

H₂-Mapping: Real-time Dense Mapping Using Hierarchical Hybrid Representation

Chenxing Jiang^{2,*}, Hanwen Zhang^{3,*}, Peize Liu², Zehuan Yu², Hui Cheng³, Boyu Zhou^{1,†}, Shaojie Shen²

Abstract—Constructing a high-quality dense map in real-time is essential for robotics, AR/VR, and digital twins applications. As Neural Radiance Field (NeRF) greatly improves the mapping performance, in this paper, we propose a NeRF-based mapping method that enables higher-quality reconstruction and real-time capability even on edge computers. Specifically, we propose a novel hierarchical hybrid representation that leverages implicit multiresolution hash encoding aided by explicit octree SDF priors, describing the scene at different levels of detail. This representation allows for fast scene geometry initialization and makes scene geometry easier to learn. Besides, we present a coverage-maximizing keyframe selection strategy to address the forgetting issue and enhance mapping quality, particularly in marginal areas. To the best of our knowledge, our method is the first to achieve high-quality NeRF-based mapping on edge computers of handheld devices and quadrotors in real-time. Experiments demonstrate that our method outperforms existing NeRF-based mapping methods in geometry accuracy, texture realism, and time consumption. The code will be released at <https://github.com/SYSU-STAR/H2-Mapping>.

Index Terms—Mapping; RGB-D Perception; Visual Learning

I. INTRODUCTION

USING robots to build highly-detailed dense maps in real-time benefits advanced robot autonomous navigation, AR/VR, and digital twins applications. These maps enable robots to perform high-level tasks and provide humans with real-time feedback on the environment, allowing them to adjust the robot’s tasks promptly as needed. Besides, high-fidelity maps serve as critical assets for AR/VR and digital twins. The automatic and faithful recreation of environments in real-time using robots can be more efficient and time-saving than manual or offline reconstruction methods.

To be suitable for real-time and high-quality robot mapping in unknown environments with limited onboard computation power, a mapping system must meet four key requirements: (1) Adaptability to growing scenes, allowing the robot to dynamically expand the map without prior knowledge of the scene; (2) High level of detail; (3) Real-time capability and high memory efficiency; and (4) Novel view synthesis ability,

Manuscript received: June, 5, 2022; Accepted: August, 22, 2023. This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments.

* Equal contribution. † Corresponding Author.

¹ School of Artificial Intelligence, Sun Yat-Sen University, Zhuhai, China.

² Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong, China.

³ School of Computer Science and Engineering, Sun Yat-Sen University, Guangzhou, China

{cjangan@connect.ust.hk {zhanghw66@mail2,zhouby23@mail}.sysu.edu.cn}

Digital Object Identifier (DOI): see top of this page.



Fig. 1. We tested our methods on a handheld device (a) and quadrotors (b). Our method builds a high-quality map in real-time on edge computers and can support robotic applications.

which allows rendering high-quality images from views apart from the sparse input views. This is particularly important for creating scenes for AR/VR applications.

In robotics, mapping has been studied for decades. Previous works utilize explicit scene representations like occupancy grids [1], TSDF [2]–[7], surfels [8, 9], and meshes [10] to achieve real-time performance. However, these methods face challenges in balancing memory consumption and mapping accuracy [11] and are weak in novel view synthesis. In recent years, implicit representations have gained popularity following the introduction of NeRF [12]. Several works [13]–[15] employ NeRF to overcome limitations associated with explicit representations and achieve better mapping results in various aspects. These NeRF-based methods can produce high-fidelity reconstructions using less memory and generate high-quality images from novel views by continuously querying the scene attributes. However, the implicit representation describes the scene as high-dimensional features and neural networks that lack physical meaning, resulting in a long time for training. As a result, these methods cannot run in real-time even on the most powerful edge computers like AGX Orin (as evaluated in Sec.IV-A6).

Aiming to design a real-time and high-quality robot mapping method that fulfills the four requirements mentioned above, we propose a NeRF-based mapping method using a hierarchical hybrid representation. Our approach accelerates

