

Symmetry detection based on multiscale pairwise texture boundary segment interactions[☆]



Max Mignotte*

DIRO, Université de Montréal, C.P. 6128, Succ. Centre-ville, Montréal, H3C 3J7, Canada

ARTICLE INFO

Article history:

Received 17 August 2015

Available online 15 February 2016

Keywords:

Berkeley Segmentation Dataset (BSDS300)

Hough-style voting approach

Line segment detector (LSD)

Local symmetry detection

Medial axis segment

Pair of lines

ABSTRACT

In this paper, we propose a new unsupervised and simple approach to local symmetry detection of ribbon-like structure in natural images. The proposed model consists in quantifying the presence of a partial medial axis segment, existing between each pair of (preliminary detected) line segments delineating the boundary of two textured regions, by a set of heuristics related both to the geometrical structure of each pair of line segments and its ability to locally delimit a homogeneous texture region in the image. This semi-local approach is finally embedded in a two-step algorithm with an amplification step, *via* a Hough-style voting approach achieved at different scales and coordinate spaces which aims at determining the dominant local symmetries present in the image and a final denoising step, *via* an averaging procedure, which aims at removing noise and spurious local symmetries. The experiments, reported in this paper and conducted on the recent extension of the Berkeley Segmentation Dataset for the local symmetry detection task, demonstrate that the proposed symmetry detector performs well compared to the best existing state-of-the-art algorithms recently proposed in the literature.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

The problem of extracting automatically symmetry axes or contours lying in the middle of elongated structures (also referred to as ribbons or ridges in the literature) from natural image is an important and still difficult mid-level computer vision task. This pre-processing step, amounting to approximate reflective symmetry of each object present in the scene, can thus be the preliminary pre-attentive and/or useful complementary step of many high-level image understanding algorithms and computer vision systems such as object localization/recognition [7], 3D reconstruction [2], indexing [11], gait analysis [27], human action recognition [33] or motion estimation problems [13,14] to name a few.

The concept of symmetry both plays an important role in human visual and perceptual system (which turns out extremely sensitive to symmetry structures) and is a prolific phenomenon that exists widely in the real world. Indeed, most man-made and biological objects exhibit symmetry to some extent. That is why symmetry is an important mid-level cue, which is receiving a growing attention in image analysis due to its inherent ability to reveal the shapes of salient structures in medical applications

[30], and also a powerful concept that facilitates object detection (or localization) and enhance recognition in many image understanding systems in which some prior knowledge of a shape to be recognized can be (jointly or subsequently) included [7,15,35]. In addition, this symmetry-based mid-level cue can be efficiently used as semi-global constraints to low-level image tasks such as segmentation and denoising problems (e.g., by simply favoring label or gray-level homogeneity for pair of sites sharing the same, preliminary detected, axis of symmetry [25] or for regulating a nonlinear diffusion process for adaptive image enhancement and denoising in fingerprint imagery [1]. This symmetry-based cue can also be exploited as a feature to improve a real-time tracking procedure of moving objects [19] such as hand tracking or a hand posture recognition system [3]. Besides, since symmetry cues are one of the possible principles of the Gestalt laws, this mid-level feature could be useful to detect the other high-level cues of the Gestalt laws of perceptual organization. In this work, we aim at extracting local symmetry axes directly from the unsegmented image. By this fact, this work is different from the silhouette-based symmetry detection approach first relying on a reliable map of segmented shapes.

In the case of the symmetry detection directly from the image, a possible approach is to use a global approach, processing the entire image as a signal from which symmetric properties are inferred often *via* frequency or Fourier analysis [16]. However, these

[☆] This paper has been recommended for acceptance by Egon L. van den Broek.

* Tel.: +1 514 343 5747; fax: +1 514 343 5834.

E-mail address: mignotte@IRO.UMontreal.ca, mignotte@iro.umontreal.ca

global approaches are limited to the detection of a single axis of symmetry and are biased by the background structure of the images.

An alternative to these global approaches consists in exploiting a local pixel or feature-based approach and more precisely, either the idea of quantifying the symmetry contribution of each (preliminary detected) edge points [10] or (its soft version, *i.e.*) each pair of pixels [31], weighted by their intensity gradient amplitude (in order to favor pairs of edge points belonging to each object contour). In the same spirit, another strategy consists of matching symmetric pairs of feature points (local features such as edge features, contours, boundary points or SIFT/SURF feature points). The “amount” of symmetry exhibited by each pair is quantified by the relative location, orientation and scale of the features in the pair. These pairwise symmetries are then either analytically determined or accumulated in a Hough-style voting space (or in an axis-parameter space) to determine the dominant symmetries present in the image [17,21]. Nevertheless, it is worth mentioning that the reliability of these local features-based techniques suffers if the feature point extraction step fails to generate sufficient number of relevant feature points, which is the case of input images with a smooth image with no texture. Besides, these latter approaches are more suited to detect global reflection symmetries, associated with each whole object present in the scene, rather than local symmetries, associated to each part of an object [28,35].

Finally, another way consists in combining a local and global or a multi-scale approach. In this latter approach, the authors in [29] generalize the medialness notion (which measures the degree of belonging of the point to the medial axis of the object) and defined it as a convolution product of the initial image with a kernel, based on normalized Gaussian derivatives of intensity. In this optic, a recent machine learning framework exploiting local and global spectral features (on several scales) has also been recently proposed in [34] to train a symmetry detector on a symmetry ground truth dataset constructed on top of the well-known Berkeley Segmentation Dataset (BSDS300) [22] and using human-assisted segment skeletonization. The availability of this new dataset (herein further referenced as the LS-BSDS300 for Local Symmetries in the BSDS300) has allowed, up to now, to measure, with good statistical confidence, the performance of several symmetry detectors proposed in the literature. Amongst the methods validated on this new dataset, we can cite the following multi-scale and/or semi-local approaches; Lindeberg [20] introduces a multi-scale ridge detector by combining multiple eigenvalues of the image Hessian. Levinshtein et al. [18] propose to detect partial medial axes after superpixel segmentation and learn an affinity function to perceptually group adjacent superpixels belonging to the same object which will then be approximated by fitting ellipses from which he finally retrieves the major axes. Finally, let us mention the statistical framework proposed in [35] in which the symmetry detection problem is formulated as a spatial Bayesian tracking task using a particle filtering approach and an adaptive semi-local geometric model of ribbons.

