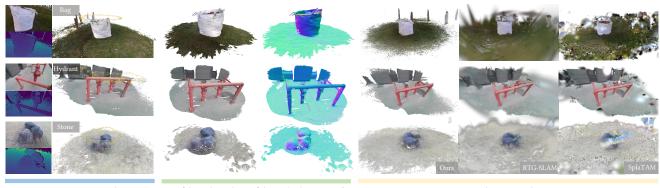
# EGG-Fusion: Efficient 3D Reconstruction with Geometry-aware Gaussian Surfel on the Fly

XIAOKUN PAN, State Key Lab of CAD&CG, Zhejiang University, China ZHENZHE LI, State Key Lab of CAD&CG, Zhejiang University, China ZHICHAO YE\*, SenseTime Research, China HONGJIA ZHAI, State Key Lab of CAD&CG, Zhejiang University, China GUOFENG ZHANG\*, State Key Lab of CAD&CG, Zhejiang University, China



Input Reconstructed Scene Surfels with High Confidence (Color/Normal) Novel View Synthesis

Fig. 1. Qualitative comparison on three real-world scenes reconstructed by our system and the state-of-the-art Gaussian Splatting based SLAM methods (RTG-SLAM[Peng et al. 2024a], SplatAM[Keetha et al. 2024]). Leveraging Gaussian surface representation and surface fusion with information filter, our approach achieves higher geometric accuracy and rendering fidelity (yellow bar). Moreover, our method supports the extraction of high-confidence scene surfaces by maintaining the information matrix of scene primitives (green bar).

Real-time 3D reconstruction is a fundamental task in computer graphics. Recently, differentiable-rendering-based SLAM system has demonstrated significant potential, enabling photorealistic scene rendering through learnable scene representations such as Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS). Current differentiable rendering methods face dual challenges in real-time computation and sensor noise sensitivity, leading to degraded geometric fidelity in scene reconstruction and limited practicality. To address these challenges, we propose a novel real-time system EGG-Fusion, featuring robust sparse-to-dense camera tracking and a geometry-aware Gaussian surfel mapping module, introducing an information filter-based fusion method that explicitly accounts for sensor noise to achieve

\*Guofeng Zhang and Zhichao Ye are Corresponding Authors

Authors' Contact Information: Xiaokun Pan, panxkun@gmail.com, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China; Zhenzhe Li, 12321093@zju.edu.cn, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China; Zhichao Ye, 11721010@zju.edu.cn, SenseTime Research, Hangzhou, China; Hongjia Zhai, zhj1999@zju.edu.cn, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China; Guofeng Zhang, zhangguofeng@zju.edu.cn, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SA Conference Papers '25, Hong Kong, Hong Kong

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-2137-3/2025/12

https://doi.org/10.1145/3757377.3763878

high-precision surface reconstruction. The proposed differentiable Gaussian surfel mapping effectively models multi-view consistent surfaces while enabling efficient parameter optimization. Extensive experimental results demonstrate that the proposed system achieves a surface reconstruction error of 0.6cm on standardized benchmark datasets including Replica and ScanNet++, representing over 20% improvement in accuracy compared to state-of-the-art (SOTA) GS-based methods. Notably, the system maintains real-time processing capabilities at 24 FPS, establishing it as one of the most accurate differentiable-rendering-based real-time reconstruction systems. Project Page: https://zju3dv.github.io/eggfusion/.

CCS Concepts: • Computing methodologies  $\rightarrow$  Point-based models; Rasterization; Reconstruction; Tracking.

Additional Key Words and Phrases: 3D Reconstruction, SLAM, 3D Gaussian Splatting, Gaussian Surfels, Differential Rendering

# **ACM Reference Format:**

Xiaokun Pan, Zhenzhe Li, Zhichao Ye, Hongjia Zhai, and Guofeng Zhang. 2025. EGG-Fusion: Efficient 3D Reconstruction with Geometry-aware Gaussian Surfel on the Fly. In SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25), December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3757377.3763878

# 1 Introduction

Real-time 3D reconstruction technology aims to generate high-precision 3D models by efficiently processing real-world data, which has been extensively studied in the fields of computer graphics and has demonstrated broad application prospects in fields such as mixed reality, autonomous driving, and robotics.

Over the past decades, significant progress has been made in RGBD-based Simultaneous Localization and Mapping (SLAM) [Dai et al. 2017; Kerl et al. 2013; Newcombe et al. 2011a, 2015, 2011b; Whelan et al. 2015; Xu et al. 2022; Zhou and Koltun 2013]. Meanwhile, researchers have proposed various scene representations to better model scenes and improve reconstruction efficiency, such as point clouds [Keller et al. 2013; Kerl et al. 2013; Zhou and Koltun 2013], surfels [Whelan et al. 2015; Xu et al. 2022], and Truncated Signed Distance Fields (TSDF) [Dai et al. 2017; Newcombe et al. 2011a, 2015]. However, these map representations have limited expressive capacity and struggle to handle complex scenes, resulting in limitations in representing precise geometric details and scalability.

In recent years, the rapid development of differentiable rendering techniques has significantly enhanced the realism and visual quality of 3D reconstruction, with notable advancements in technologies such as NeRF [Mildenhall et al. 2020] and 3DGS [Kerbl et al. 2023]. Existing studies [Lu et al. 2024; Yu et al. 2024; Zhai et al. 2024, 2025b] have demonstrated that differentiable rendering techniques can be effectively integrated into SLAM frameworks [Chen et al. 2025; Li et al. 2024; Pan et al. 2024, 2025; Qin et al. 2018]. In particular, realtime reconstruction methods based on 3DGS [Keetha et al. 2024; Peng et al. 2024a] have shown great promise due to their efficient differentiable rasterization capabilities. However, the high degree of freedom in the parameterization of 3DGS primitive introduces geometric ambiguities, which can degrade reconstruction accuracy. This issue is further exacerbated in some complex scenes by factors such as limited viewpoint coverage and sensor noise [Cao et al. 2018; Fontan et al. 2020, 2023], leading to significant reconstruction errors or even complete failure. These challenges pose serious limitations to the reliability and scalability of 3DGS-based real-time reconstruction methods in practical applications.

To address these issues, we propose EGG-Fusion, a novel framework for high-quality and real-time 3D reconstruction in real-world environments. Our approach begins by adopting Gaussian surfels [Dai et al. 2024; Huang et al. 2024b] as the scene representation, enabling consistent multi-view geometric modeling. However, Gaussian surfels alone are insufficient to improve reconstruction quality, as sensor noise remains a pervasive challenge. To address this, we introduce a novel surfel fusion strategy with information filter [Maybeck 1982; Thrun et al. 2004], which incrementally updates the geometric attributes of surfels using depth observations from each frame. This incremental fusion process jointly refines the surfel geometry and its associated information matrix, effectively enhancing the stability and fidelity of the reconstruction. The information matrix also facilitates confidence estimation for each primitive, allowing the system to extract highly reliable reconstruction results. Thanks to frame-wise geometric updates, the scene maintains high-precision geometry throughout reconstruction. Consequently, during map optimization, only minor refinements are required, leading to rapid convergence with significantly reduced computational overhead. To enable robust camera pose estimation, we design a sparse-to-dense tracking strategy that combines the robustness of sparse features with the precision of dense alignment. Additionally, we propose a geometry-aware surfel initialization method, which adaptively determines the density of scene primitives based on depth observations. This strategy ensures accurate

surface alignment while producing a compact and efficient scene representation. Experimental results demonstrate that EGG-Fusion operates in real time at 24 FPS, and consistently outperforms SOTA methods in terms of tracking accuracy, reconstruction quality, and rendering fidelity. As shown in Fig. 1, our method maintains high geometric precision and visual quality, even when tested under challenging, unseen viewpoints.

In summary, the main contributions of our work are as follows:

- We propose EGG-Fusion, a Gaussian surfel-based real-time 3D reconstruction system for high-quality scene modeling in realworld environments. We release the source code to facilitate research reproducibility and community advancement.
- We propose an information-filter-based surfel fusion method that achieves stable high-precision reconstruction by dynamically updating geometric information matrices, while significantly accelerating optimization convergence.
- Extensive experimental results demonstrate that the proposed method outperforms current GS-based SOTA methods in tracking accuracy, reconstruction quality, and rendering quality, while running in real-time at 24 FPS.

