Matching Points into Pairwise-Disjoint Noise Regions: Combinatorial Bounds and Algorithms

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

January 14, 1992

Esther M. Arkin
Department of Applied Mathematics and Statistics
State University of New York at Stony Brook
Stony Brook, NY 11794-3600
(516) 632-8363
estie@ams.sunysb.edu

Klara Kedem
Department of Computer Science
Tel-Aviv University
Tel-Aviv, 69978 Israel
klara@taurus.bitnet

Joseph S. B. Mitchell
Department of Applied Mathematics and Statistics
State University of New York at Stony Brook
Stony Brook, NY 11794-3600
(516) 632-8366
jsbm@ams.sunysb.edu

Josef Sprinzak
Department of Computer Science
The Hebrew University
Jerusalem, Israel

Michael Werman
Department of Computer Science
The Hebrew University
Jerusalem, Israel
werman@cs.huji.ac.il

Please address correspondence to Professor Arkin.

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.

Introduction

A fundamental problem in pattern matching and 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.* For example, we may know the pattern (model) of a particular constellation and want to find an instance of the constellation within an image of the nighttime sky. Or, as another example, we may have given a geometric model of an automobile and want to determine if an instance of it occurs in a given digitized image. A standard approach to this problem is to extract interesting "features" (e.g., a corner of a bumper, a handle, an antenna, etc.) in both the model and the image, using local "interest operators", and then to determine how well these sets of features can be made to correspond when we look at all possible positions of the model with respect to the image (e.g., see Huttenlocher [11]).

We are therefore motivated to study the following *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 $\tau : \Re^d \rightarrow \Re^{d'}$, 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 associated with pairs.

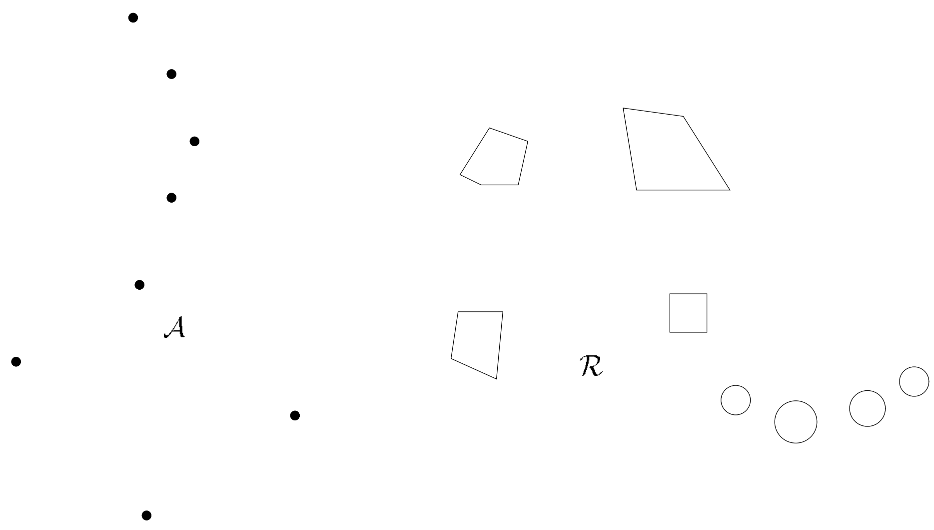

Figure 1: An inexact point matching problem.

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. An example of an inexact point matching problem is illustrated in Figure 1, where the goal may be to determine if the point set $\mathcal{A}$ can be transformed by rotation and translation so that each point lies within exactly one of the noise regions $\mathcal{R}$.

Related Work. 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, and possibly a reflection); see Alt et al.[1], earlier work by Atallah[3] and Atkinson[4], and the recent work of Sprinzak and Werman[19].

Baird[6] did some of the pioneering work on the inexact point matching problem but left open the question of obtaining polynomial-time algorithms. Alt et. al[1] applied techniques of computational geometry to give polynomial-time algorithms for a wide variety of inexact point matching problems. Several other papers have also obtained efficient algorithms for instances of the inexact point matching problem[2, 9, 10, 14, 17, 18].

Closely related to the 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 point sets has been recently solved in time $O(n^3)$ by Huttenlocher et al.[12, 13]. The problem of finding a "least squares" registration between point sets has been addressed by Zikan[20] and by Aurenhammer et. al[5].

Our Results. In this paper, we consider the inexact point matching problem. 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*, *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 focus on the equal cardinality case ($m = n$); thus, 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".

(3) The regions B_j may be unit circles, arbitrary circles, aligned unit squares, arbitrary unit squares, or a set of "normal" polygonal shapes (which are well-behaved disjoint polygons, as we define below). In this last case, we let M denote the maximum number of vertices of any one polygon B_j and let N denote the total number of vertices of all polygons B_j. We note that in most applications the noise regions are assumed to be circles or aligned squares (disks in the L_2 or L_∞ metric).

(4) Various assumptions about separation of the regions B_j can be considered. Previous results[1, 2] 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 lower-degree polynomial bounds on the number of matchings and on the time to compute all matchings. Some of our results apply under slightly stronger notions of separation: linearly separable regions and ϵ-separated regions (defined below).

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. To find a

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 $O(n^2 \log n)$	$5n-16$ $5n$ $O(n^2 \log n)$	$\Omega(n^2)$ $O(n^2)$ $O(n^2 \log n)$	$\Omega(n^2)$ $O(n^2)$ $O(nN \log N)$	$\Omega(n^2)$ $O(n^2)$ $O(nN \log N)$
TR	n $O(n^2)$ $O(n^4 \log n)$	n	n $O(n^4)$ $O(n^4 \log n)$	$\Omega(n^2)$ $O(n^4)$ $O(n^4 \log n)$	$\Omega(n^2)$ $O(n^4 M^3)$ $O(n^4 M^3 \log N)$	$O(N^6 \log N)^{(\dagger)}$
TRS	n $O(n^2)$ $O(n^5 \log n)$		n $O(n^5)$ $O(n^5 \log n)$	$\Omega(n^2)$ $O(n^5)$ $O(n^5 \log n)$	$O(n^5 M^4)$ $O(n^5 M^4 \log N)$	

Figure 2: 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. ($^{(*)}$Alt et. al [1], $^{(\dagger)}$Sprinzak[17])

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 $(a_i, b_j) \in \mu$. All of our algorithms have the feature that they find *all* matches; we know of no methods to find a single match any faster in the worst case than finding all matches. Our results are summarized in the table of Figure 2.

Some Notation. We define $\beta_r(p)$ to be the circular disk (ball) of radius r centered at p; if p is the origin, then we simply write β_r. We abuse notation by using $\beta_r(p)$ also to denote the circle of radius r about p (i.e., the boundary of the disk); the meaning should be clear from context.

We say that $\mathcal{R} = \{B_1, \ldots, B_n\}$ are *linearly separable* if for each pair of regions there exists a line separating that pair. We say that $\mathcal{R}$ 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 (implying that no two regions can be closer than distance 2ϵ).

