

RGB-D Mapping and Tracking in a Plenoxel Radiance Field

Andreas L. Teigen^{1*} Yeonsoo Park^{2*} Annette Stahl¹ Rudolf Mester¹

¹Norwegian University of Science and Technology (NTNU), Trondheim, Norway

²Mobiltech, Seoul, Republic of Korea

andreas.l.teigen@ntnu.no

yspark@mobiltech.io

Abstract

The widespread adoption of Neural Radiance Fields (NeRFs) have ensured significant advances in the domain of novel view synthesis in recent years. These models capture a volumetric radiance field of a scene, creating highly convincing, dense, photorealistic models through the use of simple, differentiable rendering equations. Despite their popularity, these algorithms suffer from severe ambiguities in visual data inherent to the RGB sensor, which means that although images generated with view synthesis can visually appear very believable, the underlying 3D model will often be wrong. This considerably limits the usefulness of these models in practical applications like Robotics and Extended Reality (XR), where an accurate dense 3D reconstruction otherwise would be of significant value. In this paper, we present the vital differences between view synthesis models and 3D reconstruction models. We also comment on why a depth sensor is essential for modeling accurate geometry in general outward-facing scenes using the current paradigm of novel view synthesis methods. Focusing on the structure-from-motion task, we practically demonstrate this need by extending the Plenoxel radiance field model: Presenting an analytical differential approach for dense mapping and tracking with radiance fields based on RGB-D data without a neural network. Our method achieves state-of-the-art results in both mapping and tracking tasks, while also being faster than competing neural network-based approaches. The code is available at: https://github.com/ysus33/RGB-D_Plenoxel_Mapping_Tracking.git.

1. Introduction

In the computer vision field, a dense map can be defined as a continuous 3D surface map generated using all observed pixels in an image set. The RGB-D sensor is a

* Authors contributed equally to this work.




Figure 1. Visualization of the generated map and estimated trajectory on Office-3 of Replica dataset. The figure illustrates the voxel grid radiance field of the map.

popular sensor choice for creating dense maps due to its rich geometric and photometric information. Dense maps are very useful for many tasks, such as path planning and collision avoidance in robotics, the interaction between real-world geometry and digital objects in *extended reality* (XR), and simply as maps for human inspection. However, dealing with the prevalence of noise and missing measurements in the sensor data and the sheer amount of data collected can be challenging. In many cases, efficient processing of such data may require the utilization of a GPU. Dense maps can also be created from only RGB images, but this requires some assumption regarding areas of the scene that contain no gradients. Although dense mapping generally requires more effort than sparse, feature-based mapping, they

contain much more information, and tracking the camera motion in a known environment with a good dense map is highly accurate as all image and model information can be used. This is opposed to sparse tracking, which only uses a subset of the available information, usually in the form of easily recognizable points in the images/model. Therefore, dense tracking is positioned to result in a smoother and more accurate trajectory.

The learned radiance field algorithms, such as *Neural Radiance Fields* (NeRF) [20], were introduced as a mathematically simple formulation for creating photorealistic dense models using RGB images. This works very well for some types of scenes, especially object-centric and inward-facing scenes and scenes with significant gradient coverage. These kinds of scenes significantly reduce the need to make assumptions in the modeling stage, resulting in highly accurate photogrammetry models. However, if you apply these algorithms on outward-facing scenes and/or scenes with sparse gradient coverage, they will be incapable of capturing the geometry, and the view synthesis will not be valid outside of the camera pose distribution that was used for the training of the model. Imbuing the NeRF optimization scheme to allow for full use of the RGB-D sensor removes ambiguities in gradient-less regions of the scene and makes it more useful for practical applications.

We base our work on the Plenoxel algorithm [11]: The analytical radiance field representation that does not use a neural network but rather a voxel grid representation. Building upon this representation, we make it more applicable for practical tasks by augmenting the algorithm to incorporate depth data. We then implement a pose optimization algorithm to track a camera throughout the scene, using volumetric, dense image-to-model alignment based on the radiance field rendering equations. We show the analytical derivations of all equations used for both optimization strategies and implement them in CUDA for fast computation times. Leveraging the inherent speed of the Plenoxel model of more than two magnitudes faster than the original NeRF algorithm, we create a very efficient mapping algorithm and a real-time, highly accurate tracking algorithm.

Our contributions are summed up as follows:

- Present a discussion on the differences between models for novel view synthesis and models for 3D reconstruction and why this might lead to problems directly attempting to use NeRF for many practical applications like robotics and XR.
- Derive the analytical derivative equations for mapping and tracking in a voxel-based radiance field based on RGB-D data for efficient optimization in CUDA.
- Showing improvement in both mapping and tracking results compared to existing radiance field mapping and tracking methods given the same time constraints.

