LOBSTR
Learning On-Board Signals for Timely Reaction

Public report

Project within FFI – Machine Learning
Author Juan Carlos Andresen, Scania CV AB
Date 2020-05-27

Content

1. Summary

The project objective is to investigate the possibility to apply and implement real time anomaly detection on temporal multivariate signals on-board the vehicle with knowledge sharing for fault detection.

Within the truck industry, severe faults are rare and very expensive to reproduce on test labs. This makes it difficult to go with traditional classification methods for fault detection. Anomaly detection is identification of rare observations that differ significantly from the majority of the data set. In this case, normal functional operations of the vehicle would be the majority of the data set, which makes this a perfect approach for fault detection. Learning and knowledge sharing between vehicles are important as one vehicle might not expand the whole space of possible operations. A new encounter might look suspicious for one vehicle but it might be a daily routine for another vehicle, hence the collective experience sharing is an important aspect of anomaly detection.

In this project, several different anomaly detection methods on recorded time series data were investigated. These recorded data include normal driving of a vehicle and faulty drives with purposely injected fault. These methods are:

- Univariate Gaussian parametric model
- Histogram-based anomaly detection
- Prediction-based anomaly detection based on recurrent neural network (RNN) with long short-term memory (LSTM)
- Compression-based anomaly detection with an LSTM autoencoder
- Mixture-based anomaly detection

The two first methods were dropped for being too simplistic and gave too high false alarm (false positive) rate. The third method was too specific and required great domain knowledge of relation between input and output signals. The aim was to have a more holistic approach rather than too specific.

Both the compression-based anomaly detection with an LSTM autoencoder and mixture-based anomaly detection showed promising results. They both can handle arbitrary signals (with no requirement of domain knowledge) and have the possibility to share model parameters that can be further trained. Simulations show that federated training (training data split among the vehicles) yield similar results compared to one vehicle training the whole data.

We managed to implement the mixture-based anomaly detection on our edge device, which is a production-ready communicator unit, in a vehicle. It was able to listen to the vehicle’s real time CAN traffic and return anomaly score.

2. Sammanfattning på svenska

Scania satsar enormt för att förbättra kundens värde genom att öka fordonets tillgängliga drifttid och underhållskvalitén. I dagens läge sparas fordonens driftdata främst kumulativt som histogram i fordonens styrenheter. Datat läses ut ur fordonets styrenheter vid verkstadsbesök eller fjärrutläses över mobilnätet. Fjärrutläsningarna kan göras mer frekvent på vecko- eller månadsbasis. Dessa kumulativa histogrammen är användbart för att t.ex. förstå fordonets användning eller modellera långsamma dynamiska förlopp som slitage av bromsklossar men har sina begränsningar inom feldetektering. Detta är på grund av brist på den temporala sambanden i signalen samt en inbyggd fördröjning i systemet när data kan analyseras.

Den övergripande frågan som projektet LOBSTR försöker besvara är om det är möjligt att tillämpa anomalidetektionsmetoder på temporala multivariata signaler för feldetektering på fordonets styrenhet. Anomalidetektion är identifiering av sällsynta händelser som inte passar in med resterande data, vilket oftast sker genom ett avståndsmått mellan datapunkterna. I detta fall, ska metoden lära sig fordonets "normala" beteende och upptäcka när något är "konstigt" med fordonet, dvs. en anomali. Då fordonets styrenheter har begränsad prestanda och kapacitet, måste metoderna vara lättviktiga samt kunna dela med sig kunskaperna fordonet lärt sig med andra fordon, så kallad federerad inlärning (eng. federated learning). Innan projektet, hade vi samlat in massa tidsseriedata från lastbilen i normala körförhållanden samt en del körningar där vi medvetet hade något fel i fordonet (t.ex. läckage i lufttrycksystemet). All detta data kan användas för att träna metoderna för det normala beteendet och verifiera på de felinjicerade körningarna.

Projektets mål som sattes upp i ansökan var att bidra med ny kunskap i följande:

- anpassning av existerade anomalidetektionsmetoder så de fungerar väl i distribuerade system (IoT) med inlärning
- evaluera och jämföra prestanda mellan olika anomalidetektionsmetoder för IoT system, detta inkluderar precision, känslighet, fotavtryck av CPU och minne, datakommunikationskostnader och tolkningsbarheten av metoderna
- kategorisering av fordonets signaler i grupper för att upptäcka meningsfulla anomalier

Projektet var väldigt framgångsrikt. Dock behövde målen och projektets omfattning ändras lite på grund av personalbyte och rotation från både Scania och RISE inom projektet. Sista punkten behövde släppas för att prioritera att ha ett fungerade system då det uppstod många hinder när vi försökte implementera metoderna i styrenheten tillsammans med Scanias molninfrastruktur samt att utvecklingen av anomalidetektionsmetoderna tog betydligt längre tid än förväntat. Andra målpunkten blev delvis uppfylld då en av metoderna som använder neurala nätverk hade för komplexa paketberoenden för att kunna kompileras i styrenheten givet projekttiden samt de förändrade arbetssituationerna med korttidsvecka och korttidspermittering pga. COVID-19. Dock kan vi fortfarande jämföra resultaten mellan metoderna. Däremot utökade vi omfattningen i första målpunkten till att ha anomalidetektionsmetoder med federerad inlärningskapabilitet.

Projektet LOBSTR undersökte fem olika metoder, nämligen:

- Univariat Gaussisk parametrisk modell
- Histogrambaserad anomalidetektion
- Prediktionsbaserad anomalidetektion av RNN med LSTM
- Komprimeringsbaserad anomalidetektion med LSTM autoencoder
- Mixture-baserad anomalidetektion

Vi beslöt att inte gå vidare med de tre första metoderna för att de var antingen för enkla (vilket gav för många falsk positiva utslag) eller behövde vara för specifika (domänkunskap behövdes för att veta kopplingen mellan in- och utdata) vilket inte uppfyllde våra behov. Vi var ute efter ett mer holistiskt tillvägagångssätt, vilket den komprimeringsbaserade och mixture-baserade anomalidetektionsmetoderna var mer lovande.

