

# QuickSplat: Fast 3D Surface Reconstruction via Learned Gaussian Initialization

Yueh-Cheng Liu   Lukas Höllein   Matthias Nießner   Angela Dai

Technical University of Munich



Figure 1. QuickSplat performs surface reconstruction of large indoor scenes from multi-view images as input. We learn strong priors for initialization of gaussian splatting optimization for surface reconstruction, as well as for per-iteration joint densification and gaussian updates. This results in high-quality mesh geometry that more accurately models flat wall structures as well as object details, while optimizing significantly faster than baselines.

## Abstract

Surface reconstruction is fundamental to computer vision and graphics, enabling applications in 3D modeling, mixed reality, robotics, and more. Existing approaches based on volumetric rendering obtain promising results, but optimize on a per-scene basis, resulting in a slow optimization that can struggle to model under-observed or textureless regions. We introduce QuickSplat, which learns data-driven priors to generate dense initializations for 2D gaussian splatting optimization of large-scale indoor scenes. This provides a strong starting point for the reconstruction, which accelerates the convergence of the optimization and improves the geometry of flat wall structures. We further learn to jointly estimate the densification and update of the scene parameters during each iteration; our proposed densifier network predicts new Gaussians based on the rendering gradients of existing ones, removing the needs of heuristics for densification. Extensive experiments on large-scale indoor scene reconstruction demonstrate the superiority of our data-driven optimization. Concretely, we accelerate runtime by 8x, while decreasing depth errors by 48% in comparison to state of the art methods.

## 1. Introduction

Surface reconstruction of large, real-world scenes is a key problem in computer vision and graphics. Reconstructing high-quality surfaces is essential to many applications, such as creating effective virtual environments that support physical reasoning, enabling accurate simulations, and imbuing robots with crucial knowledge for navigation and interaction. In particular, achieving both high fidelity as well as efficient and fast reconstruction for large scenes remains a difficult problem.

Recently, 3D Gaussian Splatting (3DGS) [27] achieves photorealistic novel-view-synthesis from multi-view images as input. It parameterizes a scene as a set of Gaussian primitives that are initialized from SfM and subsequently optimized and densified using reconstruction losses and a differentiable volumetric renderer. Subsequent works extend 3DGS to also obtain accurate surface reconstructions [9, 22, 25, 46, 55]. However, these methods typically optimize each scene separately, i.e., many iterations of gradient descent are required, which can take over 30 minutes for large-scale indoor scenes [52]. Additionally, the surface is only optimized from the observed input images, but capturing sufficiently many diverse images remains challenging for large scenes. The resulting geometry may thus contain missing or deformed regions where there is less view cov-

erage or texture information.

To this end, we propose a novel generalized prior for 3D surface reconstruction. It combines high-fidelity reconstructions based on 2D Gaussian Splatting (2DGS) [25] with the advantages of learning reconstruction priors from data. Concretely, we improve the efficiency of key elements of the optimization: initialization, Gaussian updates, and densification. This allows us to drastically accelerate the per-scene optimization time (Fig. 1, right). Our priors also guide the optimization towards high-quality indoor-scene geometry and thus overcome limitations stemming from insufficient observations or textureless regions (e.g., floating artifacts or non-straight wall geometry). In comparison, our reconstructions show higher quality than state-of-the-art baselines (Fig. 1, left).

We learn several sparse 3D CNN-based networks that jointly produce Gaussian parameters from the input posed multi-view images. Our initializer densifies the input SfM point cloud, which enables completing large holes in unobserved or textureless regions of a scene. We then propose a novel densification-optimization scheme, that grows new Gaussians and predicts update vectors for existing Gaussians. Similar to gradient-descent optimizers, we iteratively improve the Gaussians by repeating this scheme multiple times. Finally, we extract the surface from the converged 2D Gaussian primitives using TSDF fusion from rendered depth maps. We demonstrate that our method accurately reconstructs large-scale, real-world indoor environments with arbitrary many views as input. Extensive experiments on the ScanNet++ dataset [52] verify that QuickSplat reconstructs higher-quality geometry 8x faster than baselines.

To summarize, our contributions are:

- We propose a learned, generalized initializer network, that leverages scene priors to create effective Gaussian initializations for more efficient and accurate 3D surface reconstruction optimization, especially in under-observed or textureless areas of a scene.
- We employ learned, generalized priors in a *densification-optimization* loop, that jointly predict 2DGS updates and new densification locations. This heuristic-free optimization converges significantly faster and obtains more consistent geometry for large-scale scene reconstructions.

## 2. Related Work

### 2.1. Novel view synthesis

Novel view synthesis (NVS) has received significant attention in recent years [1, 3, 23, 27, 33, 34, 42]. NeRF [33] renders photo-realistic images by optimizing an MLP-based scene representation with differentiable volumetric rendering. Later, explicit or hybrid representations improved on optimization runtime [7, 20, 34, 40, 48]. 3DGS [27] further enables rendering at real-time rates by rasterizing explicit

Gaussian primitives. Recent methods improve the per-scene optimization speed of 3DGS by changing the underlying differentiable rasterizer, optimizer, or densification strategy [15, 18, 24, 31]. Others learn a data-prior for sparse-view reconstruction and predict the Gaussian primitives from a feed-forward network [6, 8, 11, 47, 56, 59]. We similarly train a data-prior that reconstructs Gaussian primitives in a feed-forward fashion. However, we are not limited to a sparse-view setting and focus on reconstructing better surface geometry.