We say that a set of simple polygons is *normal* if their convex hulls are pairwise-disjoint, each hull has an aspect ratio bounded above by a constant, ρ_1, and the ratio of the largest diameter to the smallest diameter is bounded above by a constant, ρ_2. (The *aspect ratio* of a convex figure is the ratio of its diameter to its width.) A simple but important property of normal shapes is the following: The number of shapes that intersect a disk of some fixed radius r is bounded above by some number, $K(r, \rho_1, \rho_2)$, depending on r, ρ_1, and ρ_2. (Note too that normal shapes are necessarily linearly separable.)

When we speak of the pure rotation problem ($\mathcal{T} = R$), we will assume without loss of generality that the center of rotation is the origin. We define the *orbit* of point a_i to be the circle, $\beta_{||a_i||}$, centered at the origin of radius equal to $||a_i||$, the Euclidean length of a_i.

Organization. The remainder of this paper is organized as follows: In Section 1 we discuss lower bounds on the number of matches under translation, and pure rotation. In Section 2 we give upper bounds on the number of matches in several cases. Section 3 describes algorithms for finding matches. We conclude in Section 4 with some discussion and open problems.

1 Lower Bounds on the Number of Matches

In this section we describe the construction of lower bounds of various point matching problems. Arkin et al.[2] and Sprinzak[18] 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 ($\mathcal{T} = T$), Alt et al.[1] showed that there is at most one match when the noise regions are unit circles. Thus, the trivial lower bound of 1 is a tight bound in this case. (In the next section we generalize this to the case of noise regions that are linearly separable.)

For the case of pure rotation, we obtain some simple lower bounds:

Theorem 1 *Let $\mathcal{R} = \{B_1, \ldots, B_n\}$ be an arbitrary set of n pairwise-disjoint regions. Then, there exists a set of n points $\mathcal{A}$ and a placement of the regions $\mathcal{R}$ so that, in the case of pure rotation ($\mathcal{T} = R$), there can be n distinct matches.*

Proof. Let C be a large circle centered at the origin, and let the points $\mathcal{A}$ be evenly spaced along C. Place the regions B_j 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 $\mathcal{A}$. $\Box$

Remark. In Theorem 6 we show this bound to be tight when the noise regions are circles.

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) $\mathcal{R} = \{B_1, \ldots, B_n\}$. There can be $\Omega(n^2)$ distinct matches.*

Proof. We describe the construction of an example in which there are a quadratic number of distinct matches between points and unit squares. Figure 3 shows the macroscopic placement of the squares: There are $n/2$ "inner" squares with centers equally distributed about the circle β_r of radius r (very large), and there are $n/2$ "outer" squares with centers equally distributed about the slightly larger circle $\beta_{r+1+\epsilon}$, 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. (It is helpful to think of the gap size as actually being zero, so that the squares are abutting.)

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

Figure 4 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 $\beta_{r'}$ (whose curvature is exaggerated in the figure). We label the arcs of intersection of $\beta_{r'}$ with the squares

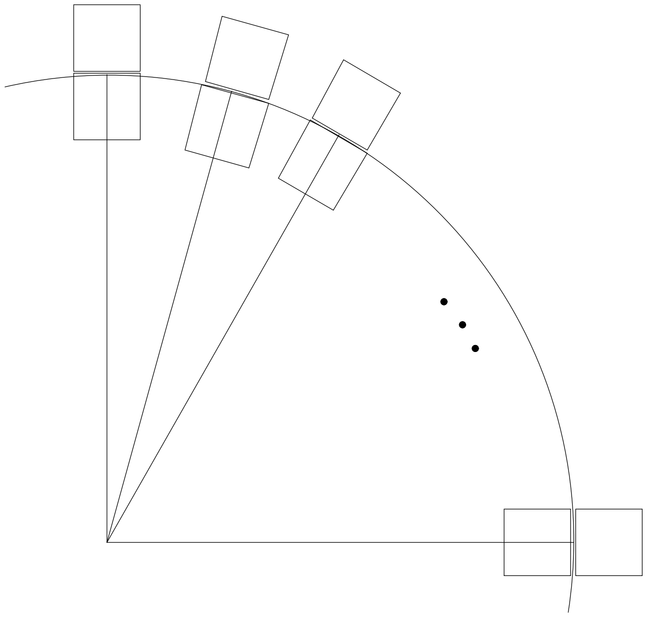

Figure 3: Lower bound construction: Unit squares at arbitrary orientations

as follows: U_kV_k is the first arc of intersection with B_{2k-1}, W_kX_k is the arc of intersection with B_{2k}, and Y_kZ_k is the second arc of intersection with B_{2k-1}. Note that, for appropriate choices of ϵ and ϵ', the arc lengths $|U_kV_k|$ and $|Y_kZ_k|$ are less than the arc length $|W_kX_k|$.

For the first pair of squares we place point a_1 in square B_1, on arc U_1V_1, at arc length $\delta < |U_1V_1|/n$ from W_1. We place point a_2 in square B_2, on arc W_1X_1, at arc length 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_kV_k, at arc length $k\delta$ from W_k, and we place point a_{2k} within square B_{2k}, on arc W_kX_k, at arc length $k\delta$ from Y_k.

A (clockwise) rotation of $\mathcal{A}$ by an amount δ will cause a_1 to move to B_2 and a_2 to B_1, leaving the matching of all other points unchanged. A further rotation by δ will cause a_3 and a_4 to switch squares while all other points remain unchanged. Continuing in this way $n/2$ 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. $\square$

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). An answer is given by the lower bound in the following theorem, together with an (almost matching) upper bound given in the next section (Theorem 7).

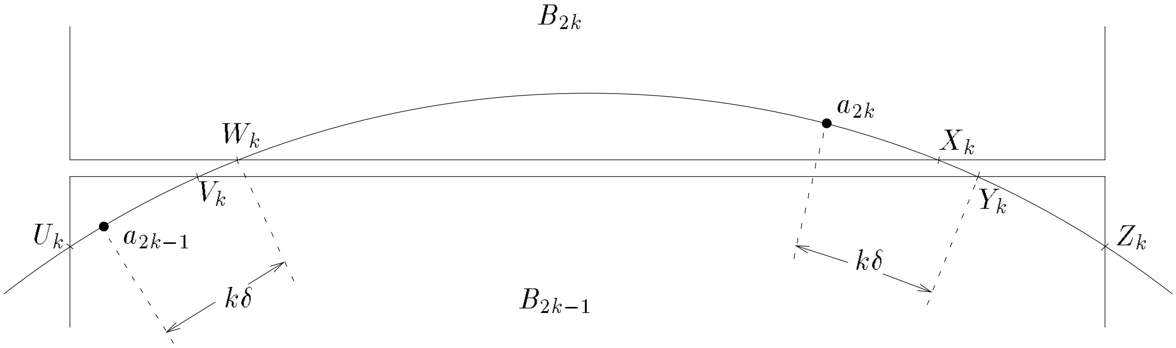

Figure 4: Microscopic placement of points

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 $5n - 16 = \Omega(n)$.*

Proof. We prove the theorem by the following construction. Centered on each of the four axes we place a pair of essentially touching squares, as in the proof of Theorem 2, such that a circle $\beta_{r'}$ of radius r' (centered at the origin) 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, for indices $i = 1, 2, 3, 4$ corresponding to the $+y$-, $+x$-, $-y$-, and $-x$-axes, respectively.

