Chapter 1

Comparing Point Sets Under Projection *

Daniel P. Huttenlocher† Jon M. Kleinberg‡

Abstract

We say that two sets of n points in the plane, P and Q, are views of the same object when there exists a set of n points $S \subset \Re^3$ and two distinct planes A and B, such that P is the projection of S onto A and Q is the projection of S onto B. In the case of orthographic projection, we provide an $O(n^3)$ algorithm for deciding whether P and Q are views of the same object. For central projection we provide an $O(n^7)$ algorithm. In both cases, the running times are a quadratic factor better than naive methods. We also show how randomization can be used to improve the expected running time by a factor of $O(n/\log n)$ when the largest subsets of P and Q that "match" are some fixed fraction of the total size of the sets.

1 Introduction

In this paper we provide polynomial-time algorithms for deciding whether two (unlabelled) sets of n points in the plane are projections of the same (unknown) set of points in $\Re^3$, under either orthographic or central projection. This is a fundamental problem arising in computer vision and robotics, where pictures that are taken at unknown locations must be compared in order to determine whether or not they are views of the same object. While there have been a number of investigations of geometric algorithms for pattern recognition (e.g., [1, 2, 3]) the objects of study have been configurations, such as sets of points in the plane, rather than projections of configurations as we study here.

The idea underlying both orthographic and central projection is that each point in space defines a line of sight which intersects an image plane producing the image of the point under projection. Orthographic and central projection differ in terms of how the line of sight, ℓ_s, is defined for each point $s \in \Re^3$. In the case of orthographic projection, ℓ_s is the line through s and normal to the image plane. We denote orthographic projection by the map $\pi_A : \Re^3 \rightarrow A$ for some given plane A. In the case of central projection, ℓ_s is the line through s and some distinguished center of projection. We denote this central projection by $\pi_A^x : \Re^3 \rightarrow A$ for some plane A and center point x.

The problem that we investigate is that of deciding whether given two sets P and Q of n points in $\Re^2$, there exists a set S of n points in $\Re^3$ such that P and Q are both projections of S. In the case of orthographic projection, this means determining whether there exist planes A and B such that $\pi_A S = P$ and $\pi_B S = Q$. In the case of central projection it means determining whether there exist planes A and B and points x and y such that $\pi_A^x S = P$ and $\pi_B^y S = Q$. In the former case, we provide an algorithm that makes use of the *circular sequence* of the sets P and Q; in the latter case, we introduce a structure which can be viewed as serving the analogous function in the projective plane. Note that the input to the algorithms is just the two-dimensional sets P and Q; they are unlabelled (the correspondence between points of P and Q is not known) and the three-dimensional point set S is unknown. Moreover, the positions of the two image planes A and B (and the centers in the case of central projection) are unknown.

2 Orthographic Projection

In order to simplify the presentation, we rule out certain degeneracies in the point sets P and Q. Assume that no three points of P (or of Q) are collinear. Moreover assume that no four points of P are an affine transformation of any four points of Q. An affine transformation of the plane is determined uniquely by three pairs of points, and thus four (or more) pairs of points that determine the same affine transformation are a degeneracy. (Note that if the set S spans a two-dimensional subspace of $\Re^3$, then the entire sets P and Q will be affine copies of each other. Here we assume that S is not "flat", but it is trivial to handle this case separately.)

The following lemma captures the basic geometric structure of the problem. Recall that $\pi_A : \Re^3 \rightarrow A$ denotes orthographic projection onto the plane A;

*This work was supported in part by National Science Foundation PYI grant IRI-9057928 and matching funds from General Electric, Kodak and Xerox, and in part by Air Force contract AFOSR-91-0328.

†Computer Science Dept., Cornell University, Ithaca NY 14853 USA

‡Laboratory for Computer Science, MIT, Cambridge MA 02139 USA. Supported by an ONR Graduate Fellowship.

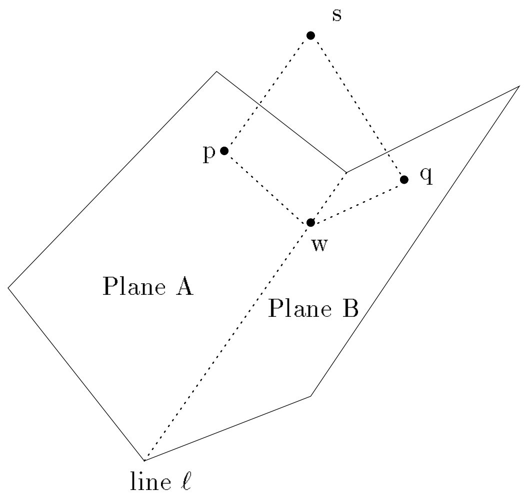

Figure 1: The orthographic projection of a point onto planes A and B.

analogously we denote the orthographic projection of points in the plane onto a line ℓ by $\pi_\ell : \Re^2 \to \ell$. In the following, let ℓ be the line $A \cap B$, where A and B are two distinct planes.