In this paper, we propose a semi-local approach which consists in exploiting the idea of quantifying the presence of a local symmetry axis segment existing between each pair of preliminary detected straight line segments delineating the boundary of two textured regions in the scene. This preliminary detection task is efficiently carried out, in our application, by an adaptive, accurate and fast (linear-time) Line Segment Detector (LSD) (introduced by Goi et al. [9]) and optimal according to the perception-based a contrario (or Helmholtz) principle (which is also adaptive to image content) and followed by a filtering process aiming to select lines that delineates the boundaries between different textured areas. The likelihood of the presence of a partial medial axis segment, exhibited by each pair of lines, is estimated from a set of

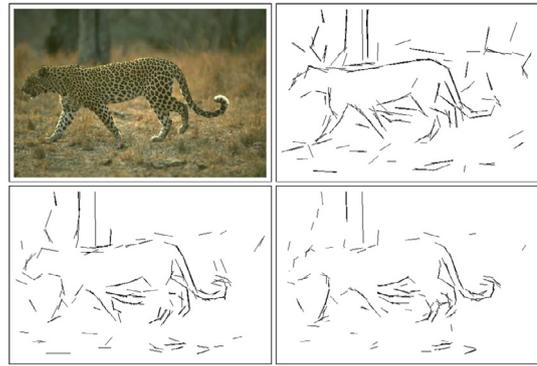


Fig. 1. From lexicographic order: original image number 134052 of the BSDS300 and Line Segment Detection ($N_l = 70$ detected straight line segments) at three different resolution levels (namely {0.25, 0.625, 1.0} filtered by the proposed soft edge potential map in the texture sense in order to detect only line segments that delineate the boundaries between different textured areas.

heuristics related both to the geometrical structure of each pair of lines and its ability to locally delimit a homogeneous texture region in the image. This semi-local approach is finally embedded in a two-stage algorithm with an amplification procedure combining both a Hough-style [8] voting scheme, achieved in different coordinate spaces, and a denoising step based on an averaging technique. The Hough-like voting scheme aims at determining the dominant local symmetries present in the image while the averaging procedure helps to remove noise and spurious local symmetries.

2. Proposed model

2.1. Texture boundary segment extraction step

Our symmetry detection approach relies mainly, first, on the reliable extraction, at different resolution scales, of the different line segments existing in the image. More conceptually, this preliminary step thus relies on the fact that man-made objects and many shapes accept an economic description in terms of straight lines. In addition, this representation style, in terms of main straight lines, is also interesting due to its appealing ability to abstract away unwanted (or useless) details, to clarify or simplify shapes and to focus on relevant features while allowing to offer a compact representation of the geometric content of the scene with few line segments [23]. To attain this goal, we rely on the fast Line Segment Detector (LSD) recently introduced in [9] which is a linear-time line segment detector that gives sub-pixel accurate detection results of a controlled pre-determined number of false detections according to the *a contrario* principle. It is worth mentioning that our symmetry detection procedure remains independent of the choice of the underlying line detector as far as the latter turns out both, as the LSD detector, reliable and fast. Indeed, in contrast to classic edge detector, this algorithm defines a line segment as a rectangular region whose points share roughly the same image gradient orientation (within a certain tolerance). In addition, the LSD algorithm has the appealing property to achieve a line segment detection for a given resolution level of the input image¹ (see Fig. 1). This multiscale analysis capability of the LSD algorithm will be fully exploited in our symmetry detection approach. This algorithm is very efficient to detect line segments based on the empirical discovery made by Burns et al. [4] showing that connected

¹ The scale of analysis is a choice left to the user and thus an (interesting and important) internal input parameters of the LSD algorithm. More precisely, in the LSD algorithm, this scaling is performed by a Gaussian sub-sampling (the image is in fact filtered with a Gaussian kernel to avoid aliasing and then sub-sampled).

Table 1
The estimation of the bin descriptor $\mathbf{h}_{x,y}$ written in pseudo-code.

Estimation of the bin descriptor $\mathbf{h}_{x,y}$	
Input:	Input RGB image I and location (x,y)
Output:	$\mathbf{h}_{x,y}$ bin descriptor or normalized re-quantized local color histogram estimated at $s=(x,y)$ (array of N_b floats $[h_{x,y}^{(1)}, \dots, h_{x,y}^{(N_b-1)}]$)
N_b	Bin descriptor length ($=q^3$, with q integer)
N_s	Set of pixels within the squared $N_w \times N_w$ neighborhood region centered at location $s=(x,y)$ of I
$\lfloor \cdot \rfloor$	Integer part of \cdot
for each pixel $p \in N_s$ with color value R_p, G_p, B_p do	
•	$k \leftarrow q^2 \cdot \lfloor \frac{R_p}{256} \rfloor + q \cdot \lfloor \frac{G_p}{256} \rfloor + \lfloor \frac{B_p}{256} \rfloor$
•	$h_{x,y}^{(k)} \leftarrow h_{x,y}^{(k)} + 1/N_w^2$

(rectangular) regions with common orientation would almost always coincide with straight edges.

However, by locating line-support regions formed by grouping adjacent pixels sharing the same gradient angle, the LSD algorithm is able to only detect areas of strong discontinuities between uniform regions (in the grey or color level sense), but not between two textured regions. In our case, in order to only keep significant lines that delineates the boundaries between different textured areas (i.e., to avoid many false line detections), we have to filter out irrelevant lines resulting from the previous LSD detection stage. To this end, we resort to a simple filtering approach exploiting a preliminary estimation of a soft edge map, in the texture sense. One possibility, which we exploit here, is the simple texture gradient based edge map proposed in [24]. Our filtering process then consists in simply estimating the average potential value located on each detected straight lines and to retain the best N_L line segments. This potential map of texture edge requires to compute, at each pixel, the following potential or gradient distance:

$$\frac{1}{2} \{ \mathcal{D}(\mathbf{h}_{x,y-(N_w/2)}, \mathbf{h}_{x,y+(N_w/2)}) + \mathcal{D}(\mathbf{h}_{x-(N_w/2),y}, \mathbf{h}_{x+(N_w/2),y}) \} \quad (1)$$

where \mathbf{h} is the N_b -bin re-quantized local color histogram, located at row y and column x (see Algorithm 1), and $\mathcal{D}(\mathbf{h}_{x-d,y}, \mathbf{h}_{x+d,y})$ is the L_1 norm between vectors (or bin descriptors) $\mathbf{h}_{x-d,y}$ and $\mathbf{h}_{x+d,y}$ computed on a squared N_w -size window centered respectively at location $(x-d,y)$ and $(x+d,y)$ [24]. This textural gradient map, normalized to lie within 0 and 1, assumes simply that the repetitive character or element of a textured image (i.e., a *texton*) is herein characterized by a mixture of colors (or more precisely by the values of the re-quantized color histogram).

