FREESPLATTER: POSE-FREE GAUSSIAN SPLATTING FOR SPARSE-VIEW 3D RECONSTRUCTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Existing sparse-view reconstruction models heavily rely on accurate known camera poses. However, deriving camera extrinsics and intrinsics from sparse-view images presents significant challenges. In this work, we present FreeSplatter, a highly scalable, feed-forward reconstruction framework capable of generating high-quality 3D Gaussians from **uncalibrated** sparse-view images and recovering their camera parameters in mere seconds. FreeSplatter is built upon a streamlined transformer architecture, comprising sequential self-attention blocks that facilitate information exchange among multi-view image tokens and decode them into pixel-wise 3D Gaussian primitives. The predicted Gaussian primitives are situated in a unified reference frame, allowing for high-fidelity 3D modeling and instant camera parameter estimation using off-the-shelf solvers. To cater to both objectcentric and scene-level reconstruction, we train two model variants of FreeSplatter on extensive datasets. In both scenarios, FreeSplatter outperforms state-ofthe-art baselines in terms of reconstruction quality and pose estimation accuracy. Furthermore, we showcase FreeSplatter's potential in enhancing the productivity of downstream applications, such as text/image-to-3D content creation.

- 1 INTRODUCTION
- 029

042

004

010 011

012

013

014

015

016

017

018

019

021

025

026 027 028

031 Recent breakthroughs in neural scene representation and differentiable rendering, e.g., Neural Radi-032 ance Fields (NeRF) (Mildenhall et al., 2021) and Gaussian Splatting (GS) (Kerbl et al., 2023), have 033 shown unprecedented multi-view reconstruction quality for densely-captured images with calibrated 034 camera poses, employing a per-scene optimization approach. However, they are not applicable to sparse-view scenarios, where classical camera calibration techniques like Structure-from-Motion (SfM) (Schonberger & Frahm, 2016) tend to fail due to insufficient image overlaps. Generalizable 036 reconstruction models (Hong et al., 2024b; Xu et al., 2024a; Charatan et al., 2024) attempt to ad-037 dress sparse-view reconstruction in a feed-forward manner by utilizing learned data priors. Despite their efficiency and generalization capabilities, these models predominantly assume access to accurate camera poses and intrinsics, or imply that they are obtained in a pre-processing step, thereby 040 circumventing the difficulty of deriving camera parameters in real application scenarios. Liberating 041 sparse-view reconstruction from known camera poses remains a significant challenge.

Prior works have explored training *pose-free* reconstruction models. PF-LRM (Wang et al., 2024a) 043 and LEAP (Jiang et al., 2024b) share a similar framework that maps multi-view image tokens to a 044 NeRF representation with a transformer. Despite the milestone they have set, their NeRF representa-045 tion suffers from inefficient volume rendering and low resolution, limiting the training efficiency and 046 scalability to complex scenes. Besides, inferring the input camera poses from such implicit represen-047 tations is not trivial. PF-LRM employs an additional branch to predict per-view coarse point clouds 048 for camera pose prediction, which introduces extra training cost and difficulties. DUSt3R (Wang et al., 2024b) presents a novel paradigm for joint 3D reconstruction and pose estimation by modeling SfM and Multi-view Stereo (MVS) as an end-to-end point-regression task. By regressing the 051 "point maps", *i.e.*, the 3D un-projection of depth maps, from a stereo image pair in a unified reference frame, it can recover the relative camera pose efficiently with a PnP (Fischler & Bolles, 1981; 052 Hartley & Zisserman, 2003) solver. DUSt3R demonstrates impressive zero-shot 3D reconstruction and pose prediction capability by training on a broad range of datasets.



Figure 1: Given **uncalibrated** sparse-view images, our FreeSplatter can reconstruct pixel-wise 3D Gaussians, enabling both high-fidelity novel view rendering and instant camera pose estimation in mere seconds. FreeSplatter can deal with both object-centric (up) and scene-level (down) scenarios.

071 However, the sparsity of point cloud makes it a suboptimal representation for many downstream applications, e.g., novel view synthesis. In comparison, 3D Gaussian Splats (3DGS) can encode 072 high-fidelity radiance fields and enable efficient view synthesis by augmenting point clouds with ad-073 ditional Gaussian attributes. We then ask, can we directly predict the "Gaussian maps" from multi-074 view images to achieve both high-quality 3D modeling and instant camera pose estimation? In this 075 work, we presents **FreeSplatter**, a feed-forward reconstruction framework that predicts per-pixel 076 Gaussians from **uncalibrated** sparse-view images and estimates their camera parameters jointly, to 077 say "yes" to this question. The core of FreeSplatter is a scalable single-stream transformer architecture that maps multi-view image tokens into pixel-aligned Gaussian maps with simple self-attention 079 layers, requiring no input camera poses, intrinsics, nor post-alignment steps. The predicted Gaussian 080 maps serve as a high-fidelity and efficient 3D scene representation, while the Gaussian locations en-081 able ultra-fast camera extrinsics and intrinsics estimation in mere seconds using off-the-shelf solvers.

082 The training and rendering efficiency of 3D Gaussians makes it possible to consider more com-083 plex scene-level reconstruction. We thus train two FreeSplatter variants on Objaverse (Deitke et al., 084 2023) and a mixture of multiple scene datasets (Yao et al., 2020; Yeshwanth et al., 2023; Reizenstein 085 et al., 2021) to handle object-centric and scene-level reconstruction scenarios, respectively. The two variants, denoted as FreeSplatter-O (object) and FreeSplatter-S (scene), share a common model 087 architecture with \sim 306 million parameters while making essential adjustments for task adaption. 088 We conduct extensive experiments to demonstrate the superiority of FreeSplatter compared to prior sparse-view reconstruction methods in terms of reconstruction quality and pose estimation accu-089 racy. To be highlighted, FreeSplatter-O outperforms previous *pose-dependent* large reconstruction 090 models by a large margin, while FreeSplatter-S achieves a pose-estimation accuracy on par with 091 MASt3R (Leroy et al., 2024), the upgraded version of DUSt3R, on ScanNet++ and CO3D bench-092 marks. Furthermore, we showcase the potential of FreeSplatter in enhancing the productivity of 093 downstream applications, such as text/image-to-3D content creation. 094

095 096

097

068

069

2 RELATED WORK

098 Large Reconstruction Models. The availability of large-scale 3D datasets (Deitke et al., 2023; 099 2024) unlocks the possibility of training highly generalizable models for open-category image-to-3D 100 reconstruction. Large Reconstruction Models (LRMs) (Hong et al., 2024b; Xu et al., 2024c; Li et al., 101 2024a) leverage a scalable feed-forward transformer architecture to map sparse-view image tokens 102 into a 3D triplane NeRF (Mildenhall et al., 2021; Chan et al., 2022) and supervise it with multi-view 103 rendering loss. Subsequent works have made efforts on adopting mesh (Xu et al., 2024a; Wei et al., 104 2024; Wang et al., 2024c) or 3D Gaussians representation (Tang et al., 2024; Xu et al., 2024b; Zhang 105 et al., 2024b) for real-time rendering speed, exploring more efficient network architectures (Zheng et al., 2024; Zhang et al., 2024a; Chen et al., 2024a; Li et al., 2024b; Cao et al., 2024), improving the 106 texture quality (Boss et al., 2024; Siddiqui et al., 2024; Yang et al., 2024a), and exploiting explicit 107 3D supervision for better geometry (Liu et al., 2024a). Despite the superior reconstruction quality

and generalization capability, LRMs require *posed* images as input and are highly sensitive to the pose accuracy, limiting their applicable scenarios.

