

# Tiny-DroNeRF: Tiny Neural Radiance Fields aboard Federated Learning-enabled Nano-drones

Ilenia Carboni<sup>1,3\*</sup>, Elia Cereda<sup>1\*</sup>, Lorenzo Lamberti<sup>1,2\*</sup>, Daniele Malpetti<sup>1</sup>, Francesco Conti<sup>3</sup>, Daniele Palossi<sup>1,2</sup>

**Abstract**—Sub-30 g nano-sized aerial robots can leverage their agility and form factor to autonomously explore cluttered and narrow environments, like in industrial inspection and search and rescue missions. However, the price for their tiny size is a strong limit in their resources, i.e., sub-100 mW microcontroller units (MCUs) delivering  $\sim 100$  GOPs/s at best, and memory budgets well below 100 MB. Despite these strict constraints, we aim to enable complex vision-based tasks aboard nano-drones, such as dense 3D scene reconstruction: a key robotic task underlying fundamental capabilities like spatial awareness and motion planning. Top-performing 3D reconstruction methods leverage neural radiance fields (NeRF) models, which require GBs of memory and massive computation, usually delivered by high-end GPUs consuming 100s of Watts. Our work introduces Tiny-DroNeRF, a lightweight NeRF model, based on Instant-NGP, and optimized for running on a GAP9 ultra-low-power (ULP) MCU aboard our nano-drones. Then, we further empower our Tiny-DroNeRF by leveraging a collaborative federated learning scheme, which distributes the model training among multiple nano-drones. Our experimental results show a 96% reduction in Tiny-DroNeRF’s memory footprint compared to Instant-NGP, with only a 5.7 dB drop in reconstruction accuracy. Finally, our federated learning scheme allows Tiny-DroNeRF to train with an amount of data otherwise impossible to keep in a single drone’s memory, increasing the overall reconstruction accuracy. Ultimately, our work combines, for the first time, NeRF training on an ULP MCU with federated learning on nano-drones.

## SUPPLEMENTARY VIDEO MATERIAL

Supplementary video can be found at:  
<https://youtu.be/-frFPUBGa0c>

## I. INTRODUCTION

Autonomous flying miniaturized nano-drones, with a few tens of grams of weight and sub-10 cm diameter, can cope with cluttered and narrow spaces, such as hazardous industrial plants [11] or post-disaster environments [12]. Thanks to their autonomy, i.e., all processing is carried onboard with no external infrastructure, they are not limited to line-of-sight operations or to a maximum distance from the base station. Nano-drones equipped with small cameras and onboard intelligence have been demonstrated in simple inspection and detection tasks, including anomaly detection [13], semantic

localization [14], and object detection [15]. Their effectiveness further increases in swarm configurations, where multiple drones cover larger areas simultaneously or collaborate toward a shared goal [16], such as survivor search. However, more complex autonomous workloads are still infeasible on this class of tiny robots. A primary reason is their lack of spatial awareness, stemming from the limited computational and sensory resources they can carry. Compared to kg-scale flying drones, a nano-drone can rely on roughly  $1000\times$  less memory and computation [17], and simpler sensors, such as low-resolution monochrome monocular cameras.

Neural radiance fields (NeRF) [1] are a well-established paradigm for dense 3D reconstruction tasks. Enabling NeRF on nano-drones would unlock the ability to autonomously perform novel view synthesis [1], SLAM [18], 3D surface reconstruction [19], or 3D semantic scene reconstruction [20]. However, current NeRF algorithms demand memory and compute resources that exceed nano-drones’ capabilities by orders of magnitude [5], due to their complexity and the need for scene-specific retraining. Therefore, they are generally limited to offline execution on GPUs, with GBs of memory available, and consuming hundreds of Watts.

In this work, we address the task of novel view synthesis using NeRF under the constraints of the limited resources of an ultra-low-power MCU suitable for miniaturized nano-drones. Taking advantage of the tiny size of nano-drones, we focus on the operation of a swarm, where each nano-drone takes a few images of the target object from different known poses and trains a separate NeRF instance onboard. Then, thanks to a federated learning approach [21], we iteratively aggregate the individual NeRF instances to create a global NeRF producing a richer 3D reconstruction than any single drone can provide. Overall, our work enables the cooperative construction of a continuous 3D representation of the scene, enabling the generation of photorealistic views from previously unseen camera poses using images captured from non-overlapping drone viewpoints.

In detail, our paper’s contributions are:

- a lightweight NeRF algorithm, called *Tiny-DroNeRF*, based on the State-of-the-Art (SoA) Instant-NGP [5]. Through hyperparameter optimization, we seek the best trade-off between 3D reconstruction quality, i.e., peak signal to noise ratio (PSNR), memory, and computation.
- A memory-efficient implementation of Tiny-DroNeRF, targeting the ultra-low-power multi-core GreenWaves Technologies (GWT) GAP9 MCU aboard a nano-drone.
- A detailed experimental analysis of *i)* the reconstruction accuracy on standard synthetic datasets [1] and on a

This work was partially supported by the SNSF RoboMix2 project (grant nb. 10004854).

\* denotes equal contributions.

<sup>1</sup>I. Carboni, E. Cereda, L. Lamberti, D. Malpetti and D. Palossi are with the Dalle Molle Institute for Artificial Intelligence, USI-SUPSI, Switzerland. Corresponding author: [ilenia.carboni@supsi.ch](mailto:ilenia.carboni@supsi.ch).

<sup>2</sup>L. Lamberti and D. Palossi are also with the Integrated Systems Laboratory, ETH Zürich, Switzerland.

<sup>3</sup>I. Carboni and F. Conti are with the Department of Electrical, Electronic and Information Engineering, University of Bologna, Italy.

TABLE I: Comparison of novel view synthesis methods on the Synthetic NeRF 360° dataset [1].

