

Performance Evaluation of State-of-the-Art Discrete Symmetry Detection Algorithms

Minwoo Park[†], Seungkyu Lee[†], Po-Chun Chen[†], Somesh Kashyap^{*}, Asad A. Butt[†] and Yanxi Liu^{†*}

[†]Department of Computer Science and Engineering, ^{*}Department of Electrical Engineering
The Pennsylvania State University

{mipark, seunlee, pocchen, asad, yanxi}@cse.psu.edu, suk203@psu.edu, <http://vision.cse.psu.edu/evaluation.htm>


Abstract

Symmetry is one of the important cues for human and machine perception of the world. For over three decades, automatic symmetry detection from images/patterns has been a standing topic in computer vision. We present a timely, systematic, and quantitative performance evaluation of three state of the art discrete symmetry detection algorithms. This evaluation scheme includes a set of carefully chosen synthetic and real images presenting justified, unambiguous single or multiple dominant symmetries, and a pair of well-defined success rates for validation. We make our 176 test images with associated hand-labeled ground truth publicly available with this paper. In addition, we explore the potential contribution of symmetry detection for object recognition by testing the symmetry detection algorithm on three publicly available object recognition image sets (PASCAL VOC'07, MSRC and Caltech-256). Our results indicate that even after several decades of effort, symmetry detection in real-world images remains a challenging, unsolved problem in computer vision. Meanwhile, we illustrate its future potential in object recognition.


1. Motivation

Symmetry is an essential concept in perception and a ubiquitous phenomenon presenting itself in all forms and scales in the real world (Figures 1 and 2), from galaxies to atomic structures [9]. Symmetry is considered a pre-attentive feature [5] that enhances object recognition. Much of our understanding of the world is based on the perception and recognition of repeated patterns that are generalized by the mathematical concept of symmetries [29].

The development of symmetry detection algorithms has a long history in computer vision. The earliest attempt at detection of a single bilateral reflection symmetry even predates computer vision itself [1]. Even though symmetries



(1) reflection symmetry group detection [17]: multiple symmetry axes of local regions are detected one-by-one.



(2) reflection (left) and rotation (right) symmetry detection [19]
Top row: input images.



(3) rotation symmetry detection [25].

Left image: input image.

Figure 1. Sample results for multiple-symmetry detection without segmentation/pre-processing from each of the three algorithms to be evaluated in this paper.

take several basic forms (rotation, translation, reflection and glide-reflection), the detection of bilateral reflection symmetry (mirror-symmetry) and its skewed version from images has been dominant in computer vision for several decades [7, 22, 14, 2, 24, 10, 12, 26, 21, 31, 28, 3, 16]. In

spite of years of effort, we are still short of a robust, widely applicable “symmetry detector” that can parallel other types of computer vision tools, such as an “edge” or “corner” detector. Furthermore, we have yet to see a systematic, quantitative evaluation and a publically available test-image database to gauge the progress in this important, widely applicable, albeit seemingly illusive research direction.

Within the past couple of years, we observe a surge of new symmetry detection papers in several related fields [13, 17, 19, 20, 23, 25, 27]. While each paper demonstrates certain experimental results of the proposed algorithm, without a systematic evaluation of different symmetry detection algorithms against a common image set under a uniform standard, our understanding of the power and limitations/pitfalls in state of the art symmetry detection algorithms remains partial. This situation seriously hinders a wide applicability of existing symmetry detection algorithms and any concrete, measurable improvements thereafter. In this paper, we make a specific effort on proposing and implementing a performance evaluation system for three state of the art discrete symmetry detection algorithms:

1. **Digital Papercutting**[17] (SIGGRAPH 2005) for reflection symmetry detection (Figure 1 (1));
2. **Detecting Symmetry and Symmetric Constellations of Features** [19] (ECCV06) for both rotation and reflection symmetry detection (Figure 1 (2));
3. **Detecting Rotational Symmetries** [25] (ICCV05) for rotation symmetry detection (Figure 1 (3)).

The main reason we choose these algorithms is that they all go beyond single bilateral reflection symmetry detection, they are all published recently (2005-2006), their source code is publically available and their functionalities differ yet are comparable. One attractive feature of all three algorithms is their claim of *multiple-symmetry-detection* from an *unsegmented* real image¹.

