

FoundationSLAM: Unleashing the Power of Depth Foundation Models for End-to-End Dense Visual SLAM

Yuchen Wu¹, Jiahe Li¹, Fabio Tosi², Matteo Poggi², Jin Zheng^{1,3}, Xiao Bai^{1*}

¹School of Computer Science and Engineering, State Key Laboratory of Complex Critical Software Environment, Jiangxi Research Institute, Beihang University

²University of Bologna

³State Key Laboratory of Virtual Reality Technology and Systems, Beijing, China
wuyuchen@buaa.edu.cn, lijiahe@buaa.edu.cn, fabio.tosi5@unibo.it

Abstract

We present FoundationSLAM, a learning-based monocular dense SLAM system that addresses the absence of geometric consistency in previous flow-based approaches for accurate and robust tracking and mapping. Our core idea is to bridge flow estimation with geometric reasoning by leveraging the guidance from foundation depth models. To this end, we first develop a Hybrid Flow Network that produces geometry-aware correspondences, enabling consistent depth and pose inference across diverse keyframes. To enforce global consistency, we propose a Bi-Consistent Bundle Adjustment Layer that jointly optimizes keyframe pose and depth under multi-view constraints. Furthermore, we introduce a Reliability-Aware Refinement mechanism that dynamically adapts the flow update process by distinguishing between reliable and uncertain regions, forming a closed feedback loop between matching and optimization. Extensive experiments demonstrate that FoundationSLAM achieves superior trajectory accuracy and dense reconstruction quality across multiple challenging datasets, while running in real-time at 18 FPS, demonstrating strong generalization to various scenarios and practical applicability of our method.

1 Introduction

Simultaneous Localization and Mapping (SLAM) is a fundamental problem in computer vision and robotics, enabling autonomous agents to perceive and navigate unknown environments. Recent advances in learning-based SLAM systems have demonstrated impressive results by leveraging dense optical flow as a unified representation for tracking and mapping. Among them, DROID-SLAM (Teed and Deng 2021) and a series of its extended variants (Zhang et al. 2023; Zhou et al. 2025; Zhang et al. 2024) have established new performance baselines in multiple benchmarks, showcasing the potential of flow-based dense SLAM approaches in theory and practice.

Despite recent progress, existing flow-based monocular dense SLAM systems still perceive and reconstruct scenes solely through pixel-wise correspondences from the 2D optical flow, resulting in a lack of geometric consistency in both

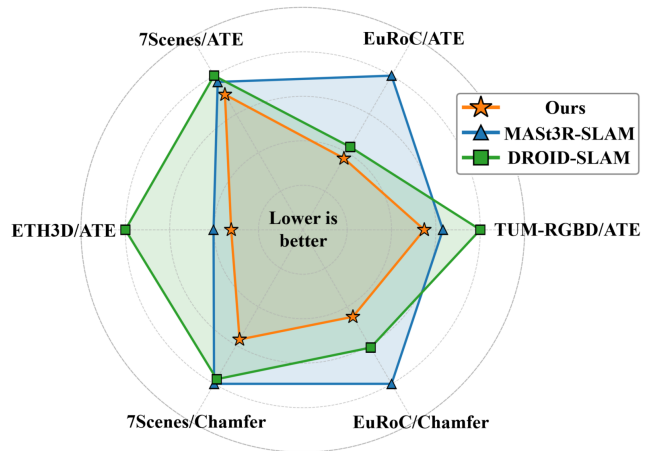


Figure 1: SLAM Performance Comparison. Radar plot shows normalized ATE (TUM, EuRoC, 7Scenes, ETH3D) and Chamfer distance (7Scenes, EuRoC). Our method (orange star) achieves optimal performance on any metrics.

tracking and mapping. As a result, the reconstructed depth may exhibit structural artifacts, layered ambiguities, or incomplete geometry, ultimately degrading pose accuracy and reconstruction quality.

Specifically, two key factors are primarily behind this limitation. First, dense correspondence estimation is performed solely in image space and lacks awareness of underlying scene geometry, leading to structurally inconsistent matches across views, especially in textureless and ambiguous regions. Second, current systems lack explicit enforcement of multi-view geometric constraints during optimization and dedicated mechanisms to refine flow predictions based on these constraints, resulting in accumulated errors during optimization. These issues call for a tightly coupled framework where geometric priors direct the correspondence estimation, and multi-view optimization in turn guides refinement.

To this end, we propose FoundationSLAM, a monocular dense SLAM framework that integrates geometric guidance with multi-view constrained optimization into a fully differentiable pipeline. Our approach leverages the strong priors encoded in foundation depth models to guide flow match-

*Corresponding author: Xiao Bai (baixiao@buaa.edu.cn)

ing under challenging conditions, and introduces a novel Bi-Consistent Bundle Adjustment Layer that jointly refines depth and pose while enforcing multi-view consistency. Furthermore, a residual-based refinement mechanism enables the system to identify unreliable predictions and use geometric priors for targeted correction, thereby achieving robust and interactive SLAM across diverse scenarios.

