

CPG3D : Cross-modal Priors Guide the 3D Object Reconstruction

Weizhi Nie, Chuanqi Jiao, Rihao Chang, Lei Qu, An-An Liu *Senior Member, IEEE*

Abstract—Three-dimensional reconstruction is a multimedia technology widely used in computer-aided modeling and 3D animation. Nevertheless, it is still hard for reconstruction methods to overcome the 3D geometry missing and the object occlusion in the single-view images. In this paper, we propose a novel method (CPG3D) for reconstructing high-quality 3D shapes from a single image under the guidance of prior knowledge. Using the single-view image as the query, prior knowledge is collected from public 3D datasets, which can compensate for missing 3D geometries and assist the 3D reconstruction network to high fidelity results. Our method consists of three parts: 1) Cross-modal 3D shape retrieval module: This part retrieves related 3D shapes based on 2D images. Here, we apply the pre-trained model to guarantee the correlation between the retrieved 3D shape and the input image. 2) Multimodal information fusion module: We propose a multimodal attention mechanism to handle the information fusing of 2D visual and 3D structural information; 3) Three-dimensional reconstruction module: We propose a novel encoder-decoder network for 3D shape reconstruction. Specifically, we employ the skip connection operation to link the target image’s visual information with the 3D model’s structural information to enhance the prediction of 3D details. During training, we employ two carefully designed loss functions to lead the multimodal learning to obtain proper modal features. On the ShapeNet and Pix3D datasets, the final experimental results reveal that our method notably increases reconstruction quality and outperforms SOTA methods.

Index Terms—3D model reconstruction, Multimodal learning, Cross-modal retrieval.

I. INTRODUCTION

Three-dimensional object reconstruction from a single-view RGB image is a widely researched multimedia task. It has been already applied in computer-assisted modeling [1], [2], 3D animation [3], [4] and robot localization [5], [6]. Unlike humans, who can easily infer the 3D shape of an object from a single image due to previously learned knowledge and an innate ability for visual understanding, computer vision systems are unable to directly reconstruct the 3D shapes from a single image because of the lack of structural information. Hence, the assistance of shape priors and the ability of shape reasoning play an indispensable and important role in 3D object reconstruction.

Lots of single-view 3D object reconstruction methods [7]–[9] rely on the massive training on large-scale datasets, which

maps 2D image features into the 3D shape feature space through convolutional neural networks. However, this learning manner constrains their reconstruction ability due to the information loss in the single-view image. Although these methods have achieved some promising results on the views of 3D objects with clean backgrounds, they still fail to produce credible reconstruction results on views with complex backdrops, such as occlusion and truncation. This condition presents a new challenge to single-view 3D object reconstruction.

A. Motivation

Prior knowledge [10]–[13] is an effective form to compensate for the 3D geometry loss, which provides the network with multimodal input for a comprehensive understanding of the target object. Consequently, the pipeline of prior-guided 3D reconstruction can be redesigned as: given a single-view image, computer vision systems should be able to automatically associate the shape prior knowledge and utilize this prior knowledge to guide the shape reconstruction. Thus, in this paper, single-view reconstruction is divided into two tasks. 1) Obtaining the prior knowledge from the target image: Humans proactively learn the knowledge and memorize it as priors, while computers similarly obtain their priors from memory units. Therefore, this task constructs prior memories and retrieves the desired knowledge. 2) Reconstructing high-quality 3D shapes under the guidance of prior knowledge: Humans can infer the 3D structure of objects depicted in images based solely on prior knowledge. In a similar approach, computers should combine previously acquired priors with image data and infer the desired shape in a learnable manner. Consequently, this paper focuses on efficient methods of information integration and shape reconstruction from the integrated features.

Based on previous analyses, we propose a novel 3D object reconstruction method that recovers the desired shape from a combined multimodal representation. First, to obtain appropriate shape priors, we utilize a cross-domain retrieval method that bridges the gap between images and 3D shapes. Given the input image, this method returns 3D shapes describing structurally homologous objects. For instance, when the network receives an image of a boss chair, the retrieval method will provide 3D chairs with structures such as large backrests and wheeled five-star feet. Second, we need to utilize the retrieved 3D shapes to make up for the missing structural information of the 2D target image. We first extract the features of images and 3D shapes under the supervision of carefully designed domain-specific loss functions, which learn the image features

Corresponding author: Rihao Chang (e-mail: changrihao@tju.edu.cn).

Weizhi Nie, Chuanqi Jiao, Rihao Chang, and An-An Liu are with the School of Electrical and Information Engineering, Tianjin University, Tianjin 300072, China. (e-mail: weizhinie@tju.edu.cn; chuanqi_097@tju.edu.cn; changrihao@tju.edu.cn; anan0422@gmail.com)

Lei Qu is with the Hisense Group Holdings Company, Ltd., Qingdao 266000, Shandong, China (e-mail: qulei1@hisense.com)

on the individual structures and the shape features on the overall structures, respectively. Then, we adopted the cross attention mechanism to explore the potential relation between the image feature and shape prior feature, based on which they are further combined into a highly integrated feature for 3D reconstruction. Finally, we expected the integrated feature to carry the overall homologous characteristics from priors and the differentiated individual characteristics from the single-view image. Thus, we adopted the 3D shape decoder to recover the 3D structural details from the integrated feature. Note that, to meet the previously desired reconstruction effect, the whole training process is under the supervision of the loss function that examines the reconstruction performance in 3D space.

B. Contributions

The contributions of this paper are as follows:

- We propose a novel approach for reconstructing high-quality 3D shapes based on a given single-view image and its retrieved 3D prior knowledge. Shape priors can effectively compensate for the unseen structures in single-view images and optimize the performance of reconstruction;
- We propose a novel multimodal knowledge fusion network based on the cross attention mechanism, which can hierarchically fuse the image's visual information with relevant 3D shape structural information;
- We propose a three-dimensional reconstruction network assisted by the skip connections from image features. Besides, we validate the performance of the proposed method on the ShapeNet and Pix3D datasets. Several ablation studies are conducted to evaluate the effectiveness of each module. All experiments demonstrated the superiority and reasonableness of our network design.

The remainder of this article is organized as follows. Section II presents related works on 3D shape reconstruction and cross-modal 3D shape retrieval. Section III provides the details of our approach. The corresponding experimental results and analysis are described in Section IV. Finally, we discuss the contributions and conclude this paper in Section V.

II. RELATED WORK

A. Three-Dimensional Object Reconstruction

Due to the wide application prospects of 3D vision, the reconstruction of 3D objects from a single-view image has attracted increasing attention. The main concern of single-view reconstruction is how to recover the missing geometric information of the input image. Traditional methods dig deep into the potential 3D representation of the input image, such as shading [14]–[17], occlusion [18], texture [19], [20] and vanishing points [21].

With the rapid development of deep learning and large-scale datasets, deep neural networks have shown their superiority. The early exploration, 3D-R2N2 [7], used a convolutional neural network with gated recurrent units to recover the full 3D structure of an object from single or several images. Unlike previous 3D reconstruction methods that rely on feature matching, 3D-R2N2 directly learns the mapping relationship

from 2D to 3D space. However, 3D-R2N2 remains computationally inefficient. Inspired by the octree structure [22]–[24], researchers of HSP [25], OGN [26] and O-CNN [27] optimized the voxel representation by the voxel-level occupancy frequency and facilitated higher resolution prediction. However, it is still difficult for these methods to generate high-quality 3D shapes. Therefore, some researchers utilized the efficient 3D representations of the polygon mesh [28]–[33] and deep implicit representation [34]–[38]. Polygon mesh represents the required shape with vertexes and faces, while implicit representation depicts the shape with a function between the 3D point position and the explicit or implicit distance from the point to the object surface. These methods can generate shapes with less redundant information. Besides, although the multi-view reconstruction task has long been dedicated to the combination of and consistency between different views, some methods [9], [39], [40] still have decent performance on single-view reconstruction tasks.

The key point of single-view reconstruction is the prior knowledge that compensates for the missing geometric information. Pontes et al. [31] represented the desired shape as a weighted combination of parameterized shape features. Yang et al. [11] first utilized memory networks [?], [?], [41], [42] to supplement 3D priors to 2D images and generate volumetric shapes. They introduced the LSTM-based shape encoder to extract shape priors from the memory network for reconstruction. Siddique et al. [12] leveraged the symmetries of 3D shapes as the prior knowledge. They considered that if the neighbors of an object partial are all symmetric, this partial is strongly judged as symmetric, and thus, the predicted shape can be optimized. Gao [13] et al. employed the domain adaptive learning for the alignment between 3D priors and 2D input images. They explored the multimodal correspondence between priors and input images, and further, they designed to divide the predicted shape according to the voxel grids and locally optimize the predicted shape. These previous works have verified the effectiveness of shape priors in 3D reconstruction. Thus, we employ cross-domain retrieval to precisely obtain desired 3D priors and design an efficient network to learn this guidance.

