

Curved Reflection Symmetry Detection with Self-validation

Jingchen Liu¹

Yanxi Liu^{1,2}

¹Department of Computer Science and Engineering

²Department of Electrical Engineering

The Pennsylvania State University

University Park, PA 16802, USA

{jingchen, yanxi}@cse.psu.edu

Abstract. We propose a novel, self-validating approach for detecting curved reflection symmetry patterns from real, unsegmented images. Our method benefits from the observation that any curved symmetry pattern can be approximated by a sequence of piecewise rigid reflection patterns. Pairs of symmetric feature points are first detected (including both inliers and outliers) and treated as ‘particles’. Multiple-hypothesis sampling and pruning are used to sample a smooth path going through inlier particles to recover the curved reflection axis. Our approach generates an explicit supporting region of the curved reflection symmetry, which is further used for intermediate self-validation, making the detection process more robust than prior state-of-the-art algorithms. Experimental results on 200+ images demonstrate the effectiveness and superiority of the proposed approach.

1 Introduction

Symmetry is pervasive in nature and man-made environments [1, 2]. It is one of the most important cues for human and machine perception of the world [1]. Automatic perception of symmetry patterns from images has been a standing research topic in computer vision. Reflection symmetry [2], as one of the four basic symmetries, is the most common and has received most attention in psychology as well as in computer vision [1]. Various applications utilize reflection symmetry such as face analysis [3], multi-target pattern analysis and tracking [4], vehicle detection [5] and medical image analysis [6].

Reflection symmetry detection algorithms dominate the literature of all types of symmetry detections [1, 7]. For example, Sun and Si [8] used histogram of gradient orientations to find the orientation of dominant reflection axis. Masuda, et. al. [9] explored edge features to measure symmetry similarity and Loy and Eklundh [10] matched feature points and then extracted reflection (and also rotation) symmetry patterns via clustering; Mitra et.al. [11] developed partial or approximate Euclidean reflection symmetry detection in subsampled 3D data.

Besides rigid reflection symmetry, Kanade in 1983 proposed the term *skewed symmetry* denoting reflection symmetry of an object going through global affine



Fig. 1. Some example images containing curved reflection pattern, including real-world/synthesized, segmented/unsegmented, nature/man-made object images.

or perspective skewing [12]. Symmetry recognition from global affinity and perspective distorted views has also been well studied in [13–17], where the reflection axis is assumed to be a straight line. In real world however, many symmetrical objects/patterns present curved reflection axes as shown in Figure 1. Automatically recognizing curved symmetry axis from unsegmented images is motivated by a wide range of applications. For example, symmetric region segmentation and curvature analysis from spine x-ray images, as well as leaves recognition and classification, can all benefit from a curved reflection detection algorithm.

Lee and Liu in [18] proposed the first, state-of-the-art curved glide-reflection symmetry extraction algorithm from real, unsegmented images. Their algorithm detects symmetric feature points first, which we refer to in this paper as symmetry ‘particles’ and then clusters these particles in the parameter space, subsequently fits a polynomial function to obtain the curved symmetry axis. The weakness of this approach is the ability against potential outliers (in some cases much more than the number of inliers) contained in the particles, which can seriously affects the robustness of clustering and curve fitting of the reflection axis. Besides, the polynomial fitting of symmetry axis misclassifies many inlier/outlier particles, as a result, the supporting region of the detected symmetry is not well defined thus making it hard to quantitatively assess the reliability of the detected pattern online.

Based on these facts and the abundance of real world curved reflection symmetries (Fig.1), we propose a curved reflection symmetry detection approach that explicitly selects symmetry ‘particles’ with local supporting regions and achieves more robust performance than [18] on curved reflection symmetry extraction by being able to effectively self-validate the detected results.

We adopt the bottom-up framework of [18, 10] that first detects and matches symmetric feature points to form symmetry particles (including both inliers and outliers), while build up symmetric regions in the higher level. The major novel advantage of detecting deformed symmetry patterns from bottom-up is that feature points are free of global deformation, meanwhile local deformation can be

handled by more sophisticated feature points such as SIFT [19], which is robust against scale change and rotation with good repeatability and high efficiency.

