

A Parameterless Line Segment and Elliptical Arc Detector with Enhanced Ellipse Fitting

Viorica Pătrăucean^{1,2}, Pierre Gurdjos¹, and Rafael Grompone von Gioi³

¹ IRIT-ENSEEIH^T*, Toulouse, France

² Military Technical Academy of Bucharest, Romania

³ CMLA, ENS Cachan, France

{viorica.patrucean,pierre.gurdjos}@enseeiht.fr,
grompone@cmla.ens-cachan.fr

Abstract. We propose a combined line segment and elliptical arc detector, which formally guarantees the control of the number of false positives and requires no parameter tuning. The accuracy of the detected elliptical features is improved by using a novel non-iterative ellipse fitting technique, which merges the algebraic distance with the gradient orientation. The performance of the detector is evaluated on computer-generated images and on natural images.

Keywords: ellipse detection, *a contrario* approach, ellipse fitting.

1 Introduction

The automatic detection of elementary geometric features (line segments, elliptical arcs) in images, is as old as computer vision itself [11], being often a prerequisite for high-level procedures, e.g. see [10]. The task is far from trivial and many research works addressed it over the decades. Nonetheless, the problem is not completely settled. Most of the existing detectors, despite the parameter tuning they demand, provide no control of false detections. Theoretical tools that guard against false detections exist in the literature [18,3] and efficient implementations for line segment detection were proposed [3,8]. The purpose of this paper is to put in such a theoretical framework the simultaneous detection of line segments and elliptical arcs, in order to obtain a low-level detector that requires no parameter tuning, and produces reliable results when applied on any kind of image, regardless of its size, content, or source.

Existing geometric feature detectors can be roughly classified into two categories: *Hough-based* and *edge chaining* methods. Most of them rely on an edge detector, e.g. like [1]. The Hough-based algorithms [21] implement variants of the Hough transform (HT) [11] together with some detection thresholds. HT accumulates in an array the votes granted by the edge points to each potential feature. Array points that exceed the threshold become detections. The critical

* The authors wish to thank Géraldine Morin and Vincent Charvillat from IRIT-ENSEEIH^T for their helpful input and comments on this work.

parameters involved are the detection thresholds and the quantisation precision of the accumulator array. Their values –usually empirically tuned– have direct impact on the number of false detections: too permissive detection thresholds or too coarse quantisation step introduce false positives, whilst over-restraining detection thresholds or too fine quantisation step cause false negatives.

A second class of detection methods relies on edge chaining techniques, which use extensively the geometric properties of the sought features, such as straightness criteria for line segments or curvature properties for ellipses [5,20]. Usually, these algorithms begin with a seed pixel (or a group of pixels), and subsequently, other pixels are added, provided they obey some geometric properties of the sought feature. At the end, a line (or an ellipse) is fitted on the gathered pixels, using either some deterministic fitting techniques (e.g. based on least-squares [6]), or randomised robust RANSAC-like approaches [13]. Although efficient in execution time (as opposed to Hough-based methods), these algorithms have difficulties in handling noisy edge maps, reporting an important number of false detections [5,2], caused by inappropriate detection thresholds.

A few works have tried to address the detection thresholds issue. In the Hough class, the *progressive probabilistic* HT (PPHT), introduced by Matas et al., stops the voting procedure for a particular feature when an excess is observed in the accumulator that could not have appeared by accident [14]. The *non-accidentalness* is assessed by using as detection threshold, a cut-off value on the probability that the observed excess occurred by chance in a noise image. This reasoning allows the rejection of false positives. Nonetheless, PPHT lacks scalability: the cut-off values are set for a predefined image size, and the guard against false positives is not ensured when larger images are analysed [8]. Similar reasonings are present in the edge chaining class, based generally on Lowe’s criterion of significance [12,20]. Albeit parameter-free, Lowe’s criterion does not assess the *overall* non-accidentalness, which causes an incomplete guard against false detections. The ellipse detector proposed by Chia et al. [2] tackles this problem by learning detection threshold values on computer-generated images. Although their results are improved comparing to existing works, their detector is still dependent on critical parameters: e.g. the learned detection thresholds favour complete ellipses; they are not able to reject small complete ellipses reported on parasite contours, while falsely rejecting some valid incomplete ellipses.

The formalisation of the false detections control within the feature detection problem was addressed by some authors: e.g. MINPRAN, proposed by Stewart for 3D alignment detection [18], but it still needs a parameter tuning.

A successful parameterless approach was introduced by Desolneux et al. [3], and it is known as the *a contrario* approach. It provides an efficient technique for the automatic computation of the detection thresholds based on what they call Helmholtz’s perception principle; informally, it states that there is no perception in white noise. This comes basically to the non-accidentalness principle, used also by PPHT and Lowe’s criterion. The major difference is that the *a contrario* statistical framework controls the overall number of false detections. Accordingly, the detection thresholds are self-tuning, guarding efficiently against

false positives. Moreover, in the spirit of the parameter-free quest, Desolneux et al.'s approach detects geometric features taking as input the original image, without a previous edge detection step. Using this framework, Grompone von Gioi et al. proposed an efficient parameterless line segment detector (LSD) [8].