We first describe how we place the points $\mathcal{A}$ with respect to these eight squares. In the pair of squares intersected by the positive y-axis we place point a_1 on the arc U_1V_1 at arc length δ from W_1, and we place point a_2 on arc W_1X_1 at arc length δ from Y_1. We choose δ to be small enough so that $4\delta < |U_1V_1|, |W_1X_1|$, and we denote the arc length from a_1 to a_2 by $\Delta\theta$. (The separation $\Delta\theta$ will be determined by n and other constraints of the construction.) On the pair intersected by the positive x-axis we put a point on the arc U_2V_2 at arc length 2δ from W_2, and a point on arc W_2X_2 at arc length 2δ from Y_2. Similarly we places pairs of points on $\beta_{r'}$ straddling the $-y$- and $-x$-axes, using offset arc lengths of 3δ and 4δ, respectively.

The remaining $n - 8$ points of $\mathcal{A}$ are distributed along the circle $\beta_{r'}$, one quarter of the points in each quadrant, such that the points are equally spaced (with spacing $\Delta\theta$) within each quadrant, with the exception of one interval within each quadrant. (We assume that n is divisible by 4.) Specifically, the arc lengths of $a_{\frac{in}{4}}a_{\frac{in}{4}+1}$ are $\Delta\theta - \delta$ (for $i = 1, 2, 3$) and $\Delta\theta + 3\delta$ (for $i = 4$). We refer to these four intervals as *special*. The parameter $\Delta\theta$ is selected in order for the sum of all arc lengths to equal the circumference of $\beta_{r'}$. Refer to Figure 5.

We place the remaining $n - 8$ squares so that each square intersects $\beta_{r'}$ in a single interval, roughly of arc length $\Delta\theta$. The squares are placed essentially abutting, with their centers alternately inside and outside the circle $\beta_{r'}$. We make certain that each of the $n - 8$ points of $\mathcal{A}$ that we distributed along $\beta_{r'}$ is within a square, at arc distance greater than 4δ from the boundary of the square. It is easy to see that, with an appropriate choice of r', we can indeed place the squares in such a manner; see Figure 6.

We now count the number of possible matches in this construction. One match corresponds to

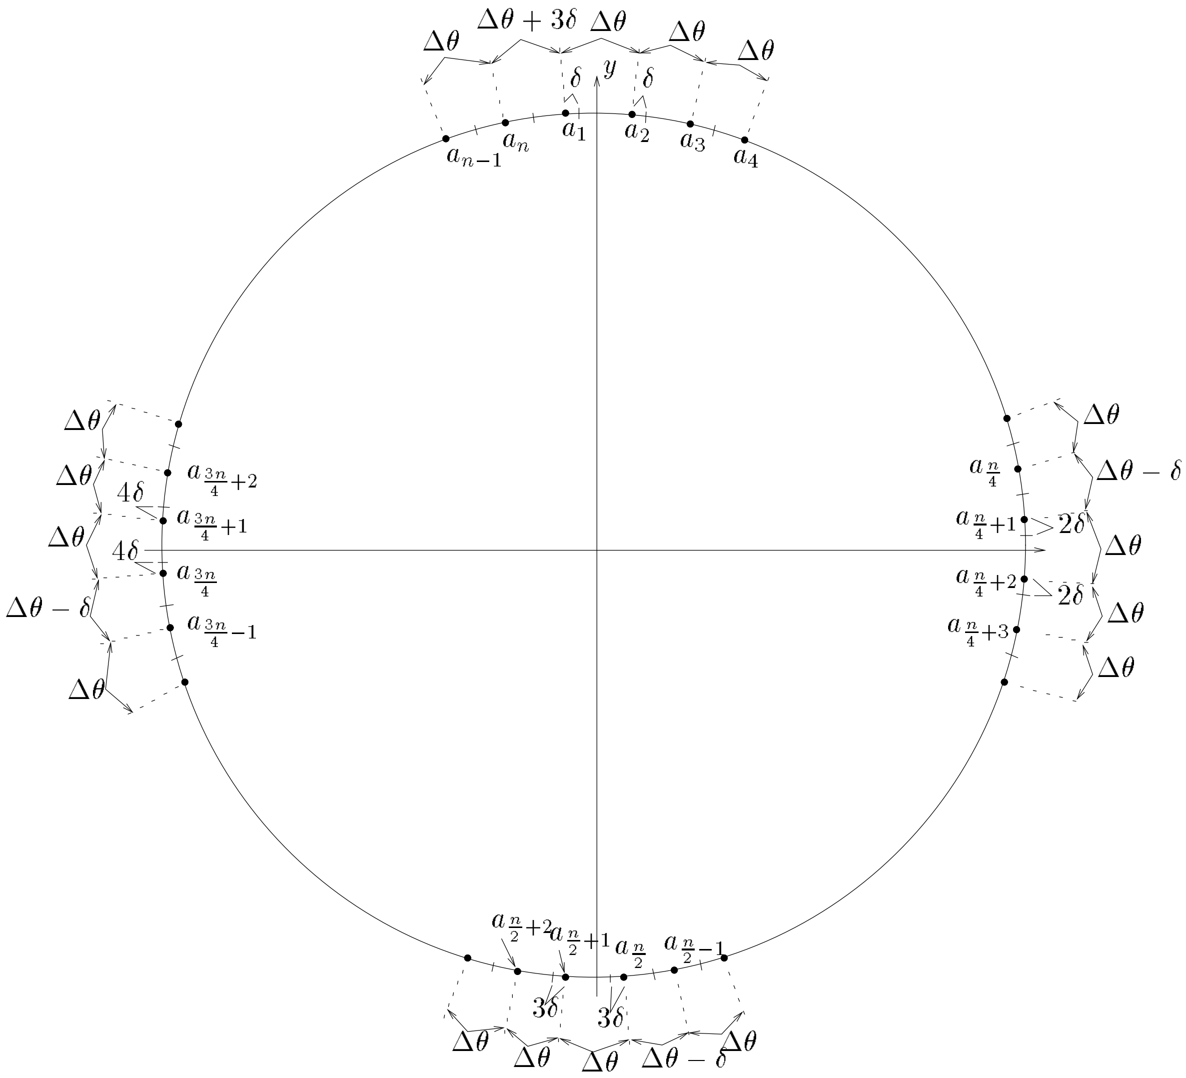

Figure 5: Positioning of points $\mathcal{A} = \{a_1, \ldots, a_n\}$ around the circle $\beta_{r'}$. The small slashes along $\beta_{r'}$ represent the boundaries of the squares $\mathcal{R}$. The points $\mathcal{A}$ are almost evenly distributed about $\beta_{r'}$, with the exception of 4 intervals.