Pose-free Reconstruction. Classical pose-free reconstruction algorithms like Structure from Mo-111 tion (SfM) (Hartley & Zisserman, 2003; Ullman, 1979; Schonberger & Frahm, 2016) begins with 112 finding pixel correspondences across multiple views first, and then perform 3D points triangulation 113 and bundle adjustment to optimize 3D coordinates and camera parameters jointly. Recent efforts on 114 enhancing the SfM pipeline encompass leveraging learning-based feature descriptors (DeTone et al., 115 2018; Dusmanu et al., 2019; Revaud et al., 2019; Wang et al., 2023c) and image matchers (Edstedt 116 et al., 2023; 2024; Sarlin et al., 2020; Lindenberger et al., 2023), as well as differentiable bundle 117 adjustment (Lin et al., 2021; Wang et al., 2023a; Wei et al., 2020). SfM-based methods works 118 well with sufficient image overlaps between nearby views, but face challenges in matching correspondences from sparse views. On the other hand, learning-based methods (Jiang et al., 2024a;b; 119 Hong et al., 2024a; Fan et al., 2023) rely on learned data priors to recover the 3D geometry from 120 input views. PF-LRM (Wang et al., 2024a) extends the LRM framework by predicting a per-view 121 coarse point cloud for camera pose estimation with a differentiable PnP solver (Chen et al., 2022). 122 DUSt3R (Wang et al., 2024b) presents a novel framework for Multi-view stereo (MVS) by model-123 ing it as a pointmap-regression problem. Subsequent works further improve it in local representation 124 capability (Leroy et al., 2024) and reconstruction efficiency (Wang & Agapito, 2024). 125

Generalizable Gaussian Splatting. Compared to NeRF's (Mildenhall et al., 2021) MLP-based 126 implicit representation, 3DGS represents a scene explicitly with a point cloud, which strikes the 127 balance between high-fidelity and real-time rendering speed. However, the per-scene optimization of 128 3DGS requires densely-captured images and sparse point cloud generated by SfM for initialization. 129 Recent works (Charatan et al., 2024; Chen et al., 2024b; Szymanowicz et al., 2024; Liu et al., 2024b; 130 Wewer et al., 2024) have explored feed-forward models for sparse-view Gaussian reconstruction 131 by capitalizing on large-scale datasets and scalable model architectures. These models typically 132 assume the access to accurate camera poses and utilize 3D-to-2D geometric projection for feature 133 aggregation, adopting model designs like epipolar lines (Charatan et al., 2024) or plane-swept cost-134 volume (Chen et al., 2024b; Liu et al., 2024b). InstantSplat (Fan et al., 2024) and Splatt3R (Smart 135 et al., 2024) leverage the pose-free reconstruction capability of DUSt3R/MASt3R for sparse-view reconstruction. The former initializes the Gaussian positions using the DUSt3R point cloud and 136 then optimizes other Gaussian paremters, while the latter trains a Gaussian head in a frozen MASt3R 137 model. Despite the impressive results, their reconstruction quality highly depends on the point cloud 138 quality generated by DUSt3R. 139

140 141 3

142

146

3 Method

Given a sparse set of N input images $\{I^n \mid n = 1, ..., N\}$ with no known camera extrinsics nor intrinsics, FreeSplatter aims to reconstruct the scene as a set of Gaussians and estimate the camera parameters of the N images simultaneously. The pipeline can be formulated as:

$$\boldsymbol{G}, \boldsymbol{P}^1, \dots, \boldsymbol{P}^N, \boldsymbol{f} = \text{FreeSplatter}\left(\boldsymbol{I}^1, \dots, \boldsymbol{I}^N\right),$$
 (1)

where G denotes the reconstructed Gaussians which is the union of N pixel-wise Gaussian maps, *i.e.*, $G = \{G^n \mid n = 1, ..., N\}$, P^n denotes the estimated camera pose for I^n , and f denotes the focal length. We assume a common focal length for all input images which is reasonable in most scenarios.

We implement FreeSplatter as a feed-forward transformer that takes multi-view image as input and predict N Gaussian maps $\{G^n \mid n = 1, ..., N\}$ in a unified reference frame. The Gaussian maps enable both high-fidelity scene modeling and ultra-fast camera parameter estimation using off-theshelf solvers (Fischler & Bolles, 1981; Hartley & Zisserman, 2003; Plastria, 2011), since the Gaussian centers have located the surface point cloud explicitly. In this section, we first briefly review the background knowledge of Gaussian Splatting (Section 3.1), and then introduce the model architecture (Section 3.2) and training details (Section 3.3) of our approach.

- 158 159
- 159 3.1 PRELIMINARY 160
- 161 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) represents a scene as a set of 3D Gaussian primitives $\{g_k = (\mu_k, r_k, s_k, o_k, c_k) \mid k = 1, ..., K\}$, each primitive is parameterized by location

Image Tokens

Self-Attn

Patchify

& Linear

Patchify

&

Linear

Reference Embed

Source Embed



162

Input Views

165 166

167 168 169



171

172 173

Figure 2: FreeSplatter pipeline. Given input views $\{I^n \mid n = 1, ..., N\}$ without any known camera extrinsics or intrinsics, we first patchify them into image tokens, and then feed all tokens into a sequence of self-attention blocks to exchange information among multiple views. Finally, we decode the output image tokens into N Gaussian maps $\{G^n \mid n = 1, ..., N\}$, from which we can render novel views, as well as recovering camera focal length f and poses $\{P^n \mid n = 1, ..., N\}$ with simple iterative solvers.

MLP

Transformer Block ($\times L$)

Image Tokens

Novel View

Renderings

Linear

&

Unpatchify

Linear

& Unpatchify Gaussian Maps

Camera Poses

PnP Solver

180 $\mu_k \in \mathbb{R}^3$, rotation quaternion $r_k \in \mathbb{R}^4$, scale $s_k \in \mathbb{R}^3$, opacity $o_k \in \mathbb{R}$, and Spherical Harmonic 181 (SH) coefficients $c_k \in \mathbb{R}^{3 \times d}$ for computing view-dependent color, with d denoting the degree of 182 SH. These primitives parameterize the scene's radiance field with an explicit point cloud and can be 183 rendered into novel views with an efficient rasterizer. Compared to the expensive volume rendering 184 process of NeRF, 3DGS achieves comparable rendering quality while being much more efficient in 185 terms of time and memory.

186 187

3.2 MODEL ARCHITECTURE

As Figure 2 shows, FreeSplatter employs a transformer architecture similar to GS-LRM (Zhang et al., 2024b). Given a set of uncalibrated images $\{I^n \mid n = 1, ..., N\}$, we first tokenize them into image tokens $\{e^{n,m} \mid n = 1, ..., N, m = 1, ..., M\}$, where M denotes the patch number of each image. Then we feed all image tokens into a sequence of self-attention blocks to exchange information among multiple views, and finally decode them into N Gaussian maps $\{G^n \mid n = 1, ..., N\}$ composed of per-pixel Gaussian primitives, from which we can render novel views and recover camera parameters with simple iterative solvers.

Image Tokenization. The only input of our model is N images $\{I^n \in \mathbb{R}^{H \times W \times 3} | n = 1, ..., N\}$. Following ViT (Dosovitskiy et al., 2021), we split the images into $p \times p$ patches (p = 8) and flatten them into 1D vectors of dimension $p^2 \cdot 3$. Then a linear layer is applied to map the vectors into d-dim image tokens (d equals to the transformer width). Denoting the m-th token of the n-th image as $e^{n,m}$, we add it with two additional embeddings before fedding it into the transformer:

$$^{n,m} = \boldsymbol{e}^{n,m} + \boldsymbol{e}^m_{\text{pos}} + \boldsymbol{e}^n_{\text{view}},\tag{2}$$

where e_{pos}^m denotes the position embedding for the *m*-th patch, and e_{view}^n is a view embedding for distinguishing the reference view and other views. Specifically, we take the first image as the reference view and predict all Gaussian in its camera frame. We use a reference embedding e^{ref} for the first view (n = 1) and another source embedding e^{src} for other views (n = 2, ..., N), both of which are learnable during training.

Feed-forward Transformer. The image tokens added with position and view embeddings are then fed into a sequence of L self-attention blocks to exchange information across multiple views. Each block consists of a self-attention layer and an MLP layer, equipping with pre-normalization and residual connection (He et al., 2016). The forward process of each block can be formulated as:

211
212
213

$$e_{\text{attn}}^{n,m} = \text{SelfAttn} (\text{LayerNorm} (e_{\text{in}}^{n,m}), \{\text{LayerNorm} (e_{\text{in}}^{n,m})\}) + e_{\text{in}}^{n,m},$$
(3)
 $e_{\text{out}}^{n,m} = \text{MLP} (\text{LayerNorm} (e_{\text{attn}}^{n,m})) + e_{\text{attn}}^{n,m}.$

213 214

Gaussian Map Prediction. For each output token
$$e_{out}^{n,m}$$
 of the last (*L*-th) block, we decode it back into p^2 Gaussians with a simple linear layer, leading to a 1D vector of dimension $p^2 \cdot q$, where q



Figure 3: **Object-centric Sparse-view Reconstruction.** We show 6 samples from the Google Scanned Objects dataset. To be noted, the results of LGM and InstantMesh (2nd and 3rd rows) are generated with ground truth camera poses (and intrinsics), while our results (4th row) are generated in a completely pose-free manner.

is the parameter number of each Gaussian primitive. Then we unpatchify the predicted Gaussian parameters into a Gaussian patch $G^{n,m} \in \mathbb{R}^{p \times p \times q}$. Finally, we spatially concatenate the Gaussian patches $G^{n,m}$ for each view, leading to N Gaussian maps $\{G^n \in \mathbb{R}^{H \times W \times q} \mid n = 1, ..., N\}$.

