

Local Feature Extraction in Log-Polar Images

Manuela Chessa^(✉) and Fabio Solari

Department of Informatics, Bioengineering, Robotics and System Engineering -
DIBRIS, University of Genoa, Via All'Opera Pia 13, 16145 Genova, Italy
manuela.chessa@unige.it

Abstract. We propose two different strategies to compute edges in the log-polar (cortical) domain. The space-variant processing is obtained by applying local operators (e.g. local derivative filters) directly on the log-polar images, or by embedding the same operators into the log-polar mapping, thus obtaining a cortical representation of the Cartesian features. The two approaches have been tested by taking into consideration three standard algorithms for edge detection (Canny, Marr-Hildreth and Harris), applied onto the BSDS500 dataset. Qualitative and quantitative comparisons show a first indication of the validity of the proposed approaches.

Keywords: Space-variant processing · Foveated images · Edge detection · Corner detection · Cortical representation · Bio-inspired visual processing

1 Introduction

The computation of local image features such as edges and corners is at the basis of many approaches for matching and recognition, which are important in image processing, computer vision and robotics applications.

In the literature, several edge or contour detection methods, based on local operators applied in the Cartesian domain, are described. Among them, some algorithms detect edges by convolving a grayscale image with local derivative filters (e.g. the Roberts, Sobel and Prewitt operators), the Marr and Hildreth method uses zero crossings of the Laplacian of Gaussian operator, and the Canny operator defines edges' detection and localization criteria based on first derivatives of a Gaussian [5]. A combined edges and corners detector operator, based on local derivatives, has been proposed by Harris [8]. More recent local approaches take into account color and texture information and make use of learning techniques for cue combination [11]. Recently, a contour detector method that combines multiple local cues into a globalization framework, based on spectral clustering, is described in [1], and a multi-scale Harris corner detector is proposed in [7].

Though a great effort in improving edge detectors in the Cartesian domain has been done, few works address the same problem in the log-polar (cortical) domain [14]. The space-variant images are promising for many image processing

and robotics applications, since they provide a high spatial resolution in the region of interest, i.e. the fovea, and a reduction of the amount of data to be processed, similarly to what happens in the mammals' visual system. Indeed, the distribution of the photoreceptors in the primates' retina is space-variant (i.e. denser in the center, the fovea, and sparser in the periphery), and the projection of such photoreceptors into the primary visual cortex can be described by a log-polar mapping [14]. However, the processing of the log-polar images is a challenging task, due to the image distortions generated by the retino-cortical transform, which often require a specific adaptation of the algorithms, in order to properly work. Primal feature extraction [10] in log-polar images has been addressed by some authors in the literature, which propose ad-hoc solutions designed to work in the cortical domain. In [12], the authors present a mechanism for computing operators such as edge detection and Hough transform directly in foveated images; in [6] the authors present an approach to extract edges, bars, blobs and ends from log-polar images, based on neural networks that learn the feature's class; and a comparison of several strategies for gradient detection in log-polar images is presented in [18].

In this paper, we propose two approaches that allow standard algorithms to work in the log-polar domain, without specific adaptation. Thus, we do not propose a new method for feature detection, but we analyze how well known state-of-the-art methods for edges and corners extraction work in the log-polar domain. In particular, the aim of the paper is to show the performances, also in terms of accuracy of the feature detection, of the two considered approaches: (i) the local operator for edge detection (i.e. derivative of Gaussian, and Laplacian of Gaussian) is applied on the cortical image, i.e. on the image that has been transformed into the log-polar image through low-pass Gaussian filters; (ii) the local operator is embedded into the log-polar transform, thus producing a cortical representation of the Cartesian derivatives of the image, on which to compute edges. Moreover, we assess the two proposed approaches by using the metrics and the BSDS500 dataset presented in [1].

