The 16th Meeting on Image Recognition and Understanding

High Accuracy Ellipse-Specific Fitting

Tomonari Masuzaki[1,a)] Yasuyuki Sugaya[1,b)] Kenichi Kanatani[2,c)]

1. Introduction

Detecting circles and ellipses in images is the first step of many computer vision applications including industrial robotic operations and autonomous navigation. To this end, point sequences constituting elliptic arcs are detected by image processing operations, and then an ellipse equation is fitted to them. In the past, many methods have been proposed for ellipse fitting [1], [3], [4]. However, most of them fit a quadratic equation in x and y, or a *conic*, to a point sequence. Usually, this produces an ellipse if the sequence forms an elliptic arc, but a hyperbola or a parabola could result when the sequence is very short and the noise is very large.

It is Fitzgibbon et al. [2] who first proposed a method that only fits an ellipse. It is an algebraic method, and the computation is very easy, but the accuracy is low. Recently, Szpak et al. [5] introduced a high accuracy ellipse-specific method based on Sampson error minimization. In this paper, we incorporate a procedure for avoiding non-ellipses to the hyper-renormalization method of Kanatani et al. [3], which is regarded as the most accurate of all currently known methods and demonstrate by experiments that our technique outperforms the method of Szpak et al. [5].

2. Ellipse fitting

Quadratic equations in x and y in the form

$$Ax^2 + 2Bxy + Cy^2 + 2f_0(Dx + Ey) + f_0^2 F = 0, \quad (1)$$

represent curves called *conics*, which include ellipses, parabolas, hyperbolas, and their degeneracies such as two lines. In Eq. (1), f_0 is a constant that has the order of the image size for stabilizing finite length numerical computation ($f_0 = 600$ in our experiments). For a point sequence (x_α, y_α), $\alpha = 1, ..., N$, we define 6-D vectors

$$\boldsymbol{\xi}_\alpha = (x_\alpha^2,\ 2x_\alpha y_\alpha,\ y_\alpha^2,\ 2f_0 x_\alpha,\ 2f_0 y_\alpha,\ f_0^2)^\top,$$
$$\boldsymbol{\theta} = (A,\ B,\ C,\ D,\ E,\ F)^\top. \quad (2)$$

Then, the condition that (x_α, y_α) satisfies Eq. (1) is written as

$$(\boldsymbol{\xi}_\alpha, \boldsymbol{\theta}) = 0, \quad (3)$$

where $(\boldsymbol{a}, \boldsymbol{b})$ denotes the inner product of vectors $\boldsymbol{a}$ and $\boldsymbol{b}$. Since vector $\boldsymbol{\theta}$ has scale indeterminacy, we normalize it to unit norm: $\|\boldsymbol{\theta}\| = 1$.

3. Proposed method

It has been observed that the accuracy of hyper-renormalization is higher than Sampson error minimization [3]. Hence, it is reasonable to retain the solution of hyper-renormalization as long as it is an ellipse. If the hyper-renormalization iterations do not converge within a fixed limit (we set it to 100 times in our experiment), or if the resulting solution is not an ellipse, we switch to random sampling: we randomly choose from the point sequence five different points and compute the conic that passes through them. If a non-ellipse results, we discard the five points and choose new five points. If an ellipse results, we compute its Sampson error. We repeat this many times (1000 times in our expriment) and choose the solution for which the Sampson error is the smallest.

4. Experiment

We considered the three point sequences shown in **Fig. 1**(a). We added independent Gaussian noise of mean 0 and standard deviation σ to the x and y coordinates of each point and fitted an ellipse. Since the computed value $\boldsymbol{\theta}$ and its true value $\bar{\boldsymbol{\theta}}$ are both unit vectors, we measure the error $\Delta\boldsymbol{\theta}$ by the component of $\boldsymbol{\theta}$ orthogonal to $\bar{\boldsymbol{\theta}}$ and evaluate the RMS error

$$D = \sqrt{\frac{1}{10000}\sum_{a=1}^{10000} \|\Delta\boldsymbol{\theta}^{(a)}\|^2}, \quad (4)$$

over 10000 trials using different noise each time (the superscript (a) indicates the value for the ath trial).

