

A Simple Operator for Very Precise Estimation of Ellipses

Jean-Nicolas Ouellet Patrick Hébert
Computer Vision and Systems Laboratory
Laval University
Quebec, Canada
{jouellet,hebert}@gel.ulaval.ca

Abstract

This paper presents a simple linear operator that accurately estimates the position and parameters of ellipse features. Based on the dual conic model, the operator avoids the intermediate stage of precisely extracting individual edge points by exploiting directly the raw gradient information in the neighborhood of an ellipse's boundary. Moreover, under the dual representation, the dual conic can easily be constrained to a dual ellipse when minimizing the algebraic distance. The new operator is assessed and compared to other estimation approaches in simulation as well as in real situation experiments and shows better accuracy than the best approaches, including those limited to the center position.

1. Introduction

Precise estimation of an ellipse in an image usually consists in accurately extracting contour points with subpixel precision before fitting ellipse parameters on the obtained set of points. Numerous methods have been proposed for fitting an ellipse from a given set of contour points [2–4, 6, 12, 14]. They differ from each other depending on their precision, accuracy, robustness to outliers or on their ability to fit parameters on partial ellipse contours. All these methods rely on a set of contour points that are extracted beforehand. Nevertheless, extracting contour points usually subsumes multiple stages including gradient estimation, non-maximum suppression, thresholding, and subpixel estimation.

Extracting contour points imposes making a decision for each of those points based on neighboring pixels in the image. The following question thus arises. Is it possible to develop a more direct method that precludes the extraction of a set of contour points? One would directly exploit the information encoded in all pixels in the ellipse neighborhood. By eliminating precise contour point extraction, the method




Figure 1. (a) Image of a calibration target. The marked region indicates the pixels contributing to the ellipse estimation. (b) Close-up view of the marked region. (c) Close-up view of the region identified in (b) where the image gradient orientation and magnitude are represented by arrows superimposed on the image. A dotted line depicts the fitted ellipse contour.

would be greatly simplified and the uncertainty on the recovered ellipse parameters will be assessed more closely to the image data. Another motivation for developing such an approach is the possibility to process low contrast images where each contour point can hardly be extracted along the ellipse.

Formally, given a set of pixels describing a complete ellipse in a real image, what are the ellipse parameters? Although the limitation to complete ellipses might appear restrictive, it is of great interest for camera calibration and precise measurement applications where it is common to use physical circular targets that project into ellipses in an image [10]. In these applications, partial ellipses lead to a less reliable estimation and are generally avoided. In Fig. 1(a), a rectangular area enclosing an ellipse is shown. We develop a method that can fit the ellipse's parameters from all the pixels that lie inside such an area or preferably nearby the ellipse's contour. More precisely, we will show how one can estimate the ellipse's parameters directly from the gradient vector field. Fig. 1(c) is a close up view where the

gradient vector field as well as the fitted ellipse contour are superimposed on the image.

The most relevant work to ours is that of Forstner who developed an operator for extracting circular and corner features in an image [5]. In this extraction process, detection is decoupled from localization. For localizing features, the detection of individual contour points is avoided by directly processing the gradient components of the image in the neighborhood of a detected feature. In this paper, we show how it is possible to extend the basic principle for the estimation of ellipse parameters. For that purpose, we will exploit the dual conic model and we will show that one can easily constrain the dual conic to a dual ellipse. The resulting linear algorithm is of remarkable simplicity and provides results as precise and accurate as the best methods. It also makes it possible to fit ellipses in images where individual contour points would be unreliable.

A review of widely used methods for estimating ellipse parameters, including those limited to center position estimation, is presented in the next section. Then, Forstner's original method is described before developing the new operator in Section 3. Finally, a series of comparative experiments for assessing the relative accuracy and precision in presence of noise and blur is presented in Section 4.

2. Related work

An ellipse can be estimated from the image intensity or its gradient. Intensity based methods are also referred to as *direct methods* [16] when they directly exploit the image intensity with no transformation or derivation. One simple method which returns the coordinates of the feature is the gray-scale centroid; for a given region, the intensity of each pixel is used to weigh its contribution to the estimation of the intensity center of mass. Integrating over all pixels in the region provides good immunity to noise. However, intensity based methods like centroid estimation are particularly sensitive to non-uniform illumination. The causes of non-uniform illumination are frequent and various, e.g. light sources, vignetting or reflections on specular surfaces.