Each pixel in the Gaussian maps contains a *q*-dim 3D Gaussian primitive. Prior pose-dependent
Gaussian LRMs (Zhang et al., 2024b; Xu et al., 2024b; Tang et al., 2024) parameterize each Gaussian's location with a single depth value. However, our model assumes no available camera poses
nor instrinsics, making it impossible to un-project the Gaussian locations from depths. Therefore, we directly predict the actual Gaussian locations in the reference frame and adopt a pixel-alignment loss to restrict Gaussians to lie on camera rays (see Section 3.3 for details).

Camera Pose Estimation. To recover the camera poses, we first estimate the focal length f from predicted Gaussian maps (see Section A.1 for details). Different from DUSt3R that can only feed 253 a pair of images into the model to predict pair-wise point maps and requires an additional align-254 ment to unify the pair-wise reconstruction results into a global point cloud, FreeSplatter predicts 255 the Gaussian maps of all views in a unified reference frame, thus can recover the camera poses 256 $\{P^n \mid n = 1, ..., N\}$ in a feed-forward manner without global alignment. For each Gaussian loca-257 tion map $X^n \in \mathbb{R}^{H \times W \times 3}$ (the first 3 channels of G^n), its pixel coordinate map $Y^n \in \mathbb{R}^{H \times W \times 2}$ 258 and a valid mask $M^n \in \mathbb{R}^{H \times W}$, we apply a PnP-RANSAC solver (Hartley & Zisserman, 2003; 259 Bradski, 2000) to estimate $P^n \in \mathbb{R}^{4 \times 4}$: 260

$$\mathbf{P}^{n} = \operatorname{PnP}\left(\mathbf{X}^{n}, \mathbf{Y}^{n}, \mathbf{M}^{n}, \mathbf{K}\right), \qquad (4)$$

where $K = [[f, 0, \frac{W}{2}], [0, f, \frac{H}{2}], [0, 0, 1]]$ denotes the estimated intrinsic matrix. The mask M^n aims to mark the pixel locations that contribute to the pose optimization, which differs in objectcentric and scene-level reconstruction scenarios, please see Section A.1 for more details.

3.3 TRAINING DETAILS

266 267

265

261

239

240

241

We train two FreeSplatter variants targeting for object-centric and scene-level pose-free reconstruc tion. They share the same model architecture and parameter size, but we made essential adjustments in training objectives and strategies to adapt them to different scenarios.

270 Staged Training Strategy. Prior pose-dependent LRMs (Hong et al., 2024b; Li et al., 2024a; Xu 271 et al., 2024b; Zhang et al., 2024b) leverage pure rendering loss for supervision. However, our model 272 assumes no known camera poses nor intrinsics and the Gaussian positions are free in 3D space, mak-273 ing it extremely challenging to predict correct Gaussian positions. Gaussian-based reconstruction 274 approaches heavily rely on the initialization of Gaussian positions, e.g., 3DGS (Kerbl et al., 2023) initializes the Gaussian positions with the sparse point cloud generated by SfM, while the param-275 eters of our model are randomly initialized at the beginning. In practice, we found it essential to 276 supervise the Gaussian positions at the beginning: 277

278

280

 $\mathcal{L}_{\text{pos}} = \sum_{n=1}^{N} \left\| \boldsymbol{M}^{n} \odot \hat{\boldsymbol{X}}^{n} - \boldsymbol{M}^{n} \odot \boldsymbol{X}^{n} \right\|,$ (5)

where $\hat{X}^n \in \mathbb{R}^{H \times W \times 3}$ is the predicted Gaussian position map, X^n is the ground truth point map unprojected from depths. $M^n \in \mathbb{R}^{H \times W}$ is a mask indicating the pixels storing valid depth values, which is exactly the foreground object mask for object-centric reconstruction. For scene-level reconstruction, M^n depends on where the depth values are defined in different datasets.

We apply \mathcal{L}_{pos} in the pre-training stage, so that the model learns to predict approximately correct Gaussian positions. In our experiments, this pre-training is **essential** to model's convergence. However, \mathcal{L}_{pos} can only supervise the pixels with valid depths, while the Gaussian positions predicted at other pixels remain unconstrained. Besides, the ground truth depths are noisy in some datasets, and applying \mathcal{L}_{pos} throughout the training leads to degraded rendering quality. To provide a more stable geometric supervision, we adopt a pixel-alignment loss to enforce each predicted Gaussian to be aligned with its corresponding pixel, which maximizes the ray cosine similarities:

292 293 $\mathcal{L}_{\text{align}} = \sum_{n=1}^{N} \sum_{i=0}^{H} \sum_{j=0}^{W} \left(1 - \frac{\hat{r}_{i,j}^{n} \cdot r_{i,j}^{n}}{\|\hat{r}_{i,j}^{n}\| \|r_{i,j}^{n}\|} \right),$ (6)

where $r_{i,j}^n$ denotes the ray from the camera origin t^n to point $X_{i,j}^n$. \mathcal{L}_{align} restricts the predicted Gaussians to be distributed on the camera rays, which improves the rendering quality. Besides, enforcing the pixel-alignment is important to camera parameter estimation, since the iterative solvers minimize the pixel-projection errors.

Loss Functions. The overall training objective of our model is:

299 300 301

 $\mathcal{L} = \mathcal{L}_{\text{render}} + \lambda_{\text{a}} \cdot \mathcal{L}_{\text{align}} + \mathbf{1}_{\text{t} \leq \text{T}_{\text{max}}} \lambda_{\text{p}} \cdot \mathcal{L}_{\text{pos}}, \tag{7}$

where the rendering loss \mathcal{L}_{render} is a combination of MSE and LPIPS loss. t and T_{max} denote the current training step and maximum pre-training step, respectively. In our implementation, we set $\lambda_{a} = 1.0, \lambda_{p} = 10.0$.

305 Occlusion in Pixel-aligned Gaussians. Prior pose-dependent Gaussian-basd LRMs (Tang et al., 306 2024; Zhang et al., 2024b; Xu et al., 2024b) parameterize the Gaussian positions with a single depth value to ensure they are pixel-aligned. Despite the simplicity, pixel-aligned Gaussians can 307 only represent the areas captured by input images. It is challenging for sparse input views to cover 308 the whole scene, resulting in missing reconstruction results in the occluded areas. Our model also 309 suffers from this problem due to the usage of \mathcal{L}_{align} . For the object-centric model, we adopt a simple 310 strategy to alleviate this problem by only calculate Equation 6 in the foreground area, while allowing 311 other Gaussians to move freely and model the occluded areas. 312

For the scene-level model trained on real-world images, all predicted Gaussians have be pixelaligned to model the complex background. We instead focus on reconstructing the observed areas. To avoid the areas in the target views that are invisible in input views provide negative guidance for training, we adopt the target-view masking strategy proposed by Splatt3R and only calculate the rendering loss in visible areas. Please refer to Smart et al. (2024) for more details.

318

4 EXPERIMENTS

319 320

Due to the space limit, we put more comprehensive experimental results in Section A.2 of the appendix, including comparison on object-level reconstruction with PF-LRM (Wang et al., 2024a) (Section A.2.1), more extensive image-to-3D generation results with FreeSplatter (Section A.2.2), cross-dataset generalization results (Section A.2.3), and additional ablation studies (Section A.3).



Figure 4: Scene-level Reconstruction on ScanNet++. The results of pixelSplat and MVSplat are obtained with ground truth input poses, while the results of Splat3R and ours are pose-free.

4.1 EXPERIMENTAL SETTINGS

Training Datasets. FreeSplatter-O assumes the input images are white-background with an object in the center, which is trained on Objaverse (Deitke et al., 2023), a large-scale dataset containing around 800K 3D assets. For each 3D asset, we normalize it into a $[-1, 1]^3$ cube, and then render 32 random views associated with depth maps around it, all in a resolution of 512×512 . FreeSplatter-S is trained on a mixture of multiple datasets, including BlendedMVS (Yao et al., 2020), ScanNet++ (Yeshwanth et al., 2023) and CO3Dv2 (Reizenstein et al., 2021), which is a subset of DUSt3R's (Wang et al., 2024b) training datasets covering the most representative scene types: outdoor, indoor, and realistic objects.

362 **Evaluation Datasets.** For object-centric reconstruction and pose-estimation, we evaluate our model 363 on OmniObject3D (Wu et al., 2023) and Google Scanned Objects (GSO) (Downs et al., 2022) 364 datasets, containing around 6K and 1K scanned 3D assets, respectively. We use the whole GSO 365 dataset and pick out 300 OmniObject3D objects in 30 categories as the evaluation set. For each 366 object, we render 20 random views and 4 additional structured input views for the sparse-view re-367 construction task, which are evenly placed around the objects with an elevation of 20° to ensure a 368 good coverage. For scene-level reconstruction and pose estimation, we mainly evaluate on the test 369 splits of ScanNet++ (Yeshwanth et al., 2023) and CO3Dv2 (Reizenstein et al., 2021).

