Pattern Recognition 48 (2015) 2198–2208

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Geometrically-constrained balloon fitting for multiple connected ellipses

Michael Kemp [a,*], Richard Yi Da Xu [b]

[a] School of Computing and Mathematics, Charles Sturt University, PO Box 883, Orange NSW 2800, Australia
[b] Faculty of Information and Technology, University of Technology Sydney, 2007 Australia

ARTICLE INFO

Article history:
Received 18 February 2014
Received in revised form
24 December 2014
Accepted 27 January 2015
Available online 7 February 2015

Keywords:
Ellipse fitting
Constrained Gaussian mixture model
Expectation maximisation
Geometric constrained fitting

ABSTRACT

This paper presents a framework to fit data to a model consisting of multiple connected ellipses. For each iteration of the fitting algorithm, the representation of the multiple ellipses is mapped to a Gaussian mixture model (GMM) and the connections are mapped to geometric constraints for the GMM. The fitting is a modified constrained expectation maximisation (EM) method on the GMM (maximising with respect to the ellipse parameters rather than Gaussian parameters). A key modification is that the precision of the chosen GMM is increased at each iteration. This is similar to slowly inflating a bunch of connected balloons and so this is called balloon fitting. Extensions of the framework to other constraints and possible pre-processing are also discussed. The superiority of balloon fitting is demonstrated through experiments on several silhouettes with noisy edges which compare other existing methods with balloon fitting and some of the extensions.

1. Introduction

Ellipses are useful for modelling and assessing shape. Although non-linear they have simple equations and so are easy to work with. By coupling several ellipses together, complicated structures with joints can be modelled, for example the human upper body [1], the human hand [2] and fish [3].

In this paper we propose a framework for fitting a silhouette to a predetermined model of multiple connected ellipses. The framework utilises a natural equivalence of multiple ellipses with a Gaussian mixture model (GMM) to determine the fit using a variant of the expectation maximisation (EM) algorithm. The entire silhouette is used to avoid noise along the edges. A possible problem of the iterative methods used in the maximisation is that the algorithm could get stuck in a local maximum. To reduce this problem, we proposed a novel balloon fitting algorithm and apply it iteratively during the E-M procedure. To explain in words, the ellipses are considered as balloons that begin mostly deflated and then are inflated each iteration of the EM algorithm until they become full size (in terms of statistics, the variance of the GMM is increased at the beginning and then slowly reduced each step).

There already exist many ellipse fitting algorithms. Wong et al.'s survey article [4] shows the variety of methods and applications of ellipse fitting. There are two main approaches, least square methods and Hough transforms. The classic work involving least square methods is direct least square fitting in [5], however improvements are still being made [6,7]. Hough transforms use bins to determine the key parameters of ellipses. Although the five parameters could make the searching of all the bins costly, practical algorithms can be achieved using the parallel processing power of GPUs [8] or randomised Hough Transforms (for example [9]). Other tricks are to use Hough transforms to instead detect tangents and then reconstruct the ellipse from its tangents [10]. Both main approaches can readily be used in a scene with multiple ellipses (for example [6,8,9]).

However both of the main approaches to ellipse fitting are susceptible to noise when fitting to silhouettes. All of the cited ellipse fitting methods fit an ellipse curve to some data points. A common scenario (as used in [6,8,9,11,12]) is to fit multiple ellipses to a silhouette. A pre-processing step is used to determine the edge points of this silhouette and then apply the ellipse fitting to these edges. Even if the majority of the silhouette is reasonable, the edges can be difficult to obtain or can be badly affected by noise. In this paper, we overcome this problem by using the whole silhouette in the algorithm.

Ellipse fitting can also be applied where the multiple ellipses are coupled together, for example [1,2,11–13]. These citations all use edge (or boundary) values and so also are susceptible to noisy edges. The coupling of ellipses is achieved through geometric constraints. Least

* Corresponding author. Tel.: +61 2 6365 7840.
E-mail addresses: mkemp@csu.edu.au (M. Kemp), yida.xu@uts.edu.au (R.Y.D. Xu).

http://dx.doi.org/10.1016/j.patcog.2015.01.026
0031-3203/

squares methods can include these constraints, but force the least squares to be solved iteratively rather than directly [12]. The downside to numeric iterative methods is that they are sensitive to the initial guess and can get stuck in local optima when performing optimisation. For example, in [2], Jeune et al. fit multiple connected 3D ellipsoids to the edge of silhouette of a hand via an iterative Levenberg–Marquadt (LM) algorithm. The paper acknowledges that "sometimes the algorithm does not find the right solution". A common cause is because the LM algorithm gets "stuck in a local minimum". Weak initial estimates can promote this issue. The paper reduces the issue by using the context of hands to detect the false solution and hence fix it, however their solution is context dependent. In our previous work [11], we reduced the issue by initially segmenting the data and then hierarchically fitting the ellipses according to the joint structure. In this paper, we improved upon our previous work by using balloon fitting to further reduce the issue.

Ellipse fitting can also be done using statistical techniques. Wong et al.'s survey article, [4] classifies these as part of a third category – titled "Other" for short. One example of this framework is [14]. Points of the curve of the ellipse are modelled by Gaussians. Then maximum likelihood estimation (MLE) is used to fit the curve to the noisy data. The maximisation essentially reduces to a least squares method – minimising the reprojection error. A downside to this particular article [14] is that the method applies to all conics and so constraints are required to guarantee an ellipse.

