AIM-TRUE: AI-driven automotive service market–towards more resource-efficient and sustainable vehicle maintenance

Nuwan Gunasekara[1,*], Abdallah Alabdallah[1], Mohammed Ghaith Altarabichi[1] and Slawomir Nowaczyk[1]

[1] *Center for Applied Intelligent System Research, Halmstad University*

Abstract

This paper introduces AIM-TRUE, a project leveraging meta-learning methods to enhance the sustainability and resource efficiency of vehicle maintenance in the automotive aftermarket. By addressing challenges such as intermittent demand and complex global supply chains, the project aims to develop innovative predictive models to optimise inventory control, reduce unnecessary transport, and minimise part scrapping. Initial results show that combining machine learning with uncertainty-aware strategies significantly improves forecasting accuracy and operational efficiency, highlighting the potential for more sustainable and cost-effective vehicle maintenance logistics.

1. Introduction

Europe's automotive industry is increasingly facing the need for concrete solutions to the challenges related to resource-efficient and sustainable transport systems. With more data available, it has become apparent that Artificial Intelligence (AI) and Machine Learning (ML) methods can also help to reduce climate emissions and energy consumption through more efficient use of resources in vehicle aftermarket operations. Volvo Group aims to deliver complete transport solutions to its customers, from personalised vehicles suited for any task at hand – be it hauling goods over thousands of kilometres or distributing them within a few city blocks – to services that keep the vehicles running efficiently throughout their lifetime. Doing it successfully and with sustainable resource utilisation requires new ML-based, flexible, and green services that reduce costs while increasing customer satisfaction and maintaining a competitive advantage.

All these goals can only be achieved by anticipating where and when a part will be needed and delivering that part to the correct region before this need even arises, thus reducing costs and increasing service levels. The AIM-TRUE (AI-driven Automotive Service Market: Towards more Resource-Efficient and Sustainable Vehicle Maintenance) project focuses on using state-of-the-art methods based on meta-learning to improve the services provided by the Service Market. In particular, more predictability enables the use of environmentally friendly transport

SAIS2025: Swedish AI Society Workshop 2025, 16-17 June 2025, Halmstad, Sweden.
*Corresponding author.
nuwan.gunasekara@hh.se (N. Gunasekara); abdallah.alabdallah@hh.se (A. Alabdallah); mohammed.ghaith.altarabichi@hh.se (M. G. Altarabichi); slawomir.nowaczyk@hh.se (S. Nowaczyk)

channels and reduces the scrapping of parts due to obsolescence.

AIM-TRUE leverages ML to better understand the factors affecting parts availability and enable individualised inventory control policies. The project's primary goal is to improve heavy-duty aftermarket resource efficiency and sustainability by reducing three aspects: urgent transport orders, back-and-forth haulage, and part scrapping. The new generation of predictive logistics provides opportunities for better system understanding, large-scale optimisation, quality monitoring, and new data-driven innovative services, all of which are prerequisites for the efficient use of resources – while providing the right parts at the right place and time.

The project started in January 2024, with three partners: Volvo Service Market Logistics, Rejmes Transportfordon, and the Center for Applied Intelligent Systems Research at Halmstad University.

We aim to create a demand forecasting model in combination with decision support tools that ensure vehicle components are not unnecessarily produced or transported. However, doing so involves issues like intermittent or sporadic demand, delayed information about the needs, imminent drops in availability, and limited supply chain capacity. Solving these challenges in a complex worldwide system that is typical for automotive OEM's dealer network requires a novel approach to reach a balance between incoming and outgoing part flows. Therefore, within the AIM-TRUE project, we explore, for the first time, state-of-the-art forecasting models based on meta-learning, allowing us to integrate new global data sources (such as predictive maintenance information) while efficiently adjusting the predictions according to specific, local conditions.

In AIM-TRUE, we will specifically address two key challenges related to meta-learning research. The first one is meta-generalisation, which refers to the problem of generalising from meta-train to meta-test tasks for a large number of tasks, as well as learning supplementary tasks and noise modelling. The majority of today's solutions assume that the task distribution is unimodal, while logistics data from Volvo SML has a highly heterogeneous task structure with relatively complex relations across a number of somewhat well-separated task families. The second challenge concerns tracking the time relationship between meta-tasks that, in reality, do not switch at random but rather follow technical and business principles. In this case, continual learning provides some techniques worth investigating with more explainable semantic meaning.

