

AIM-TRUE

AI-driven Automotive Service Market: Towards more Resource-Efficient and Sustainable Vehicle Maintenance

Contributors:
Abdallah Alabdallah, Mohammed Ghaith Altarabichi, Iulian Carpatorea, Nuwan Gunasekara, Anders Malmberg, Philip Mårtensson, Slawomir Nowaczyk, Sepideh Pashami

Project Partners:
Halmstad University, Volvo Group, Rejmes Transportfordon

FFI Fordonsstrategisk Forskning och Innovation
VINNOVA Energimyndigheten TRAFIKVERKET FKG VOLVO SCANIA VOLVO

Table of Contents

1 Summary: 2
2 Background 2
3 Objectives....3
4 Project Deliverables and Key Performance Indicators 4
5 Research Outputs 6
6 Utilisation, Implementation, and Industrial Uptake....10
7 Societal, Environmental, and Strategic Impact 11
8 Deviations and Lessons Learned....11
9 Gender Equality and Project Execution....12
10 Publication list....12
11 Participating parties and contact persons....13

1 Summary:

The AIM-TRUE project (AI-driven Automotive Service Market: Towards more Resource-Efficient and Sustainable Vehicle Maintenance) aimed to address resource inefficiencies in the heavy-duty vehicle aftermarket caused by intermittent demand for spare parts and suboptimal logistics practices. The project's objective was to improve sustainability and operational efficiency by applying advanced Artificial Intelligence (AI) and Machine Learning (ML) methods to demand forecasting, inventory control, and decision evaluation.

The project developed and evaluated AI-based forecasting models that adapt to local dealer conditions and highly irregular demand patterns. These models achieved up to a 20% improvement in forecast accuracy compared to existing statistical baselines. To bridge the gap between forecasting performance and operational outcomes, the project introduced decision-centric evaluation frameworks and uncertainty-aware inventory control policies, demonstrating potential reductions in safety stock levels, rush orders, unnecessary transport, and part obsolescence.

Results were validated using both real-world industrial data to show practical impact and simulated data to improve reproducibility and build a stable foundation for future developments in this area. In addition to technical advances at all the partners, the project produced peer-reviewed publications, open-source software tools, and dissemination activities to strengthen the Swedish logistics and automotive ecosystem. Overall, AIM-TRUE successfully demonstrated how AI-driven decision support can significantly enhance resource efficiency and sustainability in complex aftermarket supply networks.

2 Background

The automotive aftermarket is a vital component of the vehicle life cycle, ensuring maintenance and repair services that maintain customer trust and operational uptime. However, managing this sector is hindered by the extreme technical complexity of spare parts demand, which is typically intermittent, irregular, and sporadic. Analysis shows that approximately 77% of time periods record no sales, making traditional "one-size-fits-all" forecasting models highly unreliable. This unpredictability results in massive inefficiencies, driving an estimated $1.1 trillion in global supply chain waste through overstocking and the necessity of rush orders, which can be up to 50 times more resource-intensive and environmentally damaging than standard shipping modes. The problem is further intensified by the transition to e-mobility; expensive components like electric vehicle traction batteries have short shelf lives and limited supply, requiring a far more precise logistics approach than current industry practices provide.

In practice, the core challenge is not so much the forecasting in isolation but making reliable replenishment and positioning decisions in multi-echelon networks under strict service-level requirements, long lead times, and substantial dealer-to-dealer heterogeneity. Frequent assortment changes, local campaigns, and concept drift further reduce the value of static models, while the long tail of low-volume parts makes uncertainty the dominant driver of cost. This creates a need for decision-centric, uncertainty-aware analytics that can translate intermittent demand signals into robust inventory policies and can be evaluated against operational KPIs (e.g., safety stock, rush orders, obsolescence, and transport intensity), rather than accuracy metrics alone. Without such methods, organisations are forced to “buy” availability through excess inventory and carbon-intensive emergency logistics, directly undermining both profitability and sustainability targets.

