

# MaterialSeg3D: Segmenting Dense Materials from 2D Priors for 3D Assets

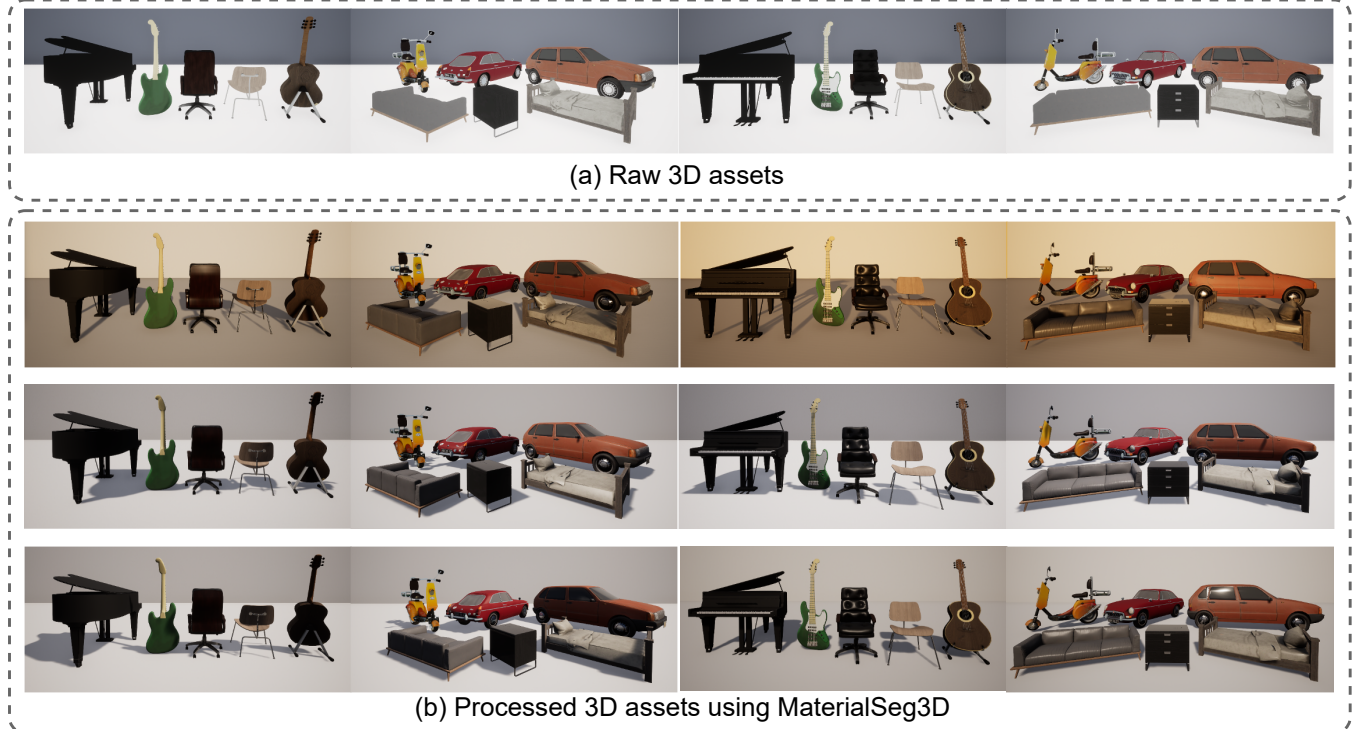


Figure 1: (a) Renderings of raw 3D assets that only have albedo information. (b) Renderings of processed assets with material information under different lighting conditions. Given a raw asset, our MaterialSeg3D can actively predict and refine dense explicit surface material based on 2D priors. Equipped with material definitions, 3D assets support physically based rendering, leading to photorealistic visual effects.

## ABSTRACT

Driven by powerful image diffusion models, recent research has achieved the automatic creation of 3D objects from textual or visual guidance. By performing score distillation sampling (SDS) iteratively across different views, these methods succeed in lifting 2D generative prior to the 3D space. However, such a 2D generative image prior bakes the effect of illumination and shadow into the texture. As a result, material maps optimized by SDS inevitably involve spurious correlated components. The absence of precise material definition makes it infeasible to relight the generated assets reasonably in novel scenes, which limits their application in

downstream scenarios. In contrast, humans can effortlessly circumvent this ambiguity by deducing the material of the object from its appearance and semantics. Motivated by this insight, we propose *MaterialSeg3D*, a 3D asset material generation framework to infer underlying material from the 2D semantic prior. Based on such a prior model, we devise a mechanism to parse material in 3D space. We maintain a UV stack, each map of which is unprojected from a specific viewpoint. After traversing all viewpoints, we fuse the stack through a weighted voting scheme and then employ region unification to ensure the coherence of the object parts. To fuel the learning of semantics prior, we collect a material dataset, named *Materialized Individual Objects (MIO)*, which features abundant images, diverse categories, and accurate annotations. Extensive quantitative and qualitative experiments demonstrate the effectiveness of our method.

Permission to make digital or hard copies of all or part of this work for personal or

Unpublished working draft. Not for distribution.

distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ACM MM, 2024, Melbourne, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nmmmmmmmmmmmm>

## CCS CONCEPTS

• Computing methodologies → Image processing; Texturing; Mesh geometry models.

## KEYWORDS

Image-based Rendering, Texture Mapping, Texture Synthesis

## 1 INTRODUCTION

3D asset creation, as a pivotal topic in computer graphics, has great application potential in virtual reality, augmented reality, games, and movies. It is a laborious workload for the artist in the traditional industrial pipeline. To create a 3D object of high quality, the artist often spends several days on sculpting geometry and drawing texture. The creation should adhere to some commonly recognized principles, such as neat polygon mesh and proper material design. This paper focuses on the material assignment of 3D assets. We follow Disney-principled BRDF and employ roughness and metallic as the primary physical properties of the material. These properties modulate the BRDF terms in the rendering equation and enable realistic re-lighting effects in different illumination conditions. With the advance of generative modeling, recent research [13, 39, 71] has achieved automatic creation of 3D objects according to textual or visual description. Most current methods resort to powerful 2D generative image models to supervise the 3D content generation. However, such 2D supervision bakes illumination. In this case, score distillation sampling inevitably leads to entangled material maps. Without precise material information, the generated assets cannot be re-lit realistically in novel scenes, which limits their application scope greatly.