370 371

372

350

351 352 353

354

4.2 Sparse-view Reconstruction

Baselines. Prior pose-free object reconstruction models like LEAP (Jiang et al., 2024b) were only
 trained on small datasets, showing very limited generalization capability in open-domain datasets.
 PF-LRM (Wang et al., 2024a) is the most comparable work to our FreeSplatter-O. Although its code
 is not publicly available, we compare with its inference results on its evaluation datasets on both
 sparse-view reconstruction and camera pose estimation tasks. Due to the space limit, we put the
 results in Section A.2.1 of the appendix, which demonstrate that FreeSplatter-O outperforms PF-

399

400

401

402



Figure 5: Scene-level Reconstruction on CO3Dv2.

LRM on its OmniObject3D benckmark and preserves visual details better. We also compare with two pose-dependent large reconstruction models, *i.e.*, LGM (Tang et al., 2024) and InstantMesh (Xu et al., 2024a), which are base on 3D Gaussians and tri-plane NeRF, respectively. We directly feed ground truth camera poses to them, while our model is totally pose-free.

For scene-level reconstruction, we compare with two state-of-the-art works on generalizable Gaussians, *i.e.*, pixelSplat (Charatan et al., 2024) and MVSplat (Chen et al., 2024b), both of which require camera poses. We fine-tune their checkpoints trained on RealEstate10K (Zhou et al., 2018) on Scan-Net++ for more fair comparison, while do not apply them on CO3Dv2 due to significant domain gap. Besides, we also compare with an up-to-date pose-free Gaussian reconstruction method, *i.e.*, Splatt3R (Smart et al., 2024), which utilizes frozen MASt3R (Leroy et al., 2024) model for Gaussian location prediction, while fine-tuning an additional head for predicting other Gaussian attributes.

410 Metrics. For both object-centric and scene-level reconstruction, We evaluate the rendering quality 411 of reconstructed Gaussians and report the standard metrics for novel view synthesis, *i.e.*, PSNR, 412 SSIM and LPIPS. All metrics are evaluated at the resolution of 512×512 .

 Results and Analysis. As demonstrated in Table 2, FreeSplatter models achieve superior performance across both object-centric and scene-level reconstruction tasks without camera poses.
 FreeSplatter-O significantly outperforms LGM and InstantMesh across all metrics even if they are pose-dependent, achieving PSNR improvements of > 5 and > 7 on previously unseen GSO and OmniObject3D datasets, respectively. Qualitative comparisons in Figure 3 further illustrate this performance gap, with InstantMesh and LGM exhibiting blurry artifacts (notably in the alphabets, 4th column), while our approach maintains clarity and detail fidelity consistent with ground truth views.

Existing Large Reconstruction Models (LRMs) assume the necessity of accurate camera poses for
high-quality 3D reconstruction, incorporating pose information through LayerNorm modulation (Li
et al., 2024a) or plucker ray embeddings (Xu et al., 2024c; Tang et al., 2024; Xu et al., 2024b).
However, FreeSplatter-O's superior performance challenges this assumption, suggesting that camera poses may not be essential for developing high-quality, scalable reconstruction models.

In scene-level reconstruction on ScanNet++, FreeSplatter-S demonstrates superior performance
across most metrics compared to pose-dependent pixelSplat and MVSplat, highlighting its robust
reconstruction quality and generalization capability. While the pose-free competitor Splatt3R leverages MASt3R's (Leroy et al., 2024) 3D point regression capabilities trained on extensive datasets,
its performance is constrained by MASt3R's prediction errors due to its frozen architecture and
Gaussian head-only training, resulting in fixed Gaussian positions. In contrast, our end-to-end
training approach with rendering supervision enables simultaneous optimization of Gaussian locations and attributes. Qualitative results on ScanNet++ and CO3Dv2 (Figure 4, Figure 5) illustrate

			_					
Mathad	OmniObject3D			GSO				
Wethod	$RRE \downarrow$	RRA@15° \uparrow	RRA@ 30° \uparrow	TE↓	$RRE \downarrow$	RRA@15° \uparrow	RRA@ 30° \uparrow	TE↓
FORGE	76.822	0.081	0.257	0.430	97.814	0.022	0.083	0.898
MASt3R	96.670	0.052	0.112	0.524	61.820	0.244	0.445	0.353
FreeSplatter-S	83.795	0.076	0.167	0.631	72.914	0.212	0.396	0.426
FreeSplatter-O	11.550	0.909	0.937	0.081	3.851	0.962	0.978	0.030
	ScanNet++				CO3Dv2			
	$RRE \downarrow$	RRA@15° \uparrow	RRA@ 30° \uparrow	TE↓	$RRE \downarrow$	RRA@15° \uparrow	RRA@ 30° \uparrow	TE↓
PoseDiffusion	-	-	-	-	7.950	0.803	0.868	0.328
MASt3R	0.724	0.988	0.993	0.104	2.918	0.975	0.989	0.112
FreeSplatter-S	0.791	0.982	0.987	0.110	3.054	0.976	0.986	0.148
	Table 2. Or antitating manufactor and manufactor							

Table 1: **Quantitative results on camera pose estimation.** We highlight the best metric as red.

Table 2. Qu	Table 2. Quantitative results on sparse-view reconstruction.							
Method		OmniObject3D			GSO			
Method	PSNR 1	SSIM ↑	LPIPS \downarrow	PSNR ↑	SSIM ↑	LPIPS \downarrow		
LGM (w/ gt pose)	24.852	0.942	0.060	24.463	0.891	0.093		
InstantMesh (w/ gt pose	24.077	0.945	0.062	25.421	0.891	0.095		
Ours (FreeSplatter-O)	31.929	0.973	0.027	30.443	0.945	0.055		
Method		ScanNet++			CO3Dv2			
	PSNR 1	SSIM ↑	LPIPS \downarrow	PSNR ↑	SSIM \uparrow	LPIPS \downarrow		
pixelSplat (w/ gt pose)	24.974	0.889	0.180	-	-	-		
MVSplat (w/ gt pose)	22.601	0.862	0.208	-	-	-		
Splatt3R	21.013	0.830	0.209	18.074	0.740	0.197		
Ours (FreeSplatter-S)	25.807	0.887	0.140	20.405	0.781	0.162		

our method's superior visual fidelity. To be noted, novel view synthesis for both FreeSplatter and Splatt3R is performed by aligning predicted cameras with target viewpoints.

4.3 CAMERA POSE ESTIMATION

Baselines. To evaluate our method's pose estimation capabilities, we benchmark against MASt3R, the current state-of-the-art in zero-shot multi-view pose estimation. Additional comparisons include
FORGE (Jiang et al., 2024a) for object-centric datasets and PoseDiffusion (Fan et al., 2023) for scene-level reconstruction, with PoseDiffusion evaluated exclusively on CO3Dv2 due to its training
scope. Traditional COLMAP-based approaches (Schonberger & Frahm, 2016) are excluded from our evaluation due to their documented high failure rates in sparse-view scenarios, a limitation consistently reported in previous studies (Wang et al., 2024a).

463 Metrics. Following established protocols Wang et al. (2024a; 2023b), we evaluate pose estima-464 tion through distinct rotation and translation metrics. Rotational accuracy is assessed via relative 465 rotation error (RRE) in degrees and relative rotation accuracy (RRA), where RRA represents the 466 percentage of rotation errors below predefined thresholds (15° and 30°). Translational accuracy is 467 measured through translation error (TE) between predicted and ground truth camera centers. For 468 scenarios with more than two input views, errors are computed pairwise and averaged across all 469 camera combinations.

As shown in Table 1, FreeSplatter-O demonstrates substantial performance improvements over ex-470 isting baselines on object-centric datasets, while FreeSplatter-S achieves comparable results with 471 MASt3R on scene-level datasets. MASt3R's poor performance on rendered object-centric datasets 472 without backgrounds can be attributed to the significant domain gap between its training data and 473 these synthetic images. Specifically, rendered object-centric images lack the complex backgrounds 474 and rich textures present in natural scenes, presenting unique challenges for pose estimation with-475 out adequate data priors. In scene-level evaluations on ScanNet++ and CO3Dv2, while MASt3R 476 achieves marginally superior metrics, FreeSplatter-S maintains competitive performance despite its 477 more limited training scope. This performance differential primarily stems from dataset scale dis-478 parity—our FreeSplatter-S utilizes only three datasets, a subset of DUSt3R/MASt3R's extensive 479 training corpus, resulting in fewer learned data priors. Future work will focus on scaling the model 480 across broader datasets to bridge this gap.

482 4.4 ABLATION STUDIES

483

481

432

452

453 454

455

Number of Input Views. We make an experiment on a GSO sample to illustrate how the number of input views influences the reconstruction quality on the object-centric scenarios. Please refer to Figure 16 in the appendix for more details.



Table 3: Ablation on the pixel-alignment loss.

