



Robust ellipse detection with Gaussian mixture models



Claudia Arellano¹, Rozenn Dahyot

School of Computer Science and Statistics, Trinity College Dublin, Ireland

ARTICLE INFO

Article history:

Received 27 August 2014

Received in revised form

30 October 2015

Accepted 19 January 2016

Available online 28 January 2016

Keywords:

Ellipse detection

L2 distance

GMM

Parameter estimation

ABSTRACT

The Euclidian distance between Gaussian Mixtures has been shown to be robust to perform point set registration (Jian and Vemuri, 2011). We propose to extend this idea for robustly matching a family of shapes (ellipses). Optimisation is performed with an annealing strategy, and the search for occurrences is repeated several times to detect multiple instances of the shape of interest. We compare experimentally our approach to other state-of-the-art techniques on a benchmark database for ellipses, and demonstrate the good performance of our approach.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Estimating the parameters of a shape is a problem that arises in different fields in computer vision. The problem is usually classified according to the parameters to be estimated, the knowledge or information about the shape and the type of observations that have been collected. Jian and Vemuri [1] proposed a method for estimating registration parameters between sets of points based on the \mathcal{L}_2 distance between density functions. Each data set is used for modelling a GMM. The transformation parameters between the two sets are computed by minimising the Euclidean distance (\mathcal{L}_2) between those two density functions. This metric has the advantage of having closed form solution when the density functions are Gaussian mixtures. We have explored this metric and the modelling of the GMM when estimating rigid transformation and for morphable model fitting [2–4]. We have shown how using non-isotropic Gaussians to represent shapes better can be beneficial for the robustness and accuracy of the results [5]. Following our previous work we propose here:

- To extend the framework based on \mathcal{L}_2 to estimate a parametric family of curves (i.e. ellipses).
- To propose a multidimensional modeling for the density functions in order to include additional information available with little impact on the computational efficiency of the approach.
- We propose a method for detecting multiple instances of an ellipse. This method is applied to ellipse detection in 2D point clouds and images. We evaluate the performance of our method with comparison to the state-of-the-art [6–10].

The remaining of the paper is structured as follows. In Section 2 we review the most relevant methods for estimating ellipses classified according to the strategy used. In Section 3 we propose our \mathcal{L}_2 -based method for estimating the parameters of an ellipse. Section 5 extends our work to higher dimensional space allowing the integration of additional shape information. A second extension of our framework is presented in Section 4.2 for estimating multiple instances of ellipses. The experiments performed and results obtained are discussed in Section 6. Comparison with the state-of-the-art techniques is also provided in order to assess our proposed method. Finally, the conclusions of our work are presented in Section 8.

2. Literature review


Ellipse fitting is a challenging problem that arises in several fields. Some examples of applications are segmentation of cells [11], study of galaxies [12], medical diagnostics [13], camera calibration and face detection among others [14,15]. As many applications as there are of fitting ellipses there are also a great number of algorithms proposing solutions to this problem [16]. They are commonly classified in three categories: Least Square based methods, Hough Transform based methods and the most recent approach known as edge contour following methods.

2.1. Least Square based methods

Least Square based methods are usually classified into two categories according to the cost function to optimise. Those categories are (1) methods minimising a *geometric* error and (2) methods minimising an *algebraic* error [17–20]. The minimisation of the geometric error are

E-mail address: arellanc@tcd.ie (C. Arellano).

¹ Tel.: +56 974351525.



Ellipse represented with $N = 20$ samples from the parametric curve (Eq. 2) with the angular distance $\Delta = \frac{2\pi}{N}$ between two points.




Fig. 1. Representation of an ellipse with GMMs computed using different values of orthogonal bandwidth h and various sampling rate N .

regarded as the most accurate methods for ellipse fitting and various computational schemes have been proposed [21]. However, those methods have several drawbacks. They need several iterations for solving the non-linear optimisation problem and they are very sensitive to noise. Therefore, the convergence of those methods is not guaranteed. Moreover, the convergence often depends on the accuracy of the initialisation.

Algebraic methods on the other hand are easier to implement and computationally efficient. However, the main problem of those methods is that they do not guarantee the result which will be an ellipse. A normalisation process is required in order to enforce the solution [22,23]. For instance, Szpak et al. [24] propose a penalty function that guaranteed an ellipse when using the Sampson distance. However, the result is then biased by the normalisation scheme chosen. Algebraic methods are less robust with respect to geometric methods when the data is coming from a small segment of an ellipse. Furthermore, accuracy of the results of those methods often depends on the initialisation. A method for

finding a reasonable starting guess for those algorithms is proposed in [25].

2.2. Hough Transforms

The Hough Transform is a well known approach to detect ellipses [26–28]. It is based on a voting system where each edge pixel of an image is considered. This voting procedure is carried out in a parameter space, from which candidate ellipses are obtained as local maxima in an accumulator space that is explicitly constructed based on the parameters of the ellipse. Several algorithms have been proposed to improve performance of the HT method [29,9,30]. Some approaches explore the inclusion of additional information such as the directional property of the pixels [31,32]. Unfortunately, those methods are easily affected by possible noise in the image. A different strategy that improves the computational complexity of the HT is to sub-sample the data set. For instance, Kiryati et al. [33] used the Probabilistic Hough Transform (PHT) where just a portion of the edge pixel of an image

is used. Xu et al. [34] on the other hand, proposed the Randomised Hough Transform (RHT) which used randomly chosen n -tuples of data points. This method was originally designed for detecting circles but it was extended to ellipses by McLaughlin et al. [10]. The Randomised Hough Transform serves as a powerful variant of the standard Hough Transform that exploits the geometric properties of ellipses in order to speed up the detection process [35,36]. Despite its simplicity and efficiency, the RHT performs poorly if the target ellipses overlap (or mutually-occlude) with each other.

A general advantage of Hough Transform based methods is that they do not require connectivity in between consecutive edge pixels. This makes these algorithms useful when the observation is a sparse set of points. The main drawback is that they are very sensitive to the choice of the quantisation of the parametric space. Incorrect quantisation leads to the detection of false ellipse or missing the true ellipse. Furthermore, the performance of the HT algorithms deteriorate when the number of ellipses in the image increases.

2.3. Edge following methods

Edge following methods exploit the connectivity between edge pixels to detect ellipses [8,7,11,6,37]. The main strategy of these methods is to detect arcs and then group them in order to detect the ellipses. Kim et al. [8] and Mai et al. [7] connect pixels by line fragments from where the arcs are computed. The grouping of arcs becomes a critical process since errors in that stage would be propagated to the ellipse detection step. Chia et al. [6] introduce some improvements by introducing a self-correction stage where

the grouping process is corrected using a feedback loop. In other words, low confidence ellipses are replaced by a set of better hypothetical ellipses obtained by combining arcs from different groups. Those methods are considered as the most successful in detecting multiples of ellipses in digital images. However, they only work when there is connectivity along the edge pixels. On the other hand, if the observations are a sparse set of data points, this strategy is inadequate.

3. Ellipse detection using the \mathcal{L}_2 distance

We propose a strategy using the Euclidian distance \mathcal{L}_2 between GMMs for robustly estimating the parameters of an ellipse occurring on a 2D point cloud or an image. Gaussian mixture models are used to represent shapes for both the observations (denoted f) and the ellipse model (denoted g_θ) with θ being the parameters that define the ellipse $\theta = [\gamma, a, b, x_o, y_o]$ (where the centre of the ellipse is (x_o, y_o) , the semi-minor and semi-major lengths a and b respectively and γ the angle of rotation of the ellipse with respect to the horizontal axis). The parameters of the ellipse can then be estimated by minimising the Euclidean distance between the two density functions f and g_θ . This can be expressed as follows:

$$\hat{\theta} = \arg \min_{\theta} \{ \mathcal{C}(\theta) = \mathcal{L}_2(\theta) = \|f - g_\theta\| \} \quad (1)$$

In the next two sections we describe the modelling of those density functions f and g_θ .