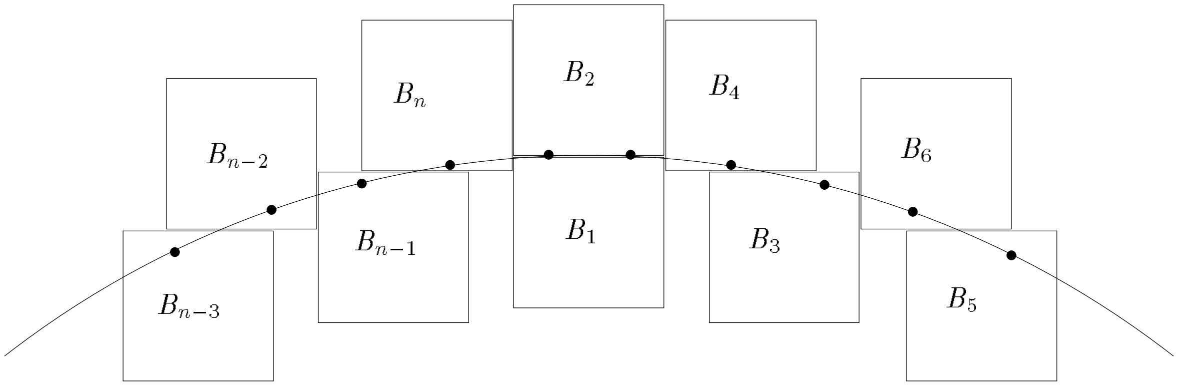

Figure 6: Aligned squares: macroscopic view

the original placement of $\mathcal{A}$, as we have just described. If we rotate $\mathcal{A}$ clockwise by an amount δ, the pair of points a_1 and a_2, which initially straddle the $+y$-axis, interchange squares, while all of the other points remain in their original squares. This yields a second match. Another incremental rotation by δ causes the points $a_{\frac{n}{4}+1}$ and $a_{\frac{n}{4}+2}$ to interchange squares, while leaving all other points a_i within their respective squares. Similarly, we can cause the points of $\mathcal{A}$ that straddle the $-y$- and $-x$-axes to interchange squares. In total, then, by microscopic rotations of $\mathcal{A}$, we get 5 distinct matches corresponding to the original "macro" placement of the points. 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) claimed bound of $5n - 16$, so we must be more careful to see the origin of the "-16" term. A macroscopic placement of $\mathcal{A}$ yields 5 distinct matches (via microscopic motions) *except* when the four intervals of $\mathcal{A}$ points being stabbed by the axes are the special intervals. It is not hard to see that when the special intervals straddle the axes, there can be only a single match (since they do not all have arc length $\Delta\theta$, as do all other intervals). There are four macroscopic positions in which the special intervals are straddled, each yielding a single match instead of 5 matches. Thus, there are $5n - 16$ distinct matches in all for our construction. $\square$

2 Upper Bounds on the Number of Matches

In this section we prove upper bounds on the number of possible matches for several cases of the matching problem, beginning with the case of pure translation. Alt et al. [1] 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 to the case of linearly separable noise regions. 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 and a set $\mathcal{R} = \{B_1, \ldots, B_n\}$ of n linearly separable regions under pure translation.*

Proof. The proof is an extension of the proof in Alt et al.[1]. 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 (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.

(b). "$<_t$" is "transitive" in the following sense: if $B_1 <_t B_2 <_t \ldots <_t B_k$ and if 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. Assume now that the claim holds for integers $3, \ldots, k-1$. Consider a set of regions satisfying $B_1 <_t B_2 <_t \ldots <_t B_k$. If there exists a line parallel to t stabbing B_1 and B_k, then we can move it (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 (1). □

We now turn to pure rotation, and consider the case of linearly separable noise regions (which clearly includes all convex cases of noise regions). The following theorem shows that the lower bound we showed in Theorem 2 is tight:

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. The orbit of a point $a_i \in \mathcal{A}$ may enter and leave any one region B_j many times (in fact, arbitrarily many times). However, the fact that the regions $\mathcal{R}$ are linearly separable implies that the orbit of a_i cannot visit a subsequence of regions of the form "$\ldots B_k \ldots B_\ell \ldots B_k \ldots B_\ell \ldots$" ($k \neq \ell$). This implies that during a full rotation of $\mathcal{A}$, the region that contains a_i can switch at most $2n - 1$ times (e.g., since $2n - 1$ is the maximum length of an order 2 Davenport-Schinzel sequence; see Hart and Sharir[10]). Applying this analysis to each of the n points yields the $O(n^2)$ bound. □

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 Theorem 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 at most n (and this bound is tight in the worst case).*

Proof. Assume first that no circle covers the origin. Define a *cut path* π as follows. 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 corresponding arcs of B_j.

If there exists a circle B_i that covers the origin, then we modify the definition of the cut π. We define π to go along a straight segment from the origin to the closest point on the boundary of B_i, then to go along a semi-circular arc of the boundary of B_i to the point of B_i farthest from the origin, and then to continue, as in the case of no circle covering the origin, along a ray away from the origin, going around any circles that are encountered by taking the shorter of the two corresponding arcs.

Note that π is a radially monotone path (i.e., monotone in the distance, r, from the origin). This implies that as we rotate the points $\mathcal{A}$ about the origin, each point will pass over π exactly once. Thus, with respect to rotations of $\mathcal{A}$, there are exactly n homotopically distinct positions of the cut π. (We think of "cutting" the plane along π; intuitively, this gets rid of the "wrap-around" effect of rotational motion.)

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 at the origin) goes from B_i to B_j without crossing (or touching) π. It is easy to check that the relation "$<_c$" satisfies the same two properties as did "$<_t$" in the proof of Theorem 4; namely, "$<_c$" is anti-symmetric and "transitive". As a result, for each of the (at most) n positions of the cut π with respect to $\mathcal{A}$, there can be at most one possible match. □

It is interesting to ask where, in the above proof, did we use the assumption that the noise regions are circles? What we needed was that the intersection of any noise region with any circle centered at the origin must consist of at most one arc. This implies that the cut π is radially monotone, since it always follows the boundary of a noise region from the point closest to the origin to the point furthest from the origin. This property of noise regions is obeyed by a broader class of shapes than just circles. When this property fails, as in the case of arbitrary unit squares, we have already seen that there can, in fact, be a quadratic number of distinct matches (Theorem 2). However, we now show that if the unit squares are *aligned*, then we can still get an $O(n)$ bound on the number of possible matches:

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 distinct matches is at most $5n$.*

Proof. If none of the noise regions $\mathcal{R}$ are cut by the coordinate axes, then any circle centered at the origin intersects each square in a single arc, so, by the discussion above, the exact same proof goes through as in the case of circular noise regions. Thus, we assume that some of the squares $\mathcal{R}$ are cut by the coordinate axes; we say that such squares are "bad".

Similar to before, we define a cut π as follows. From the origin, path π initially goes straight up; when a square is encountered, π turns left, following the bottom of the square until its lower left corner; then π goes up again, etc, until finally π goes to infinity along a vertical ray. It is easy to see that π is radially monotone (monotone in r); thus, there are exactly n positions of $\mathcal{A}$ that are homotopically distinct with respect to π.

Because of the bad squares, however, each distinct position of the points $\mathcal{A}$ with respect to the cut π may give rise to more than one match. But, *because the squares are all of equal size*, the orbit of each point $a_i \in \mathcal{A}$ can intersect at most four bad squares. Hence, each match that occurs without a change in the homotopic position of $\mathcal{A}$ with respect to π, can be charged to a (point, bad square)-pair. Thus, there are at most $4n$ such "additional" matches that occur while having no point of $\mathcal{A}$ cross the cut π. In total, then, we get an upper bound of $5n$. □

Remark. 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. Also, it is an open question to determine the maximum number of possible matches for aligned *unequal* size squares. (The difficulty comes from the fact that there may be many more than four bad squares per orbit of a point a_i.)

