


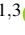




# 2D-SuGaR: Surface-Aware Gaussian Splatting for Geometrically Accurate Mesh Reconstruction

Prajwal Gupta C. R.<sup>†1</sup> , Divyam Sheth<sup>†1,2</sup> , Jinjoo Ha<sup>1</sup> , Mirela Ostrek<sup>1,3</sup>  and Justus Thies<sup>1,2,3</sup> 

<sup>1</sup>TU Darmstadt, <sup>2</sup>ELIZA, <sup>3</sup>Max Planck Institute for Intelligent Systems, <sup>†</sup>equal contribution



**Figure 1: Overview.** We propose 2D-SuGaR, a method to reconstruct 3D meshes from multi-view input images by leveraging 2D Gaussian Splatting. We initialize and regularize the 2D Gaussian primitives with pretrained normal and depth priors, and extract a mesh from this volumetric representation. We refine the mesh afterwards with additional re-rendering losses.

## Abstract

3D Gaussian Splatting enables the reconstruction of a volumetric scene representation from multi-view images that allows for real-time novel-view point synthesis, however, it struggles with recovering an accurate surface geometry. While 2D Gaussian Splatting (2DGS) addresses this through surface-aligned primitives, its performance depends critically on the initialization quality. Reliance on Structure-from-Motion (SfM) limits the initialization flexibility as well. In this work, we present two key contributions to enhance 2DGS and the extraction of a clean surface mesh. First, we incorporate monocular depth and normal priors for robust initialization, coupled with a clustering-based pruning strategy to eliminate degenerate Gaussians. Second, we introduce a joint mesh-Gaussian refinement similar to SuGaR, that relaxes the strict 2D constraint by transitioning to 3D primitives, providing stronger training signals. Evaluated on the DTU dataset, our method achieves state-of-the-art mesh reconstruction with a Chamfer Distance of 0.67, outperforming prior methods.

## CCS Concepts

• **Computing methodologies** → **Point-based models; Mesh models; Rasterization; Computer vision;**

## 1. Introduction

Reconstructing accurate 3D meshes from multi-view images remains a fundamental challenge in computer vision, with applications in AR/VR, robotics, and digital content creation. Recent advances based on Gaussian Splatting [KKLD23, HYC\*24] have enabled fast optimization and real-time photorealistic rendering. These methods represent a 3D scene in a volumetric way by a set of translucent Gaussians. The Gaussians are optimized using a differentiable renderer, where rendering proceeds by splatting and rasterizing the Gaussians in screen space. It is worth noting that this scheme is sensitive to the initialization of Gaussians. Extracting a mesh-based representation from optimized Gaussians for usage in classical computer graphics frameworks is typically performed as

a post-processing step, preventing errors introduced at this stage from being corrected. This limits the ability to capture accurate geometry and fine details.

Our work builds on top of 2D Gaussian-Splatting [HYC\*24] to reconstruct meshes from multi-view input images. To address the challenges mentioned above, we propose several contributions. First, we incorporate monocular depth and normal priors to enable *robust initialization and regularization during optimization*, which proves useful given the highly unconstrained nature of 3D reconstruction. We also introduce a *clustering-based pruning method to remove spurious Gaussians* that are not eliminated during adaptive density control. Finally, we propose a *joint mesh refinement stage* taking inspiration from SuGaR [GL24].

© 2026 The Author(s).

Proceedings published by Eurographics - The European Association for Computer Graphics. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.2312/egs.20261000

## 2. Method

We leverage 2D Gaussian primitives as an intermediate representation [HYC\*24] to recover a 3D mesh from multi-view input images, see Figure 1. To initialize the 2D Gaussians, we augment feature points from Structure-from-Motion (SfM) with estimates from monocular depth and normal predictions [HYZ\*24]. Besides initialization, the normal predictions from this general prior, are used in a new normal loss that encourages surface alignment of the 2D Gaussians. Once the 2D Gaussian-based representation is optimized, we extract a coarse mesh and further refine it using joint mesh-gaussian optimization. [GL24].

