

Real-Time LiDAR Gaussian Splatting SLAM via Geometry-Aware Covariance Coupling

SeungJun Tak^{*}, Yewon Jeon^{*}, Jaeik Hwang, Suk Min Hwang,
Seongbo Ha, and Hyeonwoo Yu

Sungkyunkwan University, Suwon, South Korea

{tmdwns8840,yw8431,jaeik0211,raph03,sobo3607,hwyu}@skku.edu

Abstract. We present a real-time LiDAR-based framework for Gaussian Splatting SLAM that tightly couples fast G-ICP registration with spherical rasterization-based dense mapping for large-scale sequences. Leveraging LiDAR geometry rather than appearance, we reuse tracking-estimated local covariances to initialize Gaussians with range-aware scales and to derive surface normals for geometry-aware map optimization. We further introduce a covariance-derived geometry score that measures local complexity and drives pruning in planar regions and selective densification in structurally rich areas, while optimized Gaussians and LiDAR-specific confidence cues are fed back to improve tracking robustness. On the Newer College dataset, our method achieves an F-score of 86.78% using purely online trajectories at real-time speed (>20 FPS), and additional experiments on other datasets confirm its stability and scalability. Code is available at github.com/Lab-of-AI-and-Robotics/LiDAR-GS-SLAM. Project Page: lab-of-ai-and-robotics.github.io/GS-SLAM-Family

Keywords: LiDAR SLAM · Gaussian Splatting · Dense Mapping

1 Introduction

LiDAR-based SLAM [8, 27, 35, 39] provides robust trajectory tracking in large-scale environments, but most systems rely on sparse point maps or voxel structures for real-time operation. While effective for localization, these representations are limited for high-fidelity reconstruction and visibility-aware reasoning, which are increasingly demanded in autonomous-driving simulation and HD map production. Recently, 3D Gaussian Splatting (3DGS) [16] has emerged as an efficient dense representation, yet integrating it into online large-scale SLAM remains challenging due to (i) *unbounded map growth* that inflates GPU memory and rendering/registration cost, and (ii) *error persistence*, where corrupted primitives bias subsequent registrations over long horizons.

^{*} These authors contributed equally to this work.




Fig. 1: Performance Overview of LiDAR GS. We compare dense mapping SLAM with respect to Map Size (MB) and Frame per Second (FPS), with F-score (%). The two panels illustrate the trade-off between efficiency and accuracy.

We address these issues with a LiDAR-based Gaussian Splatting SLAM system that enables scalable and reliable dense mapping in large-scale outdoor environments. We perform fast tracking using Generalized ICP [25] against a trackable subset of the map. For mapping, we adopt rasterization-based optimization but represent primitives as surface-oriented 2D anisotropic Gaussians [13] in spherical coordinates, which better matches LiDAR surface observations. Crucially, we tightly couple tracking and mapping by reusing geometric information across stages with minimal overhead: tracking covariances provide geometric cues for Gaussian initialization and online map control, while mapping-refined geometry reweights tracking residuals for robust registration under limited observations.

This coupling is realized through two mechanisms. First, we derive normals and anisotropic scales directly from local G-ICP covariances and store them as primitive parameters, stabilizing and accelerating map optimization without additional learning or separate geometry estimation. Second, we enforce scalability via geometry-aware online map management, performing targeted pruning in planar/redundant regions and selective densification in structurally informative areas. We further infer a reliability mask from ray-drop patterns and observation-induced normal inconsistencies, and use it consistently in both mapping (to suppress unreliable primitives) and tracking (for correspondence filtering/weighting) to improve long-term robustness.

We evaluate our method on large-scale LiDAR benchmarks including KITTI Odometry [6], Oxford Spires [29], and Newer College [24], comparing against classical and neural field-based LiDAR SLAM baselines. Results on trajectory accuracy (ATE RMSE), runtime (FPS), and map growth (number of primitives) demonstrate that our system achieves competitive accuracy while substantially improving scalability and reliability. Our main contributions are as follows:

- **System:** We propose a real-time scalable LiDAR Gaussian Splatting SLAM system that builds dense, continuous maps in large-scale outdoor scenes.
- **Mutual reuse of covariance-derived geometric cues:** Tracking covariances initialize Gaussian geometry for mapping, while mapping-refined Gaussians provide covariance-aware targets for robust registration.
- **Reliability & map management:** We present geometry-aware pruning/densification and an observation-driven reliability mask, jointly used in tracking and mapping for robust long-horizon operation.

2 Related Work

LiDAR Odometry and SLAM. To achieve real-time processing, traditional large-scale LiDAR SLAM systems abstract scenes into sparse point sets or low-resolution voxels. The field has evolved from early curvature-based feature extraction [26, 39] to tightly-coupled optimization frameworks [11, 27, 34] and continuous-time trajectory modeling [2, 4, 8] for enhanced robustness against motion and noise. However, since these methods primarily target odometry and localization, they maintain sparse map structures tailored for registration, neglecting explicit continuous surface representations or visibility-aware operations.

Dense Mapping and Implicit Representation. To overcome sparse maps, dense mapping has modeled surface continuity, from volumetric TSDF [20, 31] and surfel maps [1, 3] to neural/implicit representations [5, 23, 30, 37, 40] that enable high-quality reconstruction. However, in online large-scale SLAM, frequent updates and global re-optimization incur high cost and hinder real time. Over long trajectories, accumulated noise and drift also make map consistency difficult, degrading tracking stability. Thus, balancing geometric fidelity and online scalability remains open.

Gaussian Splatting-based SLAM. Recently, Gaussian Splatting [13, 16] has emerged as an effective representation that enables real-time rendering with explicit spatial extents. Early vision-based GS-SLAM studies [14, 15, 19, 36] optimized Gaussians mainly with dense photometric cues, while subsequent works coupled G-ICP with Gaussian mapping for RGB-D [10, 21], introduced bidirectional monocular VO-Gaussian coupling [38], and explored large-scale outdoor operation with LiDAR-camera-IMU fusion [17, 18, 32, 33]. LiDAR-only dense Gaussian mapping remains more challenging due to sparse range observations and the lack of photometric cues. Splat-LOAM [7] uses spherical LiDAR projection for geometry-only reconstruction, while NeRF-LOAM [5] is LiDAR-only but much slower in our online evaluation. In contrast to RGB-D or LiDAR-visual systems such as GS-ICP, G2S-ICP [10, 21], PINGS [22], and GSO-SLAM [38], our protocol is LiDAR-only, geometry-centered, and online.

