

Multimodal On-Device Learning for Monocular Depth Estimation on Ultralow-Power MCUs

Davide Nadalini¹, Manuele Rusci², Elia Cereda³, *Graduate Student Member, IEEE*,
Luca Benini⁴, *Fellow, IEEE*, Francesco Conti⁵, *Senior Member, IEEE*, and Daniele Palossi⁶

Abstract—Monocular depth estimation (MDE) plays a crucial role in enabling spatially aware applications in ultralow-power (ULP) Internet of Things (IoT) platforms. However, the limited number of parameters of deep neural networks (DNNs) for the MDE task, designed for IoT nodes, results in severe accuracy drops when the sensor data observed in the field shifts significantly from the training dataset. To address this domain shift problem, we present a multimodal on-device learning (ODL) technique, deployed on an IoT device integrating a Greenwaves GAP9 microcontroller unit (MCU), an 80-mW monocular camera, and an 8×8 pixel depth sensor, consuming ~ 300 mW. In its normal operation, this setup feeds a tiny 107 k-parameter μ PyD-Net model with monocular images for inference. The depth sensor, usually deactivated to minimize energy consumption, is only activated alongside the camera to collect pseudo-labels when the system is placed in a new environment. Then, the fine-tuning task is performed entirely on the MCU, using the new data. To optimize our backpropagation-based on-device training, we introduce a novel memory-driven sparse update scheme, which minimizes the fine-tuning memory to 1.2 MB, $2.2\times$ less than a full update, while preserving accuracy (i.e., only 2% and 1.5% drops on the KITTI and NYUv2 datasets). Our in-field tests demonstrate, for the first time, that ODL for MDE can be performed in 17.8 min on the IoT node, reducing the root-mean-squared error from 4.9 to 0.6 m with only 3k self-labeled samples, collected in a real-life deployment scenario.

Index Terms—Fine-tuning, Internet of Things (IoT), monocular depth estimation (MDE), multimodal, on-device learning (ODL), pseudo-labeling, sparse update, ultra low-power (ULP).

I. INTRODUCTION

A NEW generation of spatially aware miniaturized Internet of Things (IoT) devices is rapidly taking shape,

Received 2 May 2025; revised 17 October 2025; accepted 18 November 2025. Date of publication 1 December 2025; date of current version 5 February 2026. This work was supported in part by the Swiss National Science Foundation (SNSF) RoboMix2 Project under Grant 10004854 and in part by the ChipsJU Project ISOLDE under Grant 101112274. (*Corresponding author: Davide Nadalini.*)

Davide Nadalini is with the Department of Electrical, Electronic, and Information Engineering (DEI), University of Bologna, 40126 Bologna, Italy, and also with the Department of Control and Computer Engineering (DAUIN), Politecnico di Torino, 10129 Turin, Italy (e-mail: d.nadalini@unibo.it).

Manuele Rusci is with the Department of Electrical Engineering (ESAT), KU Leuven, 3000 Leuven, Belgium.

Elia Cereda is with the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, 6962 Lugano, Switzerland.

Luca Benini is with the Department of Information Technology and Electrical Engineering (D-ITET), ETH Zürich, 8092 Zürich, Switzerland, and also with the Department of Electrical, Electronic, and Information Engineering (DEI), University of Bologna, 40126 Bologna, Italy.

Francesco Conti is with the Department of Electrical, Electronic, and Information Engineering (DEI), University of Bologna, 40126 Bologna, Italy.

Daniele Palossi is with the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, 6962 Lugano, Switzerland, and also with the Integrated Systems Laboratory, ETH Zürich, 8092 Zürich, Switzerland.

Digital Object Identifier 10.1109/JIOT.2025.3636826

including palm-sized autonomous quadrotors (nano-drones) [1], mixed-reality eyewear [2], and smart cameras [3]. These novel devices feature ultralow-power (ULP) processing units that typically run artificial intelligence (AI) algorithms to analyze sensor data in real-time. A key algorithmic task is monocular depth estimation (MDE) [4], which takes data from an inexpensive single-camera sensor and extracts depth map information, crucial for a large variety of applications, ranging from human pose estimation [5], people counting [6], 3-D scene understanding [7], vehicle detection [8], and obstacle avoidance [9]. MDE is critical for low-power and low-cost IoT devices, as it avoids the need for bulky stereo systems or costly and power-hungry long-range depth sensors, such as LIDARs [10].

Recent MDE approaches show compelling results by employing deep neural network (DNN) models [11], [12], [13]. Such models typically run on high-end computing devices, ranging from GPU-equipped desktop computers to smartphones, which can afford to execute large DNNs featuring billions of parameters [13], [14], [15]. Conversely, ULP microcontroller units (MCUs), with a power envelope of a few hundreds of mW, limited on-chip memory (i.e., a few Mbyte at most) and a computational power up to a few hundred GOPS/s, make the deployment of state-of-the-art (SoA) models featuring 0.5–1 billion parameters [13], [16], [17] prohibitive. However, Peluso et al. [18] have recently introduced a lightweight DNN model for MDE, called μ PyD-Net, which requires only 0.1 M parameters ($1000\times$ less than SoA) and demonstrated its deployment on an MCU for IoT nodes. Unfortunately, this class of tiny MDE DNNs achieves acceptable accuracies only on test data distributions close to the training dataset, which is generally collected and then labeled using precise, yet bulky equipment. As we also report in our study, when exposed to environment data diverse from the training set, these tiny DNNs are prone to large accuracy drops versus SoA models because of the reduced capacity and the low generalization ability.

To tackle this fundamental domain shift problem, this work proposes a novel on-device learning (ODL) system-level method for tiny MDE models, deployed on IoT sensor nodes powered by MCUs. The ODL paradigm has recently been proposed to incrementally fine-tune DNNs directly on-device, based on locally sensed data and without relying on external compute resources [23]. In our solution, we leverage a multimodal sensing frontend to address the absence of labeled data in the field. More in detail, we couple an always-on low-power camera consuming 80 mW with a power-hungry (~ 300 mW) depth sensor, which is activated only to produce

TABLE I
COMPARISON OF DNN-BASED METHODS FOR MDE. THE SCORES REFER TO THE EVALUATION ON THE KITTI DATASET

Model	Year	Params [M]	Computation device	Device class	Power [W]	On-device learning	Image [px]	δ_1 [%]
MonoViT [17]	2022	10.0	Nvidia RTX 3090	Desktop GPU	350	✗	640×192	89.5
Depth Anything V2 [13]	2025	25.0	Nvidia V100	Desktop GPU	250	✗	518×518	97.3
Metric3D v2, ViT-S [16]	2024	22.0	Nvidia A100	Desktop GPU	250	✗	1064×616	95.0
Metric3D v2, ViT-L [16]	2024	306.0	Nvidia A100	Desktop GPU	250	✗	1064×616	98.5
PyD-Net [14]	2018	1.9	Raspberry Pi 3	Embedded GPU	5	✗	512×256	78.9
EdgeNet [19]	2022	1.7	Nvidia Jetson Nano	Embedded GPU	10	✗	912×228	85.4
Liu <i>et al.</i> [20]	2022	2.0	Raspberry Pi 3	Embedded GPU	5	✗	512×256	81.5
LightDepthNet [21]	2023	1.3	Nvidia Jetson Nano	Embedded GPU	10	✗	640×192	86.0
RT-Monodepth [22]	2024	1.2	Nvidia Jetson Nano	Embedded GPU	10	✗	640×192	84.0
μ PyD-Net [18]	2021	0.1	STM32F7	MCU	0.4	✗	48×48	73.5
Ours	2025	0.1	GWT GAP9	MCU	0.1	✓	48×48	80.8

the depth ground truth. With respect to high-end LIDAR or depth cameras, our depth sensor provides low-resolution depth maps (8×8 px) and can be integrated within an IoT node, thanks to its small form factor (6.4×3 mm). Still, its power consumption is $> 3\times$ higher than a typical ULP MCU.

When pseudo-labeled data is collected without human intervention by matching image and depth information, the deployed MDE DNN undergoes on-device fine-tuning. The learning task relies on backpropagation and is executed on a Greenwaves Technologies GAP9 multicore MCU, thanks to half-precision floating-point single-instruction multiple-data (SIMD) hardware. We propose a new sparse update logic to update the DNN parameters of only a subset of layers based on a memory-driven criterion to reduce the high cost of on-device training. Our approach minimizes the memory requirements of the backpropagation algorithm while preserving fine-tuning accuracy versus full-model fine-tuning.

We summarize our contributions as follows.

- 1) We introduce a multimodal ODL technique to mitigate domain shift in tiny DNNs deployed for MDE on MCU-powered IoT nodes.
- 2) We propose a novel memory-driven sparse update logic to minimize memory usage and training time for on-device backpropagation in tiny U-Net-like models [24].
- 3) We extensively evaluate our ODL technique on public datasets to analyze how the two primary ODL challenges—the limited number of samples and the noisy labels—affect the fine-tuning accuracy.
- 4) We collect and publicly release IDSIA- μ MDE, a custom dataset for MDE on IoT devices, which we use to evaluate our ODL method in a real-world scenario.

