

BAA-NGP: Bundle-Adjusting Accelerated Neural Graphics Primitives

Sainan Liu*
Intel Labs

sainan.liu@intel.com

Shan Lin*
UCSD

shl102@ucsd.edu

Jingpei Lu
UCSD

jil360@ucsd.edu

Alexey Supikov
Intel Labs

alexei.soupikov@intel.com

Michael Yip
UCSD

yip@ucsd.edu

Abstract

Implicit neural representations have become pivotal in robotic perception, enabling robots to comprehend 3D environments from 2D images. Given a set of camera poses and associated images, the models can be trained to synthesize novel, unseen views. To successfully navigate and interact in dynamic settings, robots require the understanding of their spatial surroundings driven by unassisted reconstruction of 3D scenes and camera poses from real-time video footage. Existing approaches like COLMAP and bundle-adjusting neural radiance field methods take hours to days to process due to the high computational demands of feature matching, dense point sampling, and training of a multi-layer perceptron structure with a large number of parameters. To address these challenges, we propose a framework called bundle-adjusting accelerated neural graphics primitives (BAA-NGP) which leverages accelerated sampling and hash encoding to expedite automatic pose refinement/estimation and 3D scene reconstruction. Experimental results demonstrate 10 to 20 \times speed improvement compared to other bundle-adjusting neural radiance field methods without sacrificing the quality of pose estimation. The github repository can be found here <https://github.com/IntelLabs/baa-ngp>.

1. Introduction

As robotics continues to push the boundaries of automation and intelligence, the incorporation of advanced computer vision techniques has become indispensable. The effectiveness of robots in the real world relies on their ability to localize, perceive, and model the 3D world from cameras

with high accuracy and high resolution. This perception challenge has been ever-present in robotics and automation research for decades. In recent years, implicit neural representation (INR) [4] has shown great success in capturing highly detailed and accurate 3D reconstructions of objects and scenes with neural networks. Given a set of known camera intrinsic and extrinsic parameters and a rough estimation of the boundary of the scene, an implicit 3D scene representation can be learned by sampling 3D points along camera rays and performing supervised learning using ground-truth color images [15]. Such INR can accurately model scene geometry, radiance, and material properties providing key data for robotics tasks like localization, navigation, and object manipulation [1, 5, 12, 23, 32].

Using INR in robotics poses unique challenges. The need for accurate camera poses is one of them. While popular methods like COLMAP [19, 20] can estimate the camera poses from input images, they often drop frames that lack distinctive features and lead to suboptimal results. Some approaches propose to simultaneously optimize pose estimations and learn the radiance field [3, 10]. Yet, their computational cost, leading to hours of training on just 100 images of size 400×400 , makes them unsuitable for real-time applications.

Recent works focused on making INR learning faster. For example, instant neural graphics primitives (iNGP) [17] uses occupancy grid and hash encodings to converge significantly faster (< 5 seconds) than other techniques with ground truth poses. However, this approach was not shown to work without or with poorly estimated camera poses, making it less applicable in many real-world, roaming camera scenarios.

In this paper, we propose a novel approach called Bundle-Adjusting Accelerated Graphics Primitives (BAA-NGP) that estimates camera poses and optimizes the ra-

*Equal contributions.

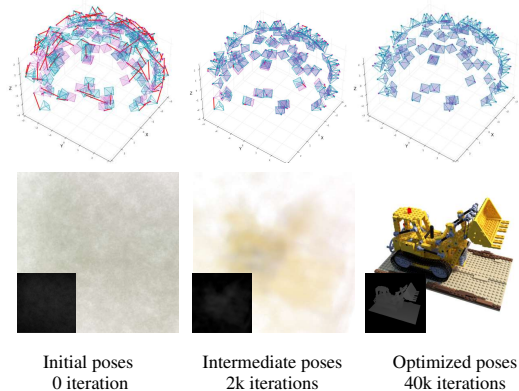


Figure 1. BAA-NGP is a neural implicit representation that captures 3D scenes from 2D images with unknown camera poses. It learns the 3D scene together with the camera poses within minutes of training, whereas previous methods would have taken hours.

diance field simultaneously using a computationally optimized approach that results in a 10 to 20 times speedup Figure 1. This fundamentally addresses the challenge of accelerated learning of INR models in unstructured settings without tracked cameras, making the method widely applicable to many real-world applications. BAA-NGP works by combining pose estimation with fast occupancy sampling and multiresolution hash encoding through a new curriculum learning strategy. We evaluate our proposed method on several benchmark datasets, including multi-view object-centric scenes [15] as well as frontal-camera video sequences of unbounded scenes [14]. We compare our results with state-of-the-art techniques and show that we achieve comparable or better performance while being significantly faster to learn. This gives BAA-NGP a range of applicability to a broad set of real-world scenarios and applications ranging from robotics and automation to virtual and augmented reality.

