

LCP-Fusion: A Neural Implicit SLAM with Enhanced Local Constraints and Computable Prior

Jiahui Wang, Yinan Deng, Yi Yang and Yufeng Yue*

Abstract—Recently the dense Simultaneous Localization and Mapping (SLAM) based on neural implicit representation has shown impressive progress in hole filling and high-fidelity mapping. Nevertheless, existing methods either heavily rely on known scene bounds or suffer inconsistent reconstruction due to drift in potential loop-closure regions, or both, which can be attributed to the inflexible representation and lack of local constraints. In this paper, we present LCP-Fusion, a neural implicit SLAM system with enhanced local constraints and computable prior, which takes the sparse voxel octree structure containing feature grids and SDF priors as hybrid scene representation, enabling the scalability and robustness during mapping and tracking. To enhance the local constraints, we propose a novel sliding window selection strategy based on visual overlap to address the loop-closure, and a practical warping loss to constrain relative poses. Moreover, we estimate SDF priors as coarse initialization for implicit features, which brings additional explicit constraints and robustness, especially when a light but efficient adaptive early ending is adopted. Experiments demonstrate that our method achieve better localization accuracy and reconstruction consistency than existing RGB-D implicit SLAM, especially in challenging real scenes (ScanNet) as well as self-captured scenes with unknown scene bounds. The code is available at <https://github.com/laliwang/LCP-Fusion>.

I. INTRODUCTION

Dense visual Simultaneous Localization and Mapping (SLAM) plays a vital role during perception, navigation and manipulation in unknown environments. In recent decades, traditional SLAM methods [1]–[3] have made significant progress in localization accuracy and real-time applications. However, due to the use of explicit scene representations like occupancy-grids [4]–[6], point cloud [7]–[9], Signed Distance Function [10]–[12], and surfels [13], which directly store and update limited scene information at fixed resolution without context, they struggle to balance memory consumption and mapping resolution, while being unable to reconstruct complete and consistent surfaces in noisy or unobserved areas.

Therefore, recent research has focused on implicit representation using neural networks [14] or radiance fields [15] to encode any points in scenes as continuous function, which can be used to extract isosurfaces at arbitrary resolution or to synthesize realistic unseen views. Utilizing the representation coherence and ability to render unseen views, numerous neural implicit SLAM systems [16]–[18] have emerged to perform high-fidelity mapping and camera tracking in various scenes. However, most of them require known scene

This work is supported by the National Natural Science Foundation of China under Grant 92370203, 62233002. (*Corresponding Author: Yufeng Yue, yueyufeng@bit.edu.cn)

¹ School of Automation, Beijing Institute of Technology, Beijing, China.

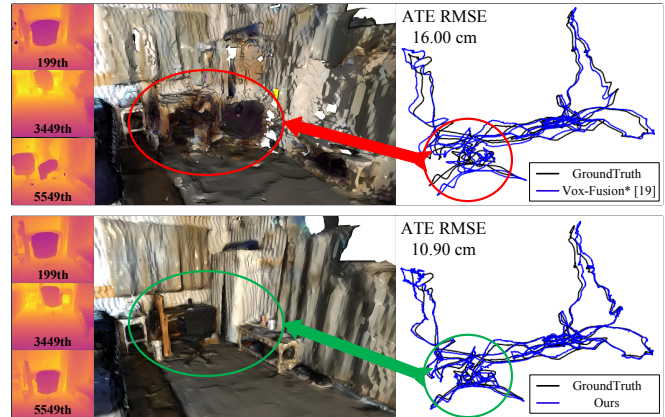


Fig. 1. Example from baseline [19] of inconsistent surfaces due to drift in potential loop-closure regions composed of frames 119, 3449 and 5549 (top). Our method can reconstruct unknown scenes with less drift utilizing enhanced local constraints and easily computable SDF priors (bottom).

bounds due to inflexible scene representation [19], resulting in performance degradation or failure in unknown scenarios.

Focusing on applications in unknown scenes, one of the mainstream solutions is to allocate implicit feature grids dynamically in surface areas utilizing flexible sparse voxel octree (SVO) [19]. Since SVO-based methods [20] only represent scenes using high-dimensional features in sparse voxel grids, they tend to be sensitive to odometry drift caused by their insufficient local constraints in loop closure regions. This may contribute to inconsistent reconstruction much further, which is shown in Fig. 1. Additionally, with an explicit SDF octree prior, hybrid methods [21] are proposed for precise mapping, but instead use a traditional visual odometry [22] as tracking module. Thus, for the unified dense SLAM system utilizing a neural implicit representation for both tracking and mapping, it is worth investigating in alleviating reconstruction inconsistency caused by localization drift in unknown scenes with potential loop closure.

To this end, we introduce LCP-Fusion (a neural implicit SLAM system with enhanced Local Constraints and Computable Prior), which can alleviate drift in potential loop-closure without other external modules. Our key ideas are as follows. **First**, to handle unknown scene boundaries, we utilize the SVO to dynamically allocate hybrid voxel grids containing coarse SDF priors and residual implicit features, which yields scene geometry and color through sparse volume rendering. **Second**, through pixel reprojection between frames, we propose a novel sliding window selection strategy based on visual overlap, which not only strengthens

local constraints but also alleviates catastrophic forgetting. In addition to only evaluating individual frame, a practical warping loss constraining relative poses is introduced to further improve localization accuracy. **Third**, to reduce the redundancy of iterations in joint optimization, we adopt an adaptive early ending without significant performance degradation owing to our proposed hybrid representation. We perform extensive evaluations on a series of RGB-D sequences to demonstrate the localization improvement of our method, as well as the applications in real scenes with unknown bounds. In summary, our contributions are:

- We present LCP-Fusion, a neural implicit SLAM system based on hybrid scene representation, which allocates hybrid voxels with implicit features and estimated SDF priors dynamically in scenes without known bounds.
- We introduce a novel sliding window selection strategy based on visual overlap and a warping loss constraining relative poses for the enhanced local constraints.
- Extensive evaluations on various datasets demonstrate our competitive performance in localization accuracy and reconstruction consistency, as well as robustness to fewer iterations and independence on scene boundaries.