3.1. Modelling the ellipse g_θ

A point $u^{(j)}$ on an ellipse with parameter θ can be computed with the following parametric equation:

$$u^{(j)} = \begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix} \begin{pmatrix} a \cos \tau_j \\ b \sin \tau_j \end{pmatrix} + \begin{pmatrix} x_o \\ y_o \end{pmatrix} \quad (2)$$

given $\tau_j \in [0; 2\pi]$. The model g_θ uses $N=20$ equally spaced points from the ellipse $\{u^{(j)}\}_{j=1, \dots, 20}$ to define a non-isotropic GMM for g_θ :

$$g_\theta(x) = \sum_{j=1}^N w_j \mathcal{N}(x; \mu_j, \Sigma_j) \quad (3)$$

where $\mathcal{N}(x; \mu_j, \Sigma_j)$ indicates the Normal density function for random vector x (here $x \in \mathbb{R}^2$), with mean μ_j and covariance matrix Σ_j . The mean μ_j can be chosen as the mean between consecutive vertices $(u^{(j)}, u^{(j+1)})$. The covariance matrix $\Sigma_j = Q_j^T \Lambda_j Q_j$ is computed using a rotation matrix $Q_j = [\vec{n}_{1j} | \vec{n}_{2j}]$ composed of two orthogonal unit vectors, \vec{n}_{2j} that has the same direction as the vector defined with consecutive vertices $(u^{(j)}, u^{(j+1)})$ on the




Fig. 2. Perspective visualisation of g_θ (also shown in Fig. 1(c)). Our Gaussian mixture model g_θ creates a ridge forming the ellipse. The height of the ridge is not quite even and varies depending on the choice of bandwidth h , and the parameters θ . H_{ellipse} provides an estimate of the height of the ridge.

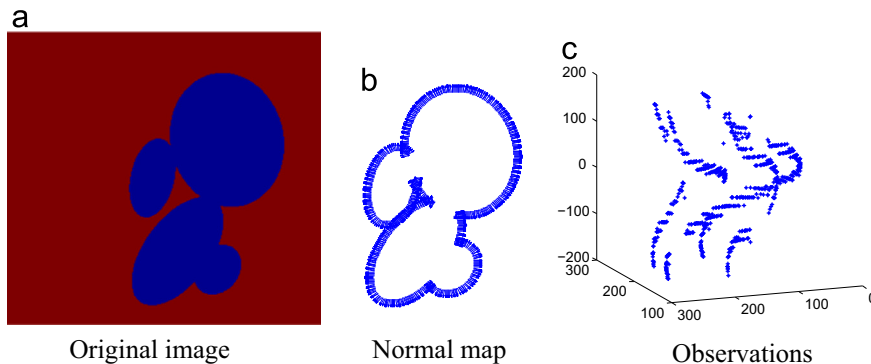


Fig. 3. Observations extracted from the image (a). We use the edge map and the gradient of the image to compute the position and orientation of the points of interest. In (b) we show the map of the normal vector associated to the edge map. In (c) we show the computed observation $\{(v_i, \psi_i)\}_{i=1, \dots, n} \in \mathbb{R}^3$.

contour, and \vec{n}_{1j} its orthogonal unit vector (see Fig. 1(a) for illustration). The diagonal matrix Λ_j is set to:

$$\Lambda_j = \begin{pmatrix} h^2 & 0 \\ 0 & h_{tj}^2 \end{pmatrix} \quad (4)$$

where parameter $h_{tj} = \|u^{(j)} - u^{(j+1)}\|$ and the parameter h that controls the fuzziness in the normal direction to the contour is set by the user. The weights are computed as $w_j \propto h_{tj} h$ subject to $\sum_{j=1}^N w_j = 1$. Fig. 1(b)–(d) shows an ellipse represented by several GMMs for various values of the orthogonal bandwidth h . Fig. 1(e)–(g) shows the impact of the number N of Gaussians: if too few points are taken, the GMM representation does not provide an even ridge along the contour of the ellipse (e.g. Fig. 1(e)). The choice $N=20$ was found to be a good choice for capturing the shape of the ellipse and adding more Gaussians in the mixture g_θ does not change much the model (e.g. Fig. 1(c) and (g)).

3.2. Modelling the observations f

The GMM f modelling the observations, depends on the information available and the structure of the observations. For instance, let's assume we have as observation a set of points $\{v^{(i)}\}_{i=1,\dots,n}$. In this case there is no information about how those points are connected. Therefore, we define the GMM using isotropic covariance matrices, with the mean of each Gaussian chosen as the observation itself $v_i = v^{(i)}$, $\forall i$:

$$f(x) = \frac{1}{n} \sum_{i=1}^n \mathcal{N}(x; v_i, h^2 I) \quad (5)$$

When observations are extracted from an image $\mathcal{I}(x, y)$, we can use the edge pixel locations $\{v^{(i)}\}_{i=1,\dots,n}$ as found by an edge detector (e.g. Canny [38]). In addition, the gradient of the image $\nabla \mathcal{I} = (I_x, I_y)^T$ is also available and can be used to model non-isotropic covariance matrices. In this case the GMM for f can be expressed as follows:

$$f(x) = \sum_{i=1}^n w_i \mathcal{N}(x; v_i, \Sigma_i) \quad (6)$$

The mean v_i is chosen as the edge pixel location $v^{(i)}$ and the covariance $\Sigma_i = Q_i^T \Lambda Q_i$ is computed with $Q_i = [\vec{n}_{1i} \mid \vec{n}_{2i}]$ where the orthogonal unit vector is available from the gradient $\vec{n}_{1i} = \frac{\nabla \mathcal{I}(v_i)}{\|\nabla \mathcal{I}(v_i)\|}$ and the tangent vector at location v_i is

$$\vec{n}_{2i} = \frac{1}{\|\nabla \mathcal{I}(v_i)\|} \begin{pmatrix} \mathcal{I}_y(v_i) \\ -\mathcal{I}_x(v_i) \end{pmatrix}$$

The diagonal matrix Λ is

$$\Lambda = \begin{pmatrix} h^2 & 0 \\ 0 & h_t^2 \end{pmatrix}$$

where h_t can be chosen proportional to the width of the pixel (i.e. $h_t=1$ is the distance between two neighbouring pixels on the image grid). h controls the fuzziness in the normal direction to the curve and is set by the user. The weights are computed as $w_i \propto h_{tj} h$ subject to $\sum_{i=1}^n w_i = 1$.




Fig. 4. Figure shows three examples of observation used in the experiments. The ellipses were computed using the target ellipse for which the parameters are $\theta = [\gamma, a, b, x_0, y_0] = [0, 3, 5, 0, 0]$ corrupted using Gaussian noise with standard deviation σ equal to 0.1, 0.3 and 0.5 respectively.

