
PatchComplete: Learning Multi-Resolution Patch Priors for 3D Shape Completion on Unseen Categories

Yuchen Rao Yinyu Nie Angela Dai

Technical University of Munich
{yuchen.rao, yinyu.nie, angela.dai}@tum.de

Abstract

While 3D shape representations enable powerful reasoning in many visual and perception applications, learning 3D shape priors tends to be constrained to the specific categories trained on, leading to an inefficient learning process, particularly for general applications with unseen categories. Thus, we propose *PatchComplete*, which learns effective shape priors based on multi-resolution local patches, which are often more general than full shapes (e.g., chairs and tables often both share legs) and thus enable geometric reasoning about unseen class categories. To learn these shared substructures, we learn multi-resolution patch priors across all train categories, which are then associated to input partial shape observations by attention across the patch priors, and finally decoded into a complete shape reconstruction. Such patch-based priors avoid overfitting to specific train categories and enable reconstruction on entirely unseen categories at test time. We demonstrate the effectiveness of our approach on synthetic ShapeNet data as well as challenging real-scanned objects from ScanNet, which include noise and clutter, improving over state of the art in novel-category shape completion by 19.3% in chamfer distance on ShapeNet, and 9.0% for ScanNet.

1 Introduction

The prevalence of commodity RGB-D sensors (e.g., Intel RealSense, Microsoft Kinect, iPhone, etc.) has enabled significant progress in 3D reconstruction, achieving impressive tracking quality [22, 15, 25, 7, 36, 12] and even large-scale reconstructed datasets [10, 3]. Unfortunately, 3D scanned reconstructions remain limited in geometric quality due to clutter, noise, and incompleteness (e.g., as seen in the objects in Figure 1). Understanding complete object structures is fundamental towards constructing effective 3D representations, that can then be used to fuel many applications in robotic perception, mixed reality, content creation, and more.

Recently, significant progress has been made in shape reconstruction, across a variety of 3D representations, including voxels [8, 13, 30], points [14, 39], meshes [35, 11], and implicit field representations [26, 18]. However, these methods tend to rely heavily on strong synthetic supervision, producing impressive reconstructions on similar train class categories but struggling to generalize to unseen categories. This leads to an expensive compute and data requirement in adapting to new objects in different scenarios, whose class categories may not necessarily have been seen during training and so must be re-trained or fine-tuned for.

In order to encourage more generalizable 3D feature learning to represent shape characteristics, we observe that while different class categories may have very different global structures, local geometric structures are often shared (e.g., a long, thin structure could represent a chair leg, a table leg, a lamp

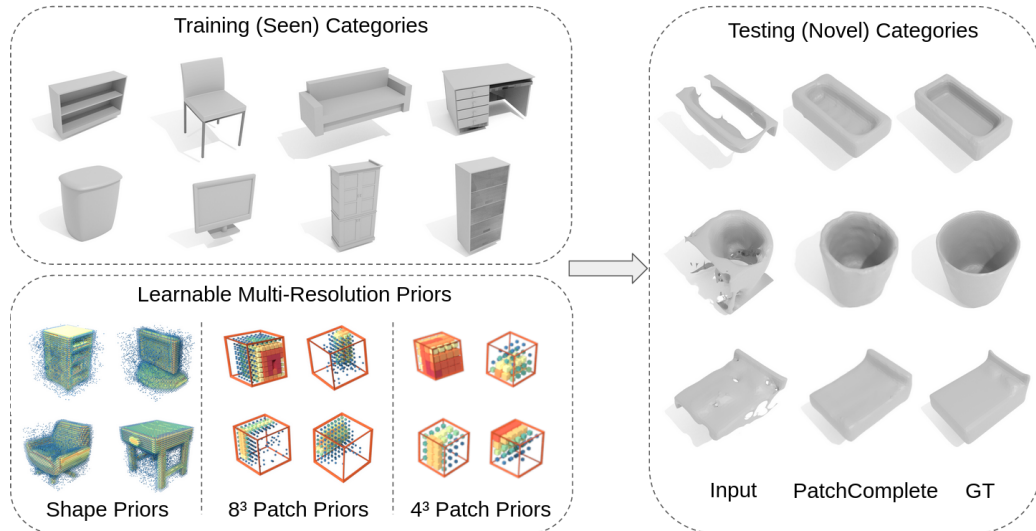


Figure 1: PatchComplete learns strong local priors for 3D shape completion by constructing multi-resolution patch priors from train shapes, which can then be applied for effective shape completion on unseen categories.

rod, etc.). We thus propose to learn a set of multi-resolution patch-based priors that captures such shared local substructures across the training set of shapes, which can be applied to shapes outside of the train set of categories. Our local patch-based priors can thus capture shared local structures, across different resolutions, that enable effective shape completion on novel class categories of not only synthetic data but also challenging real-world observations with noise and clutter.

We propose PatchComplete, which first learns patch priors for shape completion by correlating regions of observed partial inputs to the learned patch priors through an attention-based association, and decoding to reconstruct a complete shape. These patch priors are learned at different resolutions to encompass potentially different sizes of local substructures; we then learn to fuse the multi-resolution priors together to reconstruct the output complete shape. This enables learning generalizable local 3D priors that facilitate effective shape completion even for unseen categories, outperforming state of the art in synthetic and real-world observations by 19.3% and 9.0% on Chamfer Distance.

In summary, our contributions are:

- We propose generalizable 3D shape priors by learning patch-based priors that characterize shared local substructures that can be associated with input observations by cross-attention. This intermediate representation preserves structure explicitly, and can be effectively leveraged to compose complete shapes for entirely unseen categories.
- We design a multi-resolution fusion of different patch priors at various resolutions in order to effectively reconstruct a complete shape, enabling multi-resolution reasoning about the most informative learned patch priors to recover both global and local shape structures.

2 Related Work

2.1 3D Shape Reconstruction and Completion

Understanding how to reconstruct 3D shapes is an essential task for 3D machine perception. In particular, the task of shape completion to predict a complete shape from partial input observations has been studied by various works toward understanding 3D shape structures. Recently, many works have leveraged deep learning techniques to learn strong data-driven priors for shape completion, focusing on various different representations, e.g. volumetric grids [37, 13], continuous implicit functions [27, 18, 26, 6, 32], point clouds [28, 29], and meshes [11, 16]. These works tend to focus on strong synthetic supervision on a small set of train categories, achieving impressive performance on unseen shapes from train categories, but often struggling to generalize to unseen classes. We focus

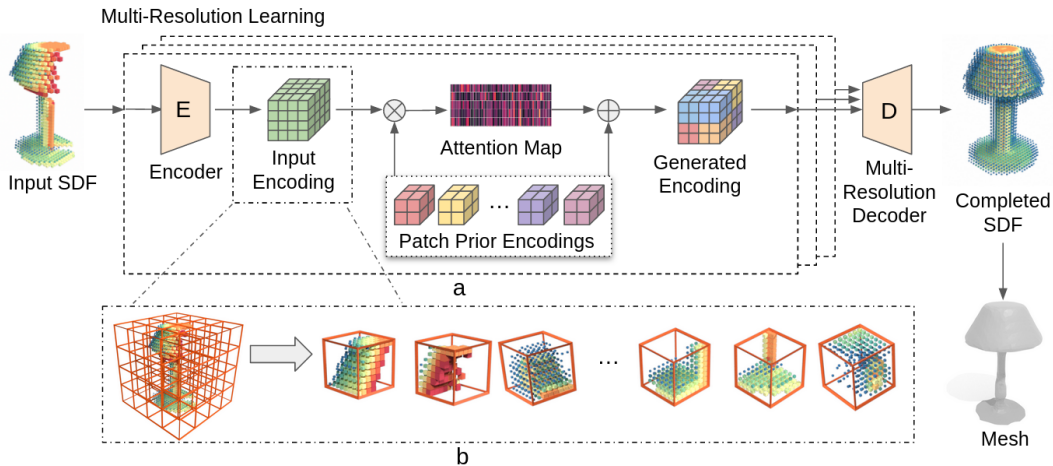


Figure 2: Overview of our approach. (a) shows our multi-resolution prior learning for shape completion. Each dotted block indicates patch prior learning for a single resolution, and the three different resolution encodings are then fused in a multi-resolution decoder that outputs a complete shape as an SDF grid. (b) illustrates our input partial observations and local input patches, from which we learn mappings to learned patch priors as to how to best compose the complete shape.

on learning more generalizable, local shape priors in order to effectively reconstruct complete shapes on unseen class categories.

