

# DiffCAD: Weakly-Supervised Probabilistic CAD Model Retrieval and Alignment from an RGB Image

DAOYI GAO, Technical University of Munich, Germany  
 DÁVID ROZENBERSZKI, Technical University of Munich, Germany  
 STEFAN LEUTENEGGER, Technical University of Munich, Germany  
 ANGELA DAI, Technical University of Munich, Germany

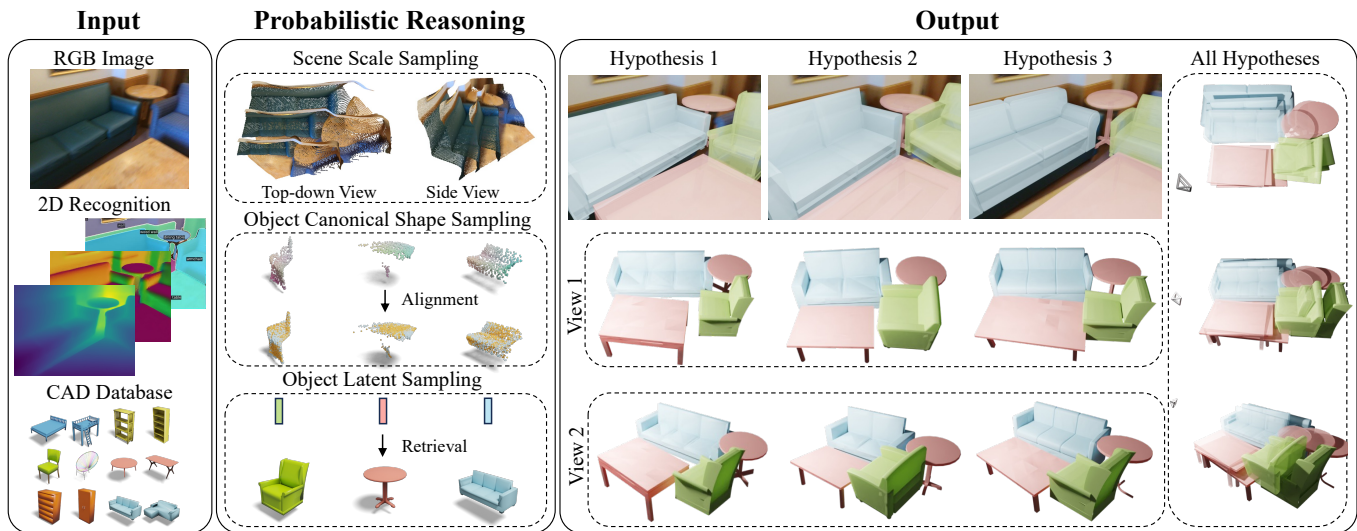


Fig. 1. We introduce DiffCAD, a weakly-supervised probabilistic approach for CAD model retrieval and alignment to a single RGB image. Unlike existing methods that require expensive CAD associations to real images for supervision, our approach is only trained on synthetic data, yet it can demonstrate superior performance on real-world datasets.

Perceiving 3D structures from RGB images based on CAD model primitives can enable an effective, efficient 3D object-based representation of scenes. However, current approaches rely on supervision from expensive yet imperfect annotations of CAD models associated with real images, and encounter challenges due to the inherent ambiguities in the task – both in depth-scale ambiguity in monocular perception, as well as inexact matches of CAD database models to real observations. We thus propose DiffCAD, the first weakly-supervised probabilistic approach to CAD retrieval and alignment from an RGB image. We learn a probabilistic model through diffusion, modeling likely distributions of shape, pose, and scale of CAD objects in an image. This enables multi-hypothesis generation of different plausible CAD reconstructions, requiring only a few hypotheses to characterize ambiguities

Authors’ addresses: Daoyi Gao, Technical University of Munich, Germany, daoyi.gao@tum.de; Dávid Rozenberszki, Technical University of Munich, Germany, david.rozenberszki@tum.de; Stefan Leutenegger, Technical University of Munich, Germany, stefan.leutenegger@tum.de; Angela Dai, Technical University of Munich, Germany, angela.dai@tum.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or 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).

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
 0730-0301/2024/7-ART106 \$15.00  
<https://doi.org/10.1145/3658236>

in depth/scale and inexact shape matches. Our approach is trained only on synthetic data, leveraging monocular depth and mask estimates to enable robust zero-shot adaptation to various real target domains. Despite being trained solely on synthetic data, our multi-hypothesis approach can even surpass the supervised state-of-the-art on the Scan2CAD dataset by 5.9% with 8 hypotheses.

CCS Concepts: • **Computing methodologies** → **Scene understanding**.

Additional Key Words and Phrases: CAD Model Retrieval and Alignment, Weak Supervision, 3D Reconstruction from a Single Image

## ACM Reference Format:

Daoyi Gao, Dávid Rozenberszki, Stefan Leutenegger, and Angela Dai. 2024. DiffCAD: Weakly-Supervised Probabilistic CAD Model Retrieval and Alignment from an RGB Image. *ACM Trans. Graph.* 43, 4, Article 106 (July 2024), 15 pages. <https://doi.org/10.1145/3658236>

## 1 INTRODUCTION

Obtaining 3D perception from 2D input is a fundamental challenge within the realm of computer graphics, carrying extensive implications for various applications in virtual environments and digital content creation, such as VR applications and the development of interactive gaming experiences.

While 2D perception from a single RGB image has achieved significant success in recent years [He et al. 2017; Kirillov et al. 2023;

Ranftl et al. 2021; Russakovsky et al. 2015], 3D perception – required for enabling tasks such as exploration and interaction with objects – remains challenging, requiring not only object recognition but reconstruction of diverse, complex objects. In 3D perception, leveraging a CAD model basis for representing 3D objects in an image enables the use of a strong 3D prior on geometric structures and generates clean, compact mesh outputs directly compatible with efficient modern rendering pipelines, thereby enhancing compatibility for subsequent graphics tasks.

While such CAD retrieval and alignment to generate an object-based 3D scene representation of real-world scene observation has shown strong promise [Avetisyan et al. 2019a; Gümeli et al. 2022; Izadinia et al. 2017; Kuo et al. 2020, 2021; Langer et al. 2022; Maninis et al. 2022], existing methods tend to rely on expensive real-world annotations that necessitate trained annotators [Avetisyan et al. 2019a]. Moreover, these annotations fail to provide exact ground truth matches regarding shape and pose, as existing CAD databases cannot tractably cover the distribution of general real objects. As shown in the bottom row of Fig. 5, the chair’s footrest is incorrectly annotated as a table. In contrast, the synthetic domain offers precise correspondences between renderings and object shapes inherent in the data without requiring manual annotations. Thus, rather than relying on limited, imperfect real-world data annotations for supervision, we adopt a synthetic-to-real domain adaptation strategy [Schwonberg et al. 2023; Ze and Wang 2022] to propose a weakly-supervised approach trained only on synthetic data, that can be applied to various real-world images.

To effectively address the problem of CAD retrieval and alignment with weak supervision, our key insight is that a probabilistic model is required. This enables capture of not only ambiguities in synthetic to real domain transfer, but also ambiguities inherent in the monocular perception task, such as depth-scale ambiguity and object shape ambiguity arising from partial visibility. Notably, previous methods overlook the inherent probabilistic nature of these challenges, opting for deterministic inference. In addressing the inherent ambiguities in single-view 3D perception systems, we model the likely distributions of scene scale, object pose, and object shape as separate and disentangled conditional generative tasks. To this end, we introduce DiffCAD, the first probabilistic approach for CAD retrieval and alignment to an RGB image without real-world supervision.

Our approach first employs a diffusion model to analyze potential scene scales based on an estimated monocular depth map. This enables considering multiple feasible solutions for object translation and scaling when subsequently solving for object poses. Given the estimated scales and depth map, we then model object pose through diffusion to predict an explicit canonical representation of objects, which we parameterize as normalized object coordinates (NOCs) [Wang et al. 2019]. The NOCs inform robust pose estimation and guide the diffusion of an object shape descriptor for CAD retrieval.

At inference time, we consider RGB images from real scenes. To help reduce the domain gap, we operate on machine-generated monocular depth and mask estimates from the RGB image (as photo-realistic RGB generation from synthetic 3D data is challenging and expensive). We can then sample from our learned distributions for

multiple plausible CAD reconstruction results. Our sampling scheme cascades from the scene-level diffusion, which offers potential scene scales. It then progresses to the explicit object representation diffusion, generating multiple NOCs that align with the scaled depth maps. Finally, it ends with the object latent diffusion, where the sampled latent vectors are used to query the CAD model database. This sampling scheme enables the generation of several CAD objects and poses that capture shape and depth-scale ambiguity in representing the 3D scene depicted in an image.

Experiments on ScanNet [Avetisyan et al. 2019a; Dai et al. 2017] images show that our learned probabilistic distributions well-capture likely CAD-based reconstructions, and with only 8 hypotheses, can even outperform fully-supervised state-of-the-art by 5.9%.

In summary, our contributions are:

- We propose the first probabilistic approach to CAD model retrieval and alignment from an RGB image, capturing inherent ambiguities due to depth-scale and lack of exact CAD matches in a database.
- We formulate our learned probabilistic model with diffusion processes that capture the distribution of scene scale, object pose, and object shape, with efficient sampling for multiple plausible hypotheses of CAD reconstructions.
- Our probabilistic approach leverages machine-estimated depth and 2D masks, enabling robust generalization to real images while training only on synthetic data.