In this work, we have tested these symmetry detection algorithms on a total of more than 800 images, including 176 images with hand-labeled ground truth and images from three publicly available databases: PASCAL VOC’07[8], MSRC Object class recognition database [30] and Caltech-256 [11] data sets.

To the best of our knowledge, our work is the first in computer vision to evaluate multiple symmetry detection algorithms systematically and quantitatively against a common image set and a set of evaluation standards. All test images used in this paper, with hand-labeled ground truth, are made publicly available for future algorithm validation (<http://vision.cse.psu.edu/evaluation.htm>). Our work establishes the first benchmark in symmetry detection. Our quantitative algorithm performance evaluation results (Sections 5) clearly indicate that symmetry detection research has yet



Figure 2. Examples of images with rotation (left column), reflection and rotation (middle column), and translation plus rotation/reflection symmetries (right column). Their symmetry groups are cyclic, dihedral and 2D crystallographic groups respectively. Top-row displays synthetic images while the bottom-row contains real-world photos.

to reach its desired goal of multiple-symmetry detection from un-segmented real images (best sensitivity is around 30-40% and best net classification rate when taking false positives into consideration is less than 20%). Meanwhile, our exploration on object recognition and object class categorization databases (Section 6) suggests a potential role a successful symmetry detector could play for object recognition in computer vision.

2. Symmetry and Symmetry Groups

Mathematically speaking, a *symmetry* g of a set S is an isometry such that $g(S) = S$ [6], i.e. the transformation g keeps S invariant as a whole while permuting its parts. S could be a purely geometrical entity or something geometric with additional attributes, like color or texture. All symmetries of S form a mathematical group [6] called the *symmetry group* of S . There are four atomic symmetries in 2D Euclidean space R^2 : translation, rotation, reflection and glide-reflection [29]. The *discrete symmetry groups* in 2D Euclidean space can be further divided into (1) *point groups* where all transformations in the group keep at least one point of S invariant. These are the cyclic groups (C_n) containing rotation symmetries only (Figure 2 left); and dihedral groups (D_n , where n is the order of its cyclic subgroup, and $2n$ is the cardinality of the dihedral group) containing both reflection and rotation symmetries (Figure 2 middle). (2) *space groups* [6] where no point in R^2 is kept invariant by all members of the group. These are the symmetry groups containing translation symmetries: the seven frieze (1D translation, reflection, rotation, glide-reflection)

¹Though [17] was designed for images of papercut patterns only.

and the 17 wallpaper symmetry groups (2D translation, rotation, reflection, glide-reflection) (Figure 2 right) [29, 6].

Euclidean symmetries and symmetry groups so defined may seem to be restrictive in computer vision applications, choosing these types of symmetry detection algorithms to evaluate in this initial validation work is justified by the fact that (1) this type of Euclidean symmetries and symmetry groups form the basis of the mathematical theory on symmetry [6, 29] thus can be immediately and formally verified for completeness systematically; (2) a large amount of such symmetries does appear in real world images (Figure 5) as shown in various image databases (PASCAL VOC'07[8], MSRC Object class recognition database [30] and CALtech-256 [11] data sets). (3) this type of rigid symmetries form a subset of more complicatedly deformed symmetries (e.g. affine or perspective skewed, or locally deformed via diffeomorphisms), thus we are testing the necessary condition of a good symmetry detection software; our results will provide practical guidance for the future development of more robust, distorted-symmetry detection algorithms.

In this paper, we focus on the evaluation of algorithms that detect discrete point-group symmetries, *i.e.* the cyclic and dihedral symmetry groups composed of reflection and rotation symmetries. We consider the simplest bilateral symmetry as a special case of the dihedral group D_n where $n = 1$ indicating an identity group as its cyclic subgroup. No particular distinction is made between cyclic and dihedral symmetries (Figure 2 left column vs middle) since existing algorithms do not discriminate these two cases.

3. Reflection and Rotation Symmetry Detection Algorithms

We briefly describe each symmetry detection algorithm evaluated in this paper (Figure 1).

