From Blobs to Spokes: High-Fidelity Surface Reconstruction via Oriented Gaussians

Diego Gomez[1⋆], Antoine Guédon[1⋆], Nissim Maruani[2], Bingchen Gong[1], and Maks Ovsjanikov[1]

[1] LIX, École Polytechnique, France
[2] Inria, Côte d'Azur, France

Fig. 1: High-fidelity surface reconstruction from 3D Gaussian Splatting. Our method, Gaussian Wrapping, reconstructs watertight and textured surface meshes of full 3D scenes given multiview RGB images, by interpreting 3D Gaussians as stochastic oriented surface elements. This figure illustrates the resulting surface meshes, with and without the RGB texture. Our meshes represent the full scene—including background geometry and extremely thin structures such as bicycle spokes where existing methods fail—with a *significantly* more compact representation than concurrent works [9, 56].

Abstract. 3D Gaussian Splatting (3DGS) has revolutionized fast novel view synthesis, yet its opacity-based formulation makes *surface extraction* fundamentally difficult. Unlike implicit methods built on Signed Distance Fields or occupancy, 3DGS lacks a global geometric field, forcing existing approaches to resort to heuristics such as TSDF fusion of blended depth maps. Inspired by the *Objects as Volumes* framework [38], we derive a principled occupancy field for Gaussian Splatting and show how it can be used to extract highly accurate watertight meshes of complex scenes. Our key contribution is to introduce a learnable oriented normal at each Gaussian element and to define an adapted attenuation formulation, which leads to closed-form expressions for both the normal and occupancy fields at arbitrary locations in space. We further introduce a novel consistency loss and a dedicated densification strategy to enforce Gaussians to *wrap* the entire surface by closing geometric holes,

⋆ Both authors contributed equally to the paper.

ensuring a complete shell of oriented primitives. We modify the differentiable rasterizer to output depth as an isosurface of our continuous model, and introduce **Primal Adaptive Meshing** for Region-of-Interest meshing at arbitrary resolution. We additionally expose fundamental biases in standard surface evaluation protocols and propose two more rigorous alternatives. Overall, our method **Gaussian Wrapping** sets a new state-of-the-art on DTU and Tanks and Temples, producing complete, watertight meshes at a fraction of the size of concurrent work—recovering thin structures such as the notoriously elusive bicycle spokes. Our project page is available here.

Keywords: 3D Gaussian Splatting · Surface Reconstruction · Meshes

1 Introduction

Reconstructing high-quality 3D surfaces from a set of 2D images is a key problem in computer vision. While recent advances in neural rendering, such as NeRF [37] and 3DGS [24], have shown impressive results on the novel view synthesis problem, extracting accurate and explicit geometry continues to be challenging. To bridge efficient volumetric rendering and surface reconstruction requires addressing their inherent differences: neural rendering is particularly suited to soft, semi-transparent representations, whereas surface reconstruction demands hard, well-defined boundaries to extract watertight geometry.

Current approaches to this inverse problem generally fall into two categories. Implicit representations, such as Signed Distance Functions (SDFs) and occupancy fields, excel at representing continuous, topologically consistent surfaces. However, they are often computationally expensive to train and tend to produce over-smoothed results, as their global solvers struggle to capture high-frequency details [28, 49]. On the other hand, explicit particle-based methods (i.e, 3DGS) exhibit real-time rendering and high-fidelity visual quality by representing scenes as discrete primitives. While they can capture fine details, there is no straightforward way to recover an ordered structure, i.e, a mesh, from such representations. Recent hybrid approaches attempt to close the gap by regularizing 3DGS [17, 18, 41, 43, 54, 56], but often resort to ad-hoc density thresholds or compromise the rendering speed and quality that makes 3DGS attractive.

A core limitation of extracting surfaces from 3DGS is the interpretation of the primitive itself. Standard approaches treat Gaussians as symmetric "blobs" of mass or density. This assumption conflicts with the nature of an orientable surface, which is an asymmetric boundary separating empty and occupied space. By modeling surface points with symmetric primitives, existing methods bias reconstruction, and make surface extraction significantly more challenging.

In this work, we address this limitation by introducing **Gaussian Wrapping**, a novel framework that reinterprets the role of the Gaussian primitive inspired by Objects as Volumes [38]. Rather than treating Gaussians as symmetric clouds, we view them as *oriented probabilistic surface elements* that represent the underlying geometry only within an oriented half-space. In this view, the

(a) Mesh w/wo texture **(b)** Close-up view of Mesh w/wo texture

Fig. 2: Examples of textured meshes reconstructed with our method. While being watertight and lighter than meshes reconstructed by state-of-the-art methods, our approach is able to reconstruct extremely fine details such as a few blades of grass—as can be seen in the close-up view of the stump scene [5].

Gaussian models the density decay on the outward-facing side of the boundary, while the inward-facing side is considered as fully occupied. This formulation allows us to derive a robust "wrapping" strategy: we identify geometric gaps where our oriented surface assumption is violated and densify the representation to obtain a watertight shell of oriented primitives. Our framework allows the derivation of closed-form geometric quantities—namely occupancy, its complement (*vacancy*), and the normal field—enabling reconstruction of thin structures without artifacts or missing geometry. Links to extracted meshes, code and more are available here. Our key contributions are as follows.

- We introduce **Oriented-Gaussians** and their associated training strategy in the multi-view setting.
- We derive a theoretical connection between 3DGS and implicit surface reconstruction by formulating Gaussians as oriented surface elements, inspired by *Objects as Volumes* [38]. Importantly, this leads to closed-form expressions for both normal and occupancy fields at arbitrary locations without any additional learnable parameters.
- We propose **Primal Adaptive Meshing**, a mesh extraction procedure that leverages the derived Gaussian fields to produce high-quality, watertight meshes at controllable resolution, enabling recovery of extremely thin structures such as bicycle spokes (Fig. 1).

2 Related work

Novel View Synthesis. NeRF [37] models scenes as continuous radiance fields whose image formation relies on ray-marching through an attenuation field (also called *density*) parameterized by an MLP. Subsequent work improves anti-aliasing [4,5] and training efficiency via hash encodings [39] or explicit volumetric structures [8,14]. 3D Gaussian Splatting [24] redefined this landscape by replacing ray-marching with rasterization of learned Gaussian primitives, achieving real-time, high-quality rendering. Further extensions address aliasing [53], scalability [25, 34], and rendering fidelity. Despite impressive visual quality, none of these methods target explicit surface mesh extraction.

Surface Reconstruction from Images. Implicit approaches coupling volume rendering with signed distance functions (NeuS [49], VolSDF [50], Neuralangelo [28]) convert a learnable SDF into a view-dependent attenuation field and yield smooth, topologically consistent surfaces but require very long optimization times. Mesh Baking [42, 51] approaches employ a two-stage pipeline that first trains a volumetric representation, then extracts a mesh to bake rendering onto it; this line of work however focuses on rendering and does not measure the geometric loyalty of the produced meshes. Discrete mesh methods [40, 44] produce strong results on isolated objects but do not scale to entire scenes.