At this stage, we have a sequence of N_L straight line segments (delineating two texture regions) at N_c different resolution scales (see Fig. 1). Each line segment l_i ($i < N_L$) is associated with a mean potential value \mathcal{P}_i (lying within [0 1]) which defines a confidence level that the line segment belongs to a texture contour.

2.2. Symmetry detection step

These candidate line segments can be efficiently further exploited, in a second step, as semi-global low-level features to extract the different partial medial axis in an image. More precisely, the second step of our model consists in quantifying the presence of a partial medial axis segment, existing between each pair of (preliminary detected) line segments, delineating the boundary of two textured regions, by a set of heuristics related both to the geometrical structure of each pair of line segments and its ability to locally delimit a homogeneous texture region in the image.

To attain this goal, the amount of symmetry exhibited by each pair of pixels belonging to two different texture boundary segments is accumulated in a voting space, via an accumulation matrix, $SY(s)$ congruent with the image domain [6], and quantify-

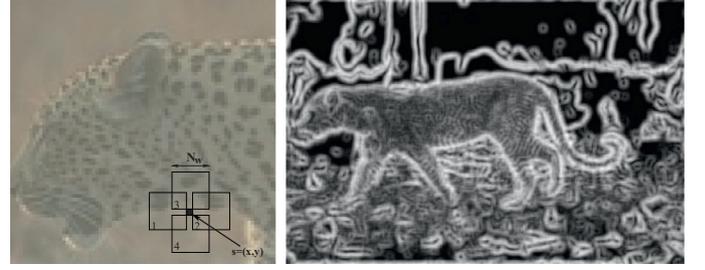


Fig. 2. Left; the potential map of texture edge is estimated for each pixel of the RGB image located at coordinates $s = (x,y)$ by computing first, the normalized re-quantized local color histograms (or descriptor bin) \mathbf{h} (see Table (1)) within the squared $N_w \times N_w$ windows respectively located at coordinates $(x - (N_w/2), y)$, $(x + (N_w/2), y)$, $(x, y - (N_w/2))$, $(x, y + (N_w/2))$ (i.e., respectively, within the window number 1, 2, 3 and 4) and second, by the formula $0.5 \cdot (|h_3 - h_4| + |h_1 - h_2|)$ (see Eq. (1)) which exploits the local spatial gradient information between these adjacent re-quantized local color histograms. Right; example of potential map of texture edge obtained on the bottom left image of Fig. 1. The filtering process consists in selecting the best (subset of) N_L straight lines, detected by the LSD, with the highest average potential (integrated on each point belonging to the line segment).

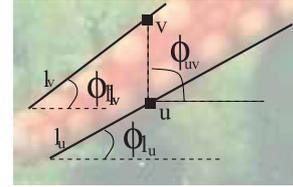


Fig. 3. A crop of the image number 12.003 of the BSDS300 and the angles $\phi_u, \phi_v, \phi_{u,v}$ used in the symmetry potential terms $\Theta_{u,v}$ associated to the geometrical structure of the pair of line segments l_u, l_v passing through the pixel u and v .

ing, after accumulation, the symmetry magnitude located at pixel $s = (x,y)$:

$$SY\left(\frac{u+v}{2}\right) = \sum_{\langle u,v \rangle} \mathcal{P}_{u,v} \Upsilon_{u,v} \mathcal{L}_{u,v} \mathcal{D}_{u,v} \Theta_{u,v} \quad (2)$$

where the summation is over all pairs of pixels (u,v) (see Fig. 3) belonging to two different line segments and $(u+v)/2$ corresponds to the midpoint of the line segment joining the pixels u and v . $\mathcal{P}_{u,v}$ and $\Upsilon_{u,v}$ are two symmetry potential terms related to the ability of each pair of line segments (i.e., the line segment passing through the pixel u and the other passing through v) to locally delimit a homogeneous texture region in the image. More precisely,

$$\mathcal{P}_{u,v} = \frac{1}{2} (\mathcal{P}_{l_u} + \mathcal{P}_{l_v}) \quad (3)$$

$$\Upsilon_{u,v} = (1 + \Gamma^2(u,v))^{-1} \quad (4)$$

with \mathcal{P}_{l_u} is a confidence value that the line segment passing through the pixel u (i.e., l_u) belongs to a texture contour (see Section 2.1 and Fig. 2). Moreover $\Gamma(u,v)$ corresponds to the number of texture boundary segments crossed by the line segment joining the pixels u and v . Therefore, the overall symmetry score also increases if there are no intervening texture boundaries across the line passing through u and v . Finally, $\mathcal{L}_{u,v}$, $\mathcal{D}_{u,v}$ and $\Theta_{u,v}$ are three symmetry potential terms related to the geometrical structure of each pair of line segments (i.e., the line segment passing through the pixel u and the other passing through v) defined as:

$$\mathcal{L}_{u,v} = \frac{(\text{Lg}[l_u] + \text{Lg}[l_v])}{2} \quad (5)$$

$$\mathcal{D}_{u,v} = \exp\left\{-\frac{\|u-v\|_2}{\sigma \bar{S}}\right\} \quad (6)$$

where $\text{Lg}[l_u]$ is the length (in pixels) of the line segment l_u and \bar{s} is the mean size of the image (at the full original resolution level). These two heuristics aim to favor a symmetry axis between u and v (by increasing the overall score) if the texture boundary contours l_u and l_v are long and the distance between u and v is close.

$$\Theta_{u,v} = \begin{cases} 1 - \cos(\phi_{l_u} + \phi_{l_v} - 2\phi_{u,v}) & \text{if } \phi_{l_u, l_v} < \tau \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

The symmetry potential term $\Theta_{u,v}$ is adapted from the phase weight function proposed in [31] and aims at favoring (by increasing the overall score) a symmetry axis between u and v if the gradients at pixels u and v are oriented in the same direction towards each other. In our application, ϕ_{l_u} , $\phi_{u,v}$ are respectively the angles of l_u and the angle of the line passing through the pixel u and v relatively to the x -axis (see Fig. 3). ϕ_{l_u, l_v} is the angle between the line segments l_u and l_v . τ is a parameter of the model which will be later estimated by a trial-and-error procedure on the training dataset (see Section 3).