Metoden för komprimeringsbaserad anomalidetektion av RNN (recurrent neural network) med LSTM (long short-term memory) går ut på att lära sig komprimera data baserad på normal data och därefter rekonstruera datat. Om rekonstruktionsfelet av indatat är för stort så anses det vara en anomali. Modellen kan fånga temporala beroenden över flera signaler och det kan tillämpas på vilka signaler som helst. Modellparametrarna kan även delas mellan fordonen för att dela inlärningskunskaperna. Simuleringsresultaten visar att federerad inlärning med flera fordon är möjligt och ger bättre resultat än om ett fordon hade lärt sig all träningsdata.

Metoden för mixture-baserad anomalidetektion utökar signalrymden och skapar Gaussiska komponenter över den gemensamma signalrymden. Om en datapunkt ligger långt ifrån en Gaussisk komponent anses datapunkten vara en anomali. Denna metod kan hantera strömmande data vilket är passande för distribuerad inlärning samt att modellparametrarna (vilket är de Gaussiska komponenterna) kan delas mellan fordonen. Simuleringsresultaten visar att federerad inlärning med flera fordon nästan lär sig lika bra som för ett fordon. Denna metod kom vi längst med under projektet, då vi även lyckades implementera in den i en styrenhet tillsammans med hela Scanias molninfrastruktur.

Resultaten som LOBSTR tagit fram går i linje med de övergripande målen för FFI programmet och FFI – maskininlärning. Att ha anomalidetektion på fordonet är en *smart funktion* för att kunna vidta åtgärder innan något kritisk händer på väg, vilket är ett steg närmare för *ökad trafiksäkerhet*. Dock har vi enbart utvecklat anomalidetektionsmetoden och inte hela processen för att kunna vidta åtgärder, men det låg utanför projektets omfattning. Möjligheten med federerad inlärning, dvs. att lära sig från varandra (andra fordon), kräver en *distribuerad arkitekturlösning* och resultaten vi tagit fram är lovande. Detta har varit ett stort framsteg för Scania och RISE för att *stärka konkurrenskraften inom svensk fordonsindustri*. Detta *samarbete har varit så givande* att Scania vill fortsätta med forskning inom detta område och kommer att starta en industriell doktorandtjänst i samarbete med RISE och KTH.

Det finns fortfarande mycket att utreda inom detta område. En enskild anomali i sig behöver inte vara kritisk för fordonet, dock kan många anomalier vara det.

Anomalidetektion blir värdefullt först när man kan skapa konkreta åtgärder. Möjliga forskningsområden kan vara att utreda hur man kan gruppera anomalier (t.ex. baserad på dess karaktäristik) kopplat till något konkret fel eller för att göra en kausalanalys för att hitta rotorsaken för att kunna skapa ett konkret åtgärdsplan. Andra forskningsområden är hur oregelbundna frånkopplingar (t.ex. p.g.a. dålig mottagning eller att fordonet slås av) påverkar anomalidetektionsmetoderna och hur lång tid fordonen kan vara frånkopplade utan att metodernas optimering divergerar från varandra. Allt detta är viktiga aspekter för att kunna förbättra fordonets drifttid samt öka trafiksäkerheten.

3. Background

Scania puts tremendous efforts toward increasing vehicle up-time and quality maintenance services for high customer satisfaction. This in the background requires a well-established fault management process. We separate the process of fault management into discovery, detection, prediction and intervention. Fault discovery deals with understanding the many possible ways in which a vehicle or part can fail, as well as understanding how to mitigate or act on each type of fault. Fault detection and prediction deal with assessing when a vehicle has failed or is about to fail, in a way previously discovered or undiscovered. Intervention is the planning of actions that restore the vehicle to a healthy state. Compared to its competitors, Scania has a well-developed multi-faceted approach to fault management.

With a frequency of weeks to months, Scania extracts and uploads to a central repository large amounts of aggregate operational data from vehicles, in the form of accumulated histograms capturing the lifetime distributions of signal readings. This accumulated histogram data is of great use for understanding vehicle use at population levels, but its application in fault detection is considered limited to specific use cases. Additionally, no features representing the temporal dynamics in the signals are recorded, the data does generally not capture multivariate patterns, and there is a built-in latency from the low sampling frequency.

This project aims to investigate how a subfield of machine learning, called anomaly detection, can be used for fault detection in large commercial fleets of heavy trucks and buses. Anomaly detection is increasingly being adopted as a general method to detect problems in a range of fields, including financial fraud, healthcare, process industry and manufacturing. On-board deployment of models is a requirement to meet low latency post-failure detection, and if possible, near-future prediction.

4. Purpose, research questions and method

The overarching question that the project aims to answer is if it is possible to apply anomaly detection on temporal multivariate signals on-board the vehicle for fault detection. The method needs to learn 'normal' behaviour of the vehicle and an anomaly should indicate that something might be wrong with the vehicle. To overcome the constraint of limited computational power and memory of an on-board unit, exchange of knowledge via a central server (as opposed to pure on-board learning) is necessary to learn patterns across the fleet.

One of the main goals of the project is to implement an anomaly detection model in the vehicle with communication to an off-board infrastructure. Research questions that needed to be answered to achieve the goal are:

- The choice of anomaly detection methods to handle real time (streaming) data with constraints on the on-board computational unit in consideration
- The design of protocol to exchange knowledge between vehicles (federated learning), with consideration to the chosen anomaly detection methods.