The AIM-TRUE project was initiated to address some of these sustainability and efficiency challenges within the heavy-duty vehicle aftermarket. The project sought to transition from traditional "siloed" supply chains to a collaborative supply network, and we demonstrated that leveraging advanced AI and ML provides significant benefits in that area. To effectively navigate the extreme technical complexity of intermittent demand forecasting, one needs tools that can utilise huge amounts of available data and capture complex real-world patterns.

3 Objectives

The AIM-TRUE project aimed to enhance resource efficiency and sustainability in the heavy-duty vehicle aftermarket by leveraging advanced Artificial Intelligence (AI) and Machine Learning (ML) methods. Its objectives were to develop accurate spare parts demand-forecasting models, implement individualised inventory control policies, and provide decision-support tools that translate forecasts into operational KPIs.

Key deliverables included:

- **WP1**: Knowledge dissemination, including scientific publications, open-source software, promotional materials, and a final workshop to transfer knowledge to industry and academia.
- **WP2**: Machine Learning based forecasting methods developed using historical data, evaluated on essential spare parts.
- **WP3**: New inventory control policies and virtual stock strategies to reduce safety stock and urgent transport.
- **WP4**: Analytical frameworks to identify key factors affecting back-and-forth shipments and parts obsolescence, supporting resource-efficient aftermarket operations.

When preparing the AIM-TRUE project proposal, we formulated these deliverables as clear targets designed to strengthen the competitiveness of Swedish industry by improving aftermarket availability while materially lowering the resource intensity of operations. For Volvo SML, more accurate, dealer-adaptive forecasting combined with uncertainty-aware inventory policies was intended to reduce tied-up working capital, safety stock and emergency logistics, while maintaining high service levels for customers operating heavy-duty fleets, thereby increasing uptime, customer loyalty and margins in a strategically important service business. For Rejmes, the same capabilities were intended to translate into more predictable replenishment, fewer stock-outs and fewer "back-and-forth" shipments, enabling more efficient workshop planning and a measurable reduction in transport, scrapping and associated emissions. Beyond direct operational gains, the project was designed to accelerate innovation at both industrial partners by introducing reusable decision-support methods, evaluation frameworks linked to operational KPIs, and a stronger data-driven culture. This will, hopefully, in time translate into creating a scalable foundation for new digital service offerings and continued capability building within Sweden's automotive and logistics ecosystem.

4 Project Deliverables and Key Performance Indicators

This section summarises the fulfilment of deliverables and Key Performance Indicators (KPIs) as defined in the original AIM-TRUE project proposal. The project has largely met or exceeded its stated quantitative and qualitative objectives. Some KPIs were not fully achieved, but it should still be viewed positively in light of the project's deliberately ambitious targets and its focus on transferring state-of-the-art AI/ML methods into a safety- and service-critical industrial setting. In addition, the work revealed previously underappreciated complexities (such as pronounced dealer-to-dealer heterogeneity, non-stationary demand driven by campaigns and assortment changes, intermittent "zero-inflated" sales with long cold-start periods for new or replaced parts, as well as data-quality and lead-time constraints in multi-echelon logistics), which both explain the remaining gaps and constitute valuable knowledge creation for future initiatives.

Primarily, the project developed a simulation pipeline to evaluate how demand forecasts translate into real business costs. It consists of a synthetic data generator that mimics realistic, intermittent demand patterns, an ML-based forecasting module, and a cost simulator modelling daily inventory operations. By simulating diverse demand scenarios, the pipeline quantifies the cost range associated with different forecasting strategies, directly linking prediction accuracy to operational outcomes. This enables safe, data-driven testing of demand predictions, including prediction intervals to assess upper and lower cost bounds, helping reduce rush orders in future real-world deployment at Volvo while ensuring economically and environmentally sound decisions.