B. Image-based 3D Object Retrieval

For input images, the compensation of the 3D structure information involves the challenge of cross-domain learning. Inspired by the prior-guided method of 3D reconstruction, cross-domain retrieval appears to be a more effective alternative to the acquisition of priors. Early methods [43]–[45] relied on a linear classifier to match the related shapes. However, the performance of these simple structures was unsatisfactory. Therefore, metric learning was introduced in many works to reinforce the performance of cross-domain retrieval. Lee et al. [46] proposed cross-view convolution to sequentially learn multi-view features for 3D shape representations. Moreover, they proposed a cross-domain triplet neural network to incorporate metric learning with cross-domain retrieval tasks. Mu et al. [47] described each input point as a point in Euclidean space and then mapped each 3D shape to a Riemannian

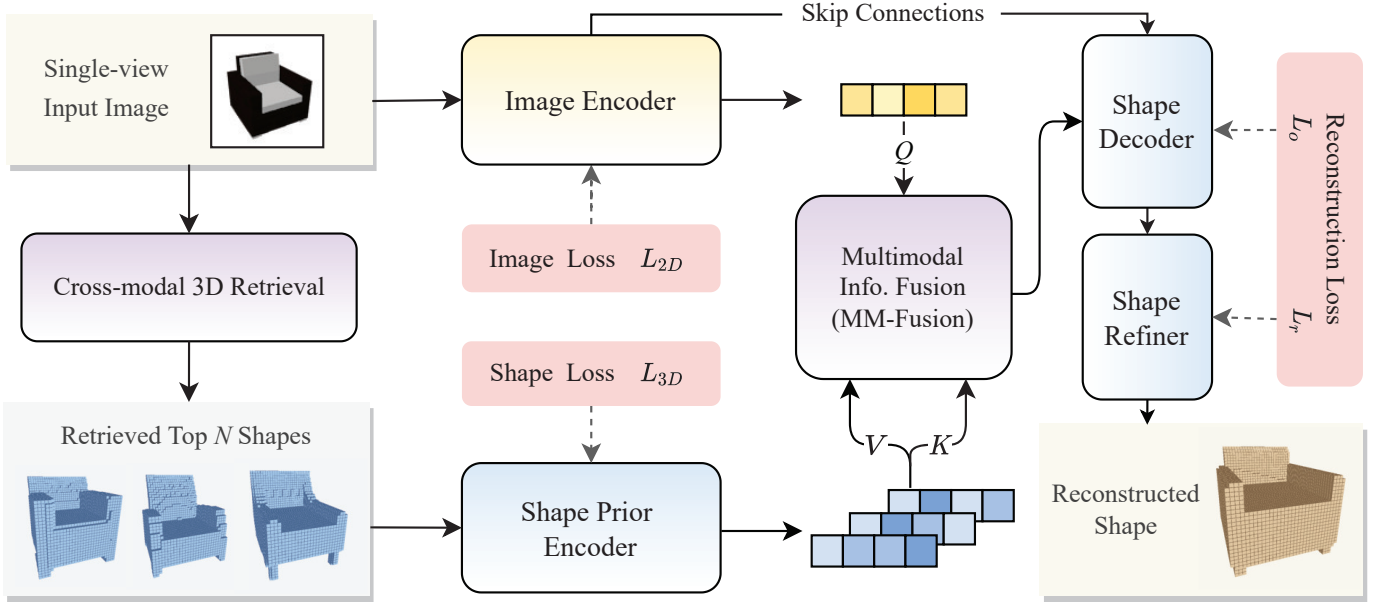


Fig. 1. The framework includes three parts: 1) Cross-modal prior knowledge extraction: input single-view image is used to retrieve the related 3D shape priors. 2) Multimodal information learning: the loss function L_{2D} and L_{3D} is utilized to supervise the feature extraction of the input image and shape priors, and these extracted features are further fused by the cross-modal attention modules. 3) Three-dimensional model reconstruction: The decoder is used to recover the 3D shape, and the refiner is adopted to optimize this reconstructed shape in detail.

manifold. To reduce the gap between the image and the 3D shape in different spaces, they bridged these two spaces to a shared high-dimensional Hilbert space, which greatly facilitated feature matching. Hence, cross-domain metric learning can effectively address the problem of complementing the 3D structure priors for the input image. More details about our network will be introduced in the following sections.

III. APPROACH

In this section, we will detail our approach. Fig.1 shows the framework of our approach, which is divided into three parts.

1) Cross-modal prior knowledge extraction: Given the target image, a cross-modal retrieval approach is adopted to extract the shape prior knowledge. To guarantee the correlation of retrieved results with the target image, we apply a pre-trained CLN model [48], which will be introduced in the following subsection. Finally, we select the top N retrieved results as the 3D prior knowledge to compensate for the missing structural information.

2) Multimodal information learning: On the one hand, for each modality, the network extracts the individual 3D information from the input image and the common 3D information from the shape priors. On the other hand, the integration of individual information and common information is generated from our cross-modal attention modules.

3) Three-dimensional shape reconstruction: The 3D shape reconstruction is solved by adopting the decoder to recover the 3D information from the fused feature. Moreover, a shape refiner [40] is employed for better performance, but in the sacrifice of network lightness.

We will detail these three parts in the following subsections.

A. Cross-modal prior knowledge extraction

Following the previous work of [48], we deploy the cross-modal retrieval method based on metric learning. In [48], Nie et al. first adopted a pose estimation network to predict the pose information of the input image. Then, this method takes a single-view image for 3D shape representation, which is rendered from a shape under the predicted pose. An extra CNN architecture is employed for 3D model feature extraction. Subsequently, this approach introduces a joint network for cross-modal feature learning, which effectively decreases the gap between modalities. In this procedure, metric learning is exploited to control the cross-modality feature distribution. The performance on the MI3DOR dataset [49] demonstrates the effectiveness of this method. Therefore, we utilized this method in our work for the retrieval of prior knowledge (3D structural information) from the input image.

B. Multimodal Information learning

In this section, we extract the features of the input single-view image and corresponding 3D shape priors, and then fuse the multimodal information for 3D shape reconstruction. Here, for effective feature learning, we extract domain-specific characteristics from each modality. The retrieved 3D shape priors are selected based on the similarity of 3D structures. Thus, this modality can provide an approximate description of the target 3D structure while still lacking individual details. Considering that the single-view image explicitly describes the target shape from the 2D perspective, we can directly capture the distinguishing characteristics from this modality to compensate for the lack of individual characteristics in the 3D modality.

Based on these analyses, we face three problems: 1) How can the distinguishing individual information be extracted during the feature learning of single-view images? 2) How can the general structure information be captured during the feature learning of the 3D shapes? 3) How can these separate learned information be combined for 3D shape reconstruction? To deal with these problems, we design different network components, which will be introduced in the following subsections.

1) *Image Encoder*: In this section, we plan to utilize the target image to guide the generation of individual 3D information. Therefore, we need to guarantee that the extracted image feature includes such distinguishing individual information. Here, we redesign the encoder network based on the VGG-16 structure [50]. The details of the image encoder are shown in Figure 2. We also employ additional branches as skip connections, which can effectively compensate for the distinguishing features of different scales during the reconstruction. The stacked convolutional layers can effectively extract the individual information of different scales. To guide the learning of this distinguishing individual information, the loss function L_{2D} is designed to amplify the individual details of the input image. The loss function is as follows:

$$L_{2D} = -\frac{1}{B} \sum_{i=1}^B \log\left(\frac{\exp(f_i f_i^T)}{\sum_{j=1}^B \exp(f_i f_j^T)}\right), \quad (1)$$

where i, j is the index of the feature vector, and f denotes the feature vector. f^T represents the image feature's transpose. B is the size of the training data batch. We compute the cross-correlation of each image feature vector in every batch of input samples. Here, each sample can be considered an individual class because the aim of feature learning is to amplify the individual characteristics of the current sample. Thus, we minimize this function to extract more individual information from the extracted image features.




Fig. 2. Architecture of the image encoder. The green, blue, yellow, and red arrows represent the skip connections of the feature maps of the input image.

2) *Three-Dimensional Shape Prior Encoder*: We are inspired by the method in [51] to design the 3D shape prior encoder. This encoder extracts the features of retrieved shape priors. Because the priors provide a general representation of the target 3D structure, the extracted features should include a union of similar structures in the 3D domain. As shown in Fig. 3, we design five 3D convolutional blocks and an extra linear layer to obtain a feature vector of length 1024 from a $32 \times 32 \times 32$ voxel grid. The kernel size of these convolutional layers is 4^3 with padding of 1. The linear layer is used only to adjust the output channels.

To capture more similar structures of the retrieved priors, we introduce the cross-entropy loss function for the parameter learning. The loss function is as follows:

$$L_{3D} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{ic}^* \log(y_{ic}), \quad (2)$$

where N is the number of training samples, and C is the number of classes. y_{ic} represents the predicted probability of class c , and y_{ic}^* is the one-hot coding value of the real label. We applied this loss function to capture the common 3D characteristics under each class and to encode these categorical shape structures into the output features.

We also tried to increase the depth of the shape prior encoder. However, additional convolutional layers do not lead to better performance. It not only increases the computation complexity and optimization difficulty, but also causes the information oblivion in shallow layers. However, too few convolution layers fail to capture sufficient structural information for the reconstruction, which is revealed by the wrong reconstruction of categorical parts on the shape. A practical case is that, an input chair with four straight legs was wrongly reconstructed as a board with no legs, which runs counter to the common understanding of the structure of chairs. In the future, we hope to introduce proper residual connections into the deeper prior encoder to further enhance the representational ability of the 3D prior features.




Fig. 3. Architecture of the 3D encoder. We applied the cross-entropy loss function to optimize the parameters of this network.

3) *Multimodal Information Fusion*: The multimodal inputs, the view, and the prior knowledge are encoded to obtain the individual feature vectors and the common feature vectors, respectively. For example, for the input view of a "chair", the view encoder usually focuses on the shape of the armrests, back and seat surface (round, square, tall, short, and hollow, etc.), and the type of legs (straight, curved, or five-star legs, etc.), while the 3D shape encoder focuses on the common characteristics of chairs, for instance, the existence of backs and legs. Obviously, these common features are composed of some individual features, and the individual features provide the structural characteristics of the target object that cannot be ignored. Therefore, we design a multimodal information fusion module, MM-Fusion, based on the attention mechanism in Transformer [52], which focuses on the information interaction between different features and cascades the correlation between them to achieve cross-modal feature fusion.

The structure of the multimodal information fusion module is shown on the left of Fig.4. The input of the module consists of two parts, image features and prior knowledge feature sets

(a collection of features of related 3D prior shapes). As shown in the figure, the fusion module consists of a series of cross-modal attention sub-blocks, and each sub-block can accept two input vectors. They are named as side input term, f_s , and bottom input term, f_b , according to the input positions in the Fig.4. Among them, in order to make full use of the prior knowledge, all the subblocks take the output of the previous level one as the f_b , except that the first subblock accepts image features as its f_b . Similarly, to prevent the network from forgetting the image details, all the subblocks take different 3D prior features as the f_s , except the f_s of the last subblock, which accepts the residual connection of image features.

