

Sparfels: Fast Reconstruction from Sparse Unposed Imagery

Shubhendu Jena*, Amine Ouasfi*, Mae Younes, Adnane Boukhayma
 Inria, Univ. Rennes, CNRS, IRISA



Figure 1. Examples of reconstruction meshes obtained within **3 minutes** from sparse pose-free images via our method. We used here **6 (left) and 3 (right) input images** respectively from scenes of datasets MVImgNet [97] and BlendedMVS [92].

Abstract

We present a method for *Sparse* view reconstruction with surface *element* splatting that runs within 3 minutes on a consumer grade GPU. While few methods address sparse radiance field learning from noisy or unposed sparse cameras, shape recovery remains relatively underexplored in this setting. Several radiance and shape learning test-time optimization methods address the sparse posed setting by learning data priors or using combinations of external monocular geometry priors. Differently, we propose an efficient and simple pipeline harnessing a single recent 3D foundation model. We leverage its various task heads, notably point maps and camera initializations to instantiate a bundle adjusting 2D Gaussian Splatting (2DGS) model, and image correspondences to guide camera optimization midst 2DGS training. Key to our contribution is a novel formulation of splatted color variance along rays, which can be computed efficiently. Reducing this moment in training leads to more accurate shape reconstructions. We demonstrate state-of-the-art performances in the sparse uncalibrated setting in reconstruction and novel view benchmarks based on established multi-view datasets. Code will be made available at <https://shubhendu-jena.github.io/Sparfels-web/>

1. Introduction

3D reconstruction is a long standing problem at the intersection of computer vision, graphics and machine learning. It is core to a myriad of downstream applications in virtual and augmented reality, autonomous driving, medical imaging, film and animation, to name a few. Classical optimization based techniques such as structure from motion and

multi-view stereo [70, 71], shape from visual cues such as shading or silhouette [28, 43], photometric stereo [83] *etc.*, have been popular for decades in this regard. More recently, deep learning has been progressively disrupting 3D vision and top-down approaches to vision more broadly. Among the latest additions, large vision and foundation models (*e.g.* DUST3R [79], MAST3R [44]) can offer robustness and fast feed-forward inference compared to classical photogrammetry [70, 71]. Given a few raw uncalibrated images of a scene, they are capable of inferring reasonable relative poses and a competitive coarse 3D structure. These predictions can lack precision and details, though, and cannot enable stand-alone novel view synthesis.

Furthermore, state-of-the-art data-prior free scene specific learning methods for reconstruction can deliver highly detailed geometry and impressive novel view synthesis capabilities from dense calibrated image grids. While implicit and some grid shape representations (*e.g.* NeuS, VolSDF, HF-NeuS, Voxurf [78, 80, 85, 94]) trained via differentiable volumetric rendering (NeRF [55]) have been powerful tools for this task for a while, Gaussian Splatting [40] based alternatives emerged with comparable novel view and reconstruction [29, 100] quality, within a fraction of the training time, and with orders of magnitude higher interactive rendering frame rates. These methods typically fail however under uncalibrated and sparse training images. Contemporary literature attempts to hedge against the underconstrained nature of the sparse view problem through regularization with a multitude of deep geometric priors (monocular depth [18, 76, 104], normal [98], point cloud [67] *etc.*), or training new priors from scratch [50, 67, 87]. Such strategies can prove to be cumbersome, impractical or costly in terms of data, memory or compute.

In this paper, we propose to marry the best of both worlds efficiently — namely pre-existing large 3D vision feed-forward models (MASt3R) and fast test-time optimization

*Equal contribution.

(2DGS) — with the aim of **addressing the underexplored problem of 3D geometric reconstruction from sparse unposed images within minutes** (less than 3 on average in our experiments) at minimal cost.

We propose to exploit a single pre-existing foundation model (MASt3R [44]) to the fullest to bootstrap our reconstruction, without needing to load or deploy any additional deep neural networks. Given a few sparse color images, MASt3R delivers a coarse point cloud and initial camera poses. We use these to initialize a bundle adjusting 2D Gaussian Splatting (2DGS) optimization via differentiable render-and-compare, where surface elements are instantiated at the point samples, oriented initially according to their normals. Differently from the closest competition [21, 22], we harness the MASt3R [44] produced image pair correspondences to guide camera optimization during 2DGS training. This is achieved via a splatted depth reprojection error loss. The resulting camera alignment improvement proves crucial for obtaining accurate shape representations.

By analyzing the surface element splatting from a statistical moment perspective, rendered pixels can be seen as color expectation over fragments following the ray termination probability. Based on this, we derive an expression of splatted color variance along the ray, and propose a novel loss tasked with reducing it. This results in improved reconstruction metrics and qualitative sharper mesh details as witnessed by our ablation studies.

The integration of these elements yields state-of-art performance in sparse pose-free 3D reconstruction, obtained in very short times on mid-range GPUs. We evaluate our method using established reconstruction and novel view datasets [1, 5, 42, 92, 97], showcasing both superior geometric reconstruction, camera pose estimation and novel view synthesis w.r.t. comparable state-of-the-art baselines.

Evaluation of 3D shape reconstruction in the posed setting relies traditionally on mesh Chamfer distance to ground-truth, as used in multi-view reconstruction benchmarks like DTU [1]. This is a mesh pose and scale sensitive metric. In the unposed setting, reconstructions are obtained naturally up-to a rigid transformation. RANSAC powered alignments between predicted and ground-truth cameras have been shown to be acceptable for pre-evaluation reconstruction mesh alignment in the dense view setting [7], and good enough for camera pose evaluation for sparse sets [75]. However, we found these camera based alignments to be too inaccurate in the sparse setting to enable reliable mesh pre-alignments for Chamfer distance evaluation. Hence, inspired by the rigid-transform invariant evaluation in [79], we compare the reconstructed mesh depth and normal rasterizations to the groundtruth in screen space in our evaluation, and show considerable improvements compared to the competition in the standard DTU benchmark




Figure 2. 3D Foundation models meet differentiable perspective accurate Splatting for fast reconstruction from pose-free sparse images. We bootstrap a bundle adjusting 2DGS with MASt3R [44] 3D point clouds and initial camera predictions. Besides standard 2DGS photometric reconstruction and geometric regularization losses, we leverage a camera reprojection error loss using MASt3R’s correspondences and a novel splatted color variance reduction loss.

