

The Rendering–Surface Gap in Gaussian Splatting: Gauge Ambiguity, Extraction Conditioning, and a Controlled DTU Study

Fariborz Taherkhani

Abstract

Gaussian Splatting can synthesize convincing novel views while yielding an inaccurate or extraction-dependent surface. This paper studies that gap as an inverse-problem ambiguity rather than proposing another unverified geometry loss. We define an ε -render gauge—the set of Gaussian parameters producing indistinguishable images—and its geometric diameter. For alpha compositing, we prove an exact optical-mass subdivision invariance; locally, we show that image-preserving geometric progress is controlled by the projection of the geometry gradient onto the rendering Jacobian’s nullspace. We also derive the condition number of level-set extraction, which explains why a fixed radiance representation can yield materially different meshes.

We test these predictions in a controlled DTU Scan 24 study with the same 42/7 train/test split. Among reproduced published systems, TSGS gives the best held-out appearance (23.480 dB PSNR, 0.9149 SSIM), whereas MLo gives the best standard-extraction geometry (2.431 mm Chamfer). Three project interventions fail their preregistered continuation gates: RayOT improves PGSR Chamfer by 5.36% but loses 0.186 dB PSNR and 0.0155 SSIM; GaugeSplat and TraceSplat preserve appearance but improve Chamfer by only 0.166% and 0.190%. A TSGS first-surface extraction diagnostic changes no rendered image yet improves Chamfer by 3.86%, trading 11.39% worse accuracy for 22.13% better completeness. These are single-scene diagnostic results, not state-of-the-art claims. Together they provide measured evidence that rendering quality, latent Gaussian geometry, and extracted-surface quality are distinct quantities and should be evaluated as such.

1 Introduction

Three-dimensional Gaussian Splatting (3DGS) [1] represents a scene by anisotropic primitives whose projected contributions are depth sorted and alpha composited. Its explicit representation enables fast optimization and photorealistic rendering, but a surface is not an explicit output of the image-formation equation. Geometry-aware variants therefore replace volumetric Gaussians by disks [2], align them to a surface [3], render plane-intersection depth [4], extract an opacity level set [5], or put a differentiable mesh in the training loop [10].

The usual experimental question is whether a new regularizer improves both novel-view metrics and a mesh metric. A prior

question is more fundamental: *how much surface information is identified by the renderer at all?* The rendered images depend on cumulative optical thickness, projected footprints, colors, and visibility. A mesh additionally depends on a chosen map from those quantities to a surface. Consequently, two parameter sets may be nearly render-equivalent but geometrically different, while two extraction rules may produce different meshes from exactly the same trained model. This is the Gaussian analogue of shape–radiance ambiguity in neural radiance fields [13].

Our project encountered this distinction empirically. Geometry-only motion with a strict image trust region barely changed external geometry; relaxing the image constraint improved Chamfer but damaged rendering; and changing only the depth-to-mesh extraction improved Chamfer without touching the rendered RGB. Rather than hide these negative results, we organize them into a mathematical and experimental diagnosis of the *rendering–surface gap*.

The paper makes four contributions:

- We define the ε -render gauge and its geometric diameter, separating representation ambiguity from extraction ambiguity.
- We prove exact optical-mass subdivision invariance and derive a first-order compatibility coefficient that quantifies how much geometric descent is available without changing rendered images.
- We derive the normal displacement induced by a perturbation of an extracted level set, exposing an extraction condition number that is independent of RGB quality.
- We report a controlled, auditable DTU Scan 24 study over reproduced PGSR, MLo, and TSGS baselines and four project diagnostics. Every failed gate remains in the results, and method ownership is explicit.

Our conclusion is deliberately limited: the study supports the existence and practical importance of the rendering–surface gap on one controlled scene. It does not establish benchmark-wide superiority for any project intervention.