We demonstrate the effectiveness of FoundationSLAM across multiple challenging benchmarks, including TUM-RGBD (Sturm et al. 2012), EuRoC (Burri et al. 2016), 7Scenes (Glocker et al. 2013), and ETH3D (Schops, Sattler, and Pollefeys 2019), outperforming existing approaches, as shown in Figure 1. Our main contributions are:

- We propose a Hybrid Flow Network that leverages geometric priors from foundation models for robust correspondence estimation. A Reliability-Aware Refinement module further improves predictions by selectively correcting unreliable flow based on optimization residuals.
- We propose a Bi-Consistent Bundle Adjustment Layer that jointly optimizes dense depth and pose with multi-view residuals. By enforcing bidirectional consistency, our formulation improves tracking and reconstruction accuracy in challenging scenes.
- Extensive experiments on standard SLAM benchmarks demonstrate that our method outperforms previous monocular dense SLAM systems in both trajectory accuracy and reconstruction quality, while running in real time at 18 FPS.

2 Related Work

Matching-based SLAM. Classical and modern SLAM systems often adopt a keyframe-based matching paradigm using inter-frame correspondences to jointly optimize camera poses and scene geometry. Early systems (e.g., ORB-SLAM (Mur-Artal and Tardós 2017), DSO (Engel, Koltun, and Cremers 2017)) rely on handcrafted pipelines involving tracking, mapping, and bundle adjustment. With the rise of deep learning, end-to-end SLAM methods such as DROID-SLAM (Teed and Deng 2021) and DPVO (Teed, Lipson, and Deng 2023) incorporate learnable feature extractors and differentiable optimization modules for dense, robust tracking. A key advantage of these methods lies in the tight coupling between front-end perception and back-end optimization, where matching predictions and optimization mutually inform each other, leading to accurate, consistent SLAM even in challenging scenarios. However, these methods still estimate pixel-wise correspondences solely based on local correlation, without incorporating explicit geometric priors. This makes them prone to matching failures in textureless or ambiguous regions. More critically, the estimated flow is not constrained to be consistent across views, leading to structural inconsistencies and degraded reconstruction quality over long sequences.

Hybrid SLAM with Global Scene Representations. To improve global consistency, several recent works incorporate explicit 3D representations such as NeRF (Sucar et al. 2021; Zhu et al. 2022; Johari, Carta, and Fleuret 2023) or 3D Gaussian Splatting (Keetha et al. 2024; Yan et al. 2024;

Matsuki et al. 2024) into SLAM pipelines (Tosi et al. 2024). These methods use a matching-based frontend for pose tracking and simultaneously optimize a global scene representation. This hybrid structure enables dense and globally consistent reconstructions and even supports loop closure. Despite their benefits, these systems suffer from key drawbacks: they often require scene-specific optimization, assume known intrinsics, and demand high computational cost. More importantly, the global representation is updated independently of the pose tracker, weakening the feedback between perception and optimization. Compared to tightly coupled systems, this loose interaction further reduces the ability to correct or guide front-end predictions using multi-view consistency.

SLAM from 3D Reconstruction Models. Foundation models for 3D reconstruction, such as DUST3R (Wang et al. 2024) and MAST3R (Leroy, Cabon, and Revaud 2024), have shown remarkable capabilities in predicting dense scene geometry and relative poses from as few as two uncalibrated images. These priors are increasingly used in SLAM pipelines such as MAST3R-SLAM (Murai, Dexheimer, and Davison 2025), VGGT-SLAM (Maggio, Lim, and Carlone 2025), and VGGT (Wang et al. 2025), where per-frame geometry predictions provide additional guidance for inter-frame matching. However, these systems typically predict geometry on a per-frame or per-pair basis, and the frontend priors are not explicitly refined or guided by the backend optimization. As a result, these systems often suffer from inaccurate priors in challenging regions and lack a mechanism to correct them through joint optimization. Other approaches (e.g., SLAM3R (Liu et al. 2025)) discard backend optimization entirely and directly fuse point clouds from foundation models, achieving efficiency but at the cost of robustness and long-term accuracy.

This work aims to unify these strengths by embedding geometry-aware priors into a tightly integrated optimization framework that enforces multi-view geometric consistency while retaining efficiency and generalization.

3 Method

FoundationSLAM addresses the lack of geometric consistency in flow-based SLAM systems through a fully differentiable and tightly coupled framework. As shown in Figure 2, our method consists of three key components: (1) a Hybrid Flow Network that injects structural awareness into correspondence estimation via geometric priors from foundation depth models, (2) a Bi-Consistent Bundle Adjustment Layer that jointly refines depth and pose while enforcing multi-view consistency, and (3) a Reliability-Aware Refinement mechanism that leverages optimization residuals to guide flow correction. Together, these components form a closed-loop system to achieve robust and consistent monocular dense SLAM.

3.1 Hybrid Flow Estimation with Geometry Prior

A core limitation of existing flow-based SLAM systems is that dense correspondences are estimated independently for each image pair, without enforcing consistency across




Figure 2: Method Overview. Given a pair of keyframes, we estimate dense optical flow using a hybrid network fusing geometry-aware features from a foundation depth model. Predicted flow is iteratively refined via Flow GRU, guided by context features and learned reliability masks. Refined flow drives a Bi-Consistent BA Layer jointly optimizing keyframe depth and pose using flow and geometry consistency residuals. Optimization feedback updates flow reliability in closed-loop manner. This process unrolls over multiple iterations to progressively improve accuracy and consistency.

views. This lack of structural awareness often leads to geometrically inconsistent matches, especially in low-texture or ambiguous regions, ultimately degrading the accuracy of downstream optimization. To address this, we propose a Hybrid Flow Network that injects geometric priors from foundation depth models into the correspondence estimation process. These priors encode global scene structure and enable the flow network to produce more consistent and reliable matches across multiple viewpoints.