The key contribution of our work is to describe a parameter-free combined line segment and elliptical arc detector, based on the *a contrario* approach, which focuses on the control of false detections. The detector obeys a 3-step scheme: candidate selection, candidate validation, and model selection. The latter two are formally sound, being grounded on statistical foundations. Only the candidate selection is heuristic, for efficiency reasons. Our second contribution is related to the candidate selection step, where the proposed detector requires the use of an ellipse fitting operator. We propose a non-iterative fitting technique, which exhibits improved performance on incomplete data, due to the efficient merge of the geometric and photometric information available in images.

The remaining paper is organised as follows. Section 2 details the stages of the proposed feature detector ELSD (Ellipse and Line Segment Detector). Section 3 presents the ellipse fitting technique used by ELSD and compares it with other fitting operators. Section 4 reflects experimentally the efficiency and the robustness of the proposed detector, and section 5 concludes the paper.

2 ELSD – Ellipse and Line Segment Detector

The ELSD algorithm is a three-stage process: (1) first feature candidates are identified using a heuristic; (2) then each candidate has to pass a validation phase. Owing to the multiple families of features addressed, (3) a model selection step is required to choose the best geometric interpretation. The candidate selection step is based on a greedy (heuristic) approach, in the same spirit as in LSD. Ideally, to avoid heuristics at this step, all possible candidates should be considered; but this is impossible in practice, due to time constraints. Therefore we consider a heuristic free of critical parameters and as permissive as possible, which does not alter the quality of the result. In contrast, the validation step is formally sound, being grounded on the probabilistic *a contrario* approach proposed by Desolneux et al. [3], which ensures an efficient guard against false positives. Equally, the model selection step is based on a selection criterion that follows the principles of the model selection theory. Algorithm 1 gives the main steps of ELSD.

2.1 Candidate Selection

Due to the *a contrario* validation, the guard against false positives is ensured. The false negatives, however, are strictly due to an over-restrictive candidate selection procedure, i.e. there exist candidates which do not even get the chance to arrive to the validation phase. Any of the existing feature detectors could carry out this step. But they usually operate on edge maps and no reliable parameterless edge detector is yet available. The main feature of the proposed candidate selection is to be free of critical parameters and as permissive as possible, in order to avoid introducing false negatives.

Algorithm 1. ELSD

```

Input: Gray-level image  $x$ , parameters: none
Output:  $\mathcal{L}_f$  – list of valid features (line segments, circular arcs, elliptical arcs).
1  $grad \leftarrow \text{compute\_gradient}(x)$ ;
2 foreach pixel  $p_i$  in  $x$  do
3    $R \leftarrow \text{region\_grow}(p_i, grad)$ ;
4    $C \leftarrow \text{curve\_grow}(R, grad)$ ;
5    $line \leftarrow \text{fit\_rectangle}(R)$ ;
6    $circle \leftarrow \text{fit\_circular\_ring}(C)$ ;
7    $ellipse \leftarrow \text{fit\_elliptical\_ring}(C)$ ;
8    $(NFA_{line}, NFA_{circle}, NFA_{ellipse}) = NFA(line, circle, ellipse)$ ;
9    $NFA_{min} \leftarrow \min(NFA_{line}, NFA_{circle}, NFA_{ellipse})$ ;
10  if  $NFA_{min} \leq 1$  then
11    | add feature corresponding to  $NFA_{min}$  in  $\mathcal{L}_f$ ;
12  end
13 end

```

For the line segment part, we use the heuristic proposed for LSD [8], i.e. starting from seed pixels, a *region growing* process recursively groups pixels into connected regions sharing the same gradient orientation up to a given precision (see figure 1, first row). For the circular and elliptical arcs, we propose a *curve growing* procedure that adds a second level of grouping compared to the LSD approach, by simply alternating region growing and region chaining operations. The ends of the rectangle given by region growing are used as seed pixels, p_{s_r} and p_{s_l} , for subsequent region growing procedures. In LSD, the region growing chains pixels sharing the same orientation; in ELSD, the curve growing chains neighbour regions produced by region growing, provided they obey some loose, elementary constraints, which characterise the elliptical shapes. Namely, we impose that the contour described by the chained regions be *convex* and *roughly*

