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

TUDelft
Delft University of Technology

Yancong Lin, Silvia Laura Pintea, and Jan C. van Gemert
TUDelft, The Netherlands
https://github.com/yanconglin/3D-mirror-symmetry-plane

BMVC 2022

Task: detecting mirror symmetry planes from single-view images
- Input: a single RGBA image
- Output: perspective mirror symmetry plane (wx+1=0)

(I) Contributions:

- Improved data-efficiency over sota NeRD [1]
  Solution space is a hemisphere - spherical convs for localization
- Improved compute efficiency (20x faster):
  Feature correlations indicate similarity between pixels and mirrors.

Mirror symmetries from single-view images:

(II) Background knowledge

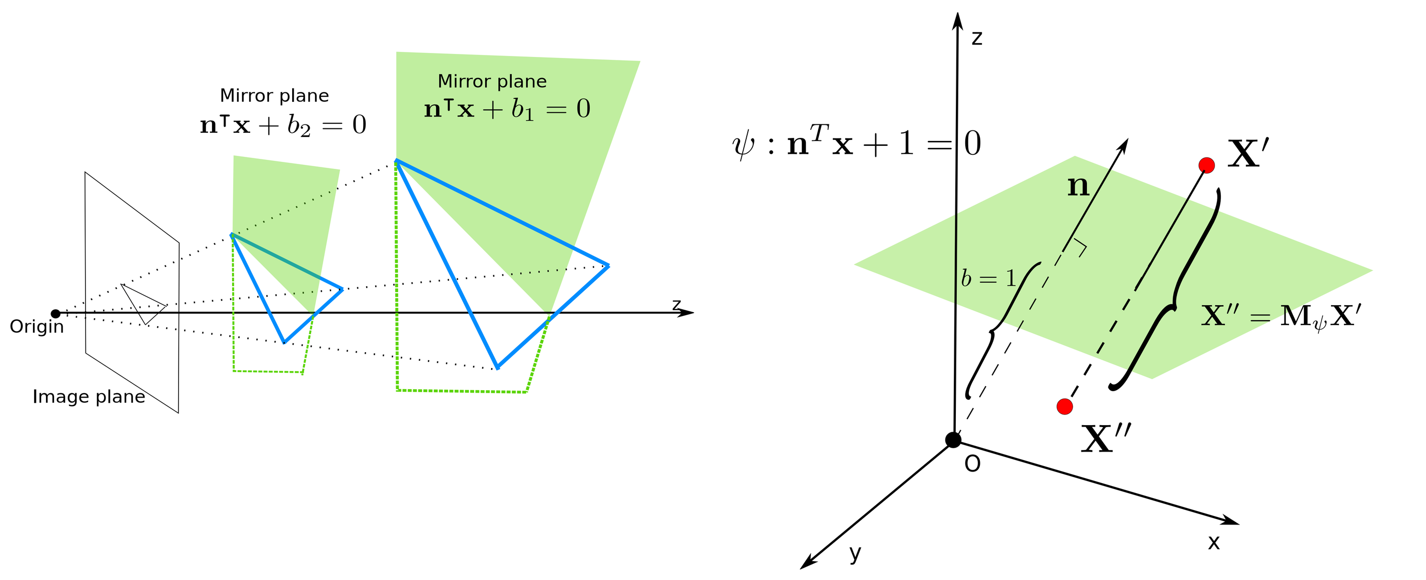

1. The bias (scale) can not be determined from an image.