2. Related Work

2.1. Dense Mapping and Tracking

Despite the usefulness of dense visual mapping and the accuracy of dense tracking [10], dense mapping and tracking methods [15, 24] have received relatively little attention compared to their sparse counterparts [9, 22]. This is mostly due to their technical and computational complexity as well as their reliance on either a depth sensor or an assumption for image regions with no gradients. Despite this, there have still been several noteworthy papers on the topic in the past few years: The DTAM algorithm [24] by Newcombe et al. first proposed the idea of performing dense *simultaneous localization and mapping* (SLAM) by separating the problem into alternating the tasks of updating a dense 3D model and tracking the camera pose by aligning the camera image to the model using randomly sampled image pixels. Kinectfusion [15] built on the premise, but distinguished itself by relying only on a depth-only sensor and representing the entire model as a truncated signed distance field, using the iterative closest point algorithm for pose optimization. A more recent RGB-D SLAM paper: Bad-SLAM [28] creates a dense map based on surfels instead of individual pixels. The surfels' position, orientation, and size are optimized with a clever bundle adjustment implementation.

The popularity of dense mapping and tracking has increased through the use of deep learning-based methods [3, 33], often by the use of pixel-level depth estimators. CodeSLAM [3] does this by training a variational auto-encoder offline and estimating pixel depth based on RGB images. DROID-SLAM [33] is a surprisingly robust method compared to other pre-trained deep learning-based SLAM algorithms, providing good results on several different datasets, even for some datasets not included in the training set.

2.2. Radiance Fields

Neural Radiance Fields (NeRF) [20] is a new technology that has taken the research community by storm. It allows for the creation of photo-realistic, dense, volumetric 3D view synthesis models of real-world objects, only requiring posed RGB images of the scene. NeRF natively stores a model in a highly compressed format as a neural network. The model is trained using ray-based volumetric rendering functions for training on images using multi-view consistency to produce a globally consistent model. Ever since NeRF was proposed, there has been a lot of research to expand its applicability:

Speed and efficiency. Despite the original NeRF's remarkable capabilities, early versions were hindered by limitations of slow convergence rates. Numerous strategies [2, 25, 38] have been proposed to improve its efficiency

for both training and rendering. Several papers [11, 21] also show that it is possible to improve the speed by several orders of magnitude by changing the model representation. Muller et al. [21] do this by utilizing a hash encoding that enables disambiguation of hash collisions and consequently allows a smaller neural network to represent a larger scene. Fridovich et al. [11] completely discard the neural network, favoring a voxel grid where each vertex is modeled by 28 float values, with samples at arbitrary locations are retrieved by linear interpolation of the eight closest voxels. Direct data access greatly improves the efficiency at the cost of increased memory consumption.

Depth guided NeRF. One direction of NeRF research has been to incorporate depth sensors to aid convergence speed [6] and geometric accuracy [7] of the NeRF model. The optimization based on the depth data uses the same underlying rendering equations as the color-only optimization and integrates seamlessly with the base model.

Pose optimization. Wang et al. [34] have demonstrated the possibility of optimizing intrinsic camera parameters along with the neural radiance fields. This was followed by Yen et al. [36], who proposed camera pose optimization based on a trained NeRF model, using image-to-model alignment. Lin et al. [19] went further and showed that given a coarse initialization of the camera poses, the pose optimization for all training images could be done simultaneously during the training of the model. However, to improve both the radius of convergence and performance, gradually, more layers of frequency encoding had to be introduced during the optimization process, which required human supervision due to scene-specific variations.

Mapping and tracking. Due to NeRF’s simple formulation of dense mapping and its small storage size, several authors have attempted to use NeRF as a map representation in dense SLAM algorithms [31, 35, 40]. The first to attempt this for real-time processing was Sucar et al. [31], who used the original NeRF model [20] and an RGB-D sensor both for speed and to solve the geometric ambiguity problem covered in Sec. 3. Although NeRFs are slow to train, the rapid convergence speed in the earlier epochs is a great help in achieving passable real-time performance. They showed promising results in a room-scale scene, but an increase in scene size would lead to catastrophic forgetting [12]. Zhu et al. [40] addressed the weakness of catastrophic forgetting by storing values from the learned model in a grid, which then would be queried by pre-trained decoder networks. Their follow-up work in [39] continued this trend of offloading tasks to pre-trained networks, replacing the depth sensor with two convolutional networks predicting a depth map and a normal map, respectively. Vox-Fusion [35]




Figure 2. Color and depth rendering from two different radiance fields trained on office-2 in the Replica dataset: One trained with an RGB-D sensor(left column) and the other one trained with only an RGB sensor(right column), both trained on a large set of images viewing the scene from several different angles.