Backbone Design. Inspired by FoundationStereo’s (Wen et al. 2025) effective integration of depth priors (Yang et al. 2025), we design a dual-branch architecture: (1) a *Geometric Prior Branch* utilizing the frozen FeatureNet encoder from FoundationStereo to provide stable geometric features learned from diverse real-world imagery, and (2) a *Task-Specific Adaptation Branch* with trainable CNN layers mirroring parts of FeatureNet, optimized for monocular SLAM data association challenges. Features from the two branches are fused via 3×3 convolution followed by residual layers, yielding the final matching descriptor. We additionally incorporate the frozen ContextNet from FoundationStereo for context with rich geometric priors. These pretrained modules are used purely as convenient sources of geometry-guided features and remain fixed during training.

Flow Estimation Process. Operating over a dynamically maintained keyframe graph, MixFeatureNet extracts fused matching features for target frames and neighbors, while ContextNet provides context feature. A Feature Aggregator processes these inputs with initial flow estimates, feeding into a Flow GRU module that iteratively predicts flow updates (ΔF) and confidence maps (ω). The integration with optimization feedback is detailed below.

3.2 Bi-Consistent Bundle Adjustment Layer

While modern SLAM systems combine flow prediction and optimization in a unified pipeline (Teed and Deng 2021;

Teed, Lipson, and Deng 2023), flow is typically estimated independently for each frame pair, without enforcing consistency across multiple views. This frame-wise decoupling limits the ability to recover globally coherent structure and leads to geometric inconsistencies.

To address this, we propose the Bi-Consistent Bundle Adjustment (BA) Layer, which introduces explicit multi-view geometric supervision into the optimization loop. By incorporating both flow alignment and geometric consistency residuals, it strengthens the connection between correspondence prediction and scene-level optimization, improving robustness and global coherence.

Flow Consistency Residual. Given a pixel \mathbf{u}_i in frame i , a 3D point is reconstructed using the estimated depth D_i , transformed to frame j using relative pose T_{ji} , and projected to the image plane:

$$\mathbf{u}_{\text{proj}} = \pi(T_{ji} \cdot \pi^{-1}(\mathbf{u}_i, D_i)). \quad (1)$$

We minimize the residual to the predicted correspondence:

$$\mathcal{L}_{\text{flow}} = \|\mathbf{u}_{\text{proj}} - (\mathbf{u}_i + F_{i \rightarrow j})\|_1. \quad (2)$$

This residual supervises the alignment between depth and flow, but only reflects consistency from frame i ’s perspective, which still lacks the multi-view geometry consistency.

Geometry Consistency Residual. To enforce explicit multi-view consistency, we introduce a symmetric constraint: for each \mathbf{u}_i in frame i , we project its reconstructed 3D point to frame j , and then sample the depth D_j at the projected location \mathbf{u}_j . Then, we check whether it supports back-projection to the original pixel to get the residual:

$$\mathbf{u}_j = \pi(T_{ji} \cdot \pi^{-1}(\mathbf{u}_i, D_i)), \quad (3)$$

$$\mathbf{u}_i^{\text{back}} = \pi(T_{ij} \cdot \pi^{-1}(\mathbf{u}_j, D_j)), \quad (4)$$

$$\mathcal{L}_{\text{geo}} = \|\mathbf{u}_i^{\text{back}} - \mathbf{u}_i\|. \quad (5)$$

Algorithm 1: Reliability-aware Flow Refinement at Iter t .

- 1: $corr = 4D$ Correlation Volume($F_{i \rightarrow j}^t + u_i$)
 - 2: $corr = corr \cdot M_i$
 - 3: $inputs = \text{FeatureAggregator}(corr, inp_i, F_{i \rightarrow j}^t, M_i)$
 - 4: $\Delta F_{i \rightarrow j}^t = \text{FlowGRU}(inputs)$
 - 5: $F_{i \rightarrow j}^{t+1} = F_{i \rightarrow j}^t + \Delta F_{i \rightarrow j}^t$
-

This residual penalizes geometric misalignment when frame j does not geometrically support the depth prediction from frame i . In practice, we compute \mathcal{L}_{geo} only for pixels with $\mathcal{L}_{\text{geo}} < \tau$ (with $\tau = 1$ pixel) to avoid enforcing agreement across occlusions or depth discontinuities.