Method	Year	Neural	Key contributions	PSNR [dB]	SSIM [-]	Params [MB]	Training time	Power class [W]	Device
NeRF [1]	2021	✓	Coordinate-based MLP	31.01	0.947	5	24 h	10 <sup>2</sup>	GPU
SNeRG [2]	2021	✓	Neural grid	30.38	0.950	1772	15 h	10 <sup>2</sup>	GPU
PlenOctrees [3]	2021	✓	Octree + Spherical harmonics	31.71	0.958	1930	15 h	10 <sup>2</sup>	GPU
Plenoxels [4]	2022	×	Voxel grid + neural network-free	31.71	0.958	778	11 min	10 <sup>2</sup>	GPU
Instant-NGP [5]	2022	✓	Parametric hash-grid	<b>33.18</b>	0.959	46	5 min	10 <sup>2</sup>	GPU
TensorRF [6]	2022	×	Voxel grid w/ tensor factorization	33.14	<b>0.963</b>	72	17 min	10 <sup>2</sup>	GPU
K-Planes [7]	2023	×	Multi-planes	32.21	0.960	385	38 min	10 <sup>2</sup>	GPU
MobileNeRF [8]	2023	✓	Rasterized mesh (baked)	30.90	0.947	538	27 h	10 <sup>2</sup>	GPU
BiRF [9]	2023	✓	3D-2D binarized feature hash encoding grid	32.03	0.954	0.5	> 5 min	10 <sup>2</sup>	GPU
CNC [10]	2024	✓	Context-based compression framework	33.19	0.964	0.418	> 5 min	10 <sup>2</sup>	GPU
Tiny-DroNeRF (ours)	2025	✓	Instant-NGP-based, tailored for MCUs	26.94	0.892	<b>0.284</b>	<b>28s</b>	10 <sup>2</sup>	GPU
				26.94	0.892	<b>0.284</b>	97 min	< <b>0.10</b>	MCU

novel real-world nano-drone dataset, and *ii*) the execution time and memory usage on the GAP9 MCU.

- A federated learning scheme for distributed training of Tiny-DroNeRF across a swarm of nano-drones.

In our experiments, Tiny-DroNeRF drastically cuts resource requirements while maintaining useful reconstruction quality. Compared to Instant-NGP [5], we reduce memory usage by 96% (from 527 MB to 21.4 MB), while losing only 5.7 dB of PSNR. We complete a full forward-backward step in 73 ms, i.e., 97 min for a complete training on a single nano-drone (10 k steps). Through our efficient hardware-aware algorithm, we succeed in making the workload compute-bound. Our implementation exhibits excellent memory locality, requiring only 3.6 MB of data transfers to/from off-chip memory per computational step. In field tests, we reconstruct a real-world scene from the nano-drone’s onboard grayscale images at up to 21.2 PSNR.

Moreover, we demonstrate that our federated learning scheme allows each nano-drone, which carries only a limited set of local images, to synthesize unseen portions of the scene by leveraging the knowledge of its peers, without exchanging raw training images. This federated approach achieves a PSNR just 0.6 dB below a centralized training with access to all images, while requiring only 2.24 s of communication overhead per merge (done every 1 k local training steps). To the best of our knowledge, this is the first demonstration of on-device NeRF tailored to highly constrained computing devices and the first application of a federated learning scheme on a swarm of resource-constrained nano-drones.

## II. RELATED WORK

### A. Novel view synthesis methods

Neural radiance fields [1] introduced a paradigm for representing 3D scenes as continuous functions, enabling photorealistic novel-view synthesis and high-fidelity scene reconstruction. This approach established implicit neural representations, in which a multi-layer perceptron (MLP) encodes a continuous volumetric scene that can be queried

at arbitrary 3D coordinates, and made this a widely adopted technique. Nevertheless, vanilla NeRF suffers from slow training and rendering, motivating extensive research aimed at accelerating both optimization and inference.

To address these limitations, subsequent work has explored explicit scene parameterizations as a means of improving efficiency. For example, PlenOctrees [3] precompute a sparse octree structure that stores radiance field parameters for fast querying, while Plenoxels [4] and TensorRF [6] replace neural networks with voxel grids or tensor factorization schemes to reduce computational overhead. Instant-NGP [5] introduces a multiresolution hash-grid positional encoding and integrates it into a CUDA/C++ library to accelerate the training and rendering of neural graphics primitives, establishing itself as a de facto standard for efficient, high-performance NeRFs.

Complementary research has investigated reducing memory requirements while preserving reconstruction fidelity. BiRF [9] and CNC [10] propose compact network architectures that trade memory footprint for additional computation. While they better accommodate memory-constrained devices, the additional computation introduced by the more complex encoding strategies increases the execution time. While instances of rendering on commodity GPUs can be found in literature, such as SNeRG [2] and MobileNeRF [8], training requires access to powerful GPUs in existing methods, a significant barrier for NeRF-based approaches on embedded robotics devices. Table I summarizes novel synthesis view methods, reporting their key contributions, average performance on the Synthetic NeRF 360° dataset [1], measured with PSNR and structural similarity index measure (SSIM), parameter count, training time, and power consumption.

### B. Federated learning for neural radiance fields

Federated learning (FL) [21] is a collaborative machine learning framework in which multiple parties jointly train a shared model without sharing their raw data. Instead of exchanging data, participants send model updates to a coordinator, which may be either a dedicated server or a