### 2.1. Depth- and Normal-Guided 2DGS

2DGS [HYC\*24] is dependent on the initialization through SfM. However, SfM fails in textureless and occluded regions, leaving gaps where reconstruction is most challenging. We address this by supplementing the SfM points with monocular depth and normal estimates from Metric3D [HYZ\*24]. To resolve the scale ambiguity of the depth, we compute per-image scale correction using SfM as reference. We then sample 1000 pixels per image over a 2D Gaussian distribution (50% in center) and backproject the depth to 3D space to initialize Gaussians, with orientations from the normal estimates and colors from the input image. Based on this initialization, we optimize the 2D Gaussian primitives using the loss:

$$\mathcal{L} = \mathcal{L}_p + \lambda_c \mathcal{L}_c + \lambda_n \mathcal{L}_n + \lambda_d \mathcal{L}_d. \quad (1)$$

Here,  $\mathcal{L}_p$  denotes the photometric loss combining an  $\ell_1$  term with the D-SSIM term from [KKLD23].  $\mathcal{L}_c$  is the normal consistency loss, and  $\mathcal{L}_d$  is the depth distortion loss (refer to Sec. 1 of the supplementary material). While the original depth-normal consistency loss  $\mathcal{L}_c$  ensures alignment between the rendered depth and rendered normals, we introduce a new term leveraging the monocular normal estimates  $\tilde{\mathbf{N}}$ :

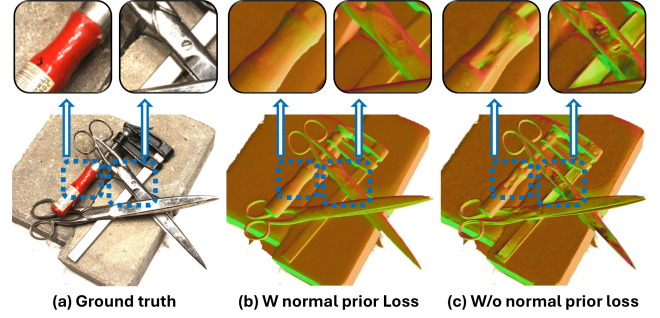
$$\mathcal{L}_n = 1 - \mathbf{N} \cdot \tilde{\mathbf{N}}. \quad (2)$$

The normal prior loss leads to normal maps that better reflect the underlying surface, as illustrated in Figure 2.

Note that errors in depth estimates may result in the initialization of spurious Gaussians. These are further amplified by the adaptive density control strategy, creating isolated Gaussian islands that produce undesirable reconstruction artifacts. To address this, we introduce a *clustering-based pruning strategy* using DBSCAN [SEKX98]. The neighborhood radius ( $\epsilon$ ) is dynamically determined using the  $k$ -distance heuristic. For each Gaussian  $p_i$ , we compute the distance  $d_i^{(k)}$  to its  $k$ -th nearest neighbor. We set  $\epsilon$  to the 90<sup>th</sup> percentile of  $d_i^{(k)}$ . We retain only the largest cluster, assuming that true geometry forms a single connected component.

### 2.2. Mesh Extraction and Refinement

Post optimization, we render depth maps from the training views. We fuse them into a Truncated Signed Distance Field (TSDF) and apply marching cubes to extract an initial mesh. We then instantiate new 3D Gaussians on the extracted mesh, following SuGaR [GL24]. We bind a single thin 3D Gaussian to each triangle of the mesh, sampled on the surface of the triangle. To ensure



**Figure 2: Effect of normal loss.** Our new normal loss reduces 2DGS reconstruction artifacts, especially in specular regions, and leads to a smoother and more faithful representation of the surface.

that the Gaussians remain bound to their respective triangles during optimization, we specify the Gaussian means in barycentric coordinates with respect to the mesh vertices. We refine it using:

$$\mathcal{L}_r = \mathcal{L}_p + \gamma \mathcal{L}_{Lap} + \delta \mathcal{L}_m, \quad (3)$$

where  $\mathcal{L}_{Lap}$  and  $\mathcal{L}_m$  are the standard mesh laplacian smoothing and mesh normal consistency losses respectively. Note that since the Gaussians in 2DGS have only two scale parameters and SuGaR is based on 3DGS, we introduce an additional dimension with a small initial value of 0.001. This modification aligns well with the final goal of joint refinement, where one dimension of each Gaussian is progressively reduced to produce a flat, surface-aligned shape.