Within the AIM-TRUE project, we work primarily with model-based meta-learning methods due to their flexibility in modelling the internal task dynamics and their broader applicability compared to most optimisation- and metric-based techniques. It is crucial for us to efficiently handle different aspects of variability, for example, across spare parts, dealers, vehicle models, geographical regions, and more. An important open question, scientifically, is whether meta-learning techniques actually learn how to perform "rapid learning" or primarily discover robust high-level features to be reused across tasks. There is some evidence for the latter, which indicates that there is room for improvement by integrating novel techniques that rely more on the former.

2. Intermittent Demand Forecasting

Intermittent demand forecasting is especially important in spare parts and logistics in the automotive industry, where demand occurs irregularly with frequent zero-demand periods. This makes statistical forecasting techniques unreliable, as they often fail to capture the sporadic nature of such demand. Effective forecasting helps reduce inventory costs while maintaining high service levels in supply chains.

We have evaluated the performance of a conventional Machine Learning algorithm (XGBoost), Neural network-based models (NN), and Kolmogorov-Arnold Networks (KAN), to predict the next month's demand based on descriptive features about the parts, the dealers, and the sales history. We also proposed a novel time-dependent NN-based model, $Forecast(t)$, to predict the demand at specific times in the future, allowing for a longer planning window. In this regard, the training data is transformed to include the next month's sales and the first sales after a period t of zero sales. We also utilised a meta-learning model based on Model-Agnostic Meta-Learning for Multimodal Task Distributions (MAML) [1], to train a model to be adaptable to future demand periods where the task is defined as a Dealer-Part-Period combination.

The evaluation of intermittent demand forecasting models can also be challenging due to the frequent zero-demand periods. An evaluation metric such as Mean Absolute Error (MAE) will be affected by the frequency of zeros, favouring models that often predict zeros. In this work, we decomposed MAE into two components by observing the MAE value for the periods for which there is a non-zero demand; we annotate this as MAE_1, while for periods with no demand, we calculate the MAE as MAE_0. The time-dependent model $Forecast(t)$, is evaluated using both MAE_0 and MAE_1 for the next month's demand, along with MAE^t for the demand at time t.

The decomposition of MAE allowed us to optimise models to excel at forecasting a certain type of sales. We carried out evolutionary feature and sample selection using XGBoost in a wrapper setting [2] while optimising for different metrics MAE_0 and MAE_1 (in Table 1 we refer to the model optimised based on MAE_0 as XGBoost w/ FS MAE_0). The feature selection optimisation is guided in these two scenarios to focus on features that help predict periods with no demand (when optimising MAE_0), or to keep good predictors of demand periods (MAE_1).

Table 1 lists the results of the compared models. We may observe that the NN model achieved the best overall MAE, however, it's clear that the good MAE performance is largely due to excelling

Table 1
Models' results, the bold indicates the best result for the metric.

Model	MAE	MAE_0	MAE_1	MAE^t
XGBoost	1.82	0.51	7.43	N\A
XGBoost w/ FS MAE	1.82	0.50	7.49	N\A
XGBoost w/ FS MAE_0	1.82	0.49	7.47	N\A
XGBoost w/ FS MAE_1	1.87	0.57	**7.39**	N\A
NN	**1.66**	0.19	7.93	N\A
KAN	1.94	0.54	7.87	N\A
DropKAN [3]	1.92	0.54	7.77	N\A
MAML	1.95	**0.10**	9.86	N\A
Forcast(t)	1.80	0.48	7.44	6.33

at predicting zero/demand periods as evident with the low $\texttt{MAE}_0$ result. XGBoost, in general, offered a better $\texttt{MAE}_1$ performance when compared to the neural-network-based models. In particular, when feature selection was carried out for $\texttt{MAE}_1$. KAN-based models KAN and DropKAN performed generally worse than XGBoost models, but their $\texttt{MAE}_1$ performance was better than the neural-network-based models. It must be stated that, unlike the black-box neural-network-based models, KAN-based models offer explainability. The meta-learning MAML demonstrated the best $\texttt{MAE}_0$ performance. However, this happens at the expense of the worst $\texttt{MAE}_1$ performance across tested methods. The Forecast(t) model showed a balanced performance between $\texttt{MAE}_0$ and $\texttt{MAE}_1$, achieving the second-best performance on the total $\texttt{MAE}$ with the added value of predicting a longer future demand window.