	RGB PSNR \uparrow	Depth diff. (m/pixel) \downarrow
RGB	30.612	0.6971
RGB-D	28.570	0.0090

Table 1. Comparison of radiance field color and geometry accuracy, including RGB and RGB-D sensor for training on Office-2 sequence of Replica dataset.

adopts a similar concept to [31] but achieves significantly reduced memory consumption by dynamically allocating sparse voxels based on an octree structure.

NeRF as visual SLAM backend. Notable but less relevant papers also use NeRF models as a visually pleasing backend on existing sparse SLAM algorithms [4, 26]. Several easy-to-use NeRF libraries [21, 32] use such an approach, using methods such as [27] as a separate front-end to calculate camera poses that are then used as input for the NeRF model. We use an entirely different method that elegantly utilizes the simple rendering equations of NeRF to optimize both the camera pose and the map with the same equation.

3. Novel View Synthesis vs 3D Reconstruction

We want to emphasize the difference between a model for novel view synthesis and a 3D reconstruction. A 3D reconstruction is a task that aims at recreating an accurate representation of the geometry of the target scene. View synthesis, on the other hand, is used to recreate an accurate representation of the appearance of the target scene. While these two statements sound similar, they are not the same. In their current form [1, 11, 20] novel view synthesis models are trained based on RGB data from a limited number of views and will only produce reasonable novel appearance renderings within the distribution of camera poses that were used for training the model. In the case of an image




Figure 3. Rendered RGB images on the same camera position from trained radiance field under ScanNet scene0207. From left to right, the order is as follows: ground truth RGB image, from ours, from Vox- Fusion, and from NICE-SLAM.

rendered from a camera pose outside of the training pose distribution, the image will most likely not correspond to the expected result. This is because the underlying geometry from a series of images is ambiguous for areas of the images with no gradients, something that has been well-known in the photogrammetry community for decades. However, although the spatial geometry of such a region is unknown, the appearance of the same image region is known, at least locally, in the training pose distribution, allowing for the production of a model that can produce highly convincing images with incorrect underlying geometry.

This is exactly the reason why radiance models can be so deceiving, but it also gives us a glimpse of our own faulty assumptions when determining scene geometry. If we are presented with an image from a novel view synthesis model that produces an accurate appearance but contains an incorrect underlying geometry, we will not perceive the incorrect model geometry but rather assume a geometry closer to the true scene geometry because that intuitively makes more sense. A good example of this can be seen in Fig. 2. Although the model trained with the RGB sensor looks good in the color rendering, we can see that the underlying geometry is very different from what we would expect compared to the radiance field trained with the RGB-D sensor.

The problem of ambiguous geometry can be mitigated either by training on a scene with color gradients covering significant parts of the scene, thereby reducing the total possible ambiguity, or by training the model with training images covering a 360-degree sphere around the model (inward-facing scene), effectively performing space carving [18]. This will encourage the model to approach the true scene geometry. For many practical applications like XR and robotics, these mitigation strategies are often infeasible and/or impossible. Consider, for instance, an example of a room with monocolored walls (outward-facing scene). There will be little to no color gradient on the walls, and capturing all objects in the room from every angle places significant requirements on the data capture process.

However, if a depth sensor is adopted, this problem can be eradicated, as this allows for the estimation of a geometrically correct model with just a very limited coverage of

the scene. A second, less ideal solution is to adopt an assumption on the geometry. For instance, assume areas with no color gradient are smooth, planar [24], or other more advanced solutions like inference from a pre-trained neural network [8]. This would reduce the ambiguity of the geometry but also induce a bias that might either help or hinder the 3D reconstruction process depending on the assumption’s correctness for any particular scene.

The RGB-D sensor is becoming more ubiquitous, providing an unbiased representation of the geometry and allowing for easy integration. Additionally, the depth signal is often a lower frequency signal than the color signal, and therefore, it also helps increase the radius of convergence when using dense image-to-model alignment. For these two reasons, we focus on the use of the RGB-D sensor in this work.

4. Volumetric Rendering Basics

All radiance field learning approaches share the volumetric rendering equations initially described in [16] and discretized in [20]. They explain how a pixel color is rendered based on a volume containing continuously valued implicit density and color functions.

Color rendering. Let \hat{C} denote the rendered RGB color value of a single pixel from the radiance field model. \hat{C} is obtained by accumulating N rendered sample values computed by the density σ_i and color values c_i for sample points \vec{p}_i along the ray $\vec{r} = \vec{\sigma} + t_i \cdot \vec{d}$, $i \in [0, 1, \dots, N]$. $\vec{\sigma}$ is the camera center, \vec{d} is the ray direction from the camera center to the pixel in the image plane and t_i is the distance from the camera center. The distances between samples are denoted as δ_i . The discrete rendering equations can then be expressed as follows:

$$\hat{C}(\vec{\sigma}, \vec{c}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i, \quad (1)$$

where,

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right). \quad (2)$$

The value T_i represents the amount of light transmitted to sample i along the ray \vec{r} . This is essentially the remaining light of a quota that is not reflected by all previous samples $j = [1, \dots, i - 1]$.

