GauSDF: Signed Distance Embedded Gaussian Surfels for 3D Reconstruction
Supplementary Material

Minsol Kim
Department of Artificial Intelligence
Sejong University, South Korea
rlaalsthf02@sju.ac.kr

Usman Ali
Departement of Computer Science and Engineering
Sejong University, South Korea
usman.ali@sejong.ac.kr

In this supplementary document, we provide additional analyses and experimental validations for our proposed method, **GauSDF**.

Specifically, we include: (1) A comparison of model memory usage. (2) Experiments analyzing the direct coupling between Gaussian Splatting (GS) and Signed Distance Function (SDF). (3) Additional qualitative results on the DF3D dataset [12]. (4) Extended experiments on the Mip-NeRF360 dataset [1]. (5) Further visualizations demonstrating mesh reconstruction quality. The code and data resources are available at https://github.com/Anonymous/GauSDF.

1. Memory Comparison

Relative Memory Utilization of Models. Figure 1 compares GPU memory usage and training time across different methods on the DTU scan 110 [4], trained on an NVIDIA RTX 4090 (24 GB). GSDF [8] exhibits the highest memory consumption, peaking at nearly 24 GB with a prolonged training duration due to its volumetric rendering pipeline. In contrast, our **GauSDF** achieves significantly more efficient memory utilization, requiring only about 6 GB of peak GPU memory—roughly four times lower than GSDF—while also training approximately twice as fast. Although slightly less memory-efficient than lightweight approaches such as 2DGS [3], GS-Pull [10], and GOF [9], GauSDF strikes an effective balance between efficiency and reconstruction quality by eliminating the need for volumetric ray marching in the SDF branch.

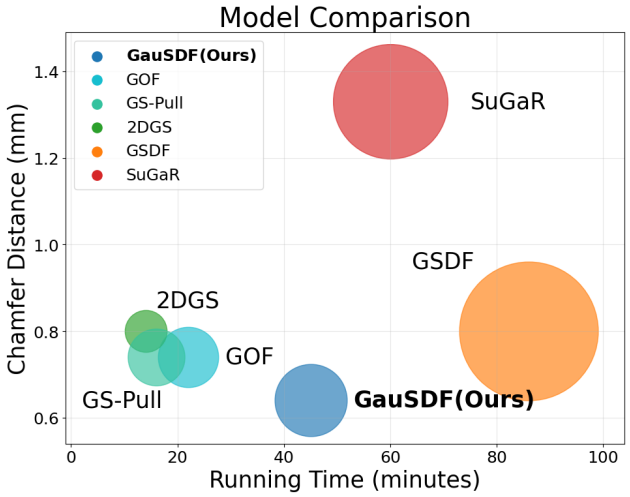

Figure 1. Bubble chart comparing different methods, including 2DGS [3], SuGaR [2], GS-Pull [10], GSDF [8], GOF [9], and GauSDF (Ours). The bubble size represents GPU memory usage, while the axes indicate training time and surface reconstruction accuracy (Chamfer Distance, mm). GauSDF achieves the highest reconstruction quality while maintaining strong computational efficiency.

2. Direct Coupling between GS and SDF

Degree of Coupling between GS and SDF. Through experimentation, we observed that directly coupling the per-surfel local SDF embedding with the global SDF network leads to unstable optimization. This relationship can be formulated as:

$$L_{\text{direct}} = \|\boldsymbol{\xi}_i - f(\boldsymbol{\mu}_i)\|_1, \tag{1}$$

where $\boldsymbol{\xi}_i \in \mathbb{R}$ represents the per-surfel local SDF embedding, $\boldsymbol{\mu}_i = (x, y, z) \in \mathbb{R}^3$ denotes the center of each surfel, and f is the global SDF network.

This direct supervision enforces agreement between local and global SDF predictions regardless of the true underlying geometry. As a result, the optimization may converge to trivial or shortcut solutions, driving the SDF field away from a valid zero-level set and leading to degraded surface reconstruction quality, as illustrated in Figure 2.

