3D Gaussian Flats: Hybrid 2D/3D Photometric Scene Reconstruction

Maria Taktasheva
Simon Fraser University
maria_taktasheva@sfu.ca

Lily Goli*
University of Toronto

Alessandro Fiorini
University of Bologna

Zhen Li
Simon Fraser University

Daniel Rebain
University of British Columbia

Andrea Tagliasacchi*
Simon Fraser University
University of Toronto

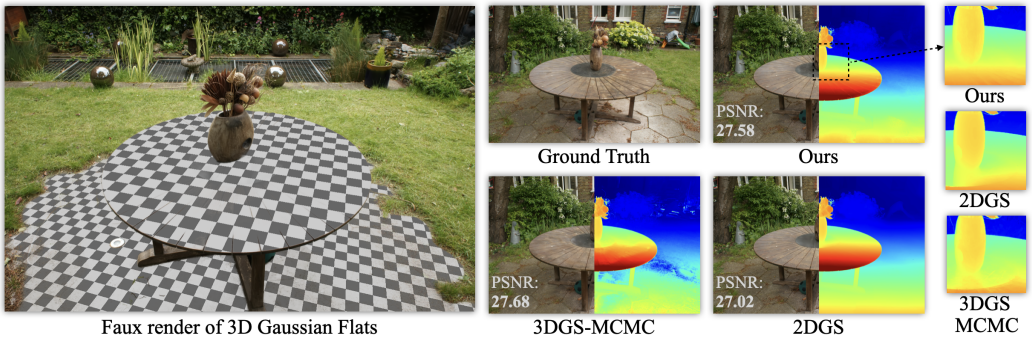

Figure 1: **Teaser** – We introduce 3D Gaussian Flats, a hybrid representation of 2D Gaussians on semantically distinct planar surfaces and 3D Gaussians elsewhere (left). Our method achieves a photorealistic quality on par with fully 3D approaches, while improving geometry over surface reconstruction methods (right) e.g. no visible hole in the middle of the 'garden' scene from Mip-NeRF360 [1].

Abstract

Recent advances in radiance fields and novel view synthesis enable creation of realistic digital twins from photographs. However, current methods struggle with flat, texture-less surfaces, creating uneven and semi-transparent reconstructions, due to an ill-conditioned photometric reconstruction objective. Surface reconstruction methods solve this issue but sacrifice visual quality. We propose a novel hybrid 2D/3D representation that jointly optimizes constrained planar (2D) Gaussians for modeling flat surfaces and freeform (3D) Gaussians for the rest of the scene. Our end-to-end approach dynamically detects and refines planar regions, improving both visual fidelity and geometric accuracy. It achieves state-of-the-art depth estimation on ScanNet++ and ScanNetv2, and excels at mesh extraction without overfitting to a specific camera model, showing its effectiveness in producing high-quality reconstruction of indoor scenes.

*Equal Advising

39th Conference on Neural Information Processing Systems (NeurIPS 2025).

1 Introduction

Recent advances in radiance fields and novel view synthesis have enabled the creation of realistic digital twins from collections of real-world photographs [2, 3]. These techniques allow for high-fidelity 3D reconstructions that capture intricate details of real-world scenes, making them invaluable for applications in virtual reality, gaming, cultural heritage preservation, and scientific visualization.

However, when optimizing for novel view synthesis on flat and texture-less surfaces (e.g. walls, ceilings, tables that are prevalent in indoor scenes), current methods struggle in producing a faithful 3D reconstruction as the problem is photometrically under-constrained [4]. Specifically, modern novel view synthesis frameworks like [5, 6], which are optimized via volume rendering, model flat surfaces with low densities, resulting in non-opaque representations of solid surfaces; see the surface of the table in Figure 1 as an example. Conversely, surface reconstruction methods that assume solid, flat surfaces avoid this limitation [7]. However, they compromise visual quality in favor of a more parsimonious 3D reconstruction; see figure 1. Our core research question is whether these seemingly conflicting objectives could be achieved simultaneously.

Some approaches have attempted to answer this questions by first creating a full 3D representation, and then – *post-training* – detecting planar surfaces to enable 3D planar reconstruction [8, 9]. However, these methods do not leverage planar assumptions during the optimization of the scene representation itself, limiting their effectiveness. Others enforce planar assumptions during training through various regularizer losses [10]. However, these losses can be hard to tune, as they are only suitable for the portion of the scene that is solid and flat, hindering the reconstruction whenever these assumptions are violated.

In contrast to these methods, we propose to look at the problem in an *end-to-end* fashion, conjoining the process of photometric to the one of planar surface reconstruction. To achieve this, we introduce a *hybrid* 2D/3D representation, where flat surfaces are modeled with 2D Gaussian splats [7] that are confined to 2D planes, while the remaining of the scene is modeled with a classical, and more expressive, 3DGS model [6]. By *jointly* optimizing planar (2D) and freeform (3D) Gaussians, our approach enables better fitting of the final representation to planar surfaces within the scene. During photometric optimization, our method dynamically detects planar regions, and adaptively grows their extent, resulting in reconstruction that *retains* high visual quality (as measured by PSNR) compared to a classical 3DGS scene, while simultaneously achieving superior geometric accuracy (as measured by depth error).

Our evaluations demonstrate that this hybrid representation achieves state-of-the-art depth estimation results on challenging indoor datasets including the new ScanNet++ dataset which was designed for dense reconstruction tasks using NeRF-based approaches, and the legacy ScanNetv2 dataset with sparser camera views. Our method delivers crisp reconstructed surfaces, while maintaining competitive visual quality compared to fully 3D representations. Beyond novel view synthesis, our approach has application in mesh extraction for planar surfaces, producing high-quality meshes and accurate mesh segmentation results across diverse capture setups (DSLR and iPhone captures), without the overfitting issues that negatively affect previous methods trained on specific camera models.

2 Related Work

Modern neural scene reconstruction methods aim to generate high-quality 3D representations from 2D images for applications like novel view synthesis [5, 6]. Despite significant progress, volumetric approaches struggle to accurately reconstruct planar surfaces [11], while surface reconstruction methods fail to recover volumetric effects [12]. Finding an approach that accurately reconstructs planar geometry without compromising the quality of the surrounding scene geometry and appearance is a key challenge.

Representations for differentiable rendering Neural Radiance Field (NeRF) [5] pioneered scene reconstruction with a 3D neural representation optimized through differentiable volumetric rendering. 3D Gaussian Splatting (3DGS) [6] overcame NeRF slow training/rendering speed by representing scenes as efficiently rasterizable 3D Gaussians, dramatically accelerating rendering while maintaining quality. The impressive speed-quality balance of 3DGS quickly established it as a standard approach,

with recent advancements such as 3DGS-MCMC [13] further enhancing its accessibility by eliminating the dependency on SfM initialization. Despite these innovations, volumetric representations still struggle with clean geometry reconstruction in flat and textureless surfaces common in indoor environments, hindering applications like mesh extraction. Our method addresses these challenges through a hybrid 2D/3D Gaussian representation that achieves superior geometric reconstruction while preserving rendering quality.

Surface representations and planar constraints While NeRF [5] and 3DGS [6] employ fully volumetric representations, alternative approaches such as [11, 14] model scenes as solid surfaces. This philosophy inspired SuGaR [15], to use a regularization term that encourages the Gaussians to align with the surface of the scene, and later 2DGS [7], which uses 2D Gaussian primitives to reconstruct surfaces outperforming prior surface reconstruction methods [11, 14, 15]. Recent work [16] uses 2D Gaussians as in 2DGS, with multi-view depth and normal regularization to improve surface quality, while RaDe-GS [17] enables depth and normal rasterization for 3D Gaussians to support similar regularization. Other works introduced more explicit primitives, including planes [18, 19], optimizable geometry through learnable opacity maps [20], and soup of planes for dynamic reconstruction [21]. While these methods excel at representing flat surfaces with clean geometry, they typically sacrifice rendering quality and struggle to model phenomena that are better explained by volumetric effects, rather than surfaces. Some methods enforce planar constraints only as regularization losses, such as Guo et al. [22] that uses Manhattan world assumptions on semantically segmented regions and Chen et al. [23] that enforces plane normal consistency in textureless regions. Although helpful, regularizers can be difficult to tune. Our approach instead explicitly detects and optimizes planes within scene reconstruction, avoiding such issues.

3D plane detection and reconstruction Another research direction *detects* planar surfaces in an initial 3D reconstruction and fits planes only to detected regions, extending single-image plane detection [24, 25] to multi-view settings. PlanarNeRF [26] adds a plane-predicting MLP branch to NeRF, supervised via ground truth labels or plane detection consistency across frames, but prevents plane MLP gradients from affecting the geometry prediction branch. PlanarRecon [8] reconstructs a sparse feature volume, which is decoded into plane features and clustered. AirPlanes [9] and NeuralPlane [27] build 3D-consistent plane embeddings per 3D point, emphasizing semantic priors for accurate detection. While we also use semantic knowledge, our method jointly detects and optimizes planes alongside scene reconstruction, allowing geometry to benefit from planar constraints. Further, unlike these methods, our approach yields full scene reconstructions suitable for novel view synthesis, vs. a coarse surface reconstruction.

