



AECC PROOF OF CONCEPT

Distributed Battery Electric Vehicle (BEV) Range Estimation via Personalized Federated Learning

By FEDML and Toyota Motor Corporation

Version 1.0

Abstract

Road vehicle passenger transportation accounts for about 15% of total global CO₂ emissions¹, and battery electric vehicles (BEVs) provide an important way to reduce this figure. However, one of the most common barriers to adoption for owning an electric vehicle is range anxiety — the fear that the battery may have insufficient charge for the car owner to reach their destination, leaving them stranded. One survey found that 58% of drivers say that range anxiety prevents them from buying an electric car².

The solution for range anxiety will need to include an intelligent prediction tool that helps drivers understand if their battery will last for the entire trip, given real-world driving conditions. This tool would need to account for factors like:

- Traffic congestion.
- The length of the specific route for the trip.
- The number and timing patterns for traffic lights along the way.
- Weather conditions (including cold temperatures that can drain batteries).
- Topography (including hills that will require more power to climb).
- The driver's habits (like speed and braking distance).

Bringing these complex factors together requires vast amounts of data and an AI system that can analyze the information to make accurate predictions. The problem — as with any connected vehicle service — is managing the data and processing given that these will require more computational power and bandwidth than vehicles and mobile networks can currently provide.

AECC's position is that distributed (edge) computing will be vital to making connected vehicle services work³. Using this approach, data can be collected and processed between the car and a centralized cloud.

The goal of this proof of concept (PoC) is to verify that machine learning (specifically federated learning and personalized federated learning) can make better predictions about battery life using data from electric vehicles and AECC hierarchical edge computing infrastructure⁴. Using federated learning not only enhances vehicle battery estimation accuracy, it enhances privacy protections as well.

If you're going to offer battery life prediction in your cars, however, you must also provide the low-cost infrastructure that can support this feature. This business fact inspired the second goal of the PoC, which was to confirm which method was the least intensive in terms of the use of network bandwidth and computing power.

The results showed that the predictions from the personalized federated learning system were more accurate because they adapted to simulated individual driving habits, temperature, elevation, and other factors. Further, federated learning required much less training time and used almost 10 times less network and computational power.

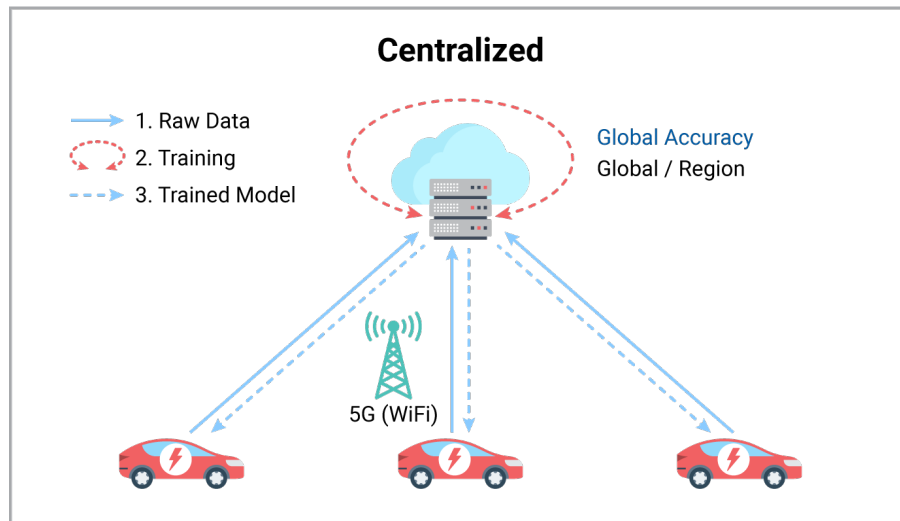
Business Strategy

The business strategy for this PoC lies in proving that very accurate battery range estimates are possible. Once deployed, this feature will, in turn, reduce buyers' range anxiety, and contribute to increased sales numbers for electric vehicles.