2 Related work

Surface-aware Gaussian Splatting. 2DGS [2] uses perspective-correct planar disks together with depth-distortion and normal consistency. SuGaR [3] encourages Gaussians to align to a surface before mesh extraction. PGSR [4]

derives unbiased plane-intersection depth and adds single- and multi-view constraints. Gaussian Opacity Fields [5] extract an adaptive opacity level set, while 3DGSR [6] and GSDF [7] couple splats to an implicit field. Quadratic Gaussian Splatting [8] increases primitive order, and MATCha Gaussians [9] organize local charts. These methods improve geometry by changing the representation, supervision, or extraction operator; our goal is to analyze why the three axes need not agree.

Mesh-in-the-loop and transparent surfaces. MILo [10] differentially constructs mesh vertices and connectivity from Gaussian pivots and enforces bidirectional Gaussian–mesh consistency. TSGS [11] targets the transparency–depth dilemma through a geometry stage with de-lighted inputs and normal priors, a frozen-opacity appearance stage, and first-surface depth extraction. The latter is especially relevant here: it makes the extraction operator an explicit design variable rather than treating expected alpha depth as a uniquely correct surface.

Appearance–geometry decoupling. View-dependent color can compensate for geometric error. Ref-NeRF [14] structures view dependence, and GStex [12] attaches high-frequency texture to individual 2D Gaussians. Such decoupling can improve rendering without proportionally improving geometry. Our gauge formulation states this observation as a property of the forward map and its Jacobian.

3 The Rendering–Surface Gap

3.1 Forward rendering and optical coordinates

Let $\theta \in \Theta$ contain Gaussian positions, covariances, opacities, and appearance. For a pixel u , sort the overlapping primitives by depth and write

$$\begin{aligned}\widehat{C}_\theta(u) &= \sum_{i=1}^n T_i(u) \alpha_i(u) c_i(u) + T_{n+1}(u) c_{\text{bg}}, \\ T_i(u) &= \prod_{j < i} (1 - \alpha_j(u)).\end{aligned}\tag{1}$$

Introduce optical thickness $\ell_i(u) = -\log(1 - \alpha_i(u)) \geq 0$. Then $T_i = \exp(-\sum_{j < i} \ell_j)$, so compositing depends on cumulative thickness, not on a unique decomposition into primitives.

Proposition 1 (Exact coincident subdivision invariance). *At a pixel, replace a primitive with color c and optical thickness ℓ by m consecutive coincident primitives with the same color and thicknesses $\ell_k \geq 0$ satisfying $\sum_{k=1}^m \ell_k = \ell$. The composited color contribution and outgoing transmittance are unchanged.*

Proof. Ignoring the common incoming transmittance, the sub-

divided contribution is

$$S = \sum_{k=1}^m e^{-\sum_{r < k} \ell_r} (1 - e^{-\ell_k}) c \tag{2}$$

$$= \sum_{k=1}^m \left(e^{-\sum_{r < k} \ell_r} - e^{-\sum_{r \leq k} \ell_r} \right) c \tag{3}$$

$$= (1 - e^{-\sum_k \ell_k}) c = (1 - e^{-\ell}) c. \tag{4}$$

The outgoing transmittance is $e^{-\sum_k \ell_k} = e^{-\ell}$. $\square$

This exact invariance is already enough to break primitive-level identifiability: the number of centers, per-primitive scale statistics, and any mesh rule that depends on primitive identity may change while the pixel does not. Approximate invariances arise when primitives move within a projected footprint, colors compensate for motion, or visibility changes only outside the observed cameras.

3.2 The ε -render gauge

Let $\mathcal{R} : \Theta \rightarrow \mathcal{I}$ concatenate the rendered images at a fixed camera set, and let d_I be an image discrepancy. Let $\mathcal{E} : \Theta \rightarrow \mathcal{S}$ be a specified surface extraction rule and d_S a surface distance.

