Fuel Efficiency in Ferry Services: GPS-Based Clustering and Explainable AI

***Robert Klar*[1, 2], *Nina Svensson*[1], *Petra Stelling*[1], *Vangelis Angelakis*[2]**

[1]Swedish National Road and Transport Research Institute (VTI)

[2]Department of Science and Technology, Linköping University (LiU)

vti

Background

- Ferry operations differ fundamentally from long-distance shipping
  - Short legs, frequent stops, and high operational variability complicate fuel analysis
- Tight schedules dominate operational decisions
  - Punctuality and service reliability leave limited margins for speed optimization
- Data availability is no longer the bottleneck
- Core challenges
  - Raw trajectories do not explain why fuel consumption differs
  - Inefficiencies are hidden in complex, noisy operational patterns

Varying fuel consumption

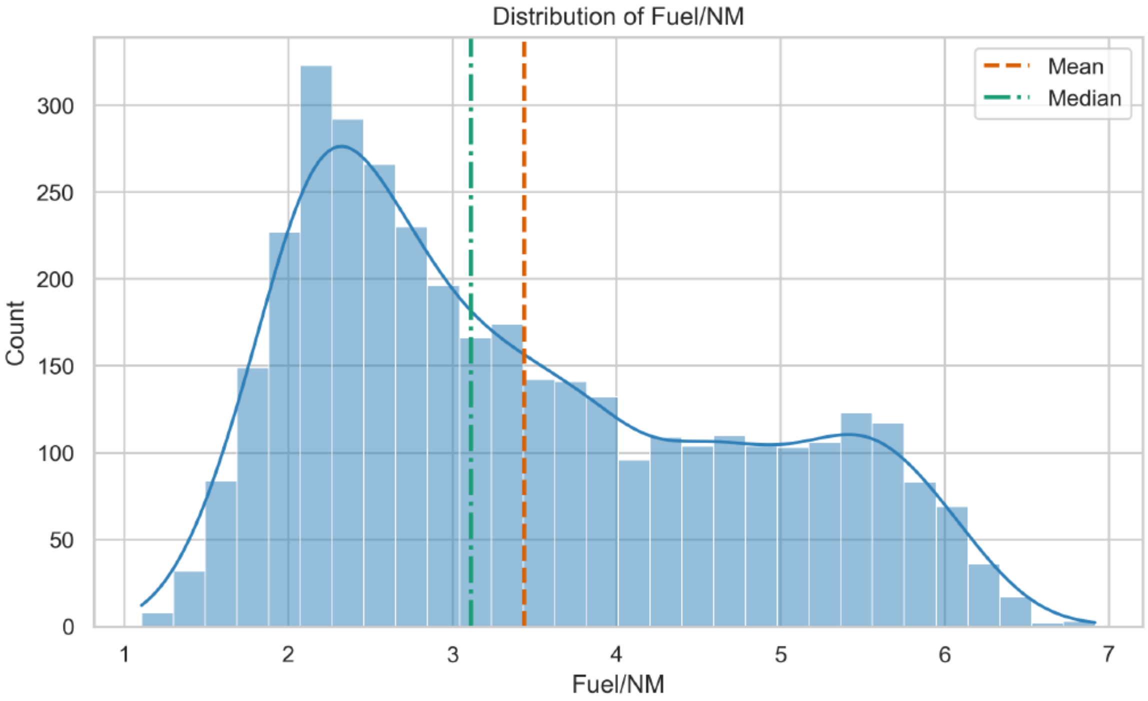

Motivation

- Limitations of existing research
  - ➢ Trajectory analysis often stops at route reconstruction
  - ➢ Fuel models prioritize predictive accuracy over interpretability
  - ➢ Most prior research is based on cargo shipping using AIS data
- Consequences
  - ➢ Operators are left with maps and black-box predictions
  - ➢ Limited support for operational decisions
- Contribution
  - ➢ A framework that discovers operational patterns, predicts fuel consumption, and explains fuel-inefficient behavior using XAI

Study context and data

- Operational setting
  - ➢Stockholm archipelago ferry operations
  - ➢Operator: Blidösundsbolaget
  - ➢Vessel: Silverö
- Data (Summer timetable 2023)
  - ➢Granular 1-second GPS data (position, fuel consumption, speed)
  - ➢Wind speed and direction from the nearest SMHI station

Photo: Liza Simonsson / Blidösundsbolaget.

