

# PT43D: A Probabilistic Transformer for Generating 3D Shapes from Single Highly-Ambiguous RGB Images

Yiheng Xiong  
yiheng.xiong@tum.de  
Angela Dai  
angela.dai@tum.de

3D AI Lab  
Technical University of Munich  
Munich, Germany




Figure 1: Our approach seeks to model the distribution of 3D shapes in latent space conditioned on a single highly ambiguous image, enabling the sampling of multiple diverse hypotheses during inference. Colored points are for visualization purposes only, indicating the overlap between ground-truth visible parts and plausible shapes.

## Abstract

Generating 3D shapes from single RGB images is essential in various applications such as robotics. Current approaches typically target images containing clear and complete visual descriptions of the object, without considering common realistic cases where observations of objects that are largely occluded or truncated. We thus propose a transformer-based autoregressive model to generate the probabilistic distribution of 3D shapes conditioned on an RGB image containing potentially highly ambiguous observations of the object. To handle realistic scenarios such as occlusion or field-of-view truncation, we create simulated image-to-shape training pairs that enable improved fine-tuning for real-world scenarios. We then adopt cross-attention to effectively identify the most relevant

region of interest from the input image for shape generation. This enables inference of sampled shapes with reasonable diversity and strong alignment with the input image. We train and test our model on our synthetic data then fine-tune and test it on real-world data. Experiments demonstrate that our model outperforms state of the art in both scenarios<sup>1</sup>.

## 1 Introduction

Generating 3D shapes from single RGB images presents a fundamental challenge in computer vision, with practical applications such as single-view 3D reconstruction, where the objective is to reconstruct the 3D structure and geometry of the objects in an observed image. To tackle this challenge, researchers have heavily relied on data-driven deep learning techniques. Earlier approaches [8, 20, 26, 41, 42, 44] focus on generating a deterministic 3D output based on an RGB image, achieving impressive results, but without considering inherent ambiguities that often arise from monocular perception. For example, an image of a chair from an upper view does not reveal the shape of its legs (see Fig. 1 top part). In contrast, recent advancements in generative AI [10, 13, 27] have led to approaches [3, 21, 31, 43] that aim to learn the entire distribution of 3D shapes conditioned on the input image, enabling the sampling of multiple 3D hypotheses at test time. Nevertheless, these approaches target scenarios where images contain a fairly clear, complete visual description of the object, while many real-world scenarios contain observations of objects that are highly occluded or truncated. For instance, a home robot operating in cluttered environments may encounter situations where parts of objects are hidden from view due to limited field-of-view of the camera or obstacles in the scene. In such cases, relying solely on methods designed for complete and unobstructed observations may lead to inaccurate understanding of the environment, limiting the robot’s ability to make informed decisions and navigate effectively.

To this end, we propose a transformer-based [25] autoregressive model to generate the probabilistic distribution of 3D shapes conditioned on a single RGB image containing potentially highly ambiguous observations of the object. To handle realistic scenarios such as occlusion or field-of-view truncation, we create simulated image-to-shape training pairs from ShapeNet [2] that enable improved fine-tuning for real-world scenarios. During training, each image can be mapped to potentially multiple ground-truth shapes due to inherent ambiguity in the input image observation. We then adopt cross-attention to effectively identify the most relevant region of interest from the input image for shape generation. This enables inference of sampled shapes that meet two fundamental criteria: reasonable diversity and strong alignment with the input image. We train and test our model on our synthetic training pairs then fine-tune and test it on real-world data from ScanNet [4, 6]. Experiments demonstrate that our model outperforms state of the art in both scenarios.

We summarize our contributions as follows:

- We propose the first probabilistic approach for 3D shape generation from a highly ambiguous RGB image, leveraging cross-attention to effectively identify the most relevant region of interest from the input image, such that diverse output hypotheses all align well with the input image.
- We introduce a synthetic data augmentation approach to enable effective multi-hypothesis generation by setting one image maps to potentially multiple ground-truth shapes dur-

<sup>1</sup>Our code is open sourced at <https://github.com/xiongyiheng/PT43D>.

ing training. Our approach with synthetic data enables pre-training that notably improves shape reconstruction results when fine-tuned on real-world data.