## 2 Related Work

#### 2.1 Dense Visual SLAM

In classical dense visual SLAM frameworks, DTAM [Newcombe et al. 2011b] pioneered the dense 3D model reconstruction of an indoor scene from monocular video by directly tracking the camera to the model using photometric error, making it robust to rapid camera movement and defocus blur. [Newcombe et al. 2011a, 2015] utilizes consumer-grade RGB-D cameras (Microsoft Kinect) to accomplish featureless camera tracking by optimizing the transformation of depth information to the TSDF target. [Keller et al. 2013] proposed a global map representation scheme based on point with radius (surfels), enabling reconstruction in dynamic scenes. BundleFusion [Dai et al. 2017] achieves robust pose estimation through an efficient hierarchical approach and attains high-quality real-time reconstruction results using a parallel optimization framework. With the assistance of neural networks, hand-crafted components in classic frameworks are replaced by end-to-end network architectures. [Koestler et al. 2021; Min et al. 2020; Teed and Deng 2021] employ prior-based dense depth or optical flow to perform direct image alignment, making photometric/geometric bundle adjustment (BA) within a keyframe sliding window feasible. Compared to traditional approaches, endto-end architectures [Bloesch et al. 2019; Matsuki et al. 2021; Zhi et al. 2019] simplify the problem's modeling complexity but demand significantly higher computational resources. In addition, these approaches are still constrained by traditional scene representations, lacking flexibility and scalability.

# 2.2 Differentiable-based Scene Representation

Since the introduction of NeRF [Mildenhall et al. 2020], neural-based differentiable rendering has emerged as a promising approach for scene representation. A series of studies have explored improvements in rendering quality [Barron et al. 2021, 2022], joint pose optimization [Lin et al. 2021; Wang et al. 2022], rendering speed [Fridovich-Keil et al. 2022; Müller et al. 2022; Sun et al. 2022],

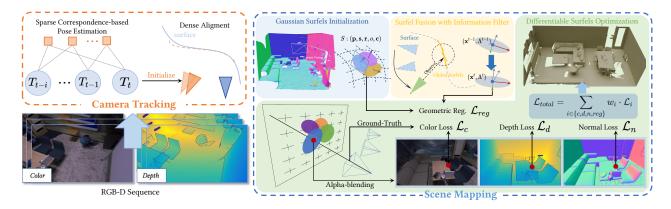


Fig. 2. Framework of EGG-Fusion. Our framework is divided into two integral components. In the scene mapping module (Sec. 3.2), Gaussian surfels are utilized as the fundamental primitives for scene representation and can achieve high-quality real-time reconstruction The camera tracking module (Sec. 3.3) employs a sparse-to-dense strategy to ensure robust estimation of camera poses.

and reconstruction quality [Li et al. 2023; Wang et al. 2023a]. However, challenges remain in terms of slow optimization convergence for scene appearance/geometry and bottlenecks in rendering speed. 3DGS represents a scene using a set of ellipsoids, significantly enhancing the rendering performance of differentiable scene representations by rasterizing these 3D ellipsoids directly in 2D image space. Subsequent works have extended its capabilities, achieving advancements in rendering quality [Lu et al. 2024; Yu et al. 2024], geometry reconstruction [Guédon and Lepetit 2024], and performance on large-scale scenes [Kerbl et al. 2024; Liu et al. 2025]. The variant 2D Gaussian surfels [Dai et al. 2024; Huang et al. 2024b] constrains the z-axis of 3D ellipsoids to enable precise geometric reconstruction of the scene.

#### 2.3 Differentiable Rendering based SLAM

The differentiable property in system-level was first discussed in ∇SLAM [Jatavallabhula et al. 2020], which framed the entire system as a differentiable computation graph. However, it retained traditional 3D representations. iMap [Sucar et al. 2021] was the first to employ differentiable scene representations for SLAM system, leading to a series of subsequent works, including those utilizing hierarchical grids [Zhu et al. 2022], voxels [Yang et al. 2022], hybrid implicit parameters and grids [Wang et al. 2023b], and point-based features [Sandström et al. 2023]. These methods leverage differentiable scene representations for online mapping but suffer from significant efficiency issues, impacting their real-time performance and scalability. More recently, SLAM systems [Keetha et al. 2024; Matsuki et al. 2023; Peng et al. 2024a; Yan et al. 2024] based on 3DGS have demonstrated impressive performance in novel view synthesis and rendering speed, thanks to the efficiency of the scene representation. However, this discrete primitive-based representation has potential limitations in terms of parameter efficiency.

# Methodology

As illustrated in Fig. 2, the system is composed of two modules: camera tracking and scene mapping. This section first introduces the

definition of Gaussian surfels and the associated notation (Sec. 3.1). Subsequently, the proposed scene mapping module is elaborated, and the approach to achieving high-quality scene surface reconstruction using Gaussian surfels as the scene representation is discussed (Sec. 3.2). Finally, a sparse-to-dense camera tracking strategy is proposed (Sec. 3.3) to achieve robust and efficient real-time pose estimation.

#### Preliminary 3.1

3.1.1 2D Gaussian Surfels. The Gaussian surfel is a primitive for learnable scene representation, with a spatial distribution approximated by a disk. Each 2D Gaussian surfel can be represented by the attributes consisting of its center  $\mathbf{p}_i \in \mathbb{R}^3$ , the scales of its two orthogonal Gaussian ellipse axes  $\mathbf{s}_i \in \mathbb{R}^2$ , its rotation in quaternion form  $\mathbf{r}_i \in \mathbb{R}^4$  (in the global coordinate system), its opacity  $o_i \in \mathbb{R}$ , and its color  $\mathbf{c}_i \in \mathbb{R}^k$ . The color attribute is explicitly encoded as coefficients of spherical harmonic basis functions, where the dimension k depends on the defined order, enabling the modeling of non-Lambertian surfaces. In our system, the scene can be represented as a collection of Gaussian surfels, which can be parameterized as  $S = \left\{ S_i : (\mathbf{p}_i, \mathbf{s}_i, \mathbf{r}_i, o_i, \mathbf{c}_i) \right\}_{i=0}^n.$ 

3.1.2 Differentiable Gaussian Splatting. In the differentiable rendering pipeline, each Gaussian surfel  $S_i$  is transformed into image space based on the camera pose, and its 2D covariance  $\Sigma_i^{2D}$  in image space are computed. After sorting by depth, alpha compositing is used for blending to obtain the final color  $\hat{C}$  with the alpha blending weight  $\alpha_i = S_i'(u; u_i, \Sigma_i^{2D}) o_i$ :

$$T_i = \prod_{j=0}^{i-1} (1 - \alpha_j), \quad \hat{C} = \sum_{i=0}^n T_i \cdot \alpha_i \cdot \mathbf{c}_i. \tag{1}$$

Similarly, the depth map  $\hat{D}$  and normal map  $\hat{N}$  are rendered by:

$$\hat{A} = \frac{1}{1 - T_{n+1}} \sum_{i=0}^{n} T_i \cdot \alpha_i \cdot a_i(u), \tag{2}$$

where  $\hat{A}$  is the rendered 2D information (*e.g.*, depth *D* or normal *N*), and *a* is the 3D property.

# 3.2 Scene Mapping with Gaussian Surfels

Given the current frame  $I_t = \{C_t, D_t\}$ , which includes a RGB image  $C_t \in \mathbb{R}^{H \times W \times 3}$  and its corresponding depth map  $D_t \in \mathbb{R}^{H \times W}$ , we pre-process the data using the intrinsic parameters of the camera  $\{f_x, f_y, c_x, c_y\}$  to obtain the normal map  $N_t \in \mathbb{R}^{H \times W \times 3}$  and the vertex map  $V_t \in \mathbb{R}^{H \times W \times 3}$  in the camera coordinate system:

$$V_{t}(\mathbf{u}) = D_{t}(\mathbf{u}) \cdot \mathbf{K}^{-1} \cdot [\mathbf{u}, 1]^{\top}, \quad N_{t} = \frac{\nabla_{x} V_{t} \times \nabla_{y} V_{t}}{|\nabla_{x} V_{t} \times \nabla_{y} V_{t}|}, \quad (3)$$

where  $\mathbf{K} \in \mathbb{R}^{3 \times 3}$  is the intrinsic matrix and  $\mathbf{u} \in \mathbb{R}^2$  is the pixel coordinate. Through the camera tracking strategy (Sec. 3.3), we acquire the camera-to-world pose  $\mathbf{T}_i \in \mathbb{R}^{4 \times 4}$ , enabling us to transform the vertex map and the normal map into the global coordinate system as  $V_t^w$  and  $V_t^w$ .

3.2.1 Geometry-aware Surfel Initialization. The proposed system incrementally expands the map by adaptively integrating Gaussian surfels from incoming RGB-D frames. Departing from previous works [Keetha et al. 2024; Matsuki et al. 2023] that uniformly sample new primitives in image space, we introduce a rendering-aware strategy that selectively activates surfels only in geometrically salient regions: (1) low-opacity zones, revealing reconstruction deficiencies, and (2) areas with positive depth disparity, indicating newly observed foreground geometry. This targeted placement ensures high-fidelity reconstruction while inherently suppressing the redundancy of uniform sampling paradigms. Each newly added Gaussian surfel is initialized by sampling geometric properties from the input depth map, including the surfel center position  $\mathbf{p} = V_t^{\mathbf{w}}(\mathbf{u})$  and surfel normal  $\mathbf{n} = N_t^{\mathbf{w}}(\mathbf{u})$ .