Work Package	Target	Outcome & Evidence / Evaluation Method	Contribution to KPIs
WP1: Dissemi-nation	≥3 journals, ≥4 conferences	Targets met, open access publications and open-source software release	Peer-reviewed publications, open-source software, and dissemination activities ensure transparency, reproducibility, and knowledge transfer within the ecosystem
WP2: Forecast accuracy	+15%	Up to +20% MAE improvement; evaluated on real Volvo SML data	AI-based forecasting methods (meta-learning, survival-based, hybrid architectures) improved forecast accuracy, a prerequisite for reducing downstream inefficiencies.
WP3: Safety stock reduction	−10%	>10% cost reduction for low-volume parts; shown via discrete-event inventory simulation	Uncertainty-aware inventory policies using prediction intervals better balance service level and stockholding risk, reducing costs while maintaining availability.
WP3: Rush order reduction	−10%	Enabled through forecasting & PI-based inventory policies; shown in simulation & qualitative dealer feedback	Multi-step forecasting enables proactive replenishment, reducing reliance on reactive ordering and resource-intensive express deliveries.
WP4: Back-and-forth shipment reduction	−15%	Drivers identified; reduction potential quantified via cost simulator + lifecycle analysis	Decision-centric simulation framework links forecast errors to operational outcomes, identifying actions to reduce unnecessary transport. Scalable grouping methods support balanced stock placement, reducing misallocated inventory.
WP4: Scrapped parts reduction	−5%	Reduction potential identified via lifetime inflexion points; evaluated using long-horizon forecasting & obsolescence analysis	Simulation framework and long-horizon forecasting help identify obsolescence risks, supporting reductions in material waste.

All AIM-TRUE project partners were initially hoping to conduct real-life deployment experiments with Volvo Group and Rejmes Transportfordon to test some of the new services. However, from a business perspective, this proved too challenging, due to a combination of technical and infrastructure limitations, as well as legal and financial constraints, making a full-scale real-world experiment impractical during the timeframe of the AIM-TRUE project.

Therefore, some of the results presented in this final report are based on the simulation platform and pilot studies rather than a full-scale industrial deployment, as explicitly stated. The table below presents a consolidated overview of the project's KPI targets and achieved outcomes.

5 Research Outputs

The AIM-TRUE project produced a suite of novel AI and ML methods tailored to the unique challenges of intermittent demand forecasting and inventory optimisation in the automotive aftermarket. Each output was validated using real industrial data from Volvo SML and Rejmes Transportfordon, as well as simulated data, to ensure practical relevance. Below is a detailed overview of the key innovations, their operational implications, and the evidence supporting their impact.

- **Intermittent Demand Forecasting with Meta-learning and Neural Networks**

The Challenge: Forecasting demand for spare parts is notoriously difficult, with 77% of time periods showing zero sales, making traditional models unreliable. This leads to costly overstocking, rush orders, and wasted resources.

Our Solution: We developed a family of machine learning models that learn globally from all available data while adapting locally to individual dealers and parts. Unlike "one-size-fits-all" approaches, our methods use **trainable embeddings** for parts and dealers, allowing the model to recognise patterns such as "expensive parts with low turnover" or "high-activity dealers in coastal regions." We also introduced a new model using **Kolmogorov–Arnold Networks (KAN)**, which replace fixed activation functions with learnable splines—making the model not only more accurate but also interpretable. In real tests on Volvo SML data, these models achieved up to **20% better accuracy** than existing statistical baselines.

Why It Matters: Better forecasts mean fewer stockouts and less excess inventory. For a single dealer, this can reduce safety stock by thousands of euros annually while maintaining service levels. The interpretability of KANs also helps logistics managers understand *why* a forecast was made, building trust in AI-driven recommendations.

Read more in:
Żarski, M., & Nowaczyk, S. (2025, May). *Balancing Performance and Scalability of Demand Forecasting ML Models*. In International Symposium on Intelligent Data Analysis (pp. 127-140). Cham: Springer Nature Switzerland.

Fukuhara, S., Alabdallah, A., Gunasekara, N., Nowaczyk, S. (2026). *Bridging Forecast Accuracy and Inventory KPIs: A Simulation-Based Evaluation Framework*. In International Symposium on Intelligent Data Analysis (Accepted).

