2022UBICOMP

DeepPCD: Enabling AutoCompletion of Indoor Point Clouds with Deep Learning

Pingping Cai, **Sanjib Sur**
sur@cse.sc.edu

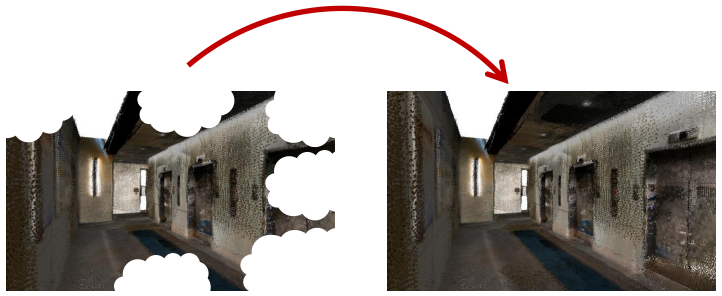

UNIVERSITY OF SOUTH CAROLINA
College of Engineering and Computing

SyReX
Systems Research on X

Point Cloud Data (PCD) and Its Applications

- Perception of surrounding 3D environments enables many ubiquitous sensing applications
- Point Cloud Data (PCD) is an efficient and popular data structure for machine representation of 3D environments

Augmented/Extended Reality

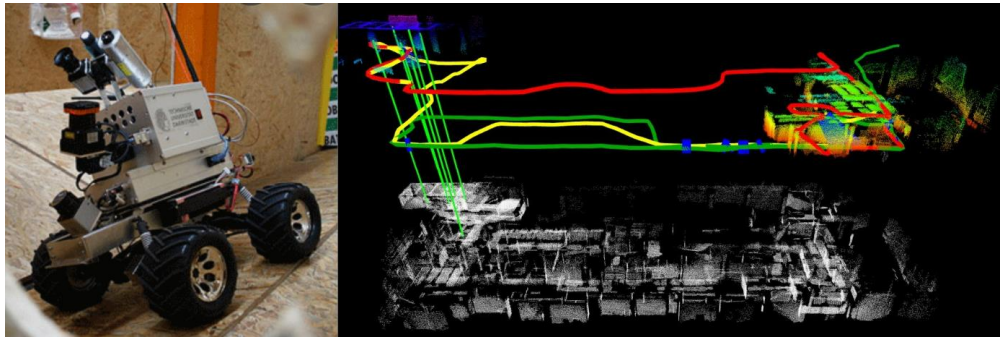

Mobile Robot Indoor SLAM

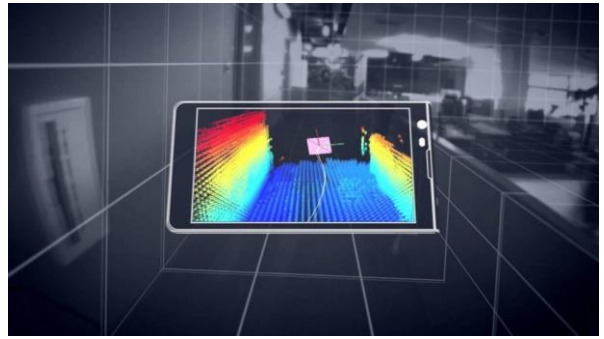

Surface and Floor Mapping

Point Cloud Data (PCD) and Its Applications

- Perception of surrounding 3D environments enables many ubiquitous sensing applications
- Point Cloud Data (PCD) is an efficient and popular data structure for machine representation of 3D environments

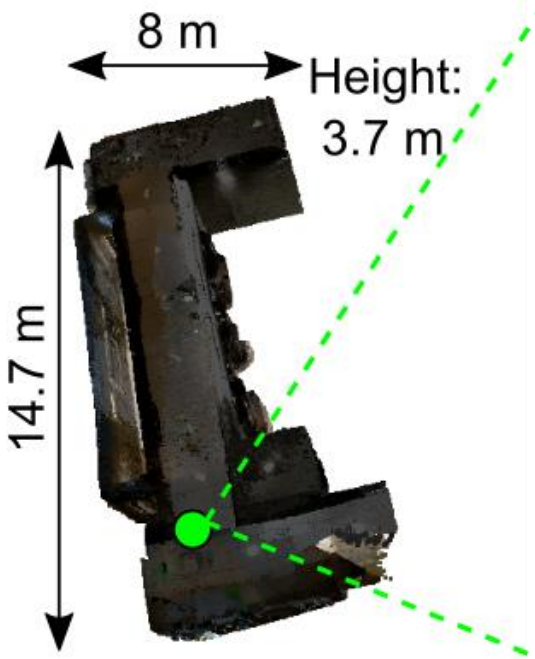

Structure Color

$\{X, Y, Z, R, G, B\}$

$\{X_1, Y_1, Z_1, R_1, G_1, B_1\}$

$\{X_n, Y_n, Z_n, R_n, G_n, B_n\}$

Constructing a PCD and Its Challenges

- RGB-D and/or LiDAR sensors are used to construct PCD
- Two key challenges
  - It requires a lot of time and effort for large-scale environments
  - It requires precise planning of the scan trajectories

Missing structures and colors

Existing Approaches for Constructing High-Quality PCD

Hardware-based

Long-range and large field-of-view camera and depth sensors

Limited by low visibility and sensor occlusion

Expensive and consumes a lot of power

Software-based

Geometrical and machine-learning models to infer shapes

Works on small PCD and mostly single objects

Reconstructs only shapes, not colors

Challenges in Constructing Indoor PCD

- Structure and color reconstructions

Hard to extract features from a large scene with different objects and preserve local and global structures

Our Proposal: DeepPCD

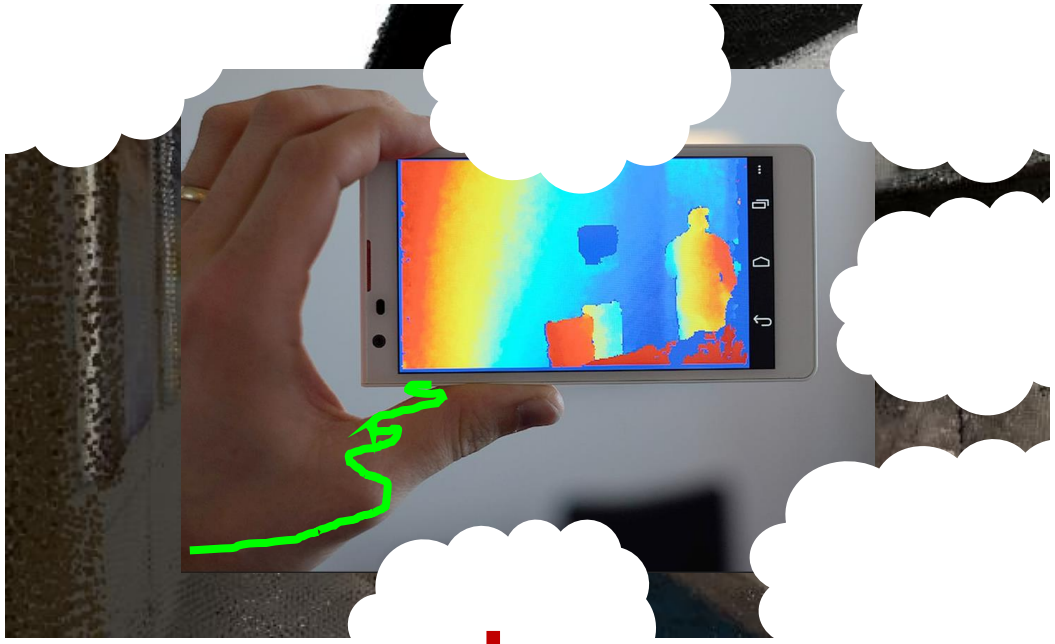

DeepPCD

Our Proposal: DeepPCD

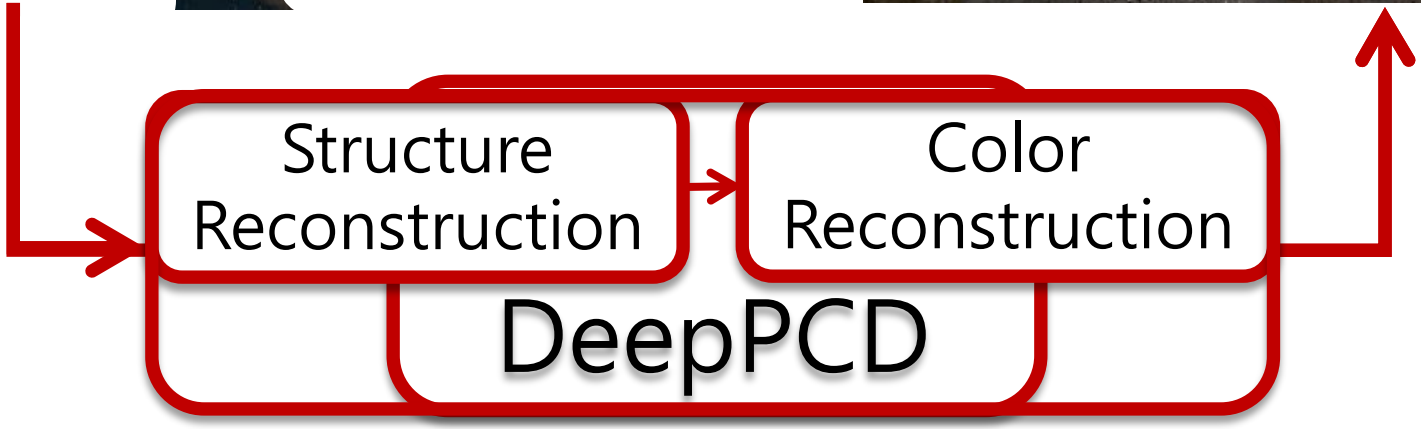

Our Proposal: DeepPCD

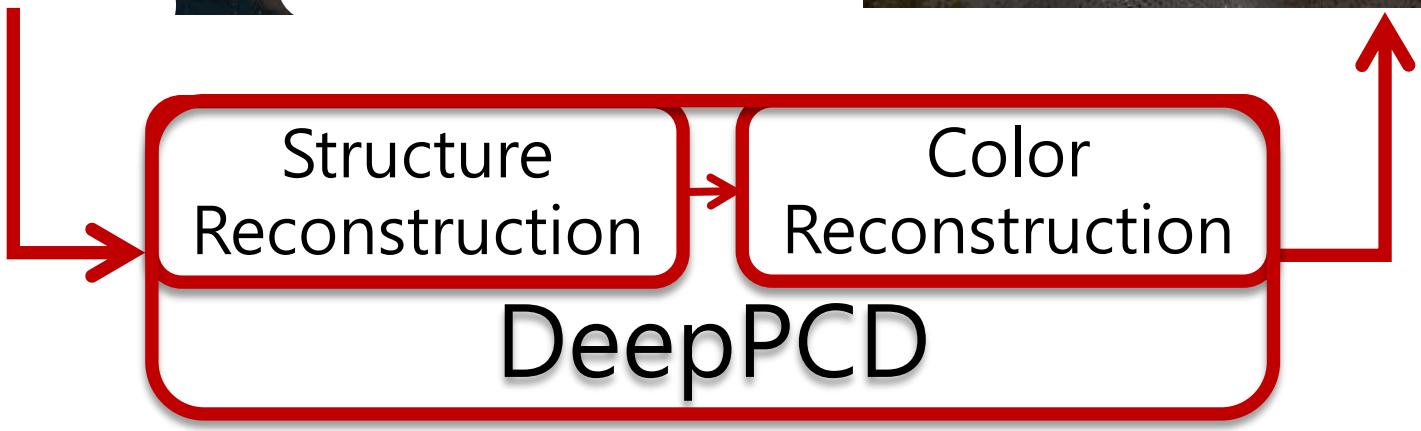

Structure Reconstruction: Key Ideas

A large PCD can be split into multiple small patches

We can reconstruct the patches and merge them iteratively

Challenge: How many patches and where to split them?

Structure Reconstruction: Key Ideas

Indoor PCD consists of geometric structures, e.g., straight walls, floors

We could approximate many points as 3D planes

Challenge: How to automatically find those planes?

