SIGGRAPH ASIA 2020 VIRTUAL

SymmetryNet: Learning to Predict Reflectional and Rotational Symmetries of 3D Shapes from Single-View RGB-D Images

Yifei Shi, Junwen Huang, Hongjia Zhang, Xin Xu,
Szymon Rusinkiewicz and Kai Xu

National University of Defense Technology

Princeton University

Motivation

- Symmetry is omnipresent in both nature and the synthetic world






Motivation

- Purely geometric symmetry detection approach

➢ Based on symmetry correspondences (counterparts) detection

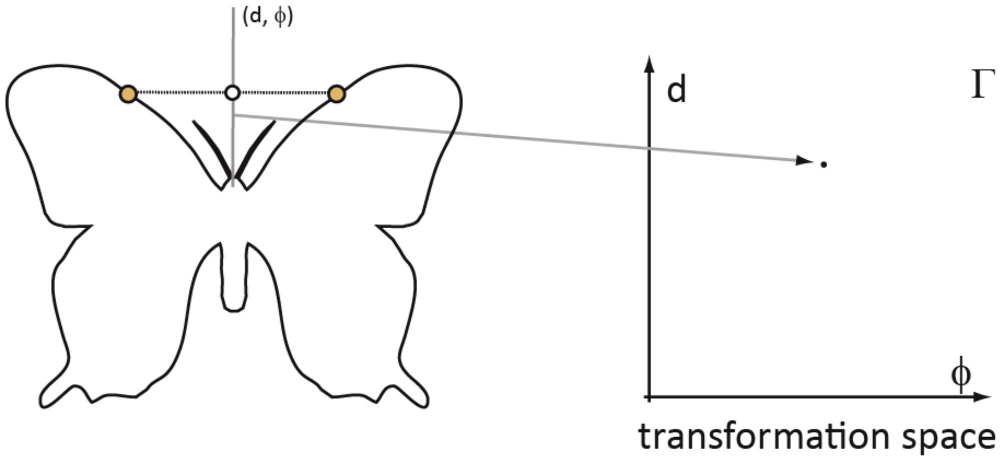

[Mitra et al. 2006]

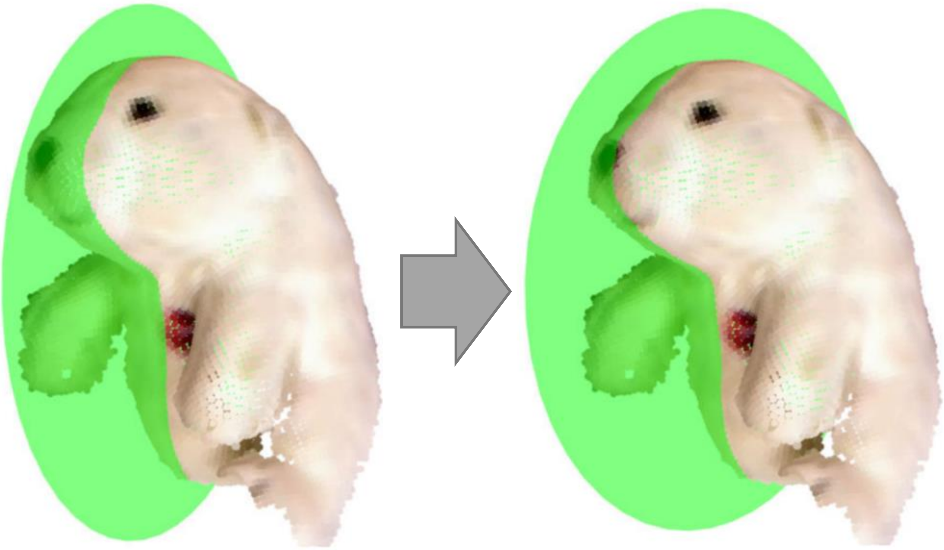

[Ecins et al. 2018]

Motivation

- Detect object symmetries from a single-view RGB-D image
  - ➢ Partial observation (self-occlusion)
  - ➢ Mutual occlusion
  - ➢ Noise
- No sufficient symmetry correspondence
- Learning-based approach?

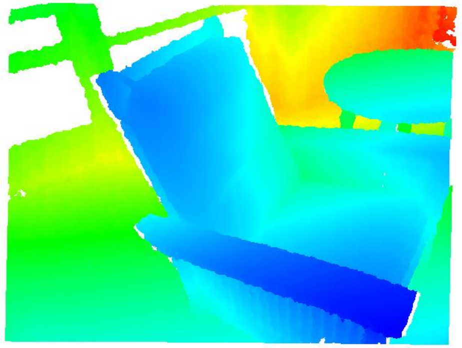

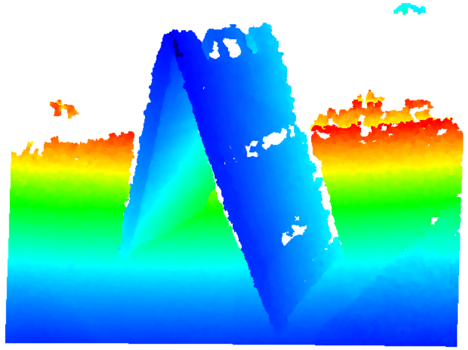

Motivation

- Naïve learning-based approach
  - ➢ Training: memorize the symmetry axes of the category
  - ➢ Testing: perform ***object classification*** & ***pose estimation***

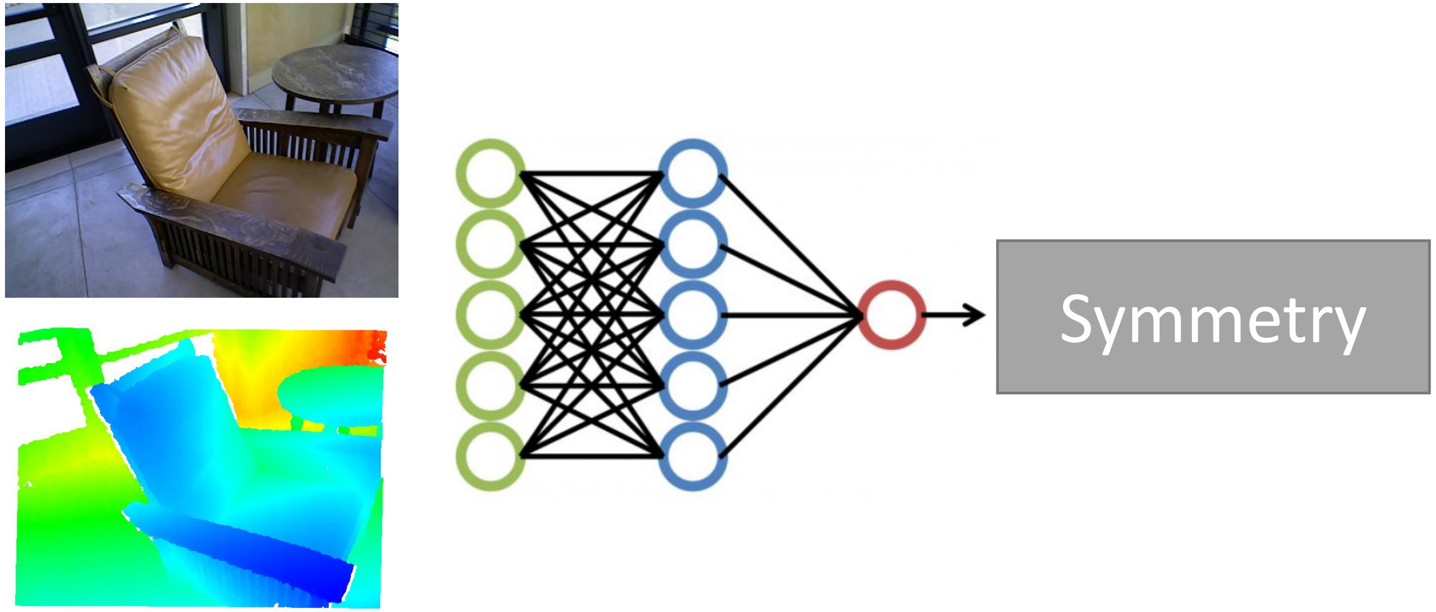

Motivation

- Symmetry is supported by local geometric cues
    - Find local correspondence
    - Aggregate and output symmetry

- The searching of local shape correspondence benefits from the global feature

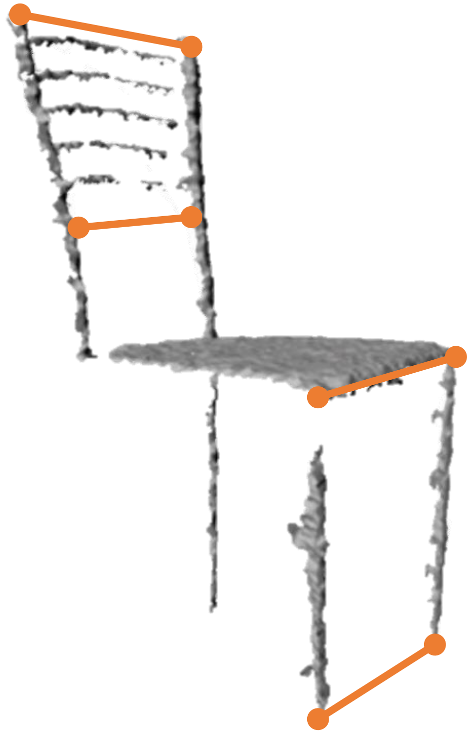

Our approach

- Multi-task learning: predict not only the ***symmetry axis (global)***, but also the ***symmetry correspondences (local)***
- Could detect ***multiple symmetries***
- Could handle both ***reflectional symmetry*** and ***rotational symmetry***
- End-to-end trainable

Related work

- 3D symmetry detection

Partial symmetry detection
[Mitra et al. 2006]

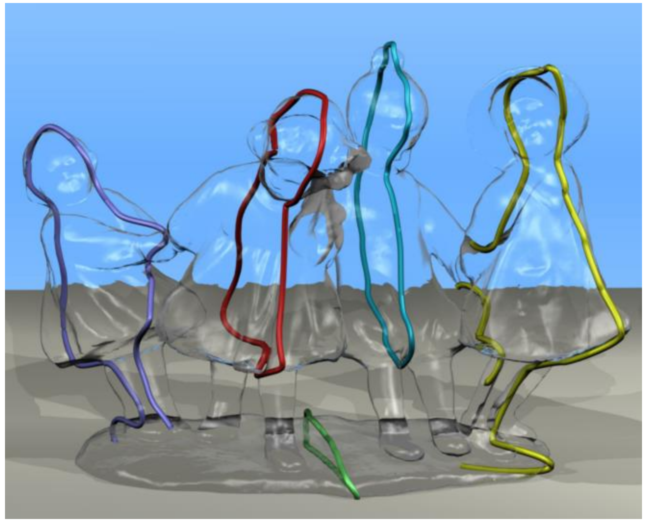

Intrinsic symmetry detection
[Xu et al. 2009]

