Stability in Truck Driving Behaviour

A Geo-Specific Analysis

Iva Surana

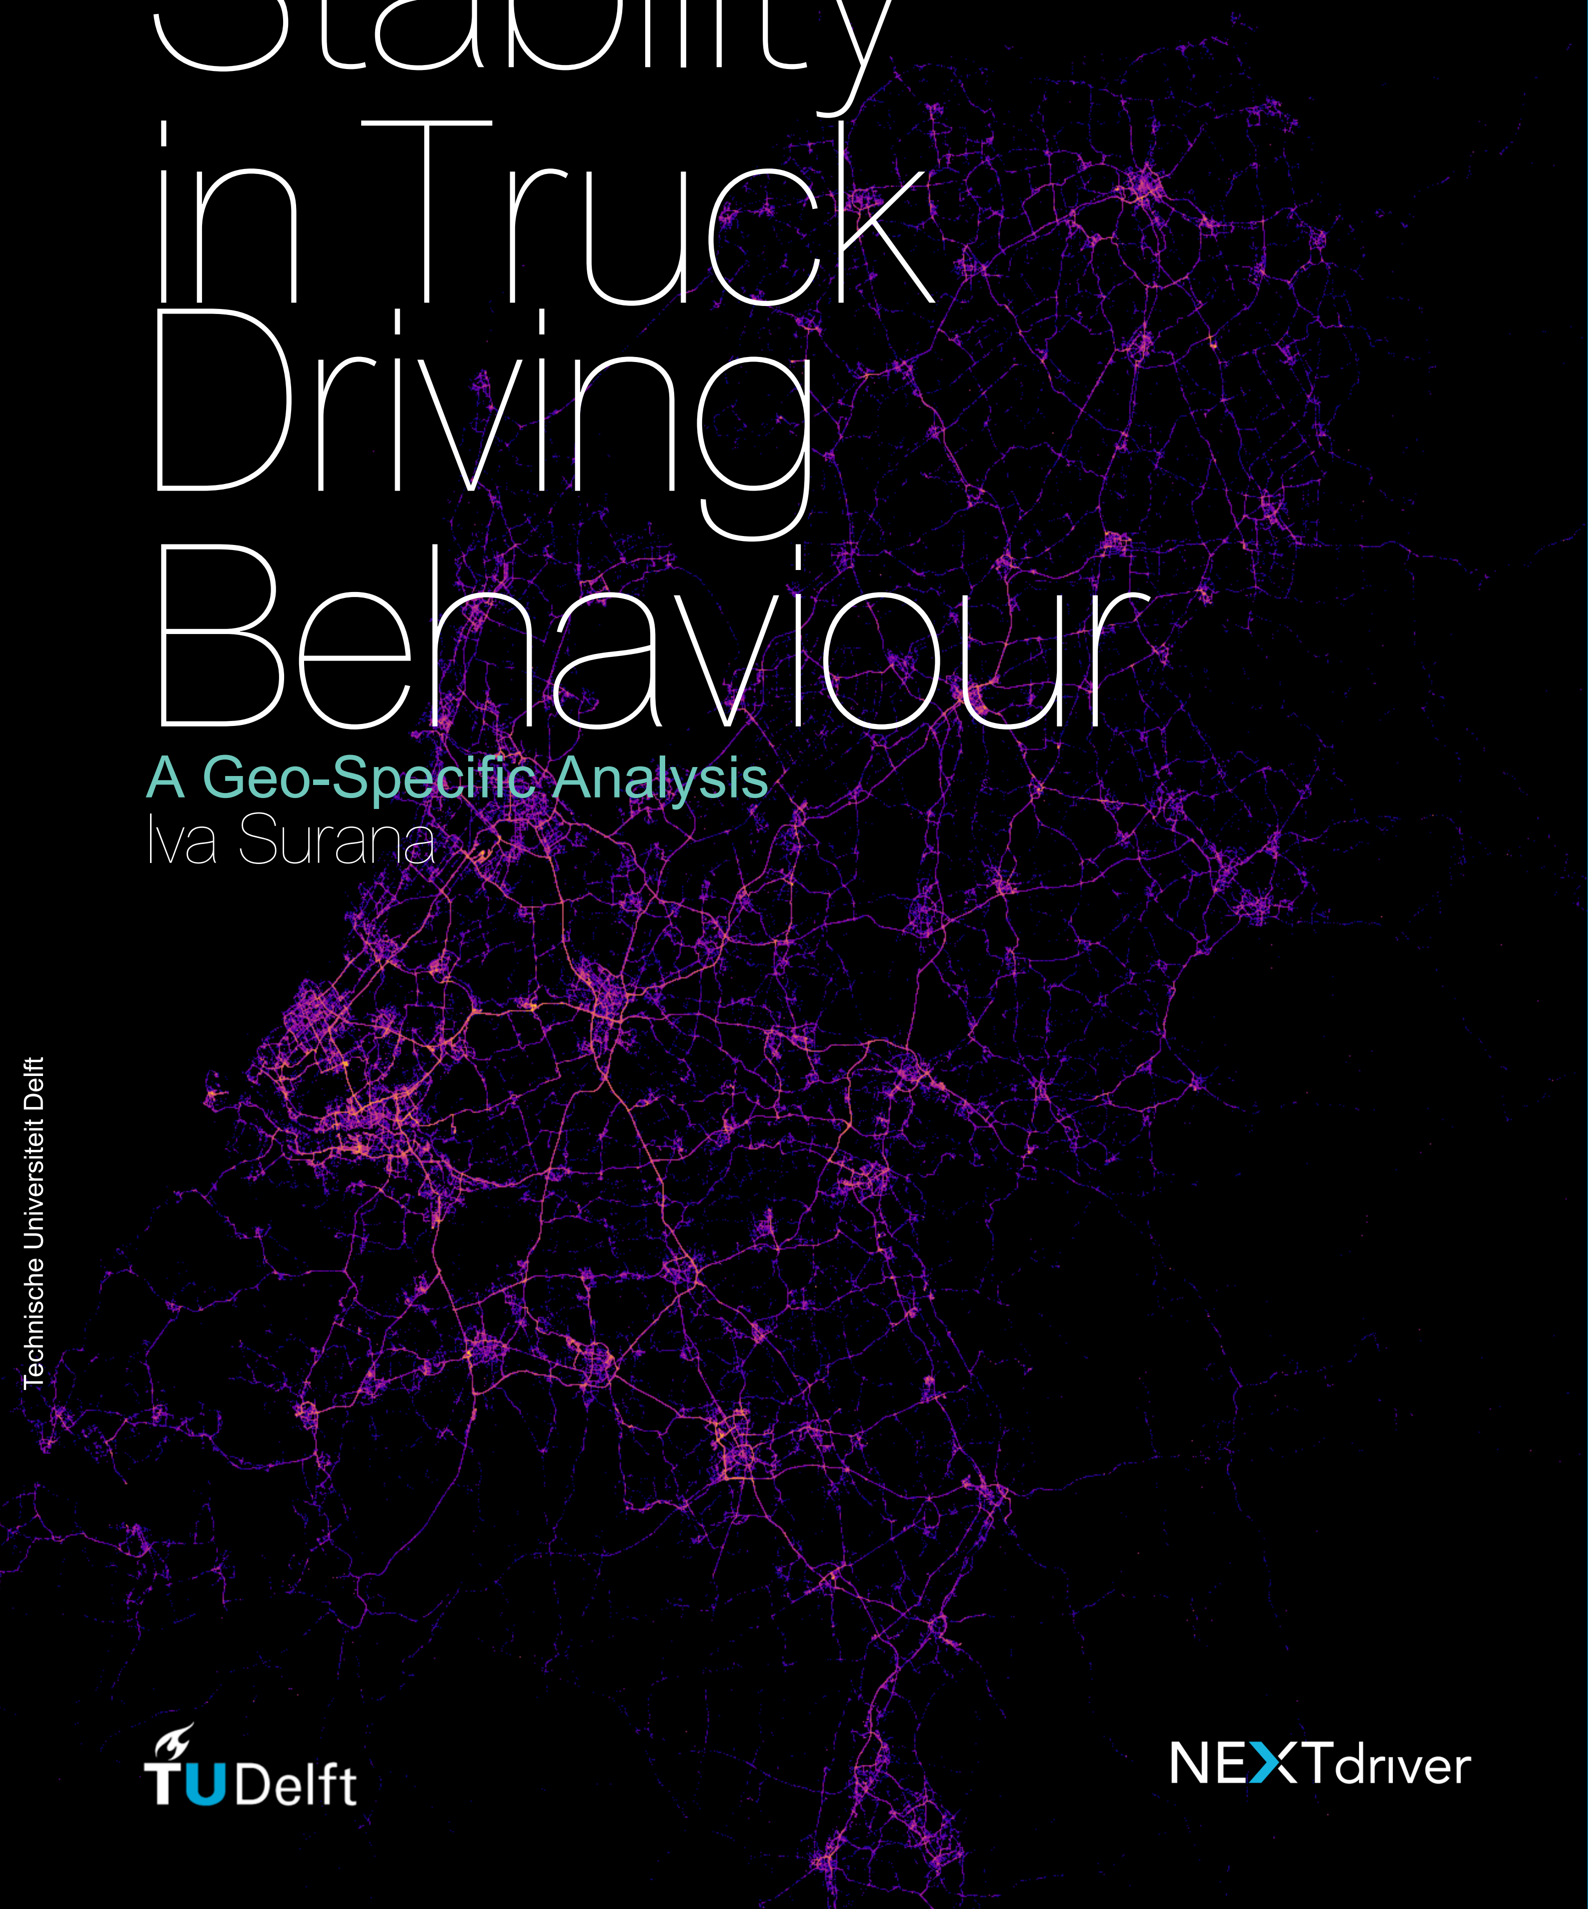

Technische Universiteit Delft

Stability in Truck Driving Behaviour

A Geo-Specific Analysis

by

Iva Surana

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday November 30, 2022 at 16:30.