LEMMA 2.1. *Let P and Q be finite subsets of planes A and B respectively. Then there is a point set $S \subset \Re^3$ satisfying $\pi_A S = P$, $\pi_B S = Q$ if and only if $\pi_\ell P = \pi_\ell Q$ (counting multiple points). See Figure 1.*

Proof. Let $S = \{s_1, \ldots, s_n\}$ be such a point set, $p_i = \pi_A s_i$, and $q_i = \pi_B s_i$. The plane determined by each p_i, q_i, s is normal to ℓ; thus $\pi_\ell p_i$ and $\pi_\ell q_i$ are both equal to the point of intersection of ℓ with this plane.

Conversely, suppose that $P = \{p_1, \ldots, p_n\}$, $Q = \{q_1, \ldots, q_n\}$, and $\pi_\ell p_i = \pi_\ell q_i = w_i$ for each i. Then the plane determined by each w_i, p_i, q_i contains the normal to A through p_i and the normal to B through q_i. Since these are two lines in the same plane that are not parallel, they intersect at a point s_i. The set $\{s_1, \ldots, s_n\}$ is the S in the statement of the lemma. ∎

Note that for any set of points, the projection of its centroid is the centroid of its projection. Thus, without loss of generality, we arbitrarily choose an origin on A and B, which we take to be the centroids of P and Q respectively. Moreover, we assume that the images of these centroids under π_ℓ are the origin of the line ℓ.

Lemma 2.1 gives a condition for P and Q to be projections of the same three-dimensional point set that depends only on the relation between P and Q, and not on their relation to S. Namely, P and Q are projections of the same set S if and only if there exist projections of P and Q to the same one-dimensional set (i.e., there exist lines ℓ and ℓ' such that $\pi_\ell P = \pi_{\ell'} Q$). Thus our decision problem can be solved by simply comparing the projections of P and Q onto a line, without considering the three-dimensional set S or the three-dimensional orientations of the planes A and B.

Note that when a point set $S \subset \Re^3$ exists such that $\pi_A S = P$ and $\pi_B S = Q$, this does not uniquely determine a particular point set S. Rather, when such an S exists it is only specified up to an affine transformation of $\Re^3$ (this question of "reconstructing" point sets from two labelled views has been considered in [5]).

As we are concerned with the projections of point sets onto a line, the underlying combinatorial structure is that of the circular sequence (cf. [4], chapter 2). For a point set P, the half-period of its circular sequence is the list of orderings of the points of $\pi_\ell(r_\theta P)$, as P is rotated from $\theta = 0$ to π (where r_θ denotes rotation about the origin by θ). We will denote this sequence of orderings by $circ(P)$. Conceptually it can be viewed as an $\binom{n}{2} \times n$ table. Each row of the table is an ordered list of the points of $\pi_\ell(r_\theta P)$, for a given interval of θ over which the ordering of the projected points does not change. A new row is listed for each (increasing) value of θ at which the line joining a pair of points of P becomes perpendicular to ℓ; at this θ two adjacent points swap order. As there are $\binom{n}{2}$ such swaps, there are $\binom{n}{2}$ rows of the table. While the table is of size $\Theta(n^3)$, $circ(P)$ can be constructed in time $O(n^2 \log n)$, by representing it implicitly in terms of the ordered sequence of swaps.

In our problem, we seek orientations θ and θ' of the two sets P and Q, respectively, such that the projections of the sets onto some line ℓ are equal (or equivalently lines ℓ and ℓ' such that $\pi_\ell P = \pi_{\ell'} Q$). One method for solving this would be to take the correspondence between P and Q determined by each pair of rows of the tables representing $circ(P)$ and $circ(Q)$, and determine whether this correspondence specifies a θ and θ' such that $\pi_\ell(r_\theta P) = \pi_\ell(r_{\theta'} Q)$. The following lemma shows that it is sufficient to consider four points from each set in order to uniquely fix θ and θ'. (Recall our non-degeneracy assumptions about P and Q which are reflected in the statement of the lemma.)