Related work

- 3D symmetry detection

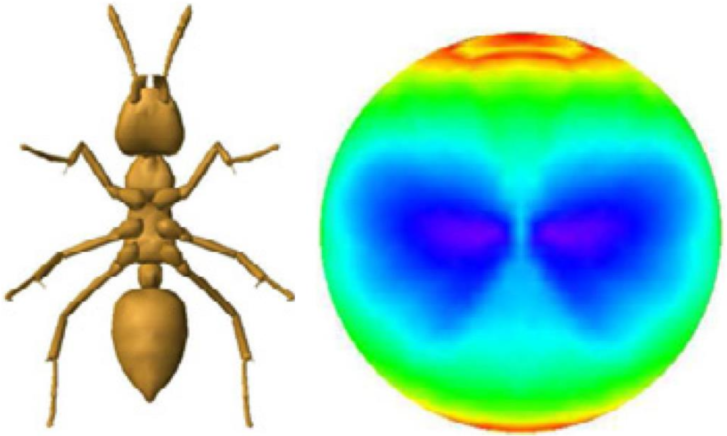

View-based symmetry detection
[Li et al. 2004]

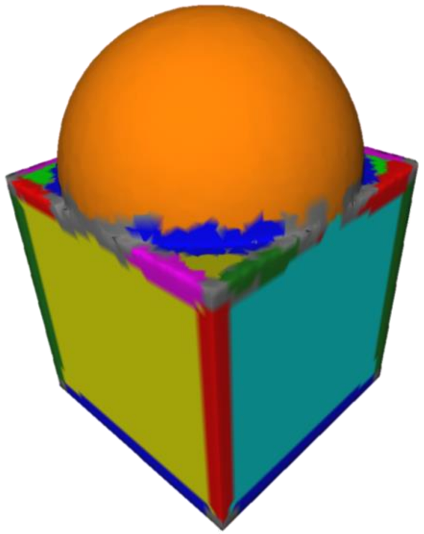

Slippage analysis
[Gelfand et al. 2004]

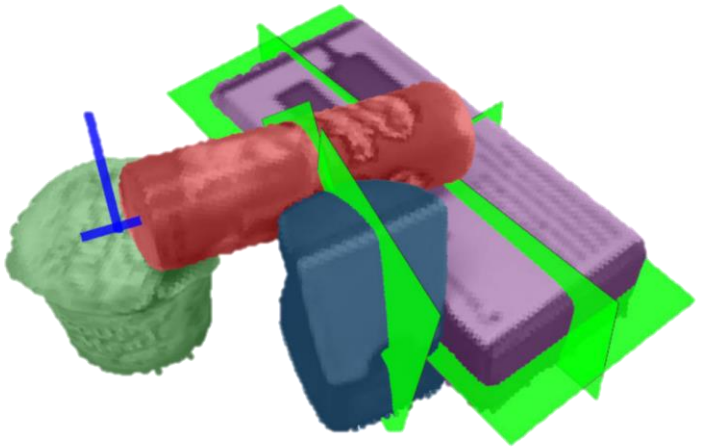

Geometric fitting
[Ecins et al. 2018]

Problem setting

- **Input:** an RGB-D image of an 3D object

➢ Object segmentation is known

- **Output:** the extrinsic reflectional and rotational symmetries

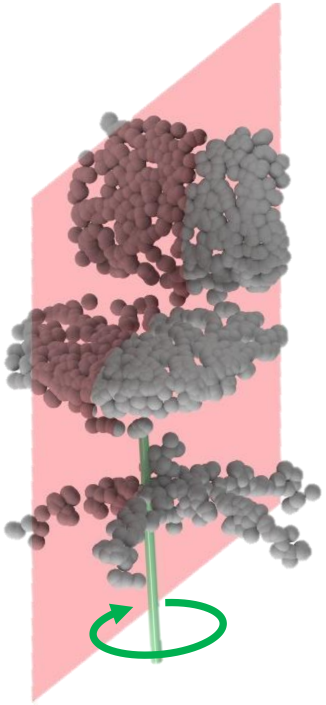

Method

- Symmetry prediction network

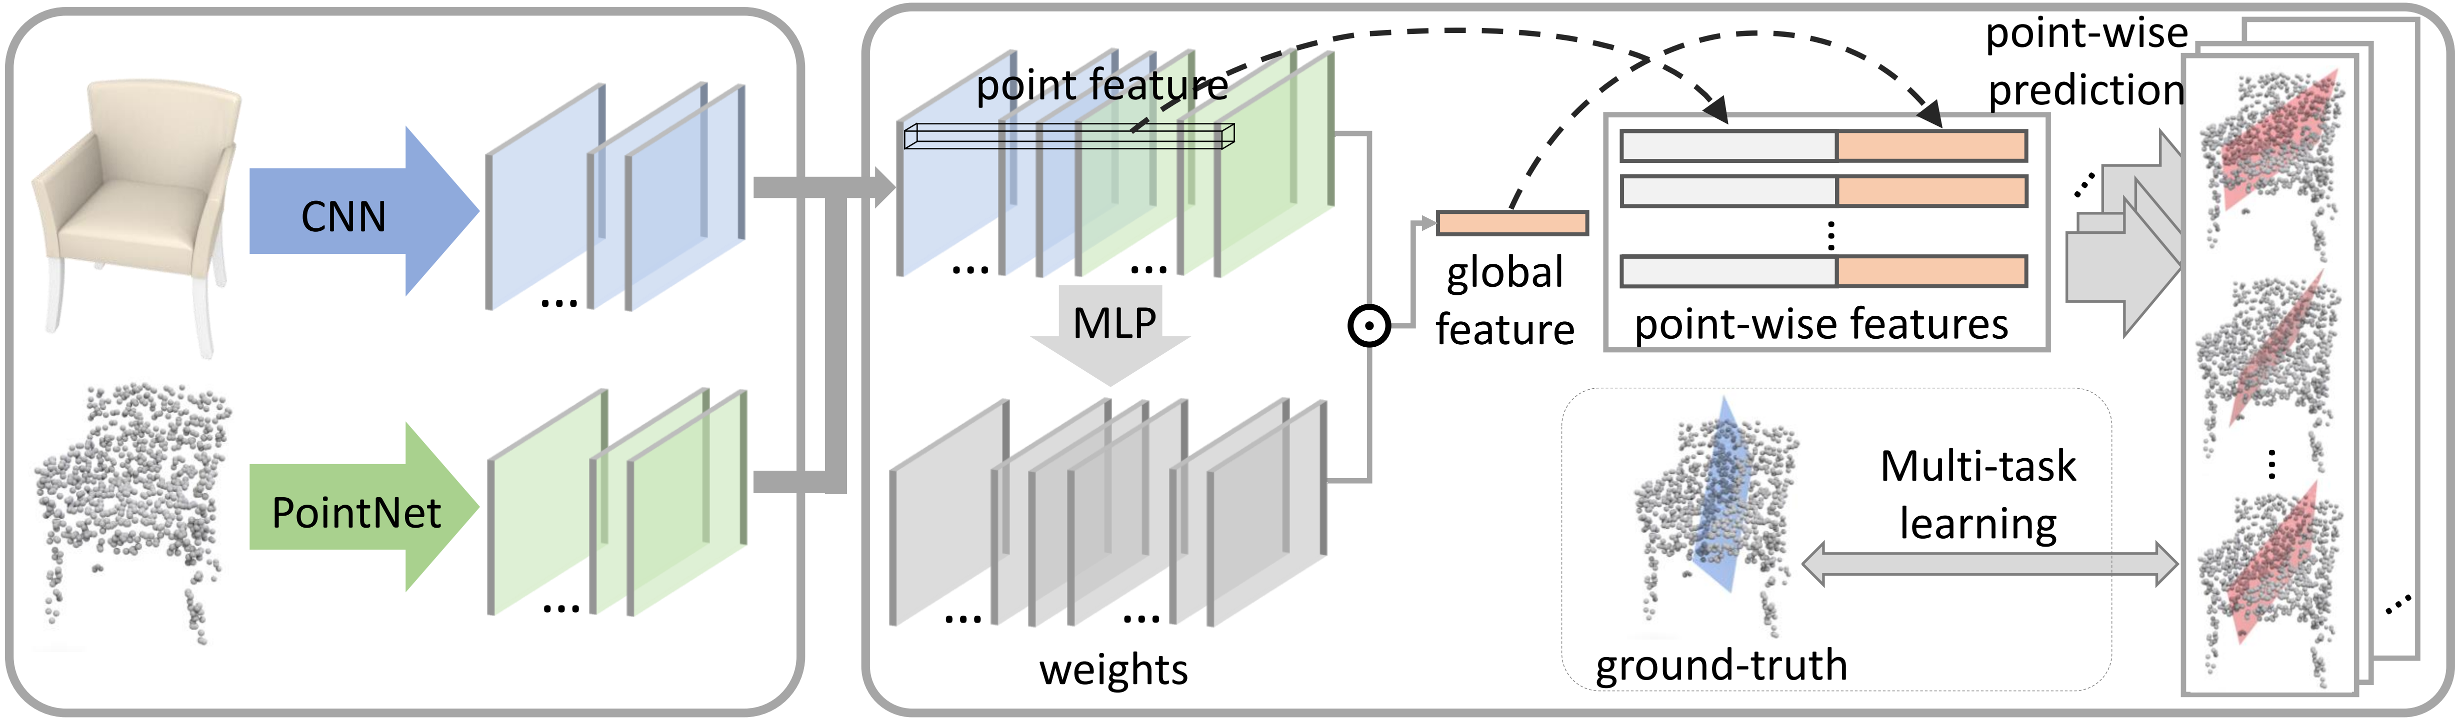

Point-wise feature extraction

Point-wise symmetry prediction

Method

- Loss function

Dense point loss: $$\mathcal{L} = \frac{1}{N}\sum_{i}^{N}\mathcal{L}_i,$$

Point-wise loss: $$\mathcal{L}_i = \mathcal{L}_i^{\text{type}} + \mathcal{L}_i^{\text{sym}},$$

ref. sym. or
rot. sym. or
no sym.

