Bilateral Symmetry Detection via Symmetry-Growing

Minsu Cho
http://cv.snu.ac.kr/~minsucho
Kyoung Mu Lee
kyoungmu@snu.ac.kr

Computer Vision Lab.
Department of EECS, ASRI,
Seoul National University, Korea
http://cv.snu.ac.kr/

Abstract

We present a novel and robust method for localizing and segmenting bilaterally symmetric patterns from real-world images. On the basis of symmetrically matched pairs of local features, our method expands and merges confident local symmetric region matches by exploiting both photometric similarity and geometric consistency via our new symmetry-growing framework. It overcomes the limitations of the previous local-feature based approaches by efficiently exploring the image space to grow symmetry beyond the detected symmetric features. The experimental evaluation demonstrates that our method successfully detects and segments multiple symmetric patterns from real-world images, and clearly outperforms the state-of-the-art methods in accuracy and robustness.

1 Introduction

Our world is full of symmetries. A variety of symmetries occur in nature, living organisms, and manufactured artifacts, and provide humans with pre-attentive cues [2] that enhance object recognition. Human beings are very good at detecting symmetry, and understand the visual world based on the perception and recognition of repeated patterns that are generalized by the mathematical concept of symmetries [21]. Wagemans [20] views symmetry as one of the most important aspects of early visual analysis, and recent psychophysical results suggest that the detection of symmetries under perspective distortion is an integral part of 3D object perception. Symmetry detection on 2D or 3D images has been an active research area for over four decades [3, 8, 9, 12, 13, 18, 19, 22]. Despite the long history of the research, the recent performance evaluation [17] shows that we are still short of a robust and widely applicable symmetry detector.

In this paper, we propose a novel and robust method for detecting and segmenting bilateral or reflective symmetry which is the most familiar form of symmetry. It enables reliable detection and segmentation of multiple symmetries from real-world cluttered images. Our work is inspired by the recent match-growing approaches [1, 5, 7, 10] used in object-recognition and image registration. On the basis of symmetrically matched feature pairs [3, 12], our method expands and merges confident symmetric region matches by exploiting both photometric similarity and geometric constraint of bilateral symmetry in our novel symmetry-growing framework. Unlike the previous local-feature based methods based

BMVC 2009 doi:10.5244/C.23.4

on voting schemes [3, 9, 12, 19, 22], it can handle both significantly low inliers and deformation by efficiently exploring the image space beyond the detected features.

2 Previous Work

The problem of symmetry detection has been extensively studied in numerous fields including visual perception, computer vision, robotics, and computational geometry. The methods can be broadly classified into global and local feature-based methods. Global methods treat the entire image as a signal from which symmetric patterns are inferred. Examples include the work of Marola [13], Keller and Shkolnisky [8], Sun and Si [18]. However, these global approaches are limited to detecting a single incidence of symmetry, and also greatly influenced by background clutters. On the contrary, local feature-based methods use local features such as edges, contours, boundary points, and regions to detect symmetry by grouping symmetric sets of local features. The standard way to accumulate symmetry hypotheses from feature information is to use voting schemes such as the Hough transform. Their main advantage is to more efficiently detect local symmetries against background clutters in images that are not globally symmetric. Moreover, the recent development of local invariant features [11, 14, 15] has brought about significant progress in this approach. Tuytelaars *et al*. [19] presented a method for the detection of regular repetitions of planar patterns under perspective skew using a geometric framework and cascaded Hough transform. Lazebnik *et al*. [9] used affine invariant clusters of features to detect symmetries. Loy and Eklundh [12] proposed an efficient method to exploit the properties of local invariant features for grouping symmetric constellations of features and detecting symmetry. Cornelius *et al*. [3] extended this approach by constructing local affine frames. Although all these local-feature based methods show more robust performance over global methods, they are largely influenced by feature detection step, and cannot exploit further information beyond the detected features.