Figure 1(b) shows the ratio of non-ellipse occurrences by hyper-renormalization, and Fig. 1(c) plots the RMS error D. The horizontal axis indicates the noise level σ divided by the average distance between neighboring points, which we call the *relative noise level*. The dotted lines indicate a theoretical accuray limit called the KCR lower bound. Interrupted plots indicate that convergence was not reached after a specified number of iterations.

As we can see, the accuracy of the method Fitzgibbon et al. [2] is generally very low, while hyper-renormalization[*1]

[1] Department of Computer Science and Engineering, Toyohashi University of Technology
[2] Okayama University
a) masuzaki@iim.cs.tut.ac.jp
b) sugaya@iim.cs.tut.ac.jp
c) kanatani2013@yahoo.co.jp

[*1] We used the code at: http://www.iim.cs.tut.ac.jp/~sugaya/

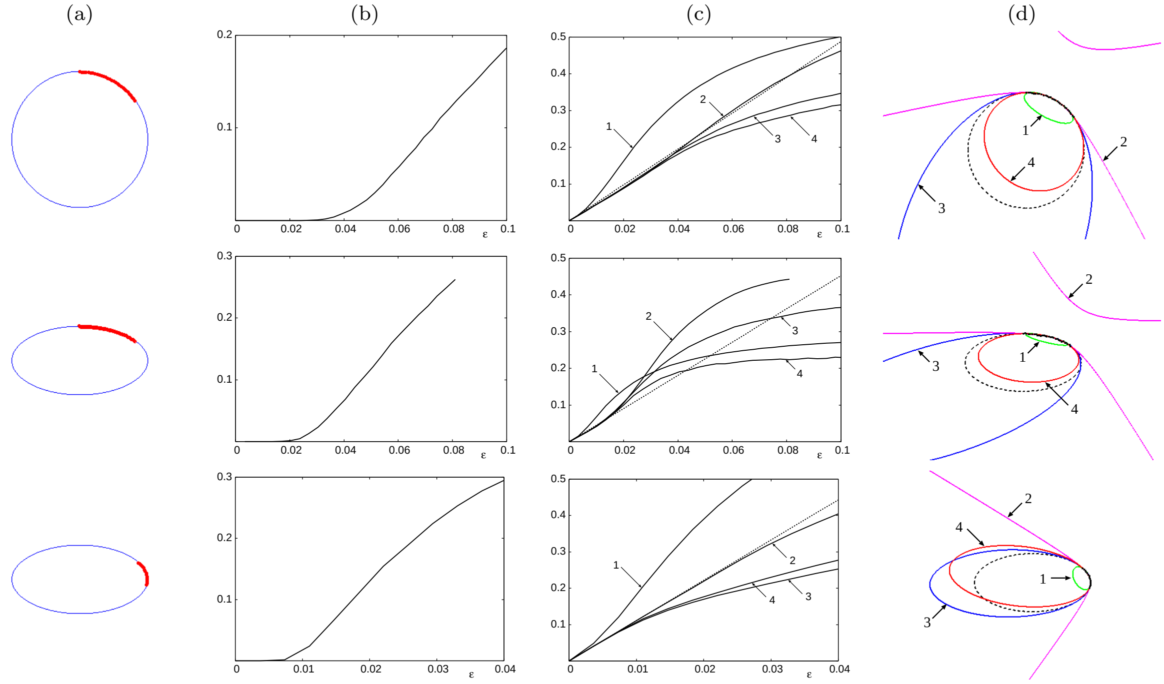

Fig. 1 (a) Point sequences for our experiment. The number of points is 30, 30, and 15 and the average distance between neighboring points is 2.96, 3.31, and 2.72, respectively. (b) The ratio of non-ellipse occurences by hyer-renormalization. The horizontal axis is for the relative noise level ϵ. (c) The corresponding RMS fitting error: 1. The method of Fitzgibbon et al. [2]. 2. hyper-renormalization. 3. The method of Szpak et al. [5]. 4. Proposed method. The dotted lines indicate the KCR lower bound. Interrupted plots indicate that convergence is not reached after a specified number of iterations above that noise level. (d) Fitting examples for a particular noise when hyper-renormalization returns a hyperbola.

