SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK 14853

TECHNICAL REPORT NO. 942

November 1990
Revised June 1991

MATCHING POINTS INTO NOISE REGIONS: COMBINATORIAL BOUNDS AND ALGORITHMS

By

Esther M. Arkin
Klara Kedem
Joseph S. B. Mitchell
Josef Sprinzak
Michael Werman

Extended abstract appears in: *Proc. Second ACM/SIAM Symposium on Discrete Algorithms*, San Francisco, pp. 42-51, January, 1991.

Matching Points into Noise Regions: Combinatorial Bounds and Algorithms *

Esther M. Arkin[(1)] Klara Kedem[(2)] Joseph S. B. Mitchell[(1)]
Josef Sprinzak[(3)] Michael Werman[(3)]

[(1)]School of Operations Research and Industrial Engineering
Cornell University, Ithaca, NY 14853
[(2)]Department of Computer Science
Tel-Aviv University, Tel-Aviv, Israel
[(3)]Department of Computer Science
The Hebrew University, Jerusalem, Israel

June 6, 1991

Abstract

We consider several cases of the *point matching problem* in which we are to find a transformation of a set of n points such that each transformed point lies in one of n given pairwise-disjoint "noise regions". We prove upper and lower bounds on the number of possible matches, under a variety of types of transformations (rotations, translations, similarity, etc.) and noise regions (circles, squares, polygons, etc.). We also give efficient algorithms for computing the set of all possible matches, along with a corresponding transformation that realizes each match.

An extended abstract of this paper appears in: *Proc. Second Annual ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, CA, January 28-30, 1991, pp. 42–51.

*E. Arkin is partially supported by NSF grants DMS-8903304, DMC-8451984, and ECSE-8857642. K. Kedem is partially supported by a fellowship from the Pikkovski-Valazzi Fund and by the Fund for Basic Research administered by the Israeli Academy of Sciences. J. Mitchell is partially supported by NSF grant ECSE-8857642, and by a grant from Hughes Research Laboratories.

1 Introduction

A fundamental problem in model-based computer vision is to design efficient algorithms to determine to what extent a set of *image features* "matches" a set of *model features*. Frequently, this problem can be formulated as a point matching problem: Given a set of n *image* points $\mathcal{A} = \{a_1, \ldots, a_n\} \subset \Re^d$ and a set of m *model points* $\mathcal{B} = \{b_1, \ldots, b_m\} \subset \Re^{d'}$, determine a *matching* μ (i.e., a list of pairs (a_i, b_j) such that no two pairs share the same first element or the same second element) and a transformation $\tau : \Re^d \to \Re^{d'}$, within an allowable class of transformations $\mathcal{T}$, such that the application of τ to point a_i brings it into "correspondence" to point b_j, for each pair $(a_i, b_j) \in \mu$. The "value" of a matching can be taken to be the number of pairs (a_i, b_j), or possibly a sum of weights.

The term "correspondence" can take on several different meanings. In the *exact point matching problem* (also known as the "image registration problem"), we require that $\tau(a_i) = b_j$ for every pair $(a_i, b_j) \in \mu$ of the matching. In the *inexact point matching problem* (also known as the "approximate congruence problem"), we only require that $\tau(a_i)$ be *close* to b_j, for each $(a_i, b_j) \in \mu$. A natural definition of closeness is to define for each model point b_j, a "noise region" B_j, and to say that $\tau(a_i)$ is "close" to b_j if $\tau(a_i) \in B_j$. We let $\mathcal{R} = \{B_1, \ldots, B_m\}$ denote the set of noise regions.

The exact point matching problem has been solved in time $O(n^{d-2} \log n)$ for $d = d'$ and $\mathcal{T}$ the set of congruences (translations and rotations); see [AMWW], earlier work by [Ata] and [Atk], and the recent work of [SW].

Closely related to the exact point matching problem is the problem of computing the "distance" between two point sets, under some appropriate definition of distance. For example, the problem of finding the best translation to minimize Hausdorf distance between two points sets has been recently solved in time $O(n^3)$ by [HK], [HKS]. The problem of finding a "least squares" registration between point sets has been addressed by [Zi].

In this paper, we consider the inexact point matching problem. Related work on this subject includes [AMWW], [AMZ], [Ba], [He], [ISI], [Sc], [Sp], etc. We focus on several variations of the inexact point matching problem, as specified by various parameters:

(1) The set of allowed transformations $\mathcal{T}$ may be *pure translation*, *pure rotation* (about a known center), *translation plus rotation* (congruence), *translation plus rotation and change of scale* (similarity), general affine transformations, etc. We denote these cases by writing $\mathcal{T} =$ T, R, TR, TRS, etc.

(2) We work almost exclusively in the case of two dimensions ($d = d' = 2$). Many of our techniques may generalize to higher dimensions.

(3) We focus on the equal cardinality case ($m = n$). Our methods also give results for the case $m \neq n$, with combinatorial bounds increasing in the difference $|m - n|$, these are briefly discussed in the conclusion.

(4) The regions B_j may be unit circles, arbitrary circles, aligned unit squares, arbitrary unit squares, identical convex polygons, arbitrary convex polygons, "normal" shapes

(defined below), or arbitrary polygons. Let N denote the total number of vertices of all polygons B_j.

(5) Various assumptions about separation of the regions B_j are considered. Previous results ([AMWW], [AMZ]) allowed for arbitrary overlap among the regions B_j, yielding in many cases very high (polynomial) bounds. In this paper we concentrate on the case in which the regions are pairwise-disjoint, and we thereby obtain much tighter bounds on the number of matchings and on the time bounds to compute matchings. We also consider stronger notions of separation: linearly separable regions and ϵ-separated regions. As might be expected, we observe the following phenomenon: the stronger the degree of separation among the B_j's, the lower the exponents in our polynomial bounds.

Since this paper concentrates on the case in which $m = n$, our interest is in finding "perfect" matchings between points a_i and regions B_j. We abbreviate the term "perfect matching" with the term "match".

In all of our problems, we are interested in both combinatorial bounds on the number of possible matches, and in time bounds on the complexity of algorithms to compute the matches. All of our algorithms have the feature that they find *all* matches (and we know of no methods to find a single match any faster in the worst case than finding all matches). To find a match means to produce a matching μ and to produce a witness transformation τ that achieves the matching: $\tau(a_i) \in B_j$ for each $(a_i, b_j) \in \mu$. Our results are summarized in the table of Figure 1. (An arrow $\longleftarrow$ ($\longrightarrow$) means that the bound in the entry to the right (left) holds.)

We now establish some basic notation and terminology. We let $a^* = \frac{1}{n}\sum_i a_i$ denote the center of mass of the points $\mathcal{A}$. Similarly, b^* denotes the center of mass of the points $\mathcal{B}$ (which frequently will be the centers of circles $\mathcal{R}$). Unless otherwise indicated, we assume (without loss of generality) that a_1 is the farthest point of $\mathcal{A}$ from a^*.