We field-test our technique on a GAP9Shield IoT node [25], which features a GAP9 MCU, a QVGA Omnivision OV5647 mono-camera, and an ST Microelectronics VL53L5CX depth sensor, which produces 8×8 px noisy depth labels. When evaluated on the public KITTI and NYUv2 datasets, a μ PyD-Net model with 107 k-parameters achieves δ_1 accuracies of, respectively, 61.0% and 47.0% after ODL with our sparse-update. These scores are +43.2% and +43.4% higher than the present baseline solutions, where the DNN is trained off-device before the deployment on the MCU [18]. On the IDSIA- μ MDE dataset, we show that the sparse ODL process

can be performed with a single epoch of training in a memory budget of 1.2 Mbyte, taking only 17.8 min and 204.9 J for data collection and fine-tuning. In this setting, a μ PyD-Net adapts to the new environment using 3k new images, and the RMSE notably reduces from 4.9 to 0.6 m. To the best of our knowledge, this is the first work to demonstrate the feasibility of ODL in MDE on IoT nodes, leveraging only low-resolution labels acquired with onboard sensors. From an application viewpoint, our solution enables the replacement of bulky and power-hungry depth or stereo cameras in IoT systems with low-cost monocular cameras, complemented by DNN-based processing that dynamically adapts to the target environment. This approach potentially makes embedded applications such as human pose estimation [5], obstacle avoidance [26], and people counting [27] more autonomous, cost- and energy-efficient. We release the code of our experiments.¹

II. RELATED WORK

A. MDE on Embedded Devices

Several recent works propose memory-efficient MDE models that target 5–10 W edge platforms [14], [18], [19], [20], [21], [22]. As reported in Table I, these models all include a few million parameters, a $10\times$ reduction from the unconstrained highest-scoring DNNs [13], [16], [17] benchmarked on popular public datasets such as KITTI [28] or NYUv2 [29]. Among the proposed embedded solutions, EdgeNet [19] introduces an encoder–decoder architecture with upsampling modules, reducing the model size down to 1.7 M parameters with pruning and enabling 55 ms inference on an NVIDIA Jetson Nano via custom GPU scheduling. PyD-Net [14] employs a 1.9 M-parameter pyramidal multiscale architecture trained with an unsupervised stereo image reconstruction loss, which achieves 2 Hz inference on a Raspberry Pi 3.

In contrast, we are interested in IoT nodes that consume less than 500 mW and can only afford to employ ULP MCUs. Peluso *et al.* [18] introduced μ PyD-Net, the first and, to this date, only MDE solution for this class of devices. Its 107 k-parameter encoder–decoder architecture can run on an STM32F7 MCU (i.e., a 400 mW power envelope) to predict 48×48 px depth maps with a latency of 651 ms.

¹Link to our code: <https://github.com/idsia-robotics/ultralow-power-monocular-depth-ondevice-learning>

Unlike the unconstrained SoA approaches that can work in any domain, thanks to the high generalization ability [13], [16], μ PyD-Net and the other edge memory-efficient works are still trained from scratch and tested on a specific domain (e.g., KITTI). To overcome this limitation, this work proposes the first end-to-end system method for fine-tuning tiny MDE models on-device, using data acquired from the deployment environment and the associated ground truths captured by a depth sensor. On the μ PyD-Net architecture, our method improves the accuracy (δ_1) of a model, pretrained on a generic dataset, by +63%, using just 2.3k images from the target environment. This corresponds to an accuracy of up to 80.8%, comparable to models 20 \times its size [14], [20].

We emphasize that, following Depth Anything V2 [13], we adopt a fully supervised training strategy, proved to achieve SoA accuracy on benchmarks such as KITTI. Specifically, we pretrain our μ PyD-Net on synthetic data with perfect ground-truth labels and fine-tune it using 8×8 pseudo-labels from the onboard low-resolution depth sensor. Nevertheless, since our ODL technique is agnostic to the pretraining paradigm, it can be readily applied to models that have been initially pretrained using self-supervised learning approaches [30].

B. Sparse ODL

The growing demand for fine-tuning DNNs on-device has recently increased interest in novel hardware and software optimizations for ODL [23]. In the context of ULP MCUs, a few hardware-optimized ODL frameworks can now produce the source code for the backpropagation-based training task [31], [32]. As a notable example, AIFES [32] uses the optimized software primitives from the ARM CMSIS-DSP library to accelerate the training kernels on ARM Cortex M CPUs. PULP-TrainLib [31] targets instead RISC-V multicore platforms and leverages parallel processing and half-precision floating-point SIMD instructions to speed up the backpropagation task. When benchmarked on a 10-core RISC-V MCU, the PULP-TrainLib primitives can reach a top efficiency of 6.62 multiply and accumulate operations per clock cycle (MAC/clock) for convolutional layers. In this work, we build on PULP-TrainLib to bring, for the first time (to the best of our knowledge), an ODL pipeline for MDE on an MCU device.

An orthogonal optimization axis concerns the usage of sparse update strategies that train (on-device) only a subset of parameters, reducing memory and computation costs versus retraining the full model weights. In this context, Ren et al. [33] proposed to add a single trainable layer on top of a frozen and quantized model. This fine-tuning task can run in milliseconds on an ARM Cortex-equipped Arduino board. TinyTL [34] introduces an extra trainable layer with a low number of parameters in parallel to the frozen convolutional blocks, achieving 12.9 \times memory savings with respect to training all parameters of the original model. Paissan et al. [35] used a structured sparse update scheme to reduce the gradient and the activation storage requirements in the backward step. This approach brings a 1.4 M-parameter convolutional neural network (CNN) to learn two unseen classes in <20 min when accelerated with PULP-TrainLib, reducing the computational and memory costs by 30% with respect to a full update.

Tiny training engine [36] (TTE) extends the sparse update approach to quantized models already deployed on MCUs. By incrementally training only a subset of weights and layers, which are identified by computing an importance score with respect to a given benchmark, TTE reduces by up to 21 \times the training memory with respect to a full update.

In contrast to the previous works that focus on DNNs tailored for classification tasks, we develop a new sparse update logic for a tiny U-Net-like model (μ PyD-Net). We propose to retrain only the first block of the decoder, which requires 2.1 \times and 1.3 \times less memory than, respectively, fine-tuning all parameters of the last decoder block. Thanks to sparse update and hardware-optimized ODL kernels, we enable the fine-tuning of μ PyD-Net with a throughput of 3.57 samples/s, consuming only 41.3 mJ/sample.

C. ODL With Unlabeled Data

SoA MDE models are generally pretrained on large image datasets and then fine-tuned to the target downstream task using human-annotated data before deployment on the target platforms [37], [38]. In contrast to this setting that requires task-specific labeled data to be available already at design time, postdeployment fine-tuning with unlabeled data has recently gained attention [39], [40]. Matsutani and Marculescu [41] suggested a custom design of a low-power hardware core for ODL in a human activity classification task. To deal with the lack of annotations during in-field operation, the system assigns pseudo-labels to new sampled data. Specifically, if the prediction confidence returned by the inference task falls below a threshold, the new data is transmitted to a high-power external mobile device. This device invokes a teacher DNN model to generate high-quality pseudo-labels. The on-device fine-tuning method for voice command classification sensors proposed by Rusci et al. [42] assigns instead pseudo-labels to new samples based on a similarity score with respect to a few human-labeled samples stored in the IoT node's memory. Differently from these works, our work addresses a more challenging regression task, that is, labels do not belong to a discrete set of values, and we rely on an additional sensor for generating the fine-tuning labels.

As a notable application of ODL to regression problems, Boldo et al. [43] proposed a teacher-student approach for human pose estimation and object detection. First, the nonreal-time teacher model is used to generate pseudo-labels. Then, the student, which is originally distilled from the teacher, is fine-tuned on-device (NVIDIA Jetson Xavier NX) using the new set of labeled data. Cereda et al. [1] addressed the domain shift observed for visual human pose estimation tasks after deployment on MCU-powered nano-drones. The proposed method leverages the temporal correlation of the images in a video to propagate human pose labels to successive frames based on the measured drone's motion. Thanks to the new labeled data, the parameters of the onboard visual DNN model are updated without relying on external compute resources.

In contrast with these works, our solution adopts an auxiliary sensor to generate the ground truths, as label propagation through time cannot be applied to our regression problem. Thanks to our multimodal sensor solution, new labels are




Fig. 1. GAP9Shield IoT node [25] mounted on a Crazyflie 2.1 (left) and the characteristics of the main components (table in the right).

produced with a throughput of 15 Hz on the MCU, $2.5\times$ faster than the teacher model in [43]. Overall, our work demonstrates, for the first time, an ODL approach to a regression MDE problem that can run entirely on a ULP IoT node.

III. BACKGROUND

We consider the GAP9Shield [25] as the reference IoT node platform to demonstrate MDE ODL. This SoA design features a GAP9 MCU, an OV5647 monocular camera, a VL53L5CX time-of-flight (ToF) depth sensor,² 32 MB of onboard PSRAM,³ and 64 MB of Flash memory. In addition to the standalone operational mode, the GAP9Shield can expand the sensor front-end of a nano-drone IoT device, that is, a Crazyflie 2.1 system⁴, as illustrated in Fig. 1. Thanks to the miniaturized form factor (a diameter of 10 cm and a total weight lower than 50 g), this commercial nano-quadcopter represents a valuable IoT platform for data acquisition in small and cluttered environments, for example, a laboratory.

The camera of the multimodal sensor front-end produces RGB images of QVGA resolution (320×240 px) at a frame rate of up to 120 frames/s. The horizontal field of view (FOV) measures 57° and the average power consumption is 80 mW in active mode. The VL53L5CX ToF is a multizone depth sensor that can record 8×8 depth maps in a range of [2 cm, 4 m]. The sensor consumes 313 mW on average with a throughput of up to 15 Hz. The horizontal FOV is 65° .

On the processing side, GAP9 includes a 32-bit RISC-V CPU, 1.5 Mbyte of on-chip RAM memory, and a wide set of peripherals to connect to the sensors or external storage memories, including a dedicated I/O DMA engine for off-chip data management. The single-core architecture is enriched by a software-programmable accelerator, denoted as Cluster, that is composed of nine RISC-V CPUs and 128 kbyte of scratchpad memory. A second DMA engine, reserved for the Cluster, can transfer data between the scratchpad memory and the larger on-chip RAM in the background of the CPU operation. The cluster cores share four floating-point units, with support for 16-bit floating-point (BFLOAT16 format) calculation. The presence of these HW blocks extends the instruction set of the CPU with SIMD instructions that can process $2 \times \text{BFLOAT16}$ vectors. When running at maximum clock speed of 370 MHz, the average power consumption is typically below 100 mW.