Structure Reconstruction Overview

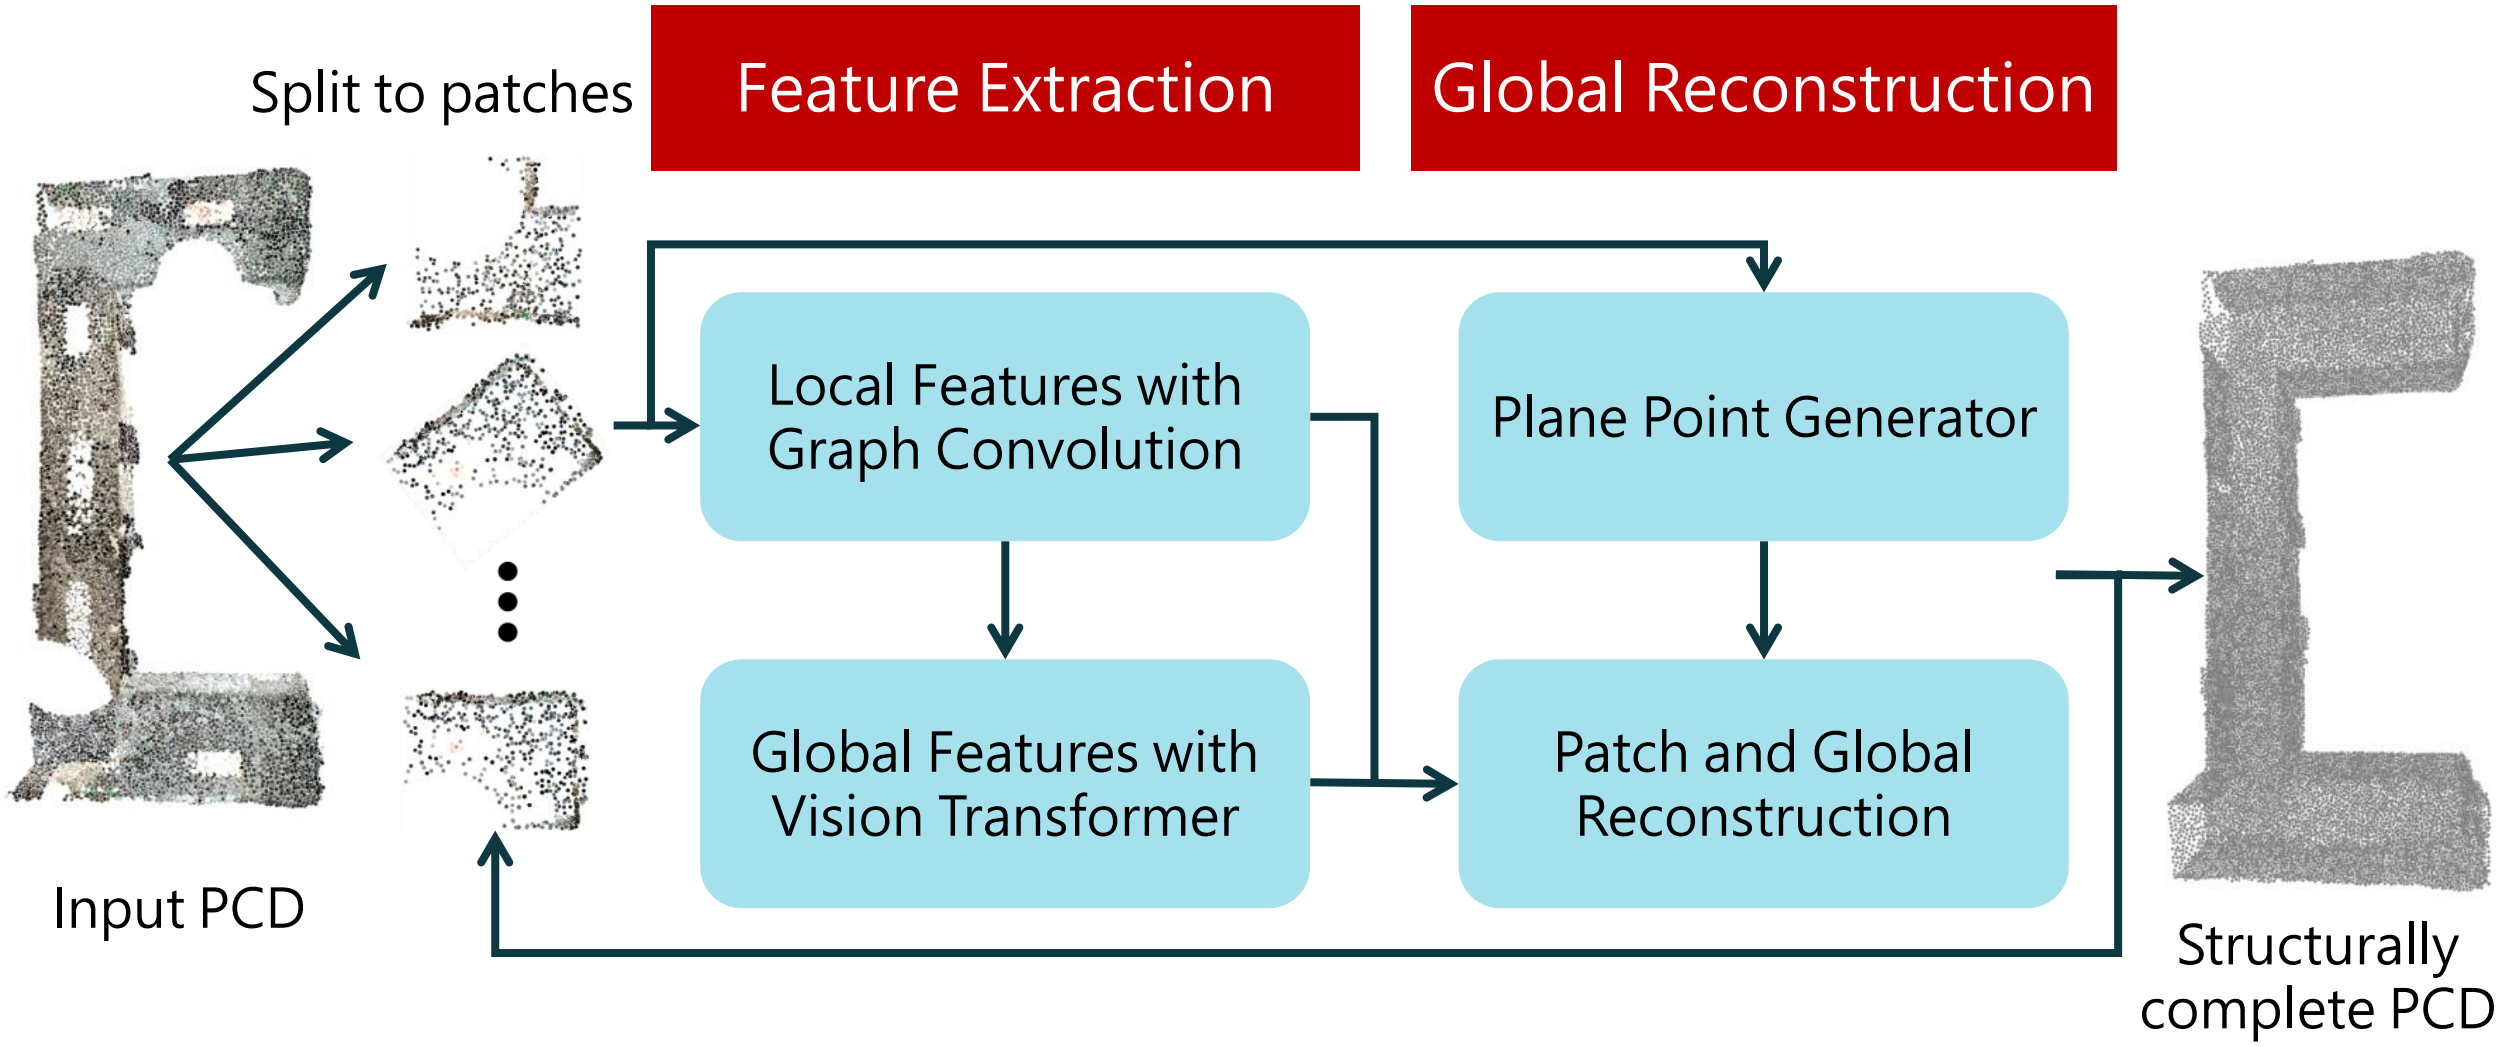

Structure Reconstruction Networks

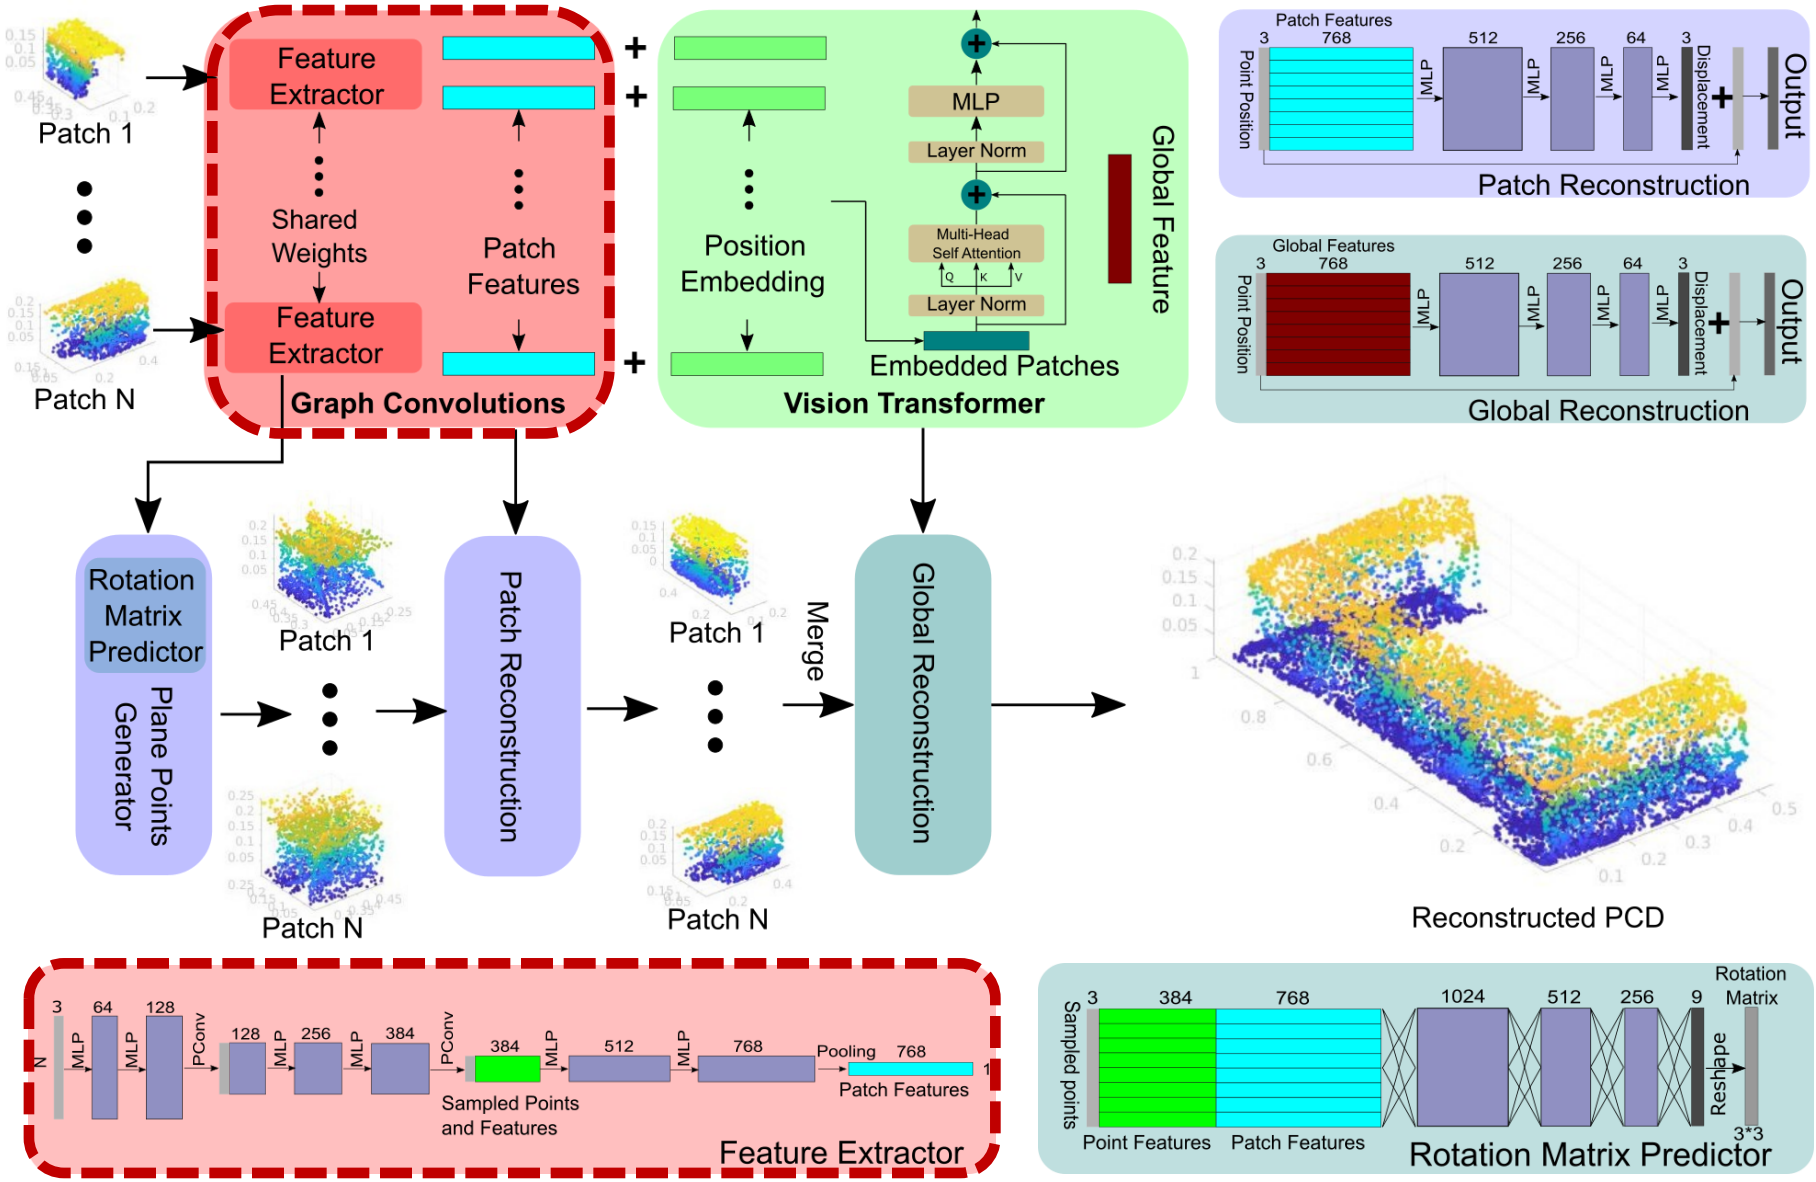

Structure Reconstruction Networks

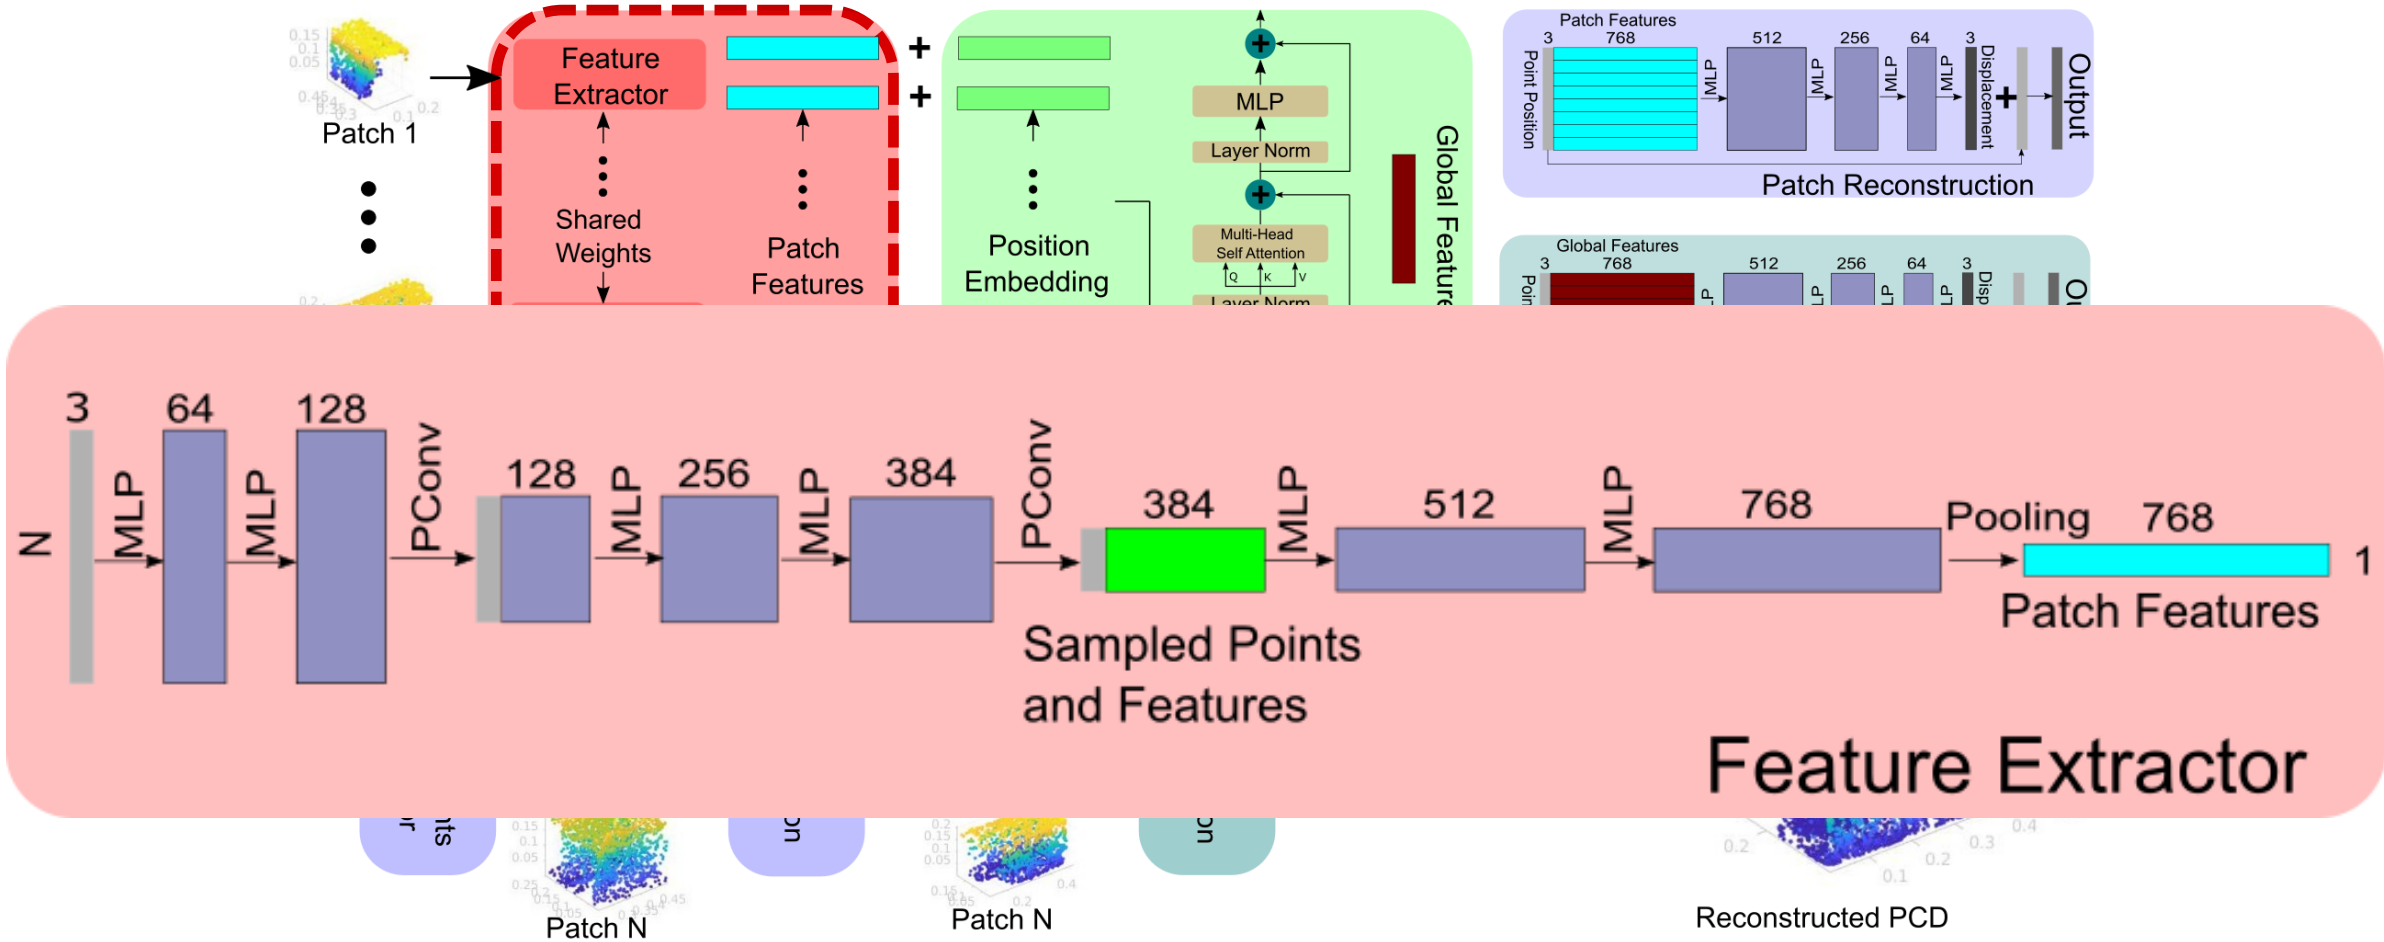

Structure Reconstruction Networks

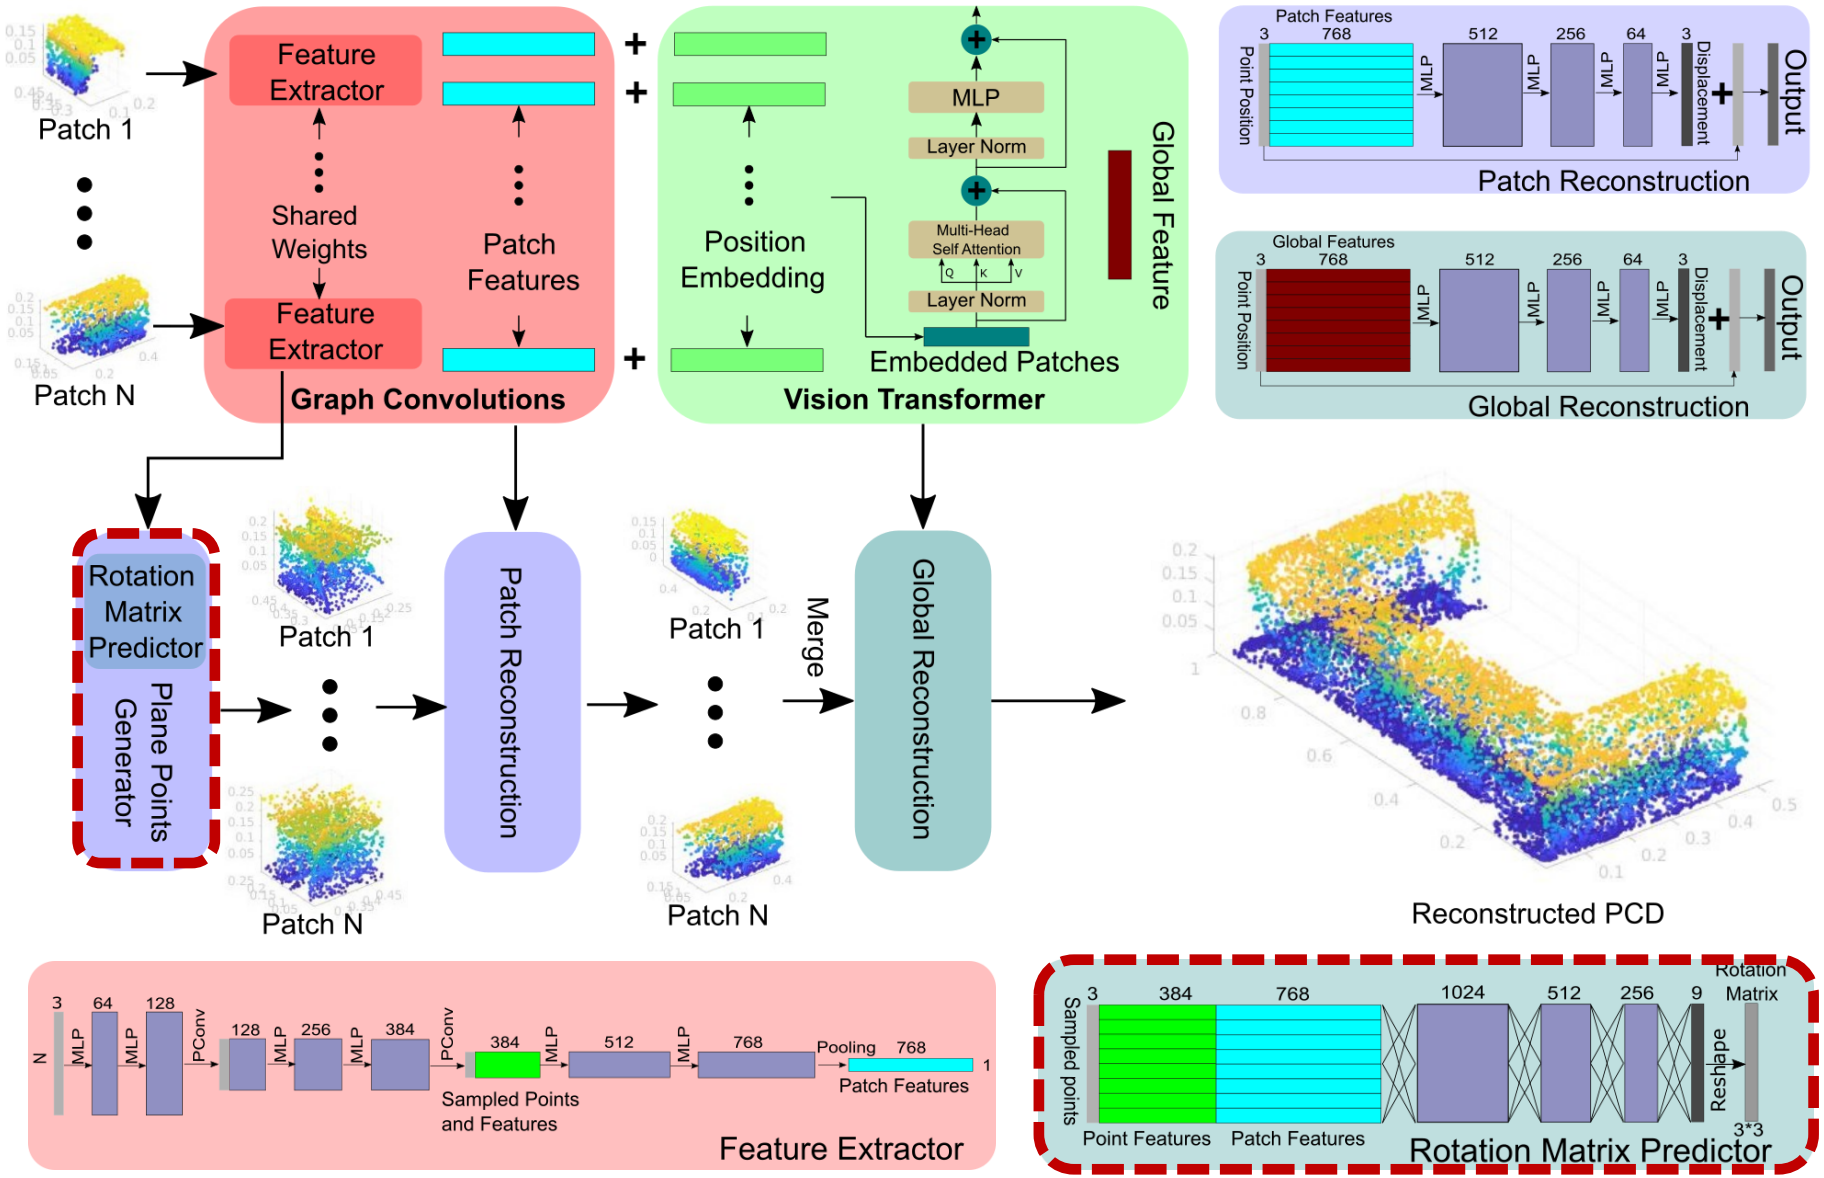

Structure Reconstruction Networks