Gradient based methods are less affected by non-uniform illumination. Nevertheless, the derivation process involved makes them more sensitive to noise. Generally, these latter methods are composed of two distinct steps; the edges of a region are first identified after non-maxima suppression and a conic is fitted to the edge points. Methods to fit a conic to points are abundant with complexity ranging from linear least-square estimation to iterative procedures with renormalization [12]. The simplest method, linear least-square, minimizes the algebraic error associated with the points. Since the linear system is homogeneous, a constraint must be imposed on the conic parameters $\Theta = [A, B, C, D, E, F]^T$ to avoid the trivial solution

$\Theta = \mathbf{0}$. Many possible constraints were studied in order to estimate the best conic from a set of points, e.g. $F = 1$, $\|\Theta\| = 1$ (see [4, 14] for a review). Nevertheless, it is essential for a conic estimator to be invariant to Euclidean transformations (translation, rotation, scale) [1]. More precisely, estimating a conic from a transformed set of points should be exactly equivalent to transforming the conic estimated from the untransformed set of points. According to this statement, the constraints $\|\Theta\| = 1$ and $F = 1$ are not invariant. On the other hand, the expression $4AC - B^2$ is a conic invariant and can be formulated as the quadratic constraint $4AC - B^2 = 1$. Under this constraint, the parameters minimizing the algebraic error, which are obtained via a generalized eigensystem, describe elliptical shapes [3]. However, in the presence of partial ellipses, this method, and all other methods based on the algebraic error, will underestimate the eccentricity of the ellipse [18]. Different iterative algorithms were introduced to reduce or compensate the bias [2, 12]. Still, given that ellipses are complete, the effect of the bias is practically nonexistent and there are no significant differences among high quality fitting methods [4]. In this situation, it is advantageous to employ the simplest method.

Estimating a conic from pixels' centers identified as local maxima is generally inaccurate. A better estimation is obtained from the edge points whose positions are known with subpixel precision. A simple way to achieve subpixel precision is to perform a parabolic interpolation of the gradient maximum along the edge direction. The solution is direct but the interpolation is sensitive to noise. An alternative is proposed in [17] where an ideal step edge is fitted to the image intensity value such that the first three sample moments are preserved. The fitting is done over a circular region centered at the gradient local-maximum which provides a larger fitting area, thus greatly reducing the sensitivity to noise. However, the location can be biased if the edge is curved inside the region. To reduce the bias in the presence of curved edges, a correction term based on the local curvature of the contour was proposed in [15] and later refined in [9]. This leads to even more complex methods. By contrast, a method directly exploiting the gradient in the region encompassing an ellipse could greatly simplify the location process.

A radically different approach was proposed by Forstner [5] for the localization of corners and most importantly, of circular features. In his approach, individual edge detection is avoided; the feature's location is estimated directly from the image gradient calculated in a region encompassing the feature. Actually, the gradient at a pixel in the neighborhood of a feature edge provides the orientation of a 2D line passing through the pixel center. The location of a feature is obtained by estimating the intersection point of a set of these lines in the feature's neighborhood. The operator can




Figure 2. (a) Lines parallel to the circle gradient vectors intersect at the circle center. (b) Lines parallel to the ellipse gradient vectors are not convergent. (c) Lines perpendicular to the ellipse gradient vectors are tangent to the ellipse.

be adapted to locate corners, where the lines are oriented perpendicularly to the gradient and intersect at the corner location, or circle centers in which case the lines are oriented parallel to the gradient as shown in Fig. 2(a). The intersection point is estimated by means of least-squares where each line contributes with a weight set proportional to the squared magnitude of the gradient [5].

Although very simple and fast, the Forstner operator is not adapted to locate elliptical feature centers; the lines defined by the gradient on the boundary of an ellipse do not converge as shown in Fig. 2(b). Since ellipses are more general and frequent in images, it is interesting to develop a new operator inspired by the same principle. Another motivation for improving the Forstner operator is to reduce the effect of the gradient orientation error hampering the operator. Actually, noise causes small angular errors in the gradient orientation. For circles, these errors cause a leverage effect due to the distance between the edges and the intersection of the lines. A small orientation error for an edge that is located farther from the estimated line intersection will produce a greater error than a closer edge. In the next section, we propose a new operator that circumvent these limitations.