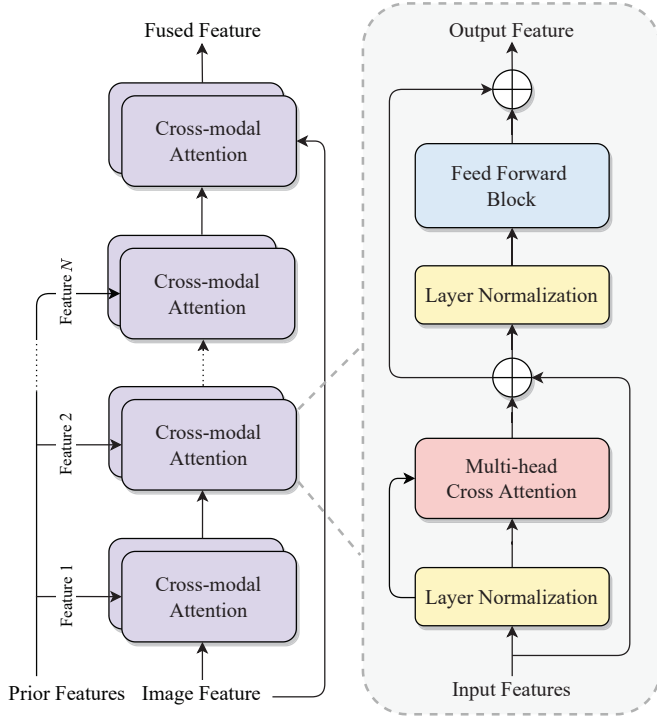


Fig. 4. Detailed structure of the Multimodal information fusion network (MM-Fusion).

Specifically, each cross-modal attention subblock contains multiple attention units, and the unit structure is shown on the right of Fig.4. Each unit mainly consists of a multi-head cross attention mechanism and a feed-forward unit. First, the f_s and f_b of each subblock need to be pre-processed by the linear layer. We name the processed results of f_b as the term q , while naming the processed results of f_s as terms k and v . Then, the cross-modal attention units can be expressed as follows.

$$f_{bn} = \text{LNorm}(f_b), f_{sn} = \text{LNorm}(f_s) \quad (3)$$

$$q = f_{bn} W_Q, k = f_{sn} W_K, v = f_{sn} W_V \quad (4)$$

$$z = \text{MHCA}(q, k, v) \quad (5)$$

$$z' = z + f_s \quad (6)$$

$$z'' = \text{FFW}(\text{LNorm}(z')) \quad (7)$$

$$f = z'' + z' \quad (8)$$

where W_Q , W_K and W_V are the linear layer weight matrices used for pre-processing. MHCA, LNorm, and FFW represent

the multi-head cross attention, layer normalization, and feed forward unit, respectively. The term z represents the multihead attention result, f represents the fusion result of this attention unit, z' and z'' is the intermediate result of the unit.

To be specific, layer normalization, LNorm, computes the mean and variance on each input feature and reduces the feature scale to speed up the convergence of training. Feed forward unit, FFW, is a multilayer perceptron with only one hidden layer and uses the GELU (Gaussian Error Linear Unit) function [53] for activation. Multi-headed cross attention can be expressed as follows,

$$\text{CrossAttn}(q, k, v) = \text{softmax}\left(\frac{qk^T}{\sqrt{d_k}}\right) v \quad (9)$$

$$\text{Head}_h = \text{CrossAttn}(q_h, k_h, v_h) \quad (10)$$

$$z = \text{Concat}(\text{Head}_1, \dots, \text{Head}_H) W_O \quad (11)$$

Here, k^T is the transpose of embedding k , d_k represents the length of the vector, Head_h represents the result of the calculation of the h^{th} attention head (the total number of attention heads is H). The function Concat splices the features on the channel, and the matrix W_O is the weight of the linear layer that processes the splice result.

C. Three-Dimensional Shape Reconstruction

In the shape decoder, we apply a structure that is symmetric with the 3D shape encoder. The decoder needs to transform the previously fused features into 3D volumes. We design a stack of 3D transposed convolutional blocks [54] in the decoder, as shown in Fig.5. The whole structure receives a combined feature of size 1×1024 and outputs a voxelized 3D shape of size $32 \times 32 \times 32$. These blocks are all composed of a transposed 3D convolution layer followed by a 3D batch normalization layer and a ReLU (Rectified Linear Unit) for optimization, except the last block, which imposes a sigmoid optimization for the voxel occupation probability.

To obtain sufficient information from the input image, we introduce the image encoder's processing information into the transposed convolutional layers in the decoder using skip connections [55], [56]. These connections are illustrated as the colored arrows in Fig.2 and Fig.5 (Arrows in the same color means they are the same skip connection). This operation is to compensate for the oblivion of the input information during 3D shape reconstruction. Note that the feature maps from the image encoder is reshaped to match the input size of its skip-connected 3D transposed convolutional layer.

Furthermore, we introduce the refiner network inspired by [57]. It aims to correct the wrongly recovered parts of a reconstructed 3D volume. Following the concept of a 3D encoder-decoder with U-net connections, the refiner preserves the local structure in the fused volume to generate optimized 3D volumes at 32^3 resolution.

1) *Optimization*: The loss function of the network is defined as the mean value of the voxelwise binary cross-entropy (BCE) between the reconstructed object and the ground truth. More formally, it can be defined as

$$L_o = \frac{1}{S} \sum_{i=1}^S [gt_i \log(p_i) + (1 - gt_i) \log(1 - p_i)] \quad (12)$$

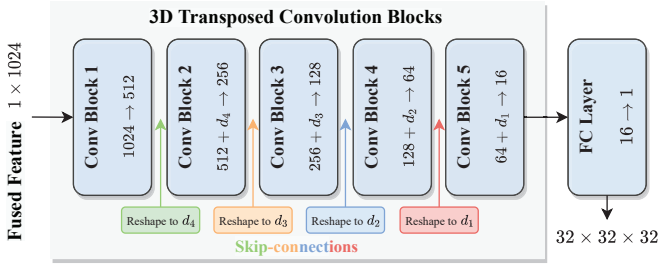


Fig. 5. Architecture of the reconstruction network (Decoder). The input of the decoder is the fused information of the image encoder and the 3D shape encoder’s output. Meanwhile, we introduce the skip connections from the image encoder which are indicated by the green, yellow blue and red arrows.

where S is the number of reconstructed voxel units per shape. gt represents the ground truth shape, and p is the predicted occupancy probability. The decrease in the L_o value demonstrates that the prediction approaches the ground truth.

Overall, the loss function L_{2D} and L_{3D} are used for updating the parameters of the image encoder and shape encoder, respectively. L_o reflects the reconstruction quality and, thus, is employed for updating the parameters of the entire network. Besides, we employ an additional BCE loss, L_r , to train the Shape Refiner. The classic Adam optimizer [58] is used to handle the optimization.

IV. EXPERIMENTS

In this section, we present extensive experimental evaluations of CPG3D on the ShapeNet [59] and Pix3D datasets [60]. We first describe the datasets and evaluation protocols. Then, the implementation details of the proposed methods are shown briefly. Above all, we report experimental evaluations of the proposed methods against state-of-the-art methods. Based on this, the ablation study and analysis further reveal the inner rationality and effectiveness of the proposed method.

A. Datasets

ShapeNet: ShapeNet [59] is announced as an ongoing effort to establish a richly annotated large-scale dataset of 3D shapes. However, the complete ShapeNet dataset is still not publicly available. For a more convenient use of the dataset, ShapNetCore V1 was released as a subset of the full ShapeNet dataset with single clean 3D models and manually verified category and alignment annotations. It covers 55 common object categories with approximately 51,300 unique 3D models, but no rendered images are provided. Thus, we use the renderings provided by Choy et al. [7]. They rendered every model into 24 views and voxelized the models into 3D voxels while following the original naming and classifying strategies. The 24 randomly rendered views of each 3D model are of size 137×137 , and voxelized 3D shapes are of size $32 \times 32 \times 32$. In addition, a uniform colored background is applied to the image during the experiment.

Pix3D: The Pix3D dataset [60] is a large-scale benchmark of diverse image-shape pairs with pixel-level 2D-3D alignment that is specially built for image-based reconstruction tasks. The dataset contains 395 real-world models and 10,069 images

of 9 object classes. Different from other previous datasets, Pix3D possesses both real-world images and precise 2D-3D alignment while maintaining the divergence of 3D models. The Pix3D dataset comes with rich information about each image-shape pair: 2D and 3D key points, voxel representation, image mask, and rendering camera intrinsic and extrinsic parameters. Moreover, Pix3D denotes the occlusion and truncation of each model as fields in the comprehensive document.

B. Evaluation Metrics

Before evaluating the reconstruction quality of the proposed method, we binarize the reconstructed probability voxels at a fixed threshold of 0.4 and thus output the standard voxel shapes. For the similarity measure between reconstructed shape and ground truth, we use intersection over union (IoU), which suits volumetric approaches best. More formally,

$$\text{IoU} = \frac{\sum_{i,j,k} \mathbb{I}(\hat{y}_{(i,j,k)} > \tau) \mathbb{I}(y_{(i,j,k)})}{\sum_{i,j,k} \mathbb{I}[\mathbb{I}(\hat{y}_{(i,j,k)} > \tau) + \mathbb{I}(y_{(i,j,k)})]} \quad (13)$$

where τ is the threshold of the previously mentioned binarization. The term y and \hat{y} represent the predicted probability and the corresponding occupancy of the ground truth at voxel position (i, j, k) , respectively. For a more robust indicator of the reconstruction performance, Tatarchenko et al. [61] proposed an F-score that calculates the harmonic mean between precision and recall. In this case, precision indicates the accuracy of the reconstruction, and recall indicates the completeness of the reconstruction. More formally,

$$\text{F-score}(d) = \frac{2P(d)R(d)}{P(d) + R(d)} \quad (14)$$

where $P(d)$ and $R(d)$ represent the precision and recall at a fixed distance threshold d , respectively. Specifically,

$$P(d) = \frac{1}{n_{\mathcal{R}}} \sum_{r \in \mathcal{R}} \left[\min_{g \in \mathcal{G}} \|g - r\| < d \right] \quad (15)$$

$$R(d) = \frac{1}{n_{\mathcal{G}}} \sum_{g \in \mathcal{G}} \left[\min_{r \in \mathcal{R}} \|g - r\| < d \right] \quad (16)$$

where \mathcal{R} and \mathcal{G} denote the reconstructed and ground truth point clouds, respectively. The term n with subscripts \mathcal{R} and \mathcal{G} represents the scale of the corresponding point clouds. The F-score works on the point clouds and concerns the distance between the points and surfaces. Therefore, for volumetric approaches, we utilize the marching cubes algorithm [62] to construct the object surface and then sample 8,192 points from this surface for the F-score calculation. For the surface-based methods without explicit points, we also use the marching cube algorithm to construct the surface and then sample 8,192 points to compute the F-Score.