3.3. Parsimonious representation of the density functions

The parameters θ of the ellipse are estimated by minimising the Euclidean distance between f and g_θ expressed as

$$\|f - g_\theta\|^2 = \|f\|^2 + \|g_\theta\|^2 - 2\langle f | g_\theta \rangle \quad (7)$$

The term $\|f\|^2$ does not depend on θ and is not computed. The terms $\|g_\theta\|^2$ and $\langle f | g_\theta \rangle$ have an analytical solution using the following result between two Normals:

$$\langle \mathcal{N}(\mu_1, \Sigma_1) | \mathcal{N}(\mu_2, \Sigma_2) \rangle = \mathcal{N}(0; \mu_1 - \mu_2, \Sigma_1 + \Sigma_2) \quad (8)$$

The computational complexity of the multivariate Normal (8) is in the order of d^3 where d is the dimension of the random vector x modelled by f and g_θ . Using the Gaussian mixture parameters (means, weights, covariances) of g_θ as explained in Section 3.1, the term $\|g_\theta\|^2$ is a sum of N^2 weighted Normals:

$$\|g_\theta\|^2 = \sum_{j=1}^N \sum_{k=1}^N w_j w_k \mathcal{N}(0; \mu_j - \mu_k, \Sigma_j + \Sigma_k) \quad (9)$$

The computational complexity of $\|g_\theta\|^2$ is about $N^2 \times d^3$. Similarly the term $\langle f | g_\theta \rangle$ can be computed as the sum of $n \times N$ weighted Normals:

$$\langle f | g_\theta \rangle = \sum_{i=1}^n \sum_{k=1}^N w_i w_k \mathcal{N}(0; \mu_i - \mu_k, \Sigma_i + \Sigma_k) \quad (10)$$

with weights w_i , means μ_i and covariances Σ_i computed as explained in Section 3.2. The computational complexity of $\langle f | g_\theta \rangle$ is about $N \times n \times d^3$ and computation cost can be limited by controlling efficiently the values of N , n and d . In Section 5, x is augmented to a higher dimensional space (\mathbb{R}^3 as an alternative to \mathbb{R}^2) and we show experimentally how this additional dimension allows more powerful detection results (cf. Section 6.2). For such low values of dimensions $d=2,3$, the computational cost remains tractable. In the case of the GMM defined for the ellipse g_θ we control the number of Gaussians used ($N=20$) without compromising much the representation of the ellipse (cf. Fig. 1). The GMM defined for the observation f on the other hand depends on the data available. This data can be sub-sampled in order to decrease the number of Gaussians. The sub-sampling can be done directly on the image (reducing resolution) or by sub-sampling the edge map. In this case we select the connected edges automatically using a standard image processing algorithm in matlab (bwboundaries). The data can then be sub-sampled using a




Fig. 5. Computing the error rate Eq. (16). Figure (a) shows in blue the true ellipse (observation) and in red dots the estimated one. The black region in figure (b) represents the area of the difference in between the two ellipses. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

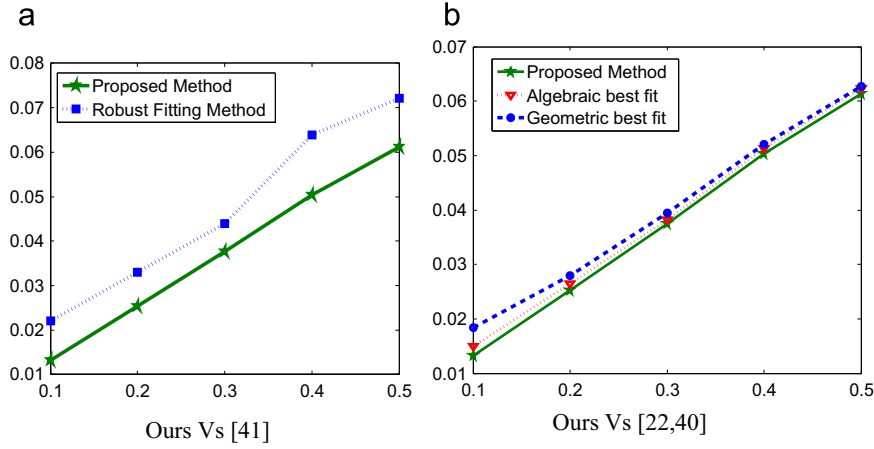


Fig. 6. Mean Error (y-axis) obtained for the data sets perturbed with 5 different standard deviation of noise (σ from 0.1 to 0.5, x-axis). In (a) we compare our proposed method (green star) versus a recent robust ellipse fitting algorithm reported by Yu et al. [41] (blue square). In (b) our proposed method (green star) is compared with two standard ellipse fitting algorithms [22,40]. The red triangle represents the algorithm which cost function minimises the algebraic distance while the blue dot is the well known Least Square method that is based on the minimisation of a geometric distance. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)




Fig. 7. Zoom on the errors reported Fig. 6 with 95% confidence intervals for each method.




Fig. 8. Robustness to 10% of outliers: mean error (green star corresponds to our method, Vs red triangle and blue dot [22,40]) is reported on the y-axis for different levels of white noise on the inliers. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

uniform rate along the connected edges in the edge map. Since the structure of the observation is preserved by the connectivity of the edge map a non-isotropic modelling can be used. This modelling with non-isotropic covariances helps to preserve the information contained on the data while reducing the number n of observations to use.

4. Algorithms for ellipse detection

We first present an algorithm to fit one ellipse in the observed data (Section 4.1). In a similar fashion as Jian and Vemuri [1], the optimisation is performed in a simulated annealing framework to avoid local solutions and limit the influence of the starting guess on the optimisation. Section 4.2 proposes a strategy for finding multiple instances of ellipse that can occur in the observations.




Fig. 9. Robustness to outliers: the red dots are the observations, the estimated ellipse with Least Squares (Blue), Direct Fit (black) and the proposed method (green). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)




Fig. 10. Results with pseudo-outliers: two (partial) ellipses perturbed with noise $\sigma = 0.1$ (observations in red). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

4.1. Ellipse detection with annealing

The only free parameter to set when computing \mathcal{L}_2 is the orthogonal bandwidth h used in the GMMs f and g_θ . This parameter plays two important roles in the proposed estimation framework. First, it affects the description of the shape. It controls the fuzziness in the normal direction to the curve and secondly it affects the convexity of the cost function. The optimisation is performed using Gradient based optimisation Algorithms (GA)

depending on the choice of the initial guess $\theta^{(0)}$ and the orthogonal bandwidth h :

$$\hat{\theta} \leftarrow \text{GA}(\mathcal{L}(\theta), \theta^{(0)}, h) \quad (11)$$

The larger the value for h the smoother the cost function. Therefore, to make our approach not sensitive to the initial guess $\theta^{(0)}$, we use a simulated annealing framework where the orthogonal bandwidth h is the temperature decreased with a geometric rate




Fig. 11. Examples from the dataset [6] used for testing the algorithm. This dataset contains 6 sets of 50 images and each set contains a different number of occluded ellipses from 4 to 24.




Fig. 12. Example of the results of our algorithm for detecting multiples ellipses using 3D GMMs (middle row) and 2D GMMs (bottom row) where orientation information is not used. The original images are shown on the top row. As can be seen, the results deteriorate significantly when the orientation information is not used.

(controlled by parameter β) starting from h_{max} and is stopped when the bandwidth is smaller than h_{min} . The use of simulated annealing helps in converging to the global solution. This optimisation is summarised in Algorithm 1.

Algorithm 1. One ellipse detection.

