

# Reflective Symmetry Detection by Rectifying Randomized Correspondences

Zhongwei Tang  
tang@cmla.ens-cachan.fr

Mariano Tepper  
mariano.tepper@duke.edu

Guillermo Sapiro  
guillermo.sapiro@duke.edu

Department of Electrical and Computer  
Engineering  
Duke University  
130 Hudson Hall  
Box 90291  
Durham, NC 27708, USA

---

## Abstract

We present a method for detecting bilateral or reflective symmetries in images. We pose the problem as an instance of a multiple model estimation problem. We build candidate symmetry models by randomly sampling minimal sets of SIFT matches. Since these symmetry models can be in non-frontal viewpoints, we rectify them, undoing the perspective effect. From the models with valid symmetric properties, we compute consensus sets by determining which SIFT matches are compatible with each symmetry model. We finally recombine these consensus sets, using a clustering algorithm. The method is able to detect single and multiple symmetries both in frontal and non fronto-parallel viewpoints, achieving state-of-the-art results.

## 1 Introduction

Symmetry is omnipresent in natural and synthetic images. Human visual perception of the world largely depends on cues provided by symmetry and repetitive patterns [3]. Automatic symmetry detection has long been an active research topic in computer vision because it also helps to enhance the performance of many vision problems, e.g., image segmentation [18], object recognition [12], person identification [6], vehicle tracking [9], and low-rank texture detection [23]. There are four types of symmetries in the 2D Euclidean space: translation, rotation, reflection and glide-reflection [24]. The human visual system is good at detecting all these symmetries, even under severely slanted viewpoints, while it remains a tough problem for computer vision. Among these, the most common form is the reflective (also called bilateral) symmetry, characterized by a line (the symmetry axis) which defines two reciprocally mirrored semi-planes. In this paper, we concentrate on detecting this type of symmetry. The concepts we propose can be nonetheless extended to the other types of symmetry.

A lot of research efforts have been dedicated to automatically detect a single, global, and dominant bilateral symmetry in images [11, 8, 13, 17]. These global methods consider the image as a whole and try to infer the global symmetry that produces the best overall fit. This poses problems when the background of the image is cluttered or the image contains multiple symmetries at different locations and scales. Although global symmetry is an important image property, encountering multiple local symmetries is far more common in practice.

Current research is thus oriented to methods based on local features [10, 20]. Finding multiple symmetries can be posed as a multiple model selection problem. Multi-RANSAC [24] and randomized Hough transform [22] are two basic model estimation tools that can be used to detect multiple symmetries. For example, a simple voting strategy was employed for bilateral symmetry detection through a fold-then-cut plane generation and synthesis [10]. Loy and Eklundh presented a more efficient method based on the Hough transform [10], where each SIFT match votes either for a bilateral symmetry axis or a center of rotational symmetry in the parameter space. Later, the work was extended to make use of the recent advance in invariant features [24] to deal with bilateral symmetry under perspective [4, 5]. Even though recent work based on a region growing scheme seems to work better [2], Loy and Eklundh’s work [10] is still considered as the baseline state-of-the-art algorithm.

The Hough transform can naturally cope with multiple symmetries. However, it is not as robust against outliers as RANSAC [7] and is sensitive to the selection and discretization of the parameter space. The J-linkage algorithm was recently proposed to detect multiple primitive geometric structures in noisy and outliers-corrupted data [19]. It combines ideas from RANSAC (robustness against outliers) and the Hough transform (multiple structures detection through voting schemes). The input of the algorithm is a set  $X$  of geometric objects (e.g., 2D points). J-linkage randomly samples minimal sets of objects from  $X$  and creates candidate models (e.g., if the model is a line, two points are sufficient to define it). It then records, for each candidate model, its consensus set, i.e., the subset of  $X$  compatible with that model. This is encoded as a binary preference matrix, whose rows indicate the sampled models an object belongs to, and whose columns indicate the consensus set for each sampled model. The objects are finally clustered using the rows as vector features via agglomerative hierarchical clustering [19]. The clusters that are populated enough, can then be used to robustly estimate the final models.

