A Hybrid Federated-Incremental Learning Framework for Vehicle Detection

Ilenia Ficili*, Marco Garofalo*, Francesco Longo*†, Antonio Puliafito*†, Giuseppe Tricomi*†
**Università di Messina,* Messina,Italy
{marco.garofalo}@studenti.unime.it;
{ilficili, gtricomi, flongo, apuliafito}@unime.it
†*CINI: National Interuniversity Consortium for Informatics, Rome, Italy*

***Abstract*—In dynamic urban environments, smart city applications increasingly rely on edge-based classification systems capable of adapting to evolving data distributions. This work proposes an enhanced learning workflow that integrates distributed incremental with federated learning syncronization feature, aiming to mitigate the well-known problem of catastrophic forgetting. The system is designed to operate in realistic, heterogeneous contexts where data is decentralized and temporally non-stationary. Several smart city classification scenarios are considered to validate the proposed method, including urban surveillance, environmental monitoring, traffic pattern detection, and public safety systems. Each scenario is characterized by continuous input evolution and localized training conditions, requiring robust learning dynamics to preserve generalization while adapting locally. Experimental results demonstrate that the combined use of local incremental updates and federated aggregation consistently sustains high accuracy, with the ability to recover from local degradation and maintain stability over time. The findings highlight the potential of this approach for long-term deployment in intelligent urban infrastructures where continual adaptation and resilience are critical.**

***Index Terms*—Incremental Learning, Federated Learning, Smart Mobility, DILoCC**

I. Introduction

In dynamic and decentralized environments, machine learning systems must continuously adapt to new data while retaining previously acquired knowledge. Traditional centralized training approaches are inappropriate to such scenarios due to their reliance on static datasets, high communication overhead, and limited scalability, especially in edge-based deployments with temporally and spatially evolving data. Incremental Learning (IL) and Federated Learning (FL) have emerged as complementary paradigms to address these challenges. IL enables models to evolve over time by incorporating new information without full retraining, while FL facilitates decentralized collaboration among distributed clients without requiring direct access to raw data. However, when applied independently, both paradigms face limitations, IL is prone to catastrophic forgetting [1], and FL often lacks mechanisms for fine-grained temporal adaptation. This work presents an integrated approach that combines the strengths of IL and FL within a distributed architecture designed to support continual learning under evolving input conditions. Each client node incrementally updates its model on local data, and periodic synchronization via federated aggregation ensures alignment and robustness across the network. The proposed framework introduces control mechanisms to detect and mitigate potential performance regressions, enabling nodes to benefit from collaborative learning while preserving autonomy and adaptability. Simulated situations, such as mobile sensing, smart infrastructure, and autonomous systems, verify the technique. The results indicate excellent classification accuracy, consistency across learning cycles, and robustness to performance decreases caused by distribution adjustments.

II. Enabling Solutions

The proposed framework is designed to address the dual challenge of continuous data acquisition and distributed learning across edge nodes operating within smart city environments. At its core, the system combines IL and FL in a coordinated process that leverages the strengths of both paradigms: IL enables each device to adapt locally to streaming data, while FL ensures periodic global coherence across the distributed network. Each edge node separately processes incoming data in sequential batches, developing a local classification model in an incremental approach. This approach is particularly suited for contexts where data is not only temporally evolving but also non-identically distributed, as commonly observed in real-world urban deployments, such as sensor-equipped streetlights, surveillance cameras, or public transport vehicles. Our approach includes DILoCC (Distributed Incremental Learning over Continual Changes) [2], a methodology designed to successfully reduce catastrophic forgetting and preserve model stability during continual local updates. DILoCC presents adaptive control algorithms that use local validation metrics to detect performance regressions. When deterioration is detected, rollback operations are initiated to return the model to a previously verified state, preventing the loss of previously learned knowledge. Furthermore, DILoCC synergizes with the federated aggregation process, which periodically collects the local models and computes a weighted average to form a global model that is redistributed to the clients. This combination ensures that while nodes preserve autonomy in local adaptation, they also benefit from the collaborative knowledge encoded in the federated model, thus enhancing resilience and generalization. The framework is designed to be lightweight and asynchronous, making it suited for low-resource edge deployments while being scalable under varying network circumstances. It also supports optional validation

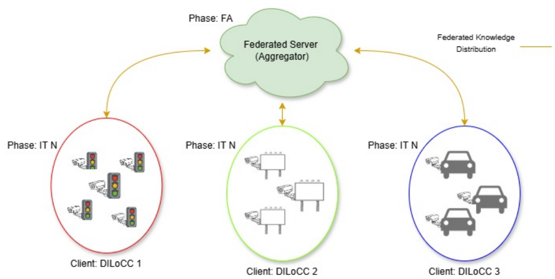

Fig. 1. Use Case Scenario.