Cover Image: Geospatial Visualization of Truck Driving Data in the Netherlands

Student Number :	5282896	
Project duration:	December 2021 – November 2022	
Supervisors :	Dr. D. (Dimitra) Dodou	TU Delft
	Dr. ir. J.C.F (Joost) de Winter	TU Delft
	Dr. P. (Pavlo) Bazilinskyy	TU Delft
	Ir. T. (Tom) Driessen	TU Delft
	Ir. J. (Jasper) Pauwelussen	NEXTdriver

An electronic version of this thesis is available at http://repository.tudelft.nl/.
The code for this thesis is available at https://github.com/bazilinskyy/truck-driver-profile.

Preface

This thesis marks my completion of the master's education in Robotics at TU Delft. The report documents the finding of my master's thesis work.

First and foremost, I would like to express my profound appreciation to all my supervisors, Tom Driessen, Pavlo Bazilinskyy, Joost de Winter, and Dimitra Dodou. This endeavour would not have been possible without their constant feedback, guidance and expertise. They are incredible supervisors with a fantastic work ethic and patience that I strive for. A special thanks to Dimitra for helping me navigate and deduce the results.

Furthermore, I want to thank my supervisor Jasper Pauwelussen for his supervision during the internship and his continued support for the thesis. This work would not have been possible without his enthusiasm and guidance. Throughout the project, he gave me the space and freedom to explore different areas. I'd also like to thank Guido Sluijsmans and the rest of the NEXTdriver team for their unwavering support.

A shout-out to all my friends in the Netherlands and India, who have made my life incredibly fun and helped me navigate through the entire project.

Lastly, I could not have undertaken this journey without my incredible parents and siblings. Their unconditional love has shaped me into who I am today. I want to thank them for believing in me, even when I doubted myself. They deserve the world!

Iva Surana
Delft, November 2022

Stability in Truck Driving Behaviour: A Geo-Specific Analysis

Iva Surana

Abstract

Naturalistic driving research with a focus on trucks has been gaining momentum in the past decade. With the advancement in sensor technology and access to big data, it becomes possible to understand driver behaviour at a more fundamental level. This can assist in mitigating the impact trucks have on the environment while enhancing safety. Several studies have worked towards examining the predictability of driving behaviour through driver profiling (i.e. scoring a driver's behaviour or classifying drivers by assigning them different categories such as aggressive/non-aggressive). However, little research still focuses on the importance and impact of individual features used to develop these models. In the current study, an analysis of driving data from 1,727 trucks recorded over one year as part of a Dutch Field Operational Test (FOT) has been performed. This FOT, to date, has not been investigated in the published academic literature. Recent studies have analysed historical location data to assess risk associated with specific routes and environments. This is being used to provide notifications to drivers around work zones to mitigate the impact of accidents.

The current thesis extends the geo-specific analysis of (truck) driving data by analysing stability in truck driving behaviour with a focus on time and location (urban areas and motorways). Here correlation analysis has been used to explore stability. Correlational analysis elucidates that metrics such as the number of headway warnings, braking events and lane departure warnings are stable over space and time. A discussion reflects on the role of vehicle characteristics (i.e. mass and engine power) towards stability.

It is concluded that in the case of spatial stability: Mean point speed has higher stability on motorways than in urban areas. It has been determined that trucks with higher mass and lower engine power tend to have lower mean speed than the norm. Contrary to mean point speed, headway warnings show higher stability in urban areas than on motorways. Braking events and lane departure warnings exhibit high stability. Secondly, a strong correlation between (t and t+1) hours over the entire day is observed for temporal stability.

This research is a precursor to building generalised models for profiling drivers and assessing various driving patterns. An in-depth understanding of different driving patterns can help driver coaching companies better understand metrics when time and location are factored in before providing targeted feedback. Apart from this can also facilitate fleet management. The code for analysing this dataset is accessible online (GitHub) and may stimulate future researchers to explore this dataset further.

Index Terms

Truck driving behaviour, heavy goods vehicles, stability, vehicle characteristics, reliability.

I. Introduction and Related Works

FOR several decades, research has focused on studying driver behaviour to aid the development of policies and infrastructure design and to improve the reliability and safety of Advanced Driver-Assistance Systems (ADAS). Conventionally, driving behaviour has been examined through questionnaires, simulator studies, and controlled on-road experiments (1; 2; 3). However, these methods fail to capture drivers' day-to-day behaviour, and decision-making process (4). With sensors becoming increasingly prevalent, there is an opportunity to understand better the complex interaction between driver, vehicle and the environment under natural conditions, for example, as part of a Field Operational Test (FOT). Next, we will discuss a few prominent FOTs and studies which utilize them to examine driving behaviour and improve the effectiveness of advanced driver-assistance systems.

A. Naturalistic Driving Studies

For over a year, a 100-car naturalistic driving study was conducted in the United States (5). This dataset has been used to investigate various aspects of driver behaviour (6; 7). Johnson et al. studied the impact of change in lane width and radius of curvature on the relationship between lateral velocity and distance (6). This research aids in the development and acceptance of Lane Departure Warning (LDW) systems.

Montgomery et al. sought to identify differences in braking behaviour between age and gender groups (8). Distinct statistical differences were observed between genders and people over and under 30, with men (between 18-20 years) having the lowest average time to collision. This can be used to tailor headway warnings according to demographic.

The Strategic Highway Research Program Naturalistic Driving Study (SHRP-NDS) is a large-scale follow-up to the previous study (9). Data was collected from over 3,400 drivers for over three years. Researchers have extensively used this dataset to study the effect of environment, driver behaviour and vehicle characteristics. For example, Ahmed et al. studied the impact of heavy rain on speed and headway behaviour (10). The researchers concluded that drivers reduced their speed by 5 km/h in rainy weather conditions to compensate for the negative impact of rain on the driving task.

Similar studies have also been conducted in Europe and Australia, i.e. UDRIVE (eUropean Naturalistic Driving and Riding for Infrastructure and Vehicle Safety and Environment) and the Australian 400-car naturalistic driving study (11; 12). Guyonvarch et al. analysed driving behaviour using the UDRIVE dataset and developed a driving style indicator to characterise risky or safe situations (13).

For a majority of these naturalistic driving studies, the focus has been on private cars. However, it is also critical to explore research that focuses on trucks. According to the National Highway Traffic Safety Administration, in the United States (US), large trucks accounted for 8.9 per cent of the vehicles in fatal crashes in 2020 (14). Despite significant safety technology advancements, crash rates, specifically for trucks, have increased over the past decade (15). In 2019, 14% of all road fatalities in the European Union (EU) involved heavy goods vehicles (16). Besides road safety, driving behaviour also impacts the overall mileage of vehicles and CO_2 emissions. Heavy duty vehicles also contribute 6% of CO_2 emissions in the EU (17). The following section discusses some naturalistic driving truck studies which focus on mitigating the problems discussed.

B. Naturalistic Driving Studies - Truck Specific

Zhou et al. researched dangerous driving behaviour using 11 months of naturalistic truck driving data in Japan (five specific districts) for over 70 drivers (18). Cai et al. collected data from around 32,000 truck drivers in the United States to analyze the correlation between crashes and safety-critical events such as headway and hard brake events (19). So far, this is one of the most extensive naturalistic driving studies among commercial truck drivers. The study concluded that a unit increase in one of the safety critical events per 10,000 miles is directly associated with an 8.4% increase in crashes per mile. From this, we can conclude that understanding truck driver behaviour is critical.

Projects such as Towards Safe Mobility for All: A Data-Driven Approach focus on examining the validity of data-driven feedback provided to drivers (20). Data-driven quantification or modelling of truck driver behaviour has recently gained momentum to aid fleet management and safer driving. This is being tackled by companies such as NEXTdriver, which use in-vehicle sensor data, not to automate the driving task but to give feedback/coaching to professional drivers (21). This can also assist insurance companies such as ABW in modelling pay-as-you-drive (PAYD) insurance schemes (22; 23). In the next section, we will consider a few papers that delve deeper into data-driven quantification of truck driving behaviour. Followed by examining the gap that exists in the literature.

C. Quantifying Truck Driver Behaviour

While quantifying driver behaviour, researchers account for the variability in internal (e.g. innate preferences and demographic factors) and external (e.g. traffic density, weather and road type) factors. This problem has been tackled in two ways.