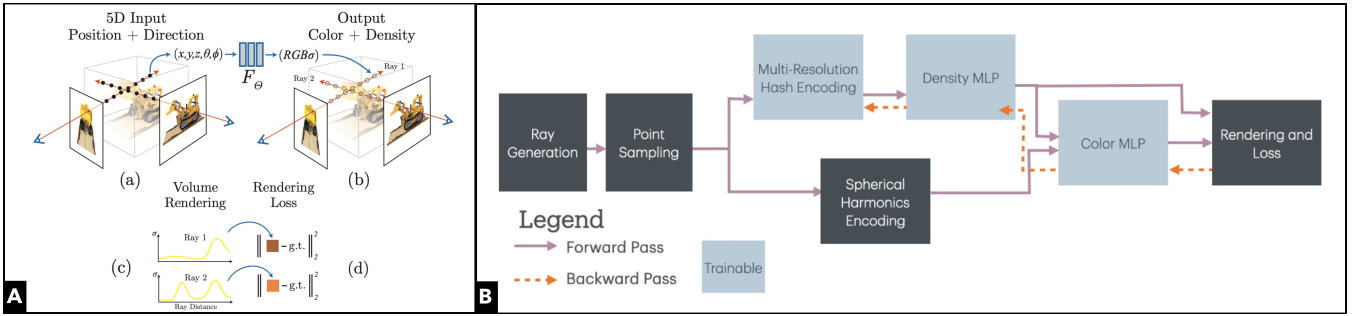


Fig. 1: A) Novel view synthesis with neural radiance fields and volumetric rendering [1]. B) The InstantNGP algorithm.

peer that also assumes the coordinating role. Training proceeds iteratively, alternating between local training phases, in which each participant updates the model on its private data, and communication phases, during which the coordinator aggregates the received updates into a global model and redistributes it to all participants to initialize the next round.

The first federated implementation of NeRF was proposed in FedNeRF [22], followed by studies investigating FL-based approaches for large-scale outdoor scene modeling [23], distributed multi-robot training with joint pose refinement [24], and planetary multi-agent mapping [25]. These contributions remain exploratory, with no reported in-field deployment or detailed resource feasibility analysis.

Across these works, the main motivations for adopting FL have been scalability to larger datasets [23], [25], [26], maintainability of a dynamic process [23], and communication efficiency [24], [25]. In contrast, in this work we employ FL primarily as a mechanism to enable a swarm of resource-constrained drones to collaboratively learn a model that outperforms what any single drone could achieve independently, thereby providing a stronger model for downstream tasks.

### III. BACKGROUND

#### A. NeRF overview

A *radiance field* represents a scene through an implicit function that maps spatial points  $\mathbf{x}$  and view direction  $\mathbf{d}$  to the volume density  $\sigma$  and the emitted color  $\mathbf{c}$ :

$$f_{\theta}(\mathbf{x}) = (\mathbf{c}, \sigma) \quad (1)$$

*NeRFs* approximate this function through two MLP networks, parametrized by  $\theta = \{\tau, \phi\}$ . Instant-NGP [5], shown in Figure 1, employs a *density MLP* to predict the volume density  $\sigma$  and a vector of hidden features  $\mathbf{h}$  as a function of the 3D location  $\mathbf{x}$ :

$$f_{\tau}(\gamma(\mathbf{x})) = (\mathbf{h}, \sigma), \quad (2)$$

with 3D location represented with a position encoding  $\gamma(\mathbf{x})$ .

A *color MLP* predicts the color  $\mathbf{c}$  as a function of hidden features and viewing direction,

$$f_{\phi}(\mathbf{h}, \gamma(\mathbf{d})) = \mathbf{c}. \quad (3)$$

As with the 3D location, also the viewing direction is encoded as  $\gamma(\mathbf{d})$ .

The main contribution of Instant-NGP is the *multi-resolution hash encoding* [5] (MRHE) to map the 3D position  $\mathbf{x}$  into a high-dimensional parametric feature vectors  $\gamma(\mathbf{x})$ . This trainable encoding is arranged into  $L$  independent levels, each subdividing space in voxel grids of increasing resolution. Each voxel is mapped to a hash table of up to  $T$  feature vectors in  $\mathbb{R}^B$ . A given location  $\mathbf{x}$  is encoded to  $\gamma(\mathbf{x}) \in \mathbb{R}^{L \times F}$ , with trilinear interpolation from each layer. As the resolution increases in deeper levels, multiple voxels map to the same feature vector in the fixed-size hash table but the subsequent density MLP learns to disambiguate the resulting aliasing. This key insight makes  $T$  a scalable parameter we can scale to fit the tight memory requirements of our nano-drone. Direction  $\mathbf{d}$  is encoded as well, albeit with a coordinate-based (i.e., non trainable) encoding function  $\gamma(\mathbf{d})$ , the *spherical harmonics*.

During training, *samples*, composed of position and view directions, are generated by marching a set number of *rays*  $R$  through random pixels of random images from the train set and sampling points along the ray. A *batch* of  $B$  samples, fed to the MLP, produces the sample colors and volume densities, which are aggregated using *volume rendering* resulting in the color value of each ray,  $\hat{\mathbf{c}}$ . Training optimizes the Huber loss of the ray color against the ground-truth pixel color using the Adam optimizer. To reduce computation wasted on empty space, Instant-NGP fast forwards ray marching across empty regions through an additional data structure, the *occupancy grid*. The occupancy grid is updated alternatively to network training, by querying density predicted by the density MLP in each cell of the reconstruction volume.

#### B. Robotic platform

Our robotic platform encompasses the Bitcraze Crazyflie 2.1 Brushless quadrotor (Figure 2-A), a commercially available nano-drone that weighs 27 g and measures 10 cm in diameter. The drone is powered by a 350 mAh LiPo battery. We extend this platform with the open-source *GAP9Shield* [27], a pluggable printed circuit board that integrates the GWT GAP9 System-on-Chip (SoC), 32 MB of external L3 HyperRAM, a ublox NINA-W102 Wi-Fi transceiver (based on ESP32) providing up to 15 Mbit/s of bandwidth, and a low-power miniaturized camera.