2.2 Few-Shot and Zero-Shot 3D Shape Reconstruction

Several works have been developed to tackle the challenging task of few-shot or zero-shot shape reconstruction, as observations in-the-wild often contain a wide diversity of objects. In the few-shot scenario where several examples of novel categories are available, Wallace and Hariharan [34] learn to refine a given category prior. Michalkiewicz et al. [19] further propose to learn compositional shape priors for single-view reconstruction.

In the zero-shot scenario without any examples seen for novel categories, Naeem et al. [21] learn priors from seen categories to generate segmentation masks for unseen categories. Zhang et al. [40] additionally proposed to use spherical map representations to learn priors for the reconstruction of novel categories. Thai et al. [31] recently developed an approach to transfer knowledge from an RGB image for shape reconstruction. We also tackle a zero-shot shape reconstruction task, by learning a multi-resolution set of strong local shape priors to compose a reconstructed shape.

Several recent works have explored learning shape priors by leveraging a VQ-VAE backbone with autoregressive prediction to perform shape reconstruction [20, 38]. In contrast, we propose to learn multi-resolution shape priors without requiring any sequence interpretation, which enables direct applicability to real-world scan data that often contains noise and clutter. Additionally, hierarchical reconstruction has shown promising results for shape reconstruction [2, 5]. Our approach also takes a multi-resolution perspective, but learns explicit shape and local patch priors and their correlation to input partial observations for robust shape completion.

3 Method

3.1 Overview

Our method aims to learn effective 3D shape priors that enable general shape completion from partial input scan data, on novel class categories not seen during training. Key to our approach is the observation that 3D shapes often share repeated local patterns – for instance, chairs, tables, and nightstands all share a support surface, and chair or table legs can share a similar structure with lamp rods. Inspired by this, we regard a complete object as a set of substructures, where each substructure represents a local geometric region. We thus propose PatchComplete to learn such local priors and assemble them into a complete shape from a partial scan. An overview of our approach is shown in Figure. 2.

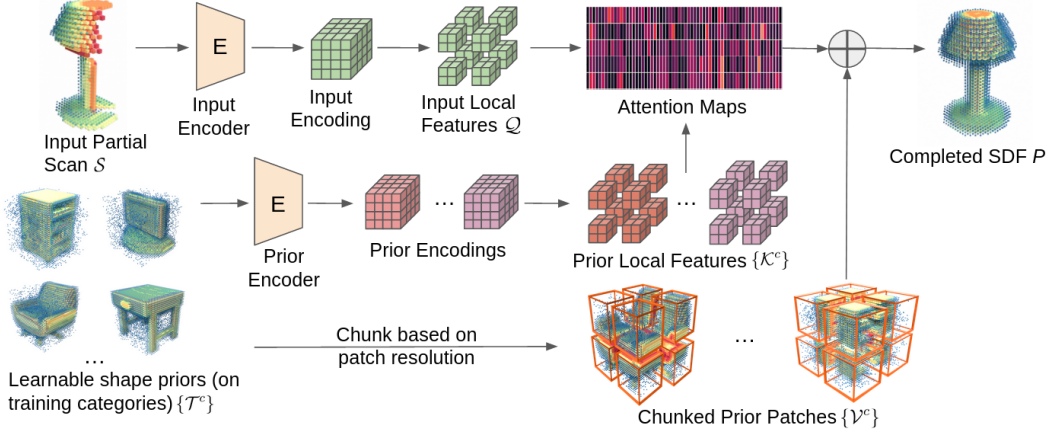


Figure 3: Network architecture for local patch-based shape prior learning under a single resolution. We learn to build mappings from local regions in a partial object scan \mathcal{S} to local priors based on complete learnable shape priors $\{\mathcal{T}^c\}$. An input encoder takes the input of \mathcal{S} to encode its local features (represented by patches \mathcal{Q}). Analogously, we adopt a prior encoder to process each prior in $\{\mathcal{T}^c\}$ to construct a local prior feature pool $\{\mathcal{K}^c\}$. In parallel, we chunk priors in $\{\mathcal{T}^c\}$ into patch volumes $\{\mathcal{V}^c\}$, which are correspondingly fused by the attention to input patches to compose the most informative patch priors: we use each incomplete local patch in \mathcal{Q} to query the keys $\{\mathcal{K}^c\}$ from complete prior patches, and assemble their corresponding patch volumes $\{\mathcal{V}^c\}$ for shape completion.

3.2 Learning Local Patch-Based Shape Priors

We first learn local shape priors from ground-truth train objects. We represent both the input partial scan \mathcal{S} and the ground-truth shape \mathcal{G} as 3D truncated signed distance field grids of size D^3 . Figure 3 illustrates the process of learning local shape priors. We learn to build mappings from local regions in incomplete input observations to local priors based on complete shapes, in order to robustly reconstruct the complete shape output.

We first aim to learn patch-based priors, which we extract from learnable global shape priors $\{\mathcal{T}^c\}$. We denote \mathcal{G}^c as the set of ground-truth train objects in the c -th category. These $\{\mathcal{T}^c\}$ priors are initialized per train category c based on mean-shift clustering within shapes in \mathcal{G}^c . Thus \mathcal{T}^c is a set of representative samples in each category, which are encoded parallel to the input scan \mathcal{S} .

Both encoders are analogously structured as 3D convolutional encoders which spatially downsample by a factor of R , resulting in a 3D encoding of size N^3 ; $N = D/R$. The input encoding (from \mathcal{S}) are uniformly chunked into patches $\{\mathcal{Q}_i\}$, $i=1, \dots, N^3$; similarly, the encoded priors (from $\{\mathcal{T}^c\}$) are chunked into patches $\{\mathcal{K}_i^c\}$.

We then use each local part \mathcal{Q}_i to query for the most informative encodings of complete local regions $\{\mathcal{K}_i^c\}$, building this mapping based on cross-attention [33]. Then for each input patch, we calculate its similarity with all local patches in representative shapes on training categories:

$$\text{Attention}(\mathcal{Q}_i, \mathcal{K}) = \text{softmax}\left(\frac{\mathcal{Q}_i \cdot \mathcal{K}^T}{d/2}\right), \mathcal{K} = \{\mathcal{K}_i^c | i = 1, \dots, N^3, c = 1, \dots, C\}, \quad (1)$$

where d is the dimension of the encoded vectors of $\mathcal{Q}_i, \mathcal{K}_i^c$; C is the category number. We then reconstruct complete shape patches $\{\mathcal{P}_i\}$ by:

$$\mathcal{P}_i = \text{Attention}(\mathcal{Q}_i, \mathcal{K}) \cdot \mathcal{V} \quad (2)$$

where $\mathcal{V} = \{\mathcal{V}_i^c\}$ is the set of N^3 chunks (of resolution R) from shapes in all categories $\{\mathcal{T}^c\}$, where each \mathcal{V}_i^c is paired with $\mathcal{K}_i^c \in \mathcal{K}$. We can then recombine $\{\mathcal{P}_i\}$ to the predicted full shape P .

Loss. We use an ℓ_1 reconstruction loss to train the learned local patch priors. Note that $\{\mathcal{T}^c\}$ are learned along with the network weights, enabling learning the most effective global shape priors for the shape completion task.