sym. parameters &
counterpart(s) location

Method

- Loss function of reflectional symmetry

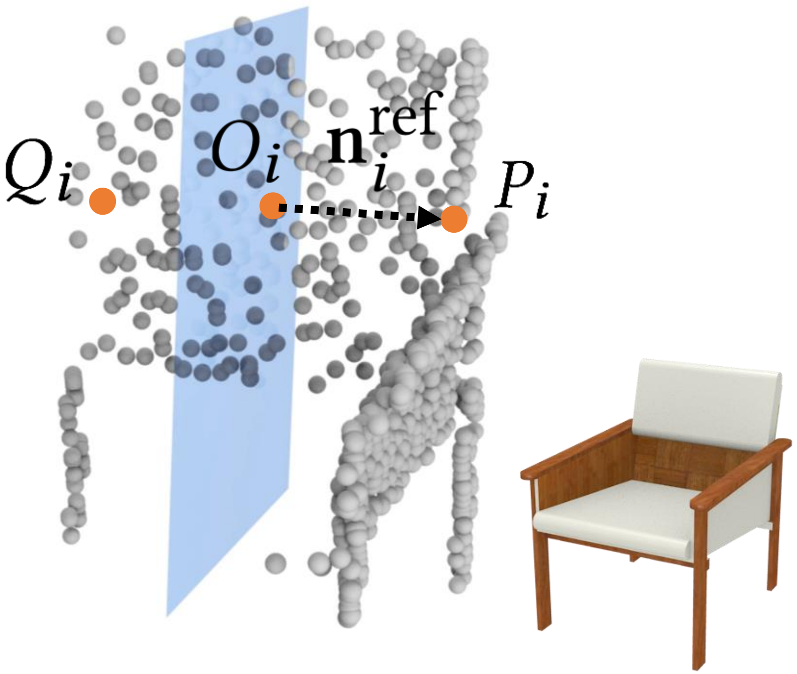

$$\mathcal{L}_i^{\mathrm{ref\_reg}} = d^2(O_i, \hat{O}_i),$$

$$\mathcal{L}_i^{\mathrm{ref\_cp}} = \frac{1}{N}\sum_j^N \mathcal{L}^{\mathrm{cls}}(p_{ij}, \hat{p}_{ij}) + d^2(Q_i, \hat{Q}_i),$$

The probability of point *j* being the counterparts of point *i*

Method

- Loss function of rotational symmetry

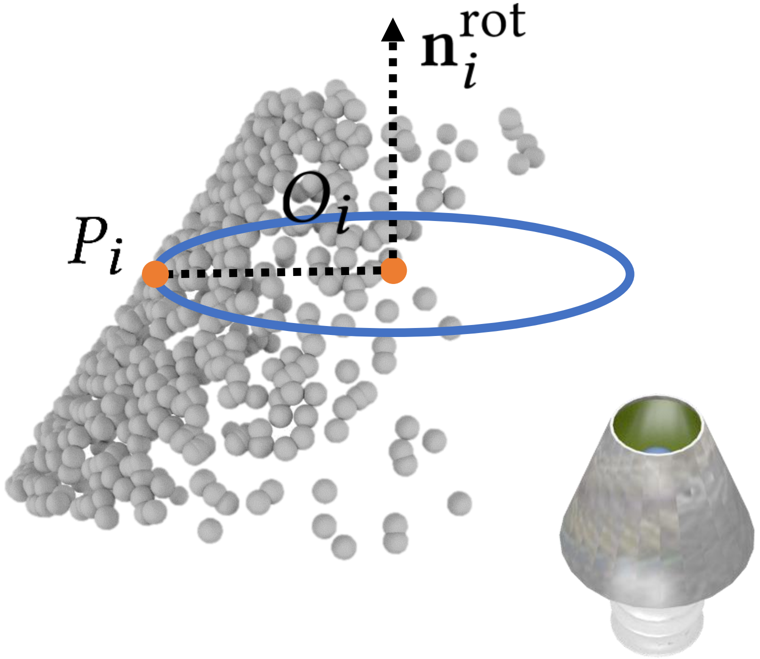

$$\mathcal{L}_i^{\mathrm{rot\_reg}} = d^2(O_i, \hat{O}_i) + \left|1 - |\mathbf{n}_i^{\mathrm{rot}} \cdot \hat{\mathbf{n}}_i^{\mathrm{rot}}|\right|,$$

$$\mathcal{L}_i^{\mathrm{rot\_cp}} = \frac{1}{N} \sum_j^N \mathcal{L}^{\mathrm{cls}}(p_{ij}^o, \hat{p}_{ij}^o)$$

The probability of point *j being the* counterparts of point *i*

Method

- Handle multiple symmetries

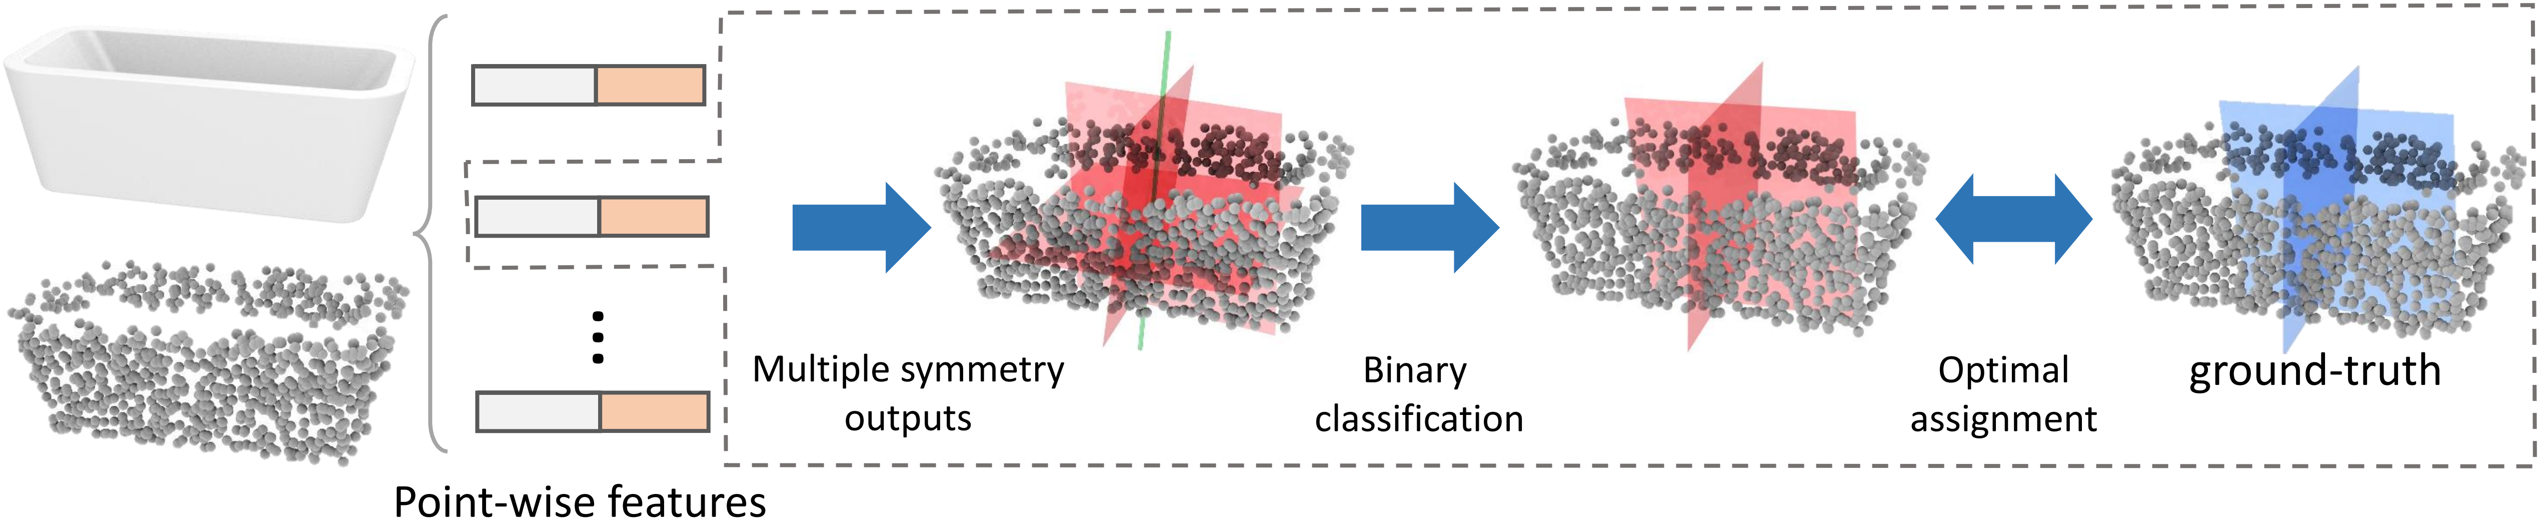

Method

- Inference

➢ Step 1: Prediction aggregation (Density-Based Spatial Clustering)