Hybrid representations Recent hybrid 2D-3D approaches have explored planar surface representation. Kim and Lim [28] integrate meshes into 3DGS for indoor scenes, using SAM [29] to detect planar surfaces and represent them with meshes while employing 3D Gaussians for other objects. Zanjani et al. [30] combine SAM segmentation with normal estimation to lift 2D plane descriptors to 3D, clustering the planar Gaussians using a tree structure. In contrast, our method offers a simpler solution by representing the scene with a mixture of 2D and 3D Gaussians. This design remains fully compatible with the 3DGS rendering pipeline, eliminating the need for complex hybrid mesh handling, or hierarchical tree structures.

3 Method

Given N posed images $\{I_c\}$ and M planar surfaces $\{P_p\}$, each specified by binary image masks $\{\mathcal{M}_{p,c}\}$, we aim to reconstruct a *hybrid* novel view synthesis method that combines a classical 3DGS model with a 2D piecewise planar representation of the scene. Our goal is to reconstruct the scene so that the planar surfaces are accurately recovered and compactly represented by 2D Gaussian primitives, while the rest of the scene is modeled with 3D Gaussians, with the key objective of avoiding the *artifacts* that can typically be seen when using 3D primitives to model planar surfaces; see Figure 1.

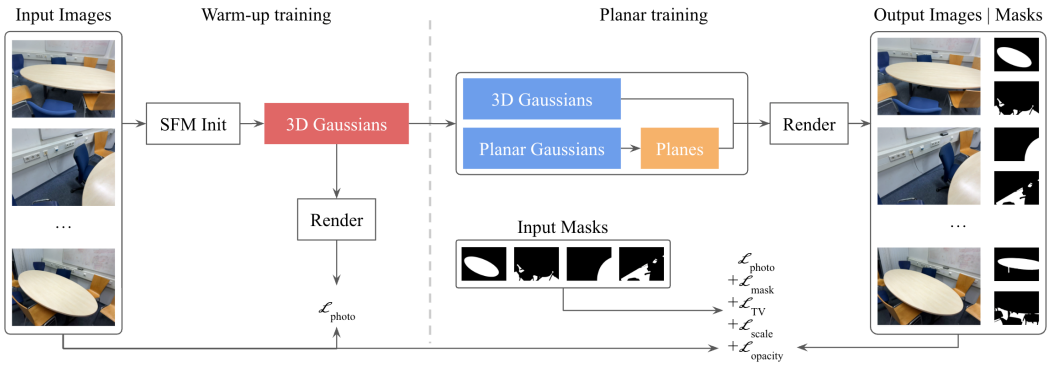

Figure 2: **Overview** – Training of our model is split into two parts: warm-up, in which 3D Gaussians are trained as in [6] using a photometric loss; and planar training, in which 3D Gaussians and planar Gaussians are trained along with the parameters of the planes to which planar Gaussians are locked. Planar training is performed in alternating phases, with Gaussian parameters frozen while plane parameters are optimized, and vice versa. Legend: learnable (warm up), learnable (Gaussian phase), learnable (plane phase).

3.1 Hybrid representation

Our representation consists of M planes $\mathcal{P}=\{P_p\}$, each characterized by its 3D origin and normal $(\mathbf{o}_p, \mathbf{n}_p)$. The geometry of each each plane P_p is represented through a set of 2D Gaussians $\mathcal{G}=\{\mathbf{g}_k\}_{k=1}^{K_k}$ such that,

$$\mathbf{g}_k = \mathcal{N}(\mu_k, \Sigma_k), \quad \mu_k \in P_k, \quad \Sigma_k \in \mathbb{R}^{2\times 2}. \tag{1}$$

Here, μ_k is the center of the k-th Gaussian on the plane P_p, and Σ_k is the 2D covariance matrix, parametrized with a 2D *in-plane* rotation $\mathbf{R}_k$ and a 2D diagonal scale matrix $\mathbf{S}_k$. The plane-to-world transformation matrix is defined as $\mathbf{T}_{\text{pw}}=\text{hom}(\mathbf{R}, \mathbf{o})$, where $\mathbf{R}$ is any rotation matrix that satisfies $\hat{z}=\mathbf{R}\mathbf{n}$ with $\hat{z}$ being the unit vector along the z-axis in the world frame. Thus, the degrees of freedom of planar Gaussians can be mapped to world coordinates through the rigid transformation:

$$\bar{\mu}_k = \mathbf{T}_{\text{pw}}[\mu_k; 0; 1], \quad \bar{\Sigma}_k = \mathbf{T}_{\text{pw}}\ \text{diag}(\Sigma_k, 1, 1)\ \mathbf{T}_{\text{pw}}^{\top} \tag{2}$$

yielding a standard 3D Gaussian primitive representation suitable for rendering. The remaining scene geometry is represented by unconstrained 3D Gaussians $\bar{\mathcal{G}}=\{\bar{\mathbf{g}}_k\}_{k=1}^{\bar{K}}$:

$$\bar{\mathbf{g}}_k = \mathcal{N}(\bar{\mu}_k, \bar{\Sigma}_k), \quad \mu_k \in \mathbb{R}^3, \quad \Sigma_k \in \mathbb{R}^{3\times 3} \tag{3}$$

All Gaussians have view-dependent colors $\mathbf{c}$ represented as Spherical Harmonics, and opacity α as in vanilla 3DGS. To reconstruct the scene with our hybrid representation, we need to optimize the degrees of freedom of planes $\mathcal{P}$, 2D planar Gaussians $\mathcal{G}$, and 3D freeform Gaussians $\bar{\mathcal{G}}$. We begin our optimization with a warm-up stage using only 3D Gaussians (for N=3500 iterations). After that, we begin our planar reconstruction where in each round of optimization we: (i) dynamically initialize plane parameters by *robustly* fitting planes to the current representation (section 3.2); (ii) alternate between optimizing plane and Gaussian parameters (section 3.2); (iii) densify our representation through a (slightly modified) MCMC densification, due to the challenges of optimizing compact-support functions (section 3.4).

3.2 Plane initialization

For compactness of notation, let us drop our indices, and consider the binary mask $\mathcal{M} \leftarrow \mathcal{M}_{c,p}$ for the p-th planar surface in the c-th view, and denote with π the function that projects a 3D point to the coordinate frame of the n-th image. We start by selecting all the Gaussians (i) whose mean projects into the mask, (ii) that are sufficiently opaque, and (iii) that lie within a shell of the expected ray termination of the n-th image:

$$\tilde{\mathcal{G}} = \{\bar{\mathbf{g}}_k \mid \pi(\bar{\mu}_k)\in\mathcal{M},\ \alpha_k>\kappa,\ |D(\pi(\bar{\mu}_k)) - d_k|<\delta\}, \tag{4}$$

where the thresholds $\alpha_{\text{th}}=0.1$ and $d_{\text{th}}=0.05$ are hyper-parameters that control this selection process, and where D is the expected ray termination map (i.e. depth map), and d_k is the depth of the

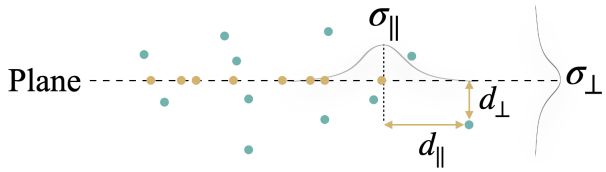

Figure 3: **Planar Relocation** – A freeform Gaussian (teal) gets relocated to the plane to become a planar Gaussian (brown), when both its distace to the plane ($d_\perp$) and along ($d_\parallel$) the plane are small.

Gaussians. We then extract a candidate plane P by RANSAC optimization on the point cloud that samples the Gaussians:

$$P, \mathcal{I} = \text{RANSAC}\big(\{x \sim \bar{\mathbf{g}} \mid \bar{\mathbf{g}} \in \tilde{\mathcal{G}}\}, \epsilon\big) \tag{5}$$

where we accept P as a viable plane candidate only whenever the mean inlier residual is lower than ϵ. The set $\mathcal{I}$ includes the indexes of Gaussians in $\tilde{\mathcal{G}}$ that are inliers of the RANSAC process. We further discard planes that are too small with set $\mathcal{I}$ having a smaller size than 100. Once a plane corresponding to $\mathcal{M}$ has been accepted, all the semantic masks for that plane p are excluded from subsequent RANSAC runs. The plane initialization process is repeated for remaining masks, after each completed round of plane and Gaussian optimization, as described in Section 3.3.

Snapping We then remove the discovered inliers from the set of 3D Gaussians $\bar{\mathcal{G}} \leftarrow \bar{\mathcal{G}} \setminus \mathcal{I}$, and add them to our set of 2D Gaussians $\mathcal{G} \leftarrow \mathcal{G} \cup \mathcal{I}$. During the latter operation, we clip 3D Gaussians to 2D to become planar by transforming to the local plane coordinates, and set the third component of their means and scales to zero. Further, only rotation about the z-axis in local plane coordinates is preserved

Active set update If the accepted plane $\mathcal{P}_i$ has an angular distance below a threshold to an already existing plane, while its origin $\mathbf{o}_i$ also has a small Euclidean distance to the closest Gaussian center on that plane, we merge the two planes. Otherwise, the plane is added as a new plane to the active set of planes $\mathcal{P}$. In merging, we assign the new plane's Gaussians to the previously found one. This allows our optimization to merge planar areas that have only been partially observed in any view.