(Tab.1,2,3). Further qualitative comparison in Blended-MVS [92], MVImgNet [97], Mip-NeRF 360 [4] illustrate the quality of our reconstructions.

2. Related work

NVS and 3D Reconstruction Novel view synthesis (NVS) aims to generate unseen views from input images [3, 54]. Neural Radiance Fields (NeRF) [55] revolutionized this field by representing 3D scenes as continuous implicit functions, though challenges remain in efficiency and quality [4–6, 66, 72, 73]. Inspired by the success of other explicit representations (*e.g.* meshes [33, 38, 77] and point clouds [2, 20]), recent advancements introduced primitive-based scene representations [13, 40, 90], notably 3D Gaussian Splatting (3DGS) [40], which uses anisotropic 3D Gaussians with differentiable splatting for efficient high-quality reconstruction. This has sparked extensions for surface reconstruction [29, 100, 101], neural representations [51], multi-resolution modeling [24, 99], and feed-forward methods [10, 12, 23, 27, 89]. In parallel, implicit representations model 3D geometry as continuous functions [16, 37, 39, 57, 59]. Methods like IDR [93], NeuS [78], and VolSDF [94] enhance geometric accuracy but often suffer from slow rendering. In contrast, 2D Gaussian Splatting (2DGS) [29] offers real-time rendering while improving view consistency for surface reconstruction through planar Gaussian primitives.

Sparse-View NVS and 3D Reconstruction Traditional NeRFs and GS methods require carefully captured images and SfM preprocessing with COLMAP [69], which could limit practical applications. The few-shot setup has been addressed for both reconstruction from Point Cloud [11, 62, 64, 65] and images. Various approaches address the challenge of limited views through regularization techniques, including depth supervision [17, 46, 58, 76], foundation model supervision [32, 84, 88], specialized architectures [91, 96], while SparseCraft [95] uses learning-free

MVS cues for SDF regularization. For Gaussian Splatting, methods like FSGS [104] and SparseGS [86] incorporate external priors. Generalizable methods address sparse-view challenges by pretraining on large datasets, enabling feed-forward inference. Inputs include images (e.g. [10, 12, 35, 45, 49]) and 3D Points (e.g. [9, 30, 52, 60, 61, 63]). For surface reconstruction from color images, approaches like SparseNeuS [50], VolRecon [68], ReTR [47], UfoRecon [56], GeoTransfer [34], and Spurries [67] leverage various mechanisms such as using image priors to form multi-view cost volumes, using transformer architectures to enhance reconstruction from limited views, as well as using reconstruction priors trained on sparse point clouds.

Pose-Free NVS and 3D Reconstruction Pose-free methods eliminate the need for accurate camera calibration. NeRFmm [82], BARF [48], SCNeRF [36], and GARF [15] jointly optimize camera parameters and scene representations. SPARF [75] and PoRF [7] address optimization with noisy poses, though often requiring lengthy training. Recent work like Nope-NeRF [8], Lu-NeRF [14], LocalRF [53], and CF-3DGS [25] leverage depth information but typically assume dense video sequences. InstantSplat [21, 22] addresses these limitations by using MAST3R [44] for initialization and jointly optimizing camera poses and 3D models. Our work builds on these foundations but focuses specifically on using 2D Gaussian Splatting for efficient 3D surface reconstruction from sparse, uncalibrated views. We initialize with MAST3R [44] and refine through joint optimization of Gaussian parameters and camera poses, incorporating a novel variance-based color loss for improved detail preservation. Unlike previous approaches requiring dense sequences, known intrinsics, or long optimization times, our method efficiently handles sparse, uncalibrated images while achieving high-quality reconstructions.

3. Method

Given a few unposed color images $\{I^i\}_{i=1}^N$, our goal is to recover a 3D triangle mesh $\mathcal{S}$ and render novel views of the observed scene efficiently. We achieve this by fitting a bundle adjusting 2DGS [29] model to the images of the scene, bootstrapped by point map and correspondence feedforward predictions from a 3D foundation model (MASt3R [44]) and carefully designed regularizations. At test time, mesh $\mathcal{S}$ can be extracted using the TSDF algorithm in Open3D [103] using the converged 2DGS model rendered depth, as recommended in [29]. We detail the workings of our method below.

3.1. Preliminaries

2D Gaussian Splatting Our method builds upon 2D Gaussian Splatting (2DGS) [29], which represents 3D scenes using planar disk primitives defined within local tan-

gent planes. Unlike volumetric 3D Gaussians [40], 2DGS is a surface-based representation that improves multi-view consistency. Each 2D Gaussian is defined as:

$$G(\mathbf{p}) = \exp\left(-\frac{u(\mathbf{p})^2 + v(\mathbf{p})^2}{2}\right) \quad (1)$$

where $u(\mathbf{p})$ and $v(\mathbf{p})$ are local UV space coordinates of a pixel $\mathbf{p}$ in the primitive’s tangent plane. This ray-splat intersection can be expressed as a function of the camera parameters and splat attributes differentially. Each primitive is parameterized by the center of the Gaussian in world coordinates $\mathbf{x} \in \mathbb{R}^3$, a quaternion representing the orientation of the primitive $\mathbf{q} \in \mathbb{R}^4$, scale factors for the two dimensions in the local UV tangent plane $\mathbf{s} \in \mathbb{R}^2$, opacity $\alpha \in \mathbb{R}$ controlling the transparency of the primitive, and Spherical harmonic coefficients $\mathbf{c}_{\text{shk}} \in \mathbb{R}^3 | k = 1, \dots, n$ encoding view-dependent appearance. Rotation $\mathbf{q}$ defines a local coordinate frame $[\mathbf{u}, \mathbf{v}, \mathbf{n}]$, $\mathbf{n}$ being the normal vector orthogonal to the primitive’s plane, and $\mathbf{u}$ and $\mathbf{v}$ spanning the tangent plane.

For rendering, 2DGS uses alpha-compositing for volumetric rendering similarly to 3DGS. Given a set of primitives along a ray sorted from front to back $G_1, G_2, \dots, G_K$, the final color C is computed as:

$$C(\mathbf{p}) = \sum_{i=1}^K \alpha'_i \mathbf{c}_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (2)$$

where $\alpha'_i = \alpha_i G(\mathbf{p})$ is the modified opacity of the i -th primitive and $\mathbf{c}_i$ is its color, computed from the spherical harmonic coefficients based on the viewing direction.