A brief overview of ML & (X)AI

Supervised ML

- uses labeled datasets to train algorithms to make accurate predictions or classifications

Unsupervised ML

- finds patterns in unlabeled data without human guidance

AI involves techniques aimed to emulate human behavior

ML uses algorithms to detect patterns in large data sets

XAI reveals the reasoning behind model predictions

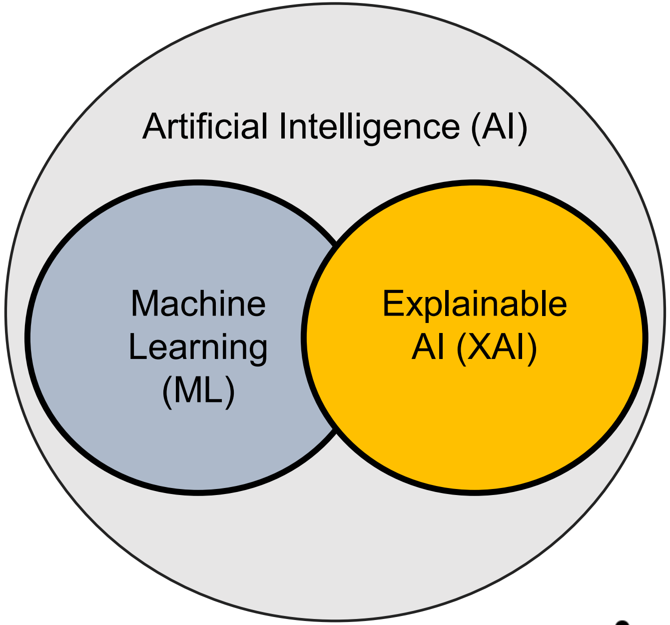

Overall methodology

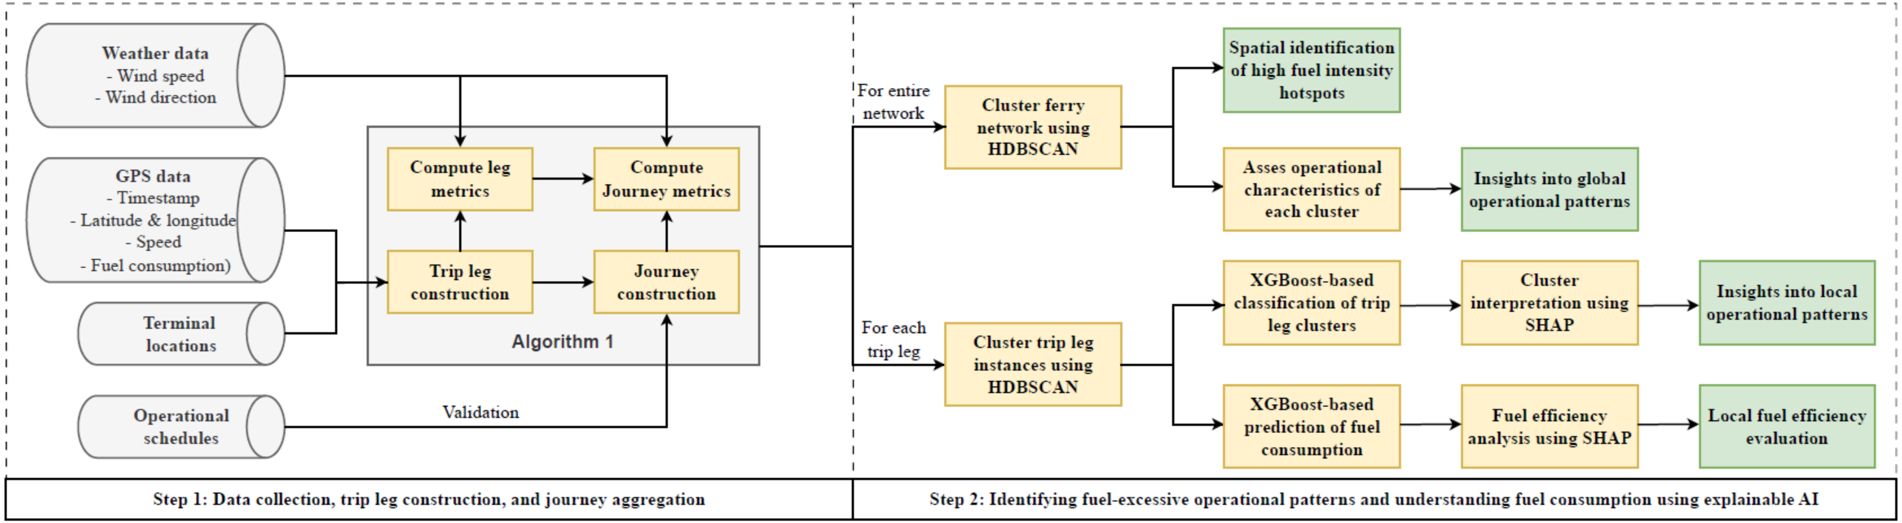

From raw GPS to meaningful operations

- Key challenge: Individual GPS points have no operational meaning
- Solution proposed in the paper
  - ➢ Detect and validate real stop visits
  - ➢ Construct trip legs as stop-to-stop operational units
  - ➢ Aggregate legs into journeys
- Enables fair comparison across routes and days

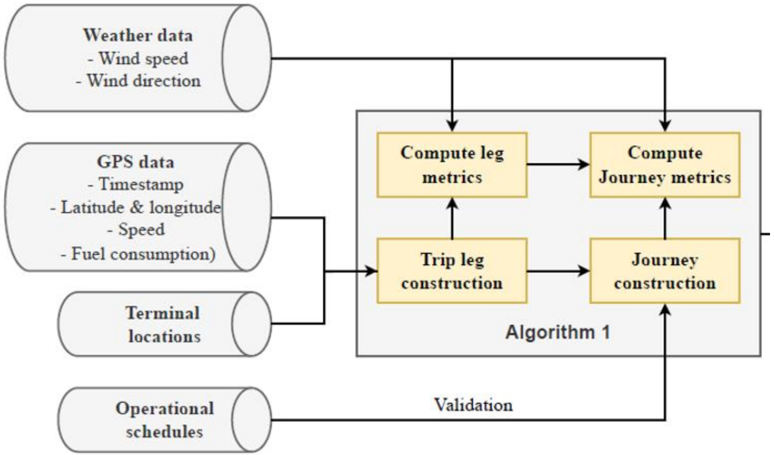

vti

Feature engineering

- Distance and time normalization to enable fair comparison
- Fuel residual relative to typical fuel use at similar speeds
- Headwind and crosswind components

Variable	Description	Unit	Typical Range	Median	Role in analysis
Distance (NM)	Total distance traveled by the vessel during a trip leg, computed from geospatial trajectory data.	NM	0.2–3.9	0.88	Normalization variable for fuel efficiency