Direct supervision fails for two key reasons, as detailed below:

(1) Range incompatibility. The valid range of the SDF embedding is governed by the pruning strategy in Adaptive Density Control (ADC) of 3DGS [5]. With an opacity prun-

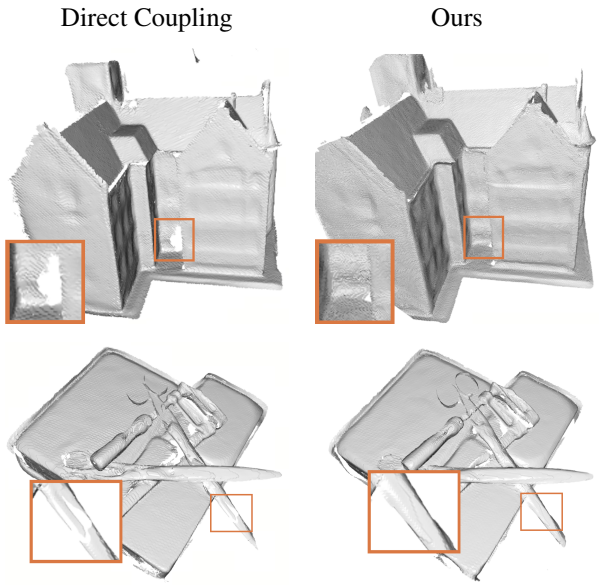

Figure 2. Drawbacks of direct coupling between GS and the global SDF. Directly enforcing consistency between local and global SDF representations leads to unstable optimization and prevents the formation of a geometrically coherent surface.

ing threshold of $\epsilon_\alpha = 0.005$, the local SDF embedding $\boldsymbol{\xi}_i$ is confined to the interval $\left(\frac{-6.7}{\beta}, \frac{6.7}{\beta}\right)$, derived by solving $\epsilon_\alpha = \frac{ke^{-\beta\xi_i}}{(1+e^{-\beta\xi_i})^2}$ for $k = 4$. Directly enforcing equality between the SDF network output and Gaussian-embedded SDF values causes numerical and geometric instabilities: the global SDF field becomes distorted, or the local embeddings $\boldsymbol{\xi}_i$ drift outside their feasible range, leading to Gaussian disappearance and surface collapse.

(2) Redundant supervision. Indirect coupling losses, such as L_{pull} and L_{surfel}, already provide sufficient guidance for coherent surface formation. As Gaussians are iteratively attracted toward the SDF's zero-level set, their centers $\boldsymbol{\mu}_i$ approach valid surface positions where $f(\boldsymbol{\mu}_i) \approx 0$. Simultaneously, local SDF embeddings $\boldsymbol{\xi}_i$ are driven toward zero to maximize opacity via photometric loss. The scaling factor β implicitly regulates the relationship between global and local SDF magnitudes. Thus, alignment $f(\boldsymbol{\mu}_i) \approx \boldsymbol{\xi}_i \approx 0$ naturally emerges during training, eliminating the need for an explicit consistency constraint.

3. DF3D Dataset

Mesh Reconstruction. Figure 3 shows qualitative comparisons on open-surface garment reconstructions. We evaluate against 2DGS [3], GOF [9], and several UDF-based methods. 2DGS exhibits incomplete object coverage, while GOF produces distorted mesh structures. NeuDF [7] yields coarse surfaces with inconsistent normals, and VRPrior [11] introduces high-frequency artifacts. GaussianUDF [6] produces sharper details but suffers from flipped normals and local discontinuities. In contrast, GauSDF reconstructs clean, stable meshes with well-preserved fine structures and minimal artifacts. Overall, our method delivers reconstruction quality competitive with UDF-based methods on the DF3D [12] dataset while maintaining robust geometric consistency.

4. Mip-NeRF360 Dataset

We further evaluate GauSDF on the Mip-NeRF360 dataset [1]. The main comparison is conducted against 2DGS [3] and GS-Pull [10], as GSDF [8] encounters out-of-memory issues on our GPU setup.

Scene-wise Quantitative Results. Table 1 reports detailed scene-wise comparisons on Mip-NeRF360. Across all scenes, GauSDF consistently outperforms prior methods in PSNR, SSIM, and LPIPS. Specifically, GauSDF surpasses 2DGS by an average of 3.6 dB in PSNR on outdoor scenes and achieves substantial gains on indoor ones. SSIM improves significantly—rising from 0.705 to 0.890 for outdoor scenes (a 26.3% relative increase)—while LPIPS values are reduced by more than half, reflecting a notable enhancement in perceptual quality. Although GS-Pull [10] struggles with degraded rendering fidelity, GauSDF attains higher visual realism and geometric coherence than the baseline [3].

