

# RO-MAP: Real-Time Multi-Object Mapping with Neural Radiance Fields

Xiao Han, Houxuan Liu, Yunchao Ding and Lu Yang

**Abstract**—Accurate perception of objects in the environment is important for improving the scene understanding capability of SLAM systems. In robotic and augmented reality applications, object maps with semantic and metric information show attractive advantages. In this paper, we present RO-MAP, a novel multi-object mapping pipeline that does not rely on 3D priors. Given only monocular input, we use neural radiance fields to represent objects and couple them with a lightweight object SLAM based on multi-view geometry, to simultaneously localize objects and implicitly learn their dense geometry. We create separate implicit models for each detected object and train them dynamically and in parallel as new observations are added. Experiments on synthetic and real-world datasets demonstrate that our method can generate semantic object map with shape reconstruction, and be competitive with offline methods while achieving real-time performance (25Hz). The code and dataset will be available at: <https://github.com/XiaoHan-Git/RO-MAP>

**Index Terms**—Mapping, SLAM, Semantic Scene Understanding

## I. INTRODUCTION

VISION-based Simultaneous Localisation and Mapping (SLAM) is an important research problem in the field of robotics, and has achieved remarkable advances in the past decade. Previous studies [1], [2], [3] concentrated on providing accurate ego-motion estimation and reconstructing scene maps. However, the sparse or dense maps constructed by these methods only contain metric information, which limits their application in complex tasks [4], [5] that require scene understanding. The development of deep learning has paved the way for introducing semantic information into SLAM, and object SLAM that incorporates detection [6] or semantic segmentation [7] has attracted the interest of many researchers.

Different from pure geometric maps, object SLAM utilizes additional semantic observations to localize and reconstruct objects in the scene, and the generated object maps can serve downstream tasks. However, a crucial issue is how to effectively represent objects. Some research that use only RGB cameras have explored simple geometric primitives, such as cuboids [8], [9], ellipsoids [10], [11], [12], and superquadrics [13]. These compact representations contain fundamental information of objects such as category, size, and pose. They can serve as semantic landmarks for localization

Manuscript received: April 12, 2023; Revised June 27, 2023; Accepted July 24, 2023. This paper was recommended for publication by Editor Markus Vincze upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by NSFC (No. 61871074).

All authors are with the School of Automation Engineering, University of Electronic Science and Technology of China, Chengdu, China. Lu Yang is the corresponding author [yanglu@uestc.edu.cn](mailto:yanglu@uestc.edu.cn).

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



Fig. 1: RO-MAP can localize and reconstruct objects in the scene online. Each object instance is represented by a NeRF, which implicitly learns dense geometry from monocular input.

and navigation [4], and have shown advantages in relocalization [14] and long-term operation of SLAM systems. However, these geometric primitives do not capture shape and texture information of objects, which poses a challenge for monocular-based methods.

Object shape reconstruction is another widely studied problem. Some studies have explored various dense object representations using additional depth sensors, such as surfels [15] and signed distance function (SDF) [16], [17]. Furthermore, it is popular to use learnable compact shape embeddings to represent objects. Recent works [18], [19], [20], [21] have used neural networks to learn category-level shape priors and optimized the object shape codes in latent space by matching image or depth observations, which are then decoded into voxel grids [18] or implicit functions [19], [20], [21]. These methods can generate dense and complete object reconstructions from partial observations, but they are limited by the categories of pre-learned priors and cannot handle arbitrary geometric shapes. A natural question is whether we can reconstruct objects using only a monocular camera and without any geometric priors. Neural Radiance Fields (NeRF) [22] are suitable object representations. With the help of volume rendering and the powerful fitting ability of MLP, NeRF can implicitly learn 3D geometry from RGB images. Recently, its successful applications [23], [24] in SLAM have demonstrated its strong potential.

In this work, we propose an online pipeline for reconstructing multiple objects from monocular videos, which consists

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

of two loosely coupled components. The first component is a lightweight object SLAM built upon the ORB-SLAM2 [1] framework. We use instance segmentation to detect objects in the scene and estimate their size and pose, and a robust data association algorithm ensures that multi-view observations are correctly associated to objects. The second component is a multi-object reconstruction system, where each object instance is represented by a NeRF and receives new observations in real-time for incremental training. We propose an efficient loss function tailored for objects to speed up convergence and reduce depth ambiguity caused by RGB-only images. Furthermore, our CUDA implementation based on the tcnn framework [25] ensures real-time performance. The average training time per object is about 2 seconds on a single GPU. Comprehensive experiments on synthetic and collected real-world datasets demonstrate the effectiveness of our method.

The contributions of this work are as follows:

- We present, to the best of our knowledge, the first 3D prior-free monocular multi-object mapping pipeline that can localize and reconstruct objects in the scene.
- We propose an efficient loss function for objects, combined with a high performance CUDA implementation, enabling the system to have real-time performance.
- We evaluate the effectiveness of the proposed method on both synthetic and real-world datasets. In addition, the code and datasets are available.