To balance the influence of the two residuals, we introduce a confidence map $\omega(\mathbf{u}_i)$ that reflects the reliability of local flow predictions. The final loss combines both terms:

$$\mathcal{L}_{\text{BA}} = \sum_{\mathbf{u}_i \in \Omega} (\omega(\mathbf{u}_i) \cdot \mathcal{L}_{\text{flow}}(\mathbf{u}_i) + (1 - \omega(\mathbf{u}_i)) \cdot \mathcal{L}_{\text{geo}}(\mathbf{u}_i)) \quad (6)$$

where Ω is the set of valid pixels satisfying the consistency threshold. We minimize \mathcal{L}_{BA} using Gauss-Newton optimization. Each iteration performs once flow update followed by twice BA, we compute Jacobians with respect to both depth and pose and solve for incremental updates ΔD and ΔT . This bidirectional formulation integrates local matching cues with multi-view geometric constraints, leading to more consistent and robust SLAM optimization, especially in scenes with wide baselines, occlusion, or low texture.

3.3 Reliability-aware Flow Refinement

Accurate optical flow estimation is critical for robust SLAM, yet it often degrades in regions affected by occlusion, low texture, repetitive patterns, or wide-baseline motion. When uncorrected, these unreliable flows introduce errors into depth and pose estimates, compromising the global consistency of the reconstruction.

To mitigate this, we propose Reliability-aware Flow Refinement, a mechanism that dynamically adapts the refinement behavior based on geometric residuals from the Bi-Consistent BA Layer. Specifically, we construct a pixel-wise reliability mask that guides the flow update process by identifying regions of high and low confidence.

Edge-wise Flow Reliability. Given a co-visible keyframe pair (I_i, I_j) , we first compute a local reliability indicator $M_{i \rightarrow j}^{\text{edge}}(\mathbf{u}) \in \{0, 1\}$ based on forward projection residual:

$$M_{i \rightarrow j}^{\text{edge}}(\mathbf{u}) = \begin{cases} 1 & \text{if } \mathcal{L}_{\text{proj}}(\mathbf{u}) < \tau_{\text{edge}}, \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where \mathbf{u} is a pixel in I_i , and a small residual indicates that the current flow prediction is geometrically consistent with the BA alignment.

Node-wise Geometric Reliability. To incorporate a more global consistency check, we compute the average geometry

residual \mathcal{L}_{geo} across all neighbors $\mathcal{N}(i)$ of keyframe I_i :

$$M_i^{\text{node}}(\mathbf{u}) = \begin{cases} 1 & \text{if } \frac{1}{n} \sum_{j \in \mathcal{N}(i)} \mathcal{L}_{\text{geo}}^{(i,j)}(\mathbf{u}) < \tau_{\text{node}}, \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

where $n = |\mathcal{N}(i)|$. This node-wise indicator provides a global confidence estimate by assessing the geometric agreement of each pixel across views.

Reliability Adaptive Refinement. We combine the local and global indicators into a unified binary reliability mask:

$$M_i(\mathbf{u}) = M_{i \rightarrow j}^{\text{edge}}(\mathbf{u}) \cdot M_i^{\text{node}}(\mathbf{u}). \quad (9)$$

During flow refinement, this mask defines two distinct update strategies:

1. *Reliable regions* ($M_i = 1$): Flow updates rely on local correlation volumes, assuming the match lies within a narrow search window. This enables efficient refinement using high-resolution correlation sampling.

2. *Unreliable regions* ($M_i = 0$): We mask out correlation features and remove them from the update pipeline. Flow updates are instead driven by contextual information containing geometry priors, enabling robust correction in ambiguous regions.

Prior works often incorporate optimization residuals into flow refinement by simply feeding them into the predictor as additional input feature (Teed and Deng 2021). Despite some optimization feedback, it does not alter the flow estimation process structurally. All regions, reliable or not, still rely on local correlation-based matching, which can lead to persistent errors in ambiguous regions.

In contrast, our approach explicitly separates the refinement behaviors through a structured, reliability-aware design. By masking out correlation features in unreliable regions, we force the network to rely solely on geometry-guided context for refinement, thus learning to handle difficult areas more effectively. For reliable regions, correlation features are retained, ensuring both precision and efficiency. This selective refinement improves learning dynamics, robustness, and overall SLAM performance.

Implementation Notes. In practice, we set both $\tau_{\text{edge}} = 5$ and $\tau_{\text{node}} = 5$, which are slightly larger than the correlation search radius of 3. This ensures that residuals outside the scope of reliable cost volume matching are excluded, aligning with the intended division of refinement responsibilities across reliable and unreliable regions.

4 Experimental Results

In this section, we evaluate FoundationSLAM on multiple public benchmarks, demonstrating its effectiveness in both localization and dense reconstruction.

Implementation Details. We train our model on 6-frame sequences sampled from the TartanAir (Wang et al. 2020) dataset, following (Teed and Deng 2021). Each sequence forms a co-visibility graph with 18 edges. This sampling strategy is only used during training and does not affect the online keyframe selection during inference. We resize input images to 512×384 . The model is trained for 300K steps using AdamW optimizer with OneCycleLR scheduling, a

