Grouped Outlier Removal for Robust Ellipse Fitting

Mang Shao
Imperial College London
ms2308@ic.ac.uk

Yoshihisa Ijiri and Kosuke Hattori
OMRON Corporation
joyport@ari.ncl.omron.co.jp
kosuke_hattori@omron.co.jp

Abstract

This paper presents a novel outlier removal method which is capable of fitting ellipse in real-time under high outlier rate, based on the phenomenon that outliers generated by ellipse edge point detector are likely to appear as groups due to real-world nuisances, such as under partial occlusion or illumination change. To confront the grouped outliers while maintaining the fitting efficiency, we introduce a proximity-based 'split and merge' approach to cluster the edge points into subsets, followed by a breath-first outlier removal process. The experiment shows that our algorithm achieves high performance under a wide range of inlier ratio and noise level with various types of realistic nuisances.

1 Introduction

Ellipse fitting is one of the fundamental problems in computer vision and robotic tasks. It is required as preprocessing modules in many higher level applications such as textureless object recognition and shape alignment. In this paper, we propose a real-time outlier removal solution for ellipse fitting which efficiently deals with outliers especially when they are contaminated consistently.

Amongst ellipse fitting algorithms presented in the last few decades, efforts have been made to approach the KCR lower bound [1], i.e. the theoretical accuracy bound for ellipse fitting without outliers. In these works, the presence of noises is considered from observational error only and thus the point-to-curve projection error is mostly assumed under Gaussian distribution. While in the real-world scenario, image or extracted edgess is likely to be contaminated in a very biased way such as under partial occlusion, specular highlight, deformation or shading. In such cases, robust fitting algorithm like random sample consensus (RANSAC) [4] is generally applied to eliminate outliers.

When the inlier rate ϵ is low, RANSAC soon becomes infeasible for many applications since the possibility of finding at least one correct ellipse model after K iterations $p = 1 - (1 - \epsilon^5)^K$ decreases drastically. Also, the model candidates generated from minimal sample sets can be greatly affected by the noise, which leads to suboptimal solutions.

A recent work [10] proposed a proximity-based outlier detection algorithm to effectively remove the isolated outliers and outlier clusters by constructing a proximity graph. However, the adjacency matrix is expensive to compute if data point set is large and some parameters need to be tuned carefully in order to achieve a good clustering result. Also, the proximity-based method fails if the outliers are connected smoothly with the inliers, as shown in Figure 1,

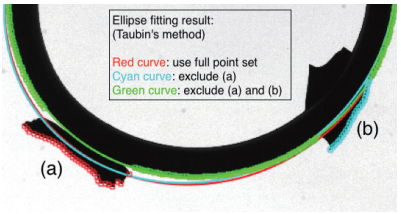

Figure 1. A challenging case for proximity-based outlier removal method. Type (a) outliers cannot be filtered by simple proximity (e.g. k-Nearest Neighbour) check. Type (b) is even more difficult since they are connected with the inlier contour.

type (b). Another work [8] presents an accurate, non-iterative method based on the geometric distance between a data point and an ellipse. However it does not take the type (b) outlier into account as well.

In this work, we demonstrate how edge points can be effectively grouped into short contours to reduce the computational cost, and further show how to eliminate the outlier contours in a breath-first manner. Our contribution is four-fold. First we introduce a split-and-merge trick to cluster data points into subsets which are likely to contain either only inliers or only outliers. Second we propose a breath-first strategy for searching the outlier contours through the combination of subsets. Third, we speed up the searching process by using the smallest generalised eigenvalue (which is a by-product of algebraic fitting algorithm) to approximate the point-to-curve projection, and then use the algebraic fitting solution as initial guess for geometric fitting algorithm to alleviate measurement noises. Finally, we proposed a synthetic dataset for evaluation.

2 Preliminaries and Related Works

Given a point set, the objective of ellipse fitting is to find a geometric parameter set that minimises the sum of inlier-points-to-ellipse-curve projection distance. In this section, we briefly introduce two broad algorithm categories: algebraic (Least-Square-based) and geometric (Maximum-Likelihood-based) method.