Require: h_{max} , h_{min} , β and $\theta^{(0)}$

```

Init  $h = h_{max}$  and  $\hat{\theta} = \theta^{(0)}$ 
repeat
   $\hat{\theta} \leftarrow \text{GA}(\mathcal{C}(\theta), \hat{\theta}, h)$ 
   $h \leftarrow \beta \times h$ 
until  $h \leq h_{min}$ 
return  $\hat{\theta}$ 

```

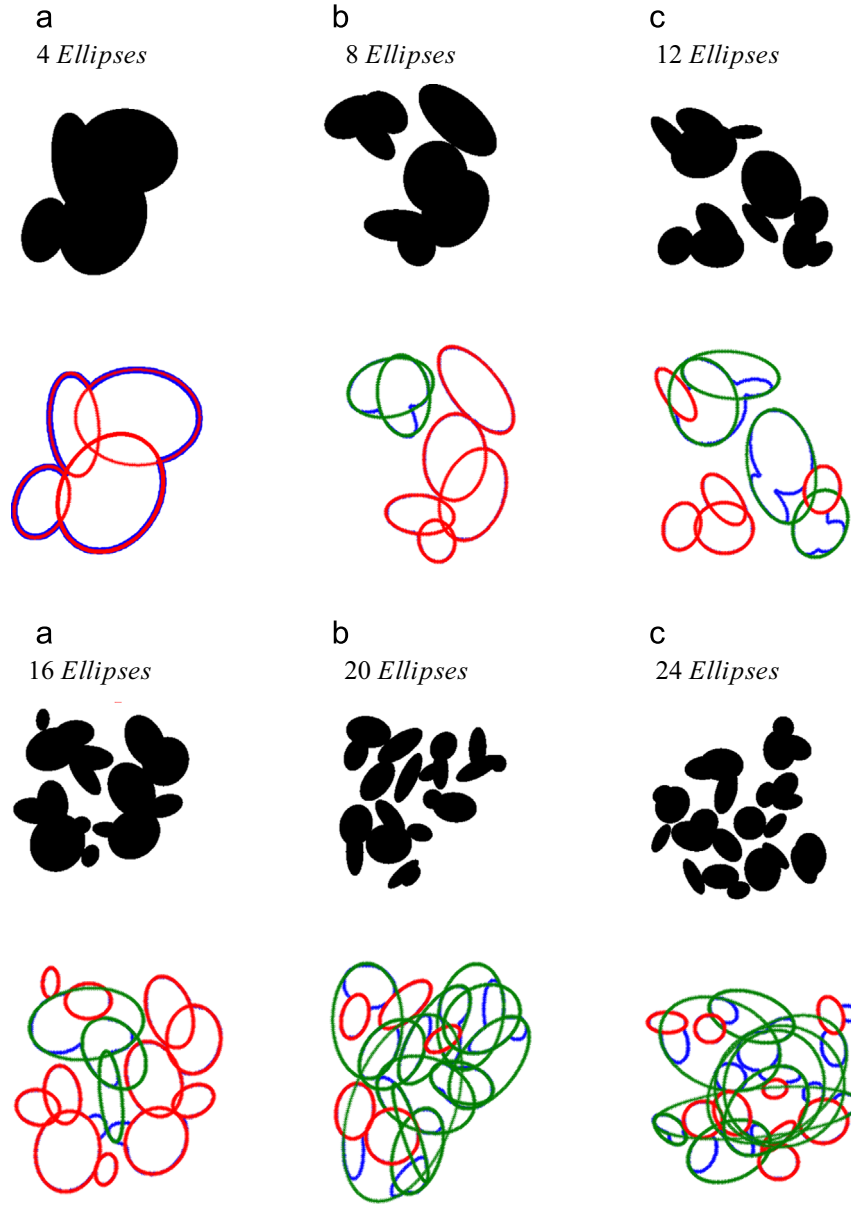


Fig. 13. Example of the results obtained when applying our algorithm to data sets containing 4, 8, 12, 16, 20 and 24 occluded ellipses respectively (observations in blue, ellipse detected in red, false positive in green). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)




Fig. 14. Testing on synthetic images containing occluded ellipses. In (a) we report the values obtained for the Recall while in (b) the results for precision. Each set (from 4 to 24) was evaluated using 50 images. For comparison we report the results obtained using the approaches proposed by Chia et al. [6], Mai et al. [7], Kim et al. [8] and the Hough transform based methods RHT and SHT proposed by [10] and [9] respectively.




Fig. 15. Testing of our algorithm on 50 synthetic images, each containing 4 occluded ellipses. The overlap error varies from 0.05 to 0.55.

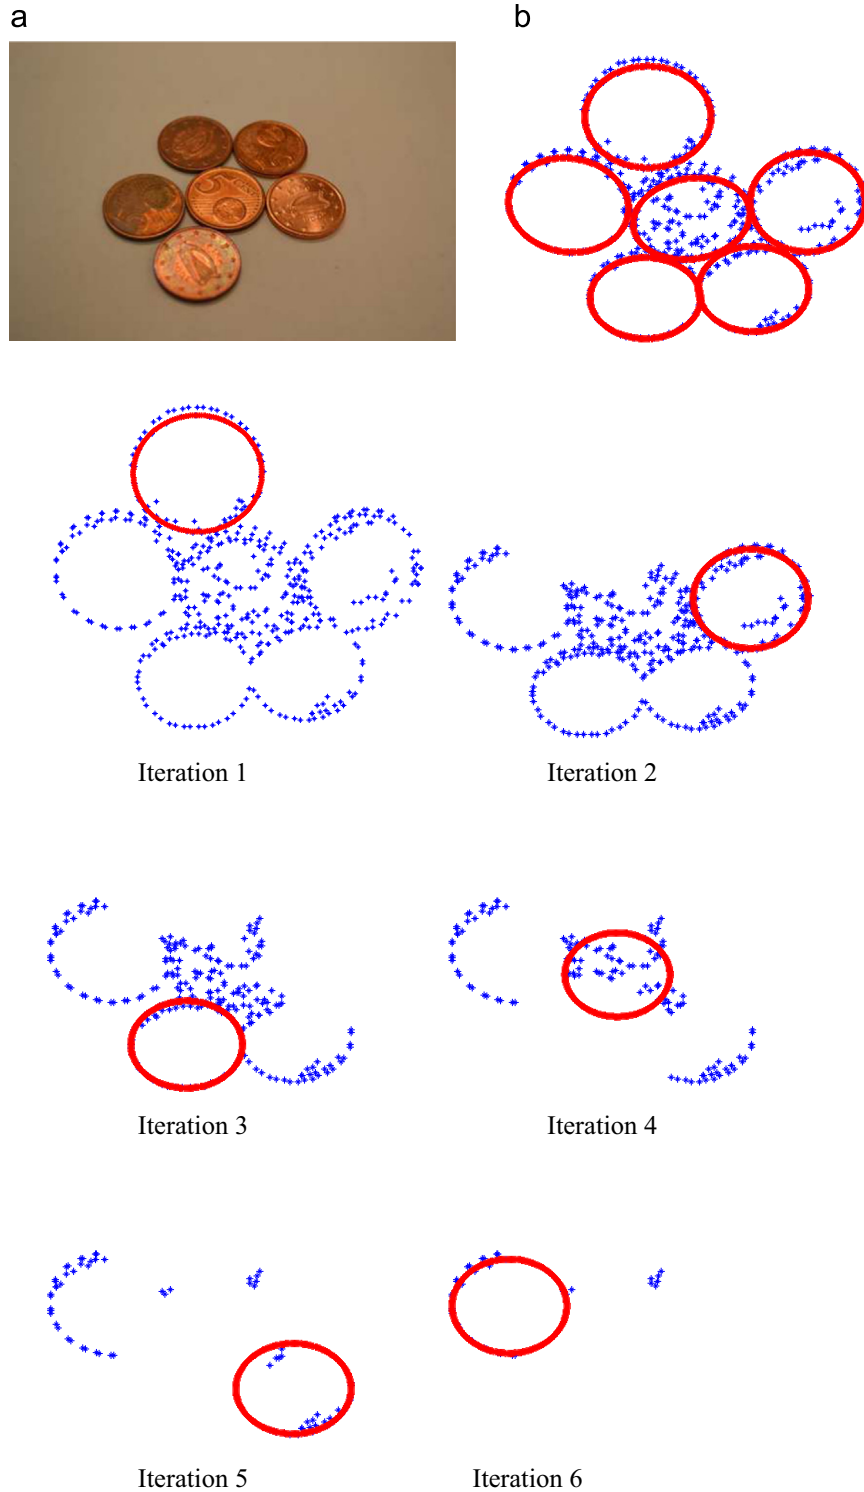


Fig. 16. Example of detection of multiples instance (i.e. coins). Illustration of the iteration of the proposed method for detecting multiple instances of a shape.