We propose a method to detect multiple bilateral symmetries at different scales and view-points. It is based on the J-linkage framework, presenting specific and novel tools in each step of the algorithm to adapt it to our problem. These tools, described in Section 2, include: (1) a rectification step that allows us to transform each symmetry to the 2D Euclidean space, in which validation can be easily performed, while coping with skewed symmetries; (2) a simple validation criterion to determine valid rectified symmetries and their consensus set; (3) a non-uniform adaptive sampling strategy, specifically designed to deal with a large set of matches corrupted by a high percentage of outliers. In Section 3 we present experimental results showing state-of-the-art results. Finally, we provide some final remarks in Section 4.

## 2 Multiple symmetry detection with J-linkage

We start by detecting affine-invariant keypoints and obtaining affine-invariant SIFT descriptors with the ASIFT algorithm [15]. We denote by  $\text{SIFT}(\mathbf{p}_j)$  the descriptor at keypoint  $\mathbf{p}_j$ . SIFT also provides the orientation  $\theta_j$  and scale  $s(\mathbf{p}_j)$  at  $\mathbf{p}_j$ . The elements in each SIFT descriptor vector  $\text{SIFT}(\mathbf{p}_j)$  are rearranged to obtain their mirrored version, where the line with orientation  $\theta_j$  is used as the mirroring axis [10]. We then match the SIFT descriptors with their mirrored versions using a simple nearest-neighborhood search. Let  $\mathcal{M} = \{m_i\}_{i=1, \dots, N}$  be this set of  $N$  matches. Each match  $m_i$  consists of two keypoints, expressed in homogeneous coordinates,  $m_i = \{\mathbf{p}_i, \mathbf{p}'_i\}$ , with  $\mathbf{p}_i = (x_i, y_i, 1)^T$  and  $\mathbf{p}'_i = (x'_i, y'_i, 1)^T$ . We denote by  $\bar{m}_i$  the line segment connecting  $\mathbf{p}_i$  and  $\mathbf{p}'_i$ . We also denote by  $\mathbf{c}_i$  the midpoint of  $\bar{m}_i$ .

The  $\mathbb{R}^{3 \times 3}$  matrix  $\mathbf{H}$  such that  $h_{33} = 1$  (where  $h_{ij}$  is the value at the  $i$ -th row and  $j$ -th

column) is a 2D homography. This type of transforms models perspective transformations. Notice that a transformed point  $\mathbf{q} = (x_q, y_q, z_q)^T = \mathbf{H}\mathbf{p}$  is normalized such that  $z_q = 1$ , although the normalization to achieve this is implicit in this paper.

**J-linkage-based symmetry detection.** Given  $N$  matches, we estimate  $K$  symmetries by randomly sampling  $K$  valid minimal seed sets from  $\mathcal{M}$ . We will see that two matches are sufficient to define a minimal seed set. The exact nature of this process will be covered in the following subsections. We thus obtain  $K$  symmetries, each with an associated consensus set (the subset of  $\mathcal{M}$  compatible with each symmetry). A binary  $N \times K$  matrix is thus built, where the entry  $(i, j)$  is 1 if the  $i$ -th match is in the consensus set of the  $j$ -th symmetry, and 0 otherwise. Each row of this matrix indicates which symmetries are preferred by each match and is considered as a binary feature vector for that match. Using these features, agglomerative hierarchical clustering based on the Jaccard distance [19] is used to cluster the matches. Finally, each, large enough, cluster corresponds to a local symmetry.

## 2.1 Estimating a fronto-parallel symmetry by rectifying match pairs

If the symmetric region only undergoes a similarity transform (rotation, translation and zoom), one match  $m_i$  is enough to determine the symmetry axis [14], which is orthogonal to the segment  $\mathbf{p}_i\overline{\mathbf{p}_i}$  and passing through  $\mathbf{c}_i$ . However, in practice, the symmetry is not necessarily observed in a frontal view and can thus undergo some perspective distortion. In this more general case, two matches  $m_1 = \{\mathbf{p}_1, \mathbf{p}'_1\}$  and  $m_2 = \{\mathbf{p}_2, \mathbf{p}'_2\}$  intersecting at the vanishing point  $\mathbf{v} = (v_x, v_y, 1)^T$  are necessary to determine the symmetric axis (see Fig. 1(a)). We will undo this perspective distortion, that is, rectify the symmetry, thus being able to determine the compatibility of the remaining matches in the simpler 2D Euclidean space.