In this regard, the efficiency of Gaussian Splatting motivates a growing family of surface-extraction approaches. Several jointly optimize Gaussians with an auxiliary neural implicit field [10, 27, 30, 52, 58], but the dual optimization limits scalability and the implicit branch remains bounded by MLP capacity. Others regularize Gaussians directly: 2DGS [20] replaces 3D primitives with lower-dimensional ones, RaDe-GS [55], PGSR [9], Quadratic Gaussian Splatting [59], and VCR-GauS [11] explore alternative rendering strategies or multi-view consistency constraints, while SuGaR [18] and GOF [54] propose specialized extraction procedures. A shared limitation is the reliance on heuristic depth definitions—typically TSDF fusion of blended depth maps—without a principled link between the Gaussian representation and a well-defined geometric field.

Objects as Volumes [38] establishes a closed-form relationship between vacancy and view-dependent attenuation under reciprocal exponential transport, grounding the theory behind Neural SDF methods [49, 50]. Since the 3DGS image formation model is not attenuation-based, this result does not apply directly—though Celarek *et al.* [7] show 3DGS approximates volumetric rendering with view-dependent attenuation. Concurrent work GGGS [56] builds upon [38] and exploits a continuous transmittance derived from Gaussians for improved depth estimation. However, their formulation suffers from surface erosion and fails to capture fine geometric details. Our method addresses this limitation by introducing oriented Gaussians that yield a reciprocal attenuation, from which closed-form vacancy and normal fields naturally follow—enabling principled, high-fidelity mesh extraction and limiting erosion.

Predicting Normals from Radiance Fields. Estimating normals from radiance fields is established practice in inverse rendering, where accurate normal maps help disentangle material properties [15, 23, 29, 33, 48]. In contrast, we predict normals not for appearance decomposition but to construct an oriented surface—directly enabling robust surface reconstruction.

Voronoi & Delaunay-based Methods. Voronoi diagrams and their dual Delaunay triangulations were first utilized as backbones for surface reconstruction with provable guarantees [1–3, 12], and later integrated into learning-based pipelines [6, 35, 36, 45–47, 57]. Within neural rendering, GOF [54] and MILo [17] build Delaunay triangulations over Gaussian-derived vertices, while other methods [16,19,32] integrate them into ray-tracing pipelines for view synthesis rather than reconstruction. In contrast, our Primal Adaptive Meshing incorporates ideas from modern reconstruction pipelines [57] to generate Delaunay-based watertight meshes at any desired resolution.

3 Background and Motivation

We focus on *fully opaque objects*, where the scene is binary: each point is either inside or outside the object. Our goal is to derive a *principled occupancy field* from 3D Gaussian Splatting, *i.e.* a view-independent scalar field encoding scene geometry, from which high-quality surface meshes can be extracted (Fig. 2).

Objects as Volumes. Reconstructing well-defined surfaces requires a properly defined geometric field such as an occupancy function $\mathcal{O} : \mathbb{R}^3 \to [0,1]$, yet volumetric rendering—the backbone of both NeRF [37] and 3DGS [24]—operates on continuous, semi-transparent volumes rather than on sharp boundaries. *Objects as Volumes* (OaV) [38] bridges this gap by interpreting the scene as a stochastic surface: the occupied region is modeled as a random set and the occupancy $\mathcal{O}$ as the probability of a point to be occupied. Standard volumetric rendering quantities then emerge as expectations over this random geometry.

Specifically, under the assumption of exponential light transport, this framework grounds the theoretical justification of Neural SDFs [49,50], and links the *vacancy* $v(\mathbf{x}) = 1 - \mathcal{O}(\mathbf{x})$ (the probability that $\mathbf{x} \in \mathbb{R}^3$ is unoccupied) to the attenuation (or density) $\sigma : \mathbb{R}^3 \times \mathcal{S}^2 \to \mathbb{R}_+$ of standard volumetric ray-marching:

$$\forall (\mathbf{x}, \mathbf{w}) \in \mathbb{R}^3 \times \mathcal{S}^2, \quad \sigma(\mathbf{x}, \mathbf{w}) = |\mathbf{w} \cdot \nabla \log v(\mathbf{x})| \,, \tag{1}$$

where attenuation is assumed to be *reciprocal*, *i.e.* $\sigma(\mathbf{x}, \mathbf{w}) = \sigma(\mathbf{x}, -\mathbf{w})$. This equation directly connects a geometric field, the vacancy, to the attenuation driving volumetric rendering. It is the cornerstone of our approach: if Gaussian Splatting can be cast as an attenuation-based model satisfying reciprocity, then Eq. 1 yields an explicit occupancy field and, in turn, a principled surface.

Gaussian Splatting. 3D Gaussian Splatting [24] represents a scene as N Gaussian primitives, each defined by $G_i(\mathbf{x}) = \alpha_i \exp(-\frac{1}{2}(\mathbf{x} - \mu_i)^T \Sigma_i^{-1}(\mathbf{x} - \mu_i))$, with opacity $\alpha_i \in [0,1]$, mean $\mu_i \in \mathbb{R}^3$, and precision matrix $\Sigma_i^{-1} \in \mathbb{S}^3_{++}$. Unlike NeRF [37], which ray-marches through a continuous attenuation field σ, 3DGS relies on *opacity*: it renders the pixel p by projecting each primitive onto the image plane and alpha-compositing contributions in depth order:

$$C(p) = \sum_i c_i \, \alpha_i G_i^*(p) \prod_{j<i} \bigl(1 - \alpha_j G_j^*(p)\bigr) \,, \tag{2}$$

where G_i^* is either a projected 2D Gaussian as in 3DGS [24], or the evaluation of the 3D Gaussian at the point of maximum contribution along the ray [54–56].

A key consequence is that each Gaussian's contribution depends only on its projected value at the pixel—not on the ray path length through the primitive. Gaussian opacities are therefore bounded in $[0, 1]$ and cannot be identified with the potentially unbounded attenuation coefficient σ. As a result, the OaV framework cannot be applied to 3DGS directly.