- **Neuroplastic Multi-Task Network for Demand Forecasting**

The Challenge: In large networks with thousands of dealer–part combinations, training a separate model for each is computationally infeasible, but a single global model fails to capture local nuances.

Our Solution: Inspired by biological neuroplasticity, we created the **Neuroplastic Multi-Task Network (NMT-Net)**, a dynamic neural architecture that grows intelligently as new tasks are introduced. When a new dealer–part combination arrives, NMT-Net identifies the most similar previously learned task and either reuses its specialised "head" or trains a new one. This allows the system to maintain high accuracy without ballooning in size.

Why It Matters: NMT-Net makes it practical to deploy personalised forecasting across thousands of SKUs without requiring a supercomputer. In simulations, it reduced the number of models needed by over **90%** while preserving forecasting performance. This is a game-changer for scaling AI in enterprise logistics.

Read more in:
Żarski, M., & Nowaczyk, S. (2026). *Neuroplasticity-inspired Dynamic ANNs for Multi-task Demand Forecasting* (Under Review).

- **Performance–Scalability Trade-offs in Large-scale Forecasting**

The Challenge: How do you balance the need for highly accurate local models with the reality of limited computational resources?

Our Solution: We systematically explored the spectrum from fully local to fully global models and introduced two smart **model grouping strategies**:

- **Clustered MLP:** We used embeddings to group similar dealer–part pairs into clusters, then trained one model per cluster. This cut the number of models from thousands to just **15**, with no loss in accuracy.
- **Grouped MLP:** An adaptive algorithm that iteratively groups pairs that can share a model without performance degradation, reducing model count by **95%**.

Why It Matters: These methods make it feasible to run high-accuracy forecasting on standard enterprise hardware, lowering the barrier to AI adoption for mid-sized logistics providers.

Read more in:
Żarski, M., & Nowaczyk, S. (2025, May). *Balancing Performance and Scalability of Demand Forecasting ML Models*. In International Symposium on Intelligent Data Analysis (pp. 127-140). Cham: Springer Nature Switzerland.

- **Hybrid Batch–Streaming Learning for Concept Drift**

The Challenge: Real-world demand patterns shift over time due to factors such as fleet changes, seasonal campaigns, and new vehicle models. Static models quickly become outdated.

Our Solution: We developed a **hybrid learning paradigm** that combines a batch-trained deep neural network with a lightweight streaming regressor. The neural network captures complex historical patterns, while the streaming component, a **SOKNL regressor**, continuously adapts to new trends by learning from prediction residuals. This setup avoids costly retraining while staying responsive to change.

Why It Matters: In tests simulating sudden demand shifts (e.g., a new service campaign), our hybrid model adapted **3× faster** than retrained static models, maintaining accuracy without operational disruption. This is especially critical in the transition to e-mobility, where component demand evolves rapidly.

Read more in:
Gunasekara, N., Nowaczyk, S., & Pashami, S. (2025, May). *Pragmatic Paradigm for Multi-stream Regression*. In International Symposium on Intelligent Data Analysis (pp. 358-372). Cham: Springer Nature Switzerland.

- **Uncertainty Quantification and Prediction Intervals for Inventory Control**

The Challenge: A single-point forecast doesn't tell you how confident the model is. Without understanding uncertainty, inventory managers either overstock (costly) or understock (risky).

Our Solution: We implemented **bootstrap-based Prediction Intervals (PIs)** that quantify forecast uncertainty by resampling historical errors. These PIs were integrated into inventory decision rules and tested in a **discrete-event simulation** of Volvo's spare parts network. We compared strategies such as using the upper bound of the PI for safety stock or reshaping the PI into a Beta-PERT distribution for reorder calculations.

Why It Matters: For low-volume, high-cost parts, PI-based policies reduced safety stock by **>10%** while maintaining service levels. The simulation also revealed a "tipping point" in demand volume where PIs stop adding value, enabling smarter portfolio segmentation.

