A New Multiresolution Blob Detector Applied to Photogrammetry

Raquel Dosil
Dpto. de Electrónica e Computación
Universidade de Santiago de Compostela
15782 Santiago de Compostela
raquel.dosil@usc.es

Xosé M. Pardo
Dpto. de Electrónica e Computación
Universidade de Santiago de Compostela
15782 Santiago de Compostela
xose.pardo@usc.es

Xosé R. Fdez-Vidal
Dpto. de Física Aplicada
Escola Politécnica Superior
Univ. de Santiago de Compostela
27002 Lugo
xose.vidal@usc.es

Antón García-Díaz
Laboratorio Oficial de Metroloxía de Galicia
Parque Tecnolóxico de Galicia
32900 Ourense
agarcia@lomg.net

Víctor Leborán
Dpto. de Electrónica e Computación
Universidade de Santiago de Compostela
15782 Santiago de Compostela
victor.leboran@usc.es

Abstract

Industrial applications of photogrammetry usually involve the detection of targets placed on the surfaces of objects. Typical targets present circular shapes, which under perspective distortion can be well approximated by an elliptical contour. In this paper we present a solution for the detection of circular targets, based on a new technique for radial symmetry detection. At a given scale, radial symmetry is defined as the product of a set of oriented features, so that only regions of space with contributions from all orientations will have strong responses. A multiresolution measure is obtained as the maximum over a set of scales. This operator detects both circular and elliptical regions and presents a sharper point spread function compared to other detectors.

1. Introduction

A key step in photogrammetry is the detection of points on the surface of objects, and reference points for image correspondence. Most commercial systems require to manually place a set of artificial landmarks, or targets, on the surface of the object to ensure their detection. These targets usually show high contrast, often retroreflective, patterns. Many of these commercial systems use circular shaped targets. Coded reference targets usually incorporate a coded ring around the central spot.

Although target detection is a crucial step, most state-of-the-art applications use quite simple approaches for their detection [1][2][5][11][12]. This is possible since controlled scenarios and illumination conditions are assumed or forced. Under such conditions, segmentation of the image by simply thresholding can successfully be used to select candidate target regions. Heuristic post-processing might then be applied for outlier rejection, based on geometric parameters of the detected blobs, like size, perimeter, aspect ratio, error fitting, etc.

In other applications, like robot navigation, artificial circular landmark detection requires more sophisticated techniques to cope with varying conditions. In general, under perspective distortion, an elliptical model is assumed to provide a good approach to the observed target shape. A common ellipse detection technique is the convolution with circular template filters. In [7], the use of the Laplacian of Gaussian (LoG) filter is proposed as a blob detector due to its radial symmetry. Multiresolution blob detection is achieved by generating a scale space of LoG and maximizing their response over scales. A different approach is that of voting based algorithms, where contour points vote an accumulator in the location of possible central positions of the shape.

Examples of this are the fast radial symmetry detector by Loy and Zelinsky [8], or the rotation and scaling invariant Hough ellipse detector by Nixon and Aguado [9].

The above mentioned algorithms present different problems. Blob detectors do not usually offer a precise shape definition, being sensitive to very different kinds of shapes. Voting based algorithms are not completely invariant to scale, since the intensity of their response depends on the number of pixels on the contour of the shape. Furthermore, they produce artifacts, since each contour point votes many possible centre locations [9]. The fast radial symmetry detector [8] solves this problem by having into account the intensity of the contours, and also by imposing strong constraints on the circularity of the shape, at the expense of a lack of robustness to perspective distortion. On the other hand, the main drawbacks of the different approaches to the Hough transform of ellipses are related to high computational cost and wide point spread functions.

To cope with these problems, we propose a novel technique to detect circular targets that is invariant to rotation and scale, robust under perspective distortion and with a narrow point spread function. The main idea behind this detector is to consider radially symmetric blobs as region features with contributions from all orientations at a given scale. To detect regions with these characteristics, we design a radially symmetric blob detector as the product of a set of features with even phase, same size, but different orientations. Like in [7], to get scale invariance, we compute a scale space of radially symmetric blobs and, subsequently, we obtain the overall symmetry response by maximization over scales.