The GAP9 SoC, depicted in Figure 2-B, is organized into two power and frequency domains: the SoC domain,

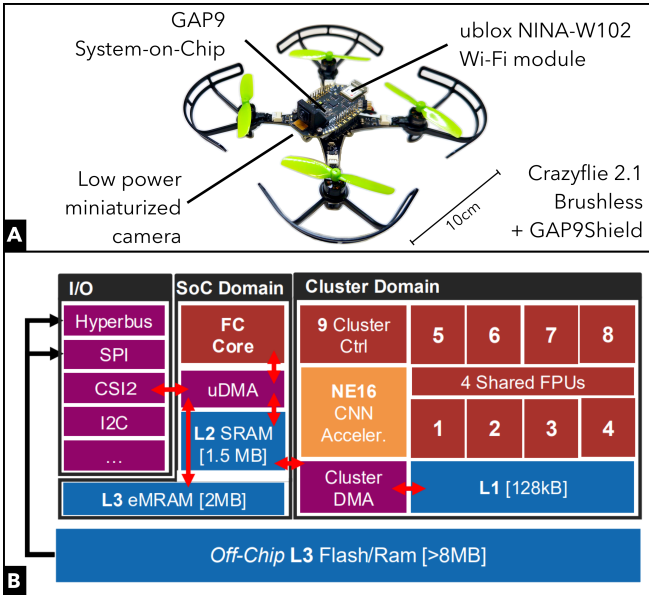


Fig. 2: Target robot platform. A) Bitcraze Crazyflie 2.1 Brushless with GAP9Shield expansion deck. B) GWT GAP9 System-on-Chip architecture.

featuring a single main core called Fabric Controller (FC), and the Cluster, featuring 9 general-purpose cores, 4 floating-point units (FPUs) supporting 16- and 32-bit data types, and a mixed-precision AI hardware accelerator (NE16). The Cluster cores support single-instruction multiple-data (SIMD) operations: the 9 general-purpose cores provide four-lane 8-bit integer SIMD, while the FPUs support two-lane 16-bit SIMD, in addition to scalar 32-bit floating-point operations. The NE16 AI hardware accelerator provides a peak throughput of 150 integer GOP/s. However, it is designed for the forward pass of convolutional neural networks with integer data types, while our workload requires backpropagation and floating-point arithmetic. In all experiments, we operate GAP9 at its peak performance configuration, clocked at 370 MHz and powered at 0.8 V, which translates to 60 mW of average power consumption.

The SoC features a hierarchical memory system with 128 kB of L1 memory shared across the Cluster and 1.5 MB of L2 RAM accessible by both the FC and the Cluster. The GAP9Shield further extends the system with 32 MB of off-chip L3 OctaSPI RAM. Cluster cores access L1 in around 1 cycle, while direct accesses to L2 incur an additional latency of about 100 cycles. Two dedicated direct memory access (DMA) engines enable asynchronous data transfers, freeing the SoC cores to perform computation: the main DMA handles L3–L2 transfers (370 MB/s), while the cluster DMA handles L2–L1 transfers (13.3 GB/s). By overlapping DMA transfers with computation, L2 data can be pre-fetched into L1 in the background, allowing the Cluster to operate exclusively on L1 and hide the L2 access latency.

## IV. METHODOLOGY

### A. Hardware-aware memory-efficient NeRF

As our robotic platform (Section III-B) relies on a resource-constrained MCU with very limited on-chip and off-chip memory, the memory and compute demands of SoA novel view synthesis algorithms must be drastically reduced. In its baseline configuration, Instant-NGP requires 527 MB of memory and 17.5 GOPs per training step. The memory footprint, in particular, includes storage for both the model’s learnable weights and the intermediate tensors needed for backpropagation. To shrink both peak memory footprint and operations per step, we perform a thorough hyperparameter optimization on Instant-NGP [5], which targets *i*) the number of images per training step (img/step), *ii*) input image resolution, *iii*) batch size ( $B$ ), and *iv*) maximum number of parameters per level of the MRHE ( $T$ ).

We first analyze the number of images per training step (img/step). At each step, the baseline samples randomly from all 100 training images of the NeRF Synthetic dataset [1]. On our MCU, storing 100 images would exceed on-chip capacity, forcing frequent off-chip accesses with much higher latency per access (see Section III-B). Thus, minimizing the number of img/step (i.e., to 1 img/step) is desirable to reduce the memory transfer latency. Next, we study the effect of input image resolution. Since our drone uses a low-resolution camera, we evaluate smaller inputs compared to the  $800 \times 800$  px images of NeRF Synthetic. While Instant-NGP’s operations and peak memory do not directly depend on input resolution, larger images increase storage requirements because they must be stored in memory (typically off-chip) for the entire training process. Reducing the input size to  $160 \times 160$  px lowers L3 memory utilization by 61.44 MB.

We then ablate different batch sizes ( $B$ ). The peak memory footprint of Instant-NGP is dominated by the intermediate feature maps of the MLPs. Reducing  $B$  linearly decreases both the number of operations per training step and the size of the output feature maps, significantly lowering peak memory requirements. Last, we evaluate the effect of  $T$ . This parameter has negligible impact on the number of operations, since the number of hash queries per coordinate remains constant. However, it directly influences memory usage, as smaller  $T$  values decrease the size of the hash tables. Overall, the batch size  $B$  is the primary factor controlling the number of operations per step, and thus total latency, while both  $B$  and  $T$  strongly affect Instant-NGP’s total memory footprint.

### B. NeRF deployment

After optimizing Instant-NGP to reduce its operations and memory footprint (a hard constraint of 32MB to deploy the model on our resource constrained MCU) through our shrinking methodology IV-A, we need an optimized embedded C implementation of the Tiny-DroNeRF algorithm. The strict resource limitations of our platform require careful structuring of computation and memory transfers to fully exploit the computational capabilities of the GAP9.

Our embedded implementation is composed of *i*) hand-written kernels for each algorithmic block, *ii*) a tiling strategy, i.e., subdividing processed data into tiles that fit in local memory, and *iii*) batching, to unlock efficient parallel execution on the multi-core cluster. We map each block of the Instant-NGP pipeline (Figure 1) to a custom hand-written kernel, which we parallelize across the GAP9 cluster. The baseline Instant-NGP uses `float16` arithmetic (except for the weight updates, accumulated in `float32` for stability [28]), which allows us to exploit GAP9’s optimized SIMD instructions for efficient computation of the MLP.