2. Mirrors can be mathamatically calculated [2].

(III) NeRD Recap

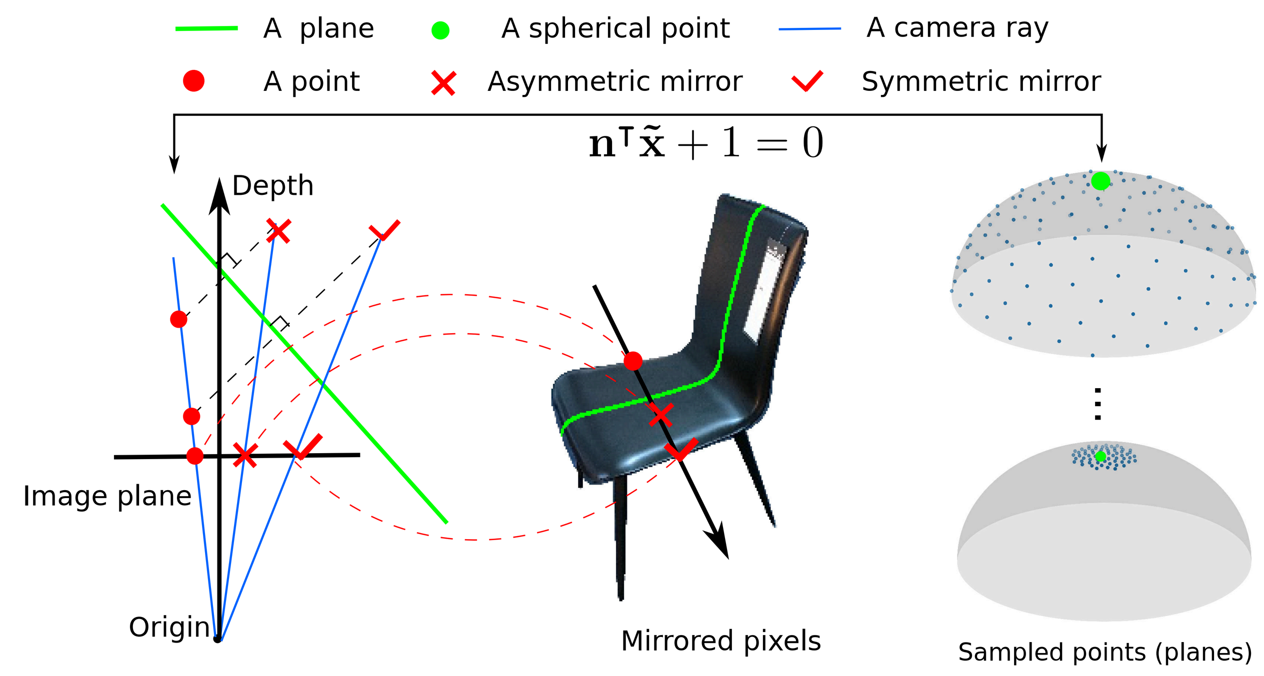

Given a candidate plane
- For each point (pixel):
  - Calculate its mirrors over depth;
  - Concatenate the feats;

Contruct an expensive 4D feature volumes [2C, D, H, W];
Apply learning and Classification (how likely the given plane is a symmetry plane);
Re-sample nearby points from the hemisphere

Iterate over 3 times (coarse-to-fine) to reach a desired resolution.

(III) NeRD++

However,
a. computationally expensive: 4D feature volumes;
b. data-hungry: massive training data/inferior results on subsets.

We propose:

a. calculate feature correlations explicitly,
   Correlation indicates similarity ([D, H, W]);

b. use another inductive prior:
   Solution space is a hemisphere, sperical convs.