In a way, this is the same idea behind the blob detector in [7]. The difference relies in the use of the product instead of the addition implicit in the LoG filter. From a set of oriented features, we can achieve a radial symmetry template resembling the LoG by summation of their responses. The overall response is strong when all orientations are strong, but can also be high in many other cases. On the other hand, the product of oriented features must be high at locations presenting radial symmetry, but will rapidly decay to zero in regions lacking any of the orientation components.

We will apply this blob detector to the problem of circular target detection in photogrammetry, under arbitrary illumination conditions. Target regions will be extracted around symmetry maxima, with size depending on the scale of the feature. Precision target location is achieved by fitting contours in the region of interest to an elliptical model. The system will be validated by showing performance results in real situation experiments, and by comparing results with other techniques, including state of the art thresholding based techniques, and other general purpose circular feature detection techniques.

The rest of the paper is organized as follows. Section 2 describes the implementation details of our radially symmetric blob detector and the proposed target detection system. Section 3 presents the experimental setup for validation, followed by the results and discussion. Conclusions will be summarized in section 5.

2. Implementation

In the following subsection, the proposed blob detector is described in detail. The complete system for target detection is presented in subsection 2.2.

2.1. Radially Symmetric Blobs

The first step of our technique for the computation of radial symmetry consists in multiresolution decomposition of the image into a set of band-passed features. To this end we choose a bank of scaled and oriented log Gabor filters [3]. Log Gabor filters show many advantageous properties [6]. In particular, we have chosen them because, unlike Gabor filters, they have a zero DC component, and also because their response is better localized in the spatial domain. The log Gabor filter does not have an analytical expression in the spatial domain. In the spectral domain, its transfer function T is usually defined in polar frequency coordinates, and parameterized by the frequency ρ_i and orientation ϕ_j of its peak and the corresponding radial and angular widths, $\sigma_{\rho i}$, and $\sigma_{\phi j}$, as follows

$$T_{ij} = \exp\left(-\frac{\log(\rho/\rho_i)^2}{2\log(\sigma_{\rho i}/\rho_i)^2} - \frac{(\phi-\phi_j)^2}{2\sigma_{\phi j}^2} \right) \quad (1)$$

The response R_{ij} of one of these filters to an input image I in the spatial domain is complex

valued, with the real and imaginary components representing even and odd phase features respectively, and with scale $\lambda_i=2/\rho_i$ and orientation ϕ_j. It can be rapidly computed in the frequency domain using the Discrete Fourier transform DFT and its inverse IDFT.

$$R_{ij} = \mathrm{IDFT}(T_{ij} \cdot \mathrm{DFT}(I)) \qquad (2)$$

The bank of filters used here is designed to evenly cover the spectral domain. We choose to use a fixed number of N_s scales, with wavelengths $\lambda=\lambda_{\min}M^{i-1}$ and $M=(\lambda_{\min}/\lambda_{\max})^{1/(Ns-1)}$. In our experiments $\lambda_{\min}=2/\rho_{i=1}=5$ pixels and $\lambda_{\max}$ is 1/8 the size of the image, while $\sigma_{\rho i}$ is computed for each scale so that the bandwidth is 2 octaves. The number of scales can be adapted to the expected sizes of the targets in the image. The number of orientations is not either a sensitive parameter of the model. Here we have chosen a number $N_o=4$ orientations, with $\sigma_{\phi j}=\pi/N_o$.

From the band-passed features R_{ij}, a scale space of radial symmetry is computed as the product of all oriented responses with same scale. Previous to product computation, the imaginary component of the responses is discarded, since the targets are assumed to be region features –even symmetric patterns. Besides, prior knowledge about contrast of the targets is exploited in order to eliminate responses due to spots with opposite contrast. Hence, the response is inverted depending on contrast and half wave rectified. Therefore, if c is the contrast of the targets (1 for white targets, -1 for black targets), then symmetry measure S for scale i is

$$S_i = \prod_j \max(0, c\Re\{R_{ij}\}) \qquad (3)$$

The overall radial symmetry is obtained from the radial symmetry scale space, by selecting the proper scale at each image pixel, as the one that maximizes symmetry over scales

$$S = S_{i*}, \quad i*(x,y) = \arg\max_i\{S_i(x,y)\} \qquad (4)$$

The method is depicted in Figure 1.