Although the configuration identified by the shrinking methodology fits within the total memory available on the nano-drone, we must ensure that each kernel’s inputs, outputs, and intermediate tensors reside in the 128 kB L1 memory. Otherwise, computation would stall on memory transfer bandwidth. To address this, we design a tiling strategy: subdividing kernel data into tiles that fit in L1 and scheduling memory transfers to overlap with computation, thereby keeping the cores consistently fed with data.

All kernels operate independently on batches of samples or rays, so memory accesses along these dimensions offer straightforward opportunities for parallelization and tiling. However, MRHE and occupancy grid would still not fit in the available L2; therefore, we leverage their hierarchical structure to move levels to L2 memory one at a time.

As the L3 memory bandwidth is about  $35\times$  lower than L2, minimizing L3 accesses is critical. Since most input data (rays and samples) is generated randomly, we restrict sampling to the maximum amount that can fit entirely within the on-chip L2 memory. However, training with excessively small batch sizes  $B$  hinders convergence: it requires more steps and typically leads to lower reconstruction quality. To address this, we employ *batch accumulation*, which decouples the size of computation batch tiles from the effective training batch size. In practice, we process batch tiles sized to fit in L2, while accumulating weight gradients over multiple steps until the desired effective batch size  $B$  is reached.

### C. Federated implementation

For the federated training of Tiny-DroNeRF, we employ a FL scheme in which one drone assumes a dual role, contributing to training with its own local data while simultaneously acting as the coordinator. The coordinator first performs a short pre-training phase of  $N_{pt}$  steps on its local dataset and subsequently shares the resulting initial global model with the rest of the swarm.

During each communication round, all drones, including the coordinator, perform local training for  $N_\ell$  steps and transmit their updated model parameters to the coordinator. We investigate two variants for parameter exchange: (i) transmitting the occupancy grid, hash table, and MLP parameters at every round, or (ii) transmitting only the hash table and MLP parameters, allowing each drone to update its occupancy grid locally.

Model aggregation is carried out using the FedAvg algorithm [21], which computes a weighted average of the

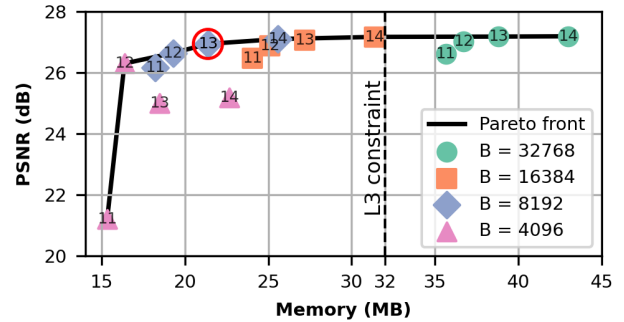


Fig. 3: Memory vs. reconstruction quality Pareto curve. Each point corresponds to a configuration of batch size  $B$  (marker) and encoding size  $\log_2(T)$  (number). Circled in red, the configuration selected for deployment.

received parameters. In our case, since all drones possess an equal number of training images in our experiments, the weighted average simplifies to an arithmetic mean. The resulting global model is then broadcast to all drones, and the procedure is repeated over multiple communication rounds.

We evaluate federated training under two data partitioning setups. In the first setup, independent and identically distributed (IID), the training images are randomly shuffled and evenly partitioned across clients, so that each drone receives a representative sample of the entire dataset. In the second setup, non-IID, we partition the training images by viewpoint sectors, so that each drone observes only a specific portion of the scene, leaving some sides completely unseen.

## V. EXPERIMENTAL RESULTS

### A. Experimental setup

We validate our models on the NeRF Synthetic 360° dataset [1], which contains 8 different objects, each one with 100 training images and 200 test images at a resolution of  $800\times 800$ . To match the resolution of our drone cameras, we decrease the image resolution to  $160\times 160$ . For evaluation, we follow prior work and report PSNR and SSIM [1].

### B. Hardware-aware memory-efficient NeRF

We assess how our shrinking methodology, through hyperparameter optimization, affects our system’s key parameters, such as total memory footprint, operations per update step, and PSNR and SSIM reconstruction metrics. Table II explores various Instant-NGP’s configurations, sweeping images per step, their resolution, the batch size ( $B$ ), and the maximum number of parameters for a single level of the hash encoding ( $T$ ). For each table row, we keep the maximum PSNR and SSIM over five independent runs. As expected, the memory required for the baseline Instant-NGP [5] (first row) would vastly exceed the total memory on our nano-drone, i.e., 527 MB vs. 32 MB at most. Similarly, its training time would exceed two days, due to the 17.52 GOPs per step.

A first optimization comes with the reduction in the number of images per step stored in the on-chip memory,

TABLE II: Results of the hyperparameter optimization starting from the Instant-NGP model [5]. In bold, the configuration selected for deployment.

Img / Res step [px]	$B$	$\log_2(T)$	Mem [MB]	Compute [GOp/step]	PSNR [dB]	SSIM [-]
100 800	264k	19	527	17.52	32.59	0.96
1 160	32k	14	43	2.64	27.19	0.89
		13	39	2.62	27.18	0.90
		12	37	2.61	27.03	0.89
		11	36	2.61	26.61	0.89
1 160	16k	14	31	1.65	27.17	0.89
		13	27	1.63	27.07	0.89
		12	25	1.62	26.87	0.89
		11	24	1.62	26.46	0.89
<b>1 160</b>	<b>8k</b>	14	26	1.15	27.10	0.89
		<b>13</b>	<b>21</b>	<b>1.14</b>	<b>26.94</b>	<b>0.89</b>
		12	19	1.13	26.64	0.89
		11	18	1.11	26.17	0.88
1 160	4k	14	23	0.91	25.18	0.78
		13	18	0.89	25.00	0.78
		12	16	0.88	26.30	0.89
		11	15	0.88	21.20	0.56

passing from 100 to 1 image reduces memory from 2.56 MB to 0.03 MB, with a slight drop in PSNR of 3.6%. To keep at least one level of the hash grid in our SoC’s L2 memory, we need to reduce  $T$  at least to  $2^{14}$ , despite all image resolution and batch size reductions. Then, starting from this configuration (second row), which drops 23% and 7% in PSNR and SSIM, respectively, compared to the baseline, we perform a more fine-grained exploration of  $B$  and  $T$ .