and the method of Szpak et al.[*2][5] are very accurate; they almost achieve the KCR lower bound when the noise is very small. However, the method Fitzgibbon et al. [2] can outperform hyper-renormalization and the method of Szpak et al. [5] when the points are chosen from a low-curvature part as in Fig. 1(middle) and the noise is very large. Since the method of Szpak et al. [5] and the proposed method both restrict the solution to be an ellipse, their RMS error is smaller than hyper-renormalization in all cases, and in most cases our method is superior to that of Szpak et al. [5].

Figure 1(d) shows fitting examples for a particular noise when hyper-renormalization returns a hyperbola. The dotted lines indicate the true shapes. We can obverve a clear contrast: the method of Fitzgibbon et al. [2] fits a small and flat ellipse, while the method of Szpak et al. [5] fits a large ellipse close to the fitted hyperbola. Our method is in between, fitting an ellipse closer to the true shape.

5. Concluding remarks

We have proposed a new method that always fits an ellipse to a point sequence extracted from images. The currently known best method is hyper-renormalization of Kanatani et al. [3], but it may return a hyperbola when the noise in the data is very large. Our proposed method incorporates random sampling so that an ellipse always results. Doing simulation, we showed that our method has higher accuracy than the method of Fitzgibbon et al. [2] and the method of Szpak et al. [5], the two currently known ellipse-specific methods. We also observed that when hyper-renormalization returns a hyperbola, the method of Szpak et al. [5] tends to fit a large ellipse close to that hyperbola while the method of Fitzgibbon et al. [2] tends to fit a small and flat ellipse. Our method fits an ellipse somewhat in between.

Acknowledgements. The authors thank Wojciech Chojnacki of the University of Adelaide, Australia, and Ali Al-Shadraqah of the University of Mississippi, U.S.A. for helpful discussions on ellipse fitting. This work was supported in part by JSPS Grant-in-Aid for Young Scientists (B 23700202) and for Challenging Exploratory Research (24650086).

[*2] We used the code at: https://sites.google.com/site/szpakz/

References

[1] W. Chojnacki, M. J. Brooks, A. van den Hengel and D. Gawley, On the fitting of surfaces to data with covariances, *IEEE Trans. Patt. Anal. Mach. Intell.*, **22**-11 (2000-11), 1294–1303.

[2] A. Fitzgibbon, M. Pilu, and R. B. Fisher, Direct least squares fitting of ellipses, *IEEE Trans. Patt. Anal. Mach. Intell.*, **21**-5 (1999-5), 476–480.

[3] K. Kanatani, A. Al-Sharadqah, N. Chernov, and Y. Sugaya, Renormalization Returns: Hyper-renormalization and Its Applications, *Proc. 12th Euro. Conf. Comput. Vis.*, October 2012, Florence, Italy, Vol. 3, pp. 385–398.
38-1 (2010-9), 1–13.

[4] Y. Leedan and P. Meer, Heteroscedastic regression in computer vision: Problems with bilinear constraint, *Int. J. Comput. Vis.*. **37**-2 (2000-6), 127-150.

[5] Z. L. Szpak, W. Chojnacki, and A. van den Hengel, Guaranteed ellipse fitting with Sampson distance, *Proc. 12th Euro. Conf. Comput. Vis.*, October 2012, Florence, Italy, Vol. 5, pp. 87–100.