Recent work [7] shows that the 3DGS image formation model can nonetheless be reinterpreted as *"a volumetric rendering model with view-dependent extinction"*, assuming primitives do not overlap in 3D. Building on this approximation, we derive an explicit attenuation-based formulation of 3DGS satisfying OaV, yielding both more accurate depth rendering and an explicit geometric field for surface extraction. Geometry field splatting [22] uses OaV to define a geometric field out of particles. However they apply their analysis to gaussian surfels, and use TSDF as their extraction method, limiting their capability to scale to full scenes. Concurrent work, Geometry-Grounded Gaussian Splatting [56], works in a similar setting as us, applying Equation 1 to derive a continuous transmittance and improve depth estimation, but does not enforce multi-view consistent depthmaps for a single Gaussian (see the supplementary material), which is essential for a well-defined occupancy field as illustrated in Section 5. We restore reciprocity and multi-view consistency, by *orienting* each Gaussian with a normal vector, and derive the following:

Definition 1 (Oriented Gaussian). *We assign an oriented normal vector parameter $n_i \in \mathcal{S}^2$ to each Gaussian G_i, and define its oriented attenuation coefficient as*

$$\bar{\sigma}_i(\boldsymbol{x}, \boldsymbol{w}) = \mathbb{1}_{n_i^T(\boldsymbol{x}-\mu_i)\geq 0} \left| \boldsymbol{w} \cdot \nabla \log(1 - G_i(\boldsymbol{x})) \right| . \tag{3}$$

It can be shown—both theoretically and qualitatively—that this formulation not only satisfies the geometric properties required for applying OaV, but is also a valid approximation of the Gaussian Splatting image formation model under specific assumptions. This attenuation-based model constitutes the basis of our approach, and allows us to derive the method below. We provide an illustration of how a simpler non-oriented attenuation [56] fails to produce accurate results in the supplementary.

In the following section, we expose our theoretical results and describe how our formulation can be used for improving surface extraction from Gaussian Splatting representations. **Please refer to the supplementary materials for extensive proofs and derivations motivating our results.**

4 Method

Method Overview: Intuitively our method operates as follows: first, as mentioned above, we endow each Gaussian with a learnable oriented normal parameter (note that these are *the only* additional learnable parameters of our framework), and introduce the corresponding attenuation coefficient, as in Definition 1. Based

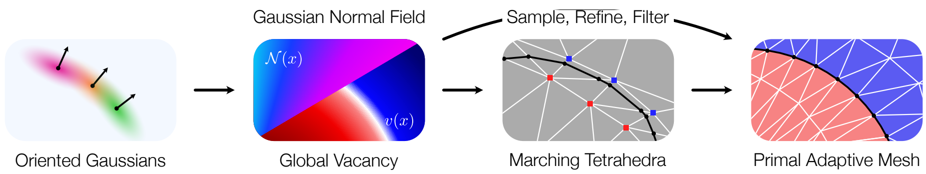

Fig. 3: Wrapping the surface with Oriented Gaussians. Our Gaussian Wrapping setting allows to evaluate the normal and vacancy fields of the scene. We first leverage the vacancy on our pivot-based marching tetrahedra meshing approach. We then leverage both the vacancy and normal fields in our Primal Adaptive Meshing.

on this construction and thanks to the Objects as Volumes formalism [38], we derive closed-form expressions for both the normal and occupancy fields for arbitrary locations in space, without any additional learnable parameters. Second, we propose a training strategy, which learns the oriented normals through an image-based consistency loss, using an adapted rasterizer and coupled with a new normal-aware densification strategy. Lastly, we exploit our normal field and occupancy estimates in a novel dedicated mesh-extraction approach, which both respects the underlying geometry, while being adaptive to details. For a pipeline figure presenting our method, please see Fig. 3.

4.1 Gaussian Wrapping

Our core intuition is simple: for accurate surface reconstruction, Gaussians must form a sealed boundary around the object. We achieve this by endowing each Gaussian with a learnable oriented normal n_i and enforcing a *wrapping* behavior: visible Gaussians near the surface should point outward, away from the occupied region and toward the camera. This wrapping creates a continuous shell of oriented primitives even around extremely thin structures—such as bicycle spokes—for which, standard methods only recover the *appearance* with a few unoriented primitives, but cannot recover the geometry. More fundamentally, we show in the supplementary material that when Gaussians properly wrap the scene, the oriented attenuation of Definition 1 becomes a valid approximation of the Gaussian Splatting image formation model, and the *Objects as Volumes* framework [38] directly yields an explicit closed-form vector field V (and, in turn, a normal field). We promote wrapping in the scene through two complementary mechanisms.

Normal alignment loss. Rendering depth and normals n_i with the Splatting rasterizer [54, 55] yields, respectively, an estimate of the surface depth and the expected orientation of Gaussians near the surface [38]. As a result, we enforce the alignment between the normals n_i and the surface during training via:

$$\mathcal{L}_{\mathrm{N}} = \sum_p 1 - N(p) \cdot \nabla D(p) \, , \tag{4}$$

where $N(p)$ is the rendered oriented normal at pixel p and $\nabla D(p)$ the image-space gradient of the rendered depth. Specifically, we modify the differentiable

CUDA rasterizer introduced in [56] to match our framework, and render depth as the exact 0.5-isosurface of our geometric field.

Densification. Gaps appear in the wrapping shell where Gaussians fail to cover all sides of a surface, causing $\mathcal{L}_{\mathrm{N}}(p)$ to rise locally. This provides a natural densification criterion: every K iterations, we compute errors $\mathcal{L}_{\mathrm{N}}(p)$ across all training views and propagate them to individual Gaussians via their blending weights. High-error Gaussians are cloned with flipped normals, closing holes and reinforcing the shell. We illustrate this process in the supplementary material.

Loss Details. During optimization, we combine our novel normal alignment objective $\mathcal{L}_{\mathrm{N}}$ in conjunction with the following set of losses: The standard photometric [24] loss $\mathcal{L}_{\mathrm{RGB}}$, depth-normal consistency $\mathcal{L}_{\mathrm{DN}}$ [54,55] and the multi-view loss $\mathcal{L}_{\mathrm{MV}} = \lambda_{\mathrm{pc}}\mathcal{L}_{\mathrm{pc}} + \lambda_{\mathrm{gc}}\mathcal{L}_{\mathrm{gc}}$, where the terms control the strength of the photometric and geometric consistency as defined in [9,56]. These have the following weights $\lambda_{\mathrm{DN}} = 0.05, \lambda_{\mathrm{N}} = 0.05, \lambda_{\mathrm{pc}} = 0.6, \lambda_{\mathrm{gc}} = 0.02$. This results in the total loss

$$\mathcal{L} = \mathcal{L}_{\mathrm{RGB}} + \lambda_{\mathrm{DN}}\mathcal{L}_{\mathrm{DN}} + \lambda_{\mathrm{N}}\mathcal{L}_{\mathrm{N}} + \mathcal{L}_{\mathrm{MV}}$$

4.2 Gaussian Vector and Normal Fields

Assuming Gaussians properly wrap the scene, applying Eq. 1 from *Objects as Volumes* [38] to our reciprocal attenuation (Eq. 3) yields our main result, expressed through the *Gaussian vector field* V.

Proposition 1 (Gaussian Vector and Normal Fields). *For N oriented Gaussians, we define the Gaussian vector field $V : \mathbb{R}^3 \to \mathbb{R}^3$ and its normalization $\mathcal{N} : \mathbb{R}^3 \to \mathcal{S}^2$ as:*

$$V(x) := \sum_{i=1}^{N} \mathbb{1}_{n_i^T(x-\mu_i)\geq 0} \nabla \log(1 - G_i(x)) \tag{5}$$

$$\mathcal{N}(x) := \frac{V(x)}{\|V(x)\|} \tag{6}$$

Under the OaV assumptions, $V = \nabla \log v$, making V a closed-form estimate of $\nabla \log v$. Consequently, $\mathcal{N}(x)$ is well-defined in a neighborhood of the surface, and coincides with the true normal field of the expected stochastic surface of M.

