

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

Images often contain multiple reflective symmetries. We propose a method to detect multiple reflective symmetries at different scales and viewpoints based on the J-linkage framework [5], which combines ideas from RANSAC (robustness against outliers) and the Hough transform (multiple models detection through voting schemes).

How J-linkage works in symmetry detection? Given N reflective matches $\mathcal{M} = \{m_i\}_{i=1, \dots, N}$ in one image, 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. We thus obtain K symmetries, each with an associated consensus set (the subset of 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 is used to cluster the matches. Finally, each, large enough, cluster corresponds to a local symmetry.

How to rectify the distorted symmetry from two matches? In practice, the symmetry can be observed in a slanted view and thus undergoes some perspective distortion. In this more general case, $m_1 = \{\mathbf{p}_1, \mathbf{p}'_1\}$ and $m_2 = \{\mathbf{p}_2, \mathbf{p}'_2\}$ intersecting at the vanishing point \mathbf{v} are necessary to determine the symmetric axis (see Fig. 1(a)). We decompose the homography, which rectifies the distorted symmetry, into three parts:

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

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). \mathbf{H}_0 is computed such that it is as close as possible to an identity matrix and sends the vanishing point to the infinity. Then, \mathbf{R} and \mathbf{A} further transform the symmetry to the fronto-parallel view.

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. This observation is used to define a simple criteria to select inliers to build the consensus set (see Fig. 1(b)).

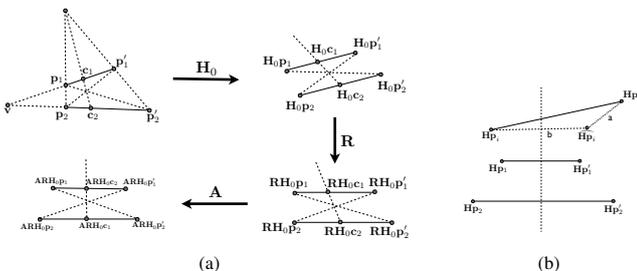


Figure 1: (a) Finding the rectification homography ($\mathbf{H} = \mathbf{A}\mathbf{R}\mathbf{H}_0$) from a pair of matches. (b) m_i is an inlier match if $\frac{a}{b}$ is smaller than some threshold η .

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. We use the following adaptive non-uniform sampling strategy. The first match is sampled according to the following mixed

	Synthetic Single				Synthetic Multiple			
	LE [3]	LHXS [2]	CL [1]	Proposed	LE [3]	LHXS [2]	CL [1]	Proposed
TP/GT	92%	62%	100%	100%	35%	28%	77%	67%
FP/GT	15%	0%	15%	0%	4%	8%	33%	10%

	Real Single				Real Multiple			
	LE [3]	LHXS [2]	CL [1]	Proposed	LE [3]	LHXS [2]	CL [1]	Proposed
TP/GT	84%	29%	94%	97%	43%	18%	68%	65%
FP/GT	68%	3%	69%	39%	44%	0%	17%	16%

Table 1: Performance comparison of several methods on the PSU dataset [4]. 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 [4] and [1].

probability, $\forall m_i = \{\mathbf{p}_i, \mathbf{p}'_i\} \in \mathcal{M}$,

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

where d_0 and l_0 indicate the scale in the descriptor domain and image domain at which we prefer to detect the symmetry, $\text{desc}(\mathbf{p}_i)$ is the descriptor of keypoint \mathbf{p}_i , Z_1 is a normalization factor, and, finally, σ_d and σ_l decide how strict the preferences are.

Since an image can contain multiple symmetries at different scales, we need to update the parameters of Eq. (6) to find all possible symmetries. 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 κ close to 1, followed by the renormalization of the probability:

$$\Pr(m_i) \leftarrow \kappa \Pr(m_i), m_i \in \mathcal{M}_0. \quad (3)$$

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

$$\Pr(m_j | m_i) = \frac{1}{Z_2} \exp \left(-\frac{1}{\sigma_c^2} (\|\mathbf{c}_i - \mathbf{c}_j\| - c_0)^2 \right), m_i \in \mathcal{M}, m_j \in \mathcal{M} \setminus \{m_i\}, \quad (4)$$

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 σ_c, c_0 control the shape of the probability function.

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.

Experiments. We compare the proposed method with three recent ones [1, 2, 3] on the PSU dataset [4], which is composed of 88 images. The results in Table 1 show that our method is much better than Loy and Eklundh's (LE [3]) and Liu *et al.*'s (LHXS [2]). Compared with Cho *et al.*'s (CL [1]) 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.

- [1] M. Cho and K. Mu Lee. Bilateral symmetry detection via symmetry-growing. In *BMVC*, 2009.
- [2] Y. Liu, J. Hays, Y-Q. Xu, and H-Y. Shum. Digital papercutting. Technical Sketch. In *SIGGRAPH*, 2005.
- [3] G. Loy and J-O. Eklundh. Detecting symmetry and symmetric constellations of features. In *ECCV*, 2006.
- [4] 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.
- [5] R. Toldo and A. Fusiello. Robust multiple structures estimation with j-linkage. *ECCV*, 2008.