In summary, the symmetry magnitude, located at mid-point of the line segment joining the pixels u and v , will be all the more high if the following conditions are met: (1) the texture boundary contours l_u and l_v delineate two different texture regions in the image (Eqs. (3) and (4)) and are long (Eq. (5)). (2) The distance between u and v is close (Eq. (6)). (3) The gradients at pixel u and v are oriented in the same direction towards each other and finally the segment $[u, v]$ is perpendicular to the medial axis of l_u and l_v (Eq. (7)).

2.3. Detection amplification/denoising step

In our application, let us note that the performance of the proposed symmetry detection algorithm is highly dependent on the first crucial step which aims at extracting reliable texture boundary segments in the input image (Section 2.1). In order to make more robust this detection step, we resort to two strategies for improving the final result of symmetry detection. These two strategies have ultimately and conceptually the same goal: to provide a more accurate final estimation of the symmetry map, by combining (with a simple averaging process) the set of estimates obtained when the original input image is represented at different resolution levels and different coordinate spaces (an estimator based on the average operation generally yields to an optimal denoised solution when the noise is uncorrelated). More precisely and as already said in Section 2.1, the first one consists in detecting the set of candidate line segments at N_c different resolution scales, thus simulating and taking into account edge detection at different detail levels in the image. This strategy can be easily implemented since the scale factor is already an intrinsic internal parameter of the LSD algorithm.

The second strategy, already used in [26] for the edge detection problem, is to detect the set of candidate line segments in different coordinate spaces, especially in log-polar domains w.r.t. N_p different equally spaced centers of reference (i.e., points of interest) within the input image. Each center of reference thus defines a log-polar mapping associated with a different point of view within the image plane. This sequence of polar mappings will possibly transform, for one of them, a given curved contour fragment into a straight line segment which will be then efficiently detected by the LSD algorithm² (see Fig. 4). In order to detect a curved con-

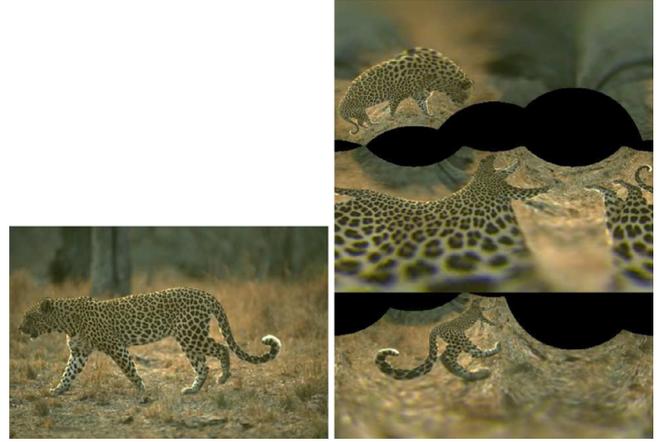


Fig. 4. From lexicographic order: original image number 134,052 of the BSDS300 and images mapped in three different log-polar domains w.r.t. three different image centers.

tour that delineates the boundary between different textured areas, we also use the filtering step (described in Section 2.1) computed in this particular polar domain.

For each of the N_L straight line segments detection achieved in a polar coordinate space for a given reference center (at N_c different resolution scales), we compute, in the same way described in Section 2.2) a new matrix S_{p_c} thus quantifying the symmetry magnitude located at pixel expressed in polar coordinate (for a given center of reference c). We finally carry out an inverse log-polar transformation of this accumulation matrix and a simple averaging technique, over all these N_p log-polar transformations, allows us to obtain, in the Cartesian domain a final symmetry accumulation map which will be averaged with the one directly obtained in the Cartesian domain. This finally defines a final soft symmetry map and a non-maximum suppression step [5] is then employed to produce thin symmetry axis (see Fig. 5 and Algorithm 3).

3. Experimental results

3.1. Parameters

In order to limit the computational requirement of our algorithm, we set the number of resolution scales to $N_c = 3$ (from the resolution level R_{\min} to the full resolution level) and the number of log-polar coordinate spaces to $N_p = 10$ (w.r.t. N_p different equally spaced centers of reference taken on the diagonal of a rectangular window, centered on the image). We will see, later in this section, that the performance measure of good symmetry detection is all better than N_p is high (up to a certain extent).

In order to ensure the integrity of the evaluation, all the internal parameters of our model are tuned on the training dataset of the LS-BSDS300, by doing a local discrete grid search, with a fixed step-size, on the parameter space and in a given range of parameter values; namely, for the number of straight lines detected at each coordinate space (i.e., Cartesian or log-polar w.r.t. different center of reference) and for each color channel, $N_L \in [20: 100]$ [step-size: 10], $\sigma \in [0: 1.0]$ [step-size: 0.1] (see Eq. (6)), $\tau \in [0: \pi]$ [step-size: 0.5] (see Eq. (7)), $R_{\min} \in [0.1: 1.0]$ [step-size: 0.05]. We have finally found that $N_L = 70$, $\sigma = 0.1$, $\tau = 2$ and $R_{\min} = 0.25$ allows us to ensure a maximal F-measure score on the training set (which was sensibly the same score of the one obtained on the test dataset).

² It is worth mentioning that a fix number N_p of log-polar maps can not obviously transform all long contour fragments into straight lines. When the polar transforms are not able to ensure this property, these contour fragments are just approximated by concatenations of small straight segments in the log-polar domain. Indeed, at least locally; a smoothly curved contour will be often efficiently approximated by concatenations of small straight segments.