## 2 RELATED WORK

### 2.1 2D Object Perception

With recent advancements in deep learning, significant progress has been achieved in various aspects of 2D object perception, including object detection [Carion et al. 2020; Fang et al. 2022; Ge et al. 2021; He et al. 2017; Lin et al. 2017], instance segmentation [Ghiasi et al. 2022; He et al. 2017; Kirillov et al. 2023; Li et al. 2022c; Liang et al. 2022; Xu et al. 2023], and reasoning about object and scene geometry, such as metric depth estimation [Bae et al. 2022; Bhat et al. 2020, 2023; Guizilini et al. 2023; Li et al. 2022a; Ramamonjisoa and Lepetit 2019], and object normal map estimation [Bae et al. 2022; Eftekhari et al. 2021]. Our approach builds upon this progress in 2D recognition, leveraging segmentation and depth estimates to build a 3D understanding of the scene through probabilistic reasoning regarding scene scale, object pose, and shape.

### 2.2 Single-View Object Reconstruction

The task of inferring object shapes from images through reconstruction has gained significant attention in recent years. Researchers have explored a variety of explicit and implicit object representations, including voxel grids [Choy et al. 2016; Wan et al. 2019], point clouds [Fan et al. 2017; Mandikal et al. 2018; Zeng et al. 2022], polygonal meshes [Gkioxari et al. 2019; Pan et al. 2019; Wang et al. 2018], and neural fields [Deng et al. 2022; Lin et al. 2022; Mescheder et al. 2019; Yu et al. 2020]. Recent methods have also leveraged such single-view object reconstruction to construct scenes on an object basis [Irshad et al. 2022; Liu et al. 2022; Nie et al. 2020; Zhang et al. 2021]. These approaches typically train on synthetic 3D shape data, producing flexible, underconstrained output representations,

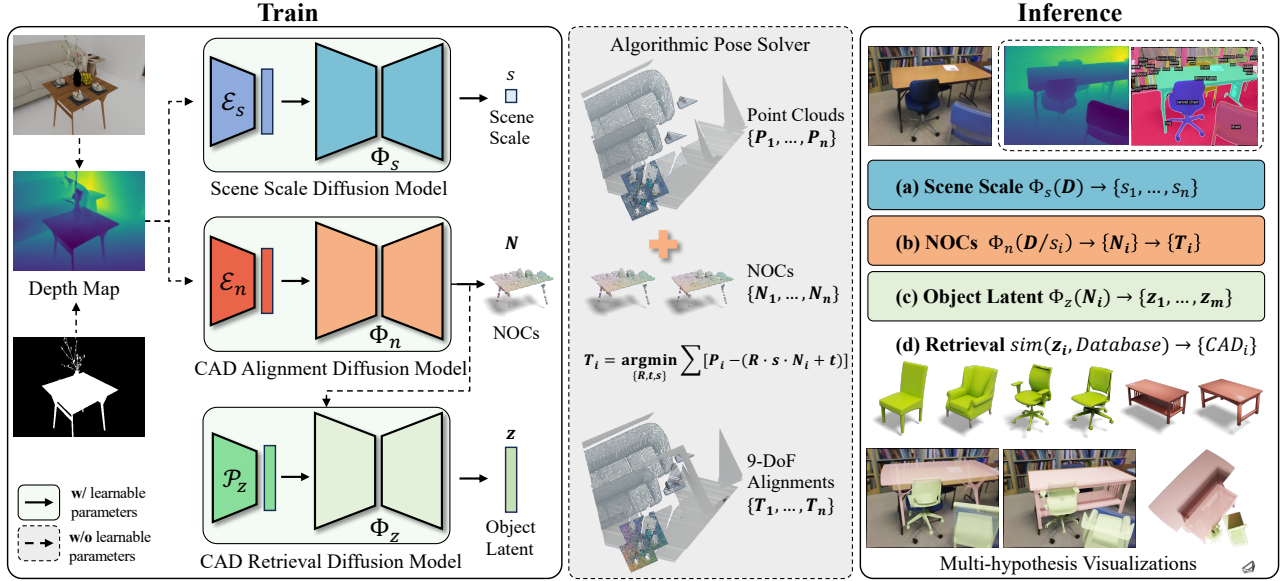


Fig. 2. **Method Overview.** To facilitate multi-hypothesis reasoning for CAD model retrieval and alignment to a single image, we employ diffusion modeling over scene scale, object pose, and shape. From an input RGB image, we employ machine-generated estimates of depth and instance segmentation. From the estimated depth, we estimate scene scales with  $\Phi_s$ .  $\Phi_n$  uses the back-projected estimated depth of each detected object to output hypotheses for its Normalized Object Coordinates (NOCs).  $\Phi_z$  then uses the estimated NOCs to predict the object shape as a latent vector that can be used for retrieval. Our probabilistic modeling also enables robust real-world CAD retrieval and alignment while training only on synthetic data.

resulting in reconstructions that are typically over-tessellated (due to Marching Cubes) and often exhibit local noise (e.g., imperfect flat surfaces) or missing finer-grained structures, due to the high dimensionality of the output representation. In contrast, our approach leverages a stronger 3D prior by employing CAD model databases to directly model scene geometry, resulting in plausible object reconstructions up to the fidelity of the CAD database, and a compact representation directly suitable for downstream applications.

### 2.3 CAD Model Retrieval and Alignment

As CAD representations enable efficient, mesh-based representations of a scene, various methods have explored CAD retrieval and alignment to real-world RGB images [Gümeli et al. 2022; Izadinia et al. 2017; Kuo et al. 2020, 2021; Langer et al. 2022] as well as RGB video [Maninis et al. 2022] and RGB-D scans [Avetisyan et al. 2019a,b, 2020; Beyer and Dai 2022; Di et al. 2023; Kim et al. 2012; Li et al. 2015; Nan et al. 2012; Shao et al. 2012].

For the challenging task of single-view RGB input, early methods have relied on the availability of ground truth scale information [Kuo et al. 2020, 2021] for full 9-DoF 3D reconstruction. Recent works have also proposed directly learning metric depth estimation within the target domain [Gümeli et al. 2022; Langer et al. 2022, 2023]. These methods all not only rely on ground truth supervision in the target domain, which is expensive and inexact to acquire, but produce deterministic outputs, whereas the task setting is inherently ambiguous due to depth-scale ambiguity and inexact CAD matches. In contrast, we propose a probabilistic approach to effectively model multiple plausible hypotheses, and maintain general applicability to real images while trained only on synthetic data.

### 2.4 Diffusion Models

Recently, diffusion models [Ho et al. 2020; Kingma et al. 2021; Rombach et al. 2021; Sohl-Dickstein et al. 2015; Song et al. 2020] have shown remarkable success in modeling image generation while employing an implicit probabilistic model. Due to their success in generative modeling of RGB images, recent works have shown more general applicability to various tasks, including segmentation [Baranchuk et al. 2021; Xu et al. 2023], keypoint matching [Hedlin et al. 2023; Luo et al. 2023; Tang et al. 2023], retrieval-based image generation [Blattmann et al. 2022; Chen et al. 2022; Sheynin et al. 2022], and 3D shape generation [Chou et al. 2022; Erkoç et al. 2023; Koo et al. 2023; Li et al. 2022b; Nam et al. 2022; Shue et al. 2023; Zeng et al. 2022; Zhang et al. 2023; Zhou et al. 2021]. We propose to employ diffusion modeling to characterize probabilistic models of scenes, along with CAD shape and pose for CAD model reconstruction of a single-view image.

## 3 METHOD

### 3.1 Overview

Given an RGB image  $I$  and a database of  $n$  CAD models  $\{O_i\}_{i \in [1,n]}$ , our goal is to represent objects in the scene by combining CAD models from the database with a set of 9-DoF transformations  $\{T_i\}_{i \in [1,m]}$  that align these models to the metric camera space, forming a compact, object-based 3D scene representation.

Figure 2 shows an overview of our method. We propose a combination of disentangled diffusion models to probabilistically model ambiguities in monocular perception and shape matching. Specifically, we model scene scale  $\Phi_s$ , pose prediction  $\Phi_n$ , and CAD model

retrieval  $\Phi_z$  as diffusion processes to effectively capture likely CAD model reconstructions of an image.

To reason robustly across various domains, we operate on machine-generated depth  $\mathbf{D}$  and semantic instance masks estimated from  $\mathbf{I}$ , denoted as  $\mathbf{D}_p$  indicating the masked depth estimate of each object (object index left out for simplicity of notation).

We start by addressing the scene-level depth-scale ambiguity. We learn the distribution of potential scales within the scene with a diffusion model  $\Phi_s$ , conditioned on  $\mathbf{D}$ . We can then sample scales  $\{s_i\}_{i \in [1, n]}$  from  $\Phi_s$ , to produce the most likely metric scenes  $\{\mathbf{D}_i = \frac{1}{s_i} \mathbf{D}\}_{i \in [1, n]}$ . This enables us to facilitate multiple reasonable transformations of objects.

We then model the object pose with diffusion model  $\Phi_n$ , which estimates the normalized object coordinates (NOCs) [Wang et al. 2019] conditioned on the back-projected point clouds  $\mathbf{P}$  of  $\mathbf{D}_p$ . This probabilistic formulation accounts for known multi-hypothesis issues for NOCs [Zhang et al. 2024], as a single observation can lead to multiple possible solutions due to the object’s symmetric structure or incomplete 2D observations, as shown in Fig. 3. The 9-DoF pose is recovered through RANSAC [Fischler and Bolles 1981] given NOC and  $\mathbf{P}$ .

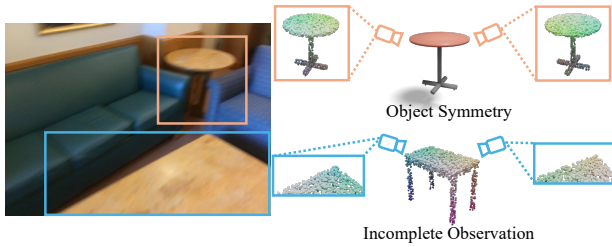


Fig. 3. **Multi-hypothesis nature of NOC.** The symmetry in object geometry and the incomplete perception can lead to multiple feasible alignments, which we characterize in our probabilistic, diffusion-based approach.