## II. RELATED WORK

### A. Object SLAM

The earliest object SLAM can be traced back to 2013 when Salas et al. [26] first treated objects in the scene as landmarks. They matched depth observations extracted by pre-trained detectors with known object models, and continuously refined camera poses and object map through pose-graph optimization. However, the requirement of having geometric models of all object instances beforehand limits the applicability of this approach. Subsequent works along the RGB-D direction turned to online reconstruction from scratch. Fusion++ [17] leverages a 2D instance segmentation network to extract depth observations belonging to objects, and fuses them across multiple views to generate TSDF reconstructions. MaskFusion [15] further extend to dynamic scenes. Some studies lie between using object instance models and reconstructing geometric shapes from scratch, exploring learning-based object shape priors. Sucar et al. [18] used CAD models of objects from the same category to train a variational encoder. The latent code used to represent shape can generate complete object reconstruction with only partial depth observation. Similarly, ObjectFusion [21] further improves the generalization of deep shape embeddings, capable of adapting to multiple object categories using only one encoder-decoder network.

Unlike the above studies that focus on dense object reconstruction, researchers have also shown interest in constructing object maps using simple geometric primitives that only contain pose and size information. Yang et al. [8] proposed CubeSLAM, which infers 3D cuboids of objects from multi-view observations using 2D detection boxes. Compared to

learning-based methods, CubeSLAM has much less computational cost. In addition, quadrics have also been used to represent objects due to their compact perspective projection model. Nicholson et al. [10] first introduced them into object SLAM, and optimized their parameters by minimizing reprojection errors in different viewpoints. In this work, we decouple the estimation of object shape and pose. Our lightweight object SLAM utilizes cuboids to represent objects, and estimates their pose and size using multi-view observations and sparse point clouds.

### B. NeRFs and NeRF-based SLAM

Neural radiance fields have emerged as popular implicit representations in computer vision. They use a multi-layer perceptron (MLP) to represent a scene and leverage volume rendering to implicitly learn the geometry and appearance information from multi-view images. Since its introduction, NeRF has been widely used in various computer vision tasks [22], [27]. These methods demonstrate remarkable performance but require a large amount of training time. Recently, Muller et al. [28] proposed a multi-resolution hash encoding, which significantly reduces the training time of NeRF, making its online application possible.

Due to the advantage of NeRF to implicitly represent 3D scenes, several works have introduced it into SLAM. iMAP [23] is the first complete SLAM system based on NeRF, which utilizes a compact MLP-based map representation and performs simultaneous tracking and mapping in two parallel threads. NICE-SLAM [24] replaces a single MLP with a hierarchical feature grid, combines with a pre-trained decoder to achieve larger scenes and faster convergence. Compared to iMAP, it updates only the visible grid features at each step, effectively solving the forgetting problem. Subsequent works [29], [30] have made further improvements, including the integration with traditional voxel grids [31] and different shape representations [32]. In contrast to these methods that focus on dense reconstruction of the scene, our approach emphasizes object instances with semantic meaning. Each object is represented by a separate NeRF model, which are trained online and in parallel to generate dense object map. The work most similar to ours is [33], but it focuses on analyzing the effect of observation quality on reconstruction results. In addition, vMAP [34] also models objects separately while reconstructing the scene, but it uses RGB-D observations.

## III. SYSTEM OVERVIEW

Fig. 2 shows an overview of the proposed method. The pipeline consists of two main components, a lightweight object SLAM and a multi-object NeRF system. Given a monocular input stream that includes RGB images and instance segmentations, our object SLAM system simultaneously estimates camera frame poses and localizes objects in the scene. We leverage both semantic information from instance segmentation and geometric information from sparse point clouds associated with objects to perform data association and object pose and size estimation. The results and original image inputs are fed into the multi-object reconstruction system. In this