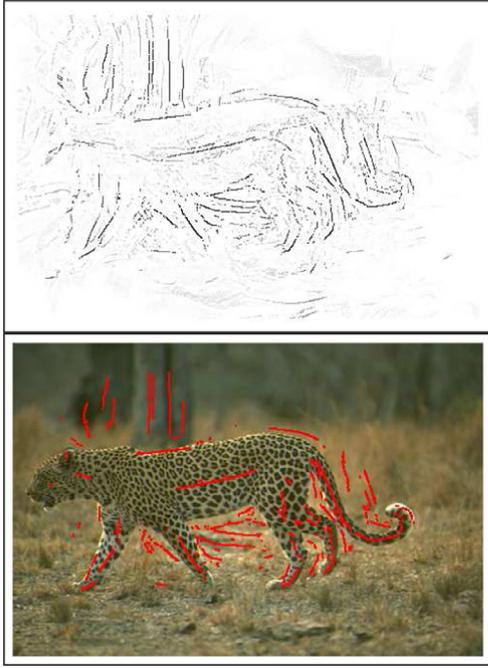


Fig. 5. Symmetry magnitude map result (in video inverse) obtained by the proposed algorithm on the Berkeley image number 134,052 (see Fig. 1) and symmetry detection results super-imposed on the original image (in red) obtained after a non-maximum suppression step [5] and the global optimal thresholding operation ensuring the maximal F-measure on the LS-BSDS300 dataset. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2

Comparison of scores (i.e., the score based on the optimal threshold obtained from each image) between the BISDA and the results obtained only in the Cartesian space and those only obtained in the N_p log-polar spaces, on the 89 images of the test set of the LS-NSDS300. Each value points out the number of images that obtain the best F score.

	Cartesian	Log-polar	BISDA
Number of best	24	15	50
F scores	59	30	-

3.2. Initial test

If we consider the symmetry magnitude map directly estimated with the Cartesian image, we obtain a F-measure³ score of 0.408 (recall: 0.51, precision: 0.33). If we consider only the average symmetry magnitude map obtained in different log-polar spaces, we obtain an F-measure score of 0.373 (recall: 0.47, precision: 0.31). By averaging the last two symmetry magnitude maps, we obtain the proposed symmetry detection model and an F-measure score of 0.422 (recall: 0.521, precision: 0.354 for the threshold: 0.142) (see Fig. 7).

In order to study the complementarity of each coordinate space, we have recorded the number of individual results obtaining the best F score (i.e., the score based on the optimal threshold obtained from each image) for respectively the three cases previously mentioned. (i.e., the symmetry detection results obtained respectively on the Cartesian space, the log-polar spaces and those obtained by the averaging of the two symmetry magnitude maps pre-

³ The F-measure correspond to the harmonic mean of the precision and recall scores computed on the 89 test images of the LS-BSDS300 for an optimal global threshold.

Table 3

The local symmetry detection algorithm written in pseudo-code.

BISDA Algorithm	
Input:	Input RGB image $I(x, y)$
Output:	A symmetry image $SY^*(x, y)$
N_p	Nb. of log-polar transf. (or centers of refer. c)
N_c	Nb. of scales for the LSD (from R_{min} to the full resolution level)
N_L	Nb. of considered line segments of the LSD with the highest potential value \mathcal{P}_i
R_{min}	Minimal resolution level for the LSD
σ, τ	Parameters of the heuristic constraints
LSD _{sc}	Line Seg. Detect. (with scale parameter sc) gives in output a set of line segments \mathcal{L}
$\mathcal{P}_{D_c(I)}$	Log-polar map of I , w.r.t. the ref. center c
$\mathcal{P}_{D_c^{-1}(\cdot)}$	Inverse log-polar w.r.t. the ref. center c
\mathcal{P}_J	Textural edge potential map of the image J
\mathcal{P}_i	Potential ($\in [0, \dots, 1]$) of a line segment l_i ($i < N_L$) of its belonging to a texture contour
$SY_{\mathcal{P}_c}$	Accumulation matrix quantifying the symmetry magnitude of $\mathcal{P}_{D_c(I)}$ in Polar domain
SY_{C_c}	idem in Cartesian domain
SY	Global accum. matrix in Cartesian domain
1. Initialization Step:	
• Build a list of N_p points $\{c_1, \dots, c_{N_p}\}$ taken on the diagonal of a rectangular window centered on I	
• $k \leftarrow 1$	
2. Detection/Amplification/Denoising Step:	
while $k \leq N_p$ do	
• Compute $\mathcal{P}_{D_c(I)}$ & its text. edge potential $\mathcal{P}_{\mathcal{P}_{D_c(I)}}$	
for each color channel and each scale factor sc do	
$\mathcal{L} \leftarrow$ LSD _{sc} on $\mathcal{P}_{D_c(I)}$	
$\mathcal{L}_{N_L} \leftarrow$ Select the N_L line segments $\in \mathcal{L}$ with the highest edge potential \mathcal{P}_i (using $\mathcal{P}_{\mathcal{P}_{D_c(I)}}$)	
for each pixel pair $\langle u, v \rangle \in$ line $l \in \mathcal{L}_{N_L}$ do	
$SY_{\mathcal{P}_c} \left[\frac{u+v}{2} \right] = \sum_{\langle u,v \rangle} \mathcal{P}_{u,v} \Upsilon_{u,v} \mathcal{L}_{u,v} \mathcal{D}_{u,v} \Theta_{u,v}$	
$SY_{C_c} \leftarrow \mathcal{P}_{D_c^{-1}}(SY_{\mathcal{P}_c})$	
$SY \leftarrow$ ACCUMULATE(SY_{C_c})	
• $k \leftarrow k + 1$	
• $SY \leftarrow$ AVERAGING (SY & SY^C) with SY^C the accumulation matrix obtained by step 2. (while ... do loop is only executing once) in Cartesian domain	
3. Non Maximum Suppression Step:	
• $SY^*(x, y) \leftarrow$ NON MAX. SUPPRESSION (SY)	

viously obtained) and also recorded the number of best obtained scores between only the first two cases and summarized all the results in Table 2. Based on these results, we can conclude that the symmetry detection strategy, in each coordinate space, performs differently and well for different images. This is not surprising since a textured color image can really be visually different in

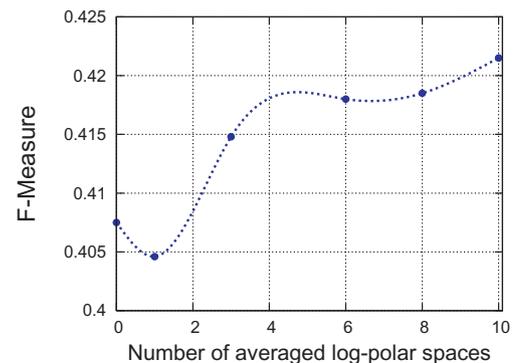


Fig. 6. Evolution of the F-score as a function of the number of log-polar transforms (the other parameters of the model have been kept identical).