Firstly, scoring or rating a driver's behaviour (on a continuous scale) and second, classifying drivers by assigning them different categories (discrete classes), formally known as Driver Behaviour Profiling (DBP) (24).

Fig 1 gives us an overview of a framework by Lu et al. that can be employed for DBP (25). Researchers extract statistical features from the raw signals. Then these signals are used as input to the model. The output of this model can either be a score or categorize drivers into classes. In order to account for the variability, internal and external features are also used as input to the model. For example, Ozgul et al. developed a model profile driver where they also took traffic flow into consideration (26). The features used to cluster drivers based on traffic flow include the number of stops, % of idling time and corresponding velocity statistics for a trip. Similarly, we can consider the weather and time of the day while developing profiles.

Valuable information can be extracted by focusing on patterns in driver behaviour and trying to profile drivers into different groups (e.g., aggressive vs non-aggressive driving). The application determines the kind of behaviours that we may investigate for different drivers. For example, the focus of insurance companies could be safety, and coaching platforms could be eco-driving.

Supervised and unsupervised models have been developed for driver behaviour profiling. Figueredo et al. collected telematics data from over 2,400 truck drivers across the United Kingdom (27). Drivers were then clustered into eight different profiles. Similarly, classification models have also been developed based on scoring algorithms (25).

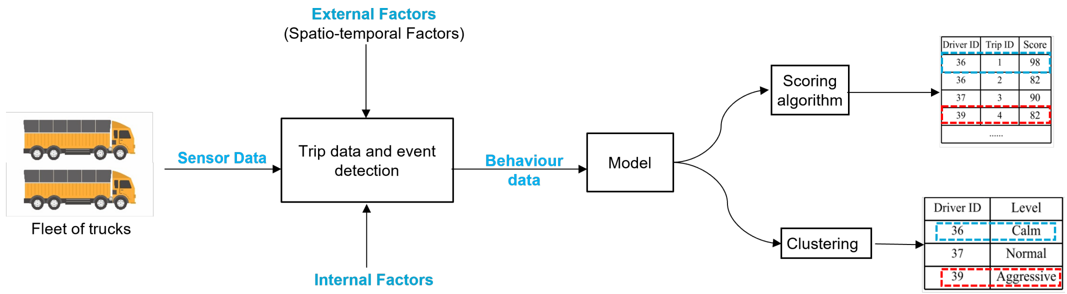

Fig. 1: Framework for Driver Behaviour Profiling (25)

Previous research used data-driven models for profiling truck drivers and generally focused on predicting a driver's behaviour (28; 29). However, we need to examine the motivation behind including or excluding certain driving features while developing such models to predict a driver's behaviour (for example, in the case of driver profiling). Also, for providing feedback/coaching to drivers, it is crucial to understand which metrics tend to be reliable (i.e. metrics which are useful for prediction) and how they are affected by **time** and **location**. Hereafter, reliable metrics will be termed as stable.

In the following section, we'll explore what a stable metric is in the context of driving behaviour with the help of an example.

D. Stability

An illustration of the relevance of stability in driving assessment is provided in Figure 2. It depicts drivers (or vehicles) at different moments in time. Consider any two time instances t_1 and t_2. Drivers who drive fast at t_1 also tend to drive fast at t_2, half an hour later. Therefore in this example, speed is stable with respect to time.

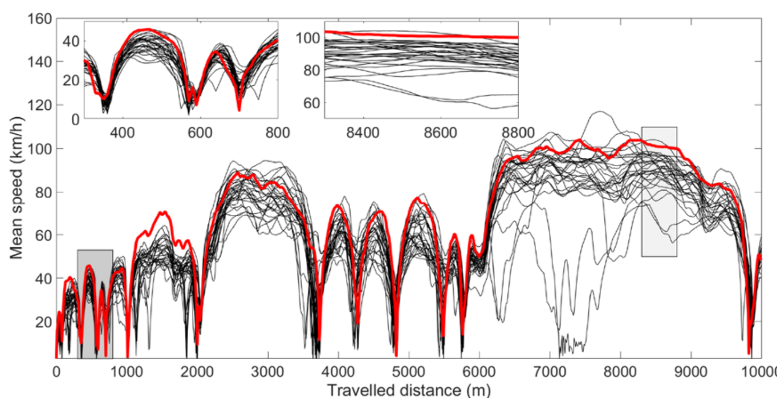

Fig. 2: Vehicle speed for 30 drivers who drove the same route through Delft, including intersections and a highway segment. The thick red line represents the driver with the fastest mean speed among the 30 drivers (from (30))

If a measure is not stable at all (drivers who drive fast at t_1 are not the same as drivers who drive fast at t_2), then the behavioural metric is not stable.

Figure 2 also illustrates that stability should ideally be assessed in a location-specific manner. For example, at intersections, speeds are low for most drivers. If just using time differences, then these effects will dilute; hence one may better assess drivers at a given location (p_1, p_2, ... p_n) in addition to just the given time. In the next section, a few studies that focus on the location-specific analysis of trucks have been discussed.

E. Location-Specific Analysis

Several researchers have identified accident hotspots using geospatial information (31; 32; 33). Desai et al. researched the relationship between harsh braking events and crashes using commercially available data in Indiana around work zones (31). A positive correlation was established between the two, and it was recommended that harsh braking events could be used as surrogates for safety in emerging work zone locations. Researchers have also developed systems to notify drivers when they are near a work zone location to mitigate risk (33).

Walnum and Simonsen studied how fuel consumption is affected by internal and external conditions for HGVs in Norway (34). The data was procured from Dynafleet (Volvo). The analysis considered driver behaviour, type of road (mountain pass or motorway), weather, and vehicle characteristics. Through multivariate regression analysis, it was concluded that vehicle characteristics and infrastructure have a more significant impact than driver-influenced variables while driving on narrow mountainous roads. Since the location-specific analysis of (truck) driving data is a relatively recent topic of scientific inquiry, the current study tries to extend it. **The present thesis takes an alternative stance and tries to examine naturalistic truck driving behaviour, focusing on the stability of driving behaviour in time and location.** The following associated research question is proposed:

- **What metrics associated with truck driving are stable across space and time?**

Driving metrics that might be predictive (or reliable) differ between urban and motorways as types of driving tasks are different and drivers are exposed to varying environments (35; 36).

To answer the proposed research question, we will use the Anti-ongevalsystemen (AOS) dataset (37). This dataset was part of a comprehensive Dutch Field Operational Test. It generated extensive behavioural data from over 1,700 trucks across Europe. It provides us with a unique insight into the behaviour of professional drivers, which is rare at such a large scale. With recent advancements in sensor technology (as part of automation systems and ADAS) and access to such large datasets, we have an opportunity to better understand driver behaviour. The unique features and insights of this dataset can be exploited to understand patterns in driving, as previously discussed, to develop driver profiles to aid driver coaching, fleet management and insurance companies (PAYD).

AOS study is well known among field-operational test (FOT) experts, but there have been no scientific publications about its results. The only analysis that has been done discusses several measures which were computed, such as mean speed, the standard deviation of speed, mean headway, etc. [Note: not available online] (38). In contrast, the SHRP2 dataset has been extensively studied and analysed. Close to 400 published papers that use SHRP2 dataset can be found (39). AOS FOT has escaped scientific scrutiny. We managed to gain access to it, and a substantial effort was devoted to reading and understanding this enormous dataset.

Since this dataset has not been used for research before, a brief description has been given in Section II, followed by the Methodology used to analyse spatial and temporal stability in Section III. After that, the results are presented in Section IV, followed by a discussion examining the influence of vehicle characteristics in Section V. Finally, a conclusion, followed by limitations and future work, is presented in Section VI. An Appendix supplements the work (Section A).

II. DATASET

A. Description

The dataset used in this study was recorded as part of a large-scale field operational test in Europe. The FOT sought to understand better how accident prevention systems can contribute to traffic safety and traffic flow on the Dutch road network. Data was collected from September 2008 – May 2009 from over 1,700 trucks in collaboration with 124 companies (40).

The data was recorded using **Clear Box** (composed of an accelerator sensor and a GPS device mounted in the vehicle) and **Mobileye [2007]** (consisting of a vision sensor and a display & control unit). The recorded data was uploaded to the OCTO server database. The geographical information such as road type, area and speed restriction was procured from the **TeleAtlas database**. The dataset consists of two types of data - Orderly Use data **(Trip data)**, and Main AOS Data **(AOS Events data)**. An overview of the data collected can be found in Figure 3.

