Point Sets Up to Rigid Transformations are Determined by the Distribution of their Pairwise Distances

Daniel Chen

December 6, 2008

Abstract

This report is a summary of [BK06], which gives a simpler, albeit less general proof for the result of [BK04].

1 Introduction

[BK04] and [BK06] study the circumstances when sets of n-points in Euclidean space are determined up to rigid transformations by the *distribution* of unlabeled pairwise distances. In fact, they show the rather surprising result that *any set of n-points in general position* are determined up to rigid transformations by the distribution of unlabeled distances.

1.1 Practical Motivation

The results of [BK04] and [BK06] are partly motivated by experimental results in [OFCD]. [OFCD] is an experimental study evaluating the use of *shape distributions* to classify 3D objects represented by point clouds. A shape distribution is a the distribution of a function on a randomly selected set of points from the point cloud. [OFCD] tested various such functions, including the angle between three random points on the surface, the distance between the centroid and a random point, the distance between two random points, the square root of the area of the triangle between three random points, and the cube root of the volume of the tetrahedron between four random points. Then, the distributions are compared using dissimilarity measures, including χ_2, Battacharyya, and Minkowski L_N norms for both the pdf and cdf for $N = 1, 2, \infty$.

The experimental results showed that the shape distribution function that worked best for classification was the distance between two random points and the best dissimilarity measure was the pdf L_1 norm. [BK04] provides a partial theoretical justification for these results by showing that there exists a general position assumption such that the distribution of pairwise distances completely determines the shape of the point set. This result, however, only considers the case when the distribution of pairwise distances are exactly the same. Whether a similar result holds when the distributions are close is an open problem.

2 Preliminaries

Definition 2.1 (General Position). *We use the definition of general position given in [M]. Let a set of* n *points in* $\mathbb{R}^d$ *be specified by a vector* $t = (t_1, t_2, \ldots, t_m)$ *for* $m = dn$*. Then, a general position condition is a condition that can be expressed as* $\bigwedge_i p_i(t) \neq 0$ *for a countable number of polynomials* p_i*.*

The intuition behind a general position condition is that configurations in general position lie arbitrarily close to any given configuration, as polynomials only have a finite number of zeros. [BK04] and [BK06] show a general position condition that also results in the point configuration being completely determined by the distribution of pairwise distances up to rigid transformations. We define rigid transformations as the following:

Definition 2.2 (Rigid Transformation). *Let* p *be a point in* $\mathbb{R}^m$*. Then a rigid transformation is a function* $R : \mathbb{R}^m \to \mathbb{R}^m$ *that can be written as* $R(p) = Mp+T$ *where* M *is an orthogonal* m*-by-*m *matrix and* T *is a column vector in* $\mathbb{R}^m$*.*

Definition 2.3 (Distribution of Pairwise Distances). *Given a set of points* P*, let the distribution of pairwise distances* $d(P)$ *be the multiset* $\{||p_i - p_j||\}_{i<j}$*.*

We also define reconstructibility from pairwise distances as the following:

Definition 2.4 (Reconstructibility from Pairwise Distances). *We say that* $p_1, \ldots, p_n \in \mathbb{R}^m$ *is reconstructible from pairwise distances if for every* $q_1, \ldots, q_n \in \mathbb{R}^m$ *with the same distribution of pairwise distances, there exists a rigid transformation* R *and a permutation* π *of* $\{1, \ldots, n\}$*, such that* $P(p_i) = q_{\pi(i)}$ *for every* $i \in 1, \ldots, n$*.*

3 Folklore Lemma

Because it is tricky to work directly with point sets under rigid transformations, we will instead compare instead the distance matrices. The Folklore Lemma states that two configurations of points have the same distance matrices if and only if there exists a rigid transformation that maps one configuration to the other. The following is an adaptation of the proof given in [BK04], however using the language of more elementary linear algebra.

Lemma 3.1 (Folklore Lemma). *Let* $p_1, \ldots, p_n$ *and* $q_1, \ldots, q_n$ *be points in* $\mathbb{R}^m$*. If* $||p_i - p_j|| = ||q_i - q_j||$ *for every* i, j *in* $1, \ldots, n$*, then there exists a rigid transformation* R *such that* $R(p_i) = q_i$ *for all* i*.*