2 Log-Polar Mapping

The log-polar mapping is a non linear transformation that maps each point of the Cartesian domain (x, y) into a cortical domain described by the coordinates (ξ, η) . In the literature, several log-polar mapping models are described [2, 4, 9]. We consider the central blind-spot model, since it is characterized by scale and rotation invariance [17].

The log-polar transformation is described by the following equations:

$$\begin{cases} \xi = \log_a \left(\frac{\rho}{\rho_0} \right) \\ \eta = q\theta, \end{cases} \quad (1)$$

where a parameterizes the non-linearity of the mapping, q is related to the angular resolution, ρ_0 is the radius of the central blind spot, and

$(\rho, \theta) = (\sqrt{x^2 + y^2}, \arctan(y/x))$ are the polar coordinates derived from the Cartesian ones. All points with $\rho < \rho_0$ are ignored, thus ρ_0 has to be small, with respect to the size of the image.

In order to deal with digital images, given a Cartesian image of $M \times N$ pixels, and defined $\rho_{max} = 0.5 \min(M, N)$, we obtain an $R \times S$ (rings \times sectors) discrete cortical image of coordinates (u, v) by taking:

$$\begin{cases} u = \lfloor \xi \rfloor \\ v = \lfloor \eta \rfloor, \end{cases} \quad (2)$$

where $\lfloor \cdot \rfloor$ denotes the integer part, $q = S/(2\pi)$, and $a = \exp(\ln(\rho_{max}/\rho_0)/R)$. Figure 1 shows the transformations through the different domains. The retinal area (i.e. the log-polar pixel) that refers to a given cortical pixel defines its *receptive field* (RF). By inverting Eq. 1 the centers of the RFs can be computed, and these points present a non-uniform distribution through the retinal plane, as in Figure 2a (green crosses). The optimal relationship between R and S is the one that optimizes the log-polar pixel aspect ratio γ , making it as close as possible to 1. It can be shown that, for a given R , the optimal rule is $S = 2\pi/(a - 1)$ [15, 17].




Fig. 1. Left: the cyan circle and the green sector in the Cartesian domain (x, y) map to vertical and horizontal stripes, respectively, in the cortical domain (ξ, η) . The red area represents a RF that is mapped in the corresponding cortical pixel. Right: an example of image transformation from the Cartesian to the cortical domain, and backward to the retinal domain. The RFs (yellow circles) are overlapping the Cartesian image. The specific choice of the mapping parameters is: $R = 80$, $S = 131$, $\rho_0 = 3$, and $\rho_{max} = 256$. The cortical image is scaled to improve the visualization.

The shape of the RFs affects both the quality of the transformation and its computational burden. In [3] the authors analyze four techniques, each characterized by a different shape for the RFs: nearest pixel, bilinear interpolation, adjacent RFs, and overlapping circular RFs. The overlapping circular RFs [2, 13] are the most biological plausible technique and they allow a better preservation of the image information [3], thus we consider this solution in the paper.

To implement the log-polar mapping, the Cartesian plane is divided in two regions: the *fovea* and the *periphery*. The periphery is defined as the part of the plane in which the distance between the centers of two RFs on the same radius

is greater than 1 pixel (undersampling). To obtain the cortical image we use overlapping Gaussian RFs, as shown in Figure 2a. The fovea (in which we have an oversampling, i.e. the distance between two consecutive RFs is less than 1 pixel) is handled by using fixed size RFs, whereas in the periphery the size of the RFs grows. The standard deviation of the RF Gaussian profile is a third of the distance between the centers of two consecutive RFs, and the spatial support is six times the standard deviation. As a consequence of this choice, adjacent RFs overlap. A cortical pixel C_i is computed as a Gaussian weighted sum of the Cartesian pixels P_j in the i -th RF: $C_i = \sum_j w_{ij} P_j$, where the weights w_{ij} are the values of a normalized Gaussian centered on the i -th RF. A similar approach is used to compute the inverse log-polar mapping that produces the *retinal image*, where the space-variant effect of the log-polar mapping is observable.