II. RELATED WORKS

While existing neural implicit SLAM systems demonstrate impressive performance in high-fidelity reconstruction, incremental consistent mapping in unknown scenes remains challenging. We mainly attribute this to the following aspects: 1) inflexible implicit scene representation; and 2) insufficient local constraints during optimization.

A. Implicit Scene Representation

Compared to earlier works [14], [23] requiring 3D ground truth for supervision in object-level reconstruction, NeRF and its variants [15], [24] encode 3D points in the weights of MLP directly from image sequences. Against earlier works [25], [26] that required well-designed encoder-decoder networks to capture scenes, NeRF-based SLAM jointly optimizes implicit maps and camera poses from end to end using shallow MLP decoders. While iMAP [16] first introduces NeRF [15] into the framework of dense SLAM, due to the limited capability and slow training speed of a single MLP, its mapping and tracking effects drop during continual learning. Thus, Nice-SLAM [17] proposes to encode local scenes with multi-resolution dense feature grids and pre-trained MLP decoder [26], which greatly enriches mapping details and speeds up the training with less catastrophic forgetting. Furthermore, Co-SLAM [18] designs a joint encoding to combine the benefits of coordinate encoding and parametric embedding, achieving both consistent and sharp real-time reconstruction.

However, these methods above assume known scene bounds in application, as they have a requirement for scene bounds to allocate dense grids or perform normalized positional encoding before running. To this end, Vox-Fusion [19] utilizes sparse voxel octree (SVO) to allocate feature grids on-the-fly and force scene learning towards surface

areas. Since [19] only encodes scene at the scale of leaf voxel, the feature drift caused by odometry drift is difficult to correct in high dimensional space. Thus, we adopt a hybrid scene representation that store an easy-to-get SDF prior at the same resolution, which can not only provide reasonable initialization but also facilitate the correction with explicit updates in low dimensional space.

B. Local Constraints during Optimization

To address reconstruction consistency during incremental mapping, most implicit SLAM systems jointly optimize implicit maps and keyframe poses [27] with a sliding window, while some introduce additional relative pose constraints.

For sliding window selection, iMAP [16] and iSDF [28] select keyframes according to their loss distribution, which is time consuming during active sampling. With pre-trained MLP decoder, Nice-SLAM [17] only selects keyframes overlapped with the current frame without considering catastrophic forgetting. For Vox-Fusion [19], without pre-trained priors, it has to randomly select keyframes from the global keyframe set, potentially weakening local constraints between overlapped frames. As for Co-SLAM [18], they randomly sample rays from the global keyframe set for global BA, which is optimized across all keyframes yet lack of local constraints. In summary, most sliding windows lack sufficient local constraints from overlapped loop-closure frames.

For relative pose constraints, Nicer-SLAM [29] introduces a sophisticated RGB warping loss among random and latest keyframes to further enforce geometry consistency, and Nope-Nerf [30] performs warping process between every consecutive frames to constrain relative poses. However, they only constrain relative poses between nearby or random frames instead of frames with visual overlap in loop closure.

Concurrently, while DIM-SLAM [31] (partially open source) employs a similar strategy to enhance local constraints, it relies exclusively on RGB inputs and multiple levels of dense feature grids. This approach may result in less precise reconstruction and increased reliance on environmental priors such as scales. Additionally, although other recent methods [32], [33] are meticulously crafted for loop closure, they are still constrained by known scene boundaries.

To this end, we further introduce warping loss for relative pose constraints between potential loop-closure frames, in a sliding window based on visual overlap.

III. LCP-FUSION

An overview of our system is shown in Fig. 2. Receiving continuous RGB-D frames of color images $C_i^{obs} \in \mathbf{R}^3$ and depth images $D_i^{obs} \in \mathbf{R}$ without poses, our expanded hybrid representation first allocates hybrid voxels (Sec. III-A) where exists valid point clouds under estimated pose from the tracking process. Then volume rendering based on ray-voxel intersection is processed through SDF prior and implicit feature grids, which yields rendered RGB \hat{C}_i , Depth \hat{D}_i and predicted SDF values s_i for optimization. By employing a sliding window selection strategy based on visual overlap (Sec. III-B), more keyframes related to the current frame

