

On-Device Federated Continual Learning on RISC-V-based Ultra-Low-Power SoC for Intelligent Nano-Drone Swarms

Lars Kröger¹, Cristian Cioflan^{2*}, Victor Kartsch² and Luca Benini^{2,3}

¹D-ITET, ETH Zurich; ²IIS, ETH Zurich; ³DEI, University of Bologna

Abstract

RISC-V-based architectures are paving the way for efficient On-Device Learning (ODL) in smart edge devices. When applied across multiple nodes, ODL enables the creation of intelligent sensor networks that preserve data privacy. However, developing ODL-capable, battery-operated embedded platforms presents significant challenges due to constrained computational resources and limited device lifetime, besides intrinsic learning issues such as catastrophic forgetting. We face these challenges by proposing a regularization-based On-Device Federated Continual Learning algorithm tailored for multiple nano-drones performing face recognition tasks. We demonstrate our approach on a RISC-V-based 10-core ultra-low-power SoC, optimizing the ODL computational requirements. We improve the classification accuracy by 24% over naive fine-tuning, requiring 178 ms per local epoch and 10.5 s per global epoch, demonstrating the effectiveness of the architecture for this task.

Introduction

The advent of energy-efficient RISC-V-based architectures is driving the development of intelligent edge devices, enabling AI-driven applications to continuously learn and autonomously adapt to environmental changes through federated On-Device Learning (ODL).

On-device Class Incremental Learning (CIL) [1] presents challenges in retaining previous knowledge when introducing new classes (catastrophic forgetting). Instead of storing samples from past classes, regularization-based strategies address this issue by incorporating explicit regularization terms that balance the model’s adaptation to old and new classes. In a Federated Learning (FL) context, nodes also require sophisticated strategies to distribute knowledge, particularly when only a single node encounters new knowledge domains. FL algorithms, such as FedProx [2], enable distributed learning across multiple devices by facilitating knowledge exchange without sharing raw data – essential for privacy-sensitive applications. However, On-Device Federated Continual Learning (ODFCL) must be optimized to the platform’s memory and computational constraints to ensure real-time, energy-efficient performance to meet the application requirements for latency and battery life.

In this work, we address these challenges within the context of CIL for face recognition, proposing a system for on-device distributed learning on a swarm of nano-drones using FedProx. Additionally, we introduce an ODFCL strategy implementing regularization-based

CIL through Mean Output Loss (MOL), effectively mitigating catastrophic forgetting. Our implementation outperforms naive fine-tuning by 24%. Furthermore, we demonstrate our methodology on a multicore RISC-V System-on-Chip (SoC), parallelizing the Deep Neural Network (DNN) execution on its RISC-V cluster and exploiting its Single Instruction Multiple Data (SIMD) extensions for the inference of the quantized backbone. With a training time of 10.5 s per federated epoch, our ODFCL approach demonstrates its suitability for extreme edge intelligent devices, made possible by the computational efficiency, flexibility, and scalability of the RISC-V architecture.

ODFCL System

We consider a 30 kparameters, pretrained, unimodal DSICNet [3] model. We split it into a frozen, `int8` part, M_F , and an `fp32` part M_T (i.e., the classifier and the last convolutional layer) trainable on-device. The DNN is distributed to N nodes and progressively exposed to K new classes over multiple learning sessions. The evaluation is performed on all classes, assessing the model’s ability to expand its domain knowledge without forgetting previously learned classes.

Compared to naive fine-tuning, we append a regularization term to the cross-entropy loss. MOL minimizes the difference between the activations of the classes learned during the current session and the activations of previously learned classes, while leaving out the output corresponding to the current target class. This encourages confident predictions and improves the stability of the network.

Each node n locally learns k classes, then the knowledge is aggregated: the parameters M_T^n are sent by each slave node to an ad hoc master, averaged into a global model, and returned. Each node also stores a

*Corresponding author: cioflan@iis.ee.ethz.ch

This work was partially supported by the Swiss State Secretariat for Education, Research, and Innovation (SERI) under the SwissChips initiative, the EU project dAIEDGE under grant No 101120726, and the ETH-Domain Joint Initiative program (project UrbanTwin).




Figure 1: GAP9Shield on the Crazyflie 2.1 nano-drone.

copy of the global model between two synchronization epochs, used to penalize large weight deviations of the local model, thus expanding the loss function with a third regularization term.

