



**HAL**  
open science

# Global Bilateral Symmetry Detection Using Multiscale Mirror Histograms

Mohamed Elawady, Cécile Barat, Christophe Ducottet, Philippe Colantoni

► **To cite this version:**

Mohamed Elawady, Cécile Barat, Christophe Ducottet, Philippe Colantoni. Global Bilateral Symmetry Detection Using Multiscale Mirror Histograms. 17th International Conference on Advanced Concepts for Intelligent Vision Systems (ACIVS 2016) , Oct 2016, Lecce, Italy. pp.14 - 24, 10.1007/978-3-319-48680-2\_2 . ujm-01387193v1

**HAL Id: ujm-01387193**

**<https://ujm.hal.science/ujm-01387193v1>**

Submitted on 25 Oct 2016 (v1), last revised 18 Oct 2017 (v2)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Global Bilateral Symmetry Detection Using Multiscale Mirror Histograms

Mohamed Elawady<sup>1</sup>, Cécile Barat<sup>1</sup>, Christophe Ducottet<sup>1</sup>, and Philippe Colantoni<sup>2</sup>

<sup>1</sup> Université Jean Monnet, CNRS, UMR 5516, Laboratoire Hubert Curien, F-42000, Saint-Etienne, France,

[mohamed.elawady@univ-st-etienne.fr](mailto:mohamed.elawady@univ-st-etienne.fr),

<sup>2</sup> Université Jean Monnet, Centre Interdisciplinaire d'Études et de Recherches sur l'Expression Contemporaine n° 3068, Saint-Étienne, France

**Abstract.** In recent years, there has been renewed interest in bilateral symmetry detection in images. It consists in detecting the main bilateral symmetry axis inside artificial or natural images. State-of-the-art methods combine feature point detection, pairwise comparison and voting in Hough-like space. In spite of their good performance, they fail to give reliable results over challenging real-world and artistic images. In this paper, we propose a novel symmetry detection method using multi-scale edge features combined with local orientation histograms. An experimental evaluation is conducted on public datasets plus a new aesthetic-oriented dataset. The results show that our approach outperforms all other concurrent methods.

**Keywords:** Symmetry Detection, Reflection Symmetry, Edge Features, Pairwise Similarity, Symmetry Histogram

## 1 Introduction

Symmetry is a very important measure in some computer vision applications such as aesthetic analysis [21], object detection and segmentation [18], depth estimation [20], or medical image processing [2]. This paper focuses on detecting a single bilateral symmetry axis by exploring the similarity between regions which are mirror reflected around an image axis.

Bilateral symmetry detection has been improved over decades since 1930's [9]. In particular, Loy and Eklundh [11] approach was a great milestone in the domain, thanks to the use of SIFT descriptor [10], an appropriate pairwise comparison, and Hough-voting scheme [6]. Since then, many methods [4,14,7,12,16,3] proposed some improvements and refinements of Loy's method, but still, they fail in cases where SIFT is not adapted (uniform texture, scale variation, etc.). Our idea is to exploit another type of local features, edges, which are the most important visual cues used by the human visual system to detect symmetry. Only few attempts have been done in this direction [13,5,19,1]. The general framework

for symmetry detection proposed by [5] provides promising results, but only edge fragments are used and it fails to account for texture information.

We propose a novel method for symmetry detection preserving the shape and texture information inside an image. This work contributes with: (1) introducing a local edge orientation histogram, (2) utilizing multi-scale edge extraction to exploit the full resolution image. In addition, we introduce groundtruth of a global symmetry axis inside images extracted from the large-scale Aesthetic Visual Analysis (AVA) dataset [15]. This dataset is composed of artistic photographs which involved in on-line professional photography competitions.

In section 2, we analyze related work of symmetry detection. Then we describe the details of the proposed algorithm in section 3. We present the experimental results on PSU and AVA datasets in section 4 and the conclusion in the last section.