C. Implementation Details

Our network inputs are set to RGB images on the scale of 224×224 , and the retrieved voxels are all in the size of $32 \times 32 \times 32$. Note that the original volumetric shapes in the Pix3D dataset are of size $128 \times 128 \times 128$. Thus, we downsampled the input voxel shapes into 32^3 to align with

the input channel of the shape prior encoder. The volumetric output is identical to the input voxels in size. For single-view images, we use pre-trained VGG-16 [50]. The retrieved voxels are encoded by the stacked 3D convolutional layers. To effectively utilize the multimodal features, we set each of the cross-modal attention blocks as two attention units with six heads. On the decoder, we exploit five transposed 3D convolutional layers to reconstruct the volumes. During the training procedure, we implement our network in PyTorch with a data batch size of 29 on a single Nvidia GTX 1080Ti GPU, and the unfull batch is discarded. In addition, we use an Adam optimizer [58] with a β_1 of 0.9 and β_2 of 0.999. The initial learning rate is set to 10^{-5} and decays to a half after 150 epochs.

D. Single-view Reconstructions

1) **Evaluation on ShapeNet:** To demonstrate the performance of our approach, we compare our method with several SOTA methods on the ShapeNet dataset. To ensure a fair comparison, all methods are evaluated under the same input images for all experiments. The final experimental results are listed in Tab.I. Higher scores of IoU and F-Score@1% indicate better reconstruction quality.

Figure 6 shows several reconstruction examples on the ShapeNet testing set. In these examples, benefiting from the retrieved shape priors, CPG3D maintains the fundamental structures of the target shape, such as the legs of the chair and the spoiler wing of the sports car. Meanwhile, CPG3D also captures more details of the input images, for instance, the shapes of sofa armrests, holes in stool legs, and wheels of the car. Compared with the implicit surface reconstruction method DISN, our method recovers better details of the target shapes, but DISN always reconstructs smoother surfaces, which is the inborn advantage of the implicit surfaces.

2) **Evaluation on Pix3D:** We also evaluated the performance on the Pix3D dataset. Pix3D aligned the 3D models precisely with the images and marked the shapes of occlusion and truncation. Thus, the results on the clear images can verify the performance on single-view reconstruction, and the results on the disturbed images demonstrate the positive effects of the retrieved shape priors. The overall quantitative results are shown in Tab.III. CPG3D achieves the best performance compared with previous SOTA methods. We visualized some reconstructed examples and corresponding retrieved priors in Fig.7. The quantitative results show that the introduction of shape priors can also improve the reconstruction qualities of real-world images. In addition, we reconstruct both the occluded images and unoccluded images in Fig.8 and Fig.9, respectively. The visualized results on unoccluded shapes show that CPG3D outperforms previous methods and recovers more details of the target shapes. On the other hand, the visualized results on occluded shapes obviously outperform the previous methods. CPG3D correctly recovers the occluded parts in the input images, while other methods are confused by these disturbances. These evaluated results indicate that CPG3D generates more compelling shapes on images with complicated backgrounds and disturbances such as occlusion.

3) **Higher-Resolution 3D Object Reconstruction:** Shapes in low resolution have better portability due to their lightness in size. However, this property also signifies their lack of high-frequency shape details. Therefore, unlike the commonly researched resolution of 32^3 , we fit our method to reconstruct shapes in higher resolution to evaluate the performance of our 3D-prior-based method in handling more complex details. Following the experiment settings in OGN [26], we report the mean IoU and F-Score@1% in the resolution of 64^3 and 128^3 in Tab.IV. The experimental results show that our method outperforms Pix2Vox, Matryoshka [69], and OGN by a large margin. Furthermore, to illustrate the reconstructing ability, we visualized several examples in Fig.10, as shown in which, our method is more capable of capturing high-resolution details than the previous method, such as the hubs, pedals, police lights, and bumpers.

E. Ablation Study

1) **Network Structure Ablation:** There are four modules in this network: image encoder, 3D model encoder, feature fusion module, and shape reconstruction module. We performed ablation experiments on these modules to demonstrate the rationality and superiority of our design. The quantitative results on the ShapeNet dataset are listed in Tab.V.

a) **Ablation on Image Encoder:** Our shape reconstruction occurs under the guidance of shape priors. Nevertheless, the retrieved shape priors focus on the recovery of structurally similar areas, which leads to unsatisfactory reconstructions. Relevant experimental results are in Tab.V.

b) **Ablation on Shape Encoder:** To demonstrate the effectiveness of shape priors, we performed an ablation study on the shape encoder, i.e., reconstruction from input images. Although the quantitative results are comparable to those of some previous methods, there is still room for improvement when compared with our prior-guided pipeline.

c) **Ablation on Refiner:** Our reconstructions from the shape decoder are further optimized in a U-net [56] manner, which is defined as refiner in the shape reconstruction module. Here, we pruned refiner to test the optimization effects. The quantitative results in Tab.V demonstrate the important role of Refiner in improving reconstruction quality.

d) **Different Batch Sizes:** Our baselines are divided into w/o Refiner, w/o Shape Encoder, and w/o Image Encoder. We conducted experiments of these baselines under different batch sizes. Under the batch size of 29, the entire pipeline achieved the best performance. According to the experimental results in Tab.V, the performance of baselines improves with increasing batch size, but the improvement trend gradually stabilizes. The results of the baselines do not exceed the entire pipeline's performance. In conclusion, the batch size indeed influences the performance of the baselines. However, the effect is limited. The baseline methods, which lack relevant modules, cannot get the best result beyond the entire pipeline of CPG3D. It indirectly proves the necessity and rationality of each module in our approach.

2) **Loss Functions:** We utilized different loss functions for each module in our network: (1) L_{2D} : Optimizing the image

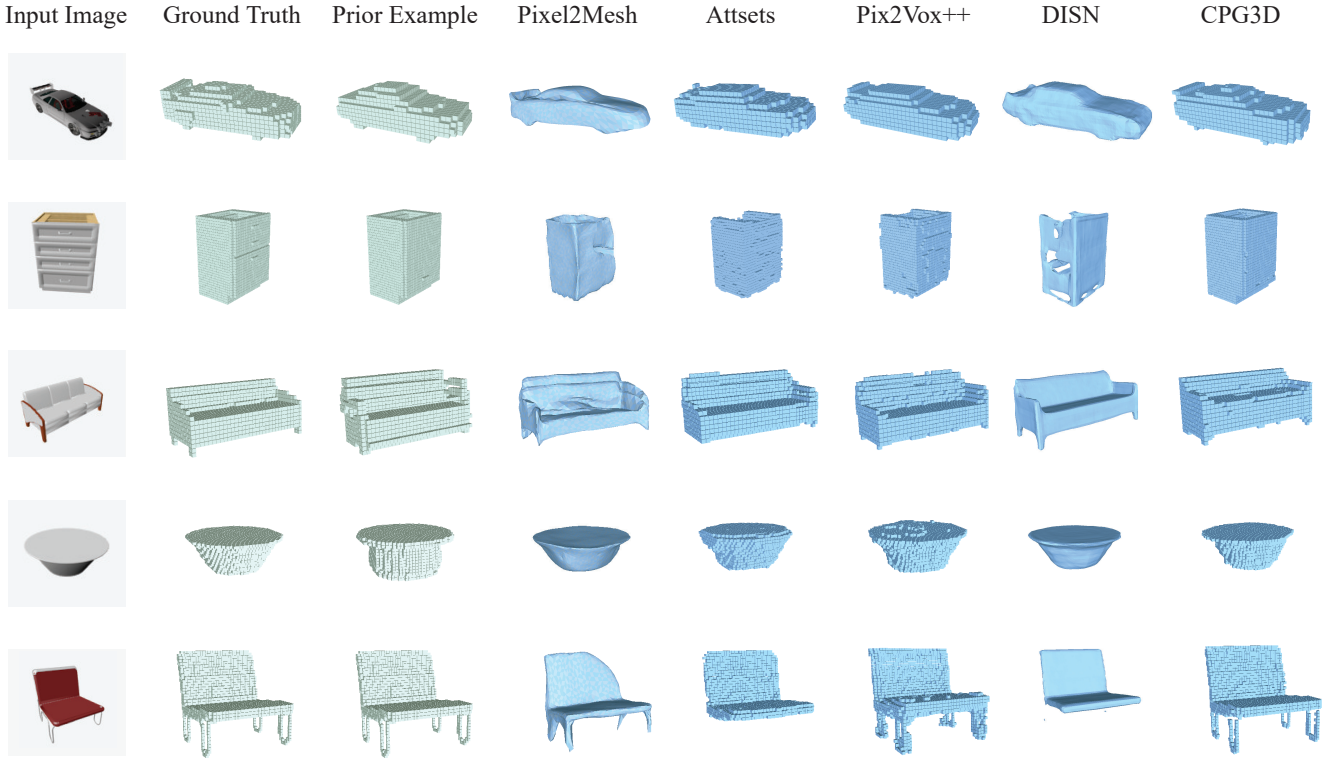


Fig. 6. Examples of single-view 3D object reconstruction on the ShapeNet dataset.

TABLE I
COMPARISON OF SINGLE-VIEW 3D OBJECT RECONSTRUCTION ON SHAPENET AT 32^3 RESOLUTION. WE REPORT THE MEAN IOU PER CATEGORY. THE BEST RESULT FOR EACH CATEGORY IS HIGHLIGHTED IN BOLD.

Category	3D-R2N2 [7]	OGN [26]	Pixel2Mesh [29]	AttSets [39]	Pix2Vox++ [40]	DASI [13]	Mem3D [11]	CPG3D
Airplane	0.513	0.587	0.508	0.594	0.674	0.701	0.767	0.772
Bench	0.421	0.481	0.379	0.552	0.608	0.625	0.651	0.659
Cabinet	0.716	0.729	0.732	0.783	0.799	0.798	0.840	0.875
Car	0.798	0.828	0.670	0.844	0.858	0.861	0.877	0.894
Chair	0.466	0.483	0.484	0.559	0.581	0.578	0.712	0.724
Display	0.468	0.502	0.582	0.565	0.548	0.552	0.631	0.648
Lamp	0.381	0.398	0.399	0.445	0.457	0.470	0.535	0.544
Speaker	0.662	0.637	0.672	0.721	0.721	0.723	0.778	0.791
Rifle	0.544	0.593	0.468	0.601	0.617	0.652	0.746	0.758
Sofa	0.628	0.646	0.622	0.703	0.725	0.723	0.753	0.770
Table	0.513	0.536	0.536	0.590	0.620	0.614	0.685	0.692
Telephone	0.661	0.702	0.762	0.743	0.809	0.801	0.823	0.836
Watercraft	0.590	0.632	0.471	0.601	0.603	0.622	0.684	0.695
Overall	0.560	0.596	0.552	0.642	0.670	0.676	0.729	0.743

encoder by enhancing the differences between samples in a batch (note that $L_{2D(CE)}$ is a cross-entropy variant of this function); (2) L_{3D} : a cross-entropy loss function concentrating on the categorical information of 3D priors; (3) L_o and L_r : two binary cross-entropy loss functions comparing the 3D occupancy between the predicted shape and corresponding ground truth. In this part, we validate the effects under different loss functions. Related experimental results are listed in Tab.VI. Note that BCE loss is indispensable for shape reconstruction, and thus our experiments reveal the effects of other losses based on L_o and L_r . For simplicity, we note $L_o + L_r$ as L_{BCEs} .