We implement our ODFCL methodology on the GAP9Shield [4], a Crazyflie¹ expansion board featuring GAP9 [5], a multi-core AI-capable SoC. We exploit its 10 RISC-V cores during inference and training stages and the SIMD extensions for the integerized backbone. All cores share access to the 128 kB Tightly-Coupled Data Memory (TCDM) memory containing activations, gradients, and the weights of the layer currently processed. The additional on-chip memory (i.e., 1.5 MB RAM and 2 MB flash) is used to store the weights of M_F , as well as copies of the global and local model M_T for regularization.

The GAP9Shield mounted on top of the Bitcraze Crazyflie 2.1 nano-drone (i.e., 27 g, 10 cm) is coupled with the Loco Positioning Deck allowing Ultra-wideband (UWB) communication during the synchronization stages. Control and communication are fully managed by the Crazyflie drone using its Cortex-M4 microcontroller. Fig. 1 shows the complete platform.

Results and Conclusions

We evaluate our ODFCL system on a split derivative of LFW [6], selecting 10 classes, with 28 samples per class for training and 7 for testing. At T0, DSICNet is trained on four classes, the remaining six distributed in two sessions of three classes each, one per node in a three-node swarm. We set the regularization factors $\mu = 2$ for MOL and $\lambda = 3.8$ for FedProx, with model synchronization taking place every training epoch.

Table 1 depicts the CIL accuracy of DSICNet, the base accuracy of 65% indicating the reduced capacity of the model. Despite the limitations, ODFCL achieves 46% on 10 classes, outperforming naive fine-tuning by 2× thanks to its regularization mechanism. Compared to an ideal scenario where all data is available beforehand, our system shows a modest 10% degradation, reducing catastrophic while enabling distributed learning on subsequent tasks.

The fp32 pointwise convolution represents the largest trainable component, with 29 kB peak memory

Table 1: Accuracy[%] on Split LFW after each session.

Session	T0	T1	T2
Naive fine-tuning	65%	29%	22%
ODFCL	65%	57%	46%
Joint training	65%	54%	56%

Table 2: Training cost per local epoch on the GAP9 SoC.

Mode	Low-Power Mode (LPM) $f = 240$ MHz $V = 650$ mV	High-Performance Mode (HPM) $f = 370$ MHz $V = 800$ mV
Latency [ms]	178.4	117.6
Power [mW]	24.3	53.1
Energy [mJ]	4.3	6.2

requirements during local learning, whereas 24 kB per node are sent over UWB for global aggregation. As shown in Table 2, 118 ms are needed for a local training epoch in HPM, as the RISC-V cluster operating at 370 MHz accelerates the training process. However, the UWB weight transmission by the Loco Deck takes 1.7 s, amounting to 10.5 s for a federated epoch on three learning nodes. In LPM we thus achieve a training latency of 178 ms (i.e., 6.4 ms per sample), sufficient for 58 local epochs without affecting the system latency, while benefiting from the ultra-low-power architecture at a computational cost per epoch of only 4.3 mJ. These results highlight the benefits of the RISC-V architecture, which allows real-time energy-efficient ODFCL processing, paving the way for intelligent sensor networks.

Ongoing work considers the deployment of larger DNNs, exhausting the remaining 400 kB of RAM and thus increasing the baseline accuracy. Furthermore, hybrid CIL approaches that combine regularization-based strategies with the storage of meaningful latent representations of past classes could further reduce forgetting without increasing the computational cost.

References

- [1] Yoga Esa Wibowo et al. “12 mJ Per Class On-Device Online Few-Shot Class-Incremental Learning”. In: *DATE*. 2024.
- [2] Tian Li et al. “Federated Optimization in Heterogeneous Networks”. In: *Proc. of Machine Learning and Systems*. 2020, pp. 429–450.
- [3] Carl Brander et al. “Improving Data-Scarce Image Classification Through Multimodal Synthetic Data Pretraining”. In: *IEEE SAS*. 2023.
- [4] Hanna Müller et al. “GAP9Shield: A 150GOPS AI-Capable Ultra-low Power Module for Vision and Ranging Applications on Nano-drones”. In: *European Robotics Forum*. 2024, pp. 292–297.
- [5] William G. Wong. *GAP9: 9 Cores Are Better than 8*. <https://www.electronicdesign.com/technologies/embedded/article/21126994/gap9-9-cores-are-better-than-8> [Accessed: 21 April 2025]. 2020.
- [6] Gary Huang et al. “Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments”. In: *Workshop on Faces in ‘Real-Life’ Images: Detection, Alignment, and Recognition*. 2008.

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