### 2.2. 3D reconstruction

Reconstructing the 3D surfaces from multi-view image observations is a long-standing goal in computer vision. Classic or neural approaches based on multi-view geometry reconstruct point clouds or mesh geometry in a multi-step pipeline based on feature matching, triangulation, and fusion [5, 13, 21, 38, 43, 50, 53, 57]. Recently, NVS-based methods were extended to model accurate geometry and allow extracting surfaces after training through the Marching Cubes algorithm [22, 25, 29, 35, 44, 51, 55]. Our method lies in between these two directions. We similarly train a neural network to predict surface geometry faster than optimization-based methods. By formulating learned priors for the initialization, densification, and optimization updates of a 2DGS scene representation, we achieve improved surface reconstruction with fast runtimes.

### 2.3. Meta learning

Predicting 3D surfaces with neural networks in a feed-forward fashion typically means a single feed-forward pass produces the output. In contrast, meta learning models the iterative optimization process with neural networks [2, 19, 28, 45]. Inspired by this, we frame surface reconstruction in an iterative optimization pipeline where a neural network produces the update steps. Recent methods that model implicit functions with coordinate-based networks successfully leverage meta learning to improve the efficiency of their optimizations [39, 41]. Most related to ours is G3R [10], which reconstructs 3DGS primitives from multi-view RGB and dense LiDAR observations. Their optimizer network iteratively refines the parameters of the 3D Gaussians, that are initialized from the LiDAR scan of the environment. We leverage a similar optimizer, but additionally learn initializer and densifier networks. This allows us to reconstruct surfaces, even for sparsely observed regions of multi-view image input.

## 3. Method

Our method reconstructs the surface of large-scale indoor scenes from posed images as input. Specifically, we predict the attributes of 2D Gaussians [25] with a novel network architecture in a feed-forward fashion (Sec. 3.1). First, we

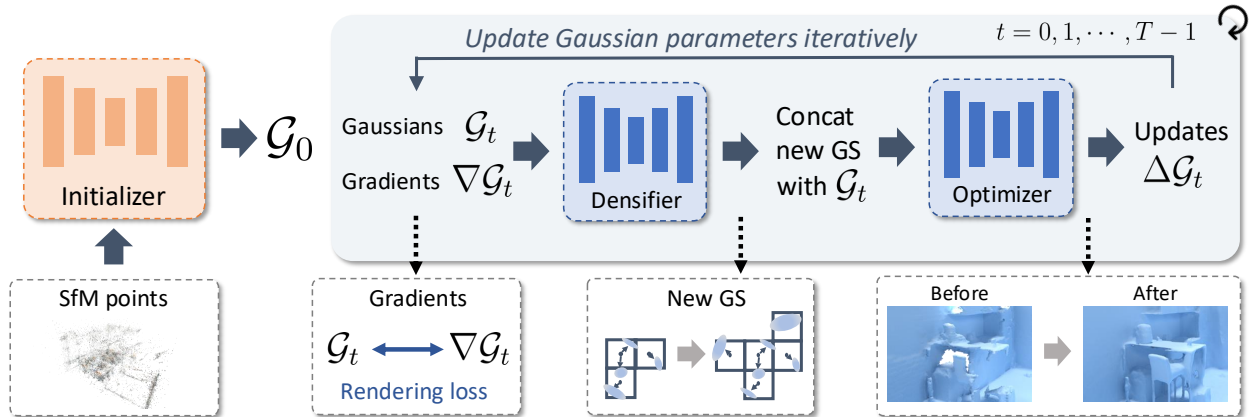


Figure 2. **Method overview.** From the SfM points of input multi-view images, our initializer network predicts an initial set of Gaussians  $\mathcal{G}_0$ . We then learn priors to improve the Gaussians during a series of optimization update steps in an iterative fashion. First, we calculate the current rendering gradients  $\nabla\mathcal{G}_t$  using all the training images. The densifier network then predicts additional Gaussians around the existing set of Gaussians  $\mathcal{G}_t$ . Finally, the optimizer network predicts an update vector  $\Delta\mathcal{G}_t$ , that we apply to create  $\mathcal{G}_{t+1}$ .

initialize the scene representation from the SfM point cloud using our initializer network (Sec. 3.2). Then, we incrementally improve the Gaussians in our proposed *densification-optimization* loop (Sec. 3.3). Concretely, we predict update vectors of scene parameters with an optimization network and grow new Gaussians with a densifier. We summarize our method in Fig. 2.