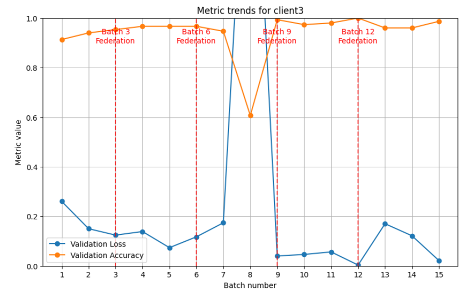

Fig. 2. Validation loss and accuracy trends across incremental batches highlighting the federated iterations for Client 3.

checkpoints, which allow to revert to prior model versions if the aggregated update worsens performance on unseen samples. The proposed framework is designed to operate effectively across a wide range of smart city scenarios, where data are generated continuously and distributed across heterogeneous nodes. This paradigm is not confined to traffic-related activities; it may be used to a variety of urban monitoring and decision-making applications. Among these, traffic scene understanding remains a particularly rich domain due to the diversity of available data sources and the dynamic nature of road environments.

In this regard, a use case deployment with three different data collecting nodes is shown in Figure 1.

- **Roadside surveillance cameras**: Combined with signal controllers, these devices continually gather traffic flow data in order to adjust signal timings, increasing vehicle throughput and minimizing congestion.
- **Smart advertising billboards**: Equipped with integrated cameras, these displays assess passing traffic—such as density, vehicle types, and movement trends—and dynamically alter advertisements. This enables context-aware and highly localized marketing strategies.
- **Vehicles with embedded vision systems**: Acting as mobile data gathering agents, these vehicles capture road conditions, traffic density, and unexpected events from various perspectives, providing vital information to enhance both traffic management and autonomous navigation systems.

These three viewpoints show the framework's adaptability to managing dispersed, constantly changing data sources.

III. Scenario and Discussion

The vehicle detection scenario is based on a Kaggle real-world dataset that includes high-resolution RGB photos tagged as vehicle-present or empty-road. The dataset mimics visual inputs common in urban settings, such as roadside cameras, smart billboards, and onboard car systems. For testing, the dataset was separated into three decentralized client partitions and one central base model division. The basic model was trained using 350 photos per class for training, and 75 for validation and testing. Each client was given 15 batches of chronologically sorted data to imitate genuine incremental updates. Each image were scaled to 64×64 and normalized to [0, 1]. A lightweight CNN was employed, featuring two convolutional layers with max pooling, followed by a dropout-regularized dense layer and a sigmoid output for binary classification. The base model initially achieved 79.3% test accuracy. Clients then refined this model locally, with federated aggregation (FedAvg) occurring every three batches. Final client models surpassed 97% accuracy, confirming the effectiveness of combining incremental and federated learning in a decentralized setting. The experimental results confirm the effectiveness of the proposed hierarchical strategy in maintaining model stability across distributed, sequential training cycles. The CNN classifier, initially trained on a centralized subset, was incrementally refined by each client node through DILoCC and synchronized periodically via federated aggregation. This hybrid structure produced a strong learning dynamic that enabled continual adaptation without compromising previously acquired knowledge. All clients achieved and retained high levels of classification accuracy, above 96%, throughout the entire training sequence, despite being exposed to distinct local data distributions and operating under isolated update regimes. The robustness of the learned representations was also reflected in the consistently low validation loss values observed after each aggregation step. These outcomes validate the system's capacity to generalize effectively under realistic non-stationary input streams, where data batches emulate temporally evolving traffic conditions. The federated phase method makes a significant contribution in recovery scenarios like the one experienced in Client 3. As illustrated in Figure 2, a sharp performance drop occurred at batch 8 due to catastrophic forgetting, caused by the model overwriting previously acquired knowledge with new data. This degradation was promptly mitigated by the federated server through periodic synchronization, which reintroduced more resilient parameter states at batch 9. The model subsequently resumed its improvement, highlighting the effectiveness of the aggregation mechanism in counteracting performance collapse.

Acknowledgement

This work was partially supported by the European Union - Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3 (PE00000014 - program "SERICS")

References

[1] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, *Continual lifelong learning with neural networks: A review*, in *Neural Networks*, vol. 113, pp. 54–71, 2019. doi: 10.1016/j.neunet.2019.01.012.

[2] G. Cicceri, G. Tricomi, Z. Benomar, F. Longo, A. Puliafito, G. Merlino, *DILoCC: An approach for Distributed Incremental Learning across the Computing Continuum*, in *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*, pp. 113–120, Irvine, CA, USA, 2021. doi: 10.1109/SMARTCOMP52413.2021.00036.