Fig. 2. The RFs considered to obtain the cortical representation of the image. (a) Gaussian RFs, used to obtain the cortical image (log-polar transform). (b) Laplacian of Gaussian RFs and (c) Derivative of Gaussian RFs (along horizontal and vertical axes, respectively), used to obtain the cortical representation of the derivatives of the image.

3 Feature Detection in the Log-Polar Domain

The cortical representation $R(\xi, \eta)$ of a space-variant processed Cartesian image $I(x, y)$ is described as follows:

$$R(\xi, \eta) = \langle g(x - x_0(\xi, \eta), y - y_0(\xi, \eta)), I(x, y) \rangle, \quad (3)$$

where $\langle \cdot \rangle$ denotes the inner product, $g(x, y)$ is the local operator that defines the weights of the log-polar mapping, and $(x_0(\xi, \eta), y_0(\xi, \eta))$ is the center of each RF.

Figure 2 shows several operators, which are used to obtain $R(\xi, \eta)$. In particular, Gaussian RFs (Fig.2a) are used to obtain the cortical image, i.e. the log-polar transform of the Cartesian image, Laplacian of Gaussian (Fig.2b-c) are used to obtain the cortical representation of the image derivatives. The choice of the RFs is at the basis of the two approaches described in the following paragraphs, and summarized in Figure 3. In particular, in the following we show that the two cortical approaches can approximate the corresponding Cartesian processing.

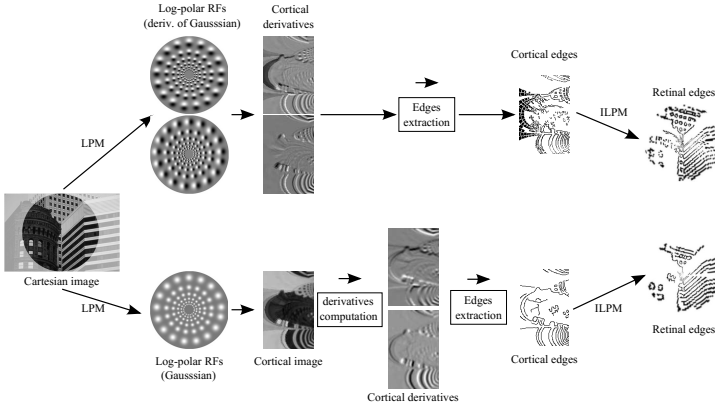


Fig. 3. Sketch of the two approaches described in the paper. Top: the Cartesian image (from the BSDS500 dataset, see Section 4) is transformed (LPM) by considering overlapping RFs embedding one of the considered local operators (in figure, the derivative of Gaussian filters) to obtain the cortical representation of the Cartesian derivatives, from which edges are extracted, thus obtaining cortical edges. Finally, through the inverse log-polar mapping (ILPM) retinal edges are computed. Bottom: the Cartesian image is transformed into the cortical image through the Gaussian RFs, then the derivatives are computed directly on the cortical image, by applying one of the considered local operators, and cortical edges are extracted. Finally, through the inverse log-polar mapping retinal edges are computed.

3.1 Direct Feature Detection in the Log-Polar Domain

To obtain a cortical representation, i.e. a cortical image, of a Cartesian image, we consider Gaussian filters as local operators that describe the RFs used to perform the log-polar mapping (see Fig.2a):

$$g(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right), \quad (4)$$

where σ is the standard deviation. Then, standard edge and corner differential local operators are applied onto the cortical image, by following the design strategies devised in [15]. In particular, to apply standard edge detection algorithms we define:

$$R_\xi(\xi, \eta) = \frac{\partial}{\partial \xi} R(\xi, \eta), \quad (5)$$

where $R_\xi(\xi, \eta)$ is the partial derivative (along ξ axis) of the cortical image, similarly for η axis.