A crucial part of our approach is to discover a smooth path going through inlier particles on the image to approximate a valid curved reflection axis. One challenge is that the set of symmetry particles detected in the first step can be misleading. This is because feature point matching only considers local patches around the feature points, and symmetry, on the other hand, is a non-local, continuous feature [20]. There always exist many outlier feature point pairs, that only appear symmetrically in a small local region. To effectively validate the symmetry particles, region-based evaluation and verification are more robust and should be adopted. It can be seen that one symmetry particle is uniquely specified by a pair of feature points, while 2 particles, consisting of 4 feature points, form a closed quadrilateral region. If we approximate the local symmetry axis using straight line within the quadrilateral, a region-based reflection symmetry evaluation step can be done easily and reliably. Therefore given any pair of symmetry particles, we can quantitatively measure the symmetry-ness within the corresponding region, which we refer to as ‘consistency between symmetry particles’, and establish a graph structure with all vertexes representing symmetry particles and the edges representing a straight reflection axis between two particles and the weight on this edge indicating the consistency or local symmetry score. (Figure 2 -C,D)

By establishing the graph of linked symmetry particles, we turn this problem of curved symmetry pattern recognition into a problem of seeking a smooth path in the graph that maximize the symmetry property along the path. We will show in Section 2.3 that this is a global optimization problem and we thus propose a multiple hypothesis path sampling and pruning approach for real world curved reflection symmetry detection. Validation results on more than 200 images of three categories show superior curved reflection symmetry detection rates of our algorithm than [18].

One important advantage of explicitly selecting symmetry particles to approximate curved reflection is that we can obtain a well-defined continuous supporting region of reflection symmetry along the curved axis. As we use thin-plate spline (TPS) warping to rectify the axis, we can evaluate quantitatively and globally how symmetric the rectified region is, thus achieving self-validation, making the algorithm more robust.

2 Our Approach

The bottom-up framework for curved reflection symmetry pattern detection starts with recognizing symmetric feature point pairs. Each pair of feature points x_{i1}, x_{i2} provides us with a symmetry axis particle $l_i = \{x_i, \alpha_i\}$, where $x_{i1}, x_{i2}, x_i \in R^2$ are the image coordinates of the two symmetric feature points and their middle point, α_i is the orientation perpendicular to the line joining the points x_{i1} and x_{i2} , with ambiguity of angle π . In the next stage, we evaluate the pair-wise consistency among all symmetry particles and establish an undirected graph

$G\{V, E\}$, with $V = \{l_i\}$ being the set of all particles and any edge $(i, j) \in E$ means the line segment joining particle x_i and x_j reflects the symmetry property locally. The recognition of curved reflection patterns thus becomes a problem of discovering a smooth path from the graph G , which goes through a subset of its vertices (particles), $(l_{i_1}, l_{i_2}, \dots, l_{i_k})$, to approximate the curved reflection axis (Figure 2).

