

Pragmatist: Multiview Conditional Diffusion Models for High-Fidelity 3D Reconstruction from Unposed Sparse Views

Songchun Zhang, Chunhui Zhao*

State Key Laboratory of Industrial Control Technology,
College of Control Science and Engineering, Zhejiang University, China

Abstract

Inferring 3D structures from sparse, unposed observations is challenging due to its unconstrained nature. Recent methods propose to predict implicit representations directly from unposed inputs in a data-driven manner, achieving promising results. However, these methods do not utilize geometric priors and cannot hallucinate the appearance of unseen regions, thus making it challenging to reconstruct fine geometric and textural details. To tackle this challenge, our key idea is to reformulate this ill-posed problem as conditional novel view synthesis, aiming to generate complete observations from limited input views to facilitate reconstruction. With complete observations, the poses of the input views can be easily recovered and further used to optimize the reconstructed object. To this end, we propose a novel pipeline *Pragmatist*. First, we generate a complete observation of the object via a multiview conditional diffusion model. Then, we use a feed-forward large reconstruction model to obtain the reconstructed mesh. To further improve the reconstruction quality, we recover the poses of input views by inverting the obtained 3D representations and further optimize the texture using detailed input views. Unlike previous approaches, our pipeline improves reconstruction by efficiently leveraging unposed inputs and generative priors, circumventing the direct resolution of highly ill-posed problems. Extensive experiments show that our approach achieves promising performance in several benchmarks.

Introduction

Dense multi-view 3D reconstruction (Schonberger and Frahm 2016; Wang et al. 2021; Barnes et al. 2009) has always been one of the core challenges in computer vision and graphics. Current 3D reconstruction techniques have achieved impressive results in various applications such as film production, e-commerce, and virtual and augmented reality. This paper focuses on the specific setting of reconstructing 3D objects from unposed sparse view observations, which is closer to real-world scenarios with only partial observations (e.g., casual capture and e-commerce scenarios).

This problem is particularly challenging from many perspectives, including lacking the explicit geometric information provided by the camera poses, and the need to infer




Figure 1: Our method reconstructs 3D geometry and texture from sparse unposed images, even across different scenes of the same object. More results are available in the supplementary material.

the complete 3D structure from only limited partial observations. To overcome these challenges, some methods have achieved promising results, either by pre-training models on large-scale datasets to introduce prior knowledge of 3D scenes (Tang et al. 2024a; Chen et al. 2024a; Yu et al. 2021; Li et al. 2023; Zhang and Zhao 2023), or by optimizing scene by scene through manually designed regularization (Zou et al. 2024; Zhou and Tulsiani 2023; Wang et al. 2023a). However, such methods rely heavily on the known accurate camera poses of the input image, which is usually only available in synthetic datasets or through privileged information in additional views. To tackle this challenge, some methods (Jiang et al. 2024a, 2023; Kani et al. 2023; Wang et al. 2023b) adopt a fully feed-forward pose-free framework by learning geometric priors from large-scale 3D objects through spatial-aware attention. However, these methods have difficulty recovering complex geometries and textures due to volumetric resolution constraints. In addition, due to the lack of geometric prior, it is challenging for these methods to converge on datasets containing complex objects. Recent methods (Han et al. 2024; Wu et al. 2023a) first estimate the sparse camera pose from images using pre-trained models, and then reconstruct the 3D model via test-time optimization, which can achieve superior quality compared to the feed-forward methods. However, these methods do not fully utilize the input view information, and the pose estimation error may lead to suboptimal 3D reconstruction

*Corresponding author.

Method	Unposed Inputs	3D Consist.	Genl.	Gen Occ.	Lev. 2D Inputs
SparseFusion (2023)	✗	✓	✓	✓	✓
PixelNeRF (2021)	✗	✓	✓	✗	✗
LEAP (2023)	✓	✓	✓	✗	✗
PF-LRM (2023b)	✓	✓	✓	✗	✗
SAP3D (2024)	✓	✓	✗	✓	✓
UpFusion (2023)	✓	✓	✓	✓	✗
Ours	✓	✓	✓	✓	✓

Table 1: **Comparison with prior methods.** *Unposed Inputs* indicates whether the method can handle unposed input images. *3D Consist.* refers to the method’s ability to maintain 3D consistency. *Genl.* denotes the generalization capability of the method. *Gen Occ.* indicates whether the method can hallucinate unseen regions. *Lev. 2D Inputs* denotes the method’s ability to leverage 2D input images.

results.

In this paper, we propose a novel framework, named *Pragmatist*, to solve the challenge of high-fidelity 3D reconstruction from unposed sparse views, as shown in Figure 1 and Table 1. Our innovation is to transform the ill-posed 3D reconstruction problem into a conditional novel view generation task. This is achieved by generating additional consistent observations via a multi-view conditional diffusion model, which effectively transforms the sparse view 3D reconstruction problem into a fully observed 3D reconstruction setting. Specifically, we first propose a pose-free multi-image conditional multi-view generation model, in which the self-attention in 2D latent features can learn 3D consistency across the conditional views and generated views, without explicitly using the camera pose information of the conditional views. In contrast to previous methods (Kong et al. 2024; Liu et al. 2023a; Gao et al. 2024), our approach does not require the camera pose of the input view or coupling the output view pose with the input view during the generation process, making it more flexible and applicable to a wider range of applications.

However, reconstructing accurate texture and geometric details directly from the generated multi-view images is quite challenging. To address this, we propose to combine the advantages of generative priors and geometric constraints. First, we use a feed-forward mesh generation model to reconstruct the object model from the generated target views. Then, we estimate the position and pose of the input views within the object coordinate system via 2D-3D matching. Finally, we further refine the texture using high-resolution input views, restoring detailed geometric structures and material textures. Comprehensive evaluations show that *Pragmatist* improves the reconstruction quality of sparse unposed inputs and can effectively reconstruct complex geometries and realistic textures of 3D objects. In a nutshell, our key contributions can be summed up as follows:

- We propose a novel paradigm for unposed sparse views 3D reconstruction by reformulating the problem as conditional novel view synthesis, aiming to generate complete observations of the given object for reconstruction.