## 2 Related Work

The baseline algorithm to detect a global bilateral symmetry axis was introduced by Loy and Eklundh [11]. The general scheme is as follows: (1) extraction of local-feature points using Lowe’s SIFT [10], (2) matching of pairs of those key-points based on the similarity of corresponding mirror SIFT descriptors, and (3) use of the pairs in Hough-voting space to find the best symmetrical candidate. Many attempts proposed refinements to Loy’s general scheme, especially Mo and Draper [14] in selecting all symmetry candidate pairs instead of finding closest matches at each point, and using less complex hough voting scheme. The improvement produced more particles than [11] for the symmetry axis, to elongate the output axis according to the global texture information of an image. These intensity-based methods depend mainly on the properties of SIFT features, and their capabilities for symmetry matching with respect to intensity and orientation in the neighboring pixels. In the presence of smooth objects with noisy background, the global symmetry information is lost due to a small number of extracted feature points inside an image.

Cicconet et al. [5] approach exploits the idea of using edge detection in a symmetry measure based on edge orientations. The method can be summarized in the following steps. First, a regular set of feature points with local edge amplitude and orientation is extracted. Second, given any pair of points, a symmetry axis candidate is defined (parametrized by angle and displacement variables) and a symmetry coefficient is associated with this axis. Third, a symmetry voting image is defined as the sum of the weighted contribution of all pairs of points for a given axis. Finally, the best candidate is extracted as the maximum over this voting image. This approach lacks neighborhood’s texture information inside the feature representation, and it is totally dependable on the scale parameter of the edge detector. Because of these drawbacks, some local symmetric characteristics are neglected.

Our method investigates Cicconet’s edge features [5] within Loy’s scheme [11] by adding neighboring-pixel information resulting a better symmetry represen-

tation in the voting space. Additionally, it overcomes the scale limitations by adding multi-scale information.

### 3 Algorithm Details

The main steps of the proposed symmetry detection algorithm is: 1) dense edge extraction, 2) multi-scale edge description with a local orientation histogram  $h_p$  as a texture feature, 3) edge triangulation to find symmetry pairs, and 4) voting based on weighted pairs to detect a global bilateral axis.

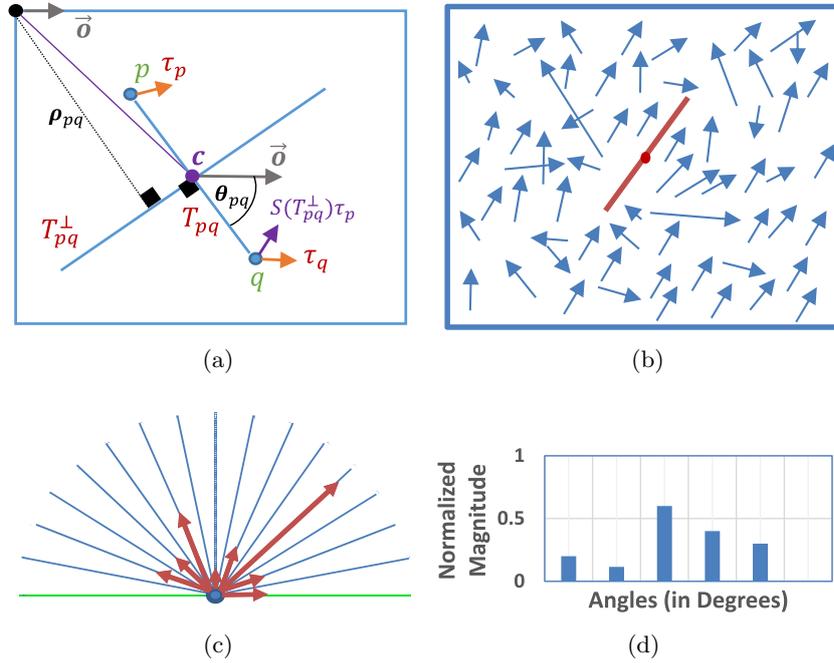


Fig. 1: Pairwise mirror symmetry and local orientation histogram. (a) Pairwise mirror symmetry. (b) Edge segment distribution within a cell. (c) Angular distribution of edge segments for 16 main orientations over 180 degrees. (d) Normalized local orientation histogram. Best seen in screen.

