Combining Convex Hull and Directed Graph for Fast and Accurate Ellipse Detection

Anonymous cvm submission

Paper ID 86

Abstract

Detecting ellipses from images is a fundamental task in many computer vision applications. However, due to the complexity of real-world scenarios, it is still a challenge to detect ellipses accurately and efficiently. In this paper, we propose a novel method to tackle this problem based on the fast computation of *convex hull* and *directed graph*, which achieves promising results on both accuracy and efficiency. We use Depth-First-Search to extract branch-free curves after adaptive edge detection. Line segments are used to represent the curvature characteristic of the curves, followed by splitting at sharp corners and inflection points to attain smooth arcs. Then the convex hull is constructed, together with the distance, length, and direction constraints, to find co-elliptic arc pairs. Arcs and their connectivity are encoded into a sparse directed graph, and then ellipses are generated via a fast access of the adjacency table. Finally, salient ellipses are selected subject to strict verification and weighted clustering. Extensive experiments are conducted on eight real-world datasets (six publicly available and two built by ourselves), as well as five synthetic datasets. Our method achieves the overall highest F-measure with competitive speed compared to representative state-of-the-art methods.

1. Introduction

As one of the most common geometric primitives, ellipses often appear in natural and artificial scenes. In particular, 3D circular or elliptic objects are usually projected as ellipses on the image. Therefore, accurate detection and localization of ellipses from images provides us with a powerful tool for pattern recognition and visual understanding [1]. Actually, ellipse detection is broadly applied in the fields of camera calibration [2, 3], industrial component inspection [4, 5], traffic sign detection [6], cell segmentation [7], pupil tracking [8], object localization for the robotic platform [9], and so on. See Fig. 1 as a reference.

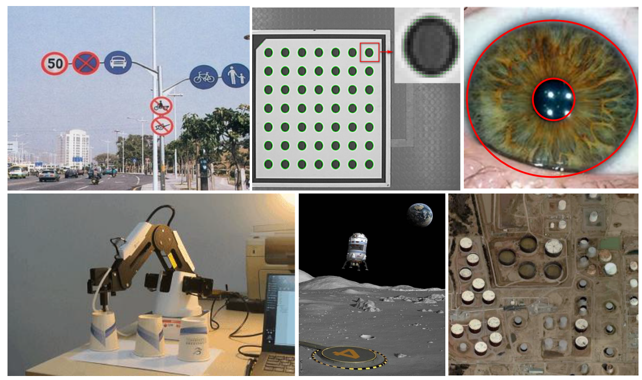

Figure 1. A wide variety of applications of ellipse detection in the real world, which provides us with a powerful tool for multiple visual understanding tasks.

Although the ellipse detection problem has gained a lot of attention in literature, it is still very challenging. The major difficulties are the presence of noise, disturbance or occlusion by other objects, image blur or flaw, and varying illuminations. These issues either break the elliptic boundaries as several low-quality arc segments, thus make the differential computations such as tangents inaccurate, or leave the ellipse partially visible, which degrades the ellipse fitting quality. Besides, the requirement of fast detection for real-time scenarios further brings the difficulty.

As a well-known geometric primitive detector, *Hough transform* (HT) is explored for ellipse detection by numerous work [10, 11, 12, 13, 14, 15]. However, due to the five-dimensional (5D) parameter space of an ellipse, HT consumes a noticeable amount of storage and time [16, 17], which seriously prevents its applications, especially for complicated images needing high-speed processing. Besides, HT suffers from the careful tuning of bin size and peak threshold, hence it may detect false ellipses or lose positive ones if the model parameters are not optimal.

The recent methods based on the *edge following* technique exhibit promising detection performance, in which the connectivity between edge pixels, continuity of arcs are used [18]. Candidate ellipses are generated by incremental least-squares fitting or arc grouping. However, direct ellipse fitting for short arcs inevitably results in errors [19]. Although other methods first group arcs together, complex arc grouping strategies are usually designed, where differential calculations or HT are introduced, hence they are more sensitive to noise or less efficient.

Different from aforementioned methods, in this paper, we introduce a new ellipse detector by a more effective arc grouping scheme, aiming to improve the detection ability in both accuracy and efficiency. We use Depth-First-Search (DFS) to extract continuous edge curves, followed by the identification of sharp corners and inflection points to attain smooth arcs. Then, the convex hull is first introduced to distinguish the convexity of arc pairs, along with the fast computation of arc distance, length, and directions. Due to the avoidance of calculations of gradients and tangents for the edge pixels, our method is more robust to noise. Based on these constraints, a sparse directed graph is built, by which arc pairs and their connectivity can be fast accessed to generate candidate ellipses. Finally, a stringent verification and a discriminative clustering are applied to further improve the detection accuracy. In a nutshell, the contributions of this work are as follows:

- a fast and accurate ellipse detector competent of detecting complicated real-world images, as well as occluded, overlapping, concentric, and concurrent ellipses;
- a novel arc grouping scheme based on the efficient computation of the convex hull and sparse directed graph, together with a more discriminative clustering criterion to depress repetitive ellipses, and
- the superior performance with less time consumption on a series of datasets compared with the representative state-of-the-art methods.

The rest of this paper is organized as follows. In Section 2, we briefly review the most related work from the perspective of ellipse generation and verification. The detailed steps of our method are presented in Section 3. Then we describe the datasets, experimental results, and performance of the proposed approach in Section 4. A general conclusion and future work are given in Section 5.

2. Related work

The significance of ellipse detection is witnessed by the large amount of work presented in the literature. In general, they can be classified as Hough transform based methods and edge following techniques.

2.1. Hough transform

Most of the traditional methods for ellipse detection rely on HT [20] to estimate the parameters, which casts the detection problem into a peak finding process. The basic principle of HT is voting each edge pixel to a 5D parameter space, and then the local peak exceeding a certain threshold is selected out as an ellipse. Although simple for implementation, it is usually unpractical to directly apply HT to ellipse detection in real images, due to the expensive storage and time load, which are $O(m^5)$ and $O(n^5)$ [17], respectively. To reduce the memory consumption, accelerate the detection, and improve the accuracy of the standard HT, a great number of variants are put forward. *Randomized HT* (RHT) [21] and *probability HT* (PHT) [22] sample subset of pixels rather than all pixels for voting, and thus a many-to-one scheme is built to replace the primary one-to-many scheme. McLaughlin [13] extends RHT to detect ellipses by randomly selecting three non co-linear points, but it is sensitive to occlusion and overlapping ellipses. Lu et al. [23] propose the iterative RHT to circumvent the noise susceptibility of RHT, but it has to divide an image into sub-images for multiple ellipse detection. On the other hand, some methods combine geometric properties of ellipses with HT to lower the voting space. Xie et al. [24] estimate the semi-axis length of the hypothetical ellipses to reduce the 5D space to 1D. Similarly, Chia et al. [25] use the foci feature to realize the same effect. Geometric symmetry is also explored to decompose the voting space, by which elliptic centers are first located and then the remaining parameters are solved [26, 27]. However, these methods are easily deteriorated by occluded or semi ellipses. Besides, we point out that HT based methods are still inefficient in practice, prone to generate false detections as the number of ellipses increasing, suffer from noise and background clutter, and take much effort to tune the required parameters such as the bin size and peak threshold [18].