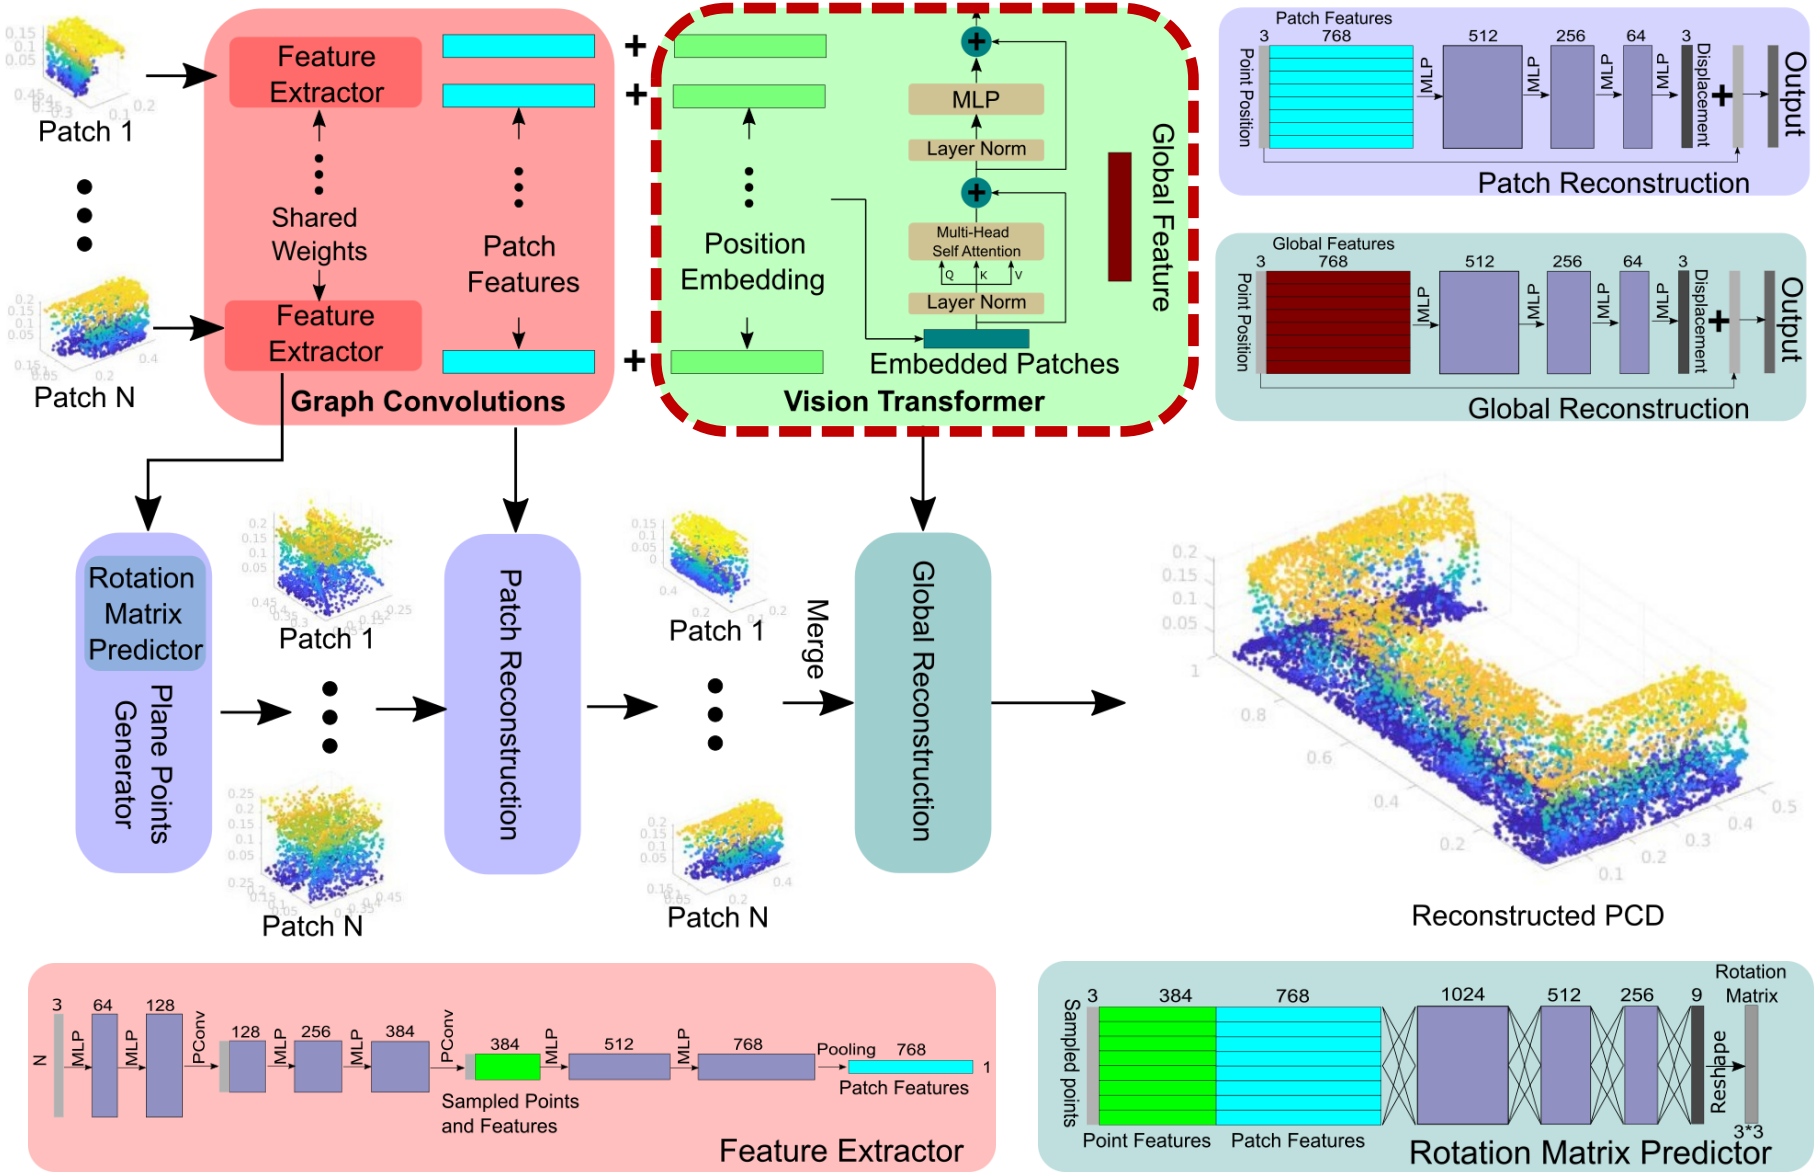

Our Proposal: DeepPCD

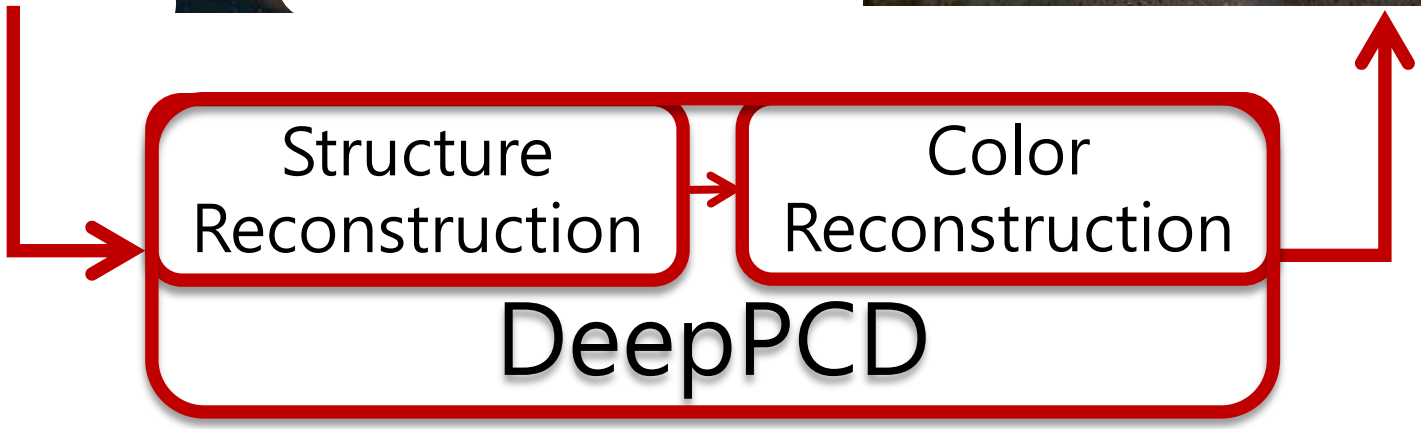

Color Reconstruction: Key Idea

Many large objects from similar environments will likely share similar colors

We could infer the missing colors from similar objects

Challenge: How to find those objects and generalize?