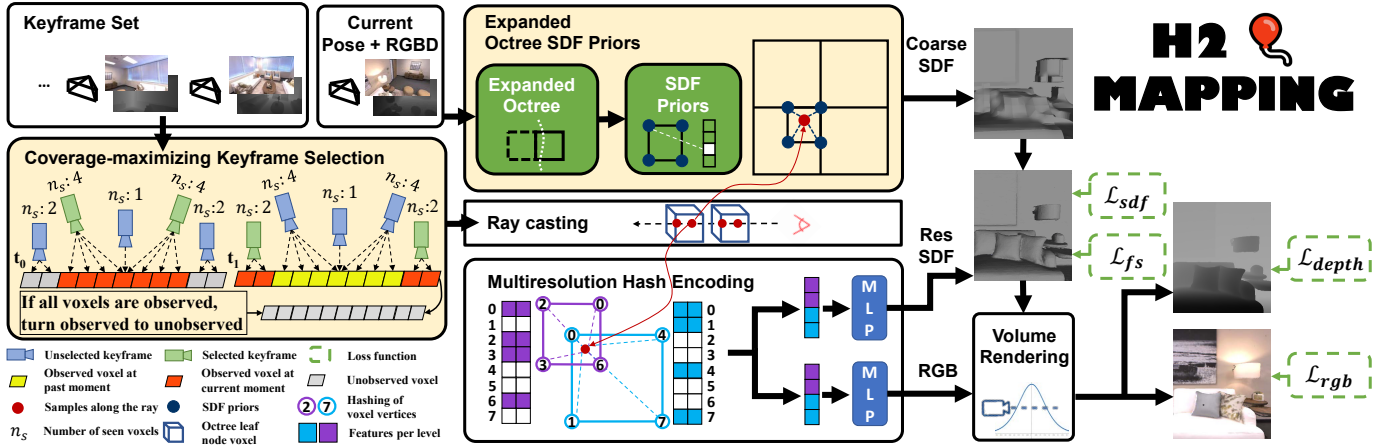


Fig. 2. The pipeline of H2-Mapping. Taking RGB-D image from sensors and pose from other tracking modules, we utilize an expanded octree SDF priors and multiresolution hash encoding to represent the scene from rough to detailed. Additionally, the proposed coverage-maximizing keyframe selection ensures quality in the edge regions.

the optimization of implicit representation with the aid of an easy-to-optimize explicit representation, describing the scene at different levels of detail. For the coarse scene geometry, we describe it with explicit octree SDF priors. Specifically, we incrementally build a sparse voxel octree with a large voxel size, where we store the optimizable SDF of each leaf node's vertex. To represent geometry details and texture, we use implicit multiresolution hash encoding [16] to encode high-resolution scene properties in a memory-efficient way. By using octree SDF priors to capture coarse geometry efficiently, the multiresolution hash encoding can focus solely on the residual geometry, which is much simpler to learn than the complete geometry, thereby improving the geometry accuracy and convergence rate.

To further speed up, we leverage a simple yet effective method to initialize the octree SDF priors. We project the voxel vertices to the depth image and calculate the associated SDF values. This initialization is based on the observation that a single measurement is usually sufficient to provide a promising estimation of coarse SDF values. Therefore, such a representation can obtain accurate geometry early on, which accelerates the optimization of texture with higher fidelity. Besides, to realize higher mapping accuracy, we propose a coverage-maximizing keyframe selection strategy to address the crucial forgetting issue in the online mapping task. Our method avoids redundant sample calculations across all keyframes [13] and ensures quality in marginal areas, without increasing the number of training samples [15].

Our method achieves faster and higher-quality NeRF-based mapping. To summarize, contributions are as follows:

- A hierarchical hybrid representation with an effective initialization technique enables real-time dense mapping with high-fidelity details and dynamical expansion ability, even on edge computers.
- An effective coverage-maximizing keyframe selection strategy that mitigates the forgetting issue and improves quality, especially in marginal areas.
- Extensive experiments show our method achieves superior mapping results with less runtime compared to existing NeRF-based mapping methods. To our knowledge,

our method is the first to run a NeRF-based mapping method onboard in real-time, as depicted in Fig.1.

II. RELATED WORKS

A. Explicit Dense Mapping

Various explicit representations have been used to store scene information for dense mapping. Octomap [1] uses probabilistic occupancy estimation to represent occupied, free, and unknown space. As a pioneer in using SDF for dense mapping, Kinect-Fusion [3] leverages volumetric SDF to enable real-time tracking and mapping. Following works improve the scalability [4, 6], the efficiency [2, 7], and the global consistency [5]. Moreover, [8] stores surfel to represent the environment, and [10] represents the robot's surrounding as a watertight 3D mesh. These methods are well known for their fast processing speed, which can be attributed to the physical meaning of explicit representations that make them easy to optimize. However, they require large amounts of memory to handle high-detailed mapping [11] and are incapable of realistically rendering from novel views.

B. Implicit Dense Mapping

Implicit representations utilize latent features and neural networks to represent a 3D scene in a high-dimensional space. DeepSDF [17] and Occupancy Networks [18] have shown the potential of implicit representation to model geometry. Recently, NeRF [12] further shows promising results in realistic novel view synthesis from sparse input views. Numerous studies [13]–[15, 19] have been inspired by NeRF [12] and utilize implicit representation for incremental dense mapping. These methods achieve more compact and accurate results than explicit representations. The NeRF-based mapping pipeline consists of two main components: (1) Scene representation; and (2) Keyframe selection strategy.

1) *Scene representation*: iMap [13] demonstrates, for the first time, that an MLP can serve as the only scene representation. To overcome the limited representation capacity of a single MLP, NICE-SLAM [14] introduces multi-resolution dense grids to store encoded features of the scene, and MLPs are used to unfold the hidden information. But the pre-allocated grids make NICE-SLAM less scalable and memory

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.