### 3.1. Surface Representation

We adopt 2D Gaussian Splatting (2DGS) as our scene representation [25], which allows us to model high-quality surfaces. Concretely, we use a set of Gaussian primitives  $\mathcal{G} = \{\mathbf{g}_i\}_{i=1}^N$  where each Gaussian  $\mathbf{g}_i \in \mathbb{R}^{14}$  is parameterized by its 3D position, scale, and rotation, a scalar opacity and diffuse RGB color. To render a scene from a given viewpoint, 2DGS computes the ray-splat intersection between all primitives and rays originating from the camera’s image plane. Then, a pixel color  $c$  for ray  $x$  is rendered using alpha-blending along the depth-sorted list of splats:

$$c(x) = \sum_{i \in \mathcal{N}} c_i \alpha_i T_i, \quad \text{with } T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (1)$$

where  $c_i$  is the color of the  $i$ -th splat along the ray,  $T_i$  is the accumulated transmittance, and  $\alpha_i = o_i \mathbf{g}_i^{2D}(\mathbf{u}(x))$  the product of opacity  $o_i$  and the 2D Gaussian value at the intersection point  $\mathbf{u}(x)$ . The Gaussians can be optimized with rendering losses against observed image colors  $C$ :

$$\mathcal{L}_c(x) = 0.8 \|c(x) - C\|_1 + 0.2(1 - \text{SSIM}(c(x), C)). \quad (2)$$

Regularizers like normal, or distortion loss [4] are also applied in addition to  $\mathcal{L}_c$ . After optimization, a mesh surface is extracted by rendering depth maps from all training views and running TSDF fusion.

We propose to predict  $\mathcal{G}$  with neural networks instead of optimizing the primitives directly with gradient descent. To this end, we align the Gaussian primitives with a grid structure to facilitate predictions, similar to [30, 36]. Specifically, we discretize the scene into a *sparse* set of 3D voxels of size  $v_d$  and assign a latent feature  $\mathbf{f} \in \mathbb{R}^{64}$  to each voxel. Our networks predict the features of these sparse voxels. We then decode them into  $v_g$  Gaussian primitives with a small MLP. Next, we explain the architecture of our networks in more detail.

### 3.2. Initialization Prior

The first step in our method is to create an initialization of all Gaussians  $\mathcal{G}$ . Following 2DGS, we initialize the primitives from the SfM point cloud and voxelize them to align with our grid structure. This results in a sparse set of Gaussians, that does not yet represent a continuous surface, but rather contains many empty regions (e.g., SfM struggles to reconstruct points in textureless areas like walls). We thus instead increase surface density by predicting additional primitives around the existing ones. Concretely, we train an initializer network  $\theta_I$  that predicts additional voxel features. Inspired by SGNN [14], this network comprises sparse 3D convolutions in an encoder-decoder architecture. Through a series of dense CNN blocks in the bottleneck followed by upsampling layers, the density of sparse voxels is gradually increased. An occupancy head acts as threshold to determine if a voxel should be allocated in the next higher resolution or not. In contrast to SGNN, which produces sparse voxel outputs, we employ a decoder MLP to interpret the densified voxel latent features as output Gaussian primitives. This employs different activation functions for each Gaussian attribute. The position  $\mathbf{g}_c \in \mathbb{R}^3$  is defined

relative to the voxel center  $v_c \in \mathbb{R}^3$  as:

$$\mathbf{g}_c = v_c + \mathbf{R}(2\sigma(x) - 1) \quad (3)$$

where  $\mathbf{R}=4v_d$  defines the radius around the voxel in which the Gaussian primitive can live and  $\sigma$  denotes the sigmoid function. Inspired by the reparameterization trick in PixelSplat [6], the opacity  $\mathbf{g}_o \in \mathbb{R}$  is the occupancy after the last upsampling layer. This allows rendering losses to backpropagate to the upsampling layers, controlling which points to keep in the last upsampling process. We utilize common functions for the remaining parameters to convert the outputs into the desired value ranges. Specifically, we utilize the softplus function for the scale parameters, sigmoid for the colors and we normalize the rotation quaternion vector.

We supervise the initializer network with multiple losses. First, the rendering loss  $\mathcal{L}_c$  (Eq. (2)) provides supervision about the quality of the predicted Gaussians. However, it does not provide signal for empty regions. In other words, the initializer should densify the Gaussians and close any holes, but the rendering loss does not backpropagate to empty voxels. To this end, we also supervise the occupancy of the voxel grid against ground truth geometry. We train the network on scenes from ScanNet++ [52], which contain mesh geometries from laser scans. This allows us to compute an occupancy loss  $\mathcal{L}_{occ}$  before every upsampling layer of the SGNN architecture. It comprises a binary cross-entropy loss on the occupancy of each voxel. This provides additional signal *where* new voxels should be allocated. Furthermore, we calculate a depth loss  $\mathcal{L}_d$ , which measure the L1 distance between rendered depth and depth of the mesh for all training views. Additionally, to make the 2D Gaussian “disk” more aligned with the ground-truth surfaces, we supervise the per-Gaussian normals  $\mathbf{n}_g$  against the mesh geometry. The normal of a primitive is defined as the direction perpendicular to 2D disk. We compare this against the mesh normals  $\mathbf{n}_m$  using cosine similarity:

$$\mathcal{L}_n = 1 - \mathbf{n}_g^T \mathbf{n}_m \quad (4)$$

The total loss for initializer training is then  $\mathcal{L}(\theta_I) = \mathcal{L}_c + \mathcal{L}_d + \mathcal{L}_{occ} + 0.01\mathcal{L}_n + 10\mathcal{L}_{dist}$  where  $\mathcal{L}_{dist}$  denotes the distortion loss [4, 25].

