

# SparseSurf: Sparse-View 3D Gaussian Splatting for Surface Reconstruction

Meiying Gu<sup>1</sup>, Jiawei Zhang<sup>1</sup>, Jiahe Li<sup>1</sup>, Xiaohan Yu<sup>2</sup>, Haonan Luo<sup>3</sup>, Jin Zheng<sup>1, 4\*</sup>, Xiao Bai<sup>1\*</sup>

<sup>1</sup>School of Computer Science and Engineering, State Key Laboratory of Complex Critical Software Environment, Jiangxi Research Institute, Beihang University, China

<sup>2</sup>Macquarie University, Australia

<sup>3</sup>School of Computing and Artificial Intelligence, Southwest Jiaotong University, China

<sup>4</sup>State Key Laboratory of Virtual Reality Technology and Systems, Beijing, China

{gumeiying, byzhangjw, lijiahe, jinzheng, baixiao}@buaa.edu.cn, xiaohan.yu@mq.edu.au, lhn@swjtu.edu.cn

## Abstract

Recent advances in optimizing Gaussian Splatting for scene geometry have enabled efficient reconstruction of detailed surfaces from images. However, when input views are sparse, such optimization is prone to overfitting, leading to suboptimal reconstruction quality. Existing approaches address this challenge by employing flattened Gaussian primitives to better fit surface geometry, combined with depth regularization to alleviate geometric ambiguities under limited viewpoints. Nevertheless, the increased anisotropy inherent in flattened Gaussians exacerbates overfitting in sparse-view scenarios, hindering accurate surface fitting and degrading novel view synthesis performance. In this paper, we propose SparseSurf, a method that reconstructs more accurate and detailed surfaces while preserving high-quality novel view rendering. Our key insight is to introduce Stereo Geometry-Texture Alignment, which bridges rendering quality and geometry estimation, thereby jointly enhancing both surface reconstruction and view synthesis. In addition, we present a Pseudo-Feature Enhanced Geometry Consistency that enforces multi-view geometric consistency by incorporating both training and unseen views, effectively mitigating overfitting caused by sparse supervision. Extensive experiments on the DTU, BlendedMVS, and Mip-NeRF360 datasets demonstrate that our method achieves the state-of-the-art performance.

**Project** — <https://miya-oi.github.io/SparseSurf-project>

## Introduction

Surface reconstruction is a long-standing problem, aiming to recover accurate 3D geometry from images. Traditional multi-view stereo (MVS) pipelines (Yao et al. 2018; Ding et al. 2022) establish dense correspondences and fuse depth maps via volumetric methods or Poisson reconstruction. Neural Radiance Field (NeRF) based methods (Wang et al. 2021; Yariv et al. 2021; Fu et al. 2022; Li et al. 2023) build neural signed distance fields and extract high-quality surfaces from the optimized fields, while they require long optimization, hindering real-world applications.

Recently, 3D Gaussian Splatting has been investigated for surface reconstruction (Huang et al. 2024a; Yu, Sattler, and

\*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

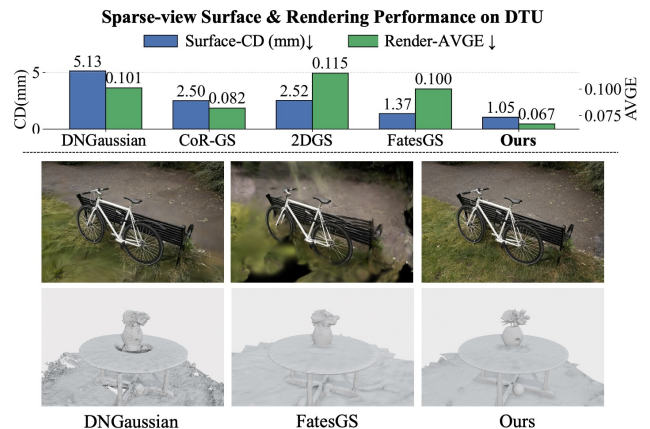


Figure 1: Comparison of sparse-view novel-view synthesis and surface reconstruction on DTU and Mip-NeRF360. Our SparseSurf achieves the best performance on both surface reconstruction and rendering in sparse-view setting.

Geiger 2024; Chen et al. 2024), efficiently estimating detailed surfaces with a sufficient collection of camera views. However, in real-world applications, the acquired views can be very sparse. Since 3DGS relies heavily on multi-view correspondences for optimization, it suffers from severe overfitting when the number of training views is limited (Li et al. 2024), leading to degraded reconstruction quality.

To address the overfitting issues in sparse-view Gaussian Splatting, recent works (Zhu et al. 2024; Li et al. 2024; Zhang et al. 2024) have introduced regularization strategies to improve the geometric quality of 3D Gaussians. However, these methods primarily focus on enhancing novel view synthesis (NVS) performance and thus apply relatively loose geometric constraints, making it challenging to reconstruct accurate surfaces. More recently, FatesGS (Huang et al. 2025) and Sparse2DGS (Wu et al. 2025) represent Gaussians as flattened 2D planes to better fit surface geometry under sparse views, while enforcing multi-view geometric consistency as regularization. Despite these improvements, the increased anisotropy of flattened Gaussians further amplifies the risk of overfitting in sparse-view settings. As shown in Figure 1, FatesGS with flattened Gaussians meets challenges

to render high-fidelity images at unseen views, indicating the flattened Gaussians meet difficulties fitting the accurate surface positions solely with sparse training view supervision, limiting the detailed surface reconstruction ability.