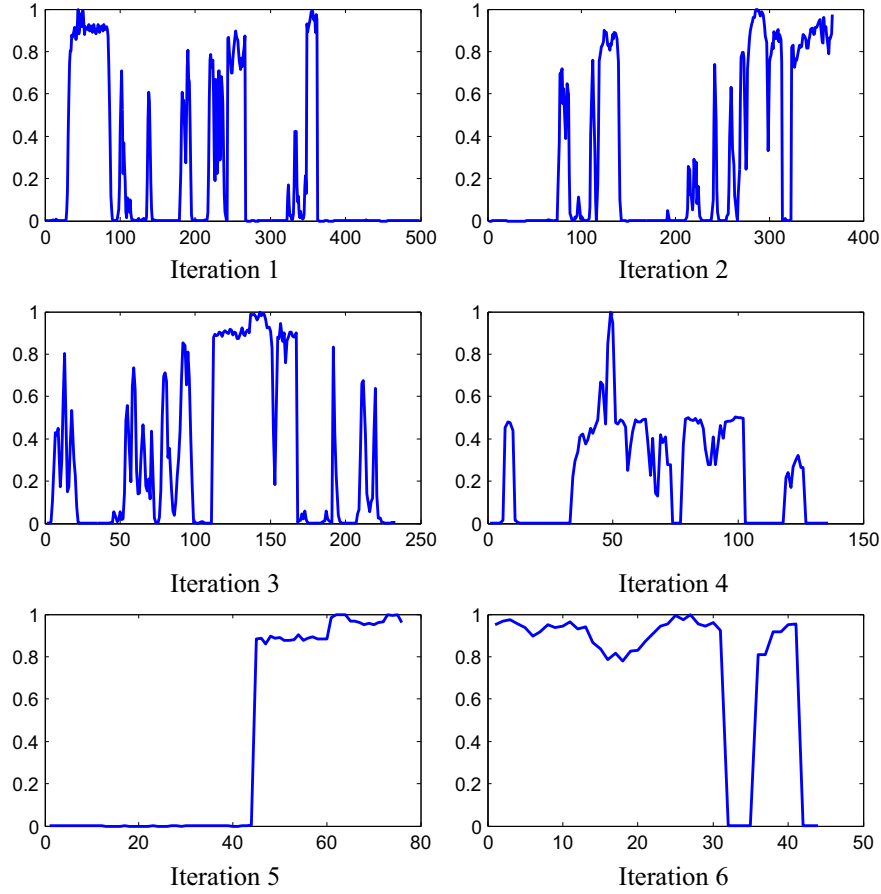


Fig. 17. Normalised score $g_{\hat{\theta}}(v_i)/H_{\text{ellipse}}$ of each observation v_i in the updated set S_f of being part of the detected shape $\hat{\theta}$ (see Eq. (12) and Algorithm 2).

4.2. Multiple ellipse detection

The goal is to estimate a set of ellipses $\mathcal{E} = \{\hat{\theta}_1, \dots, \hat{\theta}_s\}$ that represent all the s ellipses that occur in the observations. We propose an iterative algorithm based on the following steps:

1. *Global estimation*: This step estimates the parameters of the shape $\hat{\theta}$ given a set of observations by minimising the cost function proposed in Eq. (1). The general algorithm that solves this problem is described in Section 4.1 (cf. Algorithm 1).
2. *Update of the observation f*: Once an instance $\hat{\theta}$ of the shape is detected all the observations associated with that shape are removed. The remaining observations are then used to detect the following instances of the shape. The updated set of observations S_f keeps Gaussians that have mean that have poor score in the estimated model $g_{\hat{\theta}}$:

$$v_i \in S_f \quad \text{if } g_{\hat{\theta}}(v_i) < t_1 \times H_{\text{ellipse}} \quad (12)$$

$g_{\hat{\theta}}(v_i)$ is the probability density function value for the observation v_i and estimated ellipse parameters $\hat{\theta}$. t_1 is a threshold between 0 and 1 chosen by the user and H_{ellipse} is the average height of the ridge of the ellipse model:

$$H_{\text{ellipse}} = \frac{1}{N} \sum_{j=1}^N g_{\hat{\theta}}(\mu_j) \quad (13)$$

Fig. 2 illustrates this height H_{ellipse} of the ridge, and any observation v_i that are not high enough on that ridge remains in the observation set.

Those two steps are iterated until all instances of the shape are found: the algorithm is stopped when the updated set S_f has less than a given proportion t_2 of points left compared to the original set S of observations (cf. Algorithm 2). The values for the threshold t_1 and t_2 have been set to 0.3 and 0.1 respectively in this work. In other words the algorithm will stop when the remaining observations are less than 10% of the original observed data. Notation $|S|$ indicates the cardinal of set S .

5. Augmenting dimensionality d of the GMMs

We show here how to consider additional information about the shape of interest by augmenting the dimensionality of the density functions f and g . In a similar fashion as the Hough transform for finding lines in images [39], we propose to use both the vertices and the orientations of the shape. For instance in images, the means of the GMM are the vertices (edge pixel locations) and their orientations $v_i = \{(v^{(i)}, \psi_i)\}_{i=1 \dots n}$ with angle ψ_i computed with the gradient of the image:

$$\psi_i = \arctan \left(\frac{I_y(v^{(i)})}{I_x(v^{(i)})} \right) \quad (14)$$

The covariance matrix $\Sigma_i = Q\Lambda Q^T$ is similar to the one defined in Section 3.2 but augmented in one dimension such as Q :

$$Q = \begin{pmatrix} \vec{n}_{1i} & \vec{n}_{2i} & B^T \\ B & 1 \end{pmatrix}$$




Fig. 18. This figure shows results of detection of ellipses in images. In the first column (a) we show the original image. The second column (b) shows the data (edges) extracted from the image and used as input for our algorithm. In the third column (c) we show the ellipses obtained (in red) and superposed with the input data. (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

with $B = [0, 0]$. The diagonal matrix of bandwidths Λ is defined as follows:

$$\Lambda = \begin{pmatrix} h_t & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & h_\psi \end{pmatrix}$$

where h_ψ is the bandwidth associated with the orientation of the shape that has been included as an extra dimension. An example of the data extracted from an image is shown in Fig. 3. The domain of the density functions f and g is not limited to the spatial domain. It can be augmented by including extra information about the shape of interest. Increasing the dimension of the GMM can increase the amount of information about the shape to encode. We illustrate this idea using orientation as additional information. However, it can be extended to any other information about the shape such as colour, illumination, gradient, or motion among others.

Algorithm 2. Multiple ellipse detection.

Require: $t_1, t_2, \beta, h_{min}, h_{max}, S_f = S, s=0$ and $\theta^{(0)}$
repeat

s^{++}

Compute $C(\theta) = \|f - g_\theta\|$

$\hat{\theta}_s = \arg \min_{\theta} \{C(\theta)\}$ estimated with Algorithm 1

Update S_f

until $|S_f| < t_2 \mid S|$

return $\mathcal{E} = \{\hat{\theta}_1, \dots, \hat{\theta}_s\}$

6. Experiments and results

In Section 6.1 we assess the detection of one ellipse in 2D point clouds against state-of-the-art techniques [40,22,41]. In Section 6.2 we assess the problem of detecting multiples ellipses in images against state-of-the-art techniques [6–10]. The latent parameters to estimate are the centre of the ellipse (x_o, y_o) , the semi-minor and semi-major lengths a and b respectively and γ the angle of rotation of the ellipse with respect to the horizontal axis ($\theta = [\gamma, a, b, x_o, y_o]$). With the exception of the experiment reported in Section 6.1, the starting guesses $\theta^{(0)}$ for the optimisation are randomly chosen in all experiments such that the ellipse defined by $\theta^{(0)}$ should be contained in the bounding box where observations are.