It is obvious that images lack the details of unseen parts,

and the shape priors struggle to recover object-specific details. This situation requires our network to capture more individual details from input images and more generic structures from combined shape priors. The commonly used loss function $L_{2D(CE)}$ focuses on category-level features, while our L_{2D} emphasizes the importance of object-divergent features. The improvement from $L_{2D(CE)} + L_{BCEs}$ to $L_{2D} + L_{BCEs}$ demonstrates that the object-divergent information indeed benefits the reconstruction. Consistent with the above conclusions, $L_{2D} + L_{BCEs}$ outperforms $L_{3D} + L_{BCEs}$ and even $L_{2D(CE)} + L_{3D} + L_{BCEs}$ due to the incorporation of individual features. The loss term $L_{2D} + L_{3D} + L_{BCEs}$ takes both individual details and generic structures into consideration and

TABLE II
COMPARISON OF SINGLE-VIEW 3D OBJECT RECONSTRUCTION ON SHAPE3D AT 32^3 RESOLUTION. WE REPORT THE F-SCORE@1% PER CATEGORY. THE BEST RESULT FOR EACH CATEGORY IS HIGHLIGHTED IN BOLD.

Category	3D-R2N2 [7]	OGN [26]	Pixel2Mesh [29]	AttSets [39]	Pix2Vox++ [40]	DASI [13]	Mem3D [11]	CPG3D
Airplane	0.412	0.487	0.376	0.489	0.583	0.604	0.671	0.679
Bench	0.345	0.364	0.313	0.406	0.478	0.484	0.525	0.531
Cabinet	0.327	0.316	0.450	0.367	0.408	0.431	0.517	0.536
Car	0.481	0.514	0.486	0.497	0.564	0.574	0.590	0.598
Chair	0.238	0.226	0.386	0.334	0.309	0.296	0.503	0.511
Display	0.227	0.215	0.319	0.310	0.296	0.299	0.498	0.507
Lamp	0.267	0.249	0.219	0.315	0.315	0.336	0.403	0.414
Speaker	0.231	0.225	0.190	0.211	0.152	0.294	0.262	0.233
Rifle	0.521	0.541	0.340	0.524	0.574	0.604	0.626	0.632
Sofa	0.274	0.290	0.343	0.334	0.377	0.393	0.434	0.445
Table	0.340	0.352	0.502	0.419	0.406	0.392	0.569	0.571
Telephone	0.504	0.528	0.485	0.469	0.633	0.638	0.674	0.683
Watercraft	0.305	0.328	0.266	0.315	0.390	0.425	0.461	0.470
Overall	0.351	0.368	0.398	0.395	0.436	0.443	0.517	0.524

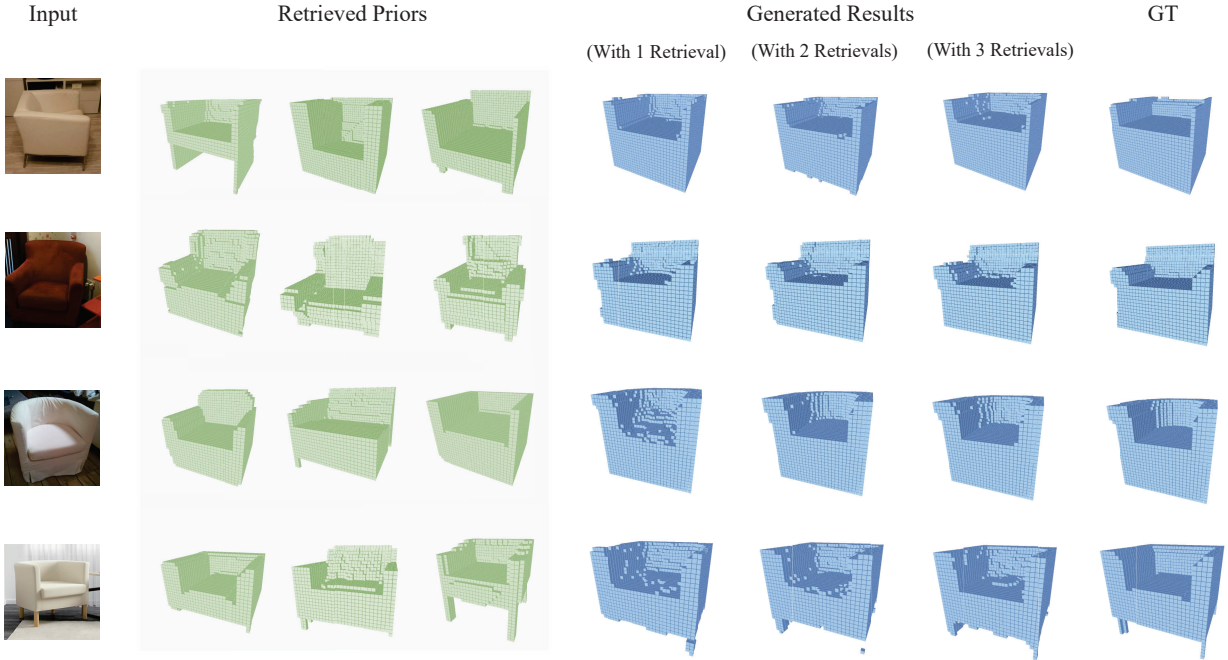


Fig. 7. Single-view 3D object reconstructions and related retrieved priors on the Pix3D dataset.

TABLE III
COMPARISON OF SINGLE-VIEW 3D OBJECT RECONSTRUCTION ON PIX3D AT 32^3 RESOLUTION USING MEAN IOU AND F-SCORE@1%. THE BEST RESULT IS HIGHLIGHTED IN BOLD.

Method	IoU	F-Score@1%
3D-R2N2 [7]	0.136	0.018
3D-VAE-GAN [63]	0.171	-
MarrNet [64]	0.231	0.026
DRC [65]	0.265	0.038
ShapeHD [66]	0.284	0.046
DAREC [67]	0.241	-
Pix3D [60]	0.267	0.041
Pix2Vox++ [40]	0.288	0.068
DASI [13]	0.310	-
FroDo [68]	0.325	-
Mem3D [11]	0.387	0.143
CPG3D	0.391	0.152

TABLE IV
COMPARISON OF SINGLE-VIEW 3D OBJECT RECONSTRUCTION ON SHAPE3D AT 64^3 AND 128^3 RESOLUTION. WE REPORT THE MEAN IOU AND F-SCORE@1%. THE BEST RESULT OF EACH METRIC IS HIGHLIGHTED IN BOLD.

Method	Resolution- 64^3		Resolution- 128^3	
	IoU	F-Score@1%	IoU	F-Score@1%
OGN [26]	0.771	0.361	0.782	0.390
Matryoshka [69]	0.784	0.380	0.794	0.426
Pix2Vox++ [40]	0.803	0.418	0.826	0.475
CPG3D	0.839	0.455	0.860	0.521

undoubtedly achieves the best performance.

3) *Skip connections*: The deep neural network is difficult to optimize due to overfitting, unstable random initialization, and gradient problems. Inspired by the designs of ResNet [70]

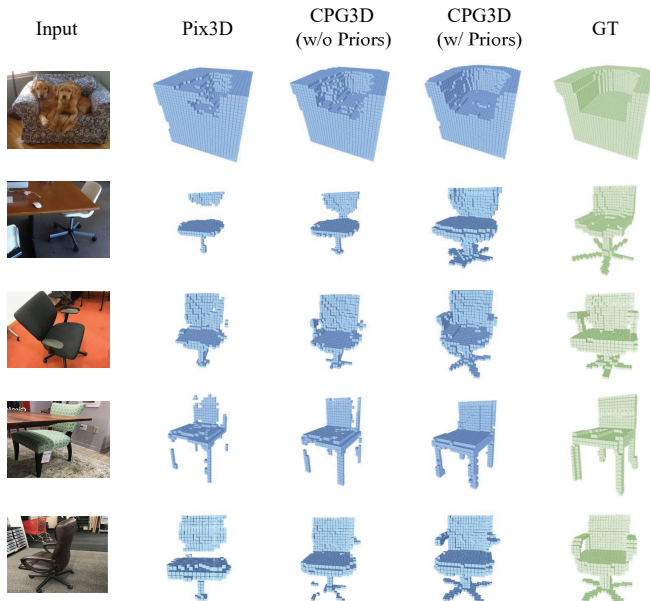


Fig. 8. Reconstructions of occluded samples on the Pix3D dataset.