Figure 6: **3D content creation with FreeSplatter.** 1st and 2nd rows: Image-to-3D results using Zero123++ (input image, Gaussian visualization, two novel views). 3rd row: Text-to-3D results using MVDream (prompt shown above; two Gaussian visualizations, two novel views).

Impact of Pixel-Alignment Loss. We study the impact of pixel-alignment loss defined in Equation 6. We remove the pixel-alignment loss from the training of both models, and report the novel view rendering metrics on GSO and ScanNet++ dataset. We can observe significant drops in all metrics without the pixel-alignment loss. We also compare the visual results in Figure 15 of the appendix, demonstrating that applying pixel-alignment loss leads to higher-fidelity renderings.

4.5 Applications

FreeSplatter can be seamlessly integrated into 3D content creation pipelines, and its pose-free nature can potentially bring great convenience to users and enhance the productivity by a large margin. In a classical 3D generation pipeline (Li et al., 2024a; Xu et al., 2024a; Tang et al., 2024; Wang et al., 2024c) that first generates multi-view images with a multi-view diffusion model and then feed them into an LRM for reconstruction, we have to figure out how the camera poses of the multi-view diffusion model are defined and carefully align them with the LRM. With FreeSplatter as the content creator, these tedious steps are not required any longer. We just feed the multi-view images into FreeSplatter and get the outputs instantly after several seconds. Figure 6 shows some results of text/image-to 3D generation with MVDream (Shi et al., 2024) and Zero123++ (Shi et al., 2023), respectively. We notice that our method can render significantly more clear views from the images than previous pose-dependent LRMs, showing stronger robustness to multi-view inconsistency.

5 CONCLUSION

CONCLUSION

In this work, we present FreeSplatter, a highly scalable framework for pose-free sparse-view reconstruction. By leveraging a single-stream transformer architecture and predicting multi-view Gaussian maps in a unified frame, FreeSplatter enables both high-fidelity 3D modeling and instant camera
pose estimation. We provide two model variants for both object-centric and scene-level reconstruction, both showing superior recontruction qulity and pose estimation accuracy. FreeSplatter also
exhibits great potential in enhancing the productivity of downstream applications like text/image-to-3D content creation, which can liberate users from dealing with complex camera poses.

Limitations. Despite the promising performance of FreeSplatter, its pre-training stage relies on depth data, making it non-trivial to be trained on datasets with no depth labels, *e.g.*, RealEstate10K
and MVImgNet (Yu et al., 2023). Besides, FreeSplatter requires two distinct model variants to deal with object-centric and scene-level reconstructions, while a unified model for both tasks would be preferred. We leave this as the future work.

540	REFERENCES
541	

550

551

552

553

559

579

580

- Mark Boss, Zixuan Huang, Aaryaman Vasishta, and Varun Jampani. Sf3d: Stable fast 3d
 mesh reconstruction with uv-unwrapping and illumination disentanglement. arXiv preprint
 arXiv:2408.00653, 2024.
- G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.
- Ang Cao, Justin Johnson, Andrea Vedaldi, and David Novotny. Lightplane: Highly-scalable components for neural 3d fields. *arXiv preprint arXiv:2404.19760*, 2024.
 - Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16123–16133, 2022.
- David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19457–19467, 2024.
- Anpei Chen, Haofei Xu, Stefano Esposito, Siyu Tang, and Andreas Geiger. Lara: Efficient large baseline radiance fields. *arXiv preprint arXiv:2407.04699*, 2024a.
- Hansheng Chen, Pichao Wang, Fan Wang, Wei Tian, Lu Xiong, and Hao Li. Epro-pnp: Generalized end-to-end probabilistic perspective-n-points for monocular object pose estimation. In *Proceed-ings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2781–2790, 2022.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear
 memory cost. *arXiv preprint arXiv:1604.06174*, 2016.
- Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*, 2024b.
- Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu,
 Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. Abo: Dataset and
 benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 21126–21136, 2022.
- Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-tated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13142–13153, 2023.
 - Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36, 2024.
- Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest
 point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 224–236, 2018.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=YicbFdNTTy.
- Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In 2022 International Conference on Robotics and Automation (ICRA), pp. 2553–2560. IEEE, 2022.

605

622

631

638

- Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-net: A trainable cnn for joint description and detection of local features. In *Proceedings of the ieee/cvf conference on computer vision and pattern recognition*, pp. 8092– 8101, 2019.
- Johan Edstedt, Ioannis Athanasiadis, Mårten Wadenbäck, and Michael Felsberg. Dkm: Dense ker nelized feature matching for geometry estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17765–17775, 2023.
- Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. Roma: Robust dense feature matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19790–19800, 2024.
- Zhiwen Fan, Panwang Pan, Peihao Wang, Yifan Jiang, Hanwen Jiang, Dejia Xu, Zehao Zhu, Dilin
 Wang, and Zhangyang Wang. Pose-free generalizable rendering transformer. *arXiv e-prints*, pp. arXiv–2310, 2023.
- ⁶⁰⁹
 ⁶⁰⁹
 ⁶¹⁰
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹³
 ⁶¹⁴
 ⁶¹⁵
 ⁶¹⁵
 ⁶¹⁶
 ⁶¹⁶
 ⁶¹⁷
 ⁶¹⁷
 ⁶¹⁸
 ⁶¹⁸
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁰
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹³
 ⁶¹⁴
 ⁶¹⁵
 ⁶¹⁵
 ⁶¹⁶
 ⁶¹⁷
 ⁶¹⁷
 ⁶¹⁸
 ⁶¹⁸
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁹
 ⁶¹⁰
 ⁶¹¹
 ⁶¹¹
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹¹
 ⁶¹²
 ⁶¹²
 ⁶¹¹
 ⁶¹¹
 ⁶¹²
 ⁶¹¹
 ⁶¹¹
 ⁶¹¹
 ⁶¹¹
 ⁶¹¹
 ⁶¹¹
 ⁶¹¹
- Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting
 with applications to image analysis and automated cartography. *Communications of the ACM*, 24 (6):381–395, 1981.
- Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jiaolong Yang, Seungryong Kim, and Chong Luo.
 Unifying correspondence pose and nerf for generalized pose-free novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20196–20206, 2024a.
- Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3d. In *The Twelfth International Conference on Learning Representations*, 2024b. URL https://openreview.
 net/forum?id=sllU8vvsFF.
- Hanwen Jiang, Zhenyu Jiang, Kristen Grauman, and Yuke Zhu. Few-view object reconstruction with unknown categories and camera poses. In 2024 International Conference on 3D Vision (3DV), pp. 31–41. IEEE, 2024a.
- Hanwen Jiang, Zhenyu Jiang, Yue Zhao, and Qixing Huang. LEAP: Liberate sparse-view 3d modeling from camera poses. In *The Twelfth International Conference on Learning Representations*, 2024b. URL https://openreview.net/forum?id=KPmajBxEaF.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- Benjamin Lefaudeux, Francisco Massa, Diana Liskovich, Wenhan Xiong, Vittorio Caggiano, Sean Naren, Min Xu, Jieru Hu, Marta Tintore, Susan Zhang, Patrick Labatut, Daniel Haziza, Luca Wehrstedt, Jeremy Reizenstein, and Grigory Sizov. xformers: A modular and hackable transformer modelling library. https://github.com/facebookresearch/xformers, 2022.
- 647 Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. arXiv preprint arXiv:2406.09756, 2024.

648 Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan 649 Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view gen-650 eration and large reconstruction model. In The Twelfth International Conference on Learning 651 Representations, 2024a. URL https://openreview.net/forum?id=21DQLiH1W4. 652 Mengfei Li, Xiaoxiao Long, Yixun Liang, Weiyu Li, Yuan Liu, Peng Li, Xiaowei Chi, Xinggun 653 Qi, Wei Xue, Wenhan Luo, et al. M-Irm: Multi-view large reconstruction model. arXiv preprint 654 arXiv:2406.07648, 2024b. 655 656 Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural 657 radiance fields. In Proceedings of the IEEE/CVF international conference on computer vision, 658 pp. 5741-5751, 2021. 659 Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. Lightglue: Local feature matching 660 at light speed. In Proceedings of the IEEE/CVF International Conference on Computer Vision, 661 pp. 17627-17638, 2023. 662 663 Minghua Liu, Chong Zeng, Xinyue Wei, Ruoxi Shi, Linghao Chen, Chao Xu, Mengqi Zhang, 664 Zhaoning Wang, Xiaoshuai Zhang, Isabella Liu, et al. Meshformer: High-quality mesh gener-665 ation with 3d-guided reconstruction model. arXiv preprint arXiv:2408.10198, 2024a. 666 667 Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao, Wei Li, and Ziwei Liu. Fast generalizable gaussian splatting reconstruction from multi-view stereo. 668 arXiv preprint arXiv:2405.12218, 2024b. 669 670 I Loshchilov. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017. 671 672 Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, 673 Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision 674 training. arXiv preprint arXiv:1710.03740, 2017. 675 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and 676 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. Communications 677 of the ACM, 65(1):99–106, 2021. 678 679 Frank Plastria. The weiszfeld algorithm: proof, amendments, and extensions. Foundations of loca-680 tion analysis, pp. 357-389, 2011. 681 682 Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and 683 David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In Proceedings of the IEEE/CVF international conference on computer 684 vision, pp. 10901–10911, 2021. 685 686 Jerome Revaud, Cesar De Souza, Martin Humenberger, and Philippe Weinzaepfel. R2d2: Reliable 687 and repeatable detector and descriptor. Advances in neural information processing systems, 32, 688 2019. 689 690 Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: 691 Learning feature matching with graph neural networks. In Proceedings of the IEEE/CVF confer-692 ence on computer vision and pattern recognition, pp. 4938-4947, 2020. 693 Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In Proceedings 694 of the IEEE conference on computer vision and pattern recognition, pp. 4104–4113, 2016. 695 696 Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, 697 Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base 698 model. arXiv preprint arXiv:2310.15110, 2023. 699 Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view 700 diffusion for 3d generation. In The Twelfth International Conference on Learning Representa-701

tions, 2024. URL https://openreview.net/forum?id=FUgrjq2pbB.