In this paper, we propose SparseSurf, which achieves accurate surface reconstruction while preserving photorealistic novel view synthesis. To compensate for the lack of geometry of sparse views, current solutions employ monocular depth priors as additional constraints (Li et al. 2024; Huang et al. 2025). However, such priors suffer from scale ambiguities and lack confidence estimation, introducing noise into the optimization process. While Gaussians can tolerate noise with sufficient training views, they become significantly susceptible under sparse-view conditions. Prior-induced noise leads to multi-view inconsistency and hinders Gaussians from fitting accurate surfaces. To address this, we introduce Stereo Geometry-Texture Alignment that provides metric supervision, enabling more reliable geometric guidance. These priors are further filtered through geometric consistency to ensure more robust constraints. As training progresses, the quality of stereo renderings improves, which in turn enhances the accuracy of the stereo priors.

Recovering fine-grained geometry often motivates the use of flattened, anisotropic Gaussian primitives, as this increases surface adherence. However, such flattening reduces the implicit regularization provided by anisotropy, thereby exacerbating overfitting under sparse views in Figure 1. Although overfitting is not readily apparent from the training views, its adverse effects become evident on novel views. To address this challenge, we propose Pseudo-Feature Enhanced Geometry Consistency that enforces multi-view geometric consistency by leveraging both sparse training views and pseudo-unseen views. To recover surface details, we enhance the geometry consistency by introducing multi-view feature representations and conducting feature distillation to efficiently render pseudo view feature representations. By jointly enforcing consistency on training and pseudo-unseen views, our approach mitigates overfitting introduced by limited views and the use of flattened Gaussian primitives, resulting in more robust and detailed surface reconstructions.

Our contributions are summarized as follows.

- We propose SparseSurf, investigating to reconstruct accurate and detailed surfaces while simultaneously improving novel-view synthesis quality of Gaussian Splatting under sparse training views.
- We propose Stereo Geometry-Texture Alignment and Pseudo-Feature Enhanced Geometry Consistency to derive metric depth supervision and enhance pseudo-view through multi-view feature consistency, alleviating overfitting issues and improving the surface reconstruction.
- We perform extensive experiments on DTU, Blended-MVS, and Mip-NeRF360 datasets. Our method achieves state-of-the-art in sparse view surface reconstruction.

## Related Work

### Neural Implicit Representation

Neural Radiance Fields (NeRF) (Mildenhall et al. 2021) represents scenes as continuous volume radiance functions,

achieving remarkably photorealistic rendering. However, the heavy MLPs causes high computational cost. Notably, 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) optimizes a set of anisotropic 3D Gaussian primitives with rasterization-based rendering, achieving fast rendering speed and high-quality performance. However, both NeRF and 3DGS focus more on view synthesis than accurate geometry, resulting in messy and noisy surfaces.

### Neural Surface Reconstruction

In obtaining a set of 2D images, neural surface reconstruction methods focus on recovering the accurate 3D geometry for the scene. Built upon NeRF, NeuS (Wang et al. 2021), VolSDF (Yariv et al. 2021), and subsequent optimization methods (Yu et al. 2022; Fu et al. 2022; Darmon et al. 2022; Li et al. 2023) represent surfaces as occupancy or signed distance fields. Despite accurate geometry, these implicit surface methods demand significant computational resources and prolonged optimization—the reconstruction time can reach up to several days per scene.

Recent works try to bring the advantages of 3DGS to surface reconstruction. Sugar (Guédon and Lepetit 2024) introduces regularization for aligning Gaussians with surface, and adopts an optional refinement to bind Gaussians to the mesh surface. 2DGS (Huang et al. 2024a) and Gaussian Surfels (Dai et al. 2024) flatten the original 3D Gaussians into 2D ellipse to recover thin surfaces. To extract continuous surfaces, GOF (Yu, Sattler, and Geiger 2024) introduces an opacity field, while GSDF (Yu et al. 2024), 3DGSR (Lyu et al. 2024), and GS-Pull (Zhang, Liu, and Han 2024) leverage distance fields. To constrain geometry, PGSR (Chen et al. 2024) proposes unbiased depth rendering method. Among them, GS2Mesh (Wolf, Bracha, and Kimmel 2024) is most relevant to our work, as it extracts meshes using a pretrained stereo matching model instead of directly relying on depth maps rendered from the vanilla 3DGS representation. GS2Mesh recovers accurate and complete geometry with dense view inputs. However, under sparse views, the render quality degrades due to overfitting, making it difficult to directly use stereo depth for surface reconstruction.

### Sparse View Gaussian Splatting

3DGS demonstrates significant advantages in both render quality and inferring speed. While given sparse views, many methods tend to reconstruct noisy and unrealistic scenes. Gaussian splatting faces the problems of overfitting and geometric inaccuracy with limited inputs. Recently, sparse-view novel view synthesis methods aims to alleviate overfitting issue for improve the novel view rendering via geometric priors (Li et al. 2024; Zhu et al. 2024; Han et al. 2024; Zheng et al. 2025) or pseudo views (Zhu et al. 2024; Zhang et al. 2024). Although these methods render high-quality appearance of scenes, they are still challenging to recover surfaces. FatesGS (Huang et al. 2025) and concurrent Sparse2DGS (Wu et al. 2025) focus on sparse view surface reconstruction. To fit surfaces, they represent 3D scenes as flattened Gaussians under multi-view depth regularization. However, with limited input views, the enhanced anisotropy of flattened Gaussians increases the overfitting risk.