Mean speed (kn)	Average vessel speed over a trip leg, derived from GPS-based speed observations.	kn	6.5–15.4	9.92	Primary operational driver; clustering and prediction feature
Fuel consumption per NM	Total fuel consumed over a trip leg normalized by travel distance; indicator of energy efficiency per distance traveled.	L/NM	1.8–5.8	3.11	Efficiency indicator; clustering feature
Headwind component	Wind component aligned with the vessel's direction of motion; positive values indicate headwind and negative values tailwind.	m/s	−3.7–3.6	0.04	Environmental influence feature
Crosswind component	Wind component perpendicular to the vessel's direction of motion; sign indicates lateral direction relative to vessel heading.	m/s	−3.6–3.7	0.01	Environmental influence feature
Relative fuel residual	Relative deviation of the fuel consumption rate from the median fuel rate observed within the same speed bin; positive values indicate higher-than-typical fuel use at a given speed.	–	−0.3–0.4	0.00	Performance anomaly indicator; clustering feature

Three-stage XAI framework

- First stage: Unsupervised learning (HDBSCAN)
  - ➢ Discover recurring operational clusters without predefined labels
- Second stage: Supervised learning (XGBoost)
  - ➢ Generalize clusters and predict cluster membership and fuel consumption
- Third stage: Explainability (SHAP)
  - ➢ Explain why clusters differ in fuel use and support operational decisions