3. Description of the new operator

3.1. Background on conic and dual conic

A conic is a planar curve formed by the intersection of a plane with a circular cone. Generally the term conic refers to a point conic, defining an equation on points. In projective geometry the role of homogeneous points and lines can be interchanged; this is known as the duality relationship. As a result, there is a dual conic (see Fig. 2(c)) which defines an equation on lines instead of points [8]. More precisely, a point, $x_i = [u_i, v_i, 1]^T$, in homogeneous coordi-

$$\text{nates lies on the conic } \mathbf{C} = \begin{bmatrix} A & B/2 & D/2 \\ B/2 & C & E/2 \\ D/2 & E/2 & F \end{bmatrix} \text{ iff} \\ x_i^T \mathbf{C} x_i = 0. \quad (1)$$

In a similar way a line $l_i = [a_i, b_i, c_i]^T$ is tangent to the dual of conic \mathbf{C} iff

$$l_i^T \mathbf{C}^* l_i = 0, \quad (2)$$

where \mathbf{C}^* is the inverse matrix of \mathbf{C} , given full rank matrix. It is interesting to note that the dual conic center $e_c = (u_c, v_c)$ is obtained directly from the parameters. Actually, if the pole of a line l is defined as a point x such that $x = \mathbf{C}^* l$, then the center of the dual conic is the pole of the line at infinity [11],

$$\lambda [u_c, v_c, 1]^T = \left[\frac{D^*}{2}, \frac{E^*}{2}, F^* \right]^T = \mathbf{C}^* l_\infty. \quad (3)$$

Given a set of lines l_i , the vector of parameters $\Theta = \{A^*, B^*, C^*, D^*, E^*, F^*\}$ of the dual conic \mathbf{C}^* can be estimated by linear least-squares. The estimation is accomplished by finding Θ for which $\Phi(\Theta)$ reaches a minimum,

$$\Phi(\Theta) = \sum_{i \in R} \omega_i (l_i^T \mathbf{C}^*(\Theta) l_i)^2, \quad (4)$$

where R contains the set of lines contributing to the dual conic estimate and ω_i is a weighting factor. Since l_i is a triplet of homogeneous coordinates, the scale of the lines is indeterminate. It can be fixed such that $\|a, b\| = 1$. Normal equations derived from Eq. (4) are linear in Θ and lead to the following form:

$$\left[\sum_{i \in R} \omega_i^2 K_i K_i^T \right] [\Theta] = \mathbf{0}, \quad (5)$$

where K_i is composed of line coefficients such that $K_i = [a_i^2, a_i b_i, b_i^2, a_i c_i, b_i c_i, c_i^2]^T$. This system can be solved under the constraint $\|\Theta\| = 1$ using the singular value decomposition (SVD). However, this constraint is not invariant to Euclidean transformations [1]. A more appropriate constraint is introduced in the next section.

3.2. From dual conic to dual ellipse

The conic discriminant can be imposed as $4AC - B^2 = 1$ as proposed in [3] to obtain an ellipse specific fitting. From the duality relationship between points and lines, we know that a similar constraint can also be derived for the dual conic. Since a conic and its dual are related by the matrix inverse operation, the dual conic parameter F^* can be expressed in terms of the parameters of the conic \mathbf{C} :

$$F^* = \frac{1}{4|\mathbf{C}|} (4AC - B^2) = \frac{1}{|\mathbf{C}|} \begin{vmatrix} A & B/2 \\ B/2 & C \end{vmatrix}. \quad (6)$$

In this expression, the factor $\frac{1}{4|\mathbf{C}|}$ can be canceled out since \mathbf{C}^* is known up to a scale factor. Hence, the conic discriminant $4AC - B^2$, which is strictly positive for ellipses [3], is intrinsically related to the term F^* . From Eq. 6, the dual ellipse condition becomes $F^* > 0$. If the scaling of \mathbf{C}^* is incorporated into the condition, one obtains the dual ellipse fitting constraint $F^* = 1$. The normal equations of the new system of five unknowns $\Theta' = (A'^*, B'^*, C'^*, D'^*, E'^*)$ hold the following matrix form:

$$\left[\sum_{i \in R} \omega_i^2 K_i' K_i'^T \right] [\Theta'] = \sum_{i \in R} -\omega_i^2 K_i' c_i^2, \quad (7)$$

where K' is composed of the first five elements of K .