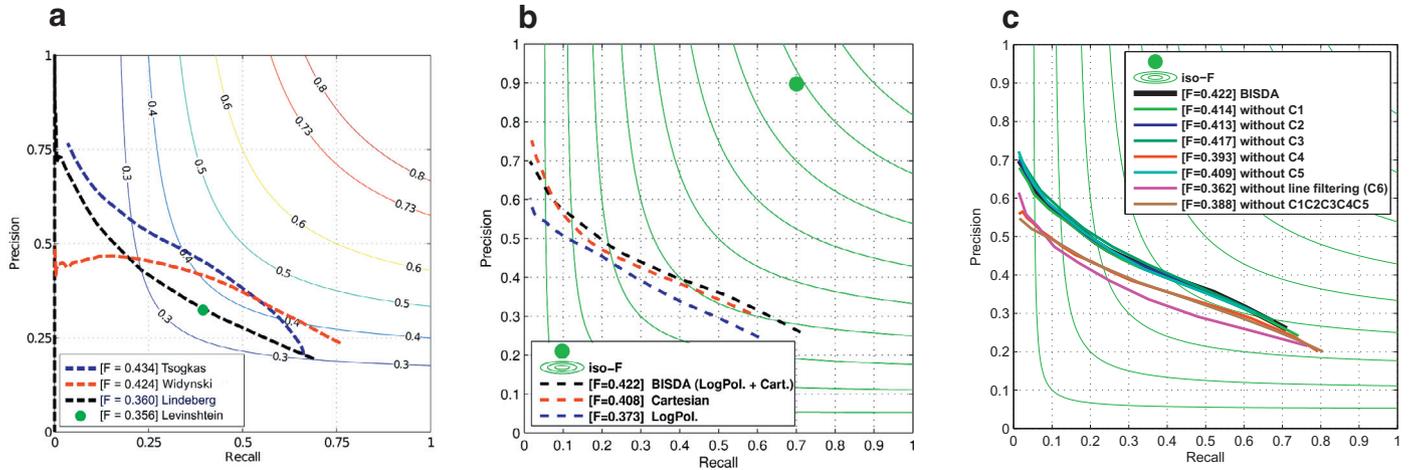


Fig. 7. (a): F-measure scores by several other state-of-the art symmetry detection methods on the LS-BSDS300 (illustration is adapted from [35]). (b): F-measure scores for different variants of our approach. Namely; F-measure curve by considering; (1) only the symmetry magnitude map directly estimated on the Cartesian image, (2) only the average symmetry magnitude map obtained by log-polar inverse transformations of the N_p different log-polar transformations of the Cartesian image, (3) the proposed symmetry detection model, obtained by averaging the two symmetry magnitude maps obtained in (1) and (2) (in cases 1 and 2, the other parameters of the model have been kept identical). (c): F-measure scores obtained by suppressing one of the six heuristics (constant factor equal to 1 for, respectively, $\mathcal{P}_{u,v}$ (C1, Eq. (3)), $\Upsilon_{u,v}$ (C2, Eq. (4)), $\mathcal{L}_{u,v}$ (C3, Eq. (5)), $\mathcal{D}_{u,v}$ (C4, Eq. (6)), $\Theta_{u,v}$ (C5, Eq. (7)), and without the filtering process) and F-measure scores obtained by considering a Hough-like accumulation (see Eq. (2)) in the voting space without the heuristics C1, C2, C3, C4 and C5, (i.e., $\mathcal{P}_{u,v} = \Upsilon_{u,v} = \mathcal{L}_{u,v} = \mathcal{D}_{u,v} = \Theta_{u,v} = 1$ but with the line filtering process).

the two spaces (see Fig. 4) and consequently the symmetry detection scheme can provide significantly different estimation results. It also means that the symmetry responses in the log-polar transformed space extract complementary information which can be efficiently combined with the Cartesian space or, otherwise said, that the log-polar transformed space complements the results obtained in the Cartesian space⁴

We have also evaluated the performance of our algorithm as a function of respectively, the number of log-polar transforms (the other parameters of the model have been kept identical). Results are summarized in Fig. 6 and shows the benefit of taking into account (up to a certain extent) several detections of texture boundary segments in different log-polar coordinate spaces⁵

Finally, in order to evaluate the contribution and importance of each separate heuristic (Eqs. (3)–(7)) and the filtering process (see Section 2.1 wherein we propose to select the N_L line segments delineating, at best, the boundary of two textures regions) to the overall performance of our symmetry detection model, we have compared the performances, in terms of F-measure and precision-recall curves, on the same plot, obtained for different combinations and different subsets of the heuristics of the proposed model. More precisely, in the first set of experiments, we have suppressed one of the six heuristics (constant factor equal to 1 for respectively, $\mathcal{P}_{u,v}$, $\Upsilon_{u,v}$, $\mathcal{L}_{u,v}$, $\mathcal{D}_{u,v}$, $\Theta_{u,v}$, and without the above-mentioned filtering process) and estimate the resulting precision recall curve and F-score (see Fig. 7(c)). In addition, we have also evaluated the F-score by considering a Hough-like accumulation (see Eq. (2)) in the voting space without the heuristics C1, C2, C3, C4 and C5, (i.e., $\mathcal{P}_{u,v} = \Upsilon_{u,v} = \mathcal{L}_{u,v} = \mathcal{D}_{u,v} = \Theta_{u,v} = 1$ but with the line filtering process). From these experiments, we can conclude that, the more eff-

icient and discriminant heuristic remains (by increasing order of significance) the line filtering process, $\mathcal{D}_{u,v}$ and $\Theta_{u,v}$. On the other hand, the heuristics $\mathcal{P}_{u,v}$, $\Upsilon_{u,v}$ and $\mathcal{L}_{u,v}$ provide (comparatively) minor improvements ($\mathcal{P}_{u,v}$ brings minor improvement certainly because the line filtering process has already retained the most significant texture boundaries).