A homography  $\mathbf{H}$  can be decomposed into

$$\mathbf{H} = \mathbf{A}\mathbf{R}\mathbf{H}_0 = \mathbf{A}\mathbf{R} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ h_{31} & h_{32} & 1 \end{bmatrix}, \quad (1)$$

where  $\mathbf{H}_0$  and  $\mathbf{A}$  are a projective transform and a shear, respectively, and  $\mathbf{R}$  is a rotation. We will compute  $\mathbf{H}$  and use it to bring the symmetry into a fronto-parallel setting. This process is depicted in Fig. 1(a).

We begin by computing  $\mathbf{H}_0$ . This transform needs to send the vanishing point  $\mathbf{v}$  to infinity, that is,

$$h_{31}v_x + h_{32}v_y + 1 = 0. \quad (2)$$

For choosing the remaining degree of freedom we add the constraint that the required scale change factor at the four keypoints of  $m_1$  and  $m_2$  is as close as to 1 as possible. Intuitively, among all possible  $\mathbf{H}_0$ , we select the one that is closer to the identity matrix, i.e., the one introducing the least perspective effect. This amounts to computing

$$\min_{h_{31}, h_{32}} \sum_{i=1,2} (h_{31}x_i + h_{32}y_i)^2 + (h_{31}x'_i + h_{32}y'_i)^2. \quad (3)$$

Geometrically, this choice can be interpreted by rotating the camera as little as possible to make the vanishing point at the infinity. By plugging Eq. (2) in this minimization, we obtain a one-variable least-square problem which can be easily solved.

For computing  $\mathbf{R}$ , we simply constrain the vanishing point (now at infinity) to lie on the  $x$  axis. Finally,  $\mathbf{A}$  is obtained by estimating the shear that aligns the two corresponding midpoints  $\mathbf{c}_1, \mathbf{c}_2$  with the  $y$  axis.




Figure 1: (a) Finding the rectification homography ( $\mathbf{H} = \mathbf{A}\mathbf{R}\mathbf{H}_0$ ) from a pair of matches. (b) Compatible matches test given a rectified symmetry.

The homography  $\mathbf{H}$  is then applied to the keypoints of all the matches  $m_i \in \mathcal{M}$ . In this rectified plane, all the matches compatible with the reflective symmetry defined by  $m_1, m_2$  should be parallel with the seed line segments  $\bar{m}_1, \bar{m}_2$  and have their midpoint on the symmetric axis. Obviously, this is an ideal scenario and in practice we need to relax this criterion. Let  $\widehat{H}_{p_i}$  (resp.  $\widehat{H}_{p'_i}$ ) be the point that is actually symmetric to  $H_{p_i}$  (resp.  $H_{p'_i}$ ). We then measure the ratio between the segment with endpoints  $H_{p'_i}, \widehat{H}_{p_i}$  (resp.  $H_{p_i}, \widehat{H}_{p'_i}$ ) and the segment with endpoints  $H_{p_i}, \widehat{H}_{p_i}$  (resp.  $H_{p'_i}, \widehat{H}_{p'_i}$ ). See Fig. 1(b) for a graphical representation. If both ratios are smaller than a precision parameter  $\eta$ , the match  $m_i$  is added to the consensus set.

## 2.2 Validating the seed matches

We need to ensure that the two randomly sampled seed matches  $m_1$  and  $m_2$  lead to a valid symmetry model. We perform the following sanity checks in order to reject invalid models:

- Since we are dealing with reflective symmetries, the segments  $\bar{m}_1$  and  $\bar{m}_2$  must not intersect before and after rectification;
- The keypoint scales must become approximately similar after rectification,

$$|\hat{s}(\mathbf{p}_i) - \hat{s}(\mathbf{p}'_i)| / \max(\hat{s}(\mathbf{p}_i), \hat{s}(\mathbf{p}'_i)) < \delta, \quad i = 1, 2, \quad (4)$$

where  $\hat{s}(\mathbf{p}_i)$  denotes the scale at  $\mathbf{p}_i$  after rectification and  $\delta$  is a precision parameter;

- The sum of the orientations  $\hat{\theta}_i, \hat{\theta}'_i$  at  $\mathbf{p}_i, \mathbf{p}'_i$  after rectification, must be approximately equal to  $\pi$  (these orientations are not very robust). We check the condition [16] (see Fig. 2)

$$1 + \cos(\hat{\theta}_i + \hat{\theta}'_i) < \varepsilon, \quad i = 1, 2, \quad (5)$$

where  $\varepsilon$  is a precision parameter.