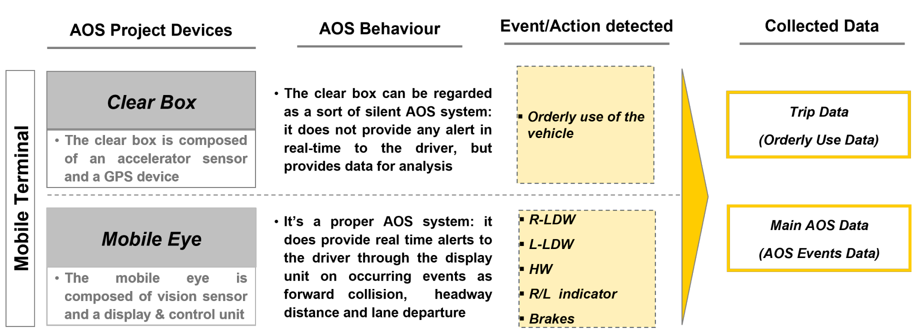

Fig. 3: Overview of data collection - AOS dataset (41)

- **Trip Data** (267×10^6 data points): This data was recorded using a clear box, and no real-time alert was provided to the driver. Features were recorded every 2 km, including number plate, latitude, longitude, point speed, timestamp, distance travelled, and time elapsed. Furthermore, information is provided by the speed limit of the road, the number of lanes, area, road type, class and form (TeleAtlas database). Figure 4 depicts trips recorded across Europe. The data was recorded for over 96×10^6 kilometres from 1,727 trucks.

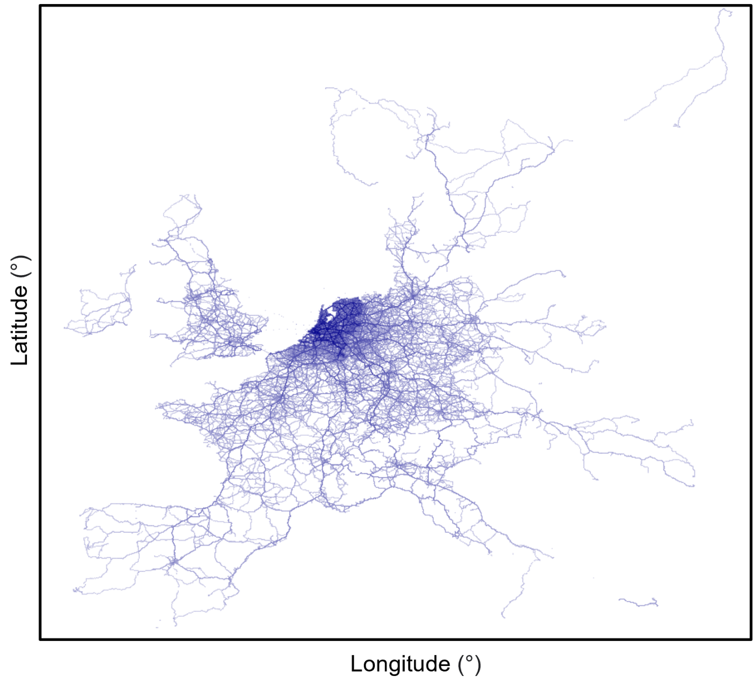

Fig. 4: Trips recorded across Europe - Trip Detail file: 267×10^6 data points.
It can be observed that the density of trips recorded in the Netherlands is higher

- **AOS Events Data** (124×10^6 data points): Recorded using Mobile Eye. Trip data described in the previous section A is recorded at a fixed frequency, whereas event data was only logged when an event triggered, which led to a warning (the driver receives a real-time alert - except braking events).

  All the events recorded include further information about the type of event/action, time stamp and GPS data (latitude, longitude, and speed) of the moment when the event occurred. The recording of events is also conditional to the mobile terminal's presence inside the Netherlands territory. Therefore, events and actions occurring outside the territory of the Netherlands are disregarded. The events have been summarized in Table I and below.

  – **Headway Warnings: (Level I/II/III-HW)**: Headway Warnings are triggered based on time to collision. The driver is alerted using audio and visual feedback. The warning is split into three different levels. The seriousness of the warning is indicated using coloured pictograms. This is based on headway time (HW): $1.1 \leq HW < 2.5$ is Level (I) - Green, $0.7 < HW < 1.1$ is Level (II) - Yellow, and $HW < 0.7$ is Level (III) - Red.

Fig. 5: Visual display of the Headway Monitoring and Warning System (HW) installed in vehicles
Three levels (Level I, II, and III from left to right) [MobileEye, 2007] (42)

  – **Right and Left Lane Departure Warnings (R/L-LDW)**: Lane Departure Warnings (LDW) occur when the vehicle has inadvertently strayed out of the lane without using an indicator. In our case, LDW was determined using a camera and a processing unit. It was triggered when the speed was above 55km/h time to line crossing fell below 0.5 seconds. The driver was alerted using a beeping sound.

Fig. 6: Visual display of the Lane Departure Warning (LDW) when a warning is on the left side [MobilEye, 2007] (42)

- **Braking Events**: A braking event is registered when a driver pushes the brake pedal. No alert was provided to the driver.

TABLE I: Summary of recorded AOS events

AOS Event	Source
Right/Left Lane Departure Warning (R/L-LDW)	MobilEye CAN Message 0x700
Headway Waning (Level I/II/III - HW)	MobilEye CAN Message 0x700
Brakes ON (Braking Event)	MobilEye CAN Message 0x765

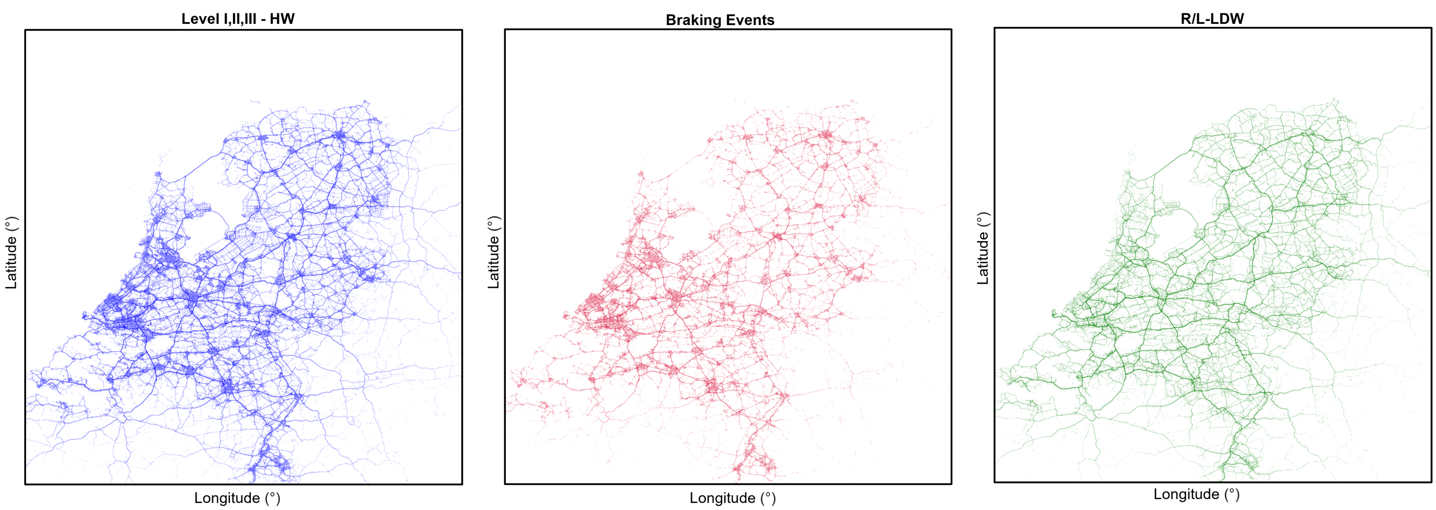

Fig. 7: Visualization of AOS Summary Events recorded (Left to Right: Headway Warnings, Braking Events and Lane Departure Warnings) AOS Summary File - HW: 41.4×10^6, Braking Events: 3.5×10^6, and LDW: 9.75×10^6 data points

Fig. 7 illustrates the different AOS events recorded in the Netherlands. It can be inferred from the figure that certain areas, for example, the Randstad region, are more prone to headway warnings. This is likely due to the higher traffic density in these areas. A similar pattern is observed in the case of braking events, which might indicate a correlation between the headway warnings and braking events.

Over 124×10^6 events were recorded in total. This includes 41×10^6 Headway Warnings (Level I/II/II-HW), 3.5×10^6 braking events and over 9.7×10^6 Right and Left Lane Departure Warnings (R/L-LDW).

The method described in Section III-A will be utilized to analyze spatial stability for different features in urban areas and motorways.

III. Methodology

The objective of this thesis is to determine metrics which are stable across space and time. In order to do this, we use number plates, hereafter referred to as a ‘truck or vehicle’, to distinguish between vehicles (and not drivers because driver information was not logged during the FOT). We have conducted analyses to examine **temporal** and **spatial** stability. Both analyses have been conducted in isolation. The spatial stability analysis has further been split on the basis of the environment into urban areas and motorways. The motivation behind this split has been discussed below.