### 2.3. Implementation Details


The proposed method builds on [HYC\*24] and [GL24], implemented in PyTorch with custom CUDA kernels for efficient execution. All our experiments are run on a single RTX 3090 GPU.

**Depth- and Normal-Guided 2DGS.** We adopt the training configuration used in 2DGS to ensure a fair and direct performance comparison. Each scene is optimized over 30K iterations, where we set  $\lambda_c = 0.05$ ,  $\lambda_n = 0.1$  and  $\lambda_d = 1000$ . During the depth-prior based initialization, we sample 1000 pixels per image in the scene. For each sampled pixel, we compute its corresponding 3D position using the monocular depth estimate and initialize a 2D Gaussian at that point. For scenes with very large number of images, the total number of samples is capped at 75000, with equal number of samples from each image.

$$N_{\text{samples}} = \min(75000, 1000 \times N_{\text{images}}) \quad (4)$$

When using the depth-prior based initialization, we apply our clustering-based pruning after 7K iterations using DBSCAN with  $\text{min\_samples} = k = 4$  and  $\epsilon = P_{90}(\{d_i^{(k)}\}_{i=1}^N)$ .

**Mesh Refinement.** We refine the extracted mesh for 8K iterations which takes approximately 8 minutes on a single RTX 3090. Even though we get good results in just 4K iterations, we got the best overall results for 8K iterations. 15K iterations did not improve the performance significantly for the time overhead it added. We set  $\gamma = 5.0$  and  $\delta = 0.1$  for all scenes as default values used in [GL24].



**Figure 3: Qualitative comparison of meshes on the DTU benchmark.** Our method produces more detailed and complete meshes as can be seen in the zoom-ins.

Method	24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean $\downarrow$	Time $\downarrow$
NeRF [MST*20]	1.90	1.60	1.85	0.58	2.28	1.27	1.47	1.67	2.05	1.07	0.88	2.53	1.06	1.15	0.96	1.49	>12h
VoISDF [YGKL21]	1.14	1.26	0.81	0.49	1.25	0.70	0.72	1.29	1.18	0.70	0.66	1.08	0.42	0.61	0.55	0.86	>12h
NeuS [WLL*21]	0.83	0.98	0.56	0.37	1.13	0.59	0.60	1.45	0.95	0.78	0.52	1.43	0.36	0.45	0.45	0.77	>12h
3DGS [KKLD23]	2.14	1.53	2.08	1.68	3.49	2.21	1.43	2.07	2.22	1.75	1.79	2.55	1.53	1.52	1.50	1.96	11.2m
SuGaR [GL24]	1.47	1.33	1.13	0.61	2.25	1.71	1.15	1.63	1.62	1.07	0.79	2.45	0.98	0.88	0.79	1.33	65m
2DGS [HYC*24]	0.48	0.91	0.39	0.39	1.01	0.83	0.81	1.36	1.27	0.76	0.70	1.40	0.40	0.76	0.52	0.80	10.9m
GOF [YSG24]	0.50	0.82	0.37	0.37	1.12	0.74	0.73	1.18	1.29	0.68	0.77	0.90	0.42	0.66	0.49	0.74	75m
RaDe-GS [ZFS*24]	0.46	0.73	0.33	0.38	0.79	0.75	0.76	1.19	1.22	0.62	0.70	0.78	0.68	0.47	0.68	0.68	11.5m
MILo [GGM*25]	0.43	0.74	0.34	0.37	0.80	0.74	0.70	1.21	1.22	0.66	0.62	0.80	0.37	0.76	0.48	0.68	35m
Ours w/o Refinement	0.50	0.67	0.33	0.44	1.24	0.70	0.63	1.24	1.13	0.63	0.49	1.29	0.43	0.56	0.44	0.72	10.5m
Ours	0.48	0.65	0.29	0.39	1.24	0.69	0.61	1.17	1.10	0.56	0.45	1.19	0.37	0.50	0.40	0.67	18.5m

**Table 1: Quantitative comparison of the mesh reconstruction quality of state-of-the-art methods on 15 scenes from the DTU Dataset.** We report the average Chamfer Distance between the ground truth mesh and the reconstructed mesh, as well as the reconstruction time.