Proof of Concept Objective

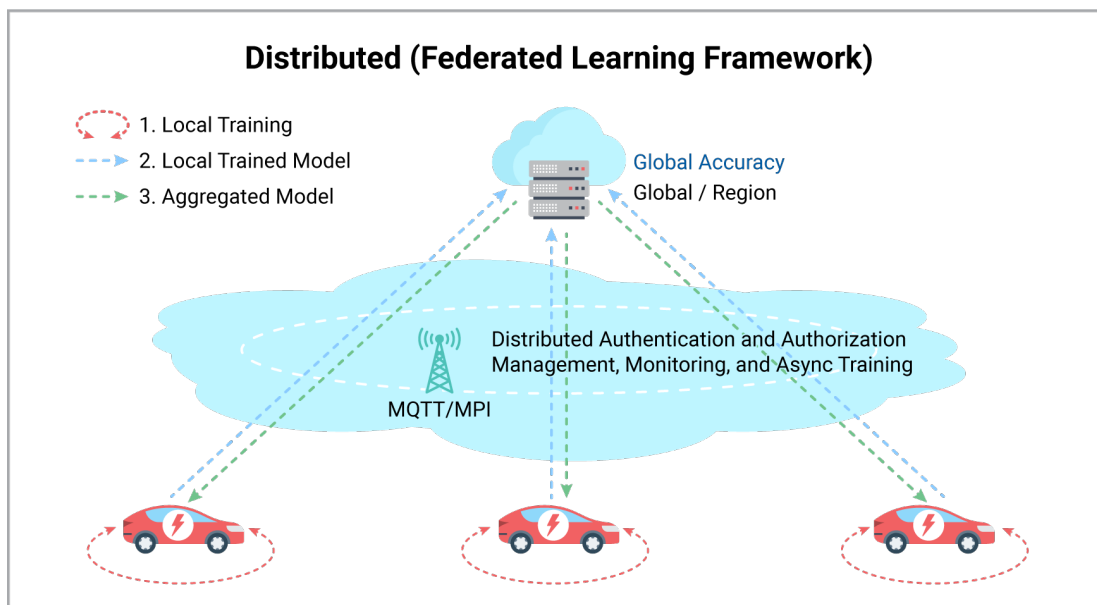
The specific objective for this PoC aims to test centralized learning (CL) against federated learning (FL) and personalized federated learning (Per-FL) for predicting battery life.

How Do the Different Models Work?



In this automotive context, centralized learning involves multiple vehicles that all generate data, which is uploaded to a central AI processing application in the cloud on an ongoing basis. The trained model is then downloaded to each vehicle and is used to predict the battery life remaining.