Regarding surfel scale initialization  $\mathbf{s}$ , previous methods [Matsuki et al. 2023; Peng et al. 2024a] employ a fixed value in 3D space. For some Gaussian ellipsoids located far from the camera, their projected area in image space becomes significantly small, necessitating dense sampling to adequately cover the corresponding image patch. To address this, we propose an adaptive scale initialization strategy: surfels located at the camera's far plane are assigned larger scales , determined by the following principle:

$$\mathbf{s} = [\alpha_s \cdot d/f_x, \alpha_s \cdot d/f_u],\tag{4}$$

where  $\alpha_s$  is a pixel scaling factor used to regulate the absolute scale across different sampling rates. The intuitive outcome is that surfels farther from the camera have larger scales, while their projections in the image space maintain a consistent distribution. This is similar to the analysis of sampling rate in [Yu et al. 2024]. Experiments demonstrate that this compact surfel initialization can achieve higher rendering quality.

3.2.2 Surfel Fusion with Information Filter. To mitigate the impact of depth noise from consumer-grade sensors, we propose an information-filter-based surface fusion method. The surfel state estimation is formulated as a Markov process, where the current state  $\mathbf{x}^t$  depends solely on its previous state  $\mathbf{x}^{t-1}$  and the latest sensor observation  $I_t$ . For each input RGB-D frame, we perform

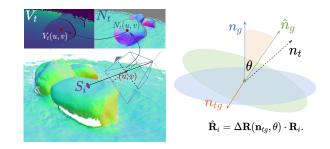


Fig. 3. Gaussian surfels fusion. We ensure that surfels can be explicitly and continuously updated with new observations, enabling them to adhere to the scene surface (left), while new observations are utilized to update normal information (right), thereby achieving more accurate surface reconstruction.

recursive Bayesian updates on reobserved surfels to refine their geometric properties. Each surface element's geometric state is defined as  $\mathbf{x}^t = [\mathbf{p}, \mathbf{n}]^\top \in \mathbb{R}^6$  and is associated with a covariance matrix  $\Sigma^t \in \mathbb{R}^{6 \times 6}$  to quantify measurement uncertainty, enabling progressive confidence accumulation through sequential observations. This probabilistic framework ensures continuous improvement in reconstruction quality, ultimately yielding high-fidelity surfaces.

Specifically, we can obtain the vertex and normal measurements  $\mathbf{z}^t = [V_t(\mathbf{u}), N_t(\mathbf{u})]^{\top} \in \mathbb{R}^6$  based on the image location  $\mathbf{u}$  of the surfel. We construct the relationship between observation and state based on the state observation equation:

$$\mathbf{z}^{t} = \mathbf{H}\mathbf{x}^{t} + \bar{\mathbf{t}} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \Sigma_{\mathbf{z}}^{t}),$$

$$\mathbf{H} = \begin{bmatrix} \mathbf{R} & \mathbf{0} \\ \mathbf{0} & \mathbf{R} \end{bmatrix} \in \mathbb{R}^{6 \times 6}, \quad \bar{\mathbf{t}} = \begin{bmatrix} \mathbf{t} \\ \mathbf{0} \end{bmatrix} \in \mathbb{R}^{6}.$$
(5)

Here **H** denotes the observation matrix, which transforms the state variables from the global coordinate system into the camera coordinate system with current camera transformation [**R**, **t**].  $\epsilon$  is the measurement noise, and  $\Sigma_{\mathbf{z}}^t$  corresponds to the associated covariance. To ensure real-time computational efficiency,  $\Sigma_{\mathbf{z}}^t$  is simplified to a diagonal matrix represented in vector form as  $\mathrm{diag}(\Sigma_{\mathbf{z}}^t) = [\sigma_p, \sigma_p, \sigma_p, \sigma_n, \sigma_n, \sigma_n]$  that  $\sigma_p$  and  $\sigma_n$  correspond to the position and normal respectively. In practice, owing to the inherent characteristics of the sensor, the noise intensity correlates with depth, and thus  $\sigma_p$  and  $\sigma_n$  are proportional to the square of the depth.

We update the information matrix  $\Lambda$  and corresponding information vector matrix  $\eta$  using the following equations:

$$\Lambda^{t} = \Lambda^{t-1} + \mathbf{H}^{\top} \Lambda_{z}^{t} \mathbf{H}, \quad \boldsymbol{\eta}^{t} = \boldsymbol{\eta}^{t-1} + \boldsymbol{\eta}_{z}^{t}, 
\Lambda^{t-1} = (\Sigma^{t-1})^{-1}, \quad \boldsymbol{\eta}^{t-1} = \Lambda^{t-1} \cdot \mathbf{x}^{t-1}, 
\Lambda_{z}^{t} = (\Sigma_{z}^{t})^{-1}, \quad \boldsymbol{\eta}_{z}^{t} = \mathbf{H}^{\top} \cdot \Lambda_{z}^{t} \cdot \mathbf{z}^{t},$$
(6)

In fact, due to the diagonal nature of the covariance matrix, the solution based on Eq. 6 can be greatly simplified, enabling efficient computation of the update. Finally, the updated geometric state of gaussian surfels and corresponding covariance are computed by

$$\hat{\mathbf{x}}^t = (\mathbf{\Lambda}^t)^{-1} \boldsymbol{\eta}^t, \quad \hat{\Sigma}^t = (\mathbf{\Lambda}^t)^{-1}. \tag{7}$$

Ultimately, the updated position and normal from states  $\mathbf{x}^t$  are applied to the attributes  $\mathbf{p}_i$  and  $\mathbf{r}_i$  of surfel  $S_i$ . However, the estimation of relative rotation based on normal updates is underconstrained. To

address this issue, we introduce an additional rotational constraint based on the normal update, resulting in a well-defined 3D rotation. As illustrated in Fig. 3, we define  $\mathbf{n}_{tq} = \mathbf{n}_q \times \mathbf{n}_t$  where  $\mathbf{n}_q$  is the original normal before update and  $\mathbf{n}_t$  is the updated normal. So the vector  $\mathbf{n}_{tq}$  represents the normal to the plane spanned by the two normals. The angle between them is  $\theta = \cos^{-1}(\mathbf{n}_q \cdot \mathbf{n}_t)$ . Based on this, we define a unique rotation represented by the rotation vector  $\Delta$ **R**( $\mathbf{n}_{tq}, \theta$ ). Therefore, for the rotation  $\mathbf{r}_i$  of surfel  $S_i$ , whose matrix representation is  $\mathbf{R}_i$ , the update is performed as follows:

$$\hat{\mathbf{R}}_i = \Delta \mathbf{R}(\mathbf{n}_{tq}, \theta) \cdot \mathbf{R}_i. \tag{8}$$

3.2.3 Differentiable Surfels Optimization. After the explicit Gaussian surfel fusion based on upcoming measurements, we implement end-to-end optimization of the surfels using the rasterization-based rendering pipeline [Kerbl et al. 2023].

Reconstruction Loss. We utilize frame batches to optimize the local map through differentiable rendering, avoiding overfitting to the current frame. Specifically, we maintain a local map composed of the most recent N<sub>batch</sub> frames. After tracking a certain number of frames, we perform batch optimizations, where each iteration randomly selects a frame from the local map. For thorough optimization, we perform a total of  $m \cdot N_{\text{batch}}$  iterations to ensure that each frame is iteratively optimized an average of *m* times. The rendered color map, depth map, and normal map are constrained by the ground truth, resulting in the following loss terms:

$$\mathcal{L}_{c} = \left| C_{k} - \hat{C}_{k} \right|, \mathcal{L}_{d} = \left| D_{k} - \hat{D}_{k} \right|, \mathcal{L}_{n} = \left| 1 - \gamma \right|, \tag{9}$$

where  $\gamma = N_k \cdot \hat{N}_k$  represents the cosine of these two unit normal vector. Eq. (9) utilizes image-space rendering to optimize Gaussian surfels through per-pixel supervision but lacks global consistency. Leveraging the geometric constraints of Gaussian surfels, we introduce per-surfel regularization constraints.