- We propose a flexible multi-view conditional generative model that produces a comprehensive set of observations in canonical coordinate, based on unposed sparse input views.
- To improve the fidelity of the reconstruction, we propose to first recover the pose of the input view by inverting the triplane representation. Then, the fine-grained information in the input view is used to optimize the surface texture of the object.

Related Work

Sparse-View Pose Estimation. Sparse-view pose estimation is a challenging task that involves determining camera poses from limited observations. Some learning-based methods (Zhang, Ramanan, and Tulsiani 2022; Lin et al. 2024) use energy-based models to compose relative rotations into sets of camera poses. Other methods (Sinha et al. 2023; Cheng, Cao, and Shan 2023) learn to refine sparse camera poses iteratively given an initial estimate, while some others (Jiang et al. 2024b) exploit synthetic data to learn camera poses. However, these methods do not generalize well to out-of-domain objects and camera viewpoints. Additionally, some approaches (Wang, Rupprecht, and Novotny 2023; Zhang et al. 2024) use a diffusion model to denoise camera poses. Recent works predict sparse poses by predicting aligned point clouds (Wang et al. 2024a, 2023b) and using PnP to recover cameras. However, these methods require large-scale synthetic data for training, and the estimated pose error can lead to suboptimal results in the subsequent reconstruction process.

Sparse-View Reconstruction. Existing reconstruction or NVS methods perform well with dense inputs but struggle to reconstruct satisfactory results when trained on sparse observations due to insufficient information. To address this challenge, some approaches have proposed to pretrain on large-scale datasets (Yu et al. 2021; Tang et al. 2024a; Chen et al. 2024a; Chan et al. 2023) to introduce prior knowledge. In contrast, other approaches focus on optimizing each scenario via regularization methods (Jain, Tancik, and Abbeel 2021; Zou et al. 2024; Zhou and Tulsiani 2023; Wang et al. 2023a; Yang, Pavone, and Wang 2023). Sparse3D (Zou et al. 2024) and SparseFusion (Zhou and Tulsiani 2023) tend to distill diffusion for sparse view reconstruction. However, these methods rely on the pose ground truth of the sparse input view, which is not practical in the real world. To address this problem, some methods have combined sparse view pose estimation (Sinha et al. 2023; Zhang et al. 2024; Lin et al. 2024) with advanced reconstruction methods, or considered methods that jointly optimize poses together with reconstruction (Jiang et al. 2024b; Han et al. 2024; Wu et al. 2023a). However, the computation of explicit poses is not always reliable, leading to poor reconstruction performance. Some methods (Jiang et al. 2023; Kani et al. 2023) use a fully pose-free model structure, but have difficulty reconstructing complex geometries and textures.

Multi-View Diffusion Models. Previous 3D generative reconstruction methods based on 2D diffusion models have poor 3D generative quality due to the lack of 3D prior

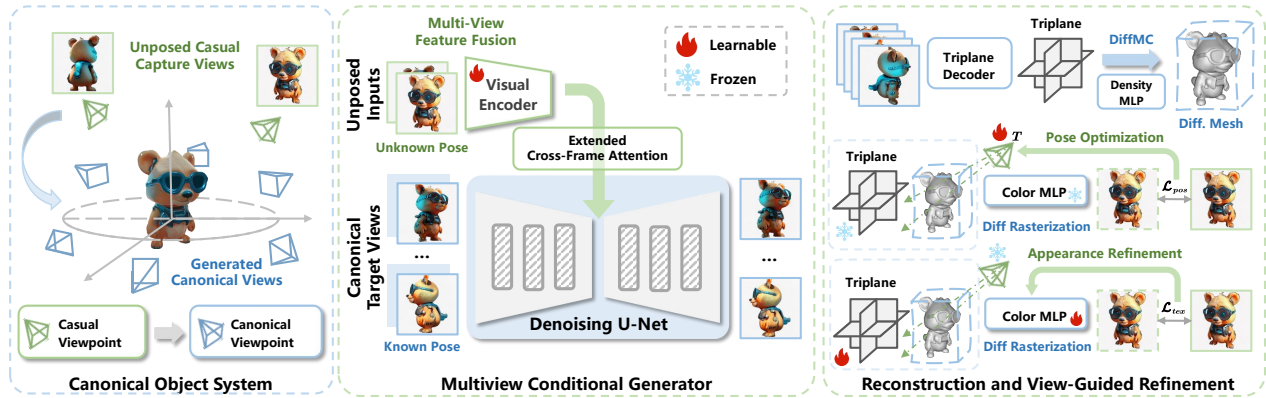


Figure 2: **Overview of our pipeline.** Our core insight is to generate additional novel views in the canonical object coordinate by conditioning with casual viewpoints to solve sparse unposed view reconstruction. The multi-view conditional generator uses a lightweight encoder to extract input view features and implicitly models the 3D consistency between the input view and the output view in the canonical coordinate system via cross-frame attention. To achieve the trade-off between the input view and the generated prior, the object mesh and triplane features are first obtained using the generated views. The input views are then aligned to the canonical coordinates via diff. rendering, and finally, the surface texture is optimized through the input views.

knowledge. Zero123 (Liu et al. 2023a) infused 3D prior knowledge into the 2D diffusion model by pre-training on large-scale synthetic 3D data, but multi-view output still cannot guarantee consistency. To solve this problem, many subsequent works (Liu et al. 2023b; Long et al. 2024; Shi et al. 2023; Voleti et al. 2024; Chen et al. 2024b; Kant et al. 2024; Tang et al. 2024b, 2023; Zheng and Vedaldi 2024) focused on improving the 3D consistency of multi-view images. These methods enhance the quality of generated multi-view images by modeling 3D generation through a joint probability distribution. However, when single-view inputs are of low quality, they can result in implausible shapes and textures. Therefore, some methods (Kong et al. 2024; Tang et al. 2024b; Gao et al. 2024; Wu et al. 2024) incorporate flexible multi-view image conditioning into the diffusion model to achieve more controllable 3D object reconstruction. However, most of them require the ground truth camera pose (which is challenging to acquire in practice) and have specific requirements for the range of viewpoints and image quality of the input images.