Given the estimated object pose characterized through NOCs, we also predict the object shape represented as a latent shape code through diffusion model  $\Phi_z$ . This probabilistic model enables capturing a distribution of possible matches, which enables more robust shape retrieval in real scenarios where no exact match from the database exists.

### 3.2 Diffusion Models

This work follows the denoising diffusion probabilistic models (DDPMs) formulation [Ho et al. 2020] to model CAD model retrieval and alignment from a single image, including the scene scale  $s$ , NOCs  $\mathbf{N}$ , and the latent object shape  $\mathbf{z}$ .

Given a data distribution  $q(x)$ , we denote a sampled data point as  $x_0$ . The *forward process* iteratively injects Gaussian noise to transform  $x_0$  to  $x_T$ , which fits an isotropic Gaussian in  $T$  timesteps, by a Markovian process [Ho et al. 2020; Song and Ermon 2019]:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

$$q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}), \quad (2)$$

where  $t \in [1, T]$  and  $\beta_t$  is a pre-defined variance schedule.

The denoising neural network  $\Phi$  learns to fit  $p_\Phi(x_{t-1}|x_t)$  that can approximate  $q(x_t|x_{t-1}, x_0)$  to and thus recover the  $x_0$  by the *reverse process*:

$$p_\Phi(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\Phi(x_t, t), \Sigma_\Phi(x_t, t)), \quad (3)$$

$$p_\Phi(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\Phi(x_{t-1}|x_t). \quad (4)$$

We can thus train our diffusion models  $\Phi_s$ ,  $\Phi_n$ , and  $\Phi_z$  following  $\epsilon$ -formulation [Ho et al. 2020; Rombach et al. 2021] or directly optimize for  $x_0$  following [Chou et al. 2023; Ramesh et al. 2022] to recover individual data distributions.

### 3.3 Scene Scale Diffusion

As monocular perception inherently contains depth-scale ambiguity, we use  $\Phi_s$  to model likely scene scales given machine-generated depth  $\mathbf{D}$ . Given the inherent bias in monocular depth estimates for both synthetic and real data, we argue that modeling the distribution of the scale factor between the prediction and the ground truth is more robust across domains. We define the target scale difference between the predicted depth map  $\mathbf{D}$  and the reference depth map  $\mathbf{D}_{gt}$  as follows:

$$s_{gt} = \text{avg} \left( \frac{\mathbf{D} \odot \mathbf{M}}{\mathbf{D}_{gt} \odot \mathbf{M}} \right). \quad (5)$$

Here,  $\mathbf{M}$  represents the mask of the target object, and  $\odot$  denotes the Hadamard product.

$\Phi_s$  is conditioned on features extracted from estimated depth map  $\mathbf{D}$  from a pre-trained ResNet-50-FCN backbone [Long et al. 2015]. Since the defined scene scale target is an isotropic scalar that zooms the initial estimates uniformly in each direction, we construct the scale diffusion target vector  $\mathbf{S}$  with each element equals to  $s_{gt}$  and has the same size as the depth feature map, and concatenate them as input to the model. We adopt the objective function proposed in [Ho et al. 2020] for training the diffusion U-Net  $\Phi_s$ :

$$\mathcal{L}_s = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I}), t} \left[ \|\epsilon - \epsilon_{\Phi_s}(t)\|_1 \right]. \quad (6)$$

At inference time, we sample  $\mathbf{S}$  by denoising a noise sample from the standard normal distribution, conditioned on  $\mathbf{D}$ , and take  $s = \text{avg}(\mathbf{S})$  as the isotropic scale factor. The re-scaled depths offer multiple hypotheses for plausible object translation and scale to bring the object from its canonical space to metric camera space.

### 3.4 CAD Alignment Diffusion

$\Phi_n$  models a distribution over likely object poses. We define the 9-DoF transformation as  $\mathbf{T} = [\mathbf{R}|\mathbf{t}|\mathbf{s}]$ , where  $\mathbf{R} \in \text{SO}(3)$  represents rotation,  $\mathbf{t} \in \mathbb{R}^3$  translation, and  $\mathbf{s} = (s_x, s_y, s_z)^T$  anisotropic scale. Instead of modeling this explicit pose representation, which must consider the different spaces for rotation, translation, and scale, we predict the normalized object coordinates (NOCs) [Wang et al. 2019]  $\mathbf{N}$  of the object. The NOC representation provides geometric correspondences between the observed object and its canonical coordinate system, enabling solving for  $\mathbf{T}$  as well as more generalized learning across different object shapes. As NOCs  $\mathbf{N}$  are structured



geometrically similarly to the back-projected points  $\mathbf{P}$  of  $\mathbf{D}_p$ , we condition their prediction on  $\mathbf{P}$ .

Features  $\mathbf{P}_f$  are extracted from condition  $\mathbf{P}$  using the 3DGC backbone [Lin et al. 2020].  $\Phi_n$  then operates on a concatenation of the target NOCs with their corresponding per-point features from  $\mathbf{P}_f$ . The objective is to minimize the difference between predicted and ground truth noise, formulated as:

$$\mathcal{L}_n = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I), t} \left[ \|\epsilon - \epsilon_{\Phi_n}(t)\|_1 \right]. \quad (7)$$

During inference, NOC candidates are sampled from Gaussian noise, conditioned on  $\mathbf{P}_f$ . We can then solve for the 9-DoF object transformation  $\mathbf{T}$  from the NOC estimates using a pose solver following a similar approach as presented in CaTGrasp [Wen et al. 2022], which is computed via RANSAC [Fischler and Bolles 1981] to find the transformation from the predicted NOCs and the observed point cloud  $\mathbf{P}$ .

### 3.5 CAD Retrieval Diffusion

Given our estimated NOCs capturing the visible geometry of the object mapped to its canonical space, we use this information to estimate the 3D shape for retrieval from a CAD database. As explicit 3D shape representations (e.g., voxels, points) are quite high-dimensional, we employ a latent space representation of objects, with latent codes  $\mathbf{z} \in \mathbb{R}^d$ .

To compress the shape into latent space, we opt for an architecture similar to ConvONet [Peng et al. 2020], denoted as  $\Phi_o$ . We add an additional MLP layer in the encoder to fuse the triplane features at the bottleneck into a global vector with a dimension of  $\mathbb{R}^d$ , and also feed it into the decoder along with interpolated triplane features. We pre-train  $\Phi_o$  to encode the CAD models with reconstruction loss following [Peng et al. 2020]. The samples in the latent space in which the  $\Phi_z$  learns to sample from are extracted from the encoding of the CAD database using  $\Phi_o$ .

To encode the information captured by the NOC estimate  $\mathbf{N}$ , we learn a positional feature embedding that maps point locations of  $\mathbf{N}$  from  $\mathbb{R}^3$  to  $\mathbb{R}^C$  together with a single-layer MLP that serves as the context of the condition for  $\Phi_z$ .  $\Phi_z$  then learns to denoise the object latent vector at each timestep. We follow recent works [Chou et al. 2023; Ramesh et al. 2022] and directly optimize for the original denoised latent vector  $\mathbf{z}_0$ :

$$\mathcal{L}_z = \|\Phi_n(\mathbf{z}_t) - \mathbf{z}_0\|_1. \quad (8)$$

During inference, we sample latent vector  $\mathbf{z}$  and query the CAD database to retrieve the nearest neighbor based on cosine similarity.

### 3.6 Synthetic Dataset Augmentation

We train our approach on a synthetic 3D scene dataset 3D-FRONT [Fu et al. 2021] to capture various scene scales, object arrangements, and shapes. However, this process can be prone to overfitting because the original 3D-FRONT scene configuration only contains 1,334 unique objects (for our target classes). To better capture a wider distribution of possible objects and arrangements, we further inject information from other existing large synthetic object databases, such as ShapeNet [Chang et al. 2015]. We consider the shape database of objects from 3D-FUTURE and those of ShapeNet, yielding

18,229 objects. We then augment the synthetic 3D scenes by replacing existing furniture with unused CAD models randomly retrieved from the same category in the database. This enables learning from more diverse object shapes and arrangements.

Our augmented synthetic dataset comprises  $\approx 300k$  images spanning 6 target categories.

### 3.7 Multi-Hypothesis Sampling from an Image

Since 3D CAD retrieval and alignment from a single view brings inherent ambiguities in terms of depth-scale, object shape, and pose, our approach addresses the ambiguities by a hierarchical multi-candidate sampling scheme. Our inference process begins by employing off-the-shelf 2D recognition backbones [Bhat et al. 2023; Eftekhar et al. 2021; Xu et al. 2023] to derive depth, mask, and normal estimates from the input image  $\mathbf{I}$ . Subsequently, we sample in a cascaded fashion from our trained diffusion models, progressing from scene-level reasoning to sampling explicit object representations and ending with the prediction of implicit object representations.

Initially, we sample a set of  $n$  potential scales, denoted as  $\{s_i\}_{i \in [1, n]}$ , from  $\Phi_s$  based on the predicted depth map  $\mathbf{D}$  from  $\mathbf{I}$ . NOC candidates are then sampled at each scene scale,  $\{\mathbf{N}_i\}_{i \in [1, n]}$ , given each scaled depth map, from  $\Phi_n$ . The corresponding 9-DoF transformations  $\{\mathbf{T}_i\}_{i \in [1, n]}$  are computed using a RANSAC-based solver [Fischler and Bolles 1981]. The possible latent candidates for objects,  $\mathbf{z}_i$ , are sampled conditionally on  $\mathbf{N}_i$  from the  $\Phi_z$ . We sample  $m$  object latent candidates given each  $\mathbf{N}_i$  and conduct a nearest neighbor search in the CAD model database to retrieve CAD models based on the cosine similarity between latent vectors.