#### 3.1 Local Orientation Histogram

Given an image  $I$  to analyze, feature points are sampled along a regular grid (stride and cell size are computed with respect to image size). A feature point  $p$  and its local edge characteristics ( $J_p, \tau_p$ ) are extracted within each cell using a

Morlet wavelet  $\psi_{k,\sigma}$  of constant scale  $\sigma$  and varying orientation  $\{\tau_k, k = 1 \dots n\}$ . The wavelet parameters are set to detect edges: the wavelength is  $4\sigma$ , and the elongation along the main orientation is half value of its orthogonal. Formally, The mother Morlet wavelet at point  $p = (p_x, p_y)$  is defined as:

$$\psi(p) = e^{jk_0 p^T} e^{-\frac{1}{2}|Ap|^2} \quad (1)$$

where  $k_0 = (\frac{\pi}{2}, 0)$  and  $A = \text{diag}(1, \frac{1}{2})$ . The wavelet function of scale  $\sigma$  and orientation  $\tau_k$  is:

$$\psi_{k,\sigma}(p) = \frac{1}{\sigma} \psi\left(\frac{\Omega_k p}{\sigma}\right) \quad (2)$$

$$\Omega_k = \begin{pmatrix} \cos(\phi_k) & \sin(\phi_k) \\ -\sin(\phi_k) & \cos(\phi_k) \end{pmatrix} \quad (3)$$

where  $\phi_k$  is the angle associated with  $\tau_k$ .

If  $J_k(p)$  denote the modulus of wavelet coefficients at point  $p$ , local edge characteristics  $J_p$  and  $\tau_p$  are obtained by seeking the maximum wavelet response and orientation over all orientations. Formally:

$$J_k(p) = |I * \psi_{k,\sigma}|(p) \quad (4)$$

$$k_p = \arg \max_k J_k(p) \quad (5)$$

$$J_p = J_{k_p}(p) \quad (6)$$

$$\tau_p = \tau_{k_p} \quad (7)$$

where  $*$  denotes the 2D convolution operation.

Although local edge orientation is important to detect symmetry, the surrounding texture around an edge segment is also an important feature especially for natural images. Thus, we introduce a local orientation histogram  $h_p$  as the weighted directional contribution of all neighboring points  $r$  in the corresponding grid cell. As described in figures (1b-1d), the histogram count at a given orientation  $\tau_k$  is:

$$h_p(k) = \sum_{r \in N(p)} J_r \delta_{\phi_k - \phi_r} \quad (8)$$

where  $\phi_k$  and  $\phi_r$  are angles associated with  $\tau_k$  and  $\tau_r$ ,  $J_r$  and  $\tau_r$  are defined the same way as for  $p$  (equations (5)-(7)) and  $\delta_x$  is the Kronecker delta.  $h_p$  is subsequently  $\ell 1$  normalized and circular shifted so as the first bin corresponds to  $\tau_p$ .

### 3.2 Multi-scale Information

In most images, relevant information about the visual content may appear at different scales. To consider this information, we propose to detect the feature points at multiple scales. For that purpose, we consider a set of regular grids at different scales and a corresponding set of wavelet scales  $\{\sigma_l, l = 1..m\}$ . The whole set of feature points is augmented by extracting wavelet local maxima at each scale and by computing their corresponding local characteristics.

### 3.3 Pairwise Symmetry Coefficient and Mirror Symmetry Histogram

As shown in figure 1a, given two feature points  $p$  and  $q$  and their corresponding local edge orientations  $\tau_p$  and  $\tau_q$ ,  $T_{pq}$  is the line passing through  $p$  and  $q$ ,  $c = (c_x, c_y)$  denotes the center of  $pq$  segment and  $T_{pq}^\perp$  represents the candidate symmetry axis defined as the line perpendicular to  $pq$  passing through  $c$ .  $\vec{\sigma}$  is the reference unit vector of x-axis of the Cartesian plane.