We define $\beta_\epsilon(p)$ to be the disk (ball) of radius ϵ centered at p. The metric will be clear from the context. We say that a set $\mathcal{R} = \{B_1, \ldots, B_n\}$ of regions is *ϵ-separated* if there exists an $\epsilon > 0$ such that no disk of radius ϵ intersects two distinct regions B_i and B_j. We say that a set of shapes is *normal* if the convex hulls of the shapes have bounded aspect ratios (meaning that they are "fat") and the ratio of the diameter of the largest shape to the diameter of the smallest shape is also bounded.

When we speak of the pure rotation problem, we will assume without loss of generality that the known center of rotation is the origin. We define the *orbit* of point a_i to be the circle centered at the origin obtained by rotating the point a_i about the origin.

The remainder of this paper is organized as follows: In Section 2 we discuss lower bounds on the number of matches under translation, and pure rotation. In Section 3 we give corresponding upper bounds on the number of matches. Section 4 describes algorithms for finding matches. We conclude in Section 5 with extensions and open problems.

Trans. $\mathcal{T}$	Unit circles	Circles	Aligned unit sq.	Arbitrary unit sq.	Normal shapes	Linearly separable
T	$1^{(*)}$ $1^{(*)}$ $O(n \log n)^{(*)}$	1 1	$1^{(*)}$ $1^{(*)}$	1 1	1 1	1 1 $O(n^2)^{(\dagger)}$
R	n n $O(n^2)$	n n $O(n^2 \log n)$	$5n - 16$ $5n$ $\longleftarrow$	n^2 $\longleftarrow$ $O(n^2 \log n)$	$\longrightarrow$ $\longleftarrow$ $\longleftarrow$	$\longrightarrow$ n^2 $O(nN \log N)$
TR	n $O(n^2)$ $O(n^4 \log n)$	$\longrightarrow$	n $\longleftarrow$ $\longleftarrow$	$\longrightarrow$ $O(n^4)$ $O(n^4 \log n)$	$\longrightarrow$ $O(nN^3)$ $O(nN^3 \log N)$	$O(N^6 \log N)^{(\dagger)}$
TRS	n $O(n^2)$ $O(n^5 \log n)$		n $\longleftarrow$ $\longleftarrow$	$\longrightarrow$ $O(n^5)$ $O(n^5 \log n)$	$O(nN^4)$ $O(nN^4 \log N)$	

Figure 1: Summary of results. Each entry consists of three lines: (1) lower bound on the number of matches; (2) upper bound on the number of matches; and (3) time complexity of an algorithm to find all matches. ($^{(*)}$[AMWW], $^{(\dagger)}$[Sp])

2 Lower Bounds on the Number of Matchings

In this section we describe the construction of lower bounds of various matching problems. [AMZ] and [Sp] have given lower bounds for the case of matching points to noise regions, where the two sets are of unequal cardinality ($m \neq n$) and/or the noise regions may overlap. The case we discuss here is that of equal cardinality ($m = n$) and of pairwise-disjoint noise regions.

For the case of pure translation ($\mathcal{T}$ = T), [AMWW] showed that there is at most one match when the noise regions are unit circles. (In the next section we generalize this for other noise regions that are linearly separable.) Thus, the trivial lower bound of 1 is a tight bound in this case.

For the case of pure rotation, we get another simple lower bound:

Fact 1 *In the case of pure rotation ($\mathcal{T}$ = R) and arbitrary pairwise-disjoint noise regions $\mathcal{R}$, there can be n distinct matches.*

Proof. Let C be a large circle centered at the origin, and let the points a_i be evenly spaced along C. Let the regions B_j be placed so that region B_i contains point a_i and no other point of $\mathcal{A}$. (This is always possible if the circle C is large enough.) Then, it is clear that there can be n distinct matches under rotation of $\mathcal{A}$. ∎

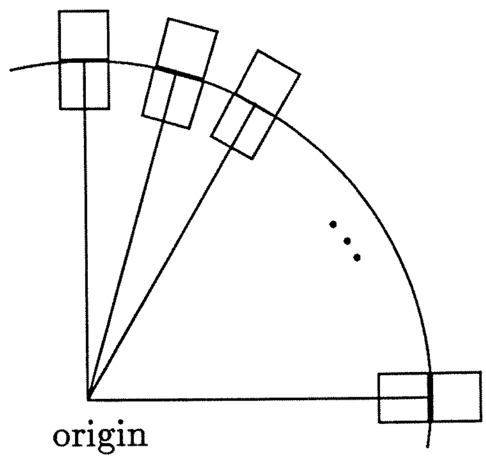

Figure 2: Unaligned unit squares

In particular, the above construction shows a lower bound of n for the number of matches under pure rotation when the noise regions are unit circles. In Theorem 6 we show this bound to be tight (for general circles).

For the case of unit square noise regions, we prove the following:

Theorem 2 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit squares (at arbitrary orientations) $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches may be $\Omega(n^2)$.*

Proof. We describe the construction of an example in which there are a quadratic number of distinct matches between points and unit squares. Figure 2 shows the macroscopic placement of the squares: There are $n/2$ "inner"squares with centers equally distributed about the circle $\beta_r(0)$ of radius r (very large), and there are $n/2$ "outer" squares with centers equally distributed about the slightly larger circle $\beta_{r+1+\epsilon}(0)$, with the inner and outer rings of squares having their centers at the same set of $n/2$ angular coordinates. Furthermore, each of the two squares whose center lies at angular coordinate θ is oriented with two of its sides at angle θ, and there is a tiny gap of $\epsilon > 0$ between the two squares.

The placement of the points $\mathcal{A} = \{a_1, \ldots, a_n\}$ is intricate. We place the points along the circle $\beta_{r'}(0)$, where $r' = r + 1/2 + \epsilon'$ for a small number $\epsilon' > \epsilon$. We pick ϵ and ϵ' so that $\beta_{r'}(0)$ intersects each of the inner squares in two disjoint arcs and each of the outer squares in exactly one arc.

Figure 3 shows a microscopic view of the point placement within the k-th pair of squares, B_{2k} (outer) and B_{2k-1} (inner). Assume that the points $\mathcal{A}$ are being rotated clockwise along $\beta_{r'}(0)$ (whose curvature is exaggerated in the figure). We label the arcs of intersection of $\beta_{r'}(0)$ with the squares as follows: U_kV_k is the first arc of intersection with B_{2k-1}, W_kX_k

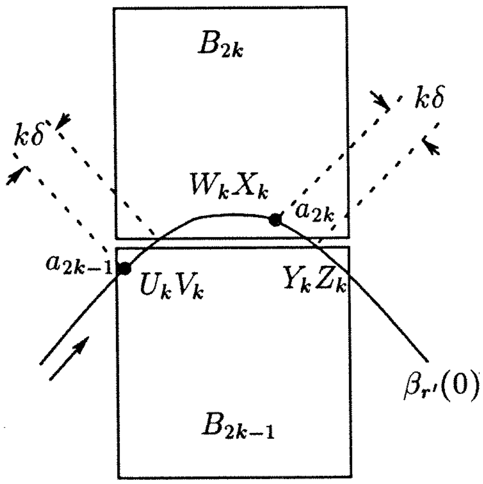

Figure 3: Microscopic placement of points

is the arc of intersection with B_{2k}, and $Y_k Z_k$ is the second arc of intersection with B_{2k-1}. Note that $|U_k V_k| = |Y_k Z_k| < |W_k X_k|$.