3.3 Optimization

We optimize our representation by block-coordinate descent, starting each round of optimization by only optimizing the plane parameters for a fixed number of 10 iterations, and then freezing these, and optimizing the Gaussian parameters (both 2D and 3D) for another 100 iterations. This alternation in optimization is critical to avoid instability; see an ablation in figure 7. In the first optimization block, within each iteration, the parameters of the p-th plane within the c-th image are optimized by the loss:

$$\underset{\mathbf{o}_p, \mathbf{n}_p}{\arg\min} = \underbrace{\|I_c - \tilde{I}_c\|_1}_{\mathcal{L}_{\text{photo}}} + \lambda_{\text{mask}} \underbrace{\|\mathcal{M}_{c,p} - \tilde{\mathcal{M}}_{c,p}\|_1}_{\mathcal{L}_{\text{mask}}}, \tag{6}$$

where $\tilde{\mathcal{M}}$ is the predicted plane mask, obtained by rendering the mixture of Gaussians with binarized color (white for planar, and black for 3D), and alpha blended using the original Gaussian opacities during the rasterization. In the second optimization block, we optimize all Gaussian parameters jointly:

$$\underset{\mathcal{G}, \bar{\mathcal{G}}}{\arg\min} \;\; \mathcal{L}_{\text{photo}} + \lambda_{\text{mask}}\mathcal{L}_{\text{mask}} + \lambda_{\text{TV}}\mathcal{L}_{\text{TV}} + \lambda_{\text{scale}}\mathcal{L}_{\text{scale}} + \lambda_{\text{opacity}}\mathcal{L}_{\text{opacity}}, \tag{7}$$

where $\mathcal{L}_{\text{TV}}$ is the total depth variation regularization from Niemeyer et al. [10], $\mathcal{L}_{\text{scale}}$ is the scale regularizer and $\mathcal{L}_{\text{opacity}}$ is the opacity regularizer from Kheradmand et al. [13] that vanishes the size of Gaussians that are unconstrained by the photometric loss. Note that planar Gaussians move *rigidly* during plane optimization (6), and move locally in the plane during Gaussian optimization (7), as only their 2D in-plane parameters are optimized.

3.4 Planar relocation

We follow 3DGS-MCMC [13] in our training dynamics. For densification of planes, we rely on relocating low-opacity Gaussians to locations of dense high-opacity Gaussians, as this allows

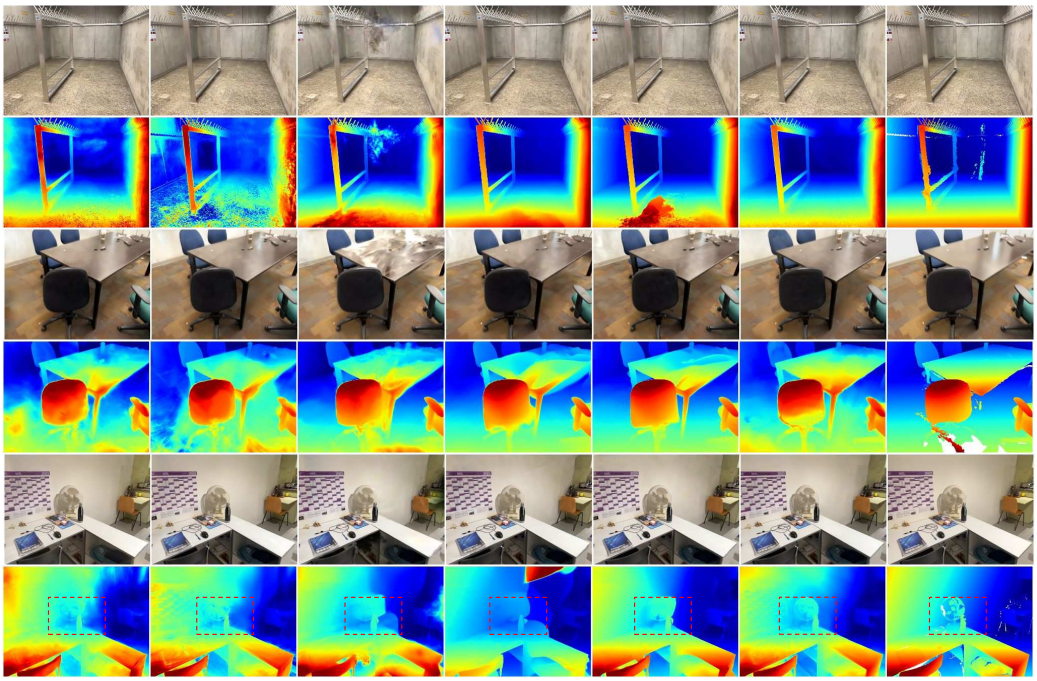

Metric / Method	RMSE↓	MAE↓	AbsRel↓	$\delta < 1.25$↑	$\delta < 1.25^2$↑	$\delta < 1.25^3$↑	PSNR↑	LPIPS↓	SSIM↑	#primitives	(%planar)
3DGS [6]	0.44	0.34	0.17	0.71	0.89	0.94	27.09	0.20	0.89	2.43M	
3DGS-MCMC [13]	0.49	0.32	0.19	0.78	0.93	0.96	**27.23**	**0.20**	**0.90**	2.43M	
RaDe-GS [17]	0.65	0.49	0.26	0.64	0.74	0.77	20.13	0.30	0.82	1.58M	
2DGS [7]	0.39	0.24	0.13	0.82	0.88	0.91	25.56	0.24	0.88	1.54M	
PGSR [16]	0.35	0.20	0.10	0.85	0.90	0.93	25.78	0.23	0.88	2.47M	
Ours	**0.27**	**0.18**	**0.10**	**0.88**	**0.96**	**0.98**	27.01	0.21	0.89	2.43M	(27.8%)

Figure 4: **Novel View Synthesis** – Quantitative and qualitative results show significant improvement in predicted depth compared to previous work, while maintaining comparable rendering quality to the full 3D representations.

transferring between 3D and 2D/planar Gaussians. However, the number of Gaussians on planes, especially when the plane has weak texture, is usually low, leading to a slow densification rate for planes / planar Gaussians. To address this issue, whenever a freeform Gaussian projects into the current mask $\pi(\bar{\mu}_k) \in \mathcal{M}$, and it is *sufficiently close* to the currently reconstruction, we stochastically re-locate it to the plane. To measure distance, we identify the 2D Gaussian with the smallest Euclidean distance to $\bar{\mu}_k$, and measure its distance in the direction of the plane normal $d_\perp$, and the one along the plane $d_\parallel$; see Figure 3. We stochastically relocate this if both distances are sufficiently small, as expressed by the following Bernoulli distribution:

$$p \sim \mathcal{B}(\beta),\ \beta = \left[1 - \Phi\left(\frac{d_\perp}{\sigma_\perp}\right)\right] \cdot \left[1 - \Phi\left(\frac{d_\parallel}{\sigma_\parallel}\right)\right], \tag{8}$$

where Φ is the cumulative distribution function of a Gaussian, and $\sigma_\perp$ and $\sigma_\parallel$ are hyper-parameters that control the stochastic re-location.

4 Results

We validate our proposed method for scene reconstruction through the novel view synthesis task on common indoor scene datasets, assessing both rendered image and depth quality metrics (section 4.1). We then show an application of our method to mesh extraction for planar surfaces (section 4.2). Finally, we validate our design choices through an ablation study on different aspects of the method (section 4.3). We provide our implementation details in the supplementary material.

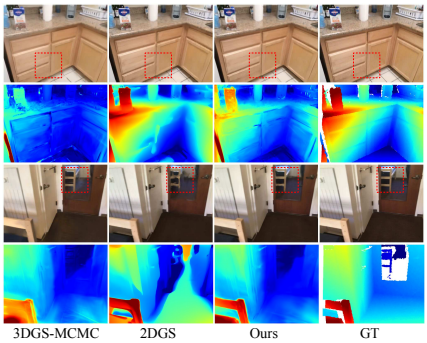

Metric	3DGS-MCMC	2DGS	Ours
RMSE↓	0.46	0.60	**0.40**
MAE↓	0.37	0.44	**0.31**
AbsRel↓	0.19	0.23	**0.16**
$\delta < 1.25 \uparrow$	0.61	0.63	**0.70**
$\delta < 1.25^2 \uparrow$	0.87	0.77	**0.90**
$\delta < 1.25^3 \uparrow$	0.95	0.83	**0.97**
PSNR↑	20.18	21.44	**21.75**
LPIPS↓	0.29	0.30	**0.27**
SSIM↑	0.83	0.85	**0.86**
# primitives	500K	809K	500K
(% planar)			(17.6%)

Figure 5: **Novel View Synthesis on ScanNetv2** – Our method outperforms baselines in image and depth quality on ScanNetv2 despite sparse camera views.

4.1 Novel View Synthesis – Figures 4 and 5

We evaluate our hybrid representation's novel view synthesis on common indoor scene reconstruction benchmarks and provide comparisons with both state-of-the-art fully 3D representations and 2D surface reconstruction approaches. We show a significant improvement in the reconstructed surface geometry while maintaining high visual quality.