The pairwise symmetry coefficient  $f(p, q)$  is defined as [5]:

$$f(p, q) = |\tau_q S(T_{pq}^\perp) \tau_p| \quad (9)$$

where  $S(T_{pq}^\perp)$  is the reflection matrix with respect to the line perpendicular to  $(pq)$ :

$$S(T_{pq}^\perp) = \begin{pmatrix} -\cos(2\theta_{pq}) & -\sin(2\theta_{pq}) \\ -\sin(2\theta_{pq}) & \cos(2\theta_{pq}) \end{pmatrix} \quad (10)$$

In addition to the geometric edge information, the symmetry degree of the two regions around  $p$  and  $q$  can be measured by comparing their corresponding local orientation histogram  $h_p$  and  $h_q$ . Due to symmetry considerations, one of the two histograms must be shifted and reversed before comparison. If  $d_I$  is the histogram intersection distance, and  $h_q^*$  the reversed histogram, the texture-based symmetry measure is given by:

$$d_I(h_p, h_q^*) = \sum_{k=1}^n \min(h_p(k), h_q^*(k)) \quad (11)$$

Given a pair of feature points  $(p, q)$ , the candidate axis  $T_{pq}^\perp$  perpendicular to  $(pq)$  is parametrized by the orientation of its normal  $\theta_{pq}$  and its distance to the origin  $\rho_{pq} = c_x \cos \theta_{pq} + c_y \sin \theta_{pq}$  where  $(c_x, c_y)$  are the coordinates of point  $c$  located at the middle of  $(p, q)$  segment (see figure 1a). A mirror symmetry histogram  $H_S(\rho, \theta)$  is defined as the sum of the contribution of all pairs of feature points such as:

$$H(c_x, c_y, \theta) = \sum_{\substack{p, q \\ p \neq q}} J_p J_q f(p, q) d_I(h_p, h_q) \delta_{(c_x, c_y) - \frac{p+q}{2}} \delta_{\theta - \theta_{pq}} \quad (12)$$

$$H_S(\rho, \theta) = \sum_{c_x, c_y} H(c_x, c_y, \theta) \delta_{\rho - \rho_{pq}} \quad (13)$$

The best bilateral symmetry axis inside the image is detected by searching for the global maximum of the mirror symmetry histogram  $H_S$ . For that purpose, we solve the orientation discontinuity problem by extending the voting space in a circular way. The spatial extension of this axis can be determined as the convex hull of the set of feature points contributing to this axis. Figures (2a-2d) show the symmetry detection for some example in terms of two voting histograms. The first  $H$  presents one major axis (A3 - red color) in the middle surrounded