The hierarchical sampling approach generates  $\{n \times m\}$  samples corresponding to various scene scales. For each scene scale  $s_i$ , a simple hypothesis ranking scheme selects the best candidate by rendering the normals of those CAD models using the solved poses  $\mathbf{T}_i$  and computing the similarity between the rendered normal and the machine-estimated normal from RGB input using LPIPS [Zhang et al. 2018]. The selection criterion is based on the model with the lowest LPIPS error, resulting in  $n$  sets of CAD models that align with the probabilistically permuted scenes.

### 3.8 Implementation Details

For  $\Phi_s$ , we use a learning rate of  $5e^{-5}$  and batch size 64, on a single RTX a6000 GPU for 2 days. For  $\Phi_n$ , we condition on subsampled point clouds with 1024 points, using learning rate  $5e^{-5}$  on a single RTX a6000 GPU with a batch size of 96, spanning 3 days per category.  $\Phi_o$  is trained on a single a100 GPU with a batch size of 128 for 3 days, with a learning rate  $1e^{-4}$ . We then train  $\Phi_z$  on an RTX a6000 GPU with batch size 128 and learning rate  $1e^{-4}$ , for 3 days per category. We detail the condition mechanisms of the diffusion models and the hyperparameters in the Appendix.

## 4 EXPERIMENTS

We validate our weakly-supervised method on real-world datasets, including ScanNet [Dai et al. 2017] and ARKit [Baruch et al. 2021]. As Scan2CAD [Avetisyan et al. 2019b] provides ShapeNet [Chang et al. 2015] annotations to ScanNet scenes, we evaluate CAD alignment accuracy and retrieval on the ScanNet25k image data, which contains




Fig. 4. **Qualitative Comparison on ScanNet images** [Avetisyan et al. 2019a; Dai et al. 2017]. Our weakly-supervised probabilistic approach produces more representative retrieval and alignment, even under strong occlusions (bottom), compared with in-domain supervised methods [Gümeli et al. 2022; Langer et al. 2022].

5k validation images. Given the challenges in ground truth object shape and pose matching in real data annotations like Scan2CAD, we introduce a probabilistic evaluation protocol for both in-domain supervised approaches and our cross-domain weakly-supervised method. In the absence of CAD annotations for ARKit, we present a qualitative-only evaluation on ARKit in the Fig. 6.

#### 4.1 Evaluation Metrics

Our evaluation protocols on ScanNet evaluate the top-n predictions, as is standard practice for evaluating probabilistic methods [Chun et al. 2021; Diller and Dai 2023; Guo et al. 2022; Neculai et al. 2022; Tevet et al. 2022] that acknowledges the probabilistic nature of the task and uncovering a distribution that adequately covers the imperfect single-mode annotation.

**4.1.1 Alignment Accuracy.** We introduce a new alignment evaluation protocol for single-view CAD alignment and retrieval. We observe that previous methods [Gümeli et al. 2022; Langer et al. 2022; Maninis et al. 2022] adopted a similar evaluation protocol as Scan2CAD, which operates on reconstructed scenes by aggregating predictions over multiple frames. Instead, we aim to evaluate the performance given only a single RGB image. We thus compute alignment accuracy following the same thresholds as in prior work [Avetisyan et al. 2019a; Gümeli et al. 2022; Langer et al. 2022; Maninis et al. 2022], but only considering a single RGB frame as input: an alignment is correct if the predicted object class is correct, translation error  $\leq 20\text{cm}$ , rotation error  $\leq 20^\circ$ , and scale ratio  $\leq 20\%$ . As multiple hypotheses should reflect a distribution from which the ground truth is likely to appear, we evaluate multiple hypotheses

from each method by evaluating how close the ground truth lies to any hypothesis, that is, the hypothesis with the minimum error as follows:

$$e = \arccos\left(\frac{\text{tr}(\hat{\mathbf{R}}\mathbf{R}_{gt}^T) - 1}{2}\right) + \|\hat{\mathbf{t}} - \mathbf{t}_{gt}\|_2 + \left\|\frac{\hat{s}}{s_{gt}} - 1\right\|_1, \quad (9)$$

where  $(\hat{\cdot})$  denotes predictions.

**4.1.2 Retrieval Similarity.** To evaluate the shape retrieval, we establish the retrieval similarity as the L1 Chamfer Distance between point clouds sampled from the retrieved mesh and the corresponding ground truth. We set the candidate CAD models for retrieval as the 1943 models that appear in the Scan2CAD training set. A total of  $10k$  points are sampled on the meshes for evaluation. To evaluate our method probabilistically, we report the candidate with the minimum L1 Chamfer Distance according to the ground truth.

**4.1.3 Metric Pose Accuracy.** We calculate average translation (in centimeters), rotation (in degrees), and scaling (in percentage) errors to directly evaluate the metric pose alignment.

**4.1.4 Average L1 Chamfer Distance for Retrieval.** Aside from the standard top-n probabilistic evaluation protocol above, we consider averaging the Chamfer Distance for the retrieval accuracy to illustrate the effectiveness of our diffusion-based retrieval. Note that this metric can penalize for plausible predictions that differ from the single ground truth annotation (e.g., ambiguity in the size of a couch that goes out of the image frame).

#### 4.2 Comparison to State of the Art

Tab. 1 compares our weakly-supervised method with state-of-the-art single-frame-based methods ROCA [Gümeli et al. 2022] and SPARC [Langer et al. 2022], both of which require full supervision, on ScanNet [Avetisyan et al. 2019a; Dai et al. 2017]. As the single image setting can contain depth-scale ambiguities, we evaluate multiple hypotheses for methods that allow so. ROCA is fully deterministic, and so only provides one hypothesis. While SPARC was not developed for multi-hypothesis reasoning, multiple samples can be drawn from different initialization angles (always including the authors' proposed initialization angle).

DiffCAD learns a much more representative distribution, surpassing fully-supervised state of the art with only 8 hypotheses, with notable improvement in both translation and rotation as in Tab. 3, reflecting the quality of our learned distribution. Performance increases with more hypotheses, with a slight saturation of around 20 hypotheses. In contrast, the performance of SPARC only improves marginally with increasing the sampled hypotheses.

While our single-hypothesis prediction slightly underperforms compared to fully-supervised approaches, this discrepancy is attributed to the scale-biased nature introduced by the domain gap of synthetic and real-world data. Given inherent imperfectly annotated real-world ground truth, our evaluation considers the likelihood of the ground truth originating from the sampled distribution.

Tab. 2 evaluates CAD retrieval similarity compared to the state-of-the-art. Since SPARC uses the same deterministic retrieval as ROCA, multiple hypotheses are not available for either method. With our

Table 1. **Alignment Accuracy on ScanNet** in comparison to state-of-the-art approaches, which require in-domain ground truth supervision. While SPARC [Langer et al. 2022] has not been developed for probabilistic reasoning, multiple hypotheses can be generated by sampling different initialization angles (always including those obtained with the officially released initialization angle). Evaluation under the same #hypotheses are colored in 8, 12, 16, and 20 hypotheses. Our probabilistic approach can outperform fully-supervised methods, while reflecting a more likely distribution of poses.

Method	in-domain supervision	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↑
ROCA [Gümeli et al. 2022]	✓	-	11.2	11.3	12.6	35.8	8.7	9.3	14.8
SPARC [Langer et al. 2022]	✓	1	26.1	21.7	27.8	47.4	25.9	19.7	28.1
SPARC [Langer et al. 2022]	✓	8	26.7	24.6	28.8	48.9	28.9	24.1	30.3
SPARC [Langer et al. 2022]	✓	16	29.2	26.1	30.5	49.6	31.9	26.4	32.3
SPARC [Langer et al. 2022]	✓	20	29.8	27.1	30.8	49.8	32.7	26.8	32.8
<b>Ours</b>	✗	1	13.0	4.4	9.6	22.4	13.3	4.9	11.3
<b>Ours</b>	✗	8	28.6	16.7	32.8	55.0	41.1	18.6	32.1
<b>Ours</b>	✗	12	30.4	20.2	35.1	58.5	41.1	21.6	34.5
<b>Ours</b>	✗	16	31.1	21.7	39.4	61.3	44.9	24.8	37.2
<b>Ours</b>	✗	20	32.9	24.1	42.1	62.5	47.5	27.1	39.4

probabilistic modeling, DiffCAD performs on par with state-of-the-art with one hypothesis and significantly improves with more hypotheses, better reflecting likely shape reconstructions, even with only weak supervision available. We further verify the effectiveness of our method by averaging the Chamfer Distance of the retrieved candidates as shown in Tab. 4; even under this metric, our model still outperforms the deterministic baselines.

In Fig. 4, we show a qualitative comparison of CAD retrieval and alignment on ScanNet images. DiffCAD achieves more accurate object retrieval and alignments across a diverse set of image views and object types due to our effective probabilistic modeling compared with in-domain supervised methods.

**4.2.1 Evaluation on ARKit.** We present qualitative results of our weakly-supervised approach on real-world ARKit [Baruch et al. 2021] data in Fig. 6, comparing against ROCA [Gümeli et al. 2022]. Our method demonstrates better robustness and accuracy in object retrieval and alignment across diverse images and object types.