²VL53L5CX: <https://www.st.com/en/imaging-and-photonics-solutions/vl53l5cx.html>

³AP Memory Pseudo-Static RAM: https://www.mouser.at/datasheet/2/1127/APM_PSRAM_OPI_Xccela_APS256XXN_OBRx_v1_0_PKG-1954780.pdf

⁴Crazyflie 2.1: <https://www.bitcraze.io/products/old-products/crazyflie-2-1>

IV. METHOD

In this section, we discuss our novel ODL system-level method aimed at improving the performance of tiny DNNs for MDE after deployment on an IoT sensor node. The sensor system described in Section III can operate in two distinct modes. Normally, it runs the MDE task based only on image data from the low-power OV5647 camera, while the low-resolution depth sensor is kept off. When placed in a new environment, the system starts the ODL routine, which comprises two phases.

- 1) *Onboard Data Collection*: The system records a new set of images from the target environment and assigns pseudo-labels based on the depth information retrieved by the other sensor. The new dataset is stored on the onboard PSRAM memory.
- 2) *Sparse ODL*: The DNN-based MDE model is fine-tuned on the MCU with respect to the collected data. Our lightweight ODL strategy adopts a sparse update rule where only a subset of the parameters is updated during training.

In the following, we detail the two main phases of our method after describing our baseline sensor system for MDE. Fig. 2 illustrates these components.

A. Baseline MDE System

Our work builds on the SoA $\mu\text{PyD-Net}$ [18] for on-device MDE. $\mu\text{PyD-Net}$, which is illustrated in Fig. 3, is a U-Net-like CNN that generates depth information with the same resolution as the input visual image. The model includes an encoder (ENC) and three decoders, which are indicated as DEC0, DEC1, and DEC2. The ENC block features six 2-D convolution layers followed by LeakyReLU activations. Every encoder layer expands the number of channels of its input tensor data and progressively halves the spatial resolution. In contrast, every decoder block upscales by $2\times$ the spatial dimension of the inputs using transposed convolutions (TrConv2D). The output features of the DEC0 and DEC1 blocks are concatenated with the encoder's intermediate results, which are propagated via the residual connections, before feeding the following decoder stage.

In the original work, $\mu\text{PyD-Net}$ is trained to predict a 48×48 px disparity map $\mathcal{I}_{\text{disp}}$, which reflects the pixel-wise displacement between the left and right images of the training data (KITTI). A depth map, that is, a 48×48 px image that encodes the pixel-level distance (in m) of the scene from the camera, is then obtained as

$$\mathcal{I}_{\text{depth}} = (f \cdot B) / \mathcal{I}_{\text{disp}} \quad (1)$$

where f and B are scalar parameters indicating, respectively, the focal length of the RGB camera used to collect images and the stereo baseline distance between the centers of projection of the stereo camera pair. Given the use of stereo images for training, the authors of $\mu\text{PyD-Net}$ employ a training loss that combines a supervised *reverse Huber* (berHu) [44] loss with an unsupervised *appearance matching* [45] stereo re-projection loss. The supervised loss computes the error between $\mu\text{PyD-Net}$'s prediction and a disparity map derived from the




Fig. 2. Diagram of the main phases of our ODL method for MDE. (a) IoT node continuously runs the inference-only task onboard and predicts the depth maps from single-camera images. In the ODL mode, (b) system collects new data from the camera and the depth sensor and then (c) performs training on-device. In every phase, system components colored in white are not active.

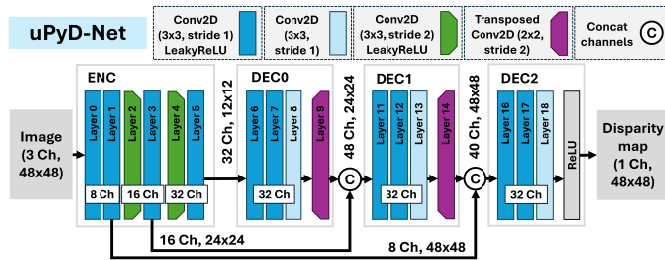


Fig. 3. μ PyD-Net [18] DNN architecture for MDE.

left and right training images. The unsupervised component, instead, returns a similarity score between one training image (e.g., the left) and the other image of the stereo pair after applying a warping transformation according to the predicted disparity map.

Since our MDE model is intended for on-device training with monocular images, we initially train μ PyD-Net on a dataset featuring single-camera images annotated with depth ground truth. Different from the original work, we use only the supervised *berHu* loss, which computes the error between μ PyD-Net’s disparity predictions and the respective depth labels, inverted to a disparity map using (1). We experimentally verified that our training recipe leads to the same accuracy as the loss function used in the original work.

After training, μ PyD-Net is deployed on the GAP9 MCU of our system [Fig. 2(a)]. In this work, we rely on the latency-optimized kernels of PULP-TrainLib [31] to estimate the latency and energy consumption of μ PyD-Net. We feed the model with single images featuring a CHW data layout. The DNN processing is performed layer-by-layer using the Cluster cores. The weights and intermediate tensors data, which are initially stored in the on-chip 1.5 Mbyte memory, are copied to the Cluster scratchpad memory before invoking the layer-wise functions. The results of the computation are then stored back to the on-chip RAM for later use using the DMA.

Thanks to the optimized deployment library, our IoT node can predict at runtime the $\mathcal{I}_{\text{depth}}$ output by feeding the network with RGB images captured by the Omnivision OV5647 monocular camera. To feed μ PyD-Net, every 320×240 px image is downsampled to 48×48 resolution using nearest neighbor interpolation. While working in this pure inference mode, the low-resolution depth sensor (the VL53L5CX ToF) is turned off to reduce the energy consumption of the node

(the depth sensor consumes $\sim 2 \times$ more than the combination of camera and MCU).

B. On-Device Data Collection

When placed in a new environment, the system records a new image dataset for ODL using the available sensors. More in detail, every image recorded by the camera is associated with the low-resolution depth map output by the VL53L5CX ToF sensor. Note that, however, this *pseudo-label* is not equivalent to the ground-truth depth because of two factors. First, the pseudo-labels are partial, as the VL53L5CX sensor is limited to a maximum range of 4 m, often constrained to ~ 2 m in real-world conditions due to nonideal lighting and surface reflectance. During fine-tuning, we address this range issue by avoiding the loss computation for every pixel marked as invalid by the depth sensor.

Second, the camera features a narrower FOV than the depth sensor, which causes a mismatch between the images’ content and the respective measured depth. Instead of aligning the two sensors by warping the depth map to match the camera’s narrower FOV—which would discard useful depth information corresponding to edge regions of the camera image—we retain the full depth map to preserve all potentially relevant spatial data. This approach produces noisy labels, which require the model to learn the sensor alignment from the data itself, during on-device fine-tuning.

The new set of labeled images is stored in the onboard PSRAM memory [Fig. 2(b)]. Thanks to the total capacity of 32 MB, our system can retain up to 4587 48×48 RGB images (6.9 kbyte each after downsampling) and the associated 8×8 depth maps (64 byte each). Given the throughput of the slowest sensor (15 Hz for the VL53L5CX versus 120 Hz for the camera), the maximum amount of data is collected in 305.8 s.

C. Sparse ODL Update

1) *Backpropagation-Based Learning for MDE*: We use the backpropagation algorithm to fine-tune μ PyD-Net on-device. The backpropagation algorithm operates in two steps, namely the *forward*, computing the DNN disparity prediction with respect to an input image \mathcal{I} , and the *backward*, which calculates the update values for the weights to be fine-tuned. During the forward step, the intermediate activations of the layers to be updated (e.g., convolutions) and of the activation functions




Fig. 4. Layer-wise breakdown of μ PyD-Net memory cost, contributing to the total fine-tuning memory.

across which the error is propagated (e.g., ReLU) are retained in memory for the backward pass.

After the forward step, we use the same loss function of the initial training to calculate the error between the model prediction (48×48 px) and the noisy label, collected in the field (8×8 px). To compute the loss, the label needs to be transformed. First, we invert the label to a disparity map, using (1). Then, we upscale it using a bilinear interpolation to 48×48 px. The backward step then backpropagates the error signal up to the first layers with trainable parameters.

The obtained gradient values are eventually averaged over a mini-batch of samples (image and depth) and then accumulated into the weight values after applying the learning rate. Since μ PyD-Net does not include any batch-dependent layer, for example, BatchNorm, the backpropagation algorithm can operate sample-by-sample, in a streaming fashion.

2) *On-Device Learning*: The backpropagation-based training task runs entirely on the GAP9 MCU, without involving any off-device (server) computing resource [Fig. 2(c)]. To improve computation and memory efficiency, we select the BFloating16 floating-point datatype to represent weights, gradients, and intermediate feature values, similar to the inference phase. This data format allows the use of SIMD-accelerated software kernels from PULP-TrainLib, which can lead to up to $1.9\times$ speed-up versus the 32-bit floating point counterpart [31] without sacrificing accuracy [46].

Our implementation runs the μ PyD-Net forward and backward passes layer-by-layer using the cluster accelerator of GAP9 and uses the 1.5 Mbyte on-chip memory as the main computation memory. Fig. 4 illustrates the total memory requirements to train μ PyD-Net on GAP9. The total fine-tuning memory (2.56 MB) is composed of two components, that is, a working memory buffer to compute every layer's activations and gradients, and a storage buffer for the backward step data.

More specifically, the working memory buffer is dimensioned for computing the largest among all layers' activations and gradients, and is shared by the forward and backward steps of all layers. The storage buffer, instead, contains the model weights, the intermediate activations, and the weight gradients

for the layers to be updated. This last cost also accounts for the additional memory required by the optimization process, for example, the buffers for the estimation of the statistical moments in the case of the Adam optimizer [47]. During ODL, while the working memory buffer is always kept on-chip, storage buffer data that exceeds the 1.5 Mbyte on-chip memory is stored in the onboard PSRAM. When requested, such data is accessed using the I/O DMA, in parallel with CPU operations.