Method

Given a set of sparse observations of an object, our task is to recover the shape and texture of the current object. Our method incorporates generative priors with geometric knowledge to alleviate the ambiguity of the ill-posed problem by generating additional novel views in canonical coordinates. The overview of our pipeline is illustrated in Figure 2. In this section, we first describe the construction of a conditional multiview generator that converts arbitrary unposed sparse views into multiview images with 3D consistency. Next, we will introduce the design of the feed-forward reconstruction model, which uses the generated images to reconstruct high-quality meshes. Finally, we will explain how to efficiently recover the camera pose of the input observation through the reconstruction results and further im-

prove the quality of the 3D model.

Multiview Conditional Multiview Generator

Existing multi-view diffusion methods usually assume that the intrinsics of the input and generated images are the same. However, in practice, the input images can come from arbitrary cameras, and when the intrinsics are quite different from the training dataset, the generated multi-view images and reconstructed 3D meshes will be distorted. In addition, these methods usually support only a single image in the near-orthogonal view or multi-view with ground truth camera poses as the condition, which greatly limits the application of the model in real-world scenarios. In contrast, our model can receive arbitrary numbers of unposed inputs and generate novel views in the canonical coordinate system for subsequent reconstruction steps. Specifically, given a sparse set of N unposed images $\mathcal{I}_{cond} = \{\mathbf{I}_i^{cond}\}_{i=1}^N$ and their corresponding camera poses \mathcal{P}_{cond} , our multiview diffusion model can generate multiple novel views $\mathcal{I}_{tgt} = \{\mathbf{I}_j^{tgt}\}_{j=1}^M$ based on given target camera poses $\mathcal{P}_{tgt} = \{\mathbf{p}_j\}_{j=1}^M$. This process can be formulated as:

$$\mathcal{I}_{tgt} \sim p(\mathcal{I}_{tgt} | \mathcal{I}_{cond}, \mathcal{P}_{cond}, \mathcal{P}_{tgt}), \quad (1)$$

where \mathcal{I}_{cond} and \mathcal{P}_{cond} are the observed images and their corresponding poses, while \mathcal{P}_{tgt} specifies the camera poses for the novel views to be generated.

Conditioning Feature Extraction. Previous methods use the frozen CLIP ViT and denoising Unet to capture the high-level and low-level information of inputs, respectively, but are only suitable for single-view input. Instead, we consider directly integrating the high- and low-level signals of multiple input views using a conditional encoder and representing them as a set of tokens for subsequent injection into a diffusion model. Specifically, we chose to fine-tune a lightweight encoder \mathcal{E}_{cond} (Woo et al. 2023) to integrate image features and use these features to construct conditional tokens.

Multiview Conditioned Generation. We extend the generative network based on the origin latent diffusion architecture. Specifically, the self-attention layer in the original stable diffusion is used to learn relationships within a single image, which we extend to learn relationships between multiple target views to ensure inter-view consistency. In addition, we also extend the cross-attention layer to learn the interactions between input and target views, thereby incorporating the information of unposed input views into the target view generation process. To enhance camera control ability, previous methods (Li et al. 2023; Hong et al. 2023) encode camera poses at the frame level, sharing the same position encoding between pixels. Aiming at more accurate pixel-level encoding, we integrate the camera conditions into the denoising network by parameterizing the rays $\mathbf{r} = (\mathbf{d}, \mathbf{o} \times \mathbf{d})$ and employing a two-layer MLP that injects a Plücker ray embedding for each pixel.

Canonical Camera System. Generating novel view images in canonical camera settings is challenging due to the sparse input view and arbitrary viewpoint, which require the model to infer the elevation and the object’s scale. To address these issues, we adopt a fixed absolute elevation angle and distance, using the relative azimuth angle as the pose for the new view. This approach eliminates uncertainties in direction and scale, removing the need for additional elevation angle estimation. Cameras are positioned at a distance of 1.5 with azimuth angles $(\beta, \beta \pm 45^\circ, \beta \pm 90^\circ, \beta \pm 135^\circ, \beta \pm 180^\circ)$ at 0° elevation.

Training Details. During training, the number of condition and target images are fixed, but the inference can support any number of images. In addition, to simulate the real-world conditional images, we perform noise perturbation and re-size operations on the condition images during training to make our model more robust to the input views.

Feed-forward Reconstruction

In our model, the feed-forward reconstruction module aims to efficiently reconstruct high-quality meshes from images generated by the multi-view conditional generator. However, direct optimization based on the mesh representation is challenging, as it is sensitive to geometry initialization and can get stuck in local minima, especially when reconstructing high-resolution meshes. To solve this problem, we first optimize the NeRF representation by volume rendering, which is used to initialize the network weights for mesh reconstruction. Then, we further optimize the surface rendering loss by differentiable rasterization, which further improves the mesh surface extraction results.

Volume Rendering Training. We use a transformer-based model to predict the triplane NeRF in a feed-forward manner from sparse view inputs. The model consists of an image encoder, an image-to-triplane decoder, and a NeRF decoder. The image encoder converts the multi-view images $\{I_i\}_{i=1}^M$ into latent codes. Then, we use concatenated image codes in the image-to-triplane decoder to predict the triplane features T of the 3D object. Unlike previous methods (Hong et al. 2023; Li et al. 2023), we use two separate MLPs to estimate color and density, because the density MLP f_d and color MLP f_c are used separately during mesh extraction

and surface rendering, and can increase the stability of the training process. The loss for volume rendering training is formulated as:

$$\mathcal{L}_{vol} = \mathcal{L}_{rgb} + \lambda_p \mathcal{L}_{LPIPS}, \quad (2)$$

where \mathcal{L}_{rgb} indicates L2 regression loss of RGB image, \mathcal{L}_{LPIPS} denotes the perceptual loss (Chen and Koltun 2017), and λ indicates the loss weight.