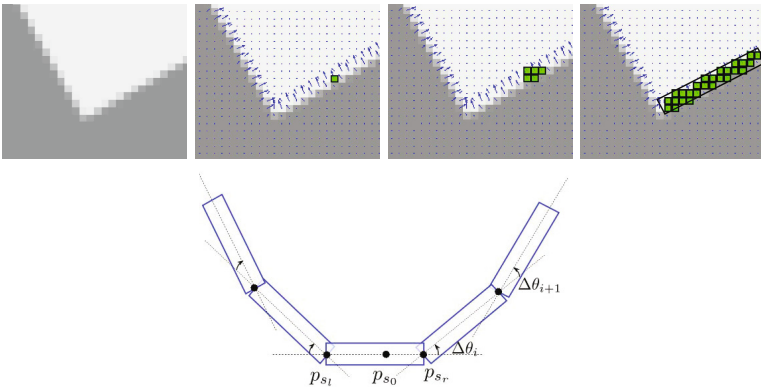


Fig. 1. First row: Region growing. Second row: Curve growing.

smooth (figure 1, second row). The **convexity rule** imposes that consecutive pairs of rectangles turn in the same direction; thus $\Delta\theta_i$ and $\Delta\theta_{i+1}$ have the same sign. The **smoothness rule** is roughly imposed by chaining only regions whose orientations differ by less than $\pi/2$.

The curve growing yields a polygonal approximation of the curve through a recursive scheme. When no more regions can be added, we compute the five parameters of the ellipse that fits the gathered pixels, using the technique detailed in section 3, as well as the three additional parameters of the elliptical ring that covers the pixels, namely the delimiting angles and the ring width. The delimiting angles are computed by sorting the gathered pixels according to their angular coordinate and taking the extremes, whilst the width is computed by summing the distances¹ of the points that are the farthest from the ellipse towards the exterior and the interior respectively. Additionally, a circle is fitted as well, together with the delimiting parameters of the circular ring. This redundancy is justified by the poor accuracy of ellipse fitting techniques when input data are sampled only along small arcs. In these cases, the curve, even if part of an ellipse, can be fairly approximated by a circular arc [17].

At the end of this stage, we are in possession of three candidates: the initial rectangle obtained after the region growing, and the circular and elliptical rings covering the pixels gathered by curve growing.

2.2 Validation

The main idea in the *a contrario* validation technique is to automatically compute the detection thresholds in a way that rejects candidates whose presence might be accidental. We remind the generic validation setup proposed by Desolneux et al. for feature detection [3] and then we state the formulation for circle and ellipse detection.

All along this study, we consider gray-level images of size $m \times n$, defined on a grid $\Gamma = [1, m] \times [1, n] \subset \mathbb{N}^2$, with values in \mathbb{R} , and we call *candidate* any region of pixels $c \subset \Gamma$, that might support a sought geometric feature. We will denote by N the total number of possible candidates of a certain type (namely line segment, circular arc, elliptical arc) in a given image.

Two main ingredients are required within the *a contrario* framework: (i) a measure function giving a score to each candidate, reflecting its *degree of structuredness*, and (ii) a model of unstructured data. A candidate will be validated if it is too structured to be expected in (ii), according to the measure of (i).

For line segment detection, Desolneux et al. give credit to the orientation of the gradient, rather than to its magnitude. Hence, the measure function $k_x(c)$ used to assess the degree of structuredness of the candidate c observed in an image x , is the number of *aligned pixels* that c contains (figure 2, left). A pixel $p \in \Gamma$ is said to be *aligned* with an oriented line segment s up to a precision σ if

$$\text{Angle}(\nabla x(p), \text{dir}_\perp(s)) \leq \sigma\pi, \quad (1)$$

¹ We use Rosin's approximation [17] to compute the Euclidean distance between a point and an ellipse.




Fig. 2. Left: A segment s containing three σ -aligned pixels. Middle and right: Aligned pixel in circle and ellipse case.

where $\nabla x(p)$ is the gradient of the image x at p and $dir_{\perp}(s)$ is the direction orthogonal to s [3].

For the proposed algorithm ELSD, additional definitions need to be stated for circle and ellipse (see figure 2, middle and right). A pixel $p \in \Gamma$ is said to be *aligned* with a circular or elliptical arc a up to a precision σ if

$$Angle(\nabla x(p), dir_{\perp}(tan_a(p))) \leq \sigma\pi, \tag{2}$$

where $\nabla x(p)$ is the gradient of the image x at p and $dir_{\perp}(tan_a(p))$ is the direction orthogonal to the tangent $tan_a(p)$ to the circle or ellipse in p .

With this choice for assessing the degree of structuredness, a convenient unstructured model, that will be denoted as the *a contrario* \mathcal{H}_0 model, is one in which gradient orientations are i.i.d. random variables, i.e.:

1. $\forall p \in \Gamma, Angle(\nabla x(p))$ is uniformly distributed over $[0, 2\pi]$;
2. $\{Angle(\nabla x(p))\}_{p \in \Gamma}$ are independent random variables.