Read more in:
Kamil, J., & Amer, M. (2025). *Prediction Intervals for ML-driven Automotive Service Market Logistics*. Master's thesis, Halmstad University, Halmstad.

- **Decision-centric Evaluation and Simulation Framework**

The Challenge: A forecasting model can achieve high statistical accuracy but still result in poor inventory outcomes. How do you evaluate models based on real business impact?

Our Solution: We built a **modular simulation framework** that connects synthetic demand generation, forecasting models, and inventory policies to measure **operational KPIs**, not just purely data-oriented metrics like MAE or RMSE. The framework includes a realistic demand generator that mimics intermittent patterns, seasonal effects, and concept drift. It allows us to test how forecast bias, noise, or intermittency affects total cost, service levels, and rush orders.

Why It Matters: In one experiment, we observed a **Simpson's paradox**: a model with worse accuracy actually yielded lower total costs due to favourable error patterns. This underscores the need for decision-centric evaluation. This solution is now being explored in Volvo SML's internal assessment processes.

Read more in:
Fukuhara, S., Alabdallah, A., Gunasekara, N., Nowaczyk, S. (2026). *Bridging Forecast Accuracy and Inventory KPIs: A Simulation-Based Evaluation Framework*. In International Symposium on Intelligent Data Analysis (Accepted).

- **Survival Analysis for Intermittent Demand**

The Challenge: Intermittent demand isn't just about "how many," but "when." Predicting the time until the next sale is crucial for proactive replenishment.

Our Solution: We applied **survival analysis** to separately model the "time-to-demand" and "demand quantity" for spare parts. Using a neural network with two heads and an evolutionarily optimised loss function, we improved the **concordance Index** (a measure of temporal ranking accuracy) and outperformed state-of-the-art methods by up to **8%** in sales quantity predictions.

Why It Matters: This approach complements volume-based forecasts by providing a probabilistic view of *when* demand will occur, enabling more precise order timing and reducing the need for safety stock.

Read more in:
Altarabichi, M. G., Alabdallah, A., Pashami, S., Rögnvaldsson, T., Nowaczyk, S., & Ohlsson, M. (2024, July). *Improving Concordance Index in Regression-based Survival Analysis: Evolutionary Discovery of Loss Function for Neural Networks*. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (pp. 1863-1869).

- **Open-source Research Software**

The Outcome: To ensure transparency and foster community adoption, we released the **Truck Parts Demand Inventory Simulator** as open-source software. This Python-based tool allows researchers and practitioners to simulate spare parts demand, test forecasting models, and evaluate inventory policies in a realistic, customizable environment.

Get the code:
https://github.com/SoFukuhara/TruckParts-Demand-Inventory-Simulator

Why It Matters: By sharing our simulation environment, we enable others to build on our work, accelerating innovation in sustainable logistics beyond the automotive sector.

Read more in:
Fukuhara, S., Alabdallah, A., Gunasekara, N., Nowaczyk, S. (2026). *Bridging Forecast Accuracy and Inventory KPIs: A Simulation-Based Evaluation Framework*. In International Symposium on Intelligent Data Analysis (Accepted).

These outputs represent a cohesive advancement toward **AI-driven, resource-efficient aftermarket operations**. Each contribution was designed not only to push the state of the art in ML, but to deliver tangible reductions in waste, cost, and emissions, proving that smarter algorithms can lead to a more sustainable supply chain.

6 Utilisation, Implementation, and Industrial Uptake

The outcomes from the AIM-TRUE project span different maturity levels, from deployed analytical tools to research-based decision support, highlighting the unique benefits of such close academia-industry collaboration. Besides research papers, the project had significant practical benefits.

Implemented / Operational Use

- Decision-centric evaluation frameworks and forecasting accuracy metrics have informed Volvo SML's internal assessment of forecasting strategies.
- Insights from prediction interval–based inventory control have influenced internal discussions on safety stock segmentation.