Eq. 5 can be approximated through *derivatives of Gaussian*, to implement the Canny edge detector in the log-polar domain (Canny_d), which is equivalent to implement it in the Cartesian domain if the following rules are satisfied:

- Local filtering operations in the log-polar domain are a good approximation of the same filtering done in the Cartesian domain, when log-polar mapping

is performed by considering the pixel aspect ratio γ close to 1, and the spatial support of the filters is small with respect to the size of the cortical image (e.g. less than 10%) [15].

- The gradients in the cortical domain are related to the ones in the Cartesian domain by the following relationship [16]:

$$\begin{bmatrix} R_\xi \\ R_\eta \end{bmatrix} = \frac{1}{\rho_0 a^\xi \ln(a)} \begin{bmatrix} \cos \eta & \sin \eta \\ -\sin \eta & \cos \eta \end{bmatrix} \begin{bmatrix} I_x \\ I_y \end{bmatrix}, \quad (6)$$

where I_x and I_y are the partial derivatives of the Cartesian image. Eq. 6 represents a rotation and a magnitude change of the Cartesian image gradient, thus cortical gradients R_ξ and R_η have the same properties of the Cartesian ones, and they can be used for edge detection (see Fig. 3 bottom).

The Marr-Hildreth edge detector (MH_d) is based on the Laplacian of Gaussian, which can be directly applied in the cortical domain, following the previously explained rules.

The Harris corner and edge detector (Harris_d) is based on derivatives of Gaussian that yield the cortical image structure tensor, and can be described as follows:

$$M(\xi, \eta) = w(\xi, \eta) * \begin{bmatrix} R_\xi^2(\xi, \eta) & R_\xi(\xi, \eta)R_\eta(\xi, \eta) \\ R_\xi(\xi, \eta)R_\eta(\xi, \eta) & R_\eta^2(\xi, \eta) \end{bmatrix}, \quad (7)$$

where $w(\xi, \eta)$ are Gaussian weights, and $*$ is the convolution operator. The edge and corner features are then derived, by computing:

$$H(\xi, \eta) = \det(M(\xi, \eta)) - k \text{tr}^2(M(\xi, \eta)), \quad (8)$$

where k is a scalar, whose values are in the interval $[0.04, 0.15]$ [7], $\det(\cdot)$ and $\text{tr}(\cdot)$ are the determinant and the trace operators, respectively.

By considering a small spatial support for the Gaussian weights $w(\xi, \eta)$, and the relationship between the cortical and the Cartesian gradients (see Eq. 6), the Harris detector in log-polar domain is a good approximation of the one in the Cartesian domain. Indeed, the cortical image structure tensor can be written in terms of a rotation of the Cartesian gradients, multiplied by a term that can be considered constant within the small spatial support of the Gaussian weights:

$$M(\xi(x, y), \eta(x, y)) = w(\xi(x, y), \eta(x, y)) * \frac{1}{\rho_0 a^{\xi(x, y)} \ln(a)} \begin{bmatrix} I_x^{\eta^2} & I_x^\eta I_y^\eta \\ I_x^\eta I_y^\eta & I_y^{\eta^2} \end{bmatrix}, \quad (9)$$

with $I_x^\eta = I_x(x, y) \cos \eta + I_y(x, y) \sin \eta$, and $I_y^\eta = -I_x(x, y) \sin \eta + I_y(x, y) \cos \eta$.

It is worth noting that the two coordinates' axes in the log polar domain have different meanings and range of values with respect to the Cartesian ones, and this affects the structure tensor. Nevertheless, by considering the design rules previously explained the error between the processing in the two domains is negligible.

3.2 Feature Detection Based on Embedded Processing in the Log-Polar Transform

By following the biological evidence, the visual processing is performed by networks of neurons described by their RFs. The RFs can be approximated by a filter bank that performs the desired visual processing. Thus, in this paper we propose to modify the log-polar transform by using as cortical mapping weights specific filters that perform the desired feature computation, in particular we consider as weights Laplacian of Gaussian and derivatives of Gaussian (see Fig. 2b-c).