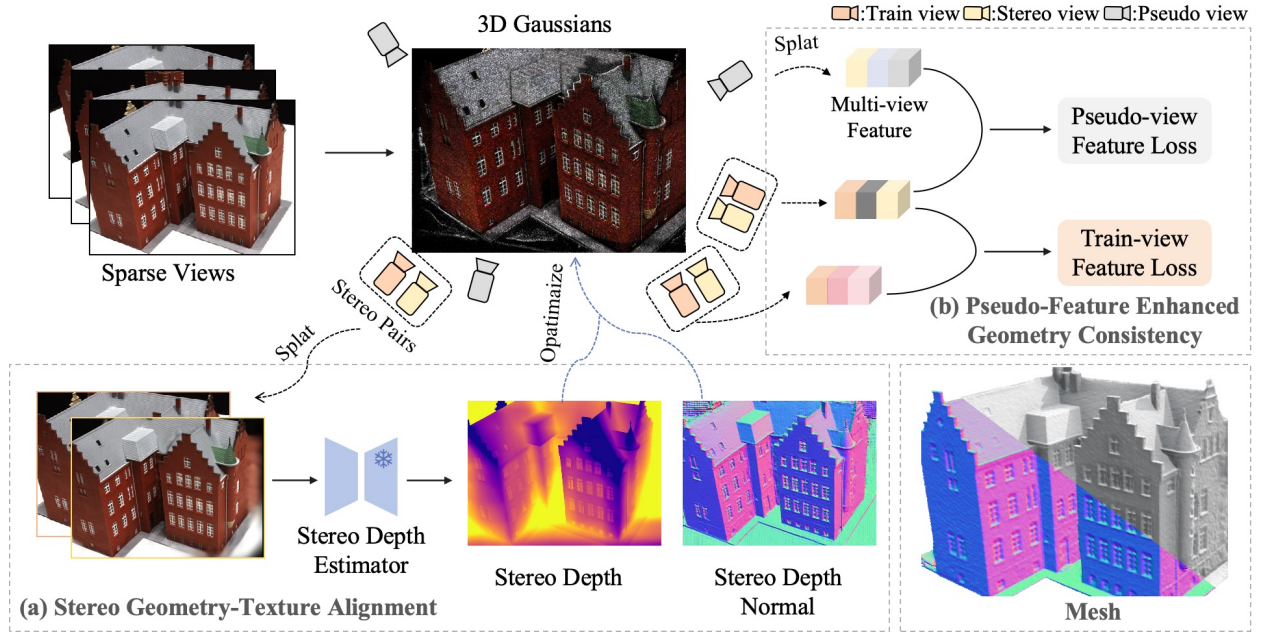


Figure 2: The framework of SparseSurf. (a) Stereo Geometry-Texture Alignment. We estimate and update stereo-view images to generate binocular priors for geometry supervision. (b) Pseudo-Feature Enhanced Geometry Consistency. To mitigate over-fitting and enhance multi-view consistency, we introduce Pseudo-view Feature Consistency and Train-view Feature Alignment.

## Method

Given sparse-view RGB images of a scene, our goal is to achieve accurate and detailed surface reconstruction and high-quality rendering. We propose SparseSurf, as illustrated in Figure 2.

### Preliminary of 3D Gaussian Splatting

**Scene Representation.** 3D Gaussian splatting (Kerbl et al. 2023) models a scene as a set of 3D Gaussians. The  $i$ -th Gaussian is defined by:

$$\mathcal{G}_i(x) = e^{-\frac{1}{2}(x-\mu_i)^T \Sigma_i^{-1}(x-\mu_i)}, \quad (1)$$

where  $\mu \in \mathbb{R}^3$ ,  $\Sigma \in \mathbb{R}^{3 \times 3}$ , and  $x$  presents its center, covariance, 3D position respectively. The covariance matrix  $\Sigma$  can be calculated from the scale  $S$  and rotation  $R$ :

$$\Sigma_i = R_i S_i S_i^T R_i^T \quad (2)$$

To accurately conform to the scene surface, we compress the Gaussian ellipsoid along the direction of the minimum scale factor, following (Chen et al. 2024; Dai et al. 2024).

**Rendering.** For Gaussian rasterization, 3D Gaussians are sorted by depth and renders through alpha-blending. Given a pixel  $u$  from one image, the color  $C \in \mathbb{R}^3$  can be obtained:

$$C = \sum_{i \in N} T_i \alpha_i c_i, \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (3)$$

where  $T_i$  is the transmittance, defined as the cumulation of the opacity values of previous Gaussians overlapping the same pixel, and  $\alpha_i$  is the blending weight.  $c_i$  is the view-dependent color.

Similarly, the rendered normal  $N$  and distance map  $\mathcal{D}$  can be accumulated as alpha-blending by:

$$N = \sum_{i \in N} R_c^T n_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad \mathcal{D} = \sum_{i \in N} d_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (4)$$

where  $R_c$  is the rotation from the camera to the global world, and  $n_i$  is the minimum scale factor.  $d_i = R_c^T (\mu_i - T_c)^T (R_c^T n_i)$  is the distance from the plane to the camera center, where  $T_c$  and  $\mu_i$  are the camera center in the world and the center of gaussian respectively. After obtaining the distance and normal of the plane through rendering, the unbiased depth map  $\mathcal{D}$  can be determined by intersecting rays with the plane (Chen et al. 2024).

### Stereo Geometry-Texture Alignment

Accurate 3D scene representation requires both reliable rendering and precise geometric reconstruction. Leveraging a binocular framework can combine the advantages of rendering and geometry reconstruction simultaneously. To this end, we render corresponding stereo-view images and feed them into pretrained stereo matching networks (Wen et al. 2025; Zhang et al. 2025; Bartolomei et al. 2024) to obtain precise priors supervising the geometry. As the rendering quality improves, the depth prior becomes more accurate.

**Stereo Prior Estimation.** 3DGS exhibits excellent rendering quality, especially for interpolated viewpoints, yet its performance degrades when rendering extrapolated views. In the case of sparse view, the camera distribution is unbalanced. Therefore, we generate the pseudo-stereo view in the direction of its nearest neighbor view. Specifically, for each

given camera pose  $\mathbf{P}_i$ , we can obtain a corresponding stereo-view camera at a horizontal baseline  $b$ . For each train view-point, we render its stereo-view image to form a stereo pair, which is then passed through a pre-trained stereo-matching network (Wen et al. 2025) to obtain a disparity map. Using the known camera baseline  $b$  and focal length, we convert disparity to depth  $\mathcal{D}^*$ . To capture local depth variations, we estimate per-pixel normals  $\mathcal{N}^*$  from the depth map  $\mathcal{D}^*$  by applying image-space depth gradients. With the limited input views, the rendered images often exhibit blurring and aliasing that can corrupt geometric priors if used directly. To mitigate this, we apply a stereo view consistency check to produce reliability masks  $\mathcal{M}^*$  that filter out unreliable pixels before supervision. Throughout training, at regular intervals, we re-render the stereo-view image, and then update the stereo depth  $\mathcal{D}^*$ , normal  $\mathcal{N}^*$ , and valid mask  $\mathcal{M}^*$ .