➢ Step 2: Visibility-based verification

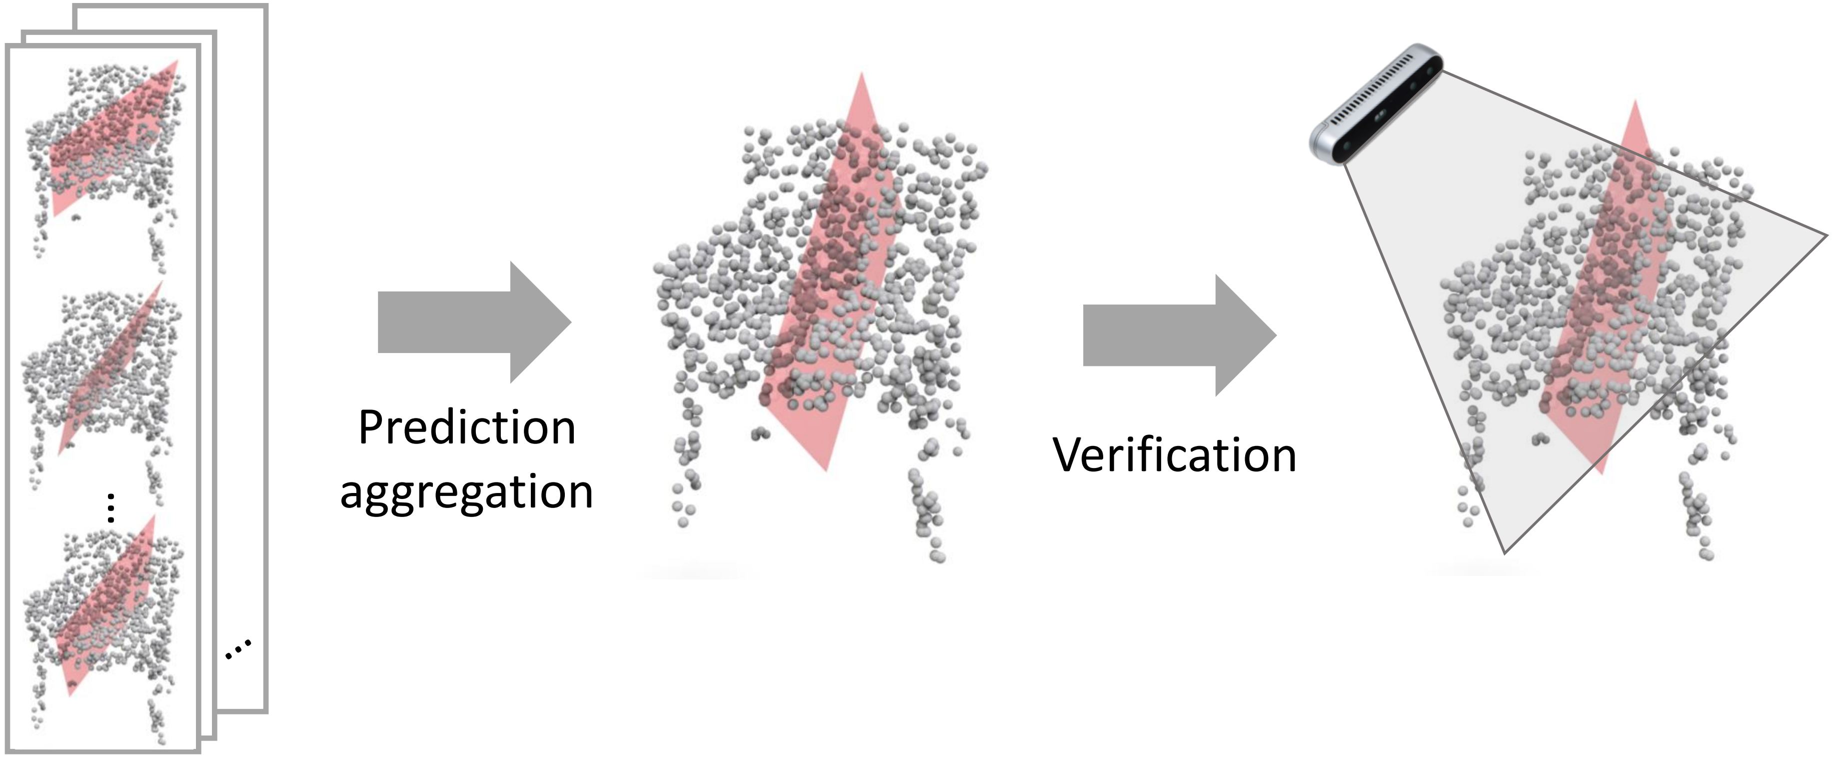

Benchmark

- We construct the 3D symmetry detection benchmark on …

➢ ShapeNet (synthetic)

➢ YCB (real)

➢ ScanNet (real)

Dataset	Subset	#View	#Object	#Scene
ShapeNet	Train	300 000	30 000	-
	Holdout view	7 200	2 400	-
	Holdout instance	7 200	2 400	-
	Holdout category	4 800	1 600	-
YCB	Train	16 189	18	80
	Test	2 949	18	12
ScanNet	Train	13 126	1 642	400
	Holdout view	4 723	1 642	400
	Holdout scene	1121	425	100

Experiments

- Qualitative results

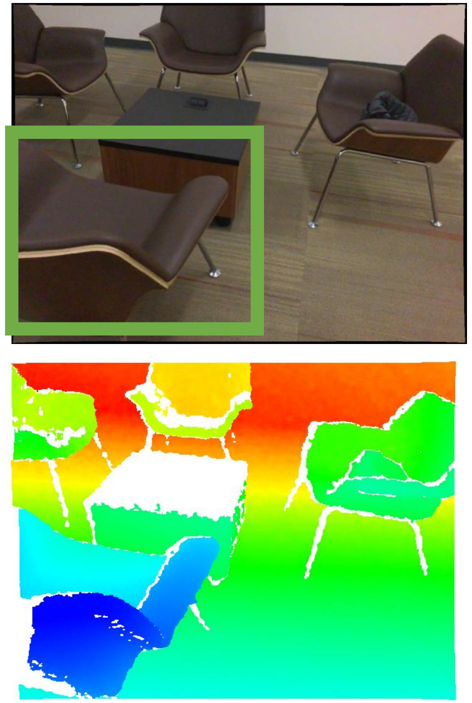

RGB-D image

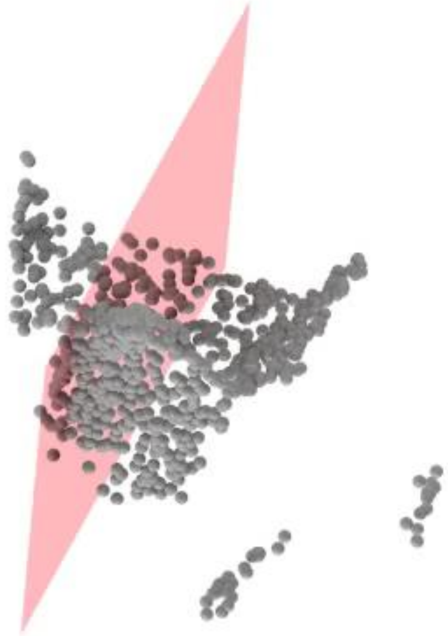

Predicted symmetry

Experiments

- Qualitative results

RGB-D image

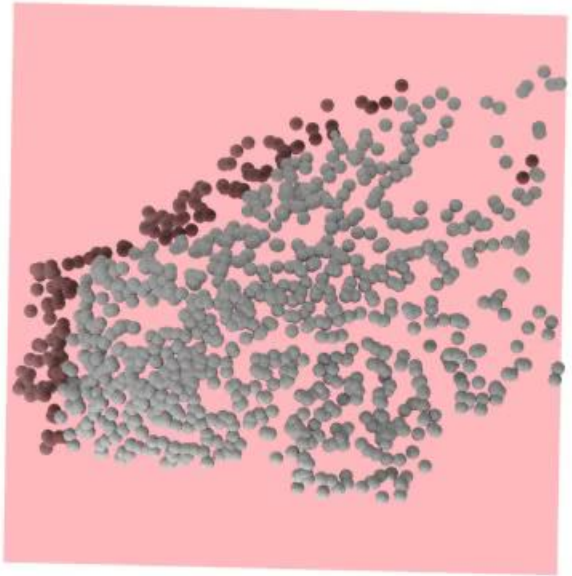

Predicted symmetry

Experiments

- Qualitative results

RGB-D image

Predicted symmetry

Experiments

- Qualitative results

RGB-D image

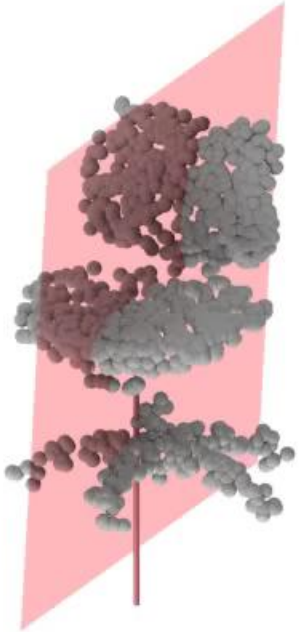

Predicted symmetry

Experiments

- Compare to baselines

➢ Geometric Fitting [Ecins et al. 2018]

➢ RGB-D Retrieval [Yang et al. 2018]

➢ Shape Completion [Liu et al. 2020]

- Evaluation metric

➢ Precision-recall curve [Funk et al. 2017]

