delta^{\mathrm{opt}}(F_1, F_2) < \frac{6}{31}(F_1 + F_2)$.

It is not known whether the fraction $9/25$ in the mentioned bound is best possible. The correct number should probably be $4/9$. (There is an example showing that the factor cannot be improved beyond $4/9$.) The proof of the $9/25$ bound uses an ingenious representation of the centroid. This technique also yields results in all higher dimensions.

The extension of Theorem 1 to higher dimensions has not been considered so far. A direct generalization of the proof of Lemma 1 to three dimensions is not possible, because there is no way to ensure that the analogous operation to Steiner symmetrization leaves the "width" w, i.e., the area of the vertical projection, unchanged. By taking into account the loss that occurs in this operation, we can show that the centroid has an approximation factor of at most $33\sqrt{3}/(8\pi)$ in three dimensions, but this bound is not tight.

Our heuristic is guaranteed to find a translation which reduces the symmetric difference between two given convex figures to within a factor of $11/3$ of the optimum. It would be nice to have a simple method for getting a solution with a better approximation guarantee, using the heuristic solution as a starting point. Techniques which have proved to be useful in similar situations include (a) testing all vectors in a sufficiently fine grid around the starting point and (b) applying the ellipsoid method for convex optimization problems. (Recall that, in the range where the intersection I is nonempty, $\sqrt{I}$ is a concave function of the translation vector.) However, one needs some a-priori knowledge about the region in which the optimal solution can lie in order to apply these methods. This question is open to further research.

As discussed in the last section, the approximate shape matching problem under rigid motions is not solved in a satisfactory way. If a good approximation algorithm for shape matching under rotations were available, it could be combined with the technique of superimposing centroids to give an approximate algorithm for rigid motions.

Acknowledgements

We would like to thank Stefan Felsner and Emo Welzl for helpful discussions and hints.

References

[AAR94] H. Alt, O. Aichholzer, and G. Rote. Matching shapes with a reference point. In *Proc. 10th Ann. Sympos. Comput. Geom.*, pp. 85–92, 1994.

[AAR97] H. Alt, O. Aichholzer, and G. Rote. Matching shapes with a reference point. *Int. J. Comput. Geom. Appl.* **7** (1997), 349–363.

[ABB91] H. Alt, B. Behrends, and J. Blömer. Approximate matching of polygonal shapes. In *Proc. 7th Ann. Sympos. Comput. Geom.*, pp. 186–193, 1991.

[ABGW90] H. Alt, J. Blömer, M. Godau, and H. Wagener. Approximation of convex polygons. In *Proc. 17th Internat. Colloq. Automata Lang. Program.*, *Lecture Notes in Computer Science*, Vol. 443, pp. 703–716. Springer-Verlag, 1990.

[AFRW96] H. Alt, U. Fuchs, G. Rote, and G. Weber. Matching convex shapes with respect to the symmetric difference. In *Algorithms — ESA '96*, Proc. Fourth Annual European Symposium on Algorithms, Barcelona, September 25–27, 1996, ed. J. Díaz and M. Serna, *Lecture Notes in Computer Science*, Vol. 1136, pp. 320–333. Springer-Verlag, 1996.

[AST94] P. K. Agarwal, M. Sharir, and S. Toledo. Applications of parametric searching in geometric optimization. *J. Algorithms* **17** (1994), 292–318.

[ABS+96] D. Avis, P. Bose, T. Shermer, J. Snoeyink, G. Toussaint, and B. Zhu. On the sectional area of convex polytopes. In *Proc. 12th Ann. Sympos. Comput. Geom.*, pp. C11–C12, 1996.

[BDK+96] M. de Berg, O. Devillers, M. van Kreveld, O. Schwarzkopf, and M. Teillaud. Computing the maximum overlap of two convex polygons under translations. In *Proc. 7th Internat. Sympos. Algorithms and Computation (ISAAC'96)*, *Lecture Notes in Computer Science*, Vol. 1178, pp. 126–135. Springer-Verlag, 1996.

[BF33] T. Bonnesen and W. Fenchel. *Theorie der konvexen Körper.* Ergebnisse der Mathematik und ihrer Grenzgebiete, Vol. 3. Springer-Verlag, Berlin 1933; repr. Chelsea Publ. Comp., New York 1948, 1971.

[CGH+93] L. P. Chew, M. T. Goodrich, D. P. Huttenlocher, K. Kedem, J. M. Kleinberg, and D. Kravets. Geometric pattern matching under Euclidean motion. In *Proc. 5th Canad. Conf. Comput. Geom.*, pp. 151–156, Waterloo, Canada, 1993.

[G83] P. M. Gruber. Approximation of convex bodies. In: Peter M. Gruber and J. M. Wills (eds.), *Convexity and its Applications.* Birkhäuser, Basel 1983, pp. 131–162.

[G93] P. M. Gruber. Aspects of approximation of convex bodies. In: Peter M. Gruber and J. M. Wills (eds.), *Handbook of Convex Geometry*, Vol. A, Chapter 1.10. North-Holland, Amsterdam 1993, pp. 319–345.

[HKS93] D. P. Huttenlocher, K. Kedem, and M. Sharir. The upper envelope of Voronoi surfaces and its applications. *Discrete Comput. Geom.* **9** (1993), 267–291.

[S82] A. Schönhage. The fundamental theorem of algebra in terms of computational complexity. Technical Report, University of Tübingen, 1982.