3 Method

In this paper, we propose a real-time LiDAR Gaussian Splatting SLAM system for large-scale outdoor environments. We normalize LiDAR scans to the

spherical domain and represent the map as 2D Gaussian primitives rasterized in the same domain. While each Gaussian stores rich parameters, memory usage and optimization cost can grow rapidly as the map scales. To address this, we reuse the registration residuals and local geometric cues already computed in G-ICP [25] tracking. Without an extra geometric estimation module, these signals initialize new primitives, set their geometric parameters, and guide densification and pruning. Conversely, geometric parameters and uncertainty refined during mapping are fed back to tracking to form a trackable subset and perform covariance-weighted G-ICP, reducing correspondence search cost and mitigating error accumulation over long trajectories. The overall pipeline is summarized in Fig. 2, after which we present the spherical-domain 2D Gaussian map, trackable Gaussian selection with geometry-weighted G-ICP tracking, tracking-informed priors for mapping, 2DGS optimization, and geometry-aware map management.

3.1 Map Representation

We represent the map as 2D Gaussian primitives $\mathcal{G} = \{(\boldsymbol{\mu}_i, \mathbf{q}_i, \tilde{\mathbf{s}}_i, \alpha_i)\}_{i=1}^N$, where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ is the mean, $\mathbf{q}_i$ is the orientation, $\tilde{\mathbf{s}}_i = [s_{i,0}, s_{i,1}]$ is the in-plane scale, and α_i is the opacity. Each primitive additionally stores per-Gaussian signals (u_i, c_i) for LiDAR confidence and map complexity control. For mapping, we build each scan as a spherical range image and render the Gaussian map under the same spherical projection $\pi_s(\cdot)$, producing $(\hat{D}, \hat{A}, \hat{\mathbf{N}})$ on a shared grid for stable geometric supervision.

3.2 G-ICP Tracking with Trackable Gaussians

Given an incoming scan at time t , we form the source set $\mathcal{P}_t^s = \{\mathbf{p}_i^s\}$ and apply voxel-grid downsampling (voxel size δ) to obtain $\mathcal{P}_t'^s$. To reduce correspondence cost and avoid stale/unstable map elements, we register only against a trackable subset:

$$\mathcal{G}_{\text{trk}} = \{g_i \in \mathcal{G} \mid m_i = 1, k_i \geq k_{\min}, \alpha_i \geq \alpha_{\min}\}, \quad (1)$$

where m_i is a trackable mask and k_i denotes the last keyframe index that updated g_i . We convert $\mathcal{G}_{\text{trk}}$ into a target point set $\mathcal{P}_t'^t$ by taking Gaussian means (and optional voxel downsampling for efficiency). Using $\mathcal{P}_t'^s$ and $\mathcal{P}_t'^t$, we estimate $\mathbf{T}_t \in SE(3)$ via G-ICP.

Covariance construction and reuse. During tracking, G-ICP estimates a local covariance for each downsampled source point $\mathbf{p}_i'^s \in \mathcal{P}_t'^s$ from its k -NN neighborhood $\mathcal{N}(\mathbf{p}_i'^s)$. We cache its eigendecomposition (principal axes and standard deviations), $(\mathbf{R}_i^s, \boldsymbol{\sigma}_i)$, which is already computed inside the tracker. When a keyframe is inserted, these source covariances are reused as low-cost priors to initialize the rotation, anisotropic in-plane scale, and confidence of newly created map Gaussians.

Conversely, for tracking we construct target covariances directly from each map Gaussian g_i : its orientation $\mathbf{q}_i$ and in-plane scales $(s_{i,x}, s_{i,y})$ define a surfel-like 3×3 covariance with a small thickness ϵ along the normal direction. This




Fig. 2: Overview of the Proposed System. We downsample each LiDAR scan and estimate per-point covariances to form the source set $\mathcal{P}_t^s$. Tracking registers $\mathcal{P}_t^s$ to a trackable target set from the map via G-ICP to estimate T_t , and the covariances produce a control score c_i for pruning/densification. Keyframes are fused into the 2D Gaussian map and optimized, while reusing stored target parameters avoids per-frame covariance re-estimation and improves robustness to geometric noise.

enables $O(1)$ target covariance construction without any map-side k -NN search. Using these target covariances, the tracker performs covariance-weighted G-ICP between $\mathcal{P}_t^s$ and the trackable target set $\mathcal{P}^t$ to estimate $T_t \in SE(3)$. As mapping optimizes $(\mathbf{q}_i, s_{i,x}, s_{i,y})$, the target covariances become more accurate, which in turn stabilizes subsequent registrations. For the very first frame (before a map exists), we fall back to the same neighborhood-based covariance estimation used for sources.

Geometry-aware Matching Weight. We leverage planar priors to improve LiDAR tracking [12, 41]. Given the in-plane Gaussian scales $s_i = [s_{i,0}, s_{i,1}]$ ($s_{i,0} > s_{i,1}$), we define an curvature

$$\kappa_i = \frac{s_{i,0}}{s_{i,1} + \epsilon}, \quad (2)$$

and normalize it to $\bar{\kappa}_i \in [0, 1]$:

$$\bar{\kappa}_i = \frac{\kappa_i - \kappa_{\min}}{\kappa_{\max} - \kappa_{\min}}. \quad (3)$$

Lower $\bar{\kappa}_i$ indicates a more planar primitive. We then solve a weighted G-ICP objective:

$$\mathbf{T}_t = \arg \min_{\mathbf{T}} \sum_i w(\bar{\kappa}_i) \mathbf{d}_i^\top (\mathbf{C}_i^t + \mathbf{T} \mathbf{C}_i^s \mathbf{T}^\top)^{-1} \mathbf{d}_i, \quad (4)$$

with

$$w(\bar{\kappa}_i) = \min \left(\frac{1}{\bar{\kappa}_i + \epsilon}, \omega_{\max} \right), \quad \omega_{\max} = 10. \quad (5)$$

This upweights stable planar structures, acting as geometric anchors and reducing orientation drift.

