FREEPORT

Federated Learning and Edge Processing for Safe and Efficient Operation

A project within FFI: Transport and mobility services

Contributors (alphabetically ordered):
Alfonso Lacovazzi, Carlos Camacho, Dennis Forslund, Erik Zeitler, Han Wang, Henrik Ydreskog, Josefin Agnas, Magnus Gedda, Nils Sahlberg, Parivash Pirasteh, Per Hasselberg, Sepideh Pashami, Slawomir Nowaczyk, Yuantao Fan, Zahra Taghiyarrenani, Zhenkan Wang, and many more colleagues at different stages of the work.

Date: Monday 20th October, 2025

FFI Fordonsstrategisk Forskning och Innovation

VINNOVA Energimyndigheten TRAFIKVERKET FKG VOLVO SCANIA VOLVO

Contents

1 Summary

Edge processing emerged as a promising solution, enabling a modernized data logging infrastructure that reduced transmission costs and lowered analytics latency. This provided clear benefits for vehicle manufacturers, fleet operators, and drivers.

The FREEPORT project was initiated to support electromobility transformation by addressing three key challenges faced by heavy-duty vehicle operators today: efficiency, safety, and uptime. The business value and use cases encompass monitoring electric components such as batteries, developing foundations for using third-party services in edge devices, energy consumption predictions to optimise charging, and improving functional safety through continuous surveillance to alert operators as needed. The project demonstrated edge-based data collection and processing in two vehicles, one operating at a customer site.

From a scientific perspective, the project focused on developing advanced edge analytics capabilities: real-time streaming anomaly detection algorithms tailored to the automotive sector, a versatile event-based data collection framework, a cybersecurity-aware architecture for real-time safety alerts, and comparative evaluation of commonly used federated learning methods. The research emphasized knowledge creation and pattern recognition in real-world fleet data.

The potential of next generation edge processing and edge learning was showcased through the AI Sweden Edge Learning Lab, which served as a platform to engage broader stakeholders beyond the core project consortium. Futher, specific edge processing capabilities have been demonstrated at the Boliden site. We have also organised a workshop on socially responsible and trustworthy AI.

The duration of the FREEPORT project was two years, running from September 1, 2023, to August 31, 2025, with a total budget of 12.5 MSEK. Volvo Group Trucks Technology (project coordinator), Volvo Trucks, Boliden, and Stream Analyze contributed 6.5 MSEK in-kind. The remaining 6 MSEK (48%) consisted of public funding allocated to Halmstad University, RISE (Research Institutes of Sweden), Lindholmen Science Park (via AI Sweden), and Stream Analyze, an SME.

2 Sammanfattning (in Swedish)

Kantbehandling framstod som en lovande lösning som möjliggjorde en moderniserad dataloggningsinfrastruktur som minskade överföringskostnaderna och analysfördröjningen. Detta gav tydliga fördelar för fordonstillverkare, flottoperatörer och förare.
FREEPORT-projektet initierades för att stödja elektromobilitetsomvandlingen genom att ta itu med tre viktiga utmaningar som tunga fordonsoperatörer står inför idag: effektivitet, säkerhet och drifttid. Affärsvärdet och användningsfallen omfattar övervakning av elektriska komponenter som batterier, utveckling av grunder för användning av tredjepartstjänster i edgeenheter, energiförbrukningsprognoser för att optimera laddning och förbättring av funktionell säkerhet genom kontinuerlig övervakning för att varna operatörer vid behov. Projektet demonstrerade edge-baserad datainsamling och -behandling i två fordon, varav ett körs hos en kund.
Ur ett vetenskapligt perspektiv fokuserade projektet på att utveckla avancerade edge-analysfunktioner: realtidsströmmande anomalidetekteringsalgoritmer skräddarsydda för fordonssektorn, ett mångsidigt händelsebaserat ramverk för datainsamling, en cybersäkerhetsmedveten arkitektur för säkerhetsvarningar i realtid och jämförande utvärdering av vanligt federerade inlärningsmetoder. Forskningen betonade kunskapsskapande och mönsterigenkänning i verkliga flottdata.
Potentialen för nästa generations kantbehandling och kantinlärning visades upp genom AI Sweden Edge Learning Lab, som fungerade som en plattform för att engagera bredare intressenter bortom kärnprojektkonsortiet. Dessutom har specifika kantbehandlingsfunktioner demonstrerats på Boliden-anläggningen. Vi har också organiserat en workshop om socialt ansvarsfull och pålitlig AI.
FREEPORT-projektet löpte i två år, från 1 september 2023 till 31 augusti 2025, med en total budget på 12,5 miljoner kronor. Volvo Group Trucks Technology (projektkoordinator), Volvo Trucks, Boliden och Stream Analyze bidrog med 6,5 miljoner kronor in natura. De återstående 6 miljoner kronor (48 %) bestod av offentlig finansiering som tilldelades Halmstads högskola, RISE (Research Institutes of Sweden), Lindholmen Science Park (via AI Sweden) och Stream Analyze, ett litet eller medelstort företag.