702 703 704 705 706	Yawar Siddiqui, Tom Monnier, Filippos Kokkinos, Mahendra Kariya, Yanir Kleiman, Emilien Gar- reau, Oran Gafni, Natalia Neverova, Andrea Vedaldi, Roman Shapovalov, et al. Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. <i>arXiv preprint</i> <i>arXiv:2407.02445</i> , 2024.
708 707 708	Brandon Smart, Chuanxia Zheng, Iro Laina, and Victor Adrian Prisacariu. Splatt3r: Zero-shot gaussian splatting from uncalibarated image pairs. <i>arXiv preprint arXiv:2408.13912</i> , 2024.
709 710 711 712	Stanislaw Szymanowicz, Chrisitian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pp. 10208–10217, 2024.
713 714 715	Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. <i>arXiv preprint arXiv:2402.05054</i> , 2024.
716 717 718	Shimon Ullman. The interpretation of structure from motion. <i>Proceedings of the Royal Society of London. Series B. Biological Sciences</i> , 203(1153):405–426, 1979.
719 720	Hengyi Wang and Lourdes Agapito. 3d reconstruction with spatial memory. <i>arXiv preprint arXiv:2408.16061</i> , 2024.
721 722 723	Jianyuan Wang, Nikita Karaev, Christian Rupprecht, and David Novotny. Visual geometry grounded deep structure from motion. <i>arXiv preprint arXiv:2312.04563</i> , 2023a.
724 725 726	Jianyuan Wang, Christian Rupprecht, and David Novotny. Posediffusion: Solving pose estimation via diffusion-aided bundle adjustment. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pp. 9773–9783, 2023b.
727 728 729 730 731	Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. PF-LRM: Pose-free large reconstruction model for joint pose and shape prediction. In <i>The Twelfth International Conference on Learning Representations</i> , 2024a. URL https://openreview.net/forum?id=noe76eRcPC.
732 733 734	Shuzhe Wang, Juho Kannala, Marc Pollefeys, and Daniel Barath. Guiding local feature matching with surface curvature. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pp. 17981–17991, 2023c.
735 736 737 738	Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Ge- ometric 3d vision made easy. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision</i> <i>and Pattern Recognition</i> , pp. 20697–20709, 2024b.
739 740 741	Zhengyi Wang, Yikai Wang, Yifei Chen, Chendong Xiang, Shuo Chen, Dajiang Yu, Chongxuan Li, Hang Su, and Jun Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction model. <i>arXiv preprint arXiv:2403.05034</i> , 2024c.
742 743 744 745	Xingkui Wei, Yinda Zhang, Zhuwen Li, Yanwei Fu, and Xiangyang Xue. Deepsfm: Structure from motion via deep bundle adjustment. In <i>Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16</i> , pp. 230–247. Springer, 2020.
746 747 748	Xinyue Wei, Kai Zhang, Sai Bi, Hao Tan, Fujun Luan, Valentin Deschaintre, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Meshlrm: Large reconstruction model for high-quality mesh. <i>arXiv</i> preprint arXiv:2404.12385, 2024.
749 750 751 752	Christopher Wewer, Kevin Raj, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. <i>arXiv preprint</i> <i>arXiv:2403.16292</i> , 2024.
753 754 755	Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, et al. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pp. 803–814, 2023.

756 757 758	Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. <i>arXiv preprint arXiv:2404.07191</i> , 2024a.
759 760 761 762	Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. <i>arXiv preprint arXiv:2403.14621</i> , 2024b.
763 764 765 766	Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, and Kai Zhang. DMV3d: Denoising multi-view diffusion using 3d large reconstruction model. In <i>The Twelfth International Conference on Learning Representa-tions</i> , 2024c. URL https://openreview.net/forum?id=H4yQefeXhp.
767 768 769 770	Fan Yang, Jianfeng Zhang, Yichun Shi, Bowen Chen, Chenxu Zhang, Huichao Zhang, Xiaofeng Yang, Jiashi Feng, and Guosheng Lin. Magic-boost: Boost 3d generation with mutli-view conditioned diffusion. <i>arXiv preprint arXiv:2404.06429</i> , 2024a.
771 772 773	Xianghui Yang, Huiwen Shi, Bowen Zhang, Fan Yang, Jiacheng Wang, Hongxu Zhao, Xinhai Liu, Xinzhou Wang, Qingxiang Lin, Jiaao Yu, et al. Hunyuan3d-1.0: A unified framework for text-to-3d and image-to-3d generation. <i>arXiv preprint arXiv:2411.02293</i> , 2024b.
774 775 776	Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In <i>Proceedings of</i> <i>the IEEE/CVF conference on computer vision and pattern recognition</i> , pp. 1790–1799, 2020.
778 779 780	Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pp. 12–22, 2023.
781 782 783 784	Xianggang Yu, Mutian Xu, Yidan Zhang, Haolin Liu, Chongjie Ye, Yushuang Wu, Zizheng Yan, Chenming Zhu, Zhangyang Xiong, Tianyou Liang, et al. Mvimgnet: A large-scale dataset of multi-view images. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> , pp. 9150–9161, 2023.
785 786 787	Chubin Zhang, Hongliang Song, Yi Wei, Yu Chen, Jiwen Lu, and Yansong Tang. Geolrm: Geometry-aware large reconstruction model for high-quality 3d gaussian generation. <i>arXiv</i> preprint arXiv:2406.15333, 2024a.
789 790	Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields, 2022.
791 792 793	Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-lrm: Large reconstruction model for 3d gaussian splatting. arXiv preprint arXiv:2404.19702, 2024b.
795 796 797	Xin-Yang Zheng, Hao Pan, Yu-Xiao Guo, Xin Tong, and Yang Liu. Mvd ² : Efficient multiview 3d reconstruction for multiview diffusion. In <i>ACM SIGGRAPH 2024 Conference Papers</i> , pp. 1–11, 2024.
798 799 800 801	Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. <i>ACM Trans. Graph</i> , 37, 2018.
802 803	
805	
806	
808 808	
809	

810 A APPENDIX

831 832 833

834 835

841

842 843

844

845

846

847

848

849

850

851

852

853

854

863

812 A.1 IMPLEMENTATION DETAILS

814 Implementation Details. FreeSplatter-O and FreeSplatter-S share the same model architecture of a 24-layer transformer with a width of 1024, containing ~ 306 million parameters. We do not 815 use bias term throughout the model except for the output linear layer for Gaussian prediction. The 816 model is trained on 256×256 input images for 500K steps and then fine-tuned on 512×512 817 input images for 100K steps on 16 NVIDIA H800 GPUs. We adopt an AdamW (Loshchilov, 2017) 818 optimizer with $\beta_1 = 0.9, \beta_2 = 0.95$, the weight decay and learning rate are set to 0.01 and 4 \times 819 10^{-5} , respectively. The maximum pre-training step $T_{\rm max} = 2 \times 10^5$, *i.e.*, we adopt the Gaussian 820 position loss in Equation 5 for the first 200K training steps. To enhance the training efficiency, 821 we utilize the xFormers (Lefaudeux et al., 2022) library for memory-efficient attention, gradient 822 checkpointing (Chen et al., 2016), bfloat16 training (Micikevicius et al., 2017), and deferred back-823 propagation (Zhang et al., 2022) for GS rendering. In each batch, We sample 4 and 2 input views 824 for FreeSplatter-O and FreeSplatter-S respectively and 4 novel views for supervision.