Geometric Regularization: We aim to ensure that the geometric properties of Gaussian surfels remain as consistent as possible during the current iterative optimization (avoiding deviations from the surface), while allowing for appearance changes introduced by new depth map. Therefore, we introduce the following geometric regularization term:

$$\mathcal{L}_{reg} = \left| \mathbf{p} - \mathbf{p}_f \right| + w_{reg}^n \cdot \left| 1 - \mathbf{n} \cdot \mathbf{n}_f \right|, \tag{10}$$

where  $\mathbf{p}_f$  and  $\mathbf{n}_f$  are derived from explicit Gaussian surfel fusion and  $w_{\text{reg}}^n$  is the weighting coefficient for the normal loss of the presurfel. Considering this per-surfel explicit geometric constraints in the end-to-end manner, the overall optimization loss is formulated

$$\mathcal{L}_{total} = \mathcal{L}_c + w_d \cdot \mathcal{L}_d + w_n \cdot \mathcal{L}_n + w_{reg} \cdot \mathcal{L}_{reg}. \tag{11}$$

It is important to note that such regularization are not applicable to 3DGS-based methods [Keetha et al. 2024; Matsuki et al. 2023; Yan et al. 2024]. This is because 3DGS relies on a volumetric representation scheme that lacks explicit constraints with the scene geometry. Additionally, the inconsistency of inferred depths across multiple viewpoints contributes to the instability in its convergence. Additionally, the proposed surfel initialization and fusion facilitate the state of the surfels to be optimized to be near convergence, thereby

Table 1. Tracking performance (ATE RMSE[cm]) on Replica.

Method	Rm0	Rm1	Rm2	Off0	Off1	Off2	Off3	Off4	Avg.
MASt3R [2024]	1.07	0.99	0.87	0.90	4.90	1.21	1.77	1.63	1.67
SLAM3R [2024]	4.56	5.88	5.72	11.17	6.32	6.15	4.95	8.09	6.61
NICE-SLAM [2022]	0.97	1.31	1.07	0.88	1.00	1.06	1.10	1.13	1.06
Vox-Fusion [2022]	0.41	0.50	0.52	0.47	0.61	0.67	0.47	0.47	0.52
Point-SLAM [2023]	0.54	0.43	0.34	0.36	0.45	0.44	0.63	0.72	0.49
SplaTAM [2024]	0.47	0.42	0.32	0.46	0.24	0.28	0.39	0.56	0.39
RTG-SLAM [2024a]	0.20	0.18	0.13	0.22	0.12	0.22	0.20	0.19	0.18
Ours	0.18	0.18	0.11	0.15	0.12	0.19	0.17	0.20	0.17

ensuring that this differentiable surfel optimization can achieve convergence performance of the map state within a few iterations.

# 3.3 Camera Tracking

Given the current frame  $I_t$  and the global model S, camera tracking module aims to estimate the current camera pose  $T_t$ . In optimization problems, we adopt the more compact and structure-preserving Lie algebra representation of camera poses as  $\xi_t \in \mathfrak{se}(3)$ . For the robustness and accuracy of camera tracking, we propose a sparseto-dense strategy.

3.3.1 Sparse-Correspondence-based Pose Initialization. Existing realtime reconstruction methods typically estimate camera poses through dense alignment between the current depth frame and global scene map. While achieving precise pose estimation, these methods are highly initial-value-dependent and prone to tracking failure during rapid camera motions. To address this issue, we introduce a sparse feature correspondence-based pose estimation strategy for pose initialization. On one hand, it provides a stable initialization for the subsequent dense alignment; on the other hand, a good initialization significantly reduces the convergence time during the refinement stage.

Given a set of 2D-3D correspondences  $\mathcal{M} = \{(\mathbf{u}_i, \mathbf{X}_i^w)\}_{i=0}^m$ , where  $\mathbf{u}_i \in \mathbb{R}^2$  denotes a 2D point in the current frame and  $\mathbf{X}_i^w \in \mathbb{R}^3$  is the corresponding 3D point in the sparse map, we formulate the pose estimation as a reprojection error minimization problem. The initial camera pose  $\xi_t^{(0)}$  is estimated via Levenberg–Marquardt (LM)

$$\boldsymbol{\xi}_{t}^{(0)} = \arg\min_{\boldsymbol{\xi}_{t}} \sum_{M} \rho \left( \left| \mathbf{u}_{i} - \Pi \left( \exp(\boldsymbol{\xi}_{t}) \cdot \mathbf{X}_{i}^{w} \right) \right|^{2} \right), \tag{12}$$

Here,  $\exp(\cdot)$  is the exponential map from the Lie algebra to the Lie group and  $\rho(\cdot)$  is a robust loss function used to mitigate the influence of outliers. A similar approach is also adopted in [Huang et al. 2024a]

3.3.2 Refine with Dense Alignment. After obtaining the initial pose, we perform dense alignment for pose refinement, which formulates a nonlinear least-squares problem to jointly optimize the camera pose based on geometric and photometric measurements. Consistent with previous work [Dai et al. 2017; Whelan et al. 2015], we align the vertex map of the current frame  $V_t$  with the vertex map  $V_G$  and normal map  $N_G$  of the global model to construct the geometric error

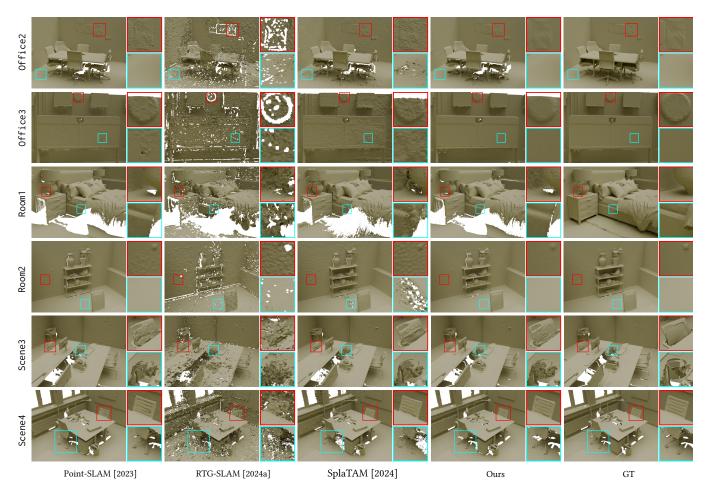


Fig. 4. TSDF-based Reconstruction Result on Replica and ScanNet++. In terms of scene reconstruction mesh details on Replica [Straub et al. 2019] and ScanNet++ [Yeshwanth et al. 2023], we outperform other methods with the overall quality and detail accuracy of the reconstructed mesh.

term, where  $(\cdot)_{\mathcal{G}}$  denote the global model:

$$E(\boldsymbol{\xi}_{t})_{\text{icp}} = \sum_{N} \left| \left( V_{\mathcal{G}} \left( \mathbf{u}_{i}^{\prime} \right) - \exp \left( \boldsymbol{\xi}_{t} \right) V_{t} \left( \mathbf{u}_{i} \right) \right) \cdot N_{\mathcal{G}} \left( \mathbf{u}_{i}^{\prime} \right) \right|_{2}^{2}, \quad (13)$$

Where  $\mathbf{u}_i' = \Pi\left(\exp(\boldsymbol{\xi}) \cdot \Pi^{-1}(\mathbf{u}_i) \cdot d_i\right)$ , and  $\mathcal{N}$  is the set of geometrically matched points after filtering. Subsequently, by comparing the rendered image with the RGB observation, the photometric error is defined as:

$$E(\boldsymbol{\xi}_{t})_{\text{photo}} = \sum_{N} \left| C_{\mathcal{G}} \left( \mathbf{u}_{i}^{\prime} \right) - C_{t} \left( \mathbf{u}_{i} \right) \right|_{2}^{2}, \tag{14}$$

Finally, we adopt a joint optimization strategy to minimize:

$$E(\xi_t)_{\text{dense}} = E_{\text{icp}} + \lambda_{\text{photo}} \cdot E_{\text{photo}}.$$
 (15)

We use  $\lambda_{photo}$  to balance the influence of photometric constraints on pose optimization. To mitigate potential degeneration in extreme cases, we further evaluate the convergence behavior of the dense alignment stage to determine whether its result should be adopted.

# 4 Experiments

## 4.1 Experimental Setup

**Datasets.** We conduct experiments on Replica [Straub et al. 2019], TUM-RGBD [Sturm et al. 2012a] and ScanNet++ [Yeshwanth et al. 2023]. Replica provides camera trajectories and scene depth, which we use to assess localization accuracy, reconstruction quality, and rendering performance. TUM-RGBD offers ground truth trajectories captured by external motion capture devices, serving as the benchmark for evaluating tracking accuracy. ScanNet++ provides highquality images and accurate depth maps, which we use to evaluate reconstruction and rendering quality. For the selection of evaluation sequences, we maintain consistency with previous works [Peng et al. 2024a; Sandström et al. 2023]. Furthermore, we captured three real-world sequences using the consumer-grade Azure Kinect sensor, enabling qualitative evaluation and comparison. This dataset contains three scenes: Bag, Hydrant, and Stone, as shown in Fig. 1. The motion pattern involves slow movements around objects within the scenes. The challenges include: 1) missing depth data in certain regions (due to exceeding the sensor range or measurement failures

Table 2. Tracking performance (ATE RMSE[cm]) on TUM-RGBD. X represents tracking failure.  $^{\ast}$  indicates that the result is taken from the original paper. † denotes the offline variants.

