

2017 ICCV Challenge: Detecting Symmetry in the Wild

Christopher Funk^{1,*} Seungkyu Lee^{2,*} Martin R. Oswald^{3,*} Stavros Tsogkas^{4,*}
Wei Shen⁵ Andrea Cohen³ Sven Dickinson⁴ Yanxi Liu¹

¹Pennsylvania State University, University Park, PA, USA

²Kyung Hee University, Seoul, Republic of Korea ³ETH Zürich, Switzerland

⁴University of Toronto, Canada ⁵Shanghai University, China

¹{funk, yanxi}@cse.psu.edu ²seungkyu@khu.ac.kr ³{moswald, acohen}@inf.ethz.ch
⁴{tsogkas, sven}@cs.toronto.edu ⁵wei.shen@t.shu.edu.cn

Abstract

Motivated by various new applications of computational symmetry in computer vision and in an effort to advance machine perception of symmetry in the wild, we organize the third international symmetry detection challenge at ICCV 2017, after the CVPR 2011/2013 symmetry detection competitions. Our goal is to gauge the progress in computational symmetry with continuous benchmarking of both new algorithms and datasets, as well as more polished validation methodology. Different from previous years, this time we expand our training/testing data sets to include 3D data, and establish the most comprehensive and largest annotated datasets for symmetry detection to date; we also expand the types of symmetries to include densely-distributed and medial-axis-like symmetries; furthermore, we establish a challenge-and-paper dual track mechanism where both algorithms and articles on symmetry-related research are solicited. In this report, we provide a detailed summary of our evaluation methodology for each type of symmetry detection algorithm validated. We demonstrate and analyze quantified detection results in terms of precision-recall curves and F-measures for all algorithms evaluated. We also offer a short survey of the paper-track submissions accepted for our 2017 symmetry challenge.

1. Introduction

The real world is full of approximate symmetries appearing in varied modalities, forms and scales. From insects to mammals, intelligent beings have illustrated effective recognition skills and smart behaviors in response to symmetries in the wild [9, 15, 38, 48], while computer vi-

sion algorithms under the general realm of *computational symmetry* [24] are still lagging behind [13, 25].

Our *Symmetry Detection in the Wild* challenge, affiliated with the International Conference in Computer Vision (ICCV) 2017 in Venice, Italy, is the third in a series of symmetry detection competitions aimed at sustained progress quantification in this important subfield of computer vision. The first symmetry detection competition [51], funded through a US NSF workshop grant, was held in conjunction with CVPR 2011, and offered the first publicly available benchmark for symmetry detection algorithms from images. The second symmetry competition, held during CVPR 2013 [22, 52], started to build comprehensive databases of real world images depicting reflection, rotation and translation symmetries respectively. In addition, a set of standardized evaluation metrics and automatic evaluation algorithms were established, solidifying the computational foundation for validating symmetry detection algorithms.

A historic overview of symmetry detection methods can be found in [25, 32]. There has been much recent work in symmetry detection since the last symmetry competition in 2013 [22], including new deterministic methods [2, 46, 49], deep-learning methods [14, 42], and other learning-based methods [45]. New applications include Symmetry reCAPTCHA [13], 3D reconstruction [8, 12, 43], image segmentation [4, 19], and rectification and photo editing [27, 37]. Many of these algorithms are featured in our challenges as baseline algorithms. Levinshtein *et al.* [18] detected straight, ribbon-like local symmetries from real images in a multiscale framework, while Lee *et al.* [17] extended the framework to detect curved and tapered local symmetries. Pritts *et al.* [37] detect reflection, rotation and translation symmetry using SIFT and MSER features. The symmetries are found through non-linear optimization and RANSAC. Wang *et al.* [49] use local affine in-

*Contributed equally, order chosen alphabetically

variant edge correspondences to make their algorithm more resilient to perspective distortion contours. Teo *et al.* [45] detect curved-reflection symmetry using structured random forests and segment the region around the curved reflection. Sawada and Pizlo [39] exploit mirror symmetry in a 2-D camera image for 3-D shape recovery.