Datasets We perform evaluations on common indoor scene benchmarks ScanNet++[31] and ScanNetv2[32], as they primarily feature indoor scenes with flat textureless surfaces suitable for the task at hand. ScanNet++ provides dense scenes with SfM camera poses and sparse point clouds, designed primarily for 3D reconstruction approaches that follow the NeRF [5] paradigm. Conversely, the legacy version of ScanNet i.e. ScanNetv2 offers sparser views without SfM information. Our method works with or without initial sparse point clouds, enabling reconstruction initialized with sparse SfM point cloud on ScanNet++ and experiments with randomly initialized point clouds on ScanNetv2. For ScanNet++, we use 11 training scenes with ground truth meshes for depth derivation, utilizing iPhone video streams, sampling every 10th frame for training at 2× downsampling and every 8th for testing. We chose the scenes that are diverse in their content and contain various planar surfaces. For ScanNet, we evaluate on 5 scenes with sufficient overlapping views of planar surfaces following the data preparation scheme of [27]. The 2D plane masks were generated using PlaneRecNet [25] and propagated through the image sequence with SAMv2 video processor [29].

Baselines We compare against SOTA reconstruction methods, both fully 3D representations and 2D surface reconstruction methods. For 3D representations, we compare with vanilla 3DGS [6], and 3DGS-MCMC [13] as it is more robust version to random initializations, and has higher rendering quality. Within photometric surface reconstruction methods, we compare to 2DGS [7] as a widely used state-of-the-art, as well as to PGSR [16] and RaDe-GS [17], which more recently report improved depth quality. All methods are trained for 30K iterations.

Metrics We use the common image quality metrics PSNR, SSIM and LPIPS for evaluating the rendered RGB. Further, we choose depth as a strong indicator for the quality of the reconstructed surface geometry. We provide depth quality metrics by computing the rendered depth as the expected ray termination at each pixel. We report RMSE, MAE and average absolute error relative to ground truth depth (AbsRel). Additionally, we provide depth accuracy percentage at different error thresholds similar to [33]. The metrics are computed only on the defined portion of the ground-truth depths. We further report the total number of primitives in our model and the percentage that are planar (and thus can be represented more compactly).

Analysis Quantitative and qualitative results across both datasets show significant improvement in depth accuracy compared to *all* baselines. Notably, our method achieves comparable image quality to SOTA 3D representations on dense ScanNet++ scenes while surpassing them in depth quality, evidenced by sharper geometry reconstruction in qualitative examples. The slight PSNR difference with 3D methods reflects a trade-off: our constrained geometry enforces correct structure, while unconstrained methods can inflate PSNR by fitting view-dependent effects with incorrect geometry.

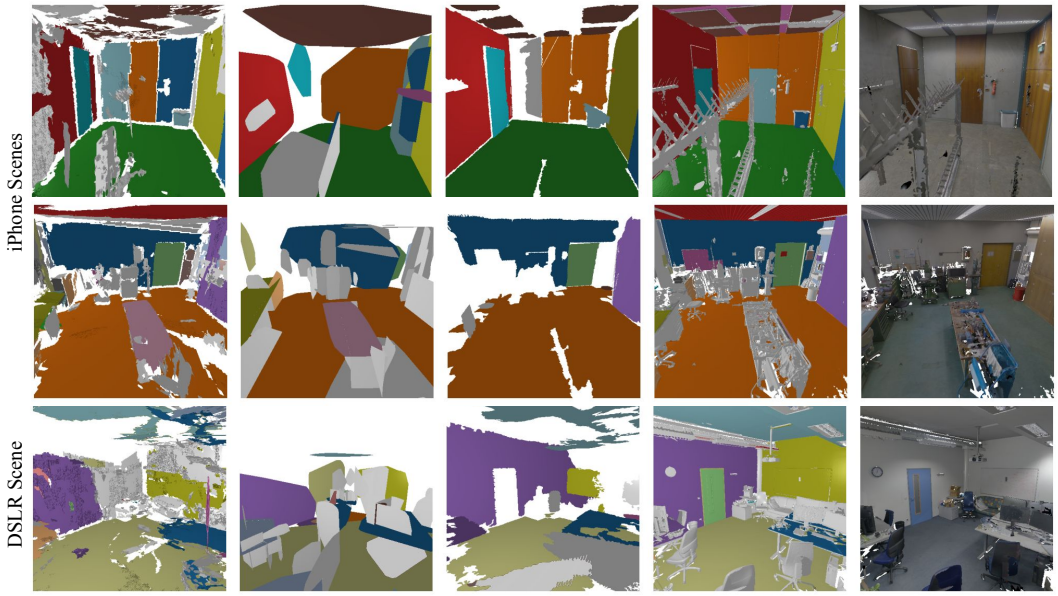

		Meshing						Segmentation		
Dataset	Metrics	Acc↓	Comp↓	Chamfer↓	Precision↑	Recall↑	F1 score↑	VOI↓	RI↑	SC↑
DSLR	Airplanes [9]	23.09	30.12	26.60	8.47	6.57	7.35	5.24	**0.64**	0.18
	PlanarRecon [8]	15.99	59.92	37.96	23.10	4.16	6.77	4.31	0.63	0.20
	Ours	**6.93**	**17.31**	**12.12**	**65.33**	**46.34**	**53.71**	**3.89**	**0.64**	**0.24**
iPhone	Airplanes [9]	7.15	**15.46**	**11.31**	48.03	38.02	41.94	4.38	**0.69**	**0.28**
	PlanarRecon [8]	8.72	30.08	19.40	50.61	30.44	36.93	4.23	0.68	0.24
	Ours	**4.60**	32.59	18.60	**75.10**	**39.12**	**50.24**	**4.08**	0.67	0.23

Figure 6: **Mesh Extraction** – Our method shows consistent results across iPhone and DSLR captures, while baselines typically overfit to one camera type. Qualitatively, our approach extracts complete meshes for most target planes with fewer inaccurate plane detections (shown in gray) compared to baselines. Target planes are shown with distinct colors on the ground truth.

In the sparser ScanNetv2 scenes, our approach delivers superior performance in both depth and image quality, leveraging the planar prior of indoor environments to overcome the geometric ambiguity that challenges pure 3D methods in sparse captures. Our method also substantially outperforms 2DGS in both image fidelity and depth accuracy metrics.

4.2 Mesh Extraction – figure 6

Our method enables mesh extraction from reconstructed 3D planar surfaces. For each plane, we un-project all 2D segmentation masks to 3D by computing ray-plane intersections, yielding a point cloud. This point cloud is downsampled using fixed-size voxels and rasterized onto plane coordinates to create an occupancy grid. We then use Marching Squares for contour extraction (We omit small contours with less than 100 points), followed by ear-clipping triangulation to produce the final mesh. We evaluate the quality of the retrieved mesh for the planar surfaces and compare our method to planar reconstruction methods.

Datasets We use ScanNet++ to extract planar surface meshes. We show results both on the subset of this dataset captured by iPhone and also the DSLR subset, showing that our method can handle different camera models, while previous methods usually overfit to one modality. For ground truth, we follow the approach of Watson et al. [9] to obtain a ground truth planar mesh. We then only consider the subset of planes in the ground truth mesh that we have annotated segmentation masks for each scene. We provide details on selecting these planes in Appendix E.

Baselines We compare against previous planar reconstruction methods AirPlanes [9] and PlanarRecon [8] that provide extracted planar mesh as output of their methods. We follow the same evaluation setting as in the original papers on the iPhone subset of the dataset. For DSLR images, we crop the images to the specified FoV in each baseline to match their training distribution.

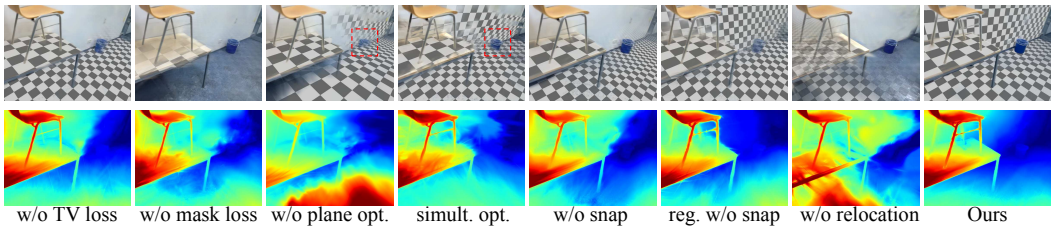

		Loss design		Optimization design		2D Gaussian design		
	Full model	w/o $\mathcal{L}_{\text{TV}}$	w/o $\mathcal{L}_{\text{mask}}$	w/o plane optimization	simult. joint optimization	w/o snapping	reg. w/o snapping	w/o relocation
PSNR↑	**26.83**	23.24	24.02	21.08	19.52	25.53	21.69	20.00
RMSE↓	**0.25**	0.34	0.62	0.54	0.40	0.38	0.36	0.59

Figure 7: **Ablation on design choices –** Loss components and optimization strategy are critical, with simultaneous plane-Gaussian optimization causing significant drops. 2D Gaussian snapping greatly improves depth accuracy compared to regularization alternatives. Similarly, Gaussian relocation is essential.

Metrics We report mesh accuracy metrics including accuracy, precision, recall, completeness and Chamfer distance as defined in Ye et al. [27]. We also provide mesh segmentation metrics that evaluate how well detected plane segments match ground truth segments following [9].