The cortical representation $R_x(\xi, \eta)$ of the Cartesian derivative $I_x(x, y)$ (computed through derivatives of Gaussian filters), along x axis, of the image is obtained as follows:

$$R_x(\xi, \eta) = \langle g_x(x - x_0(\xi, \eta), y - y_0(\xi, \eta)), I(x, y) \rangle, \quad (10)$$

where $g_x(x, y)$ is the x -axis derivative of Gaussian operator, and $(x_0(\xi, \eta), y_0(\xi, \eta))$ is the center of each RF, similarly for the y -axis.

The Canny edge detector (Canny_{RF}) can be thus implemented by using the cortical representation of the Gaussian derivatives (see Fig. 3 top). It is worth noting that this approach does not require any approximation of the filtering stage, since $R_x(\xi, \eta)$ and $R_y(\xi, \eta)$ are the actual values of the image derivatives, represented in the log-polar domain. The same principle holds for the Marr-Hildreth (MH_{RF}) edge detector, by using the Laplacian of Gaussian as weights of the log-polar mapping.

The Harris corner and edge detector (Harris_{RF}) is now based on the following equation:

$$M(\xi, \eta) = w(\xi, \eta) * \begin{bmatrix} R_x^2(\xi, \eta) & R_x(\xi, \eta)R_y(\xi, \eta) \\ R_x(\xi, \eta)R_y(\xi, \eta) & R_y^2(\xi, \eta) \end{bmatrix}. \quad (11)$$

By considering a small spatial support for $w(\xi, \eta)$, edges and corners can be computed by using Eq. 8, since the image structure tensor is now defined in terms of the Cartesian image derivatives (mapped into the log-polar domain).

Figure 4 shows the cortical derivatives, the cortical edges, and the corresponding retinal edges for a sample image, by considering the two described approaches, and the three considered algorithms. In Section 4, we will further analyze the performance of the proposed approaches.

4 Results

The qualitative and quantitative analysis of the approaches described in this paper have been performed by using the Berkeley Segmentation Data Set and Benchmarks 500 (BSDS500), described in [1]. This dataset is an extension of the BSDS300, where the original 300 images are used for training/validation and

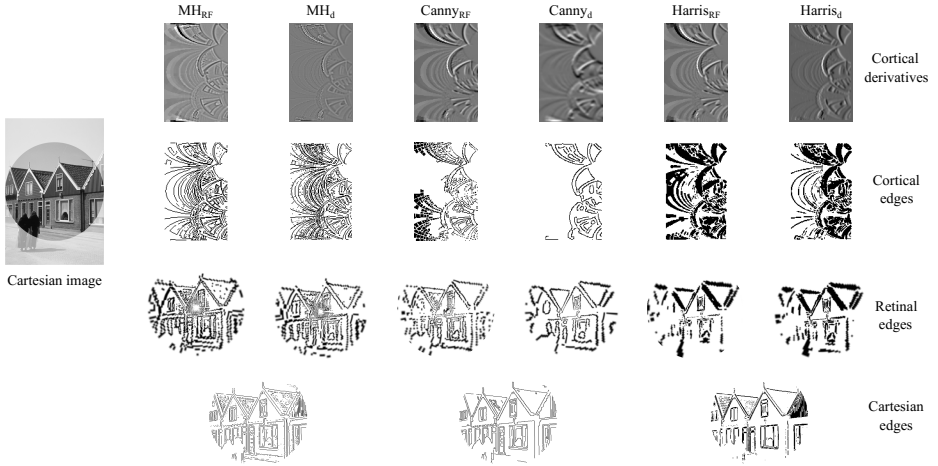


Fig. 4. Edges computed for an image of the BSDS500 dataset. The columns represent the different approaches and algorithms. First row: cortical image derivatives. Second row: cortical edges. Third row: retinal edges (computed from the cortical one through the inverse log polar mapping). Fourth row: edges computed in the Cartesian domain for comparison. Only the highlighted part of the Cartesian image is mapped into the cortical domain and analyzed (the shadow part is not considered).