Recently, match-growing approaches [1, 5, 7, 10] are proposed to address the limitation of conventional local-feature based approaches in the area of object recognition and image registration. The match-growing approaches [5, 10] expand true matches and eliminate false matches based on the fact that the true matches grow better than false ones. Cho *et al*. [1] integrate match propagation of [5] in the data-driven Monte Carlo dynamics to recognize and segment common object pairs directly from image pairs. Kannala *et al*. [7] extend the quasi-dense matching method of [10] for object recognition and segmentation.

Our symmetry-growing method overcomes the limitations of the previous local-feature based approaches by efficiently exploring the image space to exploit further information beyond the detected symmetric features. Multiple clusters of consistent symmetric feature pairs are directly detected in our growing process without conventional voting procedure of the Hough transform or RANSAC. Unlike the previous match-growing methods [1, 5, 7, 10], we propose a novel and efficient multi-layer growing algorithm for avoiding the adverse effect of outliers in symmetry detection.

3 Overview of Our Approach

Given an image, our method aims to detect and segment all the bilaterally symmetric patterns and infer their quasi-dense correspondences within the symmetric patterns. Figure 1 illustrates a brief overview of our approach. First, we extract local invariant features from

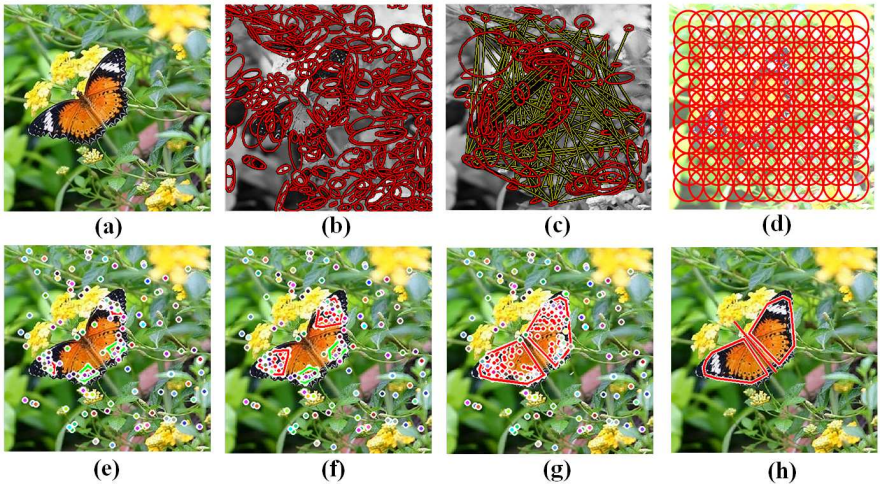

Figure 1: Overview of our symmetry-growing approach. See the tex for detailed description.

the given image (Fig. 1(a)) as shown in Fig. 1(b). Second, using the appearance around the detected features and its mirrored features, we obtain potential symmetric matching pairs of features as in Fig. 1(c). Third, starting from singleton symmetry clusters each containing a single symmetry match, we simultaneously expand and merge the symmetry clusters by exploring the image space in our symmetry-growing framework. To establish new symmetric region pairs in expansion, we use sets of element regions consisting of an overlapping circular grid of regular local regions as shown in Fig. 1(d). Using them, our algorithm gradually grows reliable symmetry clusters as shown in Fig. 1(e)-(f), where the dots with the same color represent the features in the same cluster. Finally, the reliable symmetric patterns grown well enough are chosen as in Fig. 1(h), where the detected symmetry is indicated by the convex hull of the features.

4 Bilateral Symmetry Seeds

Symmetrically matched pairs of local region features [12] are used as seeds for our symmetry-growing algorithm. Various modern local feature detectors [11, 14, 15] provide robust means for generating dense features and matching them between images. The distinctiveness of the matches obtained by local invariant features make these methods well suited for detecting pairs of symmetric features. Potential bilaterally symmetric matches can be obtained by constructing a set of mirrored feature descriptors and matching them against the original feature descriptors. The reliability of symmetry exhibited by each pair is evaluated by using the relative locations, orientations and scales of the features in the pair. The remainder of this section discusses the details of the procedure for detecting the seeds of bilateral symmetries.