Experiments

- Compare to baselines on ShapeNet

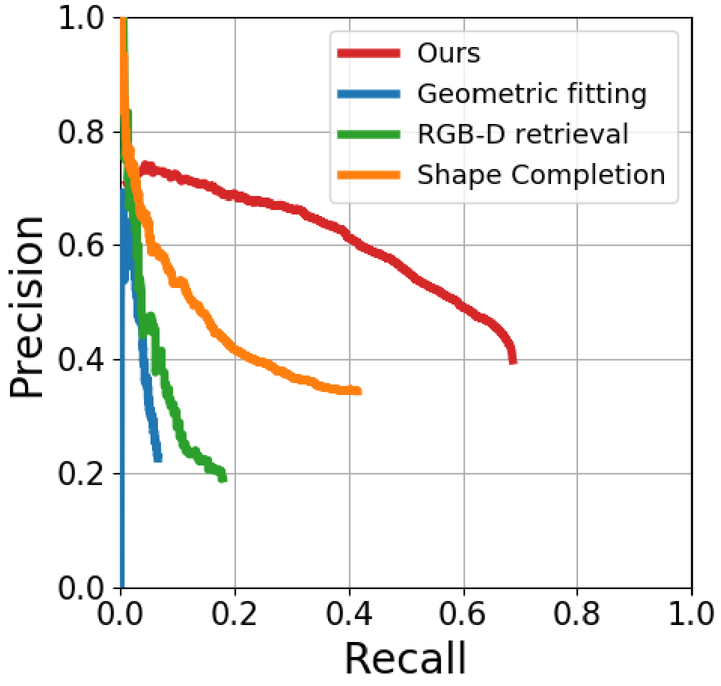

ShapeNet holdout view

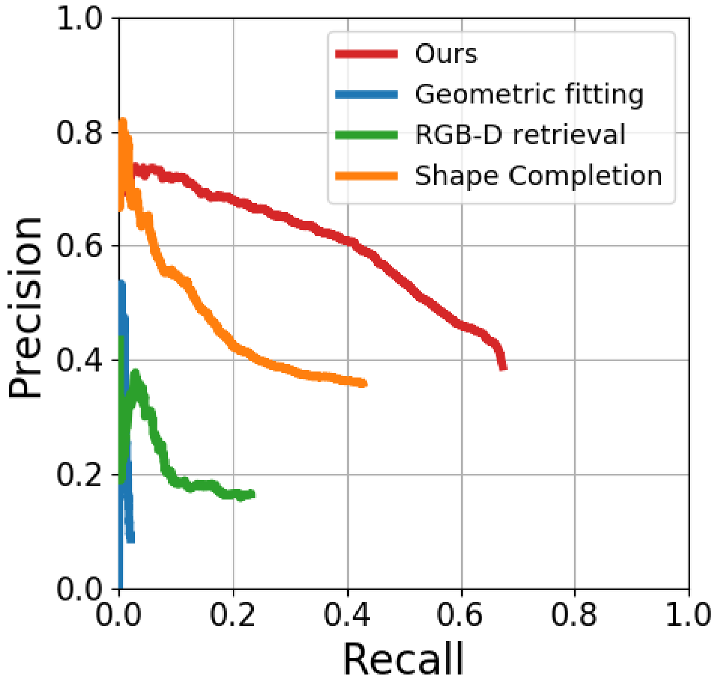

ShapeNet holdout instance

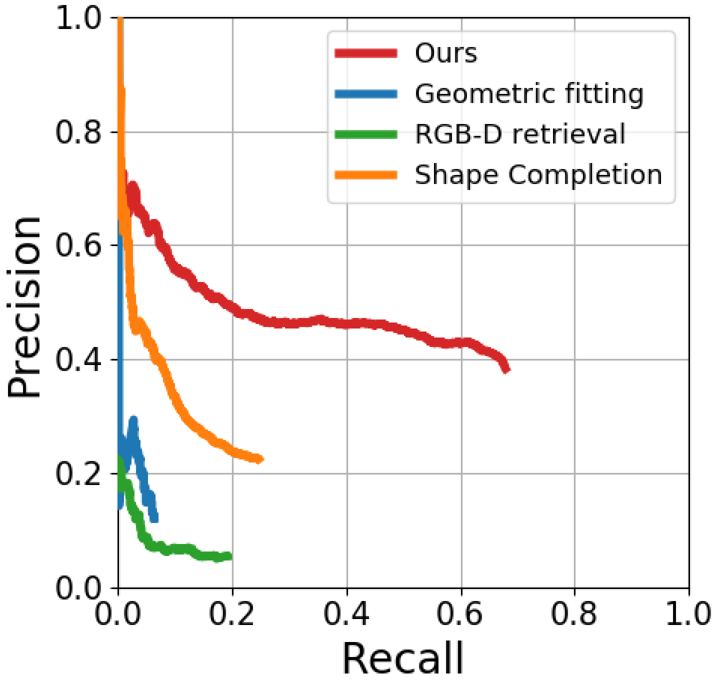

ShapeNet holdout category

Experiments

- Qualitative comparison

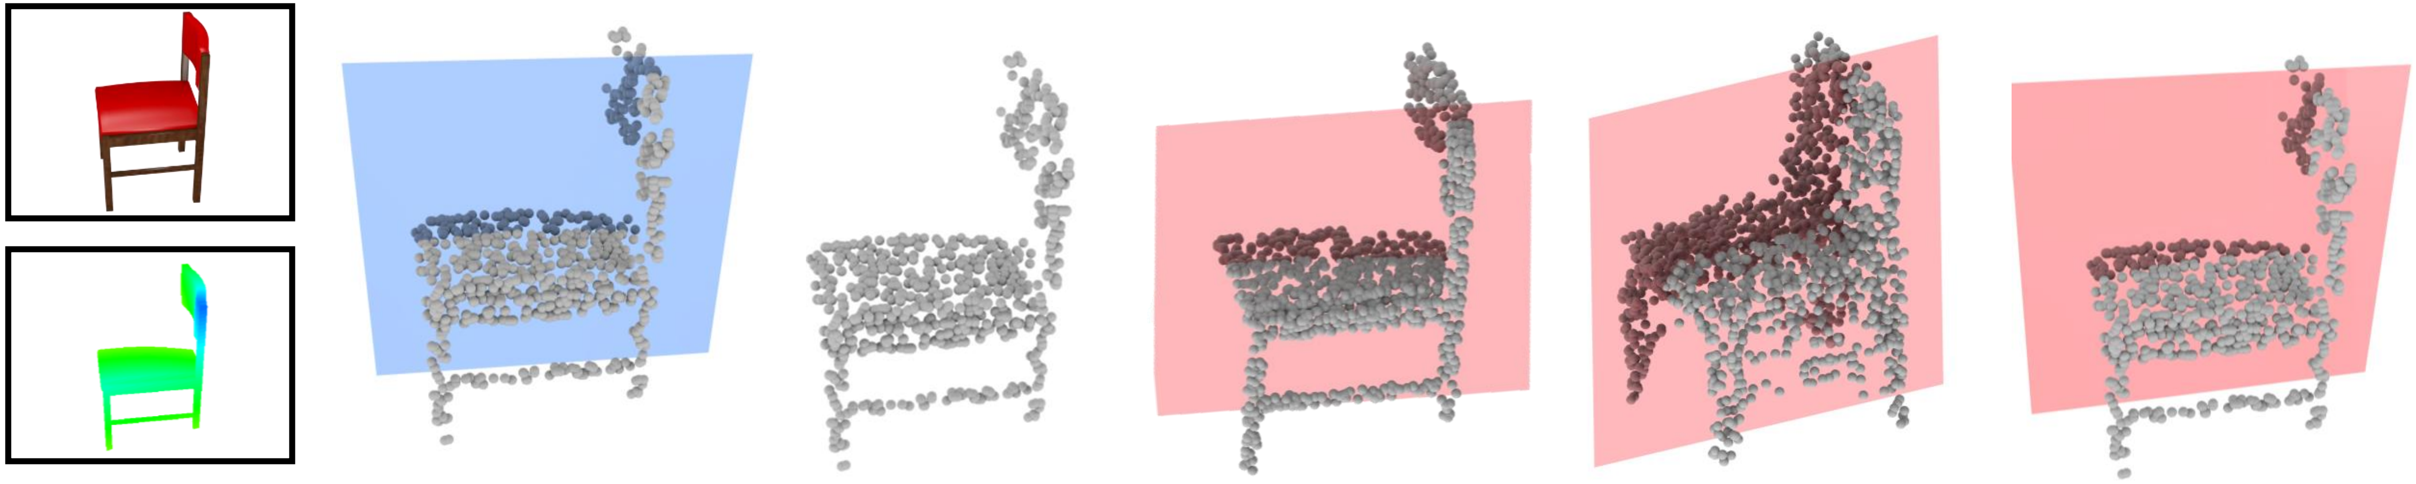

RGB-D image | Ground-truth | Geometric Fitting | RGB-D Retrieval | Shape Completion | Ours

Experiments

- Qualitative comparison

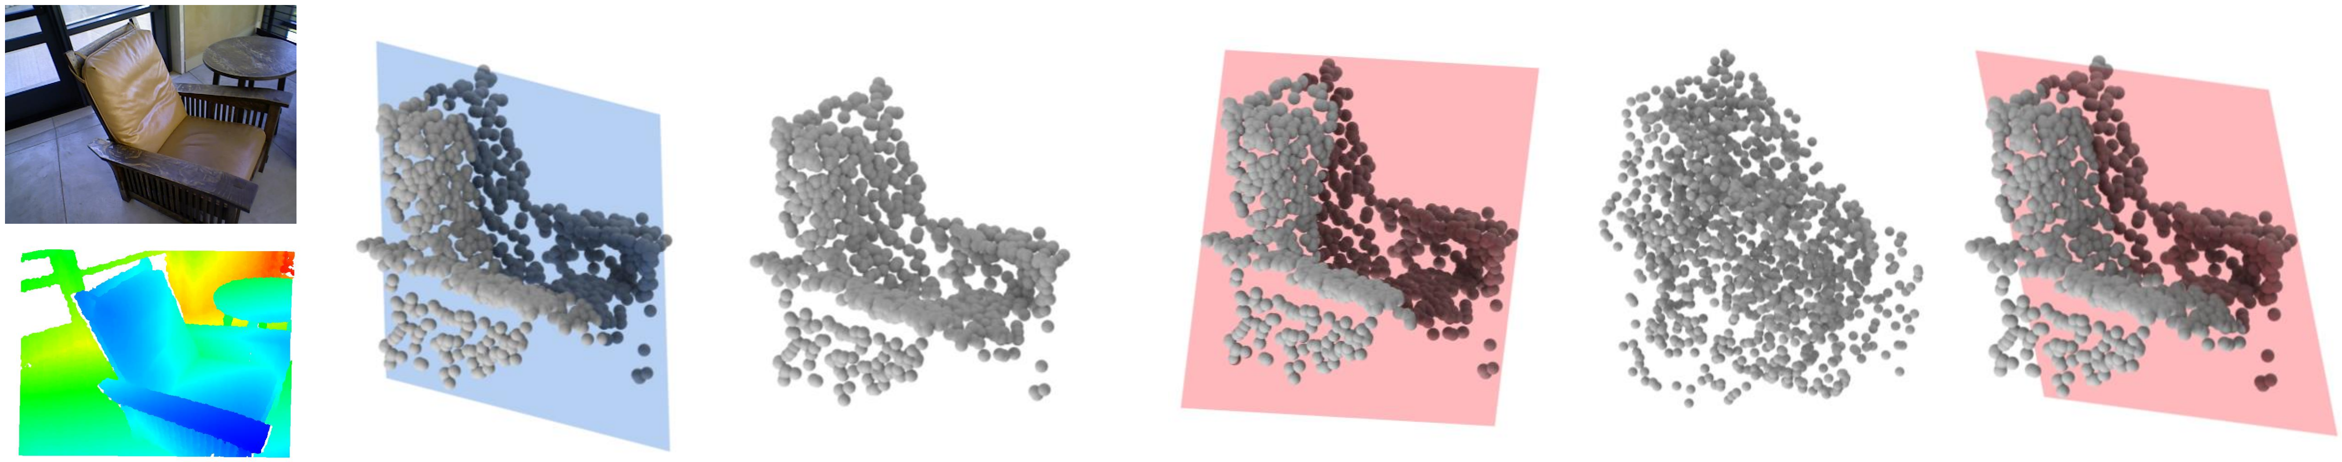

RGB-D image | Ground-truth | Geometric Fitting | RGB-D Retrieval | Shape Completion | Ours