For better usability, it is desirable to generate Physically-Based Rendering (PBR) material maps during asset creation [58]. We first investigate how the artist completes such a challenge. Given reference images of the object-of-interest, the artist can infer the material properties of each part according to the semantic information and appearance. For example, assuming an armchair with silver legs, thick black cushions, and a backrest, a human can confidently determine that the legs are metal and the seat cushion might be leather. Inspired by such a phenomenon, we point out that 2D priors' knowledge of material information can serve as powerful guidance for 3D material. Intuitively, material segmentation on 2D images is a perception-based method that can distill knowledge from labeled training images. However, existing material-related segmentation datasets such as DMS [65] or MINC [7] only provide material labels for open scenes including multiple instances, which are less reliable in dealing with single-object component segmentation. With the motivation of establishing a database to construct 2D material prior knowledge for individual objects, we collect *Materialized Individual Objects (MIO)*, a novel 2D single-object segmentation dataset consisting of dense material semantic annotations of objects with intricate semantic classes and captured camera angles. Images are (a) collected from both real-world captures and 3D asset renderings, augmenting the prior knowledge from reality and easing the domain gap; (b) sampled with various camera angles including but not limited to top and side views; (c) annotated and supervised by professional annotators. For each material class label in the dataset, we assign PBR material (Metallic, Roughness) under instructions of prior knowledge from experienced modelers. The MIO dataset contributes to establishing robust prior knowledge in material information while narrowing the distribution gap between object renderings in the application and the training data as well.

Empowered by the MIO dataset, we manage to propose *MaterialSeg3D*, a workflow that can automatically predict and generate precise surface material for 3D objects. Taking the geometry mesh and Albedo UV of an asset as input, our method first renders multi-view images of the asset with a manually and randomly selected camera pose. These multi-view renderings are then inferred by the material segmentation model, which is trained beforehand on the MIO dataset. Each predicted material result of multi-view images is further projected back onto a temporary UV map with the corresponding camera matrix. The final UV map for material labels is calculated through the voting mechanism and is further converted into a PBR material UV map including the Metallic and Roughness score for each material label assigned in the MIO dataset. As shown in Fig. 1, by absorbing 2D prior knowledge of material information from the MIO dataset, *MaterialSeg3D* can generate accurate surface material for 3D assets, resulting in vivid rendered visuals and application potential in the real world.

To summarize, the contributions of this paper are:

- We innovatively propose to utilize human prior knowledge of 2D material information in the surface material generation of 3D assets. Prior knowledge of the inherent relationship between the semantics and materials offers more reliable and precise guidance.
- We construct *MIO* dataset, which is currently the largest multiple-class single asset 2D material semantic segmentation dataset including images captured from especial camera angles and patterns, and each image is accurately annotated by a professional team.
- We introduce *MaterialSeg3D*, a novel workflow that can infer underlying material from the 2D semantic prior and accurately generate precise surface material for different parts of the 3D asset. This method can be significant in improving the quality of 3D assets from existing open-source datasets or websites.

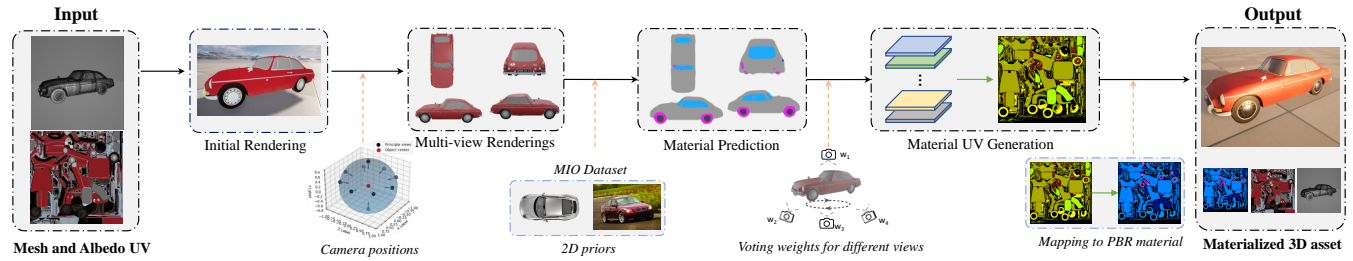
## 2 RELATED WORK

### 2.1 3D Asset Generation

Early methods in 3D asset generation often adapted existing 2D convolutional neural networks (CNNs) and generative adversarial networks (GANs) to generate 3D voxel grids [24, 30, 45, 61, 69, 74], these methods are straightforward but also difficult to generate high-quality 3D assets because they have many limitations such as high memory usage and computational complexity.

Subsequent research explored more methods such as based on point clouds [2, 46, 50, 76, 81], and implicit functions [14, 48]. The biggest problem of these 3D characterizations is the lack of compatible performance on standard computer graphics. Then to improve the quality and efficiency of 3D asset generation, the mesh-based 3D generative models [26, 29, 38, 42, 56, 64, 80] have emerged, accommodating complex topologies and shapes with varying resolutions. Importantly, the results from these models can seamlessly integrate with standard graphics engines, aligning with current industry demands for effective 3D data representation.