4.1 Features and descriptors

A set of feature regions are determined using any affine-invariant or scale-invariant detectors such as [11, 14, 15] , that detect distinctive features with good repeatability. A detected

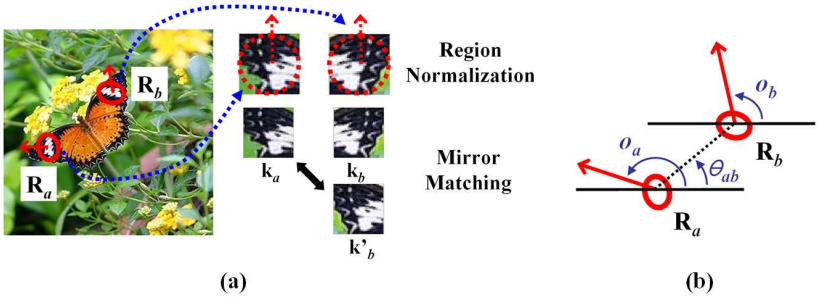

Figure 2: (a) Matching symmetric feature pairs. (b) Angle relation of a feature pair.

feature region can be represented by its location, orientation, and other region parameters. In this paper, we use the affine invariant region detector and its matrix parametrization [15]. A detected region R_a is denoted by $R_a = (\mathbf{x}_a, \Sigma_a, o_a)$ consisting of its center location, covariance matrix, and orientation, respectively. An interior pixel $\mathbf{x}$ of the region R_a satisfies $(\mathbf{x}-\mathbf{x}_a)^T \Sigma_a (\mathbf{x}-\mathbf{x}_a) \leq 1$. The orientation o_a of each feature is evaluated by its dominant gradient in the neighbor region as in [11]. Other types of features such as the scale-invariant detector [11] also can be adapted to this representation. After feature extraction, a feature descriptor $\mathbf{k}_a$ is generated for each feature region, encoding the local appearance of the feature after its normalization with respect to the orientation and affine distortion. Any feature descriptor suitable for matching can be used [16]. The experiments in this paper use MSER and Hessian affine detector [14, 15] for feature detection and the SIFT descriptor [11] for feature description.

4.2 Bilaterally Symmetric Feature Pairs

Figure 2(a) illustrates the process of matching symmetric feature pairs from an image. To describe the appearances of mirrored regions, a mirrored feature descriptor is generated for each feature. It can be produced by mirroring the original image region about the line along its orientation or by directly modifying this feature descriptor. As shown in Fig. 2(a), the mirrored feature descriptor $\mathbf{k}'_b$ describes a mirrored version of the feature region R_b associated with the feature descriptor $\mathbf{k}_b$. A set of potentially symmetric feature pairs are formed based on the similarities between the original descriptors and the mirrored descriptors. As shown in Fig. 2(a), a symmetric feature pair (R_a, R_b) is matched since the original descriptor $\mathbf{k}_a$ and the mirrored descriptor $\mathbf{k}'_b$ is similar enough. For matching, we calculate the similarities of descriptor pairs for all the possible symmetric feature pairs, and simply collect the best 300 matches allowing multiple correspondences for each feature. Usually, in this initial matching, the true match ratios become much lower than the one-to-one nearest neighbor (NN) methods, but it saves many true matches possibly eliminated by the NN matching methods in the presence of repeated patterns.

For a symmetric match $M_i = (R_a, R_b)$, its symmetry is evaluated by a function of the relative locations, orientations and scales of R_a and R_b. As in [12], based on the first component of Reisfeld's [4] phase weighting function, an angular symmetry weighting $\Phi_i \in [-1, 1]$ is computed as

$$\Phi_i = 1 - cos(o_a + o_b - 2\theta_{ab}), \tag{1}$$

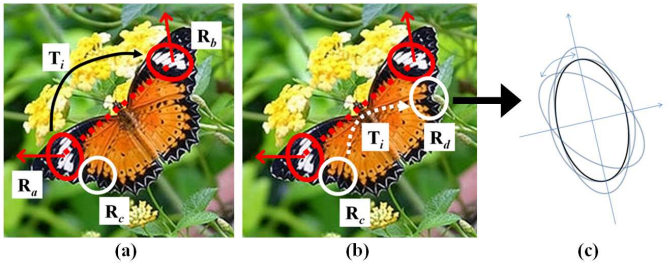

Figure 3: (a) Support match $M_i = (R_a, R_b)$ and a region R_c. (b) Propagation of the region R_c by M_i. (c) Region refinement.

