A Statistically Efficient Method for Ellipse Detection

Qiang Ji

Department of Computer Science
University of Nevada

Robert M. Haralick

Department of Electrical Engineering
University of Washington

Abstract

In this paper, we introduce a statistically efficient method for detecting ellipses in an image. Given a set of digital arc segments, we introduce geometric criteria to select possible pairs of arc segments belonging to the same ellipse. The selected arc pairs are subsequently validated or rejected based on certain statistical criteria via hypothesis testing. The advantages of our technique include: 1) the proposed criteria are scale-invariant; and 2) they can automatically adapt to the noise characteristics of each image and do not need to be adjusted empirically. Performance evaluation of our technique with real images demonstrates its good performance.

1 Introduction

Ellipses are powerful two-dimensional features for machine vision. They occur on an image through perspective projection of 3D circular or elliptical features. Their importance can be understood in two-folds. First, circular machine features represent one of the most common machine feature in industry. Detection of the circular features is therefore of great importance for industrial applications. Second, unlike points and lines, ellipses or elliptical arcs permit more constraints allowing determining the position and orientation of a part from a single ellipse or even from a partially occluded ellipse. Furthermore, 2D-3D ellipse/circle correspondence can be more easily established than points and lines. They therefore have been widely used in computer vision for matching, grouping, pose estimation, and 3D reconstruction.

Given detected edgel pixels, two classes of algorithms are currently available for segmenting the edge maps into meaningful components like lines and ellipses. The first class of algorithms involve variants of Hough transform, which work directly on the edge image. A second class of algorithms require an input of lists of chained edgel points produced from the original edge image after edge thinning and linking operation. While the first class of methods is simple to implement, they are usually computationally intense, especially for high-order features like ellipse detection. Recent algorithms [6, 4, 2, 1] employ decompositions and certain geometric properties of ellipses to reduce the number of parameters to estimate, therefore improving computational complexity. The second class of algorithms is more efficient. However it requires lists of chained edgel points from an edge image. Since the output of our edge detection are lists of connected edgels, we therefore restrict our analysis to lists of connected edgels (digital arcs) produced by edge detection.

Numerous algorithms have been developed for extracting linear and circular (or elliptical) arc segments from digital arcs. Most existing line-arc detection algorithms, however, can only detect lines and circular arcs. Additional clustering procedure is required to group circular arcs into ellipses. The algorithm recently proposed by Rosin [5] allows detection of lines, elliptical arcs and ellipses directly. Rosin's technique segments a digital arc into combinations of straight line and elliptical arc segments. The technique is simple and requires no thresholds. Given a list of connected edgels, the technique first recursively produces line approximation of the digital arc based on Lowe's significance rating measure [3]. Then in a similar manner, line segments are segmented into elliptical arc segments by fitting an ellipse to the end points of the line segments. A set of line segments are replaced with an elliptical arc segment if the replacement yields an improved significance rating. To overcome the faulty segmentation with the recursive splitting procedure, an additional stage was proposed to combine adjacent arc segments to yield a better arc approximation. This leads to an appreciable reduction in the number of line and elliptical segments. For this research, we start with the output of Rosin's technique. We introduce statistical clustering criteria to cluster arc segments belonging to the same ellipse.

2 Arc Clustering and Ellipse Detection

In spite of the efforts by Rosin to improve approximation by merging adjacent arc segments, it is noted that the ellipses detected by Rosin algorithm tends to be short elliptical arcs. For example, a hole boundary may contain multiple disconnected elliptical arc and/or line segments. These are broken up either due to edge detection or due to arc segmentation. Rosin's technique failed to merge them partially due to the inherent problem with the significance rating measure and partially due to the fact that only adjacent arc segments are examined during merging. to group together digital arc segments (either elliptical arc segments or line arc segments) belonging to the same ellipse. Arc clustering is performed in two steps. The first step, referred to as within-list-clustering, performs arc clustering by merging adjacent arc segments (replacing adjacent short arc segments with a longer elliptical arc segment). The merged arc segments normally belong to the same list of connected edgels. They are disconnected because of the arc segmentation procedure. The second step, referred to as between-list-clustering, involves merging arc segments located on different lists of connected pixels. They are separated because of the edge detection and linking procedure. Arc clustering amounts to fit an ellipse to the merged data. The goodness of the fitting is calculated and based on which a decision is made as to whether the merging should occur. For within-list-clustering, the clustering procedure iterates over the whole curve until no further improvements can be made; all adjacent combinations are tested. For between-list-clustering, pairwise arc merging is performed. A pair of arc segments are merged if certain criteria are met.