Although invariant to Euclidean transformations, the system could be sensitive to numerical error. Improving the conditioning of the equation system by means of data normalization will benefit the estimation process [7]. Line normalization is done by first shifting the origin of the coordinate system to an approximation of the center of the ellipse area and then scaling the axes such that the mean distance of the lines to the origin becomes $\sqrt{2}$.

3.3. The Localization Operator

The image gradient near a contour naturally provides the orientation of the contour normal. Therefore, given a set of pixels belonging to a neighborhood enclosing an ellipse, a set of lines can be constructed to solve for Eq. (7). More precisely, when the image gradient $\nabla I_i = [I_{ui}, I_{vi}]^T$ at a pixel x_i is not null, it defines the normal orientation of a line passing through the pixel center such that:

$$l_i = \left[\frac{I_{ui}}{\|\nabla I_i\|}, \frac{I_{vi}}{\|\nabla I_i\|}, -\frac{\nabla I_i^T x_i}{\|\nabla I_i\|} \right]^T. \quad (8)$$

To improve robustness to noise on gradient orientation in nearly constant area, the weights ω_i in Eq. (7) are set equal to the squared gradient magnitude [5],

$$\omega_i = \|\nabla I(x_i)\|^2. \quad (9)$$

If needed, the conic \mathbf{C} can then be obtained from the inverse of \mathbf{C}^* . The resulting operator requires very few steps and computations. Actually, each pixel of the region contributes to the construction of a 5x5 matrix in Eq. 7 which is then inverted.

Once the parameters are obtained, one can further estimate the variance-covariance matrix of the ellipse parameters which is particularly meaningful. The 2x2 submatrix including D^* , E^* and their covariance term, directly provides the uncertainty of the center coordinates which is an interesting characteristic of the dual representation. The

complete variance-covariance matrix, $\mathbf{S}(\Theta')$, is obtained from the following relation:

$$\mathbf{S}(\Theta') = \frac{(K\Theta)^T W (K\Theta)}{n-5} (K'^T W K')^{-1} \quad (10)$$

where W is a diagonal matrix with ω_i^2 as entry and n is the number of pixels in the region. Obtaining the covariance matrix of the dual ellipse center coordinate is direct:

$$\mathbf{S}(u_c, v_c) = \frac{1}{4} \mathbf{S}(D^*, E^*) = \mathbf{S}\left(\frac{D^*}{2}, \frac{E^*}{2}\right). \quad (11)$$

4. Experiments

Four experiments were carried out to assess the new dual ellipse operator on complete ellipses. First the accuracy is assessed and compared to four other approaches in simulation. In this part, the ground truth enables any bias of the operator to be uncovered. Then, the sensitivity to blur is assessed in simulation. In the presence of image blur, the edge transition is distributed over more pixels. This reduces the gradient magnitude of the edge hence the signal to noise ratio when noise is present. Next, the accuracy is assessed when the dual ellipse is estimated from regions covering gradually increasing proportions of the edge transition. Finally, the flexibility of the new operator is demonstrated in real conditions on low contrast images where contours are difficult to identify reliably.

In the experiments, the proposed operator is compared to the original Forstner operator for circular features (FCO), to the gray-scale centroid and to a point ellipse fitting method that is similar to the one proposed in [10]. In this last method, the edge points of the feature are detected and their positions are refined to subpixel precision [9]. An ellipse is then estimated from the points with a direct least-squares fitting procedure that imposes an ellipse constraint [3]. To show the importance of subpixel correction, the results for ellipse estimation from the uncorrected local-maxima points are also included. It is important to note that the FCO and centroid methods only extract the feature center. The experiments focus on the center estimation but the dual ellipse method also extracts the complete set of parameters. When it is required for a given approach, the image gradient is calculated using the same 5x5 Gaussian derivative filter.