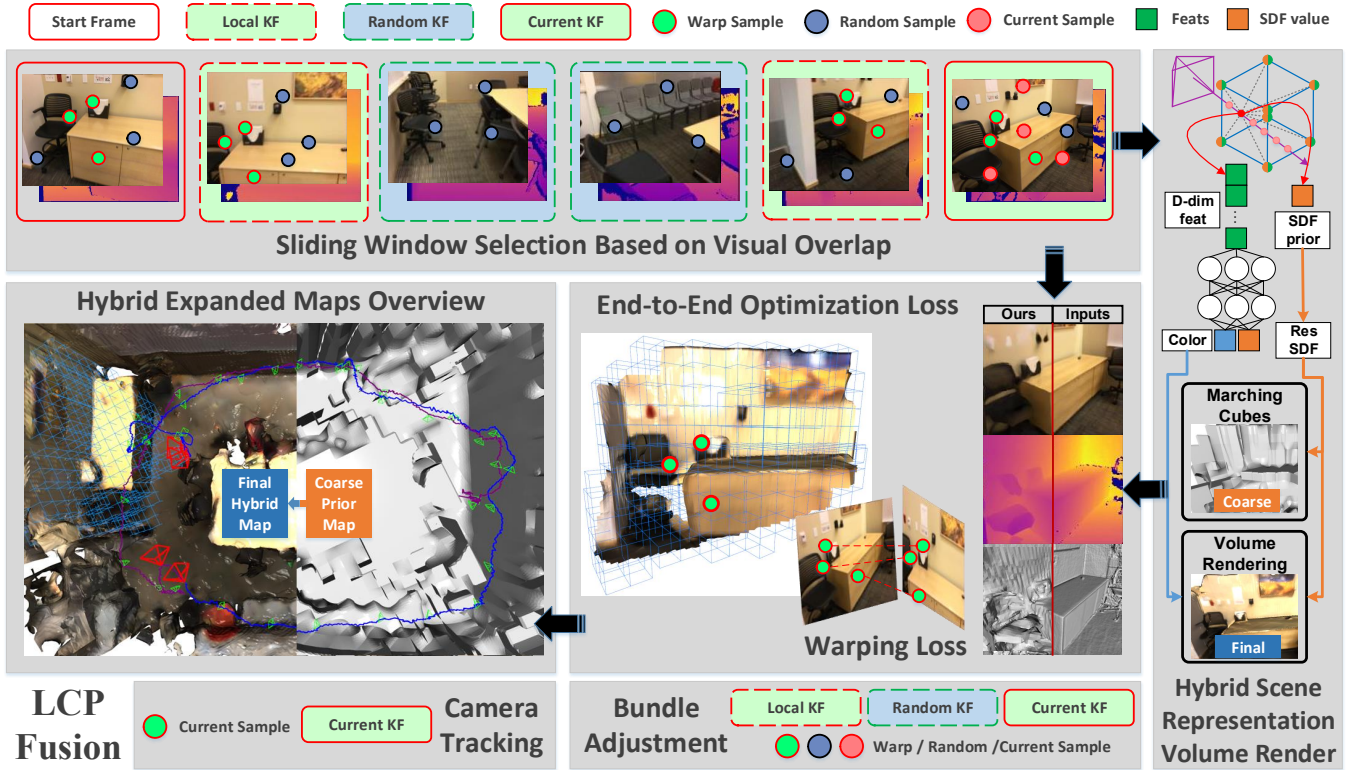


Fig. 2. Overview of our LCP-Fusion system. Receiving RGB-D inputs with initialized poses from the tracking process, we jointly optimize hybrid scene representation and camera poses among our sliding window, which contains more visual overlap with the current frame in potential loop closure regions. Additionally, the proposed warping loss can obtain sufficient constraints for relative poses with visual overlap.

are evaluated during bundle adjustment optimization. Then several loss functions include our warping loss are defined to optimize camera poses and scene representation during tracking and mapping (Sec. III-C).

A. Hybrid Scene Representation

In general, we represent unknown scene utilizing a hybrid multi-level SVO to allocate leaf voxels v^i dynamically, which store D-dim feature embeddings $\{e_{k=1\sim 8}\}_i$ and 1-dim SDF priors $\{s_{k=1\sim 8}^{prior}\}_i$ at each vertex, where SDF priors are easily computed through back-projection and provided as reasonable initialization for implicit features.

During volume rendering, sampled point p_j along rays can easily get point-wise features E_j and coarse SDF values s_j^c through trilinear interpolation $TriLerp(\cdot)$. Then E_j can be sent to MLP decoder $M_\theta(\cdot)$ for residual SDF $s_j^{res} \in R$ and color $c_j \in R^3$ at point p_j . Along with s_j^c from SDF priors, final results for volume rendering can be obtained:

$$(E_j, s_j^c) = TriLerp\left(p_j, \left\{e_{k=1\sim 8}, s_{k=1\sim 8}^{prior}\right\}_i\right) \quad (1)$$

$$(c_j, s_j^{res}) = F_\theta(E_j), s_j = s_j^c + s_j^{res} \quad (2)$$

As shown in Fig. 3, whenever enough points from new frames drop inside the leaf voxel v^i , we first estimate current SDF priors s_{curr}^{prior} at each vertex through projecting them onto current depth frame, and then update them separately like TSDF fusion [11]:

$$s_{curr}^{prior} = D(u) - d_p \quad (3)$$

$$\begin{cases} s_{fusion}^{prior} = \frac{s_{fusion}^{prior} \cdot n_{update} + s_{curr}^{prior} \cdot n_{curr}}{n_{update} + n_{curr}} \\ n_{update} = n_{update} + n_{curr} \end{cases} \quad (4)$$

where d_p is the z-axis distance of vertex p under the current camera coordinate, and $D(u)$ is the depth value at the re-projected pixel u . Due to the fact that a vertex can be shared by up to 8 surrounding voxels, s_{curr}^{prior} tends to be updated by several voxels. We directly take updated voxels number n_{curr} as current update weight to obtain final values s_{fusion}^{prior} through weighted fusion while updating n_{update} .

However, before updating s_{fusion}^{prior} with current estimate s_{curr}^{prior} , we need to eliminate the invalid estimates in the following situations:

- 1) The re-projected pixel u of vertex p locates outside current frame.
- 2) The depth value $D(u)$ of pixel u is absent due to sensor noise or depth truncation.
- 3) The estimate value s_{curr}^{prior} cannot meet $|D(u) - d_p| < \sqrt{3} \times (voxel\ size)$ as Fig. 3.

With explicit initialization and updates in low dimensional space, although at the same resolution, our system achieves improved accuracy in localization and stability over fewer iterations of optimization, as well as consistent reconstruction in noisy or inaccurate areas.

Much further, receiving inferred values (c_j, s_j) of N points from hybrid representation, sparse volume rendering

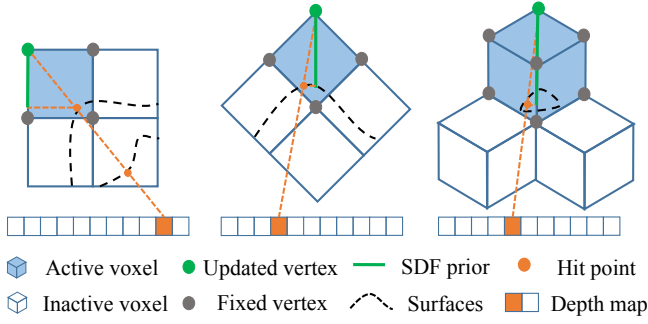


Fig. 3. Visualization of SDF prior estimates. To avoid unreasonable SDF priors due to occlusion, we indicate three extreme cases from left to right.

performs along each ray as follows:

$$w_j = \sigma\left(\frac{s_j}{tr}\right) \cdot \sigma\left(-\frac{s_j}{tr}\right)$$

$$\hat{C} = \frac{1}{\sum_{j=0}^{N-1} w_j} \sum_{j=0}^{N-1} w_j \cdot c_j, \quad \hat{D} = \frac{1}{\sum_{j=0}^{N-1} w_j} \sum_{j=0}^{N-1} w_j \cdot d_j \quad (5)$$

where $\sigma(\cdot)$ is the sigmoid function and tr is a pre-defined truncation distance. With the point-wise rendering weight w_j computed by predicted SDF s_j , rendered RGB \hat{C} and depth \hat{D} of rays are obtained by weighting predicted point-wise color c_j and sampled depth d_j respectively. The rendered RGB images, depth images and predicted SDF values will be used to evaluate with inputs for optimization in Sec. III-C.