Color Reconstruction Overview

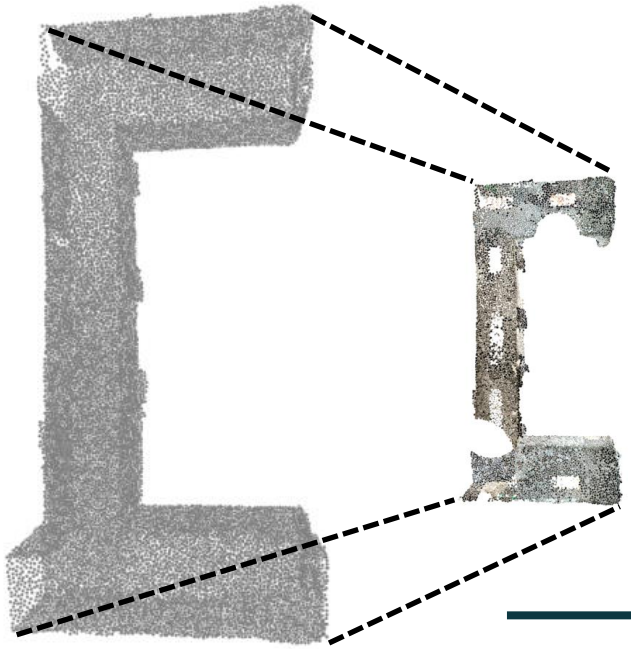

Structurally complete PCD

Input PCD

Color Inference

Conditional Feature Extraction

cGAN based Color Assignment

Fully reconstructed PCD

Color Reconstruction Networks

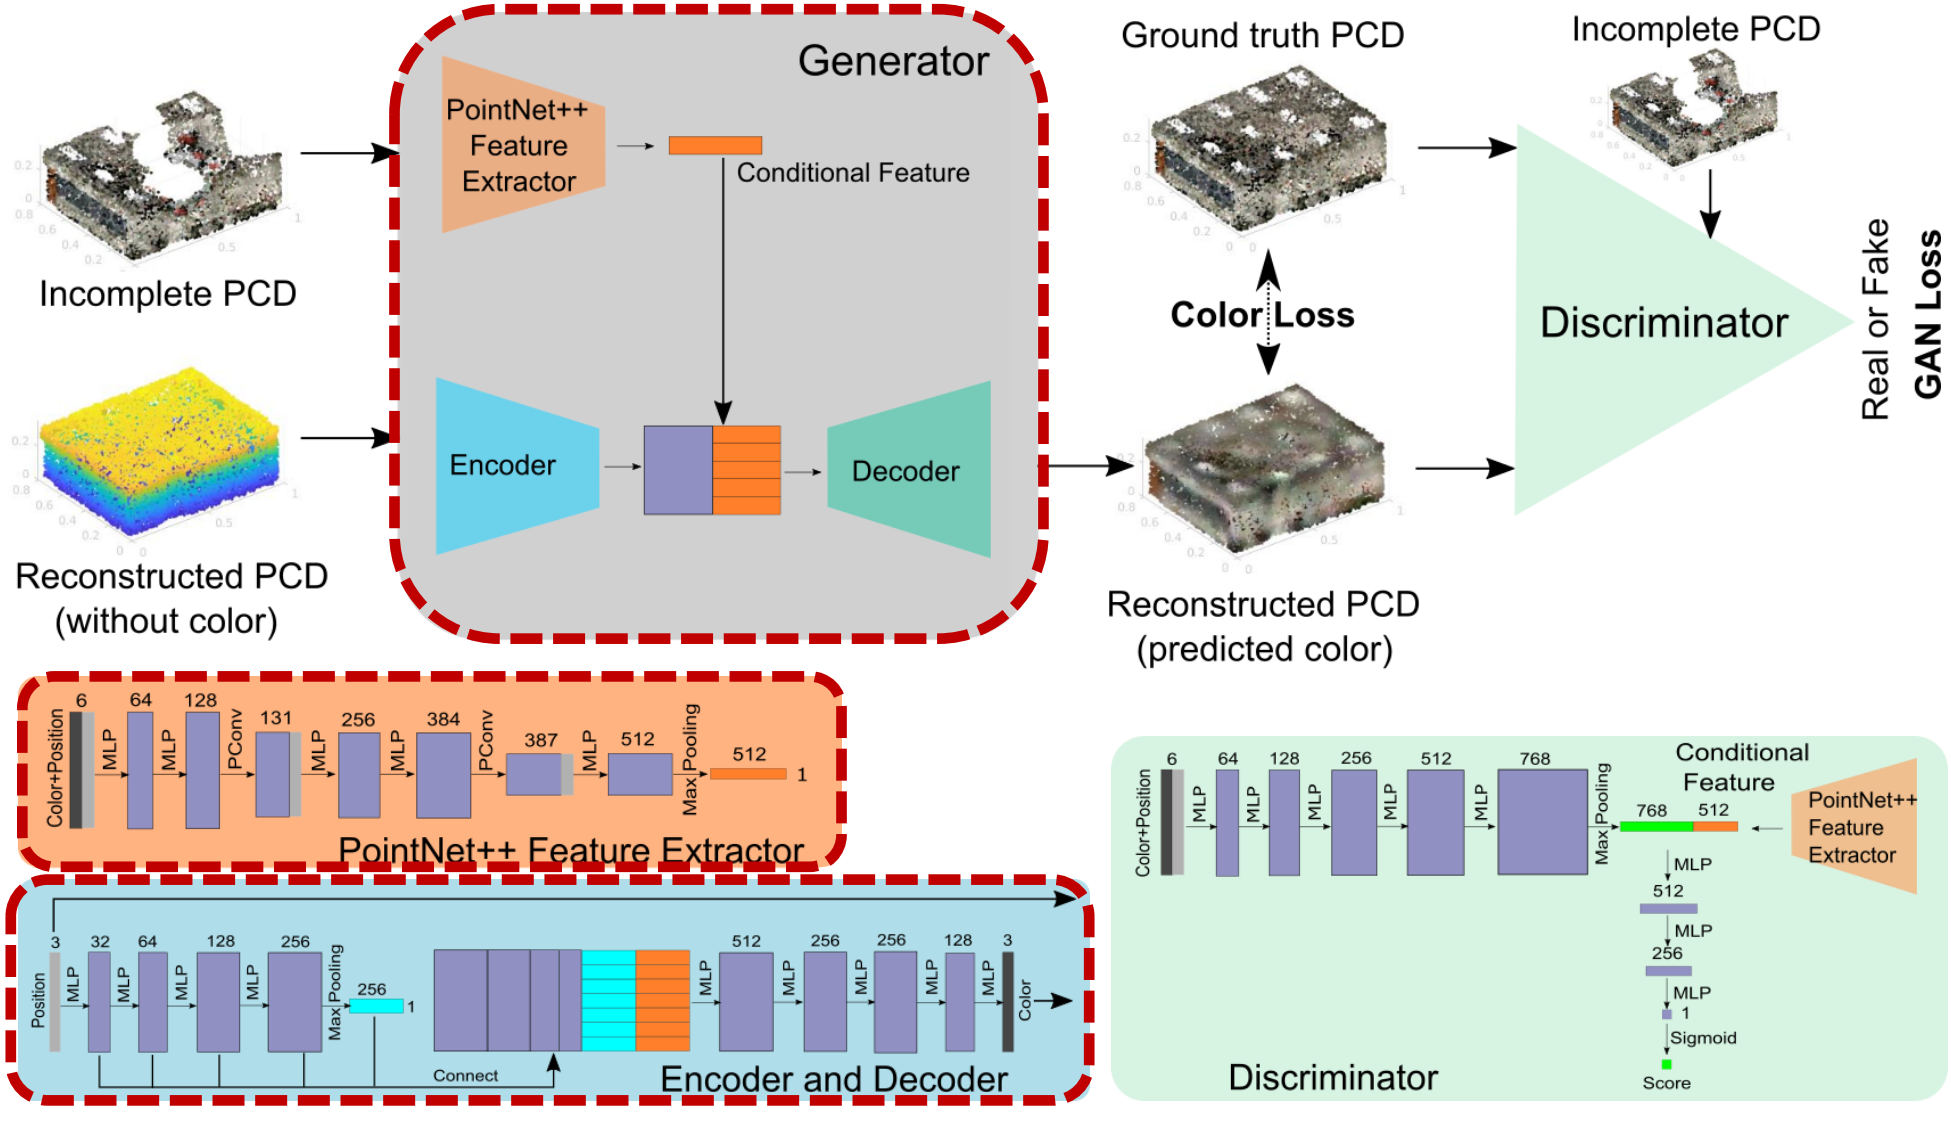

Color Reconstruction Networks

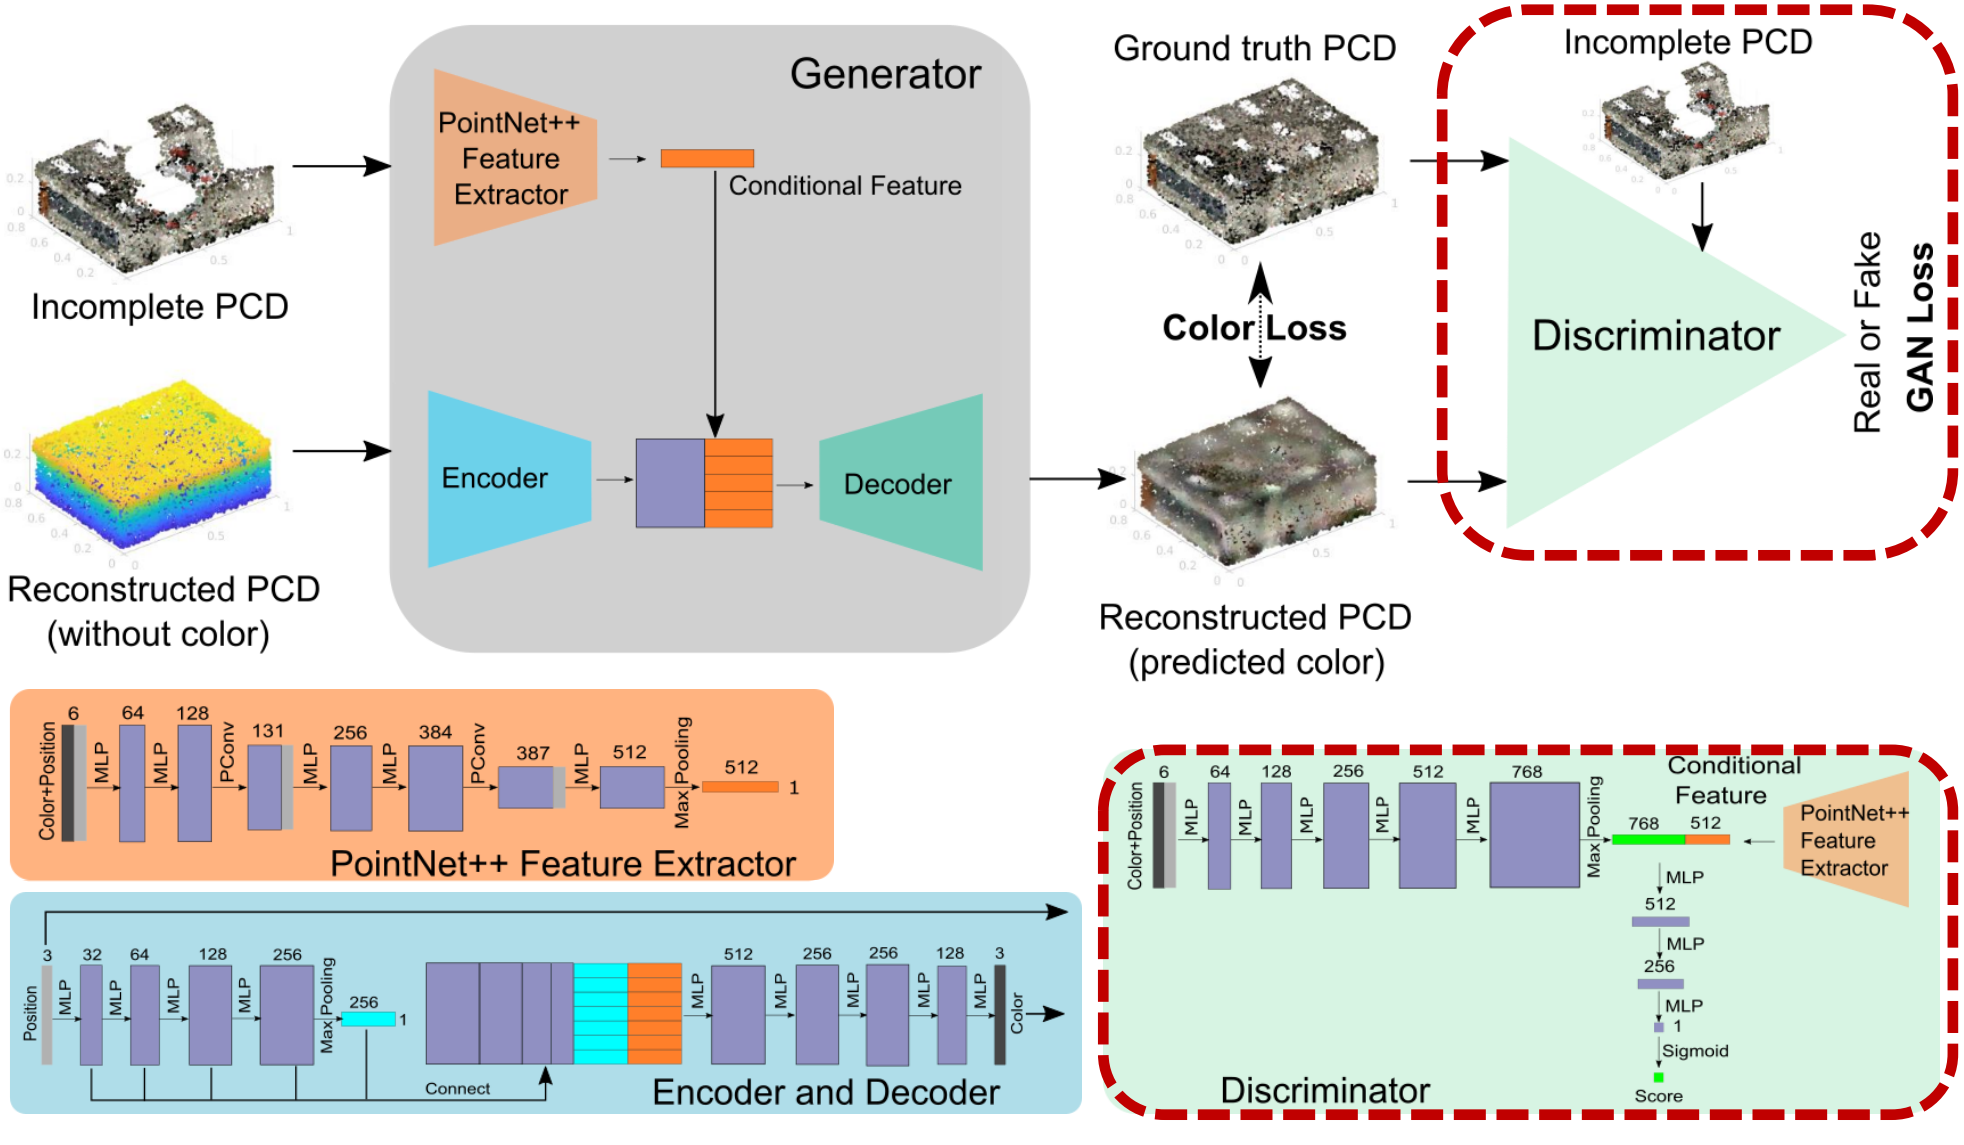

DeepPCD Summary

- DeepPCD is a two-step deep learning based Point Cloud Data (PCD) reconstruction framework
  - It automatically completes an indoor PCD by **reconstructing its structure** and **inferring missing colors**

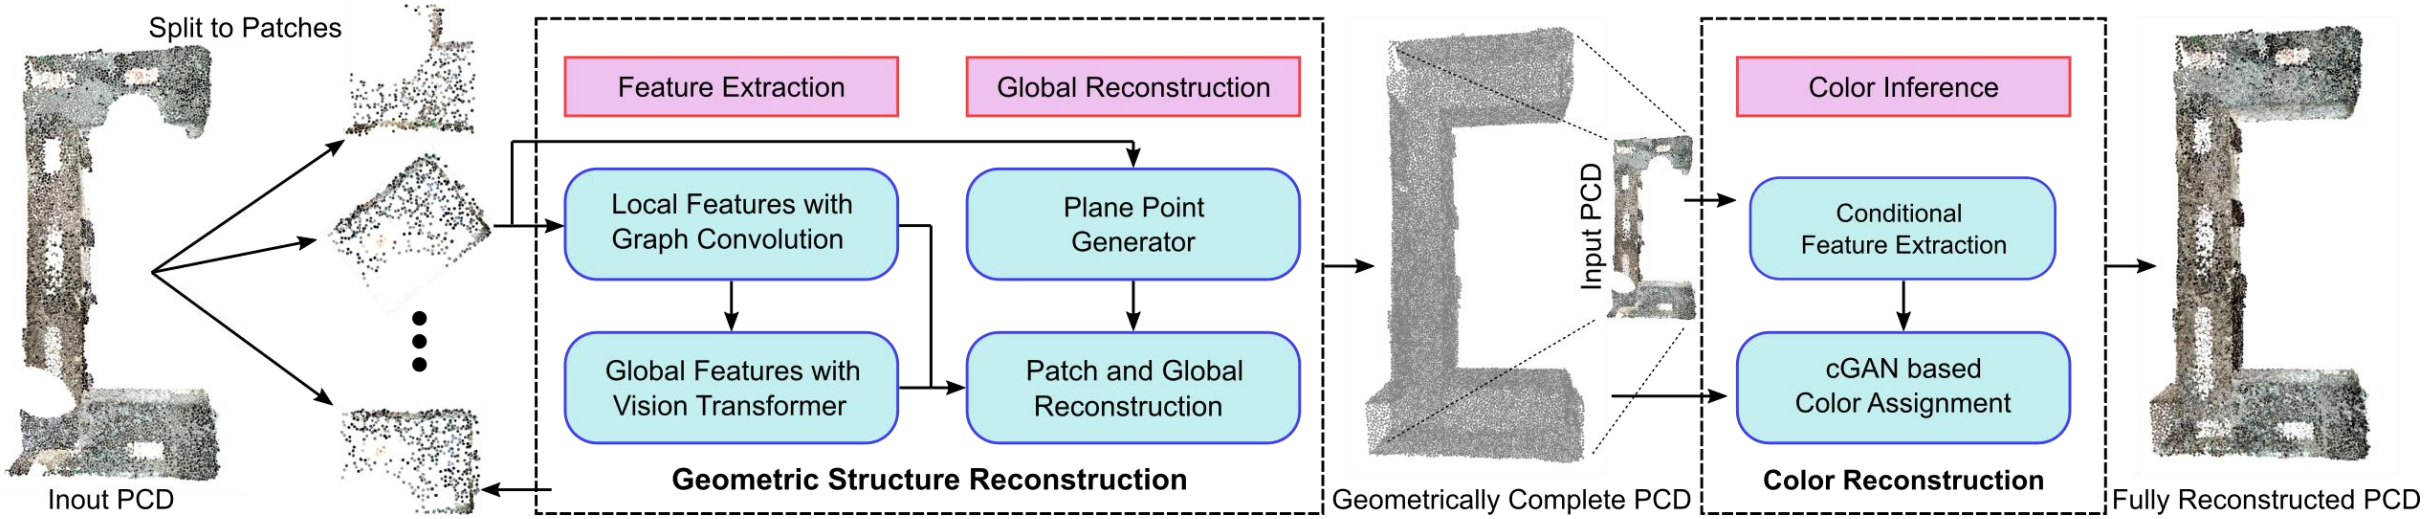

Data Collection and Implementation

- To train and test DeepPCD, we use two large-scale datasets
- Dataset 1: From an AR-capable smartphone, ASUS ZenFone AR
  - PCD from 25 large, diverse indoor environments, across 3 buildings
  - General-purpose hallways, office spaces, lobby area, etc.
  - Each ground truth PCD contains more 3 million points
- Dataset 2: From Stanford Large-Scale 3D Indoor Spaces (S3DIS)[1]
- In total, we have ~3,000 PCD (total ~45 GB)
  - 1,200 samples for training and rest of ~1,800 for testing DeepPCD