2.2. Target Detection

The location of blobs is obtained by non-maxima suppression of the radial symmetry measure, and their size is given by the corresponding scale $\lambda_{i*}(x,y)$. Maxima are thresholded to retain only the strongest peaks. The threshold is computed automatically from the output of the detector previous to non-maxima suppression, as the mean plus the standard deviation.

The region of interest of the target (ROI) is selected here as a square area around the spatial local maxima of the radial symmetry S, with size depending on the corresponding scale $\lambda_{i*}(x,y)$. To this end, we apply a standard non maxima suppression algorithm.

We are not concerned here with the precision of the location of the targets. We refer the reader to the recent work by Gutierrez and Armstrong [4]. From their review, we select an ellipse fitting algorithm due to its good tradeoff between precision and speed. Hence, we perform contour detection in the target ROI using a Canny operator. Ellipse fitting is accomplished by least squares, using a Newton-Raphson algorithm. Finally, the error of the fitting is thresholded in order to eliminate wrong detections.

3. Validation

Validation is realized by measuring its performance in real situations and comparing it with other techniques. Performance is here measured using the F_1 score statistic, which represents the balance between sensitivity and precision. *Sensitivity* is the fraction of actual targets that is correctly detected, while *precision* is the fraction of detections that are actual targets. Then,

$$sensitivity = \mathrm{TP/P} = 1-\mathrm{FN/P} = 1-\mathrm{FN/P} \qquad (5)$$

$$precision = \mathrm{TP/(TP+FP)} \qquad (6)$$

$$F_1^{-1} = (sensitivity^{-1} + precision^{-1})/2 \qquad (7)$$

where P is the actual nº of targets, FP is the nº of false positives –points incorrectly labeled as targets–, FN is the nº of false negatives –actual targets not detected– and TP is the nº of true positives –correctly detected targets–. F_1 values range from 0, in the worst case, to 1, for the best performance.

In the next subsection we present the methods used for comparison. Remainder subsections are respectively devoted to the experimental setup description, results and discussion.
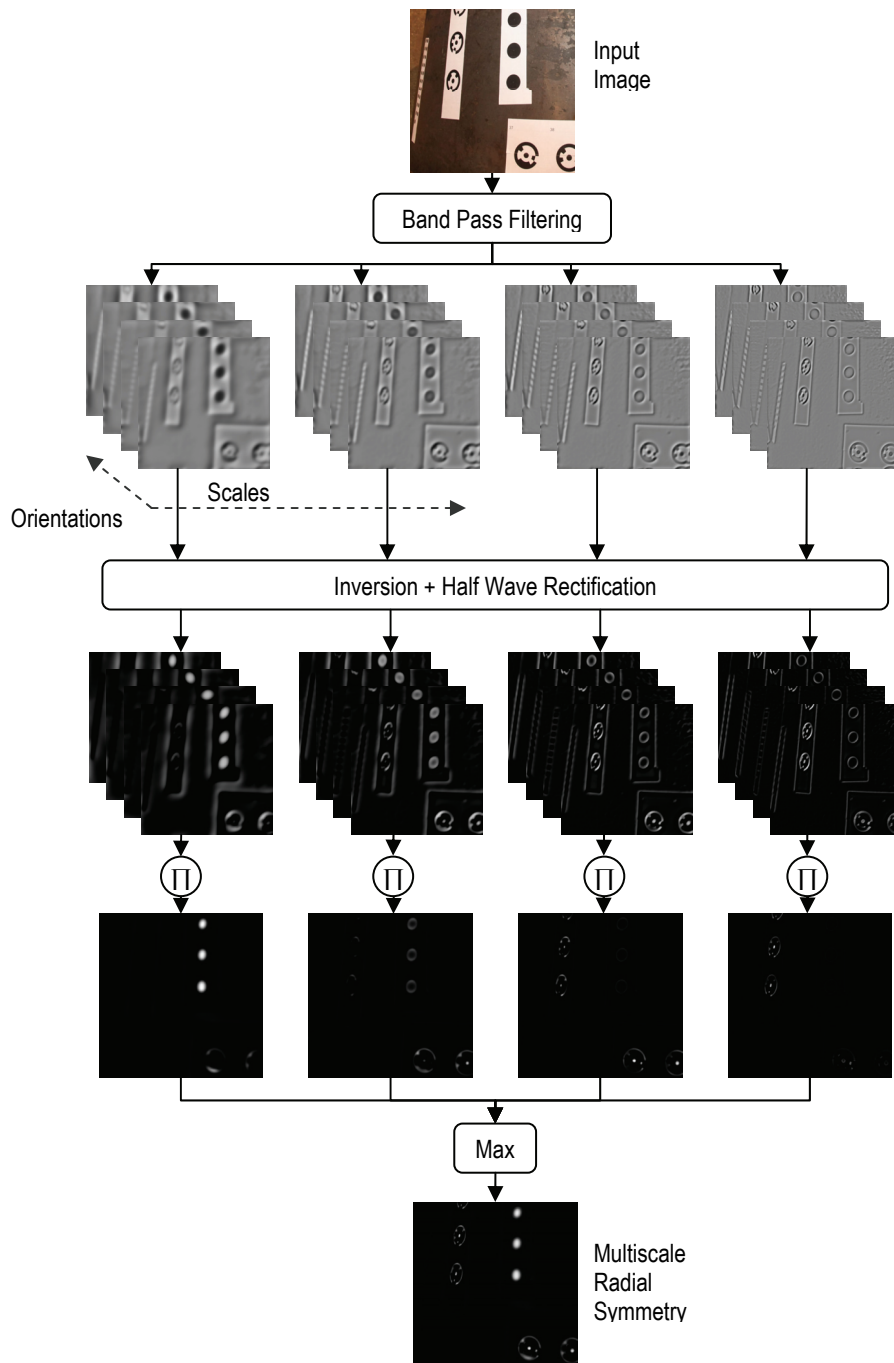
Figure 1.   Flowchart synthesizing the computation of the proposed multiscale radial symmetric blob detector.