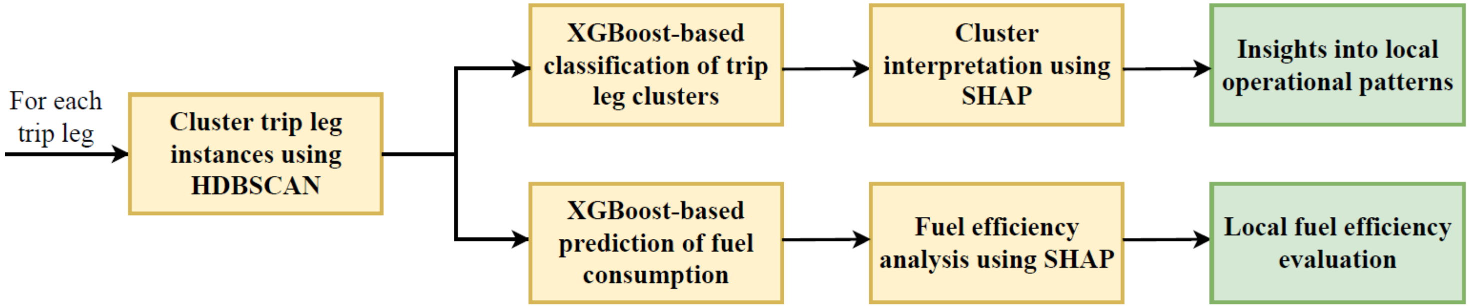

Discovering Operational Patterns (Globally)

- Approach
  - ➢ HDBSCAN applied to trip-leg feature space
- What clusters represent
  - ➢ Distinct speed and operational behavior patterns
  - ➢ Strong differences in fuel intensity between clusters

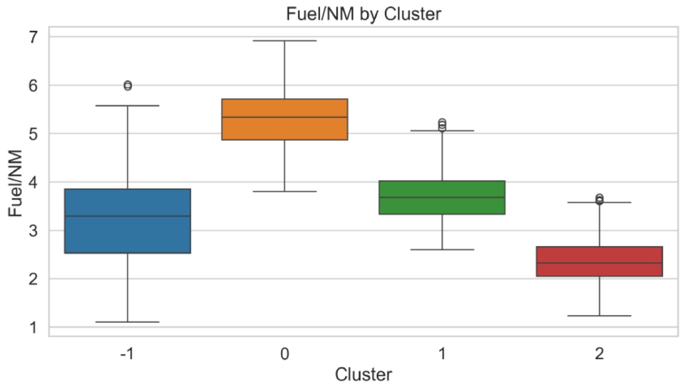

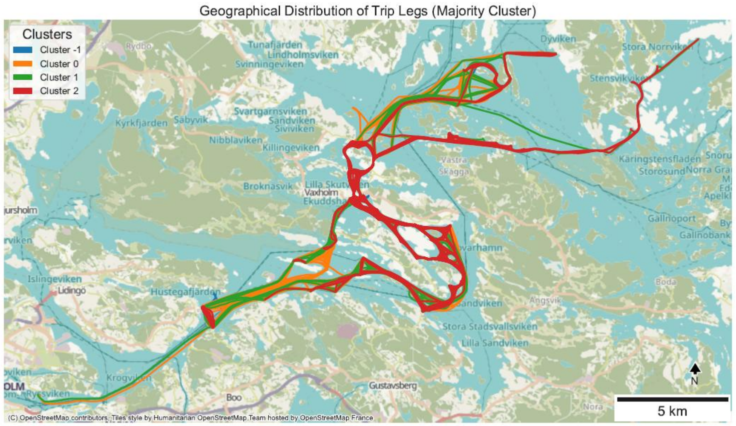

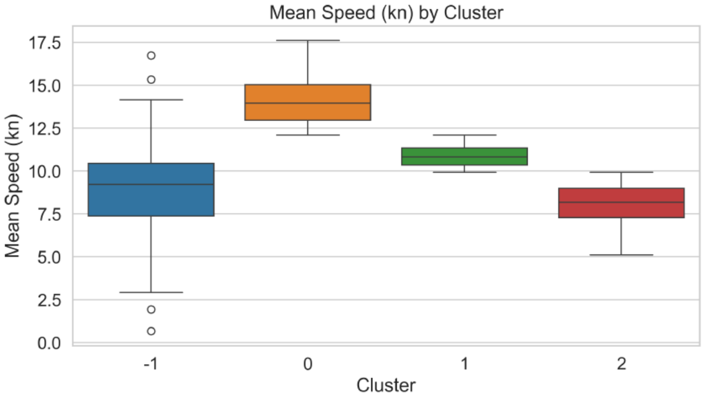

Discovering Operational Patterns (Locally)