2.2. Edge following

Different from HT working on the pixel level, edge following methods utilize continuous arcs for ellipse detection, in which edge curves are extracted and geometric characteristics such as convexity or tangents are explored. Compared with HT, edge following methods are more efficient, and currently are the benchmark among the ellipse detection field. For instance, Kim et al. [28] first extract arcs approximated by short line segments, and then frequently use the least-squares fitting to estimate elliptic parameters. Libuda et al. [29] improve the performance of [28] with less memory consumption. Mai et al. [30] inherit the idea of [28], but further link line segments to form arcs based on the adjacency and curvature constraints. However, due to the out of consideration for validating candidate ellipses, there are multiple false detections. Chia et al. [31] adopt a split and merge scheme for arcs, where co-elliptic arc pairs are grouped as an alignment problem. Nevertheless, the complex and iterative optimization process hinders its real-time usage in practice. The detector proposed by Prasad et al. [1] make use of the information of edge convexity and curva-

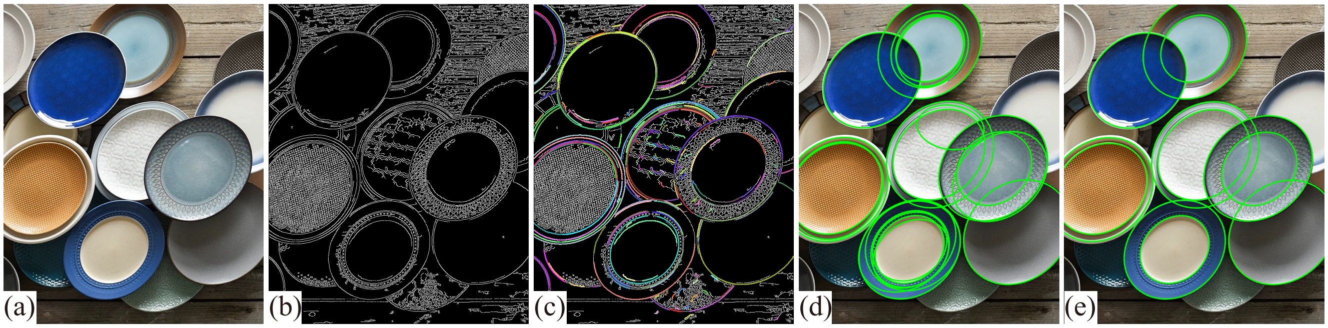

Figure 2. The workflow of our method. (a) Input image; (b) edge detection by adaptive Canny detector; (c) arc extraction via sharp corners and infection points identification; (d) candidate ellipse generation after arc grouping; (e) finally detected ellipses after validation and clustering. The proposed method is competent to detect ellipses in complex real-world images.

tures for arc grouping, although improvements are attained, it also suffers from long computational time. Fornaciari et al. [32] propose a fast ellipse detector for the embedded vision system, in which arcs are classified into four quadrants based on the gradient computation, and then parameters are estimated by the parallel chord theorem and 2D HT voting. Jia et al. [33] promote the performance of [32] by introducing a projective invariant to prune line segments and group arcs. However, both [32] and [33] encounter the same problem, that is the number of arcs for grouping must be at least three, which is impractical for occluded or semi ellipses. Dong et al. [34] take the similar scheme of [32] and incorporate the gradient analysis, but also divide the arcs into four quadrants, hence inevitably break the integrity of complete ellipses. Recently, Lu et al. [18] revisit the line detection method proposed by [35] to attain a high-quality ellipse detector, because of the iterative linking of line segments and voting for arcs, the method is much slower than [33]. Meng et al. [36] design an arc adjacency matrix (AAM) to represent the arc pair relationship, in which curvatures and tangents are computed to make AAM sparse. However, as [24, 37] pointed, curvatures and tangents are more sensitive to noise than edge points.

3. Methodology

Our method adopts a standard edge following framework, which contains three main steps: (1) edge detection and elliptic arc extraction; (2) arc grouping and candidate ellipse generation; (3) ellipse validation and clustering. The workflow of our method is shown in Fig. 2. We present details of each step in the following.

3.1. Edge detection and elliptic arc extraction

Given an input image, the very first step is to extract the edge map. Here, we implement an adaptive Canny detector [38] for this purpose, because of the efficiency and avoidance of parameter tuning. The higher threshold ensures that only 10% of the image pixels are marked as edge pixels, while the lower threshold equates 0.3 times of the higher threshold. To attain branch-free curves as shown in Fig. 3, starting from a seed point, we use the Depth-First-Search (DFS) to expand continuous curves according to the 8-connected domain of the edge points.

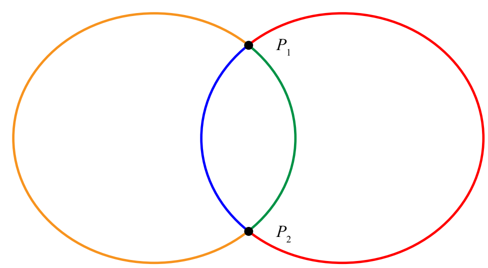

Figure 3. The two bifurcation points P_1 and P_2 separate the edge curve into four branch-free curves indicated by different colors.

After the attainment of branch-free curves, we continue to extract smooth arcs. To this end, a parameter free method [39] improved from Ramer-Douglas-Peucker (RDP) algorithm [40] is first applied to simplify curves via a series of line segments $\{l_i = P_{i-1}P_i | P_i \in \mathbb{R}^2\}_{i=1}^n$, by which we can effectively compute both the magnitude and direction of edge curvatures, as illustrated in Fig. 4(a). An angle α_i is a sharp corner, indicating the major variation in curvature magnitude, if

$$\cos\alpha_i = \frac{\overrightarrow{l}_i}{\|\overrightarrow{l}_i\|_2} \cdot \frac{\overrightarrow{l}_{i+1}}{\|\overrightarrow{l}_{i+1}\|_2} \geq \cos Th_\theta, \tag{1}$$

where $\overrightarrow{l}_*$ is the directional vector of the line segment l_*, and Th_θ is the angle threshold. Further, a point between l_i and l_{i+1} is an infection point, indicating the variation in curvature direction, if

$$\left(\frac{\overrightarrow{l}_{i-1}}{\|\overrightarrow{l}_{i-1}\|_2} \times \frac{\overrightarrow{l}_i}{\|\overrightarrow{l}_i\|_2}\right) \cdot \left(\frac{\overrightarrow{l}_i}{\|\overrightarrow{l}_i\|_2} \times \frac{\overrightarrow{l}_{i+1}}{\|\overrightarrow{l}_{i+1}\|_2}\right) = -1. \tag{2}$$

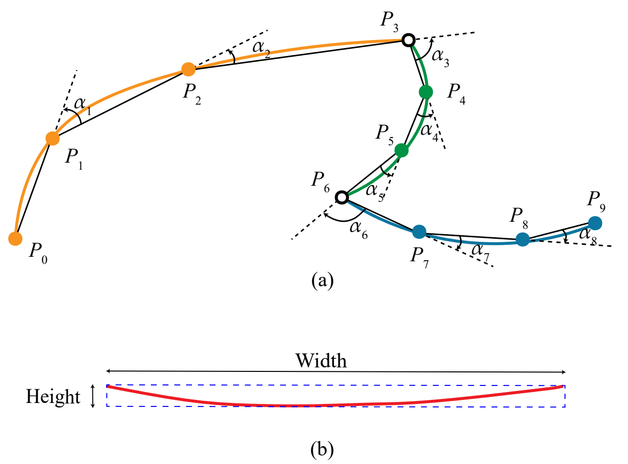

Figure 4. (a) An edge curve is approximated by nine line segments. From the inner and cross product computation, we find that α_3 is a sharp corner while P_6 is an inflection point . (b) The aspect ratio of the minimum area bounding box (dashed rectangle) is used to remove straight segments for fast detection.

Hence, α_3 is a sharp corner while P_6 is an inflection point in Fig. 4(a). Then we split curves at these points to obtain arc segments.

To speed up the following processing, we further remove straight segments based on the minimum area bounding box as illustrated in Fig. 4(b). We remove the segment if its aspect ratio

$$\frac{\max\{\text{Height}, \text{Width}\}}{\min\{\text{Height}, \text{Width}\}} > Th_r.$$

Since arc quality is critical for arc grouping, we further access each arc Arc_i by computing an inlier ratio via a fitted ellipse, which is defined as

$$I(\text{Arc}_i) = \frac{1}{|\text{Arc}_i|} \sum_{p \in \text{Arc}_i} \mathbb{1}\{dist(p, e) < \varepsilon\}. \tag{3}$$

Where $\mathbb{1}$ is the indicator function and equates to one if and only if the distance from the edge pixel p to the ellipse e is less than ε, which is equal to one pixel in default. Arcs with low inlier ratio, i.e., $I(\text{Arc}_i) < Th_{ir}$, where Th_{ir} is the threshold, are regarded as non-elliptic arcs thus are deleted. To keep consistency between different arcs, edge points of each arc are stored in the counter-clockwise order.

3.2. Arc grouping and candidate ellipse generation

Since short arcs may result in major fitting errors, we first group them from the same ellipse together by a local to global scheme. The local search aims to link adjacent arc pairs caused by noise interference, while the global process elaborates to group distant ones.