Fig. 3. Process of octree SDF priors and the reconstruction results using only the SDF priors without any optimization.

inefficient. Vox-Fusion [15], instead, only allocates voxels to the area containing the surface, forcing the network to learn more details in those regions. Nonetheless, due to the difficulty in optimizing implicit representations, it is challenging for these methods to meet real-time requirements for robotics applications. In contrast, our method utilizes a hierarchical hybrid representation for acceleration and accuracy improvement. This approach enables the implicit representation only to handle the residual geometry and texture, by taking the benefit of explicit structure. Optimizing the residual geometry is generally easier and faster.

In order to speed up, some previous works aim to accelerate geometry convergence by incorporating geometry priors. INGeo [20], for instance, scales up the initial density prediction by a factor to increase density as it approaches the surface. However, it requires manual configuration and does not provide a reasonable way to set the scaling factor. Go-surf [19] initializes its feature grid and geometry decoder to ensure that the initial SDF can represent a sphere centered at the scene origin, but this initialization process cannot adapt to map expansion and has little effect on the observed region. However, due to the hybrid representation, our method can directly initialize the explicit representation by projecting to the input depth image, which can speed up the texture optimization process with higher fidelity by providing accurate geometry in the early stage. Therefore, our method can be deployed to robots for accurate mapping in real time.

2) *Keyframe selection strategy*: iMap [13] allocates samples to every keyframe and calculates the loss distribution for selecting keyframes, which can be redundant. NICE-SLAM [14] selects optimized keyframes based on the overlap with the current frame. This strategy can keep the geometry outside the current field of view static by using a fixed, pre-trained decoder, but it cannot perform well in marginal areas that are seldom observed. Vox-Fusion [15] adds a new keyframe based on the ratio of newly allocated voxels to the currently observed voxels. All keyframes are selected to sample the same number of pixels for ray casting, leading to the increasing number of training samples over time. However, our coverage-maximizing keyframe selection strategy ensures all allocated voxels are covered with minimal iteration rounds, thereby improving the mapping quality, especially in edge regions.

III. H₂-MAPPING

In this work, we propose a real-time and high-quality mapping method, as outlined in Fig.2. Given a set of sequential poses and RGB-D frames, we utilize a hierarchical hybrid




Fig. 4. The correct sign of SDF priors (a). And different situations when using voxel expansion (c,d) or not (b) as the surface is near the voxel’s boundary.

representation (Sec.III-A) to depict the scene geometry and appearance. By employing a coverage-maximizing keyframe selection strategy (Sec.III-B), we use the volume rendering approach like NeRF [12] to obtain the depth and color of each sampled ray (Sec.III-C) and then optimize the hierarchical hybrid representation (Sec.III-D).

A. Hierarchical Hybrid Representation

To accelerate the optimization of implicit representation, we propose a hierarchical hybrid representation that explicitly stores SDF priors in an expanded octree and uses the implicit multiresolution hash encoding to only handle residual geometry and texture.

1) Expanded Octree SDF Priors:

a) *Octree SDF priors*: When a new frame is received, we allocate new voxels based on the given pose and depth image and incrementally maintain a sparse voxel octree that covers all visible areas. We only add voxels containing more than ten points to the sparse voxel octree to reduce the impact of measurement noise. The number of points is determined through experiments. For each voxel, we store the optimizable SDF in every vertex to represent the coarse geometry of the scene. The coarse SDF s^c of any sample point in a leaf node is obtained from its surrounding eight vertices through the trilinear interpolation function $TriLerp(\cdot)$:

$$s^c = TriLerp(\mathbf{p}, \{s_k^c\}), \quad k \in V, \quad (1)$$

where \mathbf{p} is the position of the sample point, s_k^c is the optimizable SDF of its surrounding vertex, and V is the set of eight vertices in the leaf node.

To accelerate the convergence rate, we provide an initial SDF to each s_k^c when allocating new voxels. As shown in the left figure of Fig.3, we project every vertex of each voxel onto the corresponding pixel in the RGB-D camera’s frame to obtain an approximate SDF at that position:

$$s_{prior}^c = \mathbf{D}(\mathbf{u}) - d_{\mathbf{p}}, \quad (2)$$

where $d_{\mathbf{p}}$ is the z-axis distance between the sensor and the vertex position \mathbf{p} , \mathbf{u} is the projected pixel, and $\mathbf{D}(\mathbf{u})$ is the depth value at the pixel \mathbf{u} . To avoid unreasonable SDF priors due to occlusion, we only provide the prior to the vertices where $(\mathbf{D}(\mathbf{u}) - d_{\mathbf{p}}) < \sqrt{6} \times (\text{VOXEL SIZE})$. $\sqrt{6}$ is chosen as the maximum distance between two adjacent voxels. The right figure in Fig.3 shows the reconstruction results using only the SDF priors without any optimization. These coarse geometry priors accelerate the geometry optimization and then enhance

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.