2.1 Algebraic Fitting

Any ellipse can be represented by a second order polynomial $F(\mathbf{u}, \mathbf{x}) = \mathbf{u} \cdot \mathbf{x} = ax^2 + bxy + cy^2 + dx + ey + f = 0$, subject to $b^2 < 2ac$. Our goal here is to estimate the parameter set $\mathbf{u} = [a, b, c, d, e, f]$ from a given point set $\mathbf{x}_i = [x_i, y_i]$ such that the sum of algebraic distance $\sum_{i=1}^{N} |F(\mathbf{u}, \mathbf{x}_i)|_2$ is minimised.

This problem is generally tackled with linear least square solvers as in several seminal approaches [9, 5, 6].

In this work, since we focus on outlier elimination, Taubin's method [9] is chosen as our algebraic ellipse fitting algorithm. From our empirical experiments, we believe that Taubin's method is still one of the most accurate and robust methods giving its efficiency, despite it has been proposed for decades. Since this method is designed for general conic fitting, it might return conics other than ellipse as well (e.g. hyperbola curve) and hence leads to incorrect fitting result.

2.2 Geometric Fitting

Geometric parameters of ellipse is of 4 elements: center position, length of major axis, minor axis and angle of tilt. Unlike algebraic methods, the linear condition has no longer held for solving the least square problem with geometric parameters. Instead, iterative optimisation methods are used to find the local optimal solution given an initial guess.

Geometric fitting, i.e. Maximum-Likelihood-based method (ML) is generally regarded as one of the most precise fitting algorithms. They do not suffer from scale indeterminacy as in the algebraic methods and are likely to achieve the local optimal solution. A detailed analysis is given in the work [6].

The main drawback of geometric fitting algorithms is their high computational cost due to iterative optimisation process. With existence of outliers or high measurement noise, they also tend to fail to converge. The initial guess is crucial as well to achieve the global optimal solution.

In this work, we apply ML as the final refinement to alleviate the measurement noise while keeping the pipeline fast by using a near-optimal algebraic solution as its initial guess.

2.3 Edge Point Detector

In this work, the edge points are collected circularly around a user-estimated center point by finding the maximal gradient change along each circle radius such that the order of the point sequence is naturally known. The center point does not necessarily need to be accurate. This setting is common in many industrial applications like shape alignment, as it provides several merits over the ordinary edge detector.

3 Proposed Approach

The proposed method consists of three stages:

- Clustering data points based on proximity so that each subset is likely to contain only inliers or outliers edge points.
- Searching through combinations of subsets to minimise algebraic fitting distance until convergence.
- Refining the algebraic solution by the geometric fitting algorithm.

3.1 Proximity-based Point Clustering

Generally a partial or full adjacent matrix needs to be calculated to determine the connections between edge points. This process roughly has computational complexity $O(N^2)$ respect to the number of data points and a adequate connection radius is also needed to be carefully chosen. The edge point detector we used in this work naturally provides the connectivity between the points, which greatly simplifies the clustering process.

Algorithm 1: Proximity-based Point Clustering

Input: $C_{full} = \{\mathbf{x}_i\}_{i=1}^{N}$
Initialisation: $t, \tau, D, C_1 = \{\}$;
for $i \leftarrow 1$ **to** N **do**
  $C_k = C_k \cup \mathbf{x}_i$;
  $d_i = |\mathbf{x}_i - \mathbf{x}_{i+1}|_2$; `/* i+1=1 if i=N */`
  if $d_i > t * \underset{i}{\text{median}}\{d_i\}$ **then**
    $k = k + 1$; $C_k = \{\}$;
  end
end
Delete any set C_k that has cardinality $|C_k| < \tau$
for $k \leftarrow 1$ **to** K **do**
  if $|C_k| > \frac{2N}{D}$ **then**
    Uniformly split C_k into $\lfloor \frac{|C_k|}{r} \rfloor$ sets;
    Replace C_k by these sets;
  end
  else if $|C_k| + |C_{k+1}| < \frac{N}{D}$ **then**
    $C_k = C_k \cup C_{k+1}$; `/* Delete set` C_{k+1} `*/`
  end
end
return $\{C_k\}_{k=1}^{K}$

The point set is ordered by the edge point detector such that $\mathbf{x}_n$ and $\mathbf{x}_{n+1}$ are next to each other. A point set is split at a point that Euclidean distance to its neighbour points d_i greater than $t * \underset{i}{\text{median}}\{d_i\}$, where t is the distance ratio threshold. If the measurement noise is too high, t should be set to a larger number to prevent over-segmentation. In this step, subset is discarded as isolated outliers if it contains less than τ points.