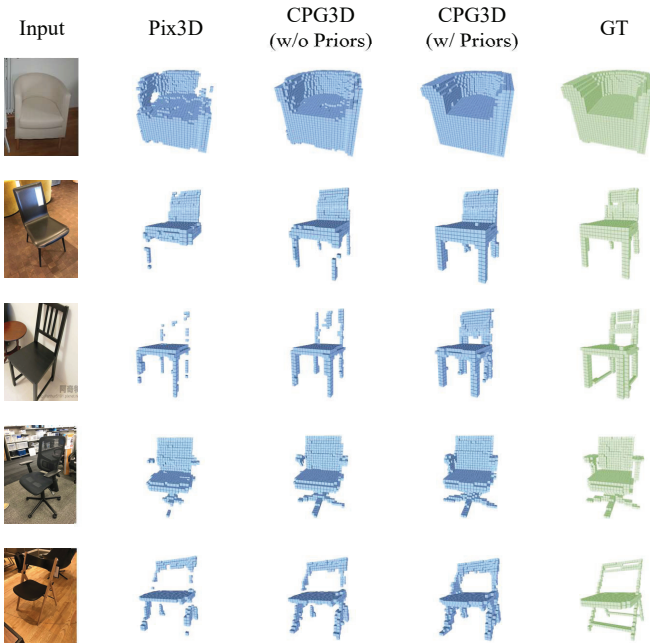


Fig. 9. Reconstructions of unoccluded samples on the Pix3D dataset.

and U-Net [71], we introduced skip connections between the image encoder and the reconstruction decoder. Although this design has widely proven to be effective for network convergence, the connection orders still remain to be further studied. To weigh the improvement benefited from skip connections against the computational consumption caused by the same scheme, we evaluated our network on different orders of such connections, and the quantitative results are listed in Tab.VII. Here, following the structure in Fig.1, Skc represents the skip connection, and the indexes ascend from shallow to deep in the image encoder. Experimental results show that the relative

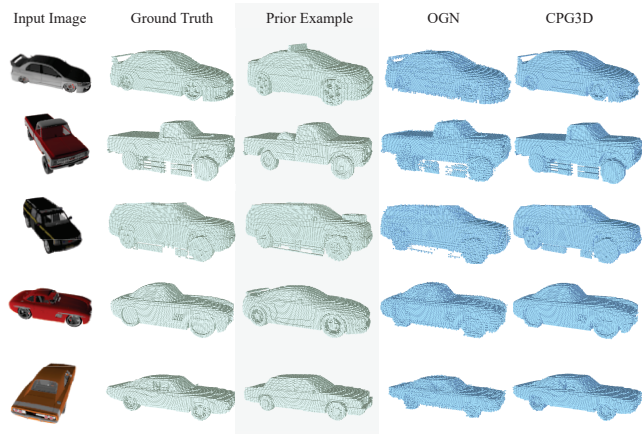


Fig. 10. Single-view reconstructions on ShapeNet-car dataset at 128^3 resolution.

TABLE V
ABLATION STUDY ON THE NETWORK STRUCTURE.

Structure	- Batchsize	IoU	F-Score@1%
CPG3D	- 29	0.743	0.524
w/o Refiner	- 35	0.723	0.498
	- 29	0.722	0.498
	- 8	0.718	0.475
w/o Shape Encoder	- 35	0.702	0.453
	- 29	0.698	0.451
	- 8	0.676	0.432
w/o Image Encoder	- 35	0.667	0.428
	- 29	0.664	0.427
	- 8	0.660	0.422

TABLE VI
EFFECTS OF DIFFERENT LOSS FUNCTIONS ON SHAPENET DATASET. (NOTE THAT WE USE “+” TO REPRESENT “AND” BUT NOT MATHEMATICAL PLUS OPERATION.)

Loss	IoU	F-Score@1%
L_{BCEs}	0.703	0.468
$L_{BCEs} + L_{2D}(CE)$	0.708	0.472
$L_{BCEs} + L_{3D}$	0.716	0.482
$L_{BCEs} + L_{2D}(CE) + L_{3D}$	0.720	0.488
$L_{BCEs} + L_{2D}$	0.731	0.512
$L_{BCEs} + L_{2D} + L_{3D}$	0.743	0.524

high-resolution features from shallow layers can compensate for the information in deep layers in the decoder and thus produce better reconstructions.

4) **Influence of Different Batch Sizes:** We optimized image loss L_{2D} to capture the divergent information between input images. Because L_{2D} is computed within a batch, we evaluated the network on different batch sizes. As illustrated in Fig.11, the size of the bubbles is proportional to the VRAM usage. A larger batch size tends to capture more divergent information while the VRAM consumption increases simultaneously. We perform the experiment on a single NVIDIA GTX 1080Ti, which supports a maximum batch size of 29, and we will explore the performance on a larger batch size in future work.

TABLE VII
EXPERIMENTS ON DIFFERENT SKIP-CONNECTION ORDERS.

Skip-connection(s)	IoU \uparrow	VRAM Usage (MB) \downarrow
<i>Skc4</i>	0.668	11031
<i>Skc3</i>	0.673	10864
<i>Skc2</i>	0.678	10680
<i>Skc1</i>	0.683	10653
<i>Skc2 + Skc3 + Skc4</i>	0.716	11124
<i>Skc1 + Skc3 + Skc4</i>	0.721	11097
<i>Skc1 + Skc2 + Skc4</i>	0.725	10943
<i>Skc1 + Skc2 + Skc3</i>	0.728	10869
<i>Skc1 + Skc2 + Skc3 + Skc4</i>	0.743	11151

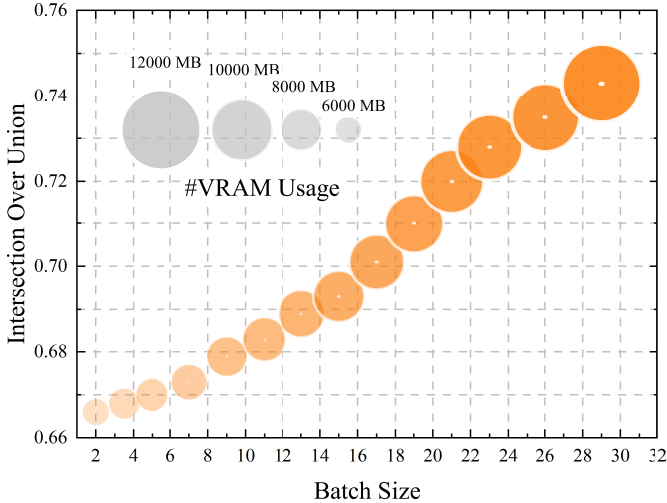


Fig. 11. IoU score and VRAM usage of different batch sizes.

F. Analysis

1) Influence of Disturbances on Retrieved Priors:

a) *Positive Disturbances*: Our cross attention module is sensitive to the size of the input feature sequence, which is experimented as a positive disturbance on the retrieved priors. In this part, we trained our network on different numbers of retrieved priors and used IoU as the criterion. The results are visualized as the orange-colored histogram in Fig.12. Considering that the function L_{3D} leads to more general representations of similar 3D structures, our newly increased shape priors can enrich these representations and thus positively influence our reconstruction. Nevertheless, this positive impact is not unlimited due to the redundancy of similar information. Our experiment shows that the increase is almost halted when there are more than five input priors.

b) *Negative Disturbance*: For some samples with unique structures, the increasing number of retrieved priors inevitably introduces priors of low retrieval confidence. Here, we artificially introduce specific numbers of unmatched priors to evaluate the robustness of the network. The qualitative results (green-colored histogram in Fig.12) show that our network can maintain the performance when the number of mismatched priors is less than 3. Note that the extreme situation of input ten mismatched priors still reached the IoU of 0.634, which also verified the robustness of our network on this disturbance.

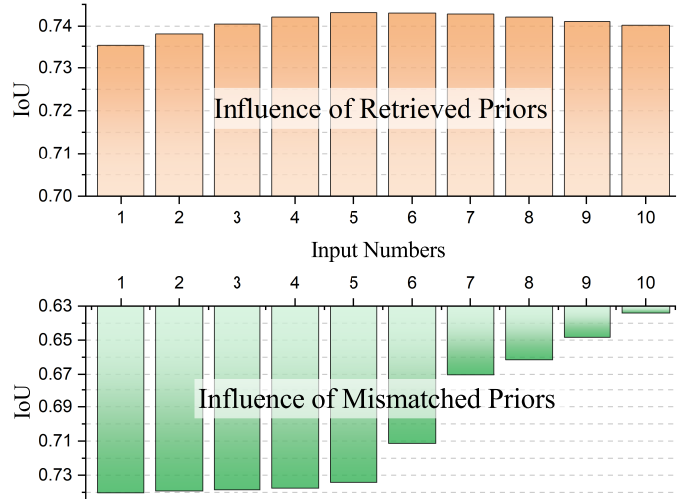


Fig. 12. Influence of Disturbances on Retrieved Priors

2) *Comparison with Best Retrieval Baseline*: As mentioned in [61], encoder-decoder structured networks easily fall into the *recognition* [61] problem, which merely pertains to the whole object and only performs the retrieval on the dataset. To test the upper bound of the retrieval methods on ShapeNet, we evaluated the Oracle Nearest Neighbour (Oracle NN) baseline, which, based on our shape encoder, finds the most similar shape from the training set for shapes in the test set and calculates IoU between the ground truth and the retrieved shape. The Oracle NN baseline demonstrates the best performance of the retrieval methods, which, on the other hand, confirms that our method is not a pure *recognition* network. The experimental results are listed in Tab.VIII.

TABLE VIII
COMPARISON WITH ORACLE NN

Category	Airplane	Bench	Cabinet	Car	Chair	Display	Lamp
CPG3D	0.772	0.659	0.875	0.894	0.724	0.648	0.544
Mem3D	0.767	0.651	0.840	0.877	0.712	0.631	0.535
OracleNN	0.540	0.341	0.576	0.778	0.291	0.412	0.218


Category	Speaker	Rifle	Sofa	Table	Telephone	Watercraft	Overall
CPG3D	0.791	0.758	0.770	0.692	0.836	0.695	0.743
Mem3D	0.778	0.746	0.753	0.685	0.823	0.684	0.729
OracleNN	0.498	0.484	0.499	0.293	0.805	0.442	0.458

3) *Analysis of Feature Fusion Methods*: To handle the fusion of multimodal information, we modified the cross attention mechanism from Transformer [52] and designed the MM-Fusion module. To verify the effect of this design, we compare it with some classic information fusion strategies. The related experimental results in Tab.IX show that MM-Fusion outperforms other methods. This structure can effectively reduce the redundant information and thus guarantee the effectiveness of the fused features.


4) *Analysis of Feature Distribution*: In this paper, we utilize the retrieved 3D priors to supplement structural information lost due to occlusion and background noise. Assume that the reconstruction target is a chair with rectangular

TABLE IX
EFFECTS OF DIFFERENT FUSION METHODS ON SHAPE3D DATASET.