To simplify the experiment procedure, the detection is decoupled from the localization. The experiment images contain contrasting ellipse features which produce a clear bimodal distribution of the gradient magnitude. This allows the setting of an automatic threshold [13] which results in a coarse identification of the ellipse edge transition region. In Fig. 3, such a region identified with gradient thresholding is displayed in white along with a plot of the gray levels




Figure 3. Regions identification procedure for the localization operator. The top row shows an individual ellipse image and the bottom is a plot of its cross-section. The solid lines represent the intensity value of the ellipse and the dashed lines depict the gradient magnitude. A threshold is computed from the gradient and is shown with a red line in (a). The region identified, shown in (b), is then dilated resulting in a wider region containing the complete edge transition, which is observable in (c).

and gradient magnitudes of the ellipse cross-sections. Identifying all the pixels where the transition occurs will benefit the operators. This can be achieved by doubling the area of the region with a dilation process. The complete procedure is illustrated in Fig. 3. For the accuracy versus noise and blur experiments, the contributing regions were identified following this procedure. For instance, point ellipses were estimated from thresholded local maxima belonging to these regions. In the case of the centroid, the ellipse interior area was added to the region. In all simulations, the region delimiting the ellipse edge transition is defined before adding noise to the image.

4.1. Accuracy versus noise

In simulation, synthetic gray-scale images of ellipses are generated with their actual positions precisely known. The pixel intensities are set inversely proportional to their areas inside the actual ellipse before smoothing using a Gaussian filter of $\sigma = 0.5$. Next, the images are contaminated with additive Gaussian noise with standard deviation increasing from $\sigma_n = 0$ to 5 percent of the image dynamic range. For each noise level, the mean and max error of the ellipse position are evaluated on 150 different single ellipse images. Actually, besides adding noise, a different ellipse is gen-


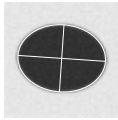
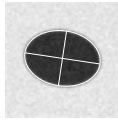
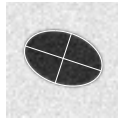
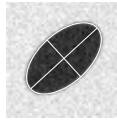
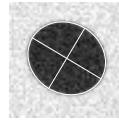
erated in each image. The ellipse center coordinates, axes and angle are generated randomly. While the ellipse center coordinates are restricted to lie within a 5.0 pixel radius of the image center, the axes and angle of the ellipse vary randomly within the range $[5.0, 15.0]$ pixels and $[-\pi/2, \pi/2]$ radians respectively. The positioning errors are compiled in Table 1. The top row includes sample images superimposed with ellipses whose parameters were obtained from dual ellipse fitting. The bottom row displays the same images with the uncertainty ellipses (1σ) on the center coordinates superimposed (enlarged by a 250x factor).

The simulation results indicate that FCO is the most sensitive to noise. Actually, although it exploits the set of local edges, it does not impose the more global constraint of an ellipse geometric model. Since the gray-scale centroid method averages over all pixels of the feature, it shows high immunity to noise. The dual ellipse fitting is the most accurate in low noise conditions. As noise increases, the estimation of the gradient orientation becomes less precise and the dual ellipse and subpixel point ellipse methods then present similar accuracy. The high error obtained from ellipse fitting on raw points is due to random leaps of position of the local maxima caused by image noise. This confirms the importance of an accurate subpixel correction step for point ellipses. Although the results are not included here, the experiment was in fact repeated with circular features where the ellipses axes were set equal. As expected, the FCO error was reduced by an average of 15% but remained far from the precision of the best methods. The leverage effect of the gradient orientation error interfering the FCO is nonexistent for the dual ellipse operator. Finally, the distribution of the error vectors confirmed the absence of bias for all of the operators.

4.2. Accuracy versus blur

Sharp and well focused images are obviously preferable when locating ellipses. However, in practice, images are rarely in perfect focus which results in edge transition being distributed over several pixels. In this section we assess how focus and blur affect the accuracy of the ellipse center estimation for the dual ellipse and subpixel point ellipse operators. For this purpose, random ellipse images are generated following the procedure given in the previous section. They are then blurred by successive convolutions with a Gaussian filter of $\sigma = 0.5$. Blurring directly affects the width of the edge transition. Next, the images are contaminated with additive Gaussian noise and the ellipses are estimated. For each convolution count, a series of 150 new images is generated and the mean error of the estimated ellipse center is evaluated. The experiment is repeated for three image noise levels ($\sigma_n = 0\%$, $\sigma_n = 2\%$, $\sigma_n = 5\%$ of the image dynamic range). The results are compiled in Fig.