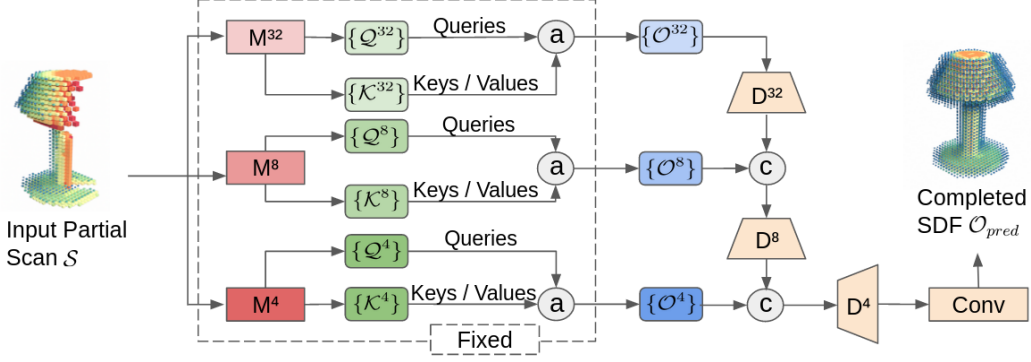


Figure 4: Network architecture for multi-resolution learning pipeline. We generate a complete shape from a partial input scan \mathcal{S} by fusing input local features and learned local priors in a multi-scale fashion. We first extract the input local features $\{Q^R\}$ and the learned local priors $\{K^R\}$ using the prior learning model (M^R) in Section 3.2 under different resolutions ($R=32, 8, 4$). Then, we use attention (Equation 3) to generate intermediate features $\{O^R\}$, which we then recursively fuse to decode a complete shape.

3.3 Multi-Resolution Patch Learning

In Section 3.2, we learn to complete shapes in a single patch resolution R . Since local substructures may exist at various scales, we learn patch priors under varying resolutions, which we then use to decode a complete shape. We use three patch resolutions ($R=4, 8, 32$), for which we learn patch priors. This results in three pairs of trained {input encoder, prior encoder} (see Figure 3). Given a partial input scan \mathcal{S} , each pair outputs a set of 1) input patch features $\{Q_i^R\}$, and 2) prior patch features \mathcal{K}^R , under the resolution of $R=4, 8, 32$. In this section, we decode a complete shape from these multi-resolution patch priors.

Since \mathcal{K}^R stores the patch priors, we use it as the (key, value) term into an attention module, where each input patch feature Q_i^R is used to query for the most informative prior features in \mathcal{K}^R under different resolutions, from which we complete partial patches in feature spaces with a multi-scale fashion. We formulate this process by

$$\mathcal{O}_i^R = \text{Concat} [Q_i^R, \text{Attention}(Q_i^R, \mathcal{K}^R) \cdot \mathcal{K}^R], R = 4, 8, 32, \quad (3)$$

where we concatenate input patch feature with the attention result to compensate information loss on observable regions. This outputs $\{\mathcal{O}_i^R\}$ as the intermediate feature of the i -th patch; $i = 1, \dots, N_R^3$. N_R denotes the number of patches under the resolution R ; $N_R = D/R$. We recombine all generated patch features into a volume $\mathcal{O}^R = \{\mathcal{O}_i^R | i = 1, \dots, N_R^3\}$. Note that the dimension of each grid feature in $\{Q_i^R\}$ and \mathcal{K}^R equals to d (see Eq. 1). Then \mathcal{O}^R is with the dimension of $N_R^3 \times 2d$.

We then sequentially use 3D convolution decoders to upsample and concatenate \mathcal{O}^R from low to high resolution to fuse all the shape features as in

$$\begin{aligned} \mathcal{O}^{R_{r+1}} &= \text{Concat} [\text{Upsample}(\mathcal{O}^{R_r}), \mathcal{O}^{R_{r+1}}], R_r |_{r=1,2,3} = 32, 8, 4, \\ \mathcal{O}_{pred} &= \text{Conv}(\text{Upsample}(\mathcal{O}^{R_3})). \end{aligned} \quad (4)$$

In Eq. 4, \mathcal{O}^{R_r} has a lower resolution than $\mathcal{O}^{R_{r+1}}$. We then use deconvolution layers to upsample \mathcal{O}^{R_r} to match the resolution of $\mathcal{O}^{R_{r+1}}$, and then concatenate them together. We recursively fuse \mathcal{O}^{R_r} into $\mathcal{O}^{R_{r+1}}$, which produces \mathcal{O}^{R_3} as the final fusion result with the resolution of $(D/4)^3$. An extra upsampling followed by a convolution layer is adopted to upsample \mathcal{O}^{R_3} into our shape prediction \mathcal{O}_{pred} with the dimension of D^3 .

In training the feature fusion, we fix all parameters from the {input encoder, prior encoder} under the three resolutions, since they are pre-trained on Section 3.2 and provide better-learned priors and attention maps under these strong constraints. The whole pipeline for this section can be found in Figure 4.

We use an ℓ_1 loss to supervise the TSDF value regression. We weight the ℓ_1 loss to penalize false predictions based on the predicted signs as Eq. 5 shows, which represents whether this voxel grid is occupied or not. It is also used in the loss function in Section 3.2.

$$\mathcal{L} = w_{occ}\ell_1(\mathcal{O}_{pred}^{occ}, \mathcal{O}_{gt}) + w_{empty}\ell_1(\mathcal{O}_{pred}^{empty}, \mathcal{O}_{gt}) + \ell_1(\mathcal{O}_{pred}^{correct}, \mathcal{O}_{gt}). \quad (5)$$

In Eq. 5, $\mathcal{O}_{pred}^{empty}$ represents the false positive TSDF values, where the ground truth has negative signs and the prediction has positive signs, which in general indicates the missing predictions. \mathcal{O}_{pred}^{occ} represents false negative TSDF values, where the ground truth has positive signs and the prediction has negative signs, which indicates extra predictions. $\mathcal{O}_{pred}^{correct}$ means those predicted TSDF values with the same signs as the ground-truth. During training, we choose the weight for false positive (w_{empty}) as 5, and the weight for false negative (w_{occ}) as 3.

3.4 Implementation Details

We train our approach on a single NVIDIA A6000 GPU, using an Adam optimizer with batch size 32 for a synthetic dataset and batch size 16 for a real-world dataset, and an initial learning rate of 0.001. We train for 80 epochs until convergence, and then decrease the learning rate by half after epoch 50. We use the same settings for learning priors and the multi-resolution decoding, which train for 4 and 15.5 hours respectively. For additional network architecture details, we refer to the supplemental material.

Note that for training for real scan data, we first pre-train on synthetic data and then fine-tune only the input encoder.

4 Experiments and Analysis

4.1 Experimental Setup

Datasets. We train and evaluate our approach on synthetic shape data from ShapeNet [4] as well as on challenging real-world scan data from ScanNet [10]. For ShapeNet data, we virtually scan the objects to create partial input scans, following [23, 13]. We use 18 categories during training, and test on 8 novel categories, resulting in 3,202/1,325 train/test models with 4 partial scans for each model.

For real data, we use real-scanned objects from ScanNet extracted by their bounding boxes, with corresponding complete target data given by Scan2CAD [1]. We use 8 categories for training, comprising 7,537 train samples, and test on 6 novel categories of 1,191 test samples.

For all experiments, objects are represented as 32^3 signed distance grids with truncation of 2.5 voxel units for ShapeNet and 3 voxel units for ScanNet. The objects in ShapeNet are normalized into the unit cube, while we keep the scaling for ScanNet objects, and save their voxel sizes separately to keep the real size information. Additionally, to train and evaluate on real data, all methods are first pre-trained on ShapeNet and then fine-tuned on ScanNet.