1. I. Armeni, et al., "3D Semantic Parsing of Large-Scale Indoor Spaces," in IEEE/CVF CVPR, 2016.

Qualitative Results: Structure

Input PCD

Ground truth

DeepPCD

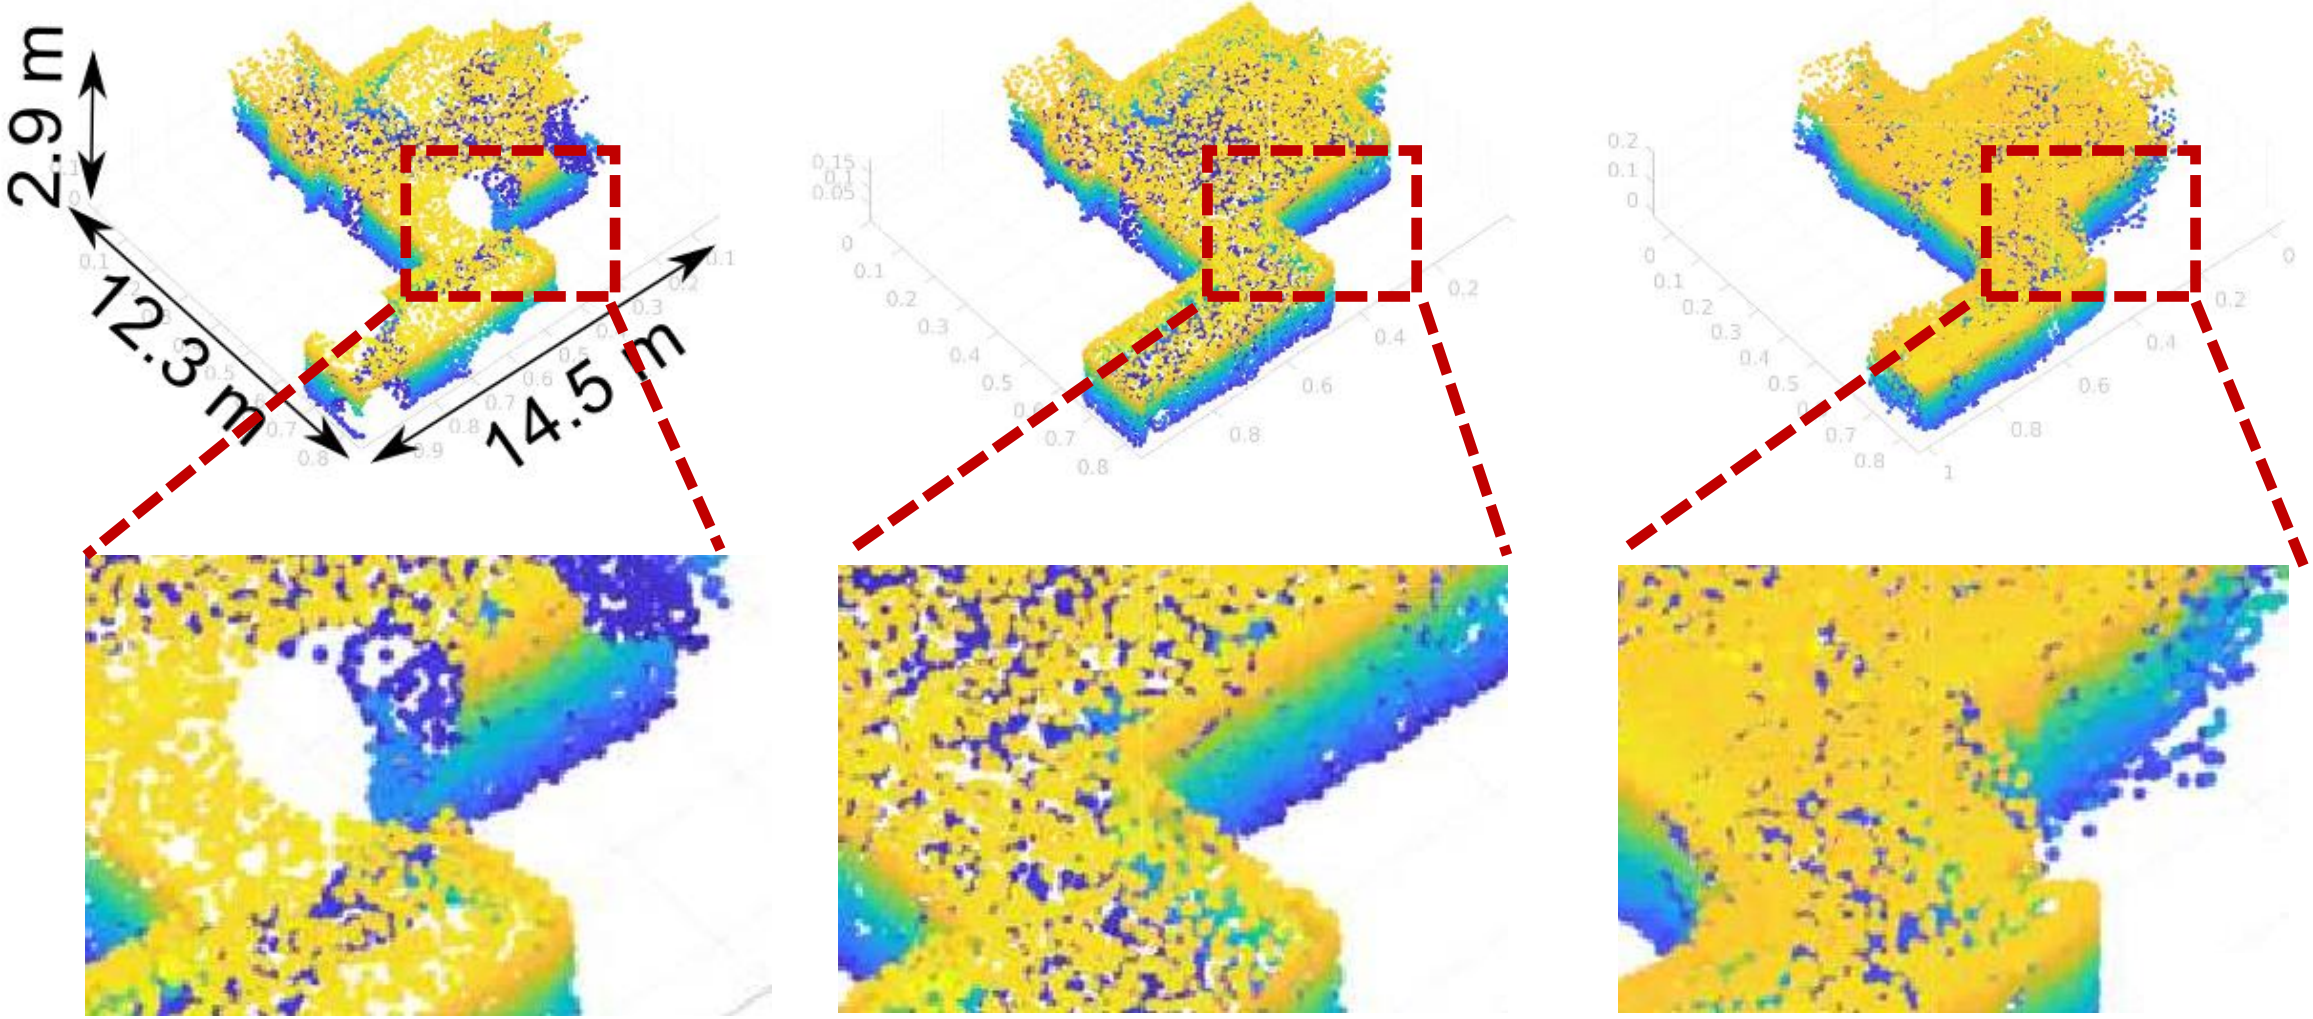

Qualitative Results: Structure

Input PCD

Ground truth

DeepPCD

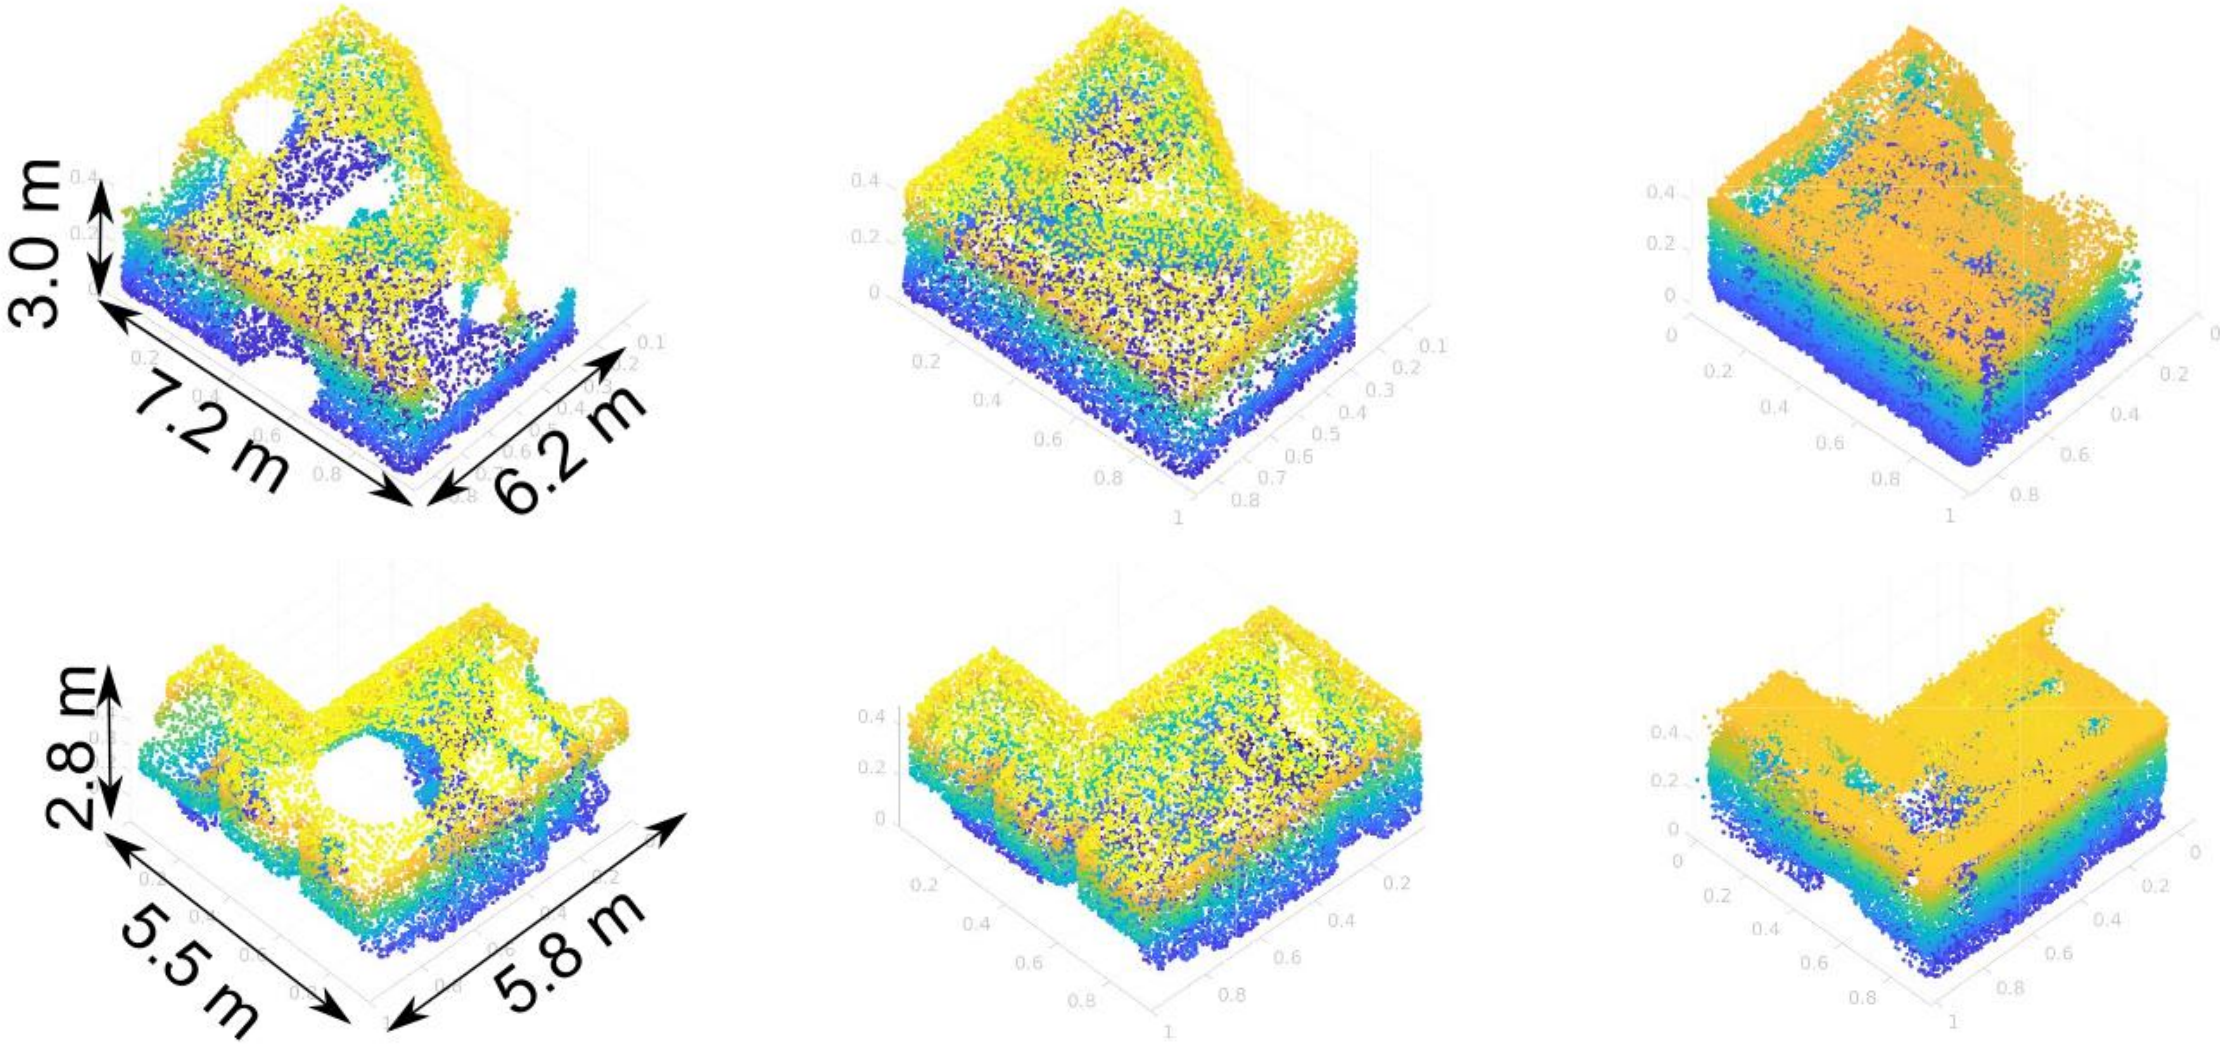

Quantitative Results: Structure

- Chamfer Distance (ChD) measured as the average squared L2-norm distance among two PCD

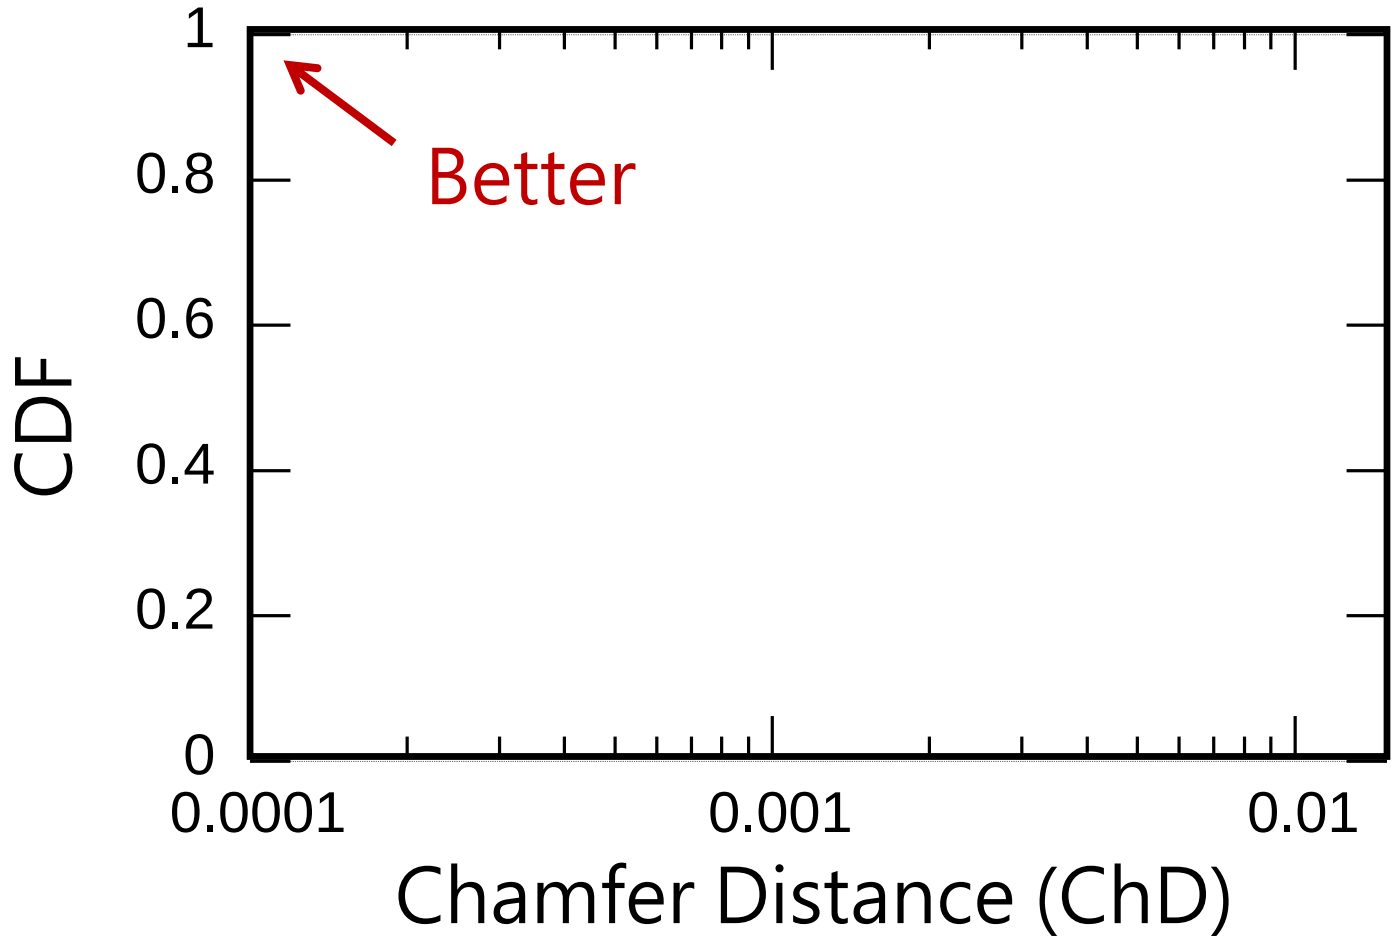

Quantitative Results: Structure

- Chamfer Distance (ChD) measured as the average squared L2-norm distance among two PCD