Definition 1 (ε -render gauge). *The render gauge around θ is*

$$\mathcal{G}_\varepsilon(\theta) = \{\theta' \in \Theta : d_I(\mathcal{R}(\theta'), \mathcal{R}(\theta)) \leq \varepsilon\}. \tag{5}$$

Its extraction-dependent geometric diameter is

$$D_S(\varepsilon; \theta, \mathcal{E}) = \sup_{\theta', \theta'' \in \mathcal{G}_\varepsilon(\theta)} d_S(\mathcal{E}(\theta'), \mathcal{E}(\theta'')). \tag{6}$$

Image-space convergence implies surface convergence only if $D_S(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$. Proposition 1 shows that this need not hold for extraction rules that are not invariant to optical-mass subdivision.

3.3 Local compatibility of image and geometry descent

Assume differentiability at θ and let J_R be the Jacobian of the render vector. Let $L_G(\theta)$ be a geometric loss with gradient g_G . The first-order image-preserving directions form $\mathcal{N}_R = \text{Null}(J_R)$. Denote the orthogonal projector onto this nullspace by P_N .

Proposition 2 (Best image-preserving geometric descent). *Among unit directions satisfying $J_R \delta = 0$, the largest first-order decrease in L_G is $\|P_N g_G\|_2$. If $P_N g_G \neq 0$, it is attained by $\delta^* = -P_N g_G / \|P_N g_G\|_2$.*

Proof. For every feasible unit vector, $-g_G^\top \delta = -(P_N g_G)^\top \delta \leq \|P_N g_G\|_2$ by Cauchy–Schwarz. Equality is attained by the stated direction. $\square$

This motivates the local compatibility coefficient

$$\rho_G(\theta) = \frac{\|P_N g_G\|_2}{\|g_G\|_2} \in [0, 1]. \tag{7}$$

When ρ_G is small, substantial geometric progress necessarily leaves the exact first-order render gauge. A tolerance-aware sensitivity is

$$\kappa_{G \leftarrow I}^{(\eta)}(\theta) = \sup_{\|\delta\|_2=1} \frac{|g_G^\top \delta|}{\|J_R \delta\|_2 + \eta}, \quad \eta > 0. \quad (8)$$

The denominator prevents unbounded ratios along an exact nullspace. Large κ means a small render change can correspond to a large geometric change; small ρ_G means the desired geometry gradient conflicts with image preservation. These are different failure modes.

3.4 Extraction is a second inverse problem

Many pipelines first construct a scalar field $F_\theta(x)$ (opacity, occupancy, signed distance, or fused TSDF) and then extract $S_\tau = \{x : F_\theta(x) = \tau\}$. The conditioning of this second map is independent of the RGB renderer.

Proposition 3 (Level-set displacement). *Let $x \in S_\tau$ be regular, $\nabla F(x) \neq 0$, and perturb the field by δF . The first-order signed displacement of the level set along the unit normal $n = \nabla F / \|\nabla F\|$ is*

$$\delta s(x) = -\frac{\delta F(x)}{\|\nabla F(x)\|_2} + O(\|\delta F\|^2). \quad (9)$$

Proof. Set $x' = x + \delta s, n$ and expand $F(x') + \delta F(x') = \tau + \text{delta } \nabla F(x)^\top n + \text{delta } F(x) + O(\delta s^2 + \|\delta F\| |\delta s|)$. Solving the first-order term for δs gives Equation (9). $\square$

Thus $1/\|\nabla F\|$ is a pointwise extraction condition number. Flat opacity or TSDF bands permit large surface motion from a small field or threshold change. Expected depth, median termination depth, first-surface depth, and TSDF fusion instantiate different $\mathcal{E}$ and can therefore move Chamfer distance while $\mathcal{R}(\theta)$ is fixed.

4 Controlled Diagnostic Interventions

We used five interventions as probes of the preceding theory. They are not presented as successful new methods.

VP-0: measure-proximal motion. A trained 2DGS checkpoint was interpreted as a discrete surface measure. A geometry-only step optimized local moment, quadratic-chart, tangent-plane, normal, and distortion energies while freezing opacity, scale, and spherical harmonic appearance. A render sentinel accepted or rolled back each block. This probe tests whether a hand-designed geometry gradient has a useful component inside a strict render gauge, as in Equation (7).