2.1 Candidates selection for merging

While performing arc clustering, decision must be made as to which two arc segments should be selected as candidates for possible merging. To be considered for merging, they must be close enough. This is measured by the proximity condition. Of the four possible Euclidean distances formed by the four end points of two arc segments, the distance between two arc segments is defined to be the smallest one. The distance of two arc segments must be less than a given threshold to satisfy the proximity condition. Proximity condition is always satisfied for within-list-clustering. For between-list- clustering, a threshold needs to be selected to define the proximity condition.

Besides proximity condition, an additional criterion is needed to ensure the combined elliptical arc segment geometrically meaningful. The additional selection criterion exploits the directions of arc segments (clockwise and counter-clockwise) as well as the configuration of their beginning and ending points. Specifically, two elliptical arc segments are selected as candidates for merging if they satisfy the proximity condition and their distance is derived between

- the start of one arc segment and the end of the other if the arc segments have the same directions.
- the end (start) of one arc segment and the end (start) of the other if the two arc segments have the opposite directions.

The two conditions listed above yield four possible configurations for the two elliptical arc segments as graphically illustrated in figure 1. Two elliptical arc segments selected for merging must belong to one of the four configurations. Arc segments belonging to other configurations should be excluded for consideration. For example, figure 2 gives an example of another possible configuration between two arc segments having the same direction. In this configuration, the distance is computed from between the start of the first arc segment and the start of the second arc segment. Though the distance meets the proximity condition, they can not be considered for merging since their merging can not yield a correct elliptical arc.

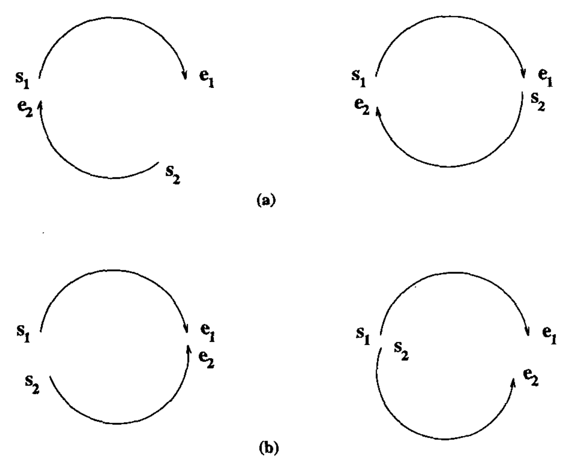

Figure 1: Possible configurations between two candidate arc segments: (a) two arc segments with the same direction; (b) two arc segments with different directions.

2.2 Merging criteria

Given two candidates of arc segments, we must establish quantitative criteria to determine whether they

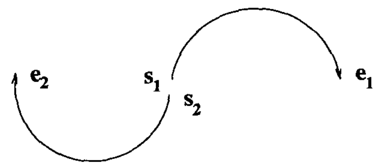

Figure 2: An example of non-fusible configuration between two arcs.

should be merged. The selected criteria must be general enough to be applied to different feature representations (line and elliptical arcs). They also need to be scale-invariant. In Rosin's algorithm, they use the significance measure proposed by Lowe [3]. The significance measure is calculated as the ratio of the length of the feature primitive divided by the maximum deviation of the curve from the feature primitive that describes the complete curve. A feature is combined with its neighboring features if the combination improves significance measure. While this measure satisfies the two requirements, it tends to over-segment a curve, yielding many false break points.

The criterion we propose to use is based on statistical testing. The statistical testing can be performed in either parameter space or in residual error space. Parameter space testing involves testing the statistical difference between the estimated parameters of the two arc segments. Two arc segments are merged if the difference between their estimated parameters is not statistically significant. This works well for linear features like lines. But it may not work for higher order feature like ellipses due to the ill-conditioned nature of the problem. Two elliptical arcs, when fitted separately, may not have the same or close parameter representation even though they originally belong to the same ellipse. This is especially true when the digital arcs to be fitted are short. An alternative way is to perform the least-squares ellipse-fitting to the combined data of the two arc segments. For a reasonably good fit, the ellipse parameters should be normally distributed with a mean parameter vector and covariance matrix. While the covariance matrix may be obtained analytically via covariance propagation, the mean parameter vector is unknown. With the unavailability of population mean parameter vector and the sample covariance matrix, we can not develop statistical testing criterion to determine whether two arc segments should merge.

In residual error space, the underlying assumption of our criterion is that for a reasonable good least-squares fit, the residual errors should be identically and independently Normally distributed with mean 0 and a variance σ^2. If this assumption holds, then the criterion can be developed as follows.