B. Sliding Window Selection

For selection of optimized keyframe in sliding windows, we propose a light yet efficient strategy to enhance local constraints and alleviate catastrophic forgetting as well. Specifically, our sliding window F_{window} consists of three parts: local keyframes F_{local} , historical keyframes F_{his} and current frame F_{curr} . For current frame F_{curr} with optimized camera pose, we first randomly sample N_{rep} pixels q_c with valid depth values, and then re-project them onto every keyframe in global keyframe list $\sum_i^{N_{kf}} F_{kf_i}$ for re-projected pixels q_{kf_i} . After counting the number of re-projected pixels that have successfully dropped inside each keyframe as $Count_{kf_i}$, we can obtain each component of sliding windows as follows:

$$F_{local} = \text{maxsort}\left(\sum_i^{N_{kf}} F_{kf_i}, \mathcal{W}/2\right)$$

$$F_{his} = \text{randsort}\left(\sum_i^{N_{kf}} F_{kf_i}, \mathcal{W}/2\right) \quad (6)$$

$$F_{window} = F_{local} + F_{his} + F_{curr}$$

where $\mathcal{W}+1$ is the width of sliding window, local keyframes F_{local} are selected among keyframes that have most points projected onto, historical keyframes F_{his} are selected randomly from global keyframe list. Both of them are essential for enhancing local constraints and alleviating forgetting, which is shown in Fig. 4. Especially, in scenes with frequent loop closures, F_{local} are modified to be selected randomly among $2\mathcal{W}$ keyframes with most re-projected points, obtaining adequate constraints at far enough intervals.

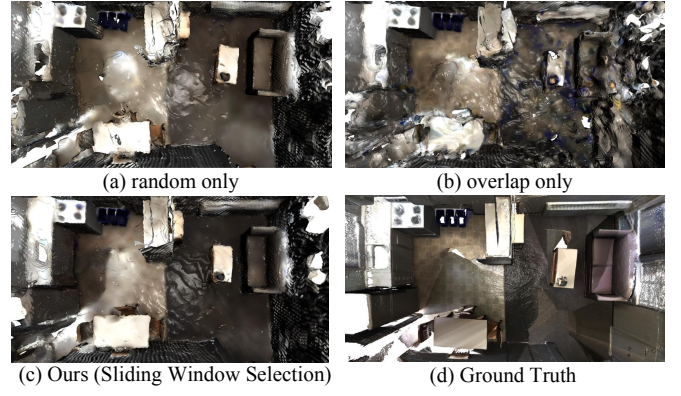


Fig. 4. Aerial view of reconstruction on scene0181 [34] using different sliding window designs. Through comparison, our sliding window selection method can not only improve localization but also alleviate forgetting.

C. End to End Optimization

1) *Loss Functions*: To supervise the learning of camera poses and hybrid representation, we apply the similar loss functions as [19]: re-render RGB loss (\mathcal{L}_{rgb}) and Depth loss (\mathcal{L}_d), SDF free-space loss (\mathcal{L}_{fs}) and estimate loss (\mathcal{L}_{sdf}) on a batch of rays ($R = R_{t,m}$) sampled randomly from current frame or sliding windows:

$$\mathcal{L}_{rgb} = \frac{1}{|R|} \sum_{r \in R} \left\| \hat{C}_r - C_r^{obs} \right\|, \quad \mathcal{L}_d = \frac{1}{|R|} \sum_{r \in R} \left\| \hat{D}_r - D_r^{obs} \right\|$$

$$\mathcal{L}_{fs} = \frac{1}{|R|} \sum_{r \in R} \frac{1}{N_r^{fs}} \sum_{n \in N_r^{fs}} (s_n - tr)^2$$

$$\mathcal{L}_{sdf} = \frac{1}{|R|} \sum_{r \in R} \frac{1}{N_r^{tr}} \sum_{n \in N_r^{tr}} (s_n - (D_r^{obs} - d_n))^2 \quad (7)$$

where photometric loss \mathcal{L}_{rgb} and \mathcal{L}_d are defined between render results (\hat{C}_r, \hat{D}_r) and observations (C_r^{obs}, D_r^{obs}); geometric loss \mathcal{L}_{fs} and \mathcal{L}_{sdf} are defined respectively on point-wise predicted SDF values s_n outside and inside truncation distance along rays.

Additionally, during bundle adjustment optimization, we introduce a warping loss defined between current frame and sliding window keyframes on another batch of rays R_w to constrain relative poses and enforce geometry consistency. Let q_c denotes a 2D pixel with valid depth value at current frame F_{curr} , we first project it into 3D space with camera parameters ($K, \tilde{R}_c | \tilde{t}_c$) and then re-project it onto sliding window keyframes $F_{w \in \mathcal{W}}$ as:

$$q_w = K \tilde{R}_w^T \left(\tilde{R}_c K^{-1} q_c^{homo} D_{q_c} + \tilde{t}_c - \tilde{t}_w \right) \quad (8)$$

where $q_c^{homo} = (u, v, 1)^T$ is the homogeneous coordinate of q_c and D_{q_c} denotes its depth observation. With the aid of depth observation and visual overlap in sliding windows, we define the RGB and depth warping loss respectively on the current frame:

$$\mathcal{L}_{warp_{R,D}} = \frac{1}{|R_w|} \sum_{r \in R_w} \sum_{w \in \mathcal{W}, w \neq c} \left| (I, D)_{q_c}^r - (I, D)_{q_w}^r \right| \quad (9)$$

It is worth noting that since the sliding window defined in Sec. III-B has more visual overlap with the current frame, we can obtain more valid re-projected points for warping constraints. The final loss function is a weighted sum of all above losses:

$$\begin{aligned} \mathcal{L}(\mathbf{P}) = & \lambda_{rgb} \mathcal{L}_{rgb} + \lambda_d \mathcal{L}_d + \lambda_{sdf} \mathcal{L}_{sdf} \\ & + \lambda_{fs} \mathcal{L}_{fs} + \lambda_{warp_R} \mathcal{L}_{warp_R} + \lambda_{warp_D} \mathcal{L}_{warp_D} \end{aligned} \quad (10)$$

where $\mathbf{P} = \{\theta, \mathbf{E}, \{\xi_t\}\}$ represent trainable MLP weights θ , feature grids \mathbf{E} and camera poses $\{exp(\xi_t^\wedge) \in SE(3)\}$. And we consider the above weights as constants without special declaration.