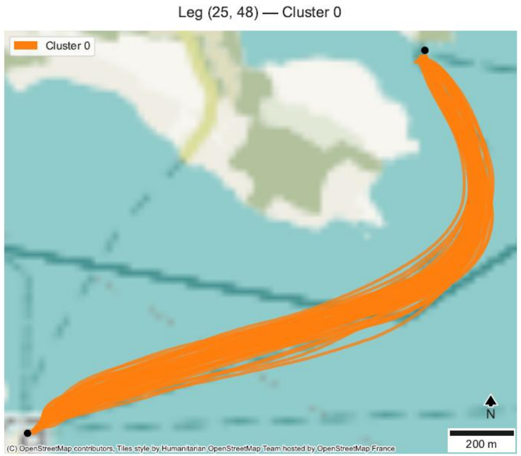

(a) Overview of trip leg instances belonging to cluster 0.

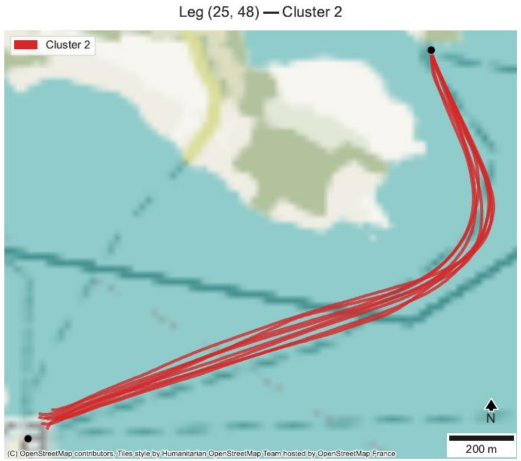

(b) Overview of trip leg instances belonging to cluster 2.

Feature	Mean	Min	Max
Fuel/NM	4.84	4.34	5.75
Mean Speed (kn)	12.83	12.14	14.54
Headwind (m/s)	-0.37	-4.3	4.49
Crosswind (m/s)	0.14	-4.08	4.41
Fuel Residual	-0.17	-0.3	0.08
Distance (NM)	1.16	1.09	1.21

Feature	Mean	Min	Max
Fuel/NM	2.59	2.49	2.74
Mean Speed (kn)	9.64	9.32	9.9
Headwind (m/s)	-1.96	-3.93	0.22
Crosswind (m/s)	2.88	1.55	4.07
Fuel Residual	0.07	0.0	0.16
Distance (NM)	1.13	1.11	1.15

vti

Training and evaluation strategy

- Why focus on individual trip legs
  - Ferry operations consist of recurring stop-to-stop trip legs
  - Analyzing legs separately avoids confounding route length and geometry effects
- Training data
  - Models are trained on the most frequently occurring trip leg
  - This provides a large, consistent dataset capturing typical operational behavior
- Evaluation data
  - Models are evaluated on the second most frequent trip leg
  - This leg is operationally similar but not identical to the training leg
- Purpose of this setup
  - Test whether learned patterns generalize beyond a single route
  - Avoid overly optimistic performance driven by repeated observations of the same leg

vti

Learning and generalizing patterns

- Supervised models used
  - ➢ XGBoost classifier for cluster interpretation
  - ➢ XGBoost regressor for fuel consumption prediction
- Why XGBoost
  - ➢ High performance on tabular data
  - ➢ Compatible with SHAP explainability
- Model performance
  - ➢ Cluster classification accuracy ≈ 94%
  - ➢ Fuel prediction $R^2 \approx 0.97$

