

FAITHFUL SHAPE REPRESENTATION FOR 2D GAUSSIAN MIXTURES

Mireille Boutin and Mary Comer

School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN, USA

ABSTRACT

It has been recently discovered that a faithful representation for the shape of some simple distributions can be constructed using invariant statistics [1, 2]. In this paper, we consider the more general case of a Gaussian mixture model. We show that the shape of generic Gaussian mixtures can be represented without any loss by the distribution of the distance between two points independently drawn from this mixture. In other words, we show that if their respective distributions of distances are the same, then there exists a rigid transformation mapping one Gaussian mixture onto the other. Our main motivation is the problem of recognizing the shape of an object represented by points given noisy measurements of these points which can be modeled as a Gaussian mixture.

Index Terms—Object recognition, shape, invariant statistics.

1. INTRODUCTION

Many applications depend on being able to identify or retrieve objects based on their shape. A good shape representation method is crucial for being able to do this quickly and effectively [3, 4]. It is often reasonable to approximate the objects to be recognized by a set of distinguished points called *landmarks* [5]. For example, the minutiae of a fingerprint form a configuration of points, or *point-set*, which can be used to decide what are the best matching candidates in a database of fingerprints. In many situations, including the case of minutiae, it may be difficult to label all the points accurately. One thus seeks a faithful representation which is invariant under a relabeling of the points. Boutin and Kemper [1] considered the deterministic problem of determining whether two point-sets are the same up to a rigid transformation. They showed that the distribution of the pairwise distances between a generic point-set is a faithful representation of the shape of this point-set. Thus, point-sets with exactly the same shape can be easily identified simply by comparing the distribution of their distances. We seek a modification of Kemper and Boutin’s method which can be applied to the case where the positions of the points are measured with some error. In particular, we seek a way to quantify the probability that the observed point samples come from distributions which have the same shape.

In a previous publication [2], it was shown that the shape of a generic mixture of 2D spherical Gaussians, each equally weighted and with the same standard deviation, can be represented without any loss by the distribution of the distance between a pair of points drawn independently from this spherical Gaussian mixture distribution. In this paper, we show how to extend this result to the case of a 2D Gaussian mixture. This is an important step towards finding an efficient shape comparison method for objects represented by landmarks. Indeed, our result implies that the difficult problem of comparing the underlying distributions of sets of point samples up

to a rigid transformation boils down to the simpler problem of comparing the underlying distributions of distances.

2. THE DETERMINISTIC CASE

In this section, we summarize the result of Boutin and Kemper [1] for the specific case that concerns us (theirs is proved in a much more general setting). It will be used in the next section for proving our result. We consider a point-set in the plane. Let us denote the points by $p_1, \dots, p_n \in \mathbb{R}^2$. Then we consider the set of all pairwise distances (squared, for simplicity) between the points:

$$\Delta_{ij} = \|p_i - p_j\|^2 \text{ for all } i, j = 1, \dots, n, i \neq j.$$

We remove the pair of indices associated to each distance and only store its value. In other words, we consider the bag of all distances, i.e. the unordered set of all Δ_{ij} ’s of the point-set p_1, \dots, p_n :

$$\{\Delta_{ij}\}_{i \neq j}.$$

It turns out that the bag of distances provides a faithful representation for the shape of generic point-sets, i.e. it faithfully represents most point-sets up to a global rotation, reflection and translation (so-called rigid transformation). Obviously, two-point sets which are related by a rigid transformation also have the same bag of distances. However, having the same bag of distances does not necessarily imply having the same shape. A counter-example, which was presented in [1], is given by the point-sets

$$\{(0, 0), (4, 0), (3, 1), (3, -1)\}, \quad (1)$$

$$\{(0, 0), (4, 0), (1, -1), (3, -1)\}, \quad (2)$$

which have different shapes but have the same bag of pairwise distances:

$$\{\sqrt{2}, \sqrt{2}, 2, \sqrt{10}, \sqrt{10}, 4\}.$$

Fortunately, such counter-examples are extremely rare because they must satisfy a polynomial equation, as stated in the following theorem.