We could have used these sanity checks to constrain more the homography  $\mathbf{H}$ . But since the scale and orientation of the keypoints is not very precise, the estimation of the homography can be instable.






Figure 2: A schematic representation of two matching keypoints. On the right, the match  $m_i = \{\mathbf{p}_i, \mathbf{p}'_i\}$  with corresponding angles  $\hat{\theta}_i, \hat{\theta}'_i$  (i.e.,  $\theta_i, \theta'_i$  transformed by the rectification).

### 2.3 Adaptively sampling seed matches

One of the keys for the success of the J-linkage algorithm is choosing a proper non-uniform sampling. The rationale behind this is the intent to oversample the *true* symmetries in the image, thus obtaining stable row features that robustify the clustering process. Among all of the sampled model seeds, it is desirable that there is at least a certain number of outlier-free model seeds sampled around each underlying true model. In [19, 24], non-uniform sampling is used to detect multiple simple geometric models in point clouds, like lines, circles, or planes, by assigning a higher probability to neighboring points.

We use the following adaptive non-uniform sampling strategy. The first match is sampled according to the following mixed probability,  $\forall m_i = \{\mathbf{p}_i, \mathbf{p}'_i\} \in \mathcal{M}$ ,

$$\Pr(m_i) = \frac{1}{Z_1} \exp \left( -\frac{1}{\sigma_d^2} (\|\text{SIFT}(\mathbf{p}_i) - \text{SIFT}(\mathbf{p}'_i)\| - d_0)^2 - \frac{1}{\sigma_l^2} (\|\mathbf{p}_i - \mathbf{p}'_i\| - l_0)^2 \right), \quad (6)$$

where  $d_0$  and  $l_0$  indicate the scale in the descriptor domain and image domain at which we prefer to detect the symmetry,  $Z_1$  is a normalization factor such that  $\sum_i \Pr(m_i) = 1$ , and, finally,  $\sigma_d$  and  $\sigma_l$  decide how strict the preferences are. Ideally, we should have  $d_0 = 0$ , but since some parts of the image match closely than others, we might end up missing some symmetries. We thus relax this constraint.

In practice, an image can contain multiple symmetries at different scales. Automatically updating the parameters of Eq. (6) to find all possible symmetries is not an easy task. This motivates us to update the sampling probability of the first match along the sampling process. Each time two seed matches are sampled and its consensus set  $\mathcal{M}_0$  is computed, we decrease the probability of all these matches by a factor  $\kappa$  close to 1, followed by the renormalization of the probability:

$$\text{probability update (followed by renormalization)} : \quad \Pr(m_i) \leftarrow \kappa \Pr(m_i), m_i \in \mathcal{M}_0. \quad (7)$$

This guarantees that the sampling will not solely focus on symmetries with “high quality” matches and “lower quality” matches will also be visited.

Once the first match is sampled following the above adaptive non-uniform sampling, we sample the second match according to the conditional probability:

$$\Pr(m_j | m_i) = \frac{1}{Z_2} \exp \left( -\frac{1}{\sigma_c^2} (\|\mathbf{c}_i - \mathbf{c}_j\| - c_0)^2 \right), m_i \in \mathcal{M}, m_j \in \mathcal{M} \setminus \{m_i\}, \quad (8)$$

where  $Z_2$  is a normalization factor such that  $\sum_j \Pr(m_j | m_i) = 1$ ,  $\mathbf{c}_i$  is the midpoint of line segment  $\bar{m}_i$ , and  $\sigma_c, c_0$  control the shape of the probability function.