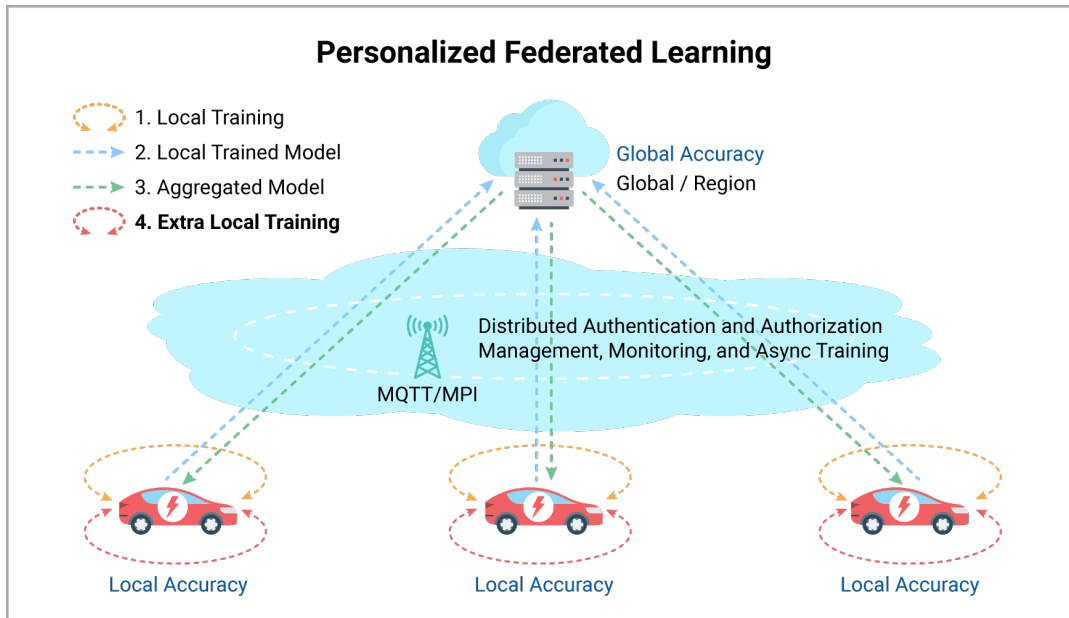
The large amount of data not only uses bandwidth during the transfer but takes up a lot of space once it gets to the cloud, which costs a lot more.



Federated learning is a type of distributed learning in which the data is generated locally and is also processed locally by separate AI models. The aggregated AI model is then shared with the cloud for wider use. Here's a step-by-step breakdown:

1. Each vehicle starts with its own copy of the AI model.
2. As the vehicle drives, it generates data which it uses to train its local model.
3. A model update (not the data it's based on) is sent to a central server. This helps keep the data private.
4. An averaged update to the global model is then calculated.
5. This aggregated model update is then shared with the local models to help improve predictions.
6. The process is repeated periodically.

Federated learning is ideal for edge computing applications and because it can leverage the data computational capacity of many devices. It also allows for the training, validating, continuous improvement, and (in one set of tests) personalized service being examined by this PoC.



In personalized federated learning, separate machine learning models are trained for individual vehicles, each with its own data set. The goal is to train personalized models collaboratively while accounting for client data disparities and reducing communication costs even further. The first few steps are the same as with federated learning, but there's an additional layer:

1. Each vehicle starts with its own copy of the AI model.
2. As the vehicle drives, it generates data which it uses to train its local model.
3. A model update is shared to a central server.
4. An averaged update to the global model is then calculated.
5. This aggregated model update is then shared with the local models to help improve predictions.
6. Each device then does additional local training, personalizing it for the individual driver's habits.
7. The local training process is repeated.

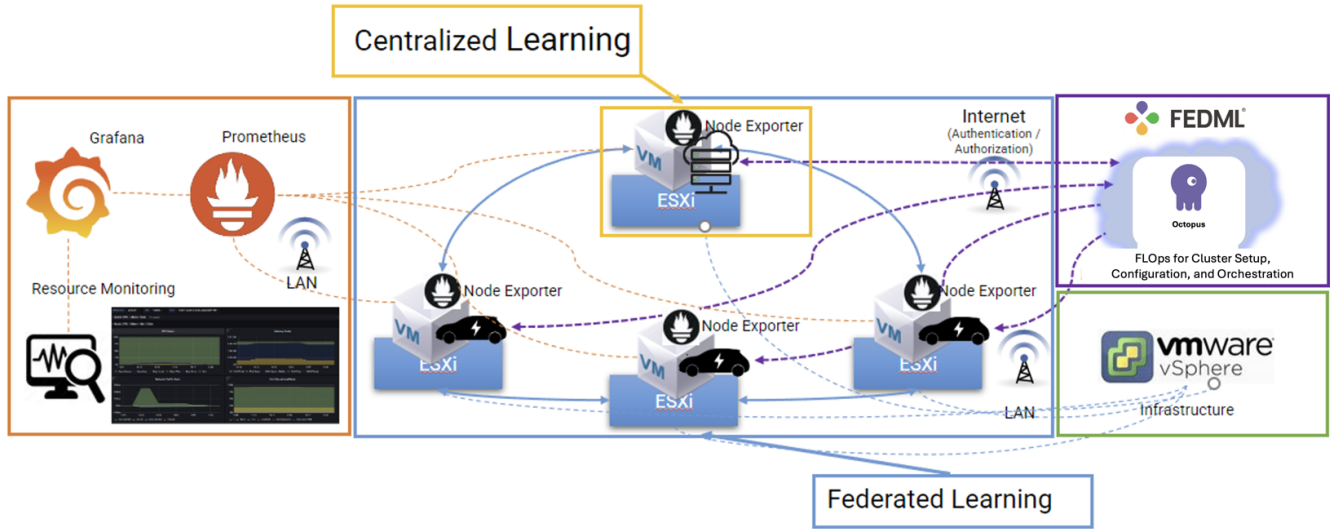
Aside from resource and cost savings benefits, constantly reconnecting to the cloud is unnecessary because you don't need a planet's worth of driving data to understand local driving situations.