Theorem 1. [6] *There exists a polynomial f in $2n$ variables such that if the points $p_1, p_2, \dots, p_n \in \mathbb{R}^2$ satisfy $f(p_1, p_2, \dots, p_n) \neq 0$, then for any other point-set $\bar{p}_1, \bar{p}_2, \dots, \bar{p}_n$ having the same bag of distances as that of p_1, p_2, \dots, p_n , there exists an orthogonal matrix $M \in \mathbb{R}^{2 \times 2}$, a translation vector $T \in \mathbb{R}^2$ and a permutation $\pi \in S_n$ such that*

$$\bar{p}_i = Mp_{\pi(i)} + T, \text{ for all } i = 1, \dots, n.$$

The point-sets which do not lie on the zero set of the polynomial f are called *generic* point-sets. What the above theorem says is that generic point-sets do not share their bag of distances with any other point-set, unless this point-set also has the same shape.

Numerical experiments were conducted in [6] to estimate the likelihood of encountering a non-generic point-set when choosing the points on a fixed grid. Although it was observed that non-generic point sets are potentially quite likely hit on a small grid, the results of the experiments suggest that they are almost never encountered when the points are chosen randomly, up to 15 digit precision, on a unit square of dimension one with a uniform distribution. The bag of distances thus appears to be a very good representation for the shape of a point-set, even when the points coordinates are specified by floating point values.

It would be tempting to try to show that bags of distances that are close, in some sense, come from point-sets which have a similar shape. Unfortunately, it is not true, even if we try to restrict the statement to generic point-sets. This is because of the presence of those counter-examples. Indeed, one can pick a generic point-set which is close to Point-set 1 and another generic point-set which is close to point-set 2 (as in Figure 1 a) and b)). Obviously, these generic point-sets would have quite different shapes, but, by construction, their bag of distances would be very similar. We thus need to follow another route to generalize the bag of distance representation method to the case of points observed under noisy conditions. The route we plan to follow is based on the result presented in this paper.

3. A FAITHFUL SHAPE REPRESENTATION FOR 2D GAUSSIAN MIXTURES

We consider the case where the points observed are drawn from a Gaussian mixture. The question we seek to answer is: "Given two sets of points, each drawn from a Gaussian mixture, determine the probability that they come from the same Gaussian mixture, up to a rigid transformation?". We now show that this is equivalent to asking: "What is the probability that their pairwise distances come from the same distribution?".

Let $\rho(x)$ be a Gaussian mixture model for a random variable $x \in \mathbb{R}^2$,

$$\rho(x) = \sum_{i=1}^n \alpha_i \rho_i(x) = \sum_{i=1}^n \frac{\alpha_i}{(2\pi)^{\frac{n}{2}} |\Sigma_i|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu_i)^T \Sigma_i^{-1} (x-\mu_i)},$$

where $\Sigma = \begin{pmatrix} \sigma_{1,i}^2 & \sigma_{12,i}^2 \\ \sigma_{12,i}^2 & \sigma_{2,i}^2 \end{pmatrix}$, $\alpha_i > 0 \forall i$ and $\sum_{i=1}^n \alpha_i = 1$.

Now let $\mathbf{p}_1, \mathbf{p}_2$ be two random variables chosen independently at random according to $\rho(x)$. Denote by Δ the square of the distance between these two points,

$$\Delta = \|\mathbf{p}_1 - \mathbf{p}_2\|^2.$$

Let us denote the pdf of Δ by $r(\Delta)$. We call $r(\Delta)$ the distribution of distances of the Gaussian mixture $\rho(x)$. Since the distance between two points is invariant under a simultaneous rigid transformation of the two points, the distribution of distances of two Gaussian mixtures which have the same shape must be the same. We now show that, for generic Gaussian mixtures, the converse statement also holds. More precisely, we prove the following theorem.