LEMMA 2.2. *Let $U = (u_1, \ldots, u_4)$ and $V = (v_1, \ldots, v_4)$ be two ordered quadruples of points in the plane, such that U is not an affine transformation of V, and no three points of U or V are collinear. Then there is at most one pair of lines ℓ and ℓ' such that $\pi_\ell U = \pi_{\ell'} V$.*

Proof. Let n and n' denote the unit tangents of the lines ℓ and ℓ'. We thus seek the solution (if any) to the system $u_i \cdot n = v_i \cdot n'$, for $i = 1, \ldots, 4$, where n and n' are unit vectors.

Let $\mathbf{U}$ be the 4×2 matrix with rows u_i, $\mathbf{V}$ the 4×2

matrix with rows v_i, and $\mathbf{N} = (n_x, n_y, n'_x, n'_y)^T$ where $n = (n_x, n_y)$ and $n' = (n'_x, n'_y)$. Then ignoring for the moment the fact that n and n' must be unit vectors, we can write this as a system of four homogeneous linear equations in the four unknowns $\mathbf{N}$,

$$\begin{pmatrix} \mathbf{U} & -\mathbf{V} \end{pmatrix} \mathbf{N} = \begin{pmatrix} 0 \end{pmatrix} .$$

If the coefficient matrix $\mathbf{M} = \begin{pmatrix} \mathbf{U} & -\mathbf{V} \end{pmatrix}$ has rank 4, then there is only the trivial solution, so there is no solution where n and n' are unit vectors. If $\mathbf{M}$ has rank 3, then the set of solutions is $\{t(a, b, c, d) : t \in \Re\}$. Since n and n' must be unit vectors, it must be that $a^2t^2 + b^2t^2 = 1$ and $c^2t^2 + d^2t^2 = 1$, so there are just two solutions (which are negations of one another and correspond to the positive and negative tangents of ℓ and ℓ').

Finally, note that $\mathbf{M}$ must have rank at least 3. The points of U and V are not collinear so the rank is clearly at least 2. But if the rank were 2, this would mean that the last two columns could be expressed as a linear combination of the first two, which contradicts the assumption of the lemma. ■

Now we return to the problem of finding θ and θ', if they exist, such that $r_\theta P$ and $r_{\theta'} Q$ project to the same one-dimensional set. In the event that there is such a pair (θ, θ'), there is a row in $circ(P)$ (specified by θ) and a row in $circ(Q)$ (specified by θ') where by definition every corresponding quadruple of points in the two rows yields the same solution to the equation of Lemma 2.2 (namely unit vectors corresponding to the orientations θ and θ'). (That is, if we choose any four column indices $i_1, \ldots, i_4$, and select the quadruple $U = (p_{i_1}, \ldots, p_{i_4})$ from the given row of $circ(P)$ and $V = (q_{i_1}, \ldots, q_{i_4})$ from the given row of $circ(Q)$ these quadruples all yield the same θ and θ'.) Thus, we need only consider four columns of $circ(P)$ and $circ(Q)$ in order to identify the θ, θ' (if any) for which P and Q have the same projection.

Rather than using any four columns, we want to find four corresponding columns of $circ(P)$ and $circ(Q)$ for which the total number of quadruples of points ever appearing in these columns is as small as possible (see Figure 2). Let s_i denote the total number of swaps of $circ(P)$ that involve column i, and define t_i analogously for $circ(Q)$. Since each swap involves two columns of the circular sequence, $\sum_{i=1}^{n}(s_i + t_i) = 2n(n-1)$. Thus if $n > 5$, it is easily verified that there are four column indices $i_1, \ldots, i_4$ such that $(s_{i_1} + t_{i_1}) + \cdots + (s_{i_4} + t_{i_4}) \leq 8(n+5)$. (Clearly a related argument can be carried out for any constant number of indices.) Moreover, given $circ(P)$ and $circ(Q)$, the columns corresponding to four such indices can be found in time $O(n^2)$ by a single pass through the circular sequences.

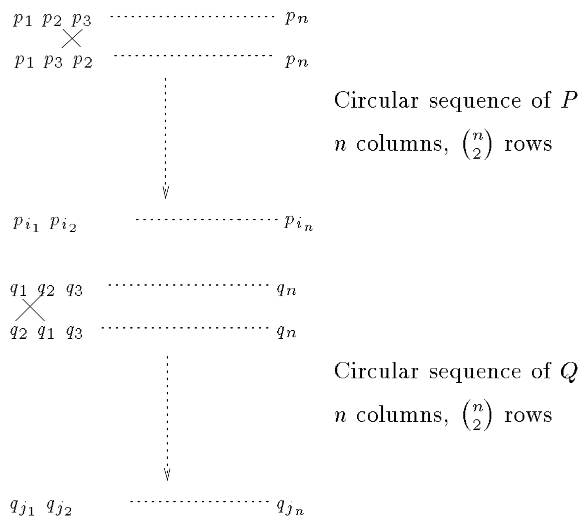

Find four columns in which the total number of swaps is $O(n)$.