Fig. 5. This figure illustrates the voxels covered by the selected keyframes in three consecutive frames. The gray bounding boxes indicate the previously selected voxels, while the red bounding boxes represent the currently chosen voxels. (a) covers most of the voxels in the scene. (b) and (c) covers most of the remaining voxels, such as the toilet and poster areas on the wall.

the scene’s appearance by providing accurate geometry in the early stage, which is evaluated in Sec.IV-B1.

b) Expanded Voxels Allocation: If the surface is close to the voxel’s boundary, the accurate SDF at the position of the vertex near the surface will be close to 0. Therefore, it is possible for the SDF priors stored in that vertex to be optimized to the wrong sign, leading to the loss of the surface (Fig.11(b)). To ensure that a surface will be created, we expand a new voxel if all the points obtained from back-projecting the depth image in the voxel are located at the edge. In Fig.4, for example, the accurate SDF priors of the upper vertices should be positive but are close to 0 (Fig.4(a)). Any slight disturbance in the optimization may cause these values to become negative, resulting in no surface being reconstructed (Fig.4(b)). However, if we allocate an extra voxel on top of it, regardless of the sign to which the vertex near the surface is optimized, a surface will always be built (Fig.4(c)(d)).

2) Multiresolution Hash Encoding: In Sec.III-A1, we efficiently obtain a coarse SDF of the scene. In order to obtain the scene’s appearance and more detailed geometry, we employed a multiresolution hash encoding approach inspired by Instant-NGP [16]. Differing from the SDF implementation in Instant-NGP [16], we only utilize the multiresolution hash encoding to handle the residual SDF which is easy to learn than the complete SDF of the scene. The multiresolution hash encoding works by arranging the surrounding voxels of a particular sample point at L resolution levels. At each level, F dimensional features are assigned to the corners of the voxels by looking up a hash table. To obtain the feature of the sample point, tri-linear interpolation is performed, and the feature at each level is concatenated. We employ two multiresolution hash encoding and shallow MLP attached to individually represent the color and residual SDF of the scene in a compact manner:

$$s = s^c + \mathcal{M}_s(\phi^s; \theta_s^w), \quad \mathbf{c} = \mathcal{M}_c(\phi^c; \theta_c^w), \quad (3)$$

where ϕ^c and ϕ^s are $L \times F$ dimensional features obtained from the multiresolution hash encoding. \mathcal{M}_s and \mathcal{M}_c , parameterized by θ_s^w and θ_c^w , are MLPs to output the residual SDF and color prediction $\mathbf{c} \in \mathbb{R}^3$. s is the final SDF prediction.

B. Coverage-maximizing Keyframe Selection

For a new input RGB-D frame, we insert this frame as a new keyframe if the ratio $(N_c \cap N_l) / (N_c \cup N_l)$ is smaller than a threshold, where N_c is the number of currently observed

voxels, N_l is the number of voxels observed at the last inserted keyframe. Our keyframe insertion strategy ensures that the frames in the keyframe set have relatively little overlap.

To select the optimized keyframes from the keyframe set, we employ a coverage-maximizing keyframe selection strategy, as illustrated in Fig. 2. At the initial time step t_0 , all voxels are labeled as unobserved. We begin by selecting K keyframes that cover the largest number of voxels from the entire keyframe set. We mark these covered voxels as observed, and then optimize these selected keyframes and the current frame jointly. In the next time step t_1 , we use the same coverage-maximizing strategy but only for voxels that are still labeled as unobserved. If all voxels have been labeled as observed, we reset the voxels that were previously marked as observed to unobserved and repeat the above process. By using this strategy iteratively, all the scene areas can be covered. As shown in Fig. 5(a), most of the voxels are covered in t_0 . In Fig. 5(b) and (c), the strategy continues to cover other remaining parts of the scene, ensuring the reconstruction quality of the edge regions. In Sec.IV-B3, we further evaluate this strategy.

C. SDF-based Volume rendering

Like Vox-Fusion [15], we only sample points along the ray that intersects with any voxel. And then get rendered color \mathbf{C} and depth D for each ray as follows:

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

$$\mathbf{C} = \frac{1}{\sum_{j=0}^{N-1} w_j} \sum_{j=0}^{N-1} w_j \cdot \mathbf{c}_j, \quad D = \frac{1}{\sum_{j=0}^{N-1} w_j} \sum_{j=0}^{N-1} w_j \cdot d_j, \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid function, s_j and \mathbf{c}_j are the predicted SDF and color obtained from the hierarchical hybrid representation described in Sec. III-A, N is the number of samples along the ray, tr is a truncation distance and d_j is the sample’s depth along the ray.