where the angles are defined as illustrated in Fig. 2(b).

The centroid of the symmetry match $\mathbf{c}_i$ and the orientation of its local symmetry axis α_i are estimated as follows.

$$\mathbf{c}_i = (\mathbf{x}_a + \mathbf{x}_b)/2, \quad \alpha_i = (o_a + o_b)/2. \tag{2}$$

All the potential symmetry matches with positive phase weights are used for the symmetry seeds in the next step.

5 Symmetry-Growing Technique

In this section, we describe how our algorithm grows symmetry from the symmetry seeds obtained at the previous step.

5.1 Symmetry Propagation

The basic building block of our symmetry-growing algorithm is the *symmetry propagation*, which is proposed on the basis of the propagation attempt and refinement in [5].

As illustrated in Fig. 3, consider a local symmetry match $M_i = (R_a, R_b)$ in the case that a red elliptical region R_a is matched to R_b. For region R_a and R_b, normalizing transformation function H_a and H_b can be calculated, respectively, which transform the pixels in the regions onto the normalized unit circle with the same orientation [6]. Then, its reflective transformation T_i from R_a to R_b is defined as follows:

$$T_i = H_b^{-1} T_M H_a, \tag{3}$$

where T_M denotes a mirroring transformation about the normalized orientation.

Suppose that a region R_c is given that is close enough to R_a and lies on the same symmetric pattern. And, now we aim to generate another symmetry match $M_j = (R_c, R_d)$ using M_i and R_c. In that case, we state that the supporter match M_i attempts to propagate the region R_c. To achieve it, we approximate the symmetric region R_d (close to R_b) by $R_d = T_i R_c$ as shown in Fig. 3(b). Then, the propagated region R_d are refined locally to find a better match. The refiner adjusts R_d to produce the maximum ZNCC (Zero-mean normalized Cross correlation) similarity with R_c by locally searching the parameter space of its affine transformation as shown in Fig. 3(c). For the details of refinement, we refer to [5].

Algorithm 1 Multi-Layer Symmetry-Growing Algorithm

Obtain symmetry seed matches $\mathscr{M} = \{M_1, M_2, ..., M_n\}$
Construct singleton clusters $C_k = \{M_k\}(k = 1, 2, .., n)$
Initialize their expansion layers $\Gamma_k = \{R_1, R_2, ..., R_m\}(k = 1, 2, .., n)$
Put all the seeds into the supporter list $\mathscr{L} = \{M_1, M_2, ..., M_n\}$
repeat
  The supporter M_i with the highest similarity is removed from the supporter list $\mathscr{L}$
  Identify the cluster $C_p(\ni M_i)$ which contains the supporter match M_i
  Try to propagate its spatial neighborhood regions $\mathscr{N}(M_i, \Gamma_p)$ by the supporter M_i
  The reliable matches are accepted and stored in the cluster C_p, and also added to the supporter list $\mathscr{L}$
  The propagated regions are eliminated from the expansion layer Γ_p
  if equivalent matches are detected between different clusters **then**
    Merge the clusters including the equivalent matches
    Combine the expansion layers of the merged clusters by intersection of the layers.
  end if
until The supporter list $\mathscr{L}$ is empty
Eliminate unreliable symmetry clusters

5.2 Multi-Layer Symmetry-Growing Algorithm

Our symmetry-growing algorithm is summarized in Algorithm 1. We start from a set of symmetry seed matches obtained in the previous section. Using the seeds, singleton symmetry clusters are constructed by assigning each symmetry seed match to each symmetry cluster. Each cluster generates its own *expansion layer* which provides a space for expansion. It consists of a set of local regions in the overlapping circular grid which covers the image plane as shown in Fig. 1(d). Then, the list of supporter matches $\mathscr{L}$ is initialized as the set of all the obtained symmetry seeds. In each iteration, the supporter match M_i with the highest similarity is removed from the supporter list. The cluster containing M_i is denoted by $C_p(\ni M_i)$.