2. RELATED WORK

Despite the relatively recent development of INRs, an expansive number of approaches have been proposed in the literature. Therefore we will focus on works related to solving for INRs with noisy/poor camera poses.

2.1. Camera Pose Estimation for 3D Perception

Structure from motion (SfM) is a classical method often used for 3D structure reconstruction from 2D images or video sequences involving estimating camera poses for each image and simultaneously recovering 3D positions of the scene points. They leveraging feature detection and matching across frames. The camera pose estimation techniques from SfM have been used in the context of INR training. COLMAP [19, 20] is a commonly used library for performing SfM, and many INR methods rely on this as part of

their pipeline. However, COLMAP and other SfM techniques have high computational and memory requirements and rely heavily on an abundance of salient features in the scene, which may result in missing frames in scenes with limited or poor feature detection/saliency.

To avoid the limitations associated with using SfM for camera pose estimation, NeRF— [25] and Self-Calibrating Neural Radiance Fields [8] learned camera pose estimates along with the neural radiance fields. These methods assumed forward-facing views and involved two-stage procedures for estimating camera poses and updating the neural radiance fields during training. Bundle-adjusting neural radiance field (BARF) [10] first proposed a direct extrinsics estimation while learning the implicit 3D scene with NeRF. A coarse-to-fine feature weighting schedule for positional encoding features was used and was critical for smoothing the signals for pose updates. G-Nerf [13] used GANs to improve generalization to large baseline captures. Gaussian-Activated Radiance Fields [3] then proposed that without positional encoding, updating the activation function of the multi-layer perceptron can further improve the pose estimation and the quality of the results.[2] proposed to further constrain pose estimation with a local deformation field, which improved the convergence quality for gradual camera movement. However, these methods still require an extensive amount of time to train (in the low tens of hours).

2.2. Approaches to Reduce INR Learning Time

To overcome the long training time of INR learning, recently, several methods have been proposed to accelerate the process from the aspect of sampling, encoding, and better hardware integration of the multilayer perceptron, including but not limited to [6, 18, 22, 26, 28]. However, these methods assume known camera poses. Methods that train on scenes without camera poses still rely on SfM, such as COLMAP, for pose estimation. [11, 27] has utilized faster architecture for pose retrieval, but their task is to align a novel view with an existing implicit 3D object, whereas we are learning the 3D representation from scratch. [21, 31] are real-time SLAM methods that utilize implicit neural rendering techniques, but they both require extra depth as part of the input. Instead, a recent paper [7] presented a fast neural radiance field learning without a camera prior that avoided classical SfM methods. This work utilizes gradient smoothing and re-implemented iNGP in PyTorch in order to use multi-level learning rate scheduling to incorporate multi-resolution hash encoding with pose estimation. Unfortunately, the code is unavailable, so the performance cannot be cross-validated.

We focus on improving the speed of the basic bundle-adjusting NeRF structure. Our work approaches the problem via a simple and novel coarse-to-fine feature

re-weighting scheme which enables pose refinement capability for hash encoding with multi-resolution occupancy sampling, resulting in $10 \times$ shorter training iterations than the concurrent work [7] and 10 to $20 \times$ speedup on two benchmark datasets comparing to BARF [10]. For unbounded scenes, we incorporate inverted sphere parameterization with hash encoding, which enables pose estimation for more general scenarios.