	360	Desk	Desk2	Floor	Plant	Room	Rpy	Teddy	Xyz	Avg.
ORB-SLAM3	-	0.017	0.210	-	0.034	-	-	-	0.009	-
DeepV2D	0.243	0.166	0.379	1.653	0.203	0.246	0.105	0.316	0.064	0.375
DeepFactors	0.159	0.170	0.253	0.169	0.305	0.364	0.043	0.601	0.035	0.233
DPV-SLAM	0.112	0.018	0.029	0.057	0.021	0.330	0.030	0.084	0.010	0.076
DPV-SLAM++	0.132	0.018	0.029	0.050	0.022	0.096	0.032	0.098	0.010	0.054
GO-SLAM	0.089	0.016	0.028	0.025	0.026	0.052	0.019	0.048	0.010	0.035
DROID-SLAM	0.111	0.018	0.042	0.021	0.016	0.049	0.026	0.048	0.012	0.038
VGGT-SLAM*	0.071	0.025	0.040	0.141	0.023	0.102	0.030	0.034	0.014	0.053
MASt3R-SLAM	0.049	0.016	0.024	0.025	0.020	0.061	0.027	0.041	0.009	0.030
FoundationSLAM (Ours)	0.055	0.015	0.028	0.020	0.014	0.038	0.015	0.018	0.009	0.024

Table 1: Tracking accuracy on TUM-RGBD dataset. *means using uncalibrated images.

	MH01	MH02	MH03	MH04	MH05	V101	V102	V103	V201	V202	V203	Avg.
ORB-SLAM3	0.071	0.067	0.071	0.082	0.060	0.015	0.020	-	0.021	0.018	-	-
DeepV2D	0.739	1.144	0.752	1.492	1.567	0.981	0.801	1.570	0.290	2.202	2.743	1.298
DeepFactors	1.587	1.479	3.139	5.331	4.002	1.520	0.679	0.900	0.876	1.905	1.021	2.040
DPV-SLAM	0.013	0.016	0.022	0.043	0.041	0.035	0.008	0.015	0.020	0.011	0.040	0.024
DPV-SLAM++	0.013	0.016	0.021	0.041	0.041	0.035	0.010	0.015	0.021	0.011	0.023	0.023
GO-SLAM	0.016	0.014	0.023	0.045	0.045	0.037	0.011	0.023	0.016	0.010	0.022	0.024
DROID-SLAM	0.013	0.014	0.022	0.043	0.043	0.037	0.012	0.020	0.017	0.013	0.014	0.022
MASt3R-SLAM	0.023	0.017	0.057	0.113	0.067	0.040	0.019	0.027	0.020	0.025	0.043	0.041
FoundationSLAM (Ours)	0.010	0.011	0.020	0.041	0.040	0.034	0.009	0.015	0.014	0.010	0.011	0.019

Table 2: Tracking accuracy on EuRoC dataset.

learning rate of 3.5×10^{-4} , and weight decay of 10^{-5} . Training takes approximately 5 days on 8 RTX 4090 GPUs with a batch size of 8. At test time, with an efficient design adopting ViT-S and running foundation encoding on half-resolution, the system runs at 18 FPS.

Baselines. We compare with state-of-the-art SLAM systems representing key paradigms: classical sparse methods (ORB-SLAM3 (Campos et al. 2021)), matching-based approaches (DeepV2D (Teed and Deng 2018), DeepFactors (Czarnowski et al. 2020), DPV-SLAM (Lipson, Teed, and Deng 2024), DROID-SLAM (Teed and Deng 2021)), the NeRF based approach GO-SLAM (Zhang et al. 2023), and recent geometry-enhanced techniques VGGT-SLAM (Wang et al. 2025), MASt3R-SLAM (Murai, Dexheimer, and Davison 2025).

4.1 Tracking Evaluation

We evaluate our system on three standard benchmarks: TUM-RGBD, EuRoC MAV, and ETH3D-SLAM. As shown in Tables 1 and 2, FoundationSLAM achieves state-of-the-art ATE RMSE across all datasets. On TUM-RGBD, it ranks first on 7 out of 9 sequences, with particularly strong performance in reflective or low-texture environments. On EuRoC, which involves grayscale drone footage with rapid motion and significant domain shift, our method outperforms both DROID-SLAM and MASt3R-SLAM, demonstrating

Method	ATE	AUC
ORB-SLAM3	0.135	16.661
DROID-SLAM	0.171	22.297
DPVO	0.137	22.628
DPV-SLAM	0.109	23.097
DPV-SLAM++	0.132	21.784
MASt3R-SLAM	0.086	23.935
FoundationSLAM (Ours)	0.069	24.775

Table 3: Tracking accuracy on ETH3D-SLAM dataset.

high robustness to viewpoint variation. ETH3D-SLAM further showcases its stability under severe motion blur and dynamic scenes. As reported in Table 3, our method achieves the highest AUC across error thresholds, indicating consistent localization quality over varying conditions.