Baselines. We evaluate our approach against various state-of-the-art shape completion methods. We compare with state-of-the-art shape completion methods 3D-EPN [13] and IF-Nets [6], which learn effective shape completion on dense voxel grids and with implicit neural field representations, respectively, without any focus on unseen class categories. We further compare to the state-of-the-art few-shot shape reconstruction approach of Wallace and Hariharan [34] (referred to as Few-Shot) leveraging global shape priors, which we apply in our zero-shot unseen category scenario. Finally, AutoSDF [20] uses a VQ-VAE module with a transformer-based autoregressive model over latent patch priors to produce TSDF shape reconstruction.

Evaluation Metrics. To evaluate the quality of reconstructed shape geometry, we use \mathcal{L}_1 Chamfer Distance (CD) and Intersection over Union (IoU) between predicted and ground truth shapes. To evaluate methods that output occupancy grids, we use occupancy thresholds used by the respective methods to obtain voxel predictions, i.e. 0.4 for [34] and 0.5 for [6]. To evaluate methods that output signed distance fields, we extract the iso-surface at level zero with marching cubes [17]. 10K points

Table 1: Quantitative comparison for shape completion on synthetic ShapeNet [4] data.

	Chamfer Distance ($\times 10^2$) \downarrow					IoU \uparrow				
	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours
Bag	5.01	8.00	4.77	5.81	3.94	0.738	0.561	0.698	0.563	0.776
Lamp	8.07	15.10	5.70	6.57	4.68	0.472	0.254	0.508	0.391	0.564
Bathtub	4.21	7.05	4.72	5.17	3.78	0.579	0.457	0.550	0.410	0.663
Bed	5.84	10.03	5.34	6.01	4.49	0.584	0.396	0.607	0.446	0.668
Basket	7.90	8.72	4.44	6.70	5.15	0.540	0.406	0.502	0.398	0.610
Printer	5.15	9.26	5.83	7.52	4.63	0.736	0.567	0.705	0.499	0.776
Laptop	3.90	10.35	6.47	4.81	3.77	0.620	0.313	0.583	0.511	0.638
Bench	4.54	8.11	5.03	4.31	3.70	0.483	0.272	0.497	0.395	0.539
Inst-Avg	5.48	9.75	5.37	5.76	4.23	0.582	0.386	0.574	0.446	0.644
	$\pm 2e^{-1}$	$\pm 9e^{-2}$	$\pm 1e^{-1}$	$\pm 3e^{-2}$	$\pm 4e^{-2}$	$\pm 9e^{-3}$	$\pm 1e^{-3}$	$\pm 4e^{-5}$	$\pm 6e^{-3}$	$\pm 1e^{-3}$
Cat-Avg	5.58	9.58	5.29	5.86	4.27	0.594	0.403	0.581	0.452	0.654
	$\pm 2e^{-1}$	$\pm 1e^{-1}$	$\pm 1e^{-1}$	$\pm 5e^{-3}$	$\pm 5e^{-2}$	$\pm 8e^{-3}$	$\pm 1e^{-3}$	$\pm 3e^{-4}$	$\pm 7e^{-3}$	$\pm 1e^{-3}$

are sampled on surfaces for CD calculation. Both Chamfer distance and IoU are evaluated on objects in the canonical system, and we report Chamfer Distance scaled by $\times 10^2$.

4.2 Evaluation on Synthetic Data

In Table 1, we evaluate our approach in comparison with prior arts on unseen class categories of synthetic ShapeNet [4] data. Our approach on learning attention-based correlation to learned local shape priors results in notably improved reconstruction performance, with coherent global and local structures, as shown in Figure 5. In Table 1, our work outperforms other baselines both instance-wise and category-wise. One of the key factors is that our method learns multi-scale patch information from seen categories to complete unseen categories with enough flexibility, while most of the other baselines are designed for 3D shape completion on known categories, which hardly leverage shape priors across categories.




Figure 5: Qualitative comparison for shape completion on synthetic ShapeNet [4] dataset.

4.3 Evaluation on Real Scan Data

Table 2 evaluates our approach in comparison with prior arts on real scanned objects from unseen categories in ScanNet [10]. Here, input scans are not only partial but often contain noise and clutter; our multi-resolution learned priors enable more robust shape completion in this challenging scenario. Results in Figure 6 further demonstrate that our approach presents more coherent shape completion

than prior methods by using cross-attention with learnable priors, which better preserves the global structures in coarse and cluttered environments.

Table 2: Quantitative comparison with state of the art on real-world ScanNet shape completion.

	Chamfer Distance ($\times 10^2$) \downarrow					IoU \uparrow				
	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours
Bag	8.83	9.10	8.96	9.30	8.23	0.537	0.449	0.442	0.487	0.583
Lamp	14.27	11.88	10.16	11.17	9.42	0.207	0.196	0.249	0.244	0.284
Bathtub	7.56	7.77	7.19	7.84	6.77	0.410	0.382	0.395	0.366	0.480
Bed	7.76	9.07	8.24	7.91	7.24	0.478	0.349	0.449	0.380	0.484
Basket	7.74	8.02	6.74	7.54	6.60	0.365	0.343	0.427	0.361	0.455
Printer	8.36	8.30	8.28	9.66	6.84	0.630	0.622	0.607	0.499	0.705
Inst-Avg	8.60	8.83	8.12	8.56	7.38	0.441	0.387	0.426	0.386	0.498
Cat-Avg	$\pm 2e^{-1}$	$\pm 2e^{-2}$	$\pm 7e^{-2}$	$\pm 2e^{-2}$	$\pm 6e^{-2}$	$\pm 2e^{-3}$	$\pm 1e^{-3}$	$\pm 3e^{-3}$	$\pm 1e^{-4}$	$\pm 9e^{-3}$
Cat-Avg	9.09	9.02	8.26	8.90	7.52	0.440	0.386	0.426	0.389	0.495
	$\pm 3e^{-1}$	$\pm 8e^{-2}$	$\pm 8e^{-2}$	$\pm 2e^{-2}$	$\pm 2e^{-2}$	$\pm 3e^{-3}$	$\pm 6e^{-3}$	$\pm 7e^{-3}$	$\pm 3e^{-4}$	$\pm 5e^{-3}$




Figure 6: Shape completion on real-world ScanNet [10] object scans. Our method to learn mappings to multi-resolution learnable patch priors enables more coherent shape completion on novel categories.

4.4 Ablation Analysis

Table 3: Ablation study on different patch resolutions. A multi-resolution approach gains benefits from both global and local reasoning.

	ShapeNet [4]				ScanNet [10]			
	Inst-CD \downarrow	Cat-CD \downarrow	Inst-IoU \uparrow	Cat-IoU \uparrow	Inst-CD \downarrow	Cat-CD \downarrow	Inst-IoU \uparrow	Cat-IoU \uparrow
Ours (32^3 priors only)	11.94	11.61	0.35	0.37	10.43	11.23	0.41	0.40
Ours (8^3 priors only)	4.86	4.89	0.61	0.62	7.64	7.81	0.48	0.49
Ours (4^3 priors only)	4.44	4.50	0.64	0.64	7.34	7.58	0.49	0.49
Ours	4.23	4.27	0.64	0.65	7.38	7.52	0.50	0.50

Does multi-resolution patch learning help shape completion for novel categories? In Table 3, we evaluate shape completion with each resolution in comparison with our multi-resolution approach. Learning only global shape priors (i.e., 32^3) tends to overfit to seen train categories, while the local

Table 4: Ablation study on attention used to learn input and patch prior correlations.