## 2 Related Work

**Single-View Object Reconstruction.** Inferring 3D shapes from monocular views is an inherently ill-posed task. Early approaches generate single 3D outputs from single RGB images in a deterministic manner, utilizing various representations such as voxels [6, 9, 26, 30, 39, 41, 42], point clouds [7, 19, 40], meshes [37, 38], and CAD retrieval [8, 17, 18]. More recent approaches also focus on implicit representations such as SDFs [16, 44], UDFs [4], and closest surface points [56], achieving impressive reconstruction fidelity. However, these approaches do not address the inherent ambiguity in a single RGB image, which often results in the co-existence of multiple plausible reconstructions that explain the input image. Our method tackles this ambiguity by treating image-based predictions as conditional distributions, allowing for generation of multiple plausible sampled reconstructions.

**Probabilistic 3D Shape Generation.** More recently, researchers have explored the use of generative neural networks [10, 13, 27, 32, 33] to model the probabilistic distribution of 3D shapes, enabling the sampling of multiple hypotheses from the learned distribution. For instance, [15, 24, 28, 45, 46] utilize transformer-based architecture to generate shapes, achieving impressive results. With recent advancements in denoising diffusion models, [11, 14, 23, 47] generate high-quality shapes by iteratively refining noise-corrupted data through a reverse diffusion process guided by learned probability distributions. For image-conditioned shape generation, AutoSDF [21] models shape distribution via an autoregressive process and combines it with a separately-trained image-conditioning model to generate shapes based on single RGB images. SDFusion [9] learns to sample from a target shape distribution by reversing a progressive noise diffusion process, generating single-view reconstructions via classifier-free image guidance [12]. These approaches typically focus on images with fairly clear and complete visual descriptions of the object, but unfortunately do not address common real-world scenarios where objects are captured under strong occlusion or truncation. Our method, instead, takes advantage of generative neural networks, analyzing common realistic scenarios where images are highly occluded or truncated.

## 3 Methodology

Our transformer-based approach seeks to model the distribution that represents potential 3D shapes conditioned on the input image. Given a shape  $X$  represented via a T-SDF, we first follow the P-VQ-VAE [22] to compress it to a low-dimensional and discrete-grid representation  $E$ . Each grid holds a 3D latent feature index  $q_i$  corresponding to the shape feature  $e_i$  from a codebook  $Z$  that is used for indexing feature embeddings. Then we learn the conditional distribution of shapes based on the input image  $c$  over this compressed latent representation.

### 3.1 Shape Compression and Discretization

For each shape  $X$ , a shape encoder  $E_\phi(\cdot)$  encodes its Truncated Signed Distance Field (T-SDF) into a low-dimensional 3D grid feature  $E \in g^3 \times D$ , where  $g$  represents the resolution of the 3D grid, and  $D$  denotes the dimension of the 3D grid features. We establish a discrete

latent space containing  $K$  embedding vectors, each with a dimension of  $D$ . For the shape feature  $e_i$  at grid position  $i$ , a vector quantization operation  $VQ(\cdot)$  is employed to find its nearest neighbor  $e'_i$  from the  $K$  embedding vectors in the discrete latent space. The feature input into the shape decoder  $D_\phi(\cdot)$  is the 3D discrete latent feature grid  $E' \in g^3 \times D$ , comprising discrete features  $e'_i$  at grid locations  $i$ . The formula is:

$$E = E_\phi(X), E' = VQ(E), X' = D_\phi(E'). \quad (1)$$

where  $X'$  is the reconstructed shape, and  $\phi$  is the set of learnable parameters of the decoder. As a result, a shape  $X$  can be represented as a 3D latent feature index grid  $Q \in g^3$ , where  $q_i$  at each location corresponds to a discrete shape feature  $e'_i$ . In particular, we adopt Patch-wise encoding VQ-VAE (P-VQ-VAE) [24] to learn this process since it allows for a better representation for local details. Given the 3D T-SDF  $X$  of the shape, we first split  $X$  into  $N$  patches of  $X_p$ . Then  $E_\phi$  encodes each patch independently and  $D_\phi$  decodes all patches jointly. Empirically, we set  $K = 512$ ,  $D = 256$ ,  $g = 8$  and  $N = 512$ . Training the network uses a combination of three losses proposed by van den Oord *et al.* [24]: reconstruction loss, the vector quantization objective and the commitment loss.