	Method	fr1/desk	fr1/desk2	fr1/room	fr2/xyz	fr3/office	Avg.
zal	ElasticFusion [2015]	2.53	6.83	21.49	1.17	2.52	6.91
Classical	ORB-SLAM2 [2017]	2.16	2.99	16.14	0.50	1.56	4.68
Cle	BAD-SLAM [2019]	1.70	×	Х	1.10	1.70	X
-	NICE-SLAM* [2022]	4.26	4.99	34.49	31.73	3.87	15.87
Differentiable	Vox-Fusion [2022]	2.54	3.95	12.51	1.40	26.01	9.28
nti	Point-SLAM* [2023]	4.34	4.54	30.92	1.31	3.48	8.92
fere	SplaTAM* [2024]	3.35	6.54	11.13	1.24	5.16	5.48
Dif	RTG-SLAM [2024a]	2.30	2.77	17.43	1.15	1.41	5.12
	Ours	2.21	3.09	14.68	0.98	1.41	4.47
91	ORB-SLAM2 <sup>†</sup> [2017]	1.53	2.21	5.38	0.40	0.91	2.09
offline	RTG-SLAM <sup>†</sup> [2024a]	1.61	2.61	4.18	0.38	1.31	2.02
6	Ours <sup>†</sup>	1.56	2.45	4.34	0.43	1.13	1.98

caused by surface reflection and transmission), 2) realistic scene lighting, and 3) continuous orbiting motion.

**Metric.** To evaluate the accuracy of camera tracking, we use ATE RMSE [Sturm et al. 2012b] as the metric. For reconstruction quality, we adopt the following metrics: accuracy, accuracy ratio [<3cm], completeness, and completeness ratio [<3cm], following [Peng et al. 2024a; Zhu et al. 2022]. Regarding rendering performance, we generate full-resolution rendered images and utilize three metrics for evaluation: PSNR [Hore and Ziou 2010], SSIM [Wang et al. 2004], and LPIPS [Zhang et al. 2018].

**Baseline.** The compared methods include differentiable renderingbased SLAM: NICE-SLAM [Zhu et al. 2022], Vox-Fusion [Yang et al. 2022], Point-SLAM [Sandström et al. 2023], SplaTAM [Keetha et al. 2024], and RTG-SLAM [Peng et al. 2024a], and dense visual SLAM: ElasticFusion [Whelan et al. 2015], ORB-SLAM2 [Mur-Artal and Tardós 2017], and BAD-SLAM [Schops et al. 2019]. At the same time, we also include partial comparisons with the latest feed-forward model-based methods including MASt3R [Leroy et al. 2024] and SLAM3R [Liu et al. 2024]. All these methods provide open-source implementations for reproducibility and comparison.

# 4.2 Main Experiments

Evaluation of Tracking. The tracking accuracy results on Replica dataset are shown in Tab. 1. All the results of other methods except Vox-Fusion are from [Peng et al. 2024a]. Compared to other methods, our method achieves the highest average accuracy and attains the best trajectory precision in most scenarios. The superiority of the proposed method in tracking accuracy stems from its sparse-to-dense tracking strategy, as well as the geometry-aware differentiable scene representation. Tab. 2 presents the tracking accuracy comparison of different methods on the TUM-RGBD dataset. We categorize the existing approaches into three groups: (1) Classical methods, (2) Differentiable scene representation based methods, and (3) Offline variants. For the Classical methods, we adopted the results from [Sandström et al. 2023; Zhu et al. 2022], which have also been widely used by the following works [Keetha et al. 2024; Matsuki et al. 2023; Peng et al. 2024a; Yang et al. 2022]. The inclusion of offline variants is motivated by the fact that the official implementations of both RTG-SLAM and ORB-SLAM2 produce globally

Table 3. Reconstruction performance on Replica/ScanNet++. N/A indicates missing data. \* indicates that the result is taken from the original paper or [Peng et al. 2024a].

	Method	Acc.↓	Acc. Ratio↑	Comp.↓	Comp. Ratio↑
	MASt3R* [2024]	4.71 / N/A	78.12 / N/A	3.36 / N/A	79.32 / N/A
	SLAM3R* [2024]	3.57 / N/A	82.23 / N/A	2.62 / N/A	81.29 / N/A
	NICE-SLAM* [2022]	2.84/4.45	84.44/74.49	2.31 /2.04	84.97/86.63
	Vox-Fusion [2022]	2.03/N/A	89.04/N/A	2.58/N/A	86.93/N/A
TSDF	Point-SLAM [2023]	0.76 / 0.67	99.80 / 99.12	2.42/ 0.68	87.46 / 98.94
TS	SplaTAM [2024]	1.12 / 1.70	96.31 / 90.87	2.49/0.83	87.52 / 98.51
	RTG-SLAM [2024a]	1.30/1.86	91.96/82.09	2.37 / 1.00	86.81/98.72
	Ours	0.90 / 1.51	97.60 / 94.20	2.25 / 0.76	87.73 / 99.20
	ElasticFusion[2015]	1.38/N/A	91.72/N/A	7.38/N/A	65.50/N/A
nts	SplaTAM [2024]	2.87/1.71	74.27/89.15	3.58/1.60	71.72/92.64
Points	RTG-SLAM [2024a]	0.80/1.06	98.52/95.34	2.88/1.22	81.92/95.78
	Ours	0.60 / 0.67	99.99 / 99.98	2.67 / 0.91	84.88 / 99.04

Table 4. Rendering performance on ScanNet++.

Methods	Novel View PSNR ↑ SSIM ↑ LPIPS ↓			Training View    PSNR↑ SSIM↑ LPIPS↓			
Point-SLAM [2023]	17.68	0.623	0.548	24.35	0.800	0.373	
SplaTAM [2024]	24.75	0.900	0.208	27.30	0.940	0.130	
RTG-SLAM [2024a]	24.77	0.882	0.255	27.54	0.925	0.184	
Ours	25.70	0.907	0.212	29.06	0.944	0.141	

optimized trajectories (i.e., in offline mode). To ensure fair comparison, we recorded their online output while also developing an offline version of EGG-Fusion as reference, which is implemented similarly to RTG-SLAM's official release, both of which are based on the ORB-SLAM2 backend. After tracking, both perform a global optimization and then output the final trajectory. Experimental results demonstrate that our method achieves the highest average tracking accuracy compared to all real-time SLAM systems in real-world scenarios.

**Evaluation of Reconstruction.** Different systems employ diverse scene representations, leading to significant variations in the format of reconstruction results. For a fair comparison across different methods, we adopt two evaluation schemes following the approach in [Peng et al. 2024a; Sandström et al. 2023]. For explicit reconstruction results, we uniformly sample a fixed number  $(1 \times 10^6)$  of points from the final reconstructed Gaussian primitives for evaluation. For implicit reconstruction result, we utilize depth maps with ground truth camera poses to perform TSDF fusion, followed by mesh extraction via the marching cubes to assess the final reconstruction accuracy [Zhu et al. 2022]. The reconstruction accuracy of different methods on Replica and ScanNet++ is shown in Tab. 3. Under both evaluation methods, our reconstruction accuracy surpasses that of other methods, with the exception of Point-SLAM. It is noteworthy that the high precision of Point-SLAM from its reliance on ground truth depth to determine rendering sampling positions, which grants it an unfair advantage in depth evaluation. In both datasets, our average Acc. Ratio approaches 100%, indicating that nearly all reconstructed surfels lie within 3cm of the true scene surface. The qualitative comparison of reconstruction results is shown in Fig. 4 and Fig. 5.

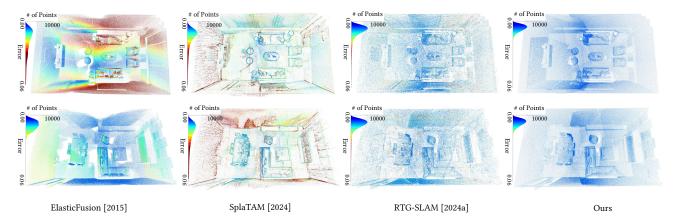


Fig. 5. Geometry accuracy of points. Our method achieved globally high-precision reconstruction. In contrast, other methods exhibited higher errors in either the details of local geometric complexity (RTG-SLAM [Peng et al. 2024a]) or the overall geometric structure (SplaTAM [Keetha et al. 2024] and ElasticFusion [Whelan et al. 2015]).

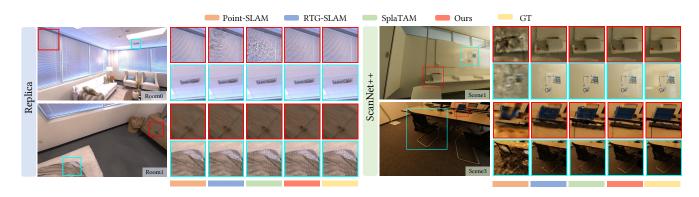


Fig. 6. Rendering Results on Replica and ScanNet++. We present a comparison of novel view synthesis results on Replica [Straub et al. 2019] and ScanNet++ [Yeshwanth et al. 2023]. Our method demonstrates superior rendering details in both training views (Replica, left) and testing views (ScanNet++, right).