### 3.3. Iterative Gaussian Optimization

The initializer network predicts denser Gaussians from the SfM points as input. While the network is supervised with a rendering loss, its input is only the geometry, i.e., the (colored) SfM points. To further improve the surface reconstruction, we aim to also include information about the quality of the current Gaussian primitives in the input.

Concretely, we render the training images and compute the gradients of the rendering loss Eq. (2) with respect to the latent voxel features, accumulated across multiple views.

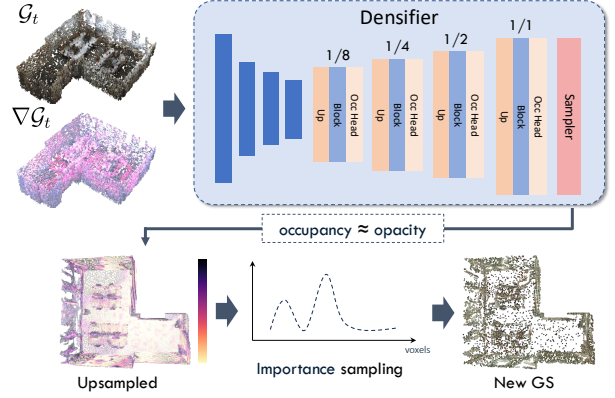


Figure 3. **Importance sampling of densified Gaussians.** Top: the densifier network predicts a pool of additional voxel features in an encoder-decoder architecture from the current Gaussians and their gradients as input. Bottom: to enable tractable memory during training, we select a subset of new Gaussians to densified. We apply importance sampling weighted by the occupancy predictions of the upsampled voxels. Since these occupancy values are then being used as the opacity of the densified Gaussians, this design encourages it to sample more “solid” Gaussians that are then passed along for the optimization updates during training.

For the ease of explanation, we denote this quantity, with a slight abuse of notation, as  $\nabla\mathcal{G}$ . This gradient contains the signal indicating how the Gaussians should be adjusted to improve.

**Optimizer** Inspired by G3R [10], we learn an optimizer network  $\theta_O$ , that predicts updates for all voxels from  $\{\mathcal{G}, \nabla\mathcal{G}\}$  as input. In other words, we learn the function  $f_{\theta_O}(\mathcal{G}_t, \nabla\mathcal{G}_t, t) = \Delta\mathcal{G}_t$ , where  $\Delta\mathcal{G}_t$  denotes the predicted update for all latent voxel features. Similar to G3R [10], we frame this process over multiple timesteps  $t$ , i.e., we iteratively calculate  $\nabla\mathcal{G}_t$ , predict  $\Delta\mathcal{G}_t$ , and update our representation as  $\mathcal{G}_{t+1} = \mathcal{G}_t + \Delta\mathcal{G}_t$ . We utilize a sparse 3D UNet architecture [12] for  $\theta_O$ . We normalize its output to lie within  $[-1, 1]$  to ensure that the updates do not overshoot.

**Densifier** Our optimizer network predicts updates on a sparse voxel grid. Although the sparse SfM point cloud gets densified by our initializer, it may still contain holes. These holes in the surface reconstruction may never be filled by  $\theta_O$ , since its predicted updates can only move Gaussians around the voxel centers, but never allocate new voxels. Empirically, we find that these holes notably reduce surface reconstruction quality (see Tab. 2).

To this end, we introduce another learnable component, the densifier network  $\theta_D$ , that predicts additional voxel features in free space. It follows the design of the initializer network  $\theta_I$  with the following key differences. First, we provide  $\nabla\mathcal{G}$  and  $t$  as additional inputs, i.e., we learn the function  $f_{\theta_D}(\mathcal{G}_t, \nabla\mathcal{G}_t, t) = \hat{\mathcal{G}}_t$ , where  $\hat{\mathcal{G}}_t$  denotes the pre-

dicted, additional voxel features. This enables using regions with large gradients to inform possible locations for densification. Second, we constrain the locations of additional voxels to be neighboring to the existing ones, by using no dense blocks in the bottleneck. We empirically find this more effective, as it better constrains where additional Gaussians should be grown.

Lastly, instead of predicting arbitrary number of new Gaussians, we perform importance sampling (see Fig. 3) to control their number and memory usage during training. The key idea is to prioritize selecting new “solid” Gaussians to grow among other candidates generated by the decoder, since they contribute more to the surface geometry, which means selecting Gaussians with higher opacity values. By leveraging the reparameterization trick (*i.e.*, interpreting occupancy as opacity), we can sample additional voxels weighted by the occupancy prediction after the last upsampling layer. We define the number of voxels that can be added in the current iteration as  $n(t)=s/(2^t)$  with  $s=20K$ . In other words, we densify more for earlier timesteps and then gradually reduce the number of new voxels. During test time, we simply select top  $n(t)$  voxels from the occupancy prediction.