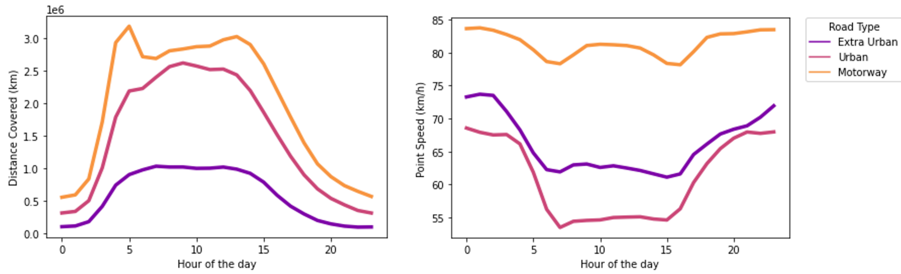

Fig. 8: Distance covered by trucks and its mean speed on different road types (Trip Detail file: 267×10^6 data points)

Fig. 8 describes the distance travelled and point speed at any given hour for different road types. It can be inferred that the distance traversed by these trucks reaches its peak between 6:00-16:00, and, at the same time, there is a sharp decrease in speed on urban and extra-urban roads. This can be attributed to an increase in traffic density. Thus, it can be postulated that the effect of traffic is more pronounced in and around urban areas than on the motorway. In other words, it can be asserted that driving conditions have an effect on driving behaviour.

Fig. 8 also depicts that a high proportion of data has been recorded on motorways and urban roads. Therefore, we will encounter a combination of short-haul and long-haul trips.

The framework to explore spatial and temporal stability will be discussed in the following sections.

A. Spatial Stability

The described method will be utilized to analyze spatial stability for different features in urban areas and motorways. The explanation below is an example of one truck on a motorway. The same can be extrapolated for other conditions.

Assume that the vehicles set (N) contains n trucks, which can be represented as :

$$N = \{n_1, n_2, n_3 \ldots n_i\} \tag{1}$$

We can define a feature matrix (X_{n_i}) for truck n_i which constitutes of $(X_{n_i})_{Trip}$ (**Trip data**) and $(X_{n_i})_{Event}$ (**AOS data**) based on the type of feature recorded (refer Sec. II). Both the files are initially filtered based on road Type, i.e. Urban and Motorways.

Assume a trip from point A to point B on a motorway. The corresponding features for each n_i have been recorded over some duration $T = \{t_1, t_2, t_3 \ldots t_d\}$. The total duration can vary for every truck included in the analysis. These matrices are further explained below.

The following features, which make up $(X_{n_i})_{Trip}$ will be used for the stability analysis, and a description of these follows in Table II.

$$(X_{n_i})_{Trip} = \begin{bmatrix} (x_{t_1})_1 & (x_{t_1})_2 & \cdots & (x_{t_1})_v \\ (x_{t_2})_1 & (x_{t_2})_2 & \cdots & (x_{t_2})_v \\ . & . & \cdots & . \\ . & . & \cdots & . \\ . & . & \cdots & . \\ (x_{t_d})_1 & (x_{t_d})_2 & \cdots & (x_{t_d})_v \end{bmatrix} \tag{2}$$

TABLE II: Features used for analysis of stability of n_i in $(X_{n_i})_{Trip}$ (Extracted from Trip Detail file) (43)

Feature	Description
Time Stamp	Time at which data is recorded (t_i)
Point Speed	(km/h): Speed recorded at time (t_i)
Meters Travelled	m: Meters travelled from previous point
Road Type	Type of road : Urban, and Motorway (according to TeleAtlas)
Latitude	GPS coordinates
Longitude	GPS coordinates
Location	Location of Truck (Standard Name for Order 8 area)
Numberplate	Number plate of truck

Similarly we can define $(X_{n_i})_{Event}$ followed by the description of the corresponding features (Table III):

$$(X_{n_i})_{Event} = \begin{bmatrix} \left(x_{t'_1}\right)_{1'} & \left(x_{t'_1}\right)_{2'} & \cdots & \left(x_{t'_1}\right)_{w} \\ \left(x_{t'_2}\right)_{1'} & \left(x_{t'_2}\right)_{2'} & \cdots & \left(x_{t'_2}\right)_{w} \\ . & . & \cdots & . \\ . & . & \cdots & . \\ . & . & \cdots & . \\ \left(x_{t'_d}\right)_{1'} & \left(x'_{t_d}\right)_{2'} & \cdots & \left(x_{t'_d}\right)_{w} \end{bmatrix} \tag{3}$$

TABLE III: Features used for analysis of stability of n_i in $(X_{n_i})_{Event}$ (Extracted from AOS Summary file) (43)

Feature	Description
Time Stamp (AOS)	Time at which an event is recorded (t'_i)
Road Type	Type of road : Urban, and Motorway (according to TeleAtlas)
Event Type	Braking event, Level I/II/III - HW, R/L - LDW and R/L-I
Latitude	GPS coordinates
Longitude	GPS coordinates
Location	Location of Truck (Standard Name for Order 8 area)
Numberplate	Number plate of truck

After the initial feature matrices have been extracted, specific data points are excluded from the analysis. This includes extreme outliers (which result from abnormal sensor recordings, for example, a point speed of 450 km/h) and when the vehicle is stationary (i.e. when the recorded point speed is 0 km/h). We can assume m rows are dropped during the filtering.

After this step, we process $(X_{n_i})_{Trip}$ to compute mean speed (Eq. 4) and total distance covered (Eq. 5).

$$(x_{n_i})_{\bar{s}} = \left(\frac{1}{d-m}\right) \sum_{d=1}^{d-m} (X_{n_i})_{dv} \tag{4}$$

$$(x_{n_i})_{SumMetersTravelled} = \sum_{d=1}^{d-m} (X_{n_i})_{dv} \tag{5}$$

$(X_{n_i})_{Event}$ undergoes similar pre-processing as described above. We can assume that p rows are dropped in the process. Then the total number of events per n_i for every type of event is calculated. These features are then normalized using Eq. 5. These events are normalized (with respect to the total distance covered) to ensure the vehicles can later be compared on the same scale. For example, normalized (Level I/II/III-HW) can be calculated as follows:

$$(x_{n_i})_{NormHW} = \frac{\sum_{d=1}^{d-p} (x_{n_i})_{dw}}{(x_{n_i})_{SumMetersTravelled}} \tag{6}$$

The features from both matrices are then processed and merged to create a new feature vector which can then be used for the analysis.

The merged vector for a particular number plate n_i can be represented by:

$$x_{n_i} = [x_1, x_2, x_3, x_4, x_5, x_6, x_7] \tag{7}$$

The vector defined in Eq. 7 comprises different features, i.e. mean point speed, normalized Level I/II/III-HW, R/L-LDW and braking events. This is followed by stacking features obtained from different trucks together.

Fig 9 visualizes the resultant matrix formed after processing and merging $(X_{n_i})_{Trip}$ and $(X_{n_i})_{Event}$. Here the matrix constitutes the feature vectors of 20 trucks that drove on the same highway from location A to B. For example, we can infer from the figure that number plate 9 has higher headway warning values and braking events when compared to other vehicles.

Numberplate	Mean Point Speed	Normalized Braking Events	Normalized Level I-HW	Normalized Level II-HW	Normalized Level III-HW	Normalized L-LDW	Normalized R-LDW
1	42.366673	0.000000	0.797515	0.086202	0.156928	0.025110	0.078954
2	43.267148	0.000000	0.308503	0.187483	0.279201	0.003502	0.003502
3	37.643894	0.435184	0.719422	0.108948	0.155869	0.025121	0.028274
4	49.218021	0.000000	0.663736	0.096230	0.136019	0.038915	0.116598
5	44.722320	0.000000	0.752859	0.137605	0.204153	0.021101	0.268176
6	49.013191	0.000000	0.634674	0.047417	0.134421	0.020907	0.199081
7	36.578629	0.146562	0.237387	0.027957	0.040135	0.007942	0.029934
8	42.576985	0.223744	0.611977	0.124808	0.215133	0.022458	0.145028
9	33.204746	0.000000	2.497552	0.321254	0.322550	0.055860	0.217312
10	50.041875	0.000000	0.669335	0.166336	0.143334	0.024283	0.094275
11	43.314655	0.000000	0.722429	0.163439	0.187784	0.070102	0.169307
12	23.631411	0.209340	0.002186	0.000000	0.273362	0.000000	0.000000
13	49.131112	0.165616	0.513025	0.058178	0.075733	0.009973	0.142547
14	37.763663	0.380097	0.637302	0.072053	0.113853	0.030629	0.043495
15	47.520439	0.000000	0.606800	0.108871	0.136177	0.029987	0.199468
16	62.127874	0.000000	0.115246	0.042684	0.013307	0.020086	0.070805
17	50.793247	0.000000	0.442680	0.067356	0.083784	0.008909	0.106879
18	53.450449	0.000000	0.440418	0.043090	0.073536	0.020231	0.167336
19	61.251846	0.000000	0.082378	0.046182	0.025587	0.008659	0.028239
20	32.221343	0.000000	0.000000	0.000000	0.194758	0.000000	0.000000

Fig. 9: Heatmap of merged features from randomly chosen trips for a reference location considered for analysis (Darker colours represent a higher value for a metric compared to a lighter one)

We follow this up by extracting number plate information based on the analysis above. These license plates are matched to extract Trip and AOS data for trips on motorways across the Netherlands. The extracted data is subject to the same method as discussed above. This is followed by correlation analysis to study spatial stability. A flowchart that visualizes the entire process can be found in Fig. 10.