Depth rendering. Let \hat{D} denote the expected depth value of a single pixel rendered from the radiance field model. Depth \hat{D} can be estimated analogously to color rendering by treating the sample distance t_i from the camera as the sample color value c_i in Eq. (1), resulting in the following:

$$\hat{D}(\vec{\sigma}) = \sum_{i=1}^N T_i(\vec{\sigma})(1 - \exp(-\sigma_i \delta_i))t_i, \quad (3)$$

with $T_i(\vec{\sigma})$ as defined in Eq. (2).

5. Method

5.1. Overview

Our proposed algorithm is based on the voxel grid representation for radiance field optimization [11, 37] and is divided into two separate parts: One offline mapping algorithm using RGB-D data with known poses to create a model of the scene, and one online tracking algorithm using the map in an image-to-model alignment scheme for pose optimization. All the partial derivatives of the loss function needed for both tasks are analytically calculated and implemented in CUDA for fast processing on GPUs.

5.2. Mapping

Model representation. Instead of using a neural network as a model representation as proposed in [20], we utilize the Plenoxel representation [11], which represents the radiance field as a sparse voxel grid utilizing trilinear interpolation to produce a continuously valued implicit density and color functions, see Fig. 4. Each vertex in the voxel grid is represented by 28 scalar values: one density value and nine values per color channel. The multiple values per color channel make up the coefficients of spherical harmonics functions, allowing for modeling view dependencies caused by Lambertian surfaces and specular reflections. The main advantage of using a voxel grid compared to a single neural network is the significant reduction in computational effort needed in both training and inference of the model. Practically the direct data access of the voxel grid representation reduces the training time by two orders of magnitude and allows for real-time image rendering [11].

RGB-D mapping. Mapping based on RGB-D data involves simultaneous optimization of the radiance field using the RGB-D sensor color values C and depth value D from the training images with known poses for supervision. The photometric color loss \mathcal{L}_p and geometric depth loss \mathcal{L}_g are defined as follows:

$$\mathcal{L}_p(\hat{C}(\sigma, c)) = \frac{1}{M} \sum_{i=1}^M \|\hat{C}(\sigma_i, c_i) - C\|^2, \quad (4)$$

$$\mathcal{L}_g(\hat{D}(\sigma)) = \frac{1}{M} \sum_{i=1}^M \|\hat{D}(\sigma) - D\|^2. \quad (5)$$

The mapping is performed by non-linear optimization based on the differentiation of Eq. (1) and Eq. (3) with respect to all color \vec{c} and density $\vec{\sigma}$ values.

Partial derivatives of the photometric loss L_p , and geometric loss L_g with respect to the variables (σ_i, c_i) gives us:

$$\frac{\partial \mathcal{L}_p(\hat{C}(\sigma, c))}{\partial(\sigma, c)} = \frac{\partial \mathcal{L}_p(\hat{C})}{\partial \hat{C}} \left(\frac{\partial \hat{C}(\sigma, c)}{\partial c_i} + \frac{\partial \hat{C}(\sigma, c)}{\partial \sigma_i} \right), \quad (6)$$

$$\frac{\partial \mathcal{L}_g(\hat{D}(\sigma))}{\partial(\sigma)} = \frac{\partial \mathcal{L}_p(\hat{D})}{\partial \hat{D}} \frac{\partial \hat{D}(\sigma)}{\partial \sigma}. \quad (7)$$

The derivatives of the rendered color \hat{C} are calculated in [37], but for the sake of completeness, we include them here:

$$\frac{\partial \hat{C}}{\partial c_i}(\sigma, c) = w_i(\sigma) = T_i(1 - \exp(-\sigma_i \delta_i)), \quad (8)$$

$$\frac{\partial \hat{C}}{\partial \sigma_i}(\sigma, c) = \delta_i \left[c_i T_{i+1}(\sigma) - \hat{C} + \sum_{j=0}^i c_j w_j(\sigma) \right] \quad (9)$$

Based on the color derivative, we derive the depth derivative \hat{D} in an analogous fashion. By interpreting the distance value t_i as the color value c_i , we can re-write the equation Eq. (9) for the depth rendering equation:

$$\frac{\partial \hat{D}}{\partial \sigma_i}(\sigma) = \delta_i \left[t_i T_{i+1}(\sigma) - \hat{D} + \sum_{j=0}^i t_j w_j(\sigma) \right] \quad (10)$$

As the partial derivatives for equations are all analytically defined, it allows us to implement them directly in a custom CUDA kernel. Building on the code developed for Plenoxel [11] to boost computational speed.