## 3.2 Method Overview

Given an input image  $c$ , an empty query sequence  $s_q$  in the length of  $n + 1$  and a start token  $\langle \text{sos} \rangle$ , our model generates the shape distribution over the latent space  $E'$  by filling the query sequence autoregressively. The complete sequence  $s_c$  comprises of discrete grids, where each location  $i$  contains the distribution  $p_i(e')$  of 3D latent features. During inference, we can sample from  $p_i(e')$  to acquire shape feature index  $q_i$  for each location and further get the corresponding shape feature  $e'_i$  from the codebook  $Z$ . Finally, a pretrained decoder  $D_\phi$  utilizes these shape features to produce the reconstruction output. The overall framework, as displayed in Fig 2, comprises three key components: the image encoder, the conditional cross-attention module and the autoregressive modeling with transformer.

**Image Encoder.** The role of the image encoder is important in the context of image-based 3D shape generation. An effective image encoder possesses the capability to extract essential features from images, playing a crucial role in facilitating the subsequent 3D shape generation process. We adopt a pretrained "ViT-B/32" CLIP [25] to extract the input image features, chosen for its robust capacity to capture rich image representations.

**Conditional Cross-Attention Module.** The conditional module establishes a connection between the image encodings and the grid sequence  $s_t$  at current autoregressive step  $t$ , ensuring that the output grid sequence incorporates the essential image features for subsequent operations. Fig. 3 depicts the architecture of it.  $s_t$  comprises of the start token  $\langle \text{sos} \rangle$ , potentially filled grids containing latent features, and the target grid with its position embedding. We perform cross-attention between  $s_t$  and the image encodings, where each grid in  $s_t$  serves as a query. The output sequence after cross-attention is essentially a weighted sum of image features. Then we project it to the original dimension and concatenate it with original embeddings.

**Autoregressive Modeling with Transformer.** After obtaining the sequence that contains the necessary image features at step  $t$ , we feed it into the transformer and predict the feature index  $\hat{q}_i$  for current target location  $l_t$ . We then fill current target location  $l_t$  with the latent feature corresponding to the predicted feature index  $\hat{q}_i$ . We learn this model by maximizing the log-likelihood of the latent representations and perform this process (including conditional cross-attention step) autoregressively until all locations are filled.

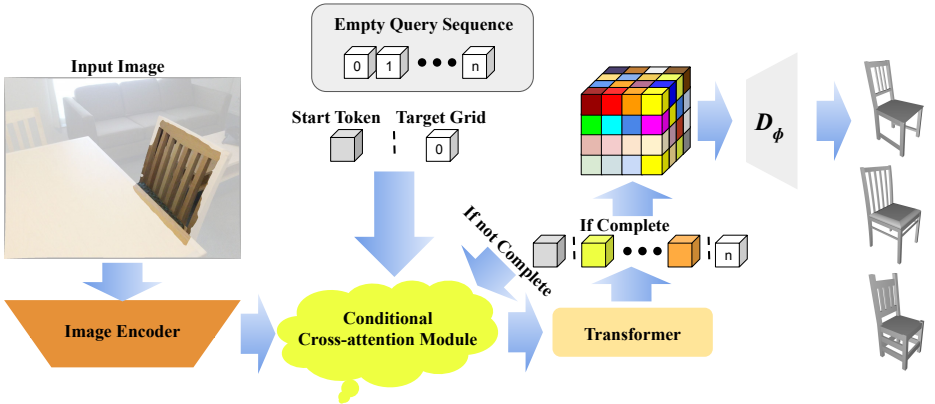


Figure 2: **Method Overview.** Our model adopts a low-dimensional and discretized latent space for shape generation. Given an input RGB image, it is first processed by an image encoder. Then the extracted image encodings are fed into the conditional module where they attend with the input sequence consisting of a start token, potentially filled grids containing latent features, and a target grid containing its position embedding via cross-attention such that the output sequence contains essential image features. Then the transformer processes this sequence and predicts the latent feature for the target grid. Conditional cross-attention and transformer are adopted autoregressively until a complete sequence is output where each cell contains the conditional distribution of latent features. During inference, we can gather the complete sequence into grids then sample and decode to multiple plausible hypotheses.