4.2 Mapping Evaluation

We evaluate the dense reconstruction performance on 7Scenes (seq-01) and EuRoC (VICON room sequences), comparing against DROID-SLAM, MASt3R-SLAM, and VGGT-SLAM. Evaluation metrics include *Accuracy* (average distance from each reconstructed point to its closest ground-truth point), *Completion* (average distance from

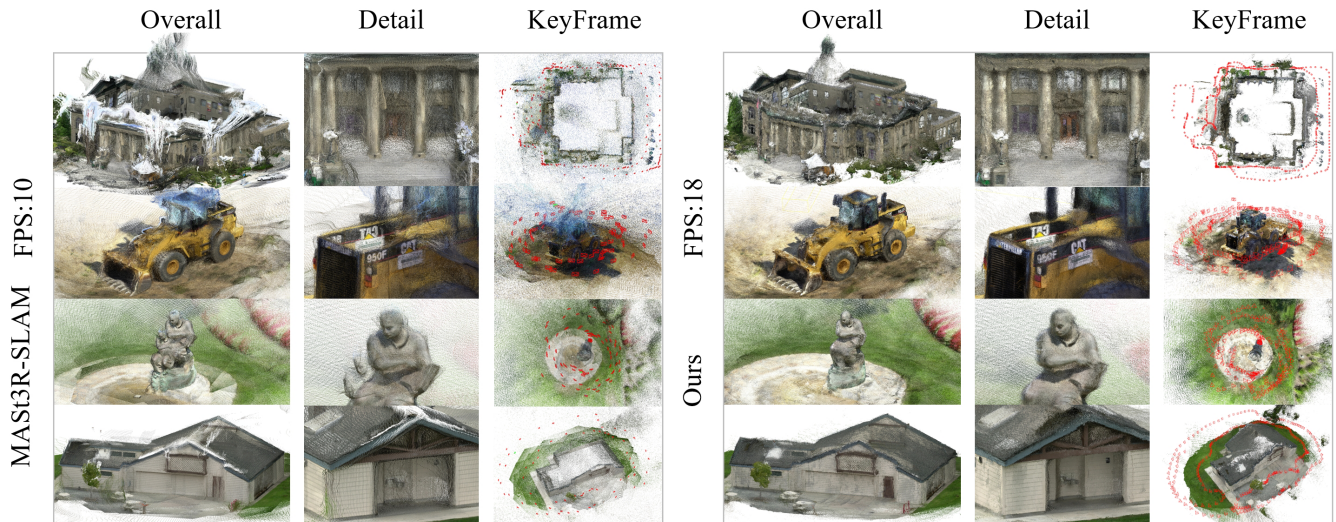


Figure 3: Qualitative Comparison on TNT Dataset. We show the qualitative results of our method and MAS3R-SLAM on the TNT (Knapsch et al. 2017) dataset, including the overall and detailed reconstruction, keyframe trajectories. Our method maintains significantly more keyframes while ensuring better geometric consistency, with less layering and artifacts.




Figure 4: Qualitative Reconstruction Comparison. Comparison with SOTA baselines on EuRoC and 7Scenes.

each ground-truth point to its nearest reconstructed point), and *Chamfer Distance* (the symmetric average of the two). For all metrics, distances are clipped at a maximum threshold of 0.5 meters to avoid the influence of large outliers.

As shown in Table 4, our method achieves the best

Method	ATE	Acc.	Comp.	Cf.
DROID-SLAM	0.049	0.052	0.076	0.064
MAS3R-SLAM	0.047	0.074	0.057	0.066
VGGT-SLAM*	0.067	0.052	0.058	0.055
FoundationSLAM (Ours)	0.043	0.039	0.055	0.047

Method	ATE	Acc.	Comp.	Cf.
DROID-SLAM	0.022	0.059	0.070	0.065
MAS3R-SLAM	0.041	0.099	0.071	0.085
FoundationSLAM (Ours)	0.019	0.035	0.063	0.048

Table 4: Tracking and mapping accuracy on 7Scenes (top) and EuRoC (bottom) datasets. *means using uncalibrated images.

DROID-SLAM	MAS3R-SLAM	VGGT-SLAM	Ours
24	10	26	18

Table 5: Comparison of frames per second on EuRoC.

overall reconstruction performance on both datasets. On the EuRoC VICON room sequences, our method surpasses DROID-SLAM in both accuracy and completion, resulting in a lower Chamfer distance. In contrast, MAS3R-SLAM shows poor reconstruction quality, likely due to the domain gap caused by its lack of training on grayscale data. These results demonstrate the robustness of our method in fast-motion, grayscale scenes with wide baselines. On the 7Scenes dataset, our approach improves reconstruction accuracy over VGGT-SLAM and MAS3R-SLAM. While all

Bi-BA	M_{node}	M_{edge}	ATE	Acc.	Comp.	Cf.
			0.021	0.051	0.070	0.061
	✓		0.021	0.047	0.067	0.057
		✓	0.021	0.048	0.068	0.058
	✓	✓	0.020	0.046	0.066	0.056
✓	✓		0.020	0.036	0.065	0.051
✓		✓	0.020	0.038	0.067	0.052
✓	✓	✓	0.019	0.035	0.063	0.048

Table 6: Ablation Studies on EuRoC datasets. Results demonstrate the contribution of each proposed component to localization and reconstruction performance in SLAM.

methods yield similar completion scores, our method still achieves the lowest Chamfer distance, indicating better overall geometry alignment. These results further validate the potential of the optical flow-based SLAM paradigm with end-to-end bundle adjustment for dense reconstruction, even when compared against methods that incorporate strong geometry priors. Figure 3 provides a direct comparison with MAST3R-SLAM on the TNT dataset, demonstrating our method’s reconstruction advantages. Additional qualitative comparisons are shown in Figure 4, highlighting sharper geometry and fewer outliers than competing methods.