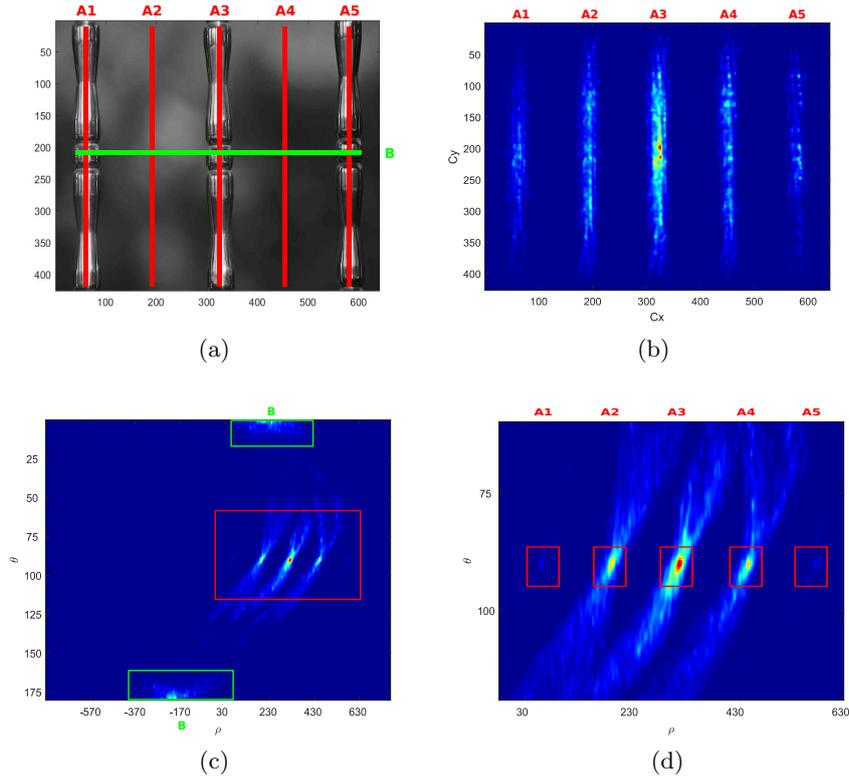


Fig. 2: Symmetry detection process: (a) Input image with local symmetry axis candidates indicated with red and green lines. (b) The output of the voting histogram  $H(c_x, c_y)$  marginalized over the orientation variable  $\theta$  and the corresponding symmetry axis IDs. (c) The output of the voting histogram  $H_S(\rho, \theta)$ . And (d) the zoomed-in region inside the red rectangle. Local maxima correspond to candidate symmetry axes (framed in red and green).

by four minor axes (A1, A2, A4, A5 - red color) vertically in figure 2b. In  $H_S$  histogram (figure 2c) and its zoom-in version (figure 2d), these axes are represented by 5 local maxima aligned along the line of orientation  $\theta = 90^\circ$ . Additionally, a horizontal axis (B - green color) appears near the boundaries of  $\theta$  axis, it describes another symmetry detection (the top half of steel and its mirror in the bottom side). The main axis is the vertical one (A3 - red color) located in the center.

## 4 Experiments and Discussions

The aim of this part is to evaluate the results of our method and to compare it to concurrent methods from the state-of-the-art. The evaluation of symmetry detection had received re-surfing interest in the main computer vision conferences. Indeed, symmetry detection challenges [17,8] were successively proposed with ECCV2010<sup>1</sup>, CVPR2011<sup>2</sup> and CVPR2013<sup>3</sup>. They all rely on a dataset proposed by Liu’s group publicly available with an associated groundtruth and evaluation method. We then use these three general purpose datasets for our evaluation and in addition, we provide an additional aesthetic-oriented dataset. To build this fourth dataset, we extracted some symmetrical images from the Aesthetic Visual Analysis (AVA) collection and generated axis groundtruth among those images.

### 4.1 Implementation and Evaluation Details

We compare our approach against three different methods (Loy2006 [11], Mo2011 [14], Cicconet2014 [5]). We found or requested their source code and ran them using default parameter values, assigned by the authors for stable performance. As described in [8], a global symmetry detection is considered to be correct if it satisfies the following conditions: (1) The angle between the detected symmetry axis and the groundtruth axis is less than 15 degrees. (2) The distance between centers of these two axes is less than 0.2 multiplied by minimum length of the axes. Accuracy represents the correct detection rate (the number of correct detection images divided by the total number of dataset images).

### 4.2 Comparison on PSU Datasets (2010, 2011 and 2013)

From Flickr photo sharing website, Liu’s vision group provides a single-symmetry groundtruth for synthetic/real images (size of 2010, 2011 and 2013 datasets are 134, 14 and 75 respectively). The quantitative comparisons are presented in Figure 3 with respect to each version of PSU dataset. Our method is ranked first for all PSU version with a difference of more than 10% among the nearest

<sup>1</sup> <http://vision.cse.psu.edu/research/symmetryCompetition/index.shtml>