For the first pair of squares we place point a_1 in square B_1, on arc $U_1 V_1$, at distance $\delta < |U_1 V_1|/n$ from W_1. We place point a_2 in square B_2, on arc $W_1 X_1$, at distance δ from Y_1. For the k-th pair of points, we place point a_{2k-1} within square B_{2k-1}, on arc $U_k V_k$, at distance $k\delta$ from W_k, and we place point a_{2k} within square B_{2k}, on arc $W_k X_k$, at distance $k\delta$ from Y_k.

A rotation of arc length larger than δ but smaller than 2δ, will cause a_1 to move to B_2 and a_2 to B_1, leaving the matching of all other points unchanged. Another small rotation by δ will cause a_3 and a_4 to switch squares while all other points remain unchanged. Continuing in this way up to $n/2$ times, we see that each small rotation by δ will cause one pair of points to switch the pair of squares they were in, leaving all other points unchanged. Hence a sequence of $n/2$ small "microscopic" motions will yield $n/2$ different matches. If we then rotate the points $\mathcal{A}$ to the next *macroscopic* position (with a_1 within B_3, at distance δ from W_2), we see that the microscopic motions can be repeated, yielding another $n/2$ distinct matches. Since there are $n/2$ macroscopic positions of a_1, each corresponding to $n/2$ microscopic switches, there are a total of $\Omega(n^2)$ distinct matches in all. ∎

Note that in the construction above we needed to have the unit squares be at $\Omega(n)$ different orientations. This raises the question of how many perfect matchings can be attained when the noise regions are *aligned* unit squares (with sides parallel to the axes). The answer is given in the following theorem:

Theorem 3 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n aligned unit squares $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches in the worst case is bounded below by $5n - 16 = \Omega(n)$.*

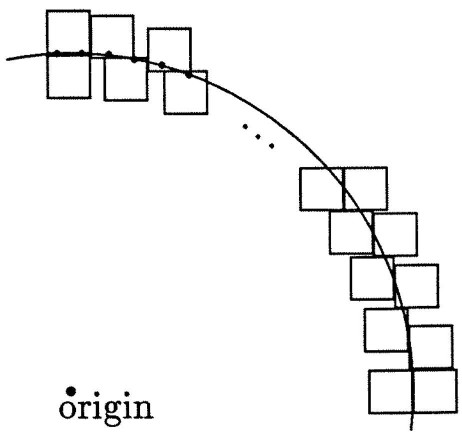

Figure 4: Aligned squares: macroscopic view

Proof. We prove the theorem by the following construction: On each of the four axes (indexed $i = 1, 2, 3, 4$) we place a pair of squares, in a similar manner to that done in the construction of Theorem 2, such that a circle $\beta_{r'}(0)$ of radius r' intersects the inner square in two arcs U_iV_i and Y_iZ_i and the outer square in a single arc W_iX_i. We first describe how we place the points in these eight squares. In the pair of squares intersected by the positive x-axis we place point a_1 on the arc U_1V_1 at distance δ from W_1, and we place point a_2 on arc W_1X_1 at distance δ from Y_1. We choose δ to be small enough so that $4\delta < |W_1X_1|$, and we denote the distance from a_1 to a_2 by $\Delta\theta$. On the pair intersected by the positive y-axis we put a point on the arc U_2V_2 at distance 2δ from W_2, and a point on arc W_2X_2 at distance 2δ from Y_2. Similarly the distances on the negative x-axis are 3δ, and on the negative y-axis 4δ.

The remaining $n - 8$ points are divided, one quarter of the points in each quadrant, such that the points are equally spaced (with spacing $\Delta\theta$) within each quadrant, *except* that the distance from the last point in a quadrant to the next-to-last point is $\Delta\theta + \delta$ (in three cases) or $\Delta\theta - 3\delta$ (in the fourth case). Thus our points are almost equally spaced (as δ is small relative to the distance between points).

We place the remaining $n - 8$ squares such that each point is within a square, at arc distance greater than 4δ, from the boundary of the square. It is easy to see that we can indeed place the squares in such a manner (see Figure 4). The reason is that all squares, other than the ones pierced by the axes, have large intersections with the circle, and thus we are free to have some "slop" in the exact location of the points in these squares.

We now count the number of possible matches in this construction. One match is the original placement which we have just described. We start rotating $\mathcal{A}$ around the origin in a clockwise direction. After rotating by δ', $\delta < \delta' < 2\delta$, the pair of points on the positive x-axis, switches squares, while all the other points remain in their original squares. This yields a second match. Another incremental rotation by δ repeats the process in the squares on the positive y-axis and, similarly, on the negative x- and y-axes. Thus we get 5 matches. This situation repeats itself with each macroscopic rotation of $\mathcal{A}$ by $\Delta\theta$.

The above construction seems to yield $5n$ distinct matches (instead of the looser bound of $5n - 16$ claimed), so we must be more careful to see the origin of the "-16" term. In the original position, a_2 is in the outer square intersecting the positive x-axis, at distance δ from Y_1. Also, a point a_i is in the outer square intersecting the negative y-axis, at distance 4δ from Y_4, while point a_{i-1} is in the inner square intersecting the negative y-axis, at distance 4δ from W_4. After a clockwise rotation of almost 90 degrees, a_i is in the outer square intersected by the positive x-axis, at distance δ from Y_1, while a_{i-1} is in the inner square intersected by the positive x-axis, at distance 4δ from W_1. Clearly, a rotation by δ will not give us a matching. In fact, we miss four matches while a_{i-1} is in the inner square intersecting the positive x-axis. This situation repeats itself when a_i gets to the inner square on the positive y-axis and negative x- and y-axes. Thus, there are exactly $5n - 16$ distinct matches for our construction. ∎

3 Upper Bounds on the Number of Matchings

In this section we prove upper bounds on the number of possible matches for several cases of the matching problem. We begin with pure translation. Alt et al [AMWW] show that there is at most one match in the case in which the noise regions are disjoint unit circles in the plane. We generalize their result for linearly separable noise regions in any dimension. In particular, our result applies to the case of pairwise-disjoint convex noise regions.

Theorem 4 *There is at most one perfect matching between a set $\mathcal{A} = \{a_1, \ldots, a_n\}$ of n points in $\Re^d$ and a set $\mathcal{R} = \{B_1, \ldots, B_n\}$ of n linearly separable regions in $\Re^d$ under pure translation.*

Proof. The proof is an extension of the proof in [AMWW]. A perfect matching is a permutation on the sequence $\{1, \ldots, n\}$. A cycle is defined as a cyclic shift of a subset of points between regions. If there are two different matchings under translation, we can assume without loss of generality that there exists a translation vector t such that

$$a_1 \in B_1,\ a_2 \in B_2, \ldots, a_k \in B_k$$

and (3.1)

$$a_1 + t \in B_2, \ldots, a_{k-1} + t \in B_k,\ a_k + t \in B_1.$$

A relation "$<_t$" can be defined on the set of noise regions as follows: $B_i <_t B_j$ if there exists a ray in the direction of t starting from B_i and passing through B_j. We claim that the relation "$<_t$" has the following two properties:

(a). "$<_t$" is anti-symmetric.