RayOT and GaugeSplat. RayOT transported a PGSR point distribution toward 92,940 train-only multi-view evidence points using sparse unbalanced optimal transport, then projected positions within a trust radius. GaugeSplat added a short appearance-compensation phase intended to return the update

to the render gauge. RayOT tests a looser image constraint; GaugeSplat tests whether compensation can keep the geometric gain.

TraceSplat. MILO was augmented with a robust trace loss from the same train-only evidence. The loss acted only on MILO’s occupancy-SDF variables; the RGB renderer did not consume those variables. This is a direct test of whether an apparently decoupled latent geometry channel improves the final extracted surface.

First-surface extraction. The trained TSGS representation was held fixed. We changed only the depth used by TSDF fusion from the standard path to the repository’s first-surface sliding-window path. This isolates $\mathcal{E}$ from $\mathcal{R}$ and is a diagnostic, not a new algorithm.

5 Experimental Protocol

Scene and split. All reported measurements use DTU Scan 24 [17]: 49 calibrated images, 42 training views, and held-out views $\{0, 8, 16, 24, 32, 40, 48\}$. Image metrics are PSNR, SSIM, LPIPS, and foreground PSNR on these seven views. Geometry is evaluated with the DTU observation mask and plane using

$$d_{d \rightarrow s} = \frac{1}{|P_d|} \sum_{p \in P_d} \min_{q \in P_s} \|p - q\|_2, \quad (10)$$

$$d_{s \rightarrow d} = \frac{1}{|P_s|} \sum_{q \in P_s} \min_{p \in P_d} \|q - p\|_2, \quad (11)$$

$$\text{CD} = \frac{1}{2}(d_{d \rightarrow s} + d_{s \rightarrow d}). \quad (12)$$

Distances are in millimeters. We refer to $d_{d \rightarrow s}$ as accuracy and $d_{s \rightarrow d}$ as completeness, matching the evaluator’s directional outputs.

Baselines and ownership. PGSR, MILO, and TSGS are published systems and are not our contributions. The project contributions in this study are the formal analysis, the controlled probes, the paired gates, and the audit. The 2DGS VP-0 probe used backend commit 335ad612 and rasterizer commit e0ed0207. The TSGS reproduction used commit 73d3f984; Lotus-2 normal inference used b05bf90c.

Preregistered continuation gates. For PGSR and MILO interventions, continuation required at least 2% relative Chamfer improvement with PSNR drop ≤ 0.1 dB, foreground-PSNR drop ≤ 0.1 dB, SSIM drop ≤ 0.005 , and LPIPS increase ≤ 0.01 . VP-0 required at least 1% median moment-defect reduction and at most 60% rollback, in addition to image preservation. We report the failed checks rather than tuning them after seeing the held-out results.

Scope boundary. This is a controlled single-scene pilot. We ran neither a multi-scene benchmark nor enough random seeds




Figure 1: **Appearance and geometry do not induce one ranking.** Every point is measured on the same Scan 24 split. The plot is diagnostic; differences between implementations prevent a causal cross-method claim.

for inferential statistics. Furthermore, our staged DTU data did not reproduce the public TSGS Scan 24 Chamfer number; the measured result is therefore treated as a local reproduction under the common staged protocol, not as a contradiction of the authors’ benchmark.

6 Results

6.1 Published baselines disagree on the winner

Table 1 contains no common winner. Relative to MILO, TSGS gains 2.915 dB PSNR and 0.1766 SSIM but has 13.23% higher Chamfer. Conversely, MILO’s stronger mesh does not produce the strongest held-out image metrics. The result is descriptive rather than causal because the systems differ in training and extraction, but it directly rules out using PSNR as a surrogate for mesh accuracy in this pilot.