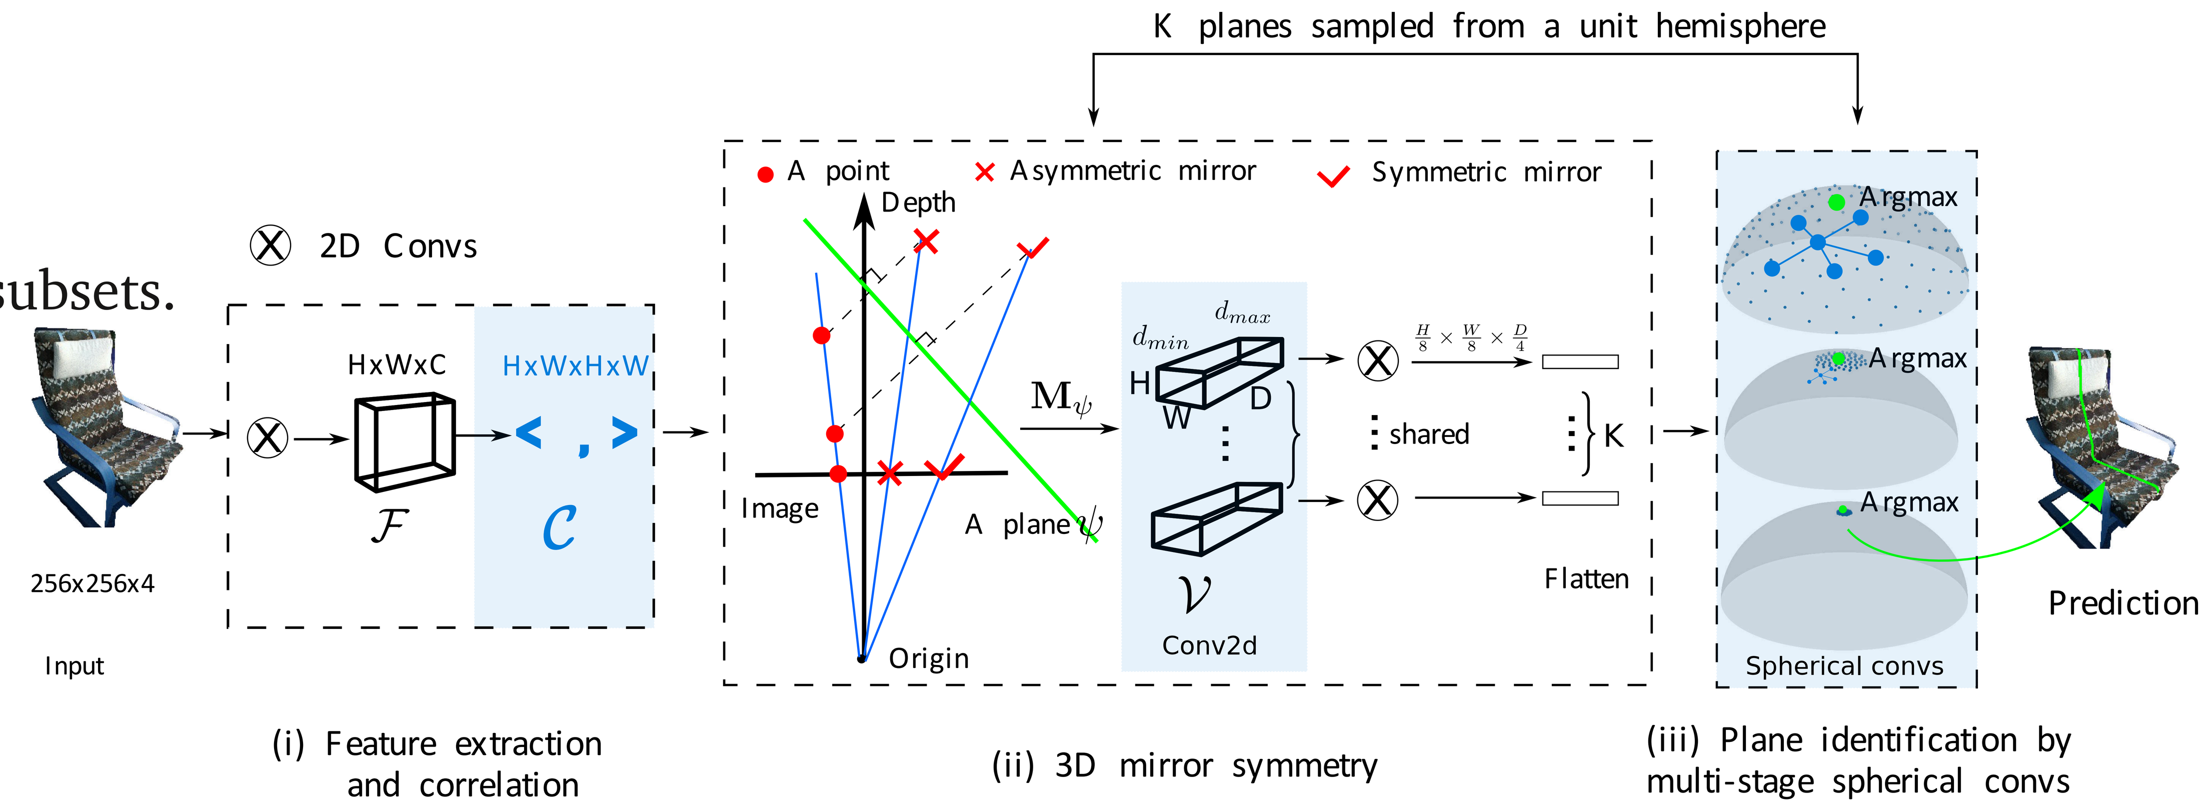

(IV) Experiments

Two datassets: ShapeNet[3], Pix3D[4]
Evaluation: angular errors in the camera space (Angular Accuracy).