Motivated by various new applications of computational symmetry in computer vision, our third challenge expands our training/testing data sets to include 3D data, establishing the most comprehensive and largest annotated datasets for symmetry detection to date. We expand the types of symmetries to cover reflection, rotation, translation [25] and medial-axis-like symmetries in 2D and 3D synthetic and real image data, respectively. We also distinguish symmetry annotations between discrete, binary pixel labels and densely-distributed, continuous firing fields reflecting a gradation in degrees of symmetry perception. We have further refined our evaluation method to include F-measures in all standardized precision-recall curves for comparison.

In terms of the symmetry challenge organization, we have established a challenge-and-paper dual track mechanism where both algorithms and articles on symmetry-related research are solicited. For the challenge track, we have quantitatively evaluated 11 challenge-track submissions against 13 baseline algorithms. In paper track, we have accepted five paper-track submissions after an extensive review process. Detailed information, datasets and results of this symmetry challenge can be found on the workshop website¹.

2. Symmetry Challenge Track

We have divided our evaluation of the datasets in the symmetry challenge track into sparse versus dense/continuous labels, as well as 2D versus 3D symmetries.

2.1. General Evaluation Methods

For all challenge track evaluations, we use the standard precision-recall and F-score evaluation metrics [14, 22, 29, 47]. *Precision* measures the number of *true positives*: this is the number of detected medial points that are actually labeled as positives in the ground-truth:

$$P = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}. \quad (1)$$

Recall measures the number of ground-truth positives that are successfully recovered by the algorithm:

$$R = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}. \quad (2)$$

Intuitively, precision is a measure of the accuracy of each detection, while recall measures detection completeness.

¹<https://sites.google.com/view/symcomp17/>

These two measures give us a quantitative means of evaluating each algorithm. To gain further insight into the differences between the algorithms, we can evaluate precision and recall by altering threshold values in the evaluation or prediction confidences to create a plot of a PR-curve that illustrates the trade-off between precision and recall.

One can summarize the performance of an algorithm (and select the optimal threshold) using the harmonic mean of precision and recall, which is called the F-measure or F-score:

$$F = \frac{2 \cdot P \cdot R}{P + R}. \quad (3)$$

F-measure offers a convenient and justified single-value score for system performance comparison.

2.2. 2D Symmetry - Sparse Evaluation

The 2D sparse evaluated symmetry competition includes challenges on the detection of *reflection* and *translation* symmetry. The total number of images and symmetries for each task is shown in Table 1.

2D Challenge Sparse	Type	# Images	# Symmetries
Reflection	single	100/100	100/100
	multiple	100/100	384/371
Translation	frieze	50/49	79/85

Table 1. The total number of images and symmetries within each 2D challenge in the training/testing sets.

2.2.1 Reflection and Translation Datasets

We have expanded and added new data sets beyond previous symmetry competitions [22]. The images are collected from the Internet and are annotated by symmetry researchers. For reflection symmetry, we divide the analysis into images containing either a single symmetry or multiple reflection symmetries. For translation, we tested the state-of-the-art algorithm(s) for 1D translation (frieze) symmetry.

The reflection symmetry annotations are line segments defined by two endpoints (Figure 1). Each translation symmetry annotation is a grid of points connected to create a lattice of quadrilaterals (Figure 2). For baselines, reflection uses Loy and Eklundh [26] and Atadjanov and Lee [3], and translation uses Wu *et al.* [54]. The evaluation metrics for each of these are similar to the previous symmetry competition [22].

Reflection Evaluation Metrics For the evaluation of reflection axis detection, we measure the angle difference between the detected and ground-truth axes and the distance from the center to the ground truth line segment. We use the same threshold values t_1 (angle difference) and t_2 (distance) used in [22]. Multiple detections for one ground truth



Figure 1. Example annotations of the 2D Reflection dataset. Blue lines are reflection symmetry axes.