These closed-form expressions link Gaussian parameters directly to surface geometry and form the basis of our mesh extraction pipeline (Sec. 4.3). Note that in practice, we approximate these quantities locally by querying only the K-nearest Gaussians around each point. Proofs are deferred to the supplementary material.

4.3 Mesh Extraction

We now leverage the vacancy as an implicit function for mesh extraction. While direct integration of V (Eq. 5) along any camera ray is possible, Gaussians that do not contribute to rendering can remain hidden inside the geometry after optimization, violating the wrapping assumption and producing artifacts if intersected along a ray. In practice, we never have access to the true vacancy v and instead compute an estimate $\tilde{v}_{\mathrm{T}}$ of it by iterating over the set of all training camera rays $\mathcal{T}_c$ with

$$\tilde{v}_{\mathrm{T}}(\mathbf{x}) = \max_{(\mathbf{o},\mathbf{w})\in\mathcal{T}_c} \left\{ \prod_{i=1}^{N} \left(1 - \bar{G}^{(i)}_{\mathbf{o},\mathbf{w}}(t)\right) : \mathbf{x} = \mathbf{o} + t\mathbf{w}, t > 0 \right\}, \tag{7}$$

where $\bar{G}^{(i)}_{\mathbf{o},\mathbf{w}}(t) = G_i(\mathbf{o} + \min(t, t^{G_i}_{\mathbf{o},\mathbf{w}})\mathbf{w})$; and $t^{G_i}_{\mathbf{o},\mathbf{w}} := \arg\max_{t\geq 0} G_i(\mathbf{o} + t\mathbf{w})$. Intuitively, this computation corresponds to finding the "most" unobstructed ray reaching point x, following our definition of the continuous transmittance detailed in the supplementary material. This result *formalizes the opacity-field meshing heuristics* empirically adopted in recent works [41,54,56]. Strictly speaking, $\tilde{v}_{\mathrm{T}}$ is a lower bound on the true vacancy v (i.e. $\tilde{v}_{\mathrm{T}}(\mathbf{x}) \leq v(\mathbf{x})$; see supplementary material for proof) and is inherently robust: since the products can only decrease along a ray, floating Gaussians hidden inside the geometry cannot inflate the vacancy estimate.

Pivot-Based Marching Tetrahedra. Pivot-based Marching Tetrahedra (PMT) is the meshing procedure introduced by GOF [54]. We adopt PMT, but make it more efficient and principled, as detailed below. Under our oriented Gaussian assumption, the surface intersects each Gaussian between its center and the low-density side, parallel to the oriented plane. We therefore spawn two *Delaunay pivots* per Gaussian: the center μ_i and $\mu_i + 3s_i n_i$, where $s_i = \|S_i R_i^T n_i\|_2$ is the ellipsoid scaling along n_i. Occupancy values at each pivot are obtained using the vacancy estimate $\tilde{v}_{\mathrm{T}}$ (Eq. 7); we then apply Marching Tetrahedra on the resulting Delaunay triangulation and refine vertices to the 0.5-isosurface via binary search. This approach is conceptually similar to GOF [54] and GGGS [56], yet contrasts with the 9 pivots per Gaussian required by prior works [17, 54–56], yielding substantially lighter, fully watertight meshes without sacrificing surface fidelity.

Primal Adaptive Meshing. Although simple and efficient, the pivot-based formulation instantiates vertices directly from Gaussian primitives rather than from the underlying scene geometry. To leverage our global vacancy estimate $\tilde{v}_{\mathrm{T}}$ and further improve mesh generation, we introduce an adaptive meshing framework inspired by primal Delaunay-based reconstruction methods [1, 13, 36, 45–47, 57]. The pipeline proceeds in four stages:

1. **Vertices Initialization.** We initialize our vertices by sampling the marching tetrahedra mesh faces, weighting each face inversely proportional to its distance from the nearest training camera.

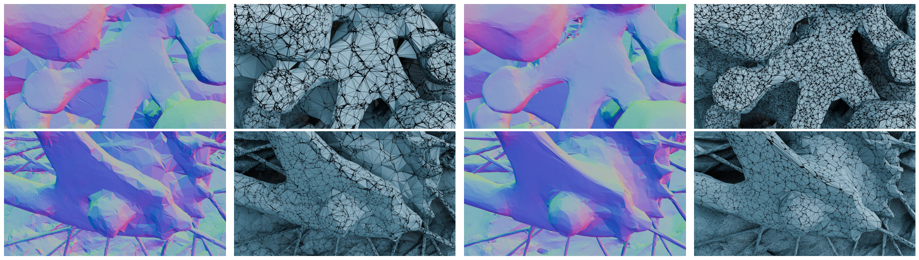

(a) Pivot-based Marching Tetrahedra **(b)** Primal Adaptive Mesh

Fig. 4: Primal Adaptive Meshing on restricted scene segments. Close-up comparisons between our standard Pivot-based MTet extraction (a) and the Primal Adaptive Mesh (b) applied to the same segment on the bicycle and bonsai scenes of the MipNeRF 360 dataset. The Primal Adaptive Mesh produces topologically clean surfaces that faithfully reflect the underlying Gaussian isosurface.

2. **Isosurface Refinement.** Vertices are projected onto the 0.5-isosurface of $\tilde{v}_{\mathrm{T}}$ via an iterative Newton update (derived in the supplementary material):

$$x_{i+1} = x_i + \frac{(0.5 - \tilde{v}_{\mathrm{T}}(x_i))\,\mathcal{N}(x_i)}{2\,\tilde{v}_{\mathrm{T}}(x_i)\,\|V(x_i)\|}, \tag{8}$$

3. **Filtering.** Vertices x with $|0.5 - \tilde{v}_{\mathrm{T}}(x)| > \epsilon$ are considered outliers and removed. Stages 1–3 iterate until no points are removed.
4. **Delaunay.** We compute the Delaunay tetrahedralization of the remaining vertices and classify each tetrahedron as inside or outside, following [57], based on the value of $\tilde{v}_{\mathrm{T}}$ at randomly sampled interior points.
5. **Mesh extraction.** The final surface mesh is obtained by extracting triangle faces that separate inside and outside tetrahedra (see Fig. 3).

This approach effectively decouples mesh resolution from Gaussian distribution, enabling extremely fine meshing of restricted segments of the scene (Fig. 4). Although global distance-based metrics remain largely unchanged —being inherently insensitive to high-frequency detail— the resulting meshes exhibit substantially improved smoothness and reduced discretization artifacts. A derivation of the Newton update (Eq. (8)) and further details on the isosurface refinement stage are provided in the supplementary material.