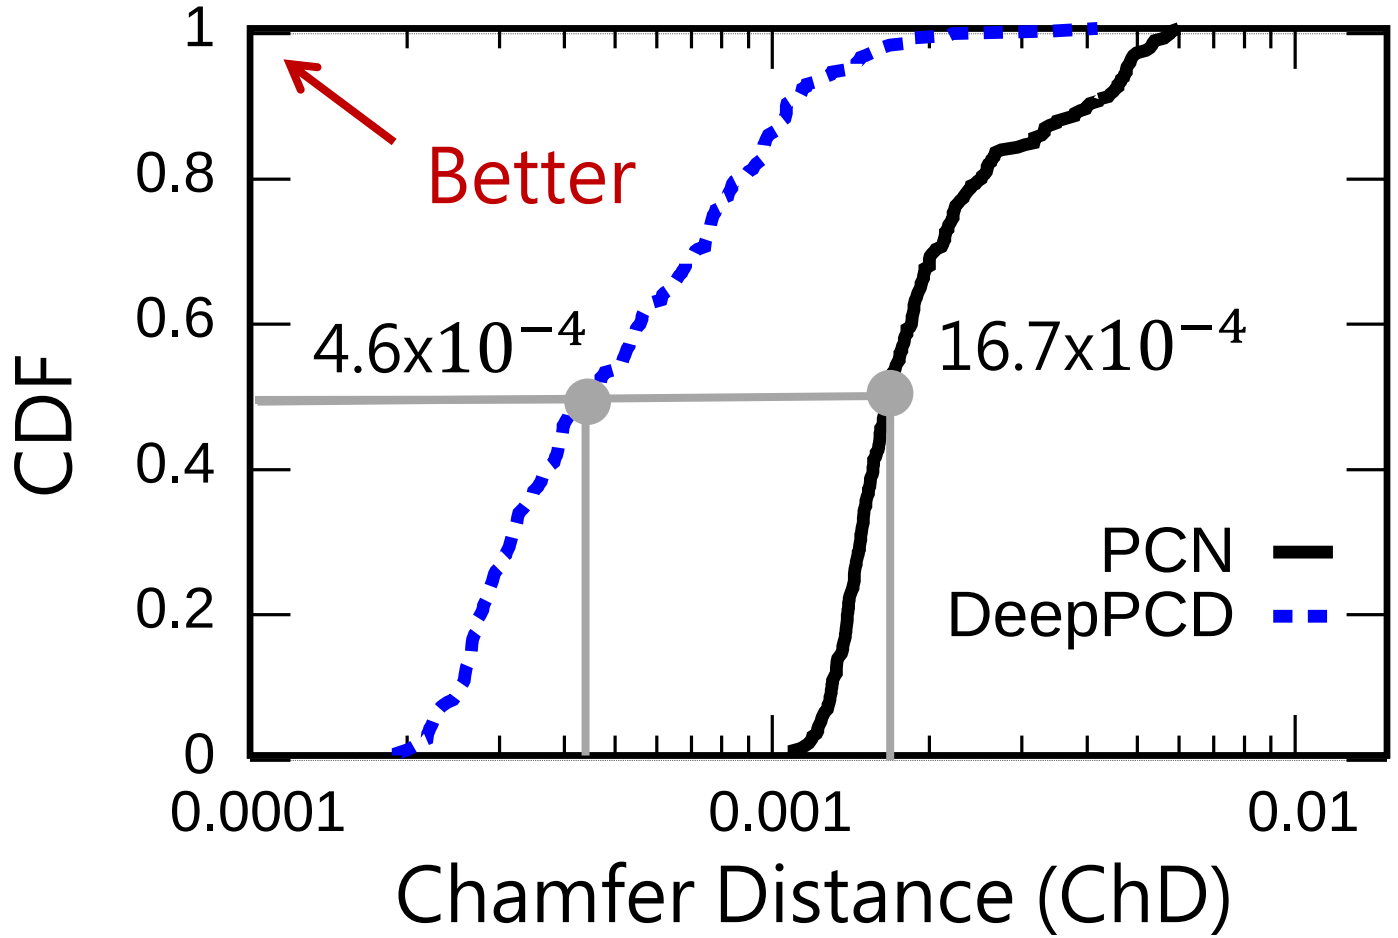

Quantitative Results: Structure

- Chamfer Distance (ChD) measured as the average squared L2-norm distance among two PCD

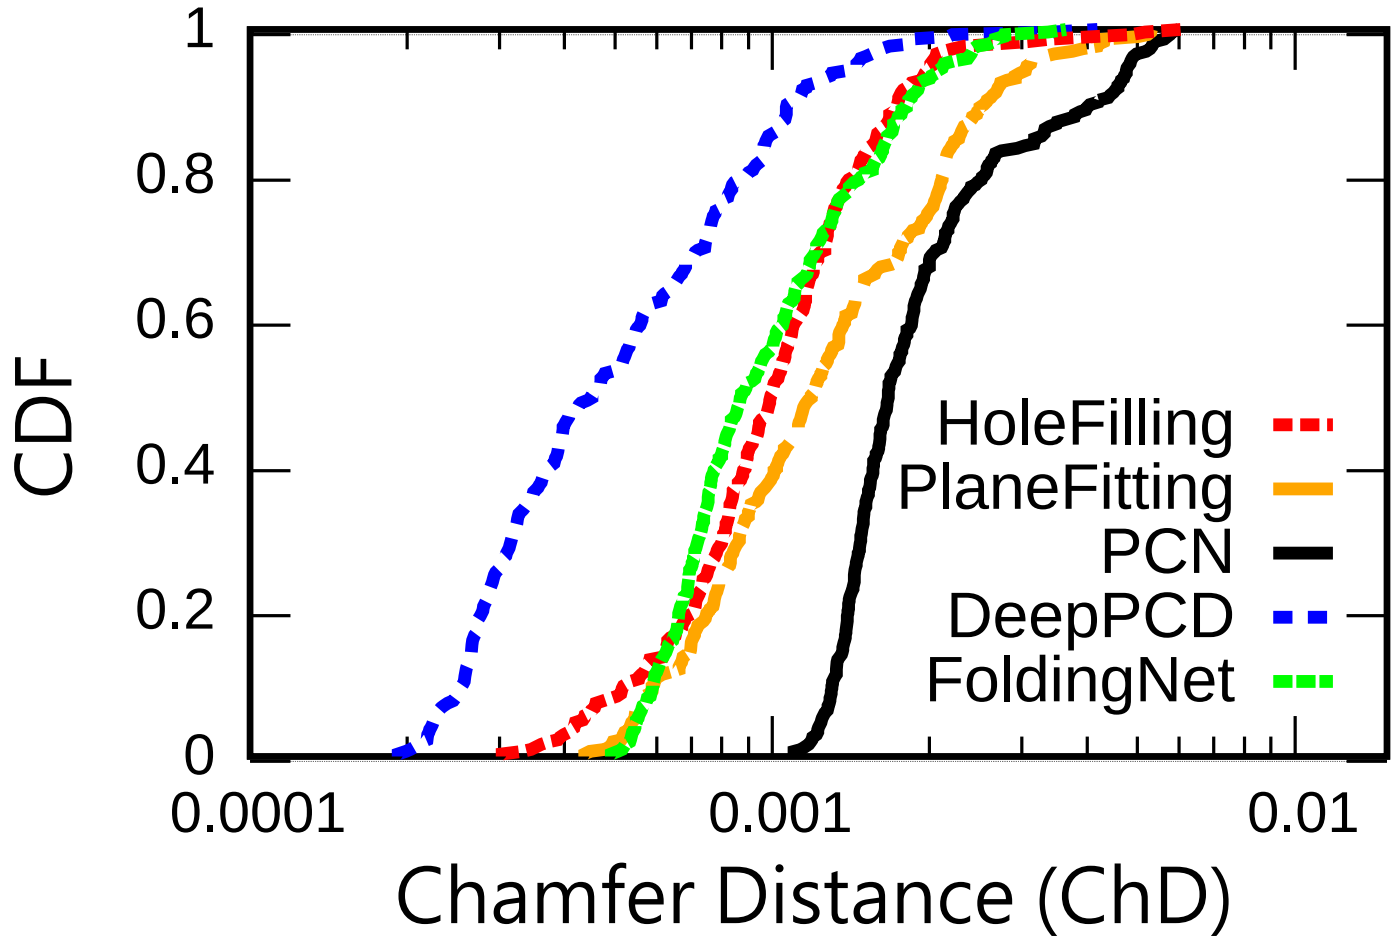

Quantitative Results: Structure

- Chamfer Distance (ChD) measured as the average squared L2-norm distance among two PCD

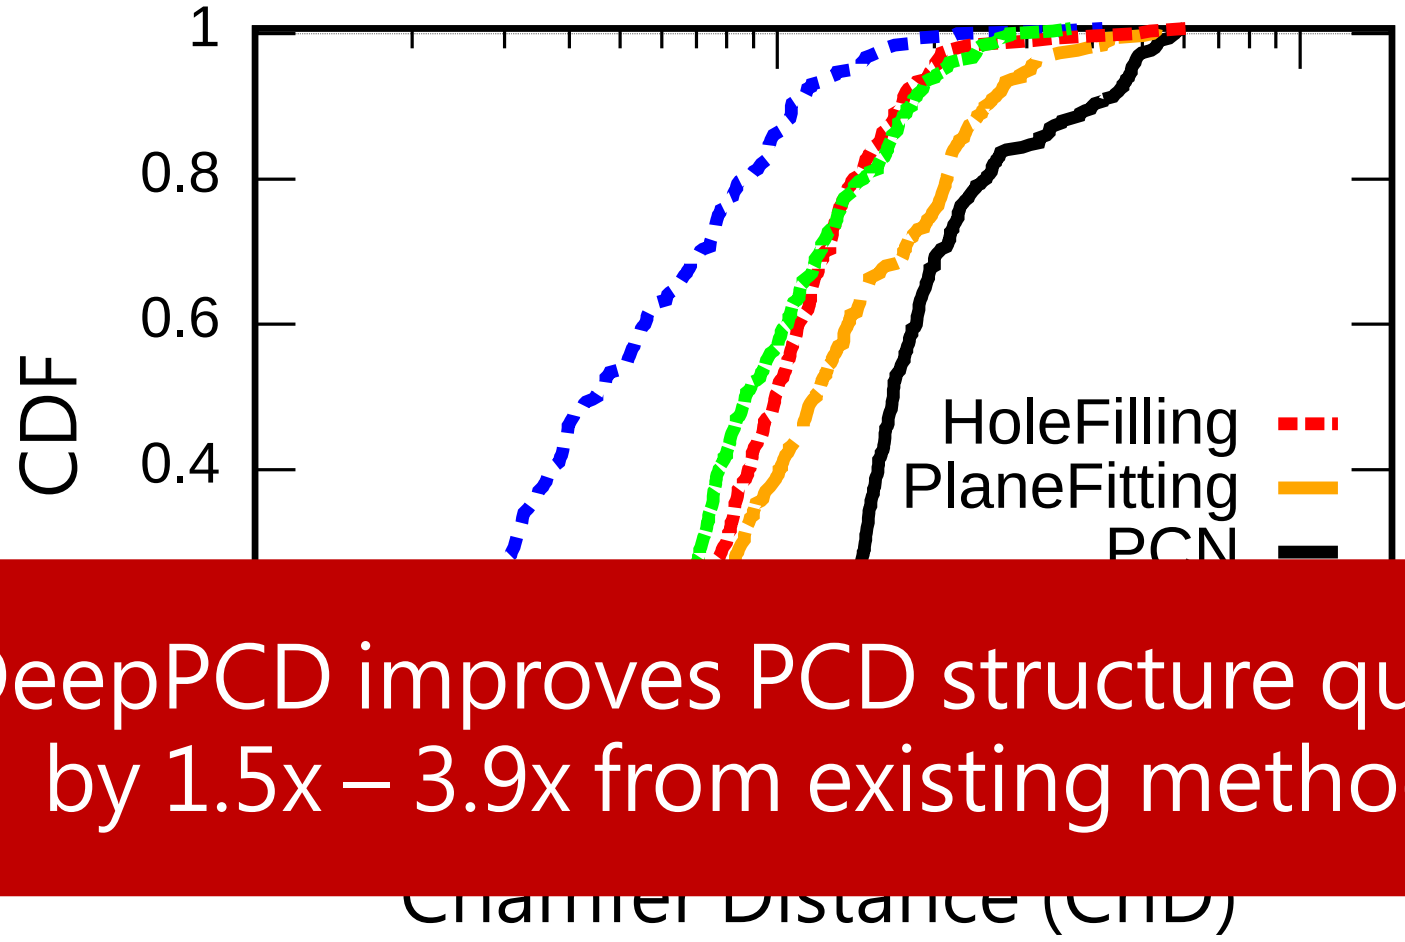

DeepPCD improves PCD structure quality by 1.5x – 3.9x from existing methods

Qualitative Results: Color

Input PCD

Ground truth

DeepPCD

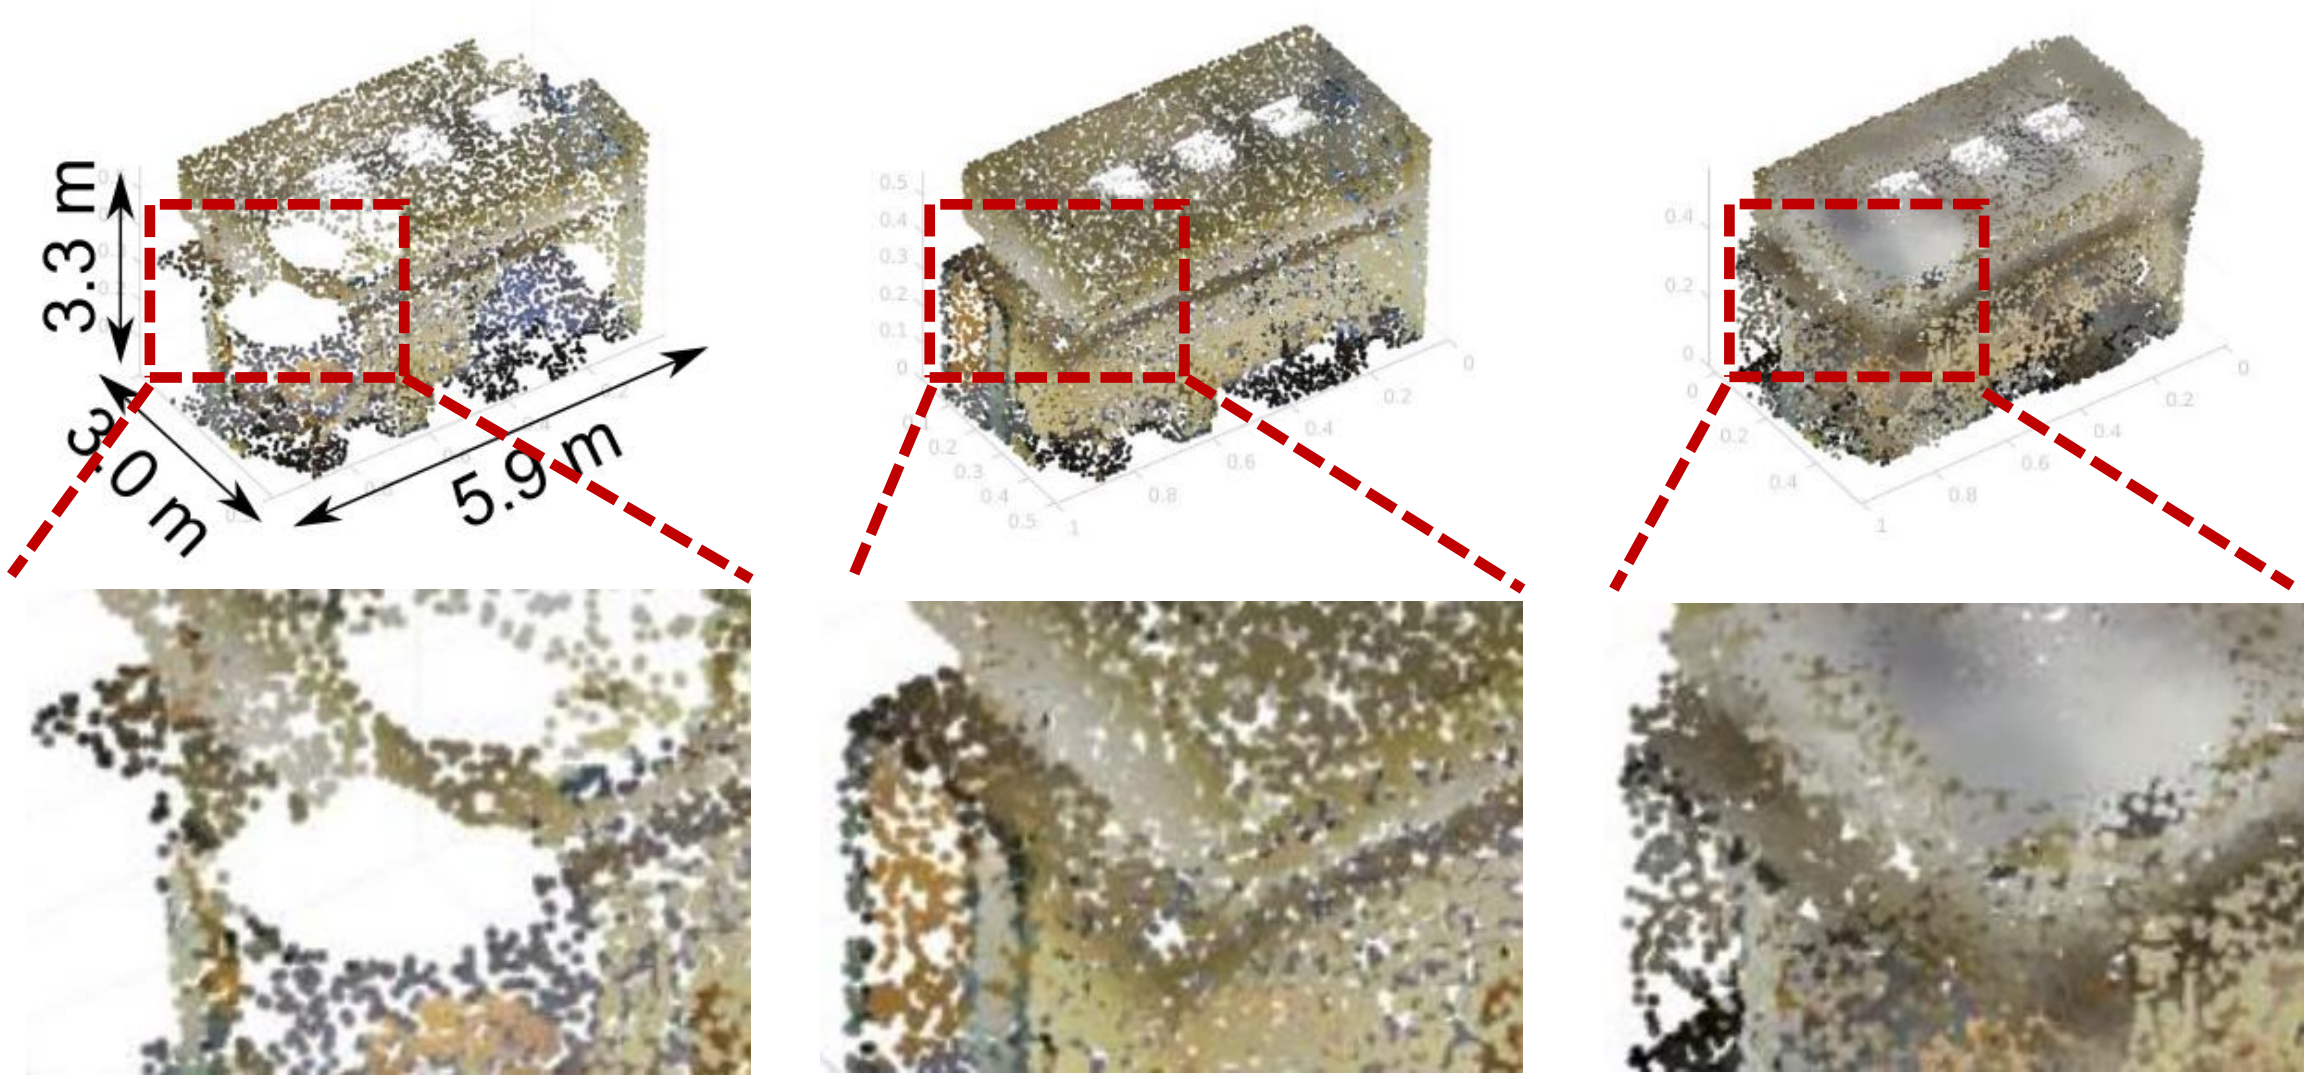

Quantitative Results: Color

- Structural Similarity Index Measure (SSIM): Average color similarity between two PCD by projecting 3D points into 3 isometric planes