3.1. Digital Papercutting [17]

This algorithm is originally designed for the analysis of images of artistic papercutting patterns. Thus, it uses edge-based features. The algorithm first exhaustively searches through the parameter space of potential reflection axes (in polar coordinates ρ, d) to identify all single reflection symmetries by voting for pairwise matches. In addition, the algorithm identifies the structures of the reflection axes to discover dihedral symmetry group D_n (if several reflection axes intersect in one point) and, frieze symmetry group (reflection axes are parallel and placed with equal distance).

3.2. Detecting Symmetry and Symmetric Constellations of Features [19]

This is a feature-based reflection and rotation symmetry detection algorithm, which takes advantage of local ori-

ented features expressed as SIFT keys [18]. The basic symmetry detection technique uses pairwise matching and voting for symmetry foci (single reflections and C_n -type symmetries) in a Hough transform fashion. It also estimates the n in cyclic group C_n but it does not make the distinction between C_n and D_n type symmetry groups.

3.3. Detecting Rotational Symmetries [25]

This algorithm filters an input color image into a gradient vector flow (GVF) field and conducts the extraction and matching of local features in the GVF field. The symmetry detection is formulated, once again, as a voting scheme for the centroids of C_n symmetries, though it does not distinguish between cyclic C_n or dihedral D_n type symmetries nor does it discover the order of the cyclic group n automatically.

4. Evaluation Methodology

Our evaluation of symmetry detection algorithms involves three major steps: (1) collect a set of test images; (2) hand label their ground truth symmetry (reflection axes, rotation centers and folds); (3) define a quantitative evaluation criteria for measuring the success of the symmetry detection results. We run each symmetry detection algorithm on all test images and compute the success rates based on a set of labeled-ground-truth (Section 4.2) and a well-defined scoring function (Section 4.3).

4.1. Test Image Sets Selection

To test the applicability of each symmetry detection algorithm, we provide a carefully selected image set that demonstrates unambiguous object symmetries, with diverse visual properties: synthesized versus real images, clean versus textured regions, frontal versus skewed views, similar versus contrasting color intensities.

















The whole 176-test-image set is divided according to two standards: (1) synthetic versus real images; and (2) images containing a single dominant symmetry versus multiple symmetries. Given two types of symmetries (rotation and reflection), we have a total of eight different subcategories of test images (Figure 3).

Furthermore, we explored the potential role a symmetry detection algorithm might play in object and object class recognition/categorization by testing and evaluating the algorithm with best performance on selected images from PASCAL VOC'07[8], MSRC Object class recognition database [30] and CALtech-256 [11] (Section 6).

4.2. Ground Truth

Labeling ground truth in complex real world images for symmetry detection evaluation is a non-trivial task both

Reflection Symmetry Detection Test Set

Image Number	Original Image	Ground Truth	Description /Challenges	Loy and Eklundh 2006	Liu et al 2005	Loy and Eklundh 2006	Liu et al. 2005
						GT/TP/FP	GT/TP/FP
Ref_SS12			Single bilateral symmetry; Highly textured lower half			1/1/0	1/1/0
Ref_SM5			Multi-reflection symmetry axes with global and local supporting regions			16/2/0	16/1/0
Ref_RS22			Real photo, off-centered single symmetry; Cluttered BG			1/0/4	1/0/0
Ref_RM27			Real photo, multiple local symmetry axes; Skewed			5/0/1	5/1/0

Rotation Symmetry Detection Test Set

















Image Number	Original Image	Ground Truth	Description /Challenges	Loy and Eklundh 2006	Prasad and Davis 2005	Loy and Eklundh 2006	Prasad and Davis 2005
						GT/TP/FP	GT/TP/FP
Rot_SS9			Synthetic image, single symmetry: C3; Textured			1/1/2	1/1/0
Rot_SM1			Synthetic multiple symmetries: D8, D16; Color variation; Global and local symmetries			5/3/0	5/0/2
Rot_RS4			Real photo, single symmetry: D24; Cluttered; Background inside the rotation region			1/1/0	1/0/1
Rot_RM5			Real photo, multiple symmetries: D13, D15; Highly textured FG and BG			14/3/0	14/1/0

Figure 3. Sample images and results from our test image set. We also provide labeled ground truth, descriptions of computational challenges, and the numbers of ground truth (GT), and detected true positive (TP) and false positives (FP). The complete test image set can be found in <http://vision.cse.psu.edu/evaluation.htm>.

conceptually and effort-wise. At a starting point, we follow some simple rules in this process: (1) only the visually obvi-

ous dominant symmetries (to human raters) are labeled; (2) only the top-level/foreground symmetries are labeled; and

(3) only the semantically meaningful, rather than accidental symmetries, are labeled.