5 Experiments

In this section, we evaluate and ablate Gaussian Wrapping on standard benchmarks. We expose a fundamental bias in the standard evaluation protocol for Gaussian-based surface reconstruction, and propose two complementary evaluations in Sec. 5.2: *Uniform Sampling* and *Virtual Scanning*. We refer the reader to the supplementary material for all of our implementation details.

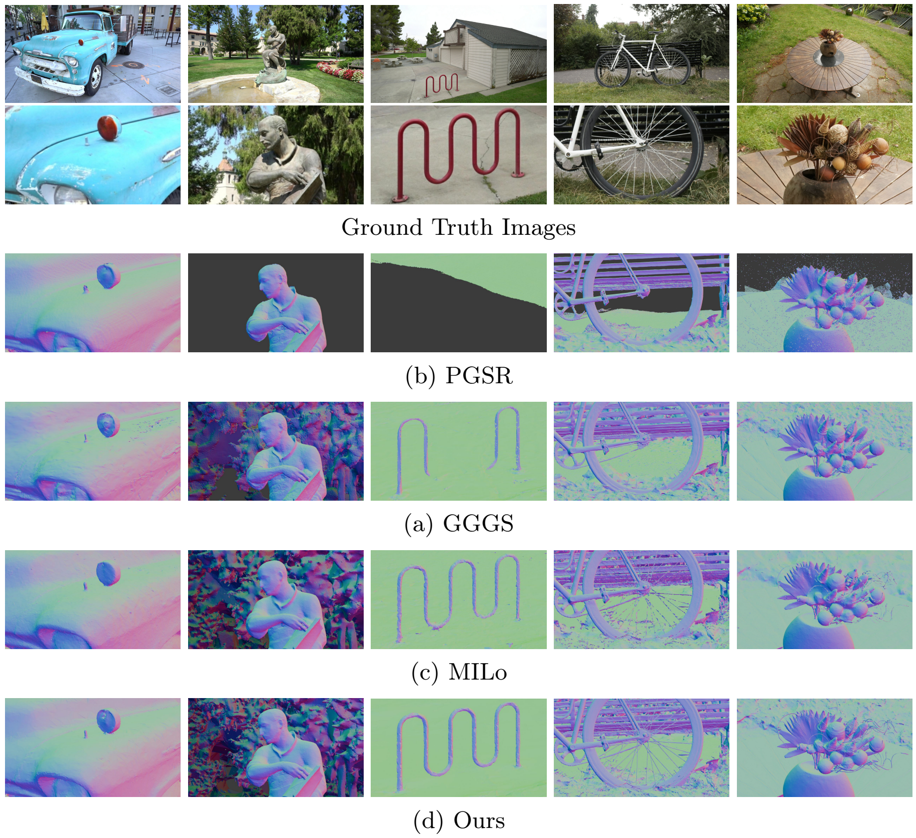

Fig. 5: Qualitative comparison on T&T and MipNerf360. Our method recovers significantly finer surface detail and thin structures compared to baselines. Close-up insets highlight regions where competing methods erode geometry or fail to close holes.

5.1 Experimental Setup

Datasets & Evaluations. We perform quantitative and qualitative evaluations on the widely used **DTU** [21] and **Tanks and Temples (T&T)** [26] datasets. Furthermore, we leverage the Mesh-Based Rendering (MBR) evaluation introduced in MILo [17] to evaluate the mesh completeness.

Baselines. We compare our method against state-of-the-art explicit surface reconstruction techniques. We separate these into two categories: Methods capable of full scene reconstruction (GOF [54], SOF [41], MILo [17], RaDe-GS [55], GGGS [56]), and methods focusing on foreground only (PGSR [9], 2DGS [20]).

Metrics. We report Chamfer Distance (CD) on the DTU dataset; F1-score on the T&T dataset; and image metrics for NVS on the Mip-NeRF 360 dataset. Moreover, we present image metrics for the MBR evaluation on T&T and Mip-NeRF 360 datasets. For *Uniform Sampling*, *Virtual Scanning*, and *Mesh-Based Rendering*, we run all baseline methods ourselves under the exact same protocol.

5.2 Revisiting Surface Evaluation for Gaussian Splatting

The standard evaluation practice on T&T and structured scan datasets is fundamentally flawed for two reasons: (1) the conventional approach builds the predicted surface point cloud from the vertices and face centers of the extracted mesh, biasing metrics toward denser tessellations; and (2) ground data (GT) acquisition, relying on laser scan, unfairly penalizes methods that correctly reconstruct occluded geometry and surfaces visible at a grazing angle. To resolve this, we introduce two robust evaluation strategies (Tab. 1).

Uniform Sampling Evaluation. To eliminate vertex-density bias, we uniformly sample a fixed number of points from the reconstructed mesh surface within the GT crop volume, ensuring a strictly fair geometric comparison.

Virtual Scanning Evaluation. Uniform sampling alone does not take into account incomplete meshes that mimic GT bias. We thus propose to simulate the original data acquisition process by rendering depth maps of the reconstructed meshes from the input camera poses and back-projecting them into the GT crop volume.

Table 1: Quantitative comparison on the Tanks & Temples Dataset [26]. F1 scores under two complementary evaluation protocols. *Uniform Sampling*: points are uniformly sampled from the mesh surface, eliminating vertex-count bias. *Virtual Scanning*: depth maps are back-projected from the original camera viewpoints, mirroring the ground truth acquisition process. We report average training time. RM and MS are the non-Gaussian baselines Radiance Meshes [32] and MeshSplatting [19] (see supplementary).

	Foreground only			Full scene extraction							
	2DGS	PGSR	**Ours (PAM)**	RM	MS	GOF	SOF	RaDe-GS	MILo	GGGS	**Ours (2p)**
Uniform Sampling Evaluation											
Barn	0.48	0.65	0.63	0.04	0.04	0.40	0.45	0.42	0.40	0.57	0.64
Caterpillar	0.24	0.45	0.38	0.11	0.02	0.23	0.20	0.22	0.24	0.41	0.41
Courthouse	0.13	0.24	0.07	0.11	0.02	0.10	0.07	0.12	0.11	0.10	0.19
Ignatius	0.51	0.81	0.76	0.13	0.17	0.55	0.68	0.66	0.72	0.79	0.77
Meetingroom	0.18	0.33	0.28	0.14	0.07	0.19	0.15	0.22	0.14	0.36	0.37
Truck	0.46	0.63	0.46	0.12	0.10	0.41	0.39	0.45	0.43	0.50	0.50
Mean	0.33	0.52	0.43	0.11	0.07	0.31	0.32	0.35	0.34	0.45	0.48
Virtual Scanning Evaluation											
Barn	0.51	0.65	0.67	0.26	0.36	0.49	0.49	0.47	0.52	0.62	0.67
Caterpillar	0.31	0.50	0.50	0.20	0.21	0.29	0.31	0.30	0.33	0.52	0.50
Courthouse	0.10	0.18	0.16	0.11	0.10	0.11	0.06	0.11	0.12	0.16	0.18
Ignatius	0.61	0.82	0.78	0.35	0.53	0.63	0.63	0.66	0.74	0.80	0.78
Meetingroom	0.19	0.35	0.41	0.15	0.18	0.24	0.23	0.25	0.23	0.42	0.41
Truck	0.54	0.66	0.63	0.25	0.40	0.50	0.48	0.57	0.56	0.65	0.63
Mean	0.38	0.53	0.53	0.22	0.30	0.38	0.37	0.39	0.42	0.53	0.53
Times	12 m	45 m	27 m	378 m	35 m	69 m	17 m	12 m	150 m	32 m	27 m