To deal with type (b) outliers in Figure 1, each subset is uniformly split if they contain more than a certain number of points. This number is determined by a pre-defined expected subset cardinality D. Larger D allows a finer segmentation of edge points, but induces higher risk of getting stuck in local optimum after the later searching stage, and also sacrifices the processing speed.

In the last step, neighbouring subsets are merged if their sum of cardinality is sufficiently small. In the end, the whole point set should be split into similar sizes and each subset is likely to contain only inliers or outliers. The pseudo code of this section is shown in the Algorithm 1.

3.2 Breath-First Searching

As shown in Figure 2, the fitting trials are performed between the combinations of subsets only. The total number of possible combinations $\{C\}_{n=1}^{D}$ is $\sum_{n=0}^{D} \frac{D!}{n!(D-n)!}$, which could be still large in number.

For each test, Taubin's fitting method [9] is used to sidestep the expensive geometric projection. We further speed up the process by approximating algebraic distance with the smallest generalised eigenvalue λ from solving the linear least square system.

The algebraic distance between an ellipse $\mathbf{u}$ and a

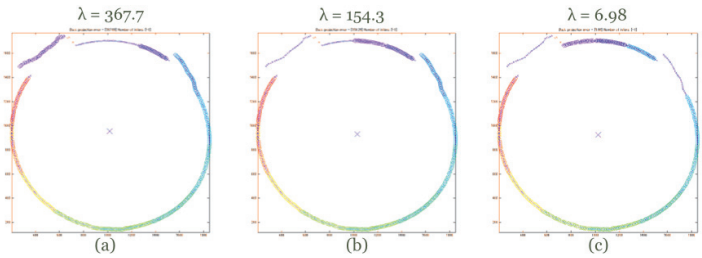

Figure 2. This figure shows an example that the value of the smallest generalised eigenvalue λ decreasing drastically after excluding the outlier subsets. The point set contains two 'outlier' subsets, (a), (b) and (c) show three cases during the iteration, while (a) does not exclude any outlier subset, (b) excludes one of them and (c) excludes all.

point set C is:

$$F(\mathbf{u}, C) = \sum_{\mathbf{x} \in C} |\mathbf{u} \cdot Z(\mathbf{x})|_2, Z(\mathbf{x}) = [x^2, xy, y^2, x, y, 1].$$

We define our energy function to approximate the algebraic distance by the minimal generalised eigenvalue solved as follows:

$$E(C) = min(\lambda), \text{ subject to } \mathbf{M}v = \lambda \mathbf{N}v$$

$$\text{where } \mathbf{M} = \frac{1}{N} \sum_{n=1}^{N} Z(\mathbf{x}^T - \bar{\mathbf{x}}^T) Z(\mathbf{x} - \bar{\mathbf{x}})$$

$$\mathbf{N} = \frac{4}{|C|} \sum_{n=1}^{|C|} \begin{pmatrix} x_n^2 & x_n y_n & 0 & f_0 x_n & 0 & 0 \\ x_n y_n & x_n^2 + y_n^2 & x_n y_n & f_0 y_n & f_0 x_n & 0 \\ 0 & x_n y_n & y_n^2 & 0 & f_0 y_n & 0 \\ f_0 x_n & f_0 y_n & 0 & f_0^2 & 0 & 0 \\ 0 & f_0 x_n & f_0 y_n & 0 & f_0^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Here we propose a searching scheme that efficiently search and exclude the outlier subsets in a breath-first manner. To minimise the energy while maximising the number of possible inliers, the combination of sets are tested by excluding one subset in each iteration. The searching process converges if excluding any subset does not reduce the energy up to certain rate σ. To prevent local optimal solution, number of S candidate sets with lowest energy are stored at each iteration, and each will be appended with another S candidates in the next iteration. In case there are much more outliers than average subset size so that removing any subset does not reduce the energy, we introduce an error threshold η as another stop condition. The pseudo code of this section is provided in Algorithm 2.