**Algorithm 1:** Non-uniform adaptive sampling.

---

**Data:** matches  $\mathcal{M} = \{m_i\}, i = 1, \dots, N$   
**Result:** preference matrix  $P$   
Initialize the sampling probability  $\Pr(m_i)$  and  $\Pr(m_j|m_i)$  according to Eq. (6) and (8) ;  
Sampling counter  $K_{counter} \leftarrow 0$  ;  
 $flag\_uniform\_sampling = 0$  ;  
Consensus set to be uniformly sampled  $\mathcal{C}_0 \leftarrow \emptyset$  ;  
**while**  $K_{counter} < K$  **do**  
     $K_{trial} \leftarrow 0$  ;  
    **while**  $K_{trial} < K_{max\_trial}$  **do**  
         $K_{trial} \leftarrow K_{trial} + 1$  ;  
        **if**  $flag\_uniform\_sampling = 0$  **then**  
            Sample two seed matches according to probability  $\Pr(m_i)$  and  $\Pr(m_j|m_i)$  ;  
        **else**  
            Sample two seed matches uniformly in the consensus set  $\mathcal{C}_0$  ;  
        Rectify  $m_1, m_2$  (see Section 2.1) ;  
        If the seeds do not form a valid model (see Section 2.2) continue sampling ;  
    If the seeds do not form a valid model (see Section 2.2) terminate the algorithm ;  
    Compute the consensus set  $\mathcal{C}$  for the rectified  $m_1, m_2$  (see Fig. 1(b)) ;  
    Add the consensus set  $\mathcal{C}$  to preference matrix  $P$  ;  
    Update the probability  $\Pr(m_i)$  according to Eq. (7) ;  
     $K_{counter} \leftarrow K_{counter} + 1$  ;  
    **if**  $flag\_uniform\_sampling = 1$  **then**  
         $K_{uniform\_sampling} \leftarrow K_{uniform\_sampling} + 1$  ;  
    **if**  $flag\_uniform\_sampling = 0 \wedge |\mathcal{C}| > T$  **then**  
         $flag\_uniform\_sampling \leftarrow 1$  ;  
         $K_{uniform\_sampling} \leftarrow 0$  ;  
        Initialize the set to be uniformly sampled:  $\mathcal{C}_0 \leftarrow \mathcal{C}$  ;  
    **if**  $flag\_uniform\_sampling = 1 \wedge |\mathcal{C}| > |\mathcal{C}_0|$  **then**  
        Update the set to be uniformly sampled:  $\mathcal{C}_0 \leftarrow \mathcal{C}$  ;  
    **if**  $flag\_uniform\_sampling = 1 \wedge K_{uniform\_sampling} = K_u$  **then**  
         $flag\_uniform\_sampling \leftarrow 0$  ;

---

If the current consensus set  $\mathcal{C}$  is big enough, i.e.,  $|\mathcal{C}| > T$  for some  $T$ , we switch to uniformly sample  $K_u$  seeds *inside*  $\mathcal{C}$ . This ensures that good models are oversampled.

The overall adaptive non-uniform sampling procedure is summarized in Alg. 1. Given the preference matrix, we run the J-linkage algorithm to cluster the matches.

### 3 Experimental results

For all experiments, we randomly select  $N = 3000$  keypoints/features from the ones provided by ASIFT [15]. These features are matched to their mirrored version by nearest neighbors matching (one neighbor per feature for single symmetry detection and four neighbors per feature for multiple symmetry detection). We initialize the probabilities in eqs. (6) and (8)

	Synthetic Single				Synthetic Multiple			
	LE [10]	LHXS [10]	CL [9]	Proposed	LE [10]	LHXS [10]	CL [9]	Proposed
TP/GT	92%	62%	100%	<b>100%</b>	35%	28%	77%	<b>67%</b>
FP/GT	15%	0%	15%	<b>0%</b>	4%	8%	33%	<b>10%</b>

	Real Single				Real Multiple			
	LE [10]	LHXS [10]	CL [9]	Proposed	LE [10]	LHXS [10]	CL [9]	Proposed
TP/GT	84%	29%	94%	<b>97%</b>	43%	18%	68%	<b>65%</b>
FP/GT	68%	3%	69%	<b>39%</b>	44%	0%	17%	<b>16%</b>

Table 1: Performance comparison of several methods on the PSU dataset [10]. TP, FP, and GT respectively denote the number of true positives, false positives, and ground truth symmetries. The percentage of the methods are taken from [10] and [9].