6.2 Image preservation alone does not yield useful geometry

The paired results in Table 2 separate two regimes. RayOT moves far enough to produce a visible 5.36% Chamfer gain, but exits the allowed render gauge. GaugeSplat and TraceSplat remain well inside the image budget, yet their Chamfer gains are only 0.166% and 0.190%, an order of magnitude below the 2% continuation threshold. In the language of Proposition 2, either the geometry gradient had little usable projection onto the render nullspace, or the implemented update failed to exploit it; the present experiment cannot distinguish those explanations.

VP-0 tells the same story internally. It froze opacity, scale, and appearance and accepted 100 of 275 attempted geometry steps. The median moment defect fell by only 0.0777%, while 63.6% of attempted blocks rolled back. The strict sentinel preserved held-out rendering, but the targeted surface statistic did not move enough to justify scaling the experiment.

6.3 Extraction changes the accuracy-completeness balance

The first-surface extractor improves aggregate Chamfer entirely by improving completeness: $d_{s \rightarrow d}$ falls by 22.13%, while accuracy $d_{d \rightarrow s}$ worsens by 11.39% (Table 3). A single Chamfer number hides this directional tradeoff. More importantly, every RGB metric is unchanged because the model is unchanged. This is direct experimental evidence for treating $\mathcal{E}$ as a separately conditioned inverse problem.

6.4 What the negative results establish

The failed probes do not prove that the proposed losses are universally ineffective. They establish narrower, reproducible facts on this pilot:

1. A strict image trust region can preserve rendering while yielding negligible external geometry change (GaugeSplat, TraceSplat) or negligible internal target change (VP-0).
2. A larger geometric correction can improve Chamfer while violating image-preservation gates (RayOT).
3. A fixed representation can move Chamfer materially through the extraction operator alone (TSGS first surface).
4. The baseline with best rendering is not the baseline with best standard-extraction geometry (TSGS versus MILO).

Together these observations match the two-factor decomposition

$$\theta \xrightarrow{\mathcal{R}} \text{images}, \quad \theta \xrightarrow{\mathcal{E}} \text{surface}, \quad (13)$$

rather than a model in which image quality is a proxy for surface quality.

7 Implications for Method Design and Evaluation

Optimize a quotient, not primitive identities. Proposition 1 implies that regularizers sensitive to Gaussian count or identity can penalize render-equivalent parameterizations differently. Optical mass, transport measures, or other split/merge invariants are preferable when the intended object is a surface rather than a particular discretization.

Measure gradient compatibility before a long run. The coefficient ρ_G in Equation (7) can be estimated with Jacobian-vector products and a nullspace projection or constrained least squares. If it is near zero, increasing a geometry loss is expected to damage rendering; if it is large but a method still makes no progress, the optimizer or parameterization is the likely bottleneck. This diagnostic can replace expensive blind sweeps.

Report an extraction family. A geometry paper should name $\mathcal{E}$, publish its threshold and fusion parameters, and report at

Table 1: **Measured published baselines on the common Scan 24 pilot.** Higher is better for PSNR/SSIM/foreground PSNR; lower is better for LPIPS and geometry. TS GS is the appearance winner, while MLo is the standard-mesh geometry winner. These methods belong to their respective authors.

Published method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FG-PSNR $\uparrow$	$d_{d \rightarrow s} \downarrow$	$d_{s \rightarrow d} \downarrow$	CD $\downarrow$
PGSR [4]	20.258	0.7210	0.2372	21.620	3.871	2.703	3.287
MLo [10]	20.565	0.7383	0.2158	20.787	2.724	2.137	2.431
TS GS [11]	23.480	0.9149	0.0987	27.574	3.001	2.504	2.752

Table 2: **Paired project interventions.** Image deltas are candidate minus parent, so positive PSNR/SSIM is favorable. Geometry gain is the relative Chamfer reduction. “Reject” means a preregistered continuation gate failed. VP-0 had no valid DTU mesh measurement and is reported using its internal target only.