Anomaly detection is identification of rare observations (data points) that differ significantly from the majority of the data set. It uses a distance metric to calculate how far a sample is from the other samples. Therefore, it is important to have a good set of data with normal behaviour of the vehicle such that the model learns it. Once a fault happens on the vehicle, the behaviour should result in data points being far away from the normal behaviour. Preferably, the drift from the normal data points should indicate the process of being a fault on the vehicle.

Prior this project, we collected time series data of normal driving the vehicle that could be used for training the models as well as collected purposely injected faults for verification. Some examples of purposely injected fault that were conducted are air leakage in the air pressure system and a broken wheel speed sensor. These collected data are used to train and evaluate the anomaly detection models.

One approach that could satisfy both requirements of handling real time data with low footprint and sharing scalable model parameters is Gaussian mixture-based anomaly detection. It is based on two ideas; first that there is structure at different time scales in the signals, and second that there are several normal states. The first step with this method is to preprocess the data by expanding each signal. The second step is to create a Gaussian Mixture Model (GMM) over the joint space of all expanded signals. A new expanded data point located far away from all mixture components (clusters) is considered an anomaly. For federated learning, a variant of the expectation maximization (EM) algorithm is used, where each iteration is synchronized with sharing models via the central server.

Another approach is using neural networks (NN), where research in this field has been reblooming in recent years due to the advancement of technology. Long short-term memory (LSTM) layers have been commonly used to handle sequences of data which is suitable for time series data. Different NN architectures could be tested out, one suggestion is prediction, based on past data predict the next data point. If the prediction is too great from the real outcome, then it is anomalous. Another NN architecture that has been found successful for anomaly detection in other fields is autoencoder. The concept of autoencoder (depicted in **Error! Reference source not found.**) is to 'squash' the input signals then reconstruct the signals back again, only to learn the essentials of the input signals. If the output deviates too much from the input, it is considered an anomaly. The distance metric could be measured using mean squared error between the input and output. The idea of federated learning for NN is depicted in Figure 2. A global model is created in the central server and its parameters are sent out to each client. For each client, a client update is performed where the local model parameters are updated using local data only. Then the parameters are sent to the central server and merged into the global model. This cycle then repeats.

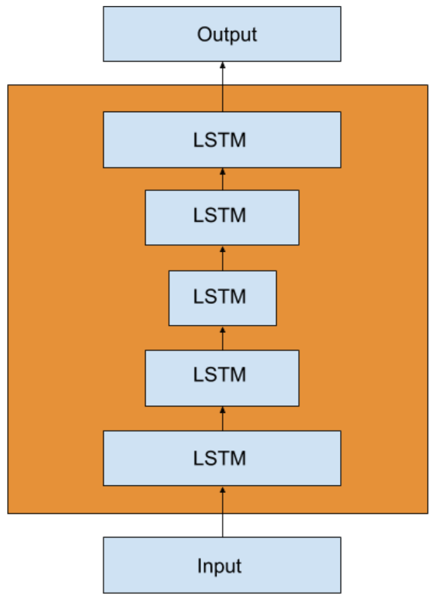

Figure 1: Concept of autoencoder, using LSTM layers squash and reconstruct the input.

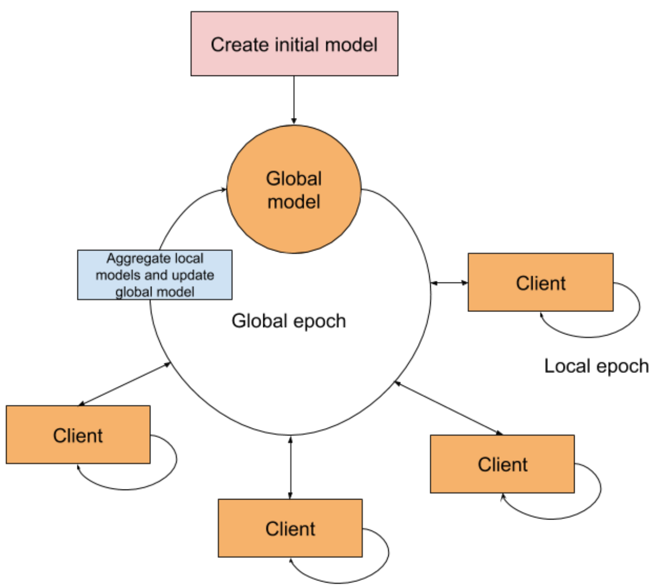

Figure 2: The idea of federated learning.

There are many signals communicated in a truck's CAN traffic and they are divided into different buses. It is unfeasible for the edge device to log and analyse all signals at once, due to limits on physical ports. Additionally, not all signals are relevant to detect anomalies based on the recorded data with purposely injected faults. Therefore, a handful of signals were selected that were large enough to be able to detect the relevant anomalies but small enough to allow for efficient experiments with the methods.

5. Objective

The main objective of this project, as stated in the research application, is to contribute the following new knowledge:

- adaptations of existing anomaly detection methods for them to work well in a distributed (IoT) online learning setting;
- performance comparisons of several anomaly detection methods in an IoT context, expected to include precision, recall, CPU footprint, memory footprint, data transfer cost and interpretability;
- categorization of on-board heavy truck signals into groups for which standard feature transformations can be defined, so that shared anomaly detections methods can be used to find meaningful anomalies.

The project was very successful. However, it did require to modify the scope slightly due to changes and rotations in personnel from both Scania and RISE sides. We had a framework ready to use in the project that enables scalability, however implementing this on the edge device (with the cloud infrastructure Scania has) proved to be much more difficult than anticipated. We had to prioritize to have a working pipeline to continue the project and therefore had to omit the last objective. Instead, we put more emphasis on the first objective and widened the scope to have anomaly detection models that can be federated. Additionally, the development of the anomaly detection methods took longer than anticipated and compiling the necessary packages required to run neural networks on the current production-ready communicator unit (that was used as an edge device) was too complex to achieve given the project time and the sudden changes in the work environment with short-time work due to COVID-19. Therefore we could not make comparisons on the footprints, however we are still able to compare the results between the anomaly detection methods.