3. Methodology

3.1. Problem Formulation

Our goal is to learn INR models from images taken without known camera poses or poorly estimated poses. The INR model is described in general terms by a function

$$F_{\Theta}(p) : \{x, y, z\} \rightarrow c, \sigma \quad (1)$$

where a 3D location in the scene $\{x, y, z\}$ maps to a 3-channel radiance color $c = \{r, g, b\}$ and density $\sigma \in [0, 1]$. In order to reconstruct an image from a camera positioned and oriented as $p \in \mathbb{S}\mathbb{E}(3)$ in the scene, we apply a pixel-by-pixel ray marching procedure, where for each pixel location $[u, v]$, we integrate all the radiance values found along the ray going from the camera origin and passing through the pixel location $[u, v]$. Therefore, the output color of a pixel along the ray R of a camera with pose p can be defined as:

$$f(p, R) = \sum_{i=1}^N \alpha_i T_i c_i \quad (2)$$

where $T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)$ is the accumulated transmittance along the ray, and c_i is the output color of sample i from $F_{\Theta}(p)$, and α for each ray segment is calculated as

$$\alpha_i = 1 - \exp(-\sigma_i \delta_i) \quad (3)$$

where δ_i is the distance between sample point i and $i + 1$. σ_i is the density output of sample point i .

For training, given M captured images $\{\mathcal{I}_i\}_{i=1}^M$ of width W and height H taken from the same scene, we optimize the INR model, and simultaneously, as a byproduct, optimize the camera poses $\{p_i\}_{i=1}^M \in \mathbb{R}^6$ associated with each image.

Similar to [10] and [3], we have no image order or sequence assumptions. We assume that camera intrinsics are known, but by definition, camera extrinsics are unknown. We also assume that the scene is static. Solving the problem involves minimizing the following loss function:

$$\min_{p_1, \dots, p_M, \Theta} \sum_{i=1}^M \sum_{j=1}^W \sum_{k=1}^H \|\hat{\mathcal{I}}_{ijk}(p_i, \Theta) - \mathcal{I}_{ijk}\|_2^2 \quad (4)$$

The camera poses $\{p_1 \dots p_M\}$ are optimization variables and are found jointly with the model parameters Θ . Previous




Figure 2. INRs use posed images from multiple viewpoints to reconstruct the scene. In our problem, we assume that a sequence of images was taken from unknown viewpoints for unbounded scenes (left) and poorly estimated viewpoints for bounded scenes (right). Purple frames are initial camera poses, gray/blue frames are ground truth camera poses, and the red line indicates a translation error.




Figure 3. Reparameterization of 3D space via spherical contraction. Each point $p = (x, y, z)$ outside unit sphere becomes $(x', y', z', 1/h)$, a quadruple that converts unbounded distances h to bounded distances $1/h$.

works involving camera pose estimation [3, 10] have considered the following two scenarios, each with different datasets and objectives (Figure 2):

- *Pose estimation (unbounded scene)*: Camera poses are reconstructed from video with frames captured in sequence relatively close to each other. This “forward-facing scene” data is often from cellphone cameras or mobile robots.
- *Pose refinement (bounded scene)*: camera pose corrections are found for available but noisy camera poses for multi-viewpoint images of a centered object.

Small variations in the model configuration are used to account for these two domain types.

3.2. Approach

3.2.1 Network Architecture

We start with an established model for learning implicit neural representations, i.e., neural radiance field, consisting of a neural network model for representing $F_{\theta}(p)$. One of the most successful models for this task is an MLP. Our final model involves the following stages for image reconstruction: inverted sphere parameterization, multi-resolution hash encoding, then an MLP decoder.

Parameterized camera poses to contracted 3D points:

Since hash grids and occupancy grids need finite bounding boxes known in advance, handling bounded and unbounded scenes requires different 3D space parameteriza-

tions. For bounded scenes with a given bounding box, we use affine transformations to bring the volume of interest into an axis-aligned bounding box $[0, 1]^3$. For unbounded scenes, we apply an inverted sphere parameterization allowing compact mapping of close and far locations [29]. A point (x, y, z) seen by a virtual camera with camera center $o \in \mathcal{R}^3$ along a ray $r = o + td$ with unit direction $d \in \mathcal{R}^3$, and with distance $t = [0, \infty]$, can be located arbitrarily far. To address this, a point within the unit sphere, centered at the origin, is unchanged. A point outside of the volume $h = \sqrt{x^2 + y^2 + z^2} > 1$ is reparameterized to quadruple $(x', y', z', 1/h)$, where $(x', y', z') = (x, y, z)/h$, and $x'^2 + y'^2 + z'^2 = 1$. Affine transformations are used to bring points within the volume into $[0, 1]^3$ and reparameterized 4D representation into $[0, 1]^4$.