3 Background

The automotive industry is undergoing a transformation that requires moving beyond traditional centralised data collection and analysis toward distributed, real-time processing at the edge of the network. As electric vehicles (EVs) generate increasingly large and complex data streams, centralised solutions create latency, high storage costs, and limited flexibility. Edge computing, by performing AI- and ML-driven analytics directly at the source device where data was collected, addresses these challenges by enabling adaptive, real-time, and customisable data use. This is particularly critical for EVs, which are inherently more service-dependent than internal combustion engine vehicles, demanding robust digital infrastructure for safety, efficiency, cost-effectiveness, and compliance with privacy regulations such as GDPR.
To achieve address these challenges, the FREEPORT project applies edge computing to address major challenges for efficient heavy-duty electric truck operation: efficiency, safety, and uptime. By enabling real-time, context-specific AI/ML solutions such as anomaly detection, the project aims to make EVs more reliable and attractive for businesses, strengthening the competitiveness of Sweden's premium vehicle industry. Unlike traditional "one size fits all" approaches, edge processing supports customised data analysis, helping OEMs and fleet operators optimise operations and resource use. Practical applications include predicting energy consumption to improve operating schedules, detecting anomalies in critical EV component and tailoring vehicle configurations to specific tasks, reducing costs and environmental impact.

4 Objectives

FREEPORT project achieved the following objectives which are well aligned with the stated objectives in the application.

4.1 Data collection and distributed learning for resource constrained devices

The project showed how a compact zero-dependency analytics software (SA Engine) could be installed in resource constrained on-board vehicle compute units, including loggers and telematics units. Each instance of SA Engine connected to a centrally hosted platform, SA Federation Services (SAFS). Users

(analysts and engineers) interacted with SAFS in order to see, understand and model real time data streams on the vehicles. On-board data sources included CANBUS, accelerometer, and positioning systems. In particular, users were able to develop, test, release, and deploy analytical and AI models to be executed on the on-board compute units in the vehicles. Using SA Engine and SAFS, several models were deployed on-board vehicle compute units, including Anomaly detection, Energy Consumption Forecasting, and Federated Learning. Furthermore, data from a vehicle was collected using SA Engine.
In addition to the data collection and analytics performed using SA Engine, data was also collected using CANBUS logger equipment.

4.2 Privacy-perserving Teechnique

Privacy-preserving methods are designed and applied to protect AI systems when trained on sensitive or confidential data. In some circumstances, ML models can memorize specific data samples, constituting a heavy privacy breach. These methods are tailored for use within the FL framework, aiming to protect against adversaries on the server (aggregator) side—especially in cases involving honest-but-curious servers that try to extract private information from client updates.

4.3 Anomaly detection

Real-time anomaly detection algorithms are developed for critical components in energy storage systems (ESS), such as battery modules, enabling early identification of abnormal behaviours that could lead to unexpected downtime, e.g., thermal events that may develop into thermal runaway. The developed techniques are expected to demonstrate their capability on board an edge computing device in real-time.

4.4 Forecasting energy consumption (ECF)

The project also addressed ECF for the auxiliary systems in heavy-duty electric vehicles, including ESS and cabin heating, air conditioning, a compressor for air brakes, and a high voltage DC to low voltage converter that supports various functions for onboard electrical and electromechanical systems.

Although these systems consume less energy compared with the propulsion system, they could have a substantial, and sometimes indirect, impact on overall energy consumption, substantially affecting vehicle operating range, especially under varying operating conditions. The project aims to develop predictive models based on high-frequency, multivariate sensor data to perform accurate short-term, within-trip forecasts of auxiliary energy consumption. These forecasts will improve range prediction and route planning, ensuring more reliable operation and improved energy management strategies for heavy-duty electric vehicles.

4.5 Federated Learning (FL)

FL has been implemented as a paradigm for collaborative model training across multiple data sources without centralising sensitive information. This approach addressed strict privacy and regulatory requirements, particularly in domains where data sharing is restricted, and reduced resource demands by avoiding the transfer of large datasets. Instead, only compact model parameters or updates were exchanged, enabling scalable and privacy-preserving learning adapted to domain-specific constraints.

4.6 Open Set Recognition (OSR)

The project advanced OSR capabilities, enabling AI systems to detect inputs that fall outside the set of classes observed during training. While related to anomaly detection, which flags irregular or abnormal patterns, OSR explicitly distinguishes between known and unknown classes, thereby improving decision-making in dynamic and evolving environments. For example, in an industrial monitoring context, a system trained to recognise known fault types may encounter a new, previously unseen fault. In such cases, AI models should be equipped with mechanisms to detect the unfamiliar pattern, distinguish it from both normal behaviour and known faults, and trigger appropriate responses.

4.7 Integrated OSR–FL Framework

The project achieved and presented an integrated OSR–FL framework to address real-world distributed scenarios where novel patterns may emerge locally at individual nodes. The framework allowed each node to detect

and manage local unknowns while collaboratively refining a global decision boundary for rejection and classification. OSR signals, such as uncertainty scores, were incorporated into federated updates, without sharing raw data, to calibrate thresholds, adapt representations, and align class boundaries across sites. This integration enhanced system reliability under distribution shifts and the emergence of novel categories across diverse, heterogeneous data sources.
The following method and results sections elaborate on each of these stated goals.