with the values

$$d_0 = \frac{1}{2N} \sum_{m_i \in \mathcal{M}} \|\text{SIFT}(\mathbf{p}_i) - \text{SIFT}(\mathbf{p}'_i)\|, \quad \sigma_d^2 = \frac{1}{10} \max_{m_i \in \mathcal{M}} \|\text{SIFT}(\mathbf{p}_i) - \text{SIFT}(\mathbf{p}'_i)\|^2, \quad (9)$$

$$l_0 = \frac{1}{10} (w^2 + h^2)^{1/2}, \quad \sigma_l^2 = \frac{1}{10} \max_{m_i \in \mathcal{M}} \|\mathbf{p}_i - \mathbf{p}'_i\|^2, \quad (10)$$

$$c_0 = \frac{1}{20} (w^2 + h^2)^{1/2}, \quad \sigma_c^2 = \frac{1}{10} \max_{m_i, m_j \in \mathcal{M}} \|\mathbf{c}_i - \mathbf{c}_j\|^2, \quad (11)$$

where  $w$  and  $h$  are the width and height of the image, respectively. The probability  $\text{Pr}(m_i)$  is updated using  $\kappa = 0.98$  in Eq. (7). We check the validity of every non-uniformly sampled match pair by setting  $\delta = 0.2$  and  $\varepsilon = 0.25$  in eqs. (4) and (5). The matches compatible with the sampled symmetry models with parameter  $\eta = 0.04$  (Fig. 2b) are considered as inliers and added to the consensus sets. We also set  $K = 4000$ ,  $T = 10$ , and  $K_u = 30$ . After the J-linkage clustering, only symmetries containing at least 10 matches are kept. We finally apply a non-maximum suppression on the clusters whose symmetry axes are close.

We first compare the proposed method with three recent ones [9, 10, 10] on the PSU dataset [10], which is composed of 88 images.<sup>1</sup> The symmetries are either frontal or slightly skewed. This dataset covers four types of images: synthetic single reflection, synthetic multiple reflection, natural single reflection, and natural multiple reflection. We should point out that even though PSU provides a reliable benchmark dataset for comparing symmetry detection algorithms, the ground truth it provides is not always complete and/or accurate. The results in Table 1 (see some examples in Fig. 3) show that our method is much better than Loy and Eklundh’s (LE [10]) and Liu *et al.*’s (LHXS [10]). Compared with Cho *et al.*’s (CL [9]) method, we have better performance for single symmetry detection. For multiple symmetries detection, our method is more conservative, sometimes detecting fewer symmetries, which explains the decrease in both the true positive and false positives rates. This conservative strategy can be explained by the strict criterion used to only create precise consensus sets. Nonetheless, the proposed method results are highly competitive. Notice that, as a post-processing, we could also adopt Loy and Eklundh’s region growing strategy [10], in order to expand and further improve the detected symmetries.

<sup>1</sup>The original PSU dataset contained 91 images [9, 10] but the original download link is broken. We obtained a version of the dataset containing 88 images at <http://vivid.cse.psu.edu/texturedb/gallery/album05>. Since code is not available for the methods in [9, 10], the results are not exactly comparable but serve as a performance indicator.

In non fronto-parallel symmetries, our algorithm is also able to obtain good results. We include several examples in Fig. 4. The algorithms that were used for comparing in the frontal setting are not designed for this type of symmetries and would only succeed in those images where the perspective effect is weak.

Recently, the TILT algorithm [23] was introduced, combining the concept of low-rank factorization with perspective estimation. Given a manually selected region of interest, it computes the homography that produces the sub-image with the lowest possible rank. In Fig. 5, we show that the detected symmetries with their consensus sets (supporting matches) can be used to automatically select input regions of interest for the TILT algorithm. This brings forth the relationship between these two concepts: symmetric regions have necessarily a low rank. Even though the regions are not completely flat, the TILT algorithm is capable of transforming the regions to be a quasi frontal view. Notice that using the rectifying homography as an initialization of the TILT optimization might also help further improve and stabilize the TILT results by obtaining a visually better local minimum.