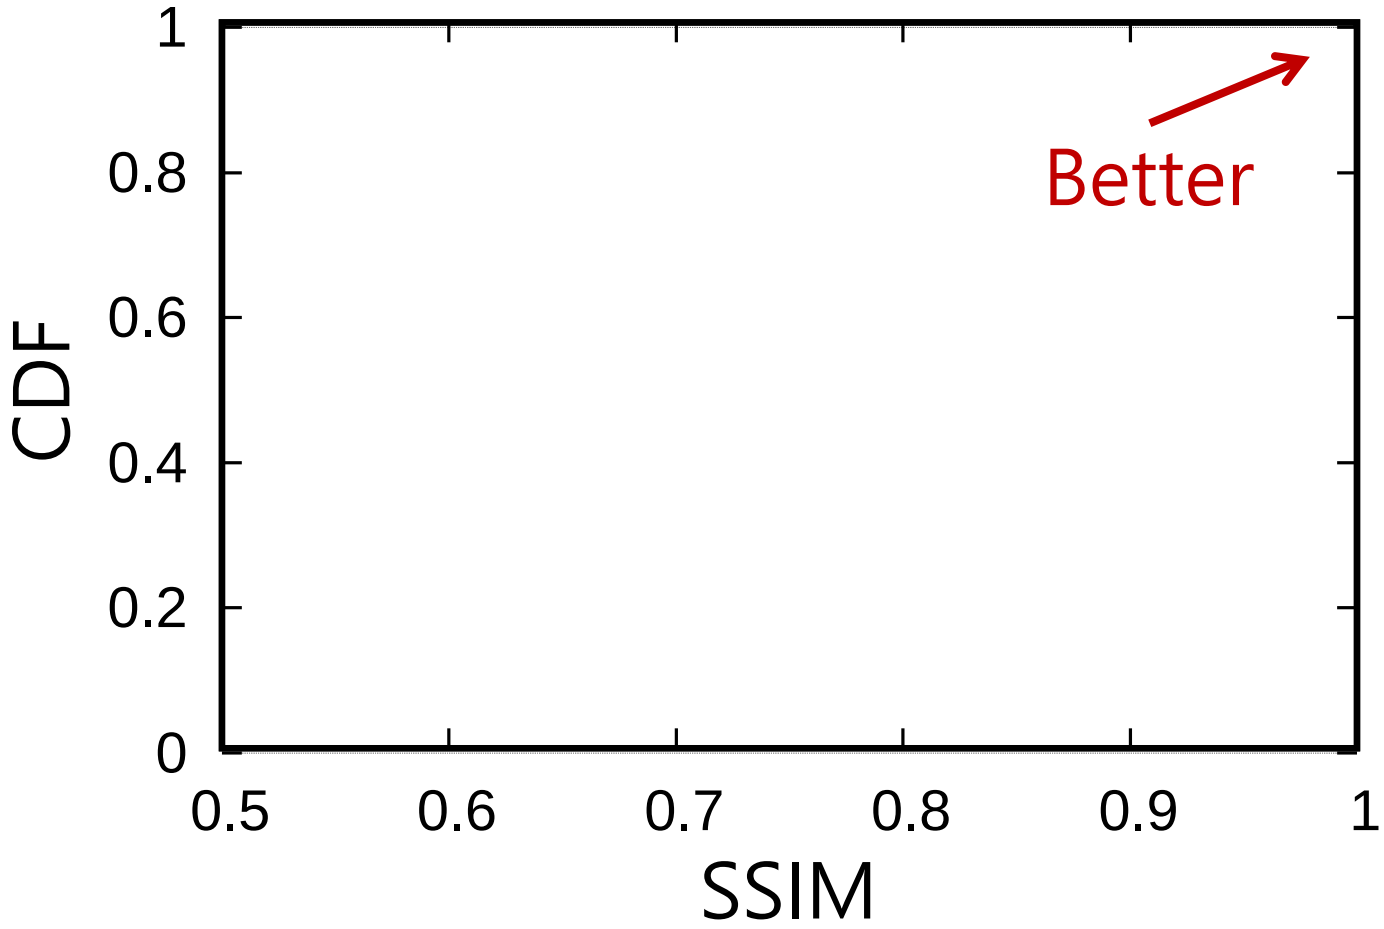

Quantitative Results: Color

- Structural Similarity Index Measure (SSIM): Average color similarity between two PCD by projecting 3D points into 3 isometric planes

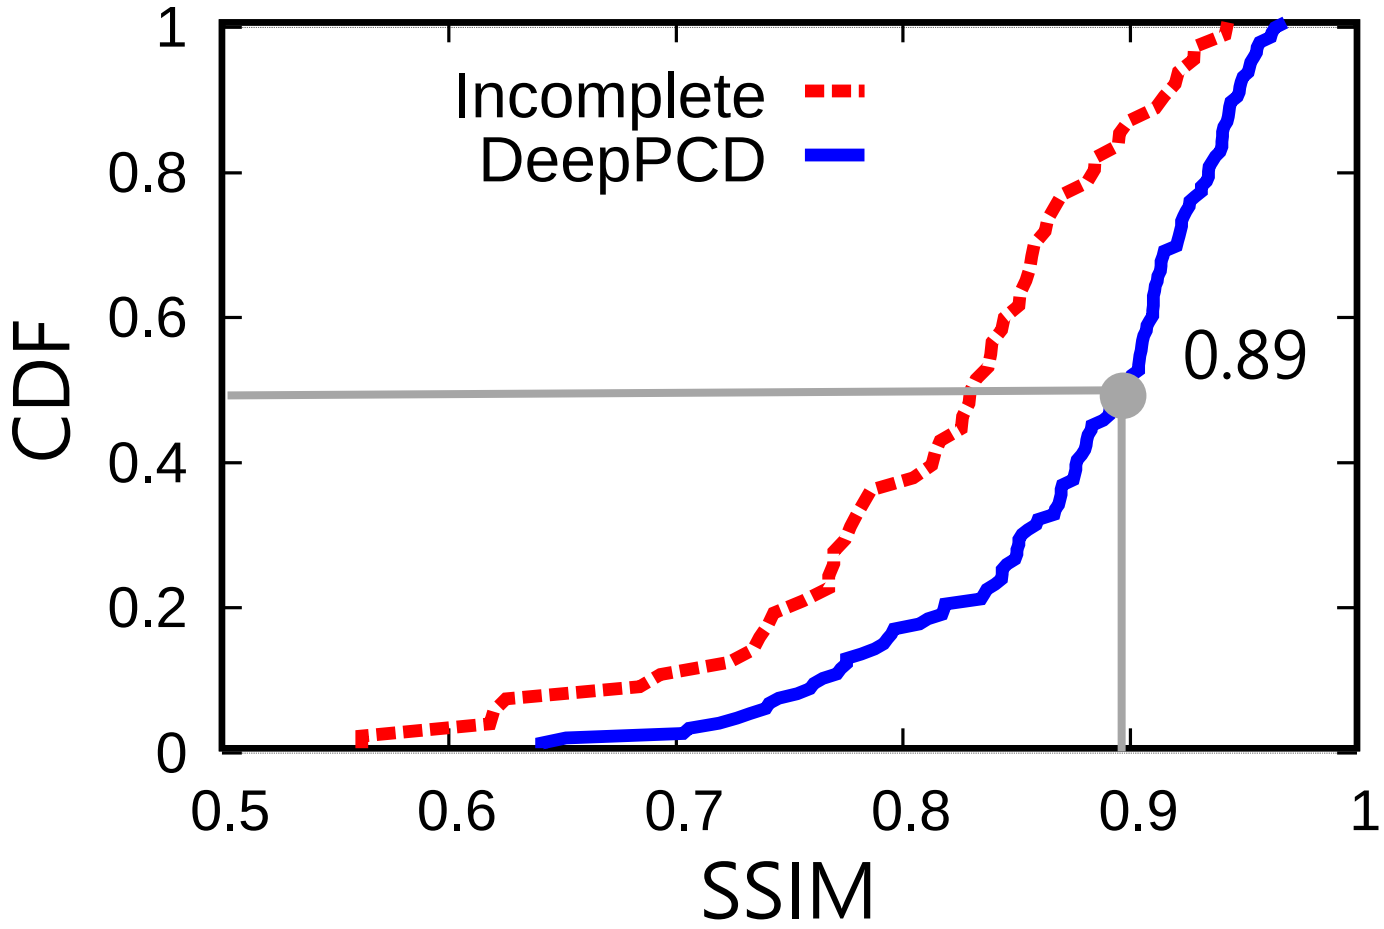

Quantitative Results: Color

- Structural Similarity Index Measure (SSIM): Average color similarity between two PCD by projecting 3D points into 3 isometric planes