<sup>2</sup> <http://vision.cse.psu.edu/research/symmComp/index.shtml>

<sup>3</sup> <http://vision.cse.psu.edu/research/symComp13/content.html>

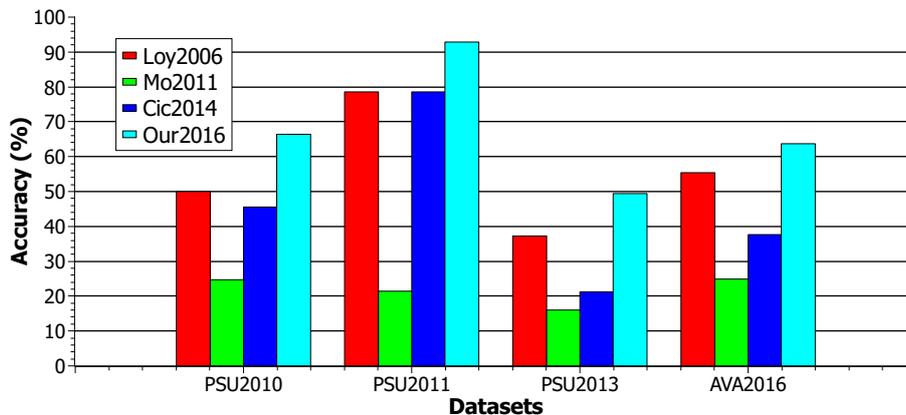


Fig. 3: Comparison of accuracy rates for our method “Our2016” against the baseline algorithm “Loy2006” [11], its improvement “Mo2011” [14] and one of the recent algorithms “Cic2014” [5]. The used datasets are: different versions of PSU (2010, 2011, 2013) with provided groundtruth [17,8] and AVA with our groundtruth.

competent Loy2006 [11], which achieves the best detection accuracy among 2011 and 2013 symmetry challenges [17,8]. Cic2014 [5] approach has ranked third due to the absence of local neighborhood information in the mirror symmetry histogram for candidate selection.

Figures (4a-4f) gives some detection examples<sup>4</sup> to present the superior performance of the proposed algorithm over others. In figures (4a,4d), our method correctly detects the symmetry axis of the bird with an extra line fitting on axis endpoints. However, [11] finds very few SIFT features generating a tiny non-meaningful detection, and [14,5] output misalignment axes representing different components within an image. Figure 4b shows a valid axis detection by our method, while [11,5] attempt to find an inaccurate global axis due to presence of the non-exact shapes inside the centered object. Inside figure 4c with different changes in environment illumination, our method perfectly aligned with the full symmetry axis and [11,14] select different secondary global axes due to surface’s mirror effect of the glass. A synthetic butterfly of figure 4e represents an easy detection of all methods except [14]. Our method has the most accurate axis output compared to the axis groundtruth. With a camera flash-on, facial figure 4f results a challenging axis detection, in which [11] shortly aligns over the groundtruth according to insufficient SIFT features. But our method gives the most appropriate axis output with a horizontal shift, such that it covers all human parts (including frontal face, neck, and hair). In all figures except 4e, [5]

<sup>4</sup> All PSU output with detected symmetry axes can be found in: <http://perso.univ-st-etienne.fr/em68594h/SupplementalFilesPSU.zip>