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

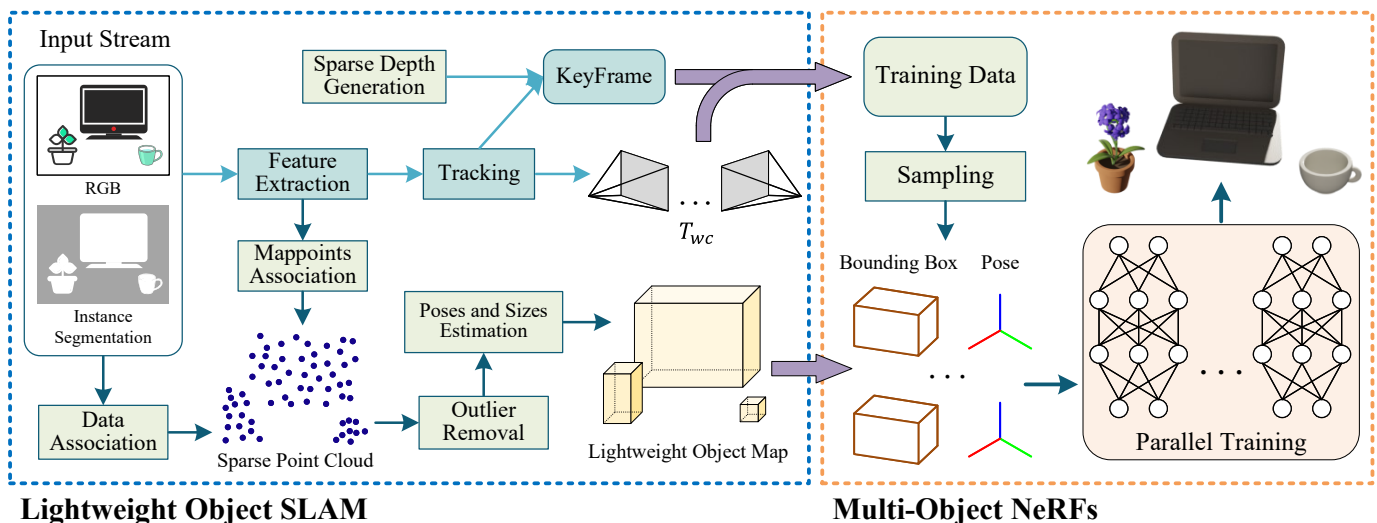


Fig. 2: Pipeline overview. RO-MAP consists of a lightweight object SLAM and a multi-object NeRF system. They correspond to the bounding box estimation (pose and size) and shape reconstruction of objects, respectively.

part, each object instance is represented by a separate NeRF model. They receive new observations in real-time and are trained in parallel. We use the marching cubes algorithm [35] to extract visual 3D meshes and transform them to the global coordinate system through object poses, thus constructing a complete dense object map.

#### IV. LIGHTWEIGHT OBJECT SLAM

Our object SLAM is implemented based on ORB-SLAM2 [1]. For monocular input, the system initially extracts ORB features and performs inter-frame matching. As we only consider static scenes, the camera pose estimation is consistent with the original ORB-SLAM2, i.e. only the traditional reprojection error is used. Simultaneously, matched image features are triangulated to generate sparse point clouds. We associate the sparse point clouds with object instance segmentation, effectively utilizing them in subsequent object data association and pose estimation. The object association strategy comprises two parts: consecutive association and non-consecutive association. The former calculates the intersection over union (IoU) of object 2D bounding boxes obtained from instance segmentation in consecutive frames. The latter utilizes a parameter statistical test based on object’s sparse point clouds to handle isolated object observations in time series and merge duplicate object landmarks. Since data association is not the focus of this paper, please refer to our previous work [13] for more details. After associating the latest object observations with landmarks, we employ a lightweight, hand-crafted method for object localization instead of learning-based methods.

##### A. Outlier Removal

After extracting image features, the feature points that are covered by the instance mask will be associated to objects. The sparse point cloud generated by triangulating these points during the tracking process is used for roughly representing

the position of objects. However, due to measurement noise and occlusion, the associated sparse point cloud often contains many outliers that do not belong to objects. We employ the Extended Isolation Forest (EIF) [36] to remove outliers and maintain a sparse point cloud that accurately fits objects. Specifically, EIF recursively partitions the sample space using a plane with a random slope, gradually reducing the number of samples in each enclosed space until each sample is isolated or the depth limit is reached. Obviously, the points that are located on the object surface after multi-view observations tend to be dense, and require more steps to be isolated. We remove those points that are isolated after very few steps, which are likely to be outliers.

##### B. Pose and Size Estimation

We represent objects using cuboids and assume that objects are always stationary and placed on a support, with roll and pitch angles fixed at zero, so that only the translation  $\mathbf{t}$  and yaw angle  $\theta$  need to be estimated. First, we directly compute the center of the filtered point cloud  $P^W$  to estimate the translation  $\mathbf{t}$  as follows:

$$\mathbf{t} = \frac{\max(P^W) + \min(P^W)}{2} \quad (1)$$

For object rotation, the simple and effective Principal Component Analysis (PCA) method is considered first. We project the 3D sparse point cloud onto the horizontal plane and then calculate its dominant orientation using PCA as the corresponding rotation matrix. However, this method performs poorly for cuboid-shaped objects such as books and keyboards, as the extracted main direction deviates significantly from the ideal orthogonal edges. This results in inaccurate object pose estimation and further affects the subsequent shape reconstruction. We combine a line feature alignment method based on object appearance to improve the robustness of rotation estimation. Specifically, we first project the three orthogonal edges  $l_i (i \in 1, 2, 3)$  of the object bounding box onto the image,