### 3.3 Training and Inference

**Training.** During training, we randomly select one view from the rendered images as input and sample one shape from the corresponding mappings as the ground-truth. We employ teacher-forcing by setting filled locations as ground-truth features during autoregressive learning. We also fine-tune the pretrained CLIP encoder to make it more adaptable to our task. The loss function is the cross entropy with respect to feature indexes for each location.

**Inference.** During inference, we complete an empty query sequence conditioned on the image in an autoregressive manner from the transformer. At step  $t$ , we have:

$$p_{l_t}(e^l) = T_\theta(\{e_{l_{t-1}}^{\wedge}, e_{l_{t-2}}^{\wedge}, e_{l_{t-3}}^{\wedge}, \dots\}, l_t, c). \quad (2)$$

where  $l_t$  represents current location and  $e_{l_t}^{\wedge}$  is generated by sampling from the distribution  $p_{l_t}(e^l)$ . We then feed the complete sequence into  $D_\phi$  to generate one plausible 3D shape.

**Implementation Details.** To train our network, we use ADAM as the optimizer with an initial learning rate of 1e-5 for the CLIP image encoder and 1e-4 for other components. We decay both learning rates by multiplying 0.9 every 10,000 iterations. The batch size is fixed to 10 and it takes around 72 hours to converge on a single GTX 1080Ti GPU.

## 4 Multi-Hypothesis Data Augmentation

As discussed in Sec. 1, we intend to create simulated image-to-shape training pairs that further enable improved fine-tuning for real-world scenarios. To this end, we render CAD models from ShapeNetCore [2] in tailored settings such that objects are captured under strong

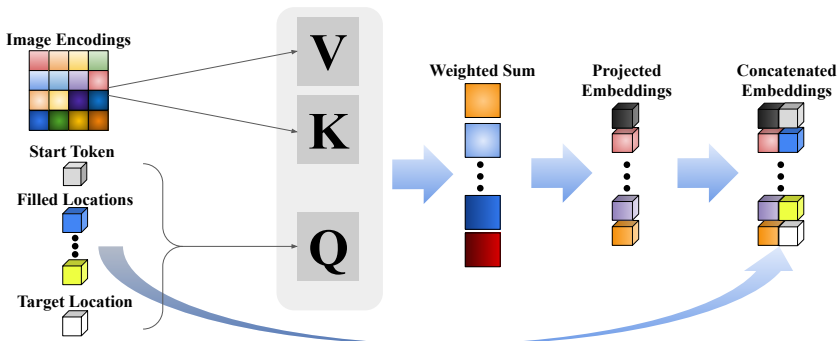


Figure 3: **Identifying Most Relevant Regions in the Image via Conditional Cross-Attention.** We perform cross-attention between image encodings and the input sequence to effectively identify the most relevant region of interest from the input image. The input sequence consisting of a start token, potentially filled locations with their own predicted latent features, and a target location with its position embedding, is multiplied with a weight matrix  $Q$  and the image encodings are multiplied with weight matrices  $V$  and  $K$  independently. The resulting sequence contains a weighted sum of image features. We then project it to the original dimension and concatenate with the original sequence embeddings.

occlusion or truncation. We focus on common indoor furniture models including *chair*, *table*, *cabinet*, *bed*, and *bookshelf*. Due to inherent ambiguities in our images, we also establish ground-truth mappings using part-level annotations from PartNet [27], which allows for multiple distinct ground-truth shapes per image.

#### 4.1 Rendering CAD Models with Ambiguity

**Occlusion.** To induce occlusion, we first randomly rotate the target object at the origin along the upward axis to capture varying levels of self-occlusion. Then, we randomly select another object to act as the occluding object, positioned between the virtual camera and the target object. This setup ensures two conditions: (1) no overlap between the occluding and target objects in 3D space, and (2) reasonable occlusion of the target object by the occluding one in the 2D image. Fig. 4(a) provides an example.




