Edge-Based Anomaly Detection in Electric Vehicle Charging Infrastructure with Continual Learning

Loic Lemoine*, Amanjot Kaur*, Nhat Pham, Omer Rana
Cardiff University
{lorentelemoinel,kaura7,phamn,ranaof}@cardiff.ac.uk

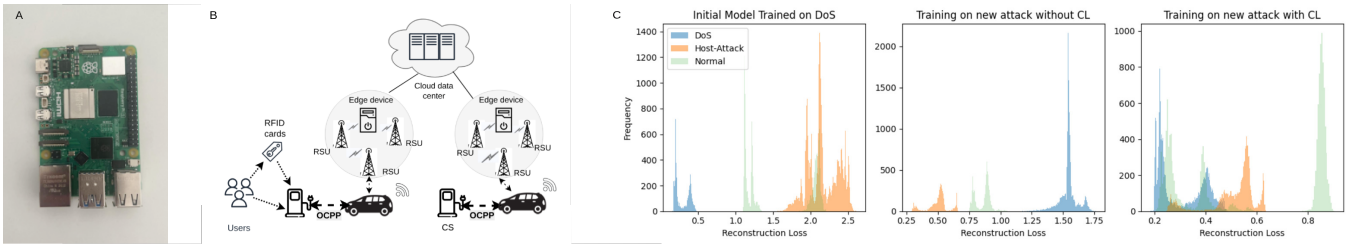

Figure 1: (A) Raspberry Pi5 edge device. (B) Conceptual architecture. (C) Local evaluation: continual learning with autoencoder.

1 Introduction

Multiple vehicle manufacturers now including Electric Vehicles (EVs) in their portfolio. This growth is supported by: (i) economic incentives to decarbonise transport; (ii) reduction in the cost of EV ecosystem. This has led to a complex ecosystem of communication protocols, including the Open Charge Point Protocol (OCPP) and ISO15118. Vulnerabilities within these communication protocols can compromise data integrity or lead to unauthorized control over vehicles or charging components.

We propose an edge-based anomaly detection framework to detect anomalous behavior in EV power consumption patterns, identifying potential cyberattacks on charging stations (often referred to as EV Charging Environments (EVCE)). We demonstrate the advantages of edge-based analysis using the CICEVSE dataset [1], differentiating benign and anomalous behavior through reconstruction loss. The proposed solution enables continual learning at the edge, facilitating adaptive, real-time detection. This approach eliminates the need to retrain the model from scratch when new attacks are discovered – reducing training overhead. Our contributions include (1) training an autoencoder on two attack types with continual learning, achieving 99% attack detection accuracy, and (2) evaluating continual learning on the edge through both edge-cloud and edge-only deployments.

2 Methods

We use the power consumption subset of the CICEVSE2024 dataset containing Reconnaissance, Denial-of-Service (DoS), Host-Attack, and Benign message types. An autoencoder is trained on attack data, enabling task-incremental learning without modifying model architecture or requiring explicit labeling. The lightweight model, which incorporates six key features (state, shunt voltage, bus voltage, current (mA), power (mW), and interface), is well-suited for edge-based EV cybersecurity applications. The experimental setup involves an edge device (Raspberry Pi 5) and a cloud system emulated on a local server. The cloud handles model training, receives data from the edge and manages communication via HTTP endpoints. After training, the updated model is sent back to the edge device. The edge performs both real-time inference and local continual learning, interacting with the cloud to exchange data, trigger cloud training, and retrieve updated models. Two operational modes are supported: (1) edge-based inference with cloud-based continual learning, followed by edge deployment, and (2) fully edge-based inference and continual learning. To evaluate continual learning, an autoencoder is trained on Denial-of-Service (DoS) attack and tested on Benign and Host-Attack data. It is then trained again on Host-Attack data with a replay mechanism that incorporates DoS data, mitigating catastrophic forgetting and enabling adaptive learning. Host-Attack is selected as the next task due to its similarity to Benign data when projected onto the first two principal components using Principal Component Analysis (PCA), posing a more challenging task for the model.

3 Evaluation & Future Work

A local evaluation was conducted in which the initial model was trained on DoS and Host-Attack. Reconstruction loss was analyzed by comparing standard training with a replay-based continual learning approach. Incorporating replayed samples effectively mitigates catastrophic forgetting. Training and inference times were measured, with a predefined set of Host-Attack samples used for retraining. The cloud-based setup (GPU) required 5.75 seconds for training, while the edge device (CPU) took 11.37 seconds. Edge inference time increased from $5.3x10^{-4}$ to $9.1x10^{-4}$ seconds during training, but remained within acceptable limits for real-time anomaly detection. Both model training utilized identical hyperparameters, resulting in a model with 99.21% attack detection accuracy. The replay-based approach successfully reduced catastrophic forgetting, yielding low reconstruction loss for DoS and Host-Attack samples and higher loss for benign samples. These findings suggest that fully edge-based operation is feasible, as the edge device can perform continual learning and real-time inference concurrently. Our future work involves continuous model training at the edge utilising triggers generated from data or the cloud platform.

References

[1] E. D. Buedi, A. A. Ghorbani, S. Dadkhah, and R. Ferreira. "Enhancing EV Charging Station Security Using A Multi-dimensional Dataset : CICEVSE2024". Submitted to ESORICS 2024 Conference. Available at: https://www.unb.ca/cic/datasets/evse-dataset-2024.html. Last accessed: April 2025

*These authors contributed equally to this work.