Loop Closure. In addition to local tracking, we leverage loop-closure to reduce long-term drift. Upon detecting a loop [9], we compute the relative pose between the source keyframe and the matched target keyframe by aligning the source keyframe with its corresponding submap. This relative pose is then introduced as a new constraint in the pose graph. Following pose-graph optimization, each Gaussian is updated by the optimized-pose delta of its source keyframe, yielding a keyframe-wise piecewise-rigid map correction rather than full global Gaussian re-optimization.

3.3 Tracking-Derived Initialization and Geometric Parameters

Mapping optimization is highly sensitive to initialization, strongly impacting stability and convergence. In large-scale settings, per-point neighbor searches or extra geometric estimation are prohibitively expensive. We therefore reuse local geometric statistics already produced by G-ICP tracking to initialize new Gaussians and derive low-cost map-management signals, achieving stable convergence with real-time efficiency. Because LiDAR provides limited appearance cues, these geometric statistics are crucial for preserving map quality and stable tracking. This section details covariance-based Gaussian initialization and the two management signals: control score and trackability confidence.

Range-adaptive Scale Initialization. G-ICP tracking provides a covariance decomposition summarized by the standard-deviation vector $\sigma_i = [\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}]$, sorted as $\sigma_{i,0} \geq \sigma_{i,1} \geq \sigma_{i,2}$. We use the two largest components as tangent-plane standard deviations $\sigma_i^{\text{tan}} = [\sigma_{i,0}, \sigma_{i,1}]$.

Because LiDAR has fixed angular resolution, surface sampling becomes sparser with range r . However, G-ICP typically normalizes covariances for numerical stability, so σ_i^{tan} mainly reflects local anisotropy rather than absolute spacing. To ensure sufficient initial coverage at long range, we initialize the 2D scale proportional to r : for $r_i = \max(\|\mathbf{p}_i\|, 1)$,

$$\mathbf{s}_i^{(0)} = \text{clip}\left(\kappa r_i \sigma_i^{\text{tan}}, 0, s_{\max}\right), \quad (6)$$

and optimize the log-scale $\tilde{\mathbf{s}}_i = \log(\mathbf{s}_i)$. This range-adaptive initialization reduces far-range holes and stabilizes subsequent map optimization.

LiDAR-specific Opacity and Trackability Confidence. Unlike image-based Gaussian splatting where opacity is optimized to explain photometric rendering, LiDAR returns provide only sparse geometry and their reliability is dominated by sensing geometry. Therefore, we reinterpret opacity as a LiDAR-specific reliability weight and initialize it from the physics-based confidence u , making it well aligned with our sensor for tracking and map management. LiDAR measurements vary in reliability depending on range and incidence angle. Given each point range $r = \|\mathbf{p}\|$, the normal $\mathbf{n}$, and the sensor ray direction $\hat{\mathbf{r}} = \mathbf{p}/\|\mathbf{p}\|$, we compute the incidence cosine term $c = |\mathbf{n}^\top \hat{\mathbf{r}}|$ and define the physics-based confidence u as

$$u = \exp\left(-\left(\frac{r}{r_0}\right)^2\right) \cdot \text{clip}\left(\frac{c-c_0}{1-c_0+\epsilon}, 0, 1\right). \quad (7)$$

When adding a new Gaussian in mapping, points with larger u_i are initialized with higher opacity:

$$\alpha_i^{(0)} = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min}) u_i. \quad (8)$$

We further define the trackability confidence as $\gamma_i = \alpha_i u_i$, which combines learned opacity and physics-based reliability.

Control Score. In large-scale sequences, Gaussians accumulate over time, causing memory and rendering costs to grow roughly linearly. To keep the map scalable, we reallocate resources by compressing planar regions and selectively densifying complex structures. To this end, we define a control score c_i from covariance-based features computed during tracking. Using the covariance eigenvalues $\lambda_{i,j} = \sigma_{i,j}^2$, we compute linearity and curvature as

$$\text{linear}_i = \frac{\lambda_{i,0} - \lambda_{i,1}}{\lambda_{i,0} + \epsilon}, \quad \text{curv}_i = \frac{\lambda_{i,2}}{\lambda_{i,0} + \lambda_{i,1} + \lambda_{i,2} + \epsilon}.$$

We also include the normalized G-ICP registration residual $\widetilde{\text{res}}_i$ and define the control score as a linear combination of normalized terms:

$$c_i = \text{clip}\left(w_l \widetilde{\text{linear}}_i + w_c \widetilde{\text{curv}}_i + w_r \widetilde{\text{res}}_i, 0, 1\right). \quad (9)$$

In our implementation, we use $w_l = 0.55$, $w_c = 0.30$, $w_r = 0.15$. A lower c_i indicates more planar structure, while a higher c_i indicates more geometrically complex structures (e.g., edges/corners). This score is used for (i) planar compression (cover-and-prune) and (ii) selective splitting of complex regions (Section 3.5).

3.4 Map Optimization

Mapping rasterizes the current Gaussian map into the range-image domain to obtain predicted depth, opacity, and normals $(\hat{D}, \hat{A}, \hat{\mathbf{N}})$, and optimizes the Gaussian parameters $\Theta = \{\boldsymbol{\mu}, \mathbf{q}, \tilde{\mathbf{s}}, \alpha\}$ by comparing them against the observations. To avoid invalid measurements, we optimize only over valid pixels using depth-valid mask $\mathbf{M}_D = \mathbb{I}[D > 0]$.

Depth & Opacity Loss. Depth and opacity losses are defined on valid LiDAR pixels using the valid-depth mask $\mathbf{M}_D$. Specifically, we use a masked L1 loss for depth and a positive-only BCE (i.e., $\text{BCE}(\hat{A}, 1)$) for opacity:

$$\mathcal{L}_{\text{depth}} = \frac{1}{|\Omega_D|} \sum_{i \in \Omega} \mathbf{M}_{D,i} |\hat{D}_i - D_i|, \quad \mathcal{L}_\alpha = -\frac{1}{|\Omega_D|} \sum_{i \in \Omega} \mathbf{M}_{D,i} \log(\hat{A}_i + \epsilon), \quad (10)$$

where Ω denotes the set of all pixels and $\Omega_D = \{i \in \Omega \mid \mathbf{M}_{D,i} = 1\}$.