The benefit of this method is that it avoids designing a densification strategy based on heuristics [17, 27, 31, 37]. By training the densifier network end-to-end with the optimizer, we instead learn to map the current state of Gaussians and their gradients into new, high-contribution Gaussians.

**Densification-Optimization Loop** We utilize both networks,  $\theta_D$  and  $\theta_O$ , in our proposed densification-optimization loop, that grows and improves the latent voxel features over multiple timesteps. First, we decode the voxels into Gaussian parameters, render all training images, and calculate  $\nabla\mathcal{G}_t$ . Then, the densifier predicts additional voxel positions  $\hat{\mathcal{G}}_t$ . We concatenate the existing and novel voxel features as  $\bar{\mathcal{G}}_t=\mathcal{G}_t\cup\hat{\mathcal{G}}_t$  and initialize the gradient of novel voxels with zeros to similarly obtain  $\nabla\bar{\mathcal{G}}_t$ . The optimizer then predicts the update  $\Delta\bar{\mathcal{G}}_t$  and we obtain  $\mathcal{G}_{t+1}=\bar{\mathcal{G}}_t+\Delta\bar{\mathcal{G}}_t$ .

**End-to-End Training** We jointly train the densifier and optimizer networks in our second training stage. The initializer network  $\theta_I$  remains frozen during this stage. We similarly calculate  $\mathcal{L}_{\text{occ}}(\theta_D)$  on the intermediate upsampling layers of the densifier. Additionally, we decode the updated voxel features  $\mathcal{G}_{t+1}$  into Gaussians and calculate the rendering loss  $\mathcal{L}(\theta_O)=\mathcal{L}_c+\mathcal{L}_d+10\mathcal{L}_{\text{dist}}$ . We run the densification-optimization loop for  $T=5$  timesteps and calculate the losses after each timestep. Similarly to G3R [10], we detach the gradient of the losses for the subsequent timesteps, *i.e.*, we optimize them separately.

## 4. Experiments

**Dataset** We train and evaluate our model on the ScanNet++ dataset [52]. We utilize 902 indoor scenes for training, after filtering out some scenes with incomplete wall structures or very large bounding extents. We train on the undistorted DSLR images at  $360\times 540$  resolution and double the resolution during evaluation to process images at  $720\times 1080$ . We evaluate our method on 20 unseen test scenes and report averaged metrics.

**Implementation Details** We set  $v_d=4\text{cm}$  and predict  $v_g=2$  Gaussians per voxel in all our experiments. The initializer and densifier network use four up/downsampling layers, and the optimizer UNet architecture follows G3R [10]. In total the networks have around 68M parameters. We set the learning rate to  $1e-4$  and train the networks for 3 days on a single Nvidia RTX A6000. In each iteration, we accumulate the gradients  $\nabla\mathcal{G}$  from 100 training images. After running our iterative optimization for  $t=5$  timesteps, we optionally refine the Gaussians for another 2000 steps of gradient descent (without adaptive density control). We denote results of our method as “w/ opt” that use this refinement and as “w/o opt” if they do not use it.

**Baselines** We compare our method with several recent 3D surface reconstruction approaches: SuGaR [22], 2DGS [25], GS2Mesh [46], PGSR [9], and MonoSDF [54]. SuGaR optimizes 3DGS [27] and regularizes the Gaussians to align with surfaces, followed by mesh extraction using Poisson surface reconstruction. Similarly, GS2Mesh also begins by optimizing 3DGS, then renders stereo pairs and predicts depth maps using an off-the-shelf stereo depth estimator (*i.e.*, DLNR [58]). 2DGS optimizes and densifies flat 2D Gaussians using gradient descent in a per-scene optimization lasting 30K iterations. PGSR renders unbiased depth maps from flattened 3D Gaussians and introduces both single-view and multi-view regularization losses to improve geometric reconstruction. MonoSDF leverages monocular depth and normal priors from a pre-trained model [16] to optimize the signed distance field (SDF) using differentiable volumetric rendering.

**Metrics** To evaluate the quality of the reconstructed geometry, we measure the error between rendered depth and the ground-truth depth maps of ScanNet++ testing frames. We calculate the absolute error, as well as the accuracy within different thresholds (2cm, 5cm, 10cm). We also calculate the Chamfer distance between the predicted and ground-truth mesh vertices. We crop the predicted vertices outside of ground-truth bounding box to prevent from penalizing false negative predictions outside of windows. Additionally, we report the optimization runtime in seconds.