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 distinct matches is $O(n^2)$.*

Proof. The proof is based on the following six 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, $-\frac{t}{\alpha}$. □

Lemma 10 *Under scale alone ($\mathcal{T} = S$, i.e., $\tau(p) = \alpha p$ for some scalar α), 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. □

Remark. Note that Lemma 10 fails for the case of aligned unit squares.

Lemma 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 (new) origin, and thus, by Lemma 10, the claim follows. □

Lemma 12 *Given a set $\mathcal{A} = \{a_1, \ldots, a_n\}$of n points in the plane. Let $a^* = \frac{1}{n}\sum_i a_i$ 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 $\tau_0(a_i)$.

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

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

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

□

Lemma 13 *Assume that the noise regions $\mathcal{R} = \{B_1, \ldots, B_n\}$ are convex. Then the set of all rotation parameters θ for which there exists a similarity transformation τ that maps point a_i to noise region B_i, for $i = 1, \ldots, n$, corresponds to a (connected) subarc of β_1.*

Proof. Consider a similarity transformation $\tau(a) = \alpha R_\theta a + t$, where $t = (x, y)$ denotes a translation vector, $\alpha > 0$ denotes scale, and R_θ denotes the rotation matrix corresponding to rotation by angle θ. (We can assume that $\alpha > 0$, since there is no similarity transformation with $\alpha = 0$ that maps $n \geq 2$ points a_i to disjoint noise regions B_i, for all $i = 1, \ldots, n$.) We identify values of θ with points on the unit circle β_1; specifically, $\theta \leftrightarrow (\cos\theta, \sin\theta)$. We consider τ to be a point in $\Re^4$ with 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 $Q = \{\tau \in \Re^4 \mid \tau(a_i) = \alpha R_\theta a + t \in B_i \quad i = 1, \ldots, n\}$ is a convex subset of $\Re^4$ (as in Baird[6]). Let Θ denote the set of all rotations θ for which there exists a similarity transformation $\tau \in Q$ with rotation matrix R_θ. We need to show that the subset of β_1 corresponding to Θ is connected.

Consider two points $\theta_1, \theta_2 \in \Theta$, and let $\tau_1, \tau_2 \in Q$ be corresponding similarity transformations. The convexity of Q implies that $\tau_\lambda = \lambda\tau_1 + (1-\lambda)\tau_2 \in Q$, for $\lambda \in [0, 1]$. We denote the coordinates of τ_λ by $(x(\lambda), y(\lambda), r_1(\lambda), r_2(\lambda))$. Letting $\alpha(\lambda)$ and $\theta(\lambda)$ denote the scale and rotation parameters for τ_λ, we see that

$$
\begin{aligned}
r_1(\lambda) &= \lambda\alpha_1\cos\theta_1 + (1-\lambda)\alpha_2\cos\theta_2 &&= \alpha(\lambda)\cos\theta(\lambda), \\
r_2(\lambda) &= \lambda\alpha_1\sin\theta_1 + (1-\lambda)\alpha_2\sin\theta_2 &&= \alpha(\lambda)\sin\theta(\lambda), \\
\alpha(\lambda) &= \sqrt{r_1^2(\lambda) + r_2^2(\lambda)}.
\end{aligned}
$$

The value $\theta(\lambda)$ lies in set Θ (by definition of Θ), and is a continuous function of λ, as seen by writing

$$\sin\theta(\lambda) = \frac{r_2(\lambda)}{\sqrt{r_1^2(\lambda)+r_2^2(\lambda)}}, \qquad \cos\theta(\lambda) = \frac{r_1(\lambda)}{\sqrt{r_1^2(\lambda)+r_2^2(\lambda)}}.$$

Then $\{(\cos\theta(\lambda), \sin\theta(\lambda)) \mid \lambda \in [0,1]\}$ gives a (continuous) path in β_1, connecting $(\cos\theta_1, \sin\theta_1)$ with $(\cos\theta_2, \sin\theta_2)$. $\square$

Remark. In fact, the statement of the above lemma can be strengthened to say that the set of similarity transformations τ that map each point a_i to region B_i is a connected subset of the (cylindrical) four-dimensional manifold that represents the set of all similarity transformations embedded in the five-dimensional space parameterized by $(x, y, \alpha, \cos\theta, \sin\theta)$.

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 $\mathcal{R} = \{B_1, \ldots, B_n\}$. Let b^* denote the center of mass of the centers, $b_1, \ldots, b_n$ of the circles $\mathcal{R}$. 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

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

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

$\square$

Remark. The above lemma can be generalized to the case in which the noise regions are normal. If we let b_i denote the center of the (smallest) circumscribing circle for the convex hull of B_i, let b^* denote the center of mass of the b_i's, and let r_{max} denote the maximum radius of all such circles, then we see that $\|\tau(a^*) - b^*\| \leq r_{max}$.

Proof of Theorem 8:

We continue to use the notation that a^* denotes the center of mass of the points $\mathcal{A}$ and that b^* denotes the center of mass of the points $\mathcal{B}$, which are the centers of circles $\mathcal{R}$. We assume (without loss of generality) that a_1 is the point of $\mathcal{A}$ farthest from a^*.

To prove the theorem, it suffices to show that if the point a_1 is constrained to have its image, $\tau(a_1)$, be inside one particular 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 radius $\frac{1}{3}\epsilon_0$. 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) is constant (for a fixed B_i).

By Lemma 12, if $\tau(a^*)$ and $\tau(a_1)$ lie within C_1 and C_2 respectively, which are of radius $\frac{1}{3}\epsilon_0$, then each point $\tau(a_i)$ is constrained to be within a circle of radius ϵ_0. By our choice of ϵ_0, a circle of radius ϵ_0 can intersect at most two unit circles in $\mathcal{R}$. Hence, during this motion, $\tau(a_i)$ can be in at most two different noise regions (unit circles). We define a bipartite graph, whose "left" set of nodes corresponds to the set $\mathcal{A}$, and whose "right" set of nodes corresponds to $\mathcal{R}$. An edge connects a point-node, a_i, with a circle-node, B_j, if there exists a transformation $\tau \in \mathcal{T}$ of $\mathcal{A}$, satisfying the constraints $\tau(a_1) \in C_1$ and $\tau(a^*) \in C_2$, such that $\tau(a_i) \in B_j$. Note that, since a_i remains within a circle (of radius ϵ_0) that can intersect at most two circles from $\mathcal{R}$, each left node in the bipartite graph is of degree at most two.

We need to bound the number of possible matches in this bipartite graph. As in Alt et al.[1], we proceed to solve the matching problem in the bipartite graph, as follows. 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.

We claim that, if there exists a perfect matching, then each node in the resulting graph will in fact have degree exactly two. To see this, consider a connected component that has m left nodes. If there exists a perfect matching, it must also have exactly m right nodes. Now, each left node has degree exactly two (since left nodes have degree at most two), so there can be at most $2m$ edges in the connected component. This implies that each right node has degree exactly two, implying that the connected components are cycles.

Consider one such cyclic component. It has exactly two possible matches — denote them by μ_1 and μ_2. Hence, there are *at most* two valid matches between the corresponding points and circles. Let θ denote the parameter of rotation of the point set $\mathcal{A}$, and define I_1 (resp., I_2) to be the set of angles θ for which the match μ_1 (resp., μ_2) is obtained for some choice of scale and translation. We identify values of θ with points of β_1. Then, Lemma 13 implies that I_1 (and, similarly, I_2) consists of a single (possibly empty) arc of the circle β_1. Lemma 11 implies that the two arcs I_1 and I_2 are disjoint.