Figure 4: **Different Types of Ambiguity.** (a) indicates ambiguity caused by occlusion. (b) and (c) represent ambiguity resulted from truncation via limited field-of-view of the camera.

**Truncation.** Truncation typically occurs when parts of an object are obscured by the limited field-of-view of the camera. We consider two scenarios: First, we move the object away from the center of the image plane as shown in Fig. 4(b). Second, we simulate scenarios

where the camera is too close to the object, as displayed in Fig. 4(c).

## 4.2 Creating Multi-Hypothesis Image-to-Shape Pairs

The inherent ambiguity in our renderings allows for the coexistence of multiple viable 3D ground-truth shapes, each corresponding to a plausible interpretation of the image’s content. Therefore, we take this fact into account to create our image-to-shape training pairs. We first classify available CAD models into similar groups. Then we construct the ground-truth mappings for each image from these groups by considering the geometry of visible parts and per-pixel part labels. Further details can be found in the supplementary material. Fig. 5 shows one sample mapping. On average, each rendering maps to two ground-truth shapes accurately, with  $l_2$  chamfer distance of visible parts below 0.04.




Figure 5: **Ground-Truth Mapping.** (a) represents a sample rendering and (b), (c) and (d) represent multiple distinct ground-truth shapes that corresponds to (a). Yellow points are for visualization purposes only, indicating the overlap between ground-truth visible parts and multiple plausible shapes.

## 5 Experiments

We first train and test our model on our synthetic dataset following splits provided by ShapeNet [2]. We then fine-tune the CLIP image encoder and conditional cross-attention module of our pretrained model on real-world data from ScanNet [10, 6]. We use the ScanNet25k image data and use ground-truth semantic instance masks to generate per-instance images for training without background. While for validation set, we adopt machine-generated masks produced by ODISE [13]. Then we align each instance image with ShapeNet CAD models using annotations provided by Scan2CAD [10]. We split the generated images according to provided train/validation scenes.

### 5.1 Evaluation Metrics

Our aim is to generate shapes that capture the distribution of possible reconstructions that explain an input image observation. Thus, we evaluate both diversity of reconstructed shapes as well as their reconstruction quality. To assess the diversity in generated shapes, we employ the Total Mutual Difference (TMD), computed as the average chamfer distance among  $k$  generated shapes. We set  $k = 6$  in our experiments, aligning with the average number of models within a similar group (refer to Sec. 4). For evaluating the generation quality, we utilize bidirectional  $l_2$  chamfer distance (CD) and F-score@1% [29] as metrics. We present the minimal CD and the highest F-score among the  $k$  generated shapes in our results.



## 5.2 Comparison to State of the Art

	Chair			Table			Cabinet			Bed			Bookshelf			Average		
	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$
Pix2Vox [10]	-	5.99	0.175	-	7.03	0.164	-	6.52	0.114	-	11.8	0.115	-	9.85	0.125	-	8.24	0.139
AutoSDF [21]	0.041	3.19	0.250	0.040	3.81	0.280	0.047	4.70	0.160	0.054	8.78	0.135	0.042	4.47	0.247	0.045	5.04	0.215
SDFusion [9]	<b>0.073</b>	3.32	0.233	<b>0.078</b>	3.93	0.249	<b>0.103</b>	4.77	0.145	0.109	<b>5.89</b>	0.119	<b>0.093</b>	4.69	0.226	<b>0.091</b>	4.52	0.194
Ours	0.052	<b>2.89</b>	<b>0.260</b>	0.058	<b>3.55</b>	<b>0.284</b>	0.066	<b>3.76</b>	<b>0.184</b>	0.075	7.54	<b>0.141</b>	0.061	<b>4.31</b>	<b>0.254</b>	0.062	<b>4.41</b>	<b>0.224</b>

Table 1: **Results Comparisons on Synthetic Data.** Our method generates higher quality shapes compared with other baselines in terms of CD and F-score, with reasonable diversity.