5 Method

5.1 Data Collection and Distributed Learning for Resource Constrained Devices

The Stream Analyze (SA) platform for interactive edge analytics was leveraged in the project. To accommodate the use cases and deliver the demonstrators, the platform was adopted and extended.
Some primitives could be easily adopted with minimal effort. For example, the SA platform included buffering capabilities, which was leveraged when network connectivity was intermittent. Furthermore, the primitives of the SA platform supported federated learning with minimum adoptions, as was demonstrated in the project.
Extensions implemented in the project include communication primitives specific for the network set-up in the project, and wrappers for the signals available in the telematics units used in the project.
A load shedding method was implemented to allow selective deletion of result streams from telematics units. This was used to prioritize critical data over less critical data when network bandwidth was limited.

5.2 Secure Group Communication

We consider the case of training models within a vehicle's electronic systems. For each onboard device, such as an electronic control unit (ECU), that transmits its model parameters to the central aggregator, securing the communication channel is essential. Without proper protection, adversaries could intercept or tamper with model updates, potentially compromising model integrity or exposing sensitive information from local datasets. Establishing secure communication is therefore critical to preserving data privacy, maintaining model reliability, and ensuring the overall trustworthiness of the FL process. In FREEPORT, we explore the use of secure group communication to enhance the communication efficiency and performance of the federated learning process, while still ensuring its security and privacy. By using group communication, we aim to reduce the model convergence time and to improve the communication efficiency and overall system performance.

5.3 Consensus Self-organizing Models

In this project we have adapted, applied and evaluated an unsupervised anomaly detection method, consensus self-organizing models (COSMO), to detect early signs of faults in ESS (focus on battery packs) across heterogeneous mission and external conditions Fan et al. (2025). We have collected on-board sensor data as a multivariate time series, including information such as voltage, current, temperature, state of charge, etc. Given the wide range of applications of heavy-duty vehicles, these signals typically exhibit extreme variability even under normal operation, making anomaly detection challenging.
We extend the baseline COSMO framework by incorporating causal discovery algorithms, enabling the construction of causal graphs that reveal latent structures underlying relationships among influential features. For each battery module, the learned causal graphs provide a consistent representation of its usage and behavior over time. Since battery modules within the same ESS are expected to behave similarly under comparable operating conditions, COSMO treats them as a homogeneous group. Battery modules whose causal graph representations deviate from the majority are then flagged as anomalous.

5.4 Fleet-based Regression (FBR)

Accurate energy consumption prediction is vital for optimizing heavy-duty electric vehicle operations, such as charging-aware route planning. Moreover, to build trust and enable real-world use, predictions must also be interpretable. Since commercial vehicles operate differently as transportation tasks, ambient, and drivers vary, a heterogeneous population is expected when building an AI system for forecasting energy consumption. The dependencies between the input features and the target values are expected to also differ across sub-populations.

In our study Fan et al. (2023), we propose and evaluate a solution involving the training of multiple regression models on various subsets of data. FBR has two key advantages. First, it allows us to identify relevant sub-patterns within the global regression task, leading to higher overall regression performance. Second, it uncovers more relevant patterns within the data, leading to more consistent Local Interpretable Model-agnostic Explanations (LIME) explanations. This method capitalizes on the idea of identifying relevant sub-populations within the overall dataset. By dividing the dataset into these groups, we ensure that the associations between the input features and the target values remain consistent within each individual group, even if they differ across separate groups. Our approach is fundamentally focused on discovering pertinent sub-patterns within the overarching regression task, which breaks down a complex problem into smaller, more manageable parts, making it easier for LIME to find suitable explanations.

5.5 Open Set Recognition (OSR)

We studied open set recognition (OSR) for time-series data, since time-series was the primary modality of interest in our use cases. Our approach was not limited to developing a single method; instead, we conducted a systematic analysis of different OSR techniques and representation learning backbones, in order to evaluate their effectiveness and efficiency when applied to time series. While many OSR approaches have been proposed in the computer vision community, their direct application to time series has not been studied in depth.

To address this gap, we adapted several representative OSR methods to the time-series domain and combined them with different backbone architectures to investigate how both the choice of representation and the choice of OSR mechanism influence performance. Among these, we also tested outlier exposure, where synthetic or auxiliary outliers are introduced during training to improve the model's ability to reject unknown inputs.

The evaluation was carried out on diverse benchmark datasets, including UCR, UEA, and HAR, covering a wide range of time-series applications. This comprehensive analysis allowed us to identify which combinations of OSR method and backbone were most effective.

5.6 Integrated OSR–FL Framework

The framework integrates OSR with FL to address the challenge of fault detection and identification in distributed time-series environments, such as fleets of vehicles. Traditional classifiers struggle when new, unseen faults appear, often misclassifying them into the closest known category. OSR remedies this by enabling models to recognize known faults while detecting previously unseen ones, effectively allowing the system to say, “I don’t know.” This evolving classification capability is critical for real-world deployments where new fault types continually emerge. Based on the comparative study of OSR methods conducted within the project, we selected the most efficient approach and embedded it into the federated framework.

To support this capability in a distributed setting, the framework incorporates FL, ensuring that edge devices (e.g., vehicles or buses) can collaboratively improve their models without sharing raw data. Each client device trains an open-set classifier locally, detects potential new faults, and communicates findings with others. When a new fault is identified, devices update their classifier heads and retrain with both the new and known samples, preventing catastrophic forgetting. Lightweight strategies such as data augmentation are used to handle situations where only a single example of a new fault is available.