MASt3R MASt3R [44] is a foundation model for 3D vision that extends DUS3R [79] with more powerful feature extraction and matching capabilities. It tackles the multi-view stereo reconstruction task without requiring prior camera calibration or poses.

Given a pair of images $I^1, I^2 \in \mathbb{R}^{W \times H \times 3}$, MASt3R outputs two corresponding pointmaps $\mathbf{P}_{1,1}, \mathbf{P}_{2,1} \in \mathbb{R}^{W \times H \times 3}$ with associated confidence maps $\mathbf{O}_{1,1}, \mathbf{O}_{2,1} \in \mathbb{R}^{W \times H}$. Each pointmap assigns a 3D point to every pixel in the image. Importantly, both pointmaps are expressed in the same coordinate frame of the first image I^1 .

The architecture of MASt3R consists of a Siamese ViT encoder that processes both images independently, Transformer decoders that exchange information via cross-attention, and regression heads that output pointmaps and confidence maps. MASt3R is trained using a confidence-weighted regression loss:

$$\mathcal{L}_{\text{conf}} = \sum_{v \in \{1,2\}} \sum_i \mathbf{O}_{v,1}^i \mathcal{L}_{\text{reg}}(v, i) - \alpha \log \mathbf{O}_{v,1}^i, \quad (3)$$

where $\mathcal{L}_{\text{reg}}(v, i)$ measures the Euclidean distance between predicted and ground-truth 3D points for pixel i , and α is a regularization parameter.

Besides pointmaps, MAST3R also provides dense pixel-wise correspondences between the input views. These correspondences establish reliable geometric constraints across different viewpoints, which we leverage in our optimization framework.

3.2. Scene Initialization

Global Geometry Alignment Given a set of input images $I^1, I^2, \dots, I^N$, we first construct a connectivity graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where each vertex $v \in \mathcal{V}$ represents an image, and each edge $e = (n, m) \in \mathcal{E}$ indicates a pair of images with shared visual content. For each edge in the graph, we run MAST3R to obtain pairwise pointmaps and correspondences. To align these pairwise predictions into a globally consistent reconstruction, we follow the optimization approach from MAST3R:

$$\min_{\chi, P, \sigma} \sum_{e \in \mathcal{E}} \sum_{v \in e} \sum_{i=1}^{HW} \mathbf{O}_{v,e}^i |\chi_i^v - \sigma_e P_e \mathbf{P}_{v,e}^i|, \quad (4)$$

where χ_i^v is the globally aligned pointmap for view v at pixel i , $P_e \in \mathbb{R}^{3 \times 4}$ the rigid transformation for edge e , $\sigma_e > 0$ is a scaling factor for edge e , $\mathbf{O}_{v,e}^i$ the confidence value for pixel i in view v for edge e , and finally $\mathbf{P}_{v,e}^i$ the original pointmap from MAST3R. This optimization produces globally aligned pointmaps $\{\chi^{v*}\}_{v=1}^N \in \mathbb{R}^{H \times W \times 3}$ and camera parameters (intrinsics and extrinsics for each view v). We stabilize the estimated focal lengths by averaging them across all training views: $\bar{f} = \frac{1}{N} \sum_{v=1}^N f_i^*$, assuming a single-camera setup.

Primitive Initialization We leverage the merged globally aligned pointmap to initialize 2D Gaussian primitives with suitable orientation, which is crucial for accurately representing surface geometry. For each point $\mathbf{p}_i$ in the globally aligned pointmap, we compute a local surface normal using Principal Component Analysis (PCA). First, we determine the k -nearest neighbors $\{\mathbf{p}_j\}_{j=1}^k$ and compute the covariance matrix $\mathbf{C} = \frac{1}{k} \sum_{j=1}^k (\mathbf{p}_j - \bar{\mathbf{p}})(\mathbf{p}_j - \bar{\mathbf{p}})^T$, where $\bar{\mathbf{p}}$ is the centroid of the neighborhood. We then perform eigen-decomposition on $\mathbf{C}$ to obtain the eigenvectors $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ and corresponding eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$, sorted such that $\lambda_1 \geq \lambda_2 \geq \lambda_3$. The eigenvector $\mathbf{v}_3$, associated with the smallest eigenvalue, is chosen as the normal direction, i.e., $\mathbf{n} = \mathbf{v}_3$. To construct an orthonormal basis, we initialize an arbitrary vector $\mathbf{a}$ orthogonal to $\mathbf{n} = (n_x, n_y, n_z)$, defined as $\mathbf{a} = (-n_y, n_x, 0)^T$. The tangent vector is then computed as $\mathbf{u} = \frac{\mathbf{n} \times \mathbf{a}}{\|\mathbf{n} \times \mathbf{a}\|}$ and the bitangent as $\mathbf{v} = \mathbf{n} \times \mathbf{u}$. Finally, the local reference frame $[\mathbf{u}, \mathbf{v}, \mathbf{n}]$ is used to define a rotation matrix, which can be converted to a quaternion $\mathbf{q}_i$ for efficient representation. Next, we initialize each

2D Gaussian primitive with position $\mathbf{x}_i$ directly from the pointmap, rotation $\mathbf{q}_i$, scale s_i set proportionally to the local pointmap density, and opacity α_i set to a default value (e.g. 0.8). Color for the lowest spherical harmonic band (DC component) is initialized from the original image colors projected onto the point, with higher-order coefficients set to zero. This initialization preserves the surface structure from the MAST3R reconstruction while enabling efficient optimization with 2D Gaussian primitives.

3.3. Joint Optimization

Correspondence Loss Using the dense correspondences provided by MAST3R, and inspired by SPARF [75], we formulate a multi-view correspondence loss that enforces geometric consistency across different viewpoints:

$$\mathcal{L}_{\text{corr}} = w_{p_n, p_m} \rho(p_m - \pi(P_m^{-1} P_n \pi^{-1}(p_n, d_n))), \quad (5)$$

where (p_n, p_m) is a correspondence pair with $p_n \in [0, 1]^2$ a pixel in view n and p_m its corresponding pixel in view m . $w_{p_n, p_m} \in [0, 1]$ is the confidence of the correspondence. ρ represents the Huber loss [31] function. $\pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is the projection operator mapping 3D points to pixel coordinates. $\pi^{-1} : \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}^3$ is the back-projection operator mapping pixel coordinates and depth to 3D points. d_n is the 2DGS model splatted depth at pixel p_n . P_n and P_m represent the camera parameters for views n and m . This loss enforces that corresponding pixels across views should project to the same 3D point when back-projected using their respective depths and camera parameters, and thus helps guide our bundle adjustment to a good minima.