Let (x_n, y_n) $n = 1, 2, \ldots, N$ be N points that are fitted to a feature primitive (like an ellipse or a line) and let e_n be the residual error for the nth point (x_n, y_n) (residual error is defined as the geometric distance between x_n and a point on the feature closest to x_n), then for a reasonably good fit, we have $e_n \sim N(0, \sigma^2)$. Let $\bar{e}$ and S^2 be the sample mean and sample variance of the residual errors. Then

$$\bar{e} = \frac{\sum_{n=1}^{N} e_n}{N} \quad S^2 = \frac{\sum_{n=1}^{N} (e_n - \bar{e})^2}{N-1}$$

Given the noise model, we have

$$\bar{e} \sim N(0, \frac{\sigma^2}{N}) \quad \frac{(n-1)S^2}{\sigma^2} \sim \chi^2_{N-1}$$

This leads to

$$\frac{N\bar{e}^2}{\sigma^2} \sim \chi^2_1$$

A test statistic T can be derived as follows

$$T = \frac{N\bar{e}^2}{S^2} \tag{1}$$

where T is distributed as $F_{1,N-1}$ if the assumed residual error model is correct. In other words, if an ellipse fits the data reasonably well, then the test statistic T derived from the residual errors should be distributed as $F_{1,N-1}$. Distribution of T is therefore the criterion used to decide whether a merge should occur.

Based on this criterion and given two arc segments S_1 and S_2, there are three possible ways to use this criterion. First, fit two ellipses to S_1 and S_2 separately and derive two test statistics T_1 and T_2. The two arc segments are merged if both T_1 and T_2 follow the desired distribution. Second, fit an ellipse to the combined data of S_1 and S_2 and compute a test statistic T. The two arc segments are merged if T follows the desired distribution. Finally, the third method fits an ellipse to the combined data like the second one. The residual errors are then computed separately for the two arc segments, yielding two test statistics T_1 and T_2. The two arc segments are merged only if both test statistics follow the same desired F distribution. Our study favors the third method. The first test can not measure the similarity between the parameters of the two fitted ellipses. Two ellipses may fit both S_1 and S_2 well as indicated by the desired distribution with T_1 and T_2, they may represent two completely different ellipses. The second method works well if two arc segments are of the similar length. If one arc segment is much longer than the other one, then the distribution of T is largely determined by the longer feature. For example, it may so happen that the longer feature fit well while the short feature doesn't fit well. In this case, two arc segments should not be merged. However, the test statistic T may indicate otherwise. In summary, we adopted the third method to decide whether two arc segments should merge. Two arc segments are merged only when both test statistics T_1 and T_2 pass the statistical test.

One problem with the third method is that the test statistic does not have any control over the variance of the residual error. The null hypothesis only requires that the residual error is distributed with zero mean and an unknown variance σ^2. In practice, for a reasonably good fit, it is generally believed that the sample variance of the residual errors should be small. Large sample variance may be a sign of poor fitting. The test statistic T, however, can not capture this intuition. One way that can partially solve this problem is to perform a test of variance. For a good fit for both arc segments, the residual errors from both arc segments should have the same or close population variances. This can be tested by statistically testing the distribution of the ratio of two sample variances. The ratio should have a $F_{N-1,M-1}$, where N and M are the numbers of points for the two arc segments respectively. Two arc segments are merged only if their respective T statistic and the ratio of variance follows the desired desired F distributions. An alternative is to individually test the variance for each feature. For example, we may test $\sigma^2 < \sigma_0^2$, where σ_0^2 is a threshold. This, however, requires the knowledge of σ_0^2, which may vary from image to image.

3 Experimental Results

We tested our technique on a wide range of images under different illumination conditions and view directions. Figures 3 and 3 show examples of the detected hole boundaries superimposed on the original image using the ellipse detection technique introduced in section 2

4 Summary

In this chapter, we introduce a technique for detecting ellipses from lists of connected edge pixels. The proposed technique improves upon Rosin's nonparametric arc segmentation technique by introducing statistical criteria for clustering arc segments belonging to the same ellipse. The improvement led to a longer elliptical arc segment for each ellipse and fewer spurious elliptical arc segments. The performance study of this technique with real images obtained under different illumination conditions and view parameters reveal the robustness of the technique.

References

[1] A. S. Aguado, M. Montiel, N. Eugenia, and M. S. Nixon. On using directional information for parameter space decomposition in ellipse detection. *Pattern Recognition*, 29(3):369–381, 1996.

[2] C. Ho and L. Chan. A fast ellipse/circle detector using geometry symmetry. *Pattern Recognition*, 28(1):117–124, 1995.

[3] D. G. Lowe. Three dimensional object recognition from single two dimensional images. *Artificial Intelligence*, 31:355–395, 1987.

[4] D. Pao and H. F. Li. A decomposable parameter space for the detection of ellipses. *Pattern recognition letters*, 14:951–958, 1993.