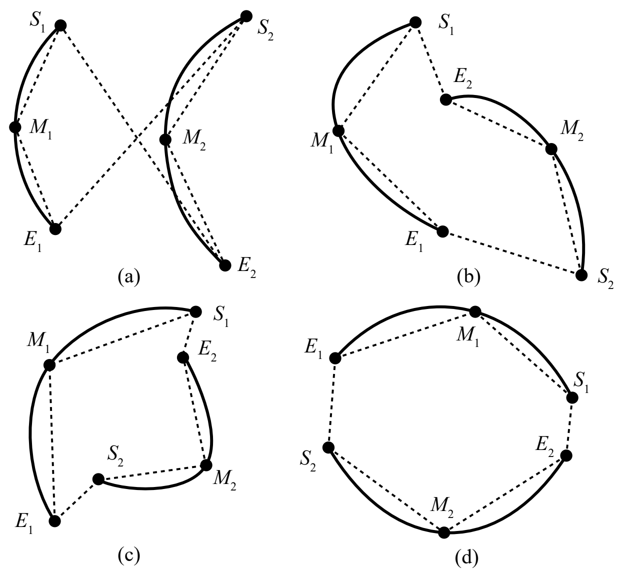

Figure 5. Grouping arc pairs based on the convex hull computation, where S_i, E_i, and M_i are the endpoints and midpoints of Arc_i. The arcs in Case (d) can be grouped, while the others cannot.

We introduce the convex hull to represent the ellipse convexity and sample the endpoints and midpoints of two arcs to define it as illustrated in Fig. 5. To check whether the polygon formed by the six points is convex, we simply judge whether the sign of the cross product of adjacent line segments are all positive. Due to the convexity of arcs themselves, there are in fact merely four computations, hence the judgement is fast enough. See Fig. 5 as reference, and two arcs Arc_i and Arc_j are said constituting a convex hull if

$$\begin{cases} \text{sgn}(\overrightarrow{M_1E_1} \times \overrightarrow{E_1S_2}) > 0, & \text{sgn}(\overrightarrow{E_1S_2} \times \overrightarrow{S_2M_2}) > 0 \\ \text{sgn}(\overrightarrow{M_2E_2} \times \overrightarrow{E_2S_1}) > 0, & \text{sgn}(\overrightarrow{E_2S_1} \times \overrightarrow{S_1M_1}) > 0 \end{cases} \tag{4}$$

From the local perspective, adjacent arcs tend to come from the same ellipse. We find arc pairs whose end-point distance are no more than one pixel, and merge them together if (1) the endpoints and midpoints of them constitute a convex hull and (2) the inlier ratio of them is larger than each arc after merging. Local grouping significantly reduces the number of arcs participating in the global grouping, hence accelerating the detection process. Because noise causes many adjacent arcs, and most of them can be merged, while other invalid arc pairs are directly skipped in the subsequent processing.

When two arcs Arc_i and Arc_j are not adjacent enough, we try to group them again by four global constraints.

Arc length constraint. Based on observation, arcs from ellipses with similar size usually have similar length. Hence

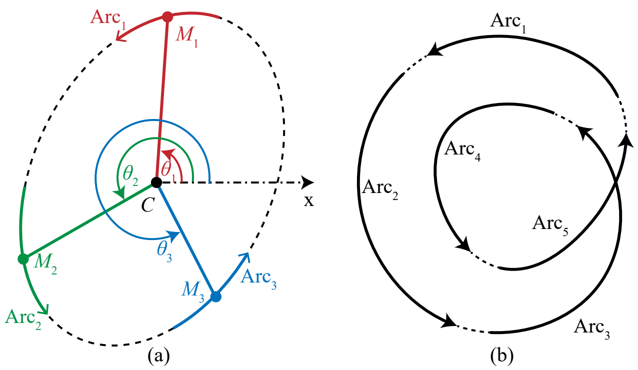

Figure 6. (a) Connection of counter-clockwise arcs and computation of rotation angles represented by θ_i. (b) A path with self-intersection, which can be effectively removed by our method.

if the length ratio of Arc_i and Arc_j satisfies

$$1/Th_{lr} < |\mathrm{Arc}_i|/|\mathrm{Arc}_j| < Th_{lr},$$

then they are checked by the subsequent constraints. Otherwise, the arc pair is invalid and ignored.

Distance constraint. Although global constraints aim to group distant arcs, two arcs apart largely are also less likely from the same ellipse. Arc_i and Arc_j are said satisfying the distance constraint if

$$\frac{dist(M_i, M_j)}{\max\{|\mathrm{Arc}_i|, |\mathrm{Arc}_j|\}} < Th_d,$$

where M_* is the middle point of Arc_*, and $dist(M_i, M_j)$ is the distance between two middle points.

Convex hull constraint. According to the convexity of ellipses, Arc_i and Arc_j can be grouped if their endpoints and midpoints form a convex hull.

Direction constraint. Arc pair $\langle \mathrm{Arc}_i, \mathrm{Arc}_j \rangle$ satisfying the above criterion are called co-elliptic, referred to as $\mathrm{Arc}_i \rightarrow \mathrm{Arc}_j$. It should also be noted that the arcs are connected in order, that is, $\mathrm{Arc}_i \rightarrow \mathrm{Arc}_j$ and $\mathrm{Arc}_j \rightarrow \mathrm{Arc}_i$ are two different situations. Arcs should be connected counter-clockwise, as shown in Fig. 6 (a), C is the center of the corresponding ellipse, and θ_i represents the rotation angle from the positive x-axis direction to the vector $\overrightarrow{CM_i}$. For example, M_1 can be connected to M_2 is that M_1 can be co-linear with C and M_2 after a rotation of no more than 180°,

$$fmod(\theta_2 - \theta_1 + 360°, 360°) < 180°, \tag{5}$$

where $fmod(x, y)$ stands for the floating point remainder of the division operation x/y. In practice, we approximate the rotation angles $\{\theta_i\}$ based on the pre-fitted ellipse in the inlier ratio step to speed up detection.

Local and global grouping discovers the relationship between any two arcs, by which we construct a directed graph to encode the relationship of all arcs. In the graph, vertices stand for arcs, and directed edges represent the connected co-elliptic arc pairs in counter-clockwise direction. Because of the above strict pairing constraints, the graph is usually sparse, and therefore we use the adjacency list instead of the adjacency table to reduce the memory usage.

By depth-first searching the directed graph, we can obtain a path

$$\mathrm{Arc}_{k_1} \rightarrow \mathrm{Arc}_{k_2} \rightarrow \cdots \rightarrow \mathrm{Arc}_{k_n},$$

which represents a group of arcs where any two adjacent arcs are co-elliptic. Benefiting from the data structure of adjacency list, we can merely visit the neighbors of a vertex without traversing the other vertices, hence greatly reduces the time consumption of the access. Note that there may exist complex paths with self-intersection as illustrated in Fig. 6 (b). In this case, we use the following criteria

$$R = \frac{1}{360°} \sum_{i=1}^{n} \Delta\theta_i$$

to filter out self-intersection paths if $R \neq 1$. Where $\Delta\theta_i$ is defined as

$$\Delta\theta_i = fmod(\theta_{i+1 \bmod n} - \theta_i + 360°, 360°).$$

Intuitively, R represents the number of circles around the center when a virtual point moves along the path. For valid paths, R is always equal to one. Through the searching process, all co-elliptic arc groups are found, and then a direct least-squares-based ellipse fitting method [41] is applied to attain candidate ellipses.

3.3. Ellipse validation and clustering

Due to the discrete properties of edge pixels, there may exist false ellipses among candidates. To further improve the detection accuracy, we execute an ellipse validation and compute the salient score $S(e)$ for each candidate ellipse e formed by the arc group G, which is defined as

$$S(e) = \frac{1}{\sum_{\mathrm{Arc} \in G} |\mathrm{Arc}|} \sum_{\mathrm{Arc} \in G} \sum_{p \in \mathrm{Arc}} \mathbb{1}\{dist(p, e) < \varepsilon\}, \tag{6}$$

where p is the edge pixel from the corresponding Arc in the same group. A candidate ellipse is validated to be true if $S(e) \geq Th_{ss}$, otherwise we remove it due to the unreliability. Let $e = (a, b, x_c, y_c, \theta)$ be the ellipse parameters, where a, b are the semi-axis length, (x_c, y_c) is the elliptic center, and θ is the rotation angle along the horizontal axis. Then, we use a weighted clustering scheme based on the Euclidean distance to evaluate the distinctiveness of two ellipses e_i and e_j

$$D(e_i, e_j) = \sqrt{\sum_{\lambda=1}^{5} k_\lambda \cdot (e_{i\lambda} - e_{j\lambda})^2}. \tag{7}$$

Ellipses e_i and e_j are clustered together if $D(e_i, e_j) < 20$ (suggested by [14]). The weight k_λ is equal to one except for the rotation angle θ and is defined as

$$k_\theta = \min\{\frac{a_i - b_i}{a_i + b_i}, \frac{a_j - b_j}{a_j + b_j}\}.$$

Note that this weighting scheme effectively eliminates the angle influence caused by the rotation symmetry of circles.