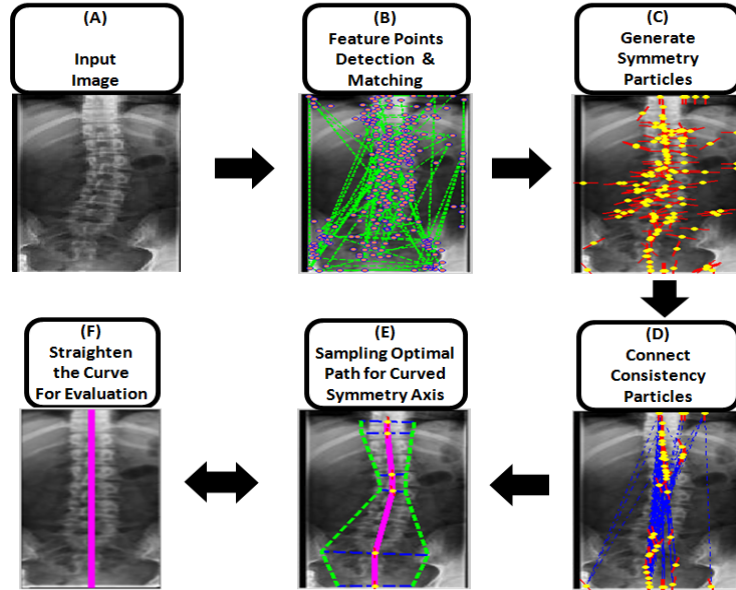


Fig. 2. The framework of our approach: (A)input image; (B)detected SIFT feature points marked as pink dots and successfully matched feature point pairs connected using green dashed lines; (C)representing feature points pairs as yellow particles with red short lines indicating the directions of potential reflection symmetry axis α_i ; (D)Maximally connected components in particle pairwise consistency graph G ; (E)Sampled optimal path from G ; (F)Rectified region via TPS warping.

2.1 Symmetry Particles Discovering

We adopt SIFT feature for effective symmetric points recognition since it is rotation and scale-invariant [19]. By rearranging the SIFT descriptor vectors v_i , we can describe the same local patch in the mirror image, denoted by $v_i^{(m)}$. The symmetry distance between two feature points is defined to be the Euclidian distance of the description vectors

$$d(i, j) = \|v_i - v_j^{(m)}\|_2 \quad (1)$$

For each point, we find top 3 best matches with smallest symmetry distance and then reject matches either having different scales or do not satisfy the angular constrains. Let a pair of SIFT feature points be $(x_{i1}, \phi_{i1}, s_{i1})$ and $(x_{i2}, \phi_{i2}, s_{i2})$, where ϕ_{i1} , ϕ_{i2} and ϕ_{i12} are the orientation angles of two feature points and the line connecting them, respectively. We specify the angular constraint that the orientation of two feature points should also be symmetric, which means $(\phi_{i1} + \phi_{i2})/2 \perp \phi_{i12}$. Each accepted pair of feature points then corresponds to a symmetry particle as illustrated in Figure 3 (a).




Fig. 3. An illustration of angle constrains and quadrilateral rectification

2.2 Generating Pairwise Consistency Graph

Given two symmetry particles, $l_i = (x_i, \alpha_i)$ and $l_j = (x_j, \alpha_j)$, and let the direction from x_j to x_i be α_{ij} , as illustrated in Figure 3 (b), we evaluate whether they form a near-symmetry region based on both of their geometric consistency and the appearance symmetry score. Assuming α_i and α_j are the tangents of the symmetry axis curve at locations x_i and x_j respectively, the geometric consistency requires that the curve be smooth, which means α_i , α_j and α_{ij} are along the similar directions, leading to the following two conditions:

$$|\alpha_i - \alpha_j| < TH_1 \quad (2)$$

$$|\alpha_{ij} - \frac{1}{2}(\alpha_i + \alpha_j)| < TH_2. \quad (3)$$

In our experiments, we set both thresholds to be $TH = \pi/8$. Once the pair of particles passed the geometric consistency, we rectify the local image patch to evaluate the appearance symmetry score by warping the quadrilateral formed by vertexes $x_{i1}, x_{i2}, x_{j1}, x_{j2}$ into an equilateral trapezoid, where the length of the parallel sides are $\|x_{i1} - x_{i2}\|_2$ and $\|x_{j1} - x_{j2}\|_2$ respectively, and the height is $\|x_i - x_j\|_2$, as illustrated in Figure 3 (c). TPS warping is used in our approach to deal with the most general transformation without assuming any specific cases like affine or perspective.

To evaluate the symmetry score of an equilateral trapezoid, we flip the trapezoid according to the middle axis and calculate normalized cross correlation

(NCC) with the original patch, which returns a score between -1 and 1 . If the NCC score is above a threshold (0.5 in our experiments), we impose an edge between particles l_i and l_j and record the NCC score as well as the area (supporting region) of the trapezoid for future use.

2.3 Multiple Hypothesis Path Sampling and Evaluation

The pairwise consistency graph can be further divided into several subgraphs based on connectivity. In cases where multiple symmetry patterns exist, each subset possibly contains one symmetry pattern. In cases of single symmetry pattern detection, we only focus on the subset with maximum number of vertices, which in most cases contains the most dominant pattern.