Expansion

For the supporter match M_i, the candidate neighborhood regions are selected in the current expansion layer of the cluster C_p for propagation. They consists of the unoccupied regions on the expansion layer, which are close to the larger region of the matching region pair in M_i. Suppose that $\mathbf{x}_i$ represents the center position of the larger region in the region pair of M_i, and $\mathbf{x}_c$ denotes the center position of a region R_c. The current expansion layer of the cluster which includes M_i is denoted by Γ_p. Then, the spatial neighborhood regions $\mathscr{N}(M_i, \Gamma_p)$ is defined as follows.

$$\mathscr{N}(M_i, \Gamma_p) = \{R_c \mid |\mathbf{x}_i - \mathbf{x}_c| \leq 2r_e, R_c \in \Gamma_p\}, \tag{4}$$

where r_e means the radius of overlapping circular regions in the expansion layer. It collects the neighborhood regions unoccupied by the cluster of the support match. The candidate regions are propagated by the supporter match M_i as in Fig. 3. Note that the direction of symmetry propagation in this manner is from the larger region of the support match to the smaller region of the support match. This propagation to the down-scale direction maintains better accuracy then the opposite direction.

The propagated symmetry matches are accepted only when both the photometric similarity and the property of local symmetry are reliable. We use the ZNCC for the photometric similarity, and the phase weight of Eq. 1 for the reliability of local symmetry. The matches exceeding both the similarity threshold δ_s and the phase weight threshold δ_Φ are accepted and included in the cluster of M_i. For further expansion, the accepted matches are also added to the supporter list $\mathscr{L}$. The regions propagated by the accepted matches are eliminated from the expansion layer Γ_p. Following this procedure, the supporter list, existing clusters, their expansion layers are all sequentially updated. The iteration continues until the supporter list becomes empty. Note that the expanded regions of different clusters can mutually overlap since each cluster has its own expansion layer. That is, the symmetry matches occupying the same region, called "region-sharing matches", can exist in our framework if any two of the matches do not belong to the same cluster. Unlike the previous single-layer match-growing methods [1, 5, 7, 10], our multi-layer approach enables to produce multiple symmetry clusters occupying the same regions and avoid the cases where expanded outlier clusters prevent inlier clusters to expand.

Merge

If a propagated match in expansion is significantly similar to existing one(s) in the same cluster of the supporter match, such an "equivalent match" is not allowed to expand for avoiding duplication. Otherwise, if the propagated match is equivalent to a match in a different cluster, the two relevant cluster are merged. In our method, two matches are determined as equivalent matches when they overlap over 50% of the area in both regions of the pair. When two clusters merge, their expansion layers are also combined into the intersection of the two layers. Thus, although our multi-layer approach multiplies the space to explore from a single-layer approach of the previous methods [1, 5, 7, 10], the merge process gradually reduces the number of the layers and guide the expansion process to concentrate on plausible expansion. Likewise, the expansion process also encourages the merge process to find compatible symmetry clusters by gradually growing them. Through these cooperative processes, our algorithm finds symmetric patterns efficiently in spite of significant outliers in initial seed matches.

Symmetry Cluster Verification

After the end of all iterations, we verify the reliability of final clusters. Typically, more reliable symmetry clusters are likely to grow larger. Thus, for a simple reliability measure of each symmetry cluster, we adopt the area of the convex hull of the local region center points in the cluster. We determine that a symmetry cluster is reliable if both of two reflective areas of the cluster are larger than $\delta_a|I|$ where $|I|$ denotes the area of the given image I.