6.1. Ellipse detection in 2D point clouds

In this section we evaluate the sensitivity to noise and the robustness of our proposed method when detecting an ellipse in 2D point clouds.

6.1.1. Sensitivity to noise

We define for this experiment a target ellipse (ground truth) centred at the origin and without rotation. The semi-major length equal to 5 and semi-minor length equal to 3. The ellipse is corrupted using five levels of Gaussian noise with standard deviation from 0.1 to 0.5. For each level of noise we take 50 points from the ellipse. For each level of noise we generate 50 data sets to perform




Fig. 19. Example of false positive: global minimum of our cost function (shown in green) does not correspond to the expected ellipse (shown in red). (For interpretation of the references to colour in this figure caption, the reader is referred to the web version of this paper.)

the experiments. Fig. 4 shows three data sets (used as observation) using level of noise 0.1, 0.3 and 0.5 respectively.

We run our algorithm using the following settings: $h_{max} = 0.9$, $h_{min} = 0.2$ and the geometric decreasing rate for the bandwidth $\beta = 0.8$. The initialisation for all the experiments is $\theta^{(0)} = [0, 1, 1, 0, 0]$, which is basically a circle centred on the origin with radius equal to 1. We compute the error between the estimated ellipse (given by $\hat{\theta}$) and the target (observation) using the normalised area of the symmetric difference [41]. This metric for error between the true ellipse E_t and the fitted ellipse E_f is defined as follows:

$$\text{Error} = \frac{S_{E_t \cup E_f} - S_{E_t \cap E_f}}{2S_{E_t}} \quad (15)$$

where $S_{E_t \cup E_f} - S_{E_t \cap E_f}$ is the area of the symmetric difference and S_{E_t} denotes the area of the true ellipse (cf. Fig. 5). In order to compute this error rate, we use the function phantom in Matlab to create an image containing an ellipse. The image is represented by assigning a value equal 1 to all pixels inside the ellipse and zero otherwise. One image I_t is created using the parameters of the target ellipse (observation). The second image I_f is created using an estimated ellipse. The symmetric difference between those two ellipses is computed as a function of the number of pixels with a value equal to 1 after adding the two images. The error rate in Eq. (15) is then re-expressed as follows:

$$\text{Error rate} = \frac{\text{pixel}(I_t + I_f = 1)}{2 \text{ pixel}(I_t = 1)} \quad (16)$$

The results obtained for all the data sets are reported in Fig. 6. In the left plot (a) we compare the results obtained using our proposed method (green stars) and the method proposed by Yu et al. [41].

In Fig. 6(b) we compare our method with the Non-Linear Least squares [40] and using the direct ellipse fitting algorithm [22]. In both cases the error between the estimated ellipse and the true ellipse is better minimised using our proposed method. The standard deviation of the error is also reported in Fig. 7.

6.1.2. Robustness to outliers

A second experiment is performed to evaluate the robustness of the algorithm to outliers. We use the same data as in the previous experiment (for all levels of noise) but adding 10% of extra points randomly distributed. As it is shown in Fig. 8, our proposed method does maintain its performance. On the contrary, the error obtained using the direct fitting and the Least Square methods increase [40,22]. Those methods are very sensitive to outliers while our proposed method keeps its robustness. Examples of the data used in the trial and the estimated ellipses are shown in Fig. 9. The results are similar when the outliers correspond to another ellipse present in the data. Examples of this case with pseudo-outliers are shown in Fig. 10. The mean error for our proposed method when 50 trials are tested is 0.0133 with a standard deviation of 0.0012. Those values are in correspondence with the values obtained when no outliers at all are presents in the data set. This shows the robustness of the proposed algorithm for fitting ellipses to noisy data sets.

6.2. Detecting multiples ellipses in images

In this section we explore the problem of detecting multiple ellipses in digital images. We use the position and orientation of the edge map of the images as information $v_i = \{(v_i^{(i)}, \psi_i)\}_{i=1 \dots n}$. This information is encoded in a multidimensional density function (as described in Section 5).

We use the benchmark data set provided by Chia et al. [6] for testing. This data contain 6 sets of synthetic images. Each set is created using a different number (4, 8, 12, 16, 20 and 24) of occluded ellipses. Each set contains 50 images with a resolution of 300×300 pixels. Fig. 11 shows a few exemplar images from that dataset. Input data for our algorithm is computed by using the Canny edge detector of Matlab with its default parameters. The normal vectors are computed using the convolution function with default parameters and applied in both directions of the image ($x-y$).

The detection evaluation is performed using the overlap error. This error is computed using the detected ellipse E_d and the true ellipse E_t as follows:

$$\text{OverlapError}(E_d, E_t) = 1 - \frac{S_{E_d \cap E_t}}{S_{E_d \cup E_t}} \quad (17)$$

An ellipse is considered detected when the overlap error is less than 0.05. A common index for detector evaluation is the F-measure. The F-measure combines the precision and recall of the detector. Precision is defined as follows:

$$P = \frac{\text{Number of correctly detected ellipses}}{\text{Total number of ellipses detected}} \quad (18)$$

and the Recall is:

$$R = \frac{\text{Number of correctly detected ellipses}}{\text{Total number of ellipses present in the test image}} \quad (19)$$

One of the advantages of our algorithm is its simplicity in terms of parameter setting. The only relevant value to set is the bandwidths used for the density function ($h_{max} = 7$ and $h_{min} = 1$) and the threshold used for updating the observation $t_1 = 0.3$ and $t_2 = 0.1$ (cf. Section 4.2). We define the covariance matrices as described in Section 5. In order to optimise the efficiency of the algorithm the




Fig. 20. GMMs (only considering 2D information) for the observation (a) and for the ellipse model as initial guess (b) for the example shown in Fig. 19. The top row was computed using bandwidth $h=10$ while at the bottom the bandwidth $h=1$ was used.

number of Gaussians in the mixture representing the observation is reduced as described in Section 3.3.

Fig. 12 shows the results obtained when applying our algorithm to images containing 4 occluded ellipses. In these cases all the instances of the ellipses are detected (second row). The third row show results of our algorithm when using only 2D Gaussian Mixtures (i.e. the information about orientation is not included in the GMMs). As can be seen, results with 3D GMMs are better than with 2D GMMs. Our modelling is flexible such that dimensions can be added to incorporate features that describe and discriminate the shape of interest better, and Fig. 12 shows how this can help improve the performance of our algorithms.

However, when increasing the number of ellipses in the image the results deteriorate. An example is shown in Fig. 13 where the white and black image corresponds to the input image (first and third row). The blue dots correspond to the observations taken from the input image. The red ellipses are the ellipses correctly detected while the green ellipses are detections that do not properly represent the observation. When more than 20 ellipses are present, the proportion of correctly detected ellipses is very small.

We analyse the metrics of recall and precision when evaluating all images in the data set. The results are reported in Fig. 14 (performance of state-of-the-art algorithms are also reported for comparison [6–10]). Fig. 14 suggests that our method outperforms methods based on the Hough Transform [9,10]. Methods based on connected edges, on the other hand, show better performance than our algorithm. The algorithm proposed by Chia et al. [6] shows outstanding performance and it is able to detect a high number of ellipses. However, it only works when the observations are a set of connected edges. In contrast, our algorithm does not

have any limitation and it can still be used when the observation is a sparse data set of points.