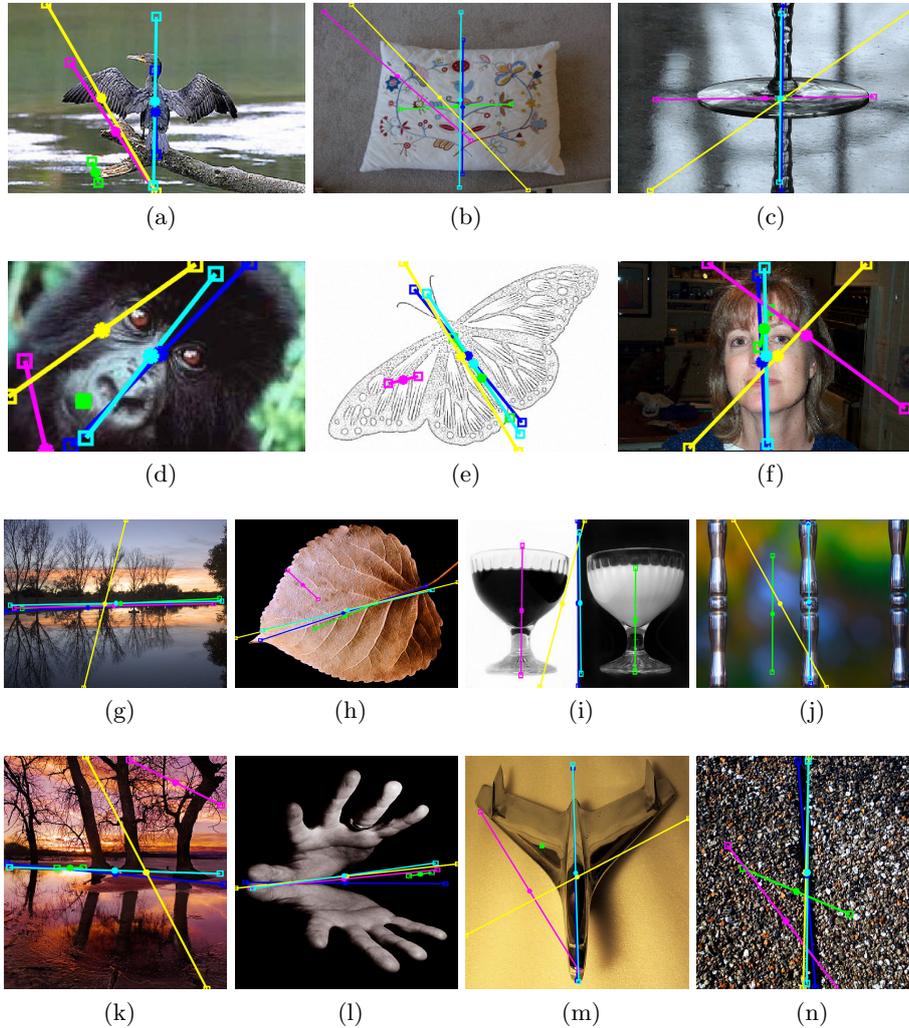


Fig. 4: Some images in (a-f) PSU datasets and (g-n) our labeled AVA dataset with groundtruth (blue), where our method (cyan) produces better or similar results over Loy2006 [11] (green), Mo2011 [14] (magenta) and Cic2014 [5] (yellow). Each axis is shown in a straight line with squared endpoints and a starred midpoint. Best seen on screen.

is very sensitive to noise in limited contrast environments and in high-texture foreground/background objects.

### 4.3 Comparison on AVA Dataset

From DPChallenge photo contest website, Murray et al. [15] introduces different annotations<sup>5</sup> (aesthetic, semantic and photographic style) for more than 250,000 images for Aesthetic Visual Analysis “AVA”. From the following photography challenges, we labeled global-axis symmetry groundtruth<sup>6</sup> for 253 out of 878 images: (1) five challenges of “Reflections Without Mirrors”: images containing bilateral representation without using mirror, (2) three challenges of “Symmetry”: photographs composing symmetrical balance. These images are selected to neglect unclear instances of ambiguity symmetry, and to represent many comparison cases (non-centering viewpoint, perspective view, blurring reflection, etc.) for detection algorithms. Figure 3 shows that our approach has better performance over other methods.

Figures (4g-4n) presents some experimental results<sup>7</sup> representing different types of artistic scene. Natural landscapes in figures (4g,4k) illustrate that [5] behaves abnormally with perspective view conditions, and [11,14] act incomplete with high shape occlusions. Abstract figures (4h,4l) demonstrate the superior performance of [5] compared to [11], in order to fit the global symmetry axis. [14] works better among less texture objects with clear contours. In figures (4i,4j), our method easily detects global symmetry axis, while [11,14] stuck with the local ones if exists. Our method only finds a clear symmetry detection in figure 4m with a high-reflected object with non-homogeneous background. Contour-based approaches, as ours and [5] ones, execute successfully within the non-exact symmetry figure 4n. In summary, our method and [5] ones align better in figures containing objects with clear edges, [11] locally detects shapes with smooth background, and [14] provides incorrect and unclear symmetry over the most visual cases.

## 5 Conclusion

This paper introduces multi-scale edge features combined with local orientation histograms, in order to develop a reliable global symmetry detection among variants of visual cues. Quantitative and qualitative comparisons show a substantial advantage for our proposed method on different types of public datasets. Our model can be improved to handle complex images with non-centered viewpoints and large degrees of perspective view. It can be extended to solve the over-fitting

<sup>5</sup> <http://www.lucamarchesotti.com/>

<sup>6</sup> Source code to generate AVA images and their symmetry labels: <http://perso.univ-st-etienne.fr/em68594h/SymAVA.zip>

<sup>7</sup> All AVA output with the detected symmetry axes can be found in: <http://perso.univ-st-etienne.fr/em68594h/SupplementalFilesAVA.zip>

of axis endpoints. The future work is to introduce a stable balance measure; describing the existence and degree of global axes in terms of both geometrical and symmetrical properties inside an image, and to integrate this measure within existing retrieval systems of visual arts.