3.3 Refinement

Experimentally we found a refinement step to be crucial to reducing the measurement noise from the edge point detector. Given the inlier point set, we convert the algebraic parameters $\mathbf{u}$ to geometric parameters H and apply a nonlinear optimiser as below to further minimise the point-to-curve (orthogonal to the ellipse tangent) projection error. The detail for calculating projection distance can be found in [3].

$$H^* = \arg\min_H ||proj(H, C)||_2$$

Since the initial geometric parameters are already estimated from the previous stages, they are close enough to the optimal solution. Thus the efficiency is maintained despite the algorithm itself is expensive because it converges within few steps in most of cases.

Algorithm 2: Preemptive searching

Input: $\{C_k\}_{k=1}^K$
Initialisation: $S, \sigma, \eta, \{O_s = \emptyset\}_{s=1}^S$;
for $m \leftarrow 1$ **to** K **do**
  $C_{test} = \bigcup_{k \in K, k \neq m} \{C_k\}$;
  $[\lambda_{m,1}, \mathbf{u}_{m,1}] = TaubinFitting(C_{test})$;
end
$O_s = \{m\}, E_s = \lambda_m$ for s^{th} least λ_m;
for $l \leftarrow K - 2$ **to** 1 **do**
  for $s \leftarrow 1$ **to** $|\{O_s\}|$ **do**
    for $m \leftarrow 1$ **to** K, $m \notin O_s$ **do**
      $C_{test} = \bigcup_{k \in K, k \neq m} \{C_k\}$;
      $[\lambda_{m,s}, \mathbf{u}_{m,s}] = TaubinFitting(C_{test})$;
    end
  end
  if $(E_1 > \sigma * \min\{\lambda_{m,s}\}) \& (E_1 < \eta)$ **then**
    return $\boldsymbol{u}_{\underset{(m,s)}{\arg\min}\{\lambda_{m,s}\}}$
  end
  else
    $S \leftarrow l$ if $S > l$;
    for $s \leftarrow 1$ **to** $|\{O_s\}|$ **do**
      replace $\{O_s\}$ by $\{O_s \cup \{m\}\}$ for S least $\lambda_{m,s}$;
    end
    update E_s;
  end
end

4 Evaluation

Our method is evaluated by both synthetic[1] and realistic datasets which consists of ellipse edge points with different types of contamination. For synthetic data, we have rendered 3 videos (2 with occlusions and 1 with background clutter) with 3D computer graphic software at 960×540 resolution. Each video contains 100 frames and each frame consists of 100 extracted ellipse edge points. The ground truth ellipse centers are fixed to the image center. To simulate the measurement noise from camera, we augmented Gaussian noise ($\sigma = 3$ pixels) to the data points. We use the following parameter setting for our method in all experiments: $t = 2$, $\tau = 5$, $D = 12$, $S = D/2$, $\eta = 5$ and $\sigma = 1.05$. The preliminary investigation suggests this parameter set adapts a wide range of inlier rate and