We then look for a smooth path within the subgraph(s) that maximizes the ‘symmetry’ along it. The symmetry score of a path can be obtained after we use TPS warping to rectify the whole path into connected equilateral trapezoids and evaluate its symmetry score the same way as we do for a single trapezoid. For the sake of computation time, it is also reasonable to approximate the path symmetry score using weighted summation of piecewise scores.

Let a path p going through N vertexes $p = (v_1, \dots, v_i, \dots, v_N)$, or $N - 1$ edges $p = (E_1, \dots, E_i, \dots, E_{N-1})$, with $E_i = (v_i, v_{i+1}, NCC_i, s_i)$, where NCC_i and s_i are the symmetry NCC score and the area of the trapezoid corresponding to the pair of particles l_i, l_{i+1} . We approximate the NCC score of a path using

$$s_p = \sum_{i=1}^{N-1} s_i \quad (4)$$

$$NCC_p = \frac{1}{s_p} \sum_{i=1}^{N-1} s_i \cdot NCC_i \quad (5)$$

To ensure the smoothness of the path, we require the turning angle at each vertex be less than a threshold of $\pi/5$, which means

$$\angle(x_i - x_{i-1}) - \angle(x_{i+1} - x_i) < \pi/5, \quad i = 2, 3, \dots, N - 1, \quad (6)$$

where $\angle()$ is the orientation of a vector, x_i correspond to the 2D coordinate of v_i in the image.

We define 2 criteria c_1 and c_2 given a path p for its ranking, one is the approximation of the path symmetry score, the other also favors paths covering more area:

$$c_1(p) = NCC_p \quad (7)$$

$$c_2(p) = NCC_p + \lambda \cdot \log(S_p) \quad (8)$$

The traditional graph solutions for finding optimal paths such as Dijkstra’s algorithm is not suitable here for the criteria in equations 7 or 8 and the smoothness condition in equation 6, all involves global information of the paths. The

enumeration of all possible pathes is also computationally unaffordable. As an alternative, we try to selectively sample the paths with high likelihood.

Once a path is initiated, we can enumerate the next valid vertexes to extend the current path. Each enumeration would generate an extended path hypothesis, each of which can be further extended recursively. Such an iterative approach forms a multiple hypothesis sampling of all possible pathes. The complexity of this sampling scheme grows exponentially and is unbounded depending on the density of the graph, thus we need to perform efficient pruning to cut unlikely paths in the first place. In each iteration after all current paths having been extended, we prune paths with low likelihood, and only keep a maximum number of K hypothesis in the pool. When $K = 1$, this becomes a greedy algorithm that starts at a random vertex and finds the local optimum; When $K = \infty$, we find global optimum by enumerating all valid paths in the graph that contains the initialization vertex. In our experiment, we take paths ranked top 100 either under criteria 1 or criteria 2. The reason we set up 2 criteria is that although we favor longer curves finally, we want to protect potential paths in the pool before they have been fully extended. This punning policy effectively bounds the computation within linear complexity meanwhile providing us with good enough solutions.

The algorithm for sampling and pruning paths recursively is illustrated in Table 1, note that once a path p is extended by a new vertex v_{N+1} , its NCC score approximation can be updated also in an efficient recursive form,

$$NCC_p^{(new)} = \frac{NCC_p^{(old)} \cdot s_p^{(old)} + NCC_N \cdot s_N}{s_p^{(old)} + s_N}. \quad (9)$$

For the complete paths (can not be extended any more) returned by the multiple hypothesis sampling, instead of using equation 5 for approximation, the one with highest score according to equation 8 is selected to be the final result.

For the final candidates returned by the multiple-hypothesis sampler, we use criteria c_2 in 8 to rank them, however, instead of using NCC score approximation as equation 5, we use TPS warping to straighten the curved reflection axis to calculate the accurate NCC score. based on the c_2 ranking, the path that produces highest score is selected as final result.

3 Experimental Results and Comparison

We test our algorithm on 210 images including 2 subcategories of the Swedish Leaf dataset [21]—one has curved reflection symmetry pattern on every single leaf (*Quercus rober*, 75 images), the other has curved reflection symmetry pattern among the leaves on a branch (*Sorbus aucuparia*, 75 images), as well as a human-spine X-ray dataset containing 30 images that we collect ourselves. We also collected a set of miscellaneous real-world and synthesized images with curved reflection symmetry patterns (30 images). In addition, we compare our approach with the method in [18] on the same image sets.