In the U-Net-like [24] architecture of μ PyD-Net, introduced in Section IV-A, DEC2 impacts the most on the fine-tuning memory cost. The size of the working memory buffer (354.9 KB) is determined by the first convolution of DEC2, which features the largest input activation of the model (40 channels with 48×48 spatial size). Similarly, the size of the storage memory buffer (2.2 MB) is determined for 43.1% by DEC2 (947.5 kB), while ENC, DEC0, and DEC1 contribute to its 18.6%, 14.1%, and 24.2%, respectively. Of the whole storage buffer, 1.3 MB (60.8%) is required to store the intermediate activations of the model, while the remaining 39.2% stores the weights and their related gradients. Such requirements are unevenly distributed among the blocks: while DEC2 contains 58.3% of all activations (778.8 kB), it possesses only 19.6% of the weights of the model. The other blocks (ENC, DEC0, and DEC1, respectively), instead, require 19.7%, 17.9%, and 4.1% of the stored activation cost. At the same time, the central blocks (DEC0 and DEC1) contain the majority of the weights (29.6% and 33.9%), while the ENC possesses only 16.9%.

3) *Sparse Update*: We use a sparse update approach to fine-tune only a subset of the μ PyD-Net parameters. Our proposed layer selection logic aims to minimize the memory requirements versus the full model update while preserving final accuracy in a regression task, that is, MDE. With respect to typical DNNs used for classification tasks, which show the lowest cost when retraining only the last layers [35], we make the fundamental observation that our U-Net-like μ PyD-Net presents the highest memory cost in correspondence to the last layers. Furthermore, the computation of DEC2's weight gradients alone constitutes 28.2% of the total cost for a backward step, making the weight update for the last layers the most computationally demanding, while early blocks of layers (ENC and DEC0) only represent 3.4% and 3.7% of such cost.

Therefore, our memory-driven sparse fine-tuning strategy updates only DEC0, which constitutes the memory-minimal solution due to the small footprint of its activations. In this configuration, the total fine-tuning memory is reduced to 1.2 MB ($2.2\times$ less than a full update). Specifically, the storage buffer (835 kB) accounts for 310.2 kB corresponding to DEC0 activations and gradients, 373.2 kB for the input activations of the ReLU and LeakyReLU layers in DEC1 and DEC2 (required for error propagation), and 151.5 kB for the weights of the remaining network blocks. This setup reduces by 24% the memory cost, compared to the standard fine-tuning approach, which updates only the final layers (DEC2 for μ PyD-Net). In this conventional scheme, the storage buffer alone requires 1.1 MB, comprising 947.5 kB for the activations, weights, and gradients of DEC2, and an additional

173 kB for the weights of the other network blocks used during the forward pass.

In Section VI-C, we experimentally show that our memory-efficient sparse configuration also retains the same accuracy score compared to retraining the full model. This result is in line with the observation of previous studies that observed no accuracy drops when incrementally updating only intermediate layers of DNNs tailored for classification tasks [48] and U-Net models for ultrasound image segmentation [49]. In compliance with the limited amount of RAM of our IoT platform, we also demonstrate that our sparse update scheme is robust to a significant reduction in the number of samples available for fine-tuning on-device.

V. EXPERIMENTAL DATA PROTOCOL

We define a new data protocol to validate our ODL method for MDE. Using data from available public datasets, we initially mimic an ideal sensor setup with a perfect match between the FOVs of the camera and the depth sensor. More in detail, the baseline μ PyD-Net is trained on TartanAir [50], and then we fine-tune it on data from the KITTI [28] or NYUv2 [29] dataset. *TartanAir* is a synthetic dataset comprising four TB of photorealistic data from structured urban, indoor, and unstructured natural environments, annotated with dense depth maps and featuring a resolution of 640×480 px. *KITTI* is a popular computer vision benchmark featuring stereo images with an average resolution of 1242×375 px, coupled with point-cloud LIDAR depth measurements, capturing scenes from urban, suburban, and rural environments. *NYUv2* is a popular Microsoft Kinect benchmark comprising more than 400 000 RGB 640×480 images, annotated with dense depth maps, collected as video sequences from various indoor environments, such as kitchens, living rooms, bedrooms, and bathrooms.

TartanAir, KITTI, and NYUv2 are preprocessed by cropping centered portions of every image and label, which are then downsampled to 48×48 resolution with nearest neighbor interpolation to feed μ PyD-Net. To keep the same downsampling ratio for all datasets, we apply the cropping size adopted for KITTI (360×360 px). This size is chosen to always crop square portions of images and labels in such a dataset, since several samples feature a vertical resolution less than the average one (375 px).

In our protocol, we pretrain μ PyD-Net on TartanAir, extracting 64 358 training images from three environments (*AbandonedFactory*, *Neighborhood*, and *SeasonsForest*), using 48×48 labels. We evaluate our fine-tuning on KITTI and NYUv2. For the KITTI dataset, we adopt the Eigen split [11], which comprises 22 600 images for training, 888 for validation, and 697 for testing. Since KITTI features ground-truth data collected with a LIDAR sensor that produces sparse measurements, we follow [18] to generate dense proxy labels from stereo images, for training. For the NYUv2 dataset, we sample 22 600 images for training, 2726 for validation, and 2948 for testing from 180 sequences of kitchens, living rooms, bedrooms, and bathrooms, each divided into 80% data for training, 10% for validation, and 10% for testing. To eventually match the specifications of the VL53L5CX depth

sensor, we simulate the 8×8 depth labels used for fine-tuning by downsampling every 48×48 label of KITTI and NYUv2 with a 6×6 min-pooling filter.

In a second phase, we assess the performance of our ODL solution on a set of real-world data, *IDSIA- μ MDE*⁵, collected with our IoT setup (Section III). *IDSIA- μ MDE* features 3035 labeled images for training, 1149 for validation, and 3171 for testing, collected from different perspectives of the same laboratory environment, featuring artificial lighting and a set of obstacles, such as recycle bins, robots, and gates. Due to the nonidealities of our IoT setup (i.e., sensor FOV mismatch and limited depth sensor range), *IDSIA- μ MDE* images are annotated with noisy labels, which we employ both for fine-tuning and testing.

VI. RESULTS

A. Experimental Setup

We evaluate our ODL solution using the PyTorch framework,⁶ and we report every result as the average over five experiments. Additionally, we profile the training task running on the GAP9 MCU by measuring the latency and power consumption of the IoT system.

1) *Metrics*: We evaluate the accuracy of MDE prediction by employing three metrics. We define as \mathcal{I}_{GT} the depth ground truth and with N the number of elements in μ PyD-Net’s prediction \mathcal{I}_{dpth} . The quality of pixel-wise predictions is measured using the accuracy under threshold 1.25^K (δ_K). δ_K expresses the ratio of pixels in which the predicted depth is within a multiplicative threshold 1.25^K of the ground truth

$$\delta_K = \frac{1}{N} \sum_{i \in N} I \left\{ \max \left(\frac{\mathcal{I}_{dpth,i}}{\mathcal{I}_{GT,i}}, \frac{\mathcal{I}_{GT,i}}{\mathcal{I}_{dpth,i}} \right) < 1.25^K \right\} \quad (2)$$

where $I\{\cdot\}$ is an indicator function that equals 1 if its argument is true and 0 otherwise. The root-mean-squared error (RMSE) measures the average prediction error, in meters, between the prediction and ground truth, with the formula

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i \in N} \|\mathcal{I}_{dpth,i} - \mathcal{I}_{GT,i}\|^2}. \quad (3)$$

The scale-invariant mean squared error in log space (SiLog) metric expresses, on a logarithmic scale, the relative scale error between the ground truth and the depth prediction. We express SiLog following the formula:

$$\text{SiLog} = \frac{1}{N} \sum_{i \in N} d_i^2 - \frac{1}{N^2} \left(\sum_{i \in N} d_i \right)^2 \quad (4)$$

where $d_i = \log(\mathcal{I}_{dpth,i}) - \log(\mathcal{I}_{GT,i})$ represents the pixel-wise error in log scale. Readers may refer to [11] for more details on the common metrics used to evaluate MDE. For the public datasets, we test the accuracy of μ PyD-Net by upsampling with nearest neighbor interpolation its 48×48 predictions to 360×360 , inverted to a depth map using (1), and by comparing

⁵IDSIA- μ MDE: <https://github.com/idsia-robotics/ultralow-power-monocular-depth-ondevice-learning>

⁶PyTorch: <https://pytorch.org/>

TABLE II

ACCURACY EVALUATION OF OUR FINE-TUNED MODELS (FULL UPDATE) VERSUS NO FINE-TUNED BASELINE AND DUMMY PREDICTORS. ALL MODELS ARE PRETRAINED ON THE TARTANAIR DATASET

Model	Dataset	Fine-tuned (Samples [k])	δ_1 [%] \uparrow	RMSE [m] \downarrow	SiLog \downarrow
Baseline	KITTI	\times (–)	17.8	19.5	6.8
(no fine-tuning)	NYUv2	\times (–)	3.6	74.9	8.7
Full update (48 \times 48 labels)	KITTI	\checkmark (22.6)	81.9	7.2	2.0
		\checkmark (2.3)	80.8	7.4	2.1
	NYUv2	\checkmark (22.6)	51.7	0.5	2.5
		\checkmark (2.3)	49.2	0.5	2.6
Full update (8 \times 8 labels)	KITTI	\checkmark (22.6)	63.2	9.3	2.3
		\checkmark (4.5)	62.3	9.2	2.3
	NYUv2	\checkmark (22.6)	50.1	0.5	2.5
		\checkmark (4.5)	49.3	0.5	2.6
Dummy predictor (48 \times 48 labels)	KITTI	\times (–)	75.3	9.8	2.8
	NYUv2	\times (–)	40.2	0.6	2.6
Dummy predictor (8 \times 8 labels)	KITTI	\times (–)	45.7	14.3	3.3
	NYUv2	\times (–)	39.9	0.6	2.6

them with 360×360 portions of the ground truth. For the real-world data, the accuracy of μ PyD-Net is tested by comparing its 48×48 predictions, inverted with (1), with the 8×8 noisy labels, upsampled to 48×48 resolution with nearest-neighbor interpolation.