Table 1. Error on the center estimation (in pixels) for five location methods given at six noise levels

						
Noise level σ_n (percent)	0	1	2	3	4	5
	<i>Mean (Max)</i>	<i>Mean (Max)</i>	<i>Mean (Max)</i>	<i>Mean (Max)</i>	<i>Mean (Max)</i>	<i>Mean (Max)</i>
<i>DualEllipse</i>	0.002 (0.006)	0.005 (0.012)	0.009 (0.022)	0.014 (0.036)	0.020 (0.047)	0.023 (0.060)
<i>PointEllipse</i>	0.008 (0.017)	0.008 (0.021)	0.011 (0.028)	0.015 (0.042)	0.021 (0.054)	0.023 (0.063)
<i>PointEllipse</i> ¹	0.170 (0.327)	0.167 (0.293)	0.161 (0.295)	0.166 (0.283)	0.169 (0.309)	0.169 (0.298)
<i>Centroid</i>	0.012 (0.026)	0.015 (0.055)	0.020 (0.050)	0.024 (0.069)	0.033 (0.091)	0.039 (0.109)
<i>FCO</i>	0.079 (0.208)	0.089 (0.215)	0.123 (0.392)	0.161 (0.620)	0.219 (0.606)	0.241 (0.734)

¹ Without subpixel correction

4. In the figure, the bottom horizontal axis reveals the convolution count whereas the top horizontal axis reveals the corresponding edge transition width in pixels. The gray region delimits the mean standard deviation estimation of the dual ellipse center coordinates.

In the absence of noise, the plot in Fig. 4(a) shows the error of the dual ellipse operator decreasing with blurring, which confirms that the dual ellipse operator is mostly not affected by image blur. Although blurring distributes the edge transition over several pixels, all pixels still contribute to the estimation. Inversely, the error of the subpixel point ellipse operator increases since the subpixel correction step does not adapt to different edge widths. In Fig. 4(b) and Fig. 4(c) we observe a sudden increase in the error progression rate for the point ellipse operator when image blur increases in the presence of noise. The sudden change is due to the appearance of noisy gradient local maxima in the edge transition region, that have a magnitude similar to those of the ellipse edge. These points are mistakenly interpreted as ellipse contour points and bias the estimation. The error progression of the dual ellipse operator remains constant since all pixels of the transition region contribute. Moreover, the increasing uncertainty of the center coordinates clearly reflects the degradation of the accuracy by noise and blur. This information can be exploited to assess the relative precision of target estimation in an image.

4.3. Accuracy versus estimation area

In the preceding experiments all the pixels of the ellipse edge transition contributed to the dual ellipse estimation.

While this is intuitively preferable, it was also suggested that a bounding box region enclosing the ellipse could be selected. In this section, we analyze how the coverage of the edge transition affects the accuracy of the dual ellipse operator. To accomplish this, random ellipse images are first generated following the procedure given in Section 4.1. They are then blurred by successive convolutions with a Gaussian filter ($\sigma = 0.5$) so that the width of the edge transition reaches six pixels. Next, the images are contaminated with additive Gaussian noise before estimating the ellipse. An example of such an image with additive Gaussian noise of $\sigma_n = 2\%$ is given in Fig. 4(b) right. The region from which the dual ellipse is initially estimated contains the gradient local maxima describing the ellipse contour extracted before adding noise. An example of such a region is shown superimposed to its corresponding image in Fig. 5(a). The region is then successively dilated with a 3x3 structuring element. The progression from a one-pixel region to an area covering a region larger than the ellipse is illustrated in Fig. 5. For each dilation count, a series of 150 new images is generated and the mean error of the estimated ellipse positions is evaluated. The experiment is repeated for three image noise levels ($\sigma_n = 0\%$, $\sigma_n = 2\%$, $\sigma_n = 5\%$ of the image dynamic range). The mean error is plotted in Fig. 6 with respect to the dilation count (bottom horizontal axis) and the corresponding annular region width (top horizontal axis). The plot in Fig. 6 depicts the mean ellipse center error after each dilation for three noise levels, $\sigma_n = 0\%$, $\sigma_n = 2\%$, $\sigma_n = 5\%$ of the image dynamic range.