Thus, the union of the endpoints of arcs I_1 and I_2 for all connected components yields an arrangement of $O(n)$ points on the circle β_1. These $O(n)$ points partition β_1 into arcs, each of which corresponds to at most one perfect matching. □

Remark. 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 Lemma 10 (and, hence, Lemma 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 have degree at most two. This leaves open the question of bounding the number of matches under similarity in the case of aligned unit square noise regions.

3 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 previously known[1, 2] for the general case.

We begin with noise regions that are linearly separable polygons (not necessarily normal) and the allowed motion is pure rotation. Theorem 5 showed that there can be at most $O(n^2)$ matches. We now see that these can be computed in time $O(nN \log N)$, where N is the total number of edges of all 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 polygonal 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. Each of the N points where the orbit of a_i crosses the boundary of some noise region corresponds to a "critical" value of the rotation parameter θ. We think of the critical values of θ as being points on the unit circle β_1. In time $O(nN \log N)$ we sort all $O(nN)$ critical values. Then, in time $O(nN)$, we advance clockwise around β_1, checking for each interval of θ whether or not there exists a corresponding perfect matching. We keep track of a matching vector, $\mu(1, \ldots, n)$, whose ith entry indicates which region contains the rotated image of point a_i at the current value of θ. (An empty ith entry indicates that the image of a_i lies in no region.) Each time we advance through a critical value of θ, some point a_i either enters or leaves some region B_j, so we update the vector μ accordingly. The perfect matchings correspond to values of μ for which there are no empty entries. $\square$

Note that the above algorithm applies also to the case in which the n noise regions are circles of possibly different radii; the time bound is then $O(n^2 \log n)$. For the case of *unit* circles we will give an improved algorithm below. We will need the following lemma:

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 $\mathcal{R} = \{B_1, \ldots, B_n\}$. Given a correspondence between points and circles, the set of all rotations that achieve this correspondence can be found 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. Let I_i denote the set of rotation angles θ for which the rotated image of point a_i lies in circle B_i; each such I_i corresponds to an arc on the unit circle, β_1. We need to determine whether or not $I_1 \cap \cdots \cap I_n = \emptyset$. We do this by incrementally computing $I_1 \cap \cdots \cap I_i$. Note that since at most one circle can contain the origin, at most one of the arcs is of length greater than π; let this arc, if it exists, be I_n, the last one added to the intersection set. Then, the set $I_1 \cap \cdots \cap I_i$ is always an arc of length less than π, so its intersection with I_{i+1} is always a single arc or empty, for $i = 1, \ldots, n-1$, $n > 1$. This implies that this incremental construction requires $O(n)$ time. $\square$

Remark. Alt et al.[1] give an $O(n \log n)$ algorithm for the case of matching under pure rotation when the noise regions are allowed to overlap arbitrarily; Iwanowski [15] shows a lower bound of $\Omega(n \log n)$ for this general case.

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. We identify points of the circle β_1 with values of the rotation parameter θ. We sort the points $\mathcal{A}$ according to distance from the origin, with a_1 being the point farthest from the origin. We construct the *cut path* π, as defined in the proof of Theorem 6. (This is done in time $O(n)$, after having sorted the points $\mathcal{A}$ in time $O(n \log n)$.)

Those values of θ for which the rotated image of a_1 lies in some circle of $\mathcal{R}$ form a collection of $O(n)$ *primary intervals*, which correspond to intersections between the orbit of a_1 and $\mathcal{R}$. The primary intervals are easily found in time $O(n)$. If the set of primary intervals is empty, then there can be no perfect matching; thus, we assume that there is at least one primary interval.

Consider the orbit of a_1 to be partitioned into pieces of arc length $\epsilon_0 = \frac{2}{3}\sqrt{3} - 1$. For rotations of $\mathcal{A}$ that keep the image of a_1 within one such piece, each other point a_i $(i \geq 2)$ must execute a motion that lies within a ball of radius ϵ_0 (since a_1 is the point furthest from the center of rotation). By the choice of ϵ_0, this implies that each point a_i will intersect at most two different noise regions (disjoint unit circles) during a motion that keeps a_1 within a single piece. There are only a constant number of pieces corresponding to one primary interval. Thus, we can subdivide the primary intervals into $O(n)$ *primary subintervals*, according to the partitioning into pieces.

There are exactly $n-1$ values of θ that correspond to the orbit of some point a_i $(i \geq 2)$ crossing the cut path π, since π is radially monotone. These θ values, which can be determined in $O(n \log n)$ time, induce a refinement of the primary subintervals into $O(n)$ *basic intervals*. Furthermore, by the arguments of Theorem 6, and by the choice of these refining θ values, we see that there can be at most one match per basic interval.

Now we proceed as in the proof of Theorem 8: For each of the $O(n)$ 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 there exists a rotation θ within the basic interval such that the (rotated image of) point a_i is in circle B_j. By the construction of the basic intervals, this bipartite graph has left nodes with degree at most 2. This will permit a simple matching algorithm, as described below.

First, we show that these bipartite graphs can be constructed efficiently. For one basic interval, we build the graph in the straightforward way, in time $O(n^2)$. In order to update the graph when we rotate $\mathcal{A}$ so that the image of a_1 moves from one basic interval to the next one (going, say, clockwise), we must examine the corresponding motion of each point a_i $(i \geq 2)$ along its orbit. We do this simply by "walking" along a_i's orbit from its subarc corresponding to the first basic interval to its subarc corresponding to the next basic interval, noting any crossings with circles $\mathcal{R}$. Since the walk is monotonic, each crossing with a circle will be encountered only once. This implies an overall time bound of $O(n^2)$, provided we have the intersection points with circles $\mathcal{R}$ in sorted order along each orbit. This is easily accomplished in $O(n)$ time per orbit, given an $O(n \log n)$ preprocessing step in which we sort the centers of circles $\mathcal{R}$ by angle about the origin: simply check each circle, in order, for intersection with a given orbit.

In order to solve the matching problem on each bipartite graph, we proceed as in the proof of Theorem 8: We delete a node of degree one, and the corresponding incident edge, and repeat this process until we are left with a graph whose nodes all have degree at least 2. By arguments in the proof of Theorem 8, this implies that the connected components are cycles, if there exists a matching. There are two possible graph matchings per cycle. Using the algorithm described in Lemma 16 for the case of labeled point matching, we can determine for each possible graph matching the (possibly empty) interval of θ for which there exists a match between the corresponding points and circles; this requires time linear in the size of the cycle. (The θ interval will be a subset of the basic interval.) For each cycle, at most one of the two possible graph matchings will have a nonempty corresponding θ interval (for the same reasons that there is at most one match per basic interval). If any cycle has both θ intervals empty, then there can be no perfect matching of points to circles for the given basic interval. Thus, assume that each cycle has exactly one nonempty θ interval. We can then intersect these θ intervals to obtain the set of all rotations that yield a perfect matching for the given basic interval. The total time per basic interval is $O(n)$, completing the proof. □

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 $O(n^4 \log n)$. Furthermore, the number of matches is $O(n^4)$.*

Proof. Let a_1 be, without loss of generality, the point in $\mathcal{A}$ farthest from the center of mass, a^*. By Lemma 14, in any match, the center of mass, a^*, of the points $\mathcal{A}$, must lie in $\beta_1(b^*)$, the unit circle around the center of mass of the centers of the circles $\mathcal{R} = \{B_1, \ldots, B_n\}$.