	ShapeNet [11]	ScanNet [12, 13]
SDFusion [9]	6.71	10.4
Ours	<b>5.41</b>	<b>5.89</b>

Table 2: **Average One-Way  $l_2$  Chamfer Distance between Visible Ground-Truth Points with Generated Hypotheses.** Our method generates shapes that align better with the input image on both synthetic and real-world dataset.

**Results on Synthetic Data.** Tab. 1 compares our method with state-of-the-art single-image based methods Pix2Vox [10], AutoSDF [21] and SDFusion [9] on synthetic dataset. Pix2Vox is fully deterministic, and so only provides one hypothesis. Our method is able to generate higher quality shapes compared with other baselines. In terms of generation diversity, SDFusion is capable of generating more diverse shapes. However, the higher diversity comes at the cost of outputs that do not align well with the input data. To validate this, we calculate the average one-way  $l_2$  chamfer distance between visible ground-truth points and generated hypotheses. The results, presented in Tab. 2, illustrate that our model produces shapes with better alignment compared to those generated by SDFusion. This suggests the lacking of plausibility among higher diverse shapes produced by SDFusion, thereby demonstrating more reasonable diversity among hypotheses produced by our method.

In Fig. 6 top part, we present qualitative comparisons of synthetic data among AutoSDF, SDFusion and our method. The visual inspection clearly demonstrates that our method excels in generating higher-fidelity shapes from a single highly ambiguous RGB image compared to other baselines, with a more reasonable diversity level.

	Chair			Table			Cabinet			Bed			Bookshelf			Average		
	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$
Pix2Vox [10]	-	7.01	0.158	-	11.1	0.157	-	12.8	0.080	-	10.2	0.161	-	7.14	0.119	-	9.68	0.135
AutoSDF [21]	0.019	5.33	0.174	0.019	9.46	0.169	0.030	10.9	0.113	0.019	8.24	0.196	0.020	7.59	0.170	0.021	8.32	0.164
SDFusion [9]	<b>0.104</b>	4.52	0.197	<b>0.152</b>	5.66	0.207	<b>0.141</b>	6.37	0.166	<b>0.135</b>	4.54	0.216	<b>0.129</b>	3.75	0.185	<b>0.132</b>	4.97	0.200
Ours	0.077	<b>3.69</b>	<b>0.222</b>	0.096	<b>3.64</b>	<b>0.234</b>	0.130	<b>4.39</b>	<b>0.183</b>	0.096	<b>3.17</b>	<b>0.319</b>	0.081	<b>3.18</b>	<b>0.207</b>	0.096	<b>3.61</b>	<b>0.228</b>

Table 3: **Results Comparisons on Real-World Data.** Our method generates higher quality shapes compared with other baselines in terms of CD and F-score, with reasonable diversity.

**Results on Real-World Data.** We then compare the experimental results obtained from ScanNet [11, 12] images. As shown in Tab. 3, our model surpasses state of the art in terms of generation quality. Additionally, the results in Tab. 2 indicate that SDFusion [9] produces nearly twice the visible chamfer distance compared to ours. This further emphasizes our model’s capability to generate shapes with more reasonable diversity.

In Fig. 6 bottom part, we present qualitative comparisons of real-world data among AutoSDF, SDFusion and our method. Our method achieves higher shape generation quality






Figure 6: **Qualitative Comparisons.** Top: results on synthetic images; bottom: results on real images. Our method generates higher quality and more plausible hypotheses.

compared to other baselines. Despite the higher diversity among hypotheses produced by SDFusion, they lack alignment with the input image.

### 5.3 Ablation Studies

	TMD $\uparrow$	CD $\downarrow$	F-score $\uparrow$
w/o mapping	0.060	<b>4.26</b>	<b>0.230</b>
w/ coarse mapping	0.066	4.43	0.216
average pooling	0.065	4.53	0.221
max pooling	0.066	4.50	0.219
w/o concatenation	<b>0.105</b>	39.1	0.032
Ours	0.062	4.41	0.224

Table 4: **Ablations on Synthetic Data.** Our model generates more diverse shapes being trained with multiple ground-truth mappings. Cross-attention is effective in generating higher quality shapes compared with average pooling and max pooling. Without final concatenation, very little is learned by the model.