Analysis Our method outperforms the baselines on DSLR images subset of the dataset. Unlike previous methods that are trained on specific modalities (i.e. phone camera) and struggle to transfer to different camera models (i.e. DSLR camera), our approach maintains consistent mesh quality due to having zero-shot mesh extraction on test scenes through photometric reconstruction. Additionally, our method outperforms PlanarRecon on iPhone data, while having competitive performance to AirPlanes. Qualitative results reveal that both PlanarRecon and AirPlanes extract extraneous planes with numerous random small fragments, resulting in unsightly and impractical meshes. In contrast, our method produces clean planar surfaces, yielding a more coherent and usable reconstruction.

4.3 Ablation – Figure 7

We ablate our design choices and additionally test our method's robustness to random point cloud initialization (in table 1).

Loss design We ablate the effect of $\mathcal{L}_{mask}$ and $\mathcal{L}_{TV}$. Although removing these losses reduces the image quality by some margin, it affects depth quality more significantly. Qualitative rendering shows that $\mathcal{L}_{mask}$ contributes significantly to detecting and growing 2D Gaussians.

Optimization design Our method is based on optimizing Gaussians and plane parameters together in an alternating fashion. We show that fixing plane parameters with no optimization degrades our results both quantitatively and qualitatively. Simultaneous joint optimization of Gaussians and planes also affects the results negatively. In Figure 7, note how the floor plane gets stuck above the ground level, as revealed by its intersection with the bin.

2D Gaussian design Using hybrid 2D/3D Gaussians is one of the main components of our design. Therefore, we ablate the necessity of having 2D Gaussians by disabling snapping as described in Section 3.2. This shows a significant drop in depth accuracy, which is also evident in qualitative results. As an alternative to snapping, we can regularize the smallest scale component in planar Gaussians. However, we find that this approach is difficult to tune and provides suboptimal results. Finally, we ablate our densification process with relocation of Gaussians to planes. Without relocation, planes are not fully detected, with the planar Gaussians comprising the plane maintaining low opacity. Furthermore, some of the Gaussians remain close to the plane while not being detected as belonging to that plane.

5 Conclusions

We introduce 3D Gaussian Flats, a hybrid 2D/3D Gaussian representation that accurately models planar surfaces without sacrificing rendering quality. Our method jointly optimizes 2D Gaussians constrained to planar surfaces alongside free-form Gaussians for the remaining scene. By leveraging semantic segmentation masks, we predict both a full 3D representation and semantically distinct planes for planar mesh extraction in indoor scenes. Our approach achieves state-of-the-art depth estimation on indoor scene benchmarks while maintaining high image quality. Additionally, our planar mesh extraction method generalizes across different camera models, overcoming domain gap limitations that typically cause previous methods to fail.

Limitations Our reliance on initial 3DGS reconstruction often generates insufficient Gaussians in flat areas with no texture, although this potentially can be addressed via more adaptive densification strategies. Further, using a weak spherical harmonics appearance model still leads to building extra geometry to compensate for view-dependent effects, which a stronger appearance model would resolve. Additionally, we depend on 2D semantic masks from SAMv2 that may contain errors, but our method will naturally improve alongside advances in semantic segmentation. Finally, our RANSAC-based approach, while robust, introduces computational overhead that extends training time. We believe our hybrid representation opens exciting new avenues for research into more efficient approaches that balance geometric precision with visual fidelity.

References

[1] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. URL `https://github.com/google-research/multinerf`. 1

[2] A. Tewari, J. Thies, B. Mildenhall, P. Srinivasan, E. Tretschk, W. Yifan, C. Lassner, V. Sitzmann, R. Martin-Brualla, S. Lombardi, T. Simon, C. Theobalt, M. Nießner, J. T. Barron, G. Wetzstein, M. Zollhöfer, and V. Golyanik. Advances in neural rendering. *Computer Graphics Forum*, 2022. 2

[3] Guikun Chen and Wenguan Wang. A survey on 3d gaussian splatting. *arXiv preprint arXiv:2401.03890*, 2025. 2

[4] Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes' Rays: Uncertainty quantification in neural radiance fields. *CVPR*, 2024. URL `https://github.com/BayesRays/BayesRays`. 2

[5] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. URL `https://github.com/bmild/nerf`. 2, 3, 7

[6] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 2023. URL `https://github.com/graphdeco-inria/gaussian-splatting`. 2, 3, 4, 6, 7

[7] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH*, 2024. URL `https://github.com/hbb1/2d-gaussian-splatting`. 2, 3, 6, 7, 1

[8] Yiming Xie, Matheus Gadelha, Fengting Yang, Xiaowei Zhou, and Huaizu Jiang. Planarrecon: Real-time 3d plane detection and reconstruction from posed monocular videos. In *CVPR*, 2022. URL `https://github.com/neu-vi/PlanarRecon`. 2, 3, 8, 6

[9] Jamie Watson, Filippo Aleotti, Mohamed Sayed, Zawar Qureshi, Oisin Mac Aodha, Gabriel Brostow, Michael Firman, and Sara Vicente. Airplanes: Accurate plane estimation via 3d-consistent embeddings. In *CVPR*, 2024. URL `https://github.com/nianticlabs/airplanes`. 2, 3, 8, 9, 6

[10] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. URL https://github.com/google-research/google-research/tree/master/regnerf. 2, 5, 7

[11] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. URL https://github.com/Totoro97/NeuS. 2, 3

[12] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive shells for efficient neural radiance field rendering. *ACM TOG.*, 2023. 2

[13] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Jeff Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. In *NeurIPS*, 2024. URL https://github.com/ubc-vision/3dgs-mcmc. 3, 5, 6, 7, 1

[14] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, 2023. URL https://github.com/NVlabs/neuralangelo. 3

[15] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *CVPR*, 2024. URL https://github.com/Anttwo/SuGaR. 3

[16] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 2024. URL https://github.com/zju3dv/PGSR. 3, 6, 7

[17] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. Rade-gs: Rasterizing depth in gaussian splatting. *arXiv preprint arXiv:2406.01467*, 2024. URL https://github.com/BaowenZ/RaDe-GS. 3, 6, 7

[18] Zhi-Hao Lin, Wei-Chiu Ma, Hao-Yu Hsu, Yu-Chiang Frank Wang, and Shenlong Wang. Neurmips: Neural mixture of planar experts for view synthesis. In *CVPR*, 2022. URL https://github.com/chih-hao-lin/neurmips. 3

[19] Bin Tan, Rui Yu, Yujun Shen, and Nan Xue. Planarsplatting: Accurate planar surface reconstruction in 3 minutes. In *CVPR*, 2025. 3

[20] David Svitov, Pietro Morerio, Lourdes Agapito, and Alessio Del Bue. Billboard splatting (bbsplat): Learnable textured primitives for novel view synthesis. *arXiv preprint arXiv:2411.08508*, 2024. URL https://github.com/david-svitov/BBSplat. 3

[21] Yao-Chih Lee, Zhoutong Zhang, Kevin Blackburn-Matzen, Simon Niklaus, Jianming Zhang, Jia-Bin Huang, and Feng Liu. Fast view synthesis of casual videos with soup-of-planes. In *ECCV*, 2024. 3

[22] Haoyu Guo, Sida Peng, Haotong Lin, Qianqian Wang, Guofeng Zhang, Hujun Bao, and Xiaowei Zhou. Neural 3d scene reconstruction with the manhattan-world assumption. In *CVPR*, 2022. URL https://github.com/zju3dv/manhattan_sdf. 3

[23] Zheng Chen, Chen Wang, Yuan-Chen Guo, and Song-Hai Zhang. Structnerf: Neural radiance fields for indoor scenes with structural hints. *IEEE TPAMI*, 2023. 3

[24] Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, and Jan Kautz. Planercnn: 3d plane detection and reconstruction from a single image. In *CVPR*, 2019. URL https://github.com/NVlabs/planercnn. 3

[25] Yaxu Xie, Fangwen Shu, Jason Rambach, Alain Pagani, and Didier Stricker. Planerecnet: Multi-task learning with cross-task consistency for piece-wise plane detection and reconstruction from a single rgb image. In *BMVC*, 2021. URL https://github.com/EryiXie/PlaneRecNet. 3, 7, 4, 6

[26] Zheng Chen, Qingan Yan, Huangying Zhan, Changjiang Cai, Xiangyu Xu, Yuzhong Huang, Weihan Wang, Ziyue Feng, Lantao Liu, and Yi Xu. Planarnerf: Online learning of planar primitives with neural radiance fields. *arXiv preprint arXiv:2401.00871*, 2023. 3

[27] Hanqiao Ye, Yuzhou Liu, Yangdong Liu, and Shuhan Shen. Neuralplane: Structured 3d reconstruction in planar primitives with neural fields. In *ICLR*, 2025. URL `https://github.com/3dv-casia/NeuralPlane`. 3, 7, 9

[28] Jiyeop Kim and Jongwoo Lim. Integrating meshes and 3d gaussians for indoor scene reconstruction with sam mask guidance. *arXiv preprint arXiv:2407.16173*, 2024. 3

[29] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *ICLR*, 2025. URL `https://github.com/facebookresearch/segment-anything`. 3, 7, 4, 6