Photometric Loss Following 3DGS and 2DGS, we use a combination of $\mathcal{L}_1$ image reconstruction loss and structural similarity (SSIM) to optimize the rendering of our system:

$$\mathcal{L}_{\text{photo}} = (1 - \lambda) \cdot \mathcal{L}_1 + \lambda \cdot \mathcal{L}_{\text{SSIM}} + \mathcal{L}_{\text{reg}}. \quad (6)$$

Notice that 2DGS comes also with geometric regularization term $\mathcal{L}_{\text{reg}}$ binding pixel depth normals to splatted normals, and reducing depth distortion, as detailed in [29].

Variance Regularization Loss A key contribution of our method is the introduction of a splatted color variance regularization loss. The fundamental insight is that volume rendering can be interpreted as computing the expected value of colors sampled along a ray [74]:

$$C = \mathbb{E}_{\mathbf{t} \sim p(\mathbf{t})}[c(\mathbf{t})] = \sum_{i=1}^K \alpha_i \mathbf{c}_i \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (7)$$

where $c(\mathbf{t})$ represents the color of a point $\mathbf{t}$ along the ray, and the expectation is taken with respect to $p(t) = \sigma(t)T(t)$ with $\sigma(t)$ being the density at t and $T(t)$ representing the transmittance function.

We hypothesize that stabilizing and robustifying our training would lead to more robust and improved reconstructions. In order to robustify the training, we would like to control how the rendered color deviates from the ground truth color when the geometry is perturbed, *i.e.* given divergence measure $\mathcal{D}$, we want to minimize the following worst-case loss along the ray:

$$\mathcal{L}_r = \sup_{q: D(q,p) < \rho} |\mathbb{E}_{\mathbf{t} \sim q(\mathbf{t})}[c(\mathbf{t})] - C_{gt}|, \quad (8)$$

where ρ controls the strength of the perturbation. By leveraging theorem 2 in [19], this loss can be upper bounded by:

$$\mathcal{L}_1 + \eta \sqrt{\text{Var}_{\mathbf{t} \sim p(\mathbf{t})}[c(\mathbf{t})]}. \quad (9)$$

Consequently, we consider minimizing the variance of the color along the ray under the distribution $p(\mathbf{t})$ that can be expressed as follows:

$$\begin{aligned} \mathcal{L}_{\text{var}} &= \text{Var}_{\mathbf{t} \sim p(\mathbf{t})}[c(\mathbf{t})] \\ &= \mathbb{E}_{\mathbf{t} \sim p(\mathbf{t})}[(c(\mathbf{t}) - C)^2] = \mathbb{E}_{\mathbf{t} \sim p(\mathbf{t})}[c(\mathbf{t})^2] - C^2. \end{aligned} \quad (10)$$

For efficient implementation, we compute the first term using a modified CUDA kernel of the 2DGS splatting algorithm, where we render the square of colors in addition to the rendering of the colors. This allows for efficient forward and backward passes during optimization. Intuitively, reducing rendered color variance reduces uncertainty in the rendering, which leads to sharper and more multi-view coherent reconstructions. This reflects positively on the qualitative (Fig 5) and quantitative (Tab 4 and Fig 6) performance of our reconstructions. We note that in practice, loss $\mathcal{L}_{\text{var}}$ is averaged across pixels for a given view.

Training Objective Our total optimization objective combines the three loss terms:

$$\mathcal{L} = \lambda_{\text{photo}} \mathcal{L}_{\text{photo}} + \lambda_{\text{corr}} \mathcal{L}_{\text{corr}} + \lambda_{\text{var}} \mathcal{L}_{\text{var}}, \quad (12)$$

where λ_{photo} , λ_{corr} , and λ_{var} are weighting factors that balance the contribution of each term. We optimize all parameters simultaneously using the Adam [41] optimizer with a learning rate scheduling following [29]. Unlike multi-stage approaches [75], our single-stage optimization allows for faster convergence while maintaining high reconstruction quality. The optimization updates both camera extrinsics $\{P\}$ per input view and 2D Gaussian primitive parameters $\{\alpha, \mathbf{c}, \mathbf{q}, \mathbf{s}\}$ according to the gradient of the total loss.

4. Experiments

We evaluate our approach across 3D reconstruction, novel view synthesis, and camera pose estimation. We outline the experimental setup, including datasets, baseline methods, and implementation details, followed by quantitative and qualitative comparisons. Ablation studies further analyze the impact of key components.

4.1. Datasets and Experimental Configuration

Our method is tested on benchmark datasets [1, 5, 42, 92, 97]. Further implementation details and additional qualitative results are provided in the supplementary material.

Few-Shot 3D Reconstruction: We evaluate sparse 3D reconstruction on **DTU** [1], **BlendedMVS** [92], **MVImgNet** [97], and **MipNeRF360** [5]. For 3-view reconstruction on **DTU**, we follow SparseNeuS [50], evaluating 15 scenes with two sets of 3-view inputs, reporting the mean per scene. We also provide qualitative comparisons for 3, 6, and 12-view settings on selected scenes from **BlendedMVS** [92], **MVImgNet** [97], and **MipNeRF360** [5].

Few-Shot Novel View Synthesis: We conduct novel view synthesis (NVS) on **Tanks and Temples** [42] (8 scenes), **MVImgNet** [97] (7 scenes), and **MipNeRF360** [5] (9 scenes), following InstantSplat [21, 22]. Results are reported under sparse input settings (3, 6, and 12 views).

Camera Pose Estimation: Camera poses are evaluated on **Tanks and Temples** [42], **MVImgNet** [97], and **MipNeRF360** [5], following InstantSplat [21, 22]. Additionally, for few-shot reconstruction, we assess camera poses on **DTU** [1] using the same 3-view evaluation protocol as in the reconstruction task.

Evaluation Metrics. For 3D reconstruction quality, we employ the Absolute Relative Error (Rel) and Normal Consistency (NC). Rel measures the accuracy of depth estimation following DUST3R [79], while NC evaluates the geometric fidelity by measuring the alignment between predicted and ground truth surface normals. We apply these metrics to our final meshes. We compute ground truth normals from the gradients of the ground truth DTU depth maps.