It has been proven by Desolneux et al. that these two assertions hold under certain conditions of subsampling if x is a Gaussian white noise image [3, p. 67]. From now on, x will refer to the analysed image and X to an (unstructured) Gaussian white noise image of the same size as x , drawn from the \mathcal{H}_0 model.

The *a contrario* model proposed by Desolneux et al. for line segment detection can be seamlessly applied for ellipse detection as well. Note that \mathcal{H}_0 is not meant to model the real noise of images; instead, it is a simple and effective model for unstructured, isotropic zones of the image, where no aligned structures are perceived. This model was thoroughly evaluated in [3] and a convincing agreement was obtained between results of line segment detectors grounded on this model and human perception.

The following *a contrario* definition, used jointly with a result from probability theory (Proposition 1), completes the *a contrario* validation setup.

Definition 1. Let $\mathcal{Z} = \{Z_1, \dots, Z_{N_{\mathcal{Z}}}\}$ be a set of $N_{\mathcal{Z}}$ integer random variables. Observing a value z_i for Z_i is an ε -meaningful event in \mathcal{Z} if its associated number of false alarms, defined by $NFA = N_{\mathcal{Z}}\mathbb{P}[Z_i \geq z_i]$, is less than or equal to ε .

Proposition 1. The expected number of ε -meaningful events in \mathcal{Z} is less than or equal to ε .

Proof. We define the set of thresholds $\kappa_i = \min\{z \in \mathbb{N} \mid N_{\mathcal{Z}}\mathbb{P}[Z_i \geq z] \leq \varepsilon\}$. With this definition, an observed value z_i is ε -meaningful if and only if $z_i \geq \kappa_i$. Then, the subset of elementary ε -meaningful events associated to Z_i is $\{Z_i \geq \kappa_i\}$, and the expected overall number of ε -meaningful events in \mathcal{Z} is given by

$$\mathbb{E} \left[\sum_{i=1}^{N_{\mathcal{Z}}} \mathbb{1}_{\{Z_i \geq \kappa_i\}} \right] = \sum_{i=1}^{N_{\mathcal{Z}}} \mathbb{P}[Z_i \geq \kappa_i] \leq \sum_{i=1}^{N_{\mathcal{Z}}} \frac{\varepsilon}{N_{\mathcal{Z}}} = \varepsilon. \quad \blacksquare$$

In our case, we deal with the events of observing in image x candidates c_i , containing $k_x(c_i)$ aligned pixels. A candidate c_i is considered a valid detection and said to be ε -meaningful, in agreement with Definition 1, when its associated number of false alarms $\text{NFA}(c_i, x) = NP[k_X(c_i) \geq k_x(c_i)]$ satisfies the *validation test*:

$$\text{NFA}(c_i, x) \leq \varepsilon. \quad (3)$$

Thus, the set \mathcal{Z} corresponds to the N random variables $k_X(c_i)$ and NFA reflects the probability of observing in a random image X (of the same size as x) candidates at least as structured as the analysed one. The smaller the NFA, the more unlikely is c_i to appear in an image X drawn from \mathcal{H}_0 ; thus, it is meaningful.

Thanks to Proposition 1, the expected number of candidates in an unstructured image X with *at least* the same degree of structuredness as in the analysed image x is ensured to be (statistically speaking) less than ε . In other words, *if a candidate is accepted as valid detection when the equation (3) holds, then the number of accidental detections, which are in fact false positives, is guaranteed to be less than the chosen ε .* This result reflects precisely the formalisation of Helmholtz's perception principle, which states that in an unstructured image, only a very small number of detections should be reported.

In practice, ε can be set as small as desired. As it is proven in [3, p. 77], the threshold values $\kappa_i = \min\{k \in \mathbb{N} \mid NP[k_X(c_i) \geq k] \leq \varepsilon\}$, such that the candidates c_i are ε -meaningful when $k_x(c_i) \geq \kappa_i$ (using a similar reasoning as in the above proof), have a logarithmic dependence on ε ; thus a simple, convenient value can be assigned to ε . The value $\varepsilon = 1$ yielded satisfactory results [8], so it will be kept for the proposed detector as well. With this choice, we assume the risk of accepting on average one false positive per image. Considering the weak dependence on the *unique* detection threshold ε fixed to 1, the approach can be considered as parameterless.

If the angular precision for declaring a pixel as aligned is $\pi\sigma$, then the probability for a pixel to be aligned under the *a contrario* model is $\frac{2\pi\sigma}{2\pi} = \sigma$. Since the gradient orientations are independent under the chosen model, $k_X(c)$ follows a binomial law with parameters $l(c)$, $k_x(c)$, and σ , where $l(c)$ is the total number of pixels in the rectangle or ring. Therefore, we can write $\text{NFA}(c, x) = N\mathcal{B}(l(c), k_x(c), \sigma)$, with $\mathcal{B}(l, k, \sigma) = \sum_{i=k}^l \binom{l}{i} \sigma^i (1-\sigma)^{l-i}$ being the binomial tail. For σ we set the value $1/8$, which proved to be satisfactory in practice [8].