**Does cross-attention help?** In our model, we employ cross-attention as the image conditional mechanism, indicating a weighted sum of image features. We compare it with two similar strategies: (1) average pooling, where all weights are the same, and (2) max pooling, where only one non-zero weight is used. Results in Tab. 4 demonstrate that our model with cross-attention generates higher-quality shapes compared with two other strategies.

**Does concatenation help?** Within our conditional cross-attention module, the raw cross-attention outputs merely contain image features. We thus concatenate the original 3D embeddings with them in the final step. To analyze its effect, we conduct an additional experiment by removing the concatenation step, using raw cross-attention outputs for subsequent

	TMD↑	CD↓	F-score↑
from scratch	0.044	6.87	0.203
Ours	<b>0.096</b>	<b>3.61</b>	<b>0.228</b>

Table 5: **Ablations on Real-World Data.** Our multi-hypothesis data augmentation with synthetic data enables pre-training that notably improves shape reconstruction results when fine-tuned on real-world data.

operations. Results in Tab. 5 reveal that the model without concatenation learns very little.

**Does one image mapping to multiple ground-truths improve generation diversity?** To analyze our mapping strategy, we conduct two additional experiments: training our model (1) without mapping and (2) with coarse mapping by setting all shapes within a similar group as ground-truths. Results in Tab. 4 show that our method trained with mapping achieves greater diversity in generated shapes compared to the version without mapping. Furthermore, opting for coarse mapping further increases diversity, demonstrating our method’s ability to capture the varying diversity of ground-truth data.

**Does pretraining on the synthetic dataset improve real-world reconstruction results?**

To assess the efficacy of our synthetic dataset, we conduct experiments by training our model from scratch on real-world data and comparing the reconstruction results with our fine-tuned version. Results from Tab. 5 demonstrate the superior performance of our fine-tuned version, thereby affirming the effectiveness of our synthetic dataset.

## 5.4 Limitations

In our image-based shape generation framework, we utilize patch-wise image encodings as conditional input, enabling robust global alignment between input object observations and output reconstructions. However, this may not fully capture the shape variance and ambiguities that lie in the local details, such as the wheels of the chair present in Fig. 5. A potential approach to overcome this limitation involves adopting pixel-wise image features that are more locally extracted. In addition, while our method achieves promising reconstruction results at object-level, generating scene-level reconstructions could provide richer perception for applications like robotics. In our future work, we consider expanding the generation scale from object-level to scene-level, leveraging monocular views.

## 6 Conclusion

In this work, we propose a new approach for generating the probabilistic distribution of 3D shape reconstructions conditioned on a highly ambiguous RGB image. We employ cross-attention to effectively identify the most relevant region of interest from the input image for shape generation. We also introduce a synthetic data augmentation approach to enable effective multi-hypothesis generation, by computing multiple potential ground-truth shapes that correspond to a single image observation during training. This synthetic augmentation enables pre-training that notably improves shape reconstruction results when fine-tuned on real-world data. Experiments demonstrate that our model outperforms state of the art in both generation quality and plausible generation diversity. We hope our work will draw increased attention to this important and realistic problem, highlighting the need for innovative solutions in 3D shape reconstruction from single ambiguous image data.

## 7 Acknowledgements

This work was supported by the ERC Starting Grant SpatialSem (101076253).

## References

- [1] Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 2614–2623, 2019.
- [2] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [3] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465, 2023.
- [4] Julian Chibane, Gerard Pons-Moll, et al. Neural unsigned distance fields for implicit function learning. *Advances in Neural Information Processing Systems*, 33:21638–21652, 2020.
- [5] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 628–644. Springer, 2016.
- [6] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [7] Haoqiang Fan, Hao Su, and Leonidas J Guibas. 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*, pages 605–613, 2017.
- [8] Daoyi Gao, Dávid Rozenberszki, Stefan Leutenegger, and Angela Dai. Diffcad: Weakly-supervised probabilistic cad model retrieval and alignment from an rgb image. *arXiv preprint arXiv:2311.18610*, 2023.
- [9] Rohit Girdhar, David F Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VI 14*, pages 484–499. Springer, 2016.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