For novel view synthesis, we use standard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [81], and Learned Perceptual Image Patch Similarity (LPIPS) [102]. We note that for this evaluation we follow InstantSplat[21, 22]: for novel test views without pre-computed camera poses, we perform an additional camera optimization step, which optimizes the test cameras while keeping the 2D Gaussian primitives fixed using the $\mathcal{L}_1$ photometric reconstruction loss.

For camera pose estimation, we report the Absolute Trajectory Error (ATE) as defined in [8, 21], using COLMAP-computed poses from dense views as ground-truth references.

Baselines. We compare our method against several state-of-the-art approaches:

- **NVS specific methods:** Pose-free methods NoPeNeRF [8], CF-3DGS [25], NeRFmm [82], SPARF [75] and InstantSplat [21, 22] (in both S and XL configurations). We also include comparisons with 3DGS [40] and




Figure 3. **Qualitative comparison on DTU dataset from 3 input images.**

FSGS [104], which use pre-computed camera parameters from Colmap [69].

- **Reconstruction-specific methods:** MAST3R [44], Colmap [69], UfoRecon [56], CasMVSNet [26], Spurries [67], SpaRP [87], SparseCraft [95], InstantSplatGOF (our implementation based on [22]) and InstantSplat2dgs [22]. Of these, similar to our approach, InstantSplatGOF (our implementation based on [22]) and InstantSplat2dgs [22] further refine cameras.

Implementation Details. Our framework is implemented in PyTorch. We set the optimization iterations to 1k for DTU dataset reconstruction experiments and between 2k-4k for other reconstruction and novel view synthesis experiments. For test-time optimization, camera parameters are jointly optimized for 1k iterations (while keeping the learned Gaussian parameters fixed). We initialize our pipeline with MAST3R [44] configured at a resolution of 512 pixels and use Open3D [103] to compute normals for the point cloud, which are used to initialize rotation parameters of the 2D Gaussians. All experiments are conducted on a single NVIDIA A6000 GPU. We set $\lambda_{\text{photo}} = 1.0$ and $\lambda_{\text{corr}} = 5 \times 10^{-5}$. The variance loss weight λ_{var} follows a cosine annealing schedule: $\lambda_{\text{var}} = 0.5 \left(1 + \cos\left(\frac{\pi \cdot t}{T}\right)\right)$, where t denotes the current training step and T the total number of steps. This schedule gradually decreases λ_{var} from 1.0 to 0.0, ensuring an adaptive regularization throughout training.

4.2. 3D Reconstruction Evaluation

For few-shot reconstruction, we follow the evaluation protocol from prior work [50, 56, 95], cleaning the generated meshes with masks from training views. We evaluate the Absolute Relative Error (Rel) and Normal Consistency (NC) between predicted depths/normals from the extracted meshes and the ground truth depths/normals from the DTU dataset [1].

Tab. 1 and 2 present the quantitative results on the DTU




Figure 4. **Qualitative comparison on BlendedMVS dataset from 3 input images.**

dataset for Rel and NC metrics, respectively. Fig 3, 4 and 7 show qualitative reconstruction comparisons, with additional qualitative comparisons with SpaRP [87] included in the supplementary material. Additionally, the supplementary material contains extended qualitative results, in the form of video comparisons demonstrating our reconstructions for several datasets. Our method achieves state-of-the-art performance, outperforming all competing approaches across the majority of scenes.

Several key observations can be made from these results. Pretrained generalizable methods (UfoRecon [56] and CasMVSNet [26]) generalize poorly when faced with noisy initial camera parameters from MAST3R [44], resulting in significantly worse performance. Despite being designed for sparse reconstruction, SparseCraft [95] struggles with the noisy camera poses and point clouds from MAST3R [44]. Spurries [67], which employs an SDF prior trained on ShapeNet, performs particularly poorly with noisy camera poses, resulting in worse mesh extractions. InstantSplat [21, 22] variants (InstSplatGOF and InstSplat2DGS) perform better than other competing approaches, but still produce overly smooth meshes lacking fine details.

Our method significantly outperforms these approaches due to several key advantages, including the variance regularization loss (Sec. 3.3) promoting sharper and more detailed reconstructions Fig 5, The correspondence loss (Sec. 3.3) and joint optimization framework effectively refine the initial noisy camera poses from MAST3R [44]. Our method’s ability to preserve fine geometric details is also evidenced by the superior Normal Consistency scores.

Table 3 presents camera pose estimation results measured by ATE. Our approach achieves significant improve-

Methods	scan24	scan37	scan40	scan55	scan63	scan65	scan69	scan83	scan97	scan105	scan106	scan110	scan114	scan118	scan122	↓ Mean
MASt3R [44]	2.49	20.48	12.02	5.39	8.86	6.3	4.71	9.01	7.70	10.33	5.86	2.50	5.0	4.56	4.95	7.34
UFORecon [56]	35.55	57.15	38.81	35.47	50.02	44.94	34.70	53.56	41.67	47.93	40.22	39.34	39.75	33.34	49.04	42.77
CasMVSNet [26]	9.76	38.00	37.03	17.30	19.39	24.10	6.57	29.57	11.84	14.92	12.38	8.38	3.93	9.87	13.44	17.10
Spurries [67]	2.96	52.15	58.42	20.25	14.66	19.24	4.97	12.44	7.50	10.64	7.36	6.26	38.61	25.84	44.19	21.70
InstSplatGOF [21]	9.54	23.14	13.74	15.62	13.30	12.33	11.19	12.21	14.69	15.61	13.61	12.30	15.31	12.90	12.05	13.84
Sparsecraft [95]	2.96	10.51	7.16	2.92	8.17	9.44	3.83	6.19	7.61	6.44	4.75	6.43	5.36	7.65	8.14	6.50
InstantSplat2DGS [22]	2.42	21.74	9.08	3.08	6.77	3.60	3.63	5.24	5.82	7.12	4.48	2.78	3.62	3.11	3.43	5.73
Ours	2.42	12.313	7.9	3.15	6.05	3.526	3.48	4.98	5.29	6.22	4.39	2.41	3.46	3.32	3.38	4.82

Table 1. **3-view reconstruction results on DTU dataset** The relative error (rel $\downarrow$) is reported for 15 scans. Best results are **bold**.