Proof of Concept Scenario

This PoC used radio-controlled cars traveling on a course in an indoor lab setting, with all the computing infrastructure set up to record, process, and share the data. Each vehicle had its own computer, a Jetson Nano, to collect data and process it.

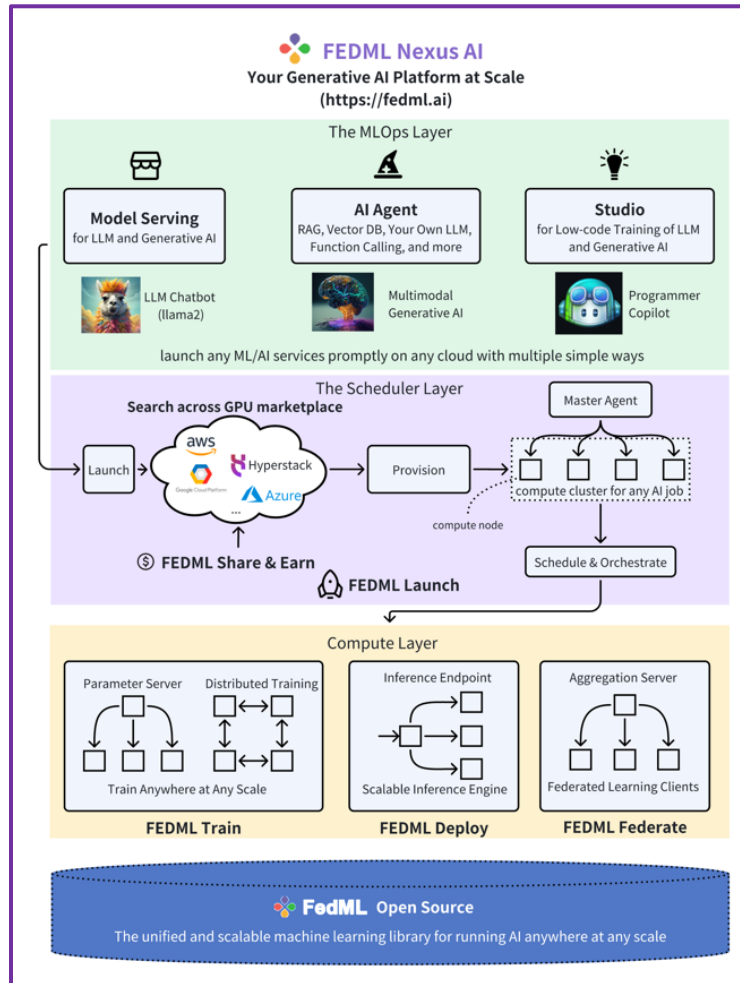
Test Architecture

Central to this PoC was the use of the FEDML⁵ Federated Learning Operations (FLOps) platform, which was chosen because it allows for flexible adjustment of parameters adjustment, monitoring of training progress, resource consumption, and continuous learning.