then extract line features [37] and select those that have similar slopes to the projected line segments as observations. The accumulative angle error between the extracted line segments  $l_{di}$  and the projected line segments  $l_{oi}$  is optimized to estimate yaw  $\theta$ . The optimization function is defined as follows:

$$\theta^* = \arg \min_{\theta} \sum_{i=1}^3 \|g(l_{oi}) - g(l_{di})\|^2 \quad (2)$$

$$l_{oi} = K T_{wc}^{-1} (R(\theta) l_i + \mathbf{t}), \quad i \in \{1, 2, 3\} \quad (3)$$

where  $g(\cdot)$  calculates the slope of line segment.  $T_{wc}$  represent the camera pose and  $K$  is the camera intrinsic matrix. A good initial value is crucial for this nonlinear optimization problem. We uniformly sample from  $-45^\circ$  to  $45^\circ$  with an interval of 5 degrees, and select the sample with the minimum error as the initial value for optimization. Finally, we obtain the sparse point cloud  $P^O$  transformed to the object coordinate system by the estimated object pose, and directly calculate the size  $\mathbf{a} = [a_x, a_y, a_z]^T$  as follows:

$$\mathbf{a} = \frac{\max(P^O) - \min(P^O)}{2} \quad (4)$$

## V. MULTI-OBJECT RECONSTRUCTION SYSTEM

After estimating bounding boxes and camera poses in object SLAM, we use NeRF to implicitly learn the dense geometry of objects. When a new object instance is detected, we initialize a new NeRF model, which consists of a multi-resolution hash encoding [28] and a single-layer MLP. Unlike some methods that reconstruct the whole scene with NeRFs, our model only needs to represent a single object, which allows us to use tiny network structures and accelerate training speed considerably. Moreover, we leverage multi-threading to train models in parallel, thus further improving system efficiency.

### A. Training

1) *Data*: Since there are small viewpoint changes between adjacent frames in SLAM, using all images for training would introduce a lot of redundant information. We only use those images that are chosen as keyframes in the tracking process. Besides the original RGB images and instance masks, we also reproject the sparse point clouds associated with objects onto the images, and the resulting sparse depth maps can serve as additional supervision during training. This facilitates the model to learn accurate geometry. For all object instances, they have different training data and different appearance times. We implement an incremental update method for the training data to handle each model separately. As shown in Fig. 3, assuming that the last updated image is  $I_m$  and the currently observed object image is  $I_n$ , we calculate their relative rotation angles with respect to the object as follows:

$$\alpha = \arccos \left( \frac{(\mathbf{t}_{I_m} - \mathbf{t}_O) \cdot (\mathbf{t}_{I_n} - \mathbf{t}_O)}{\|\mathbf{t}_{I_m} - \mathbf{t}_O\| \|\mathbf{t}_{I_n} - \mathbf{t}_O\|} \right) \quad (5)$$

If  $\alpha$  is larger than the preset threshold, then update the training data. As the viewpoints increase, the number of training iterations gradually increases to converge quickly.




Fig. 3: (a) The training data will be updated if the rotation angle between the current image  $I_n$  and the last updated image  $I_m$  is larger than  $\alpha$ . (b) Thanks to the zero density learned by background rays in image  $I_j$ , the rays pointing to the object in image  $I_k$  can quickly focus on the object surface.

2) *Parallel Training*: We adopt a thread pool approach to enable parallel training of multi-object models. Each worker thread in the thread pool has its own CUDA stream and asynchronously fetches training tasks from the work queue of the object SLAM system. When an object no longer receives new observations, the model will stop training to make more efficient use of computational resources. We configure the pool with 8 worker threads, which is sufficient for most scenes.

### B. Volume Rendering

Differentiable volume rendering is used to optimize the implicit representation of objects. We first transform the camera pose to the object coordinate system and back-project the pixels that are within the object's detection box. If the ray intersects with the 3D bounding box, we compute the truncation distance and sample  $N$  points within it. Unlike other RGB-only implicit reconstruction methods, we only perform uniform sampling and do not include the popular importance sampling. This can save the time cost of one inference process of the model, although it slightly compromises the reconstruction quality.

Since we are more interested in dense reconstruction than novel view synthesis, only the positions of the sampled points are encoded and fed into the network to estimate their density values  $\sigma$  and colors  $c$ , without including the ray direction. For a point  $\mathbf{x}_i$ , there are its occupancy probability  $o_i = 1 - \exp(-\sigma_i \delta_i)$  and the probability  $w_i = o_i \prod_{j=1}^{i-1} (1 - o_j)$  that the ray terminates at this point, where  $\delta_i = d_{i+1} - d_i$  is the difference in distance between adjacent sample points. Finally, the predicted color and depth of the corresponding ray  $\mathbf{r}$  are defined as follows:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N w_i c_i, \quad \hat{D}(\mathbf{r}) = \sum_{i=1}^N w_i d_i \quad (6)$$