Methods	scan24	scan37	scan40	scan55	scan63	scan65	scan69	scan83	scan97	scan105	scan106	scan110	scan114	scan118	scan122	↑ Mean
MASt3R [44]	0.888	0.695	0.844	0.874	0.837	0.839	0.851	0.754	0.824	0.745	0.871	0.875	0.884	0.831	0.846	0.830
UFORecon [56]	0.362	0.278	0.506	0.469	0.277	0.372	0.383	0.361	0.362	0.338	0.424	0.366	0.332	0.402	0.339	0.371
CasMVSNet [26]	0.718	0.502	0.619	0.756	0.693	0.683	0.740	0.564	0.704	0.681	0.785	0.795	0.859	0.778	0.755	0.709
Spurries [67]	0.863	0.35	0.297	0.711	0.742	0.668	0.771	0.671	0.777	0.683	0.78	0.771	0.509	0.572	0.433	0.640
InstSplatGOF [21]	0.774	0.581	0.689	0.696	0.746	0.750	0.710	0.658	0.682	0.653	0.695	0.724	0.710	0.663	0.650	0.692
Sparsecraft [95]	0.841	0.679	0.827	0.864	0.728	0.714	0.834	0.756	0.746	0.765	0.837	0.629	0.801	0.761	0.735	0.768
InstSplat2DGS [22]	0.882	0.616	0.86	0.899	0.834	0.865	0.845	0.783	0.810	0.773	0.882	0.849	0.895	0.827	0.855	0.832
Ours	0.904	0.75	0.877	0.896	0.866	0.887	0.867	0.799	0.845	0.789	0.888	0.874	0.902	0.85	0.865	0.857

Table 2. **3-view reconstruction results on DTU dataset** Normal consistency (NC $\uparrow$) is reported for 15 scans. Best results are **bold**.

ments in camera pose estimation, with a 13.1% reduction in error compared to the best-performing baseline (InstSplat-GOF [21]). This demonstrates the effectiveness of our correspondence loss in refining the initial noisy camera poses from MASt3R [44].

4.3. Novel View Synthesis Evaluation

We evaluate our method’s capability for novel view synthesis on Tanks and Temples [42], MVImgNet [97], and MipNeRF360 [5] datasets. Quantitative and qualitative result comparisons on all datasets are included in the supplementary material.

4.4. Ablation Studies

We conduct ablation studies to analyze the contribution of each component in our method, including an experiment to study the impact of varying the number of input views.

Component Ablation These experiments were performed using 3 input unposed views on 5 testing scenes from the DTU dataset using two sets of views, following the protocols established in [50, 56]. Tab 4 shows the impact of incorporating different components into our pipeline. Fig 5 shows a qualitative ablation comparison. Camera optimization using bundle adjustment improves all metrics. Incorporating the correspondence loss further enhances this alignment, resulting in improved reconstructions. Splat orientation initialization with pointmap normals improves results further. Notice that our variance loss contributes to both the final metrics and the qualitative aspect of our results. The reconstructions with variance loss exhibit significantly sharper details and better preservation of fine structures, validating the quantitative improvements observed in our metrics.

Number of Views Ablation We progressively increase the number of views and perform pose-free reconstruction with our method using scan37 from DTU. We evaluate by computing the mean metrics across both test view sets. As




Figure 5. **Qualitative ablation of the component of our method, showing also the effect of variance loss.** Left blue box: w/o $\mathcal{L}_{var}$. Right red box: w $\mathcal{L}_{var}$. (3 input images).

shown in Figure 6, our method benefits favorably from increased views. The benefits of all components of our method are also maintained, if not increased. This includes our variance loss.

5. Running time

In our reconstruction experiments, we observed that the mean training time on DTU [1] for 1000 iterations (similar to InstantSplat-XL [21, 22]) was roughly 146s on a single Nvidia A6000 GPU. For a fair comparison, we also trained InstantSplat-XL [21, 22] on the same GPU and observed that the mean training time was around 128s. Hence, our training time is in the same range as InstantSplat [21, 22] while offering significantly higher reconstruction and novel-view quality.

Methods	scan24	scan37	scan40	scan55	scan63	scan65	scan69	scan83	scan97	scan105	scan106	scan110	scan114	scan118	scan122	↓ Mean
MAS3R [44]	0.0202	0.0689	0.0094	0.0203	0.0171	0.0392	0.0291	0.0294	0.0271	0.0326	0.0167	0.0210	0.0289	0.0192	0.0242	0.0269
InstSplatGOF [21]	0.0180	0.0536	0.0089	0.0136	0.0049	0.0273	0.0263	0.0239	0.0180	0.0323	0.0128	0.0238	0.0249	0.0175	0.0149	0.0214
InstSplat2DGS [22]	0.0192	0.0895	0.0090	0.0154	0.0123	0.0338	0.0294	0.0272	0.0242	0.0322	0.0166	0.0210	0.0274	0.0163	0.0175	0.0261
Ours	0.0063	0.0450	0.0069	0.0124	0.0175	0.0294	0.0198	0.0283	0.0146	0.0273	0.0082	0.0179	0.0242	0.0112	0.0101	0.0186

Table 3. **Results on DTU dataset** Results include ATE↓ metrics for 15 scans. Best results are **bold**.

Methods	rel ↓	NC ↑	ATE ↓
2DGS	7.93	0.785	0.0327
BA-2DGS	7.79	0.788	0.0284
BA-2DGS+ $\mathcal{L}_{\text{corr}}$	7.41	0.804	0.0256
BA-2DGS+ $\mathcal{L}_{\text{corr}}$ +Norm. init.	7.03	0.829	0.0237
BA-2DGS+ $\mathcal{L}_{\text{corr}}$ +Norm. init.+ $\mathcal{L}_{\text{var}}$	6.79	0.834	0.0238

Table 4. **Ablation on DTU dataset, adding one component at a time.** Results include rel and NC metrics. Best results are **bold**.




Figure 6. **Effect of increasing the number of input views on different metrics.**

6. Limitations and Future Work

Despite the promising results, several limitations remain, which present opportunities for future research. As foundation models keep on improving, better initializations could be available in the future. Furthermore, while our approach handles a variety of scenes across different datasets, performance may be limited for highly complex geometry, very thin structures, or semi-transparent objects (observed for the glass bottle in Fig 7 room scene in MipNeRF360 [5]), which are common challenges for surface based methods especially in the sparse setting. Additionally, our method assumes a static scene under consistent lighting. Extending it to more dynamic scenarios is one of the challenges we are looking forward to exploring.