6. Results and deliverables

6.1 Project results

One of the objective of this project is to have an anomaly detection method implemented in Scania's production-ready communicator unit that would act as an edge device. The anomaly detection method needs to be able to share its knowledge from the training data to its peers, without sending the local training data, so called federated learning. During the project, we had a couple of different anomaly detection methods we wanted to investigate. As the project progressed, some methods were more promising than others and therefore we only continued with two at the end. We will describe each approach that we studied and show the results of the remaining two methods.

6.1.1. Univariate Gaussian parametric models

The univariate Gaussian parametric model is one of the simplest models to apply for anomaly detection. This model was investigated solely for the purpose of setting a baseline. The method was, as expected, suitable for only a few signals. The univariate Gaussian approach assumes that the data is normally distributed. For each signal the mean and variance of a normal distribution is estimated using a sliding window. If the signal value deviates more than a fixed number of standard deviations from the mean then it is considered an anomaly. Otherwise the signal value is used for updating the estimated distribution. This simple parametric model worked well for some signals and could detect some of the injected faults. However, the false alarm (false positive) rate was very high when applying it to a larger set of signals. This is because few signals have known simple parametric forms of their distributions. Other limitations of this method are that it can only handle anomalies in one variable at a time (as opposed to relations between variables); it only considers the instantaneous value of a signal and not its time series; and it does not consider several normal operation modes (characterized by different means and variances but all being normally distributed). This method can therefore be considered too simple for real world practical use.

6.1.2. Histogram-based anomaly detection

The second approach to anomaly detection that we investigated was a histogram-based method. Here a histogram was used to represent the distribution of signal values recorded from the CAN bus. The method considers each signal separately but integrated over time. That is, a histogram is created for a fixed length time window. If the histogram is significantly different from a "normal" situation histogram of the signal, is considered as an anomaly. Like the simple parametric model, the histogram-based model considers each signal separately. However, it accumulates each signal over time, and since it uses a histogram to model the distribution it need not make an assumption about a specific statistical distribution. Thus, it overcomes two of the limitations of the simple parametric model. The histogram-based method also worked well for some injected faults. However, it also had a high false alarm rate. Part of the reason for this was that the implementation did not consider that there could be several normal operation modes. The distribution of each signal was represented by a single histogram.

6.1.3. Prediction-based anomaly detection using a RNN with LSTM

The third approach that we studied was prediction-based anomaly detection using a Recurrent Neural Network (RNN) with Long Short-Term Memory (LSTM). Here we trained a neural network to predict one signal at a time using a number of other signals as input. If the prediction deviated much from the actual signal when applying the method on test data, then it was considered an anomaly. The LSTM prediction model did not have the limitations of the previous models described above (the univariate Gaussian parametric model and the histogram-based model). The prediction model could take into account the relations between several signals, and their development over time, and also several normal modes or complex shape in state space of the normal region. However, which signal to model would differ dependent on the type of anomalies to detect. One alternative was to make one predictive model per signal. This would however have required a large machinery with a significant memory and processing footprint. As an alternative, a compression-based approach was suggested.

6.1.4. Compression-based anomaly detection with an LSTM autoencoder

The compression-based approach uses a neural network with an LSTM-autoencoder to find a compressed representation of the normal signal state space. If the reconstructed signals are far away from the input signals, this indicates an anomaly. With a compression-based approach it suffices with a single model to catch all regularities in the normal state space, with no need for further domain knowledge. The compression-based approach was one of the two most promising methods after the initial stage of the project and it was selected for further study and federated training.

Implementing this method on our current production-ready communicator unit proved more difficult than anticipated (deep cross dependencies with TensorFlow and Keras). However we still wanted results from this method and therefore took a different approach to this, we simulated similar behaviour in our local computer. The federated learning protocol is depicted in Figure 2 in earlier section. When simulating with more vehicles, we divided the training data evenly among the vehicles. The results for one vehicle case, i.e. with no federated learning, are depicted in Figure 3. Each value is a data point from the test data that consists of both normal drive and purposely injected fault concatenated into one plot. Although it might look like there were only three different occasions of faulty drive, there are in total 27 different occasions. The threshold value is obtained from the trainingdata based on the training loss. The autoencoder finds all anomalies, however there is also a high rate of false positives.