Multiresolution Hash Encoding: We use multi-resolution hash encoding [17] to speed up the radiance field learning. It transforms every input 3D point into a higher dimension concatenated features learned from each resolution level. For each point, $x \in \mathbb{R}^d$, the method first finds the grid cell in which the point resides, then d -linearly interpolates the features at vertices of the grid cell to produce an F -dimensional feature vector for the point x . Features from multiple resolution levels are then concatenated into a single vector. Instead of explicitly storing vertex features in a regular grid of size N along each dimension, which would result in N^d features for a level, hash encoding fixes the total number of F -dimensional features per level to be T . Hence if a coarser level needs less than T features, the mapping remains 1:1, but if a finer level requires more than T features, a spatial hash function (eq. 5) [24] is then used to map them into T sets of features.

$$h(x) = \left(\bigoplus_{i=1}^d x_i \pi_i \right) \bmod T \quad (5)$$

where x denotes the vertex position in d dimension, \bigoplus indicates the bitwise XOR, and $\pi_1 = 1, \pi_2 = 2654435761, \pi_3 = 805459861$ are selected by the original authors [17].

By specifying the base resolution (N_{min}) and the maximum resolution (N_{max}) of the grid, for L levels, the grid resolution in between is chosen by $N_l = \lfloor N_{min} \cdot b^l \rfloor$, where b is calculated based on Equation 6.

$$b = \exp((\ln N_{max} - \ln N_{min}) / (L - 1)) \quad (6)$$

Value along a dimension of a point with range $[0, 1]$ can be mapped to a certain grid cell along that dimension within the l layers via $\lfloor x_l \rfloor$ and $\lceil x_l \rceil$, where $x_l = x \cdot N_l$. The d -linear interpolation weight for that dimension can then be found, for example, via $w_l = x_l - \lfloor x_l \rfloor$ for that vertex corner.

Multi-layer Perceptron (MLP): MLPs are commonly used for capturing neural radiance fields. We incorporate a

fully fused implementation of MLPs[17] with one layer for density learning and two layers for color learning, which takes the output of the density layer as well as the view directions. The MLPs takes in the encoded learnable features from multi-resolution hash encoding, concatenated with the direction encoded with spherical harmonics transform, and outputs color and density for each 3D query point.

3.3. Training Strategies

3.3.1 Occupancy Grid Sampling

For further acceleration, we adopted coarse multiscale grid sampling before passing points through the radiance field for gradient backpropagation. Our occupancy grid is defined around the origin with an axis-aligned bounding box. The density of the sampled points is first queried without gradient to check if they hit free space. Only points that reside in a coarse grid cell that do not return free space (above a certain threshold) will be passed through for backpropagation.

3.3.2 Signal Coarse-to-Fine Smoothing

A coarse-to-fine smoothing technique has been proposed [10] to shield the earlier stage of pose estimation from being affected by high-frequency information and slowly introduce such information at later stages for improved pose fine-tuning and image reconstruction quality. They use positional encoding and weigh the k -th frequency components of the inputs with w_k :

$$w_k(\alpha) = \begin{cases} 0 & \text{if } \alpha < k \\ \frac{1 - \cos((\alpha - k)\pi)}{2} & \text{if } 0 \leq \alpha - k < 1 \\ 1 & \text{if } \alpha - k \geq 1 \end{cases} \quad (7)$$

where $\alpha_i = \frac{s_i - r_s}{r_e - r_s}$, r_s is the coarse-to-fine starting point as the percentage of the training progress, r_e is the coarse-to-fine ending point as the percentage of the training progress, and s is the percentage of training progress so far at step i . This method assumes the original coordinates are concatenated with the positional encoding features, allowing initial pose estimation with just 3D points as inputs.

We propose a novel coarse-to-fine strategy better suited for learnable hash encoding features. Since hash grid features, accompanying the very small MLP, carry information about density and radiance distributions in the volume, they are not equivalent to scene-agnostic frequency encoding of the 3D coordinates. Using zero weights of the cosine window to nullify the features at finer levels results in MLP getting stuck at the early local minima and failing to learn hash grid features efficiently. Instead, starting with the coarsest level enabled, we prime the MLP by using coarse-level features in lieu of windowed-out fine-level features. We gradually replace such coarse-level estimates with the

actually learned fine-level values as the cosine window expands. Thus, the features are weighted with

$$\gamma_k(x; \alpha) = w_k(\alpha) \cdot d_k + (1 - w_k(\alpha)) \cdot d_\alpha \quad (8)$$

where d_α indicates the set of features that has the highest grid level with a nonzero weight.