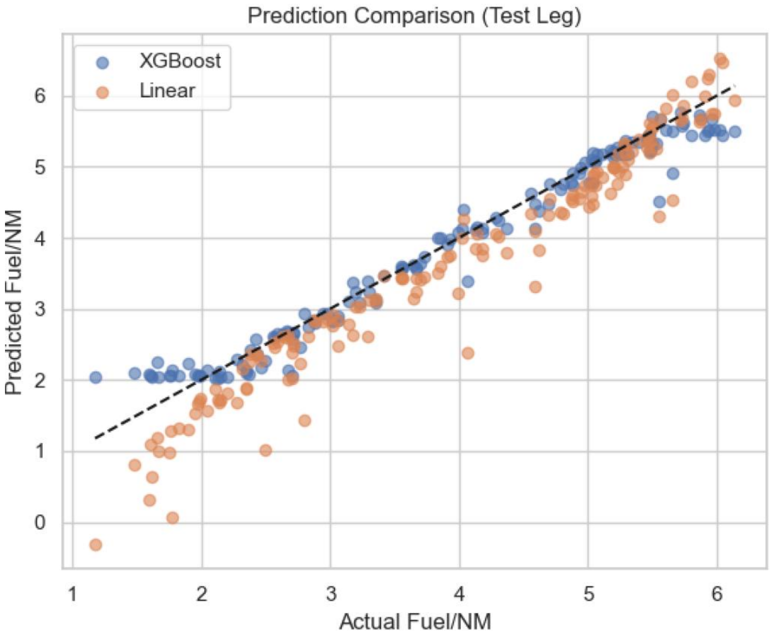

vti

Explaining Operational Clusters with SHAP

- What SHAP provides
  - ➢ Quantifies each feature's contribution to cluster formation
  - ➢ Reveals which operational characteristics define different clusters
- Cluster-wise interpretation
  - ➢ Enables comparison of feature influence across clusters
  - ➢ Allows clusters to be interpreted as distinct operational regimes

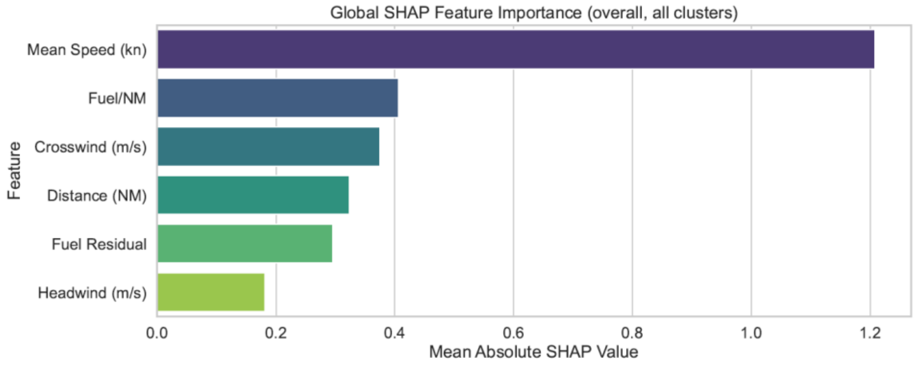

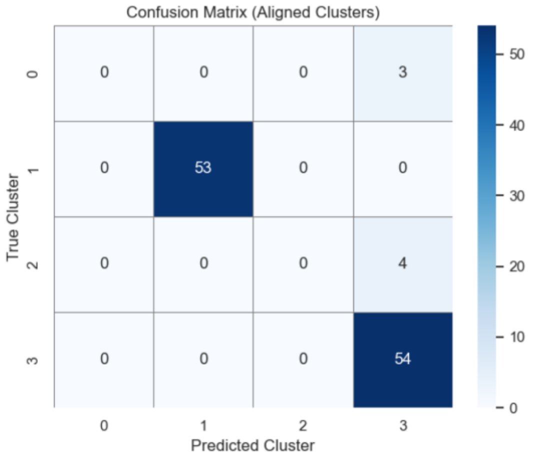

vti

Explaining fuel consumption with SHAP

- What SHAP provides
  - ➢ Quantifies each feature's contribution to fuel use
  - ➢ Allows comparison across clusters
- Key finding
  - ➢ Operational speed dominates fuel consumption
  - ➢ Speed variability also contributes significantly
- Wind effects are present but smaller
- Interpretation
  - ➢ Operational decisions outweigh environmental effects