3.1. Compared Algorithms

The most widely approach to target detection in photogrammetry is based on segmentation by thresholding. Comparisons will be made with the Adaptive Thresholding method by Wellner [13] and also with another approach to radial symmetric blob detection, the multiscale LoG described in [7]. Finally, results yielded by the commercial photogrammetric system PhotoModeler® Scanner, whose algorithm for target detection we do not know, are also shown.

To distinguish between *LoG* blob detector and ours, we will refer to them as *LoG* and *PO* (*Product of Orientations*) respectively. The procedure to extract the location of targets corresponding to the *LoG* detector will be exactly the same as the one described in section 2, but replacing the *PO* blob operator with the *LoG*.

For adaptive thresholding –we will refer to it as *AT*–, the difference relies in the way of selecting the region of interest. It now implies the detection of connected components on the thresholded image. Then, for each component, the corresponding ROI is set as the smallest square region enclosing it.

3.2. Experimental setup

In our experiments we have used targets taken from the commercial system PhotoModeler Scanner®, with black on white contrast, like the ones shown in Figure 2.
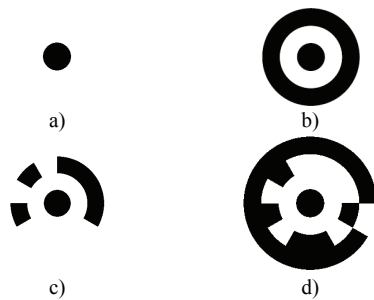

a) b)

c) d)

Figure 2.   PhotoModeler Scanner® targets. a,b) non-coded targets, c) 12 bit regular coded target, and d) RAD coded target.