4. Experiments

4.1. Dataset

We benchmarked BAA-NGP on the LLFF dataset for frontal-camera in-the-wild video sequences and the blender synthetic dataset for pose refinement. The following describes the dataset and the associated model architectures:

LLFF This video sequence dataset has images from continuous camera poses that concentrate the views at one side of the scene. The challenge of this dataset is that not many frames of each video sequence are available, and all cameras are posed on one side of the scene, giving limited coverage from multiple viewpoints. For training and evaluation, all images were resized to 640×480 .

For the network training setup, we used occupancy sampling, inverted sphere reparameterization for input points, and two sets of multiresolution hash encoding for in-sphere and outside spaces in addition to the vanilla MLP setup. Similar to [10], all cameras are initialized with the identity transformation. We use the Adam optimizer and train the models for each scene with 20K iterations. The number of randomly sampled pixel rays is initialized as 1024 and dynamically adjusted throughout the training to keep the total number of samples taken along the rays consistent. The learning rate of the network is initialized as 1×10^{-4} and linearly increased to 1×10^{-2} for the first 100 iterations, and then step decayed by a factor of 0.33 at 10000, 15000, and 18000 iterations. Following a similar pattern, the learning rate of the camera poses is initialized as 3×10^{-4} and linearly increased to 3×10^{-3} for the first 100 iterations, and then step decayed by a factor of 0.33 at 10000, 15000, and 18000 iterations.

Blender synthetic The Blender synthetic dataset is a multi-view image dataset that includes imperfect camera pose estimations, for camera pose refinement and scene reconstruction. This dataset comprises 100 training and 200 testing images, which we resized to 400×400 for training and evaluation. The images are captured from the upper hemisphere of the central object and are rendered against a white background.

In terms of network training setup, we incorporate occupancy sampling, multiresolution hash encoding ($N_{min} = 14$, $N_{max} = 4069$, $L = 16$), and signal coarse-to-fine smoothing ($r_s = 0.1$, $r_e = 0.5$), with the vanilla MLP setup. The camera poses are perturbed by adding noise $\mathcal{N}(0, 0.15\mathbf{I})$ to the ground-truth poses, similar to the setup

in Lin et al. [10]. We trained the model for each scene over 40K iterations. We used the Adam optimizer and applied an exponentially decaying learning rate schedule, starting from 1×10^{-2} and decaying to 1×10^{-4} for the network, and 1×10^{-3} decaying to 1×10^{-5} for the camera poses.

4.2. Metrics

We benchmarked our experiments for both novel view image reconstruction quality and pose estimation accuracy. For novel view synthesis, peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) [30], multi-scale structural similarity index (MS-SSIM). Learned perceptual image patch similarity (LPIPS) [30] are used to evaluate the predicted image quality against the ground truth novel views. For pose estimation, we evaluate the camera pose rotation and translation differences between the predicted and the ground truth camera set after using Procrustes analysis for alignment.

4.3. Results

We build our architecture on top of nerfacc[9], and use hash encoding from tiny-cuda-nn [16]. An NVIDIA RTX3090 GPU is used for training and evaluation.




Figure 4. Qualitative analysis of BAA-NGP on the blender synthetic dataset. BAA-NGP produces better quality in image synthesis with cleaner backgrounds and finer details than BARF with $10 \times$ less time.

4.3.1 Unposed Frontal Camera Video Sequences

Figure 5 shows sample results of BAA-NGP on the frontal camera video sequences from LLFF, demonstrating qualitatively the high amount of detail in the reconstructed images. The last column of Figure 5 shows the only case (orchids) where our method fails to converge. It could be partially explained by the suboptimal performance of iNGP

Table 1. Quantitative evaluation of pose refinement for the blender synthetic dataset. BAA-NGP is on par with BARF[10] in terms of camera pose estimation, with 10x faster training time, and with better visual synthesis quality overall.