The framework is further enhanced by leveraging pretrained time-series models (Moment) as feature extractors, reducing the need for large-scale data collection from scratch. Shallow classifiers trained on these embeddings achieve high F1-scores for known fault classification, while confidence thresholds and OOD detection techniques enable robust recognition of unknowns. Experiments across multiple clients demonstrated that the integrated OSR–FL setup achieves near-perfect classification of known faults, effective detection of unknowns, and the ability to evolve dynamically as new fault categories are introduced, all while maintaining privacy and efficiency in distributed environments.

6 Results and deliverables

6.1 WP0: Knowledge dissemination and project management

The knowledge dissemination of the project includes workshops, invited talks at conferences, public demonstrators. In addition, the project has led to publications, presented at workshops and conferences, see details in section 7.

6.2 WP1: Cybersecurity for automotive edge processing

6.2.1 DP-based Time-series data synthetization

When dealing with sensitive data—such as records of driving behavior—privacy is critical to ensuring broader acceptance of the resulting models. In some instances, neural networks may memorize individual data samples, leading to severe privacy risks Feldman (2021). A promising direction Jordon et al. (2022) involves replacing the original, potentially sensitive dataset with a synthetic one that preserves key statistical properties of the raw data. To this end, we propose a privacy-preserving data synthesis approach to mitigate these concerns in this project.
We propose AE-(dp)MERF, which is designed based on DP-MERF Harder et al. (2020). AE-(dp)MERF is an efficient all-purpose data generation algorithm based on minimizing the so-called Maximum Mean Discrepancy (MMD) between the real and the synthetic data distributions. It employs kernel mean embedding to transform the underlying probability distribution of the original data into a Hilbert space. The MMD then measures the distance between two distributions in the Hilbert space. In addition, we employ differential privacy (DP) Dwork et al. (2014), which is a formal mathematical framework that ensures information about individual records in a dataset cannot be revealed, even while allowing valuable insights about the dataset as a whole to be shared. It involves the privacy budget (ϵ), which measures the privacy loss. A smaller ϵ means stronger privacy but less accuracy. To generate time series data, we use an AE architecture that maps the time series to a compact latent representation of fixed dimension. Data points in the latent space can then be treated as tabular data with no temporal dependencies.

We evaluate the utility of synthetic data generated by AE-(dp)MERF. The utility is measured in an automotive scenario, involving vehicles with Adaptive Cruise Control systems (ACC). Detecting ACC driving from human driving behavior is valuable for identifying safety risks and near-miss scenarios unique to each driving mode. To this end, we train an Long Short-Term Memory(LSTM)-based classifier to distinguish ACC driving patterns from human driving behavior, while simultaneously aiming to preserve the privacy of individual driving habits. We synthesize the data using the AE-(dp)MERF. We are testing DP and non-private versions of the models and comparing them to a baseline model that is trained on the original data.
We compare results under both centralized and FL settings. In the FL setup, training involves three clients, each holding data collected from different vehicles, dates, and locations—resulting in a non-IID data scenario. Empirical results show that model utility is already reduced during the (non-privacy-preserving) data generation phase, with AE-MERF experiencing approximately a 5% drop in performance. Introducing DP does not significantly degrade model utility up to a certain privacy budget ϵ. The F1 score decreases by only about 3% under the lowest ϵ setting, which is 0.01. Notably, overall model performance improves in the FL scenario, which is expected given that collaboratively trained models tend to be more generalized than those trained in a centralized manner. These results also demonstrate that even when models are trained on highly privacy-preserving synthetic data generated by the AE-(dp)MERF, they can still achieve strong utility within an FL framework.

6.2.2 Secure Group communication

The Constrained Application Protocol (CoAP) Shelby et al. (2014) is a lightweight web transfer protocol designed for IoT environments, typically operating over the unreliable transport layer UDP. In one-to-one communication scenarios, CoAP messages can be secured using Datagram Transport Layer Security (DTLS) Rescorla et al. (2022) or Object Security for Constrained RESTful Environments (OSCORE) Selander et al. (2019), with the latter offering end-to-end protection even in the presence of untrusted intermediaries. In addition, CoAP over UDP supports group communication, enabling a single request message to be delivered to multiple recipients simultaneously.

6.3 WP2: Edge-based versatile data collection

The project showed how versatile data collection can be facilitated by installing the compact zero-dependency analytics software SA Engine in resource constrained on-board vehicle compute units, including loggers and telematics units.

