Generalizable Unpaired Point Cloud Completion

Mihir Garimella, Prathik Naidu

{mihirg, prathikn}@stanford.edu

Introduction

- Point clouds are a popular and versatile way to represent 3D objects
- However, raw 3D sensor data is often sparse and noisy due to occlusion, range limitations, and sensor failures, making point clouds difficult to process accurately in real-world applications
- Current methods for point cloud completion work well on synthetic data but fail on real-world data because they require *paired* datasets
- **We use a GAN for point cloud completion on unpaired data across several types of objects**

Related Work

- We build on the work of **Chen et al. (2019)**, which uses two autoencoders and a GAN to map partial to complete point clouds given unpaired data
  - **Key limitation:** not actually useful in practice because they have to train a separate network for each type of object → **we solve this**

Dataset

- **ShapeNet** (Chang et al., 2015) contains 51,300 CAD models of 55 types of common objects
- In **PCN** (Yuan et al., 2018), the authors simulate occlusion of ShapeNet models to generate partial point clouds of 28,974 objects across 8 classes
- We use PCN dataset; fully synthetic but unpaired, good starting point for multi-class completion

Technical Approach

Train AEs on distributions of partial and complete point clouds

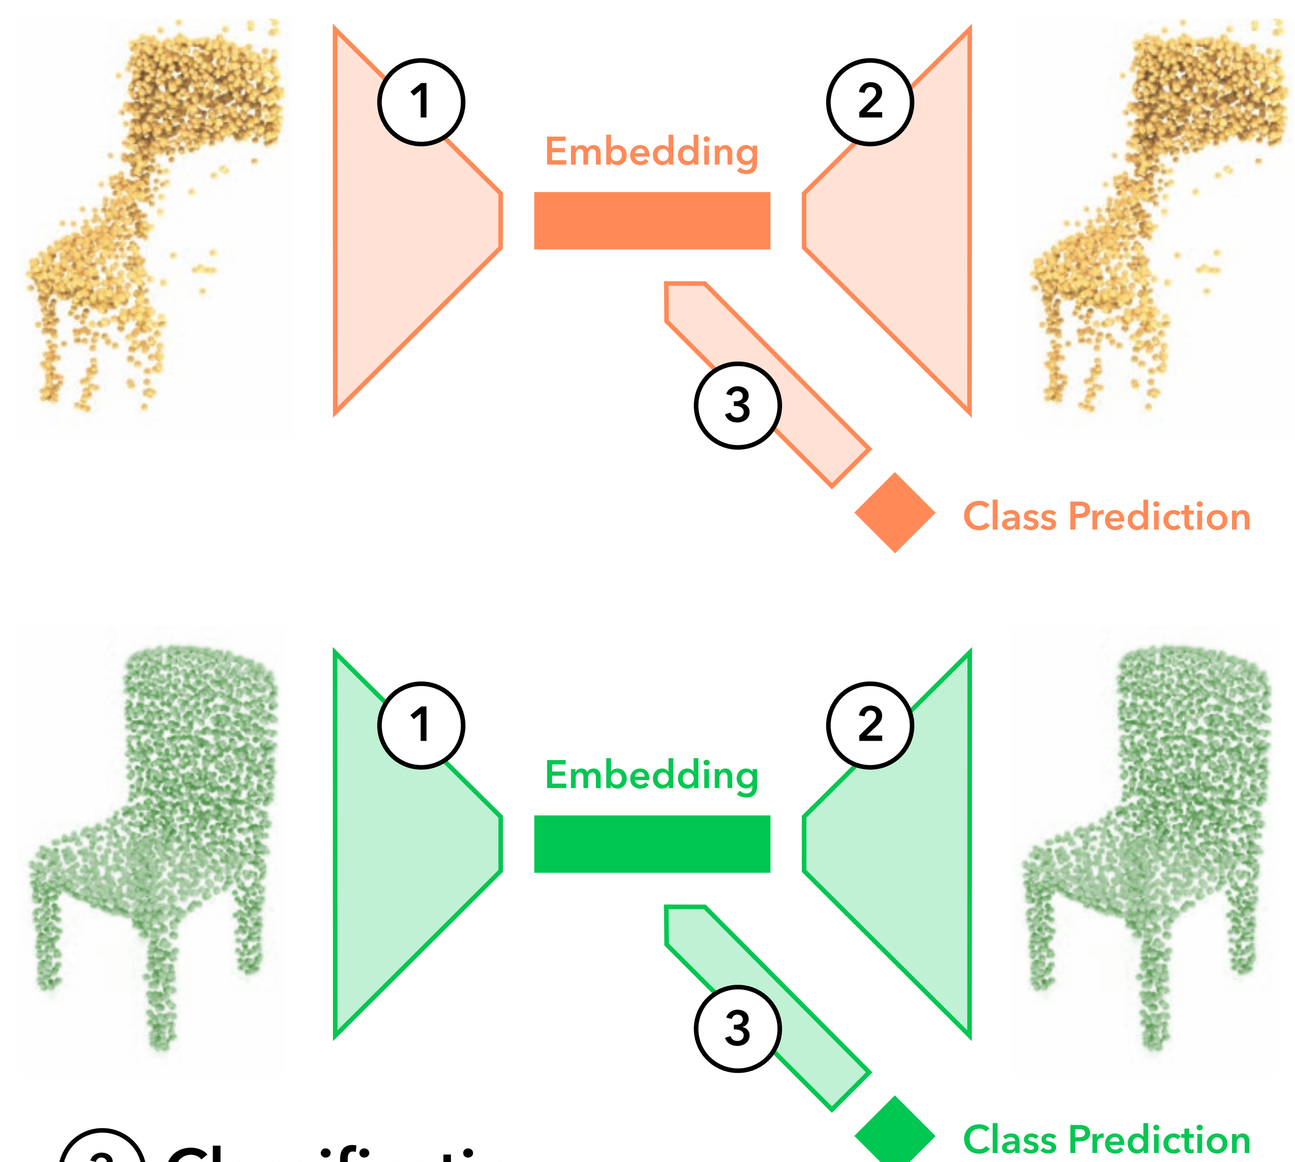

① Encoder

Using a series of repeated PointNet (Qi et al., 2017) modules to capture local and global information, we **embed each point cloud** into a single feature vector.

② Decoder

Using a **coarse-to-fine decoder** inspired by PCN (Yuan et al., 2018) + geometric reconstruction loss, we convert the low-dimensional embeddings **back to full point clouds**.

③ Classification

Using an MLP + cross entropy loss, we regress object class labels to **make different classes map to different parts of the embedding space**.