Scene	Camera pose registration						Visual synthesis quality						Training Time	
	Rotation($^{\circ}$) \downarrow		Translation \downarrow		PSNR \uparrow		SSIM \uparrow		MS-SSIM \uparrow		LPIPS \downarrow		hh:mm:ss	
	BARF	ours	BARF	ours	BARF	ours	BARF	ours	BARF	ours	BARF	ours	BARF	ours
Chair	0.093	0.093	0.405	0.658	31.15	34.36	0.952	0.985	0.990	0.996	0.074	0.026	04:26:30	00:30:09
Drums	0.046	0.029	0.202	0.134	23.91	25.03	0.895	0.934	0.954	0.971	0.147	0.072	04:28:06	00:26:10
Ficus	0.080	0.032	0.464	0.161	26.26	30.27	0.930	0.979	0.975	0.991	0.109	0.026	04:21:56	00:25:34
Hotdog	0.229	0.088	1.165	0.529	34.59	37.00	0.969	0.982	0.992	0.994	0.059	0.029	04:24:52	00:26:38
Lego	0.081	0.040	0.330	0.144	28.31	32.20	0.924	0.975	0.981	0.993	0.106	0.025	04:24:21	00:26:12
Materials	0.837	1.021	2.703	4.944	27.85	27.16	0.934	0.943	0.984	0.983	0.107	0.077	04:23:13	00:23:46
Mic	0.065	0.046	0.277	0.260	31.00	34.28	0.966	0.987	0.992	0.995	0.065	0.018	04:22:25	00:32:46
Ship	0.086	0.061	0.341	0.318	27.49	29.71	0.841	0.864	0.938	0.939	0.196	0.123	04:25:59	00:23:31
Mean	0.190	0.176	0.736	0.894	28.82	32.50	0.926	0.953	0.976	0.983	0.108	0.050	04:24:40	00:26:51

Table 2. Quantitative evaluation of pose estimation for the LLFF dataset. We show that BAA-NGP is comparable to BARF with much less training time. iNGP[17] results are included for reference.

Scene	Camera pose registration						Visual synthesis quality						Training Time						
	Rotation($^{\circ}$) \downarrow		Translation \downarrow		PSNR \uparrow		SSIM \uparrow		MS-SSIM \uparrow		LPIPS \downarrow								
	[10]	ours	[10]	ours	[10]	ours	iNGP[17]	[10]	ours	iNGP[17]	[10]	ours	iNGP[17]	[10]	ours	iNGP[17]			
Fern	0.163	3.978	0.188	1.733	23.96	19.02	25.42	0.709	0.504	0.822	0.916	0.717	0.948	0.390	0.480	0.182	05:23:29	0:10:45	0:07:33
Flower	0.224	2.258	0.233	0.530	24.07	25.52	26.78	0.712	0.811	0.851	0.891	0.935	0.952	0.379	0.157	0.139	05:24:55	0:10:39	0:08:21
Fortress	0.460	0.786	0.352	0.733	28.86	28.59	28.54	0.816	0.825	0.882	0.950	0.946	0.967	0.266	0.203	0.113	05:44:05	0:11:16	0:08:16
Horns	0.135	1.042	0.162	0.643	23.12	19.57	20.47	0.734	0.724	0.755	0.915	0.879	0.845	0.423	0.307	0.262	05:22:02	0:18:54	0:07:32
Leaves	1.274	1.249	0.253	0.342	18.67	20.14	20.80	0.529	0.687	0.738	0.831	0.902	0.922	0.474	0.264	0.243	06:47:43	0:18:37	0:08:07
Orchids	0.629	6.110	0.409	2.956	19.37	12.28	19.60	0.570	0.143	0.682	0.853	0.243	0.888	0.423	0.608	0.220	05:14:03	0:10:3	0:07:26
Rooms	0.362	1.853	0.293	1.856	31.60	29.15	34.03	0.937	0.900	0.969	0.981	0.963	0.986	0.230	0.271	0.095	05:35:49	0:20:23	0:07:55
T-rex	1.030	1.749	0.641	1.005	22.32	23.41	25.18	0.771	0.861	0.892	0.927	0.950	0.960	0.355	0.185	0.121	05:10:04	0:10:37	0:07:37
Mean	0.535	2.378	0.316	1.225	24.00	22.21	25.10	0.722	0.682	0.824	0.908	0.817	0.934	0.369	0.309	0.172	05:35:16	0:13:54	0:08:53

in this data, as shown in Table 2. Quantitatively, Table 2 shows our benchmarking data against the state-of-the-art algorithms. We show that BAA-NGP can achieve more than $20\times$ speedup with similar quality as BARF, whereas our metrics are better or on par for almost all metrics except rotation. Qualitatively, our method captures more details in the reconstructed images.