Method	Abs err↓	Acc (2cm)↑	Acc (5cm)↑	Acc (10cm)↑	Chamfer↓	Time↓
SuGaR [22]	0.2061	0.1157	0.2774	0.4794	0.2078	3130s
2DGS [25]	0.1127	0.4021	0.6027	0.7422	0.2420	1796s
PGSR [9]	0.2325	0.4795	0.6407	0.7496	0.2228	2593s
GS2Mesh [46]	0.1212	0.4028	0.6039	0.7406	0.2012	973s
MonoSDF [54]	<b>0.0569</b>	<u>0.5774</u>	<u>0.8006</u>	<u>0.8850</u>	<u>0.1450</u>	>10h
Ours (w/o opt)	0.0732	0.5263	0.7674	0.8583	0.1461	<b>26s</b>
Ours (w/ opt)	<u>0.0578</u>	<b>0.5783</b>	<b>0.8035</b>	<b>0.8887</b>	<b>0.1347</b>	<u>124s</u>

Table 1. **Quantitative comparison against baselines.** We compare the quality and optimization runtime of our reconstructed surfaces against baseline methods, and show averaged results on the test scenes in ScanNet++ [52]. Both our method without post-training (“w/o opt”) and with additional SGD iterations (“w/ opt”) obtain better geometry while achieving orders of magnitude faster runtime.

#	Initializer	Densifier	Optimizer	Abs err↓	Chamfer↓	#Gaussians
(a)	-	-	✓	0.1332	0.2881	47k
(b)	occ only	-	✓	0.0897	0.2095	276k
(c)	occ only	✓	✓	0.0844	0.2038	361k
(d)	✓	-	✓	0.0581	0.1374	184k
(e)	✓	✓	✓	<b>0.0578</b>	<b>0.1347</b>	251k

Table 2. **Ablation study.** We ablate the impact of our learned priors for initialization, densification, and optimization updates. Only using our optimizer network does not increase the number of Gaussians and thus struggles to model continuous surfaces (a). Densifying the SfM point cloud, instead of predicting initial Gaussian parameters results in less accurate geometry (“occ only”). The densifier further increases the number of Gaussians, which helps to further improve surface quality.

#### 4.1. Comparison to State of the Art

Experimental results on ScanNet++ are shown in Tab. 1 and Fig. 4. In general, our proposed QuickSplat achieves better performance: it reconstructs scenes with cleaner structures and flat surfaces that matches the ground truth compared to the baselines while maintaining similar level of details. Most importantly, via the learned initializer and densifier-optimizer, our method converges much faster.

Compared to SuGaR [22], 2DGS [25], and PGSR [9], which are optimized per scene using rendering losses, our method additionally leverages learned geometry priors from data. This is particularly useful in indoor environments, which often contain large textureless regions (e.g., white walls). Methods that rely solely on photometric errors may struggle in these areas (see Fig. 4 for examples), and therefore produce curved surfaces and floating artifacts. In contrast, our learned initializer network predicts a relatively dense Gaussian initialization, effectively addressing the issue of missing SfM points in textureless regions.

Similar to our approach, GS2Mesh [46] also utilizes geometry priors learned from data (*i.e.*, from a pretrained stereo estimator). However, since synthesized views from 3DGS may not be as realistic as actual images, a noticeable domain gap exists for stereo depth estimators when the renderings contain noise. This can lead to high-frequency artifacts in the reconstructed geometry (see Tab. 1). In con-

trast, our method learns scene priors directly in the Gaussian representation space, allowing it to perform more robustly under such conditions.

Our method matches MonoSDF in accuracy while running substantially faster. Moreover, unlike MonoSDF—which sometimes produces overly smooth surfaces and misses fine structures (e.g., the ladder in Fig. 5) when the monocular depth or normal estimator fails to capture thin geometry—our approach better preserves fine details.

After the iterative Gaussian optimization with our learned prior networks (w/o opt), QuickSplat can further improve the geometry quality by running short iterations of post-training (w/ opt). It is especially beneficial for improving fine details, which results in 5% increase in Acc (2cm). Nevertheless, even without additional per-scene optimization, we significantly reduce the runtime while maintaining state-of-the-art performance.

#### 4.2. Ablations

Our method combines three data prior networks in an iterative fashion (Fig. 2). We demonstrate the importance of each individual component.

**Only optimizer network** Without the initializer and the densifier, the optimizer network relies solely on sparse SfM point clouds. Since it has no ability to generate additional



Figure 4. **Qualitative comparison against baselines.** We show top-down views of reconstructed mesh geometries (with and without vertex colors) in comparison to the ground-truth meshes of ScanNet++ [52]. Our method more accurately models flat wall structures and objects details, while producing fewer floating geometry artifacts.

Gaussians, it is hard to accurately represent the surface geometry. As shown in Tab. 2 (a), the number of Gaussians remains too low, making the performance significantly worse. An example of the output can be seen in Fig. 6 (a), where

it struggles to model reasonable surfaces (e.g., the walls are bent and noisy).

**Importance of the initializer** To evaluate the effect of the initializer, we train a network that predicts dense point



Figure 5. **Qualitative comparison between MonoSDF [54] and ours.** Our QuickSplat achieves faster reconstruction and retains more fine details. For example, MonoSDF fails to reconstruct the ladder since the pretrained depth/normal estimator misses it.