Legacy Evaluation. For backward compatibility, we report the standard vertex-based metric. However, this metric is highly sensitive to tessellation density (e.g., extracting our meshes with 9 pivots instead of 2 artificially inflates scores). Please see the supplementary material for a detailed discussion of these flaws.

5.3 Main Results

Geometry on DTU. We report Chamfer Distance across 15 scenes of the DTU dataset in the supplementary material. Our method yields highly competitive results, outperforming most prior methods and achieving scores comparable to GGGS [56]. Importantly, our method does not exhibit the surface erosion artifacts observed in GGGS, as illustrated in Fig. 5. We note that we outperform foreground only methods such as 2DGS and PGSR on this dataset.

Geometry on T&T. We evaluate on the T&T dataset under three complementary protocols, all reported in Tab. 1.

- *Uniform Sampling Evaluation.* Our method sets a new state of the art under this unbiased protocol among the full scene extraction methods. PGSR [9] remains an apparent outlier: its use of TSDF fusion and depth filtering yields non-watertight meshes whose holes coincide with those of the ground truth scan, an artifact of the acquisition process rather than genuine geometric quality. We document this in the supplementary material. We note that Ours (PAM)'s performance takes a toll on this metric. Uniformly sampling points to reconstruct scenes such as Courthouse or Meetingroom is far from optimal. We highlight that the advantage of this meshing procedure is the meshing of regions of interest at arbitrary resolution (Fig. 4).
- *Virtual Scanning Evaluation.* Simulating the acquisition bias of the ground truth via our virtual scanning procedure, we set a new state of the art alongside GGGS, and see that PGSR has no longer an advantage when this bias is taken into account.

Mesh-Based Rendering on T&T and Mip-NeRF 360. To complement our geometric evaluation, we report Mesh-Based Rendering (MBR) metrics [17], which assesses background reconstruction quality and serve as a proxy for mesh completeness. Results are shown in Tab. 2. Despite lacking the dedicated rendering prior of MILo [17], our method performs competitively across all datasets. This evaluation also exposes a critical weakness of PGSR: its depth-filtering strategy leaves substantial holes in background geometry, resulting in severely degraded rendering quality—confirming that its strong geometric scores are an artifact of the evaluation protocol rather than true reconstruction quality.

Novel View Synthesis on Mip-NeRF 360. Although NVS is not the primary objective of our method, we achieve competitive performance on Mip-NeRF 360 [5] and set a new state of the art for outdoor scenes, reflecting the effectiveness of our representation in challenging, unbounded settings. Results are reported in the supplementary material.

Qualitative Comparisons. Fig. 5 presents visual comparisons against all baselines. Gaussian Wrapping consistently produces surfaces with sharper fine details, fewer topological artifacts, and better coverage of thin structures such as bicycle spokes, corroborating the quantitative findings.

Table 2: Mesh-Based Novel View synthesis metrics on real-world datasets. We report PSNR, SSIM, and LPIPS metrics, as well as the total number of Gaussians and vertices (in millions). Our method demonstrates superior performance on these challenging unbounded scenes.

	MipNeRF 360					Tanks & Temples				
	PSNR ↑	SSIM ↑	LPIPS ↓	#Gaussians ↓	#Verts ↓	PSNR ↑	SSIM ↑	LPIPS ↓	#Gaussians ↓	#Verts ↓
Foreground only										
2DGS	15.36	0.4987	0.4749	1.88	4.31	14.23	0.5697	0.4854	0.98	16.39
PGSR	15.63	0.5834	0.4509	3.16	22.78	14.55	0.5956	0.4819	1.47	11.79
Full scene extraction										
GOF	24.25	0.7017	0.3454	2.99	32.80	20.10	0.6475	0.4073	1.25	11.63
RaDe-GS	24.84	0.7291	0.3128	2.91	30.85	20.70	0.6767	0.3876	1.18	10.06
MILo	25.075	0.7339	0.3096	0.46	6.73	21.198	0.6908	0.3782	0.28	4.36
GGGS	24.55	0.7386	0.2986	4.08	43.46	20.70	0.6912	0.3734	1.73	21.81
Ours	25.10	0.752	0.289	4.0	11.79	20.98	0.7004	0.3674	2.07	5.80

Conclusion. Our results demonstrate that Gaussian Wrapping achieves robust state-of-the-art performance across all evaluation settings. While methods such as PGSR appear competitive under specific protocols by exploiting biases in GT data, the TSDF extraction limits their quality. Our method performs consistently across both legacy and newly proposed metrics, and is further validated by complementary rendering-based and qualitative evidence. This consistency confirms our improvements reflect genuine advances in surface reconstruction.

5.4 Ablation Studies

All ablation studies are conducted on T&T, reporting F1 scores and MBR metrics (SSIM, PSNR, LPIPS). More details are provided in supplementary material.

Normal Alignment Losses. Tab. 3 reports results with and without our normal alignment losses. F1 scores remain stagnant—our rasterizer already places Gaussians near the isosurface—but MBR metrics improve substantially: our losses align Gaussians with the surface, enabling clean mesh extraction with only 2 pivots. Without them, fine details such as thin structures and sharp edges are lost. This also serves as an illustration that the T&T dataset and ground truth is best to measure coarse geometry alignment, and fails to measure detail loyalty.

Generalizability of Gaussian Wrapping. To confirm that the Gaussian Wrapping approach is architecture-agnostic, we plug it into RaDe-GS [55]. The consistent improvement across all T&T scenes (Tab. 3; and figure in supplementary material) demonstrates that our method operates as an effective drop-in regularizer well beyond our specific pipeline. Note that the difference between the RaDe-GS + Gaussian Wrapping combination and Ours in previous tables is the use of our depth rasterization, which precisely queries the 0.5-isosurface via binary search.

6 Conclusion

We presented Gaussian Wrapping, a principled method grounded on the OaV framework. Our method reconstructs complete scenes including extremely thin structures where prior methods fail, without surface erosion artifacts, and at

Table 3: Ablation studies on Tanks & Temples. *Top*: Mean scores across all scenes for our component ablation, reporting F1-score under both evaluation protocols and Mesh-Based Rendering (MBR) metrics. *Bottom*: Per-scene Virtual Scan F1-score demonstrating that our contributions generalize as a plug-in regularizer to RaDe-GS [55]. The normal alignment loss and densification procedure constitute the Gaussian Wrapping approach.