6.3.1 Event-based data collection using high-level query language

The installation of SA Engine on loggers and telematics units enabled both interactive and autonomous control of data collection from fleets of vehicles and was demonstrated in various settings throughout the project (see Section 7 dissemination and publication).
The implemented demonstrator was event-based data collection, where data was only saved under specific conditions and then streamed immediately to the SAFS. The events were designed as aggregates of operation-critical sensor data indicative of anomalies that could result in severe damages to the vehicle and risk physical damage to the operator.
The event-based data collection facilities of SA Engine included a buffer, enabling users to specify collection of historical data from sensors of interest leading up to the event. This facilitated insights to the root cause by observing the vehicle behavior not only after the event but also before the event.
To enable detailed root cause analysis of the anomaly event, the buffered historical data was of high frequency sensor data. In the demonstration at Boliden, the collected data was streamed to the SAFS. SAFS in turn was streaming the this data in real time to data management systems and monitoring centers on multiple off-site locations both at Volvo Data Science Lab in Gothenburg as well as Boliden's monitoring center in Boliden. The real-time data was also simultaneously tapped into by a user with a client terminal who could monitor the data streams on a workstation.
Aggregation and combination of signals, or analytics that determine trends in data values, can be implemented in the SAFS using the same high-level query language as for the data collection models before relaying the results to external data centers or monitoring centers.
The on-board data sources used for data collection included CANBUS and positioning systems, and the users were able to develop, test, release, and deploy the data collection models to be executed on the on-board compute units in the vehicles.

6.4 WP3: Advanced data analytics on the edge

6.4.1 Anomaly Detection

Within the scope of the project, two anomaly detection methods Ozen et al. (2025); Fan et al. (2025), adopting graph neural networks and causal relation discovery, were developed and applied to detection anomalies in real-world dataset, resulted in scientific report being published at relevant venue. In addition, two demos were presented to partners, demostrating anomaly detection functions excuted in real-time within an edge computing envrionment. As an example, the study for detecting anomalies in battery modules was presented at the Boliden demo.

The two methods have been applied to real data from Volvo trucks, focusing on identifying deviations and abnormalities in battery modules. The results of this have shown the ability of these methods to accurately capture these abnormalities, providing valuable information on the potential early symptoms of malfunctions. By detecting these issues early, our approach facilitates more effective root cause analysis, making it easier for the industry to address and resolve potential problems before they escalate. This not only improves vehicle reliability and performance, but also contributes to more efficient maintenance strategies, supporting the operational goals of fleet management.

6.4.2 Energy Consumption Forecasting

The project has explored fleet-based regression Fan et al. (2023) for Energy Consumption Forecasting (ECF), aiming for an improved performance and more consistent explainability, and curriculum learning with auxiliary tasks for more efficient learning and adapatation for the primary forecasting tasks Fan et al. (2024b). A comprehensive evaluation Wang et al. (2025) with various conventional and deep learning models, e.g. recurrent neural networks with long short-term memory, for ECF on real-world dataset was conducted. The ECF studies resulted in three papers, presented at AI conferene and international EV Symposium.

In the demo at Boliden, an energy consumption forecasting model was demonstrated live on an electric truck going down and up the mine with different loads. The model predicted the state of charge at certain depths in the mine based on factors such as initial state of charge, truck load, regenerative braking effect, and auxiliary energy consumption.

6.4.3 Federated Learning

In the demo at AI Sweden a federated learning model was demonstrated on that used federated fault detection for geofencing of icing regions in real-time on streaming vehicle data. The data was recorded from the CANBUS and position data recorded from a fleet of buses operated by Västtrafik. We refer to Section WP4 for a detailed description of the demonstration.

6.5 WP4: Business value demonstrator

In today's discussion of AI in industrial applications, and edge analytics in particular, focus is largely on the application themselves rater than the deployment and industrialization of such. To make FREEPORT more tangible and applicable to industrialization, it was decided to not just develop the models, but to actually demonstrate the capability on a live Volvo truck in a Boliden mine. In the same manner MLOps was further shown, with Federated learning in focus, in the second demonstrator at Lindholmen, Gothenburg. This demonstration was based on continuous data from Volvo Buses operated by Västtrafik. These two demonstrations are referred to as the "Boliden on-site demonstrator", and "Edge processing solution for next generation automotive".

Through these Business value demonstrators, the FREEPORT project takes the use cases from theory and laboratory settings - to reality and industry.

Two business value demonstrator activities were performed during the FREEPORT project: *(i)* Edge processing solution for next generation automotive, and *(ii)* Boliden on-site demonstrator.

6.5.1 Edge processing solution for next generation automotive

On March 5th, 2025, AI Sweden, Halmstad University and RISE organised a half-day workshop to present the project's results regarding the electromobility transformation and the challenges faced by heavy vehicle operators.

The workshop focused on technical and practical challenges for edge processing in the automotive industry. Sessions included demonstrations of, among others, synthetic data generation, anomaly detection with graphical neural networks and federated learning, enabling improved safety. Participants gained insights into data stream processing with SA Engine and edge processing for vehicle systems on a Volvo test truck, followed by in-depth discussions. The goal was

to present the findings in an accessible way to relevant Swedish stakeholders in sectors such as automotive, healthcare and space. The event was aimed at researchers, engineers and professionals in automotive and edge technology. 50 participants were registered in person, and 30 were online.
The workshop was the publicly available demonstrator that was promised as Milestone 4. This includes a live demo of Edge analytics, open-set fault detection, a demonstration of federated denstream for geofencing icing regions, and a discussion.
The **Real-time Federated Road Condition Monitoring** demonstration set-up is shown in Figure 1, using CANBUS and position data recorded from a fleet of buses operated by Västtrafik. Data from each bus was replayed on a Host Mobility mx4 telematic units[1]. SA Engine was installed on each telematic unit. Each SA Engine instance connected to SA Federation Services (SAFS). Users interacted with SAFS using SA Studio. A screenshot of SA Studio with its real-time graph capabilities is shown on the left hand side of Figure 1.
Using SA Engine and SAFS, we showed the following functionality in the demo, while re-playing the data in each telematics unit:

- Interactive analytics on the connected telematic units, enabling remote diagnostics on the bus fleet.
- Illustrating lifecycle management of AI and analytics models over fleets of edge devices, we showed a model that performed denstream clustering over streams on each edge, and showed how to release and deploy it on the bus fleet.
- The denstream clustering model on each edge captured the geo location whenever the ABS brake system triggered. When ABS was triggered frequently in a certain area, that area was considered a cluster with slippery road conditions.
- On a regular basis, the denstream clustering model on each edge sent its clusters to a denstream clustering model on the SAFS. The SAFS denstream clustering was in turn forming super-clusters on these clusters. Once super-clusters were formed, these were distributed to all edges. That way, all edges learned from each other.

[1] https://hostmobility-eng.setek.se/host-mobility-product-mx-4-t30-fr/

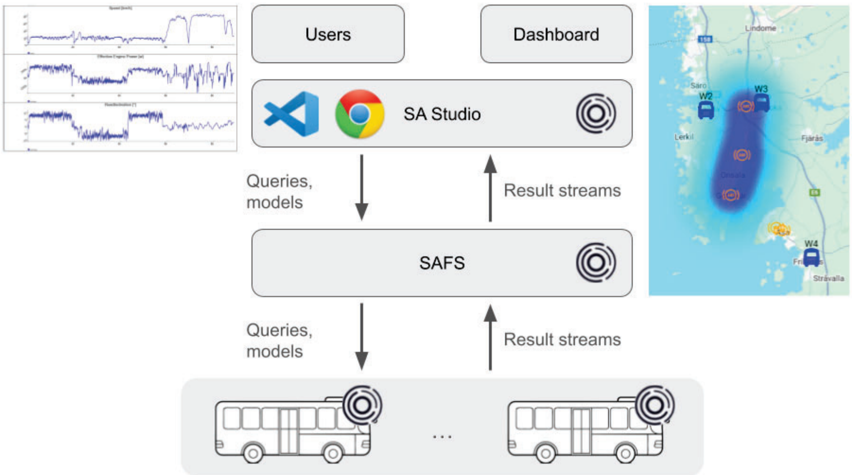

Figure 1: Demonstration set-up for Real-time Federated Road Condition Monitoring.

- For each super-cluster C_i, the model set up a geofence area A_i around C_i. We showed how to alert the driver of each bus B_j when it entered such a geofence area A_i.
- We also showed a map updated in real-time with the current location of all buses $\{B_i\}$, along with all super-clusters $\{C_i\}$ and all geofence areas $\{A_i\}$. A screenshot of the map shown in the demo is included on the right hand side of Figure 1. Such a map is useful e.g. in a bus fleet operator's dashboard.

.
The functionality shown using SAFS and SA Engine enables these important capabilities:

- Data streams from vehicles could be performed Interactively queried. This enables analysts to see, understand and model data – quickly.
- A new model can be deployed on a fleet of connected vehicles in seconds.
- Federated learning can be performed in real-time on streaming vehicle data, in vehicle telematic units.

6.5.2 Boliden on-site demonstrator

On June 3rd, 2025, the project partners organised a one-day workshop to present business value demonstrators. Figure 2 shows the set-up of the demonstrator at Boliden.
SA Engine was installed in a on a Host Mobility mx4 telematic unit in an electric truck in service for Boliden in the Kankberg mine. SA Engine had access to CAN bus and various sensors on the telematic unit. The SA Engine instance connected to SA Federation Services (SAFS) installed in the Data Science Lab (DSL) on Volvo's premises. Due to specific network configurations, connecting SA Engine to SAFS involved adoptions to traverse a tunnel on the telematics unit as well as firewalls (FW) and an application gateway (GW) at Volvo's network.
Volvo's analysts interacted with this SAFS using SA Studio or a web interface, as indicated in the lower left of Figure 2. Using SAFS, Volvo's analysts could interact in real time with the truck in operation (indicated with the curved purple arrow). This enabled real-time insights and remote diagnostics. SAFS and SA Engine also allowed Volvo's analysts to deploy AI models on SA Engine in the telematics unit with a few seconds turn-around time. For example, an ML model was deployed on the truck to estimate the state of charge (SOC) of the batteries during operation. A real-time alarm was configured to notify the user whenever the value exceeds a pre-defined threshold. Furthermore, SA Engine allowed result streams of queries and models to be streamed not only to SAFS, but also to other data infrastructure. In the demo, this capability was shown in two ways:

- SA Engine on the telematics unit was streaming model results to SAFS, which in turn streamed these results to a Volvo internal MQTT, to which a Grafana dashboard was connected.
- Using SAFS, Volvo's analysts also configured SA Engine on the telematics unit to stream selected data directly to an MQTT instance in Boliden's network. The data on this MQTT instance was then used to inform Boliden's mine operations.

This set-up showed how data could be dynamically shared between multiple parties to ensure timely access to relevant information.