Currently, mainstream 3D generative models largely rely on text guidance to create a variety of 3D assets. Some methods involved optimizing Neural Radiance Fields (NeRF) [49] through text-image



**Figure 2: Overall framework of our MaterialSeg3D workflow. The material segmentation model is trained on MIO beforehand. Multi-view renderings are first generated with pre-defined and randomly selected camera angles and are further inferred by the material segmentation model and attached to a stacked temporary UV map. Weighted voting and region unification are further applied to generate the final material UV.**

alignment using the text-image Contrastive Language-Image Pre-training (CLIP) [32, 51, 57] model. DreamFusion [55] replaced CLIP with a diffusion model and introduced the loss of Score Distillation Sampling (SDS) to extract knowledge from the denoising process. Magic3D [37] further enhanced generative performance by adopting a coarse-to-fine framework and employing grids as a 3D representation in the second stage. Additionally, some other methods [1, 10, 11, 27, 52, 53, 60, 75] have combined NeRF techniques with the diffusion-based text-to-image models, proposing NeRF-based generators, but they primarily focused on geometric generation and often overlooking appearance.

## 2.2 Surface Material Generation

Generating realistic PBR material information such as metallic and roughness on the surface of 3D assets is key to making the asset look like a real object, that will dictate how surfaces interact with incident light, determining asset surface reflective behavior and color variations.

Traditional material generation methods predominantly focus on predicting physics-based materials under given lighting conditions, often requiring intricate multi-view [4] or polarizing [21] equipment. These methods often use synthetic data to train single-view Spatially Varying Bidirectional Reflectance Distribution Function (SVBRDF) prediction networks [19], which are then combined with other single-view data [25, 47] or custom training strategies [20, 36, 67] to obtain predicted material textures. These methods generated surface material information that looks inconsistent with what we perceive in the real world.

In recent years, many works have appeared in the field of 2D material segmentation for the controllable generation of materials in the form of SVBRDF maps [34, 58, 66, 68]. Based on a similar idea, there are several new works that have also emerged in the field of 3D material generation in an attempt to estimate materials under natural light conditions. Fania3D [13] decouples geometric and appearance modeling, using Bidirectional Reflectance Distribution Function (BRDF) to generate photo-realistic textures. However, it always predicts materials entangled with environmental lights, which leads to unrealistic renderings under novel lighting conditions. PhotoScene [78] utilizes procedural graphs as a prior for materials, generating high-resolution tiled material textures for each object in a scene, along with globally consistent lighting for the entire scene. PhotoScene, DiffMat [79], and Material Palette [43] are tailored for

tiled material generation. However, the surface material of a single complex 3D asset is often not tiled, making it difficult to generate and represent the asset’s true appearance through simple tiling. MatAtlas [9] generates relightable textures for 3D models given a text prompt with GPT4-V, but its generations across different views might differ in the appearance of the details.


## 2.3 Existing 3D and 2D Datasets

When considering learning prior information about the surface material, the first step is collecting enough data to support running the training process. In recent years, there have been some large-scale 3D datasets released, one of the most representative is the Objaverse, which is divided into Objaverse-1.0 [18] and Objaverse-XL [17], with approximately 800,000 and 10 million 3D assets, respectively. However, 3D assets in Objaverse generally lack material information, posing a limitation for research on surface appearance generation. Other 3D datasets like KIT [33], YCB [8], BigBIRD [59], and pix3d [63] offer calibrated models for various household objects, but they suffer from a severe lack of scale, containing at most a few hundred objects. Larger photorealistic object datasets [23, 54] and CAD model datasets [35, 70, 72] all do not include Albedo or material information. These existing 3D datasets fail to meet the requirements for generating realistic surface materials UV maps for individual complex 3D assets.

Due to the relative ease of obtaining 2D images, 2D material segmentation has accumulated more extensive large-scale datasets than 3D in the past decades. Such as the DMS dataset [65], encompassing 44,560 indoor and outdoor images with annotations for 3.2 million dense segments. The OpenSurfaces dataset [6] contains annotations for 37 material categories on 19,000 images of residential indoor surfaces. MINC [7] hosts the largest texture recognition dataset, featuring 3 million points annotated for 23 materials across 437,000 images. While these 2D datasets are extensive, their labels are often tailored for multi-object scenarios, bringing too much training noise when learning 2D priors for single-object scenarios.

## 3 SIGNIFICANCE OF MATERIAL

Creating high-quality materials in computer graphics is a challenging and time-consuming task, which requires great expertise. 3D assets with the correct materials can present the same impressions as in the real world under various lighting conditions. Components of the asset with different PBR materials will result in various reflection



**Figure 3: Comparison of 3D assets rendered with and without PBR material information under the same lighting conditions.**



**Figure 4: Case analysis of AI-generated asset surface material. (a) shows the rendering effect with the PBR material set to a fixed value on different structural components. (b) shows the generated material information cannot be consistent within the same semantic area.**

**Table 1: Statistics on the frequency of occurrence of different material categories contained in each image.**

Material label	Number	Material label	Number
metal	935	brick	186
wood	842	porcelain	163
plastic	768	clay terracotta	154
glass	712	concrete	152
paint	626	nylon	75
rubber	524	rusty metal	53
leather	437	ston	46
fabric	391	bone	25
fruit&leaf	273	bamboo	22
flower	252	others	181

effects even under the same illumination. 3D assets without PBR material information will cause extreme distortion when rendering diversified illuminations, making these properties inapplicable for real-world demands. Visualization can be found in Fig.3.

After recognizing the importance of PBR materials for 3D assets, we have conducted our early attempts to explore the potential of existing public datasets of 3D assets. In the newly proposed large-scale 3D object datasets Objaverse [17], we have analyzed a total of more than 270,000 assets of various categories, while only about 3k assets are attached with realistic PBR material information. This lack of material information in Objaverse makes it hard to learn the distributions of material semantics from the provided 3D assets.

Although there are some of the latest 3D asset generation methods [9, 13] claimed to have provided surface materials for the AI-generated content, the surface material quality of the 3D assets generated by these methods is rather poor with obvious distortions [58, 68], mainly caused by the following two problems. One of the problems that happened in some methods is that the PBR materials (Metallic, Roughness, *etc.*) attached to the surface are pre-defined fixed values regardless of the Albedo or semantic information. As shown in Fig. 4(a), the same PBR material values are attached to the handle and the head of the hammer, but they should be two materials with significant differences. Another issue is that the generation of the PBR material lacks guidance from real-world common sense or prior knowledge. The materials attached to a continual region of the asset may be discontinuous or unconvincingly related to the actual semantics of that region. A case is shown in Fig. 4(b), the region of the back of the chair should be applied with a continual material such as *fabric* or *nylon*, but *metal* is mistakenly attached in some part of the region.