- 1. Data-efficiency on ShapeNet

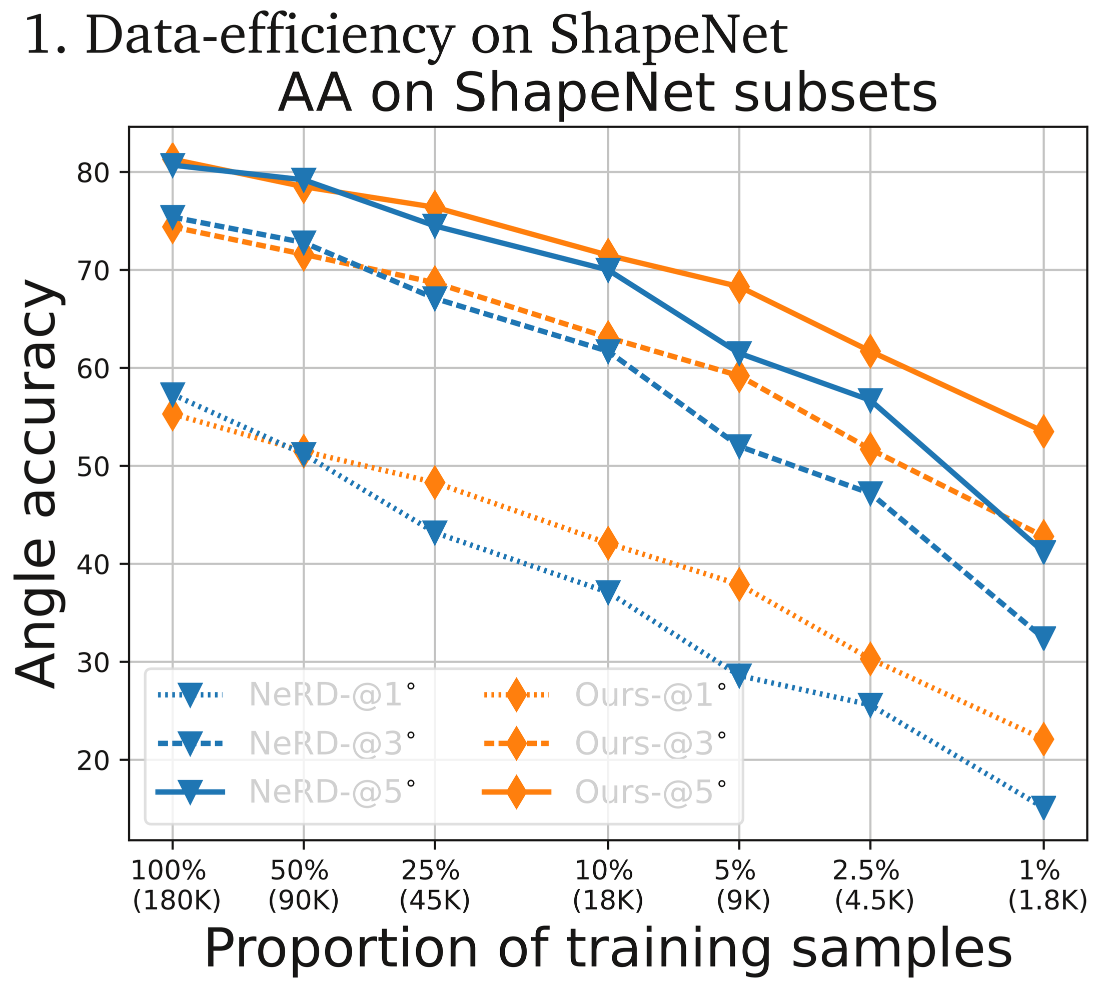

- 2. Pix3D (~5K images)