Evaluation of Novel View Synthesis. We report the rendering results on the ScanNet++ dataset. As shown in Tab. 4, we achieve the best rendering quality on both the training views and unseen test views. This demonstrates the generalization capability of our Gaussian surfel based scene representation for novel view synthesis. This success can be attributed to the proposed method's ability to tightly adhere Gaussian surfels to the scene surface, ensuring excellent geometric consistency across different viewpoints. The qualitative comparison of rendering results is shown in the Fig. 6. System Performance. We evaluated and reported the system performance of several methods. To ensure a relatively fair comparison, we used the default configuration for each method and conducted time and memory consumption statistics on the Replica office0. Specifically, we measured the time taken for each frame in the main tracking and mapping modules, reporting the average values. We also recorded the number of iterations in tracking and mapping, as well as the average time per iteration. Then we can derive the theoretical FPS for the system under a single-threaded mode. As shown in the Tab. 5, our method outperforms the current SOTA methods in tracking, mapping and overall system FPS. It is worth noting

Table 5. Comparison of time and memory performance on Replica Off0.\* denote the average cost time of sparse correspondence based pose estimation and dense alignment, respectively.

Method	Tracking [ms × it]↓	Mapping [ms × it]↓	Mapping /Frame[s]↓	Model Size [MB]↓	Mem. [GB]↓	FPS↑
NICE-SLAM[2022]	6.6 × 10	28.6× 60	1.717	48.48	9.8	2.91
Vox-Fusion[2022]	$16.5 \times 30$	$34.8 \times 15$	0.675	1.49	7.8	0.75
Point-SLAM[2023]	$10.1 \times 40$	$31.18 \times 300$	9.892	15431	9.8	0.40
SplaTAM[2024]	$45.2 \times 40$	$54.7 \times 60$	3.283	310	9.1	0.19
RTG-SLAM[2024a]	29.1 × 1	$3.5 \times 50$	0.207	51	2.7	15.73
Ours	$(14.2 + 9.4)^*$	$7.5 \times 9$	0.071	150	1.8	24.21

that, thanks to the carefully designed system, scene primitives converge rapidly with minimal updates per optimization, resulting in significantly lower average mapping time per frame compared to other methods. Moreover, our online memory consumption is also significantly lower than that of other approaches.

**Qualitative Comparison.** As shown in Fig. 1, we collected three outdoor object-centric RGB-D sequences using Azure Kinect for

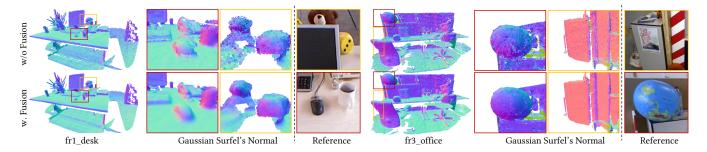


Fig. 7. Ablation of Surfel Fusion with Information Filter. The strategy of surfel fusion with information filter demonstrates robustness against noise, effectively enhancing the accuracy of surface reconstruction on the both scenes (fr1/desk, fr3/office) in TUM-RGBD [Sturm et al. 2012b] dataset.

Table 6. Ablation study of sparse-to-dense tracking.

Method	fr1/desk	fr1/desk2	fr1/room	fr2/xyz	fr3/office	Avg.
w/o. sparse	6.39	9.29	Х	1.43	X	X
full	2.21	3.09	14.68	0.98	1.41	4.47

Table 7. Ablation study of surfel fusion on ScanNet++.

Metric	TS	SDF	Points		
Wetric	w/o fusion	w/o fusion full		full	
ACC./Acc. Ratio	1.59 / 93.42	1.51 / 94.20	0.73 / 99.28	0.67 / 99.98	
Comp./Comp. Ratio	0.80 / 98.73	0.76 / 99.20	1.01 / 98.93	0.91 / 99.04	

qualitative demonstration of our method's performance and comparison with other approaches. Our method enables real-time, online, high-quality object-level reconstruction, while also generating a surface representation based on Gaussian surfels with confidence. Additionally, we render the scene from novel viewpoints. Even under large viewpoint changes compared to the training views, our geometry-aware Gaussian surfel representation achieves significantly higher rendering quality than other methods.

#### Ablation Study 4.3

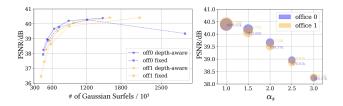


Fig. 8. Depth-aware vs. Fixed scale initialization. The relationship between the number of surfels and rendering quality under different scale initialization strategies (left). the impact of different scale scaling factors on the number of surfels and rendering quality (right).

Sparse-to-Dense Tracking. We evaluate the impact of the sparseto-dense strategy on tracking accuracy on the TUM-RGBD dataset. The results are shown in Tab. 6. The proposed sparse-to-dense tracking strategy benefits from the high-quality initialization provided

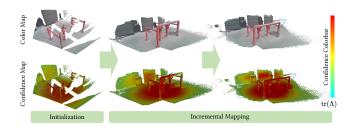


Fig. 9. Confidence Update with Incremental Mapping. We use the trace of information matrix  $tr(\Lambda)$  as the surfel's confidence and visualize it during incremental mapping on the Hydrant scene. The surfels in central part of the scene gradually gain higher confidence as observations accumulate, while peripheral surfels remain low.

by sparse tracking, allowing dense tracking to require minimal optimization. As a result, the poor performance of the w/o sparse variant in complex scenes (e.g., fr1/room, fr3/office) is understandable due to the lack of proper initialization.

**Depth-Aware Scale Initialization.** We compare different surfel scale initialization strategies on the Replica off0/off1 in terms of appearance reconstruction quality. For the fixed scale initialization strategy, we initialize the size of each surfel ranging from 0.002 to 0.010 with an interval of 0.002 (uint: cm). For the proposed depthaware scale initialization strategy, we adjust the scaling factor by tuning  $\alpha_s$  ranging from 1.0 to 3.0 with an interval of 0.5. As shown in Fig. 8, we report the number of surfels and the corresponding PSNR of the final rendered results. The proposed initialization strategy significantly improves rendering quality with the same number of surfels. To balance efficiency and accuracy, We set  $\alpha_s = 2.0$ .

Surfel Fusion with Information Filter. We conducted an ablation study on the surfel fusion strategy. As shown in Fig. 7, we performed qualitative analysis on the TUM-RGBD dataset, the proposed fusion strategy integrates multi-view depth observations, effectively mitigating sensor noise and achieving smoother and more consistent scene reconstruction, such as the smooth surfaces of the cup and mouse. For distant objects (e.g., the teddy bear) where depth measurements are more noise-prone, the proposed fusion method demonstrates particularly notable improvements. Then we evaluated its effectiveness in reconstruction tasks on ScanNet++ using the

Table 8. Ablation study of geometric regularization.

Method	ATE↓	PSNR↑	SSIM↑	LPIPS↓	Acc.↓	Comp.↓
w/o. reg	0.166	39.46	0.988	0.061	0.513	1.389
w. reg	0.158	39.75	0.991	0.058	0.509	1.384

same configuration as in Tab. 3. The results shown in Tab. 7, demonstrate that the proposed fusion strategy significantly improves the average geometric reconstruction metrics on ScanNet++ scenes. In addition, we qualitatively visualized the geometric confidence of the *Hydrant* scene from the Azure dataset in Fig. 9. Here, we use the trace of the information matrix  $\operatorname{tr}(\Lambda)$ , as a proxy for overall confidence, since larger values correspond to lower uncertainty across dimensions. As expected, regions that are continuously observed exhibit higher geometric confidence.

Geometric Regularization. We conducted an ablation study on the geometric regularization on Replica office0 and report metrics such as trajectory accuracy, rendering quality, and reconstruction quality. As shown in Tab. 8, it demonstrates that using the geometric regularization term significantly improves camera tracking accuracy as well as the appearance and geometry reconstruction of the scene.

## 5 Limitations

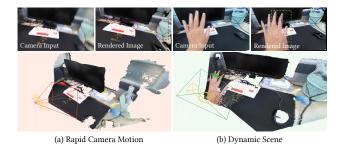


Fig. 10. Failure case of EGG-Fusion. We present the reconstruction scene and rendered images obtained with EGG-Fusion under (a) rapid camera motion and (b) the presence of dynamic objects in the scene.