The test bench of images consists of a set of three 5616x3744 pixel images, shown in Figure 3, which represent both coded and non-coded targets of different sizes, placed on the surface of three different metallic objects, with different surface reflectance properties. Pictures have been taken using the reflex camera Canon EOS 5D Mark II, with 21Mpx and 35mm full-frame sensor. The scenes correspond to three different environments, all in real industrial production plants. No special illumination setup has been used, preserving the default illumination. Images #1, #2 and #3 contain 150, 418 and 167 targets respectively.
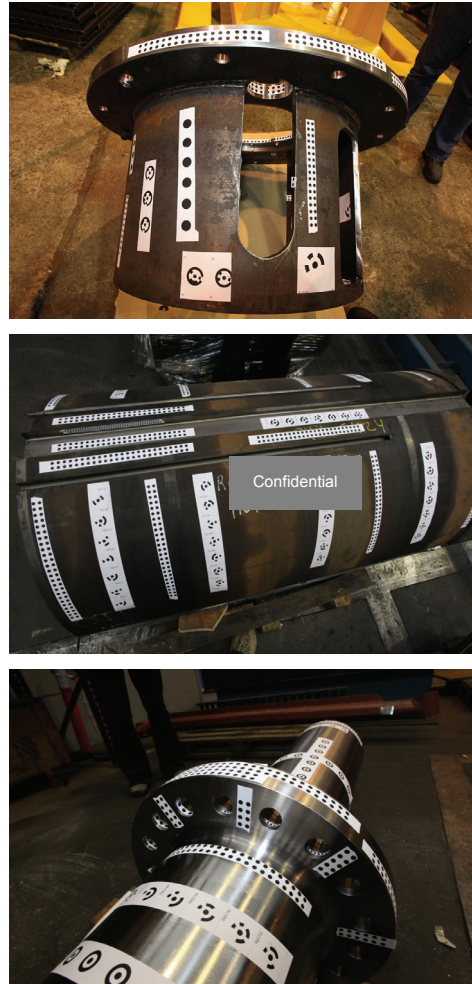

Figure 3.   The three test images.

3.3. Results

Figure 4 shows the target detection performance for each of the images and methods. It is measured as the F_1 score achieved with a range of fitting error threshold values. Figure 5 shows some of the results obtained to give an idea of the types of false positives and false negatives found.

We have also tested the performance of the commercial system PhotoModeler Scanner®. In that case the error threshold is not a controllable parameter. F_1 scores obtained are 0.34, 0.06 and 0.032 for images #1, #2 and #3 respectively.

Image #1
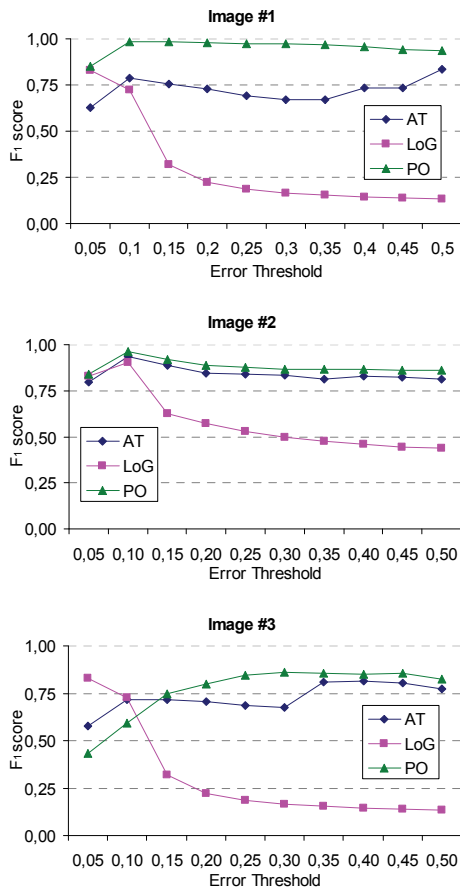

Image #2

Image #3

Figure 4. F_1 score statistic of the three different methods, as the achieved using different fitting error threshold levels.
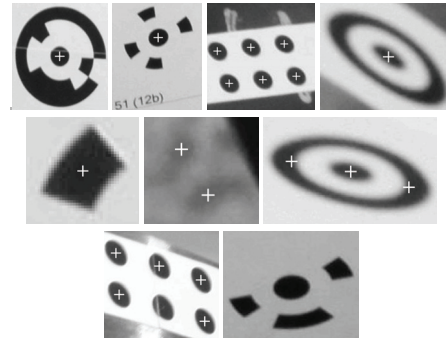

Figure 5.   Examples of: *Top,* correctly detected targets, marked with a plus sign; *Middle,* false positives; *Bottom*, false negatives.