2) *Hyperparameters Settings*: Following [14], we set the optimizer to Adam ($\beta_1 = 0.9, \beta_2 = 0.999$, and $\epsilon = 10^{-8}$), the batch size to 16, and the loss function to a *reverse Huber* (berHu) [44], both for the initial training and fine-tuning phases. The learning rate is set to 10^{-4} for models trained from scratch and 10^{-3} for fine-tuned models. We set the maximum epochs for pretraining (TartanAir) to 200, and to 120 for fine-tuning on both public datasets (KITTI and NYUv2) and data collected in the field. An early stopping mechanism selects the model belonging to the epoch where the validation loss is minimized. We augment the training data with three photometric transformations, that is, random gamma correction [0.8, 1.2], random brightness variation [0.5, 2.0], channel-wise color change [0.8, 1.2], and horizontal flipping.

B. Full Update Analysis

In this section, we analyze the accuracy of our μ PyD-Net when fine-tuning all the parameters of the model, that is, the full update scheme. As additional terms of comparison, in Table II, we also provide the results of the baseline (nonfine-tuned) and two dummy predictors. Baseline and fine-tuned models are pretrained on the TartanAir dataset (i.e., from a simulator); instead, the dummy predictors predict the pixel-wise average depth of all testing samples. For each row of Table II, we report the testing dataset on either KITTI or NYUv2, also used for fine-tuning. For both fine-tuned models and dummy predictors, we evaluate two label configurations: an ideal 48×48 and 8×8 px labels, where the latter represents the actual labels' resolution on our IoT node. Finally, in our analysis, we also consider both fine-tuning with all the dataset samples, that is, 22.6k, and with a subset constrained by the available device's PSRAM (32 Mbyte), resulting in 2.3 k and 4.5k samples, for the 48×48 and 8×8 px labels, respectively.

As can be observed from Table II, the baseline model struggles to generalize, with a limited 17.8% and 3.6% δ_1 and an RMSE of 19.5 and 74.9 m, for the KITTI and NYUv2 datasets, respectively. This performance highlights the domain shift problem between training in simulated environments and real-world testing ones. The best δ_1 accuracy comes with the models fine-tuned with 48×48 px labels on the full dataset (i.e., 22.6k samples). In this case, μ PyD-Net achieves +64.1% and +48.1% δ_1 on KITTI and NYUv2, respectively, compared to the baseline. Instead, with 8×8 px labels and the full fine-tuning datasets, the relative improvement with respect to the baseline reduces to +45.4% and +46.5% δ_1 , on KITTI and NYUv2, respectively.

While for NYUv2, we observe a similar fine-tuning accuracy with both label resolutions, μ PyD-Net achieves 18% lower accuracy in the case of KITTI, using 8×8 labels. This effect derives from downsampling the labels to emulate the VL53L5CX sensor in outdoor environments. In this case, regions featuring large depth variance in neighboring pixels are downsampled to the nearest depth, distorting the depth distribution (e.g., a region containing both a car and a portion of the sky is collapsed to the car's depth). Moving to the smaller fine-tuning datasets, the models in both label configurations show a negligible loss in all metrics, that is, a reduction of only 2.5% at most in δ_1 , for NYUv2 with 48×48 px labels. Finally, considering the dummy predictors introduced as sanity checks, they consistently score better than the baseline but significantly worse than all fine-tuned models in all configurations. These results demonstrate the full update's effectiveness in limiting the domain shift problem even when employing small labels and a limited number of samples.

In Fig. 5, we report a qualitative analysis of the various models presented in Table II with three samples for each testing dataset. As previously introduced in Section V, ground truths for the KITTI dataset [Fig. 5(a)–(c)] are produced with a stereo camera, which suffers more from occlusions, resulting in missing ground-truth depths (white pixels). Instead, the NYUv2 dataset, employing a 640×480 px Kinect depth sensor, produces dense labels [Fig. 5(d)–(f)]. The baseline model is strongly biased by the pixels' brightness, predicting high-depth values for pale-color pixels. Conversely, models fine-tuned with all 22.6 k samples, independent of their labels' resolution, accurately detect obstacles in KITTI, clearly separating cars from backgrounds. For NYUv2, instead, they capture the scene details less accurately while correctly segmenting most areas with homogeneous depth, especially in the lower part of the images.

The models fine-tuned with 2.3k and 4.5k samples and 48×48 px labels are less precise in predicting objects' edges, in particular for those on the forefront, while they are pretty accurate for far-distant objects and backgrounds. Instead, the models fine-tuned with 8×8 px labels produce blurrier predictions, resulting from the bilinear upsampling used to generate these small labels, which creates smooth transitions in the fine-tuning's upscaling process. Finally, considering the two dummy predictors that output only the pixel-wise average of the testing datasets, they capture the correlation between

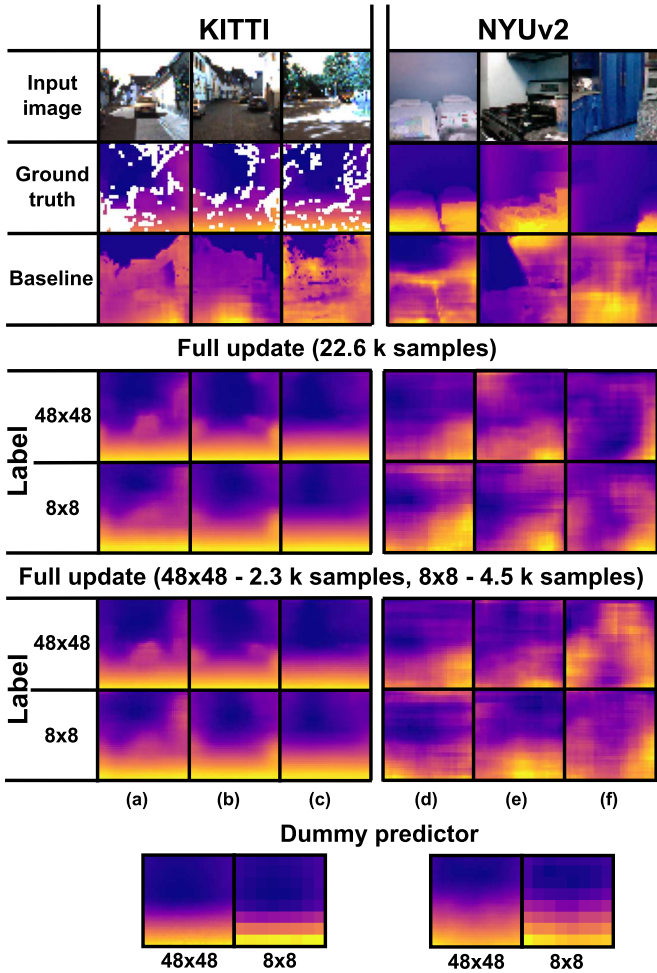


Fig. 5. Images, ground truths, and μ PyD-Net predictions.

the background in the upper part of the images (higher depth) and the foreground (lower depth) in their lower part.

C. Sparse Update Analysis

This section evaluates tradeoffs between fine-tuning memory cost, latency, and δ_1 accuracy of various sparse update schemes. Fig. 6 reports such results for all 16 sparse update schemes by combining all the possible configurations of the four blocks of μ PyD-Net, that is, a block-wise combination of ENC, DEC0, DEC1, and DEC2. Reported latency is estimated by measuring single layers of μ PyD-Net on the GAP9 MCU. As a term of comparison, we also include the nonfine-tuned baseline models, configuration A in Fig. 6. Configuration B, that is, full update, represents the scheme where all four blocks are used for fine-tuning their weights and biases. Configurations from C to F refer to those where only one block of μ PyD-Net is fine-tuned. Like in Section VI-B, all the models are pretrained on the TartanAir dataset, while the fine-tuning employs KITTI and NYUv2 datasets with 48×48 and 8×8 px labels. In the former case, the models are fine-tuned on 2260 images and the latter on 4520 ones to respect the PSRAM capacity imposed by our IoT device (see Section III). Finally, Fig. 6 highlights the Pareto frontiers of the memory

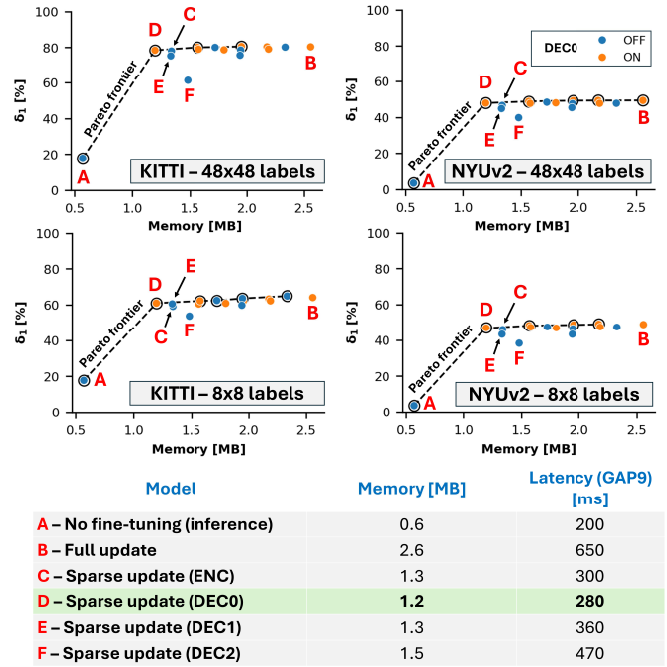


Fig. 6. Sparse update δ_1 accuracy versus memory for KITTI and NYUv2. We highlight the four most meaningful configurations.

versus accuracy, that is, the dashed line with black circled configurations.

All update schemes that imply fine-tuning one of the first three blocks of our μ PyD-Net achieve similar δ_1 accuracy, with a standard deviation of only 3.6%, over all datasets and labels. The only exceptions are represented by configurations A and F; no fine-tuning is applied in the former, while only the last block (DEC2) is fine-tuned in the latter. In CNNs, shallower layers extract low-level features, such as edges and other geometric traits, while deeper ones extract more abstract features specific to the task addressed [51]. Therefore, in the case of domain shift due to changes in the testing dataset, like in our case, the model can benefit more from fine-tuning its geometrical interpretation (shallower layers) instead of updating deeper layers responsible for the task-specific features, as the task remains the same, that is, depth estimation.