Additionally in Fig. 15 we show the performance of the detector when relaxing the overlap error in between the detected ellipse and the true ellipse. This experiment was performed using the set of images containing 4 occluded ellipses. Additional advantages of our proposed method are its feasibility including any extra information that can be added as an additional dimension of the GMM. The process of detecting multiples instances of an ellipse is shown in Fig. 16. At each step the score of each observation of being part of the estimated ellipse is computed and it is used for updating the set of observations. Fig. 17 is showing these scores at each iteration of our Algorithm 2, and as can be noticed there is a clear separation between observations belonging to the candidate ellipse (high scores with max value of 1) and the other observations that do not (scores close to 0). The threshold t_1 on these scores is chosen $t_1=0.3$ in an ad hoc fashion. Note also that because \mathcal{L}_2 is robust, it does not matter if a few observations are misclassified by our choice of threshold t_1 . More examples are shown in Fig. 18. The first two examples use the iterative algorithm for detecting multiples ellipses. The last row on the other hand shows the result when applying our algorithm for detecting faces.

7. Limitations and discussion

We have shown that our method detect well an ellipse when it is represented with a large enough proportion of the observations. Experimental results illustrate the robustness of our method to outliers, pseudo-outliers and noise. However, the performance deteriorate as more instances of ellipses (creating more pseudo-

outliers hence reducing the proportion of inliers each time [Algorithm 1](#) is run as part of [Algorithm 2](#)) occur in the observation set, as shown in [Section 6.2](#). This is a common problem that most methods fail to overcome: all methods we have compared with are decreasing their precision performances when increasing the number of ellipse instances in the observations.

On occasions, due to a particular repartition of outliers, the global solution of our cost function may not correspond to the expected ellipse but only a random ellipse formed by the observations. An example is shown in [Fig. 19](#) where the global solution does not represent the ellipse we expect to detect: our cost function for the false positive (green) is smaller (minimum of the \mathcal{L}_2 distance) than for the true ellipse (red). [Fig. 20](#) shows the corresponding 2D GMM for the observations and the ellipse model for two different bandwidth. The large difference in height between both GMMs make the Euclidian \mathcal{L}_2 distance less suited in this case. We have shown however that in general our framework with the \mathcal{L}_2 distance between GMMs compares very well with current alternative methods for ellipse detection and that its performance can be improved by adding extra dimensions carrying more information (such as edge orientation that is used in this paper) as shown in [Fig. 12](#).

8. Conclusion

In this paper we have shown an algorithm based on \mathcal{L}_2 distance between probability density functions for estimating the parameters of an ellipse. We evaluate the algorithm when detecting an ellipse on a set of 2D data points and when detecting multiples instances on an image. In the latter case we use multidimensional modelling for the GMM where the position of the contour pixels and its orientation are considered. We used the proposed algorithm to detect multiple ellipses in a benchmark data set. Results compared with state-of-the-art shows the promising performance of our algorithm. Our approach is only outperformed by recently proposed techniques that incorporate pixel connectivity in their descriptors to find ellipses. Future work will look at how this connectivity information could also be used in our representation of shapes with GMMs for robust matching with \mathcal{L}_2 .

We have proposed here a bottom-up greedy approach to multiple instance detection. This strategy was chosen due to the fact that we assume there is no information about the number of instances the shape appears in the image. However, when such information is available it could be included in the model by fixing the number of parameters to estimate.

Conflict of interest

None declared.

Acknowledgments

This work has been supported by scholarships from Trinity College Dublin in Ireland and the Government of Chile.

References

- [1] B. Jian, B. Vemuri, Robust point set registration using gaussian mixture models, *IEEE Trans. Pattern Anal. Mach. Intell.* 33 (8) (2011) 1633–1645.
- [2] C. Arellano, R. Dahyot, Mean shift algorithm for robust rigid registration between gaussian mixture models, in: 20th European Signal Processing Conference (Eusipco), Bucharest, Romania, 2012, pp. 1154–1158.
- [3] C. Arellano, R. Dahyot, Shape model fitting without point correspondence, in: 20th European Signal Processing Conference (Eusipco), Bucharest, Romania, 2012, pp. 934–938.
- [4] C. Arellano, R. Dahyot, Robust Bayesian fitting of 3d morphable model, in: Conference on Visual Media Production (CVM), London, UK, 2013.
- [5] C. Arellano, R. Dahyot, Shape model fitting using non-isotropic gmm, in: 23rd IET Irish Signals and Systems Conference, 2012.
- [6] A. Chia, S. Rahardja, D. Rajan, K. Leung, A split and merge based ellipse detector with self-correcting capability, *IEEE Trans. Image Process.* 20 (7) (2011) 1991–2006, <http://dx.doi.org/10.1109/TIP.2010.2099127>.
- [7] F. Mai, Y. Hung, H. Zhong, W. Sze, A hierarchical approach for fast and robust ellipse extraction, *Pattern Recognit.* 41 (8) (2008) 2512–2524, <http://dx.doi.org/10.1016/j.patcog.2008.01.027>, URL (<http://www.sciencedirect.com/science/article/pii/S0031320308000642>).
- [8] E. Kim, M. Haseyama, H. Kitajima, Fast and robust ellipse extraction from complicated images, in: Proceedings of the IEEE International Conference on Information Technology and Applications, 2002.
- [9] H. Yuen, J. Illingworth, J. Kittler, Detecting partially occluded ellipses using the Hough transform, *Image Vis. Comput.* 7 (1) (1989) 31–37, [http://dx.doi.org/10.1016/0262-8856\(89\)90017-6](http://dx.doi.org/10.1016/0262-8856(89)90017-6), URL (<http://www.sciencedirect.com/science/article/pii/0262885689900176>).
- [10] R.A. McLaughlin, Randomized Hough transform: improved ellipse detection with comparison, *Pattern Recognit. Lett.* 19 (34) (1998) 299–305, [http://dx.doi.org/10.1016/S0167-8655\(98\)00010-5](http://dx.doi.org/10.1016/S0167-8655(98)00010-5), URL (<http://www.sciencedirect.com/science/article/pii/S0167865598000105>).
- [11] X. Bai, C. Sun, F. Zhou, Splitting touching cells based on concave points and ellipse fitting, *Pattern Recognit.* 42 (11) (2009) 2434–2446, <http://dx.doi.org/10.1016/j.patcog.2009.04.003>.
- [12] Y. Soh, J. Bae, D. Kim, H. Kim, A new method for ellipse fitting in satellite images, in: Second International Conference on Intelligent Computation Technology and Automation, ICICTA, vol. 1, 2009, pp. 502–506, <http://dx.doi.org/10.1109/ICICTA.2009.128>.
- [13] S. Mahdavi, W. Morris, I. Spadinger, N. Chng, O. Goksel, S. Salcudean, 3d prostate segmentation in ultrasound images based on tapered and deformed ellipsoids, in: Medical Image Computing and Computer-Assisted Intervention, MICCAI, Lecture Notes in Computer Science, vol. 5762, 2009, pp. 960–967.
- [14] L. Ding, A.M. Martinez, Features versus context: an approach for precise and detailed detection and delineation of faces and facial features, *IEEE Trans. Pattern Anal. Mach. Intell.* 32 (11) (2010) 2022–2038, <http://dx.doi.org/10.1109/TPAMI.2010.28>.
- [15] D. Maio, D. Maltoni, Real-time face location on gray-scale static images, *Pattern Recognit.* 33 (9) (2000) 1525–1540.
- [16] C. Wong, S.C.F. Lin, T.R. Ren, N.M. Kwok, A survey on ellipse detection methods, in: IEEE International Symposium on Industrial Electronics, ISIE, 2012, pp. 1105–1110, <http://dx.doi.org/10.1109/ISIE.2012.6237243>.
- [17] D. Chaudhuri, A simple least squares method for fitting of ellipses and circles depends on border points of a two-tone image and their 3-d extensions, *Pattern Recognit. Lett.* 31 (9) (2010) 818–829, <http://dx.doi.org/10.1016/j.patrec.2010.01.009>.
- [18] W. Gander, G. Golub, R. Strebler, Least-squares fitting of circles and ellipses, *BIT Numer. Math.* 34 (4) (1994) 558–578, <http://dx.doi.org/10.1007/BF01934268>.
- [19] P.L. Rosin, A note on the least squares fitting of ellipses, *Pattern Recognit. Lett.* 14 (10) (1993) 799–808, [http://dx.doi.org/10.1016/0167-8655\(93\)90062-1](http://dx.doi.org/10.1016/0167-8655(93)90062-1).
- [20] J. Cabrera, P. Meer, Unbiased estimation of ellipses by bootstrapping, *IEEE Trans. Pattern Anal. Mach. Intell.* 18 (7) (1996) 752–756, <http://dx.doi.org/10.1109/34.506797>.
- [21] W. Chojnacki, M.J. Brooks, A.V.D. Hengel, D. Gawley, Fitting surfaces to data with covariances, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (2000) 1294–1303.
- [22] A. Fitzgibbon, M. Pilu, R.B. Fisher, Direct least square fitting of ellipses, *IEEE Trans. Pattern Anal. Mach. Intell.* 21 (5) (1999) 476–480, <http://dx.doi.org/10.1109/34.765658>.
- [23] K. Kanatani, Statistical bias of conic fitting and renormalization, *IEEE Trans. Pattern Anal. Mach. Intell.* 16 (3) (1994) 320–326, <http://dx.doi.org/10.1109/34.276132>.
- [24] Z. Szapak, W. Chojnacki, A. Hengel, Guaranteed ellipse fitting with the sampson distance, in: Computer Vision, ECCV, Lecture Notes in Computer Science, vol. 7576, 2012, pp. 87–100.
- [25] K. Kanatani, P. Rangarajan, Hyper least squares fitting of circles and ellipses, *Comput. Stat. Data Anal.* 55 (6) (2011) 2197–2208, <http://dx.doi.org/10.1016/j.csda.2010.12.012>.
- [26] R.O. Duda, P.E. Hart, Use of the Hough transformation to detect lines and curves in pictures, *Commun. ACM* 15 (1) (1972) 11–15, <http://dx.doi.org/10.1145/361237.361242>.
- [27] P. Hart, How the Hough transform was invented, *IEEE Signal Process. Mag.* 26 (6) (2009) 18–22, <http://dx.doi.org/10.1109/MSP.2009.934181>.
- [28] D. Ballard, Generalizing the Hough transform to detect arbitrary shapes, *Pattern Recognit.* 13 (2) (1981) 111–122, [http://dx.doi.org/10.1016/0031-3203\(81\)90009-1](http://dx.doi.org/10.1016/0031-3203(81)90009-1), URL (<http://www.sciencedirect.com/science/article/pii/0031320381900091>).
- [29] P.S. Nair, A.T. Saunders, Jr., Hough transform based ellipse detection algorithm, *Pattern Recognit. Lett.* 17 (7) (1996) 777–784, [http://dx.doi.org/10.1016/0167-8655\(96\)00014-1](http://dx.doi.org/10.1016/0167-8655(96)00014-1), URL (<http://www.sciencedirect.com/science/article/pii/0167865596000141>).
- [30] C.F. Chien, Y.C. Cheng, T.T. Lin, Robust ellipse detection based on hierarchical image pyramid and Hough transform, *J. Opt. Soc. Am.* 28 (4) (2011) 581–589, <http://dx.doi.org/10.1364/JOSAA.28.000581>.