3.4. Discussion

In Figure 4 it can be observed that the proposed *PO* blob detector shows a good tradeoff between sensitivity and precision in all three images, while the other methods present lower values in most of the experiments. The especially low performance found for PhotoModeler® Scanner, with an unexpected high level of false negatives, might be attributed to very strong constraints for outlier rejection, although this is speculative, since the actual algorithm employed is not made public.

Taking a look to the targets detected in all three cases we find out that all of the methods are able to cope with perspective distortions and different sizes of the targets. In turn, the differences must come from their sensitivity to illumination conditions, clutter, and from selectivity of the detectors. The results can be better understood by taking a look at the responses of the detectors in Figure 6.

The first column of Figure 6 shows the results of adaptive thresholding on the three test images. Due to the high contrast of the target patterns, it performs quite well in images like #2, where the surface of the object has low reflectance and the background in not highly cluttered. Its lower performance is achieved for image #1, which presents a highly cluttered background causing many false positives. Image #3, also presents this problem, but now due to the reflectance properties of the material.

The second column shows the output of the multiscale *LoG* blob detector. As can be seen,

many of the local maxima are blurred, which causes many false negatives. On the other hand, many local maxima do not correspond to target locations, causing abundance of false positives, which soon affects performance as the error threshold is increased.

Finally, the third column shows the output of the multiscale *PO* blob detector. Since it performs a product of all orientations at each scale, instead of a summation, its point spread function is much narrower than that of the *LoG*, with stronger maxima and better localization. Furthermore, its definition of radial symmetry is more constrained, thus producing fewer artifacts. It is also less sensitive to noise and clutter than thresholding. It has been observed that most of the false positives are due to the segments in the rings of the coded targets, also present in the other two methods.

It must be pointed that the procedure for outlier rejection is a very simple one. Results for all three methods could be largely enhanced, for example, by imposing heuristic constraints on geometric parameters of the targets, like average curvature of the contours, average intensity, etc. Additionally, the use of precision location procedures other than contour fitting might be also helpful for outlier rejection. For example, an intensity model would provide information, not only about the shape of the target, but also about its appearance, which should enhance results independently of the detection method.

To finish, we must also mention that none of the three methods is invariant to local contrast of the image. Some degree of invariance is achieved with the adaptive nature of the thresholding method selected. In the case of the blob detectors, they are both subject to a thresholding of the non-maxima suppressed responses based on their strength. The selection of this threshold surely affects the results. The high contrast of the target patterns helps in minimizing this effect, but some false negatives can still be caused due to local illumination conditions.

4. Conclusions

In this paper, a new measure for radial symmetry has been presented, which is easily computed from bank of log Gabor filters with multiple orientations and scales. Radial symmetric features are defined as regions of space with frequency components from all orientations at a given scale. From this definition, a radial symmetry detector is computed as the product of all oriented features with the same scale. The multiscale version of the detector is obtained by computing the maximum symmetry over scales.

The proposed measure presents several interesting characteristics. It is invariant to rotation and changes in scale. The intensity of its response does not depend on the size of the target and it is robust under strong perspective distortion. Its maxima are strong and well localized, and it does not produce many weak maxima due to non radially symmetric patterns.

The presented radial symmetric blob detector has been applied to the problem of circular photogrammetric target detection. The result of the blob detector is used to determine the region of interest of the targets. Least squares fitting to an elliptical model is applied for subpixel target location, which is a state of the art solution.

Validation of the complete system has been realized using a set of commercial targets, quite representative of the usual coding schemes in industrial photogrammetry. The test images are representative of real situations, corresponding to industrial production plants with usual illumination conditions. Results have been compared to two other algorithms. One of them, based on adaptive thresholding, is representative of the state of the art target detection techniques in photogrammetry. The other one, based on the LoG multiscale blob detector, is a common solution for radial symmetric blob detection in many computer vision applications. The results presented prove that the proposed method shows an acceptable tradeoff between sensitivity and precision, and that it is superior to the other evaluated methods.

Acknowledgements

References

[1] Ahn, S.J.; Rauh, W; Recknagel, M. Circular coded landmark for optical 3D-measurement and robot vision. Conf. on Intelligent Robots and Systems (IROS'99), pp. 1128-1133, 1999.