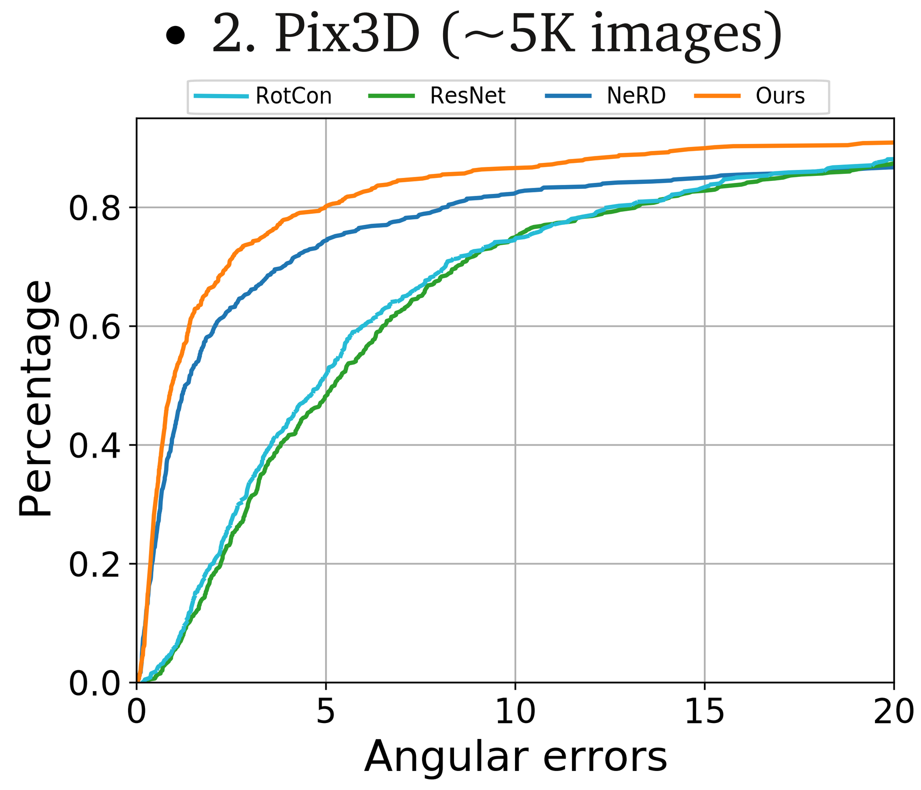

- 3. Faliure cases

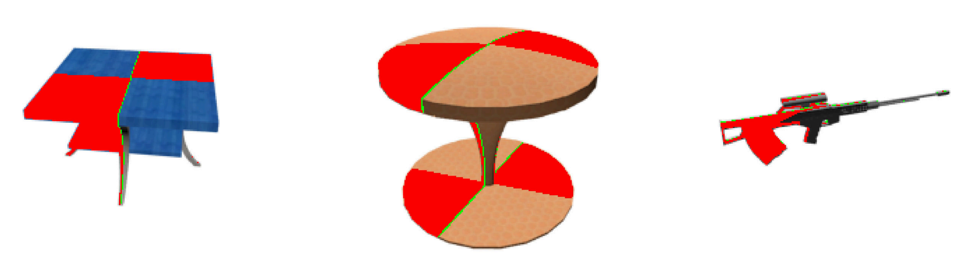

(V) Limitations

- Rely on entirely on appearance features (textureless surface/repetitive patterns);
- Iterate over all possible depth;
- Detect only the dominant symmetry plane.

(V) Future work

- How to hanle occlusion explicitly?
- Downstream tasks using symmetries (depth, shape completion, 3D reconstruction).
- Detect intrinsic symmetries (of non-rigid objects).

Reference

[1] Y Zhou et al, NeRD: Neural 3D Reflection Symmetry Detector, CVPR 2021
[2] D Cailliere et al, 3d mirror symmetry detection using hough transform, ICIP 2008.
[3] A Chang et al, Shapenet: An information-rich 3d model repository, CVPR 2015.
[4] X Sun et al, Pix3d: Dataset and methods for single-image 3d shape modeling, CVPR 2018.
[5] Y Zhou et al, On the continuity of rotation representations in neural network, CVPR 2019