

A. Supplementary Materials

A.1. Derivation of $C(\mathbf{w})$

Let $\tilde{\mathbf{x}}_0 \in \mathbb{R}^3$ be a point in the camera space and $\tilde{\mathbf{x}}_1$ be its mirror point with respect to the symmetry plane

$$\mathbf{w}^T \tilde{\mathbf{x}} + 1 = 0. \quad (9)$$

Figure 10 illustrates the process of miring a point in 2D, where the red dots are the pair of points $\tilde{\mathbf{x}}_0$ and $\tilde{\mathbf{x}}_1$ the line in the middle is the symmetry plane whose normal $\tilde{\mathbf{n}} = \frac{\mathbf{w}}{\|\mathbf{w}\|_2}$. The distance between $\tilde{\mathbf{x}}_0$ and the symmetry plane is $\frac{\mathbf{w}^T \tilde{\mathbf{x}}_0 + 1}{\|\mathbf{w}\|_2}$, according to the formula of distance from a point to a plane. Therefore, we have

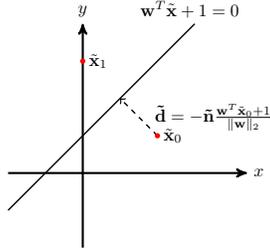


Figure 10: Illustration of reflection symmetry with two points.

$$\tilde{\mathbf{x}}_1 = \tilde{\mathbf{x}}_0 - 2 \frac{\mathbf{w}^T \tilde{\mathbf{x}}_0 + 1}{\|\mathbf{w}\|_2^2} \mathbf{w}. \quad (10)$$

We could also write this in matrix form:

$$\begin{bmatrix} \tilde{\mathbf{x}}_1 \\ 1 \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{x}}_0 \\ 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} - \frac{2\mathbf{w}\mathbf{w}^T}{\|\mathbf{w}\|_2^2} & -\frac{2\mathbf{w}}{\|\mathbf{w}\|_2^2} \\ \mathbf{0} & 1 \end{bmatrix}. \quad (11)$$

Because the transformation between the camera space and the pixel space is given by

$$\mathbf{x} = \mathbf{K} \begin{bmatrix} \tilde{\mathbf{x}} \\ 1 \end{bmatrix}, \quad (12)$$

we finally have

$$\begin{aligned} \mathbf{C}(\mathbf{w}) &= \mathbf{K} \begin{bmatrix} \mathbf{I} - \frac{2\mathbf{w}\mathbf{w}^T}{\|\mathbf{w}\|_2^2} & -\frac{2\mathbf{w}}{\|\mathbf{w}\|_2^2} \\ \mathbf{0} & 1 \end{bmatrix} \mathbf{K}^{-1} \\ &= \mathbf{K} \left(\mathbf{I} - \frac{2}{\|\mathbf{w}\|_2^2} \begin{bmatrix} \mathbf{w} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{w}^T & 1 \end{bmatrix} \right) \mathbf{K}^{-1}. \end{aligned}$$

A.2. Network Architecture

We display the NeRD’s network architecture in Figure 11.

A.3. Illustration of Coarse-to-Fine Inference

Figure 12 shows the process of coarse-to-fine inference on sampled images from ShapeNet. We display the symmetry plane with the highest confidence score in each round of coarse-to-fine inference. In the i th round, we determine the normal of symmetry plane to the accuracy of Δ_i , where $\Delta = [20.7^\circ, 6.44^\circ, 1.99^\circ, 0.61^\circ]$ are set according to the gap between near directions from the number of direction samples per round $K = 32$. The coarse-to-fine inference dramatically reduces the number of samples required to achieve