μ PyD-Net requires 0.6 Mbyte of memory to store weights and biases for the inference (i.e., forward pass). This number grows, at least, to 1.2 Mbyte for the cheapest sparse update configuration D ($2.1 \times$ less than the full update), making it possible to fit in the GAP9’s memory. Configuration D is also the fastest in the sparse update backpropagation pass, requiring only 80 ms per sample. Compared to the other update schemes on the Pareto frontier, this memory-optimal configuration D is also, at least, 14.3% faster. This results from the limited number of parameters updated in DEC0 (29.6% of the total), and the possibility of stopping the backpropagation at layer 13, as there is no need to backpropagate the gradient to the ENC block.

Relative changes in the δ_1 accuracy between testing datasets and labels’ resolution, also in this case of sparse updates, follow similar trends as those highlighted for the full update

TABLE III

ACCURACY EVALUATION WITH THE SPARSE UPDATE (CONFIGURATION D FROM FIG. 6) OF OUR FINE-TUNED MODELS

Model	Dataset	Fine-tuned (Samples [k])	δ_1 [%] \uparrow	RMSE [m] \downarrow	SiLog \downarrow
Sparse update (48×48 labels)	KITTI	\checkmark (22.6)	79.3	7.8	2.2
		\checkmark (2.3)	78.5	7.9	2.3
	NYUv2	\checkmark (22.6)	48.7	0.5	2.6
		\checkmark (2.3)	48.0	0.5	2.6
Sparse update (8×8 labels)	KITTI	\checkmark (22.6)	61.0	9.5	2.3
		\checkmark (4.5)	61.0	9.4	2.4
	NYUv2	\checkmark (22.6)	47.2	0.5	2.6
		\checkmark (4.5)	47.0	0.5	2.6

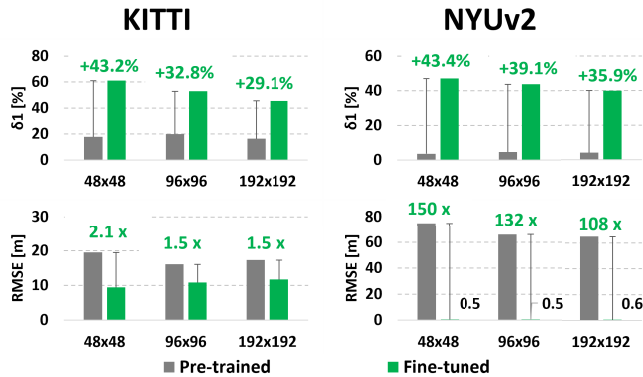


Fig. 7. Sensitivity of ODL to variable input resolution (8×8 labels).

scheme (Section VI-B), achieving the highest δ_1 accuracy for the KITTI dataset with 48×48 px labels, peaking at 80.6% and with a lower score on the NYUv2 dataset when fine-tuned with 8×8 px labels which is, at best, 49.2% δ_1 accuracy. For all these considerations on memory cost, latency, and δ_1 accuracy, we mark configuration D as the best tradeoff, showing the smallest memory footprint and fine-tuning latency, paired with an overall accuracy on par with the other sparse update schemes.

In Table III, we present an in-depth analysis of configuration D of Fig. 6. We present the regression performance of this sparse update scheme for both KITTI and NYUv2 datasets (both for fine-tuning and testing) and both 48×48 and 8×8 px labels. Additionally, like in Section VI-B, we also consider the availability of the entire fine-tuning dataset, that is, 22.6k samples or a small subset fitting the limited PSRAM memory of our IoT device. All major relative trends of the full update analysis are also confirmed in this sparse update case. For example, the best performing configuration is the one with 48×48 px labels and the entire fine-tuning dataset, and reducing the number of fine-tuning samples marginally affects performance, for example, the most significant drop in the δ_1 accuracy is lower than 0.8%. Similarly, the performance drop when using the smaller labels is significant only for the KITTI dataset (i.e., 18.3% accuracy loss) but not for the NYUv2 one, that is, 1.5% at most, like for the full update case and the same reasons.

D. Sensitivity to Image Resolution

In this section, we evaluate the sensitivity of our on-device fine-tuning strategy to input resolution, exploring its applicability to systems with higher computational resources. To this end, we fine-tune our pretrained μ PyD-Net on the KITTI and NYUv2 datasets, also using images with increasing resolutions of 96×96 and 192×192 , corresponding to $2 \times$ and $4 \times$ the 48×48 resolution used in our previous experiments. During fine-tuning, μ PyD-Net is supervised with 8×8 labels, upscaled to the target resolution via bilinear interpolation. Each experiment is carried out for 120 epochs, employing the most memory-efficient sparse update scheme (i.e., updating only DEC0), and uses 4.5k training samples. The results, summarized in Fig. 7, are compared against fine-tuning with 48×48 labels.

Consistent with Table II, the pretrained model performs poorly on higher resolutions, achieving less than 20% δ_1 accuracy on KITTI and under 5% on NYUv2. Conversely, when fine-tuned with upscaled 8×8 labels, the model exhibits substantial improvements across all resolutions, increasing its δ_1 by +43.2%, +32.8%, and +29.1% for 48×48 , 96×96 , and 192×192 inputs on KITTI, and by +43.4%, +39.1%, and +35.4% on NYUv2, respectively. Correspondingly, on KITTI (which features a depth range similar to TartanAir), the RMSE is reduced by 2.1 \times , 1.5 \times , and 1.5 \times for 48×48 , 96×96 , and 192×192 inputs, respectively. In contrast, on NYUv2, where the depth range differs substantially from TartanAir (0–5 versus 0–80 m), the RMSE improves by two orders of magnitude (150 \times , 131 \times , and 108 \times).

Despite these improvements, the relative accuracy gain decreases with increasing resolution. This effect is partially attributed to the upscaling of the 8×8 labels, which become less representative of the details of their respective images, while their resolution increases. Therefore, on KITTI, the maximum δ_1 accuracy drops from 61.0% at 48×48 to 52.8% and 45.5% at 96×96 and 192×192 , respectively. On NYUv2, the corresponding values are 47.0%, 43.7%, and 39.9%. Similarly, the RMSE after fine-tuning increases with resolution, reaching 9.4, 10.8, and 11.7 m on KITTI, and 0.5, 0.5, and 0.6 m on NYUv2.

Increasing input resolution also affects the ODL memory cost. For 96×96 inputs, fine-tuning requires 3.5 MB ($2.9 \times$ more than 48×48), including 1.35 MB for the working buffer, while for 192×192 , memory usage rises to 12.6 MB ($10.6 \times$), with 5.3 MB for the buffer. Overall, these results confirm that, even with larger input images, 8×8 labels can still provide effective in-field supervision.

E. Real-World Experiments

1) *Accuracy on IDSIA- μ Mdeidsia-Data*: This section evaluates the performance of our μ PyD-Net model in a real-world scenario, using the IDSIA- μ MDE dataset. We test μ PyD-Net, pretrained on TartanAir, in two scenarios, that is, without fine-tuning (baseline) and when fine-tuned with the memory-optimal sparse update setup selected in Section VI-C, which updates only DEC0. We report the accuracy results of these configurations in Table IV, comparing them for reference with

TABLE IV
ACCURACY EVALUATION ON IDSIA- μ MDE

Model	$\delta_1 \uparrow$ [%]	RMSE \downarrow [m]	SiLog \downarrow
Baseline (no fine-tuning)	17.0	4.9	3.9
Sparse update (DECO)	30.8	0.5	2.0
Full update	27.4	0.6	2.2
Dummy predictor	7.1	2.0	4.2

Model	Epochs	δ_1 [%]	Latency [min]	Energy [J]
Full update	120	27.4	3972	25235
	1	25.8	33.1	210.3
Sparse update (DECO)	120	30.8	1726	15038
	1	27.4	14.4	125.3




Fig. 8. Comparison between on-device sparse and full update varying the number of epochs and validation loss for sparse update on IDSIA- μ MDE.

the full update scheme and with a dummy predictor, which predicts the pixel-wise average depth of all testing samples of IDSIA- μ MDE.

From an accuracy point of view, the baseline model poorly generalizes on IDSIA- μ MDE. This can be noted by the RMSE (4.9 m), which exceeds the range of the onboard VL53L5CX depth sensor, used to collect labels in the field. Conversely, fine-tuning with both sparse and full update reduces the RMSE to 0.5 m in the former case, while in the latter to 0.6 m. However, the δ_1 accuracy is very limited on the target (up to 30.8% for the sparse update). This effect is caused by the noisy labels of IDSIA- μ MDE: while the overall scale of the target environment can be learned from a few valid pixels, learning the pixel-wise depth is limited by the mismatch between the shapes of the objects in the images and their related depth.

Furthermore, the sparse update scheme achieves better accuracy than a full update. This effect can arise because only the early layers of the network are being updated. In this scenario, characterized by noisy labels, fine-tuning just the low-level features of the CNN helps improve the estimation of the coarse depth scale. Meanwhile, the prediction of higher-level details still benefits from the pretraining phase, during which the CNN was trained on highly accurate depth labels.

2) *ODL Measurements*: Fig. 8 reports the latency and energy measurements, on GAP9, of our on-device fine-tuning. From a latency and energy point of view, fine-tuning μ PyD-Net with sparse update requires 14.4 min on the GAP9 MCU per epoch, at the cost of 125.4 J. In the case of a complete 120-epoch fine-tuning, these costs make training difficult to achieve on the MCU, requiring 66.2 h and 25.3 kJ. For our in-field setup, we can reduce such costs by minimizing the number of fine-tuning epochs performed on the MCU. To this end, we study the evolution of the validation loss to observe when the model approximates convergence on the target dataset. Fig. 8 reports this exploration over 120 epochs using sparse update. In this case, we note that μ PyD-Net converges to the average loss value (0.122) in a few epochs (<30), already approximating such value (0.123 ± 0.02) in the very first.