Algorithm 1: Multiple Hypothesis Path Sampling

Input: $G(V, E)$, with $V = \{v_i\}$, $E = \{(v_{k1}, v_{k2}, NCC_k, s_k) | k = 1, 2, \dots, K\}$;**Initialize:**Randomly pick a vertex v_i , mark it as incomplete and put it into the path pool;**while** exists incomplete paths in the pool**for** all incomplete paths in the pool**if** the path can be further extended

replace the path with all valid extend paths;

update path NCC score (Eq.9) and mark them as incomplete.

else

Mark the path as complete.

Prune paths in the path pool with low likelihood score.

Return:All paths remaining in the pool.

3.1 Our Results

A representative selection of our results on the Swedish leaf dataset, spine X-ray dataset and miscellaneous images is shown in Figures 4, 5, 6¹. For each image, we tag the detection results of curved reflection axis, as well as its support region specified by symmetric feature points. We also straighten the curved reflection axis and show the rectified image in the supporting region on the right.

(H) and (I) in Figure 6 show 2 different reflection patterns being detected from the same image. This is achieved when we separate the consistency graph G into several connected components, each of which could be checked for the existence of reflection patterns.

It can be seen from these results and rectified images that piece-wise rigid reflection is a reasonable approximation for curved reflection. Our method is effective and robust in selecting a small subset of inlier symmetry particles explicitly to represent the reflection pattern. If necessary, curve fitting can be further applied on the selected inlier particles to obtain a smoother curve.

Although we introduced several heuristic thresholds in the algorithm, e.g., $TH_1 = TH_2 = \pi/8$, $TH_{NCC} = 0.5$, they are mainly for efficiency concerns. A stricter threshold helps saving time by pruning bad hypothesis in an earlier stage; Relaxing the thresholds would result to more outliers being included in the path-sampling stage. However the final results are relatively insensitive against the threshold changes, which is because our approach has the self-validation ability, making it perform robustly in finding the correct inlier particles even when the outliers are more than inliers (which is usually the case).

In addition, we make some extra tests on images containing human perceived straight-reflection symmetries (Figure 7). The results demonstrate a more accurate capture of the slight deformations of the seemingly straight reflection axes

¹ See our project page for a complete set of Data and Results at <http://vision.cse.psu.edu/research/curvedSym/index.shtml>

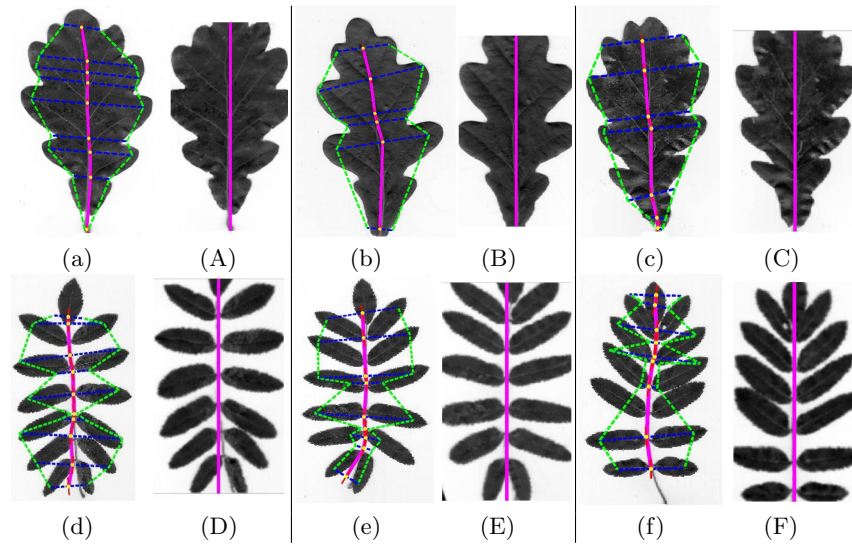


Fig. 4. Results of Swedish leaf data set. top row: single leaf with curved reflection symmetry pattern; bottom row: multiple leaves form curved reflection symmetry pattern; a-f: original images tagged with detected curved reflection axis(pink), supporting region(green); A-F: rectified images with a straightened reflection axis(pink).