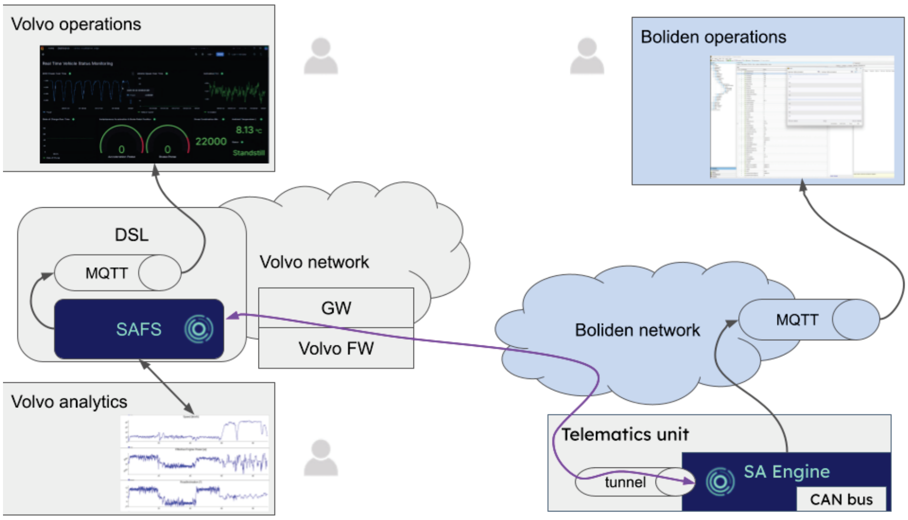

Figure 2: Demonstration set-up for final project workshop.

Figure 3: Telematics unit mounted in truck, marked in red.

7 Dissemination and publications

7.1 Dissemination

- Socially Responsible and Trustworthy AI workshop, June 18th, Halmstad University, 2025.
- Technical demonstrator of an integrated cloud-edge system in a truck in operation. Using this system, models for energy consumption prediction and critical structure monitoring were shown. June 3rd, Boliden, 2025.
- Public demonstrator on edge processing solutions for next generation automotive, March 5th, AI Sweden platform, 2025.
- Tutorial and workshop on Explainable and Robust AI for Industry 4.0 & 5.0 (X-RAI) organized by Sepideh Pashami at ECML PKDD 2024 conference.
- Sepideh Pashami presented research on Predictive Maintenance for vehicle companies at 8th intelligent maintenance conference, EPFL 2024.
- A panel session was arranged at the Underhållsmässan 2024 at Gothenburg, in March
- Erik Zeitler presented the Stream Analyze platform at Underhållsmässan 2024 at Gothenburg, March 2024.
- AI Sweden workshop is organised on "sharing is caring", 2024.

7.2 Publication

The FREEPORT project has led to the following publications and master thesis that corroborate the objectives defined in Section4.

- Accepted and/or Published:
  - Wang, Z., Fan, Y., Ydreskog, H., and Nowaczyk, S. (2025). Investigation on machine learning models for forecasting auxiliary energy consumption of hd bevs. In *The 38th International Electric Vehicle Symposium & Exhibition*. Electric Vehicle Symposium and Exhibition

  - Ozen, C., Nowaczyk, S., Tiwari, P., and Pashami, S. (2025). Assessing the graph structure learning in graph deviation networks. In *International Symposium on Intelligent Data Analysis*, pages 97–109. Springer
  - Fan, Y., Nowaczyk, S., Wang, Z., and Pashami, S. (2024b). Evaluating multi-task curriculum learning for forecasting energy consumption in electric heavy-duty vehicles. In *Workshop on Embracing Human-Aware AI in Industry 5.0 (HAII5. 0 2024) co-located with the 27TH EUROPEAN CONFERENCE ON ARTIFICIAL INTELLIGENCE (ECAI 2024)*, volume 3765. CEUR-WS
  - Fan, Y., Altarabichi, M. G., Pashami, S., Mashhadi, P. S., and Nowaczyk, S. (2024a). Invariant feature selection for battery state of health estimation in heterogeneous hybrid electric bus fleets. In *CEUR Workshop Proceedings*, volume 3765. CEUR-WS
  - Fan, Y., Wang, Z., Pashami, S., Nowaczyk, S., and Ydreskog, H. (2023). Forecasting auxiliary energy consumption for electric heavy-duty vehicles. *arXiv preprint arXiv:2311.16003*
  - Taghiyarrenani, Z., Nowaczyk, S., and Pashami, S. (2023). Analysis of statistical data heterogeneity in federated fault identification. In *PHM Society Asia-Pacific Conference*, volume 4
  - Pirasteh, P., Camacho, C., and Nowaczyk, S. (2025). Factors influencing propulsion energy consumption in battery electric heavy-duty trucks on urban and rural routes. In *The 38th International Electric Vehicle Symposium & Exhibition*. Electric Vehicle Symposium and Exhibition

- Submitted and/or ArXiv:
  - Fan, Y., Camacho, C., Pashami, S., and Nowaczyk, S. (2025). Causal graph-based anomaly detection for battery modules in electric heavy-duty vehicles. In *PHM Society Asia-Pacific Conference*

- Master Theses:
  - Martin, S. and Mohan, S. (2025). Predicting energy consumption for heavy-dutyvehicles: With an emphasis on auxiliary consumption: In collaboration with volvo trucks