Figure 2. Example annotations of the 1D translation (frieze) symmetry dataset.

axis are counted as one true positive detection, but none of them is counted as a false positive. We vary the confidence of the detections in order to create a precision-recall curve for each algorithm.

Translation Evaluation Metric We use the same distance-minimizing cost-function to align the detected and ground truth lattices as [35]. A detected quadrilateral tile is correct if it is matched with a ground truth lattice tile and the ratio of the matched tile areas is between 40% and 200%. We calculate the tile-success-ratio (TSR) [22] for each detected lattice. Similar to [22], we calculate true positives as images where a lattice is detected with a $TSR > \tau$, where τ is a threshold we vary from [0,1]. False positives are images where the best detected lattice has $TSR \leq \tau$, and false negatives are images where there was no lattice detected in the image.

2.2.2 Sparse Evaluation Results

The results for the reflection symmetry challenge are shown in Figures 3 and 4. Baseline methods achieve the best results on both single (Atadjanov and Lee [3] $F=0.52$) and multiple (Loy and Eklundh [26] $F=0.30$) reflection symmetry detections with respective highest F-measures. Loy and Eklundh [26] has shown robust performance on various datasets, including our previous competition. Atadjanov and Lee is one of the state-of-the-art methods reported in the literature. Elawady [11] shows a higher recall rate than others on the single symmetry dataset and is the top-performing




Figure 3. PR curve on 2D Single Reflection Symmetry Dataset. The algorithms are Michaelson and Arens [30], Elawady *et al.* [11], Guerrini *et al.* [16], Cicconet *et al.* [7], Loy and Eklundh [26] (baseline), and Atadjanov and Lee [3] (baseline).




Figure 4. PR curve on 2D Multiple Reflection Symmetry Dataset. The algorithms are Michaelson and Arens [30], Elawady *et al.* [11], Loy and Eklundh [26] (baseline), and Atadjanov and Lee [3] (baseline).

method among all challenge submissions on both 2D reflection symmetry detection datasets.

For 1D translation symmetry, the baseline outperformed the submission of Michaelson and Arens [30] (Figure 5). These results indicate much room for improvement for the detection of frieze patterns in natural images.

2.3. 2D Symmetry - Dense Evaluation

For dense evaluation the annotation is a binary map and the algorithm output is a confidence map, thresholded at multiple values to create the PR curve, with both of the maps have the same size as the input image (Figure 6). The evaluation is conducted per pixel rather than per reflection




Figure 8. Images and ground-truth annotations (in yellow) from the SKLARGE dataset. Only skeletons of foreground objects are annotated.



Figure 9. Image (left), ground-truth segmentation (middle) and ground-truth medial axis (right) from the BMAX500 dataset. In BMAX500, we do not distinguish between foreground and background.

Dataset	# Images	fg/bg	# Symmetries
	train/valid/test	distinction	train/test
Sym-COCO [14]	250/-/240	-	1535/1469
BMAX500 [46]	200/100/200	no	-
SKLARGE [41]	746/245/500	yes	-

Table 2. The total number of images in the dense evaluation dataset, if the datasets have foreground (fg) or background (bg) distinction in the segmentations, and the number of symmetries.

to distinguish between foreground and background, which makes BMAX500 an appropriate benchmark for evaluating a more general, class-agnostic medial axis detection framework. This is an important difference with respect to the SK-LARGE dataset, which is particularly focused on *object* skeletons.

2.3.3 Dense Evaluation Metrics

The algorithms examined in this part of the challenge output a real-valued map of probabilities or “symmetry/medial point strength” at each location in the image, rather than binary yes/no decisions. We turn this soft confidence map into a binary result of detected reflection symmetry/medial axis, by thresholding it at different values and plot the precision-recall curves as described in Section 2.1.

Detection slack. Both the ground-truth and the detected medial axes are thinned to single-pixel width in order to standardize the evaluation procedure. Now, consider two false positives returned by the tested algorithm, which are 1 pixel and 10 pixels away from a ground-truth positive. It would be unreasonable to penalize both of these false detections in the same way, since the first one is much closer to a true reflection symmetry axis/medial point.