Figure 2: The circular sequences of P and Q.

THEOREM 2.1. *Given P and Q, the question of whether there exist planes A and B and a set S of n points in $\Re^3$ such that $\pi_A S = P$ and $\pi_B S = Q$ can be decided in time $O(n^3)$.*

Proof. We present the algorithm.

1. Compute $circ(P)$ and $circ(Q)$, represented as ordered sequences of swaps. This takes time $O(n^2 \log n)$.

2. Identify four columns for which the total number of swaps in both $circ(P)$ and $circ(Q)$ is $O(n)$, as just described. From this, compute the list of the $O(n)$ quadruples of points of P and Q that ever occupy these columns (by scanning the circular sequences). This step takes time $O(n^2)$.

3. For each of the $O(n^2)$ many pairs of quadruples, determine the (θ, θ') (if any) for which the two quadruples project to the same one-dimensional set (using Lemma 2.2). We represent this as a list of pairs $(\theta_{ij}, \theta'_{ij})$. For each θ_{ij} in this list, compute $\pi_\ell(r_{\theta_{ij}} P)$. If we sort the θ_{ij}'s this can be done in time $O(n)$ for each θ_{ij}, since the ordering of the points can simply be looked up in $circ(P)$. Analogously for each θ'_{ij} and Q, for a total time of $O(n^3)$.

4. The algorithm decides "yes" if there is some $(\theta_{ij}, \theta'_{ij})$ for which $\pi_\ell(r_{\theta_{ij}} P) = \pi_\ell(r_{\theta'_{ij}} Q)$, as described above. This verification takes time $O(n)$

for each of the $O(n^2)$ many pairs.
Thus, the total running time is $O(n^3)$. ■

2.1 A Randomized Algorithm. Consider the running time of the above algorithm: it only takes time $O(n^2 \log n)$ to implicitly construct the circular sequences of P and Q; the bound of $O(n^3)$ comes about because of the possibility that we will have $\Omega(n^2)$ matches among our chosen quadruples, and that for each (or at least most) of these matches, it will take time $\Omega(n)$ to decide whether they are genuine or spurious.

In practice, however, most point sets that one compares are quite different from one another. Therefore, it seems that one should be able to devise an algorithm that runs in time $O(n^2 \log n)$ "on most inputs". We make this precise as follows. For two sets P and Q of n points in plane, we will say that they are ε-distinct if there are no subsets $P' \subset P$, $Q' \subset Q$, with

$$|P'| = |Q'| > (1-\varepsilon)n,$$

such that P' and Q' are projections of the same three-dimensional point set.

Given this definition, we can show the following.

THEOREM 2.2. *Given two sets P and Q, the question of whether there exist planes A and B and a set S of n points in $\Re^3$ such that $\pi_A S = P$ and $\pi_B S = Q$ can be decided by a randomized algorithm with expected running time $O(\frac{1}{\varepsilon} n^2 \log n)$ when P and Q are ε-distinct.*

Proof. We present the algorithm. The basic structure is similar to that of the deterministic algorithm, in that both are based on choosing "low-activity" columns from $circ(P)$ and $circ(Q)$. The main difficulty is that we no longer have enough time to construct the rows of the circular sequence explicitly.

1. As in the previous algorithm, we represent $circ(P)$ (and respectively $circ(Q)$) as an ordered list of values of the $\Theta(n^2)$ values of θ at which swaps occur; these will be called the P-critical (resp. Q-critical) values of θ. We also maintain an array that will be used to specify the order of the points of $\pi_\ell(r_\theta P)$ for a given θ. Initially this ordering is simply $\pi_\ell P$.

2. We pick four columns of $circ(P)$ at random, and compute the set of all quadruples that ever occupy these columns. We do the same for the corresponding four columns of Q. Since the expected number of swaps involving these four columns is $O(n)$, the expected number of quadruples we obtain is $O(n)$.

3. For each of the pairs of quadruples (one from P and one from Q), determine the (θ, θ') (if any) for which the two quadruples project to the same one-dimensional set (using Lemma 2.2). We represent this as a list of pairs $(\theta_{ij}, \theta'_{ij})$. The expected number of pairs produced is $O(n^2)$.