Considering a one pixel precision, in the line segment case, the number of potential candidates in an $m \times n$ image is $N_{\text{line}} = (mn)^{5/2}$, as each rectangular region has five degrees of freedom: centre (2), length (1), orientation (1), and




Fig. 3. Degrees of freedom for line segment, circular arc, and elliptical arc

width (1); see figure 3 left. We count $(mn)^{1/2}$ possible values for each degree of freedom. For circular arcs we have $N_{circle} = (mn)^3$ for six degrees of freedom: centre (2), radius (1), width (1), and delimiting angles (2); see figure 3 middle. Finally, elliptical arcs have eight degrees of freedom: centre (2), axes (2), orientation (1), width (1), and delimiting angles (2), thus $N_{ellipse} = (mn)^4$; see figure 3 right. These estimates count roughly all the possible combinations of pixels that could support a feature in the given image and they allow the detection thresholds to adapt to the image size. But in practice, only the candidates proposed by the candidate selection step will be analysed, for efficiency reasons.

All-in-all, with the chosen value for ε and the estimates for the number of candidates, the tests that the feature candidates must pass in order to be considered valid detections are:

$$\begin{aligned}
 \text{NFA}_{line} &= (mn)^{5/2} \mathcal{B}(l, k, \sigma) \leq 1 \text{ for line segments,} \\
 \text{NFA}_{circle} &= (mn)^3 \mathcal{B}(l, k, \sigma) \leq 1 \text{ for circular arcs,} \\
 \text{NFA}_{ellipse} &= (mn)^4 \mathcal{B}(l, k, \sigma) \leq 1 \text{ for elliptical arcs.}
 \end{aligned} \tag{4}$$

2.3 Model Selection

The candidates declared meaningful in the validation step compete subsequently in a *model selection* phase for the best interpretation for the given data and the winner is kept as final valid detection. The model selection theory is a central subject in statistical learning and an important number of model selection criteria is available in the literature [4]. The proposed detector ELSD entails a model selection step within a linear regression problem: given a set of pixels in an image, and three fitted models (line, circle, ellipse), decide which model is the most suitable to explain the data. We choose to use the NFA of a candidate as model selection criterion, i.e. the candidate possessing the smallest NFA is considered as most meaningful, and kept as final valid detection. This idea was suggested by Desolneux et al. [3, p. 245], but not carried out. The pertinence of this usage needs a thorough discussion. Due to space limitation, here we only point out that, qualitatively, NFA follows the *Ockham's razor principle* which guides the model selection theory, namely it contains a term illustrating the goodness of fit (the binomial tail) and a term penalising for complexity (the number of candidates N , proportional with the number of free parameters). We refer the reader to [16] for more details on this topic.

3 Conic Fitting Using Gradient Orientations

The candidate selection of ELSD relies on a conic fitting technique, restricted to the cases of circle and ellipse. Addressing circular and elliptical arcs, we seek for a closed-form solution that performs well on those being small or incomplete. Let \mathcal{C} be a general conic represented in homogeneous matrix form by:

$$\tilde{\mathbf{p}}^\top \mathbf{C}_\theta \tilde{\mathbf{p}} = 0, \quad (5)$$

where \mathbf{C}_θ is the order-3 symmetrical matrix of conic (homogeneous) coefficients, stacked in the vector $\boldsymbol{\theta}_{6 \times 1}$ (six coefficients with five degrees of freedom), and $\tilde{\mathbf{p}} = (x \ y \ 1)^\top$ is the vector of augmented Cartesian coordinates of a point.

Given n points, we consider the problem of fitting to the data an ellipse modeled by an implicit equation of the form (5). The fitting comes to minimising in a least-squares sense the objective function $F = \sum_{i=1}^n \delta^2(\boldsymbol{\theta}, \tilde{\mathbf{p}}_i)$, subject to $h(\boldsymbol{\theta}) = 0$, where $\delta(\boldsymbol{\theta}, \tilde{\mathbf{p}}_i) = \tilde{\mathbf{p}}_i^\top \mathbf{C}_\theta \tilde{\mathbf{p}}_i$ is the *error-of-fit* between the input points and the estimated conic, and $h(\cdot)$ models different constraints that need to be imposed on the conic coefficients to avoid the trivial solution or/and to ensure ellipticity [6]. Several authors proposed weighting schemes to improve the performance on incomplete data, where this technique has the tendency to grossly underestimate the eccentricity [19].