	ShapeNet [4]				ScanNet [10]			
	Inst-CD↓	Cat-CD↓	Inst-IoU↑	Cat-IoU↑	Inst-CD↓	Cat-CD↓	Inst-IoU↑	Cat-IoU↑
Ours (no attention)	4.90	4.98	0.61	0.62	7.80	8.09	0.49	0.48
Ours	4.69	4.74	0.61	0.63	7.58	7.84	0.48	0.49

patch resolutions can provide more generalizable priors. Combining all results in complementary feature learning for the most effective shape completion results.

Does cross-attention to learn local priors help? We evaluate our approach to learn both local priors and their correlation to input observations with cross-attention in Table 4, which shows that this enables more effective shape completion on unseen categories.

Table 5: Effect of synthetic pre-training on real-world ScanNet object completion vs. training from scratch.

	Inst-CD↓	Cat-CD↓	Inst-IoU↑	Cat-IoU↑
Scratch	7.61	7.73	0.50	0.50
Ours	7.38	7.52	0.50	0.50

What is the effect of synthetic pre-training for real scan completion? Table 5 shows the effect of synthetic pre-training for shape completion on real scanned objects. This encourages learning more robust priors to output cleaner local structures as given in the synthetic data, resulting in improved performance on real scanned objects.

Does learning the priors help completion? In Table 6, we evaluate our learnable priors in comparison with using fixed priors (by mean-shift clustering of train objects) for shape completion on ShapeNet [4]. Learned priors receive gradient information to adapt to best reconstruct the shapes, enabling improved performance over a fixed set of priors.

Table 6: Ablation on learnable priors in comparison with fixed priors on ShapeNet.

	Inst-CD↓	Cat-CD↓	Inst-IoU↑	Cat-IoU↑
Ours (fixed priors)	4.31	4.34	0.64	0.65
Ours	4.23	4.27	0.64	0.65

4.5 Limitations

While PatchComplete has presented a promising step towards learning more generalizable shape priors, various limitations remain. For instance, output shape completion is limited by the dense voxel grid resolution in representing fine-scale geometric details. Additionally, detected object bounding boxes are required for real scan data as input to independent shape completion predictions, while formulation considering other objects in the scene or an end-to-end framework could learn more effectively from global scene contextual information.

5 Conclusion

We have proposed PatchComplete to learn effective local shape priors for shape completion that enables robust reconstruction for novel class categories at test time. This enables learning shared local substructures across a variety of shapes that can be mapped to local incomplete observations by cross attention, with multi-resolution fusion producing coherent geometry at global and local scales. This enables robust shape completion even for unseen categories with different global structures, across synthetic as well as challenging real-world scanned objects with noise and clutter. We believe that such robust reconstruction of real scanned objects takes an important step towards understanding 3D shapes, and hope this inspires future work in understanding real-world shapes and scene structures.

Acknowledgements

This project is funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt).

References

- [1] Avetisyan, A., Dahnert, M., Dai, A., Savva, M., Chang, A.X., Nießner, M.: Scan2cad: Learning cad model alignment in rgb-d scans (2018)
- [2] Bechtold, J., Tatarchenko, M., Fischer, V., Brox, T.: Fostering generalization in single-view 3d reconstruction by learning a hierarchy of local and global shape priors (2021)
- [3] Chang, A.X., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D data in indoor environments. *International Conference on 3D Vision (3DV)* (2017)
- [4] Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository (2015)
- [5] Chen, Z., Zhang, Y., Genova, K., Fanello, S., Bouaziz, S., Haene, C., Du, R., Keskin, C., Funkhouser, T., Tang, D.: Multiresolution deep implicit functions for 3d shape representation (2021)
- [6] Chibane, J., Alldieck, T., Pons-Moll, G.: Implicit functions in feature space for 3d shape reconstruction and completion (2020)
- [7] Choi, S., Zhou, Q.Y., Koltun, V.: Robust reconstruction of indoor scenes. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5556–5565. IEEE (2015)
- [8] Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: *European conference on computer vision*. pp. 628–644. Springer (2016)
- [9] Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. pp. 303–312 (1996)
- [10] Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes (2017)
- [11] Dai, A., Nießner, M.: Scan2mesh: From unstructured range scans to 3d meshes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5574–5583 (2019)
- [12] Dai, A., Nießner, M., Zollhöfer, M., Izadi, S., Theobalt, C.: Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (TOG)* **36**(3), 24 (2017)
- [13] Dai, A., Qi, C.R., Nießner, M.: Shape completion using 3d-encoder-predictor cnns and shape synthesis (2017)
- [14] Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 605–613 (2017)
- [15] Izadi, S., Kim, D., Hilliges, O., Molyneaux, D., Newcombe, R., Kohli, P., Shotton, J., Hodges, S., Freeman, D., Davison, A., et al.: Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In: *Proceedings of the 24th annual ACM symposium on User interface software and technology*. pp. 559–568. ACM (2011)
- [16] Li, X., Liu, S., Kim, K., Mello, S.D., Jampani, V., Yang, M.H., Kautz, J.: Self-supervised single-view 3d reconstruction via semantic consistency. In: *European Conference on Computer Vision*. pp. 677–693. Springer (2020)
- [17] Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics* **21**(4), 163–169 (1987)

- [18] Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space (2019)
- [19] Michalkiewicz, M., Tsogkas, S., Parisot, S., Baktashmotlagh, M., Eriksson, A., Belilovsky, E.: Learning compositional shape priors for few-shot 3d reconstruction (2021)
- [20] Mittal, P., Cheng, Y.C., Singh, M., Tulsiani, S.: AutoSDF: Shape priors for 3d completion, reconstruction and generation. In: CVPR (2022)
- [21] Naeem, M.F., Örnek, E.P., Xian, Y., Gool, L.V., Tombari, F.: 3d compositional zero-shot learning with decompositional consensus (2021)
- [22] Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohi, P., Shotton, J., Hodges, S., Fitzgibbon, A.: Kinectfusion: Real-time dense surface mapping and tracking. In: Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on. pp. 127–136. IEEE (2011)
- [23] Nie, Y., Lin, Y., Han, X., Guo, S., Chang, J., Cui, S., Zhang, J., et al.: Skeleton-bridged point completion: From global inference to local adjustment. *Advances in Neural Information Processing Systems* **33**, 16119–16130 (2020)
- [24] Nie, Y., Lin, Y., Han, X., Guo, S., Chang, J., Cui, S., Zhang, J.: Skeleton-bridged point completion: From global inference to local adjustment. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 16119–16130. Curran Associates, Inc. (2020)
- [25] Nießner, M., Zollhöfer, M., Izadi, S., Stamminger, M.: Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (TOG)* (2013)
- [26] Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation (2019)
- [27] Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks (2020)
- [28] Stutz, D., Geiger, A.: Learning 3d shape completion under weak supervision. *International Journal of Computer Vision* **128**(5), 1162–1181 (2020)
- [29] Tang, J., Lei, J., Xu, D., Ma, F., Jia, K., Zhang, L.: Sa-convonet: Sign-agnostic optimization of convolutional occupancy networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6504–6513 (2021)
- [30] Tatarchenko, M., Dosovitskiy, A., Brox, T.: Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2088–2096 (2017)
- [31] Thai, A., Stojanov, S., Upadhyay, V., Rehg, J.M.: 3d reconstruction of novel object shapes from single images (2021)
- [32] Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Stoll, C., Theobalt, C.: Patchnets: Patch-based generalizable deep implicit 3d shape representations (2021)
- [33] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
- [34] Wallace, B., Hariharan, B.: Few-shot generalization for single-image 3d reconstruction via priors (2019)
- [35] Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 52–67 (2018)
- [36] Whelan, T., Leutenegger, S., Salas-Moreno, R.F., Glocker, B., Davison, A.J.: Elasticfusion: Dense slam without a pose graph. *Proc. Robotics: Science and Systems, Rome, Italy* (2015)

- [37] Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes (2015)
- [38] Yan, X., Lin, L., Mitra, N.J., Lischinski, D., Cohen-Or, D., Huang, H.: Shapeformer: Transformer-based shape completion via sparse representation. arXiv preprint arXiv:2201.10326 (2022)
- [39] Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4541–4550 (2019)
- [40] Zhang, X., Zhang, Z., Zhang, C., Tenenbaum, J.B., Freeman, W.T., Wu, J.: Learning to reconstruct shapes from unseen classes (2018)

In this supplementary material, we describe the data generation (Sec. A), settings for baseline methods (Sec. B), evaluation of category-wise error bars on ShapeNet and ScanNet (Sec. C), evaluation on seen categories for all the methods on ScanNet (Sec. D), network parameters and specifications (Sec. E), and additional qualitative results (Sec. F).