Surface Rendering Training. Although the trained model shows promising results in novel view synthesis, directly extracting meshes from the density field produces artifacts. Therefore, we propose to fine-tune the network using differentiable marching cubes (DiffMC) (Wei et al. 2023) and differentiable rasterization (Laine et al. 2020). We first compute the density grid using triplane features and extract grid surfaces from it using the DiffMC technique. Then, we use the differentiable rasterizer to render novel views from the extracted grid using triplane features, calculating the rendering loss to optimize the model for end-to-end mesh reconstruction. Specifically, we first obtain the 3D position of each pixel via the differentiable rasterizer, then query the corresponding triplane feature, and use f_c to obtain the color of each pixel. The pipeline is fine-tuned using MSE and LPIPS (Chen and Koltun 2017) losses on images, MSE losses on depths and masks, and a regularization loss on opacity (Wei et al. 2024). The loss for surface rendering training is formulated as:

$$\mathcal{L}_{geo} = \mathcal{L}_{rgb} + \lambda_p \mathcal{L}_{LPIPS} + \lambda_d \mathcal{L}_d + \lambda_m \mathcal{L}_m + \lambda_r \mathcal{L}_o, \quad (3)$$

where λ_p , λ_d , λ_m and λ_r denote the weights of each loss term, respectively. Since mesh rasterization is very efficient, we render the image at full resolution for supervision.

Directing Virtual Camera for Refinement

Though the 3D reconstruction quality of the feed-forward pipeline is promising, it still has trouble handling complex textures. This issue gets even worse when using unposed images as input. We argue that the rich high-resolution texture information in the input unposed view should be fully utilized. Inspired by the previous incremental SFM pipeline (Schonberger and Frahm 2016) we can treat these unposed views as new inputs and register them into the reconstruction through visual relocalization, thereby optimizing the reconstruction.

Align Camera to Canonical Coordinate. To achieve this, using the mesh from feed-forward reconstruction and the camera pose, we can obtain the corresponding image by mesh rendering. Subsequently, we use the photometric loss for backpropagation to optimize the camera pose \mathbf{p} of I^{cond} . In practice, objects are usually not centered on the screen. To align the input view with the canonical coordinate system, we use foreground segmentation to obtain the object mask and then generate an object-centered image I_v^{cond} via homography transformation. We then use the spherical coordinate system $\mathbf{p} = \{\alpha, \beta, \gamma, r\}$ to describe the input camera poses. This simplified representation reduces the search space in pose optimization. The process can be formulated as the following equation:

$$\hat{\mathbf{p}} = \arg \min_{\mathbf{p}} \mathcal{L}_{pos}(\mathcal{R}(T, \mathbf{p}, f_c, f_d), I_v^{cond}) \quad (4)$$

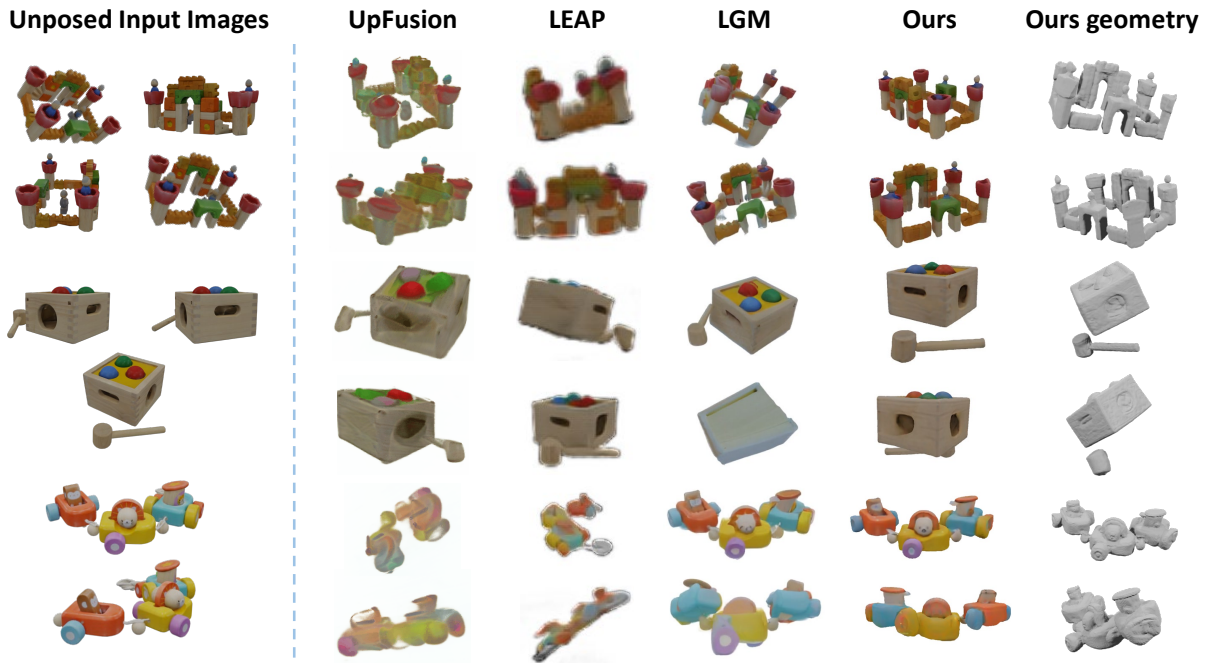


Figure 3: Qualitative comparison of unposed sparse-view 3D reconstruction results on the GSO dataset. Our method can produce more detailed and higher-quality reconstruction results than the benchmark method and can support different numbers of unposed inputs. More results can be found in the supplementary material.

During pose optimization, the parameters of the T , f_c , and f_d are frozen. To avoid local minima due to poor initial values, we start with different initial poses and choose the one with the minimum loss after convergence as the result.