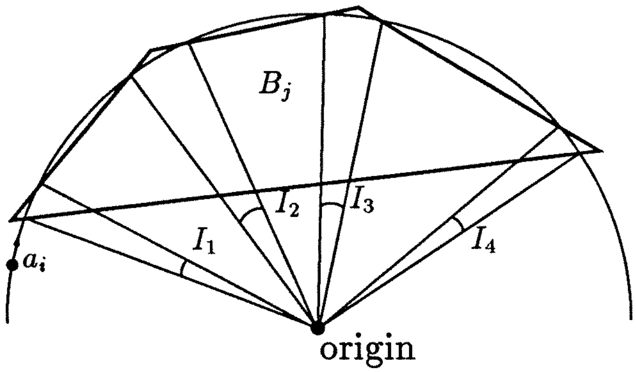

Figure 5: The intervals $\mathcal{I}_{ij} = \{I_1, I_2, I_3, I_4\}$

(b). "$<_t$" has the following property of "transitivity": if $B_1 <_t B_2 <_t \ldots <_t B_k$ and there is a line parallel to t stabbing both B_1 and B_k, then $B_1 <_t B_k$. (In other words, part of this line is a ray starting at B_1 and passing through B_k.)

The first property is an immediate consequence of the fact that the regions are linearly separable.

The second property is proved by induction: for $k = 3$ the claim is clear. It is also clear that we can move the line stabbing B_1 and B_k parallel to itself, until it stabs a third region B_i. By the induction hypothesis, $B_1 <_t B_i$ and $B_i <_t B_k$. Thus, applying the result for $k = 3$, we conclude that $B_1 <_t B_k$.

We conclude the proof of the theorem by noting that property (b) rules out the possibility of the cycle indicated in equation (3.1). ∎

We now turn to pure rotation, and consider again linearly separable noise regions. Clearly, the upper bounds we get for linearly separable noise regions apply to arbitrary unit squares and arbitrary convex polygons.

Theorem 5 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n linearly separable regions $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches is bounded above by $O(n^2)$.*

Proof. For each point a_i and for each region B_j, we define a set of angular intervals $\mathcal{I}_{ij}$, to be the set of maximally connected angular intervals, such that when $\mathcal{A}$ is rotated by an angle within an interval of $\mathcal{I}_{ij}$ then point a_i is in region B_j. Note that the number of intervals of $\mathcal{I}_{ij}$ is bounded by the number of "sides" of B_j and may be large. (See Figure 5, where it is seen that the orbit of one point may go in and out of a convex noise region.)

However, we claim that the number of times one point may switch from one region to other regions is $O(n)$. We prove this by noting that, since each pair of regions B_k and $B_l \neq B_k$ are linearly separable, a point a_i may cross the line separating them at most twice.

This implies that the sequence of regions through which a_i can pass cannot contain the alternating subsequence "$..B_k..B_l..B_k..B_l..$". Thus, this sequence is a Davenport-Schinzel sequence of order 2, implying that its length is at most $2n-1$ (see [HS]). Since this analysis applies to each of the n points a_i, we get an overall bound of $O(n^2)$ on the number of matches. ∎

Combining this theorem with Theorem 2 shows that the bound is tight. Next we consider the case in which the noise regions are disjoint circles, and prove an upper bound that matches the lower bound that we showed in Fact 1.

Theorem 6 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n circles (of any size) $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches is bounded above by n, and this bound is tight.*

Proof. Define a "cut" path π as follows: (Assume first, that no circle covers the origin.) Consider the vertical ray ρ from the origin. Define π to be the path we get by replacing portions of ρ that are interior to a circle B_j by the shorter of the two arcs of B_j. We "cut" the plane along π. Intuitively, this gets rid of the "wrap-around" effect.

The case in which there exists a circle B_i that covers the origin requires a minor modification to the definition of the cut π. (Note, there can be at most one circle covering the origin.) As before we start the cut at the origin. The cut goes (along a straight segment) from the origin to the closest point to the origin on the boundary of B_i, then along the boundary of B_i to the point of B_i farthest from the origin (halfway around the circle), and then leaves the circle (perpendicularly) along the ray from the origin. The remainder of the cut π procedes as before.

Note that π is a radially monotone path (i.e., monotone in r). This implies that as we rotate the points about the origin, each point will pass over π exactly once. Assume that the points $\mathcal{A}$ are indexed according the order in which they hit the path π. (If several points happen to hit π at the same time, then we can break ties by the distance of the points from the origin.) Thus, there are exactly n possible ways to cut the point set $\mathcal{A}$ by π.

Without loss of generality, we assume that the point set rotates clockwise. We can define a relation "$<_c$" among the circles: $B_i <_c B_j$ if a clockwise directed arc centered in the origin, goes from B_i to B_j without crossing (or touching) π. We now claim that the relation "$<_c$" satisfies the same properties as "$<_t$" in Theorem 4, namely, "$<_c$" is anti-symmetric and "transitive". As a result, if no point crosses π, there can be at most one possible matching. As there are n different possible cuts, this yields the desired result. ∎

We note that the proof of the theorem did not require the circles to be of equal size.

Theorem 7 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n aligned unit squares $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches is bounded above by $5n$.*

Proof. We begin by noting that if none of the noise regions are cut by the axes then there are at most n matches. This follows from a proof similar to the one of Theorem 6. We can again define a partial order on the squares in each quadrant and extend it to an order on all squares not cut by the axes: if a point rotates out of square B_i and into square B_j we say that $B_i <_s B_j$. This partial order is anti-symmetric and transitive. Thus, if we cut open the points A in n different ways using a monotone cut (say, go up from the origin; when you hit a square, go left until the lower left corner, and then go up again, etc), we get at most n possible matches.

We now have to count additional matches that occur when no point crosses the cut. These matches are "caused" by a square or squares that are cut by one or more of the axes (since we showed that if there are no such squares there can be at most one match as long as no point crosses the cut). Now, as each point rotates around the origin, it can enter and leave a square twice only if that square is cut by an axis. For each point there are at most four such squares (one on each of the positive and negative axes). Hence, each match that occurs, without a point having crossed the cut, can be charged to a point and one of the four squares we described. As a result there are only $4n$ such "additional" matches that occur when no point crosses the cut and in total only $5n$ possible matches. ∎

Note that the upper bound given in the above theorem matches the lower bound given in Theorem 3 up to an additive constant of 16. It is also interesting to observe that in the proof we used the fact that the squares are of equal size when we claimed that a point has at most four squares that it can enter and leave twice. This fact is not true for unequal squares. This leaves open the question of what is the number of possible matches for aligned unequal size squares.