4.3.2 Imperfect Poses Refinement for Multi-View Synthetic Images

Figure 4 shows the results of imperfect pose refinement for multi-view synthetic images. Qualitatively, BAA-NGP produces better-quality image synthesis with cleaner backgrounds and finer details. For quantitative analysis, as shown in Table 1, the results demonstrate that BAA-NGP not only yields similar pose estimation results and improved image synthesis quality but also achieves a more than 10-fold speedup. One observation from this result is that the Materials scenes in the dataset were more challenging, likely due to the limitations of hash grid feature encoding. The authors of iNGP reported similar findings when using ground truth camera poses, attributing this issue to the dataset’s high complexity and view-dependent reflections [17].

4.3.3 Homography Recovery

We also conducted a homography recovery experiment as an analogy to 3D cases to evaluate the effectiveness of homography matrix estimation using multi-resolution hash encoding architecture combined with the MLP. We can consider equation 4 as a general case for homography recovery if we define p_i as the parameterized warp transformation. We used the same setup as bundle-adjusting neural radiance field [10], where five warped patches are used for training as is shown in Figure 6 (L). 2D multi-resolution hash encoding and fully fused MLPs were used as the main network architecture, and coarse-to-fine feature weighting was used for the best performance. For this experiment, $N_{min} = 3, L = 18$, coarse-to-fine $r_s = 0.1, r_e = 0.5$, learning rate for features and parameters was $1.e - 2$ and for warp parameters was $3.e - 3$. Adam optimizer was used for both learning processes, and training was conducted over 5000 epochs.

We show in Figure 6 (R) that without signal coarse-to-fine smoothing, the learned features of hash encoding converge without learning any warping functions, which end up learning the most overlapped area. Our coarse-to-fine implementation can achieve better pose estimation comparing to vanilla coarse-to-fine masking, which can be trapped in local minimal.




Figure 5. Qualitative analysis of BAA-NGP on the LLFF dataset. We show that our results are on par with BARF’s results but converge $20\times$ faster.




Figure 6. Homography recovery on a cat image. (L) Different colors indicate five different ground truth warped patches. (R) Results of different coarse-to-fine weighting procedures. Our coarse-to-fine weighting scheme outperforms the vanilla coarse-to-fine (c2f) weighting scheme qualitatively for both 2D homography recovery.

5. Conclusions

In conclusion, we presented a method for learning INRs of scenes with unknown or poorly known camera poses while achieving learning rates in minutes. BAA-NGP, there-

fore, is a solution for addressing accelerated learning of INR models in unstructured settings. Our evaluations over several benchmark datasets, including multi-view object-centric scenes as well as frontal-camera video sequences of unbounded scenes, show that it is either comparable or outperforms state-of-the-art techniques such as BARF and COLMAP-based methods where camera poses are not known or are known imprecisely, and where training would otherwise take hours. Thus, this approach opens learning INRs for a broad set of real-world scenarios and applications ranging from virtual and augmented reality to robotics and automation, where time and unstructured image capture are vital.

One observation is that when our rendering baseline iNGP leads to suboptimal results, the corresponding results of BAA-NGP will also be influenced. In the future, we plan to improve this technique and integrate BAA-NGP into robotic systems to provide fast and accurate environment modeling for real-time applications.

Acknowledgment

We express our heartfelt appreciation to Subarna Tripathi, Ilke Demir, Saurav Sahay, along with numerous other colleagues from Intel Labs as well as Shreya Saha from UCSD. Their insightful contributions, stimulating discussions, and generous provision of computing resources have been invaluable to our project.