[2] Chen, Z; Ye, Z; Chan, D.T.W.; Peng, G. Target Recognition Based on Mathematical Morphology. CAD/Graphics'07, pp. 457-460.

[3] Field, D.J. Scale–Invariance and self-similar "wavelet" Transforms: An Analysis of Natural Scenes and Mammalian Visual Systems. In: Wavelets, fractals and Fourier Transforms. pp. 151-193, Clarendon Press, 1993.

[4] Gutierrez, J.A.; Armstrong, B.S.R. Precision Landmark Location for Machine Vision and Photogrammetry. Springer, 2008.

[5] Knyaz, V.A.; Sibiryakov, A.V. The Development of New Coded Targets for Automated point Identification and Non-contact 3D Surface Measurements, Int. Archives of Photogrammetry and Remote Sensing, 32(5):80-85, 1998.

[6] Kovesi, P. Invariant Measures of Image Features from Phase Information. Ph.D. Thesis, University of Western Australia, 1996.

[7] Lowe, D.G. Distinctive image features from scale-invariant keypoints. Int. Journal of Computer Vision, 60(2):91-110, 2004.

[8] Loy, G.; Zelinsky, A. Fast Radial Symmetry for Detecting Points of Interest, PAMI 25(8):959-973, 2003.

[9] Nixon, M.S.; Aguado, A.S. Feature Extraction and Image Processing. Academic Press, 2008.

[10] Nock R, Nielsen F. Statistical Region Merging. PAMI 26(11):1452-1458, 2004.

[11] Otepka, J.O.; Hanley H.B.; Fraser C.S. Algorithm Developments For Automated Off-Line Vision Metrology. Int. Archives of Photogrammetry & Remote Sensing, 5(34):60-67, 2002.

[12] Shortis, M.R.; Clarke, T.A.; Short, T. Comparison of some techniques for the subpixel location of discrete target images. Proc. SPIE Vol. 2350, pp. 239-250, Videometrics III, 1994.

[13] Wellner, P.D. Adaptive Thresholding for the DigitalDesk, Xerox Technical Report, 1993.
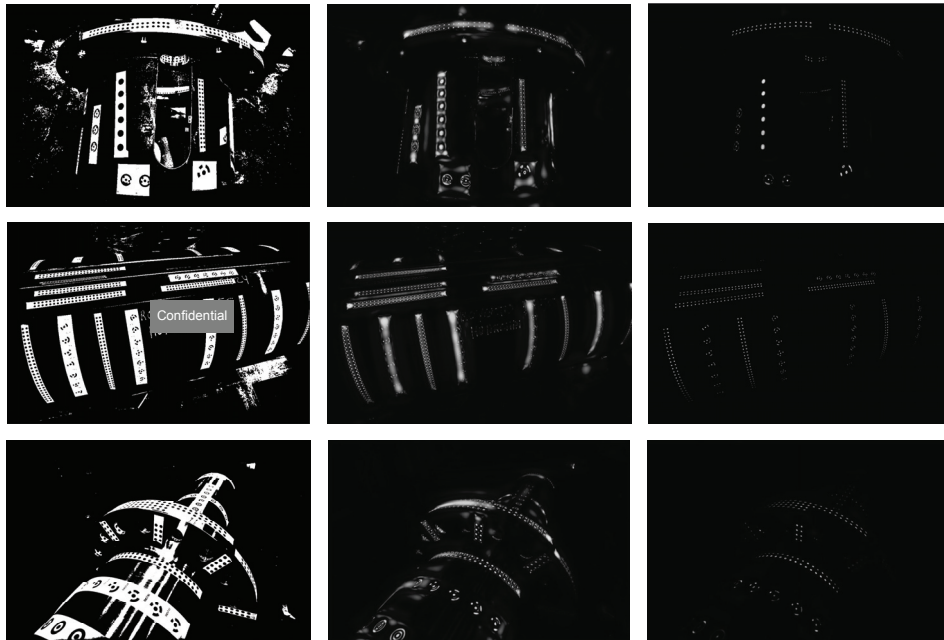
Figure 6.   Detection results. From top to bottom, images #1, #2 and #3. From left to right, adaptive thresholding, LoG blob detector, and Product of Orientations blob detector.