A few works have tried to put forward the usefulness of the derivative constraint entailed by the gradient orientation. Förstner and Gulch use exclusively the gradient orientation to accurately locate the centre of a circle, which corresponds to the intersection point (computed in a least-squares sense) of the lines supporting the gradient vectors of the pixels situated on the contour of the circle. This is due to the fact that these lines all converge towards the centre of the circular shape [7]. The radius can be computed subsequently using the positional constraints. Guennebaud and Gross applied Förstner's idea in 3D for sphere fitting, in a computer graphics application [9]. In the ellipse case, this idea cannot be seamlessly applied, as the lines supporting the gradient vectors of the pixels situated on the ellipse contour do not converge towards the ellipse centre. To overcome this, Ouellet and Hebert perform the ellipse fitting in the dual space [15], but their formulation omits useful information, namely the positional constraints. All these methods improve the fitting results while remaining computationally efficient, but they still lack accuracy on incomplete data.

We suggest that the simultaneous usage of the positional and tangential information can further improve the results. Accordingly, the fitting problem writes as: $\min_{\boldsymbol{\theta}} \sum_{i=1}^n (\delta^2(\boldsymbol{\theta}, \tilde{\mathbf{p}}_i) + \gamma^2(\boldsymbol{\theta}, \tilde{\mathbf{p}}_i, \mathbf{g}_i))$, subject to $h(\boldsymbol{\theta}) = 0$, where $\gamma(\cdot)$ returns the error-of-fit on the gradient orientation, and \mathbf{g}_i denotes the 2-vector of the image gradient at (x_i, y_i) . To obtain the expression of $\gamma(\cdot)$, we make use of the *pole-polar duality*, given by $\mathbf{C}\tilde{\mathbf{p}} \sim \mathbf{l}$, where \sim denotes the projective equality [10]. The point with vector $\tilde{\mathbf{p}}$ is called the pole of the line \mathbf{l} w.r.t. the conic \mathcal{C} , and conversely, the line with vector \mathbf{l} is the polar of $\tilde{\mathbf{p}}$. Given a point $\tilde{\mathbf{p}}$ belonging to \mathcal{C} and with gradient vector $\mathbf{g} = (g_1, g_2)^\top$, its polar \mathbf{l} is tangent to \mathcal{C} at $\tilde{\mathbf{p}}$ and is orthogonal to any line with direction \mathbf{g} . Thus, the point at infinity (denoted by

$\tilde{\mathbf{g}}_\infty^\perp$) associated to the direction orthogonal to \mathbf{g} lies on \mathbf{l} , that is: $\mathbf{l}^\top \tilde{\mathbf{g}}_\infty^\perp = 0$. If one assumes that the camera has square pixels, then $\tilde{\mathbf{g}}_\infty^\perp = (-g_2, g_1, 0)^\top$. Using the pole-polar duality, we obtain a new equation, linear in the elements of \mathbf{C}_θ :

$$\tilde{\mathbf{p}}^\top \mathbf{C}_\theta \tilde{\mathbf{g}}_\infty^\perp = 0. \quad (6)$$

Together with the positional constraints (5), the linear equation system becomes:

$$\begin{cases} \tilde{\mathbf{p}}_i^\top \mathbf{C}_\theta \tilde{\mathbf{p}}_i = 0 \\ \tilde{\mathbf{p}}_i^\top \mathbf{C}_\theta \tilde{\mathbf{g}}_\infty^\perp = 0 \end{cases} \quad i = \{1, \dots, n\}. \quad (7)$$

This way, each point contributes with two independent equations. This technique can be seamlessly applied to both circle and ellipse fitting. To improve the numerical stability, the input data are first normalised as described in [10].




Fig. 4. Mean relative error on the estimated radius for 200 tests performed on real images of printed drawn circles. Input data are sampled along arcs of ~ 45 degrees.

To show the improvement obtained by using simultaneously the positional and the tangential constraints, we compared the proposed method with Taubin’s circle fitting operator [19], which uses only positional constraints, and with Guennebaud and Gross’s method [9] mentioned above (denoted GG in figure 4). The results of the three methods on incomplete data are shown in figure 4. Table 1 presents a comparison with Ouellet and Hebert’s method [15], in fitting ellipses on incomplete data. This comparative study shows that when input data are sampled along incomplete conics, and especially when the features are small, it is crucial to exploit simultaneously all the information they possess.