## References

1. Atadjanov, I., Lee, S.: Bilateral symmetry detection based on scale invariant structure feature. In: Image Processing (ICIP), 2015 IEEE International Conference on. pp. 3447–3451. IEEE (2015)
2. Bairagi, V.: Symmetry-based biomedical image compression. Journal of digital imaging pp. 1–9 (2015)
3. Cai, D., Li, P., Su, F., Zhao, Z.: An adaptive symmetry detection algorithm based on local features. In: Visual Communications and Image Processing Conference, 2014 IEEE. pp. 478–481. IEEE (2014)
4. Cho, M., Lee, K.M.: Bilateral symmetry detection via symmetry-growing. In: BMVC. pp. 1–11. Citeseer (2009)
5. Cicconet, M., Geiger, D., Gunsalus, K.C., Werman, M.: Mirror symmetry histograms for capturing geometric properties in images. In: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on. pp. 2981–2986. IEEE (2014)
6. Duda, R.O., Hart, P.E.: Use of the hough transformation to detect lines and curves in pictures. Communications of the ACM 15(1), 11–15 (1972)
7. Kondra, S., Petrosino, A., Iodice, S.: Multi-scale kernel operators for reflection and rotation symmetry: further achievements. In: Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on. pp. 217–222. IEEE (2013)
8. Liu, J., Slota, G., Zheng, G., Wu, Z., Park, M., Lee, S., Rauschert, I., Liu, Y.: Symmetry detection from realworld images competition 2013: Summary and results. In: Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on. pp. 200–205. IEEE (2013)
9. Liu, Y., Hel-Or, H., Kaplan, C.S.: Computational symmetry in computer vision and computer graphics. Now publishers Inc (2010)
10. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International journal of computer vision 60(2), 91–110 (2004)
11. Loy, G., Eklundh, J.O.: Detecting symmetry and symmetric constellations of features. In: Computer Vision–ECCV 2006, pp. 508–521. Springer (2006)
12. Michaelsen, E., Muench, D., Arens, M.: Recognition of symmetry structure by use of gestalt algebra. In: Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on. pp. 206–210. IEEE (2013)
13. Ming, Y., Li, H., He, X.: Symmetry detection via contour grouping. In: Image Processing (ICIP), 2013 20th IEEE International Conference on. pp. 4259–4263. IEEE (2013)
14. Mo, Q., Draper, B.: Detecting bilateral symmetry with feature mirroring. In: CVPR 2011 Workshop on Symmetry Detection from Real World Images (2011)
15. Murray, N., Marchesotti, L., Perronnin, F.: Ava: A large-scale database for aesthetic visual analysis. In: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on. pp. 2408–2415. IEEE (2012)

16. Patraucean, V., von Gioi, R.G., Ovsjanikov, M.: Detection of mirror-symmetric image patches. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*. pp. 211–216. IEEE (2013)
17. Rauschert, I., Brocklehurst, K., Kashyap, S., Liu, J., Liu, Y.: First symmetry detection competition: Summary and results. Tech. rep., Technical Report CSE11-012, Department of Computer Science and Engineering, The Pennsylvania State University (2011)
18. Teo, C.L., Fermuller, C., Aloimonos, Y.: Detection and segmentation of 2d curved reflection symmetric structures. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1644–1652 (2015)
19. Wang, Z., Tang, Z., Zhang, X.: Reflection symmetry detection using locally affine invariant edge correspondence. *Image Processing, IEEE Transactions on* 24(4), 1297–1301 (2015)
20. Yang, L., Liu, J., Tang, X.: Depth from water reflection. *Image Processing, IEEE Transactions on* 24(4), 1235–1243 (2015)
21. Zhao, S., Gao, Y., Jiang, X., Yao, H., Chua, T.S., Sun, X.: Exploring principles-of-art features for image emotion recognition. In: *Proceedings of the ACM International Conference on Multimedia*. pp. 47–56. ACM (2014)