The final labeled GTs on the test images are from the consensus of all raters, involving 10 students from two separate classes of one year apart, who have an intro-level understanding of symmetry and symmetry detection algorithms. None of the raters is an author of the symmetry detection algorithms tested.

For rotation and reflection symmetries in our selected 176-image set, we have labeled the reflection axes, rotation center and the fold of the corresponding rotation symmetry group (Figure 3) though this last value is not explicitly evaluated in this paper.

For a qualitative evaluation of how a symmetry detection algorithm performs on a public object recognition data set (Section 6), we focus on the intended object only and count the number of perceived symmetries (by human) versus the number of detected symmetries by the algorithm. For example, in the case of ‘face class’ from MSRC data set [30], only one reflection symmetry centered on the face is considered as the ground truth in each image; and in the case of cars (side views) only the two visible wheels are considered having rotation symmetries. Since we only compute sensitivity in this case (disregard false positives), the results are optimistic given the minimum number of ground truths counted while reflecting the detection success for the most important symmetries of concern.

4.3. Evaluation Measurement

We propose the following formula to compute a score of symmetry detection success rate on each image:

$$S_K = \frac{(TP - K * FP)}{GT} \quad (1)$$

where 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. K is a weight that determines how strongly the false positives need to be punished for different perception tasks. For simplicity, we measure two values in this work: $S_0 = \frac{TP}{GT}$ is the commonly known *sensitivity* [15]; and the *false positive rate* $R_{FP} = \frac{FP}{GT}$.

5. Quantitative Evaluation

We quantitatively evaluate reflection and rotation symmetry detection algorithms on their respective test image sets (176 images) using the code provided by each original author. In our experiments, we use the default parameter settings without modification from image to image. Since different algorithms respond differently to image sizes, we

have tested each algorithm on four image scales (from 1 to 1/4 of the original size) and choose the best result to report.

We compare algorithms from [19] and [17] on reflection symmetry detection on a set of 91 images, and compare [19] and [25] for rotation symmetry detection on a total of 85 images. These images are divided into four categories (1) synthetic images with single symmetry, (2) synthetic images with multiple symmetry, (3) real images with single symmetry, and (4) real images with multiple-symmetry. Some sample images for reflection and rotation symmetries can be found in Figure 3. Among these images, there are 7 reflection symmetry images and 14 rotation symmetry images for which at least one of the detection algorithms fails. Thus the statistical results provided in Figure 4 are based on 84 images for reflection symmetry detection and 71 images for rotation symmetry detection.


We choose to report two quantitative measurements the sensitivity [15] S_0 (formula 1 when $K=0$), and the false positive rate R_{FP} . Since it is nearly impossible to list all the non-symmetries in an image, we use R_{FP} as an indirect measure for specificity [15]. More numerical details on the performance statistics in tabular form can be found in [4].

We have also measured the time used by each symmetry detection algorithm regardless of the output correctness. According to our log file: for two images of 530 by 531 and 800 by 600, it takes 31088 and 95114 seconds respectively for [25] to generate its output (on the same images, the other two symmetry detection algorithms take about 20-40 seconds, i.e. about three orders of magnitude faster). Based on the speed and the detection rates (Figure 4) [19] has the overall best performance out of the three algorithms tested.

6. Qualitative Evaluation on Object Recognition Image Databases

Given the universal existence of symmetry in all sizes and forms in real world objects, it is our ultimate goal to unveil the role of symmetry detection in object recognition. As an initial step, we test the reflection and rotation symmetry detection algorithm [19], the most superior algorithm from our quantitative evaluation, on three object recognition and object class categorization data sets: (1) PAS-CAL Visual Object Classes Recognition Challenge 2007 (VOC2007) [8]; (2) Caltech-256 Object Category Dataset [11]; and (3) MSRC Object class recognition database A and B1[30].