Figure 6. **Visualization of ablations.** (a) Without our initializer and densification priors during optimization, surface reconstruction of untextured regions such as walls is challenging due to the lack of SfM points. (b) With our initializer but without densification predictions, reconstruction improves, but maintains smaller-scale artifacts around regions with holes or excess geometry. (c) Our full model produces a robust reconstruction of the scene while achieving fast optimization runtimes.

clouds from SfM points and supervise it only with  $\mathcal{L}_{\text{occ}}$  (referred to as “occ only” in Tab. 2). In contrast, our full initializer predicts dense points with 2DGS attributes directly. Since the occupancy-only initializer has no knowledge over the Gaussian representation and their rendering quality, the performance is worse than our proposed initializer (see (b) vs. (d) and (c) vs. (e) in Tab. 2).

**Importance of the densifier** We compare the reconstruction performance with and without the proposed densifier. As can be seen in Fig. 6 (b) vs. (c), even with the dense initialization, the densifier is able to fill the remaining holes and predict new Gaussians adaptively, based on the current state of the Gaussian representation and their gradients. Therefore, it further improves the geometry details, such as the stairs at the corner. As shown in Tab. 2 ((b) vs. (c) and (d) vs. (e)), this results in a noticeable performance improvement.

**Gaussian attributes initialization** We further investigate the effect of predicting only a subset of the Gaussian attributes in our initializer. It can be seen in Tab. 3 that initializing each attribute of the Gaussian representation contributes to improvements. Therefore, we used all four additional Gaussian attributes for QuickSplat (next to predicting the position).

Color	Opacity	Scales	Rotation	Abs err↓	Chamfer↓
✓				0.0627	0.1500
✓	✓			0.0600	0.1421
✓	✓	✓		0.0590	0.1402
✓	✓	✓	✓	<b>0.0578</b>	<b>0.1347</b>

Table 3. **Initializer output ablation study.** We evaluate the impact of predicting different Gaussian attributes from the SfM point cloud with our initializer network. Predicting all attributes results in the best final surface reconstruction quality.

### 4.3. Limitations

Our method accelerates optimization runtime by 8x and obtains more accurate surface reconstructions from posed images in comparison to baselines. However, some drawbacks remain. First, our method struggles with mirror reflections, since the photometric loss encourages to reconstruct the reflected geometry behind the mirror, which leads to noisy artifacts. Second, we assume static environments and therefore cannot reconstruct dynamic scenes (e.g., people walking inside of a room). Lastly, even though we significantly reduce optimization runtime, our method does not yet reconstruct in real-time, but could be integrated with recent SLAM-based approaches [26, 32, 49].

### 5. Conclusion

We have presented QuickSplat, which learns several data priors to perform surface reconstruction optimization of large indoor scenes from multi-view images as input. By framing the optimization with learned prior networks for initialization, densification, and optimization updates of 2D Gaussian splats, we significantly accelerate surface reconstruction speed by 8x in comparison to baselines. Furthermore, we demonstrate that incorporating data-priors helps reduce artifacts caused by insufficiently many observations or textureless areas, that typically occur in large-scale scene reconstructions. That is, our initializer network densifies the input SfM points by exploiting learned geometry priors (e.g., flat wall structures). Then, our proposed *densification-optimization* loop refines the Gaussian attributes through a series of predicted update steps. Overall, we believe that the ability to utilize data-priors for fast and state-of-the-art reconstructions will open up further research avenues and make surface reconstructions more practical across a wide range of real-world applications.

### Acknowledgements

This work was supported by the ERC Starting Grant SpatialSem (101076253), the ERC Consolidator Grant Gen3D (101171131), and the German Research Foundation (DFG) Research Unit “Learning and Simulation in Visual Computing.”

## References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 696–712. Springer, 2020. 2
- [2] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29, 2016. 2
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5855–5864, 2021. 2
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5470–5479, 2022. 3, 4
- [5] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. Transformerfusion: Monocular rgb scene reconstruction using transformers. *Advances in Neural Information Processing Systems*, 34:1403–1414, 2021. 2
- [6] David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19457–19467, 2024. 2, 4
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pages 333–350. Springer, 2022. 2
- [8] Anpei Chen, Haofei Xu, Stefano Esposito, Siyu Tang, and Andreas Geiger. Lara: Efficient large-baseline radiance fields. In *European Conference on Computer Vision*, pages 338–355. Springer, 2024. 2
- [9] 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. 1, 5, 6
- [10] Yun Chen, Jingkang Wang, Ze Yang, Sivabalan Manivasagam, and Raquel Urtasun. G3r: Gradient guided generalizable reconstruction. In *European Conference on Computer Vision*, pages 305–323. Springer, 2024. 2, 4, 5
- [11] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*, 2024. 2
- [12] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019. 4
- [13] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)*, 36(4): 1, 2017. 2
- [14] Angela Dai, Christian Diller, and Matthias Nießner. Sg-nn: Sparse generative neural networks for self-supervised scene completion of rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 849–858, 2020. 3
- [15] Sankeerth Durvasula, Adrian Zhao, Fan Chen, Ruofan Liang, Pawan Kumar Sanjaya, and Nandita Vijaykumar. Distwar: Fast differentiable rendering on raster-based rendering pipelines. *arXiv preprint arXiv:2401.05345*, 2023. 2
- [16] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021. 5
- [17] Guangchi Fang and Bing Wang. Mini-splatting: Representing scenes with a constrained number of gaussians. In *European Conference on Computer Vision*, pages 165–181. Springer, 2024. 5
- [18] Guofeng Feng, Siyan Chen, Rong Fu, Zimu Liao, Yi Wang, Tao Liu, Zhilin Pei, Hengjie Li, Xingcheng Zhang, and Bo Dai. Flashgs: Efficient 3d gaussian splatting for large-scale and high-resolution rendering. *arXiv preprint arXiv:2408.07967*, 2024. 2
- [19] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017. 2
- [20] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5501–5510, 2022. 2
- [21] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. 2
- [22] Antoine Guédon and Vincent Lepetit. Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5354–5363, 2024. 1, 2, 5, 6
- [23] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6):1–15, 2018. 2
- [24] Lukas Höllein, Aljaž Božič, Michael Zollhöfer, and Matthias Nießner. 3dgs-lm: Faster gaussian-splatting optimization with levenberg-marquardt. *arXiv preprint arXiv:2409.12892*, 2024. 2