Probe	Parent	Δ PSNR	Δ SSIM	Δ LPIPS	Δ FG-PSNR	Geometry gain	Outcome
RayOT	PGSR	−0.1863	−0.01552	+0.01542	−0.3261	+5.358% CD	Reject: image
GaugeSplat	PGSR	−0.0006	−0.00003	−0.00004	−0.0010	+0.166% CD	Reject: geometry
TraceSplat	MLo	−0.0188	−0.00015	+0.00001	−0.0101	+0.190% CD	Reject: geometry
First-surface	TS GS	0	0	0	0	+3.860% CD	Extraction diagnostic
VP-0	2DGS	−0.0062	+0.00054	−	+0.0111	+0.0777% moment	Reject: target/rollback

Table 3: **TS GS extraction diagnostic.** The trained representation and RGB renders are identical; only depth-to-mesh extraction changes.

Extraction	$d_{d \rightarrow s}$	$d_{s \rightarrow d}$	CD
Standard	3.0007	2.5038	2.7523
First surface	3.3424	1.9497	2.6461
Relative change	+11.39%	−22.13%	−3.86%

least two plausible extraction rules when the forward representation is not itself a mesh. Directional distances should accompany Chamfer. An extraction sweep is not post-processing trivia: by Proposition 3, it probes a separately conditioned map.

Use paired Pareto gates. Geometry improvement should be reported relative to the exact parent checkpoint, together with image deltas. A method should not claim a joint improvement when geometry is obtained by spending unreported rendering quality, or when images are preserved but geometry moves below practical resolution. The gates in this study are one concrete, preregistered example; their numerical thresholds are not universal.

8 Limitations

The main limitation is scale: one DTU scene and one seed cannot establish a population effect or a state-of-the-art result. Cross-system comparisons share cameras and metrics but not identical training code, initialization, or extraction, so they are descriptive. The TS GS staged reproduction did not match its public Scan 24 number, indicating an unresolved data or protocol mismatch; we therefore use only the locally measured value and do not dispute the published benchmark. The first-surface result was a post-hoc diagnostic and must be preregistered

before any multi-scene claim.

The theory is local where it uses Jacobians and fixed visibility. Occlusion swaps, sorting discontinuities, densification, and pruning can create non-smooth changes not captured by J_R . The level-set analysis assumes a regular surface and does not cover topology changes at critical points. Finally, we did not explicitly estimate $D_S(\varepsilon)$ or ρ_G ; the experiments expose behavior consistent with these quantities, while direct measurement is future work.

9 Conclusion

Gaussian Splatting has at least three distinct objects: rendered images, latent primitive geometry, and an extracted surface. Exact compositing invariances, render-Jacobian nullspaces, and level-set conditioning prevent these objects from being identified in general. Our controlled Scan 24 study exhibits all three consequences: the appearance and geometry winners are different; image-preserving geometry updates produce only marginal Chamfer gains; and extraction alone changes Chamfer by 3.86% while leaving every RGB image fixed. The practical lesson is not that high-quality rendering is unimportant, but that it is insufficient evidence for a high-quality surface. Future Gaussian-surface papers should measure the render gauge, the extraction operator, and the appearance–geometry Pareto front explicitly.

A Audit Details

The VP-0 run used 528,731 surfels, a frozen audit subset of 65,536 nodes, and 100 accepted steps out of 275 attempted steps. Immutable hashes verified that opacity, scale, and spherical-harmonic tensors were unchanged. RayOT and GaugeSplat used 92,940 train-only evidence tracks comprising 368,834 observations, with mean track length 3.969. TraceS-



Figure 2: **Image-space near-equivalence of PGSR and GaugeSplat.** From left to right: held-out ground truth, PGSR, GaugeSplat, and the GaugeSplat error amplified by $4\times$. The candidate changes the image metrics by less than 0.001 dB PSNR and 0.00004 SSIM, yet the external Chamfer gain is also only 0.166%. This is a measured example of staying inside a tight render gauge without finding substantial geometric descent.