References

- [1] Michal Adamkiewicz, Timothy Chen, Adam Caccavale, Rachel Gardner, Preston Culbertson, Jeannette Bohg, and Mac Schwager. Vision-only robot navigation in a neural radiance world. *IEEE Robot. Autom. Lett.*, 7(2):4606–4613, 2022. 1
- [2] Yue Chen, Xingyu Chen, Xuan Wang, Qi Zhang, Yu Guo, Ying Shan, and Fei Wang. Local-to-global registration for bundle-adjusting neural radiance fields. In *CVPR*, pages 8264–8273, 2023. 2
- [3] Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Garf: Gaussian activated radiance fields for high fidelity reconstruction and pose estimation. *ECCV*, 2022. 1, 2, 3
- [4] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH*, pages 303–312, 1996. 1
- [5] Qiyu Dai, Yan Zhu, Yiran Geng, Ciyu Ruan, Jiazhao Zhang, and He Wang. GraspNeRF: multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf. In *ICRA*, pages 1757–1763, 2023. 1
- [6] Jiemin Fang, Lingxi Xie, Xinggang Wang, Xiaopeng Zhang, Wenyu Liu, and Qi Tian. Neusample: Neural sample field for efficient view synthesis. *arXiv*, 2021. 2
- [7] Hwan Heo, Taekyung Kim, Jiyoung Lee, Jaewon Lee, Soohyun Kim, Hyunwoo J. Kim, and Jin-Hwa Kim. Robust camera pose refinement for multi-resolution hash encoding. *ICML*, 2023. 2, 3
- [8] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *ICCV*, pages 5846–5854, 2021. 2
- [9] Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: Efficient sampling accelerates nerfs. *arXiv*, 2023. 5
- [10] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. *ICCV*, 2021. 1, 2, 3, 4, 5, 6
- [11] Yunzhi Lin, Thomas Müller, Jonathan Tremblay, Bowen Wen, Stephen Tyree, Alex Evans, Patricio A. Vela, and Stan Birchfield. Parallel inversion of neural radiance fields for robust pose estimation. In *ICRA*, 2023. 2
- [12] Dominic Maggio, Marcus Abate, Jingnan Shi, Courtney Mario, and Luca Carlone. Loc-NeRF: Monte Carlo localization using neural radiance fields. In *ICRA*, pages 4018–4025, 2023. 1
- [13] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *ICCV*, pages 6351–6361, 2021. 2
- [14] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM TOG*, 38(4):1–14, 2019. 2
- [15] 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. In *ECCV*, 2020. 1, 2
- [16] Thomas Müller. tiny-cuda-nn, 2021. 5
- [17] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM TOG*, 41(4):102:1–102:15, 2022. 1, 4, 6
- [18] Martin Píala and Ronald Clark. Terminerf: Ray termination prediction for efficient neural rendering. In *3DV*, pages 1106–1114, 2021. 2
- [19] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. *CVPR*, 2016. 1, 2
- [20] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. *ECCV*, 2016. 1, 2
- [21] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew Davison. iMAP: Implicit mapping and positioning in real-time. In *ICCV*, 2021. 2
- [22] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, pages 5459–5469, 2022. 2
- [23] Zhenggang Tang, Balakumar Sundaralingam, Jonathan Tremblay, Bowen Wen, Ye Yuan, Stephen Tyree, Charles Loop, Alexander Schwing, and Stan Birchfield. RGB-only reconstruction of tabletop scenes for collision-free manipulator control. In *ICRA*, pages 1778–1785, 2023. 1
- [24] Matthias Teschner, Bruno Heidelberger, Matthias Müller, Danat Pomeranets, and Markus Gross. Optimized spatial hashing for collision detection of deformable objects. In *Proceedings of VMV’03, Munich, Germany*, 2003. 4
- [25] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv*, 2021. 2
- [26] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yu-Xiong Wang, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering. In *CVPR*, pages 16200–16209, 2022. 2
- [27] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. iNeRF: Inverting neural radiance fields for pose estimation. *arxiv*, 2020. 2
- [28] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *CVPR*, 2022. 2
- [29] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *CoRR*, abs/2010.07492, 2020. 4
- [30] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CVPR*, 2018. 5
- [31] 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 *CVPR*, 2022. 2
- [32] Zhenxin Zhu, Yuantao Chen, Zirui Wu, Chao Hou, Yongliang Shi, Chuxuan Li, Pengfei Li, Hao Zhao, and Guyue Zhou. LATITUDE: Robotic global localization with truncated dynamic low-pass filter in city-scale nerf. In *ICRA*, pages 8326–8332, 2023. 1