Some sample results can be seen in Figure 5 where the reflection symmetry axes or the centers of rotation symmetry are identified. Table 1 provides a set of S_0 values for subsets of images from the three databases. These rates are obtained by counting the potential existing symmetry of the object (sometimes none) in each image (total ground truth)



S_0 :Sensitivity R_{FP} :False positive rate Real: Real Image Syn:Synthetic Image Alg#1: Loy and Eklundh 2006 Alg#2: Liu et al. 2005 Alg#3: Prasad and Davis 2005

Figure 4. The pairwise reflection and rotation symmetry detection algorithms evaluation on our 176 test-images with labeled ground truth.

and how many of those are detected by the algorithm (true positives). This is done by two raters with over 90% agreement. As observed in Table 1, for certain object classes (e.g. face) the symmetry detector is doing rather well (sensitivity $> 80\%$) in locating the midline of the face in the image regardless of background clutters, lighting variations (indoor/outdoor) and oclusions (top rows of Figure 5). This S_0 value is also consistent with the single symmetry detection S_0 rate from real image results evaluated in Figure 4. Even for the non-typically-thought-of perfect symmetrical objects, e.g. trees, the trunk of the tree is identified as the axis of reflection symmetries with sensitivity = 55-67%. In addition, the reflection or rotation symmetry of umbrella, towers car wheels, tennis rackets (local) and flowers are detected under adverse conditions (Figure 5). These symmetry cues provide great potential as unique, discriminative features for object localization, segmentation, recognition and object class categorization.

Table 1. Reflection and Rotation Symmetry Detection Evaluation (estimated sensitivity S_0 only for specific symmetry types)

Data Set	Object Class	Reflection	Rotation
CALtech256 60 selected images	Various	59%	65%
PASCAL VOC'07 58 selected images	Various	42%	45%
MSRC (A) 330 images 18 classes	Door/Window	59%	–
	Tree	55%	–
	Car	–	57%
	Bike	–	50%
MSRC (B1) 240 Images 9 classes	Face (30)	83%	–
	Tree (30)	67%	–
	Car (30)	–	40%
	Bike (30)	–	30%

7. Summary and Discussion

Overall, without counting false positives, the reflection/rotation symmetry detection algorithms on images with **single symmetry** can reach a mean sensitivity value S_0 of 92-100% on synthetic and 84-88% on real images [19]. For images with **multiple symmetries**, the best mean sensitivity S_0 of [19] drops drastically to no better than 43% (Figure 4), meaning that more than half of the symmetries in the image are missed.

A more serious problem is perhaps the low net symmetry detection rate $S_0 - R_{FP}$ in real world images (Figure 4): no more than 26% for single symmetry detection and no more than 19% for multiple-symmetry detections. Given the high false positive rates exhibited in these state of the art symmetry detection algorithms, the field is in desperate need of more discriminative symmetry detectors or higher-level reasoning tools to discard false detections.

Our exploration in Section 6 by applying [19] on images from three publicly available object recognition and object class categorization databases [8, 30, 11] demonstrates the unique potential of symmetry detection in object recognition. We have observed distinct patterns and spatial layouts of symmetry detection from images of different object classes, and high success rates on certain object classes (Figure 5, Table 1). All these suggest that, just as symmetry detection plays an important role in human perception, successful automated symmetry detection may lead to perceptually and practically enhanced object recognition, segmentation and object class categorization in high-level computer vision.

Our effort on a quantitative and qualitative symmetry detection algorithm evaluation reported in this paper is only the first step towards establishing a benchmark for researchers working in this area to make solid, measurable

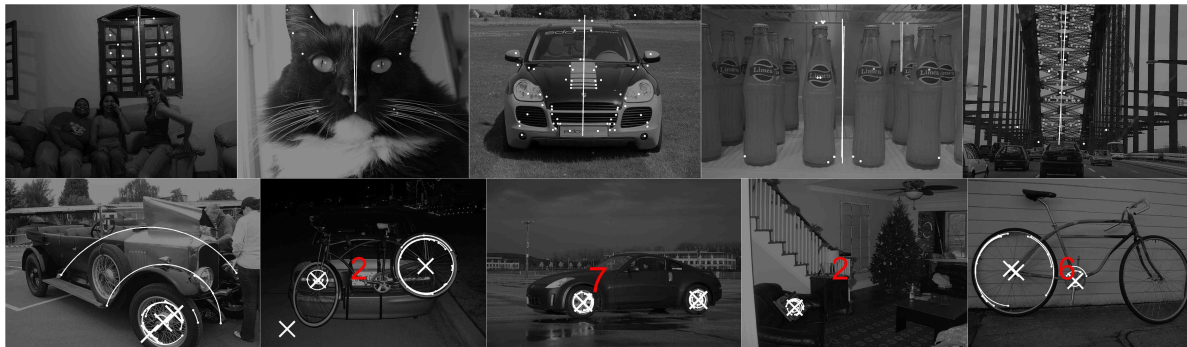
(A.1) MSRC object category dataset B1 [30]. Category: Face



