Reflection and Rotation Symmetry Detection via Equivariant Learning
— Supplemental Material —

Ahyun Seo Byungjin Kim Suha Kwak Minsu Cho
Pohang University of Science and Technology (POSTECH), South Korea
http://cvlab.postech.ac.kr/research/EquiSym

1. DENDI

The details of data annotation for the DENDI dataset are described in this section. To identify symmetry, we disregard texture and focus just on the shape of the object. The partial occlusion of boundary of symmetric object is allowed to a fourth of the object boundaries. For both reflection and rotation symmetries, we exhaustively mark all symmetries in an image, including those of parts. This policy ensures that the DENDI contains dense annotations. we present some examples in Fig.1(d) and Fig.2(d).

1.1. Reflection symmetry

A reflection symmetry axis is drawn as a line following the previous datasets [2,4,5,7,9]. The notable examples are shown in Fig.1(a). We annotate all reflection symmetry axes in an object, including non-dominant ones. Different from the existing datasets, we account for a circular object, which has an infinite number of reflection symmetry axes. Instead of an infinite number of symmetry axes to represent a circular object, we use a '4'-shaped label consisting of 5 points which are then converted to a circular mask, as shown in Fig.1(b). Note that a semantically circular object that seems to be an ellipse due to viewpoint variants is also annotated with a '4'-shaped label, as show in the first tow rows in Fig.1(c). Likewise, a skewed regular polygon due to perspective variations has the same reflection axis as a non-skewed regular polygon, as shown in the last two rows in Fig.1(c), *e.g.*, a regular STOP sign and a skewed STOP sign that are both semantically regular octagon have eight reflection symmetry axes.

Furthermore, we annotate symmetry in characters such as A, B, C, D, E, H, I, K, M, O, T, U, V, W, X, and Y, as well as the numbers 0, 1, 2, 5, and 8, except for those that are too thin or indistinct. We also annotate symmetry in the D-shaped part of characters, such as P and R. If multiple symmetry axes are overlapped, only the longest one is saved.

1.2. Rotation symmetry

For each object with rotation symmetry, we collect the coordinate of the rotation center, the boundary of the object, and the number of the rotation folds (N). We again employ the '4'-shaped labels to denote circular or elliptical objects as shown in Fig.2(a). The semantically circular object also features an infinite number of rotation folds, indicated as 0 for simplicity, in addition to the '4'-shaped labels. Similar to reflection symmetry, an semantically circular object with elliptical shape due to viewpoint variants has a rotation fold of 2, *e.g.*, the third and fourth rows in Fig.2(a). The minor axis takes precedence over the major axis when drawing '4'-shaped labels for elliptical objects. The rotation fold of a circular object, in particular, can be greater than 2 if the object contains cyclic symmetry. In the case of a non-circular object with rotation symmetry such as Fig.2 (c), we draw a convex polygon starting from the center of the object and following convex vertices. The vertex nearest to 12 o'clock takes priority among the convex vertices. Likewise, in the reflection symmetry dataset, symmetry in characters such as H, I, N, O, S, X, and Z, as well as the numbers 0 and 8, are taken into account.

2. EquiSym

The details that are omitted in the main paper are covered in this section. We show the consistency of the evaluation schemes. The implementation details are also shown in the following.

2.1. Evaluation scheme

In the main paper, we propose to use a modified evaluation scheme of blurring the ground truth rather than thinning the prediction. The primary reason is that the thinning process transforms a circular mask prediction into a single dot. The pixel matching algorithm determines whether the predicted lines are close enough to the ground truth lines within a threshold, which becomes equivalent when the ground truth itself is blurred with a radius of a threshold. In prac-

(a) multiple symmetry axes (b) circular objects (c) skewed objects (d) dense symmetry axes

Figure 1. Illustration of the examples in the reflection symmetry dataset. The samples with (a) multiple symmetry axes, (b) circular objects, (c) skewed objects, and (d) dense symmetry axes are shown in the figure. Green lines indicate the reflection axes and the yellow lines indicate the '4'-shaped reflection circle annotation. The reflection-circle annotations are then converted to masks.

method	train dataset		test dataset		mixed
	real	synth	SDRW	LDRS	
PMCNet [9]	✓		N/A	N/A	46.6
	✓	✓	**51.1**	33.7	
EquiSym-ref	✓		49.8	**36.5**	**58.0**
	✓	✓	49.2	34.9	