Qualitative Results. Figure 4 shows qualitative comparisons on *bicycle* (outdoor) and *room* (indoor) scenes. Both 2DGS and GS-Pull exhibit incomplete geometry and blurred structures, whereas GauSDF reconstructs fine details and smooth, artifact-free surfaces. These results confirm that embedding SDF parameters within Gaussian surfels enables robust alignment with true object geometry, yielding high-quality reconstructions across diverse environments.

5. Additional Mesh Visualizations

We provide further qualitative examples of reconstructed meshes from the DTU [4] and DF3D [12] datasets in Figure 5. These results highlight GauSDF's ability to produce smooth, geometrically consistent, and artifact-free surfaces across diverse object types and scene complexities.

References

[1] 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. 1, 2, 3, 4

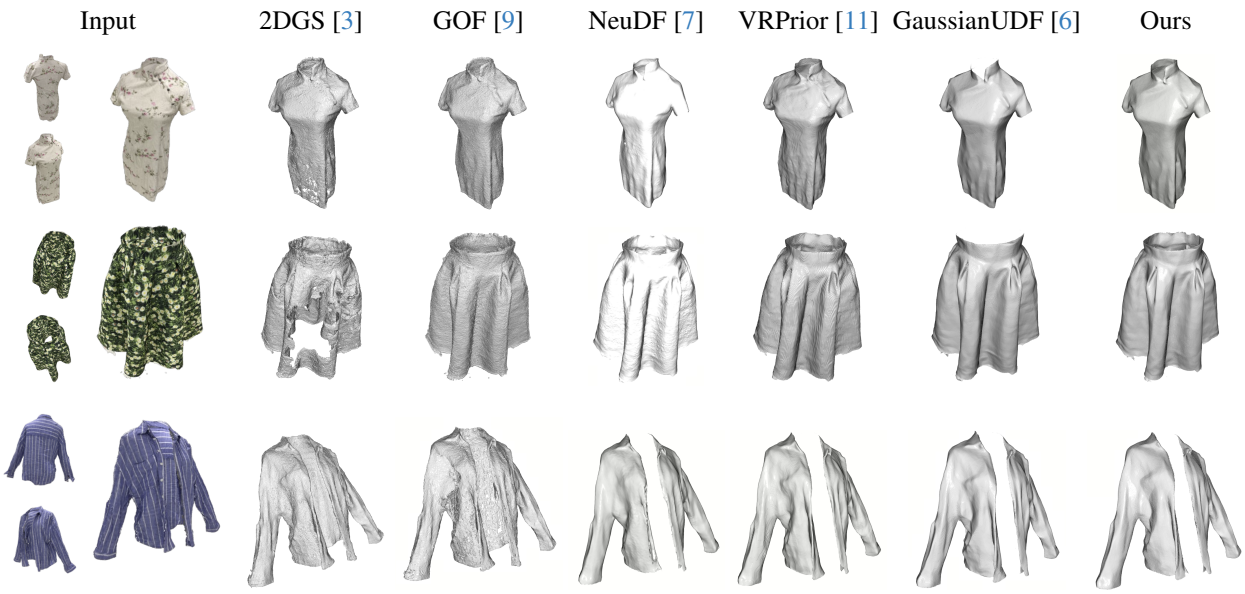

Figure 3. Qualitative comparisons on the DF3D dataset [12]. GauSDF reconstructs more accurate, artifact-free surfaces and better preserves fine structural details such as clothing folds and contours.

Table 1. Scene-wise rendering comparison between 2DGS [3], GS-Pull [10], and GauSDF (Ours) on the Mip-NeRF360 dataset [1]. GauSDF achieves better metrics across both indoor and outdoor scenes.