Experiments

- Ablation study

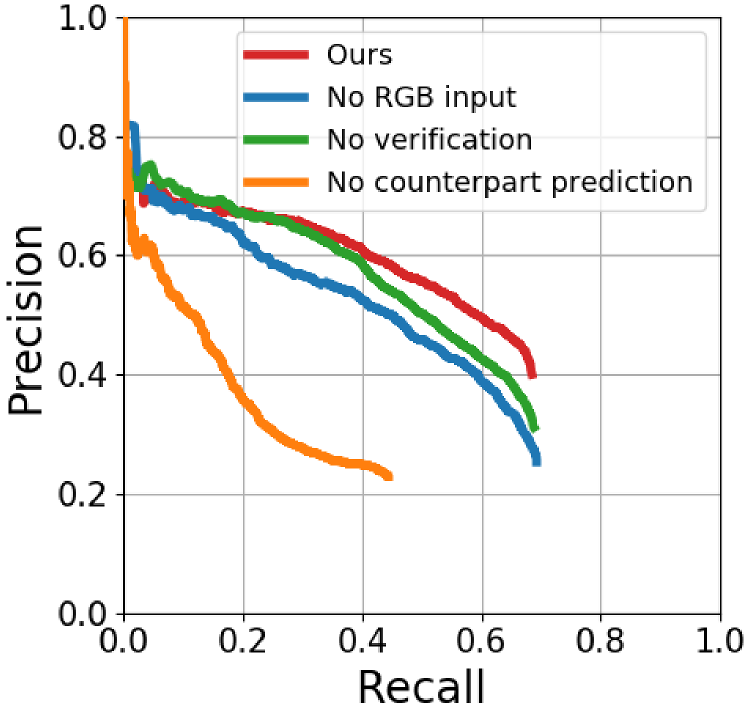

ShapeNet holdout view

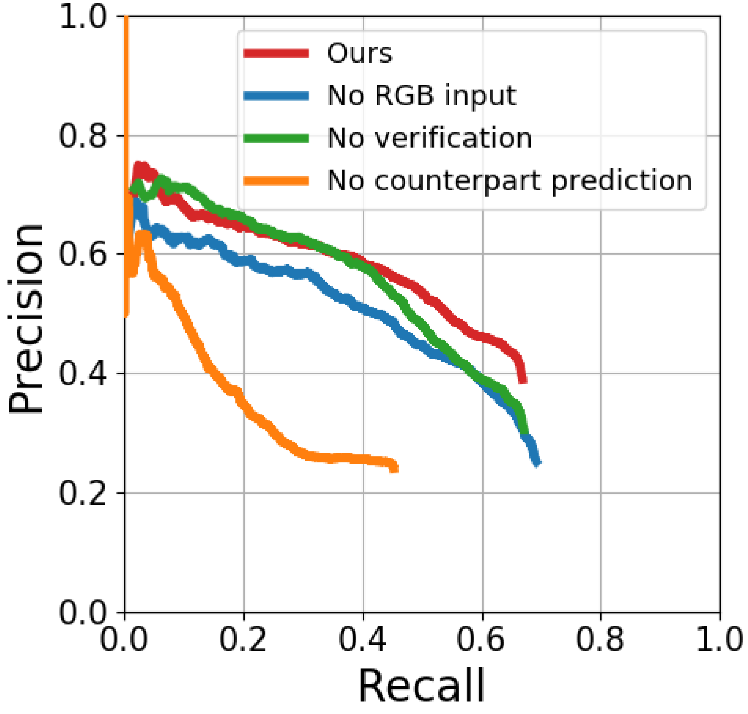

ShapeNet holdout instance

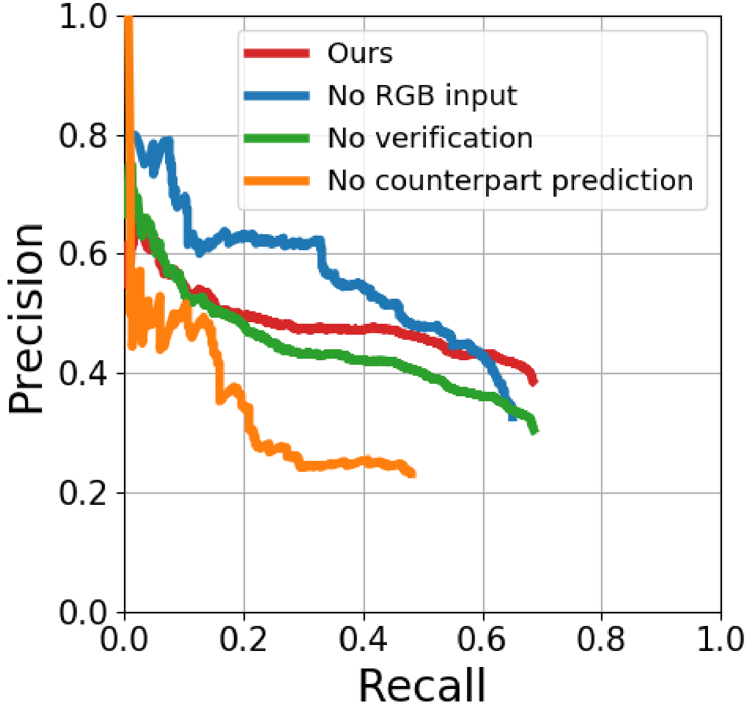

ShapeNet holdout category

Experiments

- Sensitively to occlusion

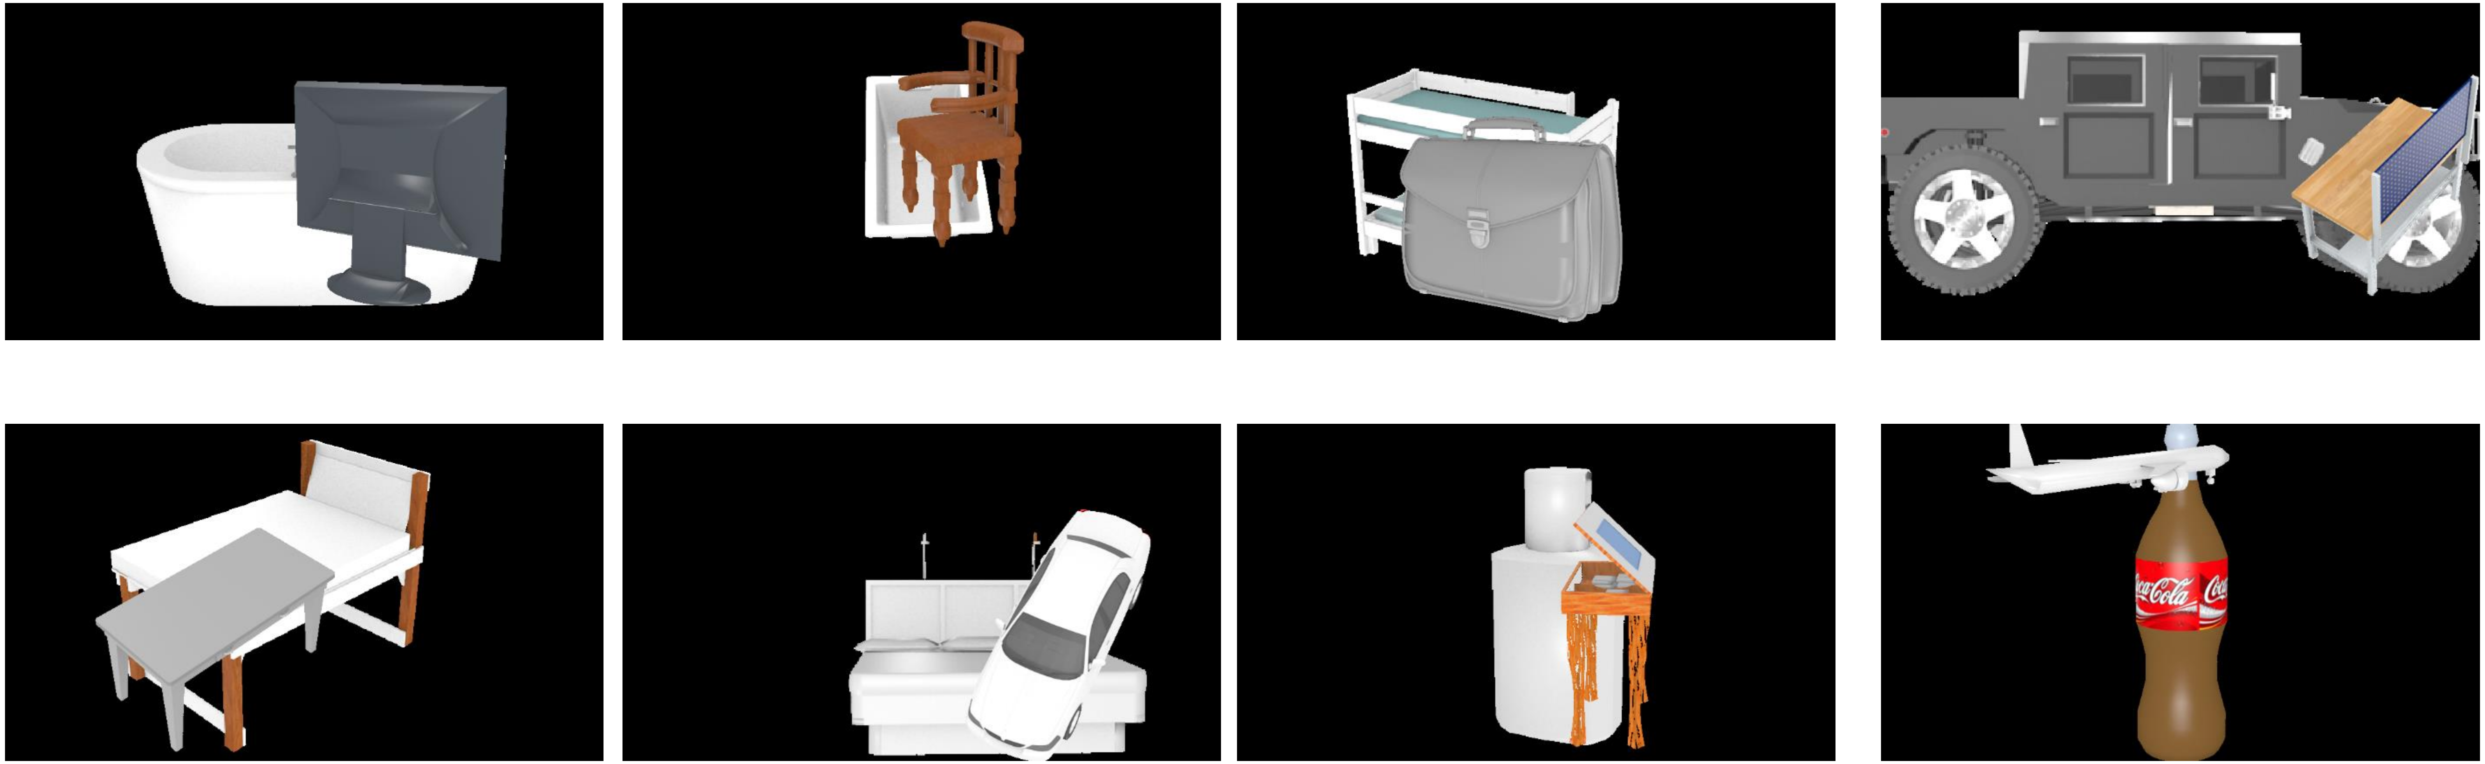

Experiments

- Sensitively to occlusion

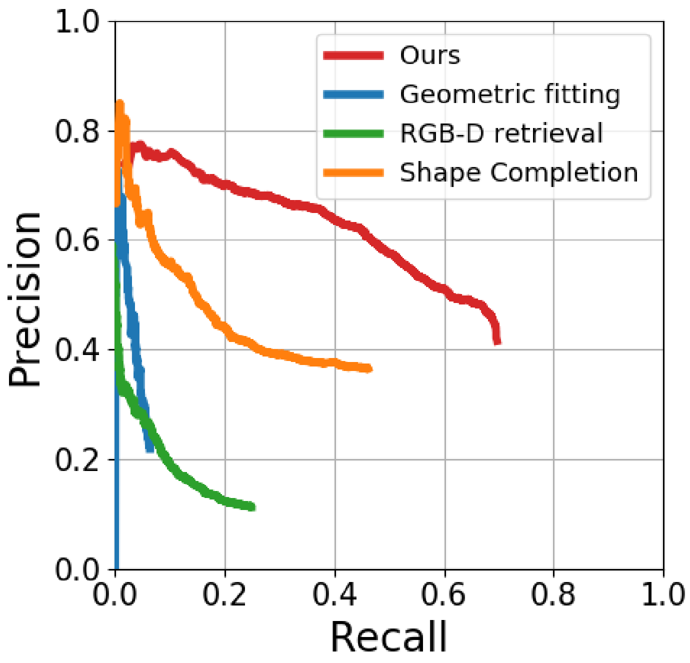

light (50-60%) occlusion