Method	IoU	F-Score@1%
Addition	0.436	0.205
Straight-Connect	0.602	0.351
Average-Pooling	0.622	0.377
Max-Pooling	0.671	0.429
MM-Fusion	0.743	0.524



(a)



(b)

Fig. 13. (a) Distribution of features extracted by the image encoder. (b) Distribution of features extracted by the shape prior encoder.

arms and four straight legs. The retrieved priors capture this structural information while they are different in details, such as the thickness of the legs or the size of the arms. In other words, these priors are clustered according to the general 3D structures. However, the goal of 3D reconstruction is to recover the exact object from the input image but not to retrieve a similar 3D structure. Thus, we apply a correlation loss L_{2D} to guide the image feature learning to capture more characteristic features. Here, we apply the tSNE [72], [73] toolbox to visualize the distribution of extracted image features and prior features in Fig.13. As illustrated in the figure, the shape priors clustered tighter according to the class to which they belong, but in some special categories, the network struggles to distinguish features such as the 'Cabinet' and 'Speaker' because their ShapeNet ground truth volumes are similar in size and also share some common characters in 3D appearance. On the other hand, the image features are loosely distributed because they are more concerned with individual-level discrimination. This observation validated the rationality

of our motivation as well as the effectiveness of the network design.

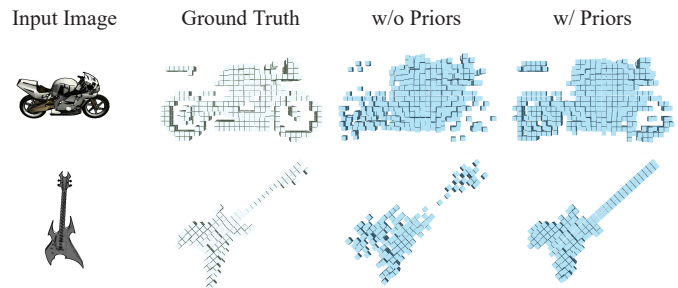


Fig. 14. Reconstructions on unseen categories on ShapeNet dataset.

5) *Analysis of Generalization*: To test the generalization of CPG3D, we processed ShapeNet categories beyond the 13-category data used in previous experiments in IV-D and performed evaluations on CPG3D. Note that our priors are also fine-tuned on novel categories. Some experimental results are visualized in Fig.14. As shown in the figure, CPG3D captures the general structures of target shapes, but there is still a gap between the reconstructed details and ground truths. Our method basically recovers the overall structure information. Furthermore, comparing reconstructed shapes with and without priors, CPG3D indeed learns essential information for better reconstructions from shape priors. This demonstrates the necessity of introducing prior information.

V. CONCLUSION

In this paper, we proposed a novel approach for reconstructing high-quality 3D objects based on a single-view image and its related 3D shape prior knowledge. We creatively utilized the image to retrieve similar 3D shapes as prior knowledge. These retrieved priors can provide effective geometric information for 3D reconstruction. Meanwhile, we proposed a novel cross-modal integration network for the joint learning of views and priors. Subsequently, we reconstruct the object volume from integrated features and introduced skip connections to optimize the reconstruction. Experimental results demonstrated that our approach significantly improves the reconstruction quality and performs favorably against state-of-the-art methods on the ShapeNet and Pix3D datasets.

REFERENCES

- [1] X. Wang, Y. Guo, Z. Yang, and J. Zhang, "Prior-guided multi-view 3d head reconstruction," *IEEE Transactions on Multimedia*, vol. 24, pp. 4028–4040, 2022. I
- [2] X. Fan, S. Cheng, K. Huyen, M. Hou, R. Liu, and Z. Luo, "Dual neural networks coupling data regression with explicit priors for monocular 3d face reconstruction," *IEEE Transactions on Multimedia*, vol. 23, pp. 1252–1263, 2021. I
- [3] X. Tu, J. Zhao, M. Xie, Z. Jiang, A. Balamurugan, Y. Luo, Y. Zhao, L. He, Z. Ma, and J. Feng, "3d face reconstruction from a single image assisted by 2d face images in the wild," *IEEE Transactions on Multimedia*, vol. 23, pp. 1160–1172, 2021. I
- [4] P. Hu, E. S. Ho, and A. Munteanu, "3dbodynet: Fast reconstruction of 3d animatable human body shape from a single commodity depth camera," *IEEE Transactions on Multimedia*, vol. 24, pp. 2139–2149, 2022. I
- [5] C. Yan, B. Shao, H. Zhao, R. Ning, Y. Zhang, and F. Xu, "3d room layout estimation from a single rgb image," *IEEE Transactions on Multimedia*, vol. 22, no. 11, pp. 3014–3024, 2020. I

- [6] Z. Zhou, F. Shi, J. Xiao, and W. Wu, "Non-rigid structure-from-motion on degenerate deformations with low-rank shape deformation model," *IEEE Transactions on Multimedia*, vol. 17, no. 2, pp. 171–185, 2015. I
- [7] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, "3d-r2n2: A unified approach for single and multi-view 3d object reconstruction," in *Proc. 14th European Conference on Computer Vision (ECCV'16)*, ser. Lecture Notes in Computer Science, vol. 9912, Amsterdam, The Netherlands, Oct.11–14, 2016, pp. 628–644. I, II-A, IV-A, I, II, III
- [8] H. Fan, H. Su, and L. 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*, 2017, pp. 605–613. I
- [9] C. Wen, Y. Zhang, Z. Li, and Y. Fu, "Pixel2mesh++: Multi-view 3d mesh generation via deformation," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV'19)*, Seoul, Korea (South), Oct.27 – Nov.2 2019, pp. 1042–1051. I, II-A
- [10] C. Liu, D. Kong, S. Wang, J. Li, and B. Yin, "DlGAN: depth-preserving latent generative adversarial network for 3d reconstruction," *IEEE Transactions on Multimedia*, vol. 23, pp. 2843–2856, 2020. I-A
- [11] S. Yang, M. Xu, H. Xie, S. W. Perry, and J. Xia, "Single-view 3d object reconstruction from shape priors in memory," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'21)*, virtual, Jun.19–25, 2021, pp. 3152–3161. I-A, II-A, I, II, III
- [12] A. Siddique and S. Lee, "Sym3dnet: Symmetric 3d prior network for single-view 3d reconstruction," *Sensors*, vol. 22, no. 2, 2022. I-A, II-A
- [13] J. Gao, D. Kong, S. Wang, J. Li, and B. Yin, "Dasi: Learning domain adaptive shape impression for 3d object reconstruction," *IEEE Transactions on Multimedia*, pp. 1–15, 2022. I-A, II-A, I, II, III
- [14] J. J. Atick, P. A. Griffin, and A. N. Redlich, "Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images," *Neural Comput.*, vol. 8, no. 6, pp. 1321–1340, 1996. II-A
- [15] R. Zhang, P. Tsai, J. E. Cryer, and M. Shah, "Shape from shading: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 8, pp. 690–706, 1999. II-A
- [16] R. Dovgand and R. Basri, "Statistical symmetric shape from shading for 3d structure recovery of faces," in *Proc. European Conference on Computer Vision (ECCV'04)*, Berlin, Heidelberg, 2004, pp. 99–113. II-A
- [17] G. Riegler, A. O. Ulusoy, H. Bischof, and A. Geiger, "Octnetfusion: Learning depth fusion from data," in *Proc. IEEE International Conference on 3D Vision (3DV'17)*, Qingdao, China, Oct.10–12, 2017, pp. 57–66. II-A
- [18] L. Guan, J. Franco, and M. Pollefeys, "3d occlusion inference from silhouette cues," in *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'07)*, Minneapolis, Minnesota, USA, Jun.18–23, 2007. II-A
- [19] J. Aliomonos and M. J. Swain, "Shape from texture," in *Proc. 9th International Joint Conference on Artificial Intelligence (IJCAI'85)*, Los Angeles, CA, USA, Aug. 1985, pp. 926–931. II-A
- [20] A. M. Loh and R. I. Hartley, "Shape from non-homogeneous, non-stationary, anisotropic, perspective texture," in *Proc. British Machine Vision Conference*, Oxford, UK, Sep. 2005. II-A
- [21] P. Parodi and G. Piccioli, "3d shape reconstruction by using vanishing points," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 2, pp. 211–217, 1996. II-A
- [22] D. Meagher, "Geometric modeling using octree encoding," *Comput. Graph. Image Process.*, vol. 19, no. 1, p. 85, 1982. II-A
- [23] A. Miller, V. Jain, and J. L. Mundy, "Real-time rendering and dynamic updating of 3-d volumetric data," in *Proc. 4th Workshop on General Purpose Processing on Graphics Processing Units (GPGPU'11)*, Newport Beach, CA, USA, Mar.5, 2011, p. 8. II-A
- [24] G. Riegler, A. O. Ulusoy, and A. Geiger, "Octnet: Learning deep 3d representations at high resolutions," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'17)*, Honolulu, HI, USA, Jun.21–26, 2017, pp. 6620–6629. II-A
- [25] C. Hane, S. Tulsiani, and J. Malik, "Hierarchical surface prediction for 3d object reconstruction," in *Proc. International Conference on 3D Vision (3DV'17)*, Qingdao, China, Oct.10–12, 2017, pp. 412–420. II-A
- [26] M. Tatarchenko, A. Dosovitskiy, and T. Brox, "Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs," in *Proc. IEEE International Conference on Computer Vision (ICCV'17)*, Venice, Italy, Oct.22–29, 2017, pp. 2107–2115. II-A, IV-D3, I, II, IV
- [27] P. Wang, C. Sun, Y. Liu, and X. Tong, "Adaptive O-CNN: a patch-based deep representation of 3d shapes," *ACM Trans. Graph.*, vol. 37, no. 6, p. 217, 2018. II-A
- [28] H. Kato, Y. Ushiku, and T. Harada, "Neural 3d mesh renderer," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2018)*, Salt Lake City, UT, USA, Jun.18–22, 2018, pp. 3907–3916. II-A
- [29] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y. Jiang, "Pixel2mesh: Generating 3d mesh models from single RGB images," in *Proc. 15th European Conference on Computer Vision (ECCV'18)*, ser. Lecture Notes in Computer Science, vol. 11215, Munich, Germany, Sep.8–14, 2018, pp. 55–71. II-A, I, II
- [30] M. Vakalopoulou, G. Chassagnon, N. Bus, R. Marini, E. I. Zacharaki, M. Revel, and N. Paragios, "Atlasnet: Multi-atlas non-linear deep networks for medical image segmentation," in *Proc. 21st Medical Image Computing and Computer Assisted Intervention (MICCAI'18)*, ser. Lecture Notes in Computer Science, vol. 11073, Granada, Spain, Sep.16–20, 2018, pp. 658–666. II-A
- [31] J. K. Pontes, C. Kong, S. Sridharan, S. Lucey, A. P. Eriksson, and C. Fookes, "Image2mesh: A learning framework for single image 3d reconstruction," in *Proc. 14th Asian Conference on Computer Vision (ACCV'18)*, ser. Lecture Notes in Computer Science, vol. 11361, Perth, Australia, Dec.2–6, 2018, pp. 365–381. II-A, II-A
- [32] J. Pan, X. Han, W. Chen, J. Tang, and K. Jia, "Deep mesh reconstruction from single RGB images via topology modification networks," in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV'19)*, Seoul, Korea (South), Oct.27 – Nov.02 2019, pp. 9963–9972. II-A
- [33] C. Lv, W. Lin, and B. Zhao, "Voxel structure-based mesh reconstruction from a 3d point cloud," *IEEE Transactions on Multimedia*, vol. 24, pp. 1815–1829, 2022. II-A
- [34] L. M. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'19)*, Long Beach, CA, USA, Jun.16–20, 2019, pp. 4460–4470. II-A
- [35] Z. Chen and H. Zhang, "Learning implicit fields for generative shape modeling," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'19)*, Long Beach, CA, USA, Jun.16–20, 2019, pp. 5939–5948. II-A
- [36] M. Michalkiewicz, J. K. Pontes, D. Jack, M. Baktashmotlagh, and A. P. Eriksson, "Deep level sets: Implicit surface representations for 3d shape inference," *CoRR*, vol. abs/1901.06802, 2019. II-A
- [37] Z. Chen and H. Zhang, "Learning implicit fields for generative shape modeling," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2019)*, Long Beach, CA, USA, Jun.14–20, 2019, pp. 5939–5948. II-A
- [38] Q. Xu, W. Wang, D. Ceylan, R. Mech, and U. Neumann, "DISN: deep implicit surface network for high-quality single-view 3d reconstruction," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS'2019)*, Vancouver, BC, Canada, Dec.8–14, 2019, pp. 490–500. II-A
- [39] B. Yang, S. Wang, A. Markham, and N. Trigoni, "Robust attentional aggregation of deep feature sets for multi-view 3d reconstruction," *Int. J. Comput. Vis.*, vol. 128, no. 1, pp. 53–73, 2020. II-A, I, II
- [40] H. Xie, H. Yao, S. Zhang, S. Zhou, and W. Sun, "Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images," *Int. J. Comput. Vis.*, vol. 128, no. 12, pp. 2919–2935, 2020. II-A, III, I, II, III, IV
- [41] J. Weston, S. Chopra, and A. Bordes, "Memory networks," in *3rd International Conference on Learning Representations (ICLR'15)*, San Diego, CA, USA, May7–9, 2015. II-A
- [42] W. Wang, R. Liu, M. Wang, S. Wang, X. Chang, and Y. Chen, "Memory-based network for scene graph with unbalanced relations," in *Proc. The 28th ACM International Conference on Multimedia (MM'20)*, Virtual Event / Seattle, WA, USA, Oct.12–16, 2020, pp. 2400–2408. II-A
- [43] M. Aubry, D. Maturana, A. A. Efros, B. C. Russell, and J. Sivic, "Seeing 3d chairs: Exemplar part-based 2d-3d alignment using a large dataset of CAD models," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'14)*, Columbus, OH, USA, Jun.23–28, 2014, pp. 3762–3769. II-B
- [44] H. Su, S. Maji, E. Kalogerakis, and E. G. Learned-Miller, "Multi-view convolutional neural networks for 3d shape recognition," in *Proc. IEEE International Conference on Computer Vision (ICCV'15)*, Santiago, Chile, Dec.7–13, 2015, pp. 945–953. II-B
- [45] B. Shi, S. Bai, Z. Zhou, and X. Bai, "Deeppano: Deep panoramic representation for 3-d shape recognition," *IEEE Signal Process. Lett.*, vol. 22, no. 12, pp. 2339–2343, 2015. II-B
- [46] T. Lee, Y. Lin, H. Chiang, M. Chiu, W. H. Hsu, and P. Huang, "Cross-domain image-based 3d shape retrieval by view sequence learning," in