4. We sort the values θ_{ij} and merge this sorted list with the sorted list of P-critical θ's. We now proceed through this list one entry at a time. If the entry is a P-critical θ, we perform the indicated swap by exchanging the two entries of the array that represents $\pi_\ell(r_\theta P)$, in constant time.

   If the entry is θ_{ij}, then we iteratively perform the following steps.

   a) Randomly choose $q_k \in Q$ and compute $\pi_\ell(r_{\theta'_{ij}} q_k)$. By performing binary search on the array representing the ordering of $\pi_\ell(r_{\theta_{ij}} P)$ (only when the array entry for p_m is examined do we compute the actual value of $\pi_\ell(r_{\theta_{ij}} p_m)$), we can determine whether there exists a $p_m \in P$ such that $\pi_\ell(r_{\theta'_{ij}} q_k) = \pi_\ell(r_{\theta_{ij}} p_m)$.

   b) If there is no such p_m, then clearly $\pi_\ell(r_{\theta_{ij}} P) \neq \pi_\ell(r_{\theta_{ij}} P)$; proceed to the next θ in the sorted list.

   c) Otherwise, repeat step (a), choosing some q_k that has not already been examined.

5. The algorithm answers "yes" if there is some $(\theta_{ij}, \theta'_{ij})$ for which $\pi_\ell(r_{\theta_{ij}} P) = \pi_\ell(r_{\theta'_{ij}} Q)$, as described above.

We analyze the running time as follows. It takes time $O(n^2 \log n)$ to build the implicit representation of the circular sequences. Producing the pairs $(\theta_{ij}, \theta'_{ij})$ takes time proportional to the expected number of pairs of quadruples, which is $O(n^2)$.

The algorithm then walks through this list of θ-values. Each of the $O(n^2)$ P-critical θ's can be handled in constant time. For a given θ_{ij}, each binary search operation takes time $O(\log n)$, to determine whether there is some $p_m \in P$ with the same projection onto ℓ as $q_k \in Q$. Since P and Q are ε-distinct, the probability that a q_k will be chosen with no matching p_m is at least ε; thus the expected number of binary searches is $O(\frac{1}{\varepsilon})$, for a total expected time of $O(\frac{1}{\varepsilon} \log n)$ on a single θ_{ij}. Thus, by linearity of expectation, the total expected running time is $O(\frac{1}{\varepsilon} n^2 \log n)$. ■

3 Central Projection

We now consider the case of central projection. Given a plane A and a point $x \in \Re^3$, recall that we denote the central projection of $S \subset \Re^3$ through x onto A by $\pi_A^x S$. Similarly, for a planar point set P, a point $a \in \Re^2$,

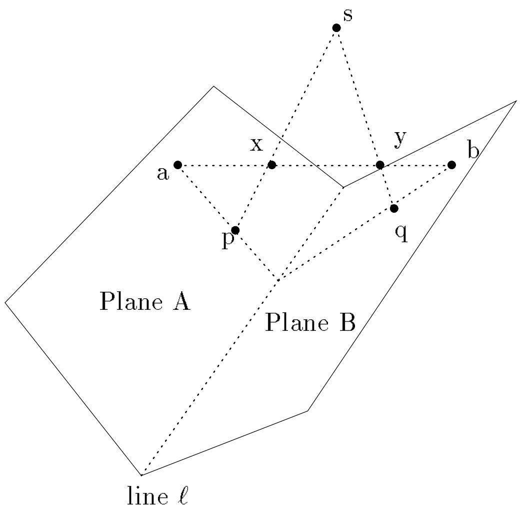

Figure 3: The central projection of a point onto planes A and B.

and a line ℓ, let $\pi_\ell^a P$ denote the central projection of P through y onto ℓ.

The decision problem we wish to answer is the following: given sets P and Q of n points in the plane, does there exist a set S of n points in $\Re^3$, two planes A and B, and two points $x, y \in \Re^3$ such that $\pi_A^x S = P$ and $\pi_B^y S = Q$? Our examination of the structure of this problem reveals a strong sense in which it is combinatorially more difficult than the orthographic version; we give a deterministic decision algorithm which runs in time $O(n^7)$, and a randomized algorithm with expected running time $O(n^6(\frac{1}{\varepsilon} + \log n))$ when P and Q are ε-distinct. These algorithms are based on a structure called the *circular partition* of a planar point set P, which we introduce below. In a sense, it is a generalization of the circular sequence of P, in that it contains the analogous structural information about the central projections of P onto a line.