2) *Tracking and Mapping*: For Tracking, $\mathbf{P} = \{\xi_{curr}\}$ and we employ final loss function without warping loss. For every new coming RGB-D frame F_{curr} , we initialize the pose of F_{curr} identical to last tracked frame, and then sample R_t pixels randomly for volume rendering in frozen representation. Finally we only optimize current pose $\xi_{curr} \in se(3)$ through loss back-propagation for $Iter_t$ iterations.

For Mapping, $\mathbf{P} = \{\theta, \mathbf{E}, \{\xi_t\}\}$ and we employ the entire final loss function. Receiving every tracked frame with coarse pose estimates, we sample R_m pixels among sliding window F_{window} and additional R_w pixels on current frame F_{curr} for joint optimization and warping constraints respectively. Finally we jointly optimize scene representation and camera poses of sliding window for $Iter_m$ iterations.

Moreover, we provide an adaptive early ending policy for mapping process, which encourages system to optimize more iterations at areas with higher loss while less iterations for others. The details and effects are demonstrated in Sec. IV-D

IV. EXPERIMENT

To evaluate the performance of our proposed method, we first compare its localization accuracy and reconstruction consistency with other RGB-D Neural implicit SLAM on the real-world ScanNet dataset [34], synthetic Replica dataset [35] and our self-captured dataset with a depth camera mounted on mobile robots. And then we demonstrate our independence of scene bounds against SOTA bounded methods [18], as well as the effects of our hybrid representation against [19]. Moreover, we also conduct ablation studies to confirm the effectiveness of each module from our method.

A. Experimental Setup

1) *Implementation Details*: For scene representation, we maintain a 8-level SVO to allocate leaf voxels of 0.2m, which store 16-D feature and 1-D SDF prior at each vertex. For MLP decoder, we use the same structure as [19] for residual SDF and color. $(R_{rep}, R_w, R_t, R_m) = 1024$ pixels are sampled randomly from each frame for optimization during tracking and mapping, $Iter_t = 30$ and $Iter_m = 15$. For mapping, we maintain a sliding window of $\mathcal{W} = 4$ and insert a new keyframe at fixed interval of 50. All experiments are carried out on a desktop system with an Intel Xeon Platinum 8255C 2.50GHz CPU and an NVIDIA Geforce RTX 3090 GPU.

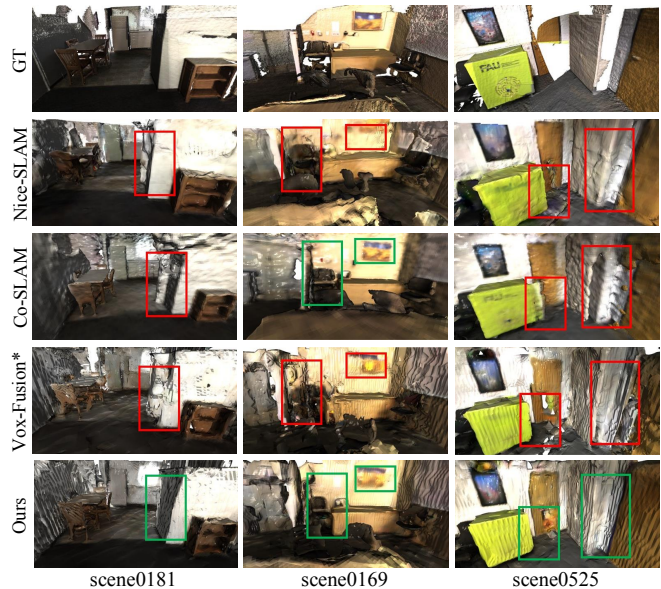


Fig. 5. Reconstructed mesh results of potential loop-closure regions in ScanNet. For comparison, we highlight the regions with inconsistent surfaces by red boxes, while green boxes for consistent surface.

2) *Baselines*: We consider Nice-SLAM [17], Co-SLAM [18] and Vox-Fusion [19] as our baselines for localization accuracy. Since [17] and [18] both require known scene bounds in advance, we provide 3D bounds for them following their instruction. Moreover, in order to demonstrate their dependence on known scene bounds, we provide coarse scene bounds for [18] in Sec. IV-C, which can be oversized or undersized. And we reproduce [19] as Vox-Fusion* from their official release for fair comparison.

3) *Metrics*: For evaluation of localization accuracy, we adopt absolute trajectory error (ATE) between estimate poses $\{P_1, P_2, \dots, P_n \in SE(3)\}$ and ground truth poses $\{Q_1, Q_2, \dots, Q_n \in SE(3)\}$. Specifically, we take the ATE RMSE for evaluation using scripts provided by [19].

B. Evaluation of Localization

1) *Evaluation on ScanNet*: To demonstrate the improvement in localization accuracy, especially in challenging scenarios from real capture. We first evaluate our system on representative sequences of ScanNet, which are captured from real scenes with potential loop-closure and contain more noises in depth inputs. The quantitative results are shown in Table. I. As can be seen, our localization accuracy surpass both bounded and unbounded methods by a large margin in these challenging scenes with loop closure. Moreover, we also qualitatively compare our system on reconstruction and render quality with baselines, shown in Fig. 5 and Fig. 6. It can be seen that owing to our localization improvement brought by local constraints and priors, there are less inconsistent surfaces in scenes, which is common in previous reconstructions, such as the wall in scene0181, the desk in scene0169 and the wardrobe in scene0525.