Table 1. Comparison of the reflection symmetry detection methods on the LDRS [9] and SDRW [7]. Note that the real dataset consists of SDRW, LDRS, and NYU [2] dataset and synthetic images generated as in [9].

tice, we construct a filter of the kernel 11×11 where the weights are set to 1 for a circle of a diameter 11 and 0 otherwise. We convolve the ground-truth heatmap of both symmetries with the filter so the heatmap is dilated to the maximum distance of 5 pixels. The true positives are then computed by pixel-wise comparisons. We re-evaluate the experiments of Tab. 3 of the main paper in Tab.1. Note that one experiment from PMCNet [9] is excluded since the trained model is not accessible. As shown in Tab.1, the rankings produced by the two evaluation schemes are consistent, while the latter is significantly faster.

2.2. Implementation details

Construction of the orientation labels. EquiSym utilizes the intermediate tasks to increase the accuracy of the symmetry detection tasks. In the case of reflection symmetry, the intermediate labels of the orientations of the reflection axes are obtained for free. The angle of the reflec-

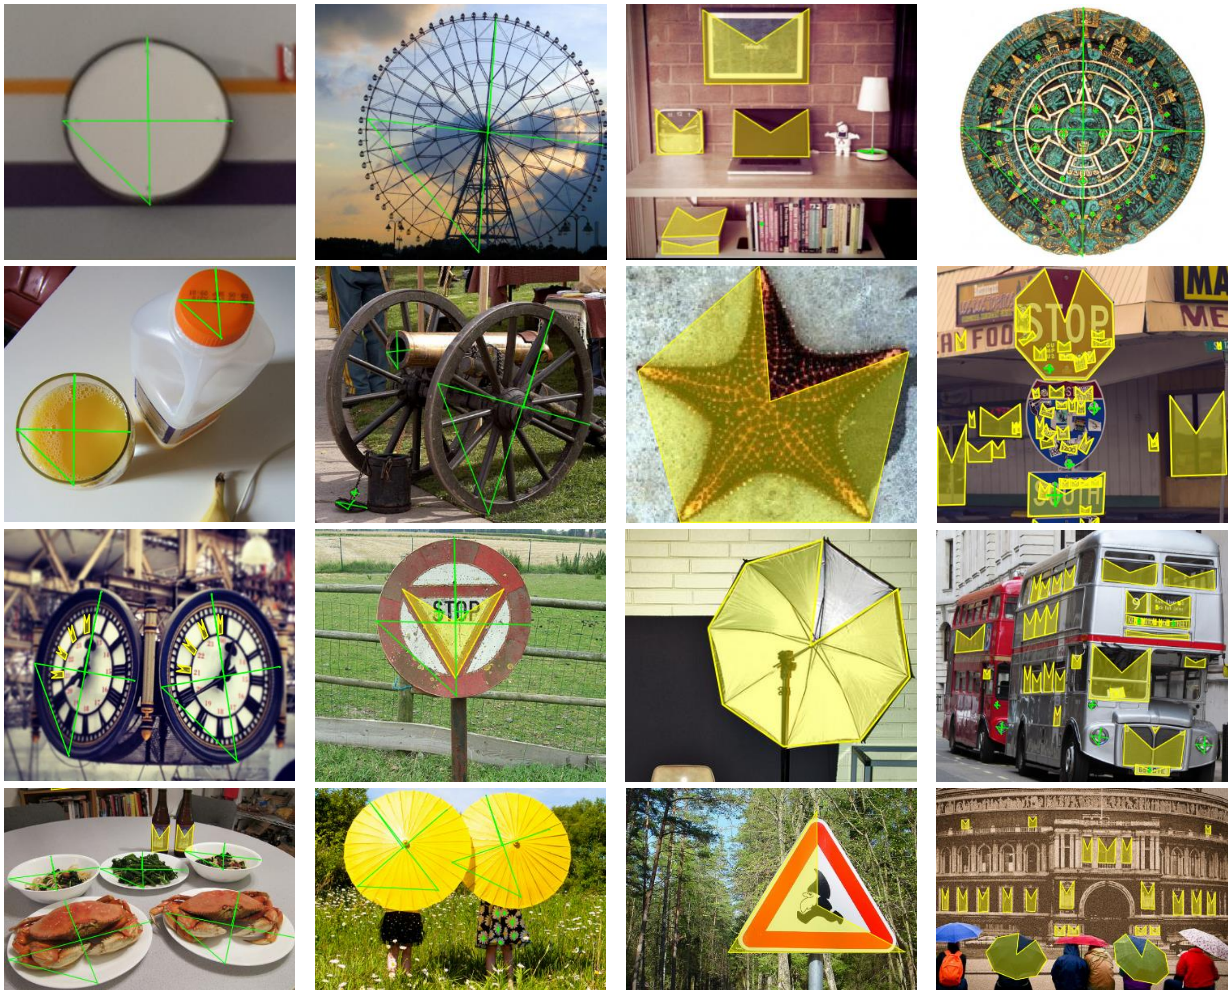

(a) circular objects | (b) circular object with non-zero fold | (c)Polygons | (d) dense symmetries

Figure 2. Illustration of the examples in the rotation symmetry dataset. The samples with (a) circular objects, (b) circular objects with folds larger than 2, (c) polygons, and (d) dense symmetries are shown in the figure. Green lines indicate the circular annotations and the yellow polygons indicate the polygon-type annotations. Only the center coordinates are used for evaluation.

tion symmetry axis in the form of a straight line can be expressed as a linear combination of the closest one or two angles among the 8 predetermined angles, which is an initial soft label. On the other hand, the circle-shaped symmetry axis has an evenly divided orientation label of 8 segments determined by the orientation of the line crossing the center. The orientation label is then quantized for training.