Theorem 2. *Suppose that two Gaussian Mixtures $\rho(x), \bar{\rho}(x)$ are such that their respective means forms a generic point-set. Then $\rho(x)$ and $\bar{\rho}(x)$ have the same distribution of distances, $r(x) = \bar{r}(x)$,*




Fig. 1. Dissimilar shapes can have a similar bag of distances. The star symbols in a) and b) represent two exceptional point-sets which have a very different shape but exactly the same bag of distances. Any pair of point-sets similar to these two exceptional ones, such as the ones graphed with the dot symbols in a) and b), respectively, will have similar bag of distances. One thousand such pairs were randomly generated and their bag of distances were compared with the Kolmogorov-Smirnov goodness of fit test to obtain the probability that these distance samples come from the same underlying distribution. In c) and d), we display a histogram of the results when comparing the exceptional point-set of a) with the one nearby, and with the one close to the exceptional point-set of b), respectively. One can see that such a test poorly discriminates between similar/dissimilar shapes when these happen to lie near an exceptional pair of shapes.

if and only if they have the same shape, i.e. if and only if there exists an orthogonal matrix $M \in \mathbb{R}^{2 \times 2}$ and a translation vector $T \in \mathbb{R}^2$ such that

$$\rho(x) = \bar{\rho}(Mx + T).$$

In other words, the distribution of distances is a faithful representation of the shape of generic Gaussian mixtures.

Proof. The moment generating function for Δ is

$$\begin{aligned} M_i(t) &= E\left(e^{t\Delta_i}\right), \\ &= \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} e^{t\|x_1-x_2\|^2} \rho(x_1)\rho(x_2) dx_1 dx_2, \\ &= \sum_{i,j=1}^n \alpha_i \alpha_j \int_{\mathbb{R}^m} \int_{\mathbb{R}^m} e^{t\|x_1-x_2\|^2} \rho_i(x_1)\rho_j(x_2) dx_1 dx_2, \\ &= \sum_{i,j=1}^n \alpha_i \alpha_j M_{ij}(t), \end{aligned}$$

where $M_{ij}(t)$ is the moment generating function for the distribution of the square of the distance between two random variables drawn independently from $\rho_i(x)$ and $\rho_j(x)$, respectively. After integrating,

we find

$$M_{ij}(t) = \frac{-\pi |\Sigma_i + \Sigma_j|}{1 - 2t(\sigma_{1,i}^2 + \sigma_{2,i}^2 + \sigma_{1,j}^2 + \sigma_{2,j}^2) - 4t^2 |\Sigma_i + \Sigma_j|} \times \frac{\|\mu_i - \mu_j\|^2 + 2t^2 (\mu_i - \mu_j)^T (\Sigma_i + \Sigma_j) (\mu_i - \mu_j)}{e^{1 - 2t(\sigma_{1,i}^2 + \sigma_{2,i}^2 + \sigma_{1,j}^2 + \sigma_{2,j}^2) - 4t^2 |\Sigma_i + \Sigma_j|}}.$$

One thus sees that $M_{ij}(t)$ is a function defined by four parameters:

$$\begin{aligned} & \|\mu_i - \mu_j\|^2, \\ & \sigma_{1,i}^2 + \sigma_{2,i}^2 + \sigma_{1,j}^2 + \sigma_{2,j}^2, \\ & (\mu_i - \mu_j)^T (\Sigma_i + \Sigma_j) (\mu_i - \mu_j), \\ & |\Sigma_i + \Sigma_j|. \end{aligned}$$

By power series expansion, one can show that functions having the same form as the functions $M_{ij}(t)$ are linearly independent from each other, unless the values of their four parameters are the same.