200 fresh images, together with human annotations, are added for testing. Each image was segmented by five different subjects on average. In this paper, we only use and present results for the testing images. The log-polar mapping is performed by setting the transform's parameters as follow: $R = 100$, $S = 155$, $\rho_0 = 3$, and $\rho_{max} = 320$, which correspond to a compression ratio of about 10 times. Edges are computed by setting the following parameters:

- Marr-Hildreth (MH_d and MH_{RF}): the threshold of the slope of the zero-crossing values is 40. The spatial support of the Laplacian of Gaussian for MH_d is 3×3 pixels.
- Canny ($Canny_d$ and $Canny_{RF}$): the spatial support of the derivatives of Gaussian for $Canny_d$ is 9×9 pixels, and the standard deviation is 1.4. The non-maximum suppression and an hysteresis threshold (whose lower and upper bound are set to 0.3 and 0.7 of the maximum gradient value, respectively) have been considered.
- Harris ($Harris_d$ and $Harris_{RF}$): the spatial support of the derivatives of Gaussian for $Harris_d$ is 3×3 pixels. The weighting function $w(\xi, \eta)$ is a Gaussian function with spatial support 3×3 pixels and standard deviation 0.5, and k is 0.05.

A qualitative comparison of edge computation is presented in Figure 4. The approaches based on direct feature detection in the log-polar domain, and the ones based on feature detection using embedded processing in the log-polar transform perform in a similar manner. This suggests that the approximation errors




Fig. 5. Edges (green) and corners (red) computed with the Harris detector, by considering derivative of Gaussian RFs embedded in the log-polar transform (first row), and directly on the cortical image (second row). Images from the BSDS500 dataset.

between computing a feature directly in the cortical domain and by embedding the local operators into the log-polar transform are negligible. The differences among the three methods are due to the different choice of the parameters, but a systematic comparison of edge detection algorithms is out of the scope of this paper. The Harris method provides us both edges and corners, as Figure 5 shows for 4 sample images of the BSDS500 dataset. Also in this case, $Harris_d$ and $Harris_{RF}$ perform very similarly. In addition to the qualitative evaluation, we have tested the performances of the implemented algorithms in terms of precision, the fraction of true positives, and recall, the fraction of ground-truth




Fig. 6. Boundary benchmark on the BSDS500 dataset. Our implemented methods, working in the log-polar domain, have been compared with respect to the human ground-truth provided by the dataset. The precision-recall curves have been computed by using the code available from <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/resources.html>. The *F-measure* for each method is shown in the legend.

boundary pixels detected [11]. We have also computed the *F-measure*, or harmonic mean of precision and recall at the optimal detector threshold, which provides a score, useful to evaluate the methods. It is worth noting that it is not possible to compare the proposed approaches with the state-of-the-art algorithm for edge detection, this is due to the following reasons. First, we compute edges only in the part of the images covered by RFs (see Fig. 2). In particular, for the BSDS500 dataset the part of the images transformed into the log-polar domain is the highlighted one (e.g. see Fig. 4). The ground-truth edges have been masked, accordingly. Moreover, working in the cortical domain yields a loss of details in the image periphery (due to the high compression ratio that is achieved), thus the comparison with respect to methods that work at full resolution is unfair.

The precision-recall graphs are reported in Figure 6. The F-measures when choosing an optimal scale for the entire dataset have almost the same values for all the algorithms, whereas F-measures per images show that Harris_{RF} has slightly a better performance ($F = 0.54$). For the sake of clarity, it is worth noting that on the same dataset Canny algorithm (in the Cartesian domain at full resolution) has a better F-measure ($F = 0.60$).