### 4.3 Ablations

**4.3.1 Benefits of Probabilistic Model.** We consider two alternative deterministic baselines: using the identical UNet architecture as

our method but without diffusion, and a PointTransformer-V3 [Wu et al. 2023] for NOC prediction. Both baselines are trained on our synthetic train data and evaluated on the real test images. Tab. 5 shows that these deterministic baselines largely suffer from the domain gap and perform inferior to our single-hypothesis model. This indicates the effectiveness and cross-domain robustness of our probabilistic modeling. We detail the class-specific performance of the non-probabilistic baselines in the Appendix.

**4.3.2 Qualitative Evaluation of Probabilistic Retrieval.** Fig. 8 demonstrates multiple candidates featuring various styles of chair and sofa that propose various possible scales. In Example 3 of Fig. 6, our probabilistic method successfully identifies potentially different sizes of sofas (single- or double-seater) and occluded cabinets, surpassing deterministic retrieval from ROCA in terms of quality. Example 4 showcases the retrieval of beds with varying lengths, again attributable to occlusion. In Fig. 5, examples 2 and 4 further illustrate probabilistic reasoning, showcasing the accommodation of varying shapes and sizes despite limited observations.

**4.3.3 Effectiveness of Our Learned Distributions.** We consider our learned probabilistic distributions in comparison with alternative

Table 2. **Retrieval Similarity on ScanNet** compared to in-domain supervised state-of-the-arts methods. Note that SPARC [Langer et al. 2022] uses the same retrieval provided by ROCA [Gümeli et al. 2022]. Evaluation under the same #hypotheses are colored in 8, 12, 16, and 20 hypotheses. Our approach matches the performance of state of the art with 1 hypothesis, and reflects a likely distribution, retrieving more accurate objects with only a few hypotheses.

Method	in-domain supervision	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↓
ROCA [Gümeli et al. 2022]	✓	-	0.087	0.089	0.130	0.100	0.099	0.132	0.106
SPARC [Langer et al. 2022]	✓	-	0.087	0.089	0.130	0.100	0.099	0.132	0.106
<b>Ours</b>	✗	1	0.117	0.108	0.115	0.115	0.093	0.145	0.116
<b>Ours</b>	✗	8	0.075	0.064	0.079	0.075	0.066	0.089	0.075
<b>Ours</b>	✗	12	0.065	0.060	0.074	0.069	0.062	0.083	0.069
<b>Ours</b>	✗	16	0.061	0.057	0.071	0.066	0.060	0.077	0.065
<b>Ours</b>	✗	20	0.058	0.055	0.069	0.064	0.058	0.072	0.063




Fig. 5. **Qualitative Results on ScanNet images.** Our probabilistic approach shows multi-feasible sets of object shape and pose pairs given the ambiguities in monocular perception. Left-bottom: The two hypotheses corresponding to the smallest and largest scene scale reconstructions follow possible depth-scale ambiguity from the camera view.






Fig. 6. **Qualitative Results on ARKit images.** Our approach presents robust retrieval and alignment to various scenes, reconstructing the scene with multi-feasible sets of object shape and pose pairs given the ambiguities in monocular perception. Dotted: The three hypotheses corresponding to the different scene scales.

Table 3. **Pose Alignment Metric.** We calculate average translation (in centimeters), rotation (in degrees), and scaling (in percentage) errors to directly evaluate the metric pose alignment. Without any training on real data, our diffusion approach achieves notable improvement in translation and rotation, with the same scale performance.

Method	#hypotheses	Translation↓	Rotation↓	Scale↓
ROCA [Gümel et al. 2022]	-	43.5	27.9	0.19
SPARC [Langer et al. 2022]	8	35.1	19.5	0.19
<b>Ours</b>	8	<b>27.5</b>	<b>18.6</b>	0.19

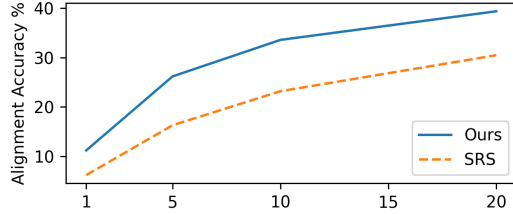


Fig. 7. **Scale Sampling Ablation.** Our learned probabilistic model for  $\Phi_s$  shows much higher sampling efficiency than non-parametric sampling.



Fig. 8. **Multi-hypothesis of Retrieval.** Our approach illustrates various styles of CAD candidates fitted to possible scene scales conditioned on the same input, capturing the solutions to the ambiguities in monocular perception and shape reasoning. Dotted: Hypotheses gathering.

parametric sampling schemes. We study this for  $\Phi_s$ , as we can employ a straightforward sampling baseline: we estimate the mean  $\mu'$  and variance  $\sigma'^2$  of scale offsets between predicted and ground truth train depth maps, and instead draw samples from  $\mathcal{N}(\mu', \sigma')$

Table 4. **Average L1 Chamfer Distance for Retrieval.** We average the Chamfer Distance for the retrieval accuracy evaluation. Our approach effectively outperform the deterministic retrieval baselines.

Method	#hypotheses	Retrieval Similarity↓
ROCA [Gümel et al. 2022]	-	0.106
SPARC [Langer et al. 2022]	-	0.106
<b>Ours</b>	8	<b>0.088</b>

Table 5. **Non-probabilistic Baselines for Alignment Accuracy on ScanNet.** We compare two non-probabilistic baselines that are trained on the synthetic domain, which include a transformer-based backbone [Wu et al. 2023] and a UNet architecture same as Ours. Our method archives much better cross-domain performance on class average alignment metric.

Method	in-domain supervision	#hypotheses	avg↑
PointTransformer [Wu et al. 2023]	✗	-	6.9
UNet (same diffusion backbone)	✗	-	10.1
<b>Ours</b>	✗	1	<b>11.3</b>

(denoted as ‘Smart Random Sampling (SRS)’ in Fig. 7). Fig. 7 illustrates that our scene scale diffusion model  $\Phi_s$  outperforms SRS with much higher sampling efficiency.

4.3.4 *Predicting NOCs Instead of Explicit Transformation Parameters* T. Tab. 6 shows that our NOC-based proxy for alignment estimation enables much more robust alignment than the direct prediction of the 9-DoF values of T. The NOC proxy estimation enables more robust estimation through its dense correspondences, which also helps mitigate the potential domain gap between synthetic training data and real-world test scenarios. We consider two alternatives to explicit T prediction: ‘ $\Phi_n + P$ ’ predicts 9-DoF values of T along with NOCs, and ‘ $\Phi_n + S$ ’ predicts only object scale along with NOCs. In ‘ $\Phi_n + P$ ’, NOCs are solely used for retrieval, whereas in ‘ $\Phi_n + S$ ’, NOCs are utilized for 6-DoF rotation and translation solving, and retrieval. We evaluate both for the explicit pose prediction as well as for the pose solved by the NOCs via the solver, and select the one with lower error according to Eq. 9. Our approach to leveraging NOCs and an algorithmic pose solver achieves significantly better cross-domain alignment accuracy.

Table 6. **Pose Estimation Ablation.** With synthetic data replacement augmentation, our NOC-based predictions from  $\Phi_n$  enable more generalized feature learning, resulting in better pose alignment.

Method	#hypotheses	Alignment Accuracy↑
$\Phi_n + P$	8	16.3
$\Phi_n + S$	8	19.9
$\Phi_n$ w/o Aug	8	25.8
<b>Ours</b>	8	<b>32.1</b>

4.3.5 *Effect of Canonical Object Representation for Retrieval.* We consider performing CAD retrieval in the camera space of the image and the canonical space of the CAD database in Tab. 7. We consider a baseline alternative, ‘ $\Phi_z + PC$ ’, which operates on features from the back-projected point cloud in camera space as the retrieval condition. Learning the shape feature embeddings in canonical space significantly improves retrieval.

4.3.6 *Object Replacement Augmentation.* We illustrate the necessity of augmenting the synthetic scene layout with extra shapes (i.e., from ShapeNet [Chang et al. 2015]) in Tabs. 6 and 7. Both  $\Phi_n$  and  $\Phi_z$

Table 7. **CAD Retrieval Ablation.** Learning retrieval in the object canonical space (vs. camera space, denoted as PC) and with synthetic augmentation enables improved retrieval performance.

Method	#hypotheses	Retrieval Similarity↓
$\Phi_z + PC$	8	0.111
$\Phi_z$ w/o Aug	8	0.082
<b>Ours</b>	8	<b>0.075</b>

trained on the augmented synthetic dataset improve significantly compared with ‘ $\Phi_n + w/o Aug$ ’ and ‘ $\Phi_z + w/o Aug$ ’ respectively, due to the increased diversity during training.

**4.3.7 Random Neighbors of Deterministic Retrieval.** We conduct experiments retrieving random 8 neighbors of the deterministic retrieval from ROCA which gives 0.093 of L1 Chamfer Distance, which falls behind our diffusion-based retrieval pipeline of 0.075 as shown in Tab. 2

#### 4.4 Limitations

While leveraging 3D perception through CAD model retrieval and alignment offers a compact representation of the scene, its efficacy for applications requiring exact reconstructions can be hindered by the absence of precise geometric matches in real-world environments. A potential approach to overcome this limitation involves deforming the extracted CAD models to enhance alignment with observations [Di et al. 2023; Ishimtsev et al. 2020; Uy et al. 2020, 2021]. Additionally, our approach does not model object relations explicitly, which could limit performance, since indoor scenes are often arranged in coherent global structures. Integrating considerations of scene context alongside object deformation represents a promising direction for achieving more accurate 3D perception.