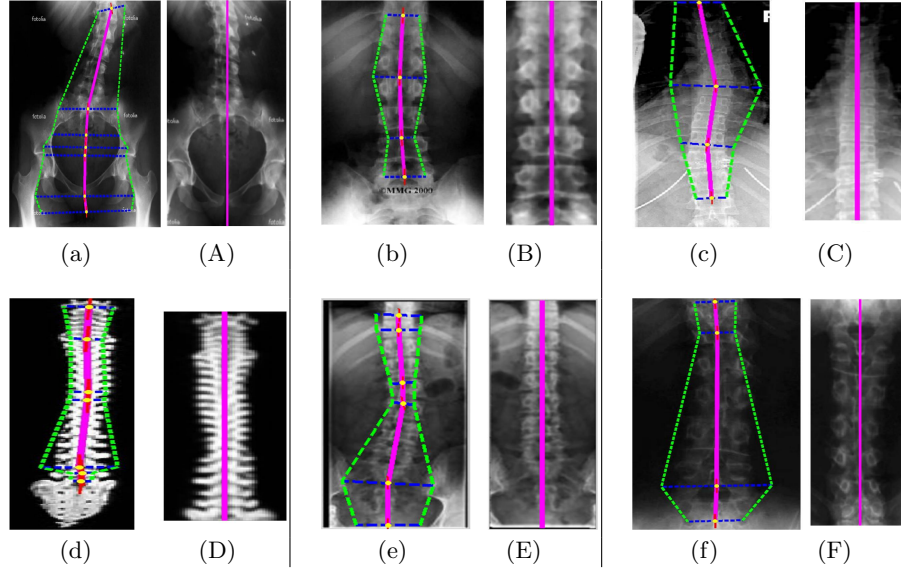


Fig. 5. Results of Spine X-ray data set. a-f: original images tagged with detected curved reflection axis(pink), supporting region(green); A-F: rectified images with a straightened reflection axis(pink).