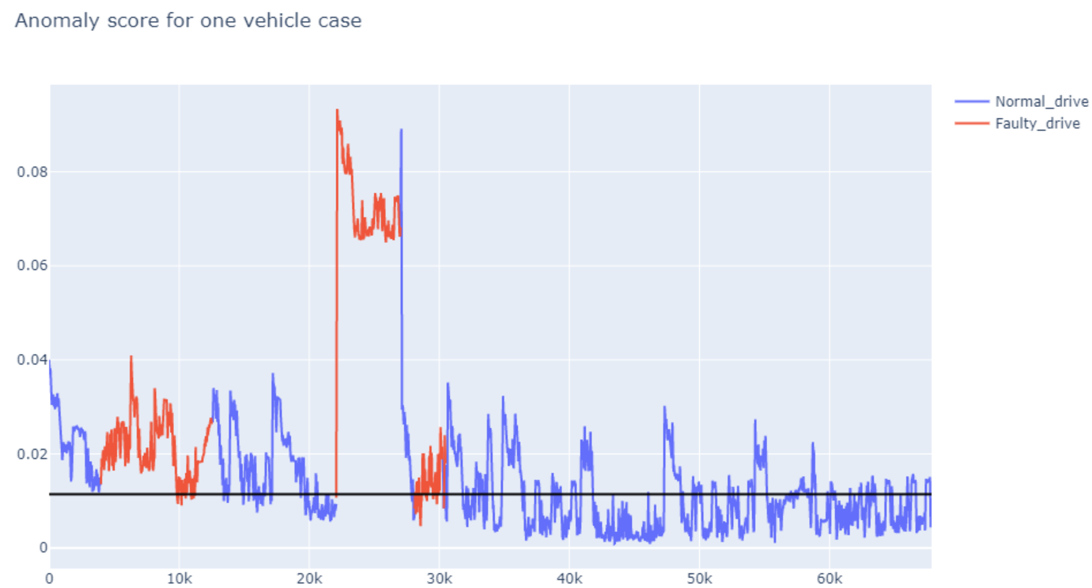

Figure 3: Anomaly score for one vehicle case. The threshold value for anomaly is 0.01144.

When we divide the training data to two or four vehicles, the results are quite similar to the one vehicle case. This can be seen in Figure 4 and Figure 5 below. There are some slight difference on the anomaly score due to the training has been split and the threshold varies a bit due to different training losses each vehicle generated. In all three figures, we can see some visible differences at 10k and 29k, some data points are below and some above their respectively threshold values.

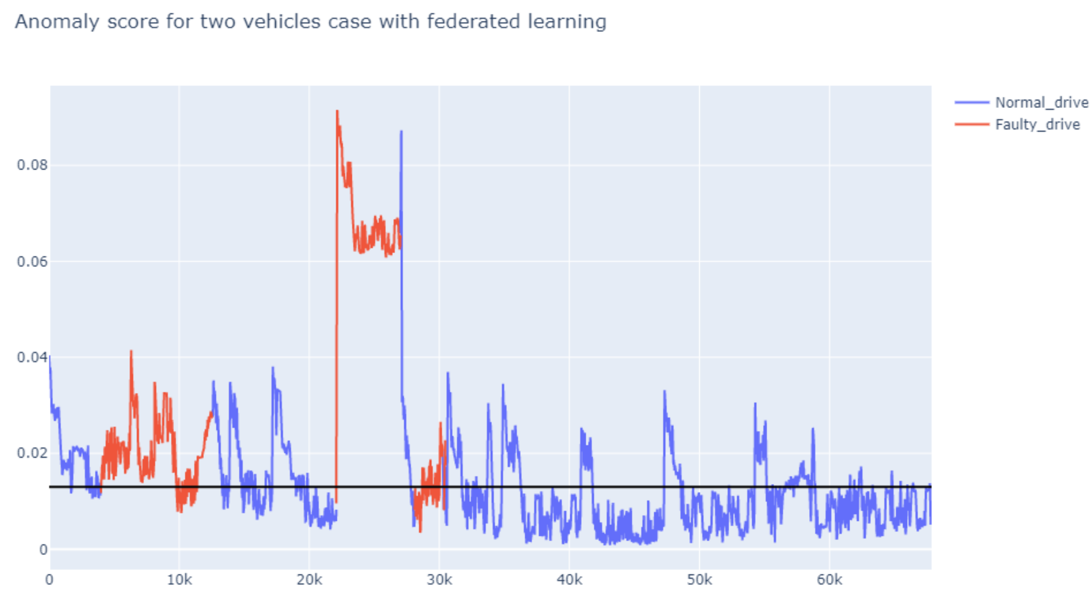

Figure 4: Anomaly score for two vehicles case using federated learning. The threshold value for anomaly is 0.013.

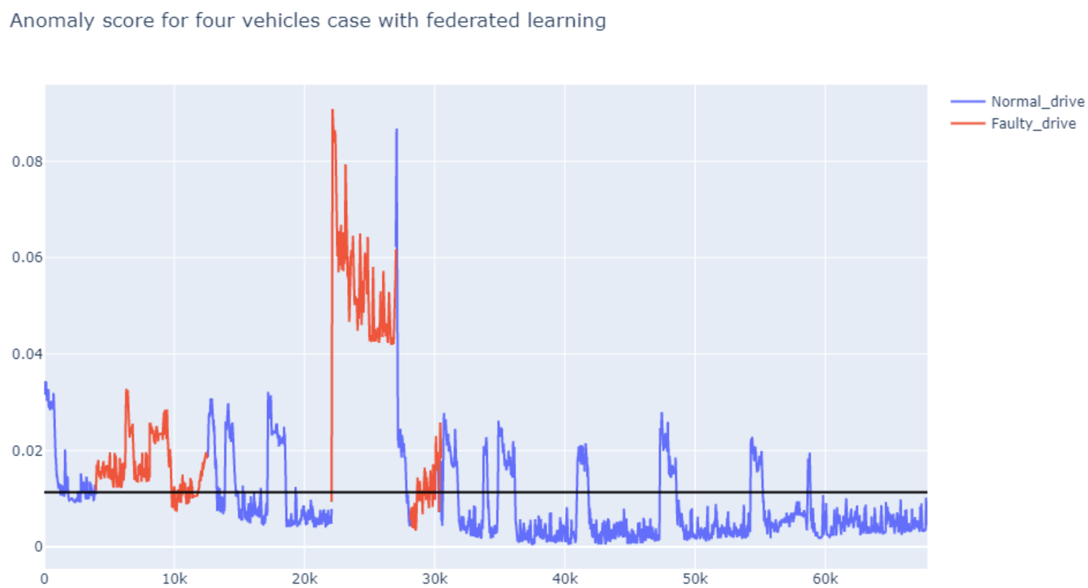

Figure 5: Anomaly score for four vehicles case using federated learning. The threshold value for anomaly is 0.01134.