Figure 10. PR curve for the Sym-COCO Reflection dataset for the 240 images and all GT labels (solid line) and for the subset of 111 reflection symmetry images with GT labels containing at least 20 labelers (dashed line), and the maximum F-measure values (dot on the line). The algorithms are the challenger Michaelsen and Arens [30] (MA), Funk and Liu [14] (Sym-VGG and Sym-Res - the baselines), Loy and Eklundh [26] (LE), Tsogkas and Kokkinos [47] (MIL), Teo *et al.* [45] (SRF), and Shen *et al.* [42] (FSDS).

We “forgive” such wrong yet reasonable detections by introducing a detection slack: all detected points within d pixels from a ground-truth positive are considered as true positives. We typically set d as 1% of the image diagonal [29].

2.3.4 Dense Evaluation Results

The Sym-COCO reflection challenge results are shown in Figure 10. The scores are the mean precision and recall, calculated among the images. The new challenger algorithm by Michaelsen and Arens [30] did not fair well in the competition and was surpassed by other algorithms. The algorithms which incorporate learning using additional images from outside this training set fared much better (all but Loy and Eklundh [26] and Michaelsen and Arens [30]) and the deep learning approaches (Funk & Liu [14] and Shen *et al.* [42]) predictably did the best. Funk & Liu [14], the baseline algorithm, took the top spot in the competition.

The results for medial axis detection are shown in Figures 11 and 12. The challenger algorithms beat the baseline algorithms for both datasets. In general, we observe that recent methods based on supervised deep learning outperform other learning-based and unsupervised, bottom-up approaches [47, 46].

2.4. 3D Symmetry

This part of the symmetry detection challenge only considers *reflection symmetries* of single 3D objects or within larger 3D scenes given as polygonal meshes.




Figure 11. **PR curves on the BMAX500 dataset.** The algorithms are: MIL-color [47] (baseline), AMAT [46] (baseline), and RSRN [21]. Human agreement on the dataset (extracted by comparing one human annotation to all the others) is also included. The performance of methods that do not involve threshold selection is represented as a single dot, which corresponds to the optimal F-score.




Figure 12. **PR curves on the SKLARGE dataset.** The algorithms are: MIL-color [47] (baseline), LMSDS [41] (baseline), and SegSkel [23]. Dots correspond to the optimal F-score.

2.4.1 Datasets and Annotation

The 3D dataset is split according to the following three different properties:

- **Synthetic vs. Real Data:** The annotation of real data is costly while large labeled synthetic datasets, which are often required for learning-based methods, can be easily obtained. Therefore, it is natural to select a combination of the two. The synthetic data is a collection of free, publicly available 3D models obtained from Archive3D [50]. The real datasets are kinect scans of real world scenes and are taken from the publicly available datasets that accompany two papers by

Choi *et al.* [5, 6] as well as from the work of Speciale *et al.* [43].

- **Global vs. Local Symmetries:** The difference between global and local symmetries is that the symmetry property respectively holds either for the entire domain or only for a local region of the domain. In the literature, local symmetries are sometimes also called partial symmetries [31, 32].
- **Training vs. Test Data:** The test and training datasets have similar properties and data statistics. Corresponding ground truth data is made publicly available for the training dataset. Furthermore, performance evaluations on the test dataset will be performed upon request and subsequently published on the workshop website [53].

Figure 13 depicts some example scenes from the 3D dataset. The total number of scenes and symmetries for each task is shown in Table 3.

3D Synthetic Data Annotation. The *synthetic data for global symmetries* only contains single objects. Ground truth symmetries were found by sampling a set of symmetry planes and rejecting the ones for which the planar reflective symmetry score [36] is below a threshold. These objects are originally axis-aligned which may bias learning-based methods. We therefore provide a dataset of over 1300 axis-aligned objects and their annotations as well as their randomly rotated counterparts. The *synthetic data for local symmetries* is a collection of scenes composed of objects from the global symmetry dataset. The scenes were gen-




Figure 13. Overview of available scenes in the 3D symmetry dataset which is split by *synthetic vs. real data* and *local vs. global symmetries*, shown with ground truth symmetry annotations.