3. Prediction Intervals for ML-driven Automotive Service Market Logistics

This work [4] explores how machine learning (ML) and prediction intervals (PIs) can enhance demand forecasting and inventory control in Volvo Group's Service Market Logistics (SML) division, specifically for spare parts distribution in Sweden. Addressing the challenge of balancing high service levels with cost efficiency, the study uses real operational data and discrete-event simulation (DES) to assess forecasting strategies under different types of forecast errors—random, systematic (bias), and intermittent [5, 6].

Traditional forecasting models like SES [7] and Croston's method [8] were compared to advanced ML models such as XGBoost [9]. To capture uncertainty, the study employed nonparametric bootstrapping to construct PIs [10], integrating them into inventory decision-making via methods such as the Triangular and Beta-PERT distribution-based estimates [11, 12]. These strategies showed that accounting for uncertainty helps reduce costs and improve service levels—particularly for low-demand, high-variability parts—while overly conservative strategies like using the PI upper bound increased inventory costs significantly.

The findings support the view that integrating ML-based forecasting with uncertainty-aware inventory strategies tailored to product characteristics leads to more sustainable, resilient, and cost-effective supply chains.

4. Pragmatic Paradigm for Multi-stream Regression

This work by Gunasekara et al. [13] addresses the challenges of time series nowcasting in dynamic multi-stream data environments, especially under concept drift. Traditional deep learning models struggle with adaptability and catastrophic forgetting when exposed to evolving data distributions. Stream learning offers adaptability but lacks batch model accuracy. The authors propose a hybrid paradigm combining a batch-trained neural network as a feature extractor and a streaming regressor (SOKNL) [14] to predict either directly or the residuals of the base model.

This hybrid method utilizes historical data to learn embeddings for stream identifiers, such as store and product IDs, which are input into a Multi-Layer Perceptron (MLP) and regressor for initial predictions. The streaming regressor is incrementally updated using different representations (raw features, embeddings, or final layer outputs), enabling adaptability during deployment. Experiments on the NZ Energy Pricing and Kaggle Demand Forecasting datasets show that residual learning approaches consistently improve performance over static models, especially when the batch model performs reasonably well [13].

Overall, the proposed approach achieves a balance between learning efficiency and prediction accuracy in evolving data streams. Future work includes comparison with online time series models like OneNet and deeper stream-level evaluations.

References

[1] R. Vuorio, S.-H. Sun, H. Hu, J. J. Lim, Model-agnostic meta-learning for multimodal task distributions, 2019. URL: https://openreview.net/forum?id=HkzOWnActX.

[2] M. G. Altarabichi, S. Nowaczyk, S. Pashami, P. S. Mashhadi, Surrogate-assisted genetic algorithm for wrapper feature selection, in: 2021 IEEE congress on evolutionary computation (CEC), IEEE, 2021, pp. 776–785.

[3] M. G. Altarabichi, Dropkan: Regularizing kans by masking post-activations, arXiv preprint arXiv:2407.13044 (2024).

[4] J. Kamil, M. Amer, Prediction Intervals for ML-driven Automotive Service Market Logistics, Master's thesis, Halmstad University, Halmstad, Sweden, 2025. Master's Programme in Artificial Intelligence.

[5] J. Banks, J. S. Carson, B. L. Nelson, D. M. Nicol, Discrete-event system simulation, 4th ed., Pearson Prentice Hall, 2005.

[6] D. Hellström, M. Johnsson, Simulation-based inventory optimization, International Journal of Physical Distribution & Logistics Management 50 (2020) 861–879.

[7] R. G. Brown, Statistical forecasting for inventory control, McGraw-Hill, New York, 1959.

[8] J. D. Croston, Forecasting and stock control for intermittent demands, Operational Research Quarterly 23 (1972) 289–303.

[9] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[10] P. Hall, A. Rieck, Bootstrap methods for time series, in: Bootstrap Methods for Time Series, Springer, 2001, pp. 1–10.

[11] C. Clark, The pert model for the distribution of an activity time, Operations Research 10 (1962) 405–406.

[12] D. Vose, Risk analysis: A quantitative guide, 3rd ed., Wiley, 2008.

[13] N. Gunasekara, S. Nowaczyk, S. Pashami, Pragmatic paradigm for multi-stream regression, in: International Symposium on Intelligent Data Analysis, Springer, 2025, pp. 358–372.

[14] Y. Sun, B. Pfahringer, H. M. Gomes, A. Bifet, Soknl: A novel way of integrating k-nearest neighbours with adaptive random forest regression for data streams, Data Min. Knowl. Discov. 36 (2022) 2006–2032.