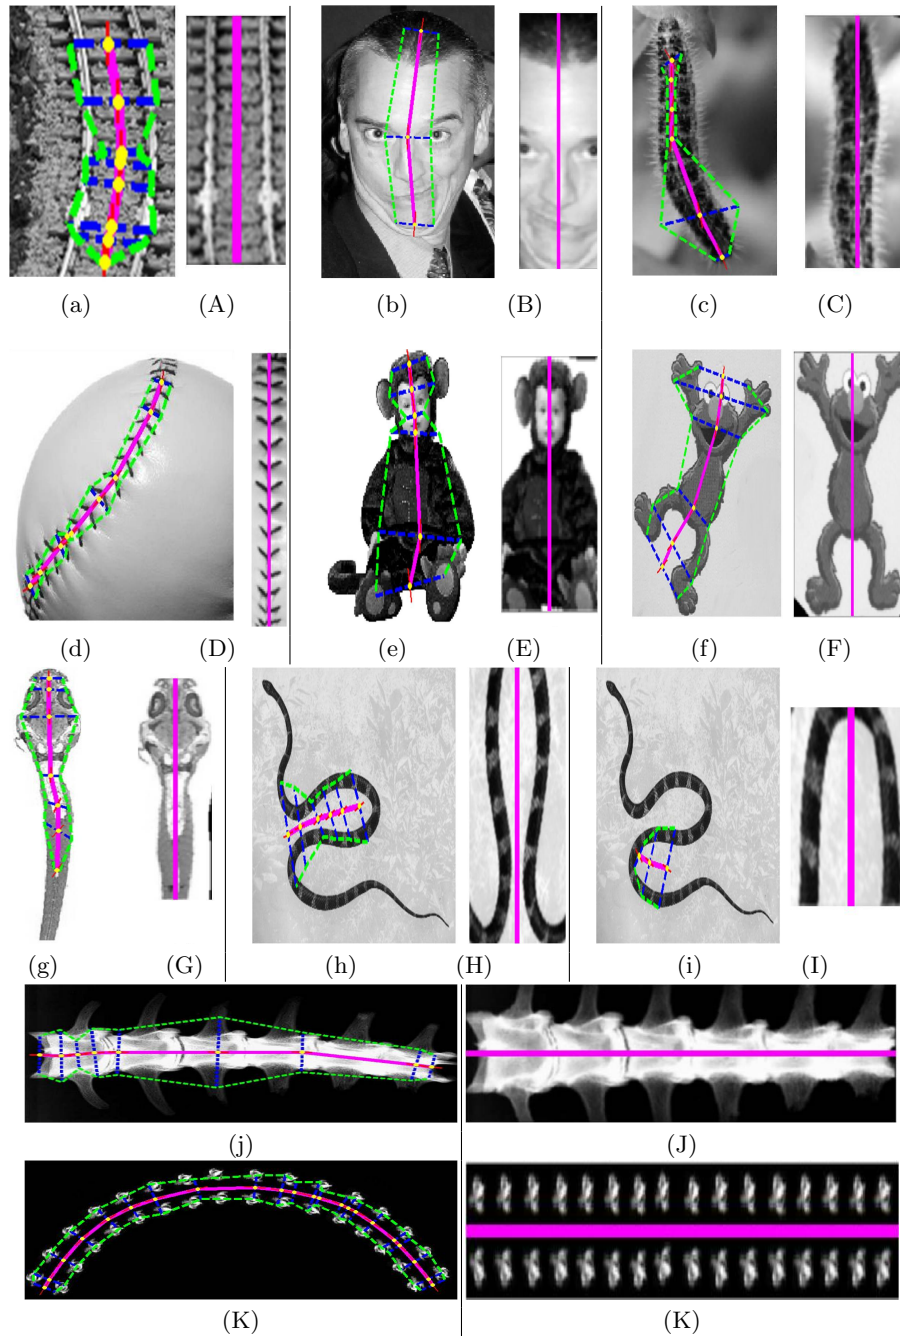


Fig. 6. Results of miscellaneous images. a-k: original images tagged with detected curved reflection axis(pink), supporting region(green); A-K: rectified images with a straightened reflection axis(pink).

by our piece-wise curve approximation algorithm, indicating the rarity of perfectly straight reflection symmetries in real world.




Fig. 7. Examples of detecting almost-straight reflection symmetries.




Fig. 8. Examples of failed cases, where (a) failed due to severe background clutter; (b) and (c) failed because not enough feature points are detected in the first place thus missed the dominant reflection pattern, while finding some local reflection symmetries.

Some failure examples are also shown in Figure 8, where (a) failed due to severe background clutter; (b) and (c) failed because not enough feature points are detected in the first place so that the dominant reflection symmetry pattern is not (fully) recognized. From the experiment, most of the failure cases are due to not enough feature points being extracted. Therefore to make our approach more robust, multiple types of feature point detection can be adopted here e.g., Harris-Laplace [22], which detects corner like points and is complement with Hessian-Laplace (blob-like) feature points.

3.2 Quantitative Evaluation and Comparison with [18]

We also apply Lee and Liu’s approach [18] to the same data sets and make quantitative comparisons. For each image, we tag it success if more than 4/5

Dataset	Leaf dataset	Spine dataset	Miscellaneous images	Overall
# images	150	30	30	210
proposed	83.3%	80.0%	73.3%	81.4%
Lee & Liu [18]	40.0%	66.7%	70.0%	48.1%

Table 1. Success rates of our proposed algorithm and Lee & Liu’s [18].

of the curved reflection axis is detected, and failure otherwise. Our method has higher success rate on all three datasets as reported in Table 1.

Some of the detection results of Lee and Liu’s in [18] are also shown here in Figure 9 for an intuitive comparison. It can be seen that by defining an explicit supporting region and TPS warping to rectify the curve, we achieve self-validation in our method thus being more robust against outliers and yield to better performance.




Fig. 9. Comparisons of our proposed approach with Lee & Liu in [18]. a-c: Lee & Liu’s approach; A-C: our approach.