5 Conclusion

In this paper, we have proposed two approaches for edge detection in the log-polar domain, which allow us to apply well-known and established techniques, previously developed in the Cartesian domain, without the need of adapting them to the cortical domain. We have presented the theoretical basis on which this is possible, by following two distinct approaches: one in which features are directly computed working on the log-polar images, the other in which the local differential operators at the basis of the considered algorithms are embedded into the log-polar mapping, thus allowing us to obtain a cortical representation of the Cartesian image processing.

The results confirm the validity of the approaches and suggest that they could be promising for achieving feature detection in space-variant images. Nevertheless, some issues are still open and will be further investigated. From one hand, a parametric analysis of the considered edge detectors should be performed, in order to find the optimal set of parameters, and more recent state-of-the-art algorithms for edge detection should be analyzed in order to describe them in the cortical domain. From the other hand, it is necessary to consider multiple log-polar mapping of the same image, by moving the *fovea* in several image points, in order to build a full-resolution representation of edges, thus mimicking the shifting focus of attention of the human visual system.

References

1. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(5), 898–916 (2011)
2. Bolduc, M., Levine, M.D.: A review of biologically motivated space-variant data reduction models for robotic vision. *Computer Vision and Image Understanding* **69**(2), 170–184 (1998)
3. Chessa, M., Sabatini, S.P., Solari, F., Tatti, F.: A quantitative comparison of speed and reliability for log-polar mapping techniques. In: Crowley, J.L., Draper, B.A., Thonnat, M. (eds.) *ICVS 2011. LNCS*, vol. 6962, pp. 41–50. Springer, Heidelberg (2011)
4. Florack, L.M.J.: Modeling foveal vision. In: Sgallari, F., Murli, A., Paragios, N. (eds.) *SSVM 2007. LNCS*, vol. 4485, pp. 919–928. Springer, Heidelberg (2007)
5. Forsyth, D.A., Ponce, J.: *Computer Vision: A Modern Approach*. Prentice Hall Professional Technical Reference (2002)
6. Gomes, H.M., Fisher, R.B.: Primal sketch feature extraction from a log-polar image. *Pattern Recognition Letters* **24**(7), 983–992 (2003)
7. Gueguen, L., Pesaresi, M.: Multi scale harris corner detector based on differential morphological decomposition. *Pattern Recognition Letters* **32**(14), 1714–1719 (2011)
8. Harris, C., Stephens, M.: A combined corner and edge detector. In: *Proceedings of the 4th Alvey Vision Conference*, pp. 147–151 (1988)
9. Jurie, F.: A new log-polar mapping for space variant imaging. Application to face detection and tracking. *Pattern Recognition* **32**, 865–875 (1999)
10. Marr, D.: *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Henry Holt and Co., Inc., New York (1982)
11. Martin, D., Fowlkes, C., Malik, J.: Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(5), 530–549 (2004)
12. Nattel, E., Yeshurun, Y.: Direct feature extraction in a foveated environment. *Pattern Recognition Letters* **23**(13), 1537–1548 (2002)
13. Pamplona, D., Bernardino, A.: Smooth foveal vision with gaussian receptive fields. In: *9th IEEE-RAS International Conference on Humanoid Robots* (2009)
14. Schwartz, E.: Spatial mapping in the primate sensory projection: Analytic structure and relevance to perception. *Biological Cybernetics* **25**, 181–194 (1977)
15. Solari, F., Chessa, M., Sabatini, S.P.: Design strategies for direct multi-scale and multi-orientation feature extraction in the log-polar domain. *Pattern Recognition Letters* **33**(1), 41–51 (2012)
16. Solari, F., Chessa, M., Sabatini, S.P.: An integrated neuromimetic architecture for direct motion interpretation in the log-polar domain. *Computer Vision and Image Understanding* **125**, 37–54 (2014)
17. Traver, V., Pla, F.: Log-polar mapping template design: From task-level requirements to geometry parameters. *Image Vision Computing* **26**(10), 1354–1370 (2008)
18. Wallace, A., McLaren, D.: Gradient detection in discrete log-polar images. *Pattern Recognition Letters* **24**(14), 2463–2470 (2003)