4. Experimental Results

In this section, the performance of the proposed method is comprehensively evaluated by a series of experiments including (1) parameter discussion, (2) comparison with six representative state-of-the-art methods regarding synthetic and real-world images, (3) robustness against ellipse variations, and (4) robustness against the intersection over union (IoU) variations. All experiments are executed on a desktop computer with Intel Core I7-7700K CPU @4.20 GHz and 32GB RAM.

4.1. Datasets

We use five synthetic datasets and eight real-world datasets to verify the general capability of the proposed ellipse detector. Fig. 7 illustrates several images from these datasets, which have different characteristics as described in the following. Code and all datasets will be released upon acceptance.

Synthetic datasets. Synthetic ellipses involving occlusion, overlapping, noise, concentric, and concurrent are tested. There are 300 images with occluded ellipses and 300 images with overlapping ones [1], in the resolution of 300×300. Each image has $\beta \in \{4, 8, 12, 16, 20, 24\}$ ellipses under the constraint that they must overlap with at least one ellipse. The complex occlusion or overlapping, especially with the number of ellipses increasing, make the detection tough enough. To test the robustness of the ellipse detector, we use the function *imnoise(img, 'salt & pepper', density)* in Matlab with density ranging from 4% to 24% at the step 4% to add salt-and-pepper noise in the images with 8 overlapping ellipses. Besides, we further test 720 images with concentric ellipses and 1200 images with concurrent ones [36] under the resolution 600×600. These images are challenging enough because of the multiple cracked arcs for grouping.

Real-world datasets. Dataset Prasad [1] has 400 images sampled from 48 categories in Caltech256 dataset [42]. However, there are only 198 images available online, and we complement the missing part named Dataset Prasad+ according to the file provided by the authors. The varying image size with cluttered background is the major challenge. Dataset Random [32] also contains 400 images up to 1280

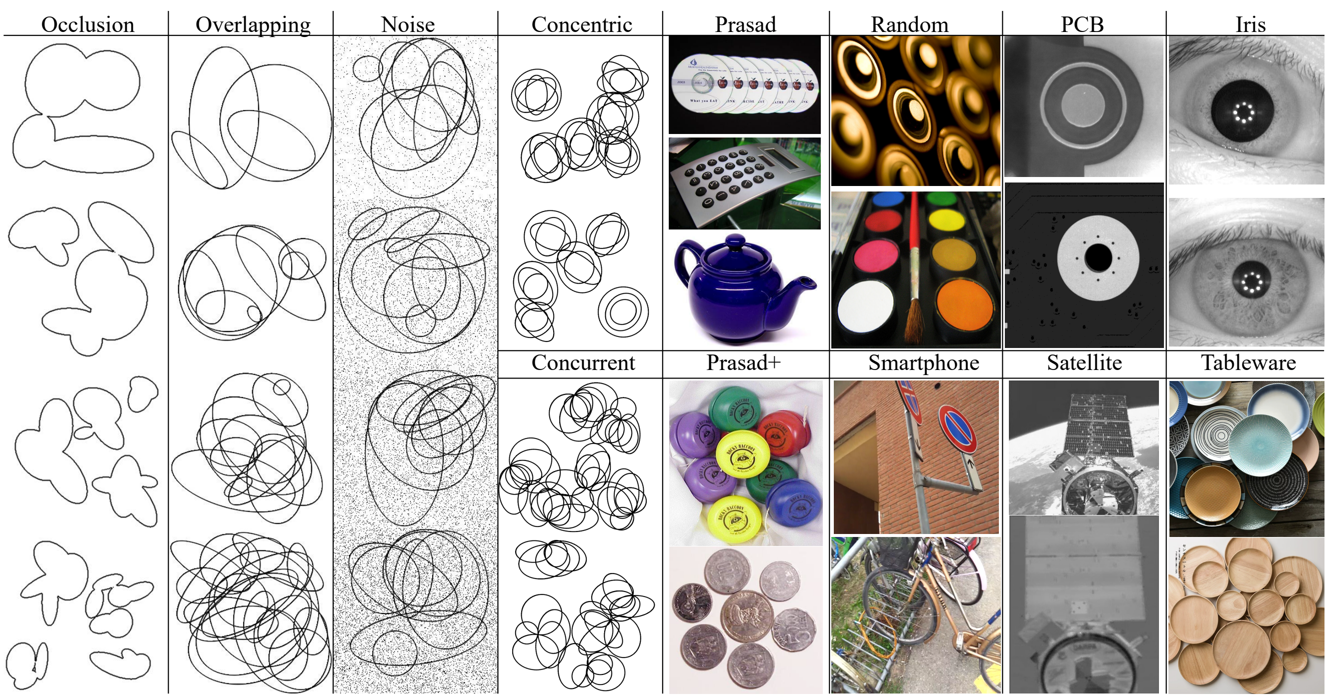

Figure 7. Example images in the test datasets. Column 1-3 show the synthetic ellipses with occlusion, overlapping, and noise, respectively. Column 4 includes synthetically concentric and concurrent ellipses. Column 5-7 are the images from datasets Prasad, Random and Smartphone, PCB and Satellite, respectively. The last column contains images from our new datasets named Iris and Tableware.

$\times$ 960 from MIRFlickr and LabelMe repositories [43] [44]. The high resolution and noisy interference dramatically degrade the detection speed and effectiveness. Dataset Smartphone [32] has 629 images collected from a video. The existence of image blur and perspective transformation are the main difficulties. Dataset PCB [45] has 100 industrial printed circuit board images. The concentric structure and substantial white noise adversely impact the performance. The satellite dataset [36] contains 757 optical images and 440 infrared images are involved, which are captured by the OEDMS camera and infrared camera of the NextSat spacecraft, respectively. The space light, camera noise, and the distant small ellipses are hard to detect. Further, we provide two new datasets named Iris and Tableware containing 100 images, respectively. Dataset Iris is used to test the detection capability for small ellipses, which are selected from CASIA Iris Database [46], while Tableware aims to simulate the robotic manipulation of cylindrical objects. All ground truth images are labeled by ourselves manually and precisely.

4.2. Evaluation metrics

To quantitatively evaluate the performance of the proposed method, three well-known metrics from information retrieval are utilized, i.e., *precision*, *recall*, and *F-measure*, which are defined as

$$\text{Precision} = \frac{|\text{TP}|}{|\text{TP} + \text{FP}|}, \text{ Recall} = \frac{|\text{TP}|}{|\text{TP} + \text{FN}|}.$$

$$\text{F-measure} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}.$$

Here, TP, FP, and FN represent the true positives, false positives, and false negatives, respectively. A detected ellipse e_d is considered to be a true positive if its *intersection over union* (IoU) regarding the ground truth e_t is no less than γ ($\gamma = 0.95$ for synthetic images and 0.8 for real images, as suggested in [32]). Otherwise, it is a false positive, and a ground truth not rightly recognized is seen as a false negative. Note that F-measure is a comprehensive performance metric. IoU is defined as

$$\text{IoU}(e_d, e_t) = \frac{area(e_d) \cap area(e_t)}{area(e_d) \cup area(e_t)},$$

where $area(e_*)$ denotes the number of pixels inside the ellipse e_*. The proposed ellipse detector is compared with six representative state-of-the-art methods including Libuda [29], Prasad[1] [1], Fornaciari [32], Jia [33], Lu [18], and Meng [36]. The source code of these methods are publicly available online, and Prasad and Lu are implemented in Matlab, while the others and ours are in C++.

[1] The implementation online for Prasad is incomplete. We re-implement the validation part based on the Section 4 in the original paper [1], as faithfully as possible.