2) *Evaluation on Replica*: We also evaluate our localization enhancement on synthetic RGB-D sequences of Replica. Since the image quality of this synthetic dataset

TABLE I
ATE RMSE (CM) ON SELECTED SCANNET SEQUENCES

Scene ID	Reference	0059	0169	0181	0207	0472	0525	Avg.
Nice-SLAM	CVPR2022	12.25	10.28	12.93	6.65	9.64	10.31	10.34
Co-SLAM	CVPR2023	12.29	6.62	13.43	7.13	12.38	11.74	10.60
Vox-Fusion*	ISMAR2022	9.06	9.64	15.38	7.74	9.98	8.42	10.04
Ours	-	7.56	5.91	10.18	6.29	8.33	5.06	7.22

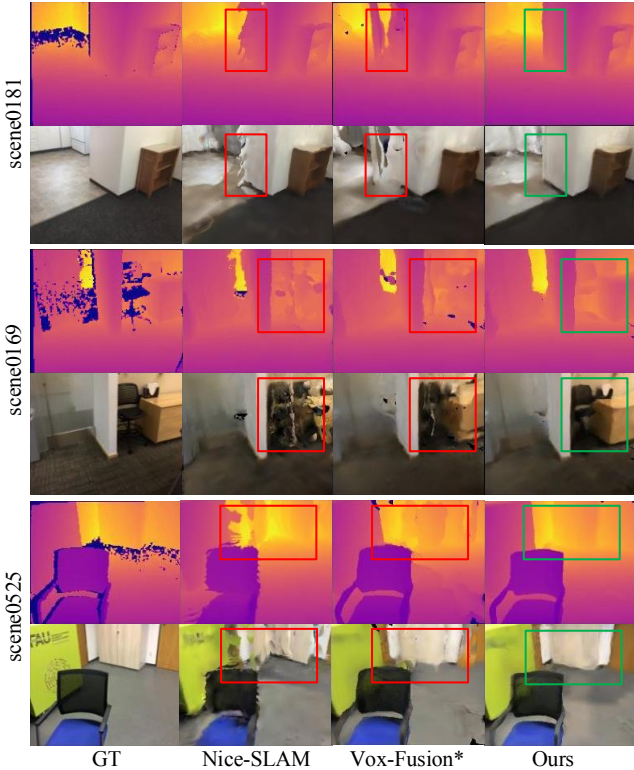


Fig. 6. Rendered Depth and Color results of potential loop-closure regions in ScanNet. Our method can improve surface consistency in highlight boxes due to enhanced local constraints and SDF priors.

is particularly pure, and each image can provide sufficient geometric and color constraints for tracking and mapping, our baselines have achieved high localization accuracy. Even so, our system can still outperform them to some extent, due to our additional constraints during bundle adjustment optimization. The quantitative results are shown in Table. II.

3) *Evaluation on Self-Captured*: To evaluate our performance in practical scenes with unknown bounds, we capture two RGB-D sequences by Azure Kinect camera mounted on mobile robots, which contains more missing depth values than ScanNet [34]. Without known scene bounds in advance, we only compare our LCP-Fusion against Vox-Fusion* [19] for they both utilize extensible SVO-based scene representations. The quantitative and qualitative results shown in Fig. 8 demonstrate that our method achieve better localization accuracy and reconstruction consistency in inaccurate or noisy regions, such as the incomplete wall in sc601 and calibration board in sc614.

C. Evaluation of Scene Representation

1) *Independence on Scene Bounds*: Since bounded neural implicit SLAM systems [18] require scene bounds as input

TABLE II
ATE RMSE (CM) ON SELECTED REPLICA SEQUENCES

Scene ID	Room0	Room1	Room2	Office1	Office2	Office3	Office4	Avg.
Nice-SLAM	1.69	2.04	1.55	0.90	1.39	3.97	3.08	2.09
Co-SLAM	0.67	1.44	1.11	0.57	2.10	1.58	0.90	1.20
Vox-Fusion*	0.64	1.36	0.84	1.15	0.98	0.72	0.92	0.94
Ours	0.54	1.02	0.78	1.08	0.92	0.66	0.85	0.84

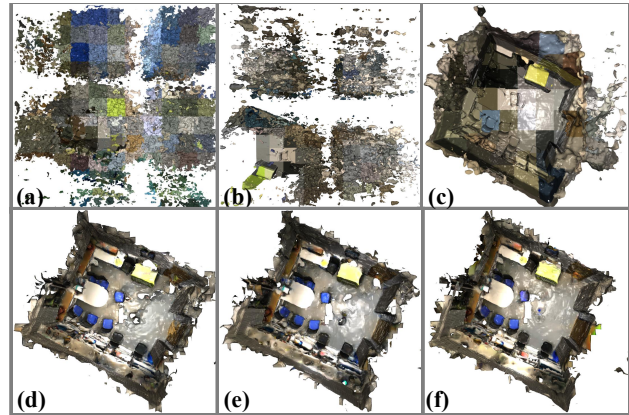


Fig. 7. Reconstructed results of [18] in scene0525 with inaccurate scene bounds (abc). Compared to degraded mesh quality and anamorphic patch-like texture, Ours (f) can perform high-fidelity mapping without any boundary priors; Compared to SVO-based baseline [19] (d) and LC-Fusion (e), Ours (f) demonstrate more consistent reconstructions.

for joint encoding and mesh extraction [36], we provide inaccurate scene bounds for [18] to magnify their reliance on scene bounds and essential difference with our LCP-Fusion. The reconstructed results given different scene bounds are shown in the first row of Fig. 7: (a) oversized for encoding and marching cubes; (b) undersized for both; (c) inaccurate for encoding but accurate for marching cubes; (f) LCP-Fusion (Ours) without boundary prior. As can be seen, inaccurate scene bounds for both in [18] will contribute to terrible mesh results. Moreover, even with accurate bounds for marching cubes, the surface texture still distorts due to the inappropriate encoding scale. For LCP-Fusion using SVO-based scene representation, we encode scenes and extract surfaces within dynamically allocated hybrid voxels, making it more suitable for scenes with unknown boundaries.