We fix a circle B_i and constrain $\tau(a_1)$ to be in B_i. There is very little freedom of motion left for the rest of the points of $\mathcal{A}$. In fact, Lemma 12 implies that (the transformation of) each point of $\mathcal{A}$ lies within a circle of radius 3, as we rotate and translate $\mathcal{A}$ subject to the constraints $\tau(a_1) \in B_i$ and $\tau(a^*) \in \beta_1(b^*)$. Specifically, $\tau(a_j)$ must lie within circle $\beta_{3\epsilon}(\tau_0(a_j))$, where τ_0 denotes the similarity transformation that maps a^* to b^* and a_1 to the center of B_i. Note that a circle of radius 3 can intersect only a *constant* number of the noise regions $\mathcal{R}$; the constant is certainly no more than 25, which is the ratio of the area of a circle of radius $3 + 2 = 5$ to the area of a unit circle. For each point a_j we determine and store the (constant number of) noise regions that intersect $\beta_{3\epsilon}(\tau_0(a_j))$; this can be done naïvely in time $O(n)$ per choice of B_i and a_j.

Let us fix attention on one pair of points, a_j and a_k, of $\mathcal{A}$. There are only a constant number of pairs of circles that can be matched to a_j and a_k. Let $B_{j'}$ and $B_{k'}$ be one such pair. If we constrain $\tau(a_j)$ and $\tau(a_k)$ to lie on the boundaries of $B_{j'}$ and $B_{k'}$, respectively, then there remains only one degree of freedom in the motion of $\mathcal{A}$. We let t denote the parameter of this motion. Then, each point a_ℓ ($\ell \neq j, k$) must be mapped onto an algebraic curve of low degree in t (degree at most 6). Such a curve can intersect a circle at most 12 times. Thus, subject to the constraints on points a_j and a_k, $\tau(a_\ell)$ executes a motion along a curve that has at most a *constant* number of crossings with boundaries of each of a *constant* number (at most 25) of noise regions. Since we have determined already the noise regions that intersect $\beta_{3\epsilon}(\tau_0(a_\ell))$, we can compute the values of t that correspond to crossings in constant time. We call the values of t corresponding to crossings the *critical* values of t.

In all, there are $O(n)$ critical values of t corresponding to the points a_ℓ, $\ell \neq j, k$. We sort the critical values of t (in time $O(n \log n)$) and examine the corresponding critical intervals of t. Within each such interval, there can be at most one match. (This is still subject to the constraint that $\tau(a_j)$ and $\tau(a_k)$ lie on the boundaries of $B_{j'}$ and $B_{k'}$, respectively.) Then, proceeding as in Theorem 15, we march through the sorted list of critical values of t, keeping track of a matching vector, $\mu(1, \ldots, n)$, and output each of the matches we find.

In summary, there are $O(n)$ choices for B_i and $O(n^2)$ choices for the pair a_j and a_k. For each of these $O(n^3)$ choices, we solve the resulting matching problem in $O(n \log n)$ time. Thus, in time $O(n^4 \log n)$ we find all matches for which two points of $\mathcal{A}$ lie on boundaries of noise regions. For any match of points to noise regions, we can always rotate and translate $\mathcal{A}$, maintaining the same point-region correspondence, so that two points lie on noise region boundaries. Thus, our algorithm finds all valid matches.

The upper bound of $O(n^4)$ on the number of matches follows immediately from the fact that there is at most one match per critical interval of t, and there are $O(n)$ critical intervals per choice of j, k, and ℓ. □

Remark. The best known algorithm for determining if a matching exists under rotation and translation, *for a given correspondence between points and noise regions*, requires time $O(n^3 \log n)$ [14]. We know that there can be as many as $\Omega(n)$ matches with distinct correspondences. If we were presented with these correspondences, one at a time, and checked each for being a matching under rotation and translation, we would spend time $O(n^4 \log n)$ to *verify* the correctness of the given correspondences. Our algorithm above achieves this same time bound for the (apparently harder) problem of *finding* all matches.

An immediate extension of Theorem 18 allows the transformation τ to be a similarity: The only change to the algorithm is that, in order to obtain a single degree of freedom, we now need to constrain *triples* of points of $\mathcal{A}$ to the boundaries of a constant number of triples of noise regions. This has the effect of increasing the time complexity by a factor of n:

Theorem 19 *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 $\mathcal{R} = \{B_1, \ldots, B_n\}$. There exists an algorithm to find all matches in time complexity $O(n^5 \log n)$. Furthermore, the number of matches is $O(n^5)$.*

Theorem 18 can be generalized to handle noise regions that form a normal set of polygons (which include the cases of aligned or arbitrary unit squares), as we now argue. By the remark after Lemma 13, for any τ that yields a match, the image of the center of mass $\tau(a^*)$ must lie in a bounded region; specifically, $\|\tau(a^*) - b^*\| \leq r_{max}$, where r_{max} is the radius of the largest circumscribing circle for any B_j. Thus, if we constrain $\tau(a_1)$ to lie in noise region B_i (which is certainly contained within a disk of radius r_{max}), then $\tau(a_j)$ lies within a disk of radius $3r_{max}$, for each of the points $a_j \in \mathcal{A}$. By the defining properties of normal regions, such a disk can intersect at most some number, $K(3\epsilon_{max}, \rho_1, \rho_2)$, of noise regions (which depends only on the constants ρ_1 and ρ_2 that characterize the normal set of regions). Then, for each choice of a pair of points a_j and a_k, we see that each of the disks containing $\tau(a_j)$ and $\tau(a_k)$ intersects with at most $O(M)$ polygon edges, implying that there are $O(M^2)$ choices for a pair of edges to which we constrain the pair of points. For each of these $O(n^2M^2)$ choices of point-edge pairs, $\tau(a_\ell)$ is constrained to lie on a curve of bounded degree, which can intersect any one polygon edge at most a constant number of times. Since $\tau(a_\ell)$ must also lie within a disk that intersects only a constant number of noise regions, this implies that there are only $O(nM)$ critical values of the curve parameter t. Continuing as before, then, we solve the resulting one-degree-of-freedom problem in time $O(nM \log nM) = O(nM \log N)$. We have thus argued the following:

Theorem 20 *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 $\mathcal{R} = \{B_1, \ldots, B_n\}$, each having at most M edges and having a total of N edges altogether. There exists an algorithm to find all matches in time complexity $O(n^4M^3 \log N)$. Furthermore, the number of matches is $O(n^4M^3)$.*

The extension to similarity transformations is again immediate:

Theorem 21 *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 $\mathcal{R} = \{B_1, \ldots, B_n\}$, each having at most M edges and having a total of N edges altogether. There exists an algorithm to find all matches in time complexity $O(n^5M^4 \log N)$. Furthermore, the number of matches is $O(n^5M^4)$.*