The final loss function we employ during mapping is represented as follows:

$$\mathcal{L} = \mathcal{L}_p + \lambda_d \mathcal{L}_g, \quad (11)$$

where λ_d is the scaling factor of the geometric loss.

5.3. Tracking




Figure 4. Illustration of trilinear interpolation on sample point p_i lying on the sample ray in a voxel base radiance field.

Up until this point, we have assumed known and fixed camera poses while treating the model parameters $(\vec{\sigma}, \vec{c})$ as variables. To estimate pose on a model build apriori, we use the same volumetric equations but consider the pose as the variable and the model parameters as static. Our objective is then to perform image-to-model alignment using the fixed volumetric color and density parameters from the radiance field, referencing the RGB-D input. To achieve this, we need to determine the camera pose by finding the derivative of the color and depth rendering with respect to the ray $\vec{r} = \vec{\sigma} + t_i \vec{d}$. We do this by leveraging the use of an alternative version of the chain rule:

$$\begin{aligned} \frac{\partial \hat{C}}{\partial \vec{r}} &= \sum_{i=1}^N \frac{\partial \hat{C}}{\partial \vec{p}_i} \frac{\partial \vec{p}_i}{\partial \vec{r}} \\ &= \sum_{i=1}^N \left(\frac{\partial \hat{C}}{\partial \sigma_i} \frac{\partial \sigma_i}{\partial \vec{p}_i} + \frac{\partial \hat{C}}{\partial c_i} \frac{\partial c_i}{\partial \vec{p}_i} \right) \frac{\partial \vec{p}_i}{\partial \vec{r}} \end{aligned} \quad (12)$$

$$\frac{\partial \hat{D}}{\partial \vec{r}} = \sum_{i=1}^N \frac{\partial \hat{D}}{\partial \sigma_i} \frac{\partial \sigma_i}{\partial \vec{p}_i} \frac{\partial \vec{p}_i}{\partial \vec{r}}. \quad (13)$$

The partial derivatives $\frac{\partial \hat{C}}{\partial \sigma_i}$, $\frac{\partial \hat{C}}{\partial c_i}$, $\frac{\partial \hat{D}}{\partial \sigma_i}$ are already given in equations 8, 9 and 10 and are shared by all volumetric radiance field methods based on equations 1 and 3, while $\frac{\partial \sigma_i}{\partial \vec{p}_i}$, $\frac{\partial c_i}{\partial \vec{p}_i}$ and $\frac{\partial \sigma_i}{\partial \vec{p}_i}$ are model specific. For our chosen representation based on the tri-linear interpolation functions, these partial derivatives produce quite messy functions, so we reserve the complete analytical equations for Appendix A.

The derivative of a sample point \vec{p}_i with respect to the ray \vec{r} can further be broken down as:

$$\frac{\partial \vec{p}_i}{\partial \vec{r}} = \begin{bmatrix} \frac{\partial \vec{p}_i}{\partial \vec{\sigma}} & \frac{\partial \vec{p}_i}{\partial \vec{d}} \end{bmatrix} = [1, t_i]. \quad (14)$$

This allows for direct optimization of the position and orientation of the camera parameters.

5.4. Optimization Details

The optimization heuristics are kept as simple as possible, and mapping is done using a random set of sample rays from all available images, while tracking is done by randomly sampling rays from the current image. The optimizers used are RMSProp [14] for the mapping, and Adam [17] for the tracking. There is no importance sampling, keyframing, or similar optimization schemes.

6. Experiments

6.1. Experimental Setup

We test both mapping and tracking algorithms on both synthetic and real indoor datasets of varying sizes and compare them to existing state-of-the-art algorithms performing similar tasks.

Datasets. Two different RGB-D datasets are used: 1) The Replica dataset [29] containing a total of 18 different synthetic indoor environments with highly accurate sensor data, 2) The ScanNet dataset [5], a real-world dataset also captured from indoor scenes, containing over 1000 unique sequences. As a real-world dataset, the latter inherently includes a significant amount of missing and imperfect depth measurements. We use a subset of eight sequences from the Replica dataset and five sequences from the ScanNet dataset. Both of these subsets have been the standard for comparison in previous works [31, 35, 40].

Comparison algorithms. To the best of our knowledge, no other algorithms are currently performing asynchronous mapping and tracking in a radiance field. Therefore, to draw comparisons with existing approaches, we select NICE-SLAM [40], and Vox-Fusion [35] as competing methods due to their status as state-of-the-art algorithms for radiance field-based simultaneous localization and mapping. To make the comparison fair, we modify these algorithms first to optimize the map using ground truth poses with a subset of the sequence images and then perform tracking based on that model.