A Data Generation

ShapeNet [4] We use ShapeNet¹ to test our performance on synthetic data. In order to generate watertight meshes as ground truth, we first normalize ShapeNet CAD models, and render depth maps under 20 different viewpoints for each model. We then use volumetric fusion [9] to generate 32³ truncated signed distance fields (TSDFs) with truncation value as 2.5 voxel units. Finally, we choose 4 views of TSDF as the input, which mimic the partial scan in real data (e.g., ScanNet). The main idea can be referred to [24]².

We split the training and testing object categories on ShapeNet as follows. The 18 training categories are *table, chair, sofa, cabinet, clock, bookshelf, piano, microwave, stove, file cabinet, trash bin, bowl, display, keyboard, dishwasher, washing machine, pots, faucet, and guitar*; and the 8 novel testing categories are *bathtub, lamp, bed, bag, printer, laptop, bench, and basket*.

ScanNet [10] We use ScanNet³ to test our method on real-world data. The inputs are directly extracted from ScanNet scenes based on the bounding box annotations from Scan2CAD [1]⁴. We keep their real scale and convert them to 32³ voxel grids with truncation value at 3 voxel units, and save their voxel size separately. These inputs could contain walls, floors, or other cluttered backgrounds, which are transformed to canonical space to be aligned with the ShapeNet model coordinate system. The ground-truths are the corresponding complete and watertight ShapeNet meshes based on Scan2CAD annotations, which are generated with the similar method as above.

We split the training and testing object categories on ScanNet as follows. The 8 training categories are *chair, table, sofa, trash bin, cabinet, bookshelf, file cabinet, and monitor*; and the 6 novel testing categories are *bathtub, lamp, bed, bag, basket, and printer*, and each category has more than 50 samples for testing.

B Baseline Comparison

We use the authors' original implementations and hyperparameters in all the baselines for fair comparisons.

3D-EPN [13] 3D-EPN is a two-stage network, which completes partial 3D scans first and then reconstructs the completed shapes to a higher resolution by retrieving priors from a category-wise shape pool. In our case, priors for novel categories are not accessible, thus, we only compare its 3D Encoder-Predictor Network (the 3D completion model) on our dataset.

Wallace and Hariharan [34] (Few-Shot) This method uses a few-shot learning strategy for single view completion with averaged shape prior for each category. For a fair comparison with other works, we adapt it to a zero-shot learning mechanism here. We pre-compute the averaged shape priors for each training category; during training, we use two voxel encoder modules in parallel for the occupied voxel grids inputs and the averaged shape priors based on the input category; in the testing step, since we cannot provide shape priors for novel categories, we average shape prior from all the training categories, and use this averaged shape prior as input to the prior encoder module, along with the testing samples for shape completion.

¹The license can be found here: <https://shapenet.org/>, received permission after registration without personally identifiable information or offensive content

²https://github.com/yinyunie/depth_renderer

³The license can be found here: <https://github.com/ScanNet/ScanNet>, filled out an agreement without personally identifiable information or offensive content

⁴The license can be found here: <https://github.com/skanti/Scan2CAD>, filled out an agreement without personally identifiable information or offensive content

Table 7: Quantitative comparisons with state of the art on ShapeNet [4].

	Chamfer Distance ($\times 10^2$) \downarrow					IoU \uparrow				
	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours
Bag	5.01 $\pm 2e^{-1}$	8.00 $\pm 4e^{-1}$	4.77 $\pm 3e^{-1}$	5.81 $\pm 2e^{-1}$	3.94 $\pm 2e^{-2}$	0.738 $\pm 5e^{-3}$	0.561 $\pm 2e^{-2}$	0.698 $\pm 2e^{-3}$	0.563 $\pm 1e^{-2}$	0.776 $\pm 3e^{-3}$
Lamp	8.07 $\pm 6e^{-1}$	15.10 $\pm 2e^{-1}$	5.70 $\pm 5e^{-1}$	6.57 $\pm 1e^{-1}$	4.68 $\pm 2e^{-2}$	0.472 $\pm 1e^{-2}$	0.254 $\pm 2e^{-3}$	0.508 $\pm 1e^{-3}$	0.391 $\pm 1e^{-2}$	0.564 $\pm 3e^{-3}$
Bathtub	4.21 $\pm 1e^{-2}$	7.05 $\pm 1e^{-1}$	4.72 $\pm 8e^{-2}$	5.17 $\pm 2e^{-2}$	3.78 $\pm 2e^{-2}$	0.579 $\pm 2e^{-2}$	0.457 $\pm 4e^{-4}$	0.550 $\pm 9e^{-3}$	0.410 $\pm 8e^{-3}$	0.663 $\pm 9e^{-4}$
Bed	5.84 $\pm 3e^{-2}$	10.03 $\pm 2e^{-1}$	5.34 $\pm 2e^{-1}$	6.01 $\pm 1e^{-1}$	4.49 $\pm 2e^{-4}$	0.584 $\pm 4e^{-3}$	0.396 $\pm 1e^{-3}$	0.607 $\pm 2e^{-3}$	0.446 $\pm 1e^{-2}$	0.668 $\pm 2e^{-3}$
Basket	7.90 $\pm 4e^{-1}$	8.72 $\pm 9e^{-2}$	4.44 $\pm 1e^{-2}$	6.70 $\pm 3e^{-1}$	5.15 $\pm 3e^{-1}$	0.540 $\pm 1e^{-2}$	0.406 $\pm 5e^{-3}$	0.502 $\pm 1e^{-4}$	0.398 $\pm 1e^{-2}$	0.610 $\pm 3e^{-3}$
Printer	5.15 $\pm 1e^{-2}$	9.26 $\pm 5e^{-2}$	5.83 $\pm 1e^{-1}$	7.52 $\pm 1e^{-1}$	4.63 $\pm 7e^{-2}$	0.736 $\pm 6e^{-3}$	0.567 $\pm 3e^{-3}$	0.705 $\pm 4e^{-3}$	0.499 $\pm 3e^{-2}$	0.776 $\pm 2e^{-4}$
Laptop	3.90 $\pm 1e^{-1}$	10.35 $\pm 3e^{-1}$	6.47 $\pm 8e^{-1}$	4.81 $\pm 2e^{-1}$	3.77 $\pm 9e^{-2}$	0.620 $\pm 6e^{-3}$	0.313 $\pm 2e^{-2}$	0.583 $\pm 1e^{-3}$	0.511 $\pm 1e^{-2}$	0.638 $\pm 8e^{-3}$
Bench	4.54 $\pm 3e^{-2}$	8.11 $\pm 8e^{-1}$	5.03 $\pm 9e^{-1}$	4.31 $\pm 5e^{-2}$	3.70 $\pm 1e^{-2}$	0.483 $\pm 1e^{-2}$	0.272 $\pm 1e^{-2}$	0.497 $\pm 4e^{-3}$	0.395 $\pm 3e^{-3}$	0.539 $\pm 3e^{-4}$
Inst-Avg	5.48 $\pm 2e^{-1}$	9.75 $\pm 9e^{-2}$	5.37 $\pm 1e^{-1}$	5.76 $\pm 3e^{-2}$	4.23 $\pm 4e^{-2}$	0.582 $\pm 9e^{-3}$	0.386 $\pm 1e^{-3}$	0.574 $\pm 4e^{-5}$	0.446 $\pm 6e^{-3}$	0.644 $\pm 1e^{-3}$
Cat-Avg	5.58 $\pm 2e^{-1}$	9.58 $\pm 1e^{-1}$	5.29 $\pm 1e^{-1}$	5.86 $\pm 5e^{-3}$	4.27 $\pm 5e^{-2}$	0.594 $\pm 8e^{-3}$	0.403 $\pm 1e^{-3}$	0.581 $\pm 3e^{-4}$	0.452 $\pm 7e^{-3}$	0.654 $\pm 1e^{-3}$

IF-Net [6] IF-Net can predict implicit shape representations conditioned on different input modalities. (e.g., voxels, point clouds). We use 128^3 occupied surface voxel grids as inputs, and use point clouds sampled from watertight ShapeNet meshes as the ground-truths for training and testing. We also normalize the ground-truth meshes from the ScanNet dataset to sample points.