The high error observed initially for low noise conditions ($\sigma_n = 0\%$ and $\sigma_n = 2\%$) in Fig. 6 decreases quickly as the

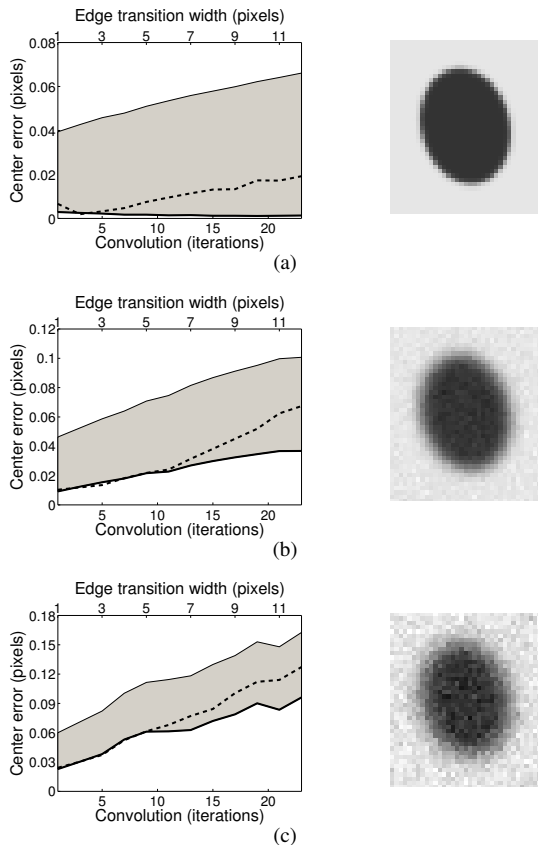


Figure 4. Error on the ellipse center estimation in pixels with respect to image blur for three noise levels. (a) $\sigma_n = 0\%$ (b) $\sigma_n = 2\%$ (c) $\sigma_n = 5\%$. The sample images are given after 1, 12 and 24 convolutions respectively. The thick solid line represents the error for the dual ellipse operator while the dashed line depicts the error of the subpixel point ellipse operator. The gray region delimits the estimated standard deviation of the dual ellipse center.




Figure 5. (a) Initial region containing the local maxima of the image gradient. (b) Initial region after performing 4 dilations. (c) Initial region after performing 20 dilations.

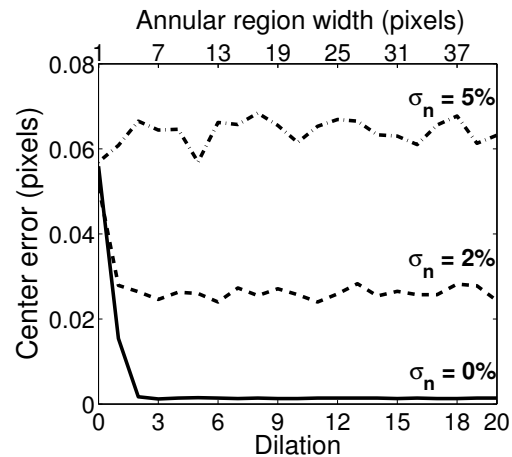


Figure 6. Error on dual ellipse centers in pixels when estimated from different proportions of the edge transition for three image noise levels ($\sigma_n = 0\%$, $\sigma_n = 2\%$ and $\sigma_n = 5\%$).

region expands and covers a higher proportion of the ellipse edge transition. The breaking point observed after 2 dilations, where the region reaches the width of the edge transition, confirms the importance of all transition pixels in the estimation. Moreover, we observe no significant loss of accuracy as the region expands beyond the edge transition. The operator weighs each pixel by the squared gradient magnitude hence it is not affected by pixels in flat regions. However, when the noise reaches 5%, the error slightly increases as the region expands. This tendency is due to the increasing importance of noisy pixels in the estimation. It is also present but non significant in the curves with noise $\sigma_n = 2\%$.