- Proc. International Conference on 3D Vision (3DV'18)*, Verona, Italy, Sep.5–8, 2018, pp. 258–266. II-B
- [47] P. Mu, S. Zhang, Y. Zhang, X. Ye, and X. Pan, “Image-based 3d model retrieval using manifold learning,” *Frontiers Inf. Technol. Electron. Eng.*, vol. 19, no. 11, pp. 1397–1408, 2018. II-B
- [48] W. Nie, Y. Zhao, J. Nie, A. Liu, and S. Zhao, “CLN: cross-domain learning network for 2d image-based 3d shape retrieval,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 3, pp. 992–1005, 2022. III, III-A, III-A
- [49] H. Zhou, A. Liu, and W. Nie, “Dual-level embedding alignment network for 2d image-based 3d object retrieval,” in *Proc. 27th ACM International Conference on Multimedia (MM'19)*, Nice, France, Oct.21–25, 2019, pp. 1667–1675. III-A
- [50] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *3rd International Conference on Learning Representations (ICLR'15)*, San Diego, CA, USA, May7–9, 2015. III-B1, IV-C
- [51] D. Maturana and S. A. Scherer, “Voxnet: A 3d convolutional neural network for real-time object recognition,” in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'15)*, Hamburg, Germany, Sep.28 – Oct.2 2015, pp. 922–928. III-B2
- [52] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proc. Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NeurIPS'17)*, Long Beach, CA, USA, Dec.4–9, 2017, pp. 5998–6008. III-B3, IV-F3
- [53] D. Hendrycks and K. Gimpel, “Bridging nonlinearities and stochastic regularizers with gaussian error linear units,” *ArXiv*, vol. abs/1606.08415, 2016. III-B3
- [54] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, “Deconvolutional networks,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'10)*, San Francisco, CA, USA, Jun.13–18, 2010, pp. 2528–2535. III-C
- [55] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox, “Flownet: Learning optical flow with convolutional networks,” in *Proc. IEEE International Conference on Computer Vision (ICCV'15)*, Santiago, Chile, Dec.7–13, 2015, pp. 2758–2766. III-C
- [56] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, “Unet++: Redesigning skip connections to exploit multiscale features in image segmentation,” *IEEE Trans. Medical Imaging*, vol. 39, no. 6, pp. 1856–1867, 2020. III-C, IV-E1c
- [57] H. Xie, H. Yao, X. Sun, S. Zhou, and S. Zhang, “Pix2vox: Context-aware 3d reconstruction from single and multi-view images,” in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV'19)*, Seoul, Seoul, Korea (South), Oct.27 – Nov.2 2019, pp. 2690–2698. III-C
- [58] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations (ICLR'15)*, San Diego, CA, USA, May7–9, 2015. III-C1, IV-C
- [59] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, “Shapenet: An information-rich 3d model repository,” *CoRR*, vol. abs/1512.03012, 2015. IV, IV-A
- [60] X. Sun, J. Wu, X. Zhang, Z. Zhang, C. Zhang, T. Xue, J. B. Tenenbaum, and W. T. Freeman, “Pix3d: Dataset and methods for single-image 3d shape modeling,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'18)*, Salt Lake City, UT, USA, Jun.18–22, 2018, pp. 2974–2983. IV, IV-A, III
- [61] M. Tatarchenko, S. R. Richter, R. Ranftl, Z. Li, V. Koltun, and T. Brox, “What do single-view 3d reconstruction networks learn?” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'19)*, Long Beach, CA, USA, Jun.16–20, 2019, pp. 3405–3414. IV-B, IV-F2, IV-F2
- [62] W. E. Lorensen and H. E. Cline, “Marching cubes: A high resolution 3d surface construction algorithm,” in *Proc. the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'87)*, Anaheim, California, USA., Jul.27–31, 1987, pp. 163–169. IV-B
- [63] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, “Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling,” in *Proc. Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016 (NeurIPS'16)*, Barcelona, Spain, Dec.5–10, 2016, pp. 82–90. III
- [64] J. Wu, Y. Wang, T. Xue, X. Sun, B. Freeman, and J. Tenenbaum, “Marrnet: 3d shape reconstruction via 2.5d sketches,” in *Proc. Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NeurIPS'17)*, Long Beach, CA, USA, Dec.4–9, 2017, pp. 540–550. III
- [65] S. Tulsiani, T. Zhou, A. A. Efros, and J. Malik, “Multi-view supervision for single-view reconstruction via differentiable ray consistency,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'17)*, Honolulu, HI, USA, Jul.21–26, 2017, pp. 209–217. III
- [66] J. Wu, C. Zhang, X. Zhang, Z. Zhang, W. T. Freeman, and J. B. Tenenbaum, “Learning shape priors for single-view 3d completion and reconstruction,” in *Proc. 15th European Conference on Computer Vision (ECCV'18)*, ser. Lecture Notes in Computer Science, vol. 11215, Munich, Germany, Sep.8–14, 2018, pp. 673–691. III
- [67] P. O. Pinheiro, N. Rostamzadeh, and S. Ahn, “Domain-adaptive single-view 3d reconstruction,” in *Proc. IEEE/CVF International Conference on Computer Vision (ICCV'19)*, Seoul, Korea (South), Oct.24 – Nov.2 2019, pp. 7637–7646. III
- [68] M. Rünz, K. Li, M. Tang, L. Ma, C. Kong, T. Schmidt, I. D. Reid, L. Agapito, J. Straub, S. Lovegrove, and R. A. Newcombe, “Frodo: From detections to 3d objects,” in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'20)*, Seattle, WA, USA, Jun.13–19, 2020, pp. 14 708–14 717. III
- [69] S. R. Richter and S. Roth, “Matryoshka networks: Predicting 3d geometry via nested shape layers,” *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1936–1944, 2018. IV-D3, IV
- [70] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'16)*, Las Vegas, NV, USA, Jun.27–30, 2016, pp. 770–778. IV-E3
- [71] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Proc. 18th International Conference of Medical Image Computing and Computer-Assisted Intervention (MICCAI'15)*, ser. Lecture Notes in Computer Science, Part III, vol. 9351, Munich, Germany, Oct.5–9, 2015, pp. 234–241. IV-E3
- [72] L. van der Maaten and G. E. Hinton, “Visualizing non-metric similarities in multiple maps,” *Mach. Learn.*, vol. 87, no. 1, pp. 33–55, 2012. IV-F4
- [73] L. van der Maaten, “Accelerating t-sne using tree-based algorithms,” *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 3221–3245, 2014. IV-F4