2) *Introduction of Hybrid representation*: As shown in earlier results, the SVO-based baseline [19] tends to suffer localization drift and inconsistent mapping in real scenes with potential loop closure, which is attributed to the lack of local constraints and difficulties in high-dimensional correction. Compared to (d) Vox-Fusion* and (e) LC-Fusion (Ours w/o hybrid rep.), (f) LCP-Fusion (Ours) yields more consistent reconstructions, which is shown in the second row of Fig. 7. Moreover, Table. IV presents quantitative comparisons with LC-Fusion that demonstrate improved localization accuracy using our hybrid representation. Therefore, in addition to our proposed local constraints, the hybrid scene representation will bring further enhancements to tracking and mapping.

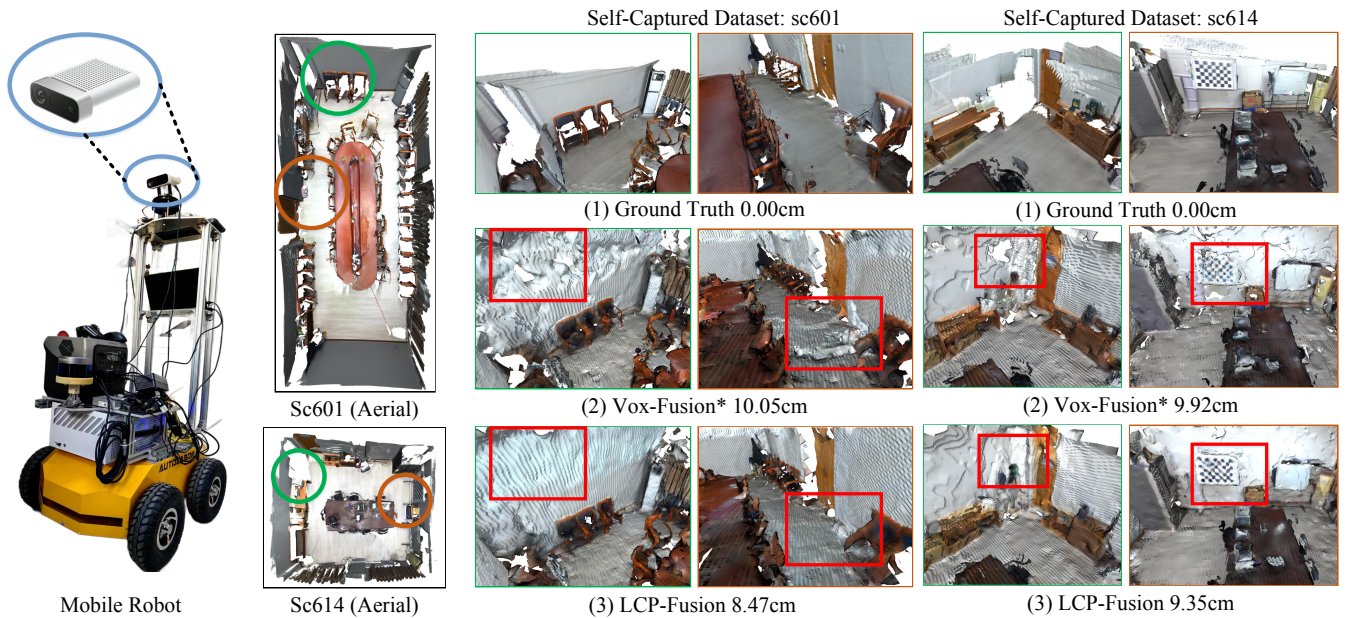


Fig. 8. Reconstructed mesh results of our self-captured dataset: sc601 and sc614. With improved localization accuracy, the highlight regions of zoomed windows indicate that our method yields more consistent surface against [19].

TABLE III
ABLATION STUDY ON SCENE0181 WITH BASELINE AND VARIANTS

Name	window selection		Warp loss	SDF prior	ATE (cm)
	random	ours			
Vox-Fusion*	✓				15.38
LCr-Fusion		✓			13.99
LCw-Fusion	✓		✓		14.21
LC-Fusion		✓	✓		11.19
LCP-Fusion		✓	✓	✓	10.18

TABLE IV
ATE RMSE (CM) AND TIME CONSUMPTION (MIN) ABLATIONS ON SELECTED SCANNET SEQUENCES

Scene ID	0059	0169	0181	0207	0472	0525	Avg.	Time
Vox-Fusion*	9.06	9.64	15.38	7.74	9.98	8.42	10.037	87.4
LC-Fusion	8.13	6.39	11.19	6.53	9.30	5.49	7.838	96.2
LCPe-Fusion	8.29	5.99	11.71	6.61	8.70	5.75	7.842	87.8
LCP-Fusion	7.56	5.91	10.18	6.29	8.33	5.06	7.222	103.2




Fig. 9. Rendered results in loop-closure regions of scene0000. Ours (closest) selects local keyframes in sliding window with most visual overlap; Ours selects randomly in a series of local keyframes with visual overlap.

D. Ablation Study

1) *Effect of enhanced constraints and priors*: In order to evaluate the effectiveness of all designed components, we conduct an ablation study on a representative scene, the quantitative results are shown in Table. III. As can be seen, our system with both local constraints and priors achieves highest localization accuracy among baseline and our incomplete variants: (LCr) LCP-Fusion with only our window selection; (LCw) LCP-Fusion with only our warping loss; (LC) LCP-Fusion without SDF priors.

2) *Effect of SDF priors with early ending*: Utilizing

SDF priors for reasonable initialization, we introduce an early ending policy for accelerating mapping process without dramatic degradation in performance. Specifically, current iterative optimization breaks in advance if there are over $Iter_m/3$ losses less than average loss on previous frames. Quantitative results for time and accuracy are shown in Table. IV and qualitative comparison in Fig. 10. It is evident that, despite the adaptive early ending, our system (LCPe-Fusion) still outperforms the main baselines [19] in terms of localization accuracy and reconstruction consistency, with a speed comparable to [19] and accuracy equivalent to our variants without SDF priors (LC-Fusion).