Proof. Let $x_i = p_i - p_n$ and $y_i = q_i - q_n$ for $i = 1, \ldots, n$. We claim that $\langle x_i, x_j \rangle = \langle y_i, y_j \rangle$ for all i, j. To see this, we have the following calculation:

$$\begin{aligned}
\langle x_i, x_j \rangle &= \langle p_i - p_n, p_j - p_n \rangle \\
&= \frac{\langle p_i - p_n, p_i - p_n \rangle + \langle p_j - p_n, p_j - p_n \rangle - \langle p_i - p_j, p_i - p_j \rangle}{2} \\
&= \frac{||p_i - p_n|| + ||p_j - p_n|| - ||p_i - p_j||}{2} \\
&= \frac{||q_i - q_n|| + ||q_j - q_n|| - ||q_i - q_j||}{2} \\
&= \frac{\langle q_i - q_n, q_i - q_n \rangle + \langle q_j - q_n, q_j - q_n \rangle - \langle q_i - q_j, q_i - q_j \rangle}{2} \\
&= \langle q_i - q_n, q_j - q_n \rangle = \langle y_i, y_j \rangle
\end{aligned}$$

Now, let $X = [x_1, \cdots, x_n]$ and $Y = [y_1, \cdots, y_n]$. We have shown above that $X^T X = Y^T Y$. Now, $X^T X$ is symmetric and positive semidefinite, and hence can be written as $Q \Lambda Q^T$ for an orthogonal Q and a non-negative diagonal matrix Λ. Since Λ, is non-negative, we can write $X^T X = Y^T Y = Q \Lambda^{1/2} \Lambda^{1/2} Q^T$. Therefore, using the singular value decomposition, we can write X as $U_X \Lambda^{1/2} Q^T$ and Y as $U_Y \Lambda^{1/2} Q^T$ for orthogonal U_X and U_Y. Then, we can write $Y = MX$ for orthogonal $M = U_Y U_X^T$. Moreover, it is easy to verify that $M x_i = y_i$ for $i = 1, \ldots, n$. To finish the proof, we note that $M(p_i - p_n) = q_i - q_n$ or $q_i = M p_i + q_n - M p_n$ so there is a rigid transformation from p_i to q_i with the orthogonal matrix M and translation vector $q_n - M p_n$. □

Now, it is obvious that the distribution of pairwise distances remains the same after any rigid transformation. We note that it is not the case that all point configurations are reconstructible from pairwise distances: [B] came up with a counterexample for a set of points in one dimension as follows:

$$P = \{0, 1, 4, 10, 12, 17\} \quad Q = \{0, 1, 8, 11, 13, 17\}$$

It can be verified that the distribution of distances is the multiset:

$$d(P) = d(Q) = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16, 17\}$$

for both point sets while it is also clear that there is no isometry from P to Q. [BK04] show a counterexample in two dimensions, see Figure 1. [BK04] and [BK06], however, show the surprising result that there is a general position condition for n points where $n \geq m + 2$ that implies the points are in fact reconstructible from pairwise distances. For concreteness, we will describe the proof for n-point configurations for $n \geq 5$ in $\mathbb{R}^2$.

4 Most Point Configurations are Reconstructible

With the Folklore Lemma in mind, our approach now is to show that two point sets with the same distribution of distances have the same distance matrices

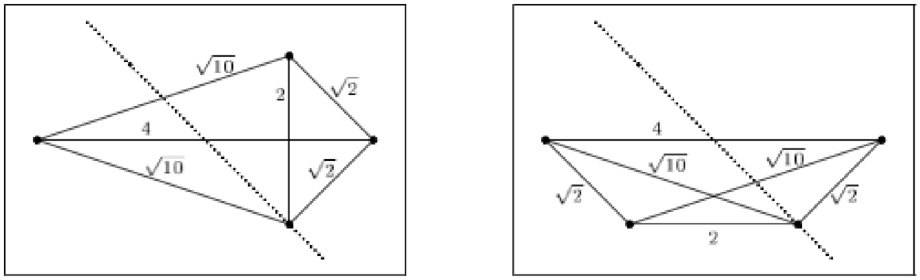

Figure 1: 2D Counterexample

up to a relabeling of points. For a set of n points, the distance matrices are determined by $\binom{n}{2}$ pairs corresponding to the distances between distinct points. For convenience, we let $\mathcal{P} = \{\{i,j\} | 1 \leq i < j \leq n\}$ be all such pairs of indices. Then, we introduce the following notation to denote distances between points in point configurations:

Definition 4.1. *Let $P = \{p_1, \dots, p_n\}$ be a labeled set of n points. Then we let $d_P : \mathcal{P} \to \mathbb{R}$ be defined as $d_P(\{i,j\}) = ||p_i - p_j||^2$.*

Our convention is to use uppercase to denote the configuration of points and the corresponding subscripted lowercase letters to denote the labeled points of that configuration, e.g., given a set of points P, p_i will refer to the ith point of P. Then, we have the following corollary:

Corollary 4.2. *If n-point configurations P and Q have the same distribution of distances, then there exists a permutation ϕ of $\mathcal{P}$ such that $d_P(\phi(\{i,j\})) = d_Q(\{i,j\})$ for all $\{i,j\} \in \mathcal{P}$.*

Equipped with this notation, we then show the following preliminary result that if two point configurations are sufficiently close and have the same distribution of distances, then they are the same up to a rigid transformation and a relabeling of points. This result is weaker than desired but gives some intuition on why point configurations might be determined by their pairwise distances:

Theorem 4.3. *Let P be a configuration of n points.Then, there exists neighborhoods of p_i such that if Q if a configuration of n points with q_i in the neighborhood of p_i and having the same distribution of distances as P, then P and Q are the same up to a rigid transformation and a relabeling of the points.*

Proof. Suppose for the purpose of contradiction that there exists a sequence of configurations $\{Q_k\}_{k=1}^{\infty}$ converging to P such that none of Q_k can be mapped to P via rigid transformation and relabeling, but there exists a sequence of permutations $\{\phi_k\}_{k=1}^{\infty}$ such that $d_P(\phi_k(\{i,j\})) = d_{Q_k}(\{i,j\})$. Since there are finitely many permutations ϕ_k, we can pick ϕ_1, for instance, and let $\{R_l\}_{l=1}^{\infty}$ be the subsequence of $\{Q_k\}_{k=1}^{\infty}$ where $\phi_k = \phi_1$. Then, taking the limit $| \to \infty$, we have $d_P(\phi_1(\{i,j\})) = \lim_{l\to\infty} d_{R_l}(\{i,j\})$. Since $\{R_l\}_{l=1}^{\infty}$ converges to P, we then have $d_P(\phi_1(\{i,j\})) = d_P(\{i,j\})$. But then, we have $d_P(\{i,j\}) = d_{Q_k}(\{i,j\})$ and by the Folklore Lemma, there is a rigid transformation from P to Q_k, so we have a contradiction. □

However, to show that such a property holds for point configurations that are far apart, we have to develop different machinery. First, we consider *relabelings*, which are permutations with the following property:

Definition 4.4 (Relabelings)**.** *A permutation ϕ of $\mathcal{P}$ is a relabeling if there exists a permutation π of $\{1,\dots,n\}$ such that $\phi(\{i,j\}) = \{\pi(i),\pi(j)\}$.*

Corollary 4.5. *If there exists a relabeling ϕ of $\mathcal{P}$ such that $d_P(\phi(\{i,j\})) = d_Q(\{i,j\})$, then there is a permutation π of $\{1,\dots,n\}$ such that $p_{\pi(1)},\dots,p_{\pi(n)}$ has the same distance matrix as $q_1,\dots q_n$, and therefore, by the Folklore Lemma, there is a rigid transformation from $p_{\pi(1)},\dots,p_{\pi(n)}$ to $q_1,\dots q_n$.*

Therefore, we want to show that most n-point configurations P and Q with the same distribution of distances in fact have a relabeling ϕ of $\mathcal{P}$ such that $d_P(\phi(\{i,j\})) = d_Q(\{i,j\})$ for all $\{i,j\} \in \mathcal{P}$. To that end, we show the following key lemma:

Lemma 4.6. *Suppose $n \neq 4$. Then a permutation ϕ of $\mathcal{P}$ is a relabeling if and only if for all pairwise distinct indices i,j,k we have $\phi(\{i,j\}) \cap \phi(\{i,k\}) \neq \emptyset$.*

Proof. For every $n \leq 3$, every ϕ is a relabeling and $\phi(\{i,j\}) \cap \phi(\{i,k\}) \neq \emptyset$. Therefore, we may assume that $n \geq 5$. The only if part of the statement is clear by the definition of a relabeling, so we will show the if direction.

Suppose we have ϕ a permutation of $\mathcal{P}$ such that for all pairwise distinct indices i,j,k we have $\phi(\{i,j\}) \cap \phi(\{i,k\}) \neq \emptyset$. We note that since ϕ is a permutation, the intersection must then contain only one element. Then, we argue that for i,j,k,l pairwise distinct, $\phi(\{i,j\}) \cap \phi(\{i,k\}) \cap \phi(\{i,l\}) \neq \emptyset$. Suppose otherwise. Then, we can write $\phi(\{i,j\}) = \{a,b\}$, $\phi(\{i,k\}) = \{a,c\}$, and $\phi(\{i,l\}) = \{b,c\}$. Now, since we have more than 4 points, we can choose m distinct from i,j,k,l. $\phi(\{i,m\})$ must intersect $\{a,b\}$, $\{a,c\}$, and $\{b,c\}$. Since $\phi(\{i,m\})$ only has two elements, it must be one of the sets $\{a,b\}$, $\{a,c\}$, and $\{b,c\}$, but then that would violate ϕ being a permutation. Therefore, $\phi(\{i,j\}) \cap \phi(\{i,k\}) \cap \phi(\{i,l\}) \neq \emptyset$.