Inspired by such case studies, we consider that human prior knowledge of witnessed categories of 3D assets can be utilized to judge or supervise the generation of the surface material. This statement also explains the logic of modelers manually assigning PBR materials to different assets, making it more convincing and logical. Further, we surveyed 100 people about the materials they thought were likely to appear in different categories of objects and showed each person 10 pictures of indoor and outdoor scenes. Each person was asked to count what materials might occur in every image, and the results are shown in Tab. 1. The survey results ensure that humans can confidently infer material information from a 2D image, and the frequency of different materials that occur in common objects is also supposed to be determined. This result greatly supports our motivation to introduce prior 2D knowledge to surface material generation of 3D assets.

## 4 MIO DATASET

### 4.1 Motivation for Establishment

Our pilot research indicates that 2D prior knowledge from humans can provide strong guidance and supervision for generating surface material on 3D assets. The following questions will be about how to employ and where to obtain such material prior knowledge related to 3D asset generation. Inspired by our relevant knowledge of the computer vision area, we figured out that perception-based methods can intuitively learn prior knowledge from training data into the models and infer the samples accordingly. Considering providing dense surface PBR material on 3D assets, segmentation is the most suitable method as it can provide pixel-wise prediction of material classes.

As aforementioned, in early attempts, we tried to collect available material information from public 3D asset datasets to build prior knowledge but ended up due to the extreme lack of material information. We subsequently notice that compared with 3D assets, 2D images are easily accessible through public websites or datasets with much wider distribution and total amounts. However, domain gaps exist between the distributions of 3D asset multi-view renderings and the existing annotated 2D image datasets, which makes learning & applying material prior knowledge less accessible. Therefore, we




Figure 5: Visual example of the material class annotations and the mapping with PBR material spheres.




Figure 6: Visual displays of samples in metaclass Cars collected in MIO dataset.

were motivated to construct a customized 2D image dataset that perfectly fits the demand of providing robust prior knowledge for surface material.

## 4.2 Data Collection and Annotation

To overcome possible domain gaps between 2D images and 3D asset renderings, we tried to collect and construct the image samples of our dataset under the following guidelines: (a) Each image sample could only contain one out-standing foreground object; (b) Image samples should be collected with similar amounts from both real-world scenes or renderings of 3D assets; (c) Image samples should be captured from diverse camera angles, including some especial angles such as the top view or bottom-side view. With the above guidelines, we ensure the gathered images share similar distributions with multi-view renderings of 3D assets, which largely guarantees the accuracy of further material predictions.

The sources of the collected images are freely accessible public datasets [31, 77] and 2D image renderings from 3D objects in website photo libraries [5]. In addition, we also procured some well-designed 3D assets that are used for game development and expanded the data collection by rendering multi-view images of these high-quality assets.

The biggest difference between our dataset and existing 2D segmentation datasets is that our customized dataset is designed to build extra alignments between semantic labels of different material classes and real PBR material values (Metallic, Roughness) for the included materials. The accuracy of the image annotation affects the overall performance of the material segmentation model trained

Table 2: Statistics of material labels occurrence in images.

Material label	Number	Material label	Number
metal	8,946	fabric	3,373
wood	7,088	fruit&leaf	1,742
plastic	6,928	flower	1,677
glass	5,802	brick	1,017
paint	5,626	porcelain	921
rubber	5,324	clay terracotta	910
leather	3,417	concrete	794

Table 3: Number of rendered images and real images of every metaclass in the MIO dataset.

Class name ( <i>abbr.</i> )	Rendered image	Real image	Total image
Furniture ( <i>fur.</i> )	4,152	5,455	9,607
Cars ( <i>car</i> )	1,935	4,117	6,052
Buildings ( <i>bui.</i> )	418	1,752	2,170
Musical Instrument ( <i>ins.</i> )	627	1,637	2,264
Plants ( <i>pla.</i> )	552	2,417	2,969

on the dataset, while the authority and rationality of the mapping between material class annotations and PBR materials influence the final rendered visualization of the assets. The number of material categories included in the dataset and their mapping relationships with PBR materials were discussed and determined by a group of nine professional 3D asset modelers. They have drawn upon their modeling expertise and considered the survey results shown in Tab. 1 to collect PBR material sphere candidates from more than 1,000 real PBR material spheres from public material libraries such as ACG [3] or Adobe Substance 3D Painter. Finally, 14 material categories, together with the mapping with the PBR materials to be the label space of our dataset.

After confirming the number of material categories in the dataset, we cooperated with a large and highly specialized annotators team to conduct pixel-wise dense annotations on the collected image samples. Based on the design of the dataset, we require that only the foreground objects contained in each picture be labeled with materials, and the background part is set to the background class regardless of the semantics. Each image sample was first annotated through an application driven by Segment Anything [34] and manual refinements, and sent to other annotators for multi-round re-annotation. Each annotator can handle approximately 50 images per day, ensuring the quality of the annotations is precise and accurate. Fig. 5 illustrates the annotations and alignments of material information in one of our samples.

## 4.3 Dataset Distribution

The dataset is named **Materialized Individual Objects (MIO)**, containing single-object image samples captured under diverse camera poses and annotated with convincing material labels and PBR material values. The MIO dataset comprises 23,062 multi-view images of individual complex objects, annotated into 14 material classes and categorized into five metaclasses: furniture, cars, buildings, musical instruments, and plants. The occurrence of each label is shown in Tab. 2. Occurrence frequency statistics of each metaclass belonging

to real images and asset-rendered images are illustrated in Tab. 3. Approximately 4,000 top-view images are included in the MIO dataset, providing a unique perspective rarely found in existing 2D datasets. Some image samples with the metaclass *cars* are displayed representing the diversity of camera poses and distributions, shown in Fig. 6.