All configurations under investigation require the use of off-chip L3 memory. Therefore, we identify the configuration with  $B = 8k$  and  $T = 2^{13}$  as the best trade-off, circled in Figure 3. We name this configuration *Tiny-DroNeRF*.

### C. Deployment

We analyze the execution time of *Tiny-DroNeRF* training on the GAP9 SoC. We hand-write all computation kernels, designing parallelization and tiling strategies according to Section IV-B and profile their individual execution time on the cycle-accurate GAP9 GVSOC simulator. We adopt  $8 \times$  *batch accumulation*, processing 1024-sample batch tiles at a time and accumulating gradients up to the selected  $B = 8k$  effective batch size. Peak on-chip memory usage across the entire pipeline remains within the tight GAP9 envelope. The maximum allocations are 1.3 MB in L2 and 112 kB in L1, fitting our on-chip memory limits (respectively 1.5 MB and 128 kB). Per-step L3 transfers amount to only 3.6 MB, confirming excellent memory locality and ensuring that the transfer latency (9.7 ms) does not dominate the step computation time (73 ms).

Figure 4 reports the execution latency (blue) and L2 memory usage (red) of each kernel on a single 1024-sample batch tile. The latency amounts to 12.8 ms for the

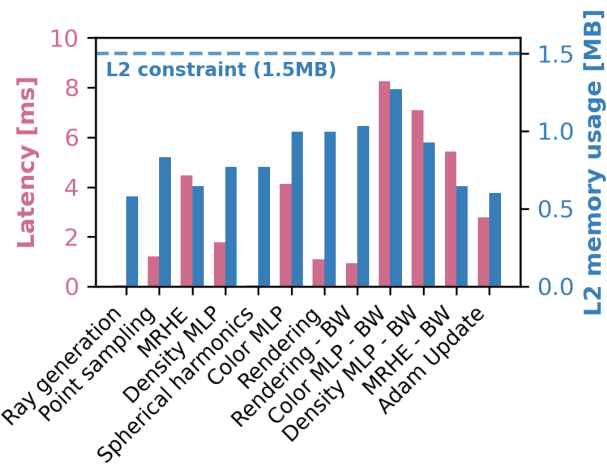


Fig. 4: Kernel latencies and peak on-chip memory usage per training step over a 1024-sample batch tile.

forward and 21.7 ms for the backward pass, while Adam optimizer updates take 22.4 ms every 8 step. The backward pass of the color MLP dominates both latency and peak L2 usage, setting the global memory bound. Other stages, such as MRHE and density MLP backward, contribute less but remain non-negligible. We also perform occupancy grid updates [5] every 256 steps, with a latency of 6.4 s (not shown in Figure 4). The amortized latency of a training step is thus 73 ms, which translates to 97 min for a complete reconstruction on a single nano-drone (8 batch tiles  $\times$  10 k training steps). On the other hand, rendering novel views of the scene (i.e., inference) requires only forward passes. The rendering latency scales with the scene-dependent samples per pixel [5]. In our experiments on NeRF Synthetic 360°, rendering a  $160 \times 160$  px image requires on average 13 samples/pixel, corresponding to 325 forward passes over 1024-sample batch tiles. Given a forward pass latency of 12.8 ms per batch tile, the total rendering time amounts to approximately 4.1 s.

Overall, our profiling highlights the effect of memory-centric optimizations, such as tiling and batch accumulation, that keep latency computation-bound rather than data-movement-bound.

### D. In-field NeRF

We evaluate *Tiny-DroNeRF* in a real deployment scenario by acquiring a new dataset with our nano-drone in an indoor environment featuring narrow walls and two traffic cones, as shown in Figure 6-A. The drone flies across the space while recording video from its onboard camera. We acquire  $160 \times 160$  px monochrome camera images at 30 Hz, streamed over Wi-Fi and post-processed with COLMAP [29] to extract ground-truth camera poses. The resulting dataset contains 1954 frames, split into 1709 train and 245 test images. Compared to the earlier experiments on NeRF Synthetic 360°, this in-field dataset introduces two additional challenges. First, images are grayscale rather than RGB, making reconstruction

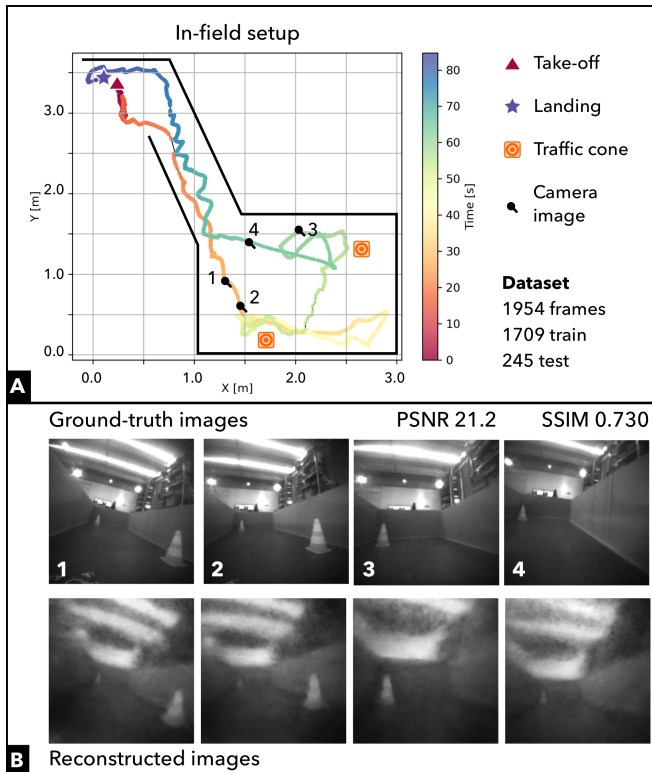


Fig. 5: A) In-field data collection setup. B) Samples of ground-truth images and the respective reconstructions.

harder. Second, the environment is cluttered and extends up to 5 m in depth, unlike the synthetic dataset which features a single object with transparent background. This requires the model to reconstruct a larger scene volume.