Object reconstruction requires suppressing the background and occluders since objects are not isolated but embedded in scenes. Instance masks contain valid spatial semantic information that is used to guide learning the geometry distribution of

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

objects and their surroundings. We follow the strategy in [33] and categorize the optimization loss of rays. Specifically, we classify the sampled rays into three types according to their instance masks. For rays  $\mathcal{R}_o$  that hit the reconstructed object, i.e., their mask value  $m$  matches with object instance  $M_o$ , we calculate their photometric loss as usual:

$$L_{rgb} = \sum_{i \in \mathcal{R}_o} \left\| \hat{C}(\mathbf{r}_i) - C(\mathbf{r}_i) \right\|_2 \quad (7)$$

For some rays  $\mathcal{R}_d$  that have depth supervision, additional depth loss is also included:

$$L_{depth} = \sum_{i \in \mathcal{R}_d} \left| \hat{D}(\mathbf{r}_i) - D(\mathbf{r}_i) \right| \quad (8)$$

Second, we expect the space outside the objects to be empty, i.e., the rays pointing to them should not terminate. To achieve this, we give those rays  $\mathcal{R}_b$  corresponding to the background varying random colors as supervision to guide them to learn zero density:

$$L_{rr} = \sum_{i \in \mathcal{R}_b} \left\| \hat{C}(\mathbf{r}_i) - C_{random} \right\|_2 \quad (9)$$

However, the convergence of volume density caused by this loss is slow and requires many rounds of optimization. We propose an efficient and aggressive loss that skips volume rendering and directly optimizes the density of sampling points as follows:

$$L_{density} = \sum_{i \in \mathcal{R}_b} \sum_{j=1}^N \left| \sigma_j^{\mathbf{r}_i} - 0 \right| \quad (10)$$

Although it looks inelegant, it quickly improves the convergence speed and helps reduce depth ambiguity caused by only monocular images. Fig. 3 demonstrates an example where rays pointing to the object in image  $I_k$  can rapidly focus on optimizing near the object due to zero density learned by background rays in image  $I_j$ . Finally, for rays that hit other occluding objects, we do not construct optimization loss because we cannot specify spatial information along their paths. Overall, the total loss for object instance is defined as:

$$L = L_{rgb} + L_{rr} + \lambda_1 L_{depth} + \lambda_2 L_{density} \quad (11)$$

where  $\lambda_1$  and  $\lambda_2$  are loss weights.

## VI. EXPERIMENTS

We evaluate the proposed pipeline on synthetic and real-world datasets. Due to the low requirement of object localization for NeRF training, i.e., the estimated bounding box only needs to loosely enclose objects, we focus on evaluating shape reconstruction. We also provide detailed runtime analysis and two ablation studies that supports our design choices. Considering the online nature of our method, please see the attached video demonstration.

TABLE I: Quantitative evaluation of object reconstruction. **Bold** and underline indicate the best and the second-best respectively.

	Acc. [cm]↓	Comp. [cm]↓	Comp.Ratio [<0.4cm %]↑	Comp.Ratio [<1cm %]↑
[33]* w/ GT depth	<b>0.259</b>	<b>0.162</b>	<b>93.69</b>	<b>99.98</b>
COLMAP	0.612	0.656	57.72	81.17
[33]*	0.476	<u>0.228</u>	83.16	99.38
<b>Ours</b>	<u>0.431</u>	0.248	80.85	98.93




Fig. 4: Qualitative results on the Cube-Diorama *room* sequence. Note that all the comparison methods run offline.


1) *Implementation Details*: Our pipeline is implemented using C++ and CUDA, and all experiments are conducted on a desktop computer with a 3.0GHz Intel Xeon 6154 CPU and an NVIDIA RTX 4090 GPU. All object instances use the same NeRF model parameters, including hash table size  $T = 2^{16}$ , finest resolution  $N_{max} = 2048$ , and hidden size 64 of the single-layer MLP. For volume rendering, we trigger 300 iterations of training each time new observations are received. Each iteration randomly samples 4096 rays from all training images, with  $N = 32$  sampling points per ray. In addition, we set the loss weights  $\lambda_1 = 0.5$ ,  $\lambda_2 = 0.01$  and the training data update threshold  $\alpha = 25^\circ$ . Marching Cubes [35] are used to extract meshes online, with the same resolution of  $64^3$  for all objects.

2) *Baselines*: We compare to the classical COLMAP [38] and an implicit object reconstruction method [33] also based on NeRF. We faithfully re-implement the latter, denoted as [33]\*. Since both of these comparison methods operate offline and do not consider object localization, we use the training data obtained from running our method online as their input to ensure a fair comparison. Additionally, we also compare our method with two other NeRF-based approaches, namely iMAP [23] and vMAP [34].

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

TABLE II: Quantitative evaluation of object reconstruction.