## 5 METHOD

### 5.1 Material Segmentation

Inspired by existing semantic segmentation methods trained under material semantic labels, we establish a material segmentation process that better fits the demands of 3D assets. Compared with current semantic datasets annotated with material information, material segmentation focuses on dense predictions of a single object under diverse poses and camera angles. Given an image  $I$  with pixel-wise RGB value  $x$  and annotated material label  $y$  as a pair  $\langle x_i, y_i \rangle$  for each pixel  $i$ , the material segmentation network encodes visual features from the input image and decodes the features into per-pixel possibility vectors  $P_i = (p^{(i,1)}, p^{(i,2)}, \dots, p^{(i,n)})$  for  $n$  different classes at pixel  $i$ . The final prediction of each pixel can be calculated from  $P_i$  through *argmax* function.

We notice that the Segment Anything Model (SAM) [34] has shown its ability to handle semantic region segmentation on single-object images in previous work [71]. Thus, we formulate the material segmentation network with a modified ViT [22] backbone using pre-trained segmentation weights from SAM-b model. The decode head follows the setting in UperNet [73] with cross-entropy loss as supervision. To prevent possible long-tail problems caused by imbalanced training data, we adopt a class-balanced sampling strategy [16] to enhance the robustness and generalization ability of the model. During training stage, the cross-entropy loss can be calculated with:

$$L = -\frac{1}{HW} \sum_{i=1}^{HW} \sum_{c=0}^{n-1} y^{(i,c)} \log(p^{(i,c)}), \quad (1)$$

where  $H, W$  denotes the shape of the input image,  $n$  denotes the number of the classes,  $y^{(i,c)}, p^{(i,c)}$  represents the ground truth value, and the predicted possibility of class  $c$  at pixel  $i$ .

### 5.2 MaterialSeg3D

In this section, we introduce a novel material generation method, named **MaterialSeg3D**, a workflow that generates precise material information for 3D assets. The proposed MaterialSeg3D includes three components: multi-view rendering, material prediction, and material UV generation, as shown in Fig. 2. Specifically, in the multi-view rendering stage, the workflow first defines diverse camera poses capturing  $360^\circ$  of the target assets. 2D rendering images can be obtained from various angles from specific camera poses. In the material prediction stage, the material segmentation model is trained beforehand and infers multi-view renderings captured in the previous stage into the predicted material labels. In the material UV generation stage, predicted results of the renderings are first projected back to temporary UV maps and are further processed through a weighted-voting mechanism to obtain the final material label UV. Pixel values of the material label UV can be further transformed into PBR material (Metallic, Roughness) with the mapping relationships

between labels and material spheres. We will introduce the details in the following subsections.

**Multi-View Rendering.** In order to provide dense material predictions on the entire surface of an object, the elevation and rotation matrices of the rendering camera should cover  $360^\circ$  of the entire asset. Therefore, we first manually define five specific camera angles with the elevation and rotation status at  $(90^\circ, 0^\circ)$ ,  $(15^\circ, 0^\circ)$ ,  $(15^\circ, 90^\circ)$ ,  $(15^\circ, 180^\circ)$ ,  $(15^\circ, 270^\circ)$ . These rendered views can provide high-quality results and serve as popular views for human inspection. Next, we equally divide the entire  $360^\circ$  rotation into 12 different directions, on which there will be three different elevation angles,  $0^\circ$  as a fixed value, and the other two will be randomly selected within the range of  $(0^\circ, \pm 30^\circ)$  respectively. Through this, the renderings can provide visual information about all surfaces of the object, including the top and bottom. The manually selected views will further present additional constraints during the ensemble stage of the material UV.

**Material Prediction.** Following the details presented in Sec. 5.1, we can obtain a material segmentation model capable of predicting accurate material labels on images captured from various views. This model is used to infer the material information of the multi-view renderings of the input object. The predicted material labels are then used to generate material UVs.

**Material UV Generation.** After acquiring the predicted material results on the multi-view renderings, we generate the PBR material UV map for the 3D asset by attaching the material information to the pixel-wise UV map. Specifically, for each rendering with the rotation and elevation angle, we assign the predicted material labels to the corresponding pixel coordinates in the Albedo UV and form a new temporary material label UV. Through this, we can obtain a group of single-angle material label UV maps  $M_{view} = M_1, \dots, M_n$ , where  $n$  represents the number of the sampled camera views mentioned in the earlier paragraph.

As each rendering view can only provide limited material label information on the entire UV map, instead of sequentially updating the material label UV [12], we introduce a weighted voting method to decide the final material label of each pixel on the UV map. As aforementioned, five manually selected views will have higher weights when voting. Thus, the voted material label UV map can be calculated as follows:

$$M_{material} = vote(\alpha(M_1, M_2, M_3, M_4, M_5), M_6, \dots, M_n), \quad (2)$$

where  $\alpha$  denotes the weighting factor of the high-value views, and we set  $\alpha = 2$  in our experiments.

While the pixel values of the material label UV map are class labels predicted from the material segmentation model, the PBR material (Metallic, Roughness) UV map used to render visual effects can be transformed from the mapping relations between class labels and material spheres defined in the dataset.

## 6 EXPERIMENTS

### 6.1 Implementations & Evaluations

Learning precise 2D material prior information is at the forefront of our MaterialSeg3D pipeline for raw 3D objects. We trained our model with SAM-b[34] pre-trained ViT [22] backbone. The optimizer is AdamW [44] with the learning rate and weight decay are



Figure 7: Detailed visual comparisons between MaterialSeg3D and previous method from three aspects: single-image-to-3D generation methods, texture generation methods, and public 3D assets.