Figure 3: **Held-out view 000000.** MILO and TraceSplat are visually almost indistinguishable, consistent with their 0.0188 dB PSNR difference. TSGS renders this view more cleanly, but its standard extracted mesh has higher Chamfer than MILO on the same pilot. Images alone do not reveal that ordering.

plat used the same evidence without held-out views; 92,934 evidence points were supported at its final map, and its trace loss acted only on occupancy-SDF variables. All paired tables were generated from the archived JSON gate files. Qualitative figures show archived held-out renders, not re-rendered or selected training views.

B Why a Better Chamfer Can Be Worse Locally

Let $a = d_{d \rightarrow s}$ and $c = d_{s \rightarrow d}$. Since $CD = (a + c)/2$, an extraction change with $\Delta a > 0$ still improves Chamfer when $-\Delta c > \Delta a$. For the TSGS diagnostic, $\Delta a = +0.3417$ mm and $\Delta c = -0.5542$ mm, hence $\Delta CD = -0.1062$ mm. The scalar score improves even though one directed notion of fidelity worsens. This algebra is simple, but it is operationally important: tuning a threshold against Chamfer alone can move a surface away from observations while filling missing regions aggressively.

References

- [1] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3D Gaussian Splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.
- [2] B. Huang, Z. Yu, A. Chen, A. Geiger, and S. Gao. 2D Gaussian Splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH Conference Papers*, 2024.
- [3] A. Guédon and V. Lepetit. SuGaR: Surface-aligned Gaussian Splatting for efficient 3D mesh reconstruction and high-quality mesh rendering. In *CVPR*, 2024.
- [4] D. Chen, H. Li, W. Ye, Y. Wang, W. Xie, S. Zhai, N. Wang, H. Liu, H. Bao, and G. Zhang. PGSR: Planar-based Gaussian Splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [5] Z. Yu, T. Sattler, and A. Geiger. Gaussian Opacity Fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics*, 43(6), 2024.
- [6] X. Lyu, Y.-T. Sun, Y.-H. Huang, X. Wu, Z. Yang, Y. Chen, J. Pang, and X. Qi. 3DGS: Implicit surface reconstruction with 3D Gaussian Splatting. *arXiv:2404.00409*, 2024.
- [7] W. Zhang, Y.-S. Liu, and Z. Han. Neural signed distance function inference through splatting 3D Gaussians pulled on the zero-level set. In *NeurIPS*, 2024.
- [8] Z. Zhang, B. Huang, H. Jiang, L. Zhou, X. Xiang, and S. Shen. Quadratic Gaussian Splatting: High quality surface reconstruction with second-order geometric primitives. In *ICCV*, 2025.
- [9] A. Guédon, T. Ichikawa, K. Yamashita, and K. Nishino. MATCha Gaussians: Atlas of charts for high-quality geometry and photo-realism from sparse views. In *CVPR*, 2025.

- [10] A. Guédon, D. Gomez, N. Maruani, B. Gong, G. Drettakis, and M. Ovsjanikov. MILO: Mesh-in-the-loop Gaussian Splatting for detailed and efficient surface reconstruction. *ACM Transactions on Graphics*, 44(6), 2025.
- [11] M. Li, P. Pang, H. Fan, H. Huang, and Y. Yang. TSGS: Improving Gaussian Splatting for transparent surface reconstruction via normal and de-lighting priors. In *ACM Multimedia*, 2025.
- [12] V. Rong, J. Chen, S. Bahmani, K. Kutulakos, and D. Lindell. GStex: Per-primitive texturing of 2D Gaussian Splatting for decoupled appearance and geometry modeling. In *WACV*, 2025.
- [13] K. Zhang, G. Riegler, N. Snavely, and V. Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020.
- [14] D. Verbin, P. Hedman, B. Mildenhall, T. Zickler, J. Barron, and P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. In *CVPR*, 2022.
- [15] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, 2021.
- [16] Z. Li, T. Müller, A. Evans, R. Taylor, M. Unberath, M.-Y. Liu, and C.-H. Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, 2023.
- [17] R. Jensen, A. Dahl, G. Vogiatzis, E. Tola, and H. Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014.