3D Challenge	Type	# Scenes	# Symmetries
Global Reflection	Synthetic	1354/441	1611/614
	Real	20/20	21/22
Local Reflection	Synthetic	200/200	1939/2239
	Real	21/21	44/46

Table 3. The total number of scenes and symmetries within each 3D challenge in the training/testing sets.

erated with a script that places a desired number of symmetric objects with arbitrary translation, rotation and scale on top of a table. The script is available on the challenge website [53] and allows for the generation of an arbitrary amount of training data. A precomputed training dataset with 200 scenes is directly available for download.

3D Real Data Annotation. The real world data was annotated manually with the help of a 3D editor and model viewer that directly overlays the semi-transparent geometry of the reflected scene. In this way, one can instantly assess the quality of the fit while adjusting the symmetry plane.

2.4.2 Evaluation Metrics

Similar to the 2D case, the 3D planar reflective symmetries are evaluated according to the position and orientation of the symmetry plane with respect to the ground truth. A correct detection (*true positive*) is credited if both the position of the symmetry plane center and the orientation are sufficiently close to the ground truth plane.

Let (c, n) denote a symmetry plane given by the center point $c \in \mathbb{R}^3$ on the plane and the plane’s normal vector $n \in \mathbb{R}^3$, and let (c_{GT}, n_{GT}) be the corresponding ground truth symmetry. The symmetry is rejected if the angular difference between the symmetry normals is above a threshold θ , i.e. if $\arccos(|n \cdot n_{GT}|) > \theta$. In practice, symmetries are given by three points $x_0, x_1, x_2 \in \mathbb{R}^3$ which span the symmetry plane and also define a parallelogram that bounds the symmetry. Center point and normal are then given by $c = \frac{1}{2}(x_1 + x_2)$, $n = (x_1 - x_0) \times (x_2 - x_0)$. A symmetry is also rejected if the distance of the tested center c to the ground truth plane restricted to the bounding parallelogram is larger than a threshold: $\|c - \Pi_{GT}(c)\|_2 > \tau$.

Precision-recall curves are generated by linearly varying both thresholds within the intervals $\theta \in [0, 45^\circ]$ and $\tau \in [0, 2] \cdot \min\{x_1 - x_0, x_1 - x_0, x_{1GT} - x_{0GT}, x_{1GT} - x_{0GT}\}$.

2.4.3 Results

Ecins *et al.* [10] submitted results for the real and synthetic test data set for local symmetries, which are shown in Figure 14 and Figure 15, respectively. The curve in Figure 14 is very short since there are not many variations due to the




Figure 14. PR curve on the real local symmetry test dataset. Results for Ecins *et al.* [10]. The dot corresponds to the optimal F-score.




Figure 15. PR curve on the synthetic local symmetry dataset. Results for Ecins *et al.* [10]. The dot corresponds to the optimal F-score.

small size of the test dataset. Generally, the method computes symmetries with high accuracy for the ones it has detected. Figure 16 presents the results by Cicconet *et al.* [7] on the synthetic global test dataset.

3. Symmetry Paper Track

In the paper track of the workshop, full-length submissions were each reviewed by two members of the Organizing and/or Advisory Committees. In all, five papers were accepted for inclusion in the proceedings and presentation at the workshop.

In “Wavelet-based Reflection Symmetry Detection via Textural and Color Histograms” [11], Elawady *et al.* address the problem of detecting global symmetries in an image, in which extracted edge-based features are used to vote for symmetry axes based on color and texture information in the vicinity of the extracted edges.

In “SymmSLIC: Symmetry Aware Superpixel Segmentation” [34], Nagar and Raman offer a new twist on the su-




Figure 16. **PR curve on the synthetic global symmetry test dataset.** Showing the results of Cicconet *et al.* [7]. The dot corresponds to the optimal F-score.

perpixel segmentation problem. If a set of corresponding pixels can be detected that exhibit approximate reflective symmetry about an axis, these points can serve as seeds for a SLIC-inspired superpixel segmentation algorithm where superpixels that encompass these symmetric seeds are simultaneously grown such that they, too, are symmetric.