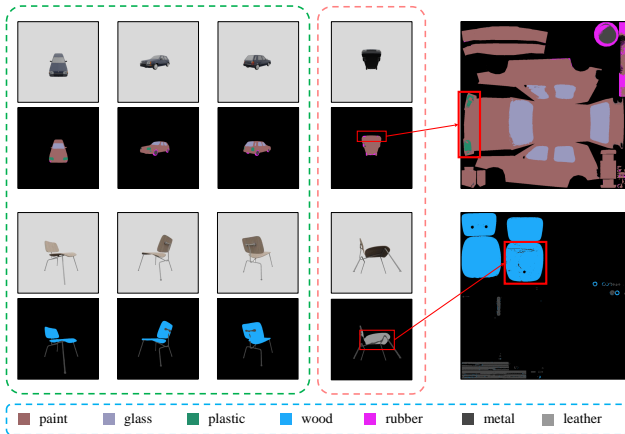
$6 \times 10^{-5}$  and  $1 \times 10^{-2}$ , respectively. We set batch size = 8 and training iterations = 80k, and images are resized to  $1024 \times 1024$ . All experiments are conducted under MMsegmentation [15] framework and on 4 80G NVIDIA A100 GPUs.

## 6.2 Compared with Previous Work

**Material Segmentation.** To evaluate the effectiveness of our proposed material segmentation method mentioned in Sec. 5.1, we apply five widely-used and state-of-the-art semantic segmentation backbones as comparisons to train segmentation models on the MIO dataset. We provide mIOU performance comparisons between these

methods on Objaverse [18] samples and the test image set of the MIO dataset. We randomly sample 50 assets with ground-truth PBR material UV from Objaverse and evaluate the accuracy of the output material label UV from MaterialSeg3D with the ground-truth. Quantitative results are shown in Tab. 4. It can be observed that our material segmentation method outperforms all the other semantic segmentation backbones, providing accurate and reliable material predictions for further renderings.

**Overall Performances.** To evaluate the effectiveness of the proposed material generation method, we compare previous approaches from the following three aspects: single-image-to-3D generation methods,



**Figure 8: Visualization of the segmentation results on multi-view rendering of 3D assets and the colored material UV map acquired from the weighted voting mechanism.**

**Table 4: Quantitative results about the performance of the semantic segmentation methods on the test split of MIO dataset / material label UV of Objaverse.**

Method	MIO Dataset (%)						Objaverse Dataset (%)					
	car	fur.	bui.	ins.	pla.	mIOU	car	fur.	bui.	ins.	pla.	mIOU
ConvNeXt[41]	71.03	74.85	69.33	72.40	76.72	72.87	75.35	76.04	72.34	76.72	78.95	75.88
HRNet[62]	75.71	79.94	76.37	80.14	81.35	78.70	78.40	78.83	76.03	82.00	81.40	79.33
ViT[22]	73.96	77.67	75.53	79.45	78.66	77.05	77.33	78.45	75.70	81.38	78.36	78.24
Swin-T[40]	75.09	79.04	78.45	80.92	81.40	78.98	78.89	79.77	78.64	82.97	82.01	80.46
MAE[28]	76.42	82.06	77.59	82.74	85.92	80.95	79.61	81.28	76.96	83.41	86.37	81.53
Ours	<b>81.83</b>	<b>85.22</b>	<b>81.76</b>	<b>84.39</b>	86.38	<b>83.92</b>	82.75	84.33	81.14	84.33	87.76	84.06

texture generation methods, and public 3D assets. The corresponding results are shown in Fig. 7. Considering single image-to-3D generation methods, we compare state-of-the-art Wonder3D [42], TripoSR [64], and OpenLRM [29] in this section. Specifically, given a reference view as input, Wonder3D, TripoSR, and OpenLRM generate a 3D object with referenced texture. We can observe that the provided MaterialSeg3D significantly outperforms the previous work owing to the adoption of well-defined 3D mesh and Albedo information. Fairly comparison, we modify existing texture generation methods like Fantasia3D [13], Text2Tex [12], and online functions provided by *Meshy*<sup>1</sup> for evaluation. Given a well-defined geometry mesh, previous work provide texturing results according to the text prompt as shown in Fig. 7. The results demonstrate our method provides much more realistic renderings under different lighting conditions. Note that for Fantasia3D, we only adopt its texture generation (Appearance Modeling) stage during comparison. Moreover, we also provide material generation results for 3D assets obtained from public websites, exempling as *tripo3d*<sup>2</sup> and *turbosquid*<sup>3</sup>. From the results in Fig. 7, we can observe the proposed MaterialSeg3D can generate precise PBR material information while significantly improving the overall quality of the assets.

Furthermore, we also provide quantitative results comparing our method and existing Image-to-3D methods including Wonder3D [42], TripoSR [64] and OpenLRM [29]. We adopt CLIP

<sup>1</sup><https://app.meshy.ai/>

<sup>2</sup><https://www.tripo3d.ai/app/>

<sup>3</sup><https://www.turbosquid.com/>

**Table 5: Quantitative evaluations from reference view and novel views on samples from Objaverse-1.0 dataset.**

Method	Evaluation view mesh	CLIP Similarity $\uparrow$		PSNR $\uparrow$		SSIM $\uparrow$	
		Reference	Novel	Reference	Novel	Reference	Novel
Wonder3D [42]		0.85	0.84	16.06	15.83	0.78	0.75
TripoSR [64]	w/o	0.93	0.90	16.93	16.14	0.79	0.76
OpenLRM [29]		0.92	0.87	16.30	15.37	0.77	0.76
Baseline	w	0.93	0.93	16.28	16.30	0.79	0.78
Baseline + Ours		<b>0.98</b>	<b>0.97</b>	<b>20.72</b>	<b>18.39</b>	<b>0.85</b>	<b>0.84</b>

Similarity [57], PSNR, and SSIM as the evaluations, and the corresponding results are shown in Table 5. We choose assets from Objaverse-1.0 dataset [18] as the test sample and randomly select three camera angles as novel views. The ground-truth reference and novel views are captured from assets with ground-truth material information and fixed lighting conditions. Given a well-defined 3D mesh and Albedo, our workflow can provide reliable PBR material, resulting in more realistic rendering visual effects.

### 6.3 Visualization on Weighted Voting

To illustrate the effectiveness of the weighted voting mechanism in the material UV generation stage, we provided visualizations of multi-view material segmentation results and the final material label UV maps, shown in Fig. 8. Although some regions might be predicted as wrong materials in some tricky angles, the correct predictions of the same region from other views will correct the final material labels through the weighted voting mechanism of the temporary UV maps.