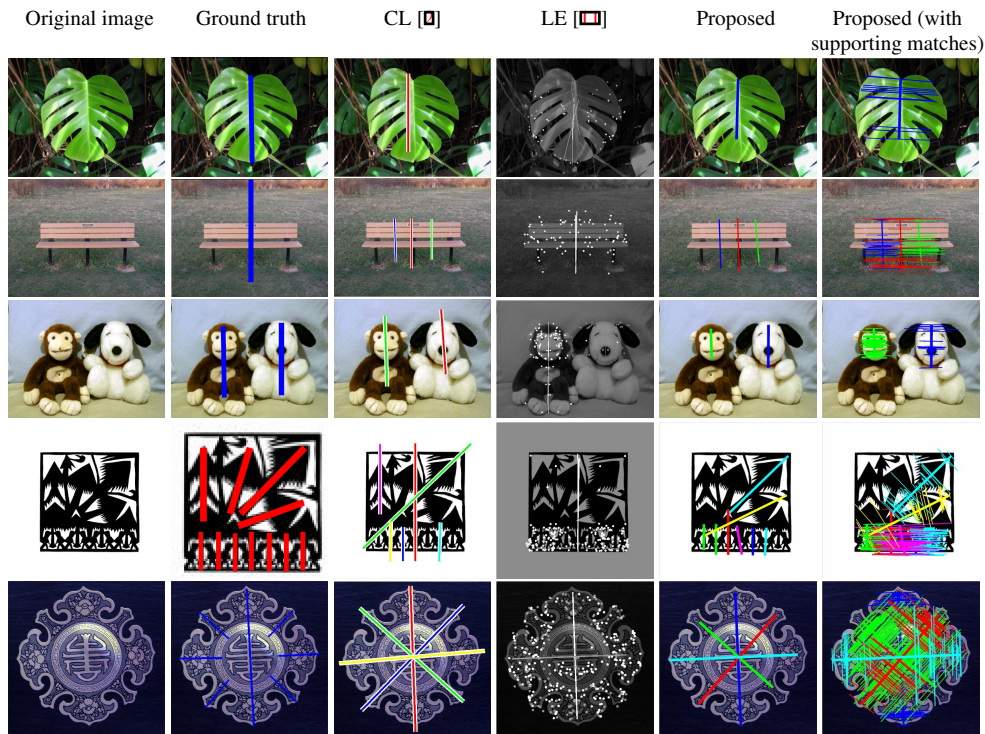


Figure 3: Several symmetry detection results on the PSU dataset. The images on the third row are borrowed from [24].

## 4 Conclusion

We presented a method for detecting bilateral symmetries in images. The method detects symmetries in a rectified image domain by sampling symmetry seeds in a non-uniform adaptive manner, and then building candidate consensus sets. Features are built from the consen-



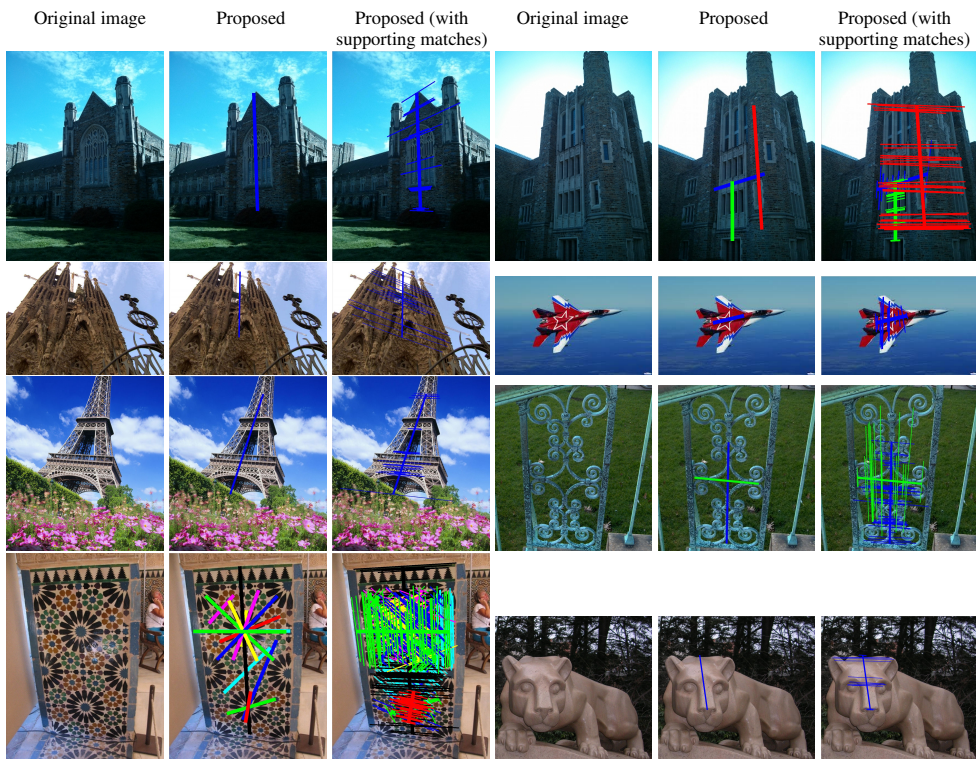


Figure 4: Several non fronto-parallel symmetry examples. Our algorithm is able to recover the symmetries, even when the perspective effect is non-negligible.