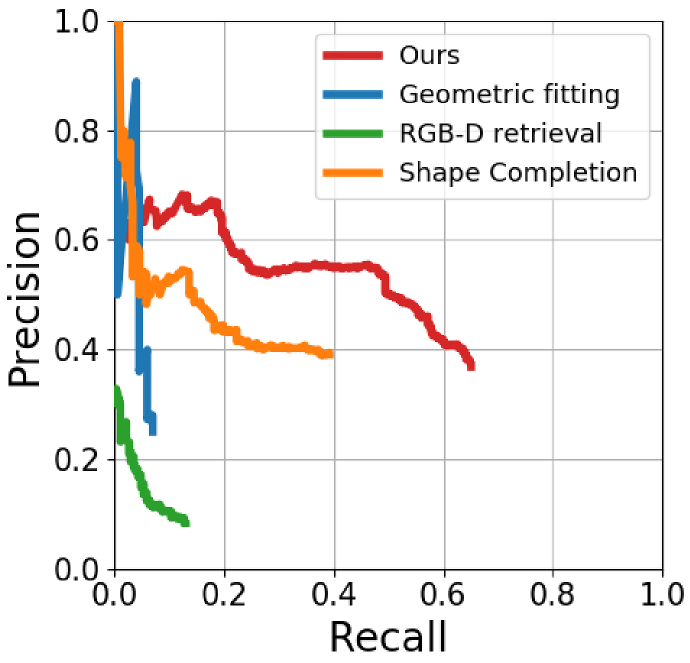

medium (60-70%) occlusion

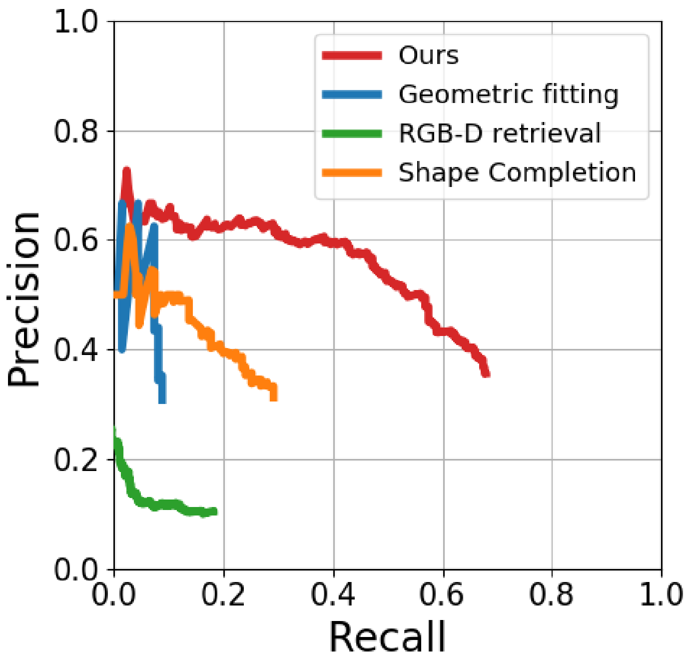

heavy (70-80%) occlusion

Experiments

- Visualization of the predicted counterparts

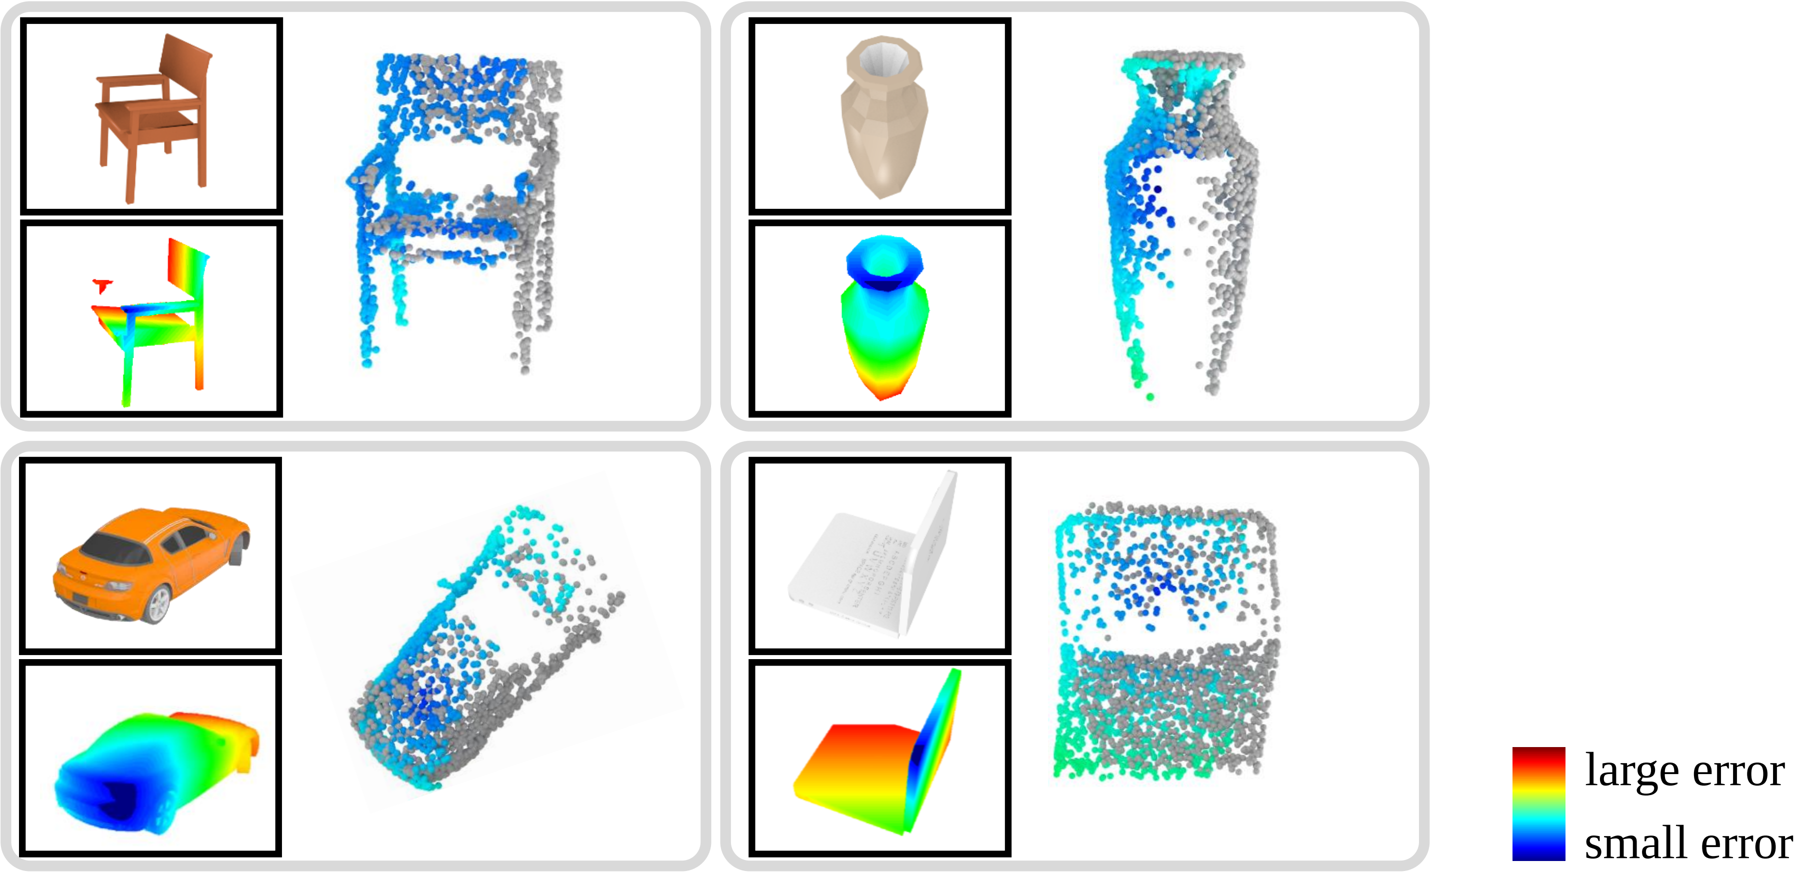

Experiments

- Runtime analysis

Dataset	Network train	Network inference	Aggregation	Verification
ShapeNet	64 h	50 ms	50 ms	40 ms
YCB	20 h	50 ms	50 ms	40 ms
ScanNet	26 h	50 ms	50 ms	40 ms