## 5 CONCLUSION

We introduce DiffCAD, the first weakly-supervised probabilistic approach for single-image CAD model retrieval and alignment. By disentangling ambiguities in the monocular perceptual system through individual distribution modeling using diffusion, we effectively address uncertainties in scene depth scale, object shape, and pose. Notably, our diffusion models are trained only on synthetic datasets, and yet outperform the supervised state-of-the-art approach on real-world Scan2CAD image data [Avetisyan et al. 2019a], achieving a 5.9% improvement with only 8 hypotheses. We envision that this advancement will spur further progress in 3D probabilistic models, as well as 3D perception without necessitating real-world labels.

## ACKNOWLEDGMENTS

This project is funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the ERC Starting Grant SpatialSem (101076253), and the German Research Foundation (DFG) Grant “Learning How to Interact with Scenes through Part-Based Understanding”. We also thank Quan Meng, Lei Li and Hanzhi Chen for the constructive discussions.

## REFERENCES

- Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. 2019a. Scan2cad: Learning cad model alignment in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*. 2614–2623.
- Armen Avetisyan, Angela Dai, and Matthias Nießner. 2019b. End-to-end cad model retrieval and 9dof alignment in 3d scans. In *Proceedings of the IEEE/CVF International Conference on computer vision*. 2551–2560.
- Armen Avetisyan, Tatiana Khanova, Christopher Choy, Denver Dash, Angela Dai, and Matthias Nießner. 2020. Scenecad: Predicting object alignments and layouts in rgb-d scans. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*. Springer, 596–612.
- Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. 2022. IronDepth: Iterative Refinement of Single-View Depth using Surface Normal and its Uncertainty. *ArXiv abs/2210.03676* (2022). <https://api.semanticscholar.org/CorpusID:252762221>
- Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrukov, and Artem Babenko. 2021. Label-Efficient Semantic Segmentation with Diffusion Models. *ArXiv abs/2112.03126* (2021). <https://api.semanticscholar.org/CorpusID:244908617>
- Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, and Elad Shulman. 2021. ARKitScenes - A Diverse Real-World Dataset for 3D Indoor Scene Understanding Using Mobile RGB-D Data. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*. [https://openreview.net/forum?id=tjZjv\\_gh\\_CE](https://openreview.net/forum?id=tjZjv_gh_CE)
- Tim Beyer and Angela Dai. 2022. Weakly-Supervised End-to-End CAD Retrieval to Scan Objects. *ArXiv abs/2203.12873* (2022). <https://api.semanticscholar.org/CorpusID:247627889>
- S. Bhat, Ibraheem Alhassim, and Peter Wonka. 2020. AdaBins: Depth Estimation Using Adaptive Bins. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), 4008–4017. <https://api.semanticscholar.org/CorpusID:218889832>
- Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. 2023. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288* (2023).
- Andreas Blattmann, Robin Rombach, Kaan Oktay, Jonas Müller, and Björn Ommer. 2022. Retrieval-augmented diffusion models. *Advances in Neural Information Processing Systems* 35 (2022), 15309–15324.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-End Object Detection with Transformers. *ArXiv abs/2005.12872* (2020). <https://api.semanticscholar.org/CorpusID:227227779>
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015).
- Wenhu Chen, Hexiang Hu, Chitwan Saharia, and William W. Cohen. 2022. Re-Imagen: Retrieval-Augmented Text-to-Image Generator. *ArXiv abs/2209.14491* (2022). <https://api.semanticscholar.org/CorpusID:252596087>
- Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. 2023. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4456–4465.
- Gene Chou, Yuval Bahat, and Felix Heide. 2022. DiffusionSDF: Conditional Generative Modeling of Signed Distance Functions. *ArXiv abs/2211.13757* (2022). <https://api.semanticscholar.org/CorpusID:254017862>
- Gene Chou, Yuval Bahat, and Felix Heide. 2023. Diffusion-sdf: Conditional generative modeling of signed distance functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2262–2272.
- Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 2016. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*. Springer, 628–644.
- Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. 2021. Probabilistic embeddings for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8415–8424.
- Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5828–5839.
- Congyue Deng, Chiyu Max Jiang, C. Qi, Xinchun Yan, Yin Zhou, Leonidas J. Guibas, and Drago Angelov. 2022. NeRDi: Single-View NeRF Synthesis with Language-Guided Diffusion as General Image Priors. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 20637–20647. <https://api.semanticscholar.org/CorpusID:254366717>
- Maximilian Denninger, Dominik Winkelbauer, Martin Sundermeyer, Wout Boerdijk, Markus Knauer, Klaus H. Strobl, Matthias Humt, and Rudolph Triebel. 2023. Blender-Proc2: A Procedural Pipeline for Photorealistic Rendering. *Journal of Open Source Software* 8, 82 (2023), 4901. <https://doi.org/10.21105/joss.04901>

- Yan Di, Chenyangguang Zhang, Ruida Zhang, Fabian Manhardt, Yongzhi Su, Jason Rambach, Didier Stricker, Xiangyang Ji, and Federico Tombari. 2023. U-RED: Unsupervised 3D Shape Retrieval and Deformation for Partial Point Clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 8884–8895.
- Christian Diller and Angela Dai. 2023. CG-HOI: Contact-Guided 3D Human-Object Interaction Generation. *arXiv preprint arXiv:2311.16097* (2023).
- Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. 2021. Omnidata: A Scalable Pipeline for Making Multi-Task Mid-Level Vision Datasets From 3D Scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10786–10796.
- Ziya Erkoç, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. 2023. Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. *arXiv preprint arXiv:2303.17015* (2023).
- Haoqing Fan, Hao Su, and Leonidas J Guibas. 2017. 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*. 605–613.
- Yuxin Fang, Wen Wang, Binhui Xie, Quan-Sen Sun, Ledell Yu Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2022. EVA: Exploring the Limits of Masked Visual Representation Learning at Scale. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 19358–19369. <https://api.semanticscholar.org/CorpusID:253510587>
- Martin A Fischler and Robert C Bolles. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24, 6 (1981), 381–395.
- Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 2021. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10933–10942.
- Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. 2021. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430* (2021).
- Golnaz Ghiasi, Xiuyi Gu, Yin Cui, and Tsung-Yi Lin. 2022. Scaling Open-Vocabulary Image Segmentation with Image-Level Labels. In *ECCV*.
- Georgia Gkioxari, Jitendra Malik, and Justin Johnson. 2019. Mesh r-cnn. In *Proceedings of the IEEE/CVF international conference on computer vision*. 9785–9795.
- Vitor Guizilini, Igor Vasiljevic, Dian Chen, Rares Ambrus, and Adrien Gaidon. 2023. Towards Zero-Shot Scale-Aware Monocular Depth Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Can Gümeli, Angela Dai, and Matthias Nießner. 2022. Roca: Robust cad model retrieval and alignment from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4022–4031.
- Chuan Guo, Shihao Zou, Xinxi Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. 2022. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5152–5161.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 2023. Unsupervised Semantic Correspondence Using Stable Diffusion. *arXiv preprint arXiv:2305.15581* (2023).
- Jonathan Ho, Ajay Jain, and P. Abbeel. 2020. Denoising Diffusion Probabilistic Models. *ArXiv abs/2006.11239* (2020). <https://api.semanticscholar.org/CorpusID:219955663>
- Muhammad Zubair Irshad, Thomas Kollar, Michael Laskey, Kevin Stone, and Zsolt Kira. 2022. CenterSnap: Single-Shot Multi-Object 3D Shape Reconstruction and Categorical 6D Pose and Size Estimation. *2022 International Conference on Robotics and Automation (ICRA)* (2022), 10632–10640. <https://api.semanticscholar.org/CorpusID:247222831>
- Vladislav Ishimtev, Alexey Bokhovkin, Alexey Artemov, Savva Ignatyev, Matthias Niessner, Denis Zorin, and Evgeny Burnaev. 2020. Cad-deform: Deformable fitting of cad models to 3d scans. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*. Springer, 599–628.
- Hamid Izadinia, Qi Shan, and Steven M Seitz. 2017. Im2cad. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5134–5143.
- Young Min Kim, Niloy J Mitra, Dong-Ming Yan, and Leonidas Guibas. 2012. Acquiring 3d indoor environments with variability and repetition. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–11.
- Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. 2021. Variational Diffusion Models. *ArXiv abs/2107.00630* (2021). <https://api.semanticscholar.org/CorpusID:235694314>
- 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).
- Juil Koo, Seungwoo Yoo, Minh Hoai Nguyen, and Minhyuk Sung. 2023. SALAD: Part-Level Latent Diffusion for 3D Shape Generation and Manipulation. *ArXiv abs/2303.12236* (2023). <https://api.semanticscholar.org/CorpusID:257663544>
- Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. 2020. Mask2cad: 3d shape prediction by learning to segment and retrieve. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 260–277.
- Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. 2021. Patch2cad: Patchwise embedding learning for in-the-wild shape retrieval from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12589–12599.
- Florian Langer, Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. 2022. SPARC: Sparse Render-and-Compare for CAD model alignment in a single RGB image. *arXiv preprint arXiv:2210.01044* (2022).
- Florian Langer, Ignas Budvytis, and Roberto Cipolla. 2023. Sparse Multi-Object Render-and-Compare. *arXiv preprint arXiv:2310.11184* (2023).
- Boyi Li, Kilian Q. Weinberger, Serge J. Belongie, Vladlen Koltun, and René Ranftl. 2022c. Language-driven Semantic Segmentation. *ArXiv abs/2201.03546* (2022). <https://api.semanticscholar.org/CorpusID:245836975>
- Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. 2022b. Diffusion-SDF: Text-to-Shape via Voxelized Diffusion. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 12642–12651. <https://api.semanticscholar.org/CorpusID:254366593>
- Yangyan Li, Angela Dai, Leonidas Guibas, and Matthias Nießner. 2015. Database-assisted object retrieval for real-time 3d reconstruction. In *Computer graphics forum*, Vol. 34. Wiley Online Library, 435–446.
- Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. 2022a. DepthFormer: Exploiting Long-range Correlation and Local Information for Accurate Monocular Depth Estimation. *Machine Intelligence Research* 20 (2022), 837 – 854. <https://api.semanticscholar.org/CorpusID:247762153>
- Feng Liang, Bichen Wu, Xiaoliang Dai, Kungpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Péter Vajda, and Diana Marculescu. 2022. Open-Vocabulary Semantic Segmentation with Mask-adapted CLIP. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 7061–7070. <https://api.semanticscholar.org/CorpusID:252780581>
- Kai-En Lin, Yen-Chen Lin, Wei-Sheng Lai, Tsung-Yi Lin, Yichang Shih, and Ravi Ramamoorthi. 2022. Vision Transformer for NeRF-Based View Synthesis from a Single Input Image. *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2022), 806–815. <https://api.semanticscholar.org/CorpusID:250450901>
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2117–2125.
- Zhi-Hao Lin, Sheng-Yu Huang, and Yu-Chiang Frank Wang. 2020. Convolution in the cloud: Learning deformable kernels in 3d graph convolution networks for point cloud analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1800–1809.
- Haolin Liu, Yujian Zheng, Guanying Chen, Shuguang Cui, and Xiaoguang Han. 2022. Towards High-Fidelity Single-view Holistic Reconstruction of Indoor Scenes. *ArXiv abs/2207.08656* (2022). <https://api.semanticscholar.org/CorpusID:250627520>
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.
- Grace Luo, Lisa Dunlap, Dong Huk Park, Aleksander Holynski, and Trevor Darrell. 2023. Diffusion Hyperfeatures: Searching Through Time and Space for Semantic Correspondence. In *Advances in Neural Information Processing Systems*.
- Priyanka Mandikal, L. NavaneethK., Mayank Agarwal, and R. Venkatesh Babu. 2018. 3D-LMNet: Latent Embedding Matching for Accurate and Diverse 3D Point Cloud Reconstruction from a Single Image. *ArXiv abs/1807.07796* (2018). <https://api.semanticscholar.org/CorpusID:49905039>
- Kevis-Kokitsi Maninis, Stefan Popov, Matthias Nießner, and Vittorio Ferrari. 2022. Vid2cad: Cad model alignment using multi-view constraints from videos. *IEEE transactions on pattern analysis and machine intelligence* 45, 1 (2022), 1320–1327.
- 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.
- Gimin Nam, Mariem Khlifi, Andrew Rodriguez, Alberto Tono, Linqi Zhou, and Paul Guerrero. 2022. 3D-LDM: Neural Implicit 3D Shape Generation with Latent Diffusion Models. *ArXiv abs/2212.00842* (2022). <https://api.semanticscholar.org/CorpusID:254220714>
- Liangliang Nan, Ke Xie, and Andrei Sharf. 2012. A search-classify approach for cluttered indoor scene understanding. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–10.
- Andrei Neculai, Yanbei Chen, and Zeynep Akata. 2022. Probabilistic compositional embeddings for multimodal image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4547–4557.
- Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jianjun Zhang. 2020. Total3DUnderstanding: Joint Layout, Object Pose and Mesh Reconstruction for Indoor Scenes From a Single Image. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), 52–61. <https://api.semanticscholar.org/CorpusID:211532831>
- Junyi Pan, Xiaoguang Han, Weikai Chen, Jiapeng Tang, and Kui Jia. 2019. Deep Mesh Reconstruction From Single RGB Images via Topology Modification Networks. *2019*