7. Conclusion

We addressed the underexplored problem of shape reconstruction from sparse uncalibrated multi-view images. Our method leverages foundation models for initial camera and geometry estimation, followed by joint optimization of 2D Gaussian primitives and camera parameters. The introduction of a variance regularization loss significantly improves the detail and sharpness of reconstructions, while our correspondence matching based framework enhances



Figure 7. **Qualitative comparison on MVImgNet and MipNeRF360 datasets from 6 and 12 input images, respectively.**

camera pose accuracy. Extensive evaluations on multiple datasets demonstrate that our approach consistently outperforms state-of-the-art methods in reconstruction quality, novel view synthesis, and camera pose estimation.

References

- [1] Henrik Aanaes, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjarholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120(2):153–168, 2016. [2](#), [5](#), [6](#), [7](#)
- [2] 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](#)
- [3] Shai Avidan and Amnon Shashua. Novel view synthesis in tensor space. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1034–1040. IEEE, 1997. [2](#)
- [4] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5460–5469, 2022. [2](#)
- [5] 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. [2](#), [5](#), [7](#), [8](#)
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19697–19705, 2023. [2](#)
- [7] Jia-Wang Bian, Wenjing Bian, Victor Adrian Prisacariu, and Philip Torr. Porf: Pose residual field for accurate neural surface reconstruction. *arXiv preprint arXiv:2310.07449*, 2023. [2](#), [3](#)
- [8] Wenjing Bian, Zirui Wang, Kejie Li, Jia-Wang Bian, and Victor Adrian Prisacariu. Nope-nerf: Optimising neural radiance field with no pose prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4160–4169, 2023. [3](#), [5](#)
- [9] Alexandre Boulch and Renaud Marlet. Poco: Point convolution for surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6302–6314, 2022. [3](#)
- [10] 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](#), [3](#)
- [11] Chao Chen, Zhizhong Han, and Yu-Shen Liu. Unsupervised inference of signed distance functions from single sparse point clouds without learning priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [2](#)
- [12] Yuedong Chen, Haoifei 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. In *European Conference on Computer Vision*, pages 370–386. Springer, 2025. [2](#), [3](#)
- [13] Zhang Chen, Zhong Li, Liangchen Song, Lele Chen, Jingyi Yu, Junsong Yuan, and Yi Xu. Neurf: A neural fields representation with adaptive radial basis functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4182–4194, 2023. [2](#)
- [14] Zezhou Cheng, Carlos Esteves, Varun Jampani, Abhishek Kar, Subhansu Maji, and Ameesh Makadia. Lu-nerf: Scene and pose estimation by synchronizing local unposed nerfs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18312–18321, 2023. [3](#)
- [15] Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Gaussian activated neural radiance fields for high fidelity reconstruction and pose estimation. In *European Conference on Computer Vision*, pages 264–280. Springer, 2022. [3](#)
- [16] François Darmon, Bénédicte Bascle, Jean-Clément Devaux, Pascal Monasse, and Mathieu Aubry. Improving neural implicit surfaces geometry with patch warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6260–6269, 2022. [2](#)
- [17] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021. [2](#)
- [18] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. [1](#)
- [19] John Duchi, Peter Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach, 2018. [5](#)
- [20] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, 2017. [2](#)
- [21] Zhiwen Fan, Wenyan Cong, Kairun Wen, Kevin Wang, Jian Zhang, Xinghao Ding, Danfei Xu, Boris Ivanovic, Marco Pavone, Georgios Pavlakos, et al. Instantsplat: Unbounded sparse-view pose-free gaussian splatting in 40 seconds. *arXiv preprint arXiv:2403.20309*, 2(3):4, 2024. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [22] Zhiwen Fan, Kairun Wen, Wenyan Cong, Kevin Wang, Jian Zhang, Xinghao Ding, Danfei Xu, Boris Ivanovic, Marco Pavone, Georgios Pavlakos, et al. Instantsplat: Sparse-view sfm-free gaussian splatting in seconds. *arXiv preprint arXiv:2403.20309*, 2024. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [23] Zhiwen Fan, Jian Zhang, Wenyan Cong, Peihao Wang, Renjie Li, Kairun Wen, Shijie Zhou, Achuta Kadambi, Zhangyang Wang, Danfei Xu, et al. Large spatial model: End-to-end unposed images to semantic 3d. *arXiv preprint arXiv:2410.18956*, 2024. [2](#)
- [24] 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](#)

- [25] Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A Efros, and Xiaolong Wang. Colmap-free 3d gaussian splatting. *arXiv preprint arXiv:2312.07504*, 2023. 3, 5
- [26] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuoqun Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504, 2020. 6, 7
- [27] Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jisang Han, Jiaolong Yang, Chong Luo, and Seungryong Kim. PF3plat: Pose-free feed-forward 3d gaussian splatting. *arXiv preprint arXiv:2410.22128*, 2024. 2
- [28] Berthold KP Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. 1970. 1
- [29] 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
- [30] Jiahui Huang, Zan Gojcic, Matan Atzmon, Or Litany, Sanja Fidler, and Francis Williams. Neural kernel surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4369–4379, 2023. 3
- [31] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pages 492–518. Springer, 1992. 4
- [32] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5885–5894, 2021. 2
- [33] Shubhendu Jena, Franck Multon, and Adnane Boukhayma. Neural mesh-based graphics. In *European Conference on Computer Vision*, pages 739–757. Springer, 2022. 2
- [34] Shubhendu Jena, Franck Multon, and Adnane Boukhayma. Geotransfer: Generalizable few-shot multi-view reconstruction via transfer learning. *arXiv preprint arXiv:2408.14724*, 2024. 3
- [35] Shubhendu Jena, Shishir Reddy Vutukur, and Adnane Boukhayma. Sparsplat: Fast multi-view reconstruction with generalizable 2d gaussian splatting. *arXiv*, 2025. 3
- [36] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5846–5854, 2021. 3
- [37] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1251–1261, 2020. 2
- [38] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018. 2
- [39] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4287–4297, 2021. 2
- [40] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 1, 2, 3, 5
- [41] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [42] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 2, 5, 7
- [43] Aldo Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on pattern analysis and machine intelligence*, 16(2):150–162, 1994. 1
- [44] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. 1, 2, 3, 6, 7, 8
- [45] Qian Li, Franck Multon, and Adnane Boukhayma. Learning generalizable light field networks from few images. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3
- [46] Qian Li, Franck Multon, and Adnane Boukhayma. Regularizing neural radiance fields from sparse rgb-d inputs. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 2320–2324. IEEE, 2023. 2
- [47] Yixun Liang, Hao He, and Yingcong Chen. Retr: Modeling rendering via transformer for generalizable neural surface reconstruction. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [48] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5741–5751, 2021. 3
- [49] Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao, Wei Li, and Ziwei Liu. Fast generalizable gaussian splatting reconstruction from multi-view stereo. *arXiv preprint arXiv:2405.12218*, 2024. 3
- [50] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, pages 210–227. Springer, 2022. 1, 3, 5, 6, 7
- [51] 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. 2
- [52] Baorui Ma, Yu-Shen Liu, Matthias Zwicker, and Zhizhong Han. Surface reconstruction from point clouds by learning predictive context priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6326–6337, 2022. 3
- [53] Andreas Meuleman, Yu-Lun Liu, Chen Gao, Jia-Bin Huang, Changil Kim, Min H Kim, and Johannes Kopf. Progressively optimized local radiance fields for robust view

- synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16539–16548, 2023. 3
- [54] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019. 2
- [55] 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. 1, 2
- [56] Youngju Na, Woo Jae Kim, Kyu Beom Han, Suhyeon Ha, and Sung-Eui Yoon. Uforecon: Generalizable sparse-view surface reconstruction from arbitrary and unfavorable sets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5094–5104, 2024. 3, 6, 7
- [57] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2
- [58] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. *arXiv preprint arXiv:2112.00724*, 2021. 2
- [59] 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
- [60] Amine Ouasfi and Adnane Boukhayma. Few’zero level set’-shot learning of shape signed distance functions in feature space. In *ECCV*, 2022. 3
- [61] Amine Ouasfi and Adnane Boukhayma. Mixing-denoising generalizable occupancy networks. *3DV*, 2024. 3
- [62] Amine Ouasfi and Adnane Boukhayma. Few-shot unsupervised implicit neural shape representation learning with spatial adversaries. *arXiv preprint arXiv:2408.15114*, 2024. 2
- [63] Amine Ouasfi and Adnane Boukhayma. Robustifying generalizable implicit shape networks with a tunable non-parametric model. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [64] Amine Ouasfi and Adnane Boukhayma. Unsupervised occupancy learning from sparse point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21729–21739, 2024. 2
- [65] Amine Ouasfi, Shubhendu Jena, Eric Marchand, and Adnane Boukhayma. Toward robust neural reconstruction from sparse point sets. *arXiv preprint arXiv:2412.16361*, 2024. 2
- [66] Eric Penner and Li Zhang. Soft 3D Reconstruction for View Synthesis. *ACM Transactions on Graphics (TOG)*, 36(6): 1–11, 2017. 2
- [67] Kevin Raj, Christopher Wewer, Raza Yunus, Eddy Ilg, and Jan Eric Lenssen. Spurfies: Sparse surface reconstruction using local geometry priors. *arXiv preprint arXiv:2408.16544*, 2024. 1, 3, 6, 7
- [68] Yufan Ren, Tong Zhang, Marc Pollefeys, Sabine Süsstrunk, and Fangjinhua Wang. Volrecon: Volume rendering of signed ray distance functions for generalizable multi-view reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16685–16695, 2023. 3
- [69] 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, 6
- [70] 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. 1
- [71] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 1
- [72] Steven M Seitz and Charles R Dyer. Photorealistic Scene Reconstruction by Voxel Coloring, 2002. US Patent 6,363,170. 2
- [73] Pratul P Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting Lighting Volumes for Spatially-Coherent Illumination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8080–8089, 2020. 2
- [74] Andrea Tagliasacchi and Ben Mildenhall. Volume rendering digest (for nerf), 2022. 4
- [75] Prune Truong, Marie-Julie Rakotosaona, Fabian Manhardt, and Federico Tombari. Sparf: Neural radiance fields from sparse and noisy poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4190–4200, 2023. 2, 3, 4, 5
- [76] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. *arXiv preprint arXiv:2303.16196*, 2023. 1, 2
- [77] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. 2
- [78] 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. 1, 2
- [79] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 1, 2, 3, 5

- [80] Yiqun Wang, Ivan Skorokhodov, and Peter Wonka. Hf-neus: Improved surface reconstruction using high-frequency details. *Advances in Neural Information Processing Systems*, 35:1966–1978, 2022. 1
- [81] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [82] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. 2021. 3, 5
- [83] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):139–144, 1980. 1
- [84] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. *arXiv preprint arXiv:2312.02981*, 2023. 2
- [85] Tong Wu, Jiaqi Wang, Xingang Pan, Xudong Xu, Christian Theobalt, Ziwei Liu, and Dahua Lin. Voxurf: Voxel-based efficient and accurate neural surface reconstruction. In *ICLR*, 2023. 1
- [86] Haolin Xiong, Sairisheek Muttukuru, Rishi Upadhyay, Pradyumna Chari, and Achuta Kadambi. Sparsegs: Real-time 360 sparse view synthesis using gaussian splatting. *arXiv preprint arXiv:2312.00206*, 2023. 3
- [87] Chao Xu, Ang Li, Linghao Chen, Yulin Liu, Ruoxi Shi, Hao Su, and Minghua Liu. Sparp: Fast 3d object reconstruction and pose estimation from sparse views. In *European Conference on Computer Vision*, pages 143–163. Springer, 2024. 1, 6
- [88] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. 2
- [89] 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
- [90] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5438–5448, 2022. 2
- [91] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8254–8263, 2023. 2
- [92] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1790–1799, 2020. 1, 2, 5
- [93] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020. 2
- [94] 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. 1, 2
- [95] Mae Younes, Amine Ouasfi, and Adnane Boukhayma. Sparsecraft: Few-shot neural reconstruction through stereopsis guided geometric linearization. In *European Conference on Computer Vision*, pages 37–56. Springer, 2024. 2, 6, 7
- [96] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 2
- [97] Xianggang Yu, Mutian Xu, Yidan Zhang, Haolin Liu, Chongjie Ye, Yushuang Wu, Zizheng Yan, Chenming Zhu, Zhangyang Xiong, Tianyou Liang, et al. Mvimgnet: A large-scale dataset of multi-view images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9150–9161, 2023. 1, 2, 5, 7
- [98] 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. 1
- [99] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. 2
- [100] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian opacity fields: Efficient and compact surface reconstruction in unbounded scenes. *arXiv preprint arXiv:2404.10772*, 2024. 1, 2
- [101] 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. 2
- [102] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5
- [103] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847*, 2018. 3, 6
- [104] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. *arXiv preprint arXiv:2312.00451*, 2023. 1, 3, 6