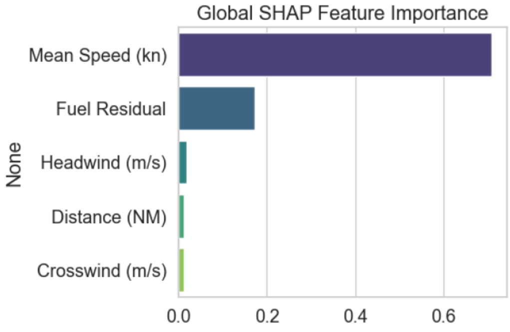

(a) Absolute Shap values

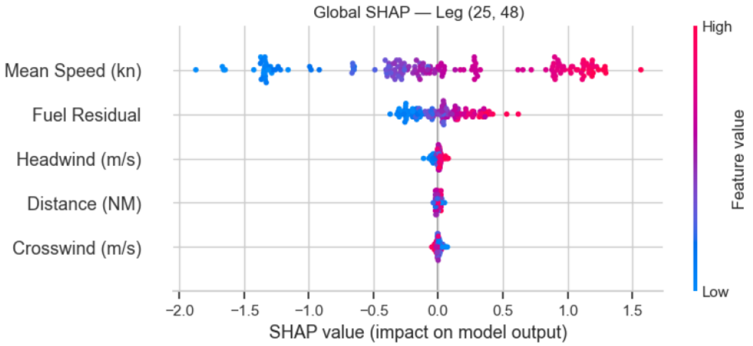

(b) Beeswarn plot.

Conclusions and Outlooks

- Main contributions
  - ➢ Identification of recurring, interpretable operational clusters in ferry operations
  - ➢ XAI analysis revealing key drivers of fuel consumption
- Key empirical insight
  - ➢ Operational factors, particularly speed, dominate differences in fuel consumption
  - ➢ High predictive performance (94% cluster accuracy, $R^2 = 0.97$ for fuel prediction)
- Limitations
  - ➢ Environmental effects are represented using land-based wind data
  - ➢ Fine-scale weather and sea-state influences may therefore be underestimated
- Future research
  - ➢ Incorporating higher-resolution onboard weather and wave measurements
  - ➢ Applying the framework across vessels, routes, and operators