**Stereo Geometry Supervision.** We enforce geometry by minimizing  $\mathcal{L}_1$  loss between rendered depth  $D$  and depth prior  $\mathcal{D}^*$ :

$$\mathcal{L}_{depth} = \mathcal{L}_1(D, \mathcal{D}^*) \quad (5)$$

To sharpen local surface orientation, we employ a cosine-based loss to align the stereo normal priors  $\mathcal{N}^*$  with both the normals  $N$  rendered by 3DGS and those  $N_d$  computed from the rendered depth.

$$\mathcal{L}_{normal} = 1 - \text{Cosine}(N, \mathcal{N}^*), \mathcal{L}_{nd} = 1 - \text{Cosine}(N_d, \mathcal{N}^*) \quad (6)$$

To transfer the inherent smoothness of our reliable stereo normals onto the rendered normals and normals from depth, we then impose an edge-aware Laplacian smoothness loss, penalizing deviations between those second-order gradients and those of the stereo priors.

$$\mathcal{L}_{smooth} = \text{Smooth}(N, \mathcal{N}^*) + \text{Smooth}(N_d, \mathcal{N}^*) \quad (7)$$

The above losses are evaluated only over pixels that pass consistency checks. The total stereo loss can be defined as :

$$\mathcal{L}_{stereo} = (\lambda_d \mathcal{L}_{depth} + \lambda_n \mathcal{L}_{normal} + \lambda_{nd} \mathcal{L}_{nd}) \mathcal{M}^* + \lambda_s \mathcal{L}_{smooth} \quad (8)$$

### Pseudo-Feature Enhanced Geometry Consistency

With limited camera views, Gaussian Splatting is prone to overfitting training views when supervised solely by ground-truth color. By introducing Stereo Geometry-Texture Alignment constrain the arrangement of Gaussian primitives helps to alleviate this effect. However, the high-quality rendering of vanilla 3DGS relies on anisotropic kernels. To enforce tighter surface adherence, we instead adopt flattened Gaussian kernels, which improve surface quality but sacrifice the innate regularization of anisotropy, bringing extra risk of overfitting in sparse-view supervision. While the overfitting issue are hardly observed from training views, it can be obvious at unseen views due to the high anisotropy of flattened Gaussians. Therefore, we proposed Pseudo-Feature Enhanced Geometry Consistency to regularize the Gaussian geometry, combining both training and unseen views.

**Pseudo-view Feature Consistency.** In sparse-view settings, pseudo-view supervision is regarded as an effective solution to alleviate overfitting because the virtual view enriches the

limited set of observed views. Previous NVS works generate pseudo views based on the positions of training cameras and supervise pseudo views with RGB colors (Zhang et al. 2024) or monocular depth maps (Zhu et al. 2024). However, they are insufficient to regularize the surface geometry of pseudo views due to the lack of multi-view constraints. To this end, we supervise pseudo views with multi-view geometry consistency and enhance the regularization effects on details through multi-view feature representations. While naively extracting features from unseen views at each training loop incurs prohibitive computational cost, we augment each Gaussian primitive with an additional feature attribute to efficiently constrain the feature consistency on unseen views.

Specifically, given a ground-truth image  $\hat{I}$ , we can extract the multi-view feature  $\mathcal{F}^*$  using a frozen feature extraction model. During training, we distill a predicted feature map  $F$  by minimizing a feature-distillation loss:

$$\mathcal{L}_f = 1 - \text{Cosine}(F, \mathcal{F}^*) \quad (9)$$

In this way, the Gaussian primitives learn to distill and reproduce rich multi-view cues directly from the extracted features, and we enforce pseudo-view supervision in a learned feature space.

Specifically, we render feature map  $F_p$  under a random pseudo-view  $\mathcal{V}_p$ . Given a reference training view  $\mathcal{V}_t$  with corresponding render depth  $D_t$  and camera pose  $\mathcal{P}_t$ , the warped feature map  $\mathcal{F}_{p2t}$  from  $\mathcal{V}_p$  to  $\mathcal{V}_t$  can be defined as:

$$\mathcal{F}_{p2t} = \text{Warp}(F_p, D_t, \mathcal{P}_t, \mathcal{P}_p) \quad (10)$$

Similarly, we warp the reference feature map  $\mathcal{F}_r$  into the pseudo-view to obtain  $\mathcal{F}_{t2p}$ . By comparing these bidirectional warps via a round-trip feature discrepancy, we derive a binary confidence mask  $\mathcal{M}_{feat}$  that marks only those pixels whose feature reprojection error falls below a threshold.

However, because pseudo-views can span wide baselines, the rendered feature maps may contain low-fidelity regions, if enforced at the pixel level, could corrupt the train-view features. To guard against this, we impose our feature-consistency loss at the patch level. Concretely, we tile both the warped feature map  $\mathcal{F}_{p2t}$  and the reference map  $\mathcal{F}_{t2p}$  into local patches, aggregate each patch into a single descriptor via average pooling, and compute a patch-wise cosine similarity. We then weight each patch loss by its geometric confidence  $M_{feat}$  before summing:

$$\mathcal{L}_{pseudo} = \sum_{i,j} \mathcal{M}_{feat}^{(i,j)} [1 - \text{Cosine}(\bar{\mathcal{F}}_{p2t}^{(i,j)}, \bar{\mathcal{F}}_r^{(i,j)})], \quad (11)$$

where  $\bar{\mathcal{F}}^{(i,j)}$  denotes the pooled descriptor of  $(i, j)$ th patch.