[1] http://bit.ly/1Dvs0id

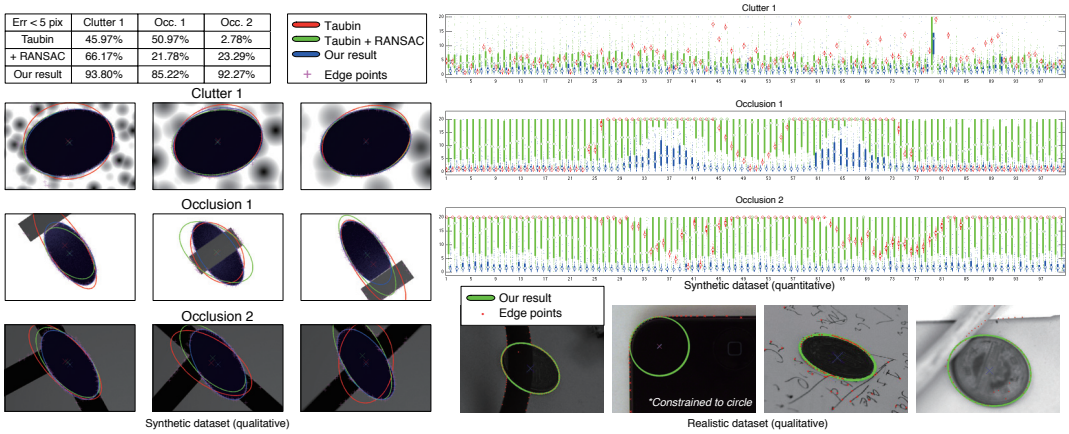

Figure 3. Evaluation result of 3 methods on synthetic and realistic datasets (best view in color). The boxplot reflects the statistic of center errors (capped at 20 pixel) from 100 runs of each method on every video frames. The table on the top left corner shows the success rate (center error less than 5 pixels) of the methods on each video. Note that for synthetic dataset, Gaussian noise is augmented to all data points in each run to simulate the noise from realistic optical sensors.

outlier type.

We compare our method with Taubin's method with/without RANSAC and summarise the evaluation result in Fig 3. In overall, our approach achieves superior accuracy than the baselines.

For Taubin's method itself, almost all fitting estimations are distorted by outliers due to the nature of least square solvers, such as in the dataset 'Occlusion 2' shown in Fig 3. However, RANSAC does not boost the accuracy significantly as expected as well. The main reason is twofold. Firstly the parameters, e.g. inlier ratio and distance threshold, are employed empirically as stop criteria. Such settings are not optimal to all situations therefore leading to an overall poor performance. Another is due to the high measurement noise within all data points. Since RANSAC picks hypotheses from random minimal samples, the impact from noises is drastically amplified comparing to estimating from all inlier samples, which also explains the motivation of our method. There are several extended RANSAC-like methods have been proposed to deal with such a problem, they either are expensive (e.g. pre-emptive RANSAC[7]) for ellipse fitting problem or suffering from poor initial estimations generated from minimal samples(e.g. locally optimised RANSAC[2]).

Our method achieves less than 20ms runtime on single core CPU with up to 360 data points, which meets the time requirement as a pre-processing module for many higher level real-time applications. With code optimisation and parallel processing, the whole pipeline can be further speeded up.

5 Conclusion

This paper has presented a new approach for efficient ellipses fitting under high outlier rate. We have demonstrated how data points can be clustered based on their proximity and how outliers can be filtered in a 'preemptive' manner. Our method has shown to be especially effective on the 'grouped' outliers due to environmental nuisances such as partial occlusion, specular highlight, deformation or shading. The runtime of the method is below 20ms for up to 360 data points, which is sufficient to achieve real-time processing speed for many higher level applications. Finally, we believe that our approach extends naturally to general conic and ellipsoid fitting.

References

[1] Nikolai Chernov and Claire Lesort. Statistical efficiency of curve fitting algorithms. *Computational statistics & data analysis*, 47(4):713–728, 2004.

[2] Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized ransac. In *Pattern Recognition*, pages 236–243. Springer, 2003.

[3] David Eberly. Distance from a point to an ellipse, an ellipsoid, or a hyperellipsoid. Technical report, Technical report, Geometric Tools, LLC, 2011. 5.1. 3, 1998.

[4] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

[5] Andrew Fitzgibbon, Maurizio Pilu, and Robert B Fisher. Direct least square fitting of ellipses. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 21(5):476–480, 1999.

[6] Kenichi Kanatani and Prasanna Rangarajan. Hyper-accurate ellipse fitting without iterations. In *VISAPP (2)*, pages 5–12, 2010.

[7] David Nistér. Preemptive ransac for live structure and motion estimation. *Machine Vision and Applications*, 16(5):321–329, 2005.

[8] Dilip K Prasad, Maylor KH Leung, and Chai Quek. Ellifit: An unconstrained, non-iterative, least squares based geometric ellipse fitting method. *Pattern Recognition*, 46(5):1449–1465, 2013.

[9] Gabriel Taubin. Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1115–1138, 1991.

[10] Jieqi Yu, Haipeng Zheng, Sanjeev R Kulkarni, and H Vincent Poor. Two-stage outlier elimination for robust curve and surface fitting. *EURASIP Journal on Advances in Signal Processing*, 2010:4, 2010.