D. Optimization Process

1) Loss Function: We apply loss functions like Vox-Fusion [15]: RGB Loss (\mathcal{L}_{rgb}), Depth Loss (\mathcal{L}_d), Free Space Loss (\mathcal{L}_{fs}) and SDF Loss (\mathcal{L}_{sdf}) on a batch of rays R .

$$\mathcal{L}_{fs} = \frac{1}{|R|} \sum_{r \in R} \frac{1}{P_r^{fs}} \sum_{p \in P_r^{fs}} (s_p - tr)^2, \quad \mathcal{L}_{sdf} = \frac{1}{|R|} \sum_{r \in R} \frac{1}{P_r^{tr}} \sum_{p \in P_r^{tr}} (s_p - s_p^{gt})^2$$

$$\mathcal{L}_d = \frac{1}{|R|} \sum_{r \in R} \|D_r - D_r^{gt}\|, \quad \mathcal{L}_{rgb} = \frac{1}{|R|} \sum_{r \in R} \|\mathbf{C}_r - \mathbf{C}_r^{gt}\| \quad (5)$$

where P_r^{fs} is a set of points on the ray r that lies between the camera and the truncation region of the surface measured by the depth sensor, P_r^{tr} is a set of points within the truncation area. (D_r, D_r^{gt}) and $(\mathbf{C}_r, \mathbf{C}_r^{gt})$ are rendered and input depth and color. s_p is the predicted SDF and s_p^{gt} is the difference between the distance to point p on the ray r and the depth measurement of that ray. The final loss function is defined as:

$$\mathcal{L} = \alpha_{sdf} \mathcal{L}_{sdf} + \alpha_{fs} \mathcal{L}_{fs} + \alpha_d \mathcal{L}_d + \alpha_{rgb} \mathcal{L}_{rgb}, \quad (6)$$

where α_{sdf} , α_{fs} , α_d , and α_{rgb} are the weighting coefficients.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

TABLE I

RECONSTRUCTION RESULTS OF 8 SCENES IN THE REPLICA DATASET [21]. COMPARED WITH NICE-SLAM [14] AND VOX-FUSION [15], OUR APPROACH YIELDS BETTER RESULTS IN ALL THE METRICS.

Metrics	Method	Room0	Room1	Room2	Office0	Office1	Office2	Office3	Office4	Avg.
Depth L1[<i>cm</i>] ↓	NICE-SLAM	1.19	0.93	1.28	1.02	1.35	1.07	1.30	1.23	1.16
	Vox-Fusion	0.51	0.47	0.95	0.65	0.87	0.71	0.81	0.72	0.71
	H₂-Mapping	0.34	0.22	0.61	0.33	0.45	0.53	0.50	0.40	0.42
Acc.[<i>cm</i>] ↓	NICE-SLAM	1.56	1.28	1.40	1.20	1.03	1.44	1.93	1.61	1.43
	Vox-Fusion	1.28	1.09	1.17	1.04	1.13	1.27	1.55	1.43	1.25
	H₂-Mapping	1.25	0.99	1.09	0.97	0.84	1.17	1.38	1.29	1.12
Comp.[<i>cm</i>] ↓	NICE-SLAM	1.62	1.34	1.63	1.12	1.24	1.59	2.31	1.78	1.58
	Vox-Fusion	1.32	1.10	1.25	1.09	1.19	1.26	1.54	1.45	1.27
	H₂-Mapping	1.26	0.97	1.13	0.98	0.84	1.16	1.38	1.34	1.13
Comp. Ratio[<5 <i>cm</i> %] ↑	NICE-SLAM	98.41	98.74	97.22	95.15	97.37	97.61	94.89	97.35	97.09
	Vox-Fusion	98.86	99.27	98.28	98.74	98.18	98.90	97.84	98.35	98.55
	H₂-Mapping	99.28	99.80	99.14	99.17	99.25	99.53	98.88	98.98	99.25
SSIM (Interpolate) ↑	NICE-SLAM	0.776	0.820	0.872	0.920	0.923	0.879	0.882	0.905	0.872
	Vox-Fusion	0.859	0.872	0.891	0.929	0.919	0.893	0.887	0.914	0.896
	H₂-Mapping	0.893	0.923	0.929	0.962	0.957	0.928	0.924	0.940	0.932
PSNR[<i>db</i>] (Interpolate) ↑	NICE-SLAM	27.19	28.88	30.28	34.14	34.82	29.11	29.43	30.51	30.55
	Vox-Fusion	27.25	28.56	28.97	33.54	32.50	28.01	28.35	29.33	29.56
	H₂-Mapping	31.76	33.61	32.89	38.67	38.92	32.68	33.13	34.34	34.49
SSIM (Extrapolate) ↑	NICE-SLAM	0.888	0.888	0.917	0.935	0.914	0.912	0.914	0.922	0.911
	Vox-Fusion	0.889	0.884	0.913	0.936	0.922	0.905	0.908	0.918	0.91
	H₂-Mapping	0.890	0.905	0.918	0.951	0.942	0.913	0.924	0.932	0.92
PSNR[<i>db</i>] (Extrapolate) ↑	NICE-SLAM	27.25	28.08	26.71	32.40	30.17	26.09	26.53	28.57	28.23
	Vox-Fusion	27.78	27.98	26.80	32.67	30.12	26.80	27.51	28.45	28.52
	H₂-Mapping	29.24	29.04	28.04	33.99	31.86	27.65	30.42	31.40	30.20