- [31] N. Guil, E. Zapata, Lower order circle and ellipse Hough transform, *J. Pattern Recognit.* 30 (1997) 1729–1744.
- [32] S.C. Zhang, Z.Q. Liu, A robust, real-time ellipse detector, *Pattern Recognit.* 38 (2) (2005) 273–287, <http://dx.doi.org/10.1016/j.patcog.2004.03.014> <http://www.sciencedirect.com/science/article/pii/S0031320304001372>.
- [33] N. Kiryati, Y. Eldar, A. Bruckstein, A probabilistic Hough transform, *Pattern Recognit.* 24 (4) (1991) 303–316, [http://dx.doi.org/10.1016/0031-3203\(91\)90073-E](http://dx.doi.org/10.1016/0031-3203(91)90073-E), URL <http://www.sciencedirect.com/science/article/pii/003132039190073E>.
- [34] L. Xu, E. Oja, P. Kultanen, A new curve detection method: randomized Hough transform (rht), *Pattern Recognit. Lett.* 11 (5) (1990) 331–338, [http://dx.doi.org/10.1016/0167-8655\(90\)90042-Z](http://dx.doi.org/10.1016/0167-8655(90)90042-Z), URL <http://www.sciencedirect.com/science/article/pii/016786559090042Z>.
- [35] T.C. Chen, K.L. Chung, An efficient randomized algorithm for detecting circles, *Comput. Vis. Image Underst.* 83 (2) (2001) 172–191, <http://dx.doi.org/10.1006/cviu.2001.0923>, URL <http://www.sciencedirect.com/science/article/pii/S1077314201909233>.
- [36] W. Lu, J. Tan, Detection of incomplete ellipse in images with strong noise by iterative randomized Hough transform (irht), *Pattern Recognit.* 41 (4) (2008) 1268–1279, <http://dx.doi.org/10.1016/j.patcog.2007.09.006>, URL <http://www.sciencedirect.com/science/article/pii/S0031320307004128>.
- [37] D.K. Prasad, M.K. Leung, S.-Y. Cho, Edge curvature and convexity based ellipse detection method, *Pattern Recognit.* 45 (9) (2012) 3204–3221, <http://dx.doi.org/10.1016/j.patcog.2012.02.014>, URL <http://www.sciencedirect.com/science/article/pii/S0031320312000763>.
- [38] L. Ding, A. Goshtasby, On the canny edge detector, *Pattern Recognit.* 34 (2001) 721–725.
- [39] R. Dahyot, Statistical Hough transform, *IEEE Trans. Pattern Anal. Mach. Intell.* 31 (8) (2009) 1502–1509, <http://dx.doi.org/10.1109/TPAMI.2008.288>.
- [40] R. Halir, J. Flusser, Numerically stable direct least squares fitting of ellipses, in: *Proceedings of 6th International Conference in Central Europe on Computer Graphics and Visualization, WSCG*, vol. 98, Citeseer, 1998, pp. 125–132.
- [41] J. Yu, S. Kulkarni, H. Poor, Robust ellipse and spheroid fitting, *Pattern Recognit. Lett.* 33 (5) (2012) 492–499, <http://dx.doi.org/10.1016/j.patrec.2011.11.025>.

Claudia Arellano received a Civil Electrical and Industrial degree from Pontificia Universidad Catolica de Chile (2003). She also hold a Master of sciences from the same university (2003). In 2013 she was awarded with a Ph.D. from the Trinity College Dublin in Ireland. Her research interest include patterns recognition, shape detection and robust statistics.

Rozenn Dahyot received Master degrees in Physics (1998, Telecom Physique Strasbourg) and Image Processing (1998, University Louis Pasteur Strasbourg) in France. She gained her Ph.D. in 2001 from the University Louis Pasteur Strasbourg. From 2002 to 2005, she was a research fellow in Trinity College Dublin, Ireland, and University of Cambridge, UK. Since 2005, she is an Assistant Professor in the School of Computer Science and Statistics, in Trinity College Dublin. Her research interest includes robust statistics, pattern detection and recognition, image and video analysis.