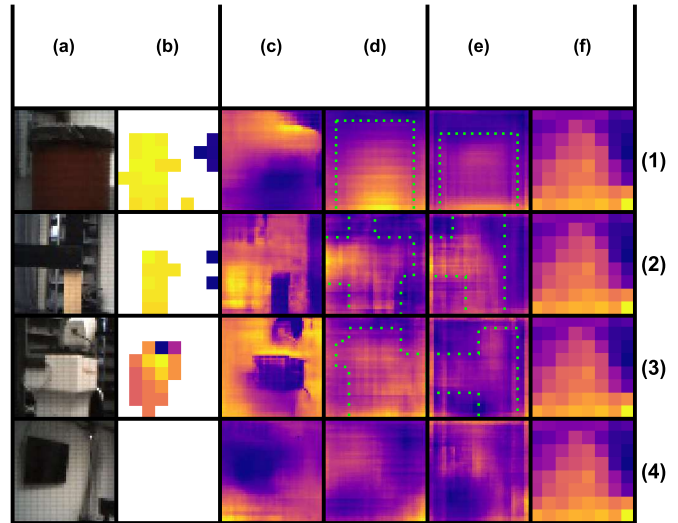


Fig. 9. Images, labels from IDSIA- μ MDE, and related μ PyD-Net predictions. (a) Image. (b) Label. (c) Baseline. (d) Sparse update (1 epoch). (e) Full update (120 epochs). (f) Dummy predictor.

Therefore, we can approximate the performance of a complete 120-epoch fine-tuning by limiting the training process to a single epoch. Even with this limited setup, μ PyD-Net achieves 0.6 m of RMSE and a δ_1 accuracy of 27.4%–3.4% less than for 120 epochs—a result comparable to 120 epochs with full update. Considering both phases of our ODL pipeline (Section IV), in which the acquisition of the 3035 training samples requires 3.37 min, our ODL routine can be completed in 17.8 min, thanks to the fine-tuning limited to a single epoch (requiring 14.4 min). Furthermore, this configuration significantly reduces the energy consumption of the MCU and the sensors to only 204.9 J for both data collection 79.5 J and fine-tuning (125.4 J).

3) *Qualitative Results*: In Fig. 9, we report a qualitative analysis of our in-field results with four labeled images [Fig. 9(a) and (b)] from IDSIA- μ MDE. As introduced in Section IV-B, the labels of IDSIA- μ MDE feature two sources of noise. First, we notice how the mismatch between the FOVs of the camera and the depth sensors produces incorrect labeling on the observed objects [e.g., depth measurements on the right side of the depth maps of Fig. 9(1) and (2)], leading to noisy training during the fine-tuning phase and a difficult pixel-wise evaluation during the testing phase. Second, we notice that the limited sensor range produces large regions with missing pixels (denoted in white), even leaving images [e.g., Fig. 9(4)] with completely invalid labels.

We show predictions of the baseline model (c) and the model fine-tuned for 1 epoch with sparse update (d), compared with the model fine-tuned for 120 epochs with full update (e) and the dummy predictor (f), highlighting detected objects with a green dotted line. Even for this dataset, the baseline model assigns a large depth to bright pixels. However, for both images 1 and 4, it also detects the central part of the main objects (recycle bin, TV) as the farthest region of each image.

In contrast, the fine-tuned models learn to distinguish the regions containing objects while imprecisely predicting details.

In particular, (d) produces coarser predictions than (e), blurring the detected object and assigning its related depth to a wider area. Overall, we observe that the predicted depth is centered in the same position as the object seen in the image, rather than shifted, as in the depth labels collected in the field. The dummy predictor, instead, only captures the correlation between the presence of obstacles in the central lower region of IDSIA- μ MDE's images, which are surrounded by free space mostly on the upper right side.

In conclusion, leveraging these fine-tuned models to estimate depth from images captured in the deployment environment provides notable advantages, even compared to the depth sensor used for generating the training labels. Although learning from noisy and sparse labels, the models produce continuous domain-specific depth maps. In contrast, the sensor is limited to detecting only nearby objects with sufficient reflectance and comes at a significant cost, being approximately $3\times$ more expensive than the camera itself.

Consequently, these fine-tuned MDE models offer a cost-effective alternative to the depth sensor, inferring depth from full scenes, including distant areas beyond the range of the sensor. Moreover, the model predictions are not limited to surfaces with detectable reflectance. Instead, they provide depth estimates for entire objects—even in areas with low reflectance—though some uncertainty remains due to the inherent limitations of the pixel-level supervision during fine-tuning. As a result, by producing dense and high-resolution depth maps from sparse sensor measurements, the system virtually expands the depth sensor's resolution, allowing richer environmental understanding.

VII. DISCUSSION

Our work presents the first ODL framework for MDE on IoT devices, addressing the challenge of autonomously determining labels for on-device adaptation. While label determination is a key step to fully exploit the potential of ODL, another critical aspect for achieving full autonomy is the detection of domain shift, that is, identifying the conditions under which our ODL routine should be triggered. In our multimodal system, domain shift could be detected by periodically activating the onboard labeler—namely, the 8×8 depth sensor—to estimate the prediction error of the deployed model during inference. To validate this approach, we conducted preliminary experiments analyzing the δ_1 accuracy distribution of our pretrained model, computed using upscaled and simulated 8×8 labels, on TartanAir (in-domain data), and KITTI and NYUv2 (out-of-domain data).

Our analysis shows that the δ_1 distributions form distinct clusters: in-domain data from TartanAir exhibit an average accuracy of $\delta_1 \approx 50.5\% \pm 21.9\%$, while out-of-domain samples from KITTI (centered at $\approx 14.7\%$) and NYUv2 ($\approx 2.3\%$) consistently fall below one standard deviation from the TartanAir mean. This clear separation indicates that a δ_1 -based in-field evaluation could effectively detect domain shifts. Consequently, the system could autonomously trigger ODL whenever the average δ_1 accuracy, estimated in the field, drops below 28.6%, a threshold that optimally distinguishes between in- and out-of-domain conditions.

Therefore, the estimation of the out-of-domain distribution can be performed during inference by activating the depth sensor at a low and constant rate (e.g., 0.1–0.5 Hz or lower), thereby ensuring uniform sampling of the data distribution while keeping the energy overhead minimal. These preliminary findings pave the way for future investigations on how to optimally tune the domain-shift detection threshold, ultimately enabling continuous learning in dynamic deployment environments.

VIII. CONCLUSION

This work presented an ODL technique to address the fundamental issue of domain shift for MDE DNNs after deployment on MCU-powered IoT nodes. Our method is based on a multimodal low-resolution sensing frontend, composed of a monocular camera, and operates in two phases. First, the system collects noisy-labeled data from the new environment using the multimodal sensors. Then, the DNN model is fine-tuned on-device to be adapted to the new deployment environment. To enable on-device fine-tuning, we introduced a novel memory-driven sparse update scheme, targeting U-Net-like models, which reduces the memory requirements of backpropagation by $2.2\times$, with respect to a full update. Thanks to the optimized deployment, based on the training kernels of PULP-TrainLib, our sparse ODL method can be performed in the field in 17.8 min, with an energy consumption of 204.9 J, making it possible to be executed on battery-operated nodes. To the best of our knowledge, our work is the first to target ODL for MDE on IoT nodes. Future work will focus on studying the conditions to detect domain shift, as well as on the criteria for the selection of the most meaningful data for optimal fine-tuning. To support the research effort on these topics, we release the code of our experiments, as well as the IDSIA- μ MDE dataset.

ACKNOWLEDGMENT

The authors thank Matteo Poggi for sharing technical details concerning μ PyD-Net.

REFERENCES

- [1] E. Cereda, A. Giusti, and D. Palossi, "Training on the fly: On-device self-supervised learning aboard nano-drones within 20 mW," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 43, no. 11, pp. 3685–3695, Nov. 2024.
- [2] P.-E. Novac, A. Pegatoquet, B. Miramond, and C. Caqueneau, "UCA-EHAR: A dataset for human activity recognition with embedded AI on smart glasses," *Appl. Sci.*, vol. 12, no. 8, p. 3849, Apr. 2022.
- [3] C. Young Lee et al., "HyperCam: Low-power onboard computer vision for IoT cameras," 2025, *arXiv:2501.10547*.
- [4] Y. Ming, X. Meng, C. Fan, and H. Yu, "Deep learning for monocular depth estimation: A review," *Neurocomputing*, vol. 438, pp. 14–33, May 2021.
- [5] V. Srivastav, A. Gangi, N. Padoy, V. Srivastav, A. Gangi, and N. Padoy, "Human pose estimation on privacy-preserving low-resolution depth images," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervent.*, 2019, pp. 583–591.
- [6] S. Sun, N. Akhtar, H. Song, C. Zhang, J. Li, and A. Mian, "Benchmark data and method for real-time people counting in cluttered scenes using depth sensors," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 10, pp. 3599–3612, Oct. 2019.
- [7] P.-Y. Chen, A. H. Liu, Y.-C. Liu, and Y.-C.-F. Wang, "Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2624–2632.