Then if we fix an i and choose any distinct j,k, $\phi(\{i,j\}) \cap \phi(\{i,k\})$ must contain an distinct element a and the above shows that a belongs to any $\phi(\{i,l\})$ where l is distinct from i,j,k. Therefore $\bigcap_{l\neq i} \phi(\{i,l\}) = a$ and we can define the map σ from $\{1,\dots,n\}$ to itself where $\sigma(i) = a$.

To show that σ is a permutation, we simply need to show that it is injective. To do this, we let M_i be the set of all pairs with i in them. Then, $\phi(M_i) \subseteq M_{\sigma(i)}$. But ϕ is a permutation and $|M_i| = |M_{\sigma(i)}|$ so therefore $\phi(M_i) = M_{\sigma(i)}$. Now consider i,j with $\sigma(i) = \sigma(j)$. Then, $M_{\sigma(i)} = M_{\sigma(j)}$ so $\phi(M_i) = \phi(M_j)$. But ϕ is a permutation so $M_i = M_j$ and therefore $i = j$ so σ is injective.

Now, consider $\phi(\{i,j\})$. By the above discussion, it contains both $\sigma(i)$ and $\sigma(j)$, so $\phi(\{i,j\})$ must be $\{\sigma(i),\sigma(j)\}$. Since σ is a permutation, ϕ is therefore a relabeling. □

The proof that most point configurations are reconstructible from pairwise distances relies on a certain determinant that is zero when four points p_i, p_j, p_k, p_l lie on a plane:

$$\det \begin{pmatrix} a & b & c \\ b & d & e \\ c & e & f \end{pmatrix} = 0$$

where $a = -2d_P(\{i,l\})$, $b = d_P(\{i,j\}) - d_P(\{i,l\}) - d_P(\{j,l\})$, $c = d_P(\{i,k\}) - d_P(\{i,l\}) - d_P(\{k,l\})$, $d = -2d_P(\{j,l\})$, $e = d_P(\{j,k\}) - d_P(\{j,l\}) - d_P(\{k,l\})$ and $f = -2d_P(\{k,l\})$.

This determinant can be expanded as a polynomial:

$$\begin{aligned} g(U,V,W,X,Y,Z) = {} & 2U^2Z + 2UVX - 2UVY - 2UVZ - 2UXW - 2UXZ + \\ & 2UYW - 2UYZ - 2UWZ + 2UZ^2 + 2V^2Y - 2VXY - \\ & 2VXW + 2VY^2 - 2VYW - 2VYZ + 2VWZ + 2X^2W - \\ & 2XYW + 2XYZ + 2XW^2 - 2XWZ \end{aligned}$$

where $U = d_P(\{i,j\})$, $V = d_P(\{i,k\})$, $W = d_P(\{i,l\})$, $X = d_P(\{j,k\})$, $Y = d_P(\{j,l\})$ and $Z = d_P(\{k,l\})$. Thus, we are ready for the main result:

Theorem 4.7. *Let $n \geq 5$ and P be a configuration of n points in $\mathbb{R}^2$. Suppose for choices of indices a,b,c,d,e,f,g,h,i,j,k such that the pairs $U = \{a,b\}$, $V = \{d,e\}, W = \{f,g\}, X = \{h,i\}, Y = \{j,k\}, Z = \{a,c\}$ are distinct, we have:*

$$g(d_P(U), d_P(V), d_P(W), d_P(X), d_P(Y), d_P(Z)) \neq 0$$

then P is reconstructible from pairwise distances.

Proof. Suppose that Q is an n-point configuration in $\mathbb{R}^2$ with the same distribution of distances as P. Then, there exists a permutation ϕ such that $d_Q(\{i,j\}) = d_P(\phi(\{i,j\}))$ for all $i \neq j$. We then try to show that ϕ^{-1} is a relabeling, and thus ϕ is also a relabeling.