Table 1 compares the different cases objectively using precision, recall, and F1 score. The four vehicle case performs better due to less false positive and false negative anomalies.

Table 1: Precision, recall, and F1 score for LSTM autoencoder.

	One vehicle	**Two vehicles with federated learning**	**Four vehicles with federated learning**
Precision	0.18	0.23	**0.35**
Recall	0.74	0.74	**0.81**
F1	0.29	0.35	**0.49**

6.1.5. Mixture-based anomaly detection

The second approach that we selected for federated training was a Gaussian mixture-based method. With this approach, the original input signals were expanded by applying gaussian smoothing over different time intervals and calculating the derivatives. A Gaussian Mixture Model (GMM) was then created over the joint space of all expanded signals. A new expanded data point located far away from all mixture components was considered an anomaly.

This approach is the one we got the furthest with as we managed to successfully implement this on our edge device. However, due to the drastic changes in our work environment with short-time work week and recommendation to work from home due to COVID-19 we were not able to extensively test it out in practise. As it is only for a vehicle, we continued using the framework to simulate virtual trucks on our local computer to investigate the federated learning part. The suggested federated learning protocol for this approach is depicted in Figure 6. With this approach, we are able to train and detect on the same data in parallel. If the model detects an anomaly while training, the training model will not learn this into the model. This is to avoid training in anomalies into the 'normal' behaviour. Similar to the LSTM autoencoder case, a central server initialises a model that is sent out to available vehicles (this is known as a vehicle registers its availability when turned on). The vehicles detect and train in parallel and after a given time the central server requests model from all vehicles. The central server merges the model and sends out the new model, and the loop repeats. Note that the models here are Gaussian components.

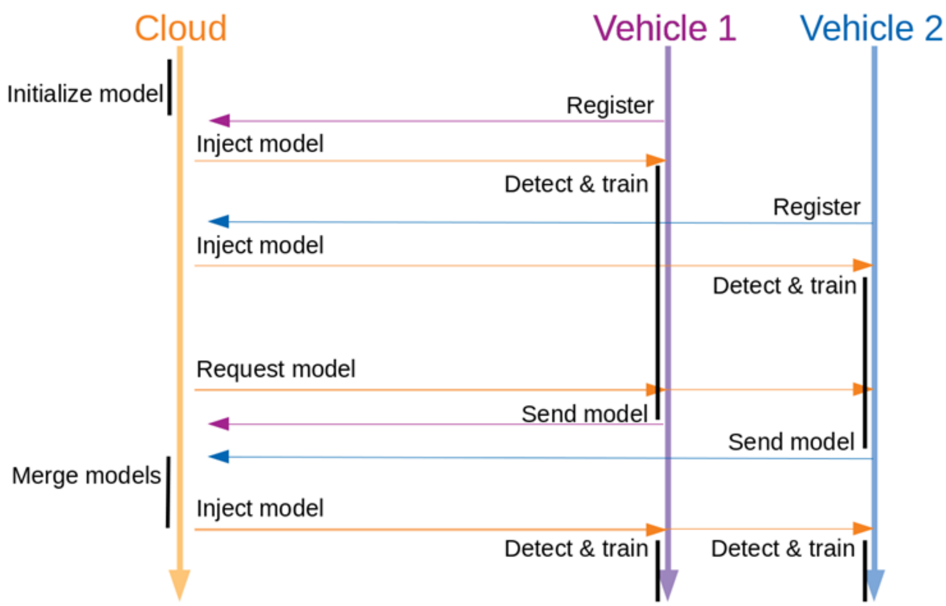

Figure 6: The proposed federated learning protocol for a fleet of vehicles.

Similar to the LSTM autoencoder case, we evaluated with different amount of vehicles; one vehicle with no federated learning and two and four vehicles using federated learning. The results are depicted in Figure 7 - Figure 9. The anomaly score is based on Mahalanobis distance from the Gaussian component. The threshold value was chosen to 100 as it looked reasonable based on training data. Note that some values exceed 1000 in anomaly score, and we just capped it there for visibility of the whole plot.

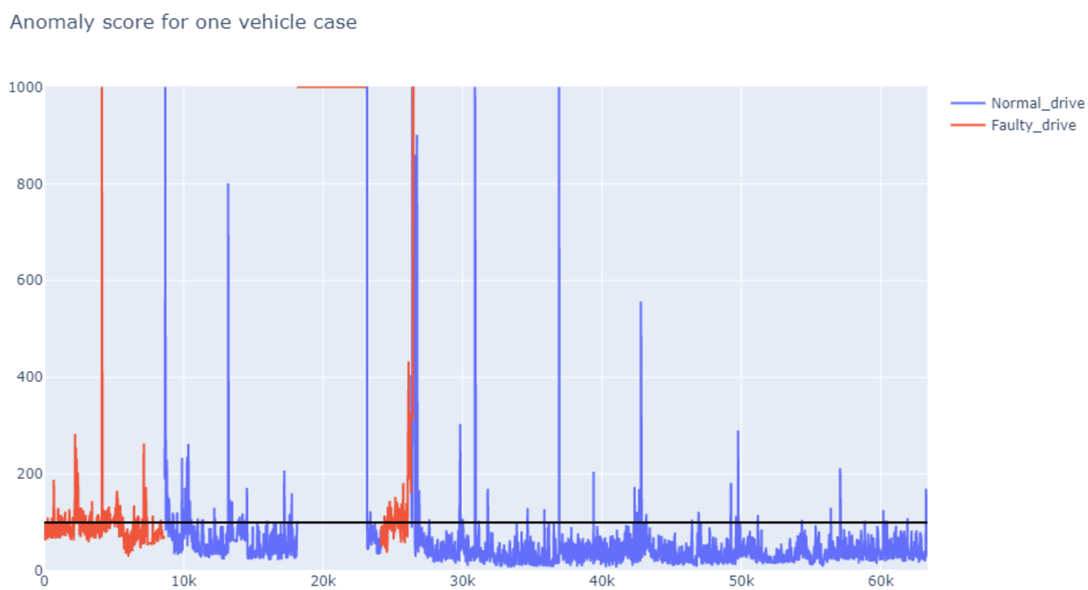

Figure 7: Anomaly score for one vehicle case with GMM. Score capped at 1000 for visibility.