Failure cases

- Spherical symmetry
- Completely missing data

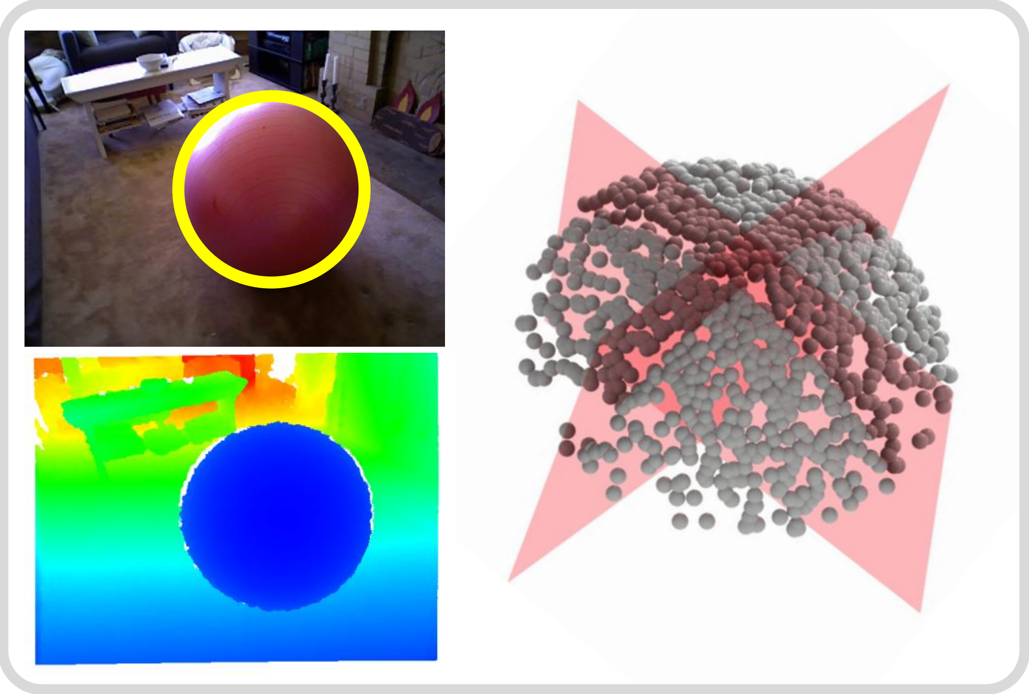

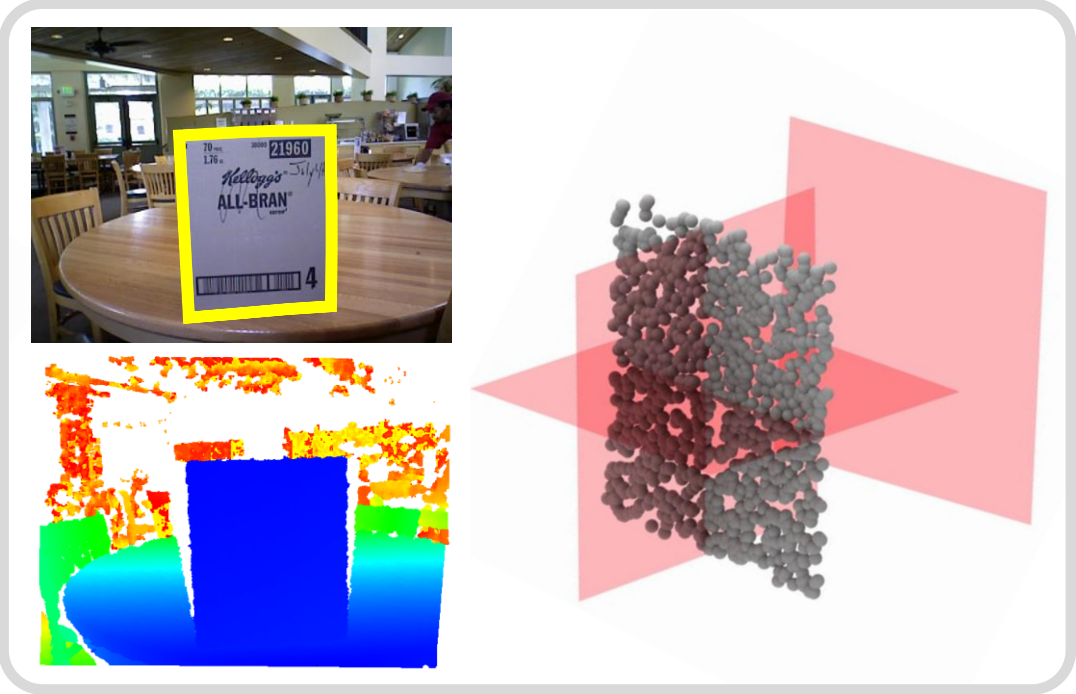

Applications

- 6D pose estimation

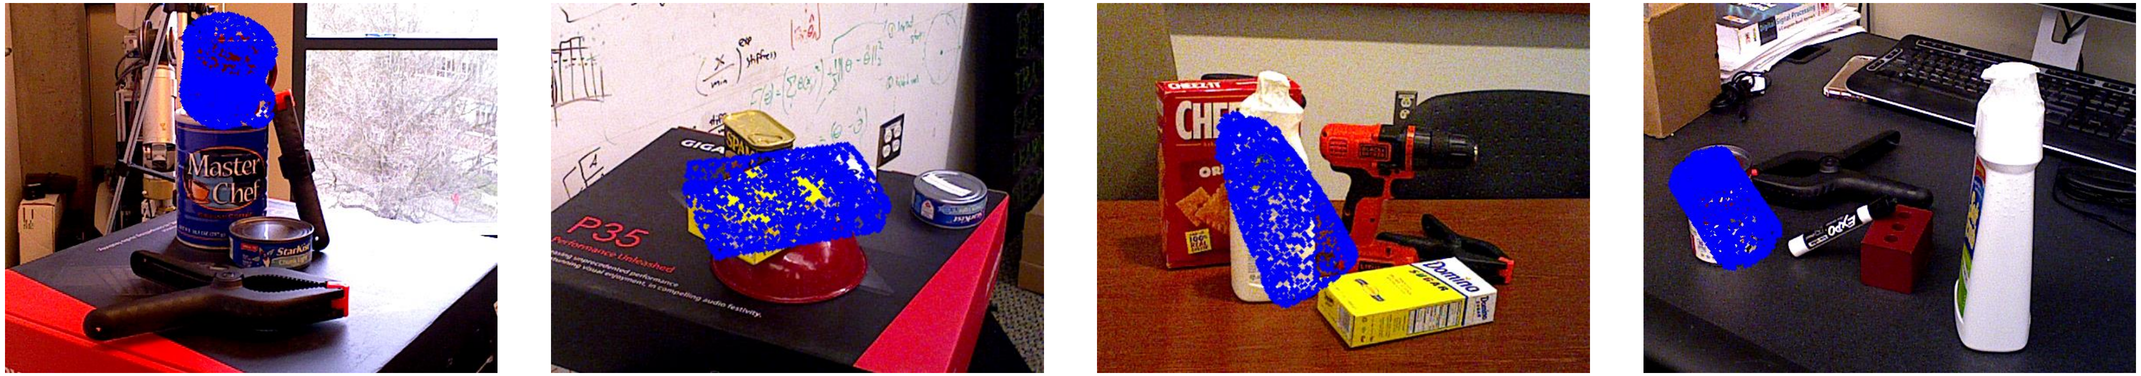

DenseFusion [Wang et al. 2019]

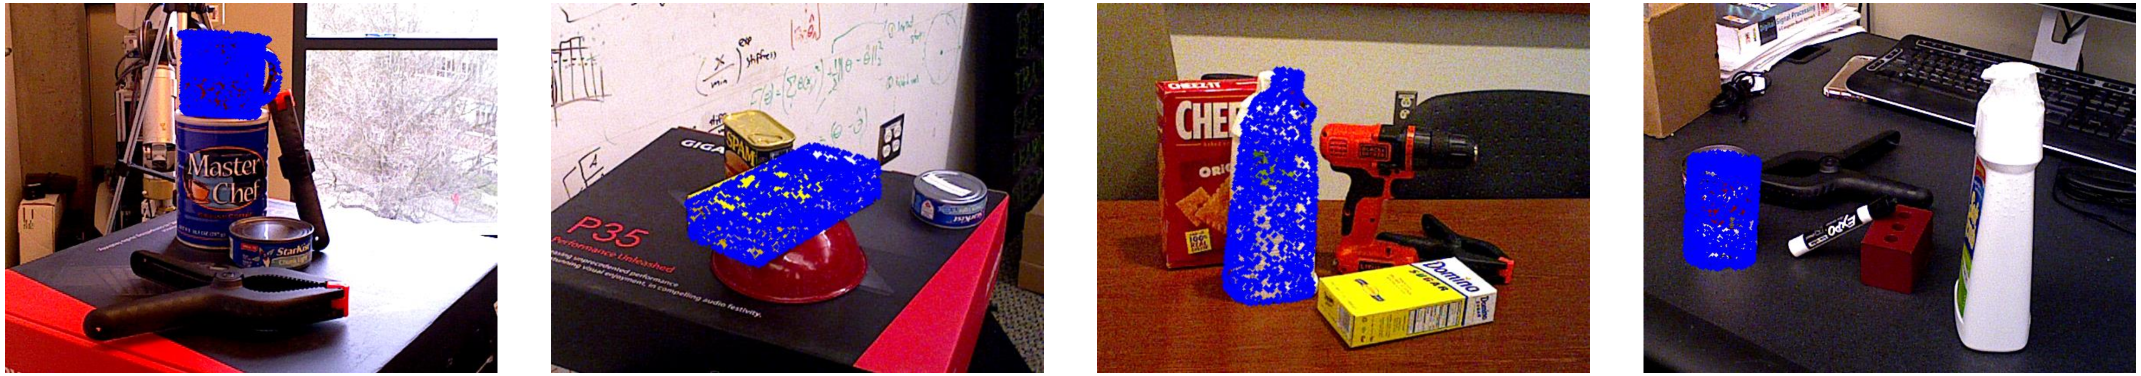

DenseFusion + Symmetry prediction

Applications

- Symmetry-based segmentation

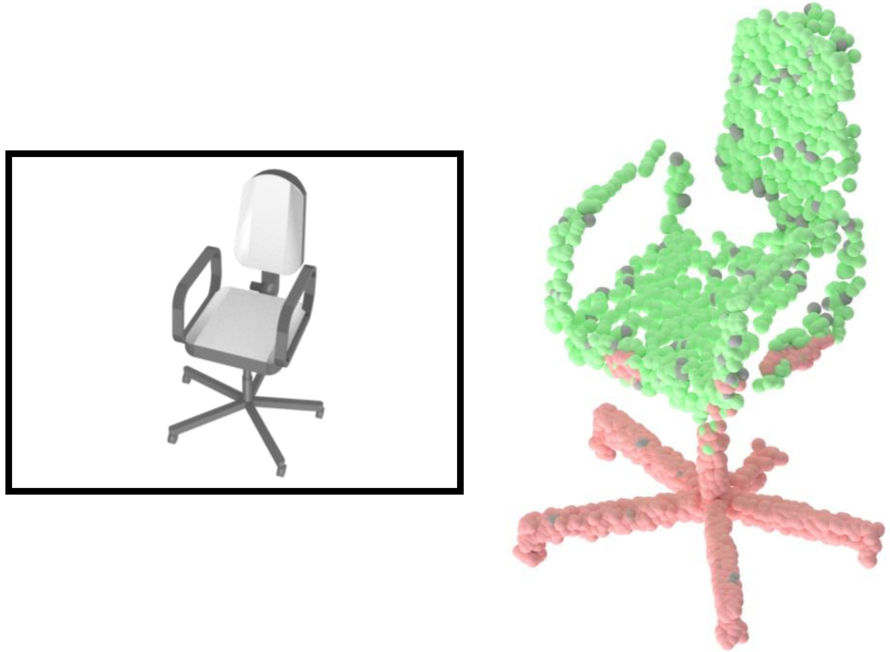

Future work

- Hierarchical symmetries
- Self-supervised approach
- Integrate symmetry prediction into *X*

Thank you