(A.2) MSRC object category dataset B1 [30]. Category: Tree



(B) PASCAL VOC'07 [8]



(C) Caltech-256 [11]

Bicycles (rotation) Tennis rackets (rotation) Umbrella (reflection)

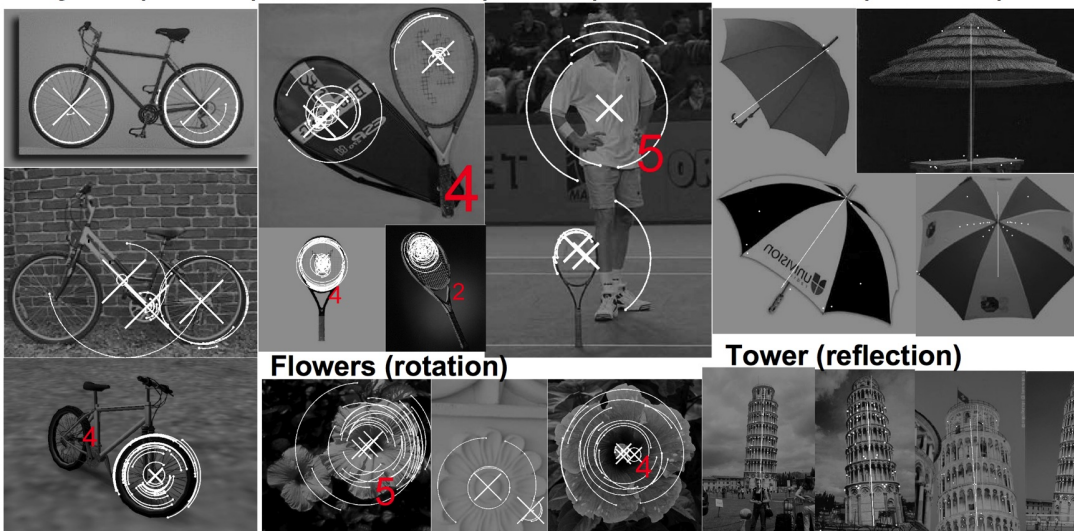


Figure 5. Sample (reflection and rotation) symmetry detection results from publicly available real image data sets. See also the quantitative results shown in Table 1. It is interesting to note the high success rate on faces, and the axes of symmetry for trees are mostly vertical.

progress in the future. As modest and limited as it is, we can already conclude that (1) symmetry detection research in computer vision still has a long way to go to become a reliable, widely usable and effective feature detector; (2) symmetry detection has a great potential to make an impact on object recognition; and (3) a much larger test image set and a finer categorization of different types of images are needed to validate the stability, the strength and the weakness of past and future symmetry detection algorithms.

Acknowledgment

We are grateful to all authors of [19, 25, 17] for contributing their original code for evaluation. We thank Daniel DeMenthon and Zhuowen Tu for suggesting the use of object recognition datasets. This paper originated from a student's homework for course CSE 598B/PSU on "Computational Symmetry" (Fall 2006, Fall 2007). This work is supported in part by NSF grants IIS-0535324, IIS-0729363 and a Fulbright Scholarship to Asad A. Butt.