2) *Adaptive Early Ending*: In the training process, the current frame and selected keyframes will be used for optimization several times. As shown in Fig. 10, the average iteration time that achieves color optimization convergence varies in different scenarios. Therefore, to adaptively choose an appropriate iteration time that can balance the time consumption and mapping precision in various scenarios, we employ an early stopping policy if the total loss exceeds twice the average total loss of the current training round, indicating that further optimization can only result in a little improvement.

IV. EXPERIMENT

To evaluate the performance of our proposed method, we compare its mapping accuracy and time consumption with other NeRF-based RGB-D mapping systems on both the synthetic Replica dataset [21] and the real-world ScanNet dataset [22]. Besides, we conduct ablation studies to evaluate the effectiveness of each module in our approach. Furthermore, we deploy our method on a handheld device and quadrotors with limited computational power to test its mapping performance.

A. Mapping and Rendering Evaluation

1) *Implementation Details*: In our method, the voxel size of the octree’s leaf node is 10*cm*, $tr = 5cm$, and the maximum number of iterations is 10. For the multi-resolution hash encoding, $L = 4$, $F = 2$, $T = 2^{19}$, and the scale difference of the adjacent level is 2. Due to octree SDF priors, we only use one-layer MLP of size 64 to decode the geometry features. The appearance decoder is a two-layer MLP of size 64. 4096 pixels are selected for each iteration to generate rays and the distance between adjacent sampled points is 1*cm*. The number of keyframes that are selected to be optimized is $K = 10$.

2) *Baselines*: We select two advanced NeRF-based dense RGB-D SLAM methods currently open-source, NICE-SLAM [14] and Vox-Fusion [15] for comparison. However, since we

solely focus on incremental mapping, we remove their tracking component and instead provide the ground truth pose. All other aspects remain unchanged.

3) *Metrics*: To evaluate scene geometry, we use the Depth L1 Error[*cm*] Accuracy[*cm*], Completion[*cm*] and Completion Ratio[< 5*cm*%] of the reconstructed mesh. Besides, We use SSIM (structural similarity index measure) [23] and PSNR (Peak signal-to-noise ratio) [24] like to evaluate scene appearance on rendered images from all training views (Interpolation) and distant novel views (Extrapolation). Only the portions of the mesh included in voxels are considered, and non-depth regions are not measured in Depth L1 Error, SSIM [23], and PSNR [24].

4) *Evaluation on Replica [21]*: In Table II, we present a quantitative comparison of the reconstruction and rendering performance of our method and the baselines. The results demonstrate that our approach outperforms the baselines for both 2D and 3D metrics. Additionally, we provide a qualitative analysis of the reconstructed mesh and rendered images. Notably, in Fig. 6, the mesh obtained by our method appears smoother in areas such as the sofa and floor, while our approach exhibits enhanced geometric details, particularly for smaller objects like chair legs and vases. Fig. 7 shows that our method can generate renderings with more realistic details from both training views and novel views.

5) *Evaluation on ScanNet [22]*: Since ground truth meshes for the ScanNet [22] are unavailable, we only provide qualitative analysis on the textured mesh of Scene0000 and Scene0050 as shown in Fig. 8. Compared to the baselines, our methods can get sharper object outline, fewer artifacts, and higher-fidelity textures like the pattern on the carpet.

6) *Runtime Analysis*: We select Room0 from Replica [21] and Scene0000 from ScanNet [22] to evaluate the runtime by comparing our method with the baselines in Table III. We report the average frame processing time (FPT) on RTX 4090, RTX 2080Ti, AGX Orin, Orin NX separately. Our method is




Fig. 6. Reconstructed mesh on the Replica dataset [21]. Our method closely approximates the ground truth mesh and can capture richer geometric details in small objects, such as chair legs and vases.




Fig. 7. Image rendering results on the Replica dataset [21]. (a) Rendering results from the input views (interpolation). (b) Rendering results from the novel views (extrapolation). The image rendered by our method has more high-fidelity details and rich textures, such as pillows and bed sheets



Fig. 8. Textured mesh on the ScanNet dataset [22]. Our method produces more accurate geometry and higher-fidelity textures.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.




Fig. 9. Ablation study about octree SDF priors on one frame. We choose the first frame in Room0 of Replica [21] to evaluate how the mapping performance of the corresponding region changes with increasing optimization iterations.