Although EGG-Fusion demonstrates efficient and stable real-time performance, it still has some limitations. First, real-time mapping using RGB sensors has broader practical applications, but the absence of depth measurements imposes higher demands on the robust estimation of surfel geometry positions and normal attributes. Second, artifacts caused by certain extreme motion patterns (e.g., wide-baseline localization due to rapid motion or rolling shutter effects) may even lead to tracking failure. Finally, our method assumes a static environment, and the presence of dynamic objects or people during scanning may affect reconstruction quality.

As shown in Fig. 10, we present two failure cases. Under rapid camera motion, tracking is highly prone to loss, resulting in significant discrepancies between the rendered image and the current frame. In addition, motion blur in such cases leads to the failure of reconstructing details, making it a critical problem to address the

degradation of system performance under extreme motion conditions [Lu et al. 2025; Peng et al. 2024b; Seiskari et al. 2024]. In dynamic scenes, hand movements introduce ghosting artifacts, which severely affect the rendered image quality. Therefore, an important future direction is to develop robust tracking and reconstruction capabilities for dynamic environments [Dai et al. 2025; Huang et al. 2025; Lei et al. 2025; Wu et al. 2024; Zhai et al. 2025a].

#### 6 Conclusion

We propose EGG-Fusion, a robust real-time RGB-D SLAM system that utilizes Gaussian surfels as the map representation, achieving accurate camera tracking and high-quality Gaussian map reconstruction at 24 FPS. This system excels in delivering high geometric precision, superior map rendering quality, and a more compact map representation. Beyond surpassing existing methods on public benchmarks, our system demonstrates exceptional performance in real-world scenarios with live data. By enabling efficient object- and scene-level map reconstruction, the proposed real-time 3D reconstruction system offering significant potential for applications in robotics, augmented reality, and autonomous driving.

# Acknowledgments

This work was partially supported by NSF of China (No. 62425209)

## References

Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*. 5855–5864.

Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 5470–5479.

Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J. Davison. 2019. CodeSLAM - Learning a Compact, Optimisable Representation for Dense Visual SLAM. arXiv:1804.00874 [cs]

Yan-Pei Cao, Leif Kobbelt, and Shi-Min Hu. 2018. Real-time high-accuracy threedimensional reconstruction with consumer RGB-D cameras. ACM Transactions on Graphics (TOG) 37, 5 (2018), 1–16.

Guyuan Chen, Xiyue Guo, Xiaokun Pan, Yujun Shen, Guofeng Zhang, Hujun Bao, and Zhaopeng Cui. 2025. DW-VIO: Deep Weighted Visual-Inertial Odometry. In 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE.

Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. 2017. BundleFusion: Real-time Globally Consistent 3D Reconstruction Using Onthe-fly Surface Re-integration. arXiv:1604.01093

Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. 2024. High-Quality Surface Reconstruction Using Gaussian Surfels. arXiv:2404.17774Pinxuan Dai, Peiquan Zhang, Zheng Dong, Ke Xu, Yifan Peng, Dandan Ding, Yujun Shen, Yin Yang, Xinguo Liu, Rynson WH Lau, et al. 2025. 4D Gaussian Videos with Motion Layering. ACM Transactions on Graphics (TOG) 44, 4 (2025), 1–14.

Alejandro Fontan, Javier Civera, and Rudolph Triebel. 2020. Information-driven direct rgb-d odometry. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 4929–4937.

Alejandro Fontan, Riccardo Giubilato, Laura Oliva Maza, Javier Civera, and Rudolph Triebel. 2023. Sid-slam: Semi-direct information-driven rgb-d slam. *IEEE Robotics* and Automation Letters 8, 10 (2023), 6387–6394.

Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. 2022. Plenoxels: Radiance Fields without Neural Networks. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, New Orleans, LA LISA 5401, 5500, doi:10.1109/CVPD52688.9322.00542.

New Orleans, LA, USA, 5491–5500. doi:10.1109/CVPR52688.2022.00542
Antoine Guédon and Vincent Lepetit. 2024. 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. 5354–5363.

Alain Hore and Djemel Ziou. 2010. Image quality metrics: PSNR vs. SSIM. In 2010 20th international conference on pattern recognition. IEEE, 2366–2369.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2024b. 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In Special Interest

- Group on Computer Graphics and Interactive Techniques Conference Conference Papers '24. ACM, Denver CO USA, 1-11. doi:10.1145/3641519.3657428
- Gan Huang, Xiaokun Pan, Hengxu Lin, Ziyang Zhang, Weiwei Xu, and Guofeng Zhang. 2025. Real-Time Consistent Monocular Depth Recovery System for Dynamic Environments. In 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE.
- Huajian Huang, Longwei Li, Hui Cheng, and Sai-Kit Yeung. 2024a. Photo-slam: Realtime simultaneous localization and photorealistic mapping for monocular stereo and rgb-d cameras. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 21584-21593.
- Krishna Murthy Jatavallabhula, Soroush Saryazdi, Ganesh Iyer, and Liam Paull. 2020. gradSLAM: Automagically Differentiable SLAM. arXiv:1910.10672 [cs]
- Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. 2024. SplaTAM: Splat, Track & Map 3D Gaussians for Dense RGB-D SLAM. arXiv:2312.02126 [cs]
- Maik Keller, Damien Lefloch, Martin Lambers, Shahram Izadi, Tim Weyrich, and Andreas Kolb. 2013. Real-Time 3D Reconstruction in Dynamic Scenes Using Point-Based Fusion. In 2013 International Conference on 3D Vision. IEEE, Seattle, WA, USA, 1-8. doi:10.1109/3DV.2013.9
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. arXiv:2308.04079
- Bernhard Kerbl, Andreas Meuleman, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, and George Drettakis. 2024. A hierarchical 3d gaussian representation for real-time rendering of very large datasets. ACM Transactions on Graphics (TOG) 43, 4 (2024), 1-15.
- Christian Kerl, Jurgen Sturm, and Daniel Cremers. 2013. Dense Visual SLAM for RGB-D Cameras. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, Tokyo, 2100-2106. doi:10.1109/IROS.2013.6696650
- Lukas Koestler, Nan Yang, Niclas Zeller, and Daniel Cremers. 2021. TANDEM: Tracking and Dense Mapping in Real-time Using Deep Multi-view Stereo. arXiv:2111.07418
- Jiahui Lei, Yijia Weng, Adam W Harley, Leonidas Guibas, and Kostas Daniilidis. 2025. Mosca: Dynamic gaussian fusion from casual videos via 4d motion scaffolds. In  $Proceedings\ of\ the\ Computer\ Vision\ and\ Pattern\ Recognition\ Conference.\ 6165-6177.$
- Vincent Leroy, Yohann Cabon, and Jérôme Revaud. 2024. Grounding image matching in 3d with mast3r. In European Conference on Computer Vision. Springer, 71-91.
- Jinyu Li, Xiaokun Pan, Gan Huang, Ziyang Zhang, Nan Wang, Hujun Bao, and Guofeng Zhang. 2024. RD-VIO: Robust Visual-Inertial Odometry for Mobile Augmented Reality in Dynamic Environments. (2024), 1-14. arXiv:2310.15072 [cs] doi:10.1109/ TVCG 2024 3353263
- Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. 2023. Neuralangelo: High-fidelity neural surface reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 8456-8465.
- Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. 2021. BARF: Bundle-Adjusting Neural Radiance Fields. arXiv:2104.06405 [cs]
- Yuzheng Liu, Siyan Dong, Shuzhe Wang, Yingda Yin, Yanchao Yang, Qingnan Fan, and Baoquan Chen. 2024. SLAM3R: Real-Time Dense Scene Reconstruction from Monocular RGB Videos. arXiv preprint arXiv:2412.09401 (2024)
- Yang Liu, Chuanchen Luo, Lue Fan, Naiyan Wang, Junran Peng, and Zhaoxiang Zhang. 2025. Citygaussian: Real-time high-quality large-scale scene rendering with gaussians. In European Conference on Computer Vision. Springer, 265-282.
- Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 20654-20664.
- Yiren Lu, Yunlai Zhou, Disheng Liu, Tuo Liang, and Yu Yin. 2025. Bard-gs: Bluraware reconstruction of dynamic scenes via gaussian splatting. In Proceedings of the Computer Vision and Pattern Recognition Conference. 16532-16542.
- Hidenobu Matsuki, Riku Murai, Paul H. J. Kelly, and Andrew J. Davison. 2023. Gaussian Splatting SLAM. arXiv:2312.06741 [cs]
- Hidenobu Matsuki, Raluca Scona, Jan Czarnowski, and Andrew J. Davison. 2021. CodeMapping: Real-Time Dense Mapping for Sparse SLAM Using Compact Scene Representations. arXiv:2107.08994 [cs]
- Peter S Maybeck, 1982. Stochastic models, estimation, and control, Vol. 3. Academic
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. arXiv:2003.08934 [cs]
- Zhixiang Min, Yiding Yang, and Enrique Dunn. 2020. VOLDOR: Visual Odometry From Log-Logistic Dense Optical Flow Residuals. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Seattle, WA, USA, 4897-4908. doi:10.1109/CVPR42600.2020.00495
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. ACM Transactions on Graphics 41, 4 (July 2022), 1-15. doi:10.1145/3528223.3530127
- Raul Mur-Artal and Juan D Tardós. 2017. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. IEEE transactions on robotics 33, 5 (2017),