[5] Paul L. Rosin and G. A. W. West. Non-parametric segmentation of curves into various representations. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(12):1140–1153, 1995.

[6] R. K. K. Yip and D. Leung P. K. S. Tam. Modification of the hough transform for circles and ellipses detection using a 2d array. *Pattern Recognition*, 25:1007–1022, 1992.

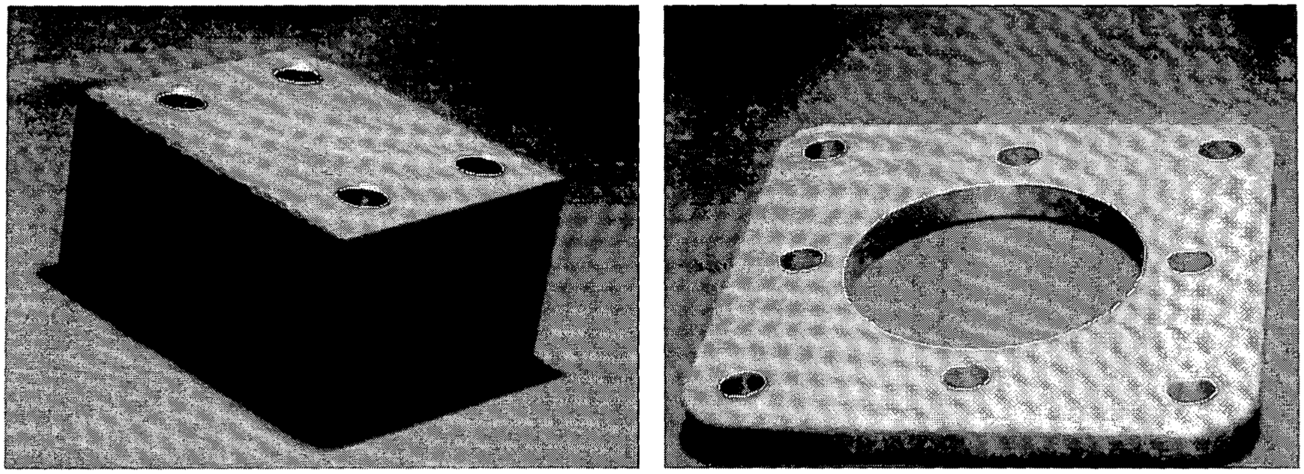

Figure 3: Detected hole boundaries superimposed on the original image.

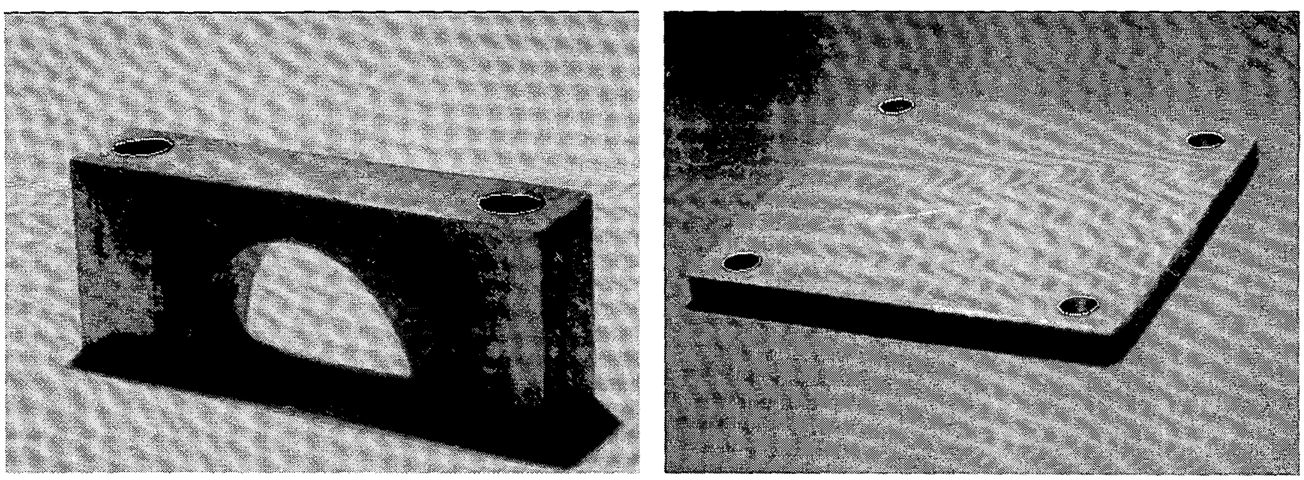

Figure 4: Detected hole boundaries superimposed on the original image.