TABLE II
RUNTIME ANALYSIS ON MULTIPLE DEVICES

	Method	Speed FTP(s) ↓			
		4090	2080Ti	AGX Orin	Orin NX
Replica Room0	NICE-SLAM	2.69	6.35	13.01	19.03
	Vox-Fusion	0.34	0.94	1.68	3.52
	H₂-Mapping	0.05	0.16	0.36	0.66
ScanNet Scene0000	Nice-slam	3.03	6.55	14.37	19.39
	Vox-Fusion	0.67	2.04	3.12	21.57
	H₂-Mapping	0.07	0.24	0.53	1.07

much faster than previous work on various devices.

B. Ablation Study

1) *Octree SDF Priors*: Octree SDF priors represent an easy-to-optimize explicit structure with a computationally efficient initialization process. As a result, the reconstructed geometry depicted in Fig. 9 is more accurate at the beginning of optimization, leading to a rapid PSNR [24] increase for the first frame during the early stage. Furthermore, by promptly providing a coarse geometry and utilizing the implicit multiresolution hash encoding exclusively for the residual part, the accuracy and completion metrics in Fig. 9 eventually converge to a lower level, leading to higher PSNR [24] for the same average iteration times on the entire sequence in Fig. 10. The PSNR [24] arises because accurate geometry ensures the gradient of color prediction mainly affects the surface region during backpropagation. Therefore, by enabling a faster and more accurate geometry reconstruction, this hybrid representation achieves the same reconstruction quality with fewer training iterations. This is particularly meaningful for robotic applications with limited computing power.

2) *Expanded Voxels Allocation*: Table IV shows the expanded voxel allocation technique has a greater impact on completion. Besides, Fig. 11 shows the visualization results about Office3 and Office4 of Replica [21]. In the left part of (a) and (b), holes are generated in regions where the surface is close to the voxel's boundary, making it easy for the SDF to be optimized to the wrong sign. However, as shown in the right part of (a) and (b), the expansion technique can reduce the holes caused by optimization sensitivity.

3) *Coverage-maximizing Keyframe Selection*: Table IV demonstrate our keyframe selection strategy can significantly improve the PSNR [24] and Acc. metrics. As shown in Fig. 12, our strategy can better optimize the ceiling area in Office2 of Replica [21]. Since the number of images corresponding to this area is low in the overall keyframe set, it is rare for previous methods to optimize this region using the random strategy. However, our keyframe selection strategy achieves




Fig. 10. Ablation study about the adaptive early ending and octree SDF priors on the entire data sequence. We compared changes in the PSNR [24] for different fixed iteration times in Room0 (left) and Room1 (right) of Replica [21]. The blue curve and green curve represent the results with and without octree SDF priors, respectively. The red star denotes the average iteration time and PSNR [23] using adaptive early ending.

TABLE III
ABLATION STUDY OF OUR DESIGN CHOICES

Expanded voxels allocation	Keyframe strategy	PSNR↑ [db]	Acc.↓ [cm]	Comp.↓ [cm]	Comp. Ratio↑ [< 5cm%]
×	✓	34.25	1.16	1.26	98.52
✓	×	33.65	1.28	1.20	99.12
✓	✓	34.49	1.12	1.13	99.25



Fig. 11. Ablation study on expanded voxels allocation technique. (a) and (b) show results about Office3 and Office4 of Replica [21], respectively. This technique can produce fewer holes during the mapping (Use: left; No use: right).




Fig. 12. Ablation study of coverage-maximizing keyframe selection strategy on Office2 of Replica [21]. It shows that our keyframe selection strategy will lead to better reconstruction performance in marginal areas such as the ceiling.

complete coverage of all the voxels in the keyframe set with minimal iteration rounds, greatly increasing the probability of optimizing the edge regions.

4) *Adaptive Early Ending*: The blue curve in Fig. 10 shows that increasing the number of iterations leads to higher PSNR [24] values, but the rate of improvement gradually slows down. As the number of iterations varies when using the adaptive early ending, we calculate the average iteration time and the corresponding PSNR [24], represented by the red star. The results demonstrate this strategy adaptively leads to different average iteration times that are close to convergence in various scenarios, which helps to reduce optimization time without compromising accuracy.

C. Real-World SLAM Demonstration

We demonstrate our mapping method with a tracking module, which completes a SLAM system on a handheld

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

device and quadrotors. Specifically, we employ the Realsense L515 as the vision sensor to provide RGB-D images, and a modified VINS-Mono [25] incorporating depth constraints as the tracking module to estimate the pose. The handheld device is powered by AGX Orin, and the quadrotor is equipped with Orin NX. All the programs are running onboard. Fig. 1 illustrates the results of our real-world experiments. The extracted mesh showcases high-quality reconstruction in both geometry and appearance. We use the handheld device to reconstruct an apartment (Mesh surface: $\approx 127m^2$) and use the quadrotors for mapping a part of the fight arena (Mesh surface: $\approx 58m^2$). The final mesh is extracted by marching cube [26] and all optimization was performed within the mapping procedure without any post-processing or additional training time. To the best of our knowledge, our method is the first to achieve high-quality NeRF-based mapping in real-time on edge computers. The specific runtime is evaluated in Sec.IV-A6 and more details can be found in the attached video¹.