View-Guided Appearance Refinement. To improve the fidelity of 3D reconstruction models to input views, we propose to use input images that have been registered to canonical coordinates to further optimize the texture of the model. Specifically, we use the predicted poses of the input views to obtain the features of the 3D surface points via the differentiable rasterizer and the obtained triplane feature T , and then use f_c to render current views. Then, we calculate the loss between the render views and the input views to optimize the color MLP and the triplane feature T , thereby improving the appearance of the model. The process can be formulated as follows

$$\min_{f_c, T} \mathcal{L}_{tex}(\mathcal{R}(T, \hat{p}, f_c, f_d), \mathbf{I}_v^{cond}). \quad (5)$$

We found that separating the density and color models helped to refine the appearance without affecting the performance of the geometry model, as we only needed to fine-tune the color MLP f_c .

Experiments

Experimental Settings

Implementation details. Our model is optimized using the AdamW (Loshchilov and Hutter 2017) with a weight decay of 0.01. The multiview generator is trained with an initial learning rate of $3e-4$, which decreases to $1e-5$ over 30,000

training steps using a cosine learning rate decay schedule. There are two stages to the training of the feedforward reconstruction model. The first stage trains at 256×256 resolution using the AdamW optimizer ($\beta_2 = 0.95$) with a weight decay of 0.05, a cosine learning rate decay of 4×10^{-4} , and 2500 warm-up iterations, for a total of $60k$ training iterations. We use a batch size of 192 and sample 128 points per ray. The refinement stage operates at 512×512 resolution with 256 ray samples and a batch size of 64. The learning rate is set to 1×10^{-5} with 1000 warm-up iterations, using DiffMC for mesh optimization over $10k$ iterations. Loss weights are $\lambda_p = 0.5$, $\lambda_d = 0.5$, $\lambda_m = 1$, and $\lambda_n = 1$. For stability, normal loss is added after convergence with the generator frozen. We randomly sample 4 input views with 4 supervision views during the NeRF training phase, expanding to 8 supervision views in the geometry refinement phase. The backbone uses a pre-LN transformer with 24 layers (1024 channels, 16 attention heads, 64 dimensions per head). GeLU-activated MLPs with 4096 intermediate features split into density (3 layers) and color (4 layers) pathways, each using 512 hidden dimensions.

Datasets. Our model is trained on a filtered subset of Objaverse with about 240k quality assets. Following the previous methods, we evaluate our model’s performance on the Google Scanned Object (GSO) (Downs et al. 2022) and OmniObject3D (Wu et al. 2023b) datasets. To further verify the generalizability of the method, we also evaluated its performance on casually captured images and images created using a text-to-image diffusion model.

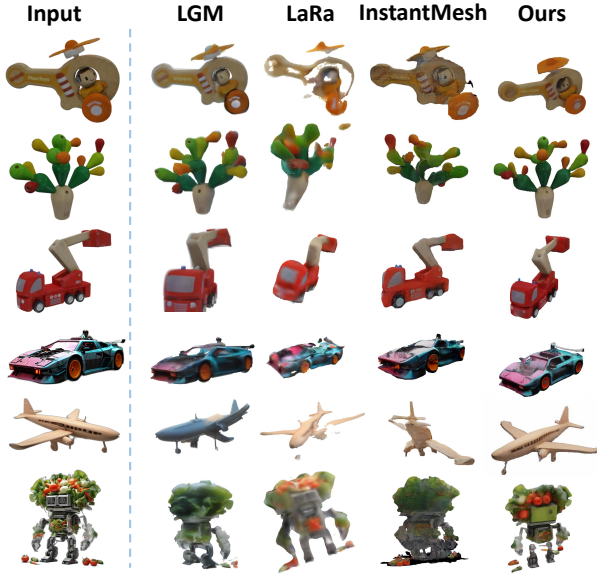


Figure 4: Qualitative comparison of single view 3D reconstruction results on the GSO and in-the-wild datasets. Though our method is not designed for single-view setting, it produces results comparable to these benchmarks and generates reasonable textures in the invisible areas.

Metrics. We evaluate our method on novel view synthesis and reconstruction tasks. Specifically, we use PSNR, SSIM, and LPIPS as novel view synthesis metrics, while Chamfer Distance (CD) and Volume IOU are used for 3D reconstruction evaluation. For 3D geometry evaluation, we align the generated meshes with the ground-truth meshes at the origin, rescale them to a unit cube, and apply ICP registration for alignment.

Unposed Sparse View Reconstruction

Baselines. We first compare our method with fully pose-free reconstruction methods, i.e., LEAP (Jiang et al. 2023), UpFusion (Kani et al. 2023), and FORGE (Jiang et al. 2024b), as they can directly receive sparse unposed inputs. In addition, many recent sparse reconstruction methods have achieved promising results, but require ground truth pose as input, such as LGM (Tang et al. 2024a), SAP3D (Han et al. 2024), SparseNeuS (Long et al. 2022), and LaRa (Chen et al. 2024a). To compare with them, we first estimate the camera poses using the pre-trained sparse pose estimators (Lin et al. 2024; Zhang et al. 2024), and then use these methods to reconstruct to compare with our results.

Comparison Results. Compared to pose-free methods, the metrics of our method are significantly improved, as shown in Figure 3 and Table 2, due to 1) avoiding a direct data-driven solution to the ill-posed problem, and 2) reconstructing in a canonical coordinate system, which alleviates the problem of scale ambiguity. In addition, our method performs better than the method that predicts the camera pose in image space and then performs multiview reconstruction. The pose estimation accuracy of such methods is prone

Method	Pose	GSO			OmniObject3D		
		PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Yu <i>et al.</i>	Pred.	11.483	0.852	0.216	10.183	0.847	0.226
	GT	14.278	0.932	0.153	11.365	0.852	0.219
FORGE	Pred.	11.432	0.754	0.762	10.164	0.849	0.228
	GT	15.164	0.867	0.191	12.456	0.856	0.211
LGM	Pred.	20.149	0.888	0.156	17.698	0.878	0.173
	GT	21.783	0.873	0.214	18.546	0.883	0.172
LaRa	Pred.	20.696	0.889	0.152	17.877	0.878	0.175
	GT	23.374	0.865	0.208	20.186	0.888	0.157
SAP3D	Pred.	17.312	0.863	0.173	15.534	0.871	0.187
LEAP	✗	21.348	0.856	0.128	19.439	0.885	0.162
UpFus.	✗	20.924	0.901	0.147	18.281	0.881	0.169
Ours	✗	22.931	0.887	0.121	20.897	0.892	0.153
Ours	Pred.	23.698	0.902	0.131	21.764	0.895	0.146

Table 2: Comparison with sparse-view reconstruction methods on novel view synthesis. We conducted experiments on two datasets, i.e., GSO (Downs et al. 2022) and OmniObject3D (Wu et al. 2023b). We denote reconstructions using predicted poses as *Pred.*, those using ground truth poses as *GT*, and reconstructions without poses as ✗.