8 Conclusion and future research

During the FREEPORT project, cutting-edge data analytics on the edge was developed, and demonstrated in use cases related to safety, efficiency, and uptime.
The FREEPORT project utilized the capabilities of the Stream Analyze platform for enabling real-time streaming, anomaly detection, versatile event-based data collection, (safety) alerts, as well as federated learning.
In the two demonstrators, we fulfilled the Technical Objective (TO) 1 of the FREEPORT project description: exhibit secure edge processing that provides real-time alerts using different network infrastructures (e.g., the customer's WiFi in a mine) at two sites, i.e., i) Boliden Kankberg mine – connected to Volvo's on-prem data center, and ii) AI Sweden.
For specification of relevant signals and models, and appropriate communication channels (e.g. MQTT integration with Boliden's operations system), we engaged with several personnel at Boliden and Volvo Trucks. This fulfills TO2.
The models deployed on the mine truck in the Boliden demonstrator included new signal combinations not previously gathered (cf TO3). The processing of signals in the telematics unit led to a substantial data reduction of data transmitted (cf technical KPI 1).
Taking the methods demonstrated in this project closer to a production setting involves work in areas such as scalable lifecycle management (including MLOps) of models over large fleets, as well as approval for use in road vehicles. In particular, the following areas are identified as the future work:

- Scaling within large, heterogenous fleets with different device and connectivity capabilities.
- Incorporate a MLOps approach to Federated learning scenarios, robust workflows for data retrieval, model version handling, automated testing, validation, and deployment processes.
- Address Asynchronous federated learning (AFL) to investigate the challenge of integrating diverse and intermittently connected edge devices. AFL should also be utilized to show that machine learning models can be refined with minimal delays and maximal accuracy, and to enable focus of computational resources on real-time data to improve efficiency and reducing redundant learning.

9 Participating parties and contact person

For more information, please contact:
Henrik Ydreskog, henrik.ydreskog@volvo.com
Sepideh Pashami, sepideh.pashami@hh.se

References

Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.

Fan, Y., Altarabichi, M. G., Pashami, S., Mashhadi, P. S., and Nowaczyk, S. (2024a). Invariant feature selection for battery state of health estimation in heterogeneous hybrid electric bus fleets. In *CEUR Workshop Proceedings*, volume 3765. CEUR-WS.

Fan, Y., Camacho, C., Pashami, S., and Nowaczyk, S. (2025). Causal graph-based anomaly detection for battery modules in electric heavy-duty vehicles. In *PHM Society Asia-Pacific Conference*.

Fan, Y., Nowaczyk, S., Wang, Z., and Pashami, S. (2024b). Evaluating multi-task curriculum learning for forecasting energy consumption in electric heavy-duty vehicles. In *Workshop on Embracing Human-Aware AI in Industry 5.0 (HAII5. 0 2024) co-located with the 27TH EUROPEAN CONFERENCE ON ARTIFICIAL INTELLIGENCE (ECAI 2024)*, volume 3765. CEUR-WS.

Fan, Y., Wang, Z., Pashami, S., Nowaczyk, S., and Ydreskog, H. (2023). Forecasting auxiliary energy consumption for electric heavy-duty vehicles. *arXiv preprint arXiv:2311.16003*.

Feldman, V. (2021). Does learning require memorization? a short tale about a long tail.

Harder, F., Adamczewski, K., and Park, M. (2020). Differentially private mean embeddings with random features (DP-MERF) for simple & practical synthetic data generation. *CoRR*, abs/2002.11603.

Jordon, J., Szpruch, L., Houssiau, F., Bottarelli, M., Cherubin, G., Maple, C., Cohen, S. N., and Weller, A. (2022). Synthetic data – what, why and how?

Martin, S. and Mohan, S. (2025). Predicting energy consumption for heavy-dutyvehicles: With an emphasis on auxiliary consumption: In collaboration with volvo trucks.

Ozen, C., Nowaczyk, S., Tiwari, P., and Pashami, S. (2025). Assessing the graph structure learning in graph deviation networks. In *International Symposium on Intelligent Data Analysis*, pages 97–109. Springer.

Pirasteh, P., Camacho, C., and Nowaczyk, S. (2025). Factors influencing propulsion energy consumption in battery electric heavy-duty trucks on urban and rural routes. In *The 38th International Electric Vehicle Symposium & Exhibition*. Electric Vehicle Symposium and Exhibition.

Rescorla, E., Tschofenig, H., and Modadugu, N. (2022). The Datagram Transport Layer Security (DTLS) Protocol Version 1.3. RFC 9147.

Selander, G., Mattsson, J. P., Palombini, F., and Seitz, L. (2019). Object Security for Constrained RESTful Environments (OSCORE). RFC 8613.

Shelby, Z., Hartke, K., and Bormann, C. (2014). The Constrained Application Protocol (CoAP). RFC 7252.

Taghiyarrenani, Z., Nowaczyk, S., and Pashami, S. (2023). Analysis of statistical data heterogeneity in federated fault identification. In *PHM Society Asia-Pacific Conference*, volume 4.

Wang, Z., Fan, Y., Ydreskog, H., and Nowaczyk, S. (2025). Investigation on machine learning models for forecasting auxiliary energy consumption of hd bevs. In *The 38th International Electric Vehicle Symposium & Exhibition*. Electric Vehicle Symposium and Exhibition.