In “SymmMap: Estimation of 2-D Reflection Symmetry Map with an Application” [33], Nagar and Raman compute a *symmetry map* representation consisting of two components: the first specifies, for each pixel, the location of its symmetric counterpart, while the second provides a confidence score for the mapping.

In “On Mirror Symmetry via Registration and the Optimal Symmetric Pairwise Assignment of Curves” [7], Cicconet *et al.* address the problem of detecting the plane of reflection symmetry in R^n by registering the points reflected about an arbitrary symmetry plane, and then inferring the optimal symmetry plane from the parameters of the transformation mapping the original to the reflected point sets.

In “Hierarchical Grouping Using Gestalt Assessments” [30], Michaelsen and Arens describe a framework for using various types of symmetry (e.g., reflection, Frieze repetition, and rotational symmetry) to organize nonaccidental arrangements (Gestalten) into hierarchies that take into account image location, scale, and orientation.

4. Conclusions

Our evaluation of 11 different symmetry challenge track submissions against 13 baseline algorithms gives us a glimpse of the state-of-the-art of symmetry detection algorithms in computer vision. It is somewhat surprising that among the six algorithms evaluated, Loy and Eklundh’s algorithm of ECCV 2006 [26] remains competitive in detecting reflection symmetries on 2D images, in particular

on detecting multiple reflection symmetries ($F=0.30$), while for detecting a single reflection symmetry in an image the baseline algorithm of Atadjanov and Lee [17] is the best ($F=0.52$). On frieze pattern detection from images, the F-scores of both baseline [54] and challenger [30] are relatively low ($F=0.19$ - 0.20). Seven algorithms are evaluated on the Sym-COCO dataset for reflection symmetry detection, for which the Funk and Liu [14] CNN baseline algorithm trained with human labels stands out ($F=0.38$ - 0.41). Moving on to medial-axis detection on real images, the good news is that the challengers RSRN [21] ($F=0.64$) and SegSkel [23] ($F=0.73$) beat the baseline algorithms as the winner on the BMAX500 and SKLARGE datasets respectively, yet they still score worse than humans ($F=0.80$). The 3D symmetry detection algorithms [7, 10] evaluated in this symmetry challenge demonstrated high promise on synthetic and real data sets respectively. But larger dataset and more algorithms are needed for a more comprehensive validation and comparison in the future.

There seems to be a general trend of an ascending order in symmetry detection performance of learning-based algorithms, deep-learning methods, and human, which suggests that we have much to learn from human perception. Though progress has been made, detecting *symmetry in the wild* has proven to be a real challenge facing the computer vision community and, more generally, the artificial intelligence community. We anticipate that future advancements on mid-level machine perception will benefit from the outcome (algorithms, labeled datasets) of our ICCV 2017 *Detecting Symmetry in the Wild* Challenge.

5. Acknowledgements

The authors would like to thank the rest of our 2017 ICCV Symmetry Challenge organizing team, all members of our advisory committee, all reviewers, and all contenders. A special thanks to Microsoft for sponsoring our 2017 Symmetry Challenge Cash awards. This work is supported in part by an US NSF CREATIV grant (IIS-1248076), NSERC Canada, and an National Natural Science Foundation of China grant (No. 61672336).