Table 1. Mean(Max) error on the estimated centre, for 500 tests performed on computer-generated images; input points are sampled along incomplete ellipses generated randomly, with axes values in the range 10 – 100 pixels

Method	75% of contour	50% of contour
Proposed	0.1136(0.4465)	0.5972(4.8615)
Ouellet	0.2018(1.0942)	3.4944(8.3421)

4 Results

The proposed detector was extensively tested on various (computer-generated and natural) images. Its performance is compared with relevant feature detectors available online: Etemadi's detector [5], which addresses the simultaneous detection of line segments and circular arcs, and Hough detectors for circles (HC) [21] and ellipses (HE) [21], respectively. All detectors were tested with their default parameters. The edge maps were obtained using Canny's edge detector available in Matlab, with default parameters. The tests carried out on computer-generated images of overlapping and non-overlapping geometric shapes, analysed the robustness against noise, and the precision at different scales. Figure 5 illustrates the typical behaviour of the above detectors when applied on these test images. Generally, Etemadi's detector is accurate in reporting the correct detections (line segments and circular arcs), but has no ability in eliminating false positives, reported on parasite contours. Equally, its precision at reduced scale is poor. The two Hough-based detectors behave reasonably well as long as the images are noise-free. When noise is present, they encounter serious difficulties and the results are not satisfactory: HC reports an important number of false positives, whereas HE reports no detection. The proposed detector gives reasonable results in all these cases: it is accurate when applied on noise-free images, even at




Fig. 5. Results on computer-generated images of geometric shapes (square, circle, ellipse) for different noise types and different scales. From top to bottom: original image, ELSD, Etemadi, HC, HE. From left to right: noise-free image, image altered by Gaussian noise, overlapping shapes, overlapping shapes altered by salt-and-pepper noise, overlapping shapes-normal scale, overlapping shapes-reduced scale.

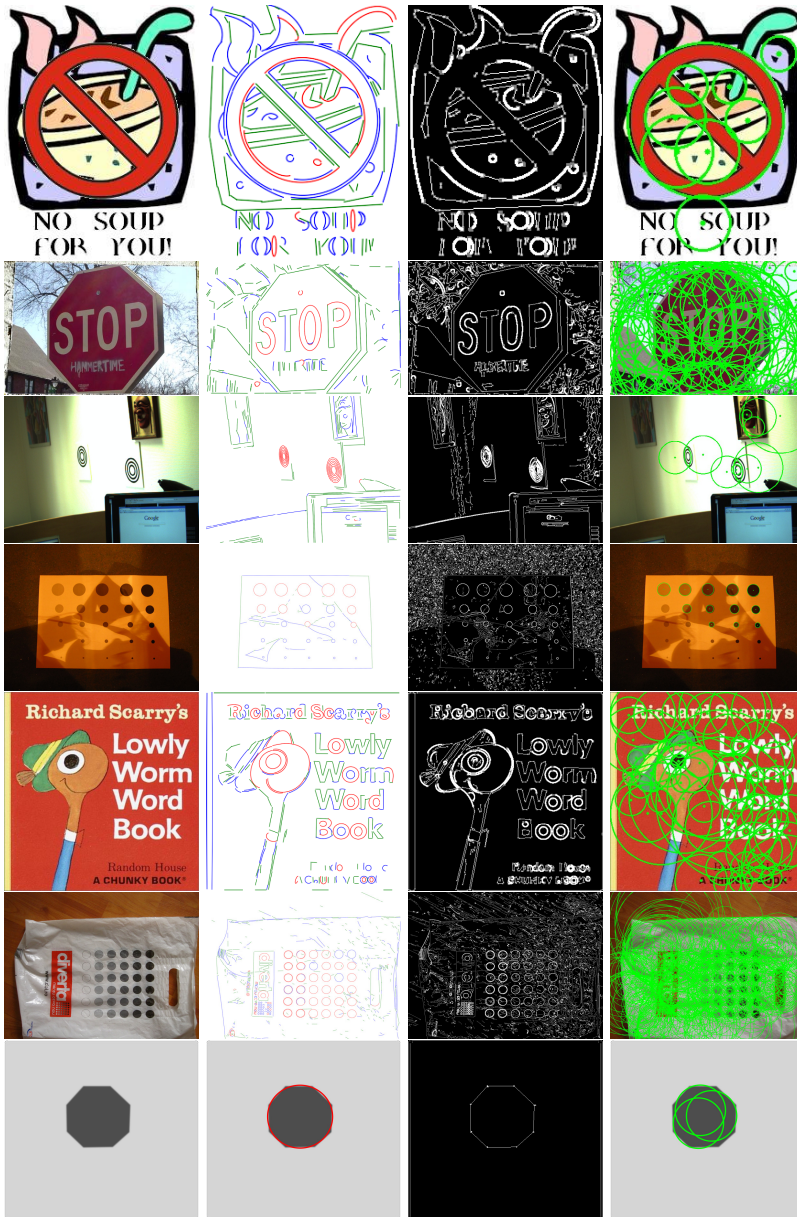


Fig. 6. Results on natural images. 1st column: original image, 2nd column: ELSD, 3rd column: Etemadi, 4th column: HC. HE reported no detection. Image size and ELSD's execution time on a regular machine for each example from top to bottom: 213×279 0.3s, 445×304 1s, 640×480 0.4s, 1600×1200 1.3s, 442×450 0.6s, 1600×1200 3.1s, 612×563 0.1s.

reduced scale, due to the enhanced conic fitting used; the noise (especially salt-and-pepper) provokes an over-segmentation of contours, but no false positive is reported, which reflects the validity of ELSD's theoretical grounds.