4.3. Parameter discussion

Our method mainly involves six parameters, which are discussed in the following. (1) The angle threshold Th_θ is used to discover sharp corners, and larger value will tolerate curves with larger curvature. Based on the elliptic curvature, we fix $Th_\theta = 46°$ as it performs well for general images. (2) Th_r is the aspect ratio of bounding box to remove straight segments, and we can speed up the detection process by setting relatively small ones. However, more arcs will also be deleted. Extensive experiments suggest that $Th_r = 10$ is a better balance between the effectiveness and efficiency. (3) Inlier ratio threshold Th_{ir} is used to attain high-quality arcs. Admittedly, larger threshold will keep better arcs, but considering the discrete pixels, we choose $Th_{ir} = 0.7$ for use. (4) In the arc grouping step, Th_{lr} is the length ratio tolerance of two arcs. Because too short arcs hardly provide rich information, we let Th_{lr} be equal to 6 to find similar arc pairs. (5) Th_d is used to evaluate the distance between two arcs, due to the limit of image size, big values less likely emerge, we set $Th_d = 10$ to incorporate as many arc pairs as possible. Since the fine performance of these parameters for hundreds of images, we fixing them as intrinsic ones without user tuning. (6) The last parameter Th_{ss} in the validation step is used to select salient ellipses. We open it as an adjustable parameter according to the practical requirement. Furthermore, to reveal the performance variation regarding different Th_{ss}, we test five datasets as illustrated in Fig. 8. As observed, with Th_{ss} increasing, precision first goes up and decreases after 0.8,

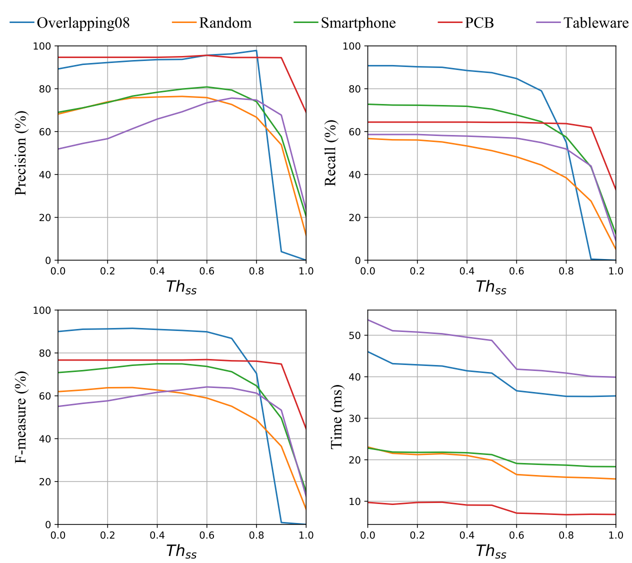

Figure 8. Investigation of the salient score parameter Th_{ss} on five datasets listed on top. A better choice of Th_{ss} falls in $[0.5, 0.7]$, considering the F-measure and time consumption.

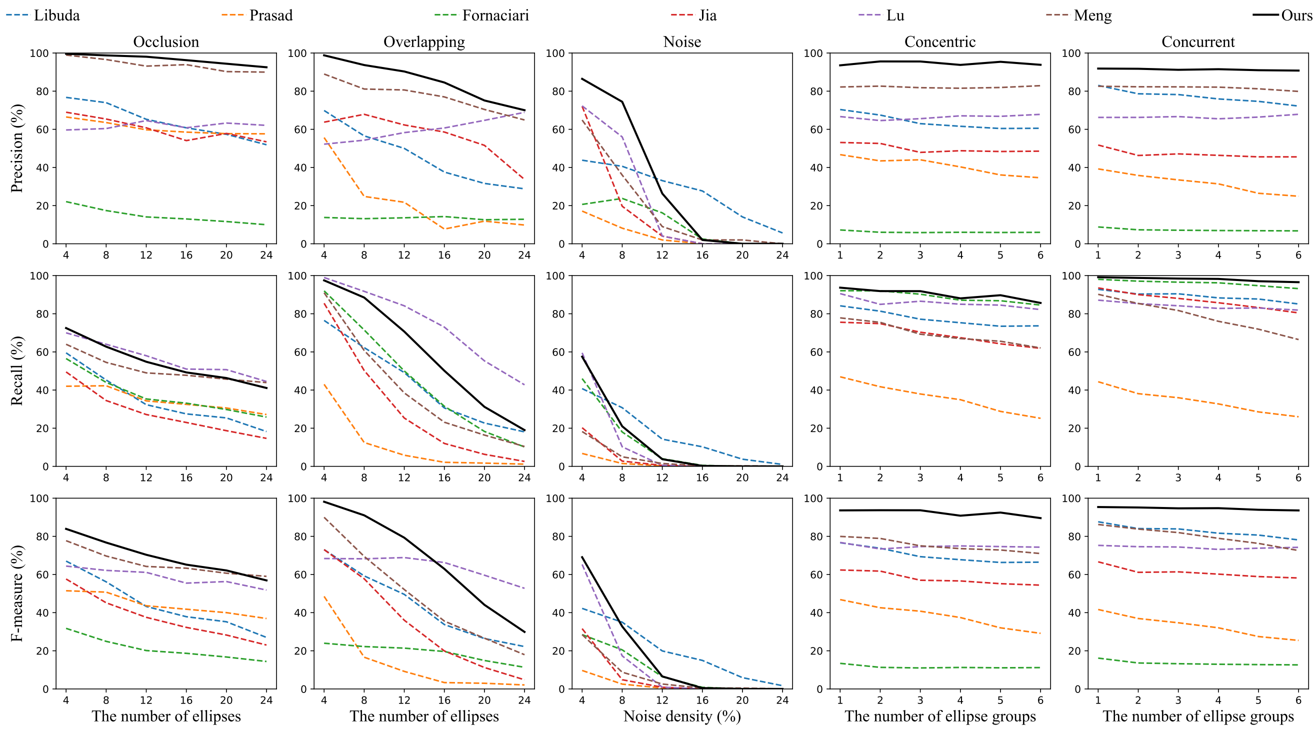

Figure 9. Ellipse detection results on synthetic datasets. Our method achieves the overall highest F-measure with superior precision.

while recall decreases beyond 0.7, which pulls down the F-measure. Taking the time consumption into consideration, we suggest $Th_{ss} \in [0.5, 0.7]$ for use.

4.4. Test on synthetic datasets

We report the detection results of synthetic images including occlusion, overlapping, noise, concentric, and concurrent in Fig. 9. As observed, the proposed detector attains the highest F-measure on datasets occlusion, concentric and concurrent, as well as the highest precision with the value more than 80%, which demonstrates its superior localization accuracy. Methods Lu and Meng share the similar performance and are lower than ours. Besides, Fornaciari has the lowest F-measure and precision on these three datasets, and Jia is better than Fornaciari, indicating the effectiveness of the added projective invariant. However, the performance of Jia and Prasad are still unsatisfactory and are lower than Libuda. Except the occlusion case, our method also achieves the highest recall on concentric and concurrent cases. For overlapping ellipses, the proposed detector has the highest F-measure when the number of ellipses is less than 20. With more ellipses, although the F-measure is lower than Lu, we still achieve the second highest one, together with the second highest recall, and Lu embraces the best recall. Nevertheless, we remain the highest precision. Note that as the number of overlapping ellipses increasing, the F-measure and recall of Prasad and Jia tend to zero, which indicates that they are subject to complex scenes. For noisy ellipses, our method returns acceptable results when the noise level is no more than 8%, and with the noise increasing, the performance of all methods decreases rapidly, this is because substantial noise breaks continuous arcs as small fragments, which adversely influences the arc grouping process. Therefore, a simple denoising step is helpful. Several detection examples are represented in Fig. 10.