**Train-view Feature Alignment.** Building on our pseudo-view feature supervision, we further leverage the high-confidence features extracted from ground-truth training views to enforce multi-view consistency at the pixel level. Concretely, let  $\mathcal{V}_s$  be one of the source views with feature  $\mathcal{F}_s$  and pose  $\mathcal{P}_s$ . We perform a one-way warp of the source features into the current training frame  $\mathcal{V}_t$  to obtain  $\mathcal{F}_{s2t}$ . We then apply a pixel-wise cosine loss to align the warped and reference features:

$$\mathcal{L}_{train} = 1 - \text{Cosine}(\mathcal{F}_{s2t}, \mathcal{F}_s) \quad (12)$$

Scan ID	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
COLMAP	2.88	3.47	1.74	2.16	2.63	3.27	2.78	3.63	3.24	3.49	2.46	1.24	1.59	2.72	1.87	2.61
SparseNeuS <sub>ft</sub>	4.81	5.56	5.81	2.68	3.30	3.88	2.39	2.91	3.08	2.33	2.64	3.12	1.74	3.55	2.31	3.34
VolRecon	3.05	4.45	3.36	3.09	2.78	3.68	3.01	2.87	3.07	2.55	3.07	2.77	1.59	3.44	2.51	3.02
UFORecon	1.52	2.58	1.85	1.44	1.55	1.81	1.06	1.52	0.96	1.40	1.19	0.94	0.65	1.25	1.29	1.40
NeuS	4.11	5.40	5.10	3.47	2.68	2.01	4.52	8.59	5.09	9.42	2.20	4.84	0.49	2.04	4.20	4.28
VolSDF	4.07	4.87	3.75	2.61	5.37	4.97	6.88	3.33	5.57	2.34	3.15	5.07	1.20	5.28	5.41	4.26
MonoSDF	3.47	3.61	2.10	1.05	2.37	1.38	1.41	1.85	1.74	1.10	1.46	2.28	1.25	1.44	1.45	1.86
NeuSurf	1.35	3.25	2.50	0.80	1.21	2.35	0.77	1.19	1.20	1.05	1.05	1.21	0.41	0.80	1.08	1.35
DNGaussian	3.40	5.58	3.28	4.62	3.18	2.70	6.19	4.84	9.04	3.29	8.10	11.28	3.90	4.86	2.69	5.13
CoR-GS	2.80	3.06	2.42	1.88	2.37	2.76	2.00	5.66	2.51	1.92	2.54	2.20	1.36	2.28	1.74	2.50
Binocular3DGS	2.13	3.33	2.76	1.33	3.55	2.66	2.48	3.79	2.30	2.74	2.32	5.79	1.53	2.03	1.98	2.71
2DGS	3.25	3.64	3.52	1.42	2.04	2.52	1.99	2.69	2.55	1.79	2.92	4.50	0.73	2.38	1.79	2.52
PGSR	5.83	4.59	4.52	3.36	4.25	3.75	2.81	5.92	4.60	4.27	3.61	6.09	1.02	2.55	2.32	3.97
FatesGS	1.32	2.85	2.71	0.80	1.44	2.08	1.11	1.19	1.33	0.76	1.49	0.85	0.47	1.05	1.06	1.37
<b>Ours</b>	1.35	2.10	1.84	0.66	0.79	1.62	0.62	1.13	1.06	0.76	0.90	0.87	0.39	0.70	0.96	1.05

Table 1: The quantitative comparisons of CD $\downarrow$  on DTU dataset (little-overlap setting). Best results are highlighted as 1st, 2nd and 3rd.

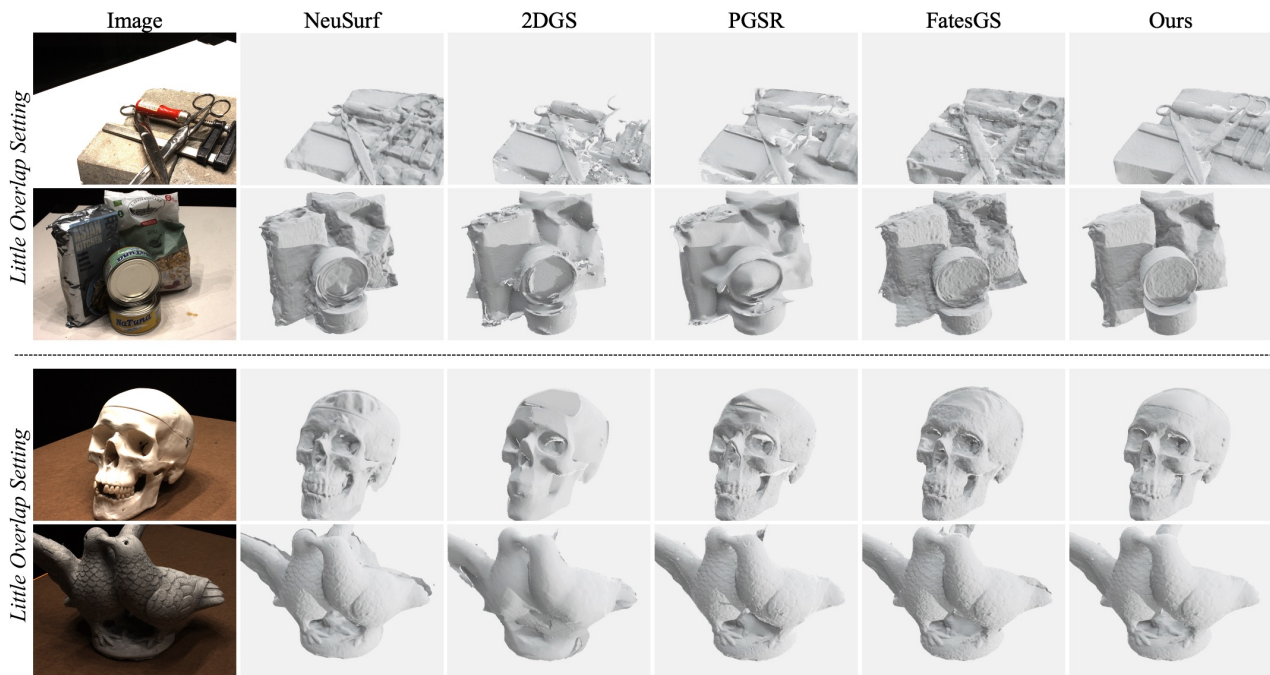


Figure 3: Qualitative comparison of reconstruction results on the DTU with little-overlap and large-overlap sparse setting.