## 7 LIMITATION

One of the limitations of our work is that current 3D asset generation methods mostly bake specific illuminations onto the generated RGB textures. Applying our workflow to the Albedo UV coupled with light reflections will lead to unrealistic visual effects under different illuminations.

Another limitation is that the quality of the input mesh will largely influence the generation of surface material and visual renderings. When applying our workflow on low-quality coarse meshes with uneven surfaces, the results are less satisfying. Detailed explanations and visualizations can be found in Supplementary Materials.

## 8 CONCLUSIONS

In this paper, we innovatively introduce the idea of adopting 2D prior knowledge of surface material in the material generation of 3D assets. We propose **MaterialSeg3D**, a novel workflow that takes a geometry mesh and Albedo UV as input, and generates dense PBR material information with the supervision of 2D prior knowledge. We also establish a 2D single-object material segmentation dataset **MIO** including images collected from diverse distributions and camera poses, thus providing strong 2D prior knowledge for the material segmentation model. Extensive experiments show the effectiveness of our proposed workflow. The workflow and the dataset show its ability to complete missing PBR material information for the public 3D assets, providing convenience for subsequent studies.



## REFERENCES

- [1] Rameen Abdal, Hsin-Ying Lee, Peihao Zhu, Menglei Chai, Aliaksandr Siarohin, Peter Wonka, and Sergey Tulyakov. 2023. 3davatar: Bridging domains for personalized editable avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4552–4562.
- [2] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. 2018. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*. PMLR, 40–49.
- [3] AmbientCG. 2024. Pbr repository. <https://ambientcg.com>.
- [4] Louis-Philippe Asselin, Denis Laurendeau, and Jean-François Lalonde. 2020. Deep SVBRDF estimation on real materials. In *2020 International Conference on 3D Vision (3DV)*. IEEE, 1157–1166.
- [5] Mathieu Aubry, Daniel Maturana, Alexei A Efros, Bryan C Russell, and Josef Sivic. 2014. Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3762–3769.
- [6] Sean Bell, Paul Upchurch, Noah Snively, and Kavita Bala. 2013. OpenSurfaces: A richly annotated catalog of surface appearance. *ACM Transactions on graphics (TOG)* 32, 4 (2013), 1–17.
- [7] Sean Bell, Paul Upchurch, Noah Snively, and Kavita Bala. 2015. Material recognition in the wild with the materials in context database. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3479–3487.
- [8] Berk Calli, Aaron Walsman, Arjun Singh, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. 2015. Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols. *arXiv preprint arXiv:1502.03143* (2015).
- [9] Duygu Ceylan, Valentin Deschaintre, Thibault Groueix, Rosalie Martin, Chun-Hao Huang, Romain Roufflet, Vladimir Kim, and Gaëtan Lassagne. 2024. MatAtlas: Text-driven Consistent Geometry Texturing and Material Assignment. *arXiv preprint arXiv:2404.02899* (2024).
- [10] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. 2022. Efficient geometry-aware 3D generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16123–16133.
- [11] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. 2021. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5799–5809.
- [12] Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. 2023. Text2tex: Text-driven texture synthesis via diffusion models. *arXiv preprint arXiv:2303.11396* (2023).
- [13] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. 2023. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873* (2023).
- [14] Zhiqin Chen and Hao Zhang. 2019. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5939–5948.
- [15] MMSegmentation Contributors. 2020. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>.
- [16] MMEngine Contributors. 2022. MMEngine: OpenMMLab Foundational Library for Training Deep Learning Models. <https://github.com/open-mmlab/mmengine>. (2022).
- [17] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. 2023. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663* (2023).
- [18] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. 2023. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13142–13153.
- [19] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. 2018. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.
- [20] Valentin Deschaintre, George Drettakis, and Adrien Bousseau. 2020. Guided fine-tuning for large-scale material transfer. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 91–105.
- [21] Valentin Deschaintre, Yiming Lin, and Abhijeet Ghosh. 2021. Deep polarization imaging for 3D shape and SVBRDF acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15567–15576.
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [23] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. 2022. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2553–2560.
- [24] Matheus Gadelha, Subhransu Maji, and Rui Wang. 2017. 3d shape induction from 2d views of multiple objects. In *2017 International Conference on 3D Vision (3DV)*. IEEE, 402–411.
- [25] Duan Gao, Xiao Li, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. 2019. Deep inverse rendering for high-resolution SVBRDF estimation from an arbitrary number of images. *ACM Trans. Graph.* 38, 4 (2019), 134–1.
- [26] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. 2022. Get3d: A generative model of high quality 3d textured shapes learned from images. *Advances In Neural Information Processing Systems* 35 (2022), 31841–31854.
- [27] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. 2021. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985* (2021).
- [28] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 16000–16009.
- [29] Zexin He and Tengfei Wang. 2023. Openlrn: Open-source large reconstruction models.
- [30] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. 2019. Escaping plato’s cave: 3d shape from adversarial rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9984–9993.
- [31] images.cv. 2024. CV image dataset. <https://images.cv>.
- [32] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. 2022. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 867–876.
- [33] Alexander Kasper, Zhixing Xue, and Rüdiger Dillmann. 2012. The kit object models database: An object model database for object recognition, localization and manipulation in service robotics. *The International Journal of Robotics Research* 31, 8 (2012), 927–934.
- [34] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643* (2023).
- [35] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. 2019. Abc: A big cad model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9601–9611.
- [36] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. 2017. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–11.
- [37] Chen-Hsuan Lin, Jun Gao, Luming Tang, Totoki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 300–309.
- [38] Minghua Liu, Chao Xu, Haiyan Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. 2023. One-2-3-4-5: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint arXiv:2306.16928* (2023).
- [39] Zexiang Liu, Yangguang Li, Youtian Lin, Xin Yu, Sida Peng, Yan-Pei Cao, Xiaojuan Qi, Xiaoshui Huang, Ding Liang, and Wanli Ouyang. 2023. UniDream: Unifying Diffusion Priors for Relightable Text-to-3D Generation. *arXiv preprint arXiv:2312.08754* (2023).
- [40] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*. 10012–10022.
- [41] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11976–11986.
- [42] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. 2023. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008* (2023).
- [43] Ivan Lopes, Fabio Pizzati, and Raul de Charette. 2023. Material Palette: Extraction of Materials from a Single Image. *arXiv preprint arXiv:2311.17060* (2023).
- [44] Ilya Loshchilov and Frank Hutter. 2018. Fixing weight decay regularization in adam. (2018).
- [45] Sebastian Lunz, Yingzhen Li, Andrew Fitzgibbon, and Nate Kushman. 2020. Inverse graphics gan: Learning to generate 3d shapes from unstructured 2d data. *arXiv preprint arXiv:2002.12674* (2020).
- [46] Shitong Luo and Wei Hu. 2021. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision*