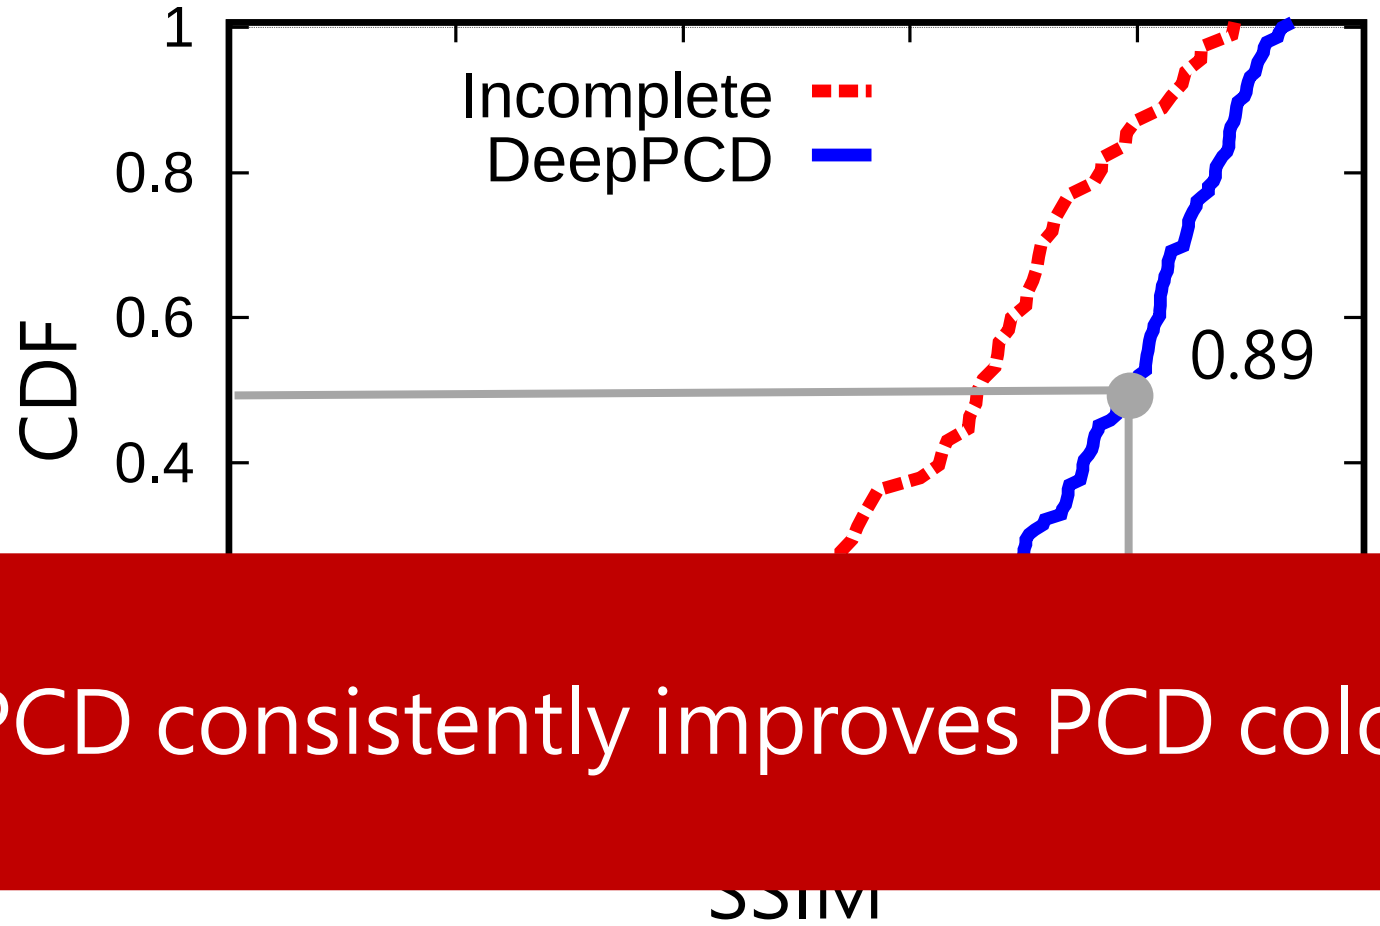

DeepPCD consistently improves PCD color quality

Application Results: Device Navigation

- Performance of vision-based device self-navigation

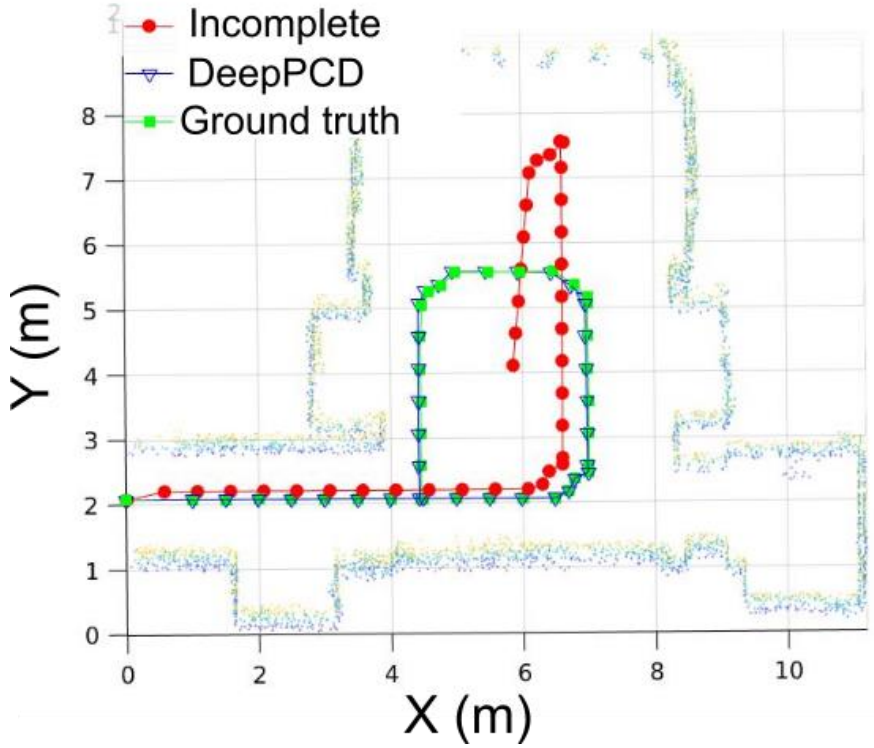

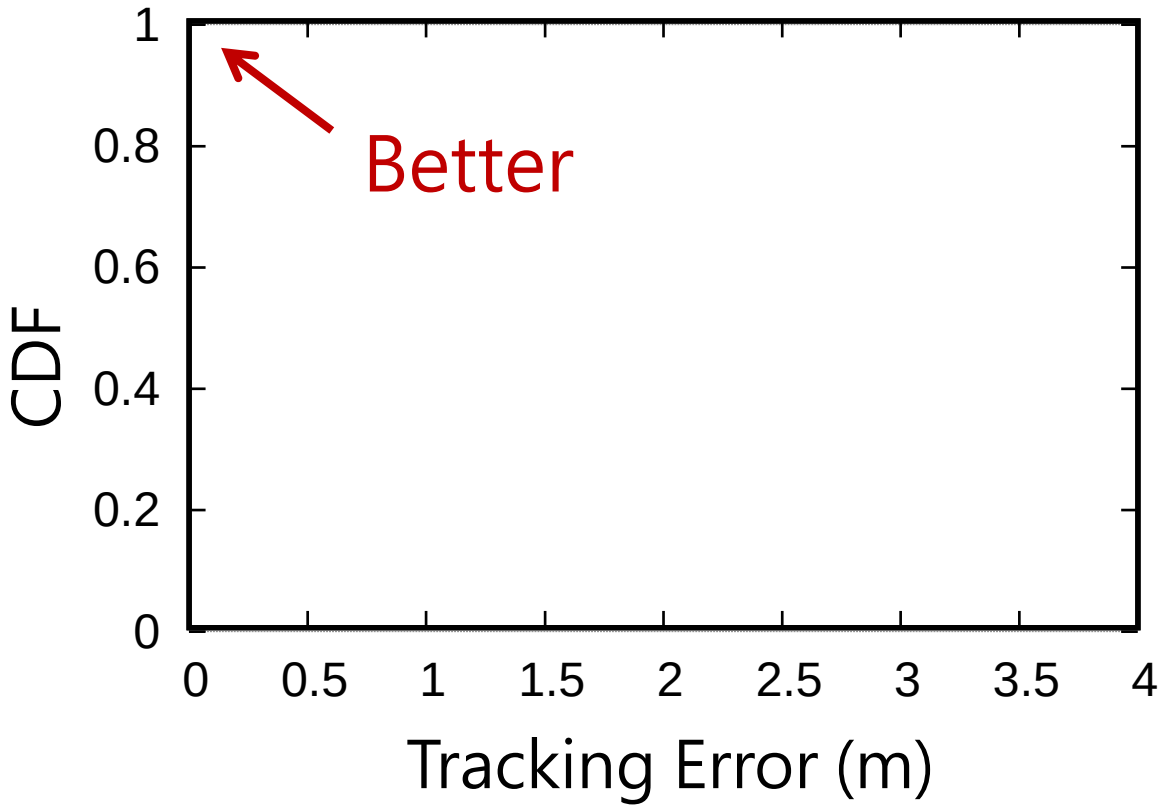

Application Results: Device Navigation

- Performance of vision-based device self-navigation

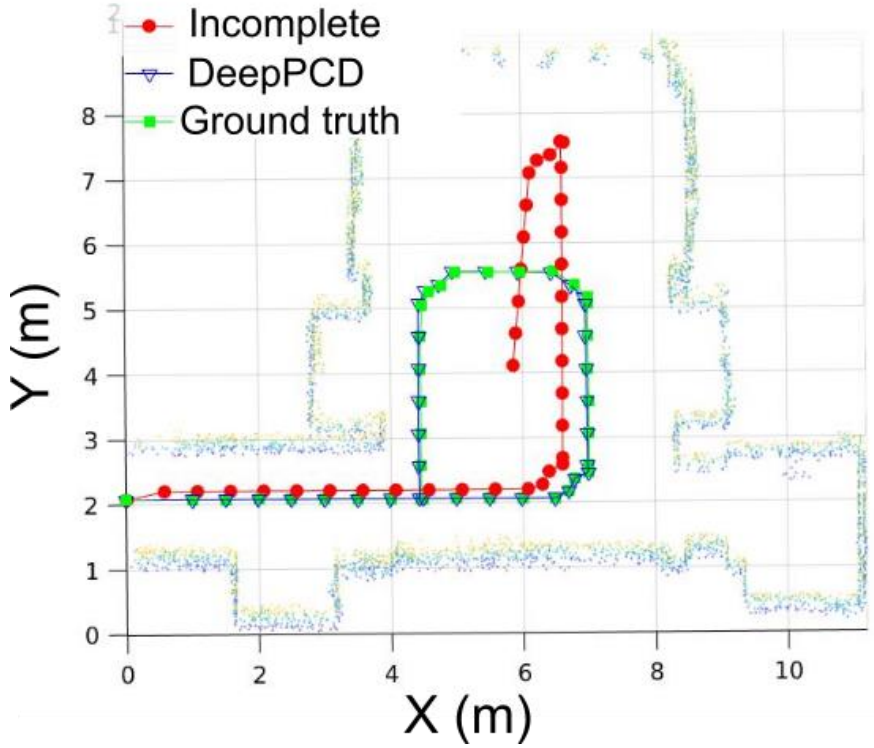

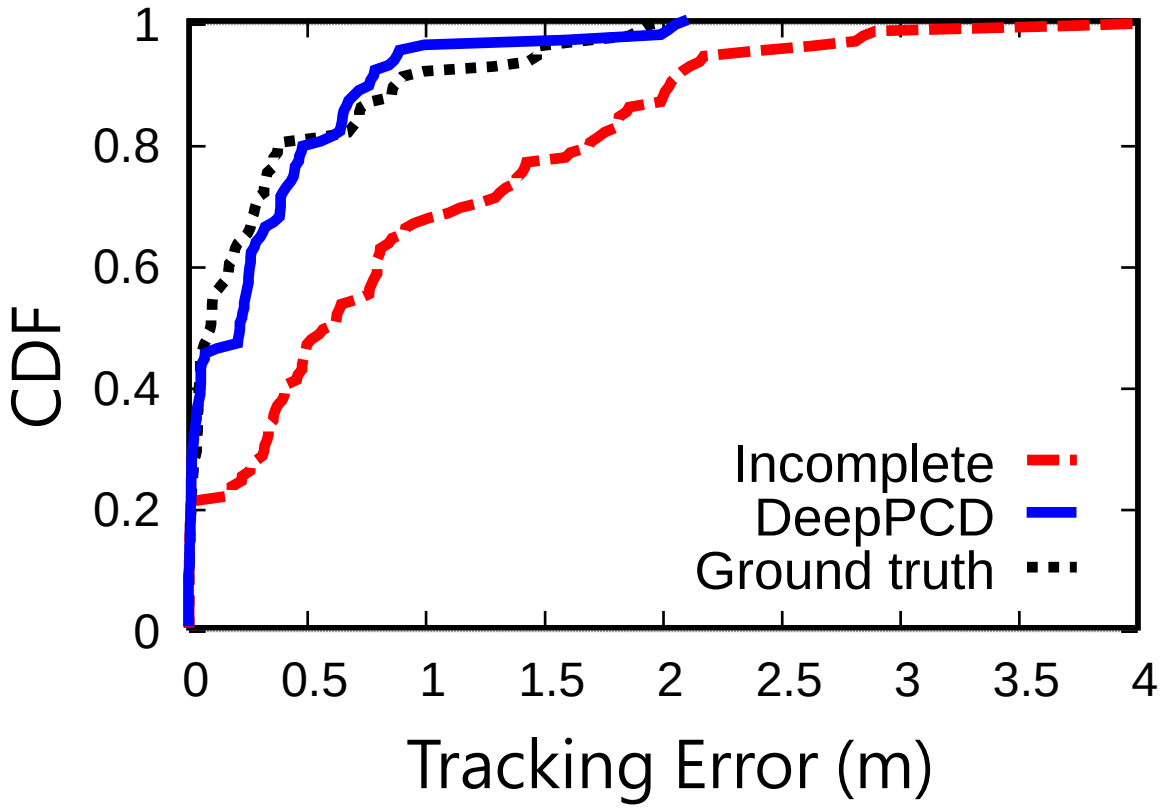

Application Results: Object Tagging

- Manually tagged ~200 objects across all PCD

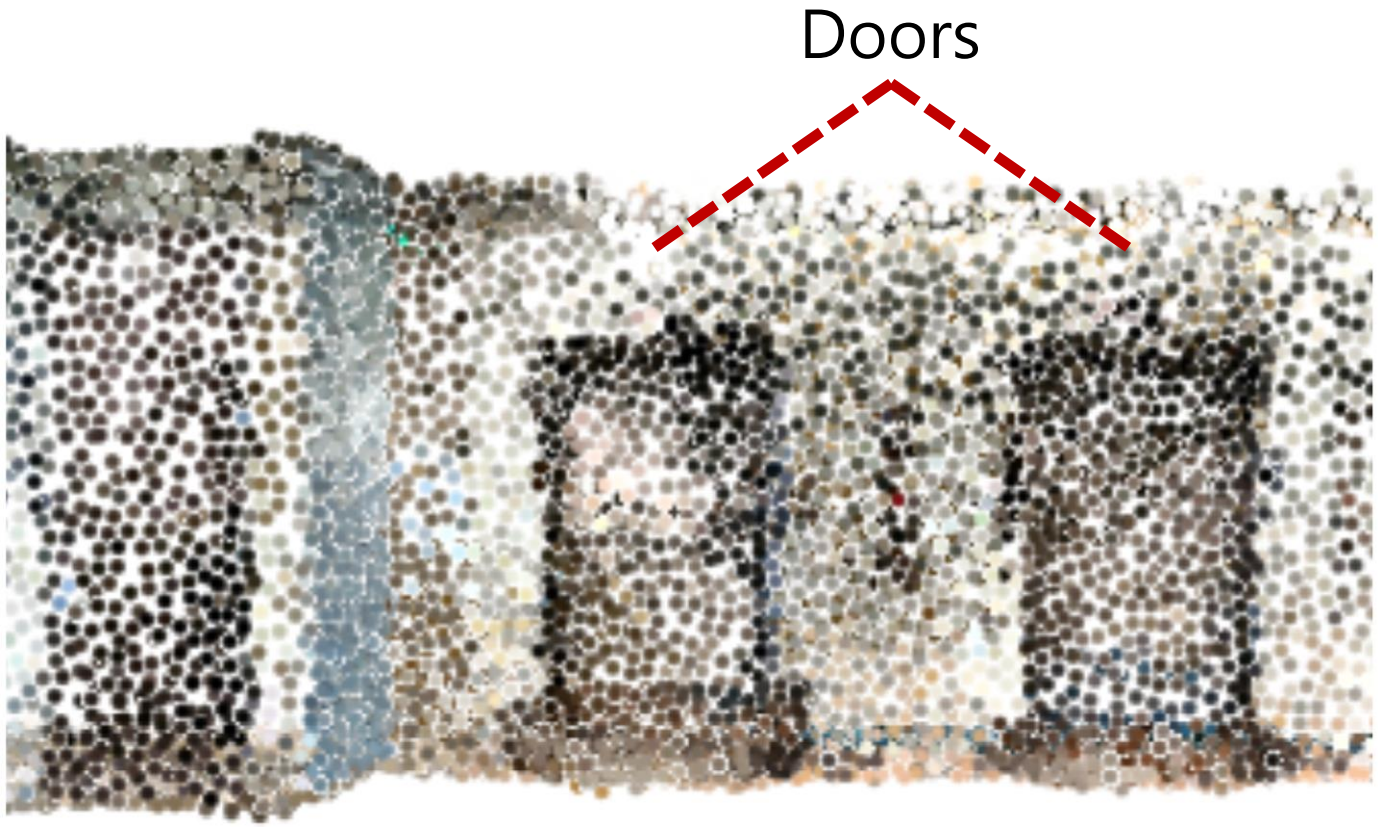

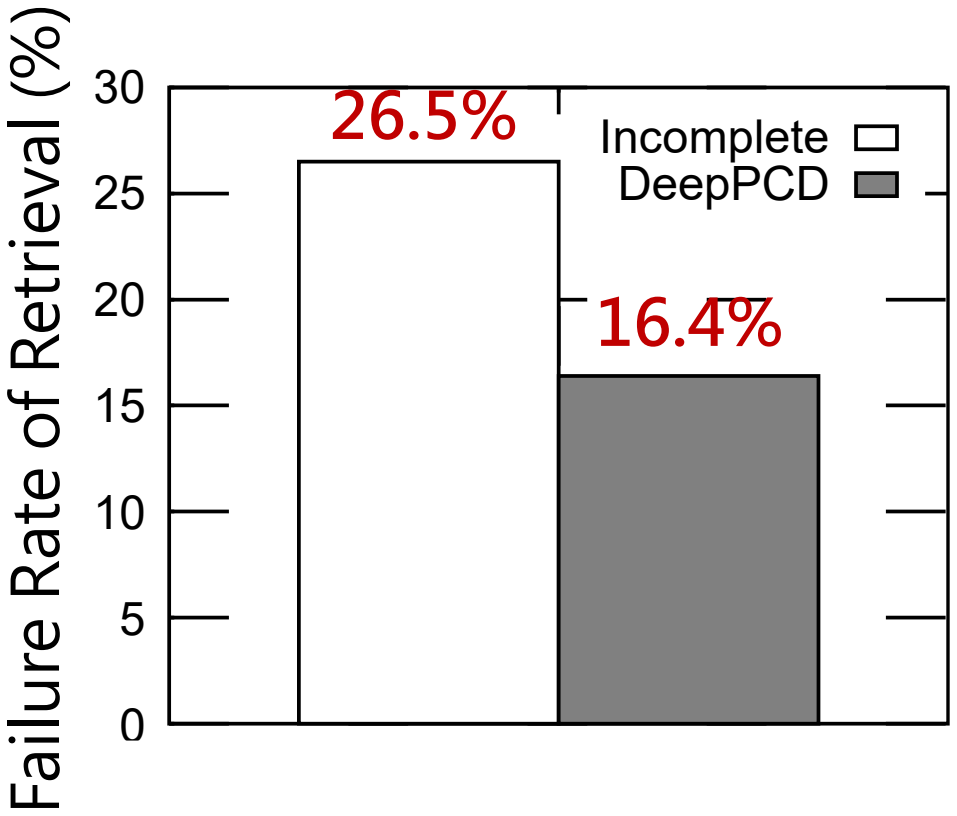

Application Results: Object Tagging

- Manually tagged ~200 objects across all PCD

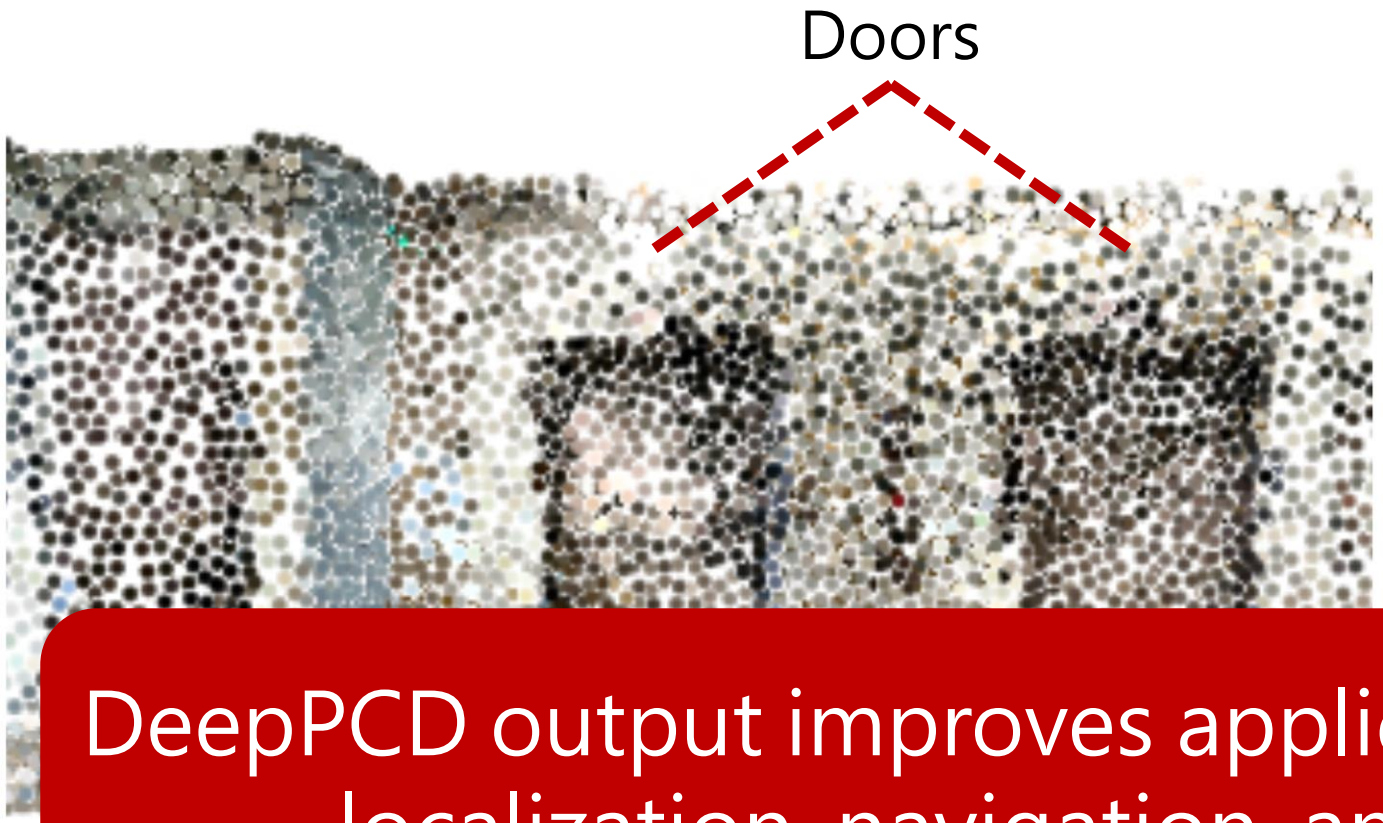

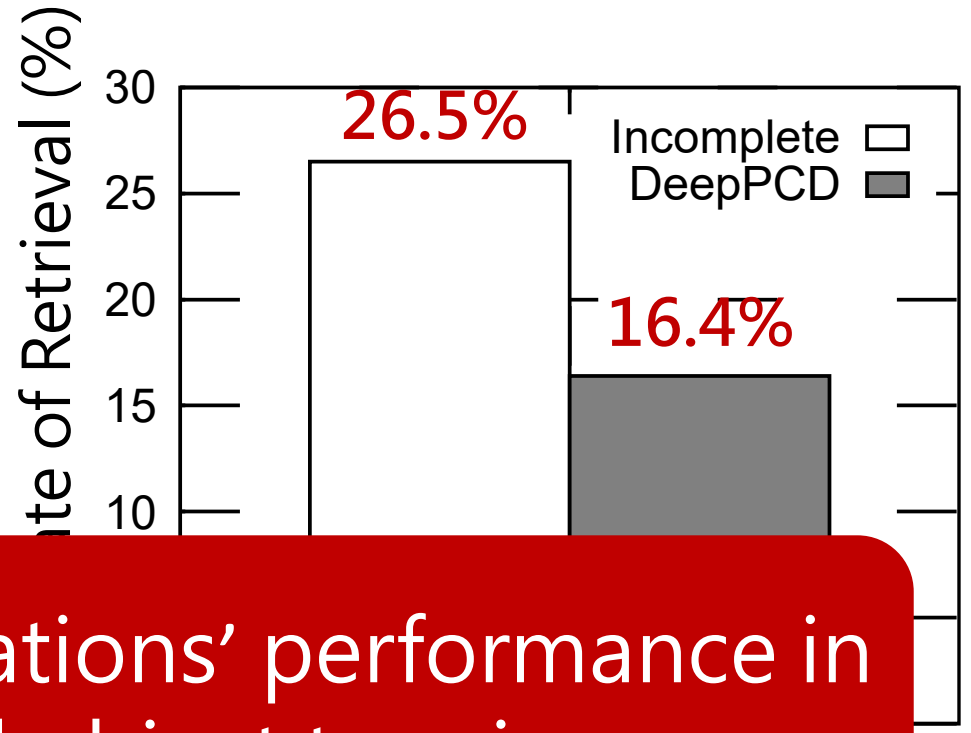

DeepPCD output improves applications' performance in localization, navigation, and object tagging

Conclusion

- DeepPCD is a deep learning based system to reconstruct the missing structure and color information of indoor PCD
- Its performance results show consistent improvement over existing methods, both in quality and quantity, for two large-scale datasets
- Its output further improves the performance of many vision-based ubiquitous sensing applications

Thank you!

Check out our project website for more results, code, and dataset

Contact:
sur@cse.sc.edu