## Experiments

### Setup

**Datasets.** We conduct our experiment on three datasets, DTU (Jensen et al. 2014), BlendedMVS (Yao et al. 2020), and MipNeRF360 (Barron et al. 2022). The DTU dataset comprises 128 scenes, each with 49 or 69 images. For surface reconstruction, we select 15 scenes to train and evaluate our model on 3 views of both the little-overlap setting (PixelNeRF) and the large-overlap (SparseNeuS) setting, following the previous work (Huang et al. 2024b, 2025).

We also show the qualitative reconstruction comparisons on the BlendedMVS and MipNeRF360 datasets. For rendering comparison, we follow the sparse NVS setting on DTU (3 views) and Mip-NeRF360 (24 views) used in previous works (Li et al. 2024; Zhu et al. 2024; Zhang et al. 2024; Han et al. 2024; Zheng et al. 2025). Consistent with previous sparse-view settings, the camera poses are assumed to be known.

**Evaluation Metrics.** To evaluate the accuracy of the reconstructed meshes, we use Chamfer Distance (CD) on DTU, using the official evaluation script. To evaluate the rendering quality, we use PSNR, SSIM, LPIPS, and AVGE scores.

Scan ID	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
COLMAP	0.90	2.89	1.63	1.08	2.18	1.94	1.61	1.30	2.34	1.28	1.10	1.42	0.76	1.17	1.14	1.52
TransMVSNet	1.07	3.14	2.39	1.30	1.35	1.61	0.73	1.60	1.15	0.94	1.34	0.46	0.60	1.20	1.46	1.35
SparseNeuS <sub>ft</sub>	1.29	2.27	1.57	0.88	1.61	1.86	1.06	1.27	1.42	1.07	0.99	0.87	0.54	1.15	1.18	1.27
VolRecon	1.20	2.59	1.56	1.08	1.43	1.92	1.11	1.48	1.42	1.05	1.19	1.38	0.74	1.23	1.27	1.38
ReTR	1.05	2.31	1.44	0.98	1.18	1.52	0.88	1.35	1.30	0.87	1.07	0.77	0.59	1.05	1.12	1.17
C2F2NeuS	1.12	2.42	1.40	0.75	1.41	1.77	0.85	1.16	1.26	0.76	0.91	0.60	0.46	0.88	0.92	1.11
GenS <sub>ft</sub>	0.91	2.33	1.46	0.75	1.02	1.58	0.74	1.16	1.05	0.77	0.88	0.56	0.49	0.78	0.93	1.03
UFORcon	0.76	2.05	1.31	0.82	1.12	1.18	0.74	1.17	1.11	0.71	0.88	0.58	0.54	0.86	0.99	0.99
NeuS	4.57	4.49	3.97	4.32	4.63	1.95	4.68	3.83	4.15	2.50	1.52	6.47	1.26	5.57	6.11	4.00
VolSDF	4.03	4.21	6.12	0.91	8.24	1.73	2.74	1.82	5.14	3.09	2.08	4.81	0.60	3.51	2.18	3.41
MonoSDF	2.85	3.91	2.26	1.22	3.37	1.95	1.95	5.53	5.77	1.10	5.99	2.28	0.65	2.65	2.44	2.93
NeuSurf	0.78	2.35	1.55	0.75	1.04	1.68	0.60	1.14	0.98	0.70	0.74	0.49	0.39	0.75	0.86	0.99
3DGS	3.38	4.19	2.99	1.76	3.38	3.80	5.21	2.91	4.29	3.18	3.23	5.18	2.78	3.48	3.32	3.54
Gaussian Surfels	3.56	5.42	3.95	3.68	4.61	2.72	4.42	5.22	4.71	3.46	4.07	5.42	2.44	3.27	4.00	4.06
2DGS	1.26	2.95	1.73	0.96	1.68	1.97	1.58	1.87	2.50	1.02	1.93	1.91	0.72	1.85	1.37	1.69
PGSR	1.20	3.09	1.97	1.36	4.99	2.13	1.18	2.43	2.17	1.81	1.36	0.83	0.43	1.37	0.96	1.82
FatesGS	0.67	1.94	1.17	0.77	1.28	1.23	0.63	1.05	0.98	0.69	0.75	0.48	0.41	0.78	0.90	0.92
Sparse2DGS	1.05	2.35	1.38	0.83	1.37	1.45	0.84	1.16	1.43	0.74	0.85	0.84	0.57	0.95	1.01	1.13
<b>Ours</b>	<b>0.72</b>	<b>1.61</b>	<b>1.17</b>	<b>0.72</b>	<b>1.11</b>	<b>1.13</b>	<b>0.57</b>	1.30	1.20	<b>0.67</b>	<b>0.70</b>	<b>0.45</b>	<b>0.36</b>	<b>0.70</b>	<b>0.87</b>	<b>0.89</b>

Table 2: The quantitative comparisons of CD $\downarrow$  on DTU dataset (large-overlap setting). Best results are highlighted as 1st, 2nd and 3rd.




Figure 4: Qualitative rendering comparison on DTU with sparse-view NVS setting.




Figure 5: Comparison on BlendedMVS.