Method	R. error↓	Acc.@15°↑	Acc.@30°↑	T. error↓
FORGE	98.471	0.024	0.036	1.122
HLoc	95.243	0.028	0.146	1.189
RelPose++	67.580	0.269	0.415	0.893
Ours	2.182	0.961	0.978	0.039

Table 3: Comparison of the accuracy of pose estimation on the GSO dataset.

to degradation when encountering uncommon input viewpoints, which leads to cumulative errors and affects the reconstruction quality. More results can be found in the supplementary material.

Pose Estimation Results. As shown in Table 3, the accuracy of our method is significantly better than that of the method that directly predicts the camera pose from the image space, and some traditional relative pose estimation methods. Our performance stems from predicting the camera pose from the obtained 3D shape rather than directly from the 2D image. Due to poor visual correspondence, traditional SfM methods tend to fail when dealing with sparse views, so we use the more robust HLoc as a benchmark for comparison.

Single View Reconstruction

Baselines. Existing single-image 3D reconstruction methods have achieved promising results, and can be divided into two main branches, namely those that rely on generative priors (Liu et al. 2023a; Long et al. 2024; Liu et al. 2023b; Wang et al. 2024b; Tang et al. 2024a) and those that do not (Hong et al. 2023; Tochilkin et al. 2024). We compare our method with these two types of methods. To ensure fairness, we use only a single image as input for our method.

Comparison Results. As shown in Table 5 and Figure 4,

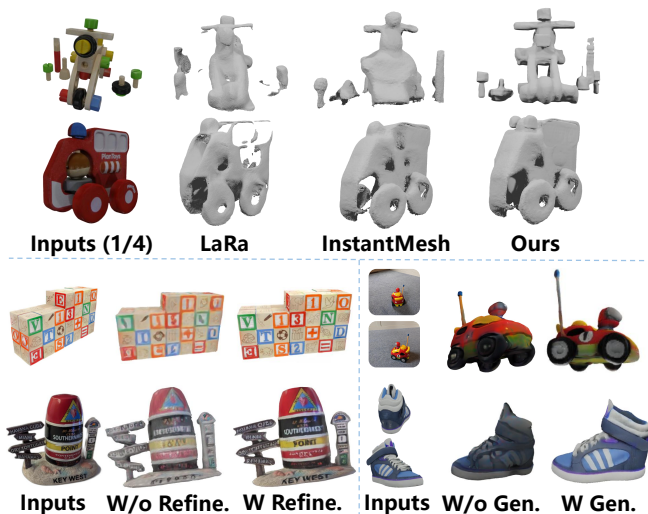


Figure 5: **Ablation Studies.** The top image illustrates the effectiveness of our reconstruction model, the bottom left image shows the effect of our refinement step, and the bottom right image demonstrates the effect of generative priors.

Method	Pose	GSO		OmniObject3D	
		CD↓	IoU↑	CD↓	IoU↑
LGM	Pred.	0.0459	0.2903	0.0532	0.4423
	GT	0.0251	0.5632	0.0305	0.5843
LaRa	Pred.	0.0448	0.3764	0.0511	0.4312
	GT	0.0246	0.5624	0.0261	0.5425
SAP3D	Pred.	0.0364	0.5346	0.0423	0.2876
Ours	\times Pred.	0.0187	0.6964	0.0202	0.6693

Table 4: Comparison with sparse-view reconstruction methods on the reconstructed mesh. We conducted experiments on two datasets, i.e., GSO (Downs et al. 2022) and OmniObject3D (Wu et al. 2023b).

our method can achieve comparable performance to the current SOTA single image-to-3D method. These methods suffer from significant performance degradation in in-the-wild scenes, mainly due to the uncertainty of the input camera view. Previous methods often rely on existing generative models that require a range of input viewpoints, resulting in lower-quality multi-view generation and suboptimal reconstruction when encountering unusual viewpoints. In contrast, our method is robust to the input image view and can use any multi-view inputs, increasing controllability and improving generation and reconstruction results. More results can be found in the supplementary material.

Ablation Studies

Ablation of Multiview Generator. We first evaluate the multi-view generator in multi-view conditions. As shown in Figure 5, when this module is removed, our method degenerates to pose-free reconstruction. Even when combined with

Method	CD↓	IoU↑	PSNR↑	SSIM↑	LPIPS↓
One-2-3-45*	0.0629	0.4086	-	-	-
SyncDreamer*	0.0261	0.5421	19.512	0.847	0.188
Wonder3D*	0.0329	0.5768	-	-	-
LGM*	0.0242	0.5618	21.418	0.861	0.214
CRM*	0.0212	0.6146	21.984	0.876	0.193
TripoSR	0.0209	0.6369	22.376	0.885	0.131
InstantMesh*	0.0187	0.6417	22.879	0.901	0.116
Era3D*	0.0175	0.6654	-	-	-
Ours* [1 View]	0.0189	0.6404	20.455	0.857	0.156
Ours* [2 Views]	0.0165	0.6949	22.763	0.898	0.123
Ours* [4 Views]	0.0146	0.7348	24.074	0.904	0.112

Table 5: Comparison with single-view reconstruction methods on the reconstructed mesh. We conducted experiments on the GSO dataset. ‘*’ denotes models that employ generative priors for 3D reconstruction.

pose estimation for end-to-end reconstruction, the results are still unsatisfactory, and the texture in the unseen region is blurred due to the lack of generative priors. More results can be found in the supplementary material.