We conclude this section with one further variation on our algorithms, for the case in which the noise regions are aligned unit squares that are ϵ-separated:

Theorem 22 *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 $O(n^2 \epsilon^{-4} \log n)$ and space complexity $O(n)$.*

Proof. The algorithm is similar to previous ones. As usual, without loss of generality, we assume that a_1 is the point of $\mathcal{A}$ farthest from a^*. From the remark after Lemma 14, we know that, for a transformation τ corresponding to a valid match, the (transformed) center of mass, $\tau(a^*)$, lies within the disk $C = \beta_{\frac{\sqrt{2}}{2}}(b^*)$ of radius $\frac{\sqrt{2}}{2}$ about b^*, the center of mass of the centers of the squares. (Here, $r_{max} = \frac{\sqrt{2}}{2}$.) For each square B_i, in turn, we consider the allowed motions of $\mathcal{A}$ when $\tau(a_1)$ is constrained to lie within B_i. We cover the circle C and the square B_i by small circles of radius $\frac{1}{3}\epsilon$. Consider a pair of (small) circles, C_1 and C_2, from the coverings of C and B_i, respectively. (Note that the number of choices for this pair is $O((1/\epsilon)^4)$.) We know by Lemma 12 that if we consider a transformation $\tau \in \mathcal{T}$ such that $\tau(a^*) \in C_1$ and $\tau(a_1) \in C_2$, then, for each point $a_j \in \mathcal{A}$, $\tau(a_j)$ lies in a disk of radius ϵ about point $\tau_0(a_j)$ (where τ_0 is the similarity transformation mapping a_1 to the center of C_1 and a^* to the center of C_2). As in the proof of Theorem 8, we build a bipartite graph, only this time the degree of a point node is at most *one*, since a disk of radius ϵ can intersect at most one of the ϵ-separated regions $\mathcal{R}$. Thus, there is at most one matching in the graph, which we can find in $O(n)$ 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 Imai, Sumino, and Imai[14].

It remains to be shown how to build each bipartite graph in time $O(n \log n)$. This is done by first constructing the Voronoi diagram of the squares $\mathcal{R}$ (treating each square as a "source"), which is done in time $O(n \log n)$ using, for example, the algorithm of Fortune [7]. We then determine the square of $\mathcal{R}$ that is closest to the point $\tau_0(a_j)$, for each j, by doing $O(n)$ point location queries, at a cost of $O(\log n)$ time each [16]. This allows us to determine the edges of the bipartite graph in time $O(n \log n)$ (for each choice of B_i). $\square$

4 Conclusion

Several directions for further study are suggested by our work:

1. 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 faster 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 region centers, fails to hold when $m \neq n$. This means that we cannot extend our proof of Theorem 8 to this case. However, it is not hard to extend the algorithms of Theorems 18, 19, 20, and 21: Instead of "pinning" the pair of points a^* and a_1, we pin the pair of points of $\mathcal{A}$ that are farthest apart. The resulting running times each increase by a factor of n.

2. It should be possible to extend many of our results to the case in which we allow *partial* overlap among the regions $\mathcal{R} = \{B_1, \ldots, B_n\}$. 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 desire complexity bounds in terms of n and K. The algorithms of Theorems 18, 19, 20, and 21 extend easily to this case. For example, the algorithm of Theorem 18 modifies so that the point $\tau(a_\ell)$ is constrained to lie on a curve that has a constant number of crossings with each of $O(K)$ other circles, for each of the $O(n^2K^2)$ choices for the pair of points a_j and a_k and the pair of circles to which they are pinned.

3. Questions are largely open in higher dimensions. For example, can one obtain nontrivial bounds on the number of matches under pure rotation (three degrees of freedom) in space? Assume, for example, that the regions $\mathcal{R} = \{B_1, \ldots, B_n\}$ are disjoint unit spheres. Our only lower bound on the worst-case number of matches is the (trivial) one — $\Omega(n)$. 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 matches are obtainable under pure translation? Unlike the corresponding 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.

Acknowledgments

We are very grateful to the referees for many useful suggestions that greatly improved the paper. 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 Eshkol grant from the Israeli Ministry of Science and Technology. J. Mitchell is partially supported by NSF grant ECSE-8857642, by a grant from Hughes Research Laboratories, and by Air Force Office of Scientific Research contract AFOSR-91-0328.

References

[1] H. Alt, K. Mehlhorn, H. Wagener and E. Welzl, 1988. Congruence, Similarity, and Symmetries of Geometric Objects, *Discrete and Computational Geometry 3*, 237–256.

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

[3] M.J. Atallah, 1985. A Matching Problem in the Plane, *J. Comput. Systems Sci. 31*, 63–70.

[4] M.D. Atkinson, 1987. An Optimal Algorithm for Geometric Congruence, *J. Algorithms 8*, 159–172.

[5] F. Aurenhammer, F. Hoffman and B. Aronov, 1991. Minkowski-Type Theorems and Least-Squares Partitioning, Technical Report, Institut für Informatik, Freie Universität Berlin, Germany.

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

[7] S.J. Fortune, 1987. A Sweepline Algorithm for Voronoi Diagrams, *Algorithmica 2*, 153–174.

[8] S. Hart and M. Sharir, 1986. Nonlinearity of Davenport-Schinzel Sequences and of Generalized Path Compression Schemes, *Combinatorica 6*, 151–177.

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

[10] P.J. Heffernan and S. Schirra, 1991. Approximate Decision Algorithms for Point Set Congruence, Technical Report MPI-I-91-110, Max-Planck-Institut für Informatik, Saarbrücken, Germany.

[11] D.P. Huttenlocher, 1988. Three-Dimensional Recognition of Solid Objects from a Two-Dimensional Image, Ph.D. thesis, Technical Report TR-1045, Artificial Intelligence Laboratory, MIT.

[12] D.P. Huttenlocher and K. Kedem, 1990. Computing the Minimum Hausdorff Distance for Point Sets Under Translation, in Proceedings of the 6th ACM Symposium on Computational Geometry, pp. 340–349.

[13] D.P. Huttenlocher, K. Kedem and M. Sharir, 1991. The Upper Envelope of Voronoi Surfaces and its Applications, in Proceedings of the 7th ACM Symposium on Computational Geometry, pp. 194–203.

[14] K. Imai, S. Sumino and H. Imai, 1989. Minimax Geometric Fitting of Two Corresponding Sets of Points, in Proceedings of the 5th ACM Symposium on Computational Geometry, pp. 266–275.

[15] S. Iwanowski, 1990. Approximate Congruence and Symmetry Detection in the Plane, Ph.D. thesis, Fachbereich Mathematik, Freie Universität, Berlin, Germany.

[16] D.G. Kirkpatrick, 1983. Optimal Search in Planar Subdivisions, *SIAM Journal on Computing 12*, 28–35.

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

[18] J. Sprinzak, 1990. Point Matching Algorithms, M.S. thesis, Department of Computer Science, Hebrew University, Jerusalem, Israel.

[19] J. Sprinzak and M. Werman, 1990. Exact Point Matching, in Proceedings of the Israeli Symposium on AI, Vision and Pattern Recognition, Elsevier Science Publishers, Amsterdam, pp 31–44.

[20] K. Zikan, 1991. Least-Squares Image Registration, *ORSA Journal on Computing 3*, 169–172.