Component Ablation (Mean across scenes)

	F1-Score ↑		MBR Metrics		
Configuration	**Virt. Scan**	**Mesh Qual.**	**PSNR ↑**	**SSIM ↑**	**LPIPS ↓**
Baseline	0.53	0.48	20.74	0.6958	0.3718
+ Normal Alignment Loss	0.52	0.48	20.78	0.6986	0.3683
+ Normal Alignment Loss + Densification	0.53	0.48	20.98	0.7004	0.3674

Gaussian Wrapping Generalization to RaDe-GS (Virtual Scan F1-Score)

Configuration	Barn	Caterpillar	Courthouse	Ignatius	Meetingroom	Truck	**Mean**
RaDe-GS	0.47	0.30	0.11	0.66	0.25	0.57	0.39
RaDe-GS + Gaussian Wrapping	0.63	0.46	0.13	0.71	0.34	0.60	0.48

a fraction of the mesh weight of competitors. We further expose fundamental biases in standard evaluation protocols and propose two alternatives. Finally, Primal Adaptive Meshing enables extraction of meshes with better topology and controllable resolution.

Limitations and Future Work. Our Primal Adaptive Meshing currently relies on uniform sampling of an initial MTet mesh, which can be suboptimal for highly detailed scenes. Developing more principled sampling strategies—for instance, guided by the Gaussian Vector Field or local surface curvature—would extend the approach to finer and more intricate geometry. Although Marching Tetrahedra guarantees watertightness, our method is not tailored to specular, transparent, or sky regions, where the reconstructed surface may be deformed or absent; we show examples in the supplementary material. Moreover, our attenuation-based formulation is not limited to 3DGS: extending it to more accurate volumetric rendering models such as EVER [31] is a natural and promising direction. Finally, the Gaussian vector field V defines a local vacancy v, that could be recovered and supervised directly through an alignment loss, a promising direction we leave to future work.

7 Acknowledgements

We are grateful to George Drettakis for insightful discussions. Parts of this work were supported by the ERC Consolidator Grant "VEGA" (No. 101087347), the ERC Advanced Grant "explorer" (No. 101097259), as well as gifts from Ansys Inc., and Adobe Research.

References

1. Amenta, N., Bern, M., Kamvysselis, M.: A new Voronoi-based surface reconstruction algorithm. In: Proceedings of the 25th annual conference on Computer graphics and interactive techniques - SIGGRAPH '98. pp. 415–421. ACM Press (1998)
2. Amenta, N., Choi, S., Kolluri, R.K.: The power crust. In: Proceedings of the sixth ACM symposium on Solid modeling and applications. pp. 249–266. SMA '01, Association for Computing Machinery (2001)
3. Aurenhammer, F.: Voronoi diagrams—a survey of a fundamental geometric data structure. ACM computing surveys (CSUR) **23**(3), 345–405 (1991)
4. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin Brualla, R., Srinivasan, P.P.: Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In: Int. Conf. Comput. Vis. pp. 5855–5864 (2021)
5. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5460–5469. IEEE (2022)
6. Binninger, A., Wiersma, R., Herholz, P., Sorkine-Hornung, O.: TetWeave: Isosurface Extraction using On-The-Fly Delaunay Tetrahedral Grids for Gradient-Based Mesh Optimization. ACM Trans. Graph. **44**(4), 1–19 (2025)
7. Celarek, A., Kopanas, G., Drettakis, G., Wimmer, M., Kerbl, B.: Does 3d gaussian splatting need accurate volumetric rendering? In: Computer Graphics Forum. vol. 44, p. e70032. Wiley Online Library (2025)
8. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: TensoRF: Tensorial Radiance Fields. In: Eur. Conf. Comput. Vis. pp. 333–350 (2022)
9. Chen, D., Li, H., Ye, W., Wang, Y., Xie, W., Zhai, S., Wang, N., Liu, H., Bao, H., Zhang, G.: PGSR: Planar-based Gaussian Splatting for Efficient and High-Fidelity Surface Reconstruction. IEEE Transactions on Visualization and Computer Graphics **31**(9), 6100–6111 (2024)
10. Chen, H., Li, C., Lee, G.H.: NeuSG: Neural Implicit Surface Reconstruction with 3D Gaussian Splatting Guidance. arXiv preprint arXiv:2312.00846 (2023)
11. Chen, H., Wei, F., Li, C., Huang, T., Wang, Y., Lee, G.H.: VCR-GauS: View Consistent Depth-Normal Regularizer for Gaussian Surface Reconstruction. In: Adv. Neural Inform. Process. Syst. (2024)
12. Dey, T.K., Goswami, S.: Tight cocone: a water-tight surface reconstructor. In: Proceedings of the eighth ACM symposium on Solid modeling and applications. pp. 127–134. ACM (2003)
13. Dey, T.K., Levine, J.A.: Delaunay meshing of isosurfaces. The Visual Computer **24**(6), 411–422 (2008)
14. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance Fields without Neural Networks. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5501–5510 (2022)
15. Gomez, D., Philip, J., Kaiser, A., Michel, É.: RRM: Relightable assets using Radiance guided Material extraction. In: Computer Graphics International Conference. pp. 17–41. Springer (2024)
16. Govindarajan, S., Rebain, D., Yi, K.M., Tagliasacchi, A.: Radiant Foam: Real-Time Differentiable Ray Tracing. In: Int. Conf. Comput. Vis. (2025)
17. Guédon, A., Gomez, D., Maruani, N., Gong, B., Drettakis, G., Ovsjanikov, M.: MILo: Mesh-In-the-Loop Gaussian Splatting for Detailed and Efficient Surface Reconstruction. ACM Trans. Graph. **44**(6), 1–15 (2025)