6 Experiments

In all our experiments, the parameter values for the algorithm are fixed as follows: the ZNCC similarity threshold $\delta_s = 0.7$, the phase weight threshold $\delta_\Phi = 0.99$, and the reliable cluster threshold $\delta_a = 0.02$. The values of δ_s and δ_Φ is used for rejecting unreliable symmetry matches in expansion, thus we can control photometric variation and geometric deformation by varying them. In generating expansion layers, we set the radius of the overlapping circular regions r_e as $l/25$ where l denotes the shorter length of the axes of the given image.

We quantitatively evaluated our method on the test dataset of 91 images used in the recent performance evaluation of symmetry detection [17]. Examples of our results are

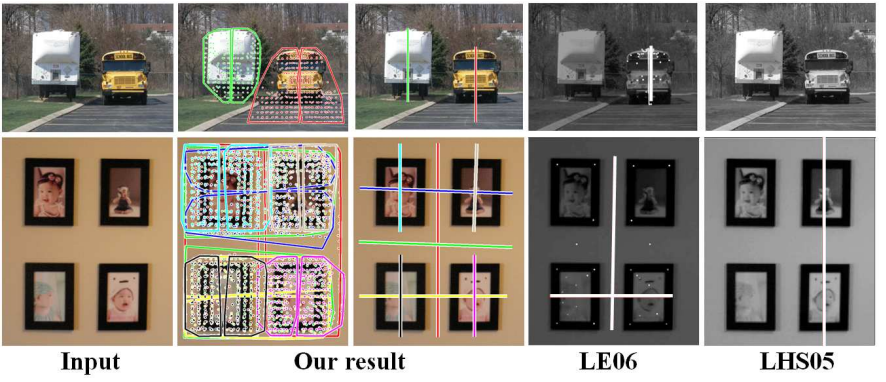

Figure 4: Symmetry detection comparison on the dataset of [17]. LE06 and LHS05 denote the methods of [12] and [22], respectively.

Image Type	Synthetic Single			Synthetic Multiple		
Algorithm	LE06	LHS05	**Ours**	LE06	LHS05	**Ours**
S_0	92%	62%	**100%**	35%	28%	**77%**
R_{FP}	15%	0%	**15%**	4%	8%	**33%**
Image Type	Real Single			Real Multiple		
Algorithm	LE06	LHS05	**Ours**	LE06	LHS05	**Ours**
S_0	84%	29%	**94%**	43%	18%	**68%**
R_{FP}	68%	3%	**69%**	44%	0%	**17%**

Table 1: Performance comparison on the dataset of 91 images from [17].

shown in Fig. 4 and 5 [1] where the convex-hull segmentation and the major axis of each symmetry cluster are visualized. Each color represents the identity of each symmetry cluster. The examples demonstrate that our symmetry-growing method detects the entire region of symmetric pattern with dense correspondences in each symmetry cluster, so that it provides more accurate and robust performance than the previous methods. Park *et al*. compared two state-of-the-art bilateral symmetry detection algorithms [12, 22] in their evaluation [17]. We compared our results with them. For comparison, sensitivity and false positive rate are measured. Suppose TP is the number of true positives: symmetries in the image that are identified correctly, FP is the number of false positives: non-symmetries detected by the algorithm as symmetries, and GT is the number of ground truth symmetries. Then, the sensitivity is defined as $S_0 = TP/GT$ and the false positive rate as $R_{FP} = FP/GT$. All the results are summarized in Table.1. Note that the results of [12] and [22] in Table.1 are obtained by testing the algorithms on four image scales (from 1 to 1/4 of the original size) and choosing the best result [17]. Whereas our results are obtained by testing only on the original size of the images. Nevertheless, for sensitivity S_0, our algorithm clearly outperforms the methods of [12, 22] in all image types.