References

- [1] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *TPAMI*, 2011. 4
- [2] I. Atadjanov and S. Lee. Bilateral symmetry detection based on scale invariant structure feature. In *Image Processing (ICIP), 2015 IEEE International Conference on*, pages 3447–3451. IEEE, 2015. 1

- [3] I. Atadjanov and S. Lee. Reflection symmetry detection via appearance of structure descriptor. In *European Conference on Computer Vision (ECCV)*, Amsterdam, 2016. 2, 3
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv:1606.00915*, 2016. 1
- [5] S. Choi, Q.-Y. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. <http://redwood-data.org/indoor/index.html>. 6
- [6] S. Choi, Q.-Y. Zhou, S. Miller, and V. Koltun. A Large Dataset of Object Scans. *ArXiv e-prints*, Feb. 2016. <http://redwood-data.org/3dscan/>. 6
- [7] M. Cicconet, D. G. C. Hildebrand, and H. Elliott. Finding mirror symmetry via registration and optimal symmetric pairwise assignment of curves. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 3, 7, 8
- [8] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Niessner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [9] J. D. Delius and G. Habers. Symmetry: can pigeons conceptualize it? *Behavioral biology*, 22(3):336–342, 1978. 1
- [10] A. Ecin, C. Fermuller, and Y. Aloimonos. Detecting reflectional symmetries in 3d data through symmetrical fitting. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 7, 8
- [11] M. Elawady, C. Ducottet, O. Alata, C. Barat, and P. Colantoni. Wavelet-based reflection symmetry detection via textural and color histograms. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 3, 7
- [12] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [13] C. Funk and Y. Liu. Symmetry reCAPTCHA. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 1, 4
- [14] C. Funk and Y. Liu. Beyond planar symmetry: Modeling human perception of reflection and rotation symmetries in the wild. *ICCV*, 2017. 1, 2, 4, 5, 8
- [15] M. Giurfa, B. Eichmann, and R. Menzel. Symmetry perception in an insect. *Nature*, 382(6590):458–461, Aug 1996. 1
- [16] F. Guerrini, A. Gnutti, and R. Leonardi. Innerspec: Technical report. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 3
- [17] T. Lee, S. Fidler, and S. Dickinson. Detecting curved symmetric parts using a deformable disc model. In *Proceedings, IEEE International Conference on Computer Vision*, Sydney, 2013. 1, 8
- [18] A. Levinshstein, C. Sminchisescu, and S. Dickinson. Multi-scale symmetric part detection and grouping. *International Journal of Computer Vision*, 104:117–134, 2013. 1
- [19] G. Li, Y. Xie, L. Lin, and Y. Yu. Instance-level salient object segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 4
- [21] C. Liu, W. Ke, J. Jiao, and Y. Qixiang. Rsrn: Rich side-output residual network for medial axis detection. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 6, 8
- [22] J. Liu, G. Slota, G. Zheng, Z. Wu, M. Park, S. Lee, I. Rauschert, and Y. Liu. Symmetry detection from real world images competition 2013: Summary and results. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 200–205. IEEE, 2013. 1, 2, 3
- [23] X. Liu and P. Lyu. Fusing image and segmentation cues for skeleton extraction in the wild. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 6, 8
- [24] Y. Liu. Computational Symmetry. In I. Hargittai and T. Laurent, editors, *Symmetry 2000*, volume 80, chapter 21, pages 231–245. Wenner-Gren International Series, Portland, London, ISBN I 85578 149 2, 2002. 1
- [25] Y. Liu, H. Hel-Or, C. S. Kaplan, and L. V. Gool. Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision*, 5(12):1–195, 2010. 1, 2
- [26] 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. 2, 3, 4, 5, 8
- [27] M. Lukáč, D. Šykora, K. Sunkavalli, E. Shechtman, O. Jamriška, N. Carr, and T. Pajdla. Nautilus: recovering regional symmetry transformations for image editing. *ACM Transactions on Graphics (TOG)*, 36(4):108, 2017. 1
- [28] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. *ICCV*, 2001. 4
- [29] D. R. Martin, C. C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE transactions on pattern analysis and machine intelligence*, 26(5):530–549, 2004. 2, 5
- [30] E. Michaelsen and M. Arens. Hierarchical grouping using gestalt assessments. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 3, 4, 5, 8
- [31] 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. 6
- [32] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3d geometry: Extraction and applications. *Comput. Graph. Forum*, 32(6):1–23, 2013. 1, 6
- [33] R. Nagar and S. Raman. SymmMap: Estimation of 2-d reflection symmetry map and its applications. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 8
- [34] R. Nagar and S. Raman. SymmSLIC: Symmetry aware superpixel segmentation. In *Proceedings, ICCV Workshop on Detecting Symmetry in the Wild*, Venice, 2017. 7