Normal Loss with Valid Masks. Normal alignment complements the weak geometric constraints of depth-only supervision, stabilizing surface orientation and boundary learning. We employ two types of normals: (i) geometric normals $\mathbf{N}_{\text{g-icp}}$ obtained from tracking (G-ICP), and (ii) surface normals $\mathbf{N}_{\text{surf}}$ computed from the rendered depth. $\mathbf{N}_{\text{g-icp}}$ is strong for local planar estimation but may




Fig. 3: Covariance-derived control score and decision map. From each tracking covariance, we compute linearity/curvature proxies from eigenvalues and fuse them with a robustly normalized G-ICP residual to obtain a control score $c_i \in [0, 1]$. Low c_i indicates planar/redundant regions and triggers cover-and-prune, while high c_i indicates geometric complexity (edges/corners) and triggers selective splitting. The visualization shows how decisions spatially correlate with scene structure.

be unstable in some regions, while $\mathbf{N}_{\text{surf}}$ provides observation-driven boundary/discontinuity cues; thus they are complementary. To avoid early-stage divergence, we activate the $\mathbf{N}_{\text{surf}}$ term only after densification via a scheduling strategy. We define the valid mask for G-ICP normals as

$$\mathbf{M}_N = \mathbf{M}_D \odot \mathbb{I}[\|\mathbf{N}_{\text{g-icp}}\| > 0],$$

and use the cosine loss $\mathcal{L}_{\cos}(\mathbf{a}, \mathbf{b}) = 1 - \mathbf{a}^\top \mathbf{b}$:

$$\mathcal{L}_n = \lambda_{n,g} \frac{1}{|\Omega_N|} \sum_{i \in \Omega_N} \mathcal{L}_{\cos}(\hat{\mathbf{N}}_i, \mathbf{N}_{\text{g-icp},i}) + \lambda_{n,s} \frac{1}{|\Omega_D|} \sum_{i \in \Omega_D} \mathcal{L}_{\cos}(\hat{\mathbf{N}}_i, \mathbf{N}_{\text{surf},i}), \quad (11)$$

where $\Omega_N = \{i \in \Omega \mid \mathbf{M}_{N,i} = 1\}$. For training stability, $\lambda_{n,s}$ is enabled only after densification begins.

Scale Regularization Loss. To prevent abnormally large Gaussians from causing geometric distortions, we penalize only when the maximum component of the 2D scale exceeds a threshold $s_{\max}$:

$$\mathcal{L}_{\text{scale}} = \sum_{i=1}^N \max(0, \max(s_{i,x}, s_{i,y}) - s_{\max}). \quad (12)$$

The final mapping optimization minimizes the following combined objective:

$$\mathcal{L} = \mathcal{L}_{\text{depth}} + \lambda_\alpha \mathcal{L}_\alpha + \lambda_n \mathcal{L}_n + \lambda_s \mathcal{L}_{\text{scale}}. \quad (13)$$

3.5 Real-time Map Budget Control for Long Sequences

Large-scale sequences continuously accumulate keyframes, leading to steady Gaussian growth and increasing memory, rendering, and optimization costs that

threaten real-time operation. Beyond opacity-based pruning/densification, we therefore introduce a tracking-derived control score c_i , enabling an efficient map representation under long sequences.

Multi-pass Planar Cover and Prune. To efficiently reduce redundant Gaussians on large planar structures, we perform a **multi-pass** grouping strategy based on projections onto orthogonal planes, instead of relying on a single 3D voxelization. In each pass, we project each Gaussian mean μ_i onto the corresponding plane and form voxel-grid groups G in projected 2D domain. For each group G , we consider only Gaussians that satisfy (i) a low control score $c_i \leq Q_{q_{plane}}(c)$ and (ii) sufficient opacity $\alpha_{min} \leq \alpha_i$. We then select the Gaussian with the largest α as the representative g_r , and expand its in-plane-scale in the local tangent-plane coordinates so that the cover region includes a q_{cover} . Among the members falling inside the cover region, planar candidates are treated as redundant and removed. To prevent excessive information loss, we cap the maximum pruning ratio per pass.

Structure-aware Adaptive Splitting. A high control score c_i indicates structurally complex regions such as edges and corners. To selectively enhance representational capacity in these areas, we perform adaptive splitting. At each map-management step, we select split candidates whose control scores are above the q_{split} , while also requiring $\alpha_{min} \leq \alpha_i$ to avoid over-splitting unstable primitives. We further limit resource growth by capping the splitting ratio relative to the total number of Gaussians N . For each split candidate g_i , we define the in-plane major axis as $a_i = R(q_i)e_0$ and generate two child Gaussians by shifting the mean in opposite directions along a_i . This symmetry-breaking initialization causes the two children to receive different gradients during subsequent optimization, allowing them to specialize to nearby geometric variations and represent fine structures. Fig. 3 visually demonstrates how the decisions derived from these techniques align with the actual scene structure.

4 Experiments

4.1 Experimental Setup

Datasets. We evaluate the proposed method on the KITTI Odometry dataset [6], the Oxford Spires dataset [29], and the Newer College dataset [24]. KITTI contains large-scale long-horizon sequences in autonomous driving scenarios, and is thus used to assess tracking stability and scalability in large-scale environments. In contrast, Newer College and Oxford Spires are captured with a handheld sensor at relatively low speed, enabling dense observations; moreover, they provide ground-truth meshes, which allow us to evaluate both map quality and local registration stability. For every dataset, we deskew each scan by interpolating the estimated ego-motion over time.