V. CONCLUSION

We propose H₂-Mapping, a novel NeRF-based dense mapping system that utilizes hierarchical hybrid representation and can be deployed on edge computers for real-time and high-quality robot mapping. The coarse geometry is represented explicitly using octree SDF priors for fast initialization and convergence, while high-resolution geometry details and texture are encoded implicitly using multiresolution hash encoding in a memory-efficient manner. Furthermore, we propose a coverage-maximizing keyframe selection strategy to improve the reconstruction quality in marginal areas.

Baseline comparisons demonstrate that our method outperforms both mapping quality and time consumption. Besides, ablation studies show that the hierarchical hybrid representation effectively accelerates geometry and texture optimization, and the proposed keyframe selection strategy guarantees reconstruction accuracy even in edge areas. However, currently, our method cannot handle dynamic objects and long-term pose drifting, and further speed-up is required.

REFERENCES

- [1] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous robots*, vol. 34, pp. 189–206, 2013.
- [2] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1366–1373.
- [3] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *2011 10th IEEE international symposium on mixed and augmented reality*. Ieee, 2011, pp. 127–136.
- [4] T. Whelan, M. Kaess, M. Fallon, H. Johannsson, J. Leonard, and J. McDonald, "Kintinuous: Spatially extended kinectfusion," 2012.
- [5] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, "Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration," *ACM Transactions on Graphics (ToG)*, vol. 36, no. 4, p. 1, 2017.
- [6] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger, "Real-time 3d reconstruction at scale using voxel hashing," *ACM Transactions on Graphics (ToG)*, vol. 32, no. 6, pp. 1–11, 2013.
- [7] O. Kähler, V. A. Prisacariu, C. Y. Ren, X. Sun, P. Torr, and D. Murray, "Very high frame rate volumetric integration of depth images on mobile devices," *IEEE transactions on visualization and computer graphics*, vol. 21, no. 11, pp. 1241–1250, 2015.
- [8] K. Wang, F. Gao, and S. Shen, "Real-time scalable dense surfel mapping," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6919–6925.
- [9] T. Whelan, R. F. Salas-Moreno, B. Glocker, A. J. Davison, and S. Leutenegger, "Elasticfusion: Real-time dense slam and light source estimation," *The International Journal of Robotics Research*, vol. 35, no. 14, pp. 1697–1716, 2016.
- [10] F. Ruetz, E. Hernández, M. Pfeiffer, H. Oleynikova, M. Cox, T. Lowe, and P. Borges, "Ovpc mesh: 3d free-space representation for local ground vehicle navigation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8648–8654.
- [11] X. Zhong, Y. Pan, J. Behley, and C. Stachniss, "Shine-mapping: Large-scale 3d mapping using sparse hierarchical implicit neural representations," *arXiv preprint arXiv:2210.02299*, 2022.
- [12] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [13] E. Sucar, S. Liu, J. Ortiz, and A. J. Davison, "imap: Implicit mapping and positioning in real-time," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6229–6238.
- [14] Z. Zhu, S. Peng, V. Larsson, W. Xu, H. Bao, Z. Cui, M. R. Oswald, and M. Pollefeys, "Nice-slam: Neural implicit scalable encoding for slam," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 786–12 796.
- [15] X. Yang, H. Li, H. Zhai, Y. Ming, Y. Liu, and G. Zhang, "Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation," in *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2022, pp. 499–507.
- [16] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics (ToG)*, vol. 41, no. 4, pp. 1–15, 2022.
- [17] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 165–174.
- [18] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.
- [19] J. Wang, T. Bleja, and L. Agapito, "Go-surf: Neural feature grid optimization for fast, high-fidelity rgb-d surface reconstruction," *arXiv preprint arXiv:2206.14735*, 2022.
- [20] C. Li, B. Wu, A. Pumarola, P. Zhang, Y. Lin, and P. Vajda, "Ingeo: Accelerating instant neural scene reconstruction with noisy geometry priors," in *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*. Springer, 2023, pp. 686–694.
- [21] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, *et al.*, "The replica dataset: A digital replica of indoor spaces," *arXiv preprint arXiv:1906.05797*, 2019.
- [22] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 5828–5839.
- [23] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [24] A. Hore and D. Ziou, "Image quality metrics: Psnr vs. ssim," in *2010 20th international conference on pattern recognition*. IEEE, 2010, pp. 2366–2369.
- [25] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [26] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," *ACM siggraph computer graphics*, vol. 21, no. 4, pp. 163–169, 1987.

¹<https://youtu.be/oR9MifL86Vw?si=YbemRqVFVLnaZGUi>