Now consider two Gaussian mixtures, say

$$\rho(x) = \sum_{i=1}^n \alpha_i \rho_i(x) \text{ and } \bar{\rho}(x) = \sum_{i=1}^n \bar{\alpha}_i \bar{\rho}_i(x),$$

where $\rho_i(x)$ and $\bar{\rho}_i(x)$ are Gaussian distributions with parameters (μ_i, Σ_i) and $(\bar{\mu}_i, \bar{\Sigma}_i)$, respectively. Assume that $\mu_1, \mu_2, \dots, \mu_n$ are distinct and that they form a generic point-set in \mathbb{R}^2 . If the distribution of distances of these Gaussian mixtures are the same, then their moment generating functions are the same. By linear independence, this means that the set of quadruples of parameters involved in the moment generating function for the first distribution, along with their coefficient $\alpha_i \alpha_j$, is the same as the set of quadruples of parameters and coefficients $\bar{\alpha}_i \bar{\alpha}_j$ for the second distribution. In particular, their set of pairwise distances is the same. By Theorem 1, and since μ_1, \dots, μ_n is assumed to be a generic point-set, this means that there exists a rigid transformation mapping the point-set μ_1, \dots, μ_n to the point-set $\bar{\mu}_1, \dots, \bar{\mu}_n$. So, after a relabeling of the components of the second mixture, we can write

$$\alpha_i \alpha_j = \bar{\alpha}_i \bar{\alpha}_j \quad (3)$$

$$|\mu_i - \mu_j| = |\bar{\mu}_i - \bar{\mu}_j| \quad (4)$$

$$\begin{aligned} \sigma_{1,i}^2 + \sigma_{2,i}^2 + \sigma_{1,j}^2 + \sigma_{2,j}^2 = \\ \bar{\sigma}_{1,i}^2 + \bar{\sigma}_{2,i}^2 + \bar{\sigma}_{1,j}^2 + \bar{\sigma}_{2,j}^2, \end{aligned} \quad (5)$$

$$\begin{aligned} (\mu_i - \mu_j)^T (\Sigma_i + \Sigma_j) (\mu_i - \mu_j) = \\ (\bar{\mu}_i - \bar{\mu}_j)^T (\bar{\Sigma}_i + \bar{\Sigma}_j) (\bar{\mu}_i - \bar{\mu}_j), \end{aligned} \quad (6)$$

$$|\Sigma_i + \Sigma_j| = |\bar{\Sigma}_i + \bar{\Sigma}_j|, \quad (7)$$

for all $i, j = 1, \dots, n$. We can actually assume, after a rigid transformation of the second Gaussian mixture, that

$$\mu_i = \bar{\mu}_i, \text{ for all } i = 1, \dots, n.$$

By solving the resulting system of Equations, we obtain

$$\alpha_i = \bar{\alpha}_i, \Sigma_i = \bar{\Sigma}_i, \text{ for all } i = 1, \dots, n,$$

and thus the Gaussian mixtures $\rho(x)$ and $\bar{\rho}(x)$ are the same, after a rigid transformation. \square




Fig. 2. In general, the Kolmogorov-Smirnov test does not accurately distinguish the underlying distributions of the samples contained in the bag of pairwise distances. Several pairs of configurations of five points were randomly generated, such as the ones illustrated with the star symbols in a) and b) respectively. Then, a point-sets similar to the first point-set was randomly generated, as illustrated with the dot symbols in a). The distance samples contained in the bag of distances were compared using the Kolmogorov-Smirnov goodness of fit test to estimate the probability that these distance samples come from the same underlying distribution. Plot c) contains a histogram of the results obtained when comparing the two nearby point sets, and d), a histogram of the results when comparing the random pair. Here we see that this method poorly discriminates between similar/dissimilar shapes in general.

4. TOWARDS A GOODNESS OF FIT TEST FOR DISTANCE SAMPLES

Now that we know that the distribution of distances can be used to faithfully represent the shape of a Gaussian mixture, the next step in this research will be to figure out how to determine the probability that two sets of distance samples come from the same underlying distribution. More particularly, we are interested in the case where we are given n point samples, each coming from a distinct Gaussian of the Gaussian mixture, and in using their pairwise distances to determine the probability that the Gaussian mixtures are the same up to a rigid transformation.

One method that has been used to attack this problem consists in comparing the pairwise distance samples using the Kolmogorov-Smirnov goodness of fit test, or another measure of similarity. As pointed out in Section 2, because of the existence of exceptional counterexamples to Theorem 1, such an approach cannot accurately distinguish between different shapes.