Implementation Details. All experiments are conducted on a desktop computer equipped with an AMD Ryzen 9 7900X CPU, 64 GB of RAM, and an NVIDIA GeForce RTX 4090 GPU. For the mapping loss (Eqn. 13), the weights

Table 1: Trajectory evaluation results (ATE RMSE [m] ↓) on Newer College, Oxford Spires, and KITTI datasets. Tracking accuracy is measured using the ATE RMSE in meters. The best, second-best, and third-best results are highlighted in **first**, second, and third colors, respectively. ‘Fail’ indicates that the method lost tracking or failed to converge on the corresponding sequence.

Method	Newer		Oxford		KITTI				
	quad	math	obs	keb	00	01	07	08	09
KISS-SLAM [8]	0.100	0.151	0.494	0.367	1.699	19.632	0.435	3.982	1.687
PIN-SLAM [23]	0.097	0.071	0.150	12.196	0.913	3.670	0.281	1.778	1.246
SuMa [1]	0.249	0.210	5.145	16.803	1.200	Fail	0.405	2.210	3.884
Splat-LOAM [7]	0.138	0.818	17.494	0.770	Fail	Fail	Fail	Fail	Fail
Ours	0.080	0.160	0.336	0.177	1.080	2.260	0.338	2.900	1.260

are set to $\lambda_\alpha = 0.1$, $\lambda_n = 1$, and $\lambda_s = 10$. For Eqn. 11, we set $\lambda_{n,g} = 0.05$ and $\lambda_{n,s} = 0.01$.

Baselines. We compare the proposed method with various map representation approaches, including point [8], surfel [1], TSDF [20, 31], neural field [23, 28], and Gaussian Splatting [7]. We evaluate our approach against mapping-only methods [20, 28, 31]. Since these methods operate using ground truth poses, they are free from the mapping degradation typically caused by pose estimation errors. Consequently, they represent an ideal scenario and serve as an upper bound for the achievable map quality.

Metrics. To evaluate tracking accuracy, we employ the Root Mean Square Error (RMSE) of the Absolute Trajectory Error (ATE). For the mapping evaluation, we assess mesh-to-mesh quality using Accuracy (Acc), Completeness (Com), Chamfer-L1 (C-L1), and the τ -F-score (%), which are based on bidirectional nearest-neighbor distances. All distance metrics are reported in centimeters (cm). In our implementation, we set $\tau = 0.2$ m, the downsampling resolution to 0.02 m, and truncate distance outliers at 0.5 m.

Table 2: Mapping quality and system speed on Newer College and Oxford Spires datasets. ‘Est.’ indicates the use of online estimated poses. Metric units are cm for Acc, Com, C-L1, and % for F-score. FPS denotes system FPS (Tracking + Mapping) or Mapping only.

Method	Pose	Newer College										Oxford Spires									
		Quad					Math Institute					Observatory					Keeble College				
		Acc↓	Com↓	C-L1↓	F↑	FPS↑	Acc↓	Com↓	C-L1↓	F↑	FPS↑	Acc↓	Com↓	C-L1↓	F↑	FPS↑	Acc↓	Com↓	C-L1↓	F↑	FPS↑
Voxblox [20]	GT	16.58	19.00	17.79	64.00	66.10	12.98	15.02	14.00	77.23	82.70	14.26	17.47	15.86	71.73	143.00	12.69	14.97	13.83	77.29	140.30
VDBFusion [31]	GT	14.88	14.09	14.49	76.27	26.44	9.04	12.19	10.61	84.59	45.40	9.88	11.97	10.93	83.68	48.16	9.11	11.29	10.20	84.99	43.69
N ³ -Mapping [28]	GT	8.73	12.33	10.48	85.66	0.68	8.71	12.42	10.57	84.83	0.68	6.02	17.73	11.87	80.55	0.75	4.82	18.25	11.53	80.87	0.79
PIN-SLAM [23]	Est.	17.11	22.83	19.97	60.59	7.62	18.06	18.29	18.18	63.78	11.04	22.44	26.40	24.42	44.15	8.20	24.16	34.22	29.19	33.93	8.62
Splat-LOAM [7]	Est.	12.23	18.91	15.57	72.06	3.91	11.70	18.33	15.02	73.66	2.54	24.79	47.88	36.34	6.14	1.52	21.48	35.05	28.26	33.97	5.77
Ours	Est.	7.62	15.02	11.32	86.78	22.15	14.43	18.54	16.48	70.35	25.25	14.48	27.18	20.83	56.89	17.99	12.71	24.57	18.64	64.02	19.58

Table 3: Ablation Results on Effect of Tracking-to-Mapping Information Transfer. We evaluate the impact of utilizing tracking-derived priors for the mapping process on the Quad and Keble College sequences. Tracking accuracy is measured with ATE RMSE [m], while map quality is assessed using distance metrics (Acc, Com, C-L1) in cm and F-score in %.

Variant	Quad					Keble College				
	ATE RMSE	Map Quality				ATE RMSE	Map Quality			
		↓	Acc↓	Com↓	C-L1↓		F-score↑	↓	Acc↓	Com↓
A. w/o Init (all identity)	0.088	8.495	35.885	22.190	48.688	0.283	15.202	32.837	24.020	44.123
B. w/o Init Range adap Scale	0.087	8.016	15.254	11.635	85.996	0.479	16.451	28.666	22.558	51.582
C. w/o Init Opacity	0.085	7.809	15.091	11.450	86.020	0.451	15.735	28.685	22.210	53.173
D. w/o G-ICP Normal	0.099	8.040	15.160	11.600	85.899	0.325	15.678	27.179	21.428	55.522
E. w/o Feat-aware	0.089	8.245	15.212	11.729	86.237	0.222	12.942	25.816	19.379	63.217
F. Full	0.079	7.624	15.020	11.320	86.776	0.212	12.713	24.568	18.640	64.024

4.2 Tracking

We evaluate tracking performance on the KITTI Odometry [6], Oxford Spires [29], and Newer College [24] datasets, as reported in Tab. 1. Our method achieves competitive ATE for real-time dense Gaussian mapping, rather than uniformly best odometry. For Splat-LOAM, we used the official implementation with the provided `kitti.yaml` and no modification; a run is marked as failed when tracking diverges or ATE RMSE exceeds 50 m. Since we share the same loop-closure and pose-graph backend as KISS-SLAM, the differences mainly reflect the front-end scan-to-map tracking and map representation.