**Implementations.** Following previous research, we use COLMAP (Schonberger and Frahm 2016) for Gaussians initialization with additional multi-view features in 8 dimensions. We render stereo-view images and feed them into the advanced stereo matching network (Wen et al. 2025) to predict geometric priors from 500 iterations and update every 300 iterations. Following previous studies (Huang et al. 2024b, 2025), We adopt Vis-MVSNet (Zhang et al. 2020) to extract feature representations.

## Comparisons

**Sparse View Surface Reconstruction.** Quantitative results for geometry reconstruction under sparse input views on both DTU sparse-view settings are reported in Table 2. Our method achieves the best average Chamfer Distance (CD), with a score of 1.05 in the little-overlap setting and 0.89 in the large-overlap setting, outperforming all baselines, including previous state-of-the-art approaches. As shown in Figure 3, our reconstructed meshes exhibit superior geome-

try in terms of accuracy, completeness, and fine-level detail. Compared with the neural implicit method NeuSurf, our approach produces sharper and more clearly defined edges.

Qualitative results on the BlendedMVS and Mip-NeRF360 datasets are presented in Figure 5 and Figure 6, respectively. Our method demonstrates higher reconstruction accuracy and fidelity on complex real-world scenes.

**Sparse View Novel View Synthesis.** We evaluate our method on DTU with the sparse-view NVS setting for the validation of rendering quality. As shown in Table 3, compared to the current state-of-the-art sparse-view methods, our approach achieves the best performance across PSNR, SSIM, LPIPS, and AVGE on DTU. Figure 4 illustrates the visual comparison. With the correct geometry and overfit relief strategy, our method has fewer artifacts and less geometric misalignment.




Figure 6: Visual comparison of reconstruction results on Mip-NeRF360 under 24 input views.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVGE $\downarrow$
Mip-NeRF	8.68	0.571	0.353	0.315
DietNeRF	11.85	0.633	0.314	0.232
RegNeRF	18.89	0.745	0.190	0.107
FreeNeRF	19.92	0.787	0.182	0.095
SparseNeRF	19.55	0.769	0.201	0.102
DS-NeRF	16.29	0.559	0.451	0.192
ViP-NeRF	10.20	0.301	0.363	0.307
SimpleNeRF	11.94	0.387	0.286	0.243
3DGS	17.65	0.816	0.146	0.102
DNGaussian	18.91	0.790	0.176	0.101
FSGS	17.14	0.818	0.162	0.110
CoR-GS	19.21	0.853	0.119	0.082
Binocular3DGS	20.71	0.862	0.111	-
NexusGS	20.21	0.869	0.102	0.071
FatesGS	17.95	0.833	0.129	0.100
<b>Ours</b>	<b>21.31</b>	<b>0.886</b>	<b>0.089</b>	<b>0.067</b>

Table 3: Quantitative evaluations of render quality on the DTU dataset(Sparse-view novel view synthesis setting).

$L_{stereo}$	$L_{pseudo}$	$L_{train}$	Accuracy $\downarrow$	Completion $\downarrow$	Average $\downarrow$
$\times$	$\times$	$\times$	1.318	2.302	1.810
$\checkmark$	$\times$	$\times$	0.822	1.612	1.217
$\checkmark$	$\checkmark$	$\times$	0.610	1.327	0.969
$\checkmark$	$\checkmark$	$\checkmark$	<b>0.533</b>	<b>1.239</b>	<b>0.886</b>

Table 4: Ablation Study for the large-overlap setting on the DTU dataset.

### Ablation Study

We conduct ablation experiments on DTU dataset with large-overlap setting. To demonstrate the effectiveness of the stereo loss  $L_{stereo}$ , pseudo-view feature loss  $L_{pseudo}$ , and train-view feature loss  $L_{train}$ , we isolate these modules and measure their impact in Table 4 and Figure 7. Additional results and analysis (e.g., efficiency) are presented in the supplementary materials.

**Stereo Geometry-Texture Alignment.** Compared to the baseline method, incorporating binocular depth supervision significantly improves the surface quality from an average distance 1.810 to 1.217. From the visualization compar-




Figure 7: Visualization results for ablating each component of our method.

isons, we observe that it reduces lots of noisy surfaces.

**Pseudo-view Feature Consistency.** Although stereo geometry-texture alignment provides geometric guidance, it only regularizes Gaussians in the sparse training views and lacks the multi-view consistency, therefore, there are still changes to encounter overfitting issues. Adding the pseudo-view feature consistency further reduces the distances from 1.217 to 0.969. From the visualization results, the erroneous reconstructions are reduced, and more surface details are reconstructed.

**Train-view Feature Alignment.** With sparse views, the pseudo-view rendering may introduce additional noise due to the inherent lack of observation. Therefore, we further add feature alignment between training views to improve the robustness of the model to pseudo-view noise and further improve the surface details from the distance of 0.969 to 0.886.

### Conclusion

This paper proposes SparseSurf, which enables accurate surface reconstruction from sparse views while simultaneously enhancing novel view synthesis. SparseSurf leverages Stereo Geometry-Texture Alignment to derive depth priors from rendered stereo pairs. We also introduce Pseudo-Feature Enhanced Geometry Consistency, which improves pseudo-view consistency through multi-view feature alignment. Extensive experiments on DTU, BlendedMVS, and Mip-NeRF360 demonstrate that SparseSurf achieves state-of-the-art results in sparse-view surface reconstruction.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No.62276016, 62372029, 62306247), Key R&D Program of Jiangxi Province, China (20232BBE50019), the Natural Science Foundation of Sichuan Province (2024NSFSC1474, 2024ZHCG0166).