The following simple set of numerical experiments should convey this fact in an unequivocal manner. We generated one thousand slight variations of the exceptional Point-set 1 and one thousand slight variations of its exceptional counterpart, Point-set 2. The variations were randomly generated by sampling points within a uniformly distributed square region around each point. The size of the region used was 0.05, thus the resulting point-sets were extremely close to each other. Using the Kolmogorov-Smirnov goodness of fit

test, we compared the variations of Point-set 1 with Point-set 1 itself. More precisely, we computed the probability that the distance samples contained in the bags of distances of the variations of Point-set 1 come from the same distribution as those of Point-set 1. A histogram of the results is plotted in Figure 1 c). The same method was used to compare the shape of Point-set 1 with that of the variations of Point-set 2. A histogram of the results is plotted in Figure 1 d). As one can see, it is impossible to distinguish which histogram corresponds to similar shapes, and which one corresponds to dissimilar ones.

To illustrate the fact that what happened in the above experiment is a generalized problem, we carried out a second set of numerical experiments where point-configurations were randomly generated and compared. For each experiment, we generated three distinct point-sets: two randomly generated configuration of 5 points (assumed to have a different shape) together with one small perturbation of the first randomly generated point-set (yielding a similar shape). An example is illustrated in Figure 2 a) b). The bag of distances of the two nearby configurations were compared using the Kolmogorov-Smirnov test. The same was done for the two point-sets with different shapes. This experiment was repeated one thousand times. A histogram of the results is plotted in Figure 2 c) and d) respectively. The second histogram shows a very high number of false matches and thus this method fails at accurately distinguishing shapes. One main reason for this is that the Kolmogorov-Smirnov test is meant to be used on independent samples, which is not the case when one uses all the pairwise distances of a point set. Moreover, the Kolmogorov-Smirnov test is a limit result, so it may or may not give an accurate solution when using merely a finite number of samples. In future work, we shall show how to address these issues.

5. CONCLUSION

We have shown that the shape of a generic Gaussian mixture can be represented without any loss by the distribution of the square of the distance between two points independently drawn from this Gaussian mixture. By generic, we mean that the means of the Gaussian mixture should not form an exceptional point-set as defined by Boutin and Kemper in [1]. This means that comparing two Gaussian mixtures up to a rigid transformation (i.e. a rotation, translation and reflection) is, in most cases, equivalent to comparing their underlying distribution of distances. The motivating application of this work is the problem of recognizing the shape of objects represented by points. We have shown that using the pairwise distances between the points as distance samples and comparing these distance samples using the Kolmogorov-Smirnov goodness of fit test is not a reliable way to identify similar shapes. In future work, we will develop a more effective statistical test to estimate the probability two given point-sets have the shape.

6. REFERENCES

- [1] Mireille Boutin and Gregor Kemper, "On reconstructing n -point configurations from the distribution of distances or areas," *Adv. Appl. Math.*, vol. 32, pp. 709–735, 2004.
- [2] Kuryung Lee, M. Boutin, and Mary Comer, "Lossless shape representation using invariant statistics: the case of point-sets," in *Proc. Asilomar Conf. Sig., Syst., and Comp.*, Pacific Grove, CA, October 29–November 1 2006.
- [3] Philip Shilane, Patrick Min, Michael M. Kazhdan, and Thomas A. Funkhouser, "The princeton shape benchmark," in *SMI*. 2004, pp. 167–178, IEEE Computer Society.
- [4] Dmitriy Bespalov, Cheuk Yiu Ip, William C. Regli, and Joshua Shaffer, "Benchmarking cad search techniques," in *Symposium on Solid and Physical Modeling*, Leif Kobbelt and Vadim Shapiro, Eds. 2005, pp. 275–286, ACM.
- [5] Luciano Da Fontoura Costa and Roberto Marcondes Cesar Jr., *Shape Analysis and Classification: Theory and Practice*, CRC, 2000.
- [6] Mireille Boutin and Gregor Kemper, "Which point configurations are determined by the distribution of their pairwise distances," *Int. J. Compt. Geometry and Appl.*, in press.