		Acc. [cm]↓	Comp. [cm]↓	Comp.Ratio [<1cm %]↑	Comp.Ratio [<5cm %]↑
iMAP	room-0	3.02	1.71	52.57	93.72
	office-1	2.62	2.58	48.93	91.09
vMAP	room-0	<b>2.18</b>	1.13	<b>74.09</b>	96.68
	office-1	<b>2.27</b>	1.77	65.24	92.94
<b>Ours</b>	room-0	3.65	<b>0.93</b>	69.25	<b>98.53</b>
	office-1	3.74	<b>1.15</b>	<b>67.93</b>	<b>97.73</b>

Fig. 5: Qualitative results on the Replica *room-0* sequence.

3) *Datasets and Metrics*: We first evaluate on the synthetic Cube-Diorama dataset [33] and Replica dataset [39], which provide ground-truth depths and instance segmentations. *Accuracy*, *Completion* and *Completion Ratio* are used for quantitative evaluation of object reconstruction. On the Cube-Diorama dataset, due to the scale ambiguity of monocular SLAM systems, we use ICP [40] to align the reconstructed meshes and the GT meshes. Subsequently, we qualitatively evaluate on a collected real-world dataset.

#### A. Object Reconstruction Evaluation

1) *Cube-Diorama*: We first evaluate the quality of object reconstruction on the *room* sequence, as shown on the left of Fig. 1, which captures four different-shaped objects on a desktop. We use ground-truth instance masks to test the upper bound of system performance and additionally show the results of [33] trained with ground-truth depth for reference. Table I presents the quantitative results. Benefiting from the powerful capabilities of implicit representation and volume rendering, our method significantly outperforms traditional COLMAP, which is difficult to handle black or textureless object surfaces. Since both our method and [33] utilize the same network




Fig. 6: Reconstruction results of all objects on the real-world sequence. Instance masks are produced by YOLOv8.

models from [25], they demonstrate comparable performance. However, [33] solely implements offline reconstruction, which requires estimating the poses of all frames and manually labeling object bounding boxes in advance. In contrast, our complete pipeline, integrating object SLAM for online operation, object pose estimation, and parallel training, achieves higher completeness and practicality. As expected, the use of additional ground-truth depth in [33] makes it easier to capture geometric information, resulting in better results. Fig. 4 shows the visualization results of all objects. Our method can generate watertight object meshes. It is worth noting that, due to the reflection and the noise in the estimated camera pose, all RGB-only methods suffer from artifacts. This remains a challenge for monocular reconstruction without 3D priors.

2) *Replica*: Since our heuristic object pose estimation method based on monocular input is difficult to handle large objects such as beds and dining tables, we restrict the evaluation of RO-MAP to the *room-0* and *office-1* sequences that mainly contain smaller objects. The number of objects in these sequences is 13 and 29, respectively. For dense reconstruction, we use RGB-D input similar to the comparison methods. The quantitative results are presented in Table II. Our method achieves higher reconstruction completion. Due to our straightforward uniform sampling strategy, the sampling number near the surface of larger objects is insufficient, which leads to some artifacts inside the generated object surfaces and consequently impacting the accuracy metrics of our method. Fig 5 illustrates the visualization results, RO-MAP has better reconstruction fineness and can construct multi-object maps with semantic information. Compared to vMAP, our method encounters challenges in dealing with heavily occluded areas, which will be a focus of future research.

3) *Real-world Sequence*: We evaluate on a real scene collected with a Realsense D455 camera, as shown on the right of Fig. 1. Noisy object masks are provided by YOLOv8, which has real-time performance. Fig. 6 presents the qualitative results. We can see that COLMAP fails to reconstruct the non-Lambertian laptop screen, resulting in large holes. Compared with [33], our method generates more complete object reconstruction and has better visual quality.



Fig. 7: (Top) The scene of the challenging real-world sequence. (Bottom) Dense object map and constrained camera motion trajectory.

TABLE III: Runtime of different system components.

Component	Tasks	Runtime(mSec)	
Object SLAM	Frontend Tracking	41.6	
	Backend optimization	204	
NeRF Model	Iteration (avg/object)	Ray Sampling	0.0932
		Forward and Volume Rendering	0.156
		Backward and Optimization	0.455
	Marching cubes	2.84	

4) *Challenging Real-world Sequence*: In practical applications of robotics or AR, it is often impossible to obtain observations from all viewpoints of objects. We provide a challenging real-world sequence, which contains eight objects of different shapes, and the camera only gives limited viewpoint observations along a constrained motion trajectory. Fig. 7 shows the scene and the qualitative results. It can be observed that the generated object meshes are well separated from the background in the observed viewpoints. However, for the unobserved regions of the objects, although the interpolation-based multi-resolution feature grid has some predictive ability, it still cannot produce satisfactory results. Reconstructing objects from partial observations [41] is an interesting direction for future work.