In order to assess the behaviour of the proposed detector in real-world applications, we have extensively tested it on natural images. Unfortunately, the literature lacks an appropriate benchmark with labeled images for evaluating geometric feature detectors. We downloaded one thousand images from flickr and used them as test images. Figure 6 gives some relevant results. The source code and an online demo of ELSD, where users can upload images and test the detector, can be found at <http://ubee.enseeiht.fr/vision/ELSD/>. A full description of the detector is given in [16]. The current limitations of ELSD are shown in figure 5, 4th column, and figure 6, last row. For the former case, a merging step should be implemented to overcome the contour over-segmentation due to noise [2]. The latter case illustrates the behaviour of the model selection step, when we deal with polygonal shapes that can be fairly approximated by curves. To improve the accuracy, polygonal candidates should be considered as well.

5 Conclusion

We proposed a parameterless line segment and elliptical arc detector, grounded on the *a contrario* approach, which controls the number of false positives. Addressing multiple feature families (line segments, circular and elliptical arcs), the proposed detector entails a model selection step. The essential quantity of the *a contrario* approach –NFA– is used for both validation and model selection purposes. To improve the accuracy of occluded shapes detection, a non-iterative circle and ellipse fitting technique was introduced; it uses simultaneously the positional and the tangential information of the image pixels. Future work will concentrate on the merging of elliptical arcs to obtain precise ellipses in case of fragmented contours, and on the model selection phase, to improve its accuracy.

References

1. Canny, J.: A computational approach to edge detection. PAMI 8, 679–698 (1986)
2. Chia, A., Rahardja, S., Rajan, D., Leung, M.K.: A split and merge based ellipse detector with self-correcting capability. TIP 20, 1991–2006 (2011)
3. Desolneux, A., Moisan, L., Morel, J.M.: From Gestalt Theory to Image Analysis: A Probabilistic Approach. Springer (2008)
4. Duda, R., Hart, P., Stork, D.: Pattern classification, 2nd edn. John Wiley & Sons (2001)
5. Etemadi, A.: Robust segmentation of edge data. In: Int. Conf. on Image Processing and its Applications, pp. 311–314 (1992)
6. Fitzgibbon, A., Pilu, M., Fisher, R.: Direct least-squares fitting of ellipses. PAMI 21, 476–480 (1999)
7. Forstner, W., Gulch, E.: A fast operator for detection and precise location of distinct points, corners and centres of circular features. In: Intercom. Conf. on Fast Processing of Photogrammetric Data, pp. 281–305 (1987)

8. Grompone von Gioi, R., Jakubowicz, J., Morel, J.M., Randall, G.: LSD: A fast line segment detector with a false detection control. *PAMI* 32, 722–732 (2010)
9. Guennebaud, G., Gross, M.: Algebraic point set surfaces. *ACM Trans. Graph.* 26, 231–239 (2007)
10. Hartley, R.I., Zisserman, A.: *Multiple View Geometry in Computer Vision*. 2nd edn. Cambridge University Press (2004)
11. Hough, P.: Method and means for recognizing complex patterns. U.S. Patent 3069654 (1962)
12. Lowe, D.G.: *Perceptual Organization and Visual Recognition*. Kluwer Academic Publishers (1985)
13. Mai, F., Hung, Y.S., Zhong, H., Sze, W.F.: A hierarchical approach for fast and robust ellipse extraction. *Pattern Recognition* 41, 2512–2524 (2008)
14. Matas, J., Galambos, C., Kittler, J.: Progressive probabilistic hough transform. In: *BMVC*, pp. 256–265 (1998)
15. Ouellet, J., Hebert, P.: Precise ellipse estimation without contour point extraction. *Machine Vision and Applications* 21, 59–67 (2009)
16. Pătrăucean, V.: Detection and identification of elliptical structure arrangements in images: Theory and algorithms. PhD thesis. University of Toulouse, France, <http://ethesis.inp-toulouse.fr/archive/00001847/>
17. Rosin, P.L.: Ellipse fitting using orthogonal hyperbolae and Stirling’s oval. *Graphical Models and Image Processing* 60, 209–213 (1998)
18. Stewart, C.V.: MINPRAN: A new robust estimator for computer vision. *PAMI* 17, 925–938 (1995)
19. Taubin, G.: Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *PAMI* 13, 1115–1138 (1991)
20. West, G., Rosin, P.: Multi-stage combined ellipse and line detection. In: *BMVC*, pp. 197–206 (1992)
21. Xu, L., Oja, E., Kultanen, P.: A new curve detection method: Randomized Hough Transform. *Pattern Recogn. Lett.* 11, 331–338 (1990)