		Indoor scenes					Outdoor scenes					
Metric	Method	room	counter	kitchen	bonsai	$\text{Avg}_{\text{indoor}}$	bicycle	flowers	garden	stump	treehill	$\text{Avg}_{\text{outdoor}}$
PSNR ↑	2DGS	30.62	28.14	30.22	31.23	30.06	24.74	21.13	26.72	26.19	22.32	24.22
	GS-Pull	30.04	26.40	28.07	28.87	28.84	23.83	20.12	26.17	24.88	22.04	23.81
	Ours	32.21	29.45	32.24	30.34	31.06	26.29	24.96	31.23	30.27	26.33	27.81
SSIM ↑	2DGS	0.907	0.893	0.916	0.931	0.912	0.733	0.575	0.843	0.758	0.618	0.705
	GS-Pull	0.899	0.858	0.888	0.906	0.888	0.698	0.508	0.822	0.675	0.647	0.670
	Ours	0.950	0.929	0.957	0.949	0.946	0.889	0.820	0.943	0.921	0.874	0.890
LPIPS ↓	2DGS	0.243	0.229	0.147	0.227	0.212	0.268	0.374	0.146	0.258	0.375	0.284
	GS-Pull	0.251	0.269	0.179	0.252	0.238	0.293	0.430	0.177	0.355	0.378	0.327
	Ours	0.103	0.107	0.061	0.095	0.092	0.122	0.186	0.052	0.100	0.154	0.123

[2] 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

[3] 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

[4] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 406–413, 2014. 1, 2, 4

[5] 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

[6] Shujuan Li, Yu-Shen Liu, and Zhizhong Han. Gaussianudf: Inferring unsigned distance functions through 3d gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27113–27123, 2025. 2, 3

[7] Yu-Tao Liu, Li Wang, Jie Yang, Weikai Chen, Xiaoxu Meng, Bo Yang, and Lin Gao. NeUDF: Leaning neural unsigned distance fields with volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 237–247, 2023. 2, 3

[8] Mulin Yu, Tao Lu, Linning Xu, Lihan Jiang, Yuanbo Xiangli, and Bo Dai. GSDF: 3DGS meets sdf for improved rendering and reconstruction. In *Advances in Neural Information Processing Systems*, 2024. 1, 2