- [25] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pages 1–11, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [26] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. [8](#)
- [27] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. [1](#), [2](#), [5](#)
- [28] Ke Li and Jitendra Malik. Learning to optimize neural nets. *arXiv preprint arXiv:1703.00441*, 2017. [2](#)
- [29] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023. [2](#)
- [30] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20654–20664, 2024. [3](#)
- [31] Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Francisco Vicente Carrasco, Markus Steinberger, and Fernando De La Torre. Taming 3dgs: High-quality radiance fields with limited resources. *arXiv preprint arXiv:2406.15643*, 2024. [2](#), [5](#)
- [32] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18039–18048, 2024. [8](#)
- [33] 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. *Communications of the ACM*, 65(1):99–106, 2021. [2](#)
- [34] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. [2](#)
- [35] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5589–5599, 2021. [2](#)
- [36] Barbara Roessle, Norman Müller, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, Angela Dai, and Matthias Nießner. L3dg: Latent 3d gaussian diffusion. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. [3](#)
- [37] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in gaussian splatting. In *European Conference on Computer Vision*, pages 347–362. Springer, 2024. [5](#)
- [38] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. [2](#)
- [39] Vincent Sitzmann, Eric Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. *Advances in Neural Information Processing Systems*, 33:10136–10147, 2020. [2](#)
- [40] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5459–5469, 2022. [2](#)
- [41] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P Srinivasan, Jonathan T Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2846–2855, 2021. [2](#)
- [42] 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 (EG STAR 2022)*, 2022. [2](#)
- [43] Fangjinhua Wang, Qingtian Zhu, Di Chang, Quankai Gao, Junlin Han, Tong Zhang, Richard Hartley, and Marc Pollefeys. Learning-based multi-view stereo: a survey. *arXiv preprint arXiv:2408.15235*, 2024. [2](#)
- [44] 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. *arXiv preprint arXiv:2106.10689*, 2021. [2](#)
- [45] Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gomez Colmenarejo, Misha Denil, Nando Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *International conference on machine learning*, pages 3751–3760. PMLR, 2017. [2](#)
- [46] Yaniv Wolf, Amit Bracha, and Ron Kimmel. GS2Mesh: Surface reconstruction from Gaussian splatting via novel stereo views. In *European Conference on Computer Vision (ECCV)*, 2024. [1](#), [5](#), [6](#)
- [47] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Depthspat: Connecting gaussian splatting and depth. *arXiv preprint arXiv:2410.13862*, 2024. [2](#)
- [48] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Pointnerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5438–5448, 2022. [2](#)
- [49] Chi Yan, Delin Qu, Dan Xu, Bin Zhao, Zhigang Wang, Dong Wang, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19595–19604, 2024. [8](#)
- [50] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo.

- In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018. [2](#)
- [51] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. [2](#)
- [52] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [53] Zehao Yu and Shenghua Gao. Fast-mvsnet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1949–1958, 2020. [2](#)
- [54] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35:25018–25032, 2022. [5](#), [6](#), [8](#)
- [55] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics (TOG)*, 43(6):1–13, 2024. [1](#), [2](#)
- [56] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. Gs-lrm: Large reconstruction model for 3d gaussian splatting. In *European Conference on Computer Vision*, pages 1–19. Springer, 2024. [2](#)
- [57] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Nerfusion: Fusing radiance fields for large-scale scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5449–5458, 2022. [2](#)
- [58] Haoliang Zhao, Huizhou Zhou, Yongjun Zhang, Jie Chen, Yitong Yang, and Yong Zhao. High-frequency stereo matching network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1327–1336, 2023. [5](#)
- [59] Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu. Long-lrm: Long-sequence large reconstruction model for wide-coverage gaussian splats. *arXiv preprint arXiv:2410.12781*, 2024. [2](#)