Camera Intrinsic Estimation. Following DUSt3R (Wang et al., 2024b), we assume centered principle point and square pixels in this work. Thanks to the point-cloud-based Gaussian representation, we can extract the Gaussian locations as "point maps" $\{X^n \in \mathbb{R}^{H \times W \times 3} | n = 1, ..., N\}$ from the predicted Gaussian maps. Since all points are predicted in the first view's camera frame, we can estimate the focal length of the first view by minimizing the pixel-projection errors:

$$f_*^1 = \arg\min_{f^1} \sum_{i=0}^W \sum_{j=0}^H \left\| (i', j') - f^1 \frac{\left(X_{i,j,0}^1, X_{i,j,1}^1 \right)}{X_{i,j,2}^1} \right\|,\tag{8}$$

where $(i', j') = (i - \frac{W}{2}, j - \frac{H}{2})$ denotes pixel coordinates relative to the principle point $(\frac{H}{2}, \frac{W}{2})$. This optimization problem can be solved by Weiszfeld algorithm (Plastria, 2011) efficiently.

Big Different from DUSt3R that only accepts binocular images as input, our model architecture naturally supports single-view reconstruction. Thus we feed each image I^n into the model individually to predict its monocular Gaussian map G^n_{mono} , and then estimate its focal length f^n from X^n_{mono} . Finally, we compute the average focal length and adopt it for all views: $f = \frac{1}{N} \sum_{n=1}^{N} f^n$.

Masks for PnP-RANSAC solver. We adopt different masks M^n in Equation 4 for object-centric and scene-level reconstruction when estimating the input camera poses:

- For object-centric reconstruction, the model is trained on white-background images rendered from 3D assets, and the predicted Gaussians at the background area are not restricted to be pixel-aligned (See Section 3.3 for details), which may influence the camera pose solving. Therefore, we assume white-background input images at inference time, and adopt the foreground mask segmented by an off-the-shelf background-removal tool¹ as M^n . The segmentation masks are not always perfect but are precise enough for the PnP-RANSAC solver to estimate the poses robustly.
- For scene-level reconstruction, we do not distinguish the foreground and background areas during training, and all the predicted Gaussians are restricted to be pixel-aligned by applying the pixel-alignment loss universally. Therefore, all visible Gaussians should contribute to pose solving and we use the Gaussian opacity map Oⁿ ∈ ℝ^{H×W} to compute Mⁿ with a minimal visibility threshold τ: Mⁿ = (Oⁿ > τ).

Camera Normalization. The definition of reconstruction frame plays an important role in model training. We take the camera frame of the first view as the reference frame for reconstruction, and normalize all cameras in a training batch into this frame. Besides, a scaling operation is required to deal with various scene scales. Denoting the original input cameras as $\{P_{in}^n = [R_{in}^n | t_{in}^n] | n = 1, ..., N\}$, we adopt different camera normalization strategies for objectcentric and scene-level reconstruction.

For object-centric reconstruction, the training cameras are sampled orbiting the object center. We first scale all input cameras to make the distance from the first camera to the object center equal to

¹https://github.com/danielgatis/rembg

2, and then transform all cameras into the reference frame so that the first view's camera pose is an identity matrix. With this strategy, we fix the object center at (0, 0, 2) in the reference frame. For scene-level reconstruction, the camera distributions are more complex. Thus we fist transform all cameras into the reference frame, and then scale them using a factor s = 1/d, d is the mean distance of all valid points to the origin:

869 870

871 872

$$d = \frac{1}{\sum_{n=1}^{N} |\mathbf{M}^{n}|} \sum_{n=1}^{N} \sum_{i=0}^{N} \sum_{j=0}^{N} M_{i,j}^{n} \cdot \|\mathbf{X}_{i,j}^{n}\|,$$
(9)

where $X^n \in \mathbb{R}^{H \times W \times 3}$ denotes the ground truth point map unprojected from depth map, $M^n \in \mathbb{R}^{H \times W}$ is the valid depth mask. To be noted, we also scale the ground truth depths when scaling the cameras.

876 Scene Scale Ambiguity. Given a set of scene images, we can not infer the actual scale of the scene, 877 since we can obtain identical rendered images when scaling the scene and cameras simultaneously, 878 *i.e.*, the scene scale ambiguity problem. For object-centric reconstruction, we have fixed the object 879 center at (0, 0, 2) in the reference frame, so there is no scale ambiguity problem. But for scene 880 images, it is hard to define a canonical center and it is reasonable for our model to reconstruct the 881 scene in arbitrary scale. Besides, the training data is a mixture of multiple datasets that vary greatly in scale, which may make the model confused on scale prediction. Then a question arises, how can 882 we align the scale of the model's prediction with the scale of the training cameras so that we can 883 render target views for supervision? 884

885 As mentioned above, we scale the camera locations in a training batch using s = 1/d (d is defined in 886 Equation 9), which can bound the scene into a range reasonable for the model to predict. However, 887 the scale factor is related to the choice of valid points, leading to different scene scales with a minimal difference in selected points. The model would get confused if we train it to reconstruct exactly in this scale since it has no idea which points we utilized to compute this scale. Inspired by 889 DUSt3R (Wang et al., 2024b), we choose to rescale the reconstructed Gaussians before rendering 890 them to deal with scene ambiguity. We first compute the scale factor $\hat{s} = 1/d$ using Equation 9 with 891 predicted Gaussian positions \ddot{X}^n and the same mask M^n . Then we adjust the *position* and *scale* of 892 each Gaussian as $\hat{\mu}_k = \hat{s} \cdot \mu_k$, $\hat{s}_k = \hat{s} \cdot s_k$. With this operation, we do not care about the absolute 893 scale of the reconstructed scene, but we expect it to align with the ground truth scene after scaling 894 them using the factors \hat{s} and s respectively. In the pre-training stage, we compute \mathcal{L}_{pos} using the 895 scaled predicted positions and ground truth positions too. 896

- A.2 ADDITIONAL EXPERIMENTS
 - A.2.1 COMPARISON WITH PF-LRM

PF-LRM is a highly-relevant work to our FreeSplatter-O model. We make a comparison with it on
 its GSO and OmniObject3D evaluation datasets, both quantitatively and qualitatively.

903 Quantitative comparison. We show the quantitative results in Table 4 and Table 5. On both sparse-904 view reconstruction and pose estimation tasks, our model achieves lower metrics than PF-LRM on 905 GSO dataset and higher metrics on Omni3D dataset. However, we point out that the images in 906 PF-LRM's GSO evaluation set are rendered with the same rendering engine and hyper-parameters 907 (e.g., light intensity) as PF-LRM's training dataset, which is beneficial for PF-LRM's evaluaton 908 metrics due to the smaller gap between its training and evaluation datasets. On the contrary, the images in the Omni3D evaluation set are not manually rendered but are from the original dataset 909 itself. Therefore, we argue that Omni3D is a more fair benchmark for comparison, and our model 910 consistently outperforms PF-LRM on this benchmark. 911

- **Qualitative comparison.** We also show qualitative comparison results on both datasets in Figure 7,
 which demonstrate that our FreeSplatter synthesizes significantly better visual details than PF-LRM.
- 914

897

899

- 915 A.2.2 MORE RESULTS ON IMAGE-TO-3D GENERATION 916
- 917 In this section, we provide a large set of visualizations on image-to-3D generation by combining our FreeSplatter-O with different multi-view diffusion models.

Method	GSO			OmniObject3D			
Withiliti	PSNR ↑	SSIM \uparrow	LPIPS \downarrow	PSNR ↑	SSIM \uparrow	LPIPS \downarrow	
		Evaluate renderings at G.T. novel-view poses					
PF-LRM	25.08	0.877	0.095	21.77	0.866	0.097	
FreeSplatter-O	23.54	0.864	0.100	22.83	0.876	0.088	
	Evaluate renderings at predicted input poses						
PF-LRM	27.10	0.905	0.065	25.86	0.901	0.062	
FreeSplatter-O	25.50	0.897	0.076	26.49	0.926	0.050	

Table 4: Quantitative comparison with PF-LRM on sparse-view reconstruction.

Table 5: Quantitative comparison with PF-LRM on pose estimation.

Mathad		G	SO		OmniObject3D			
Wiethou	$RRE \downarrow$	RRA@15° \uparrow	RRA@30°↑	TE↓	RRE↓	RRA@15° \uparrow	RRA@ 30° \uparrow	TE↓
PF-LRM	3.99	0.956	0.976	0.041	8.013	0.889	0.954	0.089
FreeSplatter-O	8.96	0.909	0.936	0.090	3.446	0.982	0.996	0.039

Comparison with pose-dependent LRMs. In Figure 8 and Figure 9, we compare FreeSplatter's image-to-3D generation results with pose-dependent LRMs, *i.e.*, InstantMesh and LGM, using Zero123++ v1.2 (Shi et al., 2023) and a recent model Hunyuan3D Std (Yang et al., 2024b) as the multi-view generator, respectively. For each input image, we fix the random seed to generate multi-view images, and visualize the novel view synthesis results of each reconstruction model. From the results, we can observe that FreeSplatter generates significantly more clear views and preserves the geometry and texture details better than other baselines.