References

- [1] G. Birkoff. *Aesthetic Measure*. Harvard University Press, Cambridge, MA, 1932.
- [2] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *The International Journal of Robotics Research*, 3(3):36–61, Fall 1984.
- [3] S. Carlsson. Symmetry in perspective. In *ECCV 1998 vol. I*, pages 249–263, 1998.
- [4] P. Chen, J. Hays, S. Lee, M. Park, and Y. Liu. A quantitative evaluation of symmetry detection algorithms. Technical Report PSU-CSE-07011 (double listed as Tech. report CMU-RI-TR-07-36), The Pennsylvania State University, State College, PA, 2007.
- [5] R. Connors and C. Ng. Developing a quantitative model of human preattentive vision. *SMC*, 19(6):1384–1407, 1989.
- [6] H. Coxeter. *Introduction to Geometry*. Wiley, New York, second edition, 1980.
- [7] L. S. Davis. Understanding shape: Angles and sides. *IEEE Trans. Computers*, 26(3):236–242, 1977.
- [8] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- [9] M. Gardner. *The new ambidextrous universe: symmetry and asymmetry, from Mirrow reflections to superstrings*. W.H. Freeman and Company, 1979.
- [10] J. Gauch and S. M. Pizer. The intensity axis of symmetry and its application to image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence archive*, 15(8):753–770, August 1993.
- [11] G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. Technical Report 7694, California Institute of Technology, 2007.
- [12] A. Gross and T. E. Boult. Analyzing skewed symmetries. *International Journal of Coomputer Vision*, 13(1):91–111, September 1994.
- [13] J. Hays, M. Leordeanu, A. Efros, and Y. Liu. Discovering texture regularity as a higher-order correspondence problem. In *European Conference on Computer Vision (ECCV'06)*, 2006.
- [14] T. Kanade. Recovery of the 3-dimensional shape of an object from a single view. *Artificial Intelligence*, 17:75–116, 1981.
- [15] M. Kutner, C. Nachtsheim, J. Neter, and W. Li. *Applied Linear Statistical Models*. McGraw-Hill Irwin, 2005.
- [16] Y. Lei and K. Wong. Detection and localisation of reflectional and rotational symmetry under weak perspective projection. *Pattern Recognition*, 32(2):167–180, 1999.
- [17] Y. Liu, J. Hays, Y. Xu, and H. Shum. Digital papercutting. In *SIGGRAPH Technical Sketch*. ACM, 2005.
- [18] D. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. of Comp. Vis.*, 60(2):91,110, 2004.
- [19] G. Loy and J. Eklundh. Detecting symmetry and symmetric constellations of features. In *European Conference on Computer Vision (ECCV'04), Part II, LNCS 3952*, pages 508,521, May 2006.
- [20] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3d geometry. *ACM Trans. Graph.*, 25(3):560–568, 2006.
- [21] D. Mukherjee, A. Zisserman, and J. Brady. Shape from symmetry—detecting and exploiting symmetry in affine images, 1995.
- [22] R. Nevatia and T. O. Binford. Description and recognition of curved objects. *Artif. Intell.*, 8(1):77–98, 1977.
- [23] J. Podolak, P. Shilane, A. Golovinskiy, S. Rusinkiewicz, and T. A. Funkhouser. A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.*, 25(3):549–559, 2006.
- [24] J. Ponce. On characterizing ribbons and finding skewed symmetries. In *Proc. Int. Conf. on Robotics and Automation*, pages 49–54, 1989.
- [25] V. Prasad and L. Davis. Detection rotational symmetries. In *IEEE International Conference on Computer Vision (ICCV)*, pages 346–352, 2005.
- [26] D. Reisfeld, H. Wolfson, and Y. Yeshurun. Context free attentional operators: the generalized symmetry transform. *The International Journal of Computer Vision*, 14:119–130, 1995.
- [27] S. Thrun and B. Wegbreit. Shape from symmetry. In *ICCV*, pages 1824–1831, 2005.
- [28] L. Van Gool, T. Moons, and M. Proesmans. Mirror and point symmetry under perspective skewing. In *The Proceedings of CVPR*, pages 285–292. IEEE Computer Society, 1996.
- [29] H. Weyl. *Symmetry*. Princeton University Press, Princeton, New Jersey, 1952.
- [30] J. Winn, A. Criminisi, and T. Minka. Object categorization by learned universal visual dictionary. In *IEEE Intl. Conf. on Computer Vision (ICCV)*, 2005.
- [31] H. Zabrodsky, S. Peleg, and D. Avnir. Symmetry as a continuous feature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(12):1154–1165, December 1995.