We train Tiny-DroNeRF on our in-field dataset and evaluate its reconstruction quality using PSNR and SSIM. Our model achieves a PSNR of 21.2 and an SSIM of 0.73. Figure 5-B illustrates qualitative results based on our supplementary video. Ground-truth images are shown on the top, while reconstructions from Tiny-DroNeRF are shown on the bottom. We show that the reconstructed views capture the overall geometry of the corridor and the two cones in the environment, demonstrating that Tiny-DroNeRF can reconstruct real-world indoor scenes from grayscale images captured onboard our nano-drone.

### E. Federated learning

We evaluated the federated implementation of Tiny-DroNeRF in two distinct setups using the NeRF 360° Lego dataset [1]. For both the IID and non-IID setups, we adhered to the dataset’s official train–test split, evenly distributing the 100 training images across four drones (25 per drone) and reserving the 200 test images exclusively for evaluation. The corresponding data distributions are shown in Figure 6-A-C.

Across both setups, we explored two pre-training configurations,  $N_{pt} = 5$  k and  $N_{pt} = 10$  k steps, to determine a pre-training duration that provides a sufficiently good initialization while avoiding delaying the start of federated training or causing unnecessary energy consumption. We also varied

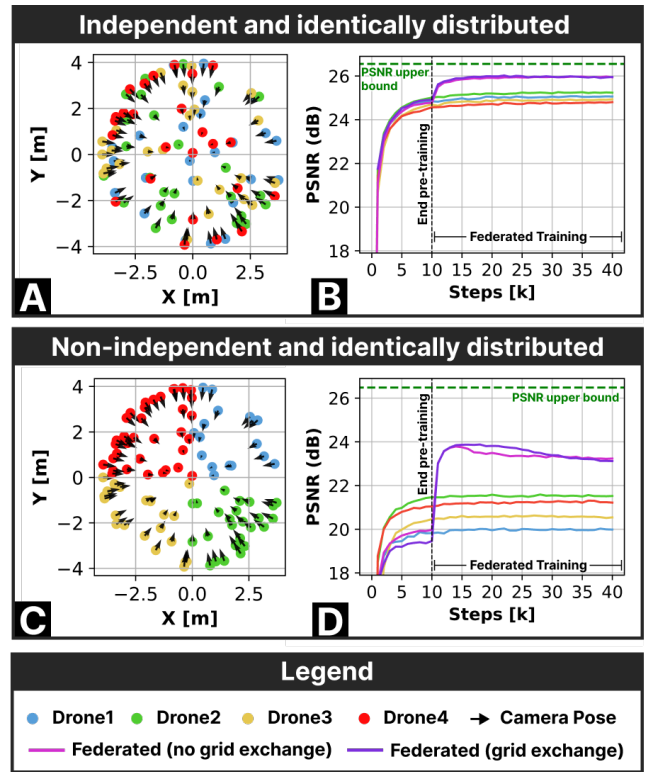


Fig. 6: A-C) 2D top-view of the IID vs. non-IID dataset distributions, B-D) Federated learning PSNR testing curves for the IID vs. non-IID dataset trainings.

the number of local training steps per round, as excessively frequent communication prevents local models from making meaningful progress, whereas overly infrequent communication leads to local overfitting, with both ultimately slowing global convergence. Based on this exploration, we selected  $N_{pt} = 10$  k for the pre-training steps and  $N_{\ell} = 1$  k for the number of local training steps per round. All results in this section are based on this configuration.

We then compared models trained independently on each drone (on their own subset of images) with the two federated variants (with and without occupancy grid exchange) under both IID and non-IID conditions, as shown in Figure 6-B-D. For reference, the figure also includes a centralized model trained on the union of all local datasets, which represents an *upper bound* on achievable performance.

Across both the IID and non-IID setups, the two federated variants, with or without occupancy grid exchange, achieve comparable results, indicating that exchanging the grid is unnecessary. This finding is particularly noteworthy, as the occupancy grid is a high-dimensional data structure and accounts for approximately 88% of the total communication volume. Including the grid requires transmitting 4.7 MB per round, resulting in a communication overhead of 20.16 s per round, whereas omitting it reduces the transmitted data to 0.52 MB and the overhead to 2.24 s per round.

Furthermore, the federated models consistently outperform the individually trained models. In the IID setup, federated

training achieves a PSNR that is 0.7 dB higher than the best single-drone model and only 0.6 dB below the centralized upper bound. In the more challenging non-IID setup, the federated models achieve performance 3.25 dB below the upper bound but still 1.7 dB above the best single-drone model, demonstrating that through the federated model the drones can infer spatial characteristics of portions of the object they have not directly observed.

## VI. CONCLUSION

Dense 3D reconstruction tasks have been unfeasible on nano-drones, as existing methods are not affordable with the available resources, i.e., sub-100 mW MCUs and a few tens of MB. We present Tiny-DroNeRF, a lightweight NeRF tailored for MCUs, which reduces memory by 96% w.r.t to Instant-NGP with only a 5.7 dB PSNR drop. We reconstruct real-world scenes from onboard grayscale images at 21.2 PSNR. We further propose a federated learning scheme that allows swarms to synthesize unseen views from limited local data, improving PSNR by 1–2 dB over single-drone training with only 2.24 s of communication every 1 k steps. To the best of our knowledge, this is the first demonstration of NeRF, with a federated learning approach, on a resource-constrained MCU.