- [11] Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Oğuz. 3dgen: Triplane latent diffusion for textured mesh generation. *arXiv preprint arXiv:2303.05371*, 2023.
- [12] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [14] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [15] Moritz Ibing, Gregor Kobsik, and Leif Kobbelt. Octree transformer: Autoregressive 3d shape generation on hierarchically structured sequences, 2021.
- [16] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfd: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1251–1261, 2020.
- [17] Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. Mask2cad: 3d shape prediction by learning to segment and retrieve. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 260–277. Springer, 2020.
- [18] Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. Patch2cad: Patch-wise embedding learning for in-the-wild shape retrieval from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12589–12599, 2021.
- [19] Priyanka Mandikal, KL Navaneet, Mayank Agarwal, and R Venkatesh Babu. 3d-lmnet: Latent embedding matching for accurate and diverse 3d point cloud reconstruction from a single image. *arXiv preprint arXiv:1807.07796*, 2018.
- [20] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019.
- [21] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 306–315, 2022.
- [22] Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 909–918, 2019.
- [23] Shentong Mo, Enze Xie, Ruihang Chu, Lanqing Hong, Matthias Niessner, and Zhenquo Li. Dit-3d: Exploring plain diffusion transformers for 3d shape generation. *Advances in Neural Information Processing Systems*, 36, 2024.

- [24] Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. Polygen: An autoregressive generative model of 3d meshes. In *International conference on machine learning*, pages 7220–7229. PMLR, 2020.
- [25] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [26] Stephan R Richter and Stefan Roth. Matryoshka networks: Predicting 3d geometry via nested shape layers. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1936–1944, 2018.
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [28] Yawar Siddiqui, Antonio Alliegro, Alexey Artemov, Tatiana Tommasi, Daniele Sirigatti, Vladislav Rosov, Angela Dai, and Matthias Nießner. Meshgpt: Generating triangle meshes with decoder-only transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19615–19625, 2024.
- [29] Maxim Tatarchenko, Stephan R. Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn?, 2019.
- [30] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2626–2634, 2017.
- [31] Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, Karsten Kreis, et al. Lion: Latent point diffusion models for 3d shape generation. *Advances in Neural Information Processing Systems*, 35:10021–10039, 2022.
- [32] Aaron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, Koray Kavukcuoglu, et al. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 12, 2016.
- [33] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016.
- [34] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning, 2018.
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- [36] Rahul Venkatesh, Tejan Karmali, Sarthak Sharma, Aurobrata Ghosh, R Venkatesh Babu, László A Jeni, and Maneesh Singh. Deep implicit surface point prediction networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12653–12662, 2021.
- [37] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 52–67, 2018.
- [38] Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. 3dn: 3d deformation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1038–1046, 2019.
- [39] Jiajun Wu, Yifan Wang, Tianfan Xue, Xingyuan Sun, Bill Freeman, and Josh Tenenbaum. Marrnet: 3d shape reconstruction via 2.5 d sketches. *Advances in neural information processing systems*, 30, 2017.
- [40] Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen. Pq-net: A generative part seq2seq network for 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 829–838, 2020.
- [41] Haozhe Xie, Hongxun Yao, Xiaoshuai Sun, Shangchen Zhou, and Shengping Zhang. Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2690–2698, 2019.
- [42] Haozhe Xie, Hongxun Yao, Shengping Zhang, Shangchen Zhou, and Wenxiu Sun. Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images. *International Journal of Computer Vision*, 128(12):2919–2935, 2020.
- [43] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023.
- [44] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *Advances in neural information processing systems*, 32, 2019.
- [45] Xingguang Yan, Liqiang Lin, Niloy J. Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Shapeformer: Transformer-based shape completion via sparse representation, 2022.
- [46] Biao Zhang, Matthias Nießner, and Peter Wonka. 3dilig: Irregular latent grids for 3d generative modeling. *Advances in Neural Information Processing Systems*, 35:21871–21885, 2022.
- [47] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023.
- [48] Xinyang Zheng, Yang Liu, Pengshuai Wang, and Xin Tong. Sdf-stylegan: Implicit sdf-based stylegan for 3d shape generation. In *Computer Graphics Forum*, volume 41, pages 52–63. Wiley Online Library, 2022.