Metric. For the offline mapping task, we report on two primary metrics to illustrate the geometrical accuracy of the constructed map: the average L1 depth loss in metric units and PSNR score of the RGB values of the model. These metrics are calculated based on randomly sampled pixels from randomly sampled images from the sequence.

To measure tracking accuracy, we use three key metrics stemming from the root mean square error (RMSE): Absolute Trajectory Error (ATE) [30], and the Relative Pose Error for both translation (RPE_t) and rotation (RPE_r). ATE reflects the global accuracy of the trajectory as it is sensitive to drifting over time. We put a greater emphasis

Methods	Metric	Room-0	Room-1	Room-2	Office-0	Office-1	Office-2	Office-3	Office-4	Avg.	time(avg.)(ms)
Vox-Fusion [35]	ATE[m]↓	0.0042	0.0036	0.0090	0.6539	0.0029	0.0038	0.0042	0.0046	0.0046	581.35
	RPE _t [m]↓	0.00391	0.0039	0.0088	0.0500	0.0032	0.0036	0.0039	0.0050	0.0046	
	RPE _r [°]↓	0.1076	0.1275	0.4034	2.6541	0.1501	0.1331	0.1316	0.1471	0.1715	
Ours ¹	ATE[m]↓	0.0017	0.0036	0.0020	0.0077	0.0020	0.0027	0.0022	0.0027	0.0031	419.86
	RPE _t [m]↓	0.0008	0.0059	0.0013	0.0034	0.0012	0.0009	0.0008	0.0011	0.0019	
	RPE _r [°]↓	0.0191	0.0736	0.0483	0.1564	0.0440	0.0282	0.0199	0.0278	0.0522	
NICE-SLAM [40]	ATE[m]↓	0.0119	0.0220	0.0334	0.0100	0.0042	0.0093	0.0862	0.0573	0.0293	149.57
	RPE _t [m]↓	0.0170	0.0229	0.0215	0.0143	0.0059	0.0131	0.0445	0.0305	0.0212	
	RPE _r [°]↓	0.3267	0.5103	1.4032	0.2700	0.2037	0.3247	1.0076	0.5058	0.5690	
Ours ²	ATE[m]↓	0.0093	0.0061	0.0179	0.0117	0.0027	0.0056	0.0072	0.0061	0.0083	145.20
	RPE _t [m]↓	0.0065	0.0090	0.0108	0.0055	0.0025	0.0033	0.0046	0.0039	0.0058	
	RPE _r [°]↓	0.1494	0.1903	0.4833	0.2034	0.1196	0.1116	0.1201	0.1113	0.1861	

Table 2. Trajectory estimation results on the Replica dataset. The lowest errors are indicated in bold. Time here stands for tracking time per frame. We provide the results for two different configurations of our method, varying in number of samples per frame.

	RGB PSNR↑	Depth diff. (m/pixel)↓
Vox-Fusion [35]	19.379	0.0705
NICE-SLAM [40]	18.455	0.0514
Ours	24.411	0.0469

Table 3. comparison of map geometry accuracy on ScanNet. The values are average through 5 sequences reported in Tab. 4.

on the RPE_t and RPE_r due to local accuracy being a more relevant metric when tracking based on a map known a priori. We measure RPE in 1-meter intervals. These tracking metrics were implemented using the "evo" Python package for odometry and SLAM evaluation [13].

Implementation details and parameter selection. All experiments were conducted on a machine equipped with an Intel Core i9-11900KF CPU and an NVIDIA RTX 3090 graphics card, with custom CUDA code for efficient runtimes. For all experiments, our method used a dynamic voxel grid with a peak voxel resolution of 512^3 , irrespective of the size of the scene. While all images are used for tracking, only every 10th image is used for mapping. Pixels with invalid depth measurements are disregarded in both mapping and tracking.

For mapping, we allot our algorithm roughly 6 seconds per frame for mapping, while the competing methods are given up to 18 seconds per frame. The parameters of the competing methods are adjusted to take full advantage of this extra time.

For tracking, the number of sample rays and the number of iterations varied between each method and each dataset. We adhere to the default configurations for these parameters provided by NICE-SLAM and Vox-Fusion for each dataset, while we adjusted the parameters for our model to match the per-frame tracking time of the two comparison models.

6.2. Mapping Results

RGB vs RGB-D mapping. To first demonstrate the point highlighted in Sec. 3, we test the difference in mapping performance between the original Plenoxel and our RGB-D modified mapping algorithm on the Office-1 scene of the Replica dataset. The results are shown in Fig. 2 and Tab. 1. Despite the PSNR for RGB rendering actually being higher for the RGB-only reconstruction, the average L1 depth error per pixel is 70cm, compared to the RGB-D L1 depth error of just 9mm. Confirming that although RGB alone might look convincing for use in novel view synthesis, it is insufficient for learning the underlying geometrical information, especially in areas with weaker gradients.