vti

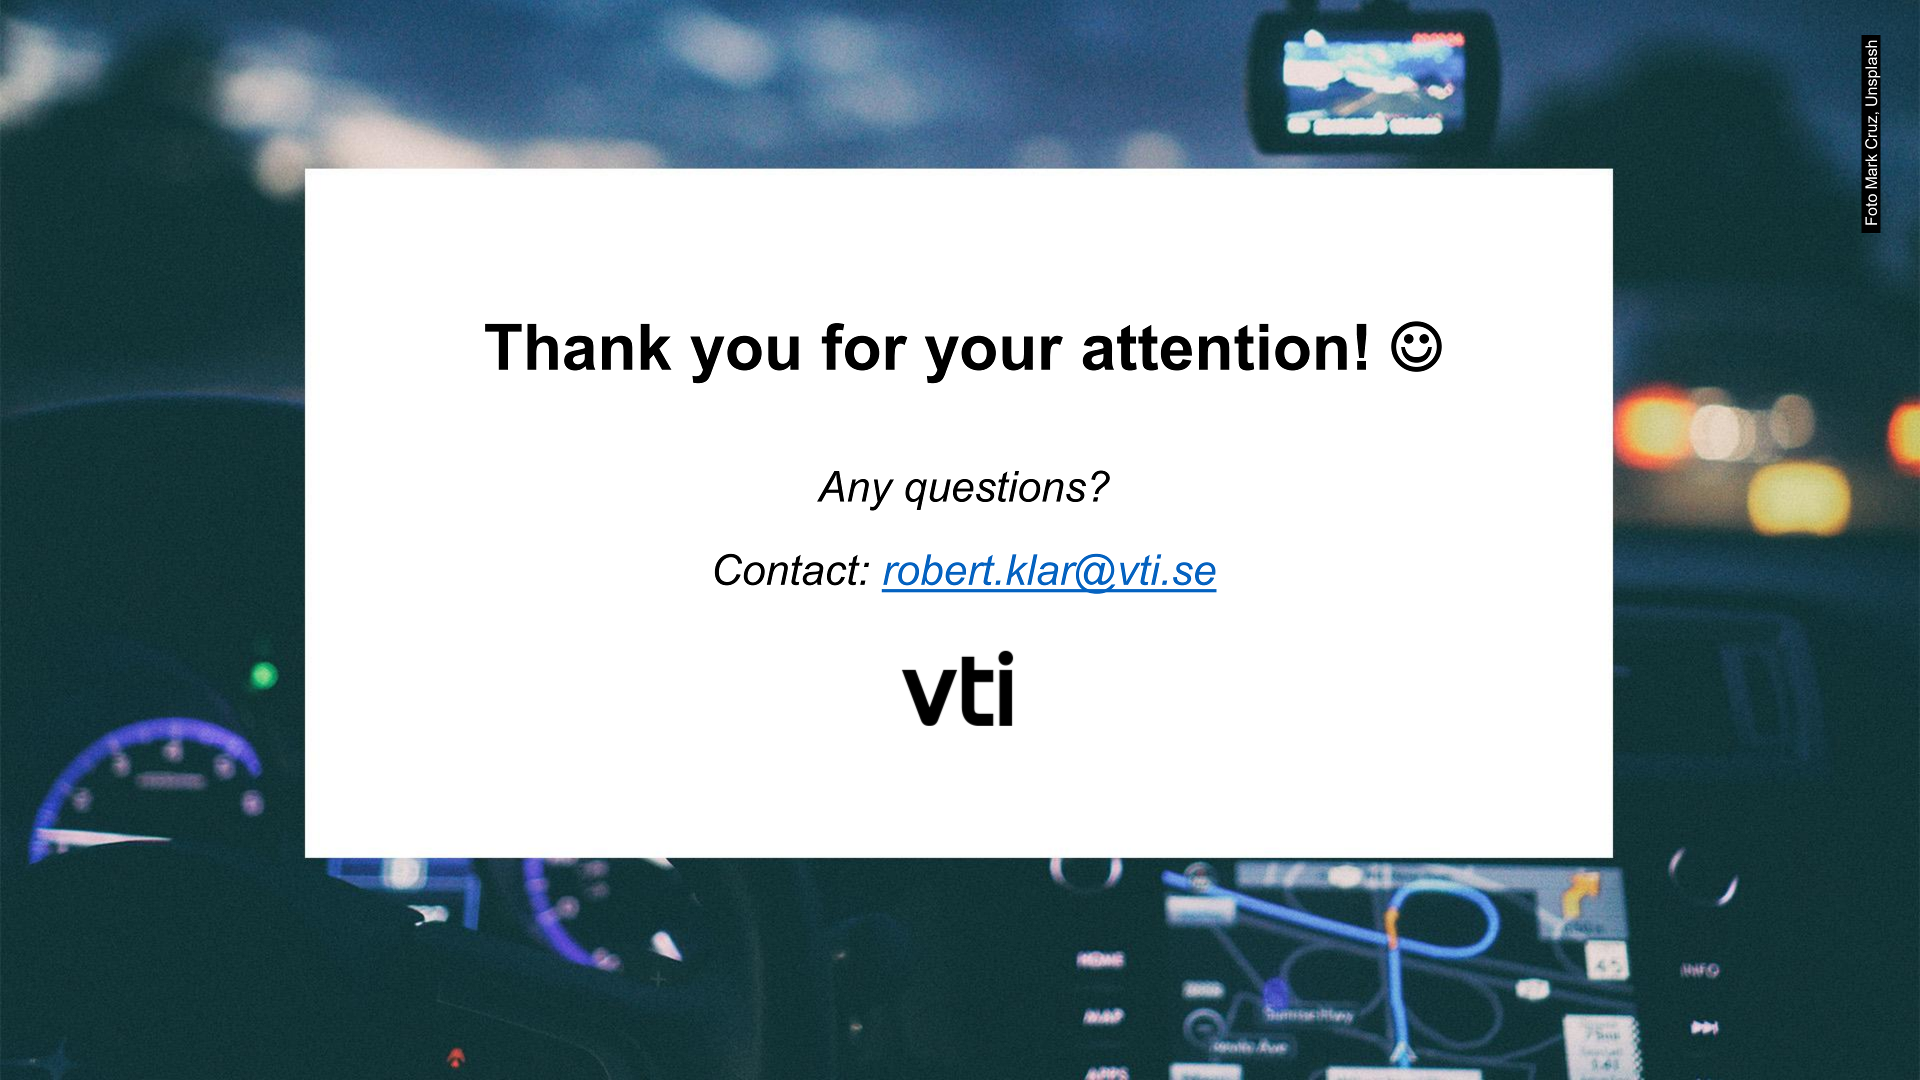

Foto Mark Cruz, Unsplash

Thank you for your attention! ☺

Any questions?

Contact: robert.klar@vti.se

vti