## REFERENCES

- [1] 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,” *Commun. ACM*, vol. 65, no. 1, p. 99–106, Dec. 2021.
- [2] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec, “Baking Neural Radiance Fields for Real-Time View Synthesis,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 5875–5884.
- [3] A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and A. Kanazawa, “PlenOctrees for Real-time Rendering of Neural Radiance Fields,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct. 2021, pp. 5732–5741.
- [4] S. Fridovich-Keil, A. Yu, M. Tancik, Q. Chen, B. Recht, and A. Kanazawa, “Plenoxels: Radiance Fields without Neural Networks,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2022, pp. 5491–5500.
- [5] T. Müller, A. Evans, C. Schied, and A. Keller, “Instant Neural Graphics Primitives with a Multiresolution Hash Encoding,” *ACM Trans. Graph.*, vol. 41, no. 4, pp. 102:1–102:15, Jul. 2022.
- [6] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su, “TensorRF: Tensorial Radiance Fields,” in *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*. Berlin, Heidelberg: Springer-Verlag, 2022, p. 333–350.
- [7] S. Fridovich-Keil, G. Meanti, F. R. Warburg, B. Recht, and A. Kanazawa, “K-planes: Explicit radiance fields in space, time, and appearance,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 12 479–12 488.
- [8] Z. Chen, T. Funkhouser, P. Hedman, and A. Tagliasacchi, “MobileNeRF: Exploiting the Polygon Rasterization Pipeline for Efficient Neural Field Rendering on Mobile Architectures,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2023, pp. 16 569–16 578.
- [9] S. Shin and J. Park, “Binary radiance fields,” in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS ’23. Red Hook, NY, USA: Curran Associates Inc., 2023.
- [10] Y. Chen, Q. Wu, M. Harandi, and J. Cai, “How Far can we Compress Instant-NGP-Based NeRF?” in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, Jun. 2024, pp. 20 321–20 330.
- [11] M. J. Anderson, J. G. Sullivan, J. L. Talley, K. M. Brink, S. B. Fuller, and T. L. Daniel, “The “Smellicopter,” a Bio-Hybrid Odor Localizing Nano Air Vehicle,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6077–6082.
- [12] A. Birk, B. Wiggerich, H. Bülow, M. Pfingsthorn, and S. Schwertfeger, “Safety, Security, and Rescue Missions with an Unmanned Aerial Vehicle (UAV),” *Journal of Intelligent & Robotic Systems*, vol. 64, no. 1, pp. 57–76, Oct. 2011.
- [13] M. Zauli, M. Pirazzi, F. Zonzini, and L. De Marchi, “Exploiting Nano Aerial Vehicles as Sensor Nodes for Wireless Vibration Monitoring,” in *2024 IEEE Sensors Applications Symposium (SAS)*, Jul. 2024, pp. 1–6.
- [14] N. Zimmerman, H. Müller, M. Magno, and L. Benini, “Fully Onboard Low-Power Localization with Semantic Sensor Fusion on a Nano-UAV using Floor Plans,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 11 913–11 919.
- [15] L. Lamberti, L. Bompani, V. J. Kartsch, M. Rusci, D. Palossi, and L. Benini, “Bio-Inspired Autonomous Exploration Policies with CNN-based Object Detection on Nano-drones,” in *2023 Design, Automation & Test in Europe Conference (DATE)*, Apr. 2023, pp. 1–6.
- [16] M. Pourjabar, M. Rusci, L. Bompani, L. Lamberti, V. Niculescu, D. Palossi, and L. Benini, “Multi-sensory Anti-collision Design for Autonomous Nano-swarm Exploration,” in *2023 30th IEEE International Conference on Electronics, Circuits and Systems (ICECS)*, Dec. 2023, pp. 1–5.
- [17] L. Lamberti, E. Cereda, G. Abbate, L. Bellone, V. J. K. Morinigo, M. Barciś, A. Barciś, A. Giusti, F. Conti, and D. Palossi, “A Sim-to-Real Deep Learning-Based Framework for Autonomous Nano-Drone Racing,” *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1899–1906, Feb. 2024.
- [18] H. Wang, J. Wang, and L. Agapito, “Co-SLAM: Joint Coordinate and Sparse Parametric Encodings for Neural Real-Time SLAM,” *Proceedings of the IEEE international conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13 293–13 302, 2023.
- [19] Z. Li, T. Müller, A. Evans, R. H. Taylor, M. Unberath, M.-Y. Liu, and C.-H. Lin, “Neuralangelo: High-Fidelity Neural Surface Reconstruction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8456–8465.
- [20] L. Roldão, R. de Charette, and A. Verroust-Blondet, “3D Semantic Scene Completion: A Survey,” *International Journal of Computer Vision*, vol. 130, no. 8, pp. 1978–2005, Aug. 2022.
- [21] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. PMLR, 20–22 Apr 2017, pp. 1273–1282.
- [22] L. Holden, F. Dayoub, D. Harvey, and T.-J. Chin, “Federated Neural Radiance Fields,” 2023.
- [23] T. Suzuki, “Federated Learning for Large-Scale Scene Modeling with Neural Radiance Fields,” 2024.
- [24] M. Asadi, K. Zareinia, and S. Saeedi, “Di-NeRF: Distributed NeRF for Collaborative Learning With Relative Pose Refinement,” *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 10 527–10 534, 2024.
- [25] T.-I. Sztamari and A. Cauligi, “Federated Multi-Agent Mapping for Planetary Exploration,” 2025.
- [26] Z. Tasneem, A. Dave, A. Singh, K. Tiwary, P. Vepakomma, A. Veeraghavan, and R. Raskar, “DecentNeRFs: Decentralized Neural Radiance Fields from Crowdsourced Images,” in *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LIX*. Berlin, Heidelberg: Springer-Verlag, 2024, p. 144–161.
- [27] H. Müller, V. Kartsch, and L. Benini, “GAP9Shield: A 150GOPS AI-capable Ultra-low Power Module for Vision and Ranging Applications on Nano-drones,” in *European Robotics Forum*. Springer, 2024, pp. 292–297.
- [28] P. Mikićević, S. Narang, J. Alben, G. F. Diamos, E. Elsen, D. García, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu, “Mixed Precision Training,” in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- [29] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.