AutoSDF [20] AutoSDF learns latent patch priors using VQ-VAE along with a transformer-based autoregressive model for 3D shape completion, and manually picks the unknown patches during testing. Following their settings, we apply their method by using ground-truth SDFs as the training data; during testing on the ShapeNet data, we choose the patches that have more than 400 voxel grids (each patch has 8^3 voxel grids) with negative signs as the unknown patches (unseen parts) that need to be generated.

Note that since AutoSDF work focuses on multi-model shape completion and produces multiple output possibilities, we report the performance of only the best prediction among the nine given an oracle to indicate the best (highest IoU value with respect to ground truth).

Furthermore, as there are no absolutely unknown patches for ScanNet scans because of the cluttered environments, we use the pipeline of their single view reconstruction task. We first replace their *ResNet* in *resnet2vq_model* with three 3D encoders (the same as 3D-EPN encoders) to extract the encoding features of desired dimensions; then we train this modified model along with the pre-trained *pvq_vae_model* with partial ScanNet inputs; finally we test our partial ScanNet inputs along with all the pre-trained models: *resnet2vq_model*, *pvq_vae_model*, and *rand_tf_model*.

C Category-wise Evaluations with Error Bars

Table 7 and Table 8 show the category-wise error bars on ShapeNet and ScanNet respectively; each method is run $n = 2$ times to obtain the error bars.

Table 8: Quantitative comparisons with state of the art on ScanNet [10].

	Chamfer Distance ($\times 10^2$) \downarrow					IoU \uparrow				
	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours
Bag	8.83	9.10	8.96	9.30	8.23	0.537	0.449	0.442	0.487	0.583
	$\pm 1e^{-1}$	$\pm 2e^{-1}$	$\pm 1e^{-1}$	$\pm 3e^{-2}$	$\pm 4e^{-2}$	$\pm 1e^{-2}$	$\pm 6e^{-3}$	$\pm 5e^{-3}$	$\pm 1e^{-4}$	$\pm 8e^{-3}$
Lamp	14.27	11.88	10.16	11.17	9.42	0.207	0.196	0.249	0.244	0.284
	$\pm 2e^0$	$\pm 3e^{-1}$	$\pm 2e^{-1}$	$\pm 3e^{-2}$	$\pm 1e^{-2}$	$\pm 4e^{-2}$	$\pm 5e^{-4}$	$\pm 4e^{-3}$	$\pm 9e^{-4}$	$\pm 2e^{-2}$
Bathtub	7.56	7.77	7.19	7.84	6.77	0.410	0.382	0.395	0.366	0.480
	$\pm 7e^{-2}$	$\pm 1e^{-1}$	$\pm 5e^{-2}$	$\pm 1e^{-2}$	$\pm 1e^{-1}$	$\pm 7e^{-3}$	$\pm 3e^{-3}$	$\pm 4e^{-3}$	$\pm 1e^{-3}$	$\pm 2e^{-3}$
Bed	7.76	9.07	8.24	7.91	7.24	0.478	0.349	0.449	0.380	0.484
	$\pm 8e^{-2}$	$\pm 1e^{-1}$	$\pm 1e^{-2}$	$\pm 5e^{-2}$	$\pm 1e^{-1}$	$\pm 7e^{-3}$	$\pm 1e^{-2}$	$\pm 9e^{-3}$	$\pm 2e^{-3}$	$\pm 3e^{-3}$
Basket	7.74	8.02	6.74	7.54	6.60	0.365	0.343	0.427	0.361	0.455
	$\pm 1e^{-1}$	$\pm 3e^{-1}$	$\pm 4e^{-2}$	$\pm 2e^{-2}$	$\pm 1e^{-1}$	$\pm 9e^{-3}$	$\pm 5e^{-3}$	$\pm 4e^{-3}$	$\pm 2e^{-3}$	$\pm 3e^{-3}$
Printer	8.36	8.30	8.28	9.66	6.84	0.630	0.622	0.607	0.499	0.705
	$\pm 7e^{-1}$	$\pm 3e^{-1}$	$\pm 2e^{-1}$	$\pm 2e^{-2}$	$\pm 2e^{-1}$	$\pm 4e^{-2}$	$\pm 7e^{-4}$	$\pm 1e^{-2}$	$\pm 1e^{-4}$	$\pm 2e^{-2}$
Inst-Avg	8.60	8.83	8.12	8.56	7.38	0.441	0.387	0.426	0.386	0.498
	$\pm 2e^{-1}$	$\pm 2e^{-2}$	$\pm 7e^{-2}$	$\pm 2e^{-2}$	$\pm 6e^{-2}$	$\pm 2e^{-3}$	$\pm 1e^{-3}$	$\pm 3e^{-3}$	$\pm 1e^{-4}$	$\pm 9e^{-3}$
Cat-Avg	9.09	9.02	8.26	8.90	7.52	0.440	0.386	0.426	0.389	0.495
	$\pm 3e^{-1}$	$\pm 8e^{-2}$	$\pm 8e^{-2}$	$\pm 2e^{-2}$	$\pm 2e^{-2}$	$\pm 3e^{-3}$	$\pm 6e^{-3}$	$\pm 7e^{-3}$	$\pm 3e^{-4}$	$\pm 5e^{-3}$

Table 9: Quantitative comparison with state of the art on real-world ScanNet [10] shape completion for seen categories. We bold the best results and underline the second best results in the table.

	Chamfer Distance ($\times 10^2$) \downarrow					IoU \uparrow				
	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours	3D-EPN [13]	Few-Shot [34]	IF-Nets [6]	Auto-SDF [20]	Ours
Trash Bin	5.03	5.65	5.23	<u>4.48</u>	4.44	0.61	0.70	0.62	0.66	<u>0.68</u>
Chair	9.99	<u>6.88</u>	7.93	6.00	7.14	0.40	<u>0.46</u>	0.43	0.49	0.45
Bookshelf	4.87	<u>4.33</u>	5.17	<u>4.12</u>	3.80	0.53	0.65	0.58	<u>0.61</u>	<u>0.61</u>
Table	8.74	7.13	10.15	<u>6.72</u>	6.60	0.47	<u>0.50</u>	0.46	0.49	0.54
Cabinet	4.60	<u>4.36</u>	5.64	4.53	4.17	0.76	0.80	0.74	0.78	<u>0.79</u>
Sofa	4.94	4.28	7.87	4.58	<u>4.53</u>	0.69	0.75	0.67	0.72	<u>0.73</u>
Monitor	5.75	<u>4.98</u>	6.39	5.92	4.74	0.52	0.59	0.53	0.49	<u>0.56</u>
Inst Avg	7.94	6.18	7.65	5.68	<u>6.02</u>	0.50	0.56	0.51	<u>0.55</u>	<u>0.55</u>
Cat Avg	6.27	5.37	6.91	<u>5.20</u>	5.06	0.57	0.63	0.58	0.61	<u>0.62</u>

D Evaluation on Seen Categories

Table 9 shows the comparisons on seen train categories with state of the art on real-world data from ScanNet [10]. We evaluate 1060 samples for 7 seen categories including: *chair*, *table*, *sofa*, *trash bin*, *cabinet*, *bookshelf*, and *monitor*; categories are selected as those which have more than 50 test samples.