We come now to one of our main results:

Theorem 8 *Consider the problem of matching under translation, rotation, and scale ($\mathcal{T}$ = TRS) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. The number of matches is bounded above by $O(n^2)$.*

Proof. The proof is based on the following five lemmas:

Lemma 9 *A transformation τ which is a translation and scale can be written as a pure scale with respect to some origin.*

Proof. The proof is trivial: Let $\tau(p) = \alpha p + t$, for translation vector t, and scale parameter $\alpha \neq 0$. We can write $\tau(p) = \alpha(p + \frac{t}{\alpha})$, which is a pure scale with respect to a new origin $O^* = -\frac{t}{\alpha}$. ∎

Lemma 10 *Under scale alone ($\mathcal{T}$ = S), there can be at most one match between a set of n points and n unit circles.*

Proof. Without loss of generality, the scale factor is $\alpha > 1$. A point will move from circle B_i to circle B_j only if the center of B_i is closer to the origin than the center of B_j. Thus a cycle of points switching positions is impossible since the point in the farthest circle of the cycle has no circle to which it can be matched. ∎

Corollary 11 *Under combined translation and scale there is at most one match between n points and n unit circles.*

Proof. Under translation only or scale only there is at most a single match (Theorem 4 and Lemma 10). By Lemma 9, a scale plus translation can be written as a pure scale with respect to some origin point, and thus using Lemma 10 we have shown the claim. ∎

Lemma 12 *Given a set $\mathcal{A} = \{a_1, \ldots, a_n\}$ of n points in the plane. Let a^* denote the center of mass of $\mathcal{A}$, and assume that a_1 is the point of $\mathcal{A}$ farthest from a^*. Let τ be a similarity transformation such that $\tau(a^*) \in C^*$ and $\tau(a_1) \in C_1$, where C^* and C_1 are disks of radius ϵ. Then, $\tau(a_i)$ lies within a disk of radius 3ϵ, for each $i = 1, \ldots, n$.*

Proof. Let τ be a similarity transformation such that $\tau(a^*) \in C^*$ and $\tau(a_1) \in C_1$. We can write $\tau(a) = \alpha R a + t$, where α is the scale factor, R is the rotation matrix, and t is the translation vector corresponding to τ. Let τ_0 be the similarity transformation that maps point a^* to the center of circle C^* and maps point a_1 to the center of circle C_1; we write $\tau_0(a) = \alpha_0 R_0 a + t_0$. We will show that $\tau(a_i)$ must lie within the disk of radius 3ϵ centered at point $\tau_0(a_i)$.

$$\begin{aligned}
\|\tau(a_i) - \tau_0(a_i)\| &= \|\alpha R a_i + t - \alpha_0 R_0 a_i - t_0\| \\
&= \|(\alpha R a^* + t - \alpha_0 R_0 a^* - t_0) + (\alpha R - \alpha_0 R_0)(a_i - a^*)\| \\
&\leq \|\tau(a^*) - \tau_0(a^*)\| + \|(\alpha R - \alpha_0 R_0)(a_i - a^*)\| \\
&\leq \epsilon + \|(\alpha R - \alpha_0 R_0)(a_i - a^*)\| \\
&= \epsilon + \|\beta \overline{R}(\alpha R - \alpha_0 R_0)(a_1 - a^*)\|,
\end{aligned}$$

where β is a scale factor, and $\overline{R}$ is a rotation matrix, such that $(a_i - a^*) = \beta \overline{R}(a_1 - a^*)$. Since a_1 is further from a^* than a_i is, we know that $0 \leq \beta \leq 1$. Thus,

$$\begin{aligned}
\|\tau(a_i) - \tau_0(a_i)\| &\leq \epsilon + \beta \|\overline{R}(\alpha R - \alpha_0 R_0)(a_1 - a^*)\| \\
&= \epsilon + \beta \|(\alpha R a_1 + t - \alpha_0 R_0 a_1 - t_0) - (\alpha R a^* + t - \alpha_0 R_0 a^* - t_0)\| \\
&\leq \epsilon + \|\tau(a_1) - \tau_0(a_1)\| + \|\tau(a^*) - \tau_0(a^*)\| \\
&\leq 3\epsilon.
\end{aligned}$$

∎

Lemma 13 *Consider the set of all similarity transformations τ to be points (x, y, α, θ) in a four-dimensional cylindrical parameter space $\mathcal{C}$, where x and y denote translation, α denotes scale, and θ denotes the angle of rotation. Then, the set $Q = \{\tau \in \mathcal{C} \mid \tau(a_i) \in B_i \quad i = 1, \ldots, n\}$ is a connected subset of $\mathcal{C}$.*

Proof. We can consider the set $\mathcal{C}$ of all similarity transformations to be a four-dimensional (cylindrical) surface embedded in $\Re^5$, so that we can account for the "wrap-around" effect of rotation, while restricting θ to lie within $[0, 2\pi]$.

If, instead, we parameterize the transformations τ by the four coordinates (x, y, r_1, r_2), where $r_1 = \alpha\cos\theta$ and $r_2 = \alpha\sin\theta$, it is easy to see that the set $Q' = \{\tau \in \mathcal{C} \mid \tau(a_i) \in B_i \;\; i = 1, \ldots, n\}$ is a convex subset of $\Re^4$ (as in Baird [Ba]).

Consider two points $p_1, p_2 \in Q$. We wish to show that there exists a path connecting them in Q. Look at their corresponding points $q_1, q_2 \in Q'$. Since Q' is convex we know that the line segment connecting q_1 and q_2 lies in Q'. Describe this segment by a continuous function f. By plugging $\alpha\cos\theta$ for r_1 and $\alpha\sin\theta$ for r_2 we write f in $\mathcal{C}$, using a continuous function from p_1 to p_2. Thus we have a path connecting p_1 and p_2. This path consists of points entirely in Q, since their images are all in Q'. ∎

Recall that when $\mathcal{R} = \{B_1, \ldots, B_n\}$ denotes a set of circles, then the center of circle B_i is denoted by b_i, and b^* is the center of mass of the centers of the circles.

Lemma 14 *Consider the problem of matching under translation, rotation, and scale ($\mathcal{T}$ = TRS) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. If $\tau \in \mathcal{T}$ = TRS is a transformation that yields a (perfect) matching, then $\|\tau(a^*) - b^*\| \leq 1$.*

Proof. Without loss of generality we assume that $\tau(a_i) \in B_i$ and thus $\|\tau(a_i) - b_i\| \leq 1$ so

$$\|\tau(a^*) - b^*\| = \|\tau(\frac{1}{n}\sum_{i=1}^{n} a_i) - (\frac{1}{n}\sum_{i=1}^{n} b_i)\| =$$

$$\|\frac{1}{n}\sum_{i=1}^{n}(\tau(a_i)) - \frac{1}{n}\sum_{i=1}^{n} b_i\| \leq \frac{1}{n}\sum_{i=1}^{n}\|\tau(a_i) - b_i\| \leq 1$$

∎

The above lemma can be generalized to the case in which the noise regions are normal. Let r_{max} be the maximum over all regions of the distance between the center of the region and its boundary. It is easy to show that $\|\tau(a^*) - b^*\| \leq r_{max}$.