- 1255-1262
- Richard A. Newcombe, Andrew Fitzgibbon, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, and Steve Hodges. 2011a. KinectFusion: Real-time Dense Surface Mapping and Tracking. In 2011 10th IEEE International Symposium on Mixed and Augmented Reality. IEEE, Basel, 127-136. doi:10.1109/ISMAR.2011.6092378
- Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. 2015. DynamicFusion: Reconstruction and Tracking of Non-Rigid Scenes in Real-Time. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Boston, MA, USA, 343-352. doi:10.1109/CVPR.2015.7298631
- Richard A. Newcombe, Steven J. Lovegrove, and Andrew J. Davison. 2011b. DTAM: Dense Tracking and Mapping in Real-Time. In 2011 International Conference on Computer Vision. IEEE, Barcelona, Spain, 2320-2327. doi:10.1109/ICCV.2011.6126513
- Xiaokun Pan, Gan Huang, Ziyang Zhang, Jinyu Li, Hujun Bao, and Guofeng Zhang. 2024. Robust collaborative visual-inertial slam for mobile augmented reality. IEEE Transactions on Visualization and Computer Graphics (2024).
- Xiaokun Pan, Zhenzhe Li, Tianxing Fan, Hongjia Zhai, Hujun Bao, and Guofeng Zhang. 2025. Scalable Multi-Session Visual SLAM in Large-Scale Scenes with Subgraph Optimization. In 2025 IEEE International Conference on Robotics and Automation (ICRA), IEEE
- Cheng Peng, Yutao Tang, Yifan Zhou, Nengyu Wang, Xijun Liu, Deming Li, and Rama Chellappa. 2024b. Bags: Blur agnostic gaussian splatting through multi-scale kernel modeling. In European Conference on Computer Vision. Springer, 293-310.
- Zhexi Peng, Tianjia Shao, Yong Liu, Jingke Zhou, Yin Yang, Jingdong Wang, and Kun Zhou. 2024a. RTG-SLAM: Real-time 3d reconstruction at scale using gaussian splatting. In ACM SIGGRAPH 2024 Conference Papers, 1-11.
- Tong Qin, Peiliang Li, and Shaojie Shen. 2018. VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator. 34, 4 (2018), 1004-1020. doi:10.1109/ TRO.2018.2853729
- Erik Sandström, Yue Li, Luc Van Gool, and Martin R. Oswald. 2023. Point-SLAM: Dense Neural Point Cloud-based SLAM. arXiv:2304.04278 [cs]
- Thomas Schops, Torsten Sattler, and Marc Pollefeys. 2019. BAD SLAM: Bundle Adjusted Direct RGB-D SLAM. In IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE, Long Beach, CA, USA, 134–144. doi:10.1109/CVPR.2019.00022
- Otto Seiskari, Jerry Ylilammi, Valtteri Kaatrasalo, Pekka Rantalankila, Matias Turkulainen, Juho Kannala, Esa Rahtu, and Arno Solin. 2024. Gaussian splatting on the move: Blur and rolling shutter compensation for natural camera motion. In European conference on computer vision. Springer, 160-177.
- Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. 2019. The Replica Dataset: A Digital Replica of Indoor Spaces. arXiv:1906.05797
- Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. 2012a. A Benchmark for the Evaluation of RGB-D SLAM Systems. In 2012 IEEE/RS3 International Conference on Intelligent Robots and Systems. 573-580. doi:10.1109/ IROS.2012.6385773
- Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. 2012b. A benchmark for the evaluation of RGB-D SLAM systems. In 2012 IEEE/RSJ international conference on intelligent robots and systems. IEEE, 573-580.
- Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J. Davison. 2021. iMAP: Implicit Mapping and Positioning in Real-Time. In 2021 IEEE/CVF International Conference on Computer Vision (ICCV). IEEE, Montreal, QC, Canada, 6209-6218. doi:10.1109/ ICCV48922.2021.00617
- Cheng Sun, Min Sun, and Hwann-Tzong Chen. 2022. Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, New Orleans, LA, USA, 5449-5459. doi:10.1109/CVPR52688.2022.00538
- Zachary Teed and Jia Deng. 2021. DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras. In Advances in Neural Information Processing Systems, Vol. 34. Curran Associates, Inc., 16558-16569. https://proceedings.neurips.cc/paper/ 2021/hash/89fcd07f20b6785b92134bd6c1d0fa42-Abstract.html
- Sebastian Thrun, Yufeng Liu, Daphne Koller, Andrew Y Ng, Zoubin Ghahramani, and Hugh Durrant-Whyte. 2004. Simultaneous localization and mapping with sparse extended information filters. The international journal of robotics research 23, 7-8
- Hengyi Wang, Jingwen Wang, and Lourdes Agapito. 2023b. Co-SLAM: Joint Coordinate and Sparse Parametric Encodings for Neural Real-Time SLAM. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Vancouver, BC, Canada, 13293-13302. doi:10.1109/CVPR52729.2023.01277
- Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2023a. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. arXiv:2106.10689

- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image* processing 13, 4 (2004), 600–612.
- Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. 2022. NeRF-: Neural Radiance Fields Without Known Camera Parameters. arXiv:2102.07064
- Thomas Whelan, Stefan Leutenegger, Renato Salas Moreno, Ben Glocker, and Andrew Davison. 2015. ElasticFusion: Dense SLAM Without A Pose Graph. In Robotics: Science and Systems XI. Robotics: Science and Systems Foundation. doi:10.15607/ RSS.2015.XI.001
- Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 2024. 4d gaussian splatting for real-time dynamic scene rendering. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 20310–20320.
- Yabin Xu, Liangliang Nan, Laishui Zhou, Jun Wang, and Charlie C. L. Wang. 2022. HRBF-Fusion: Accurate 3D Reconstruction from RGB-D Data Using on-the-Fly Implicits. ACM Transactions on Graphics 41, 3 (June 2022), 1–19. arXiv:2202.01829 [cs] doi:10. 1145/3516521
- Chi Yan, Delin Qu, Dan Xu, Bin Zhao, Zhigang Wang, Dong Wang, and Xuelong Li. 2024. GS-SLAM: Dense Visual SLAM with 3D Gaussian Splatting. arXiv:2311.11700 [cs]
- Xingrui Yang, Hai Li, Hongjia Zhai, Yuhang Ming, Yuqian Liu, and Guofeng Zhang. 2022. Vox-Fusion: Dense Tracking and Mapping with Voxel-based Neural Implicit Representation. In 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR). 499–507. arXiv:2210.15858 [cs] doi:10.1109/ISMAR55827.2022.00066
- Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. 2023. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12–22.
- Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. 2024. Mipsplatting: Alias-free 3d gaussian splatting. In Proceedings of the IEEE/CVF Conference

- on Computer Vision and Pattern Recognition. 19447-19456.
- Hongjia Zhai, Gan Huang, Qirui Hu, Guanglin Li, Hujun Bao, and Guofeng Zhang. 2024. Nis-slam: Neural implicit semantic rgb-d slam for 3d consistent scene understanding. IEEE Transactions on Visualization and Computer Graphics (2024).
- Hongjia Zhai, Hai Li, Zhenzhe Li, Xiaokun Pan, Yijia He, and Guofeng Zhang. 2025a. Panogs: Gaussian-based panoptic segmentation for 3d open vocabulary scene understanding. In Proceedings of the Computer Vision and Pattern Recognition Conference. 14114–14124.
- Hongjia Zhai, Boming Zhao, Hai Li, Xiaokun Pan, Yijia He, Zhaopeng Cui, Hujun Bao, and Guofeng Zhang. 2025b. NeuraLoc: Visual Localization in Neural Implicit Map with Dual Complementary Features. (2025).
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In Proceedings of the IEEE conference on computer vision and pattern recognition. 586–595.
- Shuaifeng Zhi, Michael Bloesch, Stefan Leutenegger, and Andrew J. Davison. 2019. SceneCode: Monocular Dense Semantic Reconstruction Using Learned Encoded Scene Representations. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Long Beach, CA, USA, 11768–11777. doi:10.1109/CVPR. 2019.01205
- Qian-Yi Zhou and Vladlen Koltun. 2013. Dense Scene Reconstruction with Points of Interest. ACM Transactions on Graphics 32, 4 (July 2013), 1–8. doi:10.1145/2461912.
- Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R. Oswald, and Marc Pollefeys. 2022. NICE-SLAM: Neural Implicit Scalable Encoding for SLAM. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, New Orleans, LA, USA, 12776–12786. doi:10.1109/ CVPR52688.2022.01245