Piloted / Evaluated

- Machine Learning–based forecasting models were evaluated on real Volvo SML data.
- Inventory control policies integrating uncertainty quantification were tested in simulation and validated through interviews with Rejmes Transportfordon personnel.

Research Outputs

- Novel architectures (e.g., NMT-Net, Forecast(t)) and evaluation frameworks were developed and validated at the research and pilot levels.
- Open-source software (TruckParts-Demand-Inventory-Simulator) enables reuse and extension beyond the project consortium.

7 Societal, Environmental, and Strategic Impact

The AIM-TRUE project contributed to national sustainability and competitiveness objectives by addressing resource inefficiencies in the heavy-duty vehicle aftermarket. Reductions in rush orders, unnecessary transport, and scrapped parts directly translate into lower emissions, reduced energy use, and more efficient material utilisation.

Beyond Volvo SML, the methods and tools developed are applicable to other sectors characterised by intermittent demand and high service-level requirements, such as rail maintenance, industrial equipment servicing, and energy infrastructure. By releasing key software components as open source and publishing results in peer-reviewed venues, the project supports broader knowledge diffusion within the Swedish and international logistics ecosystem.

8 Deviations and Lessons Learned

No major deviations from the project plan occurred. However, some KPIs related to physical logistics flows (e.g., back-and-forth shipments) could not be fully validated through live operational deployment within the project timeframe. Instead, these were evaluated using our simulation environments and historical data.

A key lesson learned is that improvements in statistical forecast accuracy do not automatically translate into operational gains. Decision-centric evaluation and uncertainty-aware inventory policies are critical to achieving real-world resource efficiency. These topics are at the core of the follow-up project AIM-FLIX.

9 Gender Equality and Project Execution

The project followed the gender equality principles outlined in the proposal. Academic and industry participation ensured inclusive collaboration, and no gender-related barriers were identified during project execution. While the project's technical nature limited measurable gender-specific outcomes, equality considerations were respected throughout staffing, supervision, and dissemination activities.

10 Publication list

The main research outputs and their associated publications and software are:

- Altarabichi, M. G., Alabdallah, A., Pashami, S., Rögnvaldsson, T., Nowaczyk, S., & Ohlsson, M. (2024, July). **Improving Concordance Index in Regression-based Survival Analysis: Evolutionary Discovery of Loss Function for Neural Networks**. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (pp. 1863-1869).
- Żarski, M., & Nowaczyk, S. (2025, May). **Balancing Performance and Scalability of Demand Forecasting ML Models**. In *International Symposium on Intelligent Data Analysis* (pp. 127-140). Cham: Springer Nature Switzerland.
- Gunasekara, N., Nowaczyk, S., & Pashami, S. (2025, May). **Pragmatic Paradigm for Multi-stream Regression**. In *International Symposium on Intelligent Data Analysis* (pp. 358-372). Cham: Springer Nature Switzerland.
- Kamil, J., & Amer, M. (2025). **Prediction Intervals for ML-driven Automotive Service Market Logistics**. *Master's thesis, Halmstad University*, Halmstad.
- Fukuhara, S., Alabdallah, A., Gunasekara, N., Nowaczyk, S. (2026). **Bridging Forecast Accuracy and Inventory KPIs: A Simulation-Based Evaluation Framework**. In *International Symposium on Intelligent Data Analysis* (Accepted, 2026).
- Żarski, M., & Nowaczyk, S. (2026). **Neuroplasticity-inspired Dynamic ANNs for Multi-task Demand Forecasting** (*Under Review*).

Software: Open-source release of the TruckParts-Demand-Inventory-Simulator.

11 Participating parties and contact persons

Halmstad University Contact person: Slawomir Nowaczyk Slawomir.Nowaczyk@hh.se		
Volvo Group Service Market Logistics (SML) Contact person: Iulian Carpatorea Iulian.Carpatorea@volvo.com		
Rejmes Transportfordon Contact person: Anders Malmberg Anders.Malmberg@rejmesbil.se		