Proof of Theorem 8: Denote by $\beta_1(b^*)$ the unit circle around the the center of mass of the centers of the circles. It suffices to prove that each time the point a_1, which is the farthest point from a^*, is constrained to be inside one unit circle $B_i \in \mathcal{R}$, then there are at most $O(n)$ matches.

Let $\epsilon_0 = \frac{2}{3}\sqrt{3} - 1$. Fix one circle B_i, and cover $\beta_1(b^*)$ and B_i with small circles of size $\frac{1}{3}\epsilon_0$. Choose a small circle C_1 from the covering of $\beta_1(b^*)$, and C_2 from the covering of B_i. We will prove that if we constrain a^* to be inside C_1 and a_1 to be inside C_2, we get at most n different matches. This proves the theorem since the number of pairs (C_1, C_2) (for each B_i) is constant.

By Lemma 12, if a^* and a_1 perform a motion within C_1 and C_2 respectively, which are of radius $\frac{1}{3}\epsilon_0$, each point a_i is constrained to be within a circle of size ϵ_0. Because of the choice of ϵ_0, a circle of that radius can intersect at most two unit circles in $\mathcal{R}$. Hence, during this motion, a_i can be in at most two different unit circles. We can build a bipartite graph, whose "left" set of nodes is the set $\mathcal{A}$, and whose "right" set of nodes is $\mathcal{R}$. An edge connects a point-node, a_i, with a circle-node, B_j, if during the motion of a_1 and a^* in the chosen circles C_1 and C_2 point a_i is in circle B_j. We now look for all the possible matches on this graph (if there are any). Note that since each point can intersect at most two circles, during this motion, each left node in the graph is of degree at most two.

In order to solve the matching problem, we proceed as in [AMWW]: For each node u of degree one, we delete the node u, add the corresponding edge (u, v) to the matching, and delete the other edges incident to v. We repeat this process until all nodes have degree at least two. But the special property of our graph implies that each node must have degree exactly two (otherwise, no perfect match exists), so the connected components in the resulting graph must be cycles.

A cyclic component has two possible matches in the graph, hence, there are at most two valid matches between the points and the circles in this component. If the number of disjoint cyclic components is k then the number of possible different matches can be 2^k. However, we show that this is not the case here. Denote by μ_1 and μ_2 the two possible (graph) matches derived from one cyclic component. Let θ be the amount of rotation between the initial position of the point set $\mathcal{A}$ and its transformed position, and define I_1 to be the set of maximally connected angular intervals where a match μ_1 is obtained for $\theta \in I_1$ and for some scale and translation. Similarly define I_2 for μ_2.

We claim that I_1 (I_2) consists of a single angular interval. This is due to the fact that the set of transformations yielding μ_1 (μ_2) is a connected set in the parameter space $\Re^4$, by Lemma 13. Corollary 11 implies that the two intervals are disjoint.

Superimposing the $2k$ intervals, subdivides $[0, 2\pi]$ into $4k + 1 = O(n)$ subintervals. Within each subinterval at most one match is possible; hence, for the motion described above, there are $O(n)$ possible matches. This concludes the proof. ∎

It is interesting to note that the proof of the preceding theorem breaks down for the case of aligned unit squares. There are two problems in trying to apply this proof. The first is that Corollary 11 is not true for aligned unit squares. The second problem is that there does not exist a small enough region that is guaranteed to intersect at most two disjoint unit squares, and hence when we build the graph as above, it will not have the property that all nodes on the left has degree at most two. This leaves open the question of proving a good bound for the number of matches under similarity in the case of aligned unit squares.

4 Algorithms for Finding Perfect Matchings

In this section we describe algorithms for finding matchings between a point set $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of noise regions $\mathcal{R} = \{B_1, \ldots, B_n\}$. By exploiting the special nature of pairwise-disjoint noise regions, we obtain running times better than those of [AMWW, AMZ], which apply to the general case.

We begin with noise regions that are convex, not necessarily normal, and the allowed motion is pure rotation. We let N be the total number of edges of the noise regions.

Theorem 15 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n linearly separable noise regions $\mathcal{R} = \{B_1, \ldots, B_n\}$,* *with N edges altogether. There exists an algorithm to find all matches in time complexity $O(nN \log N)$ and space complexity $O(nN)$.*

Proof. The algorithm proceeds as follows: For each region B_i and each point a_j we find the set of angular intervals $\mathcal{I}_{ij}$ for which point a_j is in region B_i. Recall these intervals were defined in Theorem 5. The total number of such intervals is bounded by the number of "sides" of the regions in $\mathcal{R}$. Next we sort the endpoints of these intervals, and get at most $2N$ new intervals. It is now simple to march through these intervals and check which of them constitute a complete matching. ∎

Note that the above algorithm applies also to the case in which the noise regions are circles of possibly different sizes (yielding a time bound of $O(n^2 \log n)$ and a space bound of $O(n^2)$). For unit circles we give an improved algorithm below. We first present an algorithm for the labeled matching problem for unit circles, improving the algorithm of [AMWW] for our case of disjoint noise regions:

Lemma 16 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. If we are given the correspondence between points and circles, then we can find the set of all rotations that achieve this correspondence in time $O(n)$.*

Proof. We assume without loss of generality that the correspondence matches point a_i to circle B_i for each i. The algorithm is simple. Proceeding as in [AMWW] we find the set of intervals $I_i \subset [0, 2\pi]$ for which point a_i is in circle B_i. We now need to determine whether $I_1 \cap \cdots \cap I_n \neq \emptyset$. [AMWW] solve this by sorting the endpoints of the intervals and thus get an $O(n \log n)$ time algorithm. We perform an incremental intersection of the intervals. Note that since at most one circle can contain the origin, at most one of the intervals is of length greater than π, let that circle be B_n, without loss of generality. Thus we can, in constant time, find the intersection of two of the intervals say I_1 and I_2 and get one new interval say, I'_1. We can continue now by intersecting I'_1 with I_3 and so on, until we either intersect all intervals or get an empty interval. Thus in time $O(n)$ we can determine whether the intervals have a nonempty intersection, and hence solve the problem. ∎

Remark. In [AMWW] an $O(n \log n)$ algorithm is given for the general case of matching under pure rotation when the noise regions are allowed to overlap arbitrarily. Their algorithm is optimal, since they must compute the intersection of a set of intervals on a circle, and when there are noise regions covering the origin, there can be intervals whose extent is greater than π. The simple method just presented for our pairwise-disjoint case applies also to the more general case, *provided* that there are only a bounded number of noise regions covering the origin.