A useful way to get a ellipse specific statistical method is to use the equivalence between ellipses and Gaussians. The articles [1,15] use this equivalence in 3D to model 3D ellipsoids using 3D Gaussians. In our previous work [11], we used this equivalence to cluster the data and improve the initial guess before using an iterative least squares method to refine the process. More precisely, multiple ellipses are equivalent to a GMM. We temporarily ignored the connections of the ellipses then used unconstrained EM to fit the GMM to the silhouette. We then connected the ellipses and did the constrained optimisation in the final step. One improvement given in this paper is to combine these steps by constraining the GMM and doing the constrained optimisation in the M-step of the EM algorithm rather than an entirely separate stage. Phrasing the entire algorithm in terms of a constrained GMM allows the variability to be modified during the optimisation. In a similar way to simulated annealing this can help the algorithm not get stuck in local optimums (which are not global optimums).

The EM algorithm has been used to fit a constrained GMM using maximum likelihood estimation (MLE) in a variety of ways. One simple technique to handle constraints is to run the M-step as normal and then project onto the space of feasible models if the solution does not satisfy a constraint (this approach does not guarantee the full maximum, but it is easy to implement). This is seen for *boundary constraints* where one wishes to constrain the parameters within certain intervals. Alternatively, some constraints merely require a change in the update equations for the M-step. For example with *linear subspace constraints* parameters are restricted to be within a linear subspace. This has been used for the covariance [16] and precision matrices [17]. Another example this applies for is *positive equivalence constraints* [18], where points are labelled as coming from the same class. However, other constraints require an numerical iterative M-step rather than a direct update of the unconstrained GMM. This includes *affine subspace constraints* [19] and *negative equivalence constraints* [18].

Another approach of incorporating constraints with GMM is to use Bayesian priors. Then the EM algorithm is used to find the maximum a posterior (MAP). This is equivalent to adding a penalty term in the M-step. One example of this avoids a singular covariance matrix by adding a prior to make the determinant of the covariance matrix be 0 with a low or zero probability. Another example is with *spatial constraints* [20]. Spatially adjacent points are deemed more probable to be from same mixture. This is achieved by adding a Markov Random Field (MRF) prior. Instead of having a hard constraint, like $\mathbf{c}=\mathbf{0}$, priors allow softer constraints where we only insist that the constraint is close to 0 with high probability. An example of this is the kinematic constraints in [15]. These methods often make the M-step not closed and so numerical optimisation is required.

Similar work to this paper are the kinematically constrained GMMs given in [1,15]. Kinematics is how joints move in a body. For the 2D case, kinematic constraints refer to the position of joints. In [1,15] they have a collection of connected ellipsoids. The size of the ellipsoids is determined manually at the start of the process, whereas the ellipsoids are allowed to be translated and rotated. They have kinematic constraints that join the ellipsoids at fixed locations and then allow movement at the joints (there are several options in a 3D model). The constraints in [15] are achieved by adding a prior so that each constraint equation is not precisely zero, but rather a Gaussian with zero mean. The variance for the constraints is fixed and not modified during the algorithm. Alternatively in [1] the M-step is modified to be a constrained optimisation. For both, the means and covariances (actually just the rotational part of the covariance matrices) of the components are learned using a modified EM algorithm. Although some parts have a closed form, a gradient ascent algorithm is required to handle all the constraints. This paper improves upon the ideas in these articles by getting the EM algorithm to learn the size of the ellipses as well and location and orientation. As with any iterative optimisation method, these previous papers can have issues with getting stuck in local optimums. This paper introduces the balloon fitting modification to reduce this issue. This modification is similar in purpose to simulated annealing where early stages of the algorithm have higher variability so that local optimums can be escaped.

The rest of this article is ordered as follows. In Section 2.1 we introduce our mapping between multiple ellipses and a GMM. Our general framework works regardless of the precise constraints used. However, to make the article concrete we present a reliable way to specify geometric constraints on the ellipses in Section 2.2. We then explain how expectation maximisation (Section 2.3) with the balloon fitting modification (Section 2.4) is used to fit the ellipse structure to an image silhouette. We then complete the method section by looking at some variations to the framework that could be used depending on the context in Sections 2.5 and 2.6. In Section 3 we demonstrate the effectiveness of the framework by comparing balloon fitting against [11,1,15].

2. Our method

2.1. Mapping ellipses to Gaussians

To fit multiple connected ellipses, we need a mapping from an ellipse to a Gaussian distribution. This mapping allows us to convert the ellipse fitting into a maximum likelihood estimate (MLE) for a GMM. We are interested in connected ellipses. Thus we will use the following notation.

We assume that there are M ellipses (or Gaussian mixture components), indexed by l. Each ellipse is parameterised by the following five dimensional set of parameters: $\Theta_E^l=(\mathbf{x}_v^l,\mathbf{x}_c^l,b^l)$. The first parameter, $\mathbf{x}_v^l=(x_v^l,y_v^l)$, corresponds to the location of a vertex of the ellipse. The second parameter, $\mathbf{x}_c^l=(x_c^l,y_c^l)$, corresponds to the centre of the ellipse. The last parameter corresponds to the length of the other semi-axis. For example, if $\mathbf{x}_v^l$ is a vertex on the major axis, b is the length of the semi-minor axis. Of course, this parameterisation depends on the choice of vertex.

Many other parameterisations of ellipses are possible and have been used elsewhere. However, we have selected this one to best suit the constrained optimisation we will perform later. When performing optimisation numerically, all parameters should ideally have a similar