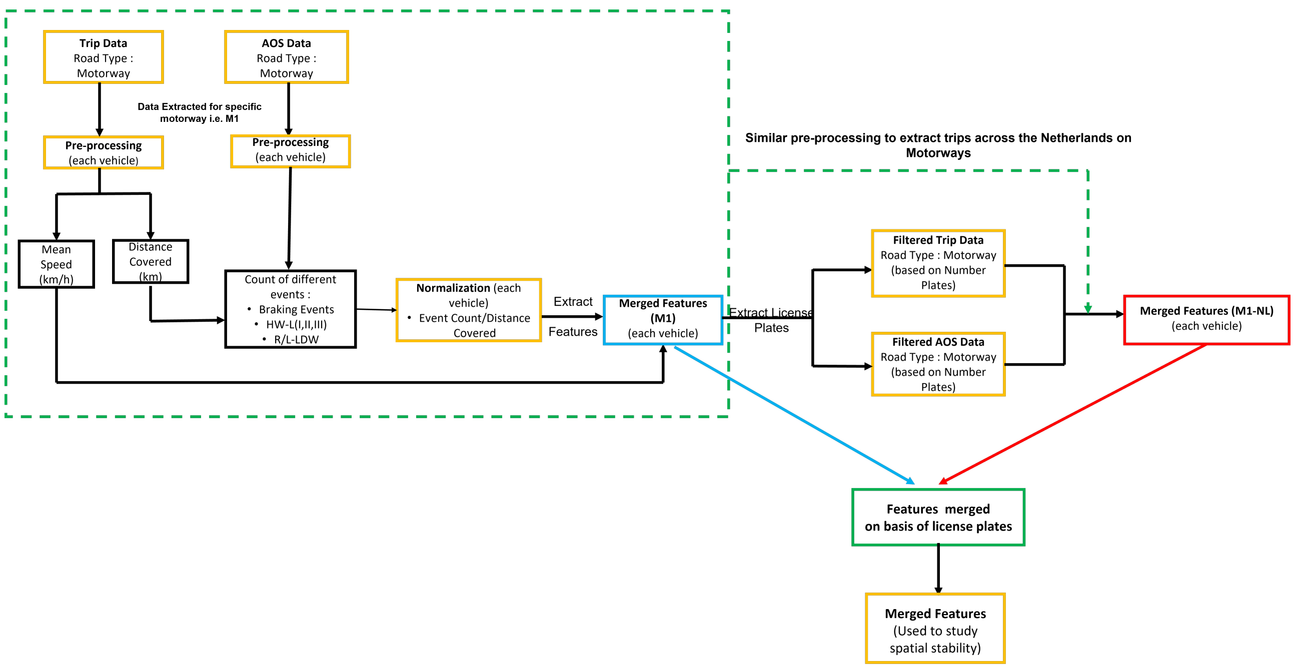

Fig. 10: Illustration of the process of merging Trip and AOS data for a specific motorway and motorways across the NL

B. Temporal Stability

Temporal stability is defined as the stability of a metric over time, i.e. understanding the influence of time on driving behaviour and if a correlation exists over time (t) for a feature. This correlation can be studied over hours, days, and months. Here, we will examine the stability of features over hours. In order to study this, we compute the mean of a metric over different hours for each vehicle. This will be followed by determining the correlation coefficient for that feature over time (hours).

For example, in Fig. 11, we see a clear trend, i.e., a sharper drop in speed during the day and on weekdays because of increased traffic density. This impact is more pronounced in urban areas than on motorways. However, we will not be looking into a combination of spatio-temporal stability. In the following section, we will discuss the findings of stability analysis.

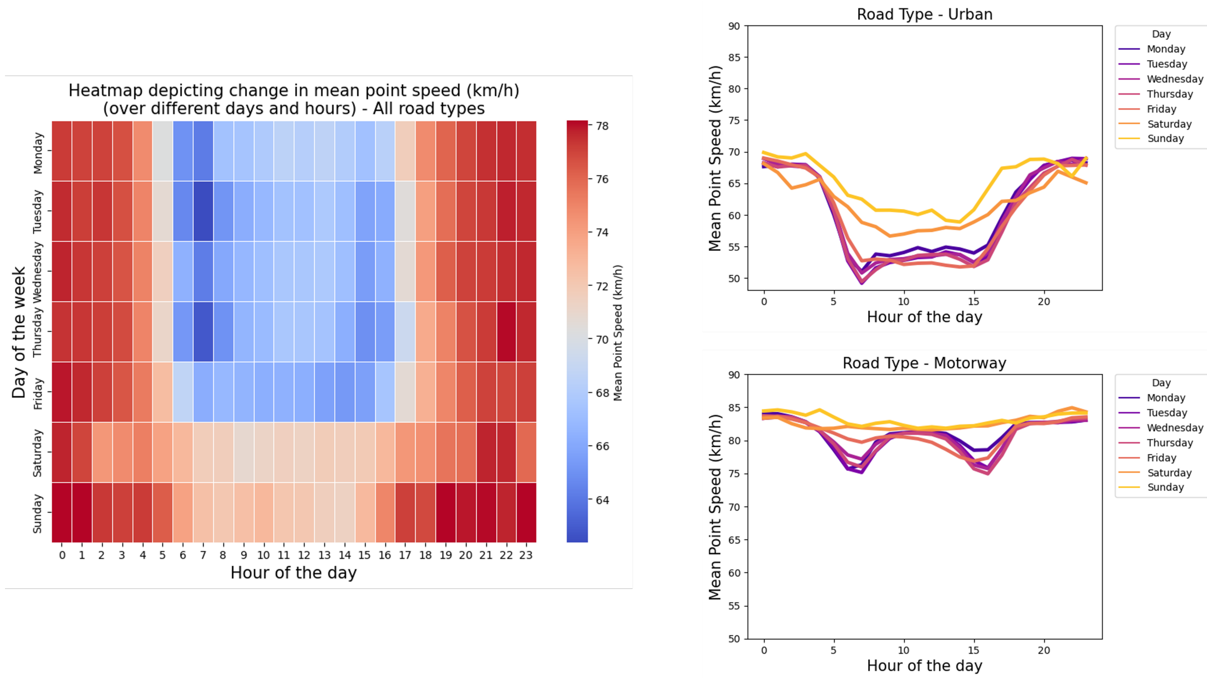

Fig. 11: Mean Point Speed of trucks over time (Trip Detail file: 267×10^6 data points.)

IV. RESULTS

In this section, the results of the experiments have been elucidated. The entire problem of interpreting stability can be treated by how closely a particular metric recorded in a reference location or time is related to the same metric on a broader scale. For example, how the mean point speed of a vehicle in Amsterdam relates to its mean point speed across cities in the Netherlands. This can be examined with the help of correlation analysis. In order to examine if variables are linearly correlated, the Pearson correlation coefficient has been used. Table IV describes the strength of the correlation coefficient and its interpretation.

TABLE IV: General guideline to describe the strength of correlation coefficient

Correlation Coefficient (r)	Description
0.0 to 0.2	Very weak + or no association
0.2 to 0.4	Weak + association
0.4 to 0.6	Moderate + association
0.6 to 0.8	Strong + association
0.8 to 1.0	Very strong + association
1.0	Perfect positive + association

Correlation analysis provides a general overview of the stability of a metric. To further understand how stable different ranges are for a particular metric, clustering analysis has been used. For example, mean point speed, on the whole, might not be stable. However, we might find that a group of trucks with a lower mean point speed than the norm is more stable compared to other groups. K-Means clustering has been used to understand how different groups of trucks interact and if the performance of a particular group tends to be stable across space (44).

The metrics used for analysis are mean point speed, normalized braking events, normalized headway warnings (level-I, II and III) and normalized lane departure warnings (right/left). To recall, metrics have been normalized based on the distance travelled in the reference location.

In the following section, spatial stability (urban areas in Section IV-A and motorways in Section IV-B) has been explored, followed by the analysis of temporal stability in Section IV-C. We have analysed features individually in all the sections.

A. Spatial Stability - Urban Areas

In order to analyze the stability of features in Urban areas, we have considered three cities in the Netherlands, i.e. Amsterdam (753 trucks), Rotterdam (891 trucks), and Zwolle (713 trucks). To make an informed decision about which cities should be included in the analysis, the busiest cities were selected based on the number of data points logged. For the exact number of data points used during the analysis, refer to the Appendix (Section A).

For a vehicle to be included in the analysis, the truck must have driven at least 50 kilometres within that city (any time of the day). The location (or city) a vehicle is being driven in is available in the AOS dataset (according to TeleAtlas, the standard name for the Order 8 area). The data points for the corresponding city in the case of the Netherlands were excluded from the analysis by creating a polygon along the edge points of the city.

Then the corresponding behaviour of that truck is analysed on urban roads across the Netherlands (excluding the corresponding city). An example can be seen in Figure 12.