4.4. Low contrast image

The last experiment shows the strength of the new method in a low contrast situation where meaningful edge points can hardly be extracted. Such a situation is shown in Fig. 7. The image was obtained by aiming a desktop projector to a planar object which created a blurry and non-uniform image. The thresholded gradient local-maxima obtained after smoothing the image can be observed in Fig. 7(a) where it is superimposed onto the image. The local maxima are sparse, unconnected and do not belong to a well defined contour. However the global structure of an ellipse is present. The dual ellipse estimated from all the pixels of the image can be seen superimposed upon the original image in Fig. 7(b) and onto the gradient in Fig. 7(c). This type of low contrast image can also be handled with active contours methods [18]. However, the proposed operator is non




Figure 7. Images of a light spot illustrating a difficult case of ellipse estimation. (a) Individual edge points. (b) The estimated ellipse superimposed onto the original image. (c) The estimated ellipse superimposed onto the image gradient magnitude. Brighter pixels represent higher gradient magnitudes.

iterative and imposes an ellipse constraint on the recovered shape which improves the quality of the solution.

5. Conclusion

The new linear operator combines the advantages of the most precise ellipse fitting methods with the simplicity of the Forstner operator as well as its capability to work on the raw image gradient. The absence of precise contour point extraction makes it faster and simpler than point conic methods and allows the uncertainty of the parameters to be evaluated directly from the raw gradient image. Although the dual ellipse is estimated from pixels present in the ellipse edge transition, it needs not to be limited to that region. Weighing with the squared gradient magnitude enables the operator to process all the pixels that lie in a region enclosing a single ellipse. Furthermore, the operator efficiently constrains the dual conic to a dual ellipse when minimizing the algebraic error.

More robust estimation techniques can also be easily applied in replacement of the least-squares if that were necessary for a given application. Although we have not considered partial ellipses, it might be interesting for some applications requiring less accuracy, to investigate this kind of more direct method.

Acknowledgements

The authors express their gratitude to the supporting organizations including CRSNG-NSERC and Precarn incorporated.

References

[1] F. L. Bookstein. Fitting conic sections to scattered data. *Computer Graphics and Image Processing*, 9(1):56–71,

- 1979.
- [2] J. Cabrera and P. Meer. Unbiased estimation of ellipses by bootstrapping. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(7):752–756, 1996.
- [3] A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(5):476–480, 1999.
- [4] A. W. Fitzgibbon and R. B. Fisher. A buyer’s guide to conic fitting. In *Proceedings of the 6th British conference on Machine vision*, volume 2, pages 513–522, Surrey, UK, UK, 1995. BMVA Press.
- [5] W. Forstner and E. Gulch. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. Intercommission Conf. on Fast Processing of Photogrammetric Data*, pages 281–305, 1987.
- [6] W. Gander, G. H. Golub, and R. Strelbel. Fitting of circles and ellipses, least square solution. Technical Report 1994TR-217, ETH Zurich, Institute of Scientific Computing, 1994.
- [7] R. I. Hartley. In defense of the eight-point algorithm. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(6):580–593, 1997.
- [8] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [9] J. Heikkila. Moment and curvature preserving technique for accurate ellipse boundary detection. In *Proceedings of the 14th International Conference on Pattern Recognition*, volume 1, page 734, Washington, DC, USA, 1998. IEEE Computer Society.
- [10] J. Heikkila. Geometric camera calibration using circular control points. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 22(10):1066–1077, 2000.
- [11] K. Kanatani. *Geometric computation for machine vision*. Oxford University Press, Inc., New York, NY, USA, 1993.
- [12] K. Kanatani. Statistical bias of conic fitting and renormalization. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(3):320–326, 1994.
- [13] N. Otsu. A threshold selection method from gray-level histograms. *IEEE Trans. on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.
- [14] P. L. Rosin. Assessing error of fit functions for ellipses. *Graphical models and image processing: GMIP*, 58(5):494–502, 1996.
- [15] R. Safaee-Rad, I. Tchoukanov, B. Benhabib, and K. C. Smith. Accurate parameter estimation of quadratic curves from grey-level images. *CVGIP Image Understanding*, 54(2):259–274, 1991.
- [16] R. Szeliski and S. B. Kang. Direct method for visual scene reconstruction. In *Proceedings of the IEEE Workshop on Representation of Visual Scenes*, pages 26–33, Washington, DC, USA, 1995. IEEE Computer Society.
- [17] A. Tabatabai and O. Mitchell. Edge location to subpixel values in digital imagery. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6(2):188–200, March 1984.
- [18] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1998.