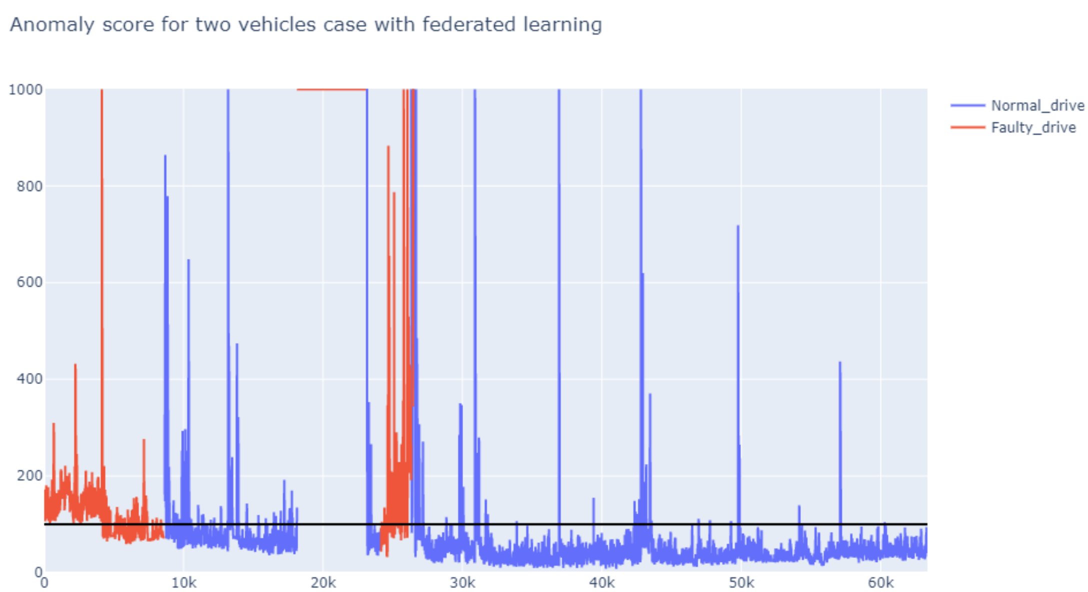

Figure 8: Anomaly score for two vehicles case with GMM using federated learning. Score capped at 1000 for visibility.

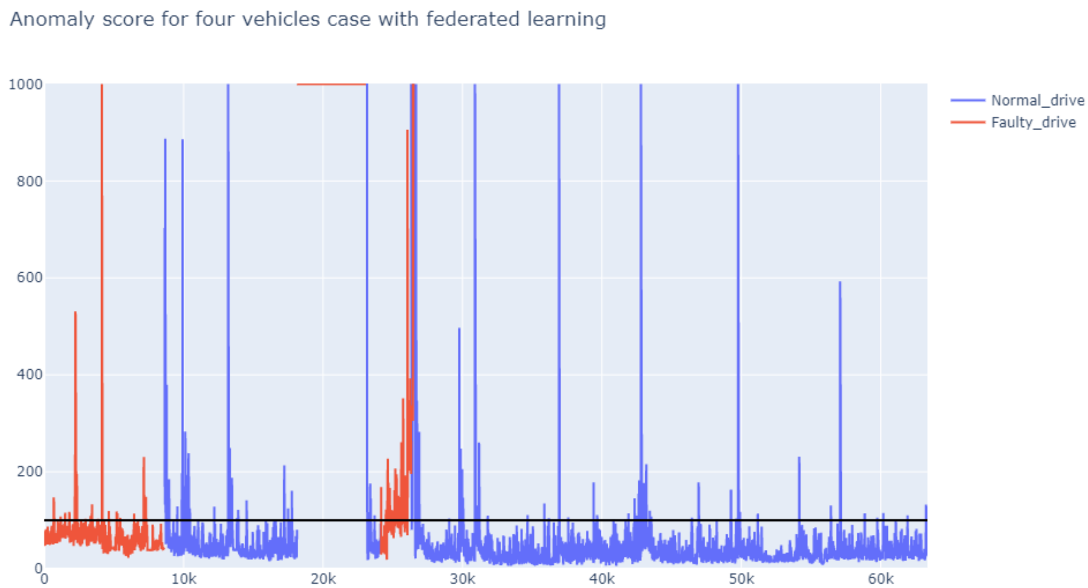

Figure 9: Anomaly score for four vehicles case with GMM using federated learning. Score capped at 1000 for visibility.

Table 2 compares is objectively using precision, recall, and F1 score. The best case is with one vehicle case with no federated learning done. If you compare between methods (LSTM autoencoder against GMM), then their respectively best cases are almost identical.

Table 2: Precision, recall, and F1 score for GMM.

	One vehicle	Two vehicles with federated learning	Four vehicles with federated learning
Precision	**0.36**	0.29	0.32
Recall	**0.81**	0.74	0.48
F1	**0.5**	0.42	0.38

6.1.6. Discussion

We tried several different methods, where we continued to elaborate more with two of the approaches, namely; compression-based anomaly detection with an LSTM autoencoder and mixture-based anomaly detection. We went the furthest with the latter approach, where we managed to implement it on our edge device on one vehicle, which was one of our main objectives. We continued with both approaches and evaluated their performance with federated learning through simulations of many vehicles. The results indicate great potential with federation, as we can in practice learn from many vehicles with small datasets instead of sending all data into one central server and learn one model based on the large dataset. More tweaking and hypermeter tuning may be required.