- [35] M. Park, Y. Liu, and R. Collins. Deformed lattice detection via mean-shift belief propagation. In *Proceedings of the 10th European Conference on Computer Vision (ECCV'08)*, 2008. 3
- [36] 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. 6
- [37] J. Pritts, O. Chum, and J. Matas. Detection, rectification and segmentation of coplanar repeated patterns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2973–2980, 2014. 1
- [38] I. Rodríguez, A. Gumbert, N. Hempel de Ibarra, J. Kunze, and M. Giurfa. Symmetry is in the eye of the ‘beeholder’: innate preference for bilateral symmetry in flower-naïve bumblebees. *Naturwissenschaften*, 91(8):374–377, Aug 2004. 1
- [39] T. Sawada and Z. Pizlo. Detecting 3-d mirror symmetry in a 2-d camera image for 3-d shape recovery. *Proceedings of the IEEE 102*, pages 1588–1606, 2014. 2
- [40] W. Shen, X. Bai, Z. Hu, and Z. Zhang. Multiple instance subspace learning via partial random projection tree for local reflection symmetry in natural images. *Pattern Recognition*, 2016. 4
- [41] W. Shen, K. Zhao, Y. Jiang, Y. Wang, X. Bai, and A. L. Yuille. Deepskeleton: Learning multi-task scale-associated deep side outputs for object skeleton extraction in natural images. *IEEE Transactions on Image Processing*, 26(11):5298–5311, 2017. 4, 5, 6
- [42] W. Shen, K. Zhao, Y. Jiang, Y. Wang, Z. Zhang, and X. Bai. Object skeleton extraction in natural images by fusing scale-associated deep side outputs. In *CVPR*, 2016. 1, 4, 5
- [43] P. Speciale, M. R. Oswald, A. Cohen, and M. Pollefeys. A symmetry prior for convex variational 3d reconstruction. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 6
- [44] A. Telea and J. Van Wijk. An augmented fast marching method for computing skeletons and centerlines. *Eurographics*, 2002. 4
- [45] C. L. Teo, C. Fermüller, and Y. Aloimonos. Detection and Segmentation of 2D Curved Reflection Symmetric Structures. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1644–1652, 2015. 1, 2, 4, 5
- [46] S. Tsogkas and S. Dickinson. Amat: Medial axis transform for natural images. In *International Conference on Computer Vision*, 2017. 1, 4, 5, 6
- [47] S. Tsogkas and I. Kokkinos. Learning-based symmetry detection in natural images. In *ECCV*, 2012. 2, 4, 5, 6
- [48] L. von Fersen, C. S. Manos, B. Goldowsky, and H. Roitblat. Dolphin detection and conceptualization of symmetry. In *Marine mammal sensory systems*, pages 753–762. Springer, 1992. 1
- [49] Z. Wang, Z. Tang, and X. Zhang. Reflection Symmetry Detection Using Locally Affine Invariant Edge Correspondence. *IEEE Transactions on Image Processing*, 24(4):1297–1301, 2015. 1
- [50] Website. Archive3d - 3d model repository. <https://archive3d.net/>. 6
- [51] Website. Detecting symmetry in the wild challenge 2011. <http://vision.cse.psu.edu/research/symmComp/index.shtml>. 1
- [52] Website. Detecting symmetry in the wild challenge 2013. <http://vision.cse.psu.edu/research/symmComp13/index.shtml>. 1
- [53] Website. Detecting symmetry in the wild challenge 2017. <https://sites.google.com/view/symmcomp17/>. 6, 7
- [54] C. Wu, J.-M. Frahm, and M. Pollefeys. Detecting large repetitive structures with salient boundaries. *Computer Vision—ECCV 2010*, pages 142–155, 2010. 2, 4, 8