4.3 System Speed and Quality of Reconstructed Map

To quantitatively evaluate the mapping quality, we convert the reconstructed GS maps into meshes and measure the Accuracy, Completeness, and F-score against the ground truth (GT) meshes. Tab. 2 and Fig. 4 present qualitative and quantitative results of the reconstructed scenes respectively. Among simultaneous tracking and mapping methods, our approach offers the best balance of quality, speed, and memory. Averaged over four sequences, it achieves an F-score of 69.51 at 21.24 FPS. This outperforms PIN-SLAM (50.61, 8.87 FPS) and Splat-LOAM (46.46, 3.44 FPS) by +18.90 and +23.05 F-score, with $2.39\times$ and $6.18\times$ speedups, respectively. Our map is also compact (7.23 MB on average), 61.6% smaller than PIN-SLAM (18.83 MB), while delivering higher reconstruction quality. These results indicate that our map management is not mere compression: it preserves structurally important primitives while removing redundancies, enabling high-fidelity, real-time mapping on large-scale sequences.

4.4 Ablation Study

Effect of Tracking-to-Mapping Information Transfer. Tab. 3 analyzes tracking-derived signals used in the mapping module. In Tab. 3.A, removing tracking covariances causes the largest degradation: without covariance-based




Fig. 4: Qualitative mesh reconstruction comparison. The figure shows reconstruction results for the quad-easy sequence using estimated poses from the Newer College dataset. Colors represent surface normals mapped to the RGB color space. Our method recovers geometry with high fidelity. PIN-SLAM shows high levels of noise, and Splat-LOAM shows areas of under-reconstruction and over-smoothing.

orientation and scale priors, Gaussians initialized with identity rotations and unaligned scales lead to poor surface coverage and misoriented accumulation, destabilizing optimization under a fixed budget. Consequently, the F-score drops sharply on Quad, with clear deterioration in Com and C-L1. Range-adaptive scale initialization is crucial for scenes with many distant returns; disabling it (Tab. 3.B) reduces far-range coverage and degrades both tracking ATE and mapping quality (e.g., Keble College). Removing opacity initialization (Tab. 3.C) consistently lowers mapping quality, while its effect on tracking ATE varies by dataset, since ATE is also affected by correspondence distribution, visibility, and dynamic-object suppression. Without the G-ICP normal loss (Tab. 3.D), mapping quality changes little but tracking ATE worsens, indicating that normal supervision rapidly stabilizes tangent-plane orientations during mapping and yields more reliable target covariances/orientations for scan-to-map registration. Finally, disabling feature-aware map management (Tab. 3.E) has minimal short-term impact on mapping quality; because the control score mainly reallocates representation budget over time, we examine the long-horizon quality-resource trade-off in Tab. 4. Overall, using all components provides the best mapping quality and tracking accuracy.

Effect of Mapping-to-Tracking Information Transfer. This ablation study evaluates the effect of feeding back the optimized GS map from the mapping module to the tracking process. In Tab. 5, the w/o GS Opt. configuration disables GS map optimization (resulting in a frozen map) to compare the tracking accuracy and processing speed under an identical tracking pipeline. Consequently, the Full system achieves a lower ATE RMSE and a higher tracking FPS compared to the frozen map baseline. This improvement indicates that the updated

Table 4: Ablation Results on Map Management. We evaluate the impact of map management methods on map fidelity (Acc, Com, C-L1 in cm; F-score in %) and resource consumption (#GS, Map size, and Peak VRAM).

Variant	Quad							Kebble College						
	Map Quality				#GS	Map	VRAM	Map Quality				#GS	Map	VRAM
	Acc ↓	Com ↓	C-L1 ↓	F-score ↑	↓	↓	↓	Acc ↓	Com ↓	C-L1 ↓	F-score ↑	↓	↓	↓
Baseline (opacity prune)	7.691	14.964	11.327	87.179	239,175	9.6MB	2.9 GB	13.219	24.664	18.942	64.863	546,400	21.9MB	4.1GB
Split	7.597	14.504	11.051	88.040	248,060	9.9MB	4.8 GB	15.718	27.220	21.469	56.681	591,567	23.7MB	4.2GB
Plane Prune	7.708	15.100	11.404	86.521	119,285	4.8MB	4.5 GB	14.694	26.163	20.429	59.327	341,435	13.7MB	4.1GB
Split + Plane Prune	7.624	15.020	11.320	86.776	129,885	5.0MB	4.5 GB	12.713	24.568	18.640	64.024	344,473	13.8MB	4.3GB

Table 5: Ablation Results on Effect of Mapping-to-Tracking Information Transfer. ATE RMSE is reported in meters. The *Full (w/ feedback)* system achieves not only the lowest tracking drift but also the highest tracking FPS, demonstrating that optimized Gaussians serve as superior registration targets.

Variant	ATE RMSE ↓	Tracking FPS ↑
w/o GS Opt. (frozen / no feedback)	0.099	17.69
w/o geometric weight	0.088	22.66
w/o Conf/Filtering	0.085	20.47
Full (w/ feedback)	0.079	23.70

Table 6: Ablation Results on Mapping Losses. We evaluate the impact of individual loss terms on tracking accuracy and mapping quality using the Quad and Kebble College sequences. ATE RMSE is reported in meters, distance metrics (Acc, Com, C-L1) in cm, and F-score in %.

Variant	Quad					Kebble College				
	ATE RMSE	Map Quality				ATE RMSE	Map Quality			
		Acc ↓	Com ↓	C-L1 ↓	F-score ↑		Acc ↓	Com ↓	C-L1 ↓	F-score ↑
w/o G-ICP Normal	0.097	8.119	15.247	11.683	85.844	0.346	16.274	27.222	22.248	54.206
w/o Scale	0.086	7.788	15.859	11.823	85.732	0.262	21.017	35.305	28.161	33.924
w/o Opacity	0.084	8.681	16.230	12.456	83.217	0.328	16.283	27.678	21.980	52.893
Full	0.079	7.624	15.020	11.320	86.776	0.212	12.713	24.568	18.640	64.024

rotation, scale, and opacity of the target Gaussians from mapping optimization provide more stable covariance-based constraints during scan-to-map registration. As a result, the number of registration iterations and the ratio of failed correspondences decrease, yielding a distinct advantage in effective processing speed. Ultimately, this demonstrates that the coupling in our system is not a disruptive correction mechanism where "mapping interferes with tracking," but rather a purposefully designed feedback structure where the optimized GS map inherently serves as a superior registration target.