To have an on-board system for anomaly detection is an *intelligent function* to take action before a severe component fails on the road, which is a step closer to *increasing road safety*.

However, we have only developed the anomaly detection methods and have not developed the full process for taking action as it was outside the scope of the project. The ability to federate, that is learning and knowledge transfer between vehicles, requires a *distributed architecture* and results from it has been very promising. This achievement has been a great step for Scania and RISE to keep being forefront in this field as well as *strengthening the Swedish vehicle industry*. Furthermore, *the collaboration between Scania and RISE in this project has been very fruitful* that Scania wanted to continue with the research in this field and is going to start an industrial Ph.D. position in collaboration with RISE and KTH.

6.2 Deliverables

In the research application, we stated the following deliverables:
D1: Report on evaluation of current methods for anomaly detection on pre-recorded data
D2: Report from workbench tests, simulating a vehicle
D3: Report on developed method for anomaly detection and publication ready to submit
D4: Report from vehicle demonstration
D5: Knowledge transfer documentation finalised

We managed to achieve majority of the deliverables. The deliverables D1, D2, and D5 were completed. D3 was partially achieved, however the development of the anomaly detection methods took longer than anticipated. We still plan to make a publication of our results. D4 was not achieved, mainly due COVID-19 and the economic situation leading to a sudden change in our work environment with short-time work weeks and recommendation to work from home. Since the demonstration was planned to be a hands-on on the truck demo, it was not possible to make it as an online demo. We still plan to present it internally at Scania once the situation improves.

7. Dissemination and publications

7.1 Dissemination

How are the project results planned to be used and disseminated?	Mark with X	Comment
Increase knowledge in the field	X	The project results has increased the knowledge within Scania about statistical methods and neural networks for time series anomaly detection on edge devices. A federated learning protocol has been implemented with minimum impact on the edge device in mind.
Be passed on to other advanced technological development projects	X	The project results has invoked great interest from different department within Scania CV AB. This interest will likely be formulated into future development projects. Some findings have already been passed on as requirements in other projects.

Be passed on to product development projects	x	Lessons on how to run models on edge devices applied in design of communication and edge devices in pre-development phases at Scania.
Introduced on the market		
Used in investigations / regulatory / licensing / political decisions		

7.2 Publications

The results achieved by the project are very promising, nevertheless, additional work and research is needed for an article submission. As mentioned above, Scania is interested in continuing researching in this direction. These results will be included in a future and more comprehensive publication.

8. Conclusions and future research

The overarching objective of this project was to answer if it is possible to apply anomaly detection on temporal multivariate signals on-board the vehicle for fault detection with knowledge transfer between vehicles. We did not want to aim for a specific component for but rather have a holistic approach to have a generalised anomaly detection method. We were set on using Scania's production-ready communicator unit as our edge device because this unit connects the vehicle to Scania's back-end infrastructure.

Several anomaly detection methods were proposed during the project. We omitted the univariate Gaussian parametric model, the histogram-based method, and the prediction-based anomaly detection using RNN with LSTM. These suggestions were either too simple (with too high false alarm rate) or required to be too specific (domain knowledge required to design outputs with inputs) and did not satisfy our requirements.

One promising approach was the compression-based anomaly detection with LSTM autoencoder. The model learns how to compress data based on normal drive and then reconstruct it. If the reconstruction error of the input data is too large, then it is an anomaly. This model can capture temporal dependencies over multiple signals and can be applied to any arbitrary signals. The model parameters can be shared, from the weights of the neural network, inducing knowledge sharing between vehicles. Results have shown that federated learning, where each vehicle synchronises its model with a central server, is possible and with better results than if one vehicle did the whole training. This is with the assumption that the data needs to be trained batch-wise due to limitations of the edge device.

Another promising approach was the mixture-based anomaly detection. The model takes the input signals, expands the space of each signal, and creates a GMM over the joint space of all expanded signals. If a data point is too far away from a Gaussian component, then it is considered an anomaly. This model can handle streaming data which is suitable for online learning, and the model parameters (which are the Gaussian components) can be shared among vehicles which makes it suitable for federated learning. Results have shown that federated learning performs almost as good as for the one vehicle case.

Both the compression-based anomaly detection with LSTM autoencoder and mixture-based anomaly detection are worth to continue pursuing. Both methods have been developed to be light with edge device in mind and have the possibility to be shared among different vehicles without the need to send local data, which is one of our main objectives. This project has only shown good initial results of these models, however, further investigations are required. Tweaking and hypermeter tuning are required to improve the results. The aim is a generalized anomaly detection, and therefore we also need to test how both models fair with unseen faults.

An anomaly itself may not be severe for the vehicle, however many anomalies might be. Anomaly detection is only a step towards fault detection and it is only valuable with concrete action to act upon. Therefore, possible future research worth investing in could be to investigate how to group anomalies (e.g. based on their characteristics) to a more concrete fault or causality analysis of anomalies to find affected system, in order to make concrete action plans. Other possible research could be to investigate how these models are affected by intermittent disconnected vehicles and how long a vehicle can be disconnected without the models diverging from each other. These are all important aspects to pursue to improve uptime of our vehicles, and increased road safety.

9. Participating parties and contact persons

Scania CV AB

Juan Carlos Andresen

ECCI
Scania CV AB
151 87 Södertälje
Sweden

Tel: +46 8 553 835 16
Email: juan-carlos.andresen@scania.com

RISE AB

Anders Holst

RISE AB
Box 1263
164 29 Kista
Sweden

Tel: +46 10 228 43 13
Email: anders.holst@ri.se