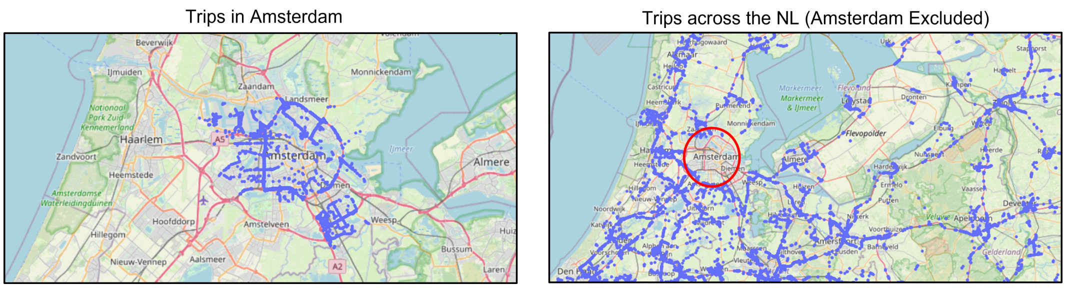

Fig. 12: Trips in Amsterdam (left) and corresponding trips of trucks across the Netherlands (right)
Trip Detail: 10^4 (left) and 10^5 (right) data points

A metric is spatially stable when both areas (i.e. the reference city and all other cities across the NL) have a high correlation. For example, if a truck has a high number of harsh braking events on urban roads in a city, is it likely to follow a similar trend in urban areas across the Netherlands?

The distribution of features can be found in Fig. 13.

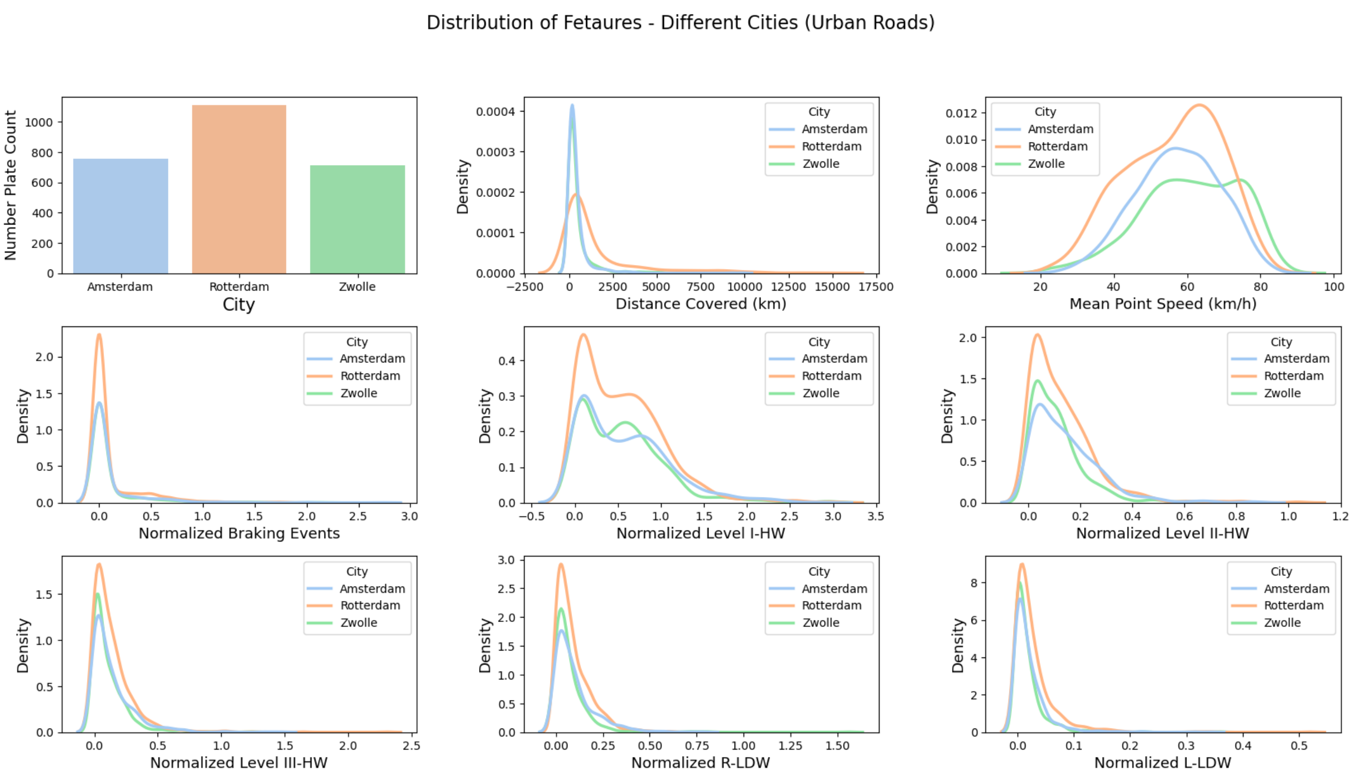

Fig. 13: Distribution of Different Features in Urban Areas (Amsterdam, Rotterdam and Zwolle)

1) Mean Point Speed: As we can see in Fig. 14 the stability of Mean Point Speed is **weak**. The value of the correlation coefficient is higher in Amsterdam (and the rest of the Netherlands) and Zwolle (and the rest of the Netherlands). On the other hand, $r_{rotterdam}$ is very low (Table VI).

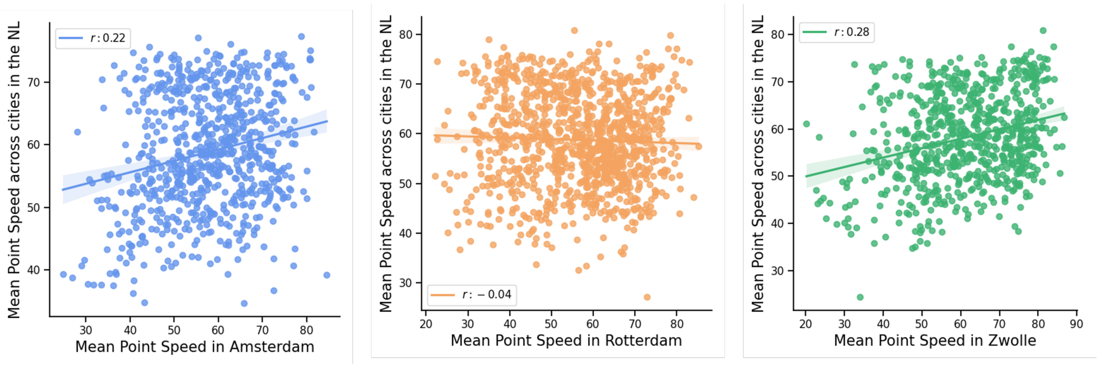

Fig. 14: Scatter Plots of Mean Point Speed in Cities
(Left to Right: Amsterdam, Rotterdam and Zwolle) - Lower Stability (overall)

2) Normalized Braking Events: We notice that braking events generally have a **strong** correlation and association in all the cities. However, the distribution is uneven because many trucks have zero recorded braking events. If we exclude the zero values, the r drops slightly but still has a **strong** correlation (Figure 15).

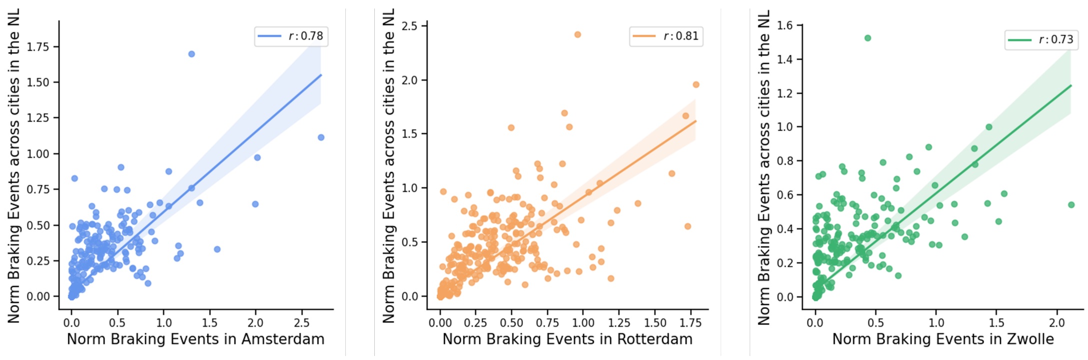

Fig. 15: Scatter Plots of Normalized Braking Events in Cities
(Left to Right: Amsterdam, Rotterdam and Zwolle) - Higher Stability (overall)

3) Normalized Headway Warnings Level-I, II, and III: Figure 16 shows the stability for normalized HW-L(I). To recall, the number of headway warnings for a truck is normalized by the total distance covered in a particular trip. We can infer from the figures that the stability of this metric on average is **moderately strong** for all three levels and across cities.