Figure 5: Relationship between low rank and symmetry in two examples of images with detected symmetries. The red bounding boxes represent the regions of the detected symmetry, which are used as automatic input to the TILT algorithm. The green bounding boxes represent the low-rank sub-images obtained with the TILT algorithm, projected back to the original image. We also show the low-rank components before re-projection.

sets and the final symmetries are detected via an agglomerative clustering algorithm. The method is able to detect single and multiple symmetries both in frontal and skewed (non fronto-parallel) viewpoints, achieving state-of-the-art results. We plan on extending the method to detect other symmetry types. We are also investigating a more unified way to combine the concept of symmetry and low-rank.

## References

- [1] G. Birkoff. *Aesthetic Measure*. Harvard University Press, Cambridge, MA, 1932.
- [2] M. Cho and K. Mu Lee. Bilateral symmetry detection via symmetry-growing. In *BMVC*, 2009.
- [3] R.W. Connors. Developing a quantitative model of human preattentive vision. *IEEE Trans Syst Man Cybern*, 19(6):1384–1407, 1989.
- [4] H. Cornelius and G. Loy. Detecting bilateral symmetry in perspective. In *POCV*, 2006.
- [5] H. Cornelius, M. Perdoch, J. Matas, and G. Loy. Efficient symmetry detection using local affine frames. In *SCIA*, 2007.
- [6] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In *CVPR*, 2010.
- [7] M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun ACM*, 24(6): 381–395, 1981.
- [8] Y. Keller and Y. Shkolnisky. An algebraic approach to symmetry detection. In *ICPR*, 2004.
- [9] A. Kuehnle. Symmetry-based recognition of vehicle rears. *Pattern Recognit Lett*, 12 (4):249–258, 1991.
- [10] Y. Liu, J. Hays, Y-Q. Xu, and H-Y. Shum. Digital papercutting. Technical Sketch. In *SIGGRAPH*, 2005.
- [11] G. Loy and J-O. Eklundh. Detecting symmetry and symmetric constellations of features. In *ECCV*, 2006.
- [12] S. Lee M. Park, P-C. Chen, S. Kashyap, A. Butt, and Y. Liu. Performance evaluation of state-of-the-art discrete symmetry detection algorithms. *CVPR*, 2008.
- [13] G. Marola. On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *IEEE Trans Pattern Anal Mach Intell*, 11(1):104–108, 1989.
- [14] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *BMVC*, 2002.
- [15] J.M. Morel and G. Yu. ASIFT: A new framework for fully affine invariant image comparison. *SIAM J Imaging Sci*, 2(2):438–469, 2009.

- [16] D. Reissfeld, H. Wolfson, and Y. Yeshurun. Context-free attentional operators: The generalized symmetry transform. *Int J Comput Vis*, 14:119–130, 1995.
- [17] C. Sun and D. Si. Fast reflectional symmetry detection using orientation histograms. *Real-Time Imaging*, 5:63–74, 1999.
- [18] Y. Sun and B. Bhanu. Reflection symmetry-integrated image segmentation. *IEEE Trans Pattern Anal Mach Intell*, 34(9):1827–1841, 2012.
- [19] R. Toldo and A. Fusiello. Robust multiple structures estimation with j-linkage. *ECCV*, 2008.
- [20] T. Tuytelaars, A. Turina, and L. Van Gool. Noncombinatorial detection of regular repetitions under perspective skew. *IEEE Trans Pattern Anal Mach Intell*, 25(4):418–432, 2003.
- [21] H. Weyl. *Symmetry*. Princeton University Press, 1952.
- [22] L. Xu, E. Oja, and P. Kultanen. A new curve detection method: Randomized Hough transform (RHT). *Pattern Recognit Lett*, 11(5):331–338, 1990.
- [23] Z. Zhang, A. Ganesh, X. Liang, and Y. Ma. TILT: Transform invariant low-rank textures. *Int J Comput Vis*, 99(1):1–24, 2012.
- [24] M. Zuliani, C.S. Kenney, and B.S. Manjunath. The multiransac algorithm and its application to detect planar homographies. In *ICIP*, 2005.