Ablation of Feed-Forward Reconstruction. We evaluated the effect of our multi-view reconstruction model. We compared our method with several other existing end-to-end reconstruction methods. As shown in Figure 5, our end-to-end reconstruction method can achieve reconstruction quality comparable to the optimization method with high efficiency. In addition, we can use differentiable marching cube to quickly extract higher-quality meshes.

Ablation of Virtual Camera Guided Refinement. We evaluated the effectiveness of the refinement step. As shown in Figure 5, after removing the refinement step, the reconstruction of complex textured objects becomes worse. During the refinement process, the use of an independent MLP can improve the reconstruction performance while maintaining geometric quality. In addition, the experimental results show that optimizing the triplane feature and color MLP at the same time achieves better results than optimizing the color MLP alone during the appearance refinement process.

Conclusions

In this paper, we introduced *Pragmatist*, a novel pipeline for inferring 3D structures from sparse, unposed observations by reformulating the problem as a task of conditional novel view synthesis. Our approach uses a multi-view conditional diffusion model to generate complete object observations, which are then used by a large reconstruction model to generate high-quality meshes. Furthermore, we enhance reconstruction accuracy by recovering input view poses and optimizing texture details using available input views. Extensive experiments demonstrate that *Pragmatist* significantly outperforms existing methods in several benchmarks. By utilizing unposed inputs and generative priors, our pipeline can efficiently solve this ill-posed problem and ensure generalizability across diverse scenes.

Acknowledgments

This research was supported by the National Natural Science Foundation of China (No. 62450020) and the National Natural Science Foundation of China (No. 62125306).

References

- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3): 24.
- Chan, E. R.; Nagano, K.; Chan, M. A.; Bergman, A. W.; Park, J. J.; Levy, A.; Aittala, M.; De Mello, S.; Karras, T.; and Wetzstein, G. 2023. Generative novel view synthesis with 3d-aware diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4217–4229.
- Chen, A.; Xu, H.; Esposito, S.; Tang, S.; and Geiger, A. 2024a. LaRa: Efficient Large-Baseline Radiance Fields. *arXiv preprint arXiv:2407.04699*.
- Chen, Q.; and Koltun, V. 2017. Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE international conference on computer vision*, 1511–1520.
- Chen, Z.; Wang, Y.; Wang, F.; Wang, Z.; and Liu, H. 2024b. V3d: Video diffusion models are effective 3d generators. *arXiv preprint arXiv:2403.06738*.
- Cheng, W.; Cao, Y.-P.; and Shan, Y. 2023. ID-Pose: Sparse-view Camera Pose Estimation by Inverting Diffusion Models. *arXiv preprint arXiv:2306.17140*.
- Downs, L.; Francis, A.; Koenig, N.; Kinman, B.; Hickman, R.; Reymann, K.; McHugh, T. B.; and Vanhoucke, V. 2022. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, 2553–2560. IEEE.
- Gao, R.; Holynski, A.; Henzler, P.; Brussee, A.; Martin-Brualla, R.; Srinivasan, P.; Barron, J. T.; and Poole, B. 2024. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv preprint arXiv:2405.10314*.
- Han, X.; Gao, Z.; Kanazawa, A.; Goel, S.; and Gandselman, Y. 2024. The More You See in 2D the More You Perceive in 3D. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20912–20922.
- Hong, Y.; Zhang, K.; Gu, J.; Bi, S.; Zhou, Y.; Liu, D.; Liu, F.; Sunkavalli, K.; Bui, T.; and Tan, H. 2023. Lrm: Large reconstruction model for single image to 3d. In *International Conference on Learning Representations*.
- Jain, A.; Tancik, M.; and Abbeel, P. 2021. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5885–5894.
- Jiang, H.; Jiang, Z.; Grauman, K.; and Zhu, Y. 2024a. Few-view object reconstruction with unknown categories and camera poses. *3DV*.
- Jiang, H.; Jiang, Z.; Grauman, K.; and Zhu, Y. 2024b. Few-view object reconstruction with unknown categories and camera poses. In *2024 International Conference on 3D Vision (3DV)*, 31–41. IEEE.
- Jiang, H.; Jiang, Z.; Zhao, Y.; and Huang, Q. 2023. Leap: Liberate sparse-view 3d modeling from camera poses. *arXiv preprint arXiv:2310.01410*.
- Kani, B. R. N.; Lee, H.-Y.; Tulyakov, S.; and Tulsiani, S. 2023. UpFusion: Novel View Diffusion from Unposed Sparse View Observations. *arXiv preprint arXiv:2312.06661*.
- Kant, Y.; Siarohin, A.; Wu, Z.; Vasilkovsky, M.; Qian, G.; Ren, J.; Guler, R. A.; Ghanem, B.; Tulyakov, S.; and Gilitschenski, I. 2024. SPAD: Spatially Aware Multi-View Diffusers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10026–10038.
- Kong, X.; Liu, S.; Lyu, X.; Taher, M.; Qi, X.; and Davison, A. J. 2024. EscherNet: A Generative Model for Scalable View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9503–9513.
- Laine, S.; Hellsten, J.; Karras, T.; Seol, Y.; Lehtinen, J.; and Aila, T. 2020. Modular Primitives for High-Performance Differentiable Rendering. *ACM Transactions on Graphics*, 39(6).
- Li, J.; Tan, H.; Zhang, K.; Xu, Z.; Luan, F.; Xu, Y.; Hong, Y.; Sunkavalli, K.; Shakhnarovich, G.; and Bi, S. 2023. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. In *International Conference on Learning Representations*.
- Lin, A.; Zhang, J. Y.; Ramanan, D.; and Tulsiani, S. 2024. RelPose++: Recovering 6D Poses from Sparse-view Observations. In *3DV*.
- Liu, R.; Wu, R.; Van Hoorick, B.; Tokmakov, P.; Zakharov, S.; and Vondrick, C. 2023a. Zero-1-to-3: Zero-shot One Image to 3D Object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9298–9309.
- Liu, Y.; Lin, C.; Zeng, Z.; Long, X.; Liu, L.; Komura, T.; and Wang, W. 2023b. SyncDreamer: Learning to Generate Multiview-consistent Images from a Single-view Image. *arXiv preprint arXiv:2309.03453*.
- Long, X.; Guo, Y.-C.; Lin, C.; Liu, Y.; Dou, Z.; Liu, L.; Ma, Y.; Zhang, S.-H.; Habermann, M.; Theobalt, C.; et al. 2024. Wonder3d: Single image to 3d using cross-domain diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9970–9980.
- Long, X.; Lin, C.; Wang, P.; Komura, T.; and Wang, W. 2022. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, 210–227. Springer.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.