Mapping Loss. Tab. 6 analyzes the impact of each term in the mapping objective on map quality and tracking performance. The Full configuration utilizes all loss terms. *w/o G-ICP Normal* removes the normal supervision transferred from the tracking process, leading to a distinct degradation in ATE across both sequences. We attribute this to the fact that the depth consistency loss alone provides weak local surface orientation constraints. Consequently, the tangent plane




Fig. 5: Geometry-aware map budget control driven by covariance-derived control scores. The map remains compact in large planar areas while preserving fine structures.

orientations of target surfels are inadequately reflected, biasing the optimization towards simply generating overly smooth surfaces. This lack of geometric detail ultimately disadvantages scan-to-map registration. *w/o Opacity* ablates the term that induces opacity from valid observation pixels, resulting in the most severe degradation in map quality. Without this term, uncertain observations are not sufficiently suppressed, or conversely, coverage in necessary areas weakens. This leads to an increase in holes and ghosting artifacts, simultaneously deteriorating both precision and completeness. *w/o Scale* disables scale regularization. Under this setting, certain Gaussians can become excessively large, over-covering surrounding surfaces. This phenomenon causes a sharp decline in map quality metrics, particularly in the Keble College sequence. In summary, the proposed loss combination simultaneously enhances depth-based registration (geometry), orientation constraints (normal), and representation stability (opacity and scale), continuously providing a highly favorable registration target for the tracking process.

Gaussian Map Management. Tab. 4 evaluates how our map-management strategy affects the long-horizon trade-off between map quality and resource usage (number of Gaussians). *Baseline* applies only opacity-based pruning. *Plane Prune* performs control-score-based planar compression via a cover-and-prune scheme, while *Split* reinforces complex structures through splitting. *Split + Plane Prune* is our final configuration combining both.

Plane Prune removes redundant Gaussians in large planar regions, substantially reducing the total Gaussian count and map size, at the cost of a slight quality drop due to losing some fine details. Conversely, *Split* increases the number of Gaussians but improves representational capacity in complex areas, often boosting the F-score. Using both mechanisms strikes the best balance. Planar compression keeps map growth in check, while splitting restores fine detail where it matters, leading to the strongest quality-resource trade-off and more reliable registration targets over long sequences. Fig. 5 visually confirms this behavior, showing a sparse distribution of Gaussians in planar regions and a dense concentration within complex geometric structures.

We further analyze scalability on KITTI sequence 00. Since KITTI does not provide ground-truth meshes, we report only resource consumption. With our management strategy, the number of Gaussians decreases from 2,482,505 to 1,510,033, and the persistent map storage drops from 61.3 MB to 42.1 MB, demonstrating a substantial reduction in stored map size in a practical large-scale environment.

5 Conclusion

We presented a real-time LiDAR-based Gaussian Splatting SLAM system for large-scale outdoor environments. The core idea is to reuse local covariance information across tracking and mapping: tracking covariances provide geometric priors for Gaussian initialization and map management, while mapping-refined Gaussian geometry provides covariance-aware targets for scan-to-map registration. Together with the proposed control-score-based pruning and densification strategy, this coupling allocates map capacity to structurally informative regions while reducing redundant Gaussians in simpler areas. As a result, the system achieves real-time performance with compact persistent map storage and competitive peak VRAM in full-system comparisons. These results suggest that covariance-coupled tracking and map management are effective design choices for scalable LiDAR Gaussian Splatting SLAM.

Limitations. Our system relies on local geometric estimates from G-ICP. When the scene has weak geometric constraints, sparse returns, vegetation, or many dynamic objects, the estimated covariances and normals can become less reliable and may affect tracking and mapping. Several parameters, such as voxel resolution, Gaussian scale bounds, and keyframe thresholds, are tied to sensor resolution and scene scale. Loop closure also updates associated Gaussians in a keyframe-wise rigid manner rather than through full global Gaussian re-optimization, which may leave local inconsistencies after large corrections. KITTI is useful for large-scale evaluation but has known ground-truth limitations.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant (no. RS-2024-00359937), MOTIE (RS-2024-00508191 and P008922), KEIT (RS-2025-25449082) and IITP (RS-2025-10692981, RS-2025-25442569 and RS-2025-25463654).

References

1. Behley, J., Stachniss, C.: Efficient surfel-based slam using 3d laser range data in urban environments. In: Robotics: science and systems. vol. 2018, p. 59 (2018)

2. Chen, K., Lopez, B.T., Agha-mohammadi, A.a., Mehta, A.: Direct lidar odometry: Fast localization with dense point clouds. *IEEE Robotics and Automation Letters* **7**(2), 2000–2007 (2022)
3. Chen, X., Milioto, A., Palazzolo, E., Giguere, P., Behley, J., Stachniss, C.: Suma++: Efficient lidar-based semantic slam. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4530–4537. IEEE (2019)
4. Dellenbach, P., Deschaud, J.E., Jacquet, B., Goulette, F.: Ct-icp: Real-time elastic lidar odometry with loop closure. In: 2022 international conference on robotics and automation (ICRA). pp. 5580–5586. IEEE (2022)
5. Deng, J., Wu, Q., Chen, X., Xia, S., Sun, Z., Liu, G., Yu, W., Pei, L.: Nerf-loam: Neural implicit representation for large-scale incremental lidar odometry and mapping. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8218–8227 (2023)
6. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
7. Giacomini, E., Di Giammarino, L., De Rebotti, L., Grisetti, G., Oswald, M.R.: Splat-loam: Gaussian splatting lidar odometry and mapping. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 27630–27639 (2025)
8. Guadagnino, T., Mersch, B., Gupta, S., Vizzo, I., Grisetti, G., Stachniss, C.: KISS-SLAM: A Simple, Robust, and Accurate 3D LiDAR SLAM System With Enhanced Generalization Capabilities. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5363–5370 (2025). <https://doi.org/10.1109/IROS60139.2025.11246613>
9. Gupta, S., Guadagnino, T., Mersch, B., Vizzo, I., Stachniss, C.: Effectively Detecting Loop Closures using Point Cloud Density Maps. In: IEEE International Conference on Robotics and Automation (ICRA) (2024)
10. Ha, S., Yeon, J., Yu, H.: Rgb-d gs-icp slam. In: European conference on computer vision. pp. 180–197. Springer (2024)
11. He, D., Xu, W., Chen, N., Kong, F., Yuan, C., Zhang, F.: Point-lio: robust high-bandwidth light detection and ranging inertial odometry. *Advanced Intelligent Systems* **5**(7), 2200459 (2023)
12. Hsiao, M., Westman, E., Zhang, G., Kaess, M.: Keyframe-based dense planar slam. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 5110–5117 (2017). <https://doi.org/10.1109/ICRA.2017.7989597>
13. 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)
14. Huang, H., Li, L., Cheng, H., Yeung, S.K.: Photo-slam: Real-time 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. pp. 21584–21593 (2024)
15. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat track & map 3d gaussians for dense rgb-d slam. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 21357–21366 (2024)
16. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>