4.5. Test on real-world datasets

Besides synthetic test, we further report the test results on the eight real-world datasets. The F-measure and time consumption are given in Table 1 and 2, respectively, where the red and blue colors indicate the two best F-measure. From Table 1, we can see that the proposed method attains the highest F-measure on five datasets and achieves the best detection effectiveness in general. Lu achieves the second best F-measure, but its detection speed is much slower than ours as shown in Table 2. Meng gets the third place along with the fastest speed, which benefits from its optimization operation. Jia also has the relatively small execution time, but the F-measure is a little low. Although Libuda has the fifth highest F-measure on the whole, it performs well on small ellipses, which can be concluded from the dataset Iris. But the time consumption of Libuda is very expensive and is much more than ours. Methods Fornaciari and Prasad share the similar F-measure, but Prasad takes significantly long

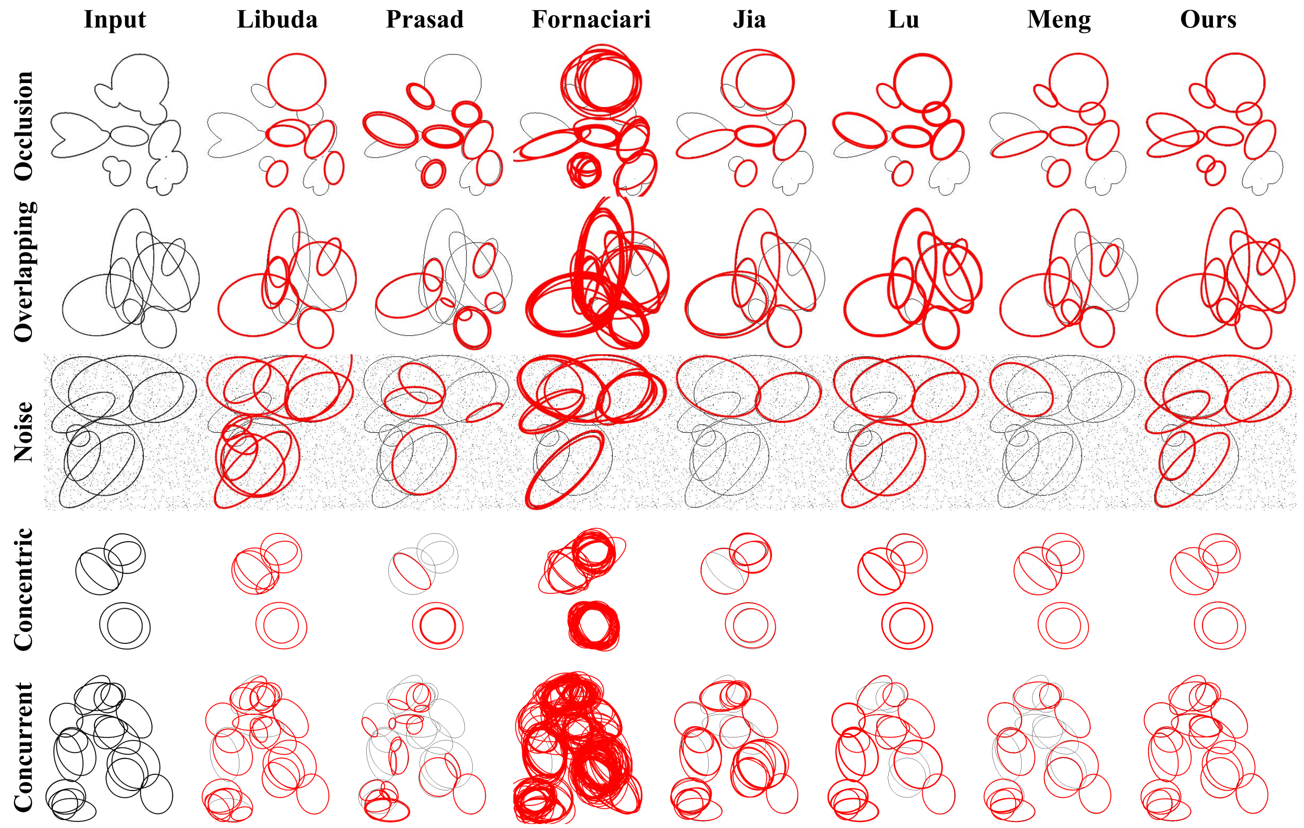

Figure 10. Ellipse detection examples on synthetic images with occlusion, overlapping, noise, concentric, and concurrent. Our method detects most of the true positives while making less false positives.

time, even 100 times than ours, which suffers from the process of complex arc grouping and HT voting. However, the F-measure of both Fornaciari and Prasad are far from satisfactory, especially for complicated images with occlusion or noise, such as the images in datasets Tableware and Satellite. As a whole, the proposed method embraces the highest F-measure with very competitive running time. Several ellipse detection examples are presented in Fig. 13. Note that the execution time of our method indicates that we can work on general camera video with 30Hz rate.

4.6. Robustness to ellipse variations

To further investigate the robustness of our method for ellipse variations regarding size, orientation, and incompleteness, we generate three datasets with the image size 512 $\times$ 512. The first dataset has 2,0000 images with the semi-major axis length varying from 1 to 200 pixels, at the same time, the axis ratio increases from 0.01 to 1 at the step 0.01. To evaluate the robustness against rotation angles, we build the second dataset by rotating the ellipse from 1° to 180° at the step 1°, fixing the semi-major axis equal to 200 pixels and varying the axis ratio from 0.01 to 1 at the step 0.01, hence there are 1,8000 images for test. The last dataset involving 3,6000 images aims to check the capacity for incomplete ellipses, where the angular coverage is from 1° to 360° at the step 1° and the axis ratio ranges from 0.01 to 1 at the step 0.01.

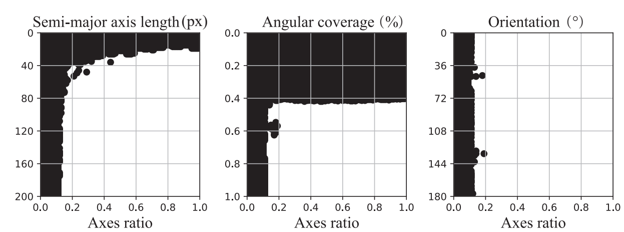

Figure 11. Robustness test results under different ellipse variations. The horizontal axis indicates the axes ratio of semi-minor axis to semi-major one, ranging from 0.01 to 1 at the step 0.01. The vertical axes are the semi-major axis length in pixel, angular coverage of ellipse arc, and ellipse orientation, respectively. Our method embraces a wide range of successful area indicated by the white region.

The results of ellipse variations are reported in Fig. 11, where the white region indicates the correctly detected el-

Table 1. Comparison on the eight real-world datasets of six different methods in terms of F-measure (%). Red and blue colors indicate the best two performance, respectively. Our method achieves the overall highest F-measure.

Method	Prasad	Prasad+	Random	Smartphone	PCB	Satellite	Iris	Tableware
Libuda	30.82	40.86	37.49	40.09	61.22	31.74	64.81	16.65
Prasad	28.78	21.35	29.1	22.25	56.11	6.81	55.52	33.07
Fornaciari	28.88	31.34	30.62	19.18	55.89	28.79	57.44	15.74
Jia	33.42	48.96	50.15	52.21	74.84	22.21	58.57	54.74
Lu	50.91	65.39	60.02	64.02	80.22	45.03	66.37	54.59
Meng	43.81	54.67	50.05	56.5	70.79	56.65	66.25	53.06
Ours	45.58	66.78	61.12	74.89	79.46	47.76	75.36	62.9

Table 2. Time (ms) comparison on the eight real-world datasets of six different methods. The proposed method can be used for camera video processing of 30Hz.

Method	Prasad	Prasad+	Random	Smartphone	PCB	Satellite	Iris	Tableware
Libuda	12.38	20.36	32.04	52.07	26.94	8.92	14.88	95.11
Prasad	1870.99	5222.84	5153.76	11743	533.97	1074.05	1451.36	16294.7
Fornaciari	3.88	10.9	11.73	16.84	5.08	2.77	2.92	74.93
Jia	3.47	7.18	9.6	12.6	4.87	2.44	2.87	40.02
Lu	78.67	277.92	334.1	618.25	54.53	17.19	27.77	4607.49
Meng	3.19	5.25	8.24	11.55	3.33	2.61	3.01	26.35
Ours	7.94	13.15	16.38	19.16	7.1	4.67	4.89	40.98

lipses and the black region means the failure cases. From Fig. 11(a), we conclude that our detector has a wide range of successful area and can detect small ellipses with the semi-major axis around 25 pixels and axis ratio slightly below 0.2. Fig. 11(b) shows that our method is able to detect incomplete ellipses with angular coverage about 150°. Furthermore, we can improve the robustness to incomplete ellipses by slightly lowering down the salient score in the validation step. The black region distributes vertically in Fig. 11(c), indicating that our method is very robust to ellipse orientation, which is a basic nature for high-quality ellipse detector.

