

NeRD++: Improved 3D-mirror symmetry learning from a single image

Supplementary material

Yancong Lin
Silvia L. Pintea
Jan C. van Gemert

Computer Vision Lab,
Delft University of Technology,
Delft, Netherlands

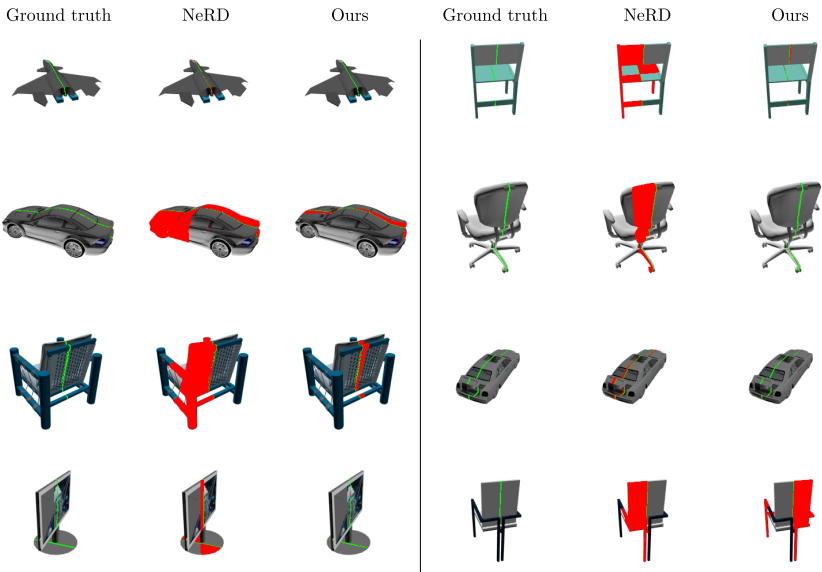


Figure 1: **Qualitative results on ShapeNet when training on the 1% subset.** We compare the ground truth, predictions from NeRD and our results when training on the 1% subset only. We highlight the ground truth symmetry axis in green and prediction errors in red. In general, our model shows superior results over NeRD in small data regime. However, both models fail in the bottom-left example due to the lack of correspondences.

We compare the predictions from NeRD [9] and our model qualitatively on the ShapeNet dataset [10]. We rely on ground truth depth for visualizing the mirror symmetry on the image plane. Fig. 1 compares NeRD (middle) and our model (right), when training on the 1% subset only. The green line shows the ground truth symmetry axis on the image plane. We highlight in red the errors between the prediction and the ground truth. The comparison verifies the advantage of our model in small data regime where only 1% data is available.

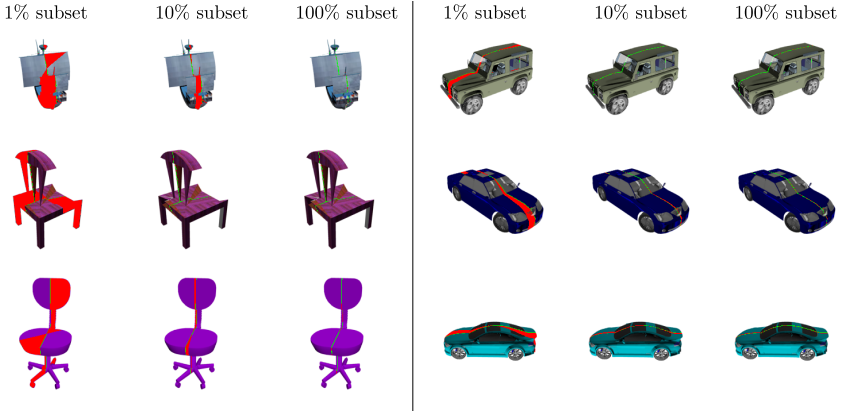


Figure 2: **Qualitative results on ShapeNet when adding more training data.** We compare the predictions from our models trained on the 1%, 10% and 100% (sub)sets. The green line represents the ground truth symmetry axis while the red region indicates errors. Our model can effectively take advantage of additional data and thus recovers from making substantially wrong predictions.

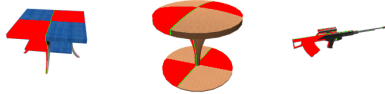


Figure 3: **Failure cases on ShapeNet.** Our model fails in certain scenarios due to the presence of multiple symmetries and lack of reliable correspondences. We highlight the ground truth in green and prediction errors in red.

Fig. 2 shows the impact of adding more training data from 1%, 10% to 100%. In the leftmost examples, our model fails on the 1% subset, as it predicts a plane orthogonal to the ground truth. However, our model can make effective use of additional training data and recover from substantial mistakes. The examples on the right show that also for our model adding data improves precision qualitatively.

We show the failure cases of our model in Fig. 3. The main challenge in these examples is that an object may admit multiple symmetries, thus leading to ambiguity. Moreover, it can be hard to find symmetric correspondences for certain objects, such as the gun example.

We also display results from the real-world Pix3D dataset [2] in Fig. 4. In general, our model is able to detect the dominant mirror symmetry accurately from images unless multiple symmetries are present.



Figure 4: **Qualitative results on Pix3D.** Our model is able to detect the dominant mirror symmetry from real-world images in most cases. However it fails to handle multiple symmetries (shown in the rightmost column). The ground truth symmetry axis is in green, while the prediction errors are in red.

References

- [1] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [2] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3d: Dataset and methods for single-image 3d shape modeling. In *CVPR*, 2018.
- [3] Yichao Zhou, Shichen Liu, and Yi Ma. NeRD: Neural 3d reflection symmetry detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15940–15949, 2021.