3.3. Comparison with state-of-the-art methods

We now compare our approach with other state-of-the-art local symmetry detection methods, particularly the ones that have been already validated on the LS-BSDS300 [18,20,34,35]. We can notice (see Fig. 7), that our symmetry detection model (F-measure = 0.422) compares favorably to the one proposed by Levinshstein et al. [18] (F-measure = 0.356) or by Lindeberg [20] (F-measure = 0.360) and the proposed method achieves a score nearly identical to the approach proposed by Tsogkas [34] (F-measure = 0.434) or identical to the method proposed by Widynski (once rounded to the nearest thousandth) [35] (F-measure = 0.422). Fig. 8 illustrates a few symmetry detection results obtained by the proposed algorithm after a non-maximum suppression step and the optimal thresholding operation. Finally, Fig. 9 shows qualitative comparison and symmetry detection results for the proposed method and four other state-of-the-art local symmetry detectors (from left to right: [18,20,34,35] and the proposed method).

Let us also remark that the methods proposed by Widynski et al. [35] or Tsogkas and Kokkinos [34] are, by nature, very different from the method proposed in this paper (see the Section 1). It means that these three methods perform differently and well for different images and thus extract complementary symmetry information and consequently, could be efficiently combined together to achieve better results.

3.4. Algorithm

The proposed symmetry detection procedure takes, on average (per image), 1 min for an Intel®Core i7, 3.2 GHz, 6403 bogomips

⁴ A good example is given by the symmetry detection result obtained on the image number 300,091 for which the F-score greatly improves from $F = 0.619$ (in the Cartesian space) to $F = 0.755$ after averaging with the log-polar response (these images are available on the website of this paper).

⁵ Above $N_p = 10$ log polar maps, the F-measure score reaches a plateau and fluctuates slightly around $F = 0.420$. This phenomenon may be due to the sampling of the reference centers of the polar coordinate space which are spatially distributed on the diagonal of the image. A potential source of enhancement of the model would be to adaptively choose the N_p appropriate centers of reference within the image (possibly with selected SIFT points) or with a better sub-sampling strategy.



Fig. 8. Example of several local symmetry detection results obtained by our algorithm and super-imposed (in red) on some images from the BSDS300. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

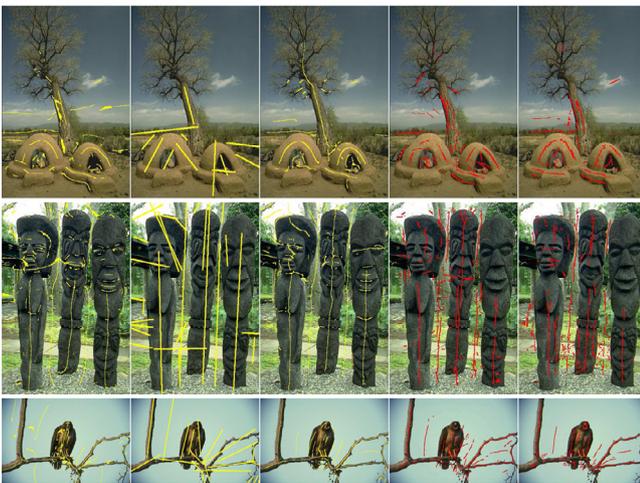


Fig. 9. Qualitative comparison and symmetry detection results against three other state-of-the-art local symmetry detectors (from left to right: [18,20,34,35]) and the proposed method.

and non-optimized code running on Linux⁶. Let us add that our algorithm and especially the set of log-polar transforms, along with the line segment detection process can be easily computed in parallel, for example by using the parallel abilities of a graphic processor unit (GPU) (embedded on most graphics hardware nowadays available on the market) [12]. The code, data, and all that is necessary for reproduction of the results are freely accessible on the author's website (directory: ResearchMaterial/bisda.html).

4. Conclusion

In this paper, we propose a multi-scale and multi-coordinate estimation approach to symmetry detection in natural images. The proposed approach mainly consists in exploiting the idea of quantifying the presence of a local symmetric part of an object existing between each pair of preliminary detected straight line segments corresponding to the boundary of two textured regions in

the scene. This estimation process is made reliable due to the inherent property that man-made objects and many shapes accept an economic description in terms of straight lines and, thus, that this constraint-based and compact image representation can be used to efficiently describe the geometric content of the scene. This estimation process is also made reliable because of our multi-coordinate approach which will possibly transform, for one of the considered log-polar mapping, a given curved contour fragment into a straight line segment which will be then exploited to detect a partial medial axis in this particular log-polar domain. We have evaluated our method on the LS-BSDS300 and demonstrate both its efficiency compared to existing works on symmetry detection and also the importance of our multi-coordinate strategy. A potential source of enhancement of the model would be to adaptively choose the N_p appropriate centers of reference within the image (possibly with selected SIFT points) or with a better sub-sampling strategy. Another direction worth exploring would be to adaptively estimate the parameter N_l according to the complexity of the image or finally to add other heuristics related both to the geometrical structure of each pair of line segments and their ability to locally delimit a homogeneous texture region in the image, such as the parametric gradient profile model proposed by Sun et al. [32] that does not use the gradient value of a texture contour, but the symmetry and the monotonicity of its distribution profile.