Throughout this section, P and Q will denote sets of n points in the plane. We assume that no three points of either set are collinear, as well as making a certain mild non-degeneracy assumption to be described below. A and B will be used to denote non-parallel planes in $\Re^3$, ℓ the line $A \cap B$, and x, y will denote two points in $\Re^3$ not on A or B. Finally, S will denote a set of n points in $\Re^3$. We begin with an introductory fact that is similar to Lemma 2.1.

Lemma 3.1. *There exist $x, y \in \Re^3$ and $S \subset \Re^3$ satisfying $\pi_A^x S = P$, $\pi_B^y S = Q$ if and only if there exist points $a \in A$, $b \in B$ such that $\pi_\ell^a P = \pi_\ell^b Q$. ($\ell$ is defined as above; see Figure 3).*

Proof. Let $x, y \in \Re^3$ and $S = \{s_1, \ldots, s_n\}$ be as in the statement of the lemma, and suppose $p_i = \pi_A^x s_i$, $q_i = \pi_B^y s_i$. Let $a = \pi_A^x y$ and $b = \pi_B^y x$; then the plane determined by s_i, x, y intersects A in the line ap_i and intersects B in the line bq_i. Thus the intersection of the plane s_i, x, y with ℓ is the common value of $\pi_\ell^a p_i$ and $\pi_\ell^b q_i$.

Conversely, suppose that $P = \{p_1, \ldots, p_n\}$, $Q = \{q_1, \ldots, q_n\}$, and there exist $a \in A$, $b \in B$ for which $\pi_\ell^a p_i = \pi_\ell^b q_i = w_i$ for each i. Choose two points x, y on the line ab such that neither x nor y lies on planes A or B, and none of the pairs of lines xp_i, yq_i is parallel. For each i, lines xp_i, yq_i lie in the same plane and are not parallel; hence they intersect at a point s_i. The set $\{s_1, \ldots, s_n\}$ is the S in the statement of the lemma. ∎

As in the orthographic case, this provides a condition for P and Q to be views of the same three-dimensional point set that depends only on the relation between P and Q, and not on their relation to S. Namely, P and Q are projections of the same set S if and only if there exist projections of P and Q to the same one-dimensional set (i.e., there exist lines ℓ and ℓ' and points a and b such that $\pi_\ell^a P = \pi_{\ell'}^b Q$).

At a high level, the algorithms for central projection are similar to the ones for orthographic projection. We partition the set of all projections of P (resp. Q) into equivalence classes for which the ordering of points in the projection remains unchanged. Under orthographic projection, these equivalence classes were given by the circular sequence of P (resp. Q). In the case of central projection there is a new structure which we introduce below; we refer to it as the *circular partition* induced by P (resp. Q). Each ordering of P and Q in these partitions defines a correspondence of points in P with points in Q, up to a cyclic shift. We use fixed-size "witness" sets $U \subset P$ and $V \subset Q$ to determine whether U and V (with a given correspondence) are both projections of the same set W. When they are, we check whether the entire sets P and Q are projections of the same set S. In the case of orthographic projection the witness sets contained four points (i.e., this was sufficient to uniquely determine the two orthographic projections). For central projection, one can take witness sets of size eight, as we now argue.

Let $U = \{u_1, \ldots, u_k\}$, $V = \{v_1, \ldots, v_k\}$, and consider a solution to the equation $\pi_\ell^a U = \pi_{\ell'}^b V$. If we apply suitably chosen similarity transformations φ, φ' (rigid motions composed with scale), then we can obtain sets $\varphi U, \varphi' V$ for which the points φx and $\varphi' y$ are the origin, and the lines $\varphi \ell$, and $\varphi' \ell'$ are both $y = -1$. We therefore obtain the equations

$$\frac{\varphi u_i(x)}{\varphi u_i(y)} = \frac{\varphi v_i(x)}{\varphi v_i(y)}, \quad i = 1, 2, \ldots, k \tag{3.1}$$

where $\gamma(x)$ denotes the x-coordinate of a point γ. Each of φ, φ' consists of the four unknown parameters of the similarity transformation, so this can be written as a system of k polynomial equations in 8 unknowns. We therefore make the following non-degeneracy assumption about our sets P and Q: if $k = 8$, then for some absolute constant c, (3.1) has at most c solutions in φ, φ'. Thus eight points from P and Q are sufficient for the equation $\pi^a_\ell U = \pi^b_{\ell'} V$ to have finitely many solutions.