4 Conclusions

In this paper, we propose a bottom-up curved-reflection symmetry detection approach, starting from recognizing symmetric points pairs (particles) in the bottom level and extract a consistent structure among the particles to form the symmetry pattern in the higher level. Multiple-hypothesis sampling and pruning method is shown to be effective in discovering the optimal curved structures from real world images. As a by-product, we obtain the supporting regions from selected particles and use them for self-validation. Quantitative evaluation and comparison against state-of-the-art algorithm on 210 real images confirm the superior robustness of our proposed approach.

Acknowledgement. We thank Lee and Liu [18] for providing their source code. This work is supported in part by an NSF grant IIS-0729363 and a gift grant to Dr. Liu from Northrop Grumman Corporation.

References

1. Liu, Y., Hel-Or, H., Kaplan, C.S., Gool, L.V.: Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision* **5** (2010) 1–195
2. Weyl, H.: *Symmetry*. Princeton University Press (1952)
3. Mitra, S., Liu, Y.: Local facial asymmetry for expression classification. In: *Proceedings of the 2004 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'04)*. Volume 2. (2004) 889 – 894
4. Liu, J., Liu, Y.: Multi-target tracking of time-varying spatial patterns. In: *Computer Vision and Pattern Recognition Conference (CVPR '10)*. (2010)
5. Kuehne, A.: Symmetry-based recognition of vehicle rears. *Pattern Recogn. Lett.* **12** (1991) 249–258
6. Mancas, M., Gosselin, B., Macq, B.: Fast and automatic tumoral area localisation using symmetry. *Acoustics, Speech, and Signal Processing, 2005. Proceedings. (ICASSP '05)*. *IEEE International Conference on* **2** (2005) 725–728
7. Park, M., Lee, S., Chen, P.C., Kashyap, S., Butt, A.A., Liu, Y.: Performance evaluation of state-of-the-art discrete symmetry detection algorithms. In: *Computer Vision and Pattern Recognition Conference (CVPR)*. (2008) 1–8
8. Sun, C., Si, D.: Fast reflectional symmetry detection using orientation histograms. *Real-Time Imaging* **5** (1999) 63–74
9. Masuda, T., Yamamoto, K., Yamada, H.: Detection of partial symmetry using correlation with rotated-reflected images. *Pattern Recognition* **26** (1993)
10. Loy, G., Eklundh, J.: Detecting symmetry and symmetric constellations of features. In: *European Conference on Computer Vision*. (2006) II: 508–521
11. Mitra, N., Guibas, L., Pauly, M.: Partial and approximate symmetry detection for 3d geometry. *ACM Transactions on Graphics* **25** (2006) 560–568
12. Kanade, T., Kender, J.R.: Mapping image properties into shape constraints: skewed symmetry, affine-transformable patterns, and the shape-from-texture paradigm. In: *Human and Machine Vision*. (1983) 237–257
13. Shen, D., Ip, H., Teoh, E.: Robust detection of skewed symmetries. In: *International Conference on Pattern Recognition vol.3*. (2000) 1010–1013
14. Van Gool, L., Proesmans, M., Moons, T.: Mirror and point symmetry under perspective skewing. In: *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition*. (1996) 285–292
15. Carlsson, S.: Symmetry in perspective. In: *European Conference on Computer Vision vol.1*. (1998) 249–263
16. Lei, Y., Wong, K.: Detection and localisation of reflectional and rotational symmetry under weak perspective projection. *Pattern Recognition* **32** (1999) 167–180
17. Cornelius, H., Loy, G.: Detecting bilateral symmetry in perspective. In: *CVPRW*. (2006) 191
18. Lee, S., Liu, Y.: Curved glide-reflection symmetry detection. In: *Computer Vision and Pattern Recognition Conference (CVPR '09)*. (2009) 1046 – 1053
19. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60** (2004) 91–110
20. Zabrodsky, H., Peleg, S., Avnir, D.: Symmetry as a continuous feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17** (1995) 1154–1166
21. Söderkvist, O.J.O.: *Computer vision classification of leaves from swedish trees*. Master's thesis, Linköping University (2001)
22. Mikolajczyk, K., Schmid, C.: Scale & affine invariant interest point detectors. *International Journal of Computer Vision (IJCV)* **1** (2004) 63–86