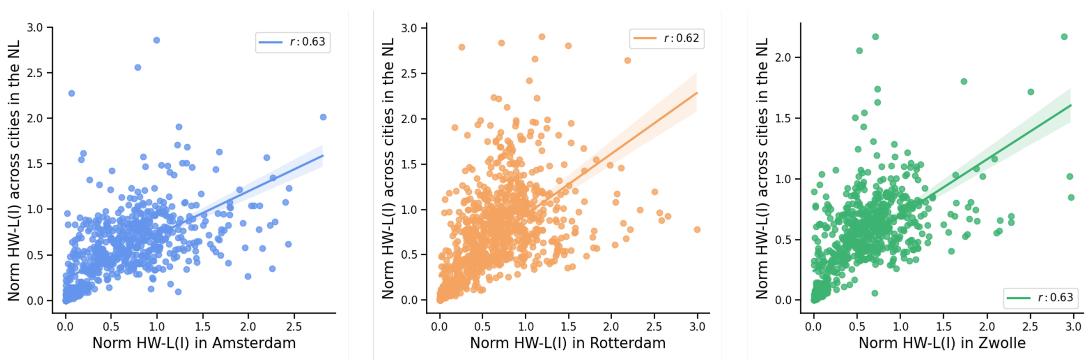

Fig. 16: Scatter Plots of Normalized Headway Warning Level - I in Cities
(Left to Right : Amsterdam, Rotterdam and Zwolle) - Higher Stability (overall)

As discussed earlier, K-Means is used to further analyse stability. K-means splits the individual feature (for example, norm HW-L(I)) into a pre-defined number of clusters. The algorithm randomly initializes centroids. Then each data point is assigned a cluster based on the nearest centroid. Then the mean of each cluster is calculated, and data points are re-clustered based on the new mean values (i.e. updated centroids). To reiterate, a cluster consists of different trucks with similar behaviour determined based on a single feature.

Fig. 17 depicts the clustering of the Norm HW-L(I) feature (in Amsterdam) using K-Means. Over 61% trucks are assigned the same cluster in Amsterdam and for their trips across cities in the Netherlands. Breaking this down further, over 51% of these vehicles belong to cluster 0, i.e. the cluster assigned to vehicles with the lowest value of Norm HW-L(I), followed by cluster 1.

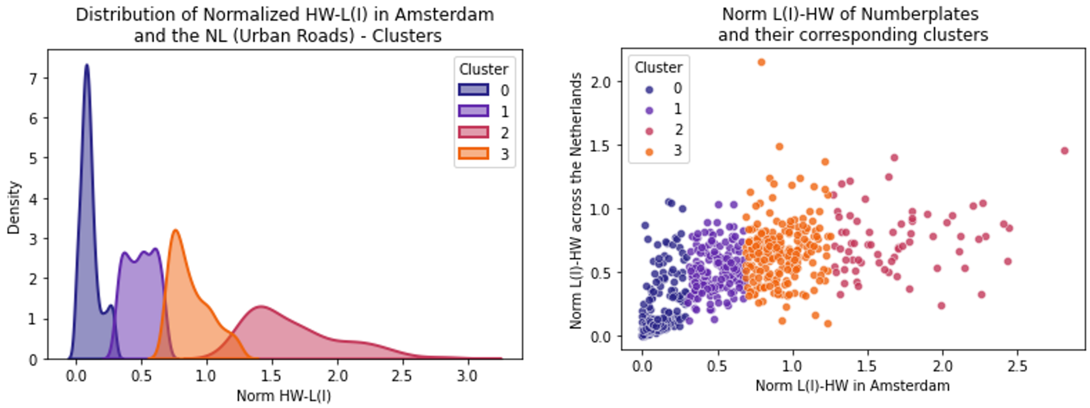

Fig. 17: Splitting Normalized Headway Warning - Level (I) into Clusters - Amsterdam

Fig. 18 also follows a similar trend, and we can see that cluster 0 has median values close to each other (median value in Amsterdam = 0.89 and median value in the NL = 0.66). We can conclude that, on average, trucks with fewer headway warnings per km maintain more headway distance when compared to others during all their trips.

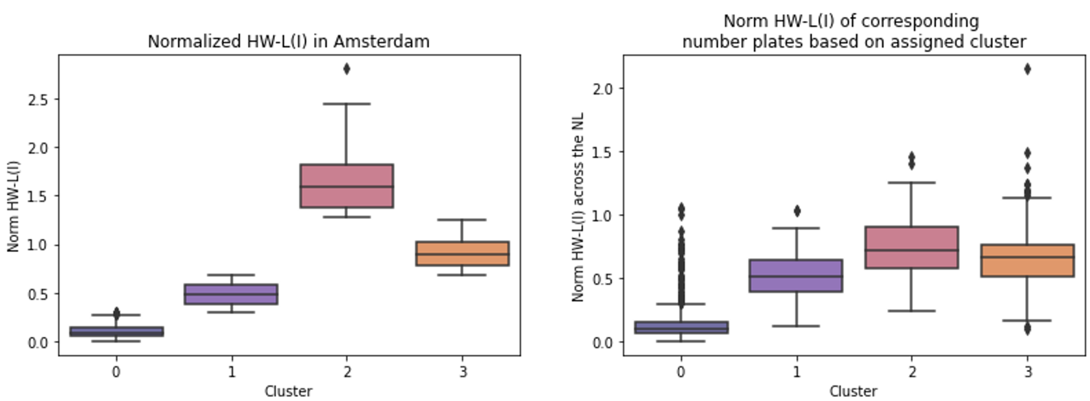

Fig. 18: Box plots representing clusters - Normalized Headway Warning - Level (I)
In Amsterdam (left) and corresponding vehicles in the Netherlands (right)

The same is observed for normalized HW-L (II and III). Each feature has over 53.5% and 59.9% vehicles assigned to the same cluster, respectively. A majority of these are vehicles assigned to either cluster 0 or 1. The breakdown of these clusters can be found in Table V.

TABLE V: Percentage of trucks assigned the same cluster for Norm HW-L(I, II, III)

Cluster	Norm HW-L(I) (%)	Norm HW-L(II) (%)	Norm HW-L(III) (%)
0	51.8	66.6	62.5
1	25.9	26.2	29.3

Based on Table V, we can conclude that maximum stability is displayed by trucks assigned to clusters 0 and 1, which represent trucks with lower metrics values or trucks that tend to maintain higher headway distance. Consequently, we can approximately predict their behaviour on different urban roads. We found comparable results in other cities, which corroborates our claim.

4) Normalized Lane Departure Warnings-Right/Left: We can conclude from Fig. 19 that normalized lane departure warnings have a **strong** correlation, indicating stability. While analyzing the R-LDW recorded in Amsterdam and their corresponding trips across the Netherlands, it is noticed 66% of vehicles are assigned to the same cluster, and the majority of these trucks are from clusters 0 and 1.

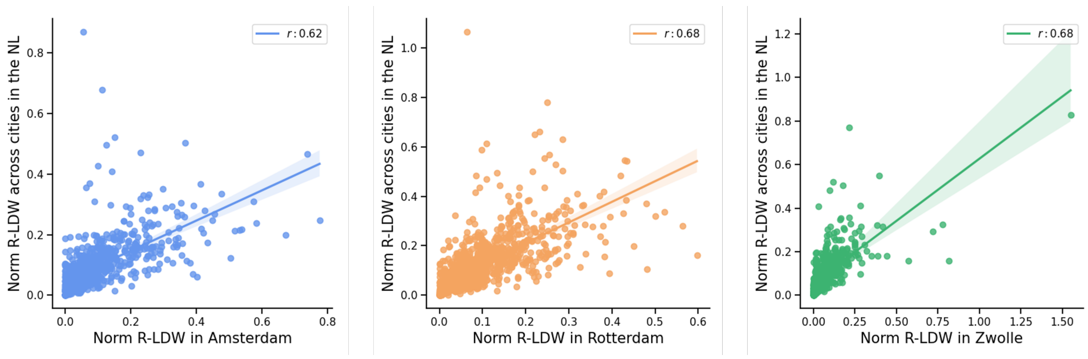

Fig. 19: Scatter Plots of Normalized Right Lane Departure Warnings in Cities (Left to Right: Amsterdam, Rotterdam and Zwolle) - Higher Stability (overall)

We can hypothesize from Fig. 21 that warnings in Amsterdam are predictive of warnings across the Netherlands, specifically for vehicles assigned clusters 0 and 1.
A summary of the stability results for all cities can be found in Table VI.

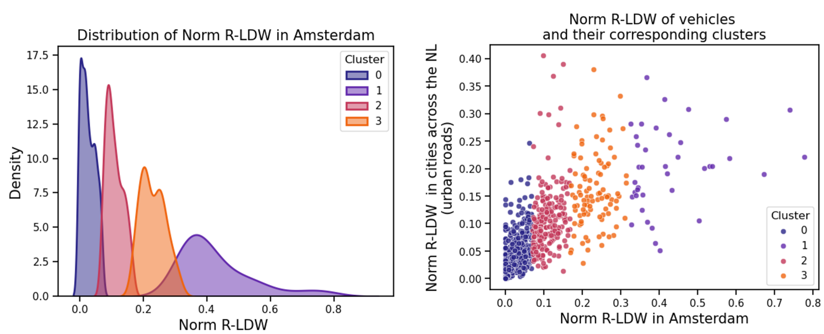

Fig. 20: Splitting Normalized Right Lane Departure Warnings into Clusters (Amsterdam)