Theorem 17 *Consider the problem of matching under pure rotation ($\mathcal{T} = R$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. There exists an algorithm to find all ($\leq n$) matches in time complexity $O(n^2)$ and space complexity $O(n^2)$.*

Proof. To get the improved algorithm for unit circles, we partition the interval $\mathcal{I} = [0, 2\pi]$ into subintervals that we can process efficiently. Call the point a_j that is farthest from the origin the "reference point" p. Along p's orbit, p will pass through circles in certain places. We subdivide $\mathcal{I}$ according to the intervals in which p is in a circle or out of any circle. Call these "primary intervals" (inside the circles). Define a cut π as in the proof of Theorem 6. Use this cut to further subdivide the intervals, whenever one of the n points crosses the cut. Note that this refined subdivision still has only $O(n)$ intervals. Now, for a small enough angle ϵ, rotation of p by ϵ will cause all the other points to execute arcs that are small enough that they do not intersect more than two unit circles. We can compute such an ϵ by noting that a circle of radius $\frac{2}{3}\sqrt{3} - 1$ can intersect at most two disjoint unit circle. Placing the center of this small circle on the orbit of p, ϵ can be defined as the size of the interval of angles for which p is in the small circle. Such an ϵ is a small constant. Now, take each refined interval that is larger than ϵ and further refine it into pieces of size at most ϵ. We still have only $O(n)$ intervals – call them "basic" intervals.

Now we proceed as in [AMWW] and the proof of Theorem 8: For each of $O(n)$ basic intervals, we define a bipartite graph whose left set of nodes is the set of points $\mathcal{A}$ and whose right set of nodes is the set of circles $\mathcal{R}$. An edge connects a point-node, a_i, to a circle node, B_j, if during the motion of p within this basic interval, point a_i is in circle B_j. Because of the way we built the basic intervals, this bipartite graph has left nodes with degree at most 2, which will permit a simple matching algorithm, described below.

We now show that these graphs can be built efficiently. For the first basic interval we build the graph by brute force in time $O(n^2)$. From each graph advance to the next in total time $O(n^2)$. We do so by noting that as p moves by some angle θ from the midpoint of its current basic interval to the next basic interval, we need to advance each of the other points along their orbits also by θ (to find the midpoint of the new arc, of whatever size). We can advance easily, since we know the intersection points of the circles along the orbits. (These are easy to find by sorting the centers of the unit circles angularly about the origin.) Now, for each orbit, march through the sorted list, checking each unit circle for intersection with the orbit.

In order to solve the matching problem on each resulting bipartite graph, we proceed (as in the proof of Theorem 8): we delete nodes of degree one, and the relevant corresponding edges, until we are left with a graph whose connected components are cycles. In time proportional the the number of nodes in each component we can find whether there is a match of the points into the circles by the algorithm described in Lemma 16 for the case of labeled point matching. Now, since by the construction of the basic intervals, we know that there is only one match per basic interval and hence only one of the possible matches of the cycle components yields a match of the points into circles. Thus we can solve the matching problem for each component separately and combine the results into a possible match for the whole graph. So, in time $O(n)$ we solve the problem for each basic interval, and since there are only $O(n)$ intervals, we have shown an $O(n^2)$ algorithm. ∎

Theorem 18 *Consider the problem of matching under translation and rotation ($\mathcal{T} = TR$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n unit circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. There* *exists an algorithm to find all matches in time complexity $O(n^4 \log n)$.*

Proof. Let $\beta_1(b^*)$ be, as above, the unit circle around the center of mass of the centers of the circles $\mathcal{R} = \{B_1, \ldots, B_n\}$. In any match, the center of mass, a^*, of the points $\mathcal{A}$, must be within $\beta_1(b^*)$. As usual, we let a_1 be, without loss of generality, the point in $\mathcal{A}$ farthest from a^*. Fix a circle B_i. If a_1 is constrained to be in B_i, there is a very little freedom of motion left for the rest of the points of $\mathcal{A}$, (since a^* must always be in $\beta_1(b^*)$).

By Lemma 12, each other point of $\mathcal{A}$ executes a motion ("track") in the plane which is contained within a circle of radius 3, as we rotate and translate $\mathcal{A}$ subject to the above constraints ($a_1 \in B_i$ and $a^* \in \beta_1(b^*)$). This circle of radius 3 can intersect only a constant number ($\leq \lceil 49/\pi \rceil$) of unit circles of $\mathcal{R}$.

Hence, a pair of points a_j and a_k of $\mathcal{A}$, has only a constant number of circles from $\mathcal{R}$ where these points may reside. Choose two such circles $B_{j'}$ and $B_{k'}$, from the circles that cover the tracks of a_j and a_k, and "pin" a_j and a_k to lie on their boundaries correspondingly. This leaves us with only one degree of freedom (parameterized by t) for the remaining points, which can be described as a small number of low-degree (degree 6) polynomial tracks.

A point moving along one of the tracks can enter and exit a circle only a constant number of times (at most 12). For each track of point a_l, $l \neq i, j, k$, parametrized by t, the number of intervals $\{I_t\}$ of t in which a_l is in one of its corresponding circles is constant.

Thus $\{I_t\}$ is subdivided into $O(n)$ subintervals, each subinterval corresponds to a matching of one a_l, $l \neq i, j, k$, into one of the unit circles B_i. By the method we used in Theorem 15 we find the interval sets $\{I_t\}$ for all a_l, $l \neq i, j, k$, we sort all the subintervals and find all the possible matches, for a_1 in B_i and for the pair a_j and a_k in $B_{j'}$ and $B_{k'}$ respectively.

The complexity of the above algorithm is: $O(n)$ for choice of B_i, $O(n^2)$ choices for choosing a pair a_j and a_k, and $O(n \log n)$ per choice of placing all a_l, $l \neq i, j, k$. Thus the overall time complexity is $O(n^4 \log n)$. ∎

Remark. In a sense, the time bound on this algorithm is the best we can expect since if we were given the n or more possible correspondences between the point set and the noise regions, checking each such labeled matching takes time $O(n^3 \log n)$ ([ISI]).

The above theorem and algorithm can easily be generalized to handle polygonal noise regions that are normal. It suffices to note that for normal regions (not necessarily polygonal) the center of mass of the centers of the noise regions must lie within a bounded region around the center of mass of the points. Constraining a_1 to lie in each noise region, in turn, restricts the motion of all other points to some bounded region that can intersect at most a constant number of the regions in $\mathcal{R}$. We can continue as in the previous theorem by pinning two other points to the boundary of two regions and doing a one-parameter search. We have thus argued the following:

Theorem 19 *Consider the problem of matching under translation and rotation ($\mathcal{T} = TR$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n normal polygonal regions $\mathcal{R} = \{B_1, \ldots, B_n\}$, with N edges altogether. There exists an algorithm to find all matches in time complexity $O(nN^3 \log N)$.*

Remark. An important consequence of the above algorithm is that it implies an upper bound of $O(nN^3)$ $(O(n^4))$ on the number of matches under translation and rotation in the case of normal noise regions (unit squares).

Another immediate extension of Theorem 18 allows the transformation between the point set and the noise regions to be a similarity, namely allow scaling as well as translation and rotation. The only change to the algorithm is that we now need to consider triples of points a_j, a_k, and a_l, (instead of pairs of points) to get a set of intervals, parametrized by t in which other points are in some region. This implies that the complexity of the algorithm increases by a factor of n.