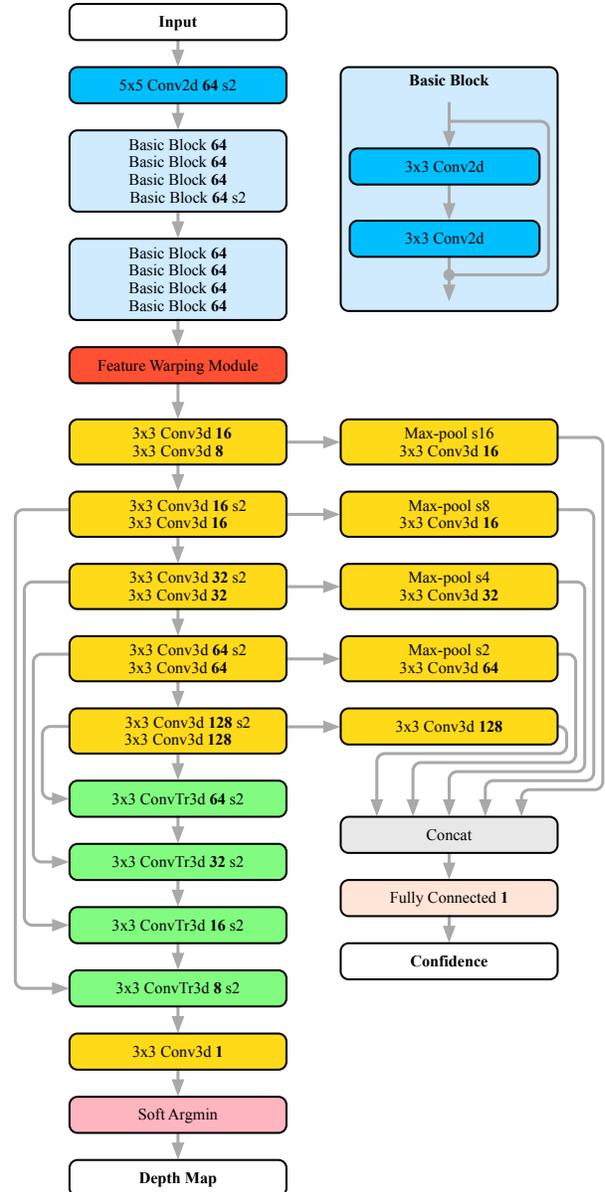


Figure 11: Illustration of NeRD’s network architecture. We show the resulting number of channels after each operator in **bold**. “s2” represents stride-2 operators.

a certain level of accuracy. As shown in the figure, the precision of the predicted plane increases with the number of rounds in the coarse-to-fine inference.

A.4. Failure Cases

Figure 13 shows sampled failure cases on ShapeNet. We categorize those cases into three classes: lack of correspondence, the existence of multiple symmetry planes, and asymmetric objects. For the first category, e.g., the firearm shown in Figure 13(a), it is hard to accurately find the symmetry plane from the geometry cues because for most pixels,

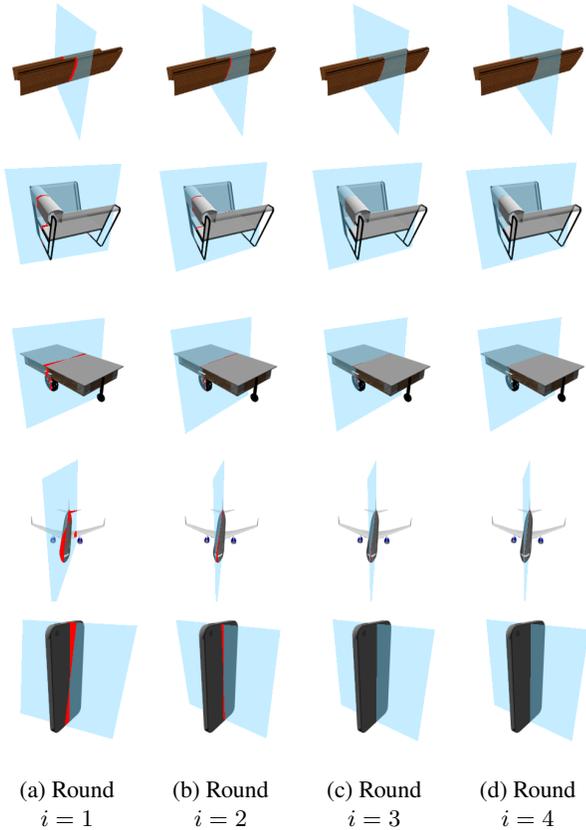


Figure 12: Illustration of the coarse-to-fine inference on sampled images from ShapeNet. The symmetry plane with the highest confidence score in each round of coarse-to-fine inference is drawn. In the i th round, we determine the normal of symmetry plane to the accuracy of Δ_i , where $\Delta = [20.7^\circ, 6.44^\circ, 1.99^\circ, 0.61^\circ]$ are set according to the gap between nearby directions from the number of direction samples per round $K = 32$.

their corresponding points are occluded and invisible in the picture. For the second category, objects in shapes such as squares and cylinders admit multiple reflection symmetry, and NeRD may return the reflection plane that differs from the symmetry plane of the ground truth. For the third category, some objects in ShapeNet are not symmetric. Thus, the detected symmetry plane might be different from the “ground truth symmetry plane” computed by applying \mathbf{R}_t to the Y-Z plane in the world space.