Inference Speed. We evaluate the inference speed on EuRoC dataset in Table 5 on a single 4090 GPU. MAST3R-SLAM maintains significantly fewer keyframes, as shown in Figure 3, while VGGT-SLAM outputs tracking results per submap rather than per frame, contributing to their speed advantages. Our system achieves real-time inference at 18 FPS, striking a balance between performance and efficiency.

4.3 Ablation Studies

In Table 6, we conduct ablation studies on the EuRoC dataset to evaluate the impact of our Bi-Consistent BA Layer and Reliability-Aware Refinement components.

Bi-Consistent BA Layer. As shown in Figure 5, incorporating the Bi-Consistent Bundle Adjustment Layer leads to significantly improved geometric consistency across keyframes. This consistency helps produce more coherent depth and pose estimates, which in turn improves both localization and mapping performance. The quantitative improvements in Table 6 confirm that enforcing multi-view geometric alignment is critical for dense SLAM.

Reliability-Aware Flow Refinement. Our refinement module introduces two reliability modeling strategies: Node-wise Geometric reliability and Edge-wise Flow reliability. Both contribute individually, and their combination yields further gains, as reflected in Table 6. Figure 6 provides qualitative examples, where ours corrects errors in reflective, low-texture, or ambiguous regions that challenge correlation-based methods. These results highlight the importance of reasoning about flow reliability under difficult conditions.

Full Model. Combining Bi-consistent BA with both reliability modules yields the best overall performance. This demonstrates the benefit of tight integration between flow

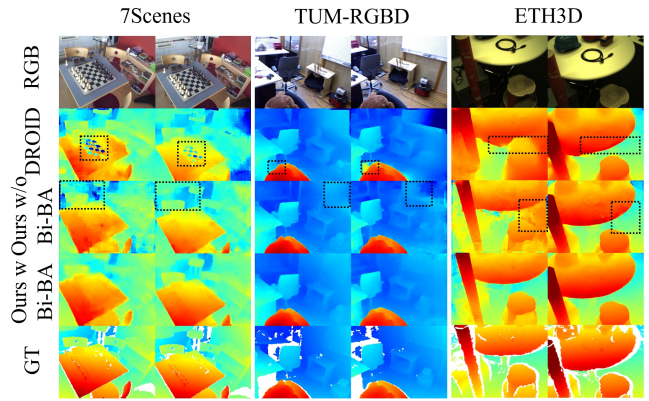


Figure 5: Impact of Bi-Consistent BA Layer. We visualize the depth maps of neighboring keyframes to reveal inconsistencies caused by unreliable optical flow in baseline systems. Our Bi-Consistent BA Layer significantly improves inter-frame geometric alignment.

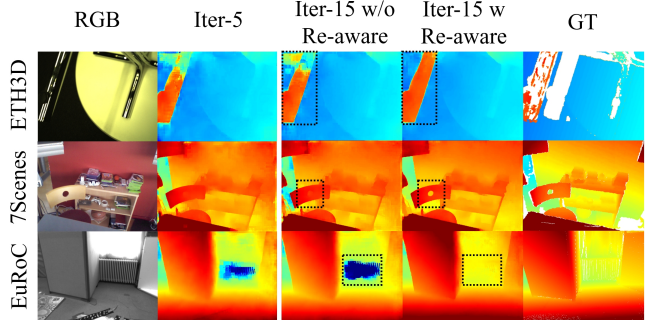


Figure 6: Flow Refinement on Challenging Areas. We visualize the flow refinement process in challenging regions across three datasets. Effect of the Reliability-Aware Refinement mechanism for correcting difficult-to-match areas such as holes, reflections, and repeated textures.

estimation and geometry reasoning, and highlights the importance of enforcing consistency and reliability throughout the SLAM optimization loop.

5 Conclusions

We present FoundationSLAM, an end-to-end monocular dense SLAM system that integrates geometry-guided optical flow estimation with multi-view consistent optimization. Our framework unifies three key components: a Hybrid Flow Network enhanced by foundation depth priors, a Bi-Consistent Bundle Adjustment Layer enforcing cross-view consistency, and a Reliability-Aware Refinement mechanism guided by geometric residuals. By tightly integrating optical flow estimation with geometry-aware optimization, FoundationSLAM achieves state-of-the-art performance in tracking and mapping across challenging benchmarks. Running in real time on monocular RGB input, our method demonstrates strong generalization and robustness, enabling practical use in diverse real-world scenarios.

Acknowledgements

This work is supported by the Key R&D Program of Jiangxi Province, China (20232BBE50019), and the National Natural Science Foundation of China (62276016, 62372029).