[30] Farhad G Zanjani, Hong Cai, Hanno Ackermann, Leila Mirvakhabova, and Fatih Porikli. Planar gaussian splatting. In *WACV*, 2025. 3

[31] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *ICCV*, 2023. URL `https://kaldir.vc.in.tum.de/scannetpp`. Licensed under the ScanNet++ Terms of Use. 7, 17, 1, 3, 4

[32] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. URL `http://www.scan-net.org`. Licensed under the MIT License. 7, 17

[33] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024. URL `https://github.com/LiheYoung/Depth-Anything`. 7

[34] Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1

[35] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics*, 2024. 1

[36] Matias Turkulainen, Xuqian Ren, Iaroslav Melekhov, Otto Seiskari, Esa Rahtu, and Juho Kannala. Dn-splatter: Depth and normal priors for gaussian splatting and meshing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025. 1

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: A discussion of current state of the literature can be found in section 2, proposed method is defined in section 3 and the discussion of results in section 4 support the claims made in the abstract and introduction.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Limitations of the work are listed in section 5.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not contain theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All experiments descriptions in the paper follow the same structure, discussing datasets, baselines, metrics and analysing the results with provided hyperparameter settings and data sampling methodology. In Appendix E, we additionally discuss data preparation procedures.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code is provided with the supplementary material, the used datasets are public.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: The hyperparameters and data splits are provided in section 4 and Appendix F. Hyperparameter choice justification is discussed.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: It is infeasible to report statistical significance for all compared methods due to the amount of compute required and availiable. The observed variance in the conducted experiments was negligible.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Information is provided in Appendix F.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The authors reviewed and followed the NeurIPS Code of Ethics. In particular, used datasets are anonymized.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: The method improves the quality and efficiency in the tasks of Novel View Synthesis and planar mesh extraction – however, these improvements do not introduce conceptually new capabilities that would require revision of societal impact when compared to prior work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: Used datasets are credited with licences ([31] and [32]), competing methods are cited after the method short name in every table, and the citation includes code access links crediting license. Used codebases are credited in Appendix E.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The code contains README file with all the information to reproduce the paper results.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: LLM policy is reflected in OpenReview submission. The paper does not involve LLMs as any important, original, or non-standard components.

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.

A Full mesh extraction results – Figures 8 to 10

We evaluate our hybrid representation on the task of full mesh extraction using the method from [7], we do it in addition to the planar-only mesh extraction experiments presented in Section 4.2, concatenating the two meshes together and comparing them to common benchmarks from Section 4.1.

Datasets We evaluate on ScanNet++ [31], a common indoor scene benchmark, as well as on subset of suitable indoor/outdoor scenes from ETH3D [34], which provides high quality mesh, and is more challenging because of sparse image supervision.

Baselines For ScanNet++ we reuse the models trained on iPhone data stream and evaluated on the task of NVS in Section 4.1 to access mesh quality reconstruction. On ETH3D, in addition, we evaluate Gaussian Opacity Fields (GOF) [35], an extention of 2DGS for higher quality mesh reconstruction, and DNSplatter [36], a method supervising 3DGS with mono-depth [2]

To obtain the mesh, we use TSDF fusion with the median depth estimate for 3DGS, 2DGS, DNSplatter and ours, rather than the expected ray termination as in default settings (i.e., average depth). For PGSR we use their proposed unbiased depth computation, and for Gaussian Opacity Fields we extract the mesh using the level set of the Gaussians, hence the mesh is not colored.

Metrics We use the same metrics as for meshing task in planar mesh experiments Section 4.2. We compute the F1-score at 5 cm threshold. For both of the datasets, we use every 8th image as a test image.

Analysis We provide full mesh renders along with the metrics on ScanNet++ in Figure 8. For ETH3D, in addition to mesh renders in Figure 10, we provide rendered novel views from the test set in Figure 9. Note that captured planar surfaces are unbiased and outline well the structures of the scenes. Moreover, on in the sparse view setting on ETH3D dataset we achieve a notable rendering quality improvement.

B Additional ablations – Tables 1 and 2

Random initialization We analyze the effect of having sparse point cloud initialization versus random initialization in our method on 11 DSLR scenes from ScanNet++ [31]. for random initialization we do 5000 iterations in our warmup stage, as opposed to the usual 3500. We show that our method maintains the robustness to random initialization similar to 3DGS-MCMC [13], and despite a drop in number of planar Gaussians, it achieves comparable depth and image quality metrics to our method when initialized with SfM sparse point cloud.

Table 1: **Ablation on initialization –** Our method is robust to random initialization and achieves comparable performance to when initialized with SfM point cloud.

Method	PSNR↑	SSIM↑	LPIPS↓	RMSE↓	MAE↓	AbsRel↓	#primitives	(%planar)
3DGS-MCMC (SfM)	23.38	**0.87**	**0.24**	0.41	0.24	0.26	1.13M	
Ours (SfM)	**23.42**	**0.87**	**0.24**	**0.20**	**0.13**	**0.12**	1.13M	(31%)
Ours (Random)	23.30	0.86	0.25	0.21	0.14	0.13	1.13M	(21%)

Full metrics set for ablation on design choices We provide the full set of metrics for ablation on design choices (described in section 4.3) in the table 2.

[2]Note that the released codebase for DNSplatter does not support multiple camera models (different camera intrinsics) for aligning mono-depth to SfM points, therefore we cannot easily report the metrics for 'Electro' and 'Terrace' scenes.

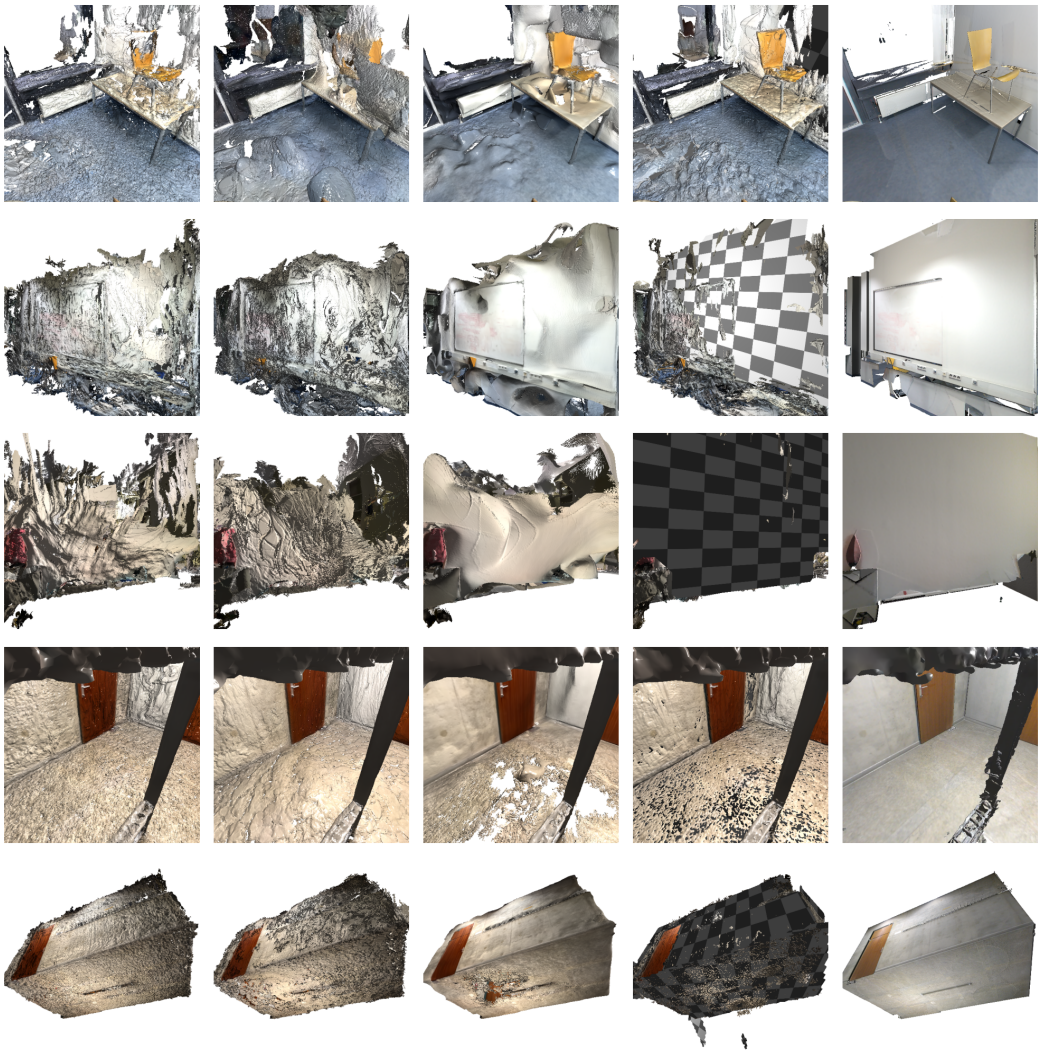

3DGS 2DGS PGSR Ours Ground Truth

	Acc↓	Comp↓	Chamfer↓	F1↑
3DGS	0.14	0.12	0.1274	0.5639
2DGS	0.27	0.15	0.2082	0.5280
PGSR	**0.13**	**0.15**	**0.1404**	**0.5981**
Ours	0.25	0.12	0.1833	0.5820

Figure 8: **Full Mesh Extraction Results on ScanNet++** – Out method achieves competitive performance for surface reconstruction, while mainatining the rendering quality. Checkered surfaces indicate different planes, planes are usually behind the TSDF-extracted mesh as they represent unbiased surfaces. Some of the meshes are shown from outside of the indoor scene to highlight the planar alignment.