Now we turn to the issue of defining the combinatorially equivalent central projections of a set of points in the plane. Recall that by this we mean those projections for which the ordering of the projected points is the same. Under central projection, the natural ordering of a set of points is with respect to the center of projection (in contrast, for orthographic projection it was the ordering along the line ℓ). Given a set $P \subset \Re^2$ and a center point $a \in \Re^2$, the pencil of lines $\eta(P, a) = \{p_i a | p_i \in P\}$ defines a cyclic ordering of the points of P with respect to a. The central projection $\pi^a_\ell P$ is then simply the intersection of the line ℓ with this pencil $\eta(P, a)$. Thus, for any line ℓ the ordering of points along that line is the same up to a cyclic shift for a given $\eta(P, a)$ (i.e., for a given center of projection, a).

For some set P we say that a and a' are combinatorially equivalent centers of projection exactly when $\eta(P, a)$ and $\eta(P, a')$ have the same cyclic ordering. This induces a partition of $\Re^2$ into equivalence classes, which we call the *circular partition* of the plane induced by P, and denote by $\kappa(P)$.

We can associate $\kappa(P)$ with a subdivision of the plane as follows. For points $p_i, p_j \in P$, let $(p_i p_j)^*$ denote the line through p_i and p_j with the segment $\overline{p_i p_j}$ deleted. That is, it is a union of two rays. Let $\mathcal{S}^*(P)$ denote the polygonal subdivision of the plane induced by the arrangement of rays $\{(p_i p_j)^* : p_i, p_j \in P\}$. Now if a point a is constrained to move in a single face of this arrangement, it is easy to see that the cyclic ordering of $\eta(P, a)$ cannot change; thus we have the following.

LEMMA 3.2. *If a and a' lie in the same face of $\mathcal{S}^*(P)$, then they are combinatorially equivalent centers of projection. As a corollary, the number of combinatorially distinct centers of projection of any point set P is $O(n^4)$.*

It is not difficult to show that this bound is tight in general; i.e. there are generally $\Theta(n^4)$ distinct cyclic orderings of $\eta(P, a)$ as a ranges over $\Re^2$. For simplicity of presentation in what follows, we will assume that the centers of projection of the sets P and Q induce induce non-degenerate pencils; i.e. they lie in one of the cells (2-faces) of $\mathcal{S}^*(P)$ (resp. $\mathcal{S}^*(Q)$), rather than on one of the edges or vertices. It is clear how to modify the following algorithm to handle this kind of degeneracy.

Remarks: We believe that $\kappa(P)$ possesses a considerable amount of structure that is interesting in itself, beyond what we make use of in the algorithm below. Consider the relation to the case of orthographic projection. The orthographic projections $\pi_\ell P$ can be thought of as central projections $\pi^a_\ell P$ in which a is a "point at infinity" in the projective plane. Thus, the point a lies in one of the unbounded faces of $\mathcal{S}^*(P)$; indeed, a circular traversal of these unbounded faces, following the line at infinity, yields the circular sequence $circ(P)$. This gives some insight into why the decision problem for central projection appears to be fundamentally more difficult than that for orthographic projection. There are only $O(n^2)$ unbounded cells of $\mathcal{S}^*(P)$, or equivalently entries of $circ(P)$, whereas in the current case we must search an arrangement with $\Theta(n^4)$ cells.

Consider a fixed constant k ($k = 8$ in our case, but the following definitions can be made in general), and a set U of k points from P. For a given point $a \in \Re^2$, the lines of $\eta(U, a)$ are a subsequence of $\eta(P, a)$; let us say that the *pattern* associated with U in this pencil is the cyclic k-tuple $I = (i_1, \ldots, i_k)$, in which the i_j are the number of lines of $\eta(P \backslash U, a)$ lying between adjacent lines of $\eta(U, a)$. Conversely, U will be called an *instance* of pattern I. Observe that if there is a solution to the equation $\pi^a_\ell P = \pi^b_{\ell'} Q$, then $\pi^a_\ell U$ is matched with some $\pi^b_{\ell'} V$, where $V \subset Q$ is an instance of the pattern associated with U. Thus, our algorithm will pick a set U of 8 points from P, enumerate all sets $V \subset Q$ that are instances of patterns associated with U, solve the equation $\pi^a_\ell U = \pi^b_{\ell'} V$, and see if this can be extended to a solution for all of P and Q.

The bound on the running time is therefore based largely on the following lemma.