## References

- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5470–5479.
- Bartolomei, L.; Tosi, F.; Poggi, M.; and Mattocchia, S. 2024. Stereo Anywhere: Robust Zero-Shot Deep Stereo Matching Even Where Either Stereo or Mono Fail. *arXiv preprint arXiv:2412.04472*.
- Chen, D.; Li, H.; Ye, W.; Wang, Y.; Xie, W.; Zhai, S.; Wang, N.; Liu, H.; Bao, H.; and Zhang, G. 2024. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*.
- Dai, P.; Xu, J.; Xie, W.; Liu, X.; Wang, H.; and Xu, W. 2024. High-quality Surface Reconstruction using Gaussian Surfels. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery.
- Darmon, F.; Bascle, B.; Devaux, J.-C.; Monasse, P.; and Aubry, M. 2022. Improving neural implicit surfaces geometry with patch warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6260–6269.
- Ding, Y.; Yuan, W.; Zhu, Q.; Zhang, H.; Liu, X.; Wang, Y.; and Liu, X. 2022. TransMVSNet: Global context-aware multi-view stereo network with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8585–8594.
- Fu, Q.; Xu, Q.; Ong, Y. S.; and Tao, W. 2022. GeoNeus: Geometry-Consistent Neural Implicit Surfaces Learning for Multi-view Reconstruction. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 3403–3416. Curran Associates, Inc.
- Guédon, A.; and Lepetit, V. 2024. SuGaR: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Han, L.; Zhou, J.; Liu, Y.-S.; and Han, Z. 2024. Binocular-Guided 3D Gaussian Splatting with View Consistency for Sparse View Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024a. 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery.
- Huang, H.; Wu, Y.; Deng, C.; Gao, G.; Gu, M.; and Liu, Y.-S. 2025. FatesGS: Fast and Accurate Sparse-View Surface Reconstruction Using Gaussian Splatting with Depth-Feature Consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Huang, H.; Wu, Y.; Zhou, J.; Gao, G.; Gu, M.; and Liu, Y.-S. 2024b. NeuSurf: On-Surface Priors for Neural Surface Reconstruction from Sparse Input Views. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2312–2320.
- Jensen, R.; Dahl, A.; Vogiatzis, G.; Tola, E.; and Aanaes, H. 2014. Large scale multi-view stereopsis evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 406–413.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4): 1–14.
- Li, J.; Zhang, J.; Bai, X.; Zheng, J.; Ning, X.; Zhou, J.; and Gu, L. 2024. DNGaussian: Optimizing Sparse-View 3D Gaussian Radiance Fields with Global-Local Depth Normalization. In *CVPR*.
- Li, Z.; Müller, T.; Evans, A.; Taylor, R. H.; Unberath, M.; Liu, M.-Y.; and Lin, C.-H. 2023. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8456–8465.
- Lyu, X.; Sun, Y.-T.; Huang, Y.-H.; Wu, X.; Yang, Z.; Chen, Y.; Pang, J.; and Qi, X. 2024. 3dgsr: Implicit surface reconstruction with 3d gaussian splatting. *ACM Transactions on Graphics (TOG)*, 43(6): 1–12.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- 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.
- 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.
- Wen, B.; Trepte, M.; Aribido, J.; Kautz, J.; Gallo, O.; and Birchfield, S. 2025. FoundationStereo: Zero-Shot Stereo Matching. *CVPR*.
- Wolf, Y.; Bracha, A.; and Kimmel, R. 2024. GS2Mesh: Surface Reconstruction from Gaussian Splatting via Novel Stereo Views. In *European Conference on Computer Vision*.
- Wu, J.; Li, R.; Zhu, Y.; Guo, R.; Sun, J.; and Zhang, Y. 2025. Sparse2DGS: Geometry-Prioritized Gaussian Splatting for Surface Reconstruction from Sparse Views. *arXiv preprint arXiv:2504.20378*.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, 767–783.

Yao, Y.; Luo, Z.; Li, S.; Zhang, J.; Ren, Y.; Zhou, L.; Fang, T.; and Quan, L. 2020. BlendedMVS: A Large-scale Dataset for Generalized Multi-view Stereo Networks. *Computer Vision and Pattern Recognition (CVPR)*.

Yariv, L.; Gu, J.; Kasten, Y.; and Lipman, Y. 2021. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34: 4805–4815.

Yu, M.; Lu, T.; Xu, L.; Jiang, L.; Xiangli, Y.; and Dai, B. 2024. Gsdf: 3dgs meets sdf for improved neural rendering and reconstruction. *Advances in Neural Information Processing Systems*, 37: 129507–129530.

Yu, Z.; Peng, S.; Niemeyer, M.; Sattler, T.; and Geiger, A. 2022. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35: 25018–25032.

Yu, Z.; Sattler, T.; and Geiger, A. 2024. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics (TOG)*, 43(6): 1–13.

Zhang, J.; Li, J.; Huang, L.; Luo, H.; Yu, X.; Gu, L.; Zheng, J.; and Bai, X. 2025. Investigating Synthetic-to-Real Transfer Robustness for Stereo Matching and Optical Flow Estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zhang, J.; Li, J.; Yu, X.; Huang, L.; Gu, L.; Zheng, J.; and Bai, X. 2024. CoR-GS: Sparse-View 3D Gaussian Splatting via Co-Regularization. In *ECCV*.

Zhang, J.; Yao, Y.; Li, S.; Luo, Z.; and Fang, T. 2020. Visibility-aware Multi-view Stereo Network. *British Machine Vision Conference (BMVC)*.

Zhang, W.; Liu, Y.-S.; and Han, Z. 2024. Neural Signed Distance Function Inference through Splatting 3D Gaussians Pulled on Zero-Level Set. In *Advances in Neural Information Processing Systems*.

Zheng, Y.; Jiang, Z.; He, S.; Sun, Y.; Dong, J.; Zhang, H.; and Du, Y. 2025. NexusGS: Sparse View Synthesis with Epipolar Depth Priors in 3D Gaussian Splatting. *arXiv preprint arXiv:2503.18794*.

Zhu, Z.; Fan, Z.; Jiang, Y.; and Wang, Z. 2024. FSGS: Real-Time Few-shot View Synthesis using Gaussian Splatting. In *ECCV*.