References

- [1] A. Almans, T. Lindeberg, Fingerprint enhancement by shape adaptation of scale-space operators with automatic scale selection, *IEEE Trans. Image Process.* 9 (12) (2000) 2027–2042.
- [2] S. Benameur, M. Mignotte, F. Destremes, J.D. Guise, Three-dimensional biplanar reconstruction of scoliotic rib cage using the estimation of a mixture of probabilistic prior models, *IEEE Trans. Biomed. Eng.* 52 (10) (2005) 2041–2057.
- [3] L. Bretzner, I. Laptev, T. Lindeberg, Hand gesture recognition using multi-scale colour features, hierarchical models and particle filtering, *Proceedings of the International Conference on Automatic Face and Gesture Recognition (ICAFGR'02)*, IEEE Computer Society, Washington D.C., 2002, pp. 423–428.
- [4] J.B. Burns, A.R. Hanson, E.M. Riseman, Extracting straight lines, *IEEE Trans. Pattern Anal. Mach. Intell.* 8 (4) (1986) 425–455.
- [5] J. Canny, A computational approach to edge detection, *IEEE Trans. Pattern Anal. Mach. Intell.* 8 (6) (1986) 679–698.
- [6] L. Chmielewski, Detection of non-parametric lines by evidence accumulation: Finding blood vessels in mammograms, in: *Proceedings of the International Conference on Computer Vision and Graphics (ICCV'04)*, 2004, pp. 373–380.
- [7] F. Destremes, M. Mignotte, Localization of shapes using statistical models and stochastic optimization, *IEEE Trans. Pattern Anal. Mach. Intell.* 29 (9) (2007) 1603–1615.
- [8] R. Duda, P.E. Hart, Use of the hough transformation to detect lines and curves in pictures, *Commun. ACM* 15 (1972) 11–15.
- [9] R.G.V. Gioi, J. Jakubowicz, J. Morel, G. Randall, LSD: A fast line segment detector with a false detection control, *IEEE Trans. Pattern Anal. Mach. Intell.* 32 (4) (2010) 722–732.
- [10] J.B. Hayfron-Acquah, M.S. Nixon, J.N. Carter, Automatic gait recognition by symmetry analysis, *Pattern Recognit. Lett.* 24 (13) (2003) 2175–2183.
- [11] A. Id-Oumohmed, M. Mignotte, J.Y. Nie, Semantic-based cross-media image retrieval, in: *Proceedings of the Third International Conference on Advances in Pattern Recognition (ICAPR'05)*, in: *Lecture Notes in Computer Science (LNCS)*, vol. 3686, 2005, pp. 414–423. *Pattern Recognition and Data Mining, Proc. Part 2*
- [12] P.M. Jodoin, M. Mignotte, Markovian segmentation and parameter estimation on graphics hardware, *J. Electron. Imaging* 15 (3) (2006) 1–15.
- [13] P.M. Jodoin, M. Mignotte, J. Konrad, Statistical background subtraction methods using spatial cues, *IEEE Trans. Circuits Syst. Video Technol.* 17 (12) (2007) 1758–1764.
- [14] P.M. Jodoin, M. Mignotte, C. Rosenberger, Segmentation framework based on label field fusion, *IEEE Trans. Image Process.* 16 (10) (2007) 2535–2550.
- [15] I. Kokkinos, A.L. Yuille, Unsupervised learning of object deformation models, in: *Proceedings of the IEEE Eleventh International Conference on Computer Vision (ICCV'07)*, Rio de Janeiro, Brazil, October 14–20, 2007, pp. 1–8.
- [16] S. Lee, R. Collins, Y. Liu, Rotation symmetry group detection via frequency analysis of frieze-expansions, in: *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR'08)*, 2008, pp. 1–8.
- [17] S. Lee, Y. Liu, Curved glide-reflection symmetry detection, *IEEE Trans. Pattern Anal. Mach. Intell.* 34 (2) (2012) 266–278.
- [18] A. Levinstein, S. Dickinson, C. Smichisescu, Multiscale symmetric part detection and grouping, in: *Proceedings of the International Conference on Computer Vision (ICCV'09)*, 2009, pp. 2162–2169.

⁶ More precisely and by decreasing order, the Hough voting procedure takes $\approx 90\%$ of the total computational cost of the algorithm, the estimation of the potential maps in the texture sense: $\approx 5\%$, the segment filtering procedure: $\approx 2\%$, the log and inverse log-polar transformations: $\approx 1.5\%$ and finally the LSD procedure: $\approx 1.5\%$.

- [19] W.H. Li, L. Kleeman, Real time object tracking using reflectional symmetry and motion, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'06, 2006, pp. 2798–2803.
- [20] T. Lindeberg, Edge detection and ridge detection with automatic scale selection, *Int. J. Comput. Vis. (IJCV)* 30 (2) (1998) 117–156.
- [21] G. Loy, J.O. Eklundh, Detecting symmetry and symmetric constellations of features, in: Proceedings of the European Conference on Computer Vision (ECCV'06), 2006, pp. 508–521.
- [22] D. Martin, C. Fowlkes, D. Tal, J. Malik, A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics, in: Proceedings of the Eighth International Conference on Computer Vision (ICCV'01), 2001, pp. 416–423.
- [23] M. Mignotte, Unsupervised statistical sketching for non-photorealistic rendering models, in: Proceedings of the International Conference on Image Process (ICIP'03), 2003, pp. 573–577.
- [24] M. Mignotte, A de-texturing and spatially constrained K-means approach for image segmentation, *Pattern Recognit. Lett.* 32 (2) (2011) 359–367.
- [25] M. Mignotte, MDS-based segmentation model for the fusion of contour and texture cues in natural images, *Comput. Vis. Image Underst.* 116 (9) (2012) 981–990.
- [26] M. Mignotte, A biologically-inspired framework for contour detection, *Pattern Anal. Appl.* (2015) 1–17.
- [27] A. Moevus, M. Mignotte, J. de Guise, J. Meunier, Evaluating perceptual maps of asymmetries for gait symmetry quantification and pathology detection., in: Proceedings of the IEEE International Conference on Engineering in Medicine and Biology Society (EMBC'14), 2014, pp. 3317–3320.
- [28] M. Park, S. Lee, P.C. Chen, S. Kashyap, A.A. Butty, Y. Liuy, Performance evaluation of state-of-the-art discrete symmetry detection algorithms, in: Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR'08), 2008, pp. 1–8.
- [29] S.M. Pizer, D. Eberly, B.S. Morse, D.S. Fritsch, Zoom invariant vision of figural shape: The mathematics of cores, *Comput. Vis. Image Underst.* 69 (1) (1998) 55–71.
- [30] S.M. Pizer, S.C. Joshi, P.T. Fletcher, M. Styner, G. Tracton, J.Z. Chen, Segmentation of single-figure objects by deformable m-reps, in: Proceedings of the Fourth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'01), 2001, pp. 862–871.
- [31] D. Reifeld, H. Wolfson, Y. Yeshurun, Context free attentional operators: The generalized symmetry transform, *Int. J. Comput. Vis.* 14 (2) (1995) 119–130. Special issue on qualitative vision.
- [32] J. Sun, Z. Xu, H. Shum, Gradient profile prior and its applications in image super-resolution and enhancement, *IEEE Trans. Image Process.* 20 (6) (2011) 1529–1542.
- [33] R. Touati, M. Mignotte, MDS-based multi-axial dimensionality reduction model for human action recognition, in: Proceedings of the Eleventh Conference on Computer and Robot Vision (CRV'14), Montreal, Quebec, Canada, 2014, pp. 262–267.
- [34] S. Tsogkas, I. Kokkinos, Learning-based symmetry detection in natural images, in: Proceedings of the European Conference on Computer Vision (ECCV'12), 2012, pp. 41–54.
- [35] N. Widynski, A. Moevus, M. Mignotte, Local symmetry detection in natural images using a particle filtering approach, *IEEE Trans. Image Process.* 23 (12) (2014) 5309–5321.