### B. Runtime Analysis

Table III shows the detailed breakdown of average computation time for each main component. For object SLAM, our hand-crafted object pose estimation method only introduces a small amount of time consumption to the original ORB-SLAM2. For a single NeRF model, our parallel implementation requires only 0.7ms for one iteration of training. The number of iterations for training different object instances depends on the size of the observed view angle. On average, each object takes about 2 seconds. Compared to the whole scene reconstruction, representing and optimizing a single object allows us to use tiny networks and simple sampling strategies, which help reduce the branch divergence problem in parallel computing and improve the training speed. When the scene contains a large number of objects, resulting in a queuing situation in the NeRF thread pool, our system’s unidirectional

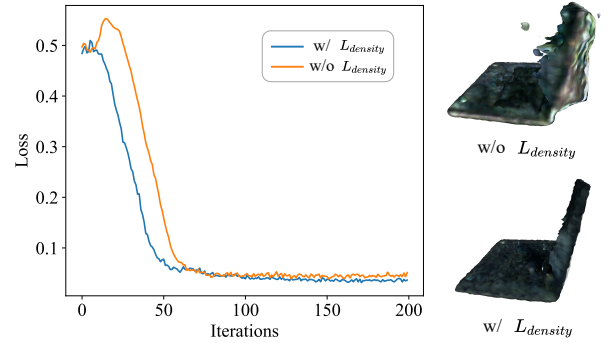


Fig. 8: Ablation Study. (Left) Convergence plot of training with all data on the *room* sequence. (Right) Intuitive results of the laptop from the challenging sequence.

TABLE IV: Ablation study of object model sizes.

Hash table size	14	16	18	20	16	16
MLP layers	1	1	1	1	2	3
Acc. [cm]↓	0.435	0.424	0.451	0.450	0.409	0.442
Comp. [cm]↓	0.227	0.233	0.236	0.243	0.241	0.243
Iteration time [ms]	0.658	0.706	0.921	1.495	0.862	1.052

data flow ensures that it does not block the SLAM tracking process. It only slightly increases the training time per iteration to approximately 0.83ms.

### C. Ablation Study

1) *Losses*: The depth ambiguity caused by monocular input has a significant impact on learning the geometry of textureless or smooth objects, such as leading to slow convergence and scattered artifacts. The convergence plot in Fig. 8 shows the comparison results on the *room* sequence. It can be observed that the training guided only by the random color loss is difficult to converge in the early stage, which indicates that directly regressing the voxel density of empty regions is helpful for the disambiguation of the optimization process. The right side of Fig. 8 shows an intuitive example.

2) *Object Model*: We investigate the impact of different object models on the reconstruction quality, including the size of the hash encoding table and the number of layers in MLP. In Table IV, we can see that using larger models did not improve the reconstruction quality. This can be attributed to the limited number of iterations caused by online processing, which is insufficient to adequately train model parameters. In contrast, models with fewer parameters are already capable of representing individual objects and have faster training speed.

## VII. CONCLUSIONS

We present RO-MAP, a real-time multi-object mapping pipeline that only uses monocular input and does not rely on 3D priors. The method employs neural radiance fields as implicit shape representations, and combines them with lightweight object SLAM, to localize and reconstruct objects in a scene and generate dense object map with semantic information. Our high-performance implementation allows

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

creating separate implicit models for each object, which can be incrementally trained and converge quickly. Comprehensive experiments demonstrate the effectiveness and advantages of RO-MAP. In the future, we are interested in how to utilize implicit object maps for downstream tasks such as robot navigation, grasping, and relocalization.

### VIII. ACKNOWLEDGMENTS

We thank Jad Abou-Chakra for his support on the dataset and Quei-An Chen for helpful discussions.

### REFERENCES

- [1] R. Mur-Artal and J. D. Tardós, “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.
- [2] J. Engel, T. Schöps, and D. Cremers, “Lsd-slam: Large-scale direct monocular slam,” in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part II 13*. Springer, 2014, pp. 834–849.
- [3] A. Pumarola, A. Vakhitov, A. Agudo, A. Sanfeliu, and F. Moreno-Noguer, “PI-slam: Real-time monocular visual slam with points and lines,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 4503–4508.
- [4] R. Martins, D. Bersan, M. F. Campos, and E. R. Nascimento, “Extending maps with semantic and contextual object information for robot navigation: a learning-based framework using visual and depth cues,” *Journal of Intelligent & Robotic Systems*, vol. 99, pp. 555–569, 2020.
- [5] Y. Wu, Y. Zhang, D. Zhu, X. Chen, S. Coleman, W. Sun, X. Hu, and Z. Deng, “Object slam-based active mapping and robotic grasping,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 1372–1381.
- [6] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [7] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [8] S. Yang and S. Scherer, “Cubeslam: Monocular 3-d object slam,” *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 925–938, 2019.
- [9] Y. Wu, Y. Zhang, D. Zhu, Y. Feng, S. Coleman, and D. Kerr, “Eao-slam: Monocular semi-dense object slam based on ensemble data association,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4966–4973.
- [10] L. Nicholson, M. Milford, and N. Sünderhauf, “Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam,” *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 1–8, 2018.
- [11] R. Tian, Y. Zhang, Y. Feng, L. Yang, Z. Cao, S. Coleman, and D. Kerr, “Accurate and robust object slam with 3d quadric landmark reconstruction in outdoors,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1534–1541, 2022.
- [12] Z. Liao, Y. Hu, J. Zhang, X. Qi, X. Zhang, and W. Wang, “So-slam: Semantic object slam with scale proportional and symmetrical texture constraints,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4008–4015, 2022.
- [13] X. Han and L. Yang, “Sq-slam: Monocular semantic slam based on superquadric object representation,” *arXiv preprint arXiv:2209.10817*, 2022.
- [14] Z. Qian, J. Fu, and J. Xiao, “Towards accurate loop closure detection in semantic slam with 3d semantic covisibility graphs,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2455–2462, 2022.
- [15] M. Runz, M. Buffier, and L. Agapito, “Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects,” in *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2018, pp. 10–20.
- [16] B. Xu, W. Li, D. Tzoumanikas, M. Bloesch, A. Davison, and S. Leutenegger, “Mid-fusion: Octree-based object-level multi-instance dynamic slam,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5231–5237.
- [17] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, “Fusion++: Volumetric object-level slam,” in *2018 international conference on 3D vision (3DV)*. IEEE, 2018, pp. 32–41.
- [18] E. Sucar, K. Wada, and A. Davison, “Nodeslam: Neural object descriptors for multi-view shape reconstruction,” in *2020 International Conference on 3D Vision (3DV)*. IEEE, 2020, pp. 949–958.
- [19] M. Shan, Q. Feng, Y.-Y. Jau, and N. Atanasov, “Ellipsdf: joint object pose and shape optimization with a bi-level ellipsoid and signed distance function description,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5946–5955.
- [20] J. Wang, M. Rünz, and L. Agapito, “Dsp-slam: object oriented slam with deep shape priors,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 1362–1371.
- [21] Z.-X. Zou, S.-S. Huang, T.-J. Mu, and Y.-P. Wang, “Objectfusion: Accurate object-level slam with neural object priors,” *Graphical Models*, vol. 123, p. 101165, 2022.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] T. Müller, “tiny-cuda-nn,” 4 2021. [Online]. Available: <https://github.com/NVlabs/tiny-cuda-nn>
- [26] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. Kelly, and A. J. Davison, “Slam++: Simultaneous localisation and mapping at the level of objects,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 1352–1359.
- [27] A. Kundu, K. Genova, X. Yin, A. Fathi, C. Pantofaru, L. J. Guibas, A. Tagliasacchi, F. Dellaert, and T. Funkhouser, “Panoptic neural fields: A semantic object-aware neural scene representation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 12 871–12 881.
- [28] 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.
- [29] A. Rosinol, J. J. Leonard, and L. Carlone, “Nerf-slam: Real-time dense monocular slam with neural radiance fields,” *arXiv preprint arXiv:2210.13641*, 2022.
- [30] Z. Zhu, S. Peng, V. Larsson, Z. Cui, M. R. Oswald, A. Geiger, and M. Pollefeys, “Nicer-slam: Neural implicit scene encoding for rgb slam,” *arXiv preprint arXiv:2302.03594*, 2023.
- [31] 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.
- [32] Y. Ming, W. Ye, and A. Calway, “idf-slam: End-to-end rgb-d slam with neural implicit mapping and deep feature tracking,” *arXiv preprint arXiv:2209.07919*, 2022.
- [33] J. Abou-Chakra, F. Dayoub, and N. Sünderhauf, “Implicit object mapping with noisy data,” *arXiv preprint arXiv:2204.10516*, 2022.
- [34] X. Kong, S. Liu, M. Taher, and A. J. Davison, “vmap: Vectorised object mapping for neural field slam,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 952–961.
- [35] 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.
- [36] S. Harii, M. C. Kind, and R. J. Brunner, “Extended isolation forest,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 4, pp. 1479–1489, 2019.
- [37] C. Akinlar and C. Topal, “Edlines: A real-time line segment detector with a false detection control,” *Pattern Recognition Letters*, vol. 32, no. 13, pp. 1633–1642, 2011.
- [38] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] 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.
- [40] D. Girardeau-Montaut, “Cloudcompare,” *France: EDF R&D Telecom ParisTech*, vol. 11, 2016.
- [41] M. Lee, K. Kang, and H. Yu, “Just flip: Flipped observation generation and optimization for neural radiance fields to cover unobserved view,” *arXiv preprint arXiv:2303.06335*, 2023.