Table 9 shows that our performance on seen categories is on par with state of the art, particularly when evaluating category averages, as our learned multiresolution priors maintain robustness across categories. Note that similar to the previous evaluation, AutoSDF results are reported as the best among their nine predictions with the highest IoU value given an oracle to indicate the best choice. Our method thus achieves performance on par with state of the art on seen categories, and notably improves shape completion for unseen categories.

E Model Architecture Details

Figure 7 details our model architecture. Figure 7 (a), (b) and (c) respectively present the submodule for learning patch priors at resolutions 32^3 , 8^3 , and 4^3 . The network in Figure 7 (d) shows our

multi-resolution patching learning stage. Inputs are partial scans and the learnable shape priors, and the outputs are completed shapes. The specifications of encoder and decoder blocks in these models are shown in Figure 8.

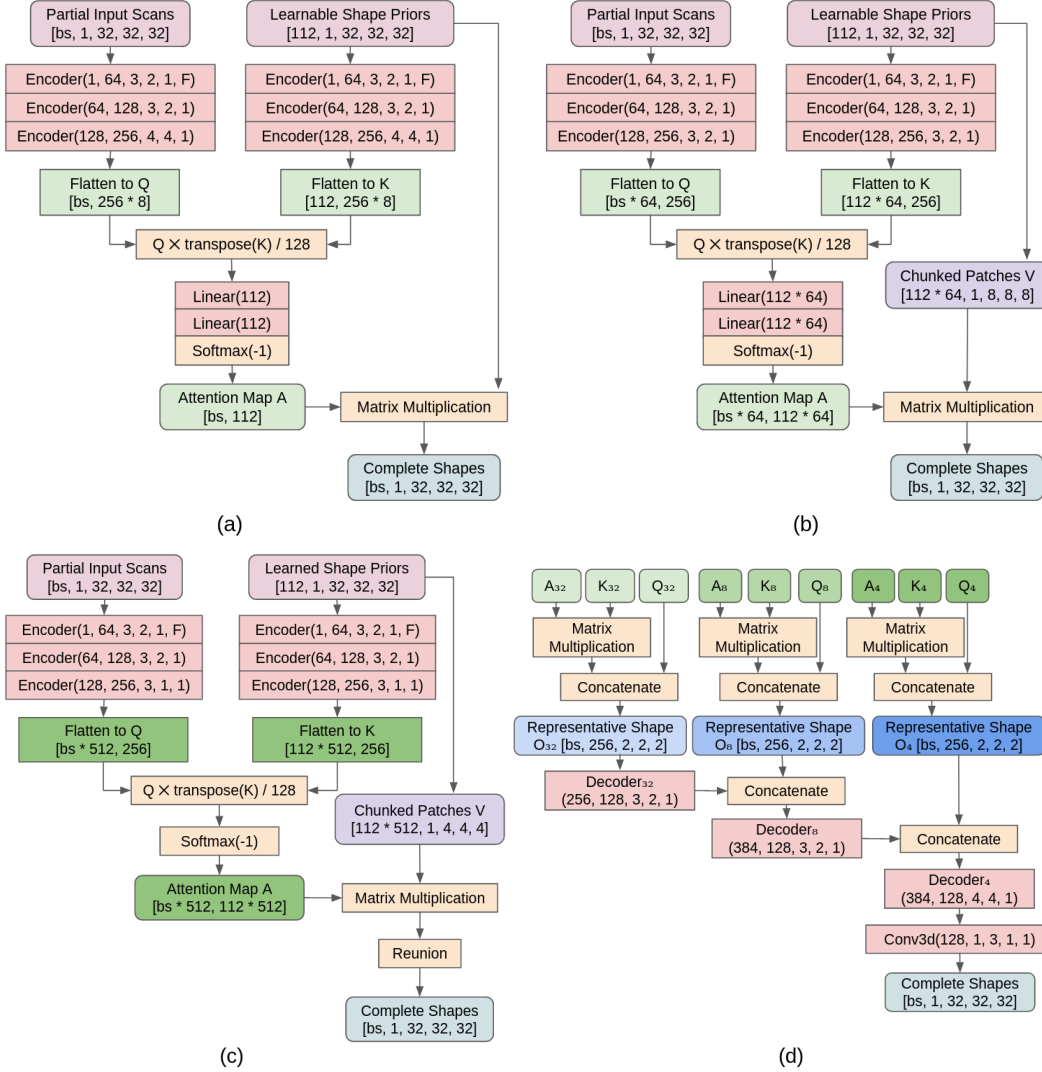


Figure 7: Model specifications in our method. (a) represents the patch learning model structure for resolution at 32^3 ; (b) represents the patch learning model structure for resolution at 8^3 ; (c) represents the patch learning model structure for resolution at 4^3 ; (d) represents the multi-resolution model structure. In figure (d), A_i represents the obtained attention map, Q_i represents the input local features, and K_i represents the learned prior patch features, where i is the resolution for patch learning model, $i = 32, 8, 4$.

F Additional Qualitative Results

Figure 9 shows more examples for qualitative results on ShapeNet, and Figure 10 shows more examples for qualitative results on ScanNet scans.

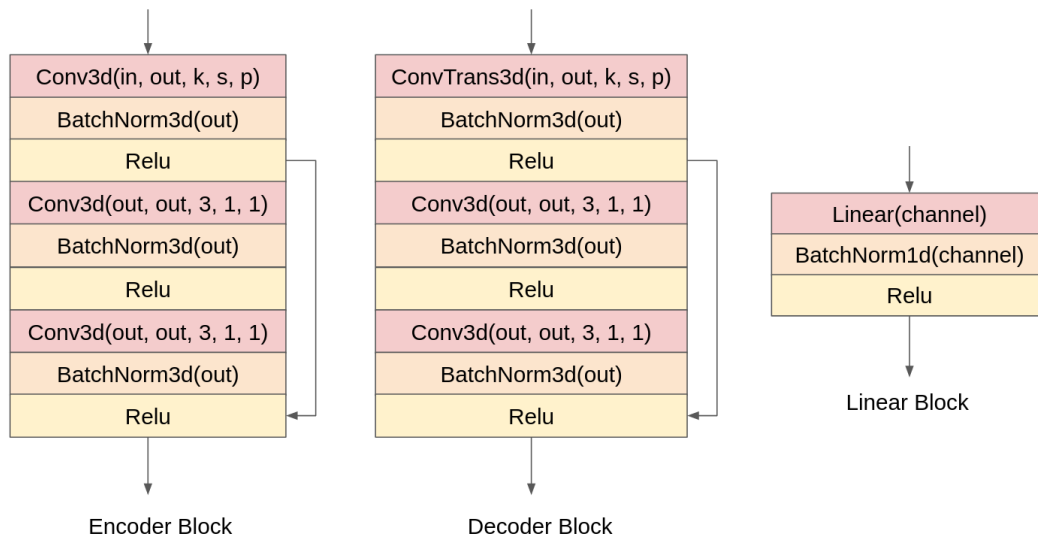


Figure 8: Layer specifications in our model. During the process of learning patch priors in a single resolution, we use encoder blocks to encode partial input scans and learnable shape priors to local features, and then use linear blocks to post-process the obtained attention map. The decoder block is used for decoding complete shapes in a multi-resolution patch learning module.




Figure 9: Qualitative comparisons with state of the art on ShapeNet [4].




Figure 10: Qualitative comparisons with state of the art on ScanNet [10].