18. Guédon, A., Lepetit, V.: SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5354–5363 (2024)
19. Held, J., Son, S., Vandeghen, R., Rebain, D., Gadelha, M., Zhou, Y., Cioppa, A., Lin, M.C., Van Droogenbroeck, M., Tagliasacchi, A.: MeshSplatting: Differentiable Rendering with Opaque Meshes. arXiv preprint arXiv:2512.06818 (2025)
20. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In: ACM SIGGRAPH 2024 conference papers. pp. 1–11 (2024)
21. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 406–413 (2014)
22. Jiang, K., Sivaram, V., Peng, C., Ramamoorthi, R.: Geometry Field Splatting with Gaussian Surfels. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5752–5762 (2025)
23. Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: TensoIR: Tensorial Inverse Rendering. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 165–174 (2023)
24. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Trans. Graph. **42**(4) (July 2023)
25. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3D gaussian splatting as markov chain monte carlo. Adv. Neural Inform. Process. Syst. **37**, 80965–80986 (2024)
26. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. ACM Transactions on Graphics **36**(4) (2017)
27. Li, S., Liu, Y.S., Han, Z.: Gaussianudf: Inferring unsigned distance functions through 3d gaussian splatting. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 27113–27123 (2025)
28. Li, Z., Müller, T., Evans, A., Taylor, R.H., Unberath, M., Liu, M.Y., Lin, C.H.: Neuralangelo: High-fidelity neural surface reconstruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 8456–8465 (2023)
29. Liang, Z., Zhang, Q., Feng, Y., Shan, Y., Jia, K.: GS-IR: 3D Gaussian Splatting for Inverse Rendering. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 21644–21653 (2024)
30. Lyu, X., Sun, Y.T., Huang, Y.H., Wu, X., Yang, Z., Chen, Y., Pang, J., Qi, X.: 3DGSR: Implicit Surface Reconstruction with 3D Gaussian Splatting. ACM Trans. Graph. **43**(6), 1–12 (2024)
31. Mai, A., Hedman, P., Kopanas, G., Verbin, D., Futschik, D., Xu, Q., Kuester, F., Barron, J.T., Zhang, Y.: Ever: Exact volumetric ellipsoid rendering for real-time view synthesis. In: Int. Conf. Comput. Vis. pp. 4930–4939 (2025)
32. Mai, A., Hedstrom, T., Kopanas, G., Kontkanen, J., Kuester, F., Barron, J.T.: Radiance meshes for volumetric reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8267–8277 (2026)
33. Mai, A., Verbin, D., Kuester, F., Fridovich-Keil, S.: Neural Microfacet Fields for Inverse Rendering (2023)
34. Mallick, S.S., Goel, R., Kerbl, B., Steinberger, M., Carrasco, F.V., De La Torre, F.: Taming 3DGS: High-quality Radiance Fields with Limited Resources. In: SIGGRAPH Asia Conference Proceedings. pp. 1–11 (2024)
35. Maruani, N., Klokov, R., Ovsjanikov, M., Alliez, P., Desbrun, M.: VoroMesh: Learning Watertight Surface Meshes with Voronoi Diagrams. In: Int. Conf. Comput. Vis. pp. 14565–14574 (2023)

36. Maruani, N., Ovsjanikov, M., Alliez, P., Desbrun, M.: PoNQ: A Neural QEM-Based Mesh Representation. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3647–3657 (Jun 2024)
37. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: Eur. Conf. Comput. Vis. pp. 99–106 (2020)
38. Miller, B., Chen, H., Lai, A., Gkioulekas, I.: Objects as Volumes: A Stochastic Geometry View of Opaque Solids. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 87–97 (June 2024)
39. Müller, T., Evans, A., Schied, C., Keller, A.: Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. ACM Trans. Graph. **41**(4), 102:1–102:15 (Jul 2022)
40. Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022)
41. Radl, L., Windisch, F., Deixelberger, T., Hladky, J., Steiner, M., Schmalstieg, D., Steinberger, M.: SOF: Sorted Opacity Fields for Fast Unbounded Surface Rconstruction. In: SIGGRAPH Asia Conference Proceedings (2025)
42. Reiser, C., Garbin, S., Srinivasan, P.P., Verbin, D., Szeliski, R., Mildenhall, B., Barron, J.T., Hedman, P., Geiger, A.: Binary Opacity Grids: Capturing Fine Geometric Detail for Mesh-Based View Synthesis. SIGGRAPH Conference Proceedings (2024)
43. Ren, H., Yan, Q., Lu, M., Lu, R., Zhu, Z.: 2DGS-R: Revisiting the Normal Consistency Regularization in 2D Gaussian Splatting. arXiv preprint arXiv:2510.16837 (2025)
44. Shen, T., Munkberg, J., Hasselgren, J., Yin, K., Wang, Z., Chen, W., Gojcic, Z., Fidler, S., Sharp, N., Gao, J.: Flexible Isosurface Extraction for Gradient-Based Mesh Optimization. ACM Trans. Graph. **42**(4) (2023)
45. Son, S., Gadelha, M., Zhou, Y., Fisher, M., Xu, Z., Qiao, Y.L., Lin, M.C., Zhou, Y.: DMesh++: An Efficient Differentiable Mesh for Complex Shapes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 26590–26599 (2025)
46. Son, S., Gadelha, M., Zhou, Y., Xu, Z., Lin, M., Zhou, Y.: DMesh: A Differentiable Mesh Representation. Advances in Neural Information Processing Systems **37**, 12035–12077 (2024)
47. Sulzer, R., Landrieu, L., Marlet, R., Vallet, B.: Scalable Surface Reconstruction with Delaunay-Graph Neural Networks. In: Computer Graphics Forum. vol. 40, pp. 157–167. Wiley Online Library (2021)
48. Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. IEEE Conf. Comput. Vis. Pattern Recog. (2022)
49. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: Adv. Neural Inform. Process. Syst. (2021)
50. Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. In: Adv. Neural Inform. Process. Syst. (2021)
51. Yariv, L., Hedman, P., Reiser, C., Verbin, D., Srinivasan, P.P., Szeliski, R., Barron, J.T., Mildenhall, B.: BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. In: ACM SIGGRAPH 2023 conference proceedings. pp. 1–9 (2023)

52. Yu, M., Lu, T., Xu, L., Jiang, L., Xiangli, Y., Dai, B.: GSDF: 3DGS Meets SDF for Improved Rendering and Reconstruction. In: Adv. Neural Inform. Process. Syst. (2024)
53. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-Splatting: Alias-free 3D Gaussian Splatting. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 19447–19456 (2024)
54. Yu, Z., Sattler, T., Geiger, A.: Gaussian opacity fields: Efficient adaptive surface reconstruction in unbounded scenes. ACM Transactions on Graphics (ToG) **43**(6), 1–13 (2024)
55. Zhang, B., Fang, C., Shrestha, R., Liang, Y., Long, X., Tan, P.: RaDe-GS: Rasterizing Depth in Gaussian Splatting. ACM Trans. Graph. (2024)
56. Zhang, B., Jiang, C., Li, H., Shen, S., Tan, P.: Geometry-Grounded Gaussian Splatting. arXiv preprint arXiv:2601.17835 (2026)
57. Zhang, C., Wang, W., Li, X., Liao, X., Su, W., Tao, W.: High-Fidelity Lightweight Mesh Reconstruction from Point Clouds. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11739–11748 (Jun 2025), iSSN: 2575-7075
58. Zhang, W., Liu, Y.S., Han, Z.: Neural Signed Distance Function Inference through Splatting 3D Gaussians Pulled on Zero-Level Set. In: Adv. Neural Inform. Process. Syst. (2024)
59. Zhang, Z., Huang, B., Jiang, H., Zhou, L., Xiang, X., Shen, S.: Quadratic Gaussian Splatting for Efficient and Detailed Surface Reconstruction. In: Int. Conf. Comput. Vis. (2025)