Theorem 20 *Consider the problem of matching under translation, rotation, and scale ($\mathcal{T} = TRS$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n normal polygonal regions $\mathcal{R} = \{B_1, \ldots, B_n\}$, with N edges altogether. There exists an algorithm to find all matches in time complexity $O(nN^4 \log N)$.*

Remark. The above result implies an upper bound of $O(nN^4)$ $(O(n^5))$ on the number of matches under similarity in the case of normal noise regions (unit squares).

In the case that the regions are aligned unit squares that are ϵ-separated, we can get an improved algorithm as stated in the following theorem:

Theorem 21 *Consider the problem of matching under translation and rotation ($\mathcal{T} = TR$) a set of n points $\mathcal{A} = \{a_1, \ldots, a_n\}$ and a set of n aligned unit squares $\mathcal{R} = \{B_1, \ldots, B_n\}$ that are ϵ-separated. There exists an algorithm to find all matches in time complexity $O(n^2 \log n)$ and space complexity $O(n)$.*

Proof. The algorithm is similar to previous ones. From the comment after Lemma 14 we know that the center of mass of the points, a^*, lies within a bounded region of the center of mass of the centers of the squares. We let C be a circle of radius $\frac{\sqrt{2}}{2}$ around the center of mass of the centers of the squares. Without loss of generality let a_1 be the point farthest away from the center of mass of the points. For each square B_i in turn, we do the following. Cover the circle C and the square B_i by small circles of radius $\frac{1}{3}\epsilon$. Let C_1 and C_2 be circles in the cover of C and B_i. Note that the number of choices for this pair is constant. We know by Lemma 12 that while a^* lies in C_1 and a_1 lies in C_2 all other points are restricted to a circle of size ϵ. As in Theorem 8, we build a bipartite graph, only this time the degree of point nodes is at most one. Thus, there is at most one matching in the graph, which we can find in linear time. Using this correspondence, we determine in time $O(n \log n)$ whether this correspondence constitutes a matching of the points to the squares, by the algorithm of [ISI]. It remains to be shown how to build each bipartite graph in time $O(n \log n)$. This is done by a standard trick: Build the Voronoi diagram of the squares and do point location of the centers of each of the size-ϵ circles in which each point lies (during the constraint of a^* to C_1 and a_1 to C_2). ∎

5 Conclusion

Several directions for further study are suggested by our work:

Questions are largely open in higher dimensions. Can we get tight bounds on the number of matches under pure rotation in space? Assume, for example, that the regions $\mathcal{R} = \{B_1, \ldots, B_n\}$ are disjoint unit spheres. Note that pure rotation in space has *three* degrees of freedom. Our best lower bound at the moment is $\Omega(n)$ for the number of matches. Another interesting open question has to do with pure translation: Given n points and n pairwise-disjoint convex regions (e.g., lines) in $\Re^3$, how many different perfect matches are obtainable under pure translation? Unlike the problem in two dimensions, it is now possible to have more than one match: consider three lines whose projection onto a plane yields a "cycle" of overlaps.

We have focused on the case in which the cardinalities of the sets of image points and model points are equal ($m = n$). It is more realistic to consider the case of *un*equal cardinalities. If we have n points and $m = n + k$ regions, then we can obtain complexity bounds in terms of k and n. Note that in many applications k is likely to be small. As an example of the increased complexity in the unequal cardinality case, note that even for one degree of translational freedom, there is a lower bound of $\Omega(nk)$ on the number of matches, in contrast to the *unique* match that exists when $m = n$. One of the means used to obtain upper bounds and fast algorithms, namely Lemma 14, which relates the transformation of a^*, the center of mass of the points, to b^*, the center of mass of the regions centers, fails to hold when $m \neq n$. As a result we cannot extend Theorem 8 to get an upper bound on the number of matches when the regions are unit circles. As for algorithms, it is not hard to extend those described in Theorem 19, for translation and rotation, and Theorem 20 for translation, rotation and scale. Instead of pinning a matching using a^* and a_1, we use the two points farthest apart, and obtain algorithms whose running time is increased by a factor of n, to $O(n^2 N^3 \log N)$ and $O(n^2 N^4 \log N)$ respectively.

We can extend many of our results to the case in which we allow *partial* overlap among the regions $\mathcal{R} = \{B_1, \ldots, B_n\}$. For example, when we know that no point of the plane is covered by more than k of the sets $\mathcal{R} = \{B_1, \ldots, B_n\}$, we can obtain efficient complexity bounds in terms of n and k.

References

[AMWW] Alt, H., Mehlhorn, K., Wagener, H. and Welzl, E., "Congruence, Similarity, and Symmetries of Geometric Objects", *Discrete and Computational Geometry* **3** (1988), pp. 237-256.

[AMZ] Arkin, E.M., Mitchell, J.S.B., and Zikan K., "Algorithms for Generalized Matching of Point Sets", Manuscript, 1988. Presented at the CORS/ORSA/TIMS National Meeting, Vancouver, B.C., May 1989.

[Ata] Atallah, M.J. "A Matching Problem in the Plane", *J. Comput. Systems Sci.* **31** (1985), pp. 63–70.

[Atk] Atkinson, M.D. "An Optimal Algorithm for Geometric Congruence", *J. Algorithms* **8** (1987), pp. 159–172.

[Ba] Baird, H.S., *Model-Based Image Matching Using Location*, MIT Press, Distinguished Dissertation Series, 1984.

[HS] Hart, S. and Sharir, M., "Nonlinearity of Davenport-Schinzel Sequences and of Generalized Path Compression Schemes", *Combinatorica* **6** (2) (1986), pp. 151–177.

[He] Heffernan, P., "The Translation Square Map and Approximate Congruence", Technical Report 940, School of Operations Research and Industrial Engineering,Cornell University, 1990.

[HK] Huttenlocher, D.P. and Kedem, K., "Computing the Minimum Hausdorff Distance for Point Sets Under Translation", *Proc. 6th ACM Symposium on Computational Geometry*, pp. 340-349, 1990.

[HKS] Huttenlocher, D.P., Kedem, K. and Sharir, M., "The Upper Envelope of Voronoi Surfaces and its Applications", *Proc. 7th ACM Symposium on Computational Geometry*, 1991.

[ISI] Imai, K., Sumino, S., and Imai, H., "Minimax Geometric Fitting of Two Corresponding Sets of Points", *Proc. 5th ACM Symposium on Computational Geometry*, pp. 266-275, 1989.

[Sc] Schirra, S., "Approximate Algorithms for Approximate Congruence", Technical Report A21/90, Universität des Saarlandes, Saarbrücken, Germany, 1990.

[Sp] Sprinzak, J., "Point Matching Algorithms", MS Thesis, Hebrew University, Department of Computer Science, 1990.

[SW] Sprinzak, J. and Werman, M., "Exact Point Matching", *Proc. Israeli Symposium on AI, Vision and Pattern Recognition*, Tel Aviv, 1990.

[Zi] Zikan, K., "Least-Squares Image Registration", *ORSA Journal on Computing* **3** (1991), No. 2, pp. 169-172.