- IEEE/CVF International Conference on Computer Vision (ICCV)* (2019), 9963–9972. <https://api.semanticscholar.org/CorpusID:202121070>
- William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4195–4205.
- Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. 2020. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 523–540.
- Michael Ramamonjisoa and Vincent Lepetit. 2019. SharpNet: Fast and Accurate Recovery of Occluding Contours in Monocular Depth Estimation. *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)* (2019), 2109–2118. <https://api.semanticscholar.org/CorpusID:160009795>
- Aditya Ramesh, Pratul Dharwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* 1, 2 (2022), 3.
- René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*. 12179–12188.
- Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. 2020. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501* (2020).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. [arXiv:2112.10752 \[cs.CV\]](https://arxiv.org/abs/2112.10752)
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115 (2015), 211–252.
- Manuel Schwöbner, Joshua Niemeijer, Jan-Aike Termöhlen, Jörg P Schäfer, Nico M Schmidt, Hanno Gottschalk, and Tim Fingscheidt. 2023. Survey on unsupervised domain adaptation for semantic segmentation for visual perception in automated driving. *IEEE Access* (2023).
- Tianjia Shao, Weiwei Xu, Kun Zhou, Jingdong Wang, Dongping Li, and Baining Guo. 2012. An interactive approach to semantic modeling of indoor scenes with an rgb-d camera. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–11.
- Shelly Sheynin, Oron Ashual, Adam Polyak, Uriel Singer, Oran Gafni, Eliya Nachmani, and Yaniv Taigman. 2022. Knn-diffusion: Image generation via large-scale retrieval. *arXiv preprint arXiv:2204.02849* (2022).
- J Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 2023. 3d neural field generation using triplane diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20875–20886.
- Jascha Narain Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *ArXiv abs/1503.03585* (2015). <https://api.semanticscholar.org/CorpusID:14888175>
- Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* 32 (2019).
- Yang Song, Jascha Narain Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. *ArXiv abs/2011.13456* (2020). <https://api.semanticscholar.org/CorpusID:227209335>
- David Stutz and Andreas Geiger. 2020. Learning 3d shape completion under weak supervision. *International Journal of Computer Vision* 128 (2020), 1162–1181.
- Luming Tang, Menglin Jia, Qianqian Wang, Cheng Peng Phoo, and Bharath Hariharan. 2023. Emergent Correspondence from Image Diffusion. *arXiv preprint arXiv:2306.03881* (2023).
- Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. 2022. Human motion diffusion model. *arXiv preprint arXiv:2209.14916* (2022).
- Mikaela Angelina Uy, Jingwei Huang, Minhyuk Sung, Tolga Birdal, and Leonidas Guibas. 2020. Deformation-aware 3d model embedding and retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*. Springer, 397–413.
- Mikaela Angelina Uy, Vladimir G Kim, Minhyuk Sung, Noam Aigerman, Siddhartha Chaudhuri, and Leonidas J Guibas. 2021. Joint learning of 3d shape retrieval and deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11713–11722.
- Qun Wan, Yidong Li, Haidong Cui, and Zheng Mao Feng. 2019. 3D-MaskGAN: Unsupervised Single-View 3D Object Reconstruction. *2019 6th International Conference on Behavioral, Economic and Socio-Cultural Computing (BESCC)* (2019), 1–6. <https://api.semanticscholar.org/CorpusID:210888106>
- He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. 2019. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2642–2651.
- Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, W. Liu, and Yu-Gang Jiang. 2018. Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. *ArXiv abs/1804.01654* (2018). <https://api.semanticscholar.org/CorpusID:4633214>
- Bowen Wen, Wenzhao Lian, Kostas Bekris, and Stefan Schaal. 2022. CaTGrasp: Learning Category-Level Task-Relevant Grasping in Clutter from Simulation. *ICRA 2022* (2022).
- Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. 2023. Point transformer v3: Simpler, faster, stronger. *arXiv preprint arXiv:2312.10035* (2023).
- Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaocong Wang, and Shalini De Mello. 2023. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2955–2966.
- Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. 2020. pixelNeRF: Neural Radiance Fields from One or Few Images. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), 4576–4585. <https://api.semanticscholar.org/CorpusID:227254854>
- Yanjie Ze and Xiaocong Wang. 2022. Category-level 6d object pose estimation in the wild: A semi-supervised learning approach and a new dataset. *Advances in Neural Information Processing Systems* 35 (2022), 27469–27483.
- Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. 2022. LION: Latent Point Diffusion Models for 3D Shape Generation. *ArXiv abs/2210.06978* (2022). <https://api.semanticscholar.org/CorpusID:252872881>
- Biao Zhang, Jiapeng Tang, Matthias Nießner, and Peter Wonka. 2023. 3DShape2VecSet: A 3D Shape Representation for Neural Fields and Generative Diffusion Models. *ACM Transactions on Graphics (TOG)* 42 (2023), 1–16. <https://api.semanticscholar.org/CorpusID:256358401>
- Cheng Zhang, Zhaopeng Cui, Yinda Zhang, Bing Zeng, Marc Pollefeys, and Shuaicheng Liu. 2021. Holistic 3D Scene Understanding from a Single Image with Implicit Representation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021), 8829–8838. <https://api.semanticscholar.org/CorpusID:232185507>
- Jiyao Zhang, Mingdong Wu, and Hao Dong. 2024. Generative Category-level Object Pose Estimation via Diffusion Models. *Advances in Neural Information Processing Systems* 36 (2024).
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- Linqi Zhou, Yilun Du, and Jiajun Wu. 2021. 3D Shape Generation and Completion through Point-Voxel Diffusion. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (2021), 5806–5815. <https://api.semanticscholar.org/CorpusID:233182041>

## A ADDITIONAL RESULTS

We present additional ablations and category-specific results for our diffusion-based scene scale, object shape and pose experiments, as well as for the non-probabilistic baselines.

### A.1 Scene Scale Ablation

Tab. 8 compares our scene scale diffusion module  $\Phi_S$  to a baseline approach without the feature extractor  $\mathcal{E}_S$ . We evaluate the scale prediction accuracy defined as  $e_s = \min_{i \in \{N\}} \|s_{gt} - s_i\|_1$ . Our design achieves a lower scale error compared to the baseline, indicating improved scale prediction accuracy.