### 3. Results

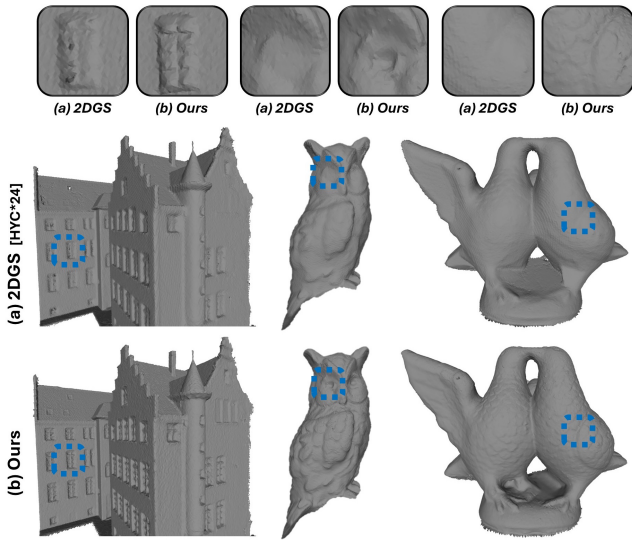
We evaluate on 15 DTU scenes [JDV\*14], using masks for bounded reconstructions following [HYC\*24], with 49 or 64 images per scene. In Table 1, we compare against several state-of-the-art approaches based on the Chamfer Distance (CD) and training time. Some of these works learn an implicit representation of the scene [MST\*20], [YGKL21], [WLL\*21], while others learn an explicit representation of the scene [KKLD23], [HYC\*24], [GL24]. Our method achieves the lowest Chamfer Distance among all compared approaches. However, the mesh refinement stage is time-consuming compared to other explicit representation methods. Even without the mesh refinement stage, our method remains more accurate and is also faster than other methods. The final scene representation also contains fewer Gaussians than [HYC\*24].

In Figure 3, we show a qualitative comparison to 2DGS and the concurrent work MILo [GGM\*25] which also aims at recovering meshes from multi-view images. As shown, our method recovers more detailed and smooth surface representations without noticeable boundary artifacts. Other methods oversmooth the surface and compensate by baking details into the texture. In contrast, our ap-

proach preserves geometric detail in the mesh itself (see zoom-ins in the top row). Furthermore, our meshes are more complete, with fewer holes and gaps in the reconstructed geometry (e.g. the scissors).

**Ablations – Depth- and Normal-Guided 2DGS.** In Table 2, we study the impact of each new component by sequentially adding them to 2DGS and evaluating the quality of mesh reconstruction using the DTU dataset. 2DGS + Normal Init. refers to the 2DGS pipeline where the orientation of the Gaussians is initialized using the normal priors, resulting in a modest performance improvement. Next, we incorporate the normal prior loss leading to a reasonable improvement in performance. The depth-prior based initialization only results in a modest improvement in performance on the DTU dataset. However, it can be highly effective in scenes where it is difficult to get a good initialization through SfM. Finally, we add the mesh refinement step which produces a significant performance improvement.

**Ablations – Mesh Refinement.** In Table 3, we compare two joint refinement strategies applied on top of vanilla 2DGS meshes to validate our relaxation approach. Relaxed 3D adds a small third



**Figure 4: Qualitative comparison to 2DGS.** Our meshes have fine geometric details such as windows, eyes, and feathers.

dimension ( $s_z = 10^{-3}$ ) to 2D surfels and refines using 3D Gaussian rasterization with losses from [KKLD23] and [GL24]. Strict 2D maintains rigid surfels using a surfel rasterizer. Without additional losses, results degrade as primitives drift along depth axis. With the regularization losses from [HYC\*24], it performs similar to the baseline but underperforms compared to Relaxed 3D. The results confirm our key insight that pure 2D primitives over-constrain mesh adaptation, while the small third dimension enables surface-normal adjustments without losing alignment.