Use GAN to map between these distributions

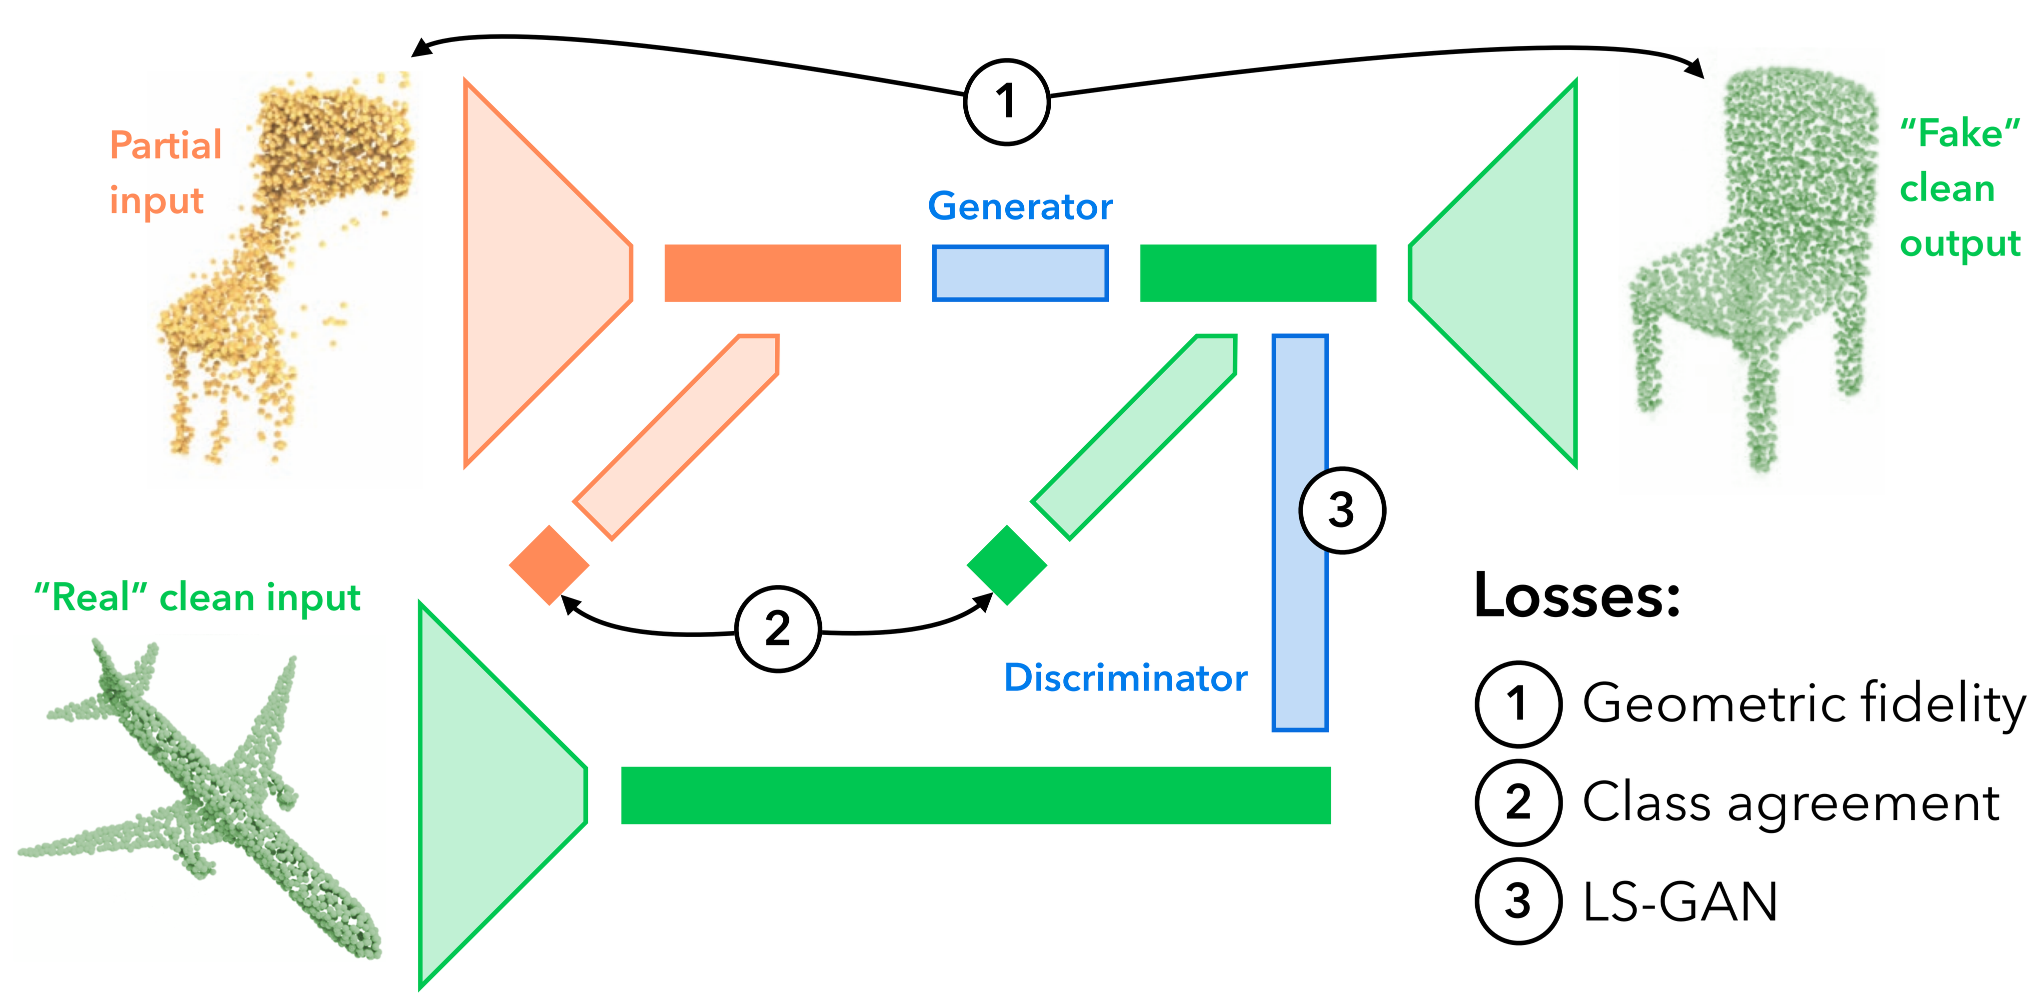

Losses:

① Geometric fidelity
② Class agreement
③ LS-GAN

Results

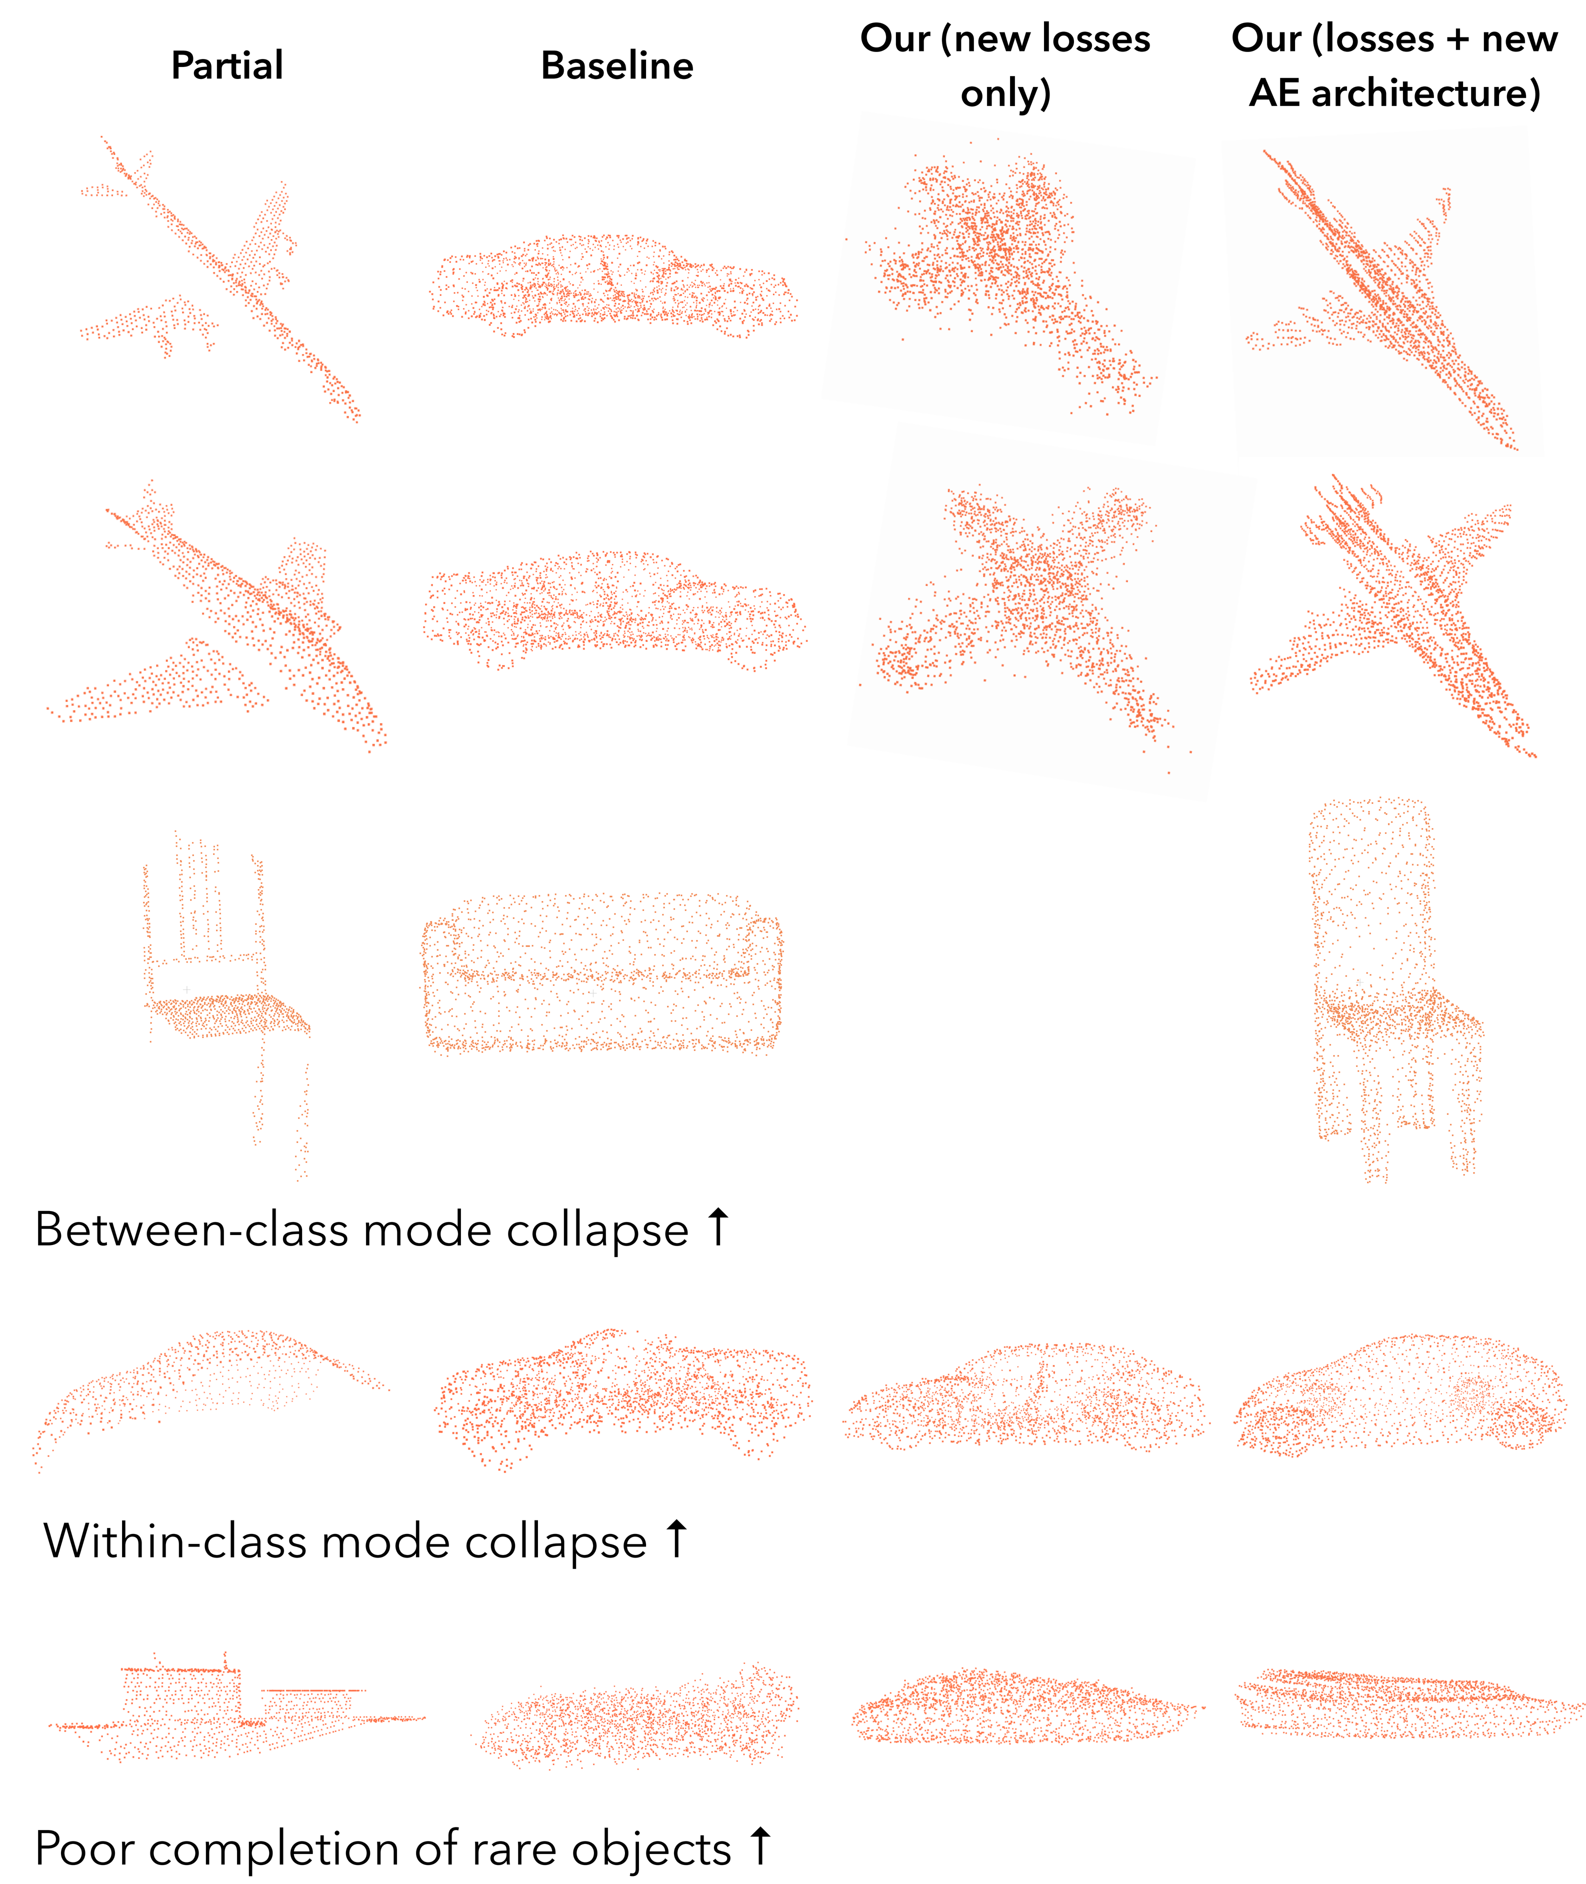

	Partial	Baseline	Our (new losses only)	Our (losses + new AE architecture)
Mean Hausdorff distance:		0.1217m	0.0804m	0.0766m

Future Work

- Fine-tune AEs while training GAN to produce more useful embeddings
- Evaluate with real-world partial + synthetic complete data
- Evaluate with combined data from more classes