In the sense of false positive rate R_{FP}, our method appears not clearly better than two

[1] The dataset, the ground truth, and the result images of [12] and [22] are borrowed from http://vision.cse.psu.edu/evaluation.html

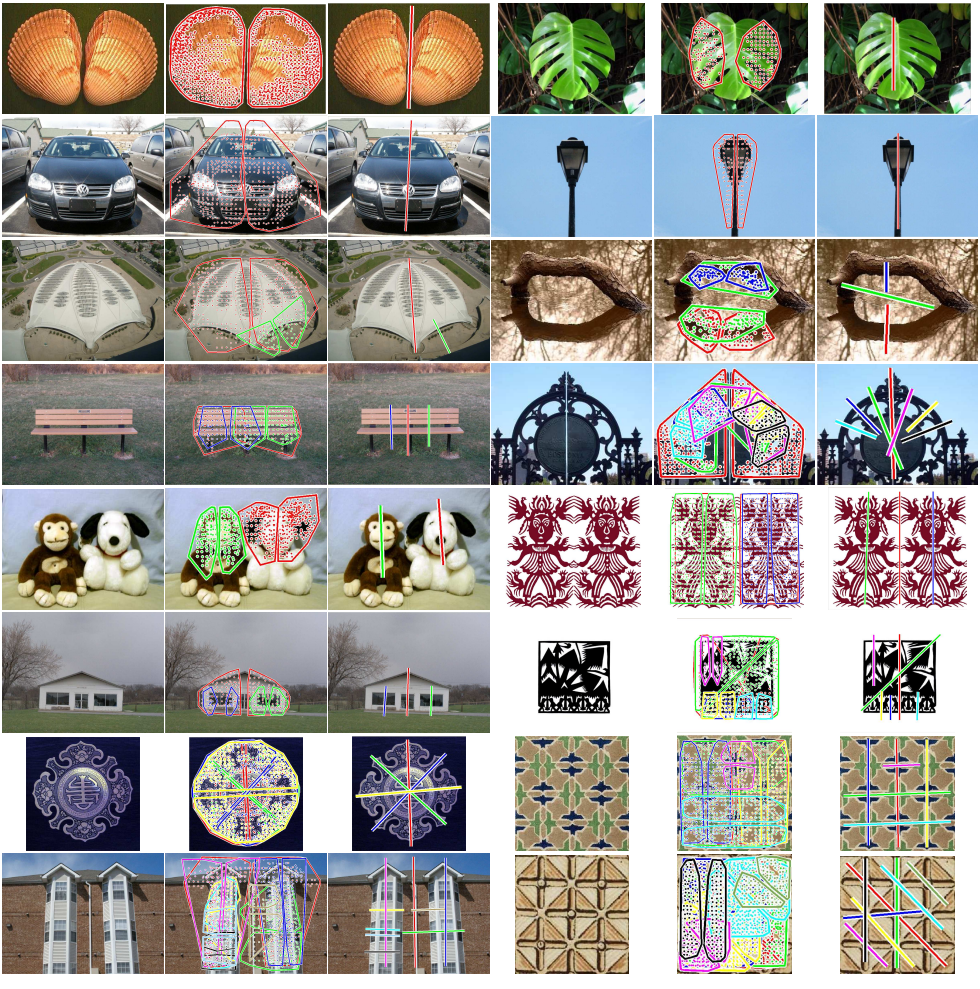

Figure 5: Examples of symmetry detection and segmentation on the dataset of [17].

other methods except for the case of real images with multiple symmetry. However, the reason is that our method detects more true symmetric patterns than the ground truth data of the dataset. That is, a great portion of our false positive detections are not false in fact. For examples, each image at the 3rd and 4th row in Fig. 5 has only one symmetric pattern according to the ground truth data, but our method detects even more true bilateral symmetric patterns. Of course, it also happens in multiple symmetry dataset as images in 8th row. These cases were prevalent in our results. Therefore, the experiments demonstrate that our method is highly robust and accurate and is widely applicable to the real-world complex images.