LEMMA 3.3. *Let k be a fixed constant. Then a set $U \subset P$ of cardinality k is an instance of $O(n^2)$ patterns. Dually, a fixed pattern I on k elements has $O(n^4)$ instances among the collection of all $U \subset P$, $|U| = k$.*

Proof. Consider a set $U \subset P$ of cardinality k. Let L be the collection of all lines passing through at least one point of U and at least one other point of P. For a given cell of the arrangement of L, the pattern associated with U is fixed. Since $|L| = O(n)$, there are $O(n^2)$ such cells, so we have the first bound.

To obtain the second bound, let I be a given pattern, and consider the dual graph G of the arrangement $\mathcal{S}^*(P)$. Since G is a planar graph on $O(n^4)$ vertices, it has $O(n^4)$ edges. Double each edge of G so as to make it Eulerian; the resulting Eulerian tour of G traverses each edge exactly twice. Consider traversing G in this manner, counting instances of I as we go. In the first cell, there are $O(n)$ instances of I that we have not seen before. Each time we traverse an edge, precisely two lines swap order in the pencil of P; this can give rise to at most $2k$ new instances of the pattern I. As there are $O(n^4)$ steps in this tour of G, we have the desired bound. ■

THEOREM 3.1. *Given P and Q, the question of whether there exist planes A and B, points $x, y \in \Re^3$, and a set S of n points in $\Re^3$ such that $\pi_A^x S = P$ and $\pi_B^y S = Q$ can be decided in time $O(n^7)$.*

Proof. We present the algorithm.

1. Construct the arrangements $\mathcal{S}^*(P)$ and $\mathcal{S}^*(Q)$, and preprocess them for point location queries.

2. Choose a set U of 8 points from P and compute the set of all $O(n^2)$ patterns that are associated with it.

3. For each pattern, enumerate the set of all $V \subset Q$ that are instances of it. This gives a list of $O(n^6)$ possible subsets V.

4. For each V, solve the equation $\pi_\ell^a U = \pi_{\ell'}^b V$, obtaining at most some constant c solutions.

5. For each solution, perform a point location query for a in $\mathcal{S}^*(P)$ and $b \in \mathcal{S}^*(Q)$; from the cell in which each lies, we can read off the ordering of the points in $\pi_\ell^a P$ and $\pi_{\ell'}^b Q$ (this saves us the cost of sorting). Determine whether in fact $\pi_\ell^a P = \pi_{\ell'}^b Q$.

We analyze the running time as follows. Building and preprocessing the arrangements takes time $O(n^4 \log n)$. The proof of Lemma 3.3 provides an algorithm for obtaining all the pattern information in time $O(n^6)$. For each V in this list, we solve the equation $\pi_\ell^a U = \pi_{\ell'}^b V$ in constant time, and check whether it extends to all of P and Q in $O(n)$ time. Thus, the total running time is $O(n^7)$. ■

Making use of the definition of ε-distinctness as in the orthographic case, we can modify this algorithm by analogy with that case to show the following.

THEOREM 3.2. *Given sets P and Q, the question of whether there exist planes A and B, points $x, y \in \Re^3$, and a set S of n points in $\Re^3$ such that $\pi_A^x S = P$ and $\pi_B^y S = Q$ can be decided by a randomized algorithm with expected running time $O(n^6(\frac{1}{\varepsilon} + \log n))$ when P and Q are ε-distinct.*

4 Conclusion

These algorithms leave open a number of interesting questions. There are tight bounds on the running time of many "classical" point set comparison problems such as congruence in $\Re^2$ and $\Re^3$ [3]. What are the corresponding bounds for comparing point sets under projection? It would also be interesting to provide an algorithm that handles uncertainty, such as when the locations of points are known only approximately (as in [1]). A related problem is that of deciding whether two given point sets $S, T \subset \Re^3$ can produce the same projection into $\Re^2$. That is, do there exist planes A and B such that $\pi_A S = \pi_B T$?

5 Acknowledgements

The authors would like to thank Paul Chew, Robert Kleinberg, and Bernd Sturmfels for helpful discussions and comments on the topic of this paper.

References

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

[2] M. Atallah, "Similarity of Planar Figures", *Int. J. Comput. Inf. Sci.*, vol. 13, pp. 279–290, 1984.

[3] M.D. Atkinson, "An Optimal Algorithm for Geometrical Congruence", *J. Algorithms*, 8, pp. 159–172, 1987.

[4] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Berlin: Springer-Verlag, 1987.

[5] J.J. Koenderink and A.J. vanDoorn, "Affine Structure from Motion" *J. Optical Soc. of Amer.*, 8(10), pp. 377–385, 1991.