- Shi, Y.; Wang, P.; Ye, J.; Mai, L.; Li, K.; and Yang, X. 2023. MVDream: Multi-view Diffusion for 3D Generation. *arXiv:2308.16512*.
- Sinha, S.; Zhang, J. Y.; Tagliasacchi, A.; Gilitschenski, I.; and Lindell, D. B. 2023. Sparsepose: Sparse-view camera pose regression and refinement. In *CVPR*.
- Tang, J.; Chen, Z.; Chen, X.; Wang, T.; Zeng, G.; and Liu, Z. 2024a. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. *arXiv preprint arXiv:2402.05054*.
- Tang, S.; Chen, J.; Wang, D.; Tang, C.; Zhang, F.; Fan, Y.; Chandra, V.; Furukawa, Y.; and Ranjan, R. 2024b. Mvdifffusion+: A dense high-resolution multi-view diffusion model for single or sparse-view 3d object reconstruction. *arXiv preprint arXiv:2402.12712*.
- Tang, S.; Zhang, F.; Chen, J.; Wang, P.; and Yasutaka, F. 2023. MVDiffusion: Enabling Holistic Multi-view Image Generation with Correspondence-Aware Diffusion. *arXiv preprint 2307.01097*.
- Tochilkin, D.; Pankratz, D.; Liu, Z.; Huang, Z.; Letts, A.; Li, Y.; Liang, D.; Laforte, C.; Jampani, V.; and Cao, Y.-P. 2024. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*.
- Voleti, V.; Yao, C.-H.; Boss, M.; Letts, A.; Pankratz, D.; Tochilkin, D.; Laforte, C.; Rombach, R.; and Jampani, V. 2024. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*.
- Wang, G.; Chen, Z.; Loy, C. C.; and Liu, Z. 2023a. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9065–9076.
- Wang, J.; Rupperecht, C.; and Novotny, D. 2023. Posediffusion: Solving pose estimation via diffusion-aided bundle adjustment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9773–9783.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems*, 34: 27171–27183.
- Wang, P.; Tan, H.; Bi, S.; Xu, Y.; Luan, F.; Sunkavalli, K.; Wang, W.; Xu, Z.; and Zhang, K. 2023b. Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. In *International Conference on Learning Representations*.
- Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; and Revaud, J. 2024a. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20697–20709.
- Wang, Z.; Wang, Y.; Chen, Y.; Xiang, C.; Chen, S.; Yu, D.; Li, C.; Su, H.; and Zhu, J. 2024b. Crm: Single image to 3d textured mesh with convolutional reconstruction model. *arXiv preprint arXiv:2403.05034*.
- Wei, X.; Xiang, F.; Bi, S.; Chen, A.; Sunkavalli, K.; Xu, Z.; and Su, H. 2023. Neumanifold: Neural watertight manifold reconstruction with efficient and high-quality rendering support. *arXiv preprint arXiv:2305.17134*.
- Wei, X.; Zhang, K.; Bi, S.; Tan, H.; Luan, F.; Deschaintre, V.; Sunkavalli, K.; Su, H.; and Xu, Z. 2024. Meshlrm: Large reconstruction model for high-quality mesh. *arXiv preprint arXiv:2404.12385*.
- Woo, S.; Debnath, S.; Hu, R.; Chen, X.; Liu, Z.; Kweon, I. S.; and Xie, S. 2023. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16133–16142.
- Wu, C.-H.; Chen, Y.-C.; Solarte, B.; Yuan, L.; and Sun, M. 2023a. iFusion: Inverting Diffusion for Pose-Free Reconstruction from Sparse Views. *arXiv preprint arXiv:2312.17250*.
- Wu, R.; Mildenhall, B.; Henzler, P.; Park, K.; Gao, R.; Watson, D.; Srinivasan, P. P.; Verbin, D.; Barron, J. T.; Poole, B.; et al. 2024. Reconfusion: 3d reconstruction with diffusion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21551–21561.
- Wu, T.; Zhang, J.; Fu, X.; Wang, Y.; Ren, J.; Pan, L.; Wu, W.; Yang, L.; Wang, J.; Qian, C.; et al. 2023b. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 803–814.
- Yang, J.; Pavone, M.; and Wang, Y. 2023. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8254–8263.
- Yu, A.; Ye, V.; Tancik, M.; and Kanazawa, A. 2021. pixelnerf: Neural radiance fields from one or few images. In *CVPR*.
- Zhang, J. Y.; Lin, A.; Kumar, M.; Yang, T.-H.; Ramanan, D.; and Tulsiani, S. 2024. Cameras as rays: Pose estimation via ray diffusion. *arXiv preprint arXiv:2402.14817*.
- Zhang, J. Y.; Ramanan, D.; and Tulsiani, S. 2022. Relpose: Predicting probabilistic relative rotation for single objects in the wild. In *ECCV*.
- Zhang, S.; and Zhao, C. 2023. Cross-Video Contextual Knowledge Exploration and Exploitation for Ambiguity Reduction in Weakly Supervised Temporal Action Localization. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Zheng, C.; and Vedaldi, A. 2024. Free3D: Consistent Novel View Synthesis without 3D Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9720–9731.
- Zhou, Z.; and Tulsiani, S. 2023. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12588–12597.
- Zou, Z.; Cheng, W.; Cao, Y.-P.; Huang, S.-S.; Shan, Y.; and Zhang, S.-H. 2024. Sparse3d: Distilling multiview-consistent diffusion for object reconstruction from sparse views. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7900–7908.