Proof of Concept Deployment via FEDML FLOPs Platform

The FEDML FLOPs platform is part of FEDML Nexus AI⁶, a unified AI platform for executing, managing, monitoring and deploying any AI job. FEDML Nexus AI handles all the complex operations related to job scheduling, resource provisioning and management. The figure below demonstrates the complete stack of the FEDML Nexus AI platform.



In this PoC, we used the FEDML FLOps platform to deploy and test the FL scenarios described in the previous section. FEDML FLOps facilitates the training and serving of AI models to edge servers, devices and smartphones, while enhancing data privacy compliance and optimizing development costs. A short overview of the deployment procedure of an FL job on the FEDML FLOps platform, is as follows:

1. Build, package and upload the job.
2. Bind the devices over which the job will be run.
3. Create a project.
4. Submit the job.
5. Monitor the overall system and model performance.

For more information, please refer to the FEDML documentation⁷.

Proof of Concept Results

Overview

While the federated learning model was trained in about half the time as the centralized model and with almost 10 times fewer resources, it was almost as accurate. The personalized federated learning model was even more accurate but had the training time and resource use advantages of federated learning.

The personalized federated learning model was more accurate because it adapted to individual (simulated) “driving habits”. It also resolves privacy concerns, archives personalized BEV range estimations, and reduces infrastructure costs.

Specific Results

The chart below shows that training is more than twice as fast with federated learning.

Training Time		
	Centralized Learning	Federated Learning
(Minutes: Seconds)	06:08	2:52

In measuring the accuracy, the team used an error measurement known as “test loss.” In machine learning model assessment, test loss refers to a measure of how inaccurate the model was with new data (as opposed to training loss, which is a measure of how inaccurate it was when compared to the data it was trained on.)

Accuracy Results for the AI Models					
Factor (Lower is Better for All Numbers)	Centralized Learning	Federated Learning (Average)	Personalized Federated Learning: Vehicle 1 Last Round	Personalized Federated Learning: Vehicle 2 Last Round	Personalized Federated Learning: Vehicle 3 Last Round
Test Loss	0.0025	0.0028	0.0022	0.0020	0.0023

Note: in general, the more data you have and the more training iterations you run, the more accurate your model will be. However, with federated learning, the model is trying to cover all vehicles and therefore the accuracy will be effectively averaged.

Resource Consumption		
	Centralized Learning	Federated Learning
Server Compute: The Number of Seconds the CPU Has Been in Intensive Use	152.88	17.01 (117.95 for all vehicles)
Bandwidth (MB)	274.69 MB	7.9 MB (20.66 MB for all vehicles)

The chart above reflects the fact that most of the processing is happening at the edge node/vehicle, and only the refined data is passed to the server.

Next Steps

This PoC was designed to be an initial set of tests at a small scale. Also, the dataset and vehicle resource prediction training requirements are limited in this situation, making it a good reference for AECC members.

In the next step, the team plans to introduce a couple of well-known personalized federated learning approaches to see how they can further improve individual vehicle accuracy.

They will also scale up the number of vehicles up to at least 100 and set up a larger and more difficult lab environment to see how the vehicles handle more intense terrain.

References

1. <https://ourworldindata.org/co2-emissions-from-transport>
2. <https://www.statista.com/chart/27974/reasons-for-not-buying-an-electric-vehicle/>
3. Distributed Computing in an AECC System, Version 1.0.0, August 18, 2021
4. General Principle and Vision (White Paper) Version 3.0, January 31, 2020
5. FEDML, Inc. (<https://fedml.ai>)
6. FEDML Nexus AI (<https://blog.fedml.ai/fedml-nexus-ai-the-next-gen-cloud-services-for-llms-and-generative-ai/>)
7. <https://doc.fedml.ai/>