V. CONCLUSIONS

We propose LCP-Fusion (a neural implicit SLAM system with enhanced Local Constraints and Computable Prior). Utilizing the SVO-based hybrid scene representation, we show that jointly optimizing scene representation and poses in a novel sliding window consisting of local overlapped and historical keyframes, as well as constraining relative poses and geometry with warping loss, achieves accurate localization and consistent reconstruction in real scenes with noise and potential loop-closure. Furthermore, our introduction of computable SDF prior provides reasonable initialization for parametric encoding to further improve and stabilize perfor-

mance even when mapping iteration decreases. Compared to our baselines, we've observed a 28.1% improvement in localization accuracy on real datasets and 10.6% on synthetic datasets, further validated under our self-captured dataset. However, LCP-Fusion is still limited to basic spatial understanding of geometry and color from RGB-D inputs. Noting the impressive advances in web-pretrained visual language models (VLMs), neural implicit visual language SLAM for robotic downstream tasks can be our future work.




Fig. 10. SLAM results with adaptive early ending on scene0181. Compared with theoretical iterations of BA, VoXe-Fusion*(top) is trained for 72.99% iterations, while LCPe-Fusion (bottom) for 69.94%. Owing to the introduction of SDF priors, our system still perform stably with less iterations.

REFERENCES

- [1] Y. Deng, et al., "S-mki: Incremental dense semantic occupancy reconstruction through multi-entropy kernel inference," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3824–3829. IEEE, 2022.
- [2] R. Mur-Artal, et al., "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE transactions on robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [3] J. Engel, et al., "Direct sparse odometry," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.
- [4] A. Hornung, et al., "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous robots*, vol. 34, pp. 189–206, 2013.
- [5] Y. Deng, et al., "See-csom: Sharp-edged and efficient continuous semantic occupancy mapping for mobile robots," *IEEE Transactions on Industrial Electronics*, vol. 71, no. 2, pp. 1718–1728, 2023.
- [6] Y. Deng, et al., "Hd-ccsom: Hierarchical and dense collaborative continuous semantic occupancy mapping through label diffusion," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2417–2422. IEEE, 2022.
- [7] Y. Deng, et al., "Opengraph: Open-vocabulary hierarchical 3d graph representation in large-scale outdoor environments," *IEEE Robotics and Automation Letters*, pp. 1–8, 2024.
- [8] Y. Tang, et al., "Multi-view robust collaborative localization in high outlier ratio scenes based on semantic features," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11 042–11 047. IEEE, 2023.
- [9] Y. Tang, et al., "Ssgm: Spatial semantic graph matching for loop closure detection in indoor environments," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9163–9168. IEEE, 2023.
- [10] H. Oleynikova, et al., "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1366–1373. IEEE, 2017.
- [11] R. A. Newcombe, et al., "Kinectfusion: Real-time dense surface mapping and tracking," in *2011 10th IEEE international symposium on mixed and augmented reality*, pp. 127–136. Ieee, 2011.
- [12] Y. Deng, et al., "Macim: Multi-agent collaborative implicit mapping," *IEEE Robotics and Automation Letters*, 2024.
- [13] J. Behley, et al., "Efficient surfel-based slam using 3d laser range data in urban environments," in *Robotics: science and systems*, vol. 2018, p. 59, 2018.
- [14] L. Mescheder, et al., "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4460–4470, 2019.
- [15] B. Mildenhall, et al., "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [16] E. Sucar, et al., "imap: Implicit mapping and positioning in real-time," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6229–6238, 2021.
- [17] Z. Zhu, et al., "Nice-slam: Neural implicit scalable encoding for slam," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12 786–12 796, 2022.
- [18] H. Wang, et al., "Co-slam: Joint coordinate and sparse parametric encodings for neural real-time slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13 293–13 302, 2023.
- [19] X. Yang, et al., "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.
- [20] E. Vespa, et al., "Efficient octree-based volumetric slam supporting signed-distance and occupancy mapping," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1144–1151, 2018.
- [21] C. Jiang, et al., "H_{2}-mapping: Real-time dense mapping using hierarchical hybrid representation," *IEEE Robotics and Automation Letters*, 2023.
- [22] T. Qin, et al., "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE transactions on robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [23] J. J. Park, et al., "DeepSDF: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 165–174, 2019.
- [24] Y. Yue, et al., "Lgsdf: Continual global learning of signed distance fields aided by local updating," *arXiv preprint arXiv:2404.05187*, 2024.
- [25] S. Lionar, et al., "Neuralblox: Real-time neural representation fusion for robust volumetric mapping," in *2021 International Conference on 3D Vision (3DV)*, pp. 1279–1289. IEEE, 2021.
- [26] S. Peng, et al., "Convolutional occupancy networks," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pp. 523–540. Springer, 2020.
- [27] Y. Tang, et al., "Robust large-scale collaborative localization based on semantic submaps with extreme outliers," *IEEE/ASME Transactions on Mechatronics*, 2023.
- [28] J. Ortiz, et al., "isdf: Real-time neural signed distance fields for robot perception," *arXiv preprint arXiv:2204.02296*, 2022.
- [29] Z. Zhu, et al., "Nicer-slam: Neural implicit scene encoding for rgb slam," in *2024 International Conference on 3D Vision (3DV)*, pp. 42–52. IEEE, 2024.
- [30] W. Bian, et al., "Nope-nerf: Optimising neural radiance field with no pose prior," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4160–4169, 2023.
- [31] H. Li, et al., "Dense rgb slam with neural implicit maps," *arXiv preprint arXiv:2301.08930*, 2023.
- [32] M. M. Johari, et al., "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*, pp. 17 408–17 419, 2023.
- [33] Y. Zhang, et al., "Go-slam: Global optimization for consistent 3d instant reconstruction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3727–3737, 2023.
- [34] A. Dai, et al., "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5828–5839, 2017.
- [35] J. Straub, et al., "The replica dataset: A digital replica of indoor spaces," *arXiv preprint arXiv:1906.05797*, 2019.
- [36] W. E. Lorensen, et al., "Marching cubes: A high resolution 3d surface construction algorithm," *ACM SIGGRAPH Computer Graphics*, vol. 21, no. 4, pp. 163–169, 1987.