### A.2 CAD Alignment Ablation

We present category-specific evaluation results in Tab. 9, corresponding to Tab.3 in the main paper. Tab. 9 demonstrates that with synthetic data augmentation and object NOC representation, our CAD alignment diffusion model  $\Phi_N$  achieves superior alignment performance across domains for all categories.

### A.3 CAD Retrieval Ablation

In Tab. 10, corresponding to Tab.4 in the main paper, we highlight the advantages of learning the retrieval condition in the canonical space, along with the benefits of synthetic data augmentation for the CAD retrieval diffusion model  $\Phi_Z$ .

#### A.4 Non-probabilistic Baselines for Alignment

In Tab. 11, corresponding to Tab.5 in the main paper, the non-probabilistic baselines suffer from a significant domain gap compared with our approach, verifying the robustness of our diffusion-based weakly supervised strategy.

#### A.5 Confidence Measure of Probabilistic Retrieval.

We estimate model confidence by sample variance, and find that higher-confidence predictions yield improved performance: for object retrieval, performance improves from 0.075 to 0.065 in terms of top-n Chamfer Distance evaluation when considering only high-confidence predictions. We determine high-confidence predictions by setting a  $2\text{-}\sigma$  threshold based on the variance.

Table 8. **Scene scale ablation.** Encoding the feature from the depth map using  $\mathcal{E}_s$  enables better scene scale modeling for the target categories.

Method	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↓
$\Phi_s$ w/o $\mathcal{E}_s$	8	0.172	0.202	0.369	0.284	0.297	0.341	0.278
<b>Ours</b>	8	<b>0.133</b>	<b>0.195</b>	<b>0.331</b>	<b>0.251</b>	<b>0.274</b>	<b>0.312</b>	<b>0.249</b>

Table 9. **Pose estimation ablation.** With synthetic data replacement augmentation, our NOC-based representation from  $\Phi_n$  achieves more generalized feature learning, resulting in better pose alignment for all target categories.

Method	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↑
$\Phi_n + P$	8	20.5	9.4	12.3	22.9	22.1	10.5	16.3
$\Phi_n + S$	8	20.5	14.3	12.9	29.3	27.8	14.6	19.9
$\Phi_n$ w/o Aug	8	23.0	12.8	24.5	45.3	37.3	12.0	25.8
<b>Ours</b>	8	<b>28.6</b>	<b>16.7</b>	<b>32.8</b>	<b>55.0</b>	<b>41.1</b>	<b>18.6</b>	<b>32.1</b>

Table 10. **CAD retrieval ablation.** Learning retrieval in the object canonical space (vs. camera space, denoted as ‘PC’) and with synthetic augmentation enables improved retrieval performance for our classes of interest.

Method	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↓
$\Phi_z + PC$	8	0.105	0.088	0.155	0.115	0.093	0.108	0.111
$\Phi_z$ w/o Aug	8	0.080	0.077	0.081	0.085	0.069	0.097	0.082
<b>Ours</b>	8	<b>0.075</b>	<b>0.064</b>	<b>0.079</b>	<b>0.075</b>	<b>0.066</b>	<b>0.089</b>	<b>0.075</b>

Table 11. **Non-probabilistic Baselines for Alignment Accuracy on ScanNet.** We compare two non-probabilistic baselines that are trained on the synthetic domain, which include a transformer-based backbone [Wu et al. 2023] and a UNet architecture same as Ours. Our method archives much better cross-domain performance on class average alignment metric.

Method	in-domain supervision	#hypotheses	bed	bkshlf	cabinet	chair	sofa	table	avg↑
PointTransformer [Wu et al. 2023]	✗	-	5.38	3.14	1.66	19.89	10.44	1.09	6.93
UNet (same diffusion backbone)	✗	-	12.72	<b>6.99</b>	1.42	19.54	<b>19.68</b>	0	10.06
<b>Ours</b>	✗	1	<b>13.0</b>	4.4	<b>9.6</b>	<b>22.4</b>	13.3	<b>4.9</b>	<b>11.3</b>




Fig. 9. **Overview of the Depth Feature Extractor  $\mathcal{E}_s$ .** The feature extractor takes as input the predicted depth map, and processes it with hierarchical convolution blocks following [Long et al. 2015], and outputs the extracted feature with the same spatial dimension as input. We indicate the feature dimension under each block.

## B DATA PREPARATION

### B.1 Object Mesh Pre-processing

To encode CAD models into a compact latent space using  $\Phi_o$ , we first canonicalize the original meshes from 3D-FUTURE [Fu et al. 2021] and ShapeNet [Chang et al. 2015] to have a normalized scale and consistent orientation. We save the scaling factor between the original and the canonicalized object for the synthetic data augmentation. Subsequently, we transform them into watertight meshes following the mesh-fusion process proposed in [Stutz and Geiger 2020].

### B.2 Synthetic Dataset Augmentation

We leverage the shape databases of objects from both 3D-FUTURE and ShapeNet, with a total of 18,229 objects. We substitute the original furniture of our target class within the 3D-FRONT layout. To achieve this, we randomly select an unused CAD model from the corresponding category in the database. Then, we scale the chosen object using the scaling factor of the original object, ensuring the preservation of object size balance in the augmented scenes. This augmentation process facilitates the learning from a more diverse set of object shapes and arrangements.

### B.3 Synthetic Data Rendering

We employ BlenderProc [Denninger et al. 2023] for our synthetic data rendering, generating RGB images, depth maps, and masks. Camera views are sampled by considering ray intersections with a minimum coverage of 15% of the object of interest, highlighting our focus on object-centric learning. Additionally, an off-the-shelf depth estimator, such as Zoedepth [Bhat et al. 2023], is used to obtain depth estimates over the rendered color images. This information is necessary for calculating the target scale  $s$  by comparing predicted  $D$  and rendered depth values  $D_{gt}$  regarding Eq.5.

## C FURTHER IMPLEMENTATION DETAILS

### C.1 Network Architecture

**C.1.1 Depth feature extractor  $\mathcal{E}_s$ .** Fig. 9 diagrams the depth feature extractor network. We adopt the FCN-ResNet50 [Long et al. 2015] and adjust its first convolution layer to take a single channel depth




Fig. 10. **Overview of the Object Point Cloud Feature Extractor  $\mathcal{E}_n$ .** The feature extractor takes the back-projected object point cloud as input, and extracts the per-point feature using GCN3D [Lin et al. 2020]. We indicate the per-point feature dimension below each 3D graph convolution block.

map as input, and use the feature before the final classification layer as the condition for the scene scale diffusion model  $\Phi_s$ .

**C.1.2 Object point cloud feature extractor  $\mathcal{E}_n$ .** We utilize the GCN3D [Lin et al. 2020] as feature extraction backbone given the object point cloud  $P$  back-projected from masked depth map  $D_p$ . Fig. 10 shows the object per-point feature extraction process, the feature then serves as input to the CAD alignment diffusion model  $\Phi_n$ .

**C.1.3 NOC point embedding  $\mathcal{P}_z$ .** Given that NOCs inherently capture the observed object in its canonical space, we focus on learning the features of NOCs capable of distinguishing intra-class shape variations. To achieve this objective, we directly learn the feature embedding of NOCs through trigonometric mappings. Subsequently, a single-layer MLP is employed to project the per-point features into a 512-dimensional space. This NOC feature then serves as a condition for the CAD retrieval diffusion model  $\Phi_z$ .

**C.1.4 Condition Mechanism.** For  $\Phi_s$  and  $\Phi_n$ , we concatenate the condition feature with the noised target vector as input to the diffusion UNet. For  $\Phi_z$ , we learn the positional feature embedding given the condition and use cross-attention to inject the feature at every diffusion UNet block following [Cheng et al. 2023] which maps explicit shape geometry with object code in the latent space.

**C.1.5 Network Architecture Choice.** Our design can also be extended to diffusion-based autoregressive backbone such as DiT [Peebles and Xie 2023]. An essential modification involves properly tokenizing the target. By initially encoding the target vector and adopting the patchify operation employed in DiT, our approach aligns with the principles of the aforementioned architecture.

## C.2 Test Split Generation

We introduce a new evaluation protocol for single-view CAD model retrieval and alignment on the ScanNet [Avetisyan et al. 2019a; Dai et al. 2017]. Our per-frame test split is designed based on the validation set of the ScanNet25k image dataset, aligning with previous methods [Gümeli et al. 2022; Langer et al. 2022]. For each target category, we first filter out the frames in which the centroid of the target object falls beyond the image plane. To avoid penalizing the

heavily occluded objects, frames where the visible mask of the target object occupies less than 10% of the total image plane are also excluded. The visible masks are obtained by comparing the rendered depth map using ground truth pose with the original sensory depth map. In total, our test split contains around 2.5k images across the 6 target categories.

## C.3 Hyperparameters for Pose Solver

For RANSAC pose, we use OpenCV estimateAffine3D with confidence threshold 0.999 and RANSAC threshold 0.005. We randomly select 2000 4-point pairs to solve for alignment, using the transformation with the largest inlier ratio [Wen et al. 2022].

## C.4 Runtime Analysis

Our method takes 12.09s for an image on a single RTX 3080 GPU, with 0.64s for scene scale sampling, 7.63s for NOC sampling (58ms for RANSAC alignment solver), and 3.82s for object latent sampling (1.52 ms for nearest neighbor retrieval).

## D USED OPEN-SOURCE LIBRARIES

Our data pre-processing, model training and inference paradigm leverage several open-source libraries. Blenderproc [Denninger et al. 2023] is employed for rendering synthetic data. The implementation of our model uses PyTorch. For the multi-hypothesis ranking scheme during inference, which includes rendering the retrieved CAD models using predicted poses, we utilize a PyTorch3D renderer [Ravi et al. 2020].