4.7. Robustness to IoU variations

The last experiment tests the robustness of different methods to different IoU. To this end, we vary IoU from 0.5 to 0.95 at the step of 0.05 on three datasets. Admittedly, higher IoU brings more stricter constraint of an ellipse being regarded as a true positive. The detection results are reported in Fig. 12. From which, we can see that our method achieves the highest precision on all datasets. Although our recall is slightly lower, we still has the best comprehensive metric F-measure, which demonstrates the high quality performance of our detector. In contrast, Fornaciari attains the highest recall, however, due to the lowest precision, its F-measure is far from satisfactory. With the value of IoU increasing, all methods show descending trend, but our method keeps the F-measure higher than 60% when IoU $<= 0.8$. When IoU $= 0.95$, although the F-measure of some methods drop below 10% such as Prasad and Fornaciari on dataset Smartphone, we still has the F-measure more than 20%, which indicates the robustness of the proposed detector to IoU variations.

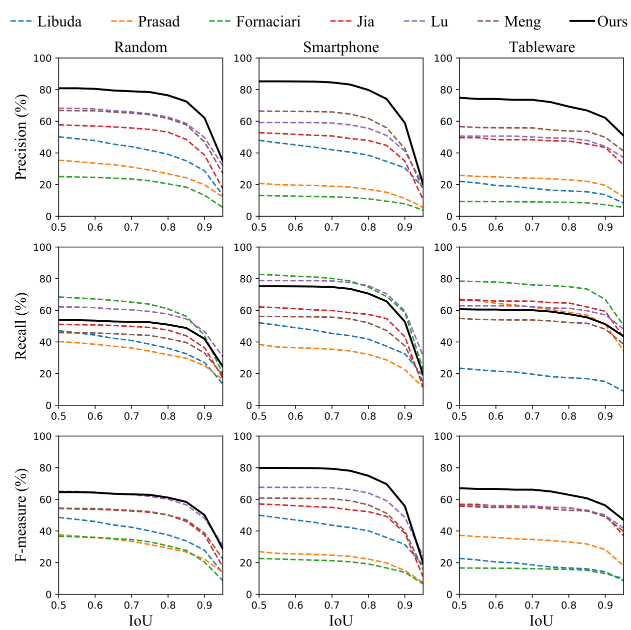

Figure 12. Robustness test results by varying different IoU values. The proposed method achieves the highest F-measure.

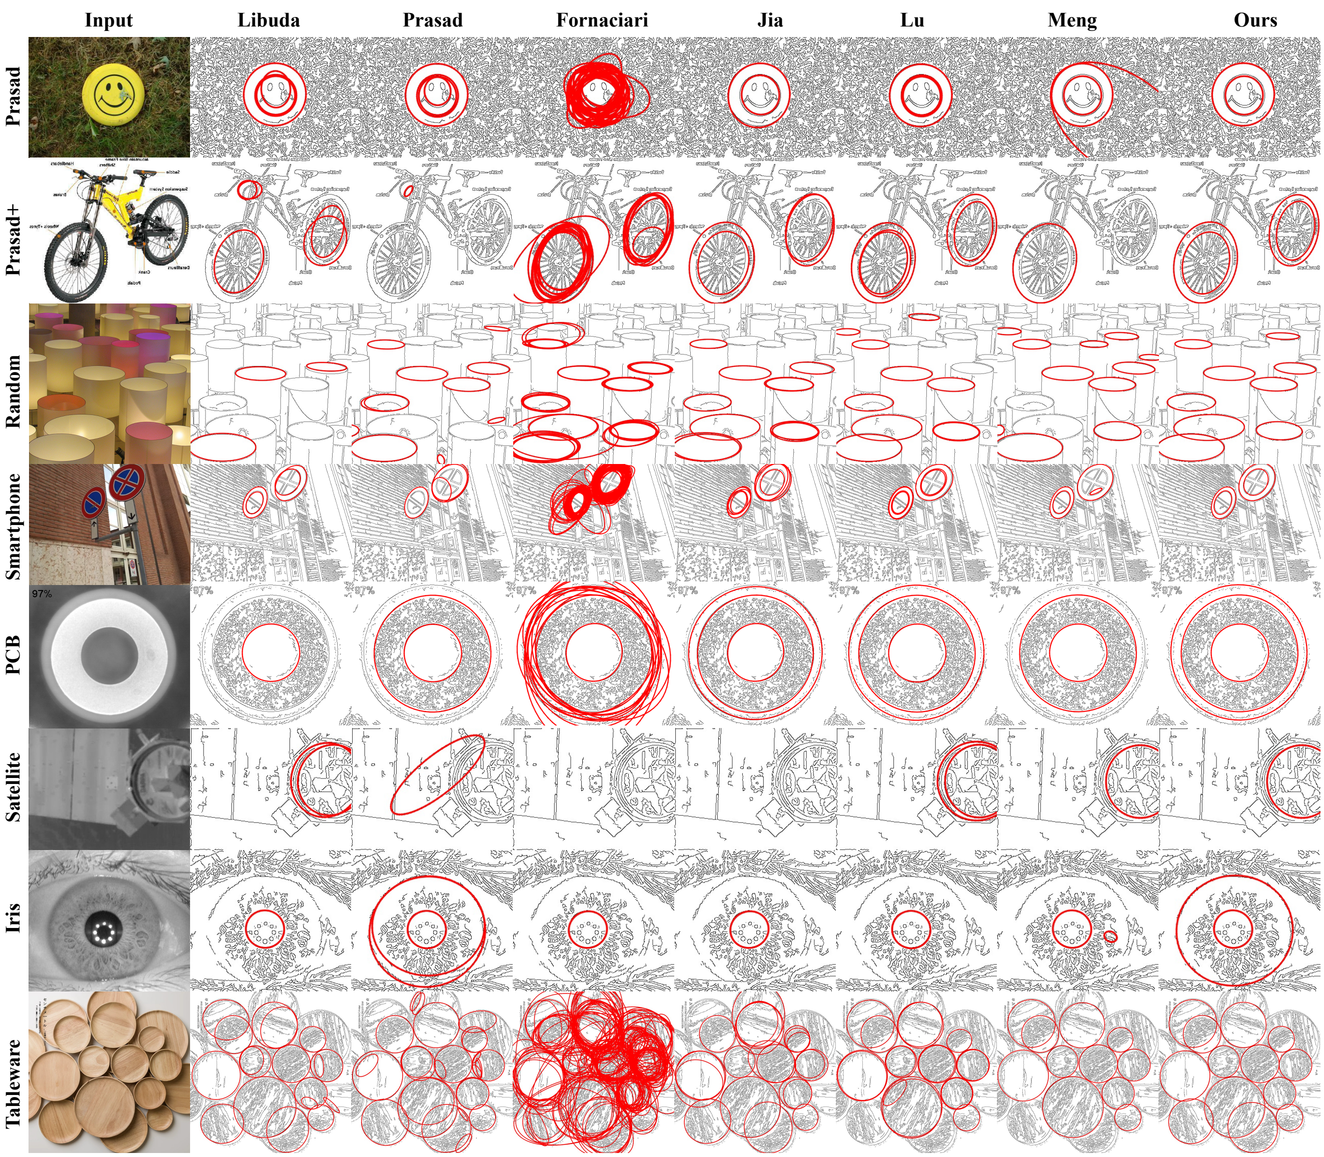

Figure 13. Sampled ellipse detection results on real-world images. The first column presents the input images from the eight datasets, and detection results of different methods are presented in the second to last columns. The proposed method detects the most true positives without false positives.

5. Conclusions

In this paper, we have presented a novel ellipse detection method by introducing the convex hull and directed graph, which performs accurately and efficiently for versatile synthetic and real-world images. We have made innovative improvements compared with previous ones. Smooth arcs are extracted by the identification of sharp corners and inflection points based on the immediate computation of inner and cross products. According to the ellipse convexity, we use convex hull to judge the convexity between arc pairs, since merely four cross products are needed, the computation is fast. By incorporating other constraints, a local to global arc grouping strategy is established. The relationship between arc pairs is encoded in a directed graph, by which all arcs from the same ellipse are found to generate candidate ellipses. Moreover, a rigorous verification and weighted clustering further enhance the accuracy by rejecting false positives and repetitive ones.

Extensive experiments on 13 datasets compared with 6 representative state-of-the-art methods demonstrate the superior performance of our method, which also has a good potential for video stream processing. In the future, we plan to apply our detector to more dedicated tasks such as camera calibration and robotic grasping.