17. Lang, X., Lv, J., Tang, K., Li, L., Huang, J., Liu, L., Liu, Y., Zuo, X.: Gaussian-lic2: Lidar-inertial-camera gaussian splatting slam. *arXiv preprint arXiv:2507.04004* (2025)
18. Liu, J., Wan, Y., Wang, B., Zheng, C., Lin, J., Zhang, F.: Gs-sdf: Lidar-augmented gaussian splatting and neural sdf for geometrically consistent rendering and reconstruction. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 19391–19398. IEEE (2025)
19. Matsuki, H., Murai, R., Kelly, P.H., Davison, A.J.: Gaussian splatting slam. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 18039–18048 (2024)
20. Oleynikova, H., Taylor, Z., Fehr, M., Siegwart, R., Nieto, J.: Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2017)
21. Pak, G., Cho, H.M., Kim, E.: G2s-icp slam: Geometry-aware gaussian splatting icp slam. *arXiv preprint arXiv:2507.18344* (2025)
22. Pan, Y., Zhong, X., Jin, L., Wiesmann, L., Popović, M., Behley, J., Stachniss, C.: Pings: Gaussian splatting meets distance fields within a point-based implicit neural map. *arXiv preprint arXiv:2502.05752* (2025)
23. Pan, Y., Zhong, X., Wiesmann, L., Posewsky, T., Behley, J., Stachniss, C.: Pinslam: Lidar slam using a point-based implicit neural representation for achieving global map consistency. *IEEE Transactions on Robotics* **40**, 4045–4064 (2024)
24. Ramezani, M., Wang, Y., Camurri, M., Wisth, D., Mattamala, M., Fallon, M.: The newer college dataset: Handheld lidar, inertial and vision with ground truth. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4353–4360 (2020). <https://doi.org/10.1109/IROS45743.2020.9340849>
25. Segal, A., Haehnel, D., Thrun, S., et al.: Generalized-icp. In: Robotics: science and systems. vol. 2, p. 435. Seattle, WA (2009)
26. Shan, T., Englot, B.: Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In: 2018 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 4758–4765. IEEE (2018)
27. Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., Rus, D.: Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In: 2020 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 5135–5142. IEEE (2020)
28. Song, S., Zhao, J., Huang, K., Lin, J., Ye, C., Feng, T.: N³-mapping: Normal guided neural non-projective signed distance fields for large-scale 3d mapping. *IEEE Robotics and Automation Letters* **9**(6), 5935–5942 (2024)
29. Tao, Y., Muñoz-Bañón, M.Á., Zhang, L., Wang, J., Fu, L.F.T., Fallon, M.: The oxford spires dataset: Benchmarking large-scale lidar-visual localisation, reconstruction and radiance field methods. *International Journal of Robotics Research* (2025)
30. Turki, H., Ramanan, D., Satyanarayanan, M.: Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12922–12931 (2022)
31. Vizzo, I., Guadagnino, T., Behley, J., Stachniss, C.: Vdbfusion: Flexible and efficient tsdf integration of range sensor data. *Sensors* **22**(3), 1296 (2022)
32. Xiao, R., Liu, W., Chen, Y., Hu, L.: Liv-gs: Lidar-vision integration for 3d gaussian splatting slam in outdoor environments. *IEEE Robotics and Automation Letters* **10**(1), 421–428 (2024)

33. Xie, Y., Huang, Z., Wu, J., Ma, J.: Gs-livm: Real-time photo-realistic lidar-inertial-visual mapping with gaussian splatting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 26869–26878 (2025)
34. Xu, W., Cai, Y., He, D., Lin, J., Zhang, F.: Fast-lio2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics* **38**(4), 2053–2073 (2022)
35. Xu, W., Zhang, F.: Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters* **6**(2), 3317–3324 (2021)
36. Yan, C., Qu, D., Xu, D., Zhao, B., Wang, Z., Wang, D., Li, X.: Gs-slam: Dense visual slam with 3d gaussian splatting. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 19595–19604 (2024)
37. Yang, X., Li, H., Zhai, H., Ming, Y., Liu, Y., Zhang, G.: Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation. In: 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR). pp. 499–507. IEEE (2022)
38. Yeon, J., Ha, S., Yu, H.: Gso-slam: Bidirectionally coupled gaussian splatting and direct visual odometry. *IEEE Robotics and Automation Letters* (2026)
39. Zhang, J., Singh, S., et al.: Loam: Lidar odometry and mapping in real-time. In: *Robotics: Science and systems*. vol. 2, pp. 1–9. Berkeley, CA (2014)
40. Zhong, X., Pan, Y., Behley, J., Stachniss, C.: Shine-mapping: Large-scale 3d mapping using sparse hierarchical implicit neural representations. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 8371–8377. IEEE (2023)
41. Zhou, L., Wang, S., Kaess, M.: π -lsam: Lidar smoothing and mapping with planes. In: 2021 IEEE international conference on robotics and automation (ICRA). pp. 5751–5757. IEEE (2021)