Using input image to enhance the results. In Figure 10, we show another interesting use case of FreeSplatter, i.e., using the input image to enhance the image-to-3D generation results. For multi-view diffusion models like Zero123++ v1.1 (Shi et al., 2023), it generates 6 views from an input image at pre-defined poses, but the pose of the input image is unknown (its azimuth is known to be 0° but elevation is unknown). In this case, classical pose-dependent LRMs cannot leverage the input image for reconstruction, but FreeSplatter is able to do this! As Figure 10 shows, using the input image apart from generated views can significantly enhance the reconstruction results in many cases, especially for contents that Zero123++ struggles to generate, e.g., human faces. Compared to generated views, the input image is often more high-quality and contains richer visual details. The pose-free nature of FreeSplatter make it capable of exploiting these precious information in the reconstruction process.

Recovering the predefined camera poses of multi-view diffusion models. In Figure 11, we show
 that FreeSplatter can faithfully recover the pre-defined camera poses of existing multi-view diffusion
 models from its reconstructed Gaussian maps. This demonstrates its robustness to generated multi-view images which may contain inconsistent contents.



Figure 7: Comparison with PF-LRM on its evaluation datasets. The test samples in the first 3 1020 rows are from the GSO evaluation set of PF-LRM, while the samples in the last 4 rows are from the OmniObject3D evaluation set of PF-LRM. Our FreeSplatter-O synthesizes significantly better 1022 visual details than PF-LRM.

1021

1024











Figure 10: Use input image to enhance the image-to-3D generation results with Zero123++ v1.1 (Shi et al., 2023). The input image is often more high-quality and contains richer visual details than the generated views, but its camera pose is unknown, making it impossible for pose-dependent LRMs to leverage it. The capability of using the input image alongside generated views of our FreeSplatter is particularly valuable for challenging content like human faces, where Zero123++ often struggles to generate.



Figure 11: Our FreeSplatter can faithfully recover the pre-defined camera poses of existing multi-view diffusion models. We use gray pyramids to visualize the ground truth pre-defined camera poses of the diffusion models, and colorful pyramids to visualize the estimated poses. φ and θ denote the pre-defined azimuth and elevation angles, respectively. Since Zero123++ v1.2 and Hunyuan3D Std generate images at pre-defined fixed Field-of-View (fov), we mark their pre-defined fov angles and corresponding focal lengths (in pixel units) too.

1242 A.2.3 CROSS-DATASET GENERALIZATION

FreeSplatter-O generalizes to other object datasets. To demonstrate the cross-dataset generalization ability of our FreeSplatter-O, we visualize the object-centric reconstruction results on OmniObject3D and ABO (Collins et al., 2022) datasets in Figure 12. Both datasets are outside the training scope of FreeSplatter-O. We can observe that FreeSplatter-O generates very high-quality renderings.

FreeSplatter-S generalizes to RealEstate10K. Although our FreeSplatter-S model was not trained on RealEstate10K (Zhou et al., 2018), we directly evaluate its zero-shot view synthesis capability on this dataset. We visualize the qualitative results in Figure 13, showing that our model can faithfully reconstruct the input views at the estimated input poses, while the rendered novel views align well with the ground truth. We also report the quantitative results of both sparse-view reconstruction and pose-estimation tasks in Table 6 and Table 7, respectively. In summary, pixelSplat and MVS-plat achieve better metrics on the sparse-view reconstruction task, which is foreseeable since they utilize RealEstate10K as training dataset and ground truth camera poses as input, while our model is not trained on this dataset and leverages no camera information. However, our model signifi-cantly outperforms another pose-free baseline Splatt3R which is built upon MASt3R trained on a larger dataset, demonstrating the effectiveness of our framework. For camera pose estimation on RealEstate10K, MASt3R achieves the best Relative Rotation Error (RRE) while our FreeSplatter-S performs the best in all other metrics.

Table 6: Quantitative sparse-view reconstruction results on RealEstate10K.

Method	PSNR ↑	SSIM ↑	LPIPS \downarrow
pixelSplat (w/ GT poses)	24.469	0.829	0.224
MVSplat (w/ GT poses)	20.033	0.789	0.280
Splatt3R (no pose)	16.634	0.604	0.422
FreeSplatter-S (no pose)	18.851	0.659	0.369

Table 7: Quantitative pose estimation results on RealEstate10K.

Method	$RRE \downarrow$	RRA@15° \uparrow	RRA@30°↑	TE↓
PoseDiffsion	14.387	0.732	0.780	0.466
RayDiffsion	12.023	0.767	0.814	0.439
RoMa	5.663	0.918	0.947	0.402
MASt3R	2.341	0.972	0.994	0.374
FreeSplatter-S	3.513	0.982	0.995	0.293

FreeSplatter-S generalizes to various other datasets. We demonstrate FreeSplatter-S's broad generalization capabilities through extensive evaluation across diverse datasets in Figure 14. Our results span scanned objects (DTU dataset), large-scale outdoor scenes (Tanks & Temples), multiview object collections (MVImgNet), and AI-generated synthetic contents (Sora-generated videos²). These results validate our framework's robustness across varying scene types and capture conditions.

²https://openai.com/index/sora/



Figure 12: Zero-shot pose-free reconstruction results on ABO and OmniObject3D. Both datasets are unseen for FreeSplatter-O. Our model can faithfully estimate the input camera poses and render high-fidelity novel views.



Figure 13: **Zero-shot pose-free reconstruction and view synthesis results on RealEstate10K.** Our FreeSplatter-S model was not trained on RealEstate10K, we directly utilize it for zero-shot view synthesis on this dataset. We can observe that our model can faithfully reconstruct the input views at the estimated input poses, and the rendered novel views align well with the ground truth.

1349

1314

1315





A.3 ABLATION STUDIES

1406 Model architecture. Our FreeSplatter model is a 24-layer transformer with a patch size of 8. In our initial experiments of this project, we have trained object-level reconstruction models with different 1407 layers and patch size. We compare these models with our final model to evaluate the influence of 1408 model architecture on the performance. Specifically, we evaluate on the Google Scanned Objects 1409 dataset and report the results in Table 8. The results demonstrate that, using more layers and smaller 1410 patch size can consistently improve the model's performance. This is reasonable since more trans-1411 former layers leads to more parameters, while smaller patch size leads to less memory usage (less 1412 image tokens for attention computation) but higher information compression ratio. 1413

Table 8: Ablation study on model architectures. The results are evaluated on GSO dataset. *L* and *P* denote number of transformer layers and patch size, respectively.

Architecture	PSNR ↑	SSIM \uparrow	LPIPS \downarrow
L = 16, P = 16	25.417	0.896	0.088
L = 16, P = 8	28.945	0.934	0.064
L = 24, P = 16	28.622	0.927	0.063
L = 24, P = 8	30.443	0.945	0.055

1421
1422 View embedding addition. We also evaluate the effectiveness of adding view embeddings to images tokens. In FreeSplatter, we add the multi-view image tokens with a reference view embedding e_{ref} or a source view embedding e_{src} before feeding them into the transformer to make the model identify the reference view, so that it can reconstruct Gaussians in the reference view's camera frame. We present the results in Figure 17 of the revised paper.

1427 Specifically, for 4 input views V_i (i = 1, 2, 3, 4), we try different combinations of \mathbf{e}_{ref} and \mathbf{e}_{src} when 1428 adding them to the image tokens, and then render the reconstructed Gaussians with an "identity 1429 camera pose" C = [[1, 0, 0, 0], [0, 1, 0, 0], [0, 0, 1, 0], [0, 0, 0, 1]] to see the results. We can observe 1430 that, when adding the *j*-th view's tokens with \mathbf{e}_{ref} and other views' tokens with \mathbf{e}_{src} , the rendered 1431 image is exactly the *j*-th view. This means that the model successfully identify the *j*-th view as 1432 the reference view and reconstruct Gaussians in its camera frame. In comparison, all other view 1433 embedding combinations leads to degraded reconstruction.

Qualitative comparison of adding pixel-alignment loss. Apart from the quantitative results in Table 3, we visualize the influence of pixel-alignment loss on sparse-view reconstruction results in Figure 15 to better demonstrate its effectiveness. The results show that the model trained without pixel-alignment loss leads to blurry novel view renderings, while the model trained with pixel-alignment loss demonstrates significantly better visual details.



Figure 15: Ablation on pixel-alignment loss. We show two samples from the GSO dataset.

1455 1456

1457



Figure 16: Illustration on the influence of input view number. We show the visual comparison of FreeSplatter-O results with varying numbers of input views (n = 1 - 6). From left to right: input views, reconstructed Gaussians, and rendered target views at 4 fixed viewpoints. Additional input views increase Gaussian density and improve previously uncovered regions, with diminishing returns beyond n=4 when object coverage becomes sufficient.