7 Conclusion

In this paper we presented a novel approach to bilaterally symmetric patterns detection and segmentation, which deals with multiple symmetry in real-world images. On the basis of

a local-feature based approach, it efficiently explores the image plane to exploit further information beyond detected symmetric features. In this strategy, our method overcomes the limitations of the previous symmetry detection methods with a novel multi-layer symmetry-growing framework which detects overlapping symmetries and avoiding interference of outliers. As demonstrated in the experiments, our method clearly outperforms state-of-the-art methods and provides robust detection on real-world complex images. Our approach can be extended for a variety of pattern analysis tasks.

Acknowledgements

This work was supported in part by the Korea Research Foundation Grant funded by the Korean Government (MOEHRD) (KRF-2008-314-D00377), and in part by the IT R&D program of MKE/IITA (2008-F-030-01, Development of Full 3D Reconstruction Technology for Broadcasting Communication Fusion).

References

[1] Minsu Cho, Young Min Shin, and Kyoung Mu Lee. Co-recognition of image pairs by data-driven monte carlo image exploration. *ECCV*, pages IV: 144–157, 2008.

[2] R. W. Conners and C. T. Ng. Developing a quantitative model of human preattentive vision. *IEEE Trans. Systems, Man and Cybernetics*, 19(6):1384–1407, November 1989.

[3] Hugo Cornelius, Michal Perd'och, Jiří Matas, and Gareth Loy. Efficient symmetry detection using local affine frames. *SCIA*, pages 152–161, 2007.

[4] Reisfeld D., Wolfson H., and Yeshurun Y. Context free attentional operators: the generalized symmetry transform. *Int. J. of Computer Vision, Special Issue on Qualitative Vision*, 1994.

[5] Vittorio Ferrari, Tinne Tuytelaars, and Luc Gool. Simultaneous object recognition and segmentation from single or multiple model views. *IJCV*, 67(2):159–188, 2006.

[6] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[7] J. Kannala, E. Rahtu, S.S. Brandt, and J. Heikkila. Object recognition and segmentation by non-rigid quasi-dense matching. *CVPR*, pages 1–8, 2008.

[8] Y. Keller and Y. Shkolnisky. An algebraic approach to symmetry detection. *ICPR*, 2004.

[9] S. Lazebnik, C. Schmid, and J. Ponce. Semi-local affine parts for object recognition. *BMVC*, 2004.

[10] M. Lhuillier and L. Quan. Match propagation for image-based modeling and rendering. *PAMI*, pages 1140–1146, 2002.

[11] David G. Lowe. Object recognition from local scale-invariant features. *ICCV*, pages 1150–1157, 1999.

[12] G. Loy and J. O. Eklundh. Detecting symmetry and symmetric constellations of features. *ECCV*, pages II: 508–521, 2006.

[13] G. Marola. On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *PAMI*, 11:104–108, 1989.

[14] Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pajdla. Robust wide baseline stereo from maximally stable extremal regions. *BMVC*, 2002.

[15] K. Mikolajczyk and C. Schmid. Scale and affine invariant interest point detectors. *IJCV*, 60(1):63–86, October 2004.

[16] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *PAMI*, 27(10):1615–1630, October 2005.

[17] Minwoo Park, Seungkyu Lee, Po-Chun Chen, Somesh Kashyap, Asad A. Butt, and Yanxi Liu. Performance evaluation of state-of-the-art discrete symmetry detection algorithms. *CVPR*, 2008.

[18] C. M. Sun and D. Si. Fast reflectional symmetry detection using orientation histograms. *Real-Time Imaging*, 5(1):63–74, February 1999.

[19] Tinne Tuytelaars, Andreas Turina, and Luc J. Van Gool. Noncombinatorial detection of regular repetitions under perspective skew. *PAMI*, 25(4):418–432, 2003.

[20] J. Wagemans. Detection of visual symmetries. *Spatial Vision*, 9(9), 1995.

[21] H. Weyl. *Symmetry*. Princeton University Press, Princeton, 1952.

[22] Ying-Qing Xu Yanxi Liu, James H. Hays and Heung-Yeung Shum. Digital papercutting. *Technical Sketch, SIGGRAPH*, 2005.