- [8] C. Shen, X. Zhao, Z. Liu, T. Gao, and J. Xu, "Joint vehicle detection and distance prediction via monocular depth estimation," *IET Intell. Transp. Syst.*, vol. 14, no. 7, pp. 753–763, Jul. 2020.
- [9] X. Yang et al., "Fast depth prediction and obstacle avoidance on a monocular drone using probabilistic convolutional neural network," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 1, pp. 156–167, Jan. 2019.
- [10] S. Lee, D. Lee, P. Choi, and D. Park, "Accuracy–power controllable LiDAR sensor system with 3D object recognition for autonomous vehicle," *Sensors*, vol. 20, no. 19, p. 5706, Oct. 2020.
- [11] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, p.27.
- [12] C. Godard, O. M. Aodha, M. Firman, and G. Brostow, "Digging into self-supervised monocular depth estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 3828–3838.
- [13] J. Feng et al., "Depth anything V2," in *Proc. Adv. Neural Inf. Process. Syst.*, 2024, pp. 21875–21911.
- [14] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia, "Towards real-time unsupervised monocular depth estimation on CPU," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2018, pp. 5848–5854.
- [15] Y. She et al., "EViTBins: Edge-enhanced vision-transformer bins for monocular depth estimation on edge devices," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 12, pp. 20320–20334, Dec. 2024.
- [16] M. Hu et al., "Metric3DV2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation," 2024, *arXiv:2404.15506*.
- [17] C. Zhao et al., "MonoViT: Self-supervised monocular depth estimation with a vision transformer," in *Proc. Int. Conf. 3D Vis. (3DV)*, Sep. 2022, pp. 668–678.
- [18] V. Peluso et al., "Monocular depth perception on microcontrollers for edge applications," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 3, pp. 1524–1536, Mar. 2022.
- [19] S. Liu, L. T. Yang, X. Tu, R. Li, and C. Xu, "Lightweight monocular depth estimation on edge devices," *IEEE Internet Things J.*, vol. 9, no. 17, pp. 16168–16180, Sep. 2022.
- [20] S. Liu, S. Zhao, P. Zhang, and J. Cheng, "Real-time monocular depth estimation for low-power embedded systems using deep learning," *J. Real-Time Image Process.*, vol. 19, no. 5, pp. 997–1006, Oct. 2022.
- [21] Q. Liu and S. Zhou, "LightDepthNet: Lightweight CNN architecture for monocular depth estimation on edge devices," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 71, no. 4, pp. 2389–2393, Apr. 2024.
- [22] C. Feng, C. Zhang, Z. Chen, W. Hu, and L. Ge, "Real-time monocular depth estimation on embedded systems," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2024, pp. 3464–3470.
- [23] S. Zhu, T. Voigt, F. Rahimian, and J. Ko, "On-device training: A first overview on existing systems," *ACM Trans. Sensor Netw.*, vol. 20, no. 6, pp. 1–39, Nov. 2024.
- [24] R. Azad et al., "Medical image segmentation review: The success of U-Net," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 12, pp. 10076–10095, Dec. 2022.
- [25] 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*. Cham, Switzerland: Springer, 2024, pp. 292–297.
- [26] H. Zheng, S. Rajadnya, and A. Zakhor, "Monocular depth estimation for drone obstacle avoidance in indoor environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2024, pp. 10027–10034.
- [27] A. Burbano, S. Bouaziz, and M. Vasiliu, "3D-sensing distributed embedded system for people tracking and counting," in *Proc. Int. Conf. Comput. Sci. Comput. Intell. (CSCI)*, Dec. 2015, pp. 470–475.
- [28] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, Sep. 2013.
- [29] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from RGBD images," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, vol. 7576, 2012, pp. 746–760.
- [30] M. Oquab et al., "DINOv2: Learning robust visual features without supervision," 2023, *arXiv:2304.07193*.
- [31] D. Nadalini, M. Rusci, L. Benini, and F. Conti, "Reduced precision floating-point optimization for deep neural network on-device learning on microcontrollers," *Future Gener. Comput. Syst.*, vol. 149, pp. 212–226, Dec. 2023.
- [32] L. Wulfert et al., "AlFES: A next-generation edge AI framework," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 6, pp. 4519–4533, Jun. 2024.
- [33] H. Ren, D. Anicic, and T. A. Runkler, "TinyOL: TinyML with online-learning on microcontrollers," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2021, pp. 1–8.
- [34] H. Cai, C. Gan, L. Zhu, and S. Han, "TinyTL: Reduce memory, not parameters for efficient on-device learning," in *Proc. Int. Conf. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 11285–11297.
- [35] F. Paissan et al., "Structured sparse back-propagation for lightweight on-device continual learning on microcontroller units," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Alberto Ancilotto, France, Jun. 2024, pp. 2172–2181.
- [36] J. Lin, L. Zhu, W.-M. Chen, W.-C. Wang, C. Gan, and S. Han, "On-device training under 256KB memory," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 22941–22954.
- [37] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 10371–10381.
- [38] S. F. Bhat, R. Birkel, D. Wofk, P. Wonka, and M. Müller, "ZoeDepth: Zero-shot transfer by combining relative and metric depth," 2023, *arXiv:2302.12288*.
- [39] K. Sunaga, T. Yamada, and H. Matsutani, "A sequential approach to detect drifts and retrain neural networks on resource-limited edge devices," *IEICE Trans. Inf. Syst.*, vol. E107.D, no. 6, pp. 741–750, 2024.
- [40] S. Lee and S. Nirjon, "Learning in the wild: When, how, and what to learn for on-device dataset adaptation," in *Proc. 2nd Int. Workshop Challenges Artif. Intell. Mach. Learn. Internet Things*, Nov. 2020, pp. 34–40.
- [41] H. Matsutani and R. Marculescu, "A tiny supervised ODL core with auto data pruning for human activity recognition," in *Proc. IEEE 20th Int. Conf. Body Sensor Netw. (BSN)*, Oct. 2024, pp. 1–4.
- [42] M. Rusci, F. Paci, M. Fariselli, E. Flamand, and T. Tuytelaars, "Self-learning for personalized keyword spotting on ultralow-power audio sensors," *IEEE Internet Things J.*, vol. 12, no. 8, pp. 10210–10221, Apr. 2025.
- [43] M. Boldo, M. De Marchi, E. Martini, S. Aldegheri, and N. Bombieri, "Domain-adaptive online active learning for real-time intelligent video analytics on edge devices," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 43, no. 11, pp. 4105–4116, Nov. 2024.
- [44] L. Zwald and S. Lambert-Lacroix, "The BerHu penalty and the grouped effect," 2012, *arXiv:1207.6868*.
- [45] C. Godard, O. M. Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 270–279.
- [46] D. Kalamkar et al., "A study of BFLOAT16 for deep learning training," 2019, *arXiv:1905.12322*.
- [47] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [48] Y. D. Kwon, R. Li, S. Venieris, J. Chauhan, N. D. Lane, and C. Mascolo, "TinyTrain: Resource-aware task-adaptive sparse training of DNNs at the data-scarce edge," in *Proc. Int. Conf. Mach. Learn.*, PMLR, Jul. 2024, pp. 25812–25843.
- [49] M. Amiri, R. Brooks, and H. Rivaz, "Fine-tuning U-Net for ultrasound image segmentation: Different layers, different outcomes," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 67, no. 12, pp. 2510–2518, Dec. 2020.
- [50] W. Wang et al., "TartanAir: A dataset to push the limits of visual SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2020, pp. 4909–4916.
- [51] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 818–833.



Davide Nadalini received the National Ph.D. degree in artificial intelligence from Politecnico di Torino, Turin, Italy, and the University of Bologna, Bologna, Italy, in 2025.

He is currently a Post-Doctoral Researcher at the University of Bologna. His research activity focuses on the hardware-efficient deployment of artificial intelligence on embedded devices and on-device continual learning.



Manuele Rusci received the Ph.D. degree in electronic engineering from the University of Bologna, Bologna, Italy, in 2018.

He has been a Post-Doctoral Fellow at the University of Bologna, and a Marie Skłodowska-Curie Actions Post-Doctoral Grantee at KU Leuven, Leuven, Belgium. He is currently a tenure-track Assistant Professor at KU Leuven. His main research interests include low-power, AI-powered smart sensors, and on-device continual learning.



Elia Cereda (Graduate Student Member, IEEE) received the double M.Sc. degree from the University of Milano-Bicocca, Milan, Italy, and USI, Lugano, Switzerland, in 2021. He is currently pursuing the Ph.D. degree with the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Lugano.

He researches robust AI-based autonomous navigation for pocket-sized robotic platforms, with a focus on self-supervised on-device learning to bridge the gap between the lab and the real world. His

work resulted in more than ten peer-reviewed publications, including Best Paper Awards at the IEEE ICCE'18 conference and ACM EWSN'23 SPICES workshop. He was part of the winning team of the first "Nanocooper AI Challenge" hosted at the IMAV'22 International Conference.



Luca Benini (Fellow, IEEE) received the Ph.D. degree from Stanford University, Stanford, CA, USA.

He holds the Chair of digital circuits and systems at ETH Zürich, Zürich, Switzerland, and a Full Professor at the Università di Bologna, Bologna, Italy. His research interests are in energy-efficient parallel computing systems, smart sensing microsystems, and machine learning hardware.

Dr. Benini is a fellow of ACM and a member of the Academia Europaea and Italian Academy of

Engineering and Technology.



Francesco Conti (Senior Member, IEEE) received the Ph.D. degree in electronic engineering from the University of Bologna, Bologna, Italy, in 2016.

He holds the position of an Associate Professor of electronics at the Department of Electrical, Electronic, and Information Engineering, University of Bologna. He worked as a Post-Doctoral Researcher at ETH Zürich, Zürich, Switzerland, from 2016 to 2020. Over his career, he has contributed to over 100 international conference presentations and journal articles, earning him multiple accolades. His

research is centered on hardware acceleration in ultralow power and highly energy efficient platforms, with a particular focus on system-on-chips for artificial intelligence applications.

Dr. Conti received the 2020 IEEE Transactions on Circuits and Systems Darlington Best Paper Award and the 2018 ESWEEK CODES+ISSS Best Paper Award, during his career.



Daniele Palossi received the Ph.D. degree in information technology and electrical engineering from ETH Zürich, Zürich, Switzerland.

He is currently a Senior Researcher and a Lecturer at the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Lugano, Switzerland, where he leads the NanoRobotics Research Group, and a Scientific Assistant at the Integrated Systems Laboratory (IIS), ETH Zürich. His work has resulted in 50+ peer-reviewed publications in international conferences and journals. His research stands at the

intersection of artificial intelligence, ultralow-power embedded systems, and miniaturized robotics.

Dr. Palossi was a recipient of multiple grants from the Swiss National Science Foundation (SNSF) and the Second Prize at the Design Contest held at the ACM/IEEE ISLPED'19, several Best Paper Awards, and led the winning team of the first "Nanocooper AI Challenge" hosted at the IMAV'22 International Conference.