C Additional video and 3D mesh results

We provide video renderings of RGB and depth for our method compared to baselines in https://theialab.github.io/3dgs-flats. Video results best capture the significant enhancement of our approach over baselines in depth estimation and accurately modeling scene geometry.

Method	Electro			Terrace			Delivery area		
	PSNR↑	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑
3DGS	16.45	0.38	0.72	20.77	0.27	0.78	19.48	0.29	0.83
2DGS	16.40	0.41	0.72	20.82	0.29	0.79	19.26	0.35	0.81
GOF	17.34	0.36	0.71	20.80	0.27	0.75	19.40	0.33	0.79
PGSR	–	–	–	–	–	–	16.64	0.41	0.69
DNSplatter	–	–	–	–	–	–	19.56	0.24	0.77
Ours	**18.72**	**0.31**	**0.75**	**22.57**	**0.22**	**0.81**	**22.56**	**0.21**	**0.87**

Figure 9: **Rendering Results on ETH3D Scenes** – Our method outperforms the baselines in terms of rendering quality on this set of sparse view outdoor/indoor scenes, and the planar representation is crucial for achieving good novel view synthesis in sparse scenarios.

Table 2: **Ablation on design choices –** Loss components and optimization strategy are critical, with simultaneous plane-Gaussian optimization causing significant drops. 2D Gaussian snapping greatly improves depth accuracy compared to regularization alternatives. Similarly, Gaussian relocation is essential.

	PSNR↑	LPIPS↓	SSIM↑	RMSE↓	MAE↓	AbsRel↓
Full model	**26.83**	**0.27**	**0.86**	**0.25**	**0.18**	**0.09**
Loss design:						
w/o $\mathcal{L}_{\text{TV}}$	23.24	0.34	0.82	0.34	0.24	0.13
w/o $\mathcal{L}_{\text{mask}}$	24.02	0.32	0.83	0.62	0.53	0.29
Optimization design:						
w/o plane optimization	21.08	0.37	0.80	0.54	0.43	0.24
simult. joint optimization	19.52	0.38	0.79	0.40	0.32	0.18
2D Gaussian design:						
w/o snapping	25.53	0.31	0.84	0.38	0.31	0.17
reg. w/o snapping	21.69	0.35	0.81	0.36	0.28	0.15
w/o relocation	20.00	0.37	0.80	0.59	0.50	0.28

D Additional qualitative results – Figures 11 and 12

We provide more qualitative evidence for the performance of our method compared to 2DGS [7], 3DGS [6] and 3DGS-MCMC [13] baselines on the ScanNet++ [31] dataset in figure 11. The results show how baselines particularly struggle with reconstructing accurate geometry for the textureless areas while our method significantly improves upon these methods in depth estimation and keeps the visual quality of images.

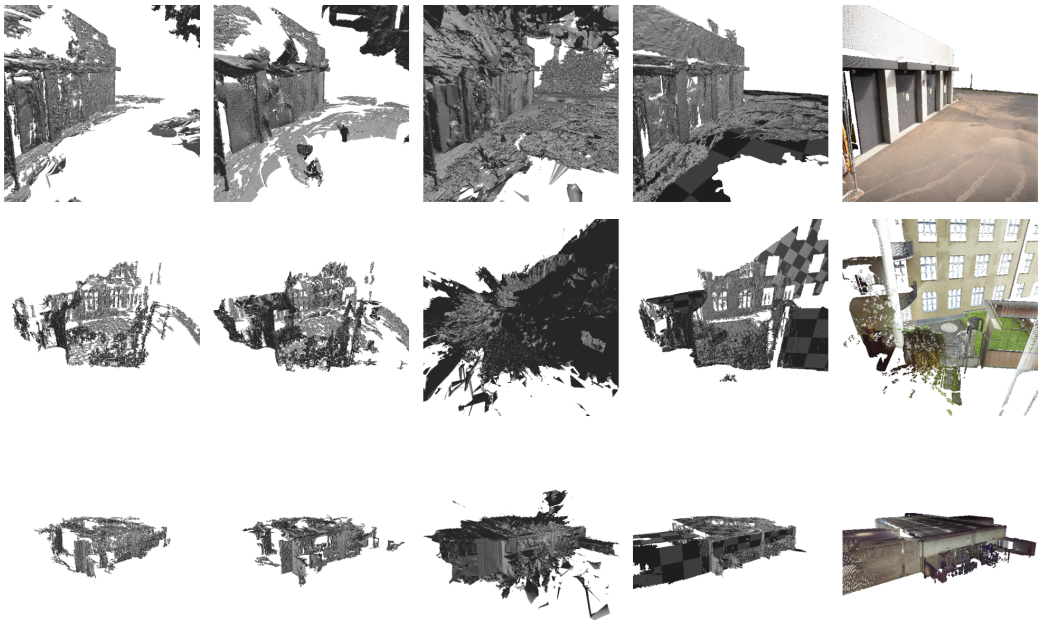

3DGS 2DGS GOF Ours Ground Truth

Method	Electro		Terrace		Delivery area	
	Chamfer↓	F1↑	Chamfer↓	F1↑	Chamfer↓	F1↑
3DGS	0.6524	0.2511	0.3258	0.4517	0.3064	0.2335
2DGS	0.5873	0.2570	0.3312	0.4036	0.3265	0.2366
GOF	0.5371	0.2991	0.2107	0.4045	0.2939	0.3131
PGSR	–	–	–	–	0.4266	**0.4287**
DNSplatter	–	–	–	–	0.2488	0.2516
Ours	**0.4062**	**0.3009**	**0.1480**	**0.5033**	**0.1825**	0.3313

Figure 10: **Full Mesh Extraction Results on ETH3D Scenes** – Our method outperforms the baselines.

Further, we provide more visualization for our estimated planes on ScanNet++ [31] dataset, showcasing the perfect alignment of our planes with the detected planar surfaces in figure 12.

E Input planar masks

2D semantic masks Our method relies on input consistent 2D segmentation masks of planar surfaces. To obtain these masks, we can either annotate each image collection manually or automate the process for larger scenes. To automatically obtain the 2D segmentation masks, we employ PlaneRecNet [25] and SAMv2 video segmentation model [29], to create an annotation pipeline. We first input images to PlaneRecNet to obtain 2D plane annotations that are not semantically consistent across the image collection. We set the plane probability threshold to 0.5. While this method works well on iPhone images, it produces fewer plane annotations for DSLR images, that are out of distribution for its network trained on iPhone data. We then input these unmatched masks as seed to SAMv2. In order to do that, we order image collections that are not already sampled from a video. We propagate masks from the initial frame in 16-frame chunks of the sequence to the next 15 frames, and match SAMv2's prediction with any subsequent 2D masks output from [25], using Hungarian matching with an IoU metric. Although largely effective, this process is prone to error accumulation through mask propagation. However, we assume resultant masks are semantically consistent across the image sequence. We provide sample segmentation of an input sequence in the supplementary video and on the website.