ImageNet pretrain. To be consistent with experiments on the vanilla ResNet [6] pre-trained on ImageNet [3], we pretrain the ReResNet50 on ImageNet-1K for the image classification. While ReResNet50 from ReDet is implemented with C_8 group, we use D_8 group instead. Furthermore, we adjust the stride and dilation of each layer in ReResNet50 to obtain a higher resolution feature map with a larger receptive field than the original one, which is a common procedure in the semantic segmentation [1]. The learning rate starts at 0.1 and decreases by 0.1 every 30 epochs, for a total of 100 epochs. We use a batch size of 512. The pretrained ReResNet50 achieves 69.06% top-1 and 87.25% top-5 accuracy on the ImageNet *val*.

Implementation details. Following [9], The hyperparameters α and β of the focal loss are set to 0.95 and 2, respectively. For training, we resize input images so that the maximum length of the width or the height is 417. The background weight w of $\mathcal{L}_{\mathrm{cls}}$ is set to 0.01 and 0.001 for reflection and rotation symmetries. We use the PyTorch [8] and e2cnn [10] framework to build our model.

References

[1] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *CoRR*, abs/1706.05587, 2017. 3

[2] M. Cicconet, V. Birodkar, M. Lund, M. Werman, and D. Geiger. A convolutional approach to reflection symmetry. http://arxiv.org/abs/1609.05257, 2016. New York. 1, 2

[3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *cvpr*, pages 248–255. Ieee, 2009. 3

[4] Christopher Funk, Seungkyu Lee, Martin R Oswald, Stavros Tsogkas, Wei Shen, Andrea Cohen, Sven Dickinson, and Yanxi Liu. 2017 iccv challenge: Detecting symmetry in the wild. In *Int. Conf. Comput. Vis. Worksh.*, pages 1692–1701, 2017. 1

[5] Christopher Funk and Yanxi Liu. Beyond planar symmetry: Modeling human perception of reflection and rotation symmetries in the wild. In *Int. Conf. Comput. Vis.*, pages 793–803, 2017. 1

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *cvpr*, pages 770–778, 2016. 3

[7] Jingchen Liu, George Slota, Gang Zheng, Zhaohui Wu, Minwoo Park, Seungkyu Lee, Ingmar Rauschert, and Yanxi Liu. Symmetry detection from realworld images competition 2013: Summary and results. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 200–205, 2013. 1, 2

[8] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Adv. Neural Inform. Process. Syst.*, pages 8024–8035. Curran Associates, Inc., 2019. 3

[9] Ahyun Seo, Woohyeon Shim, and Minsu Cho. Learning to discover reflection symmetry via polar matching convolution. In *Int. Conf. Comput. Vis.*, 2021. 1, 2, 3

[10] Maurice Weiler and Gabriele Cesa. General E(2)-Equivariant Steerable CNNs. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 3