Added Component	Accuracy ↓	Completion ↓	Average ↓
2DGS (Base)	0.701	0.818	0.759
+ Normal Init.	0.707	0.803	0.755
+ Normal Prior Loss	0.692	0.749	0.72
+ Depth Init.	0.689	0.741	0.715
+ Mesh Refinement	0.655	0.693	0.674

**Table 2: Quantitative ablation study evaluating the contribution of each component of our method for mesh extraction.**

Method	Mean CD ↓
2DGS (Baseline)	0.759
Relaxed 3D (Ours)	<b>0.713</b>
Strict 2D (Base)	0.787
Strict 2D (Regularized)	0.756

**Table 3: Mesh refinement with different primitives.**

#### 4. Conclusion

We have proposed a novel mesh reconstruction scheme based on 2D and 3D Gaussian Splatting. We improve the 2DGS [HYC\*24] reconstruction quality using off-the-shelf monocular depth and normal estimation models for initialization of the Gaussians as well as

regularization during optimization. Based on the optimized 2DGS representation, we extract a mesh and jointly refine it by constraining Gaussians to the mesh surfaces [GL24]. Through a series of experiments, we show the strengths and weaknesses of each introduced methodological aspect of our proposed method and are able to demonstrate superior quality in comparison to state-of-the-art approaches.

In the future, it will be interesting to investigate the effects of additional depth-based regularization. Improving the adaptive density control strategy is another exciting direction that may lead to better reconstructions.

**Acknowledgment** Divyam Sheth is supported by ELIZA (Konrad Zuse School of Excellence) via DAAD, funded by the Federal Ministry of Education and Research. The project is supported by the ERC Starting Grant 101162081 “LeMo” and the DFG Excellence Strategy—EXC-3057, Project No. 533677015.

#### References

- [GGM\*25] GUÉDON A., GOMEZ D., MARUANI N., GONG B., DRETTAKIS G., OVSIJANIKOV M.: Milo: Mesh-in-the-loop gaussian splatting for detailed and efficient surface reconstruction. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–15. 3
- [GL24] GUÉDON A., LEPETIT V.: 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* (2024), pp. 5354–5363. 1, 2, 3, 4
- [HYC\*24] HUANG B., YU Z., CHEN A., GEIGER A., GAO S.: 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers* (2024), pp. 1–11. 1, 2, 3, 4
- [HYZ\*24] HU M., YIN W., ZHANG C., CAI Z., LONG X., CHEN H., WANG K., YU G., SHEN C., SHEN S.: Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024). 2
- [JDV\*14] JENSEN R., DAHL A., VOGIATZIS G., TOLA E., AANÆS H.: Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2014), pp. 406–413. 3
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* 42, 4 (2023), 139–1. 1, 2, 3, 4
- [MST\*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV* (2020). 3
- [SEKX98] SANDER J., ESTER M., KRIEGEL H.-P., XU X.: Density-based clustering in spatial databases: The algorithm gdbscan and its applications. *Data mining and knowledge discovery* 2, 2 (1998), 169–194. 2
- [WLL\*21] WANG P., LIU L., LIU Y., THEOBALT C., KOMURA T., WANG W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv* (2021). [arXiv:2106.10689](https://arxiv.org/abs/2106.10689). 3
- [YGKL21] YARIV L., GU J., KASTEN Y., LIPMAN Y.: Volume rendering of neural implicit surfaces. *Advances in neural information processing systems* 34 (2021), 4805–4815. 3
- [YSG24] YU Z., SATTLER T., GEIGER A.: Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics (ToG)* 43, 6 (2024), 1–13. 3
- [ZFS\*24] ZHANG B., FANG C., SHRESTHA R., LIANG Y., LONG X., TAN P.: Rade-gs: Rasterizing depth in gaussian splatting. *arXiv* (2024). [arXiv:2406.01467](https://arxiv.org/abs/2406.01467). 3