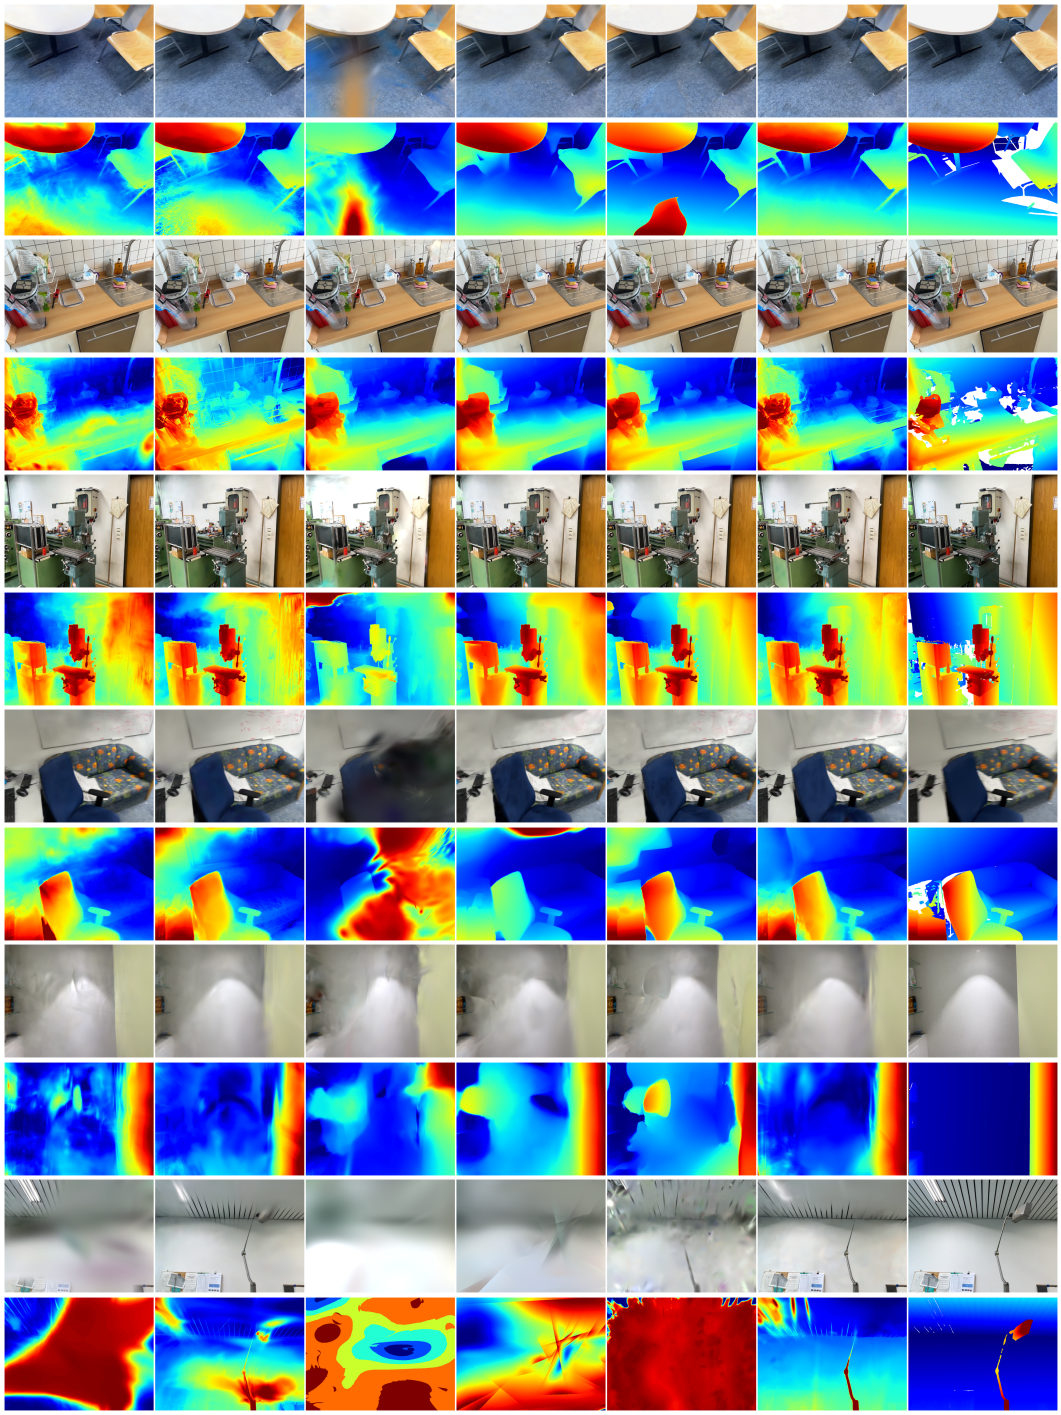

Figure 11: **Novel view synthesis and depth –** Qualitative results on ScanNet++ iPhone dataset show our superior performance in both image quality and depth estimation in novel views. Note the limitation of the quality of Gaussian Splatting based methods for textureless surfaces, which makes the plane fitting procedure less robust.

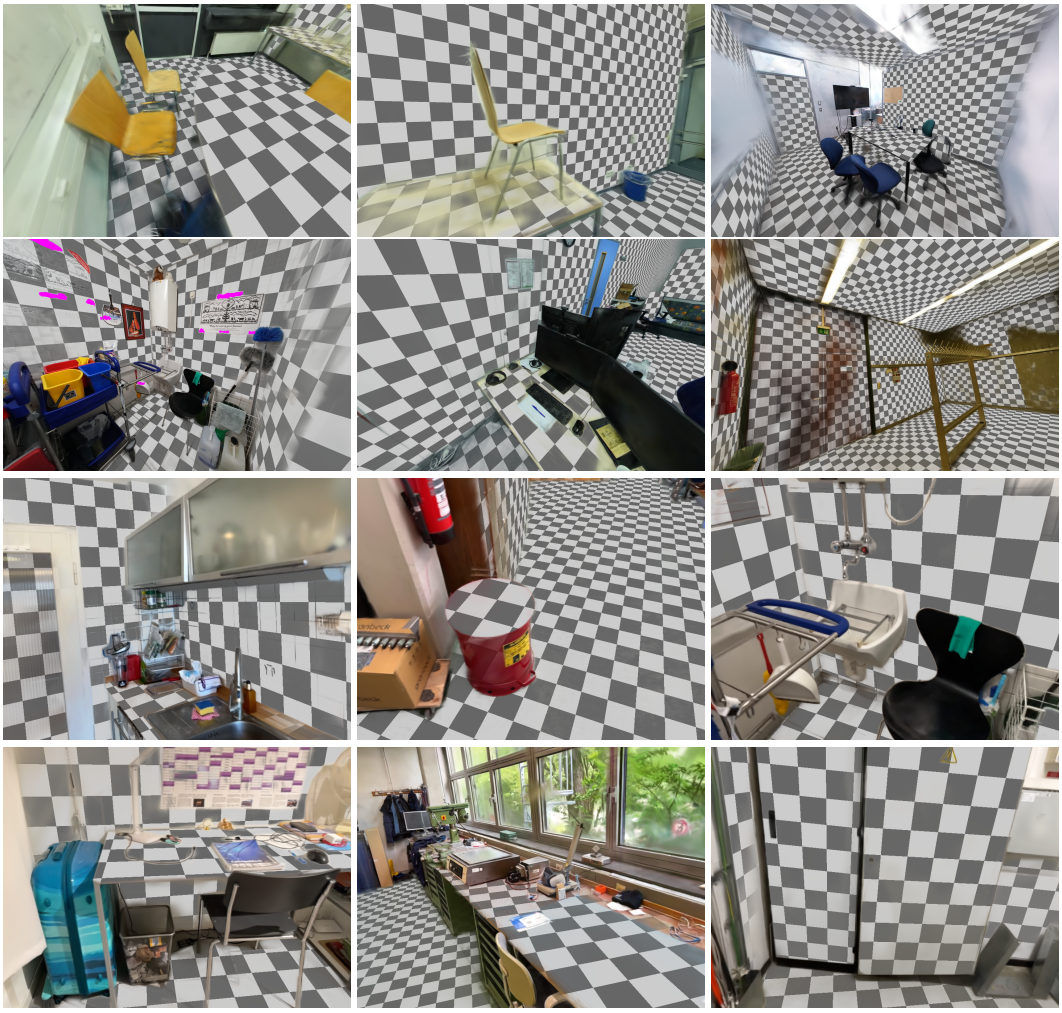

Figure 12: We provide visualizations of our output planes on the rendered test views of ScanNet++ DSLR streams (top 2 rows) and iPhone stream (bottom 2 rows). Pink markings are due to the anonymization of the original ScanNet++ dataset. While some planar surfaces are missed due to lack of manual 2D planar mask annotation, the captured planes are reconstructed faithfully.

Masked ground truth meshes For the planar mesh extraction task, we only consider planes with annotated segmentation masks from the ground truth mesh, as the 2D plane segmentation task is orthogonal to our method. To identify the relevant subset of planes, we unproject points from the ground truth depth maps that correspond to each plane according to its segmentation mask. We then fit a plane to each resulting point cloud using RANSAC and compile these fitted planes into a set S. We match planes from the ground truth mesh to those in set S by applying two criteria: the normal cosine distance must be less than 0.99 to at least one plane in S, and the distance between their closest points must be less than 0.1. Doing this allows for computational efficiency and increased robustness to missing or undersegmented planes in the input 2D annotations.

Code We release our code[3] publicly for reproducibility purposes and to facilitate future research in this area. We base our code on the 3DGS-MCMC paper [13] and additionally use SAMv2 [29], and PlaneRecNet [25] to generate masks. The baselines are evaluated using their official released code [7, 6, 16, 17, 13, 8, 9]. We further utilize AirPlanes [9] code to compute meshing metrics.

[3] `https://github.com/theialab/3dgs-flats`

F Hyperparameters settings

We use $\sigma_{\perp}$ and $\sigma_{\parallel}$ as hyper-parameters that control the stochastic re-location. These parameters are chosen depending on the metric scale of the dataset, and are defined in millimeters. For both datasets we used $\sigma_{\perp} = 0.01$ and $\sigma_{\parallel} = 0.3$. We observe that setting $\lambda_{\text{mask}} = 0.1$, yields best results empirically. For regularizers, we use λ_{TV}=0.1, λ_{scale}=0.01 and $\lambda_{\text{opacity}} = 0.01$ following [10] and [13]. We use the same scheduling policy for learning plane origin and normal (rotation) as for the Gaussian means the vanilla 3DGS. All experiments were conducted on a single A6000 ADA GPU, with 46GB memory. The method runs for approximately 1 hour for a single ScanNet++/ScanNetV2 scene, which is comparable to PGSR [16], the second best method for geometric quality according to our experiments and 1.5× longer than 3DGS-MCMC [13], the best method for Novel View Synthesis. The training time is increased due to the RANSAC overhead and block-coordinate descent optimization of planar parameters. Additionally, mesh extraction takes ~ 3 minutes and SAM mask propagation is on average 7 minutes long, depending on the scene type. We believe that the training time can be reduced in future work with addition of customized CUDA kernels.