References

[1] Dilip K Prasad, Maylor KH Leung, and Siu-Yeung Cho. Edge curvature and convexity based ellipse detection method. *Pattern Recognition*, 45(9):3204–3221, 2012. 1, 2, 6, 7

[2] Janne Heikkila. Geometric camera calibration using circular control points. *IEEE Transactions on pattern analysis and machine intelligence*, 22(10):1066–1077, 2000. 1

[3] Haifei Huang, Hui Zhang, and Yiu-ming Cheung. The common self-polar triangle of concentric circles and its application to camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4065–4072, 2015. 1

[4] Zhi-Yong Liu and Hong Qiao. Multiple ellipses detection in noisy environments: A hierarchical approach. *Pattern Recognition*, 42(11):2421–2433, 2009. 1

[5] Songlin Chen, Renbo Xia, Jibin Zhao, Yueling Chen, and Maobang Hu. A hybrid method for ellipse detection in industrial images. *Pattern Recognition*, 68:82–98, 2017. 1

[6] Cheng-Chin Chiang, Ming-Che Ho, Hong-Sheng Liao, Andi Pratama, and Wei-Cheng Syu. Detecting and recognizing traffic lights by genetic approximate ellipse detection and spatial texture layouts. *International Journal of Innovative Computing, Information and Control*, 7(12):6919–6934, 2011. 1

[7] Xiangzhi Bai, Changming Sun, and Fugen Zhou. Splitting touching cells based on concave points and ellipse fitting. *Pattern recognition*, 42(11):2434–2446, 2009. 1

[8] Lech Świrski, Andreas Bulling, and Neil Dodgson. Robust real-time pupil tracking in highly off-axis images. In *Proceedings of the Symposium on Eye Tracking Research and Applications*, pages 173–176, 2012. 1

[9] Huixu Dong, Ehsan Asadi, Guangbin Sun, Dilip K Prasad, and I-Ming Chen. Real-time robotic manipulation of cylindrical objects in dynamic scenarios through elliptic shape primitives. *IEEE Transactions on Robotics*, 35(1):95–113, 2018. 1

[10] Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972. 1

[11] Saburo Tsuji and Fumio Matsumoto. Detection of ellipses by a modified hough transformation. *IEEE transactions on computers*, pages 777–781, 1978. 1

[12] Dana H Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern recognition*, 13(2):111–122, 1981. 1

[13] Robert A McLaughlin. Randomized hough transform: improved ellipse detection with comparison. *Pattern Recognition Letters*, 19(3-4):299–305, 1998. 1, 2

[14] Cosmin A Basca, Mihai Talos, and Remus Brad. Randomized hough transform for ellipse detection with result clustering. In *EUROCON 2005-The International Conference on" Computer as a Tool"*, volume 2, pages 1397–1400. IEEE, 2005. 1, 6

[15] YG Chen and Yan Yang. Two improved algorithms for ellipse detection based on hough transform. *Semiconductor Optoelectronics*, 38(5):745–750, 2017. 1

[16] Zhengyou Zhang. Parameter estimation techniques: A tutorial with application to conic fitting. *Image and vision Computing*, 15(1):59–76, 1997. 1

[17] Priyanka Mukhopadhyay and Bidyut B Chaudhuri. A survey of hough transform. *Pattern Recognition*, 48(3):993–1010, 2015. 1, 2

[18] Changsheng Lu, Siyu Xia, Ming Shao, and Yun Fu. Arc-support line segments revisited: An efficient high-quality ellipse detection. *IEEE Transactions on Image Processing*, 2019. 1, 2, 3, 7

[19] Kenichi Kanatani, Yasuyuki Sugaya, and Yasushi Kanazawa. Ellipse fitting for computer vision: implementation and applications. *Synthesis Lectures on Computer Vision*, 6(1):1–141, 2016. 1

[20] Paul VC Hough. Method and means for recognizing complex patterns, December 18 1962. US Patent 3,069,654. 2

[21] Lei Xu, Erkki Oja, and Pekka Kultanen. A new curve detection method: randomized hough transform (rht). *Pattern recognition letters*, 11(5):331–338, 1990. 2

[22] Nahum Kiryati, Yuval Eldar, and Alfred M Bruckstein. A probabilistic hough transform. *Pattern recognition*, 24(4):303–316, 1991. 2

[23] Wei Lu, Jinglu Tan, and Randall Floyd. Automated fetal head detection and measurement in ultrasound images by iterative randomized hough transform. *Ultrasound in medicine & biology*, 31(7):929–936, 2005. 2

[24] Yonghong Xie and Qiang Ji. A new efficient ellipse detection method. In *Object recognition supported by user interaction for service robots*, volume 2, pages 957–960. IEEE, 2002. 2, 3

[25] Alex Yong Sang Chia, Maylor K. H. Leung, How Lung Eng, and Susanto Rahardja. Ellipse detection with hough transform in one dimensional parametric space. In *2007 IEEE International Conference on Image Processing*, 2007. 2

[26] Chun-Ta Ho and Ling-Hwei Chen. A fast ellipse/circle detector using geometric symmetry. *Pattern recognition*, 28(1):117–124, 1995. 2

[27] Yiwu Lei and Kok Cheong Wong. Ellipse detection based on symmetry. *Pattern recognition letters*, 20(1):41–47, 1999. 2

[28] Euijin Kim, Miki Haseyama, and Hideo Kitajima. Fast and robust ellipse extraction from complicated images. In *Proceedings of IEEE information technology and applications*. Citeseer, 2002. 2

[29] Lars Libuda, Ingo Grothues, and Karl-Friedrich Kraiss. Ellipse detection in digital image data using geometric features. In *Advances in computer graphics and computer vision*, pages 229–239. Springer, 2007. 2, 7

[30] Fei Mai, YS Hung, Huang Zhong, and WF Sze. A hierarchical approach for fast and robust ellipse extraction. *Pattern Recognition*, 41(8):2512–2524, 2008. 2

[31] Alex Yong-Sang Chia, Susanto Rahardja, Deepu Rajan, and Maylor Karhang Leung. A split and merge based ellipse detector with self-correcting capability. *IEEE Transactions on Image Processing*, 20(7):1991–2006, 2010. 2

[32] Michele Fornaciari, Andrea Prati, and Rita Cucchiara. A fast and effective ellipse detector for embedded vision applications. *Pattern Recognition*, 47(11):3693–3708, 2014. 3, 6, 7

[33] Qi Jia, Xin Fan, Zhongxuan Luo, Lianbo Song, and Tie Qiu. A fast ellipse detector using projective invariant pruning. *IEEE Transactions on Image Processing*, 26(8):3665–3679, 2017. 3, 7

[34] Huixu Dong, Dilip K Prasad, and I-Ming Chen. Accurate detection of ellipses with false detection control at video rates using a gradient analysis. *Pattern Recognition*, 81:112–130, 2018. 3

[35] Viorica Pătrăucean, Pierre Gurdjos, and Rafael Grompone von Gioi. Jointa contrarioellipse and line detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):788–802, 2016. 3

[36] Cai Meng, Zhaoxi Li, Xiangzhi Bai, and Fugen Zhou. Arc adjacency matrix-based fast ellipse detection. *IEEE Transactions on Image Processing*, 29:4406–4420, 2020. 3, 6, 7

[37] Mingyang Zhao, Xiaohong Jia, and Dong-Ming Yan. An occlusion-resistant circle detector using inscribed triangles. *Pattern Recognition*, 109:107588, 2020. 3

[38] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, PAMI-8(6):679–698, 1986. 3

[39] Dilip K Prasad, Chai Quek, Maylor KH Leung, and Siu-Yeung Cho. A parameter independent line fitting method. In *The First Asian Conference on Pattern Recognition*, pages 441–445. IEEE, 2011. 3

[40] Urs Ramer. An iterative procedure for the polygonal approximation of plane curves. *Computer graphics and image processing*, 1(3):244–256, 1972. 3

[41] Andrew Fitzgibbon, Maurizio Pilu, and Robert B Fisher. Direct least square fitting of ellipses. *IEEE Transactions on pattern analysis and machine intelligence*, 21(5):476–480, 1999. 5

[42] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. *CalTech Report*, 03 2007. 6

[43] Mark J. Huiskes and Michael S. Lew. The mir flickr retrieval evaluation. In *MIR '08: Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval*, New York, NY, USA, 2008. ACM. 7

[44] B Russell. Labeme:a database and web-based tool for image annotation. *Ijcv*, 77, 2008. 7

[45] Changsheng Lu, Siyu Xia, Wanming Huang, Ming Shao, and Yun Fu. Circle detection by arc-support line segments. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 76–80. IEEE, 2017. 7

[46] CASIA Iris Image Database. http://biometrics.idealtest.org/. Accessed May 2020. 7