A.5. More Visualization

Figure 14 and Figure 15 show the visual quality of the detected symmetry planes of NeRD on **random sampled** testing images from ShapeNet and Pix3D, respectively. For the results of Pix3D, we note that artifacts in visualization such as acentric symmetry planes are due to misalignment of CAD models and real-world images. We find that for most

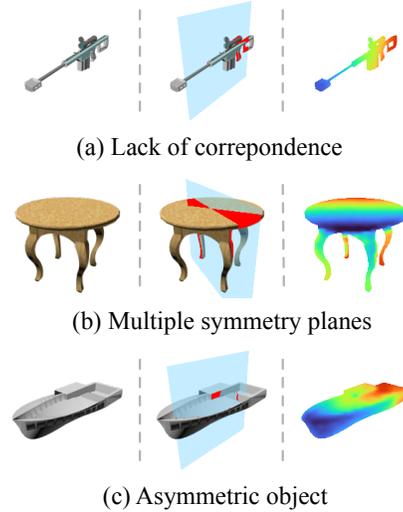


Figure 13: Sampled failure cases of NeRD on ShapeNet.

images, NeRD is able to determine the normal of symmetry plane accurately by utilization of the geometric constraints from symmetry. The errors (red pixels) are sub-pixel in most cases.

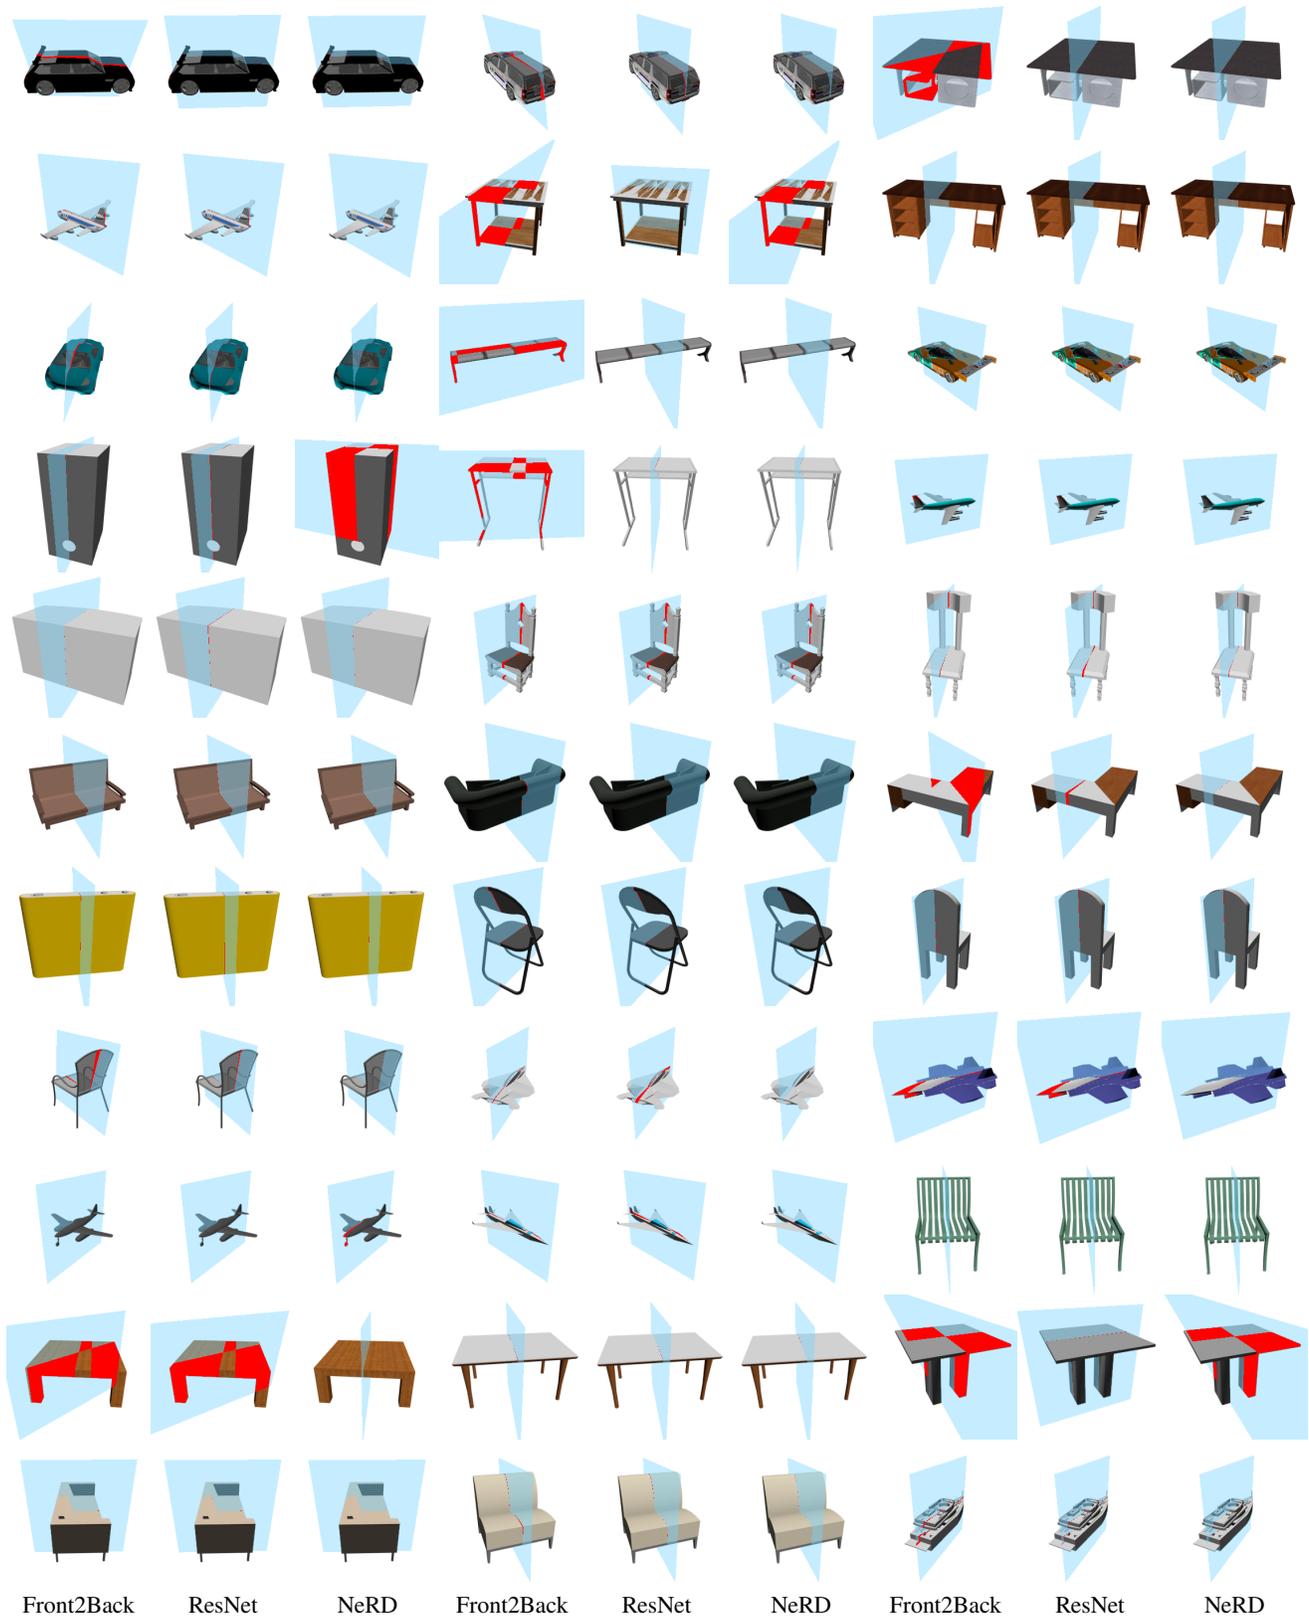


Figure 14: Detected symmetry planes of NeRD on *random sampled images* from ShapeNet. Errors of symmetry planes, i.e., pixels between the predicted plane and the ground truth plane, are **highlighted**.



Figure 15: Detected symmetry planes of NeRD on *random sampled images* from Pix3D. Errors of symmetry planes, i.e., pixels between the predicted plane and the ground truth plane, are **highlighted**. Artifacts in visualization are due to the misalignment of between the models and images in Pix3D.