References

- Burri, M.; Nikolic, J.; Gohl, P.; Schneider, T.; Rehder, J.; Omari, S.; Achtelik, M. W.; and Siegwart, R. 2016. The EuRoC micro aerial vehicle datasets. *The International Journal of Robotics Research*, 35(10): 1157–1163.
- Campos, C.; Elvira, R.; Rodríguez, J. J. G.; Montiel, J. M.; and Tardós, J. D. 2021. Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam. *IEEE transactions on robotics*, 37(6): 1874–1890.
- Czarnowski, J.; Laidlow, T.; Clark, R.; and Davison, A. J. 2020. Deepfactors: Real-time probabilistic dense monocular slam. *IEEE Robotics and Automation Letters*, 5(2): 721–728.
- Engel, J.; Koltun, V.; and Cremers, D. 2017. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3): 611–625.
- Glocker, B.; Izadi, S.; Shotton, J.; and Criminisi, A. 2013. Real-time RGB-D camera relocalization. In *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 173–179. IEEE.
- Johari, M. M.; Carta, C.; and Fleuret, F. 2023. Eslam: Efficient dense slam system based on hybrid representation of signed distance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17408–17419.
- Keetha, N.; Karhade, J.; Jatavallabhula, K. M.; Yang, G.; Scherer, S.; Ramanan, D.; and Luiten, J. 2024. Splatam: Splat track & map 3d gaussians for dense rgb-d slam. In *CVPR*, 21357–21366.
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and temples: benchmarking large-scale scene reconstruction. *ACM Trans. Graph.*, 36(4).
- Leroy, V.; Cabon, Y.; and Revaud, J. 2024. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, 71–91. Springer.
- Lipson, L.; Teed, Z.; and Deng, J. 2024. Deep Patch Visual SLAM. In *European Conference on Computer Vision*.
- Liu, Y.; Dong, S.; Wang, S.; Yin, Y.; Yang, Y.; Fan, Q.; and Chen, B. 2025. Slam3r: Real-time dense scene reconstruction from monocular rgb videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 16651–16662.
- Maggio, D.; Lim, H.; and Carlone, L. 2025. VGGT-SLAM: Dense RGB SLAM Optimized on the SL(4) Manifold. *arXiv preprint arXiv:2505.12549*.
- Matsuki, H.; Murai, R.; Kelly, P. H.; and Davison, A. J. 2024. Gaussian splatting slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18039–18048.
- Mur-Artal, R.; and Tardós, J. D. 2017. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics*, 33(5): 1255–1262.
- Murai, R.; Dexheimer, E.; and Davison, A. J. 2025. MAST3R-SLAM: Real-Time Dense SLAM with 3D Reconstruction Priors. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, 16695–16705.
- Schops, T.; Sattler, T.; and Pollefeys, M. 2019. Bad slam: Bundle adjusted direct rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 134–144.
- Sturm, J.; Engelhard, N.; Endres, F.; Burgard, W.; and Cremers, D. 2012. A benchmark for the evaluation of RGB-D SLAM systems. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 573–580. IEEE.
- Sucar, E.; Liu, S.; Ortiz, J.; and Davison, A. J. 2021. imap: Implicit mapping and positioning in real-time. In *ICCV*, 6229–6238.
- Teed, Z.; and Deng, J. 2018. Deepv2d: Video to depth with differentiable structure from motion. *arXiv preprint arXiv:1812.04605*.
- Teed, Z.; and Deng, J. 2021. DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras. In *Advances in neural information processing systems*.
- Teed, Z.; Lipson, L.; and Deng, J. 2023. Deep Patch Visual Odometry. *Advances in Neural Information Processing Systems*.
- Tosi, F.; Zhang, Y.; Gong, Z.; Sandström, E.; Mattocchia, S.; Oswald, M. R.; and Poggi, M. 2024. How nerfs and 3d gaussian splatting are reshaping slam: a survey. *arXiv preprint arXiv:2402.13255*, 4: 1.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025. VGGT: Visual Geometry Grounded Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; and Revaud, J. 2024. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20697–20709.
- Wang, W.; Zhu, D.; Wang, X.; Hu, Y.; Qiu, Y.; Wang, C.; Hu, Y.; Kapoor, A.; and Scherer, S. 2020. Tartanair: A dataset to push the limits of visual slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 4909–4916. IEEE.
- Wen, B.; Trepte, M.; Aribido, J.; Kautz, J.; Gallo, O.; and Birchfield, S. 2025. FoundationStereo: Zero-Shot Stereo Matching. *CVPR*.
- Yan, C.; Qu, D.; Xu, D.; Zhao, B.; Wang, Z.; Wang, D.; and Li, X. 2024. Gs-slam: Dense visual slam with 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19595–19604.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2025. Depth anything v2. *NeurIPS*, 37: 21875–21911.
- Zhang, W.; Sun, T.; Wang, S.; Cheng, Q.; and Haala, N. 2024. HI-SLAM: Monocular Real-Time Dense Mapping

With Hybrid Implicit Fields. *IEEE Robotics and Automation Letters*, 9(2): 1548–1555.

Zhang, Y.; Tosi, F.; Mattoccia, S.; and Poggi, M. 2023. Go-slam: Global optimization for consistent 3d instant reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3727–3737.

Zhou, H.; Guo, Z.; Ren, Y.; Liu, S.; Zhang, L.; Zhang, K.; and Li, M. 2025. MoD-SLAM: Monocular Dense Mapping for Unbounded 3D Scene Reconstruction. *IEEE Robotics and Automation Letters*, 10(1): 484–491.

Zhu, Z.; Peng, S.; Larsson, V.; Xu, W.; Bao, H.; Cui, Z.; Oswald, M. R.; and Pollefeys, M. 2022. Nice-slam: Neural implicit scalable encoding for slam. In *CVPR*, 12786–12796.