[9] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian

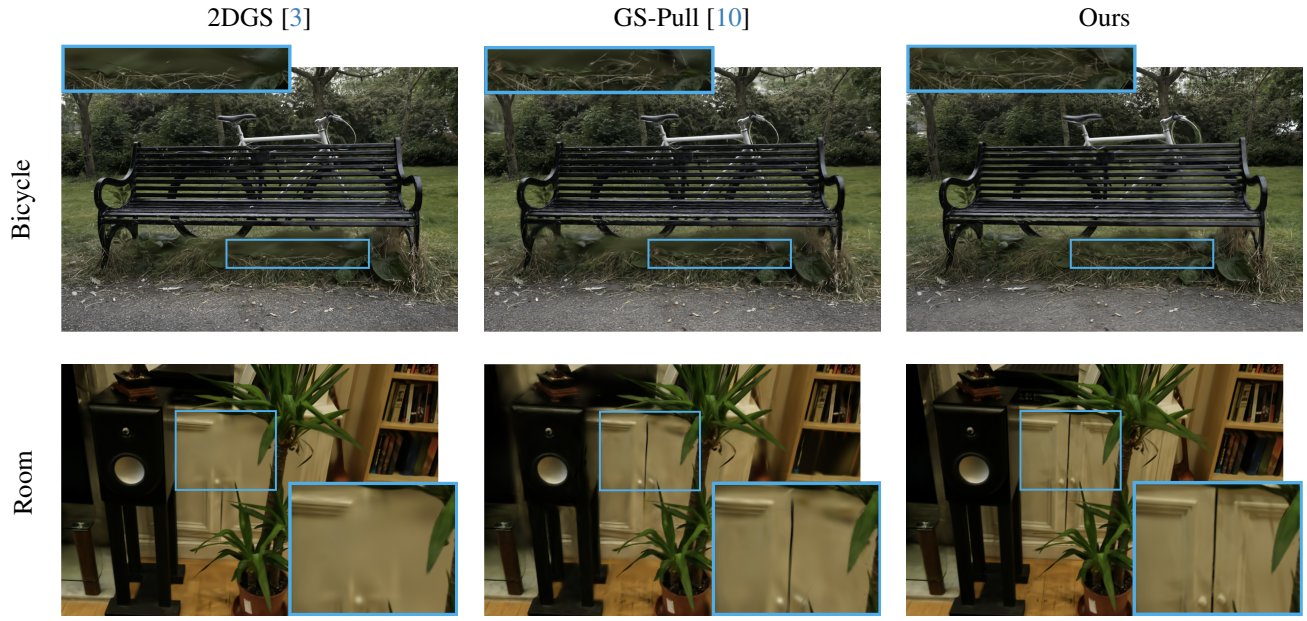

Figure 4. Qualitative comparisons on the Mip-NeRF360 dataset [1]. Ours (GauSDF) reconstructs sharper, geometry-consistent surfaces with fewer artifacts compared to existing methods.

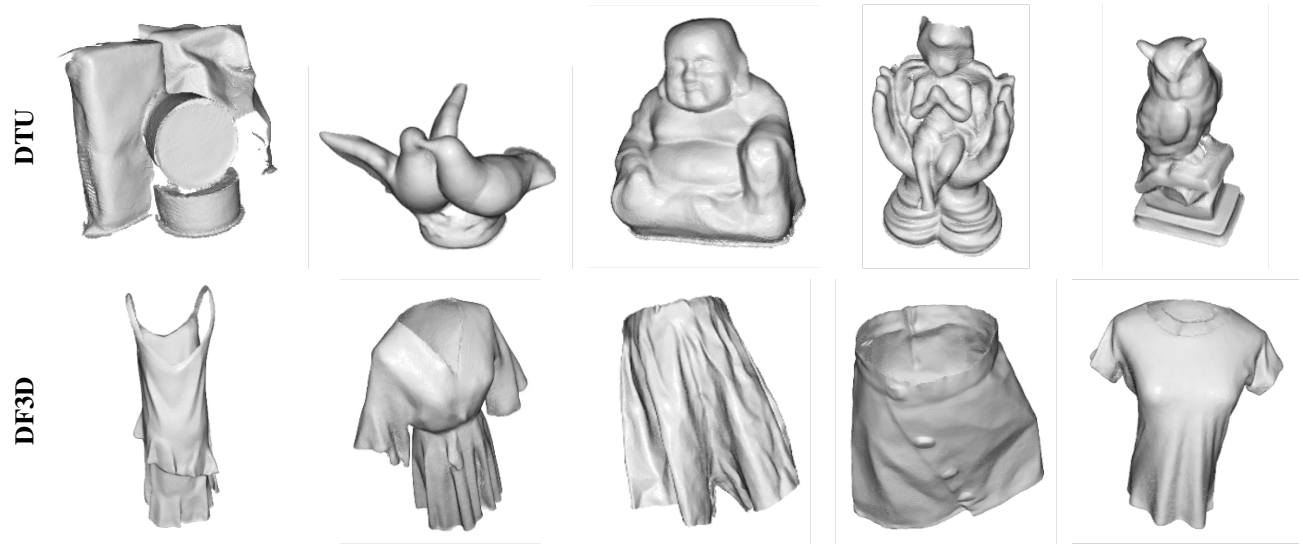

Figure 5. Additional qualitative results on the DTU [4] and DF3D [12] datasets. GauSDF produces clean, high-fidelity surfaces with well-preserved geometric structures across both watertight and open-surface scenes.

opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Trans. Graph.*, 2024. 1, 2, 3

[10] Wenyuan Zhang, Yu-Shen Liu, and Zhizhong Han. Neural signed distance function inference through splatting 3D Gaussians pulled on zero-level set. In *Advances in Neural Information Processing Systems*, 2024. 1, 2, 3, 4

[11] Wenyuan Zhang, Kanle Shi, Yu-Shen Liu, and Zhizhong Han. Learning unsigned distance functions from multi-view images with volume rendering priors. In *Proceedings of the European Conference on Computer Vision*, 2024. 2, 3

[12] Heming Zhu, Yu Cao, Hang Jin, Weikai Chen, Dong Du, Zhangye Wang, Shuguang Cui, and Xiaoguang Han. Deep Fashion3D: A dataset and benchmark for 3D garment reconstruction from single images. In *Proceedings of the European Conference on Computer Vision*, pages 512–530. Springer, 2020. 1, 2, 3, 4