- and Pattern Recognition. 2837–2845.
- [47] Rosalie Martin, Arthur Roullier, Romain Rouffet, Adrien Kaiser, and Tamy Boubekeur. 2022. MaterIA: Single Image High-Resolution Material Capture in the Wild. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 163–177.
- [48] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4460–4470.
- [49] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- [50] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas J Guibas. 2019. Structrnet: Hierarchical graph networks for 3d shape generation. *arXiv preprint arXiv:1908.00575* (2019).
- [51] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. 2022. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 conference papers*. 1–8.
- [52] Michael Niemeyer and Andreas Geiger. 2021. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11453–11464.
- [53] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. 2022. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13503–13513.
- [54] Keunhong Park, Konstantinos Rematas, Ali Farhadi, and Steven M Seitz. 2018. Photoshape: Photorealistic materials for large-scale shape collections. *arXiv preprint arXiv:1809.09761* (2018).
- [55] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022).
- [56] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. 2023. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843* (2023).
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [58] Sam Sartor and Pieter Peers. 2023. Matfusion: a generative diffusion model for svbrdf capture. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.
- [59] Arjun Singh, James Sha, Karthik S Narayan, Tudor Achim, and Pieter Abbeel. 2014. Bigbird: A large-scale 3d database of object instances. In *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 509–516.
- [60] Ivan Skorokhodov, Aliaksandr Siarohin, Yinghao Xu, Jian Ren, Hsin-Ying Lee, Peter Wonka, and Sergey Tulyakov. 2023. 3d generation on imagenet. *arXiv preprint arXiv:2303.01416* (2023).
- [61] Edward J Smith and David Meger. 2017. Improved adversarial systems for 3d object generation and reconstruction. In *Conference on Robot Learning*. PMLR, 87–96.
- [62] Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. 2019. High-resolution representations for labeling pixels and regions. *arXiv preprint arXiv:1904.04514* (2019).
- [63] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. 2018. Pix3d: Dataset and methods for single-image 3d shape modeling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2974–2983.
- [64] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. 2024. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151* (2024).
- [65] Paul Upchurch and Ransen Niu. 2022. A dense material segmentation dataset for indoor and outdoor scene parsing. In *European Conference on Computer Vision*. Springer, 450–466.
- [66] Giuseppe Vecchio, Rosalie Martin, Arthur Roullier, Adrien Kaiser, Romain Rouffet, Valentin Deschaintre, and Tamy Boubekeur. 2023. ControlMat: A Controlled Generative Approach to Material Capture. *arXiv preprint arXiv:2309.01700* (2023).
- [67] Giuseppe Vecchio, Simone Palazzo, and Concetto Spampinato. 2021. Surfnet: Adversarial svbrdf estimation from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12840–12848.
- [68] Giuseppe Vecchio, Renato Sortino, Simone Palazzo, and Concetto Spampinato. 2023. Matfuse: Controllable material generation with diffusion models. *arXiv preprint arXiv:2308.11408* (2023).
- [69] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. 2016. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *Advances in neural information processing systems* 29 (2016).
- [70] Rundui Wu, Chang Xiao, and Changxi Zheng. 2021. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6772–6782.
- [71] Tong Wu, Zhibing Li, Shuai Yang, Pan Zhang, Xingang Pan, Jiaqi Wang, Dahua Lin, and Ziwei Liu. 2023. HyperDreamer: Hyper-Realistic 3D Content Generation and Editing from a Single Image. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.
- [72] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1912–1920.
- [73] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. 2018. Unified perceptual parsing for scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*. 418–434.
- [74] Jianwen Xie, Zilong Zheng, Ruiqi Gao, Wenguan Wang, Song-Chun Zhu, and Ying Nian Wu. 2018. Learning descriptor networks for 3d shape synthesis and analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8629–8638.
- [75] Yinghao Xu, Menglei Chai, Zifan Shi, Sida Peng, Ivan Skorokhodov, Aliaksandr Siarohin, Ceyuan Yang, Yujun Shen, Hsin-Ying Lee, Bolei Zhou, et al. 2023. DisCoScene: Spatially Disentangled Generative Radiance Fields for Controllable 3D-aware Scene Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4402–4412.
- [76] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. 2019. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4541–4550.
- [77] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. 2015. A large-scale car dataset for fine-grained categorization and verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3973–3981.
- [78] Yu-Ying Yeh, Zhengqin Li, Yannick Hold-Geoffroy, Rui Zhu, Zexiang Xu, Miloš Hašan, Kalyan Sunkavalli, and Manmohan Chandraker. 2022. Photoscene: Photorealistic material and lighting transfer for indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18562–18571.
- [79] Liang Yuan, Dingkun Yan, Suguru Saito, and Issei Fujishiro. 2024. DiffMat: Latent diffusion models for image-guided material generation. *Visual Informatics* (2024).
- [80] Song-Hai Zhang, Yuan-Chen Guo, and Qing-Wen Gu. 2021. Sketch2model: View-aware 3d modeling from single free-hand sketches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6012–6021.
- [81] Linqi Zhou, Yilun Du, and Jiajun Wu. 2021. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5826–5835.