Now pick any pairwise distinct indices r,s,t,u. Now since q_r, q_s, q_t, q_u lie on a plane,

$$g(d_Q(\{r,s\}), d_Q(\{r,t\}), d_Q(\{r,u\}), d_Q(\{s,t\}), d_Q(\{s,u\}), d_Q(\{t,u\})) = 0$$

But then, it follows that

$$g(d_P(\phi(\{r,s\})), d_P(\phi(\{r,t\})), d_P(\phi(\{r,u\})), d_P(\phi(\{s,t\})), d_P(\phi(\{s,u\})), d_P(\phi(\{t,u\}))) = 0$$

Therefore, it follows that $\phi(\{r,s\})$ and $\phi(\{t,u\})$ are disjoint, otherwise the pairs would satisfy the conditions of the U,V,W,X,Y,Z stated above and

$$g(d_P(U), d_P(V), d_P(W), d_P(X), d_P(Y), d_P(Z)) \neq 0$$

so we have show that ϕ maps disjoint sets $\{r, s\}$ and $\{t, u\}$ to disjoint sets. We take the contrapositive and note that if $\phi(\{r, s\})$ and $\phi(\{t, u\})$ intersect, then $\{r, s\}$ and $\{t, u\}$ also necessarily intersect. Thus, for all i, j, k we have

$$\phi^{-1}(\{i, j\}) \cap \phi^{-1}(\{i, k\}) \neq \emptyset$$

and hence ϕ^{-1} is a relabeling. But then there exists a permutation π of $\{1, \ldots, n\}$ such that $\phi^{-1}(\{i, j\}) = \{\pi(i), \pi(j)\}$. Then, clearly, $\phi(\{i, j\}) = \{\pi^{-1}(i), \pi^{-1}(j)\}$ and thus ϕ is also a relabeling. Thus, by Corollary 4.5, there is a rigid transformation and a relabeling that maps P to Q. □

Corollary 4.8. *We note that the above proof gives a general position condition such that when P satisfies the condition, P is reconstructible from pairwise distances.*

This proof can be generalized for points in $\mathbb{R}^m$ for $n \geq m + 2$ Similarly, $m + 1$ points in $\mathbb{R}^m$ form a volume of zero in $m + 1$ dimensions, so there is a determinant with the pairwise distances that equals zero. This determinant is then, as before, expressed as a polynomial and it is possible to show that if for all choices of indices satisfying the same condition as before, the polynomial is nonzero, then there exists a relabeling for the point configuration and thus the point configuration is reconstructible from pairwise distances.

5 Computation

[BK06] also experimented with using the above proof to check if certain point sets are reconstructible. Note that this test does not always work because there are many point sets which are reconstructible but do not satisfy the general position condition in the proof. However, they use an exhaustive search of indices $a, b, c, d, e, f, g, h, i, j, k$, resulting in an $O(n^{11})$ algorithm which is not very practical. Even for $n = 8$, their test takes 58375 seconds to complete. An interesting open problem is if there exists a polynomial time algorithm for checking reconstructibility.

6 Extensions

[BK03] extend the results to matching under oriented rigid transformations, that is rotations and translations, and also the case of rotations, translations, and scalings. The case for oriented rigid transformations is shown using a different invariant that works for rotations and translations, and the case for scalings is done by rescaling the distribution of distances by normalizing the maximum distance. [BK07] considers the case of graphs with edge weights and notes that if all edge weights are distinct, then graphs are determined up to isomorphism by the distribution of their sub-triangles, which is the distribution of unordered triples of distances between three nodes.

7 Open Problems

- The first open problem that follows directly from the results of [BK06] is if there actually exists a polynomial time algorithm to test for reconstructibility.
- A second open problem, motivated by the practical approach in [OFCD] is to show that for most point configurations, if the distance between the distributions of pairwise distances is small, then the Gromov-Hausdorff distance between the point sets is also small. This is slightly related to Theorem 4.3.

References

[B] G. S. Bloom. "A Counterexample to a Theorem of S. Piccard," *J. Comb. Theory, Ser. A 22(3)*, 1977.

[BK04] M. Boutin and G. Kemper. "On Reconstructing n-Point Configurations from the Distribution of Distances or Areas," *Advances in Applied Mathematics*, 2004.

[BK06] M. Boutin and G. Kemper. "Which Point Configurations are Determined by the Distribution of their Pairwise Distances," 2006.

[BK07] M. Boutin and G. Kemper. "Lossless Representation of Graphs using Distributions," 2007.

[M] J. Matoušek "Lectures on Discrete Geometry" *Springer-Verlag*, 2002.

[OFCD] R. Osada, T. Funkhouser, B. Chazelle and D. Dobkin. "Shape Distributions," *ACM Transactions on Graphics*, 2002.