Baseline comparison. The results of computing the average PSNR and L1 difference for five sequences of ScanNet are presented in Tab. 3. Even with significantly less training time than the competing methods, our model achieves a significantly better RGB PSNR and produces the best 3D reconstruction by exhibiting the smallest L1 depth loss, keeping within a range of 5cm. The qualitative difference in RGB estimation can be observed in Fig. 3, where our method shows a visually more correct image compared to the over-smoothed results from baselines. This builds a strong foundation for achieving better tracking accuracy.

6.3. Tracking Results

Replica dataset. For the replica dataset, we report two different sampling configurations of our system. One that roughly matches the processing time of Vox-Fusion, and one that matches NICE-SLAM. Tab. 2 shows the estimated pose tracking accuracy for the entire sequence on the Replica dataset. NICE-SLAM estimates the pose with efficient use of samples, with relatively few rays. This leads to a much faster processing speed than Vox-Fusion but still with a comparative reduction in overall accuracy. When

Table 4. Trajectory estimation results on the ScanNet dataset. The lowest errors are indicated in bold. Time here stands for tracking time per frame. Both NICE-SLAM and ORB-SLAM2 failed to localize camera pose in the middle of sequence 0181. The average errors of these methods were computed excluding frame 0181 and are indicated with (*) marks.

Methods	Metric	0000	0106	0169	0181	0207	Avg.	time(avg.)(ms)
Vox-Fusion [35]	ATE[m]↓	0.0274	0.2424	0.0315	0.0924	0.0323	0.0852	1021.19
	RPE _t [m]↓	0.0222	0.0763	0.0284	0.0443	0.0303	0.0403	
	RPE _r [°]↓	0.6157	2.7116	0.8805	1.8274	1.0733	1.4217	
NICE-SLAM [40]	ATE[m]↓	0.0405	0.1188	0.1952	-	0.0491	0.1009*	360.70
	RPE _t [m]↓	0.0360	0.0666	0.0643	-	0.0457	0.0531*	
	RPE _r [°]↓	1.0862	2.3725	1.4703	-	1.4871	1.6040*	
Ours	ATE[m]↓	0.0246	0.0387	0.0165	0.0373	0.0259	0.0286	350.84
	RPE _t [m]↓	0.0177	0.0209	0.0160	0.0264	0.0230	0.0208	
	RPE _r [°]↓	0.5030	0.6216	0.4484	0.7872	0.7231	0.6166	
ORB-SLAM2 [23]	ATE[m]↓	0.0779	0.0838	0.1038	-	0.0898	0.0888*	36.30
	RPE _t [m]↓	0.0933	0.0516	0.0625	-	0.0848	0.0730*	
	RPE _r [°]↓	2.8584	2.9895	2.3596	-	3.5897	2.9493*	

we matched our method to NICE-SLAM’s tracking speed, our method outperformed NICE-SLAM across almost all metrics and sequences. Notably, NICE-SLAM showed instability in the rotational aspect of relative pose accuracy, as it relied on neural network processing with a relatively large voxel size, leading to challenges in precisely estimating subtle rotational differences. NICE-SLAM marginally outperformed our method in the Office-0 sequence in terms of ATE but is beaten on the relative pose metrics. Vox-Fusion showed remarkable tracking performance, taking advantage of its considerably longer tracking time. It trailed only slightly behind our proposed method in most of the sequences. However, like NICE-SLAM, it also tended to have lower accuracy for the orientation of the relative pose. For an analysis of our model’s speed vs tracking accuracy trade-off, please see Appendix B.

ScanNet dataset. Tab. 4 displays the comparison of tracking performance on the five ScanNet sequences. Our method excels in all metrics across all sequences when compared to the comparison algorithms while demanding significantly less computation time. Our method even outperforms Vox-Fusion, using only a third of the computation time. Since our method relies on directly optimizing small voxels, we initially presumed the significant amount of invalid depth measurements in ScanNet data might lead to holes in the final model, but our method successfully filled these gaps and created a seamless, hole-free model.

To ground our work with pre-NeRF research, we additionally provide comparisons with a more traditional RGB-D SLAM method that does not use a radiance field model, ORB-SLAM2 [23], which has long been the standard baseline in SLAM systems. We do not separate the mapping and tracking tasks for ORB-SLAM2 as we have done for

the other methods because we observed worse performance when doing this, but we have included it as a point of reference. Although a feature-based system like ORB-SLAM2 produces good results in significantly less time than any of the other algorithms, it cannot beat the accuracy of any of the dense methods. This primarily comes down to three factors: 1) The offline mapping gives a major advantage to the dense systems, 2) the dense systems use all image information, and 3) the ScanNet sequences have a significant amount of repetitive patterns on carpets and walls that make the data correspondence task difficult for ORB-SLAM2.

7. Conclusion

We have presented the analytical augmentation of the Plenoxel algorithm to allow for RGB-D mapping and tracking based on the radiance field equations. We have also argued for and shown both qualitatively and quantitatively the need for RGB-D sensors to accurately and reliably reconstruct outward-facing scenes when modeling using methods from the current paradigm of learning-based radiance field algorithms. Our method achieves superior results with much less time on both the mapping and tracking tasks compared to state-of-the-art radiance field-based SLAM methods when modified to perform offline mapping based on ground truth pose data.

Acknowledgement

This work is financially supported by the Korea Agency for Infrastructure Technology Advancement (KAIA) grant funded by the Ministry of Land, Infrastructure and Transport (Grant RS-2021-KA160637) and the Norwegian Research Council in the project Autonomous Robots for Ocean Sustainability (AROS), project number 304667.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. [3](#)
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. [2](#)
- [3] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J Davison. Codeslam—learning a compact, optimisable representation for dense visual slam. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2560–2568, 2018. [2](#)
- [4] Chi-Ming Chung, Yang-Che Tseng, Ya-Ching Hsu, Xiang-Qian Shi, Yun-Hung Hua, Jia-Fong Yeh, Wen-Chin Chen, Yi-Ting Chen, and Winston H Hsu. Orbeez-slam: A real-time monocular visual slam with orb features and nerf-realized mapping. *arXiv preprint arXiv:2209.13274*, 2022. [3](#)
- [5] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. [6](#)
- [6] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. [3](#)
- [7] Arnab Dey, Yassine Ahmine, and Andrew I Comport. Mip-nerf rgb-d: Depth assisted fast neural radiance fields. *arXiv preprint arXiv:2205.09351*, 2022. [3](#)
- [8] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10786–10796, 2021. [4](#)
- [9] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):611–625, 2017. [2](#)
- [10] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2016. [2](#)
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. [2](#), [3](#), [5](#)
- [12] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. [3](#)
- [13] Michael Grupp. evo: Python package for the evaluation of odometry and slam. <https://github.com/MichaelGrupp/evo>, 2017. [7](#)
- [14] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. *Cited on*, 14(8):2, 2012. [6](#)
- [15] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 559–568, 2011. [2](#)
- [16] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3):165–174, 1984. [4](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [18] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38:199–218, 2000. [4](#)
- [19] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741–5751, 2021. [3](#)
- [20] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [2](#), [3](#), [4](#), [5](#)
- [21] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. [3](#)
- [22] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. [2](#)
- [23] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics*, 33(5):1255–1262, 2017. [8](#)
- [24] Richard A Newcombe, Steven J Lovegrove, and Andrew J Davison. Dtam: Dense tracking and mapping in real-time. In *2011 international conference on computer vision*, pages 2320–2327. IEEE, 2011. [2](#), [4](#)
- [25] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14335–14345, 2021. [2](#)
- [26] Antoni Rosinol, John J Leonard, and Luca Carlone. Nerf-slam: Real-time dense monocular slam with neural radiance fields. *arXiv preprint arXiv:2210.13641*, 2022. [3](#)
- [27] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [3](#)

- [28] Thomas Schops, Torsten Sattler, and Marc Pollefeys. Bad slam: Bundle adjusted direct rgb-d slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 134–144, 2019. 2
- [29] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 6
- [30] Jürgen Sturm, Nikolas Engelhard, Felix Endres, Wolfram Burgard, and Daniel Cremers. A benchmark for the evaluation of rgb-d slam systems. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 573–580. IEEE, 2012. 6
- [31] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021. 3, 6
- [32] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–12, 2023. 3
- [33] Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *Advances in neural information processing systems*, 34:16558–16569, 2021. 2
- [34] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 3
- [35] Xingrui Yang, Hai Li, Hongjia Zhai, Yuhang Ming, Yuqian Liu, and Guofeng Zhang. Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 499–507. IEEE, 2022. 3, 6, 7, 8
- [36] Lin Yen-Chen, Pete Florence, Jonathan T Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. inerf: Inverting neural radiance fields for pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1323–1330. IEEE, 2021. 3
- [37] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. 5
- [38] Wenyuan Zhang, Ruofan Xing, Yunfan Zeng, Yu-Shen Liu, Kanle Shi, and Zhizhong Han. Fast learning radiance fields by shooting much fewer rays. *arXiv preprint arXiv:2208.06821*, 2022. 2
- [39] Zihan Zhu, Songyou Peng, Viktor Larsson, Zhaopeng Cui, Martin R Oswald, Andreas Geiger, and Marc Pollefeys. Nicer-slam: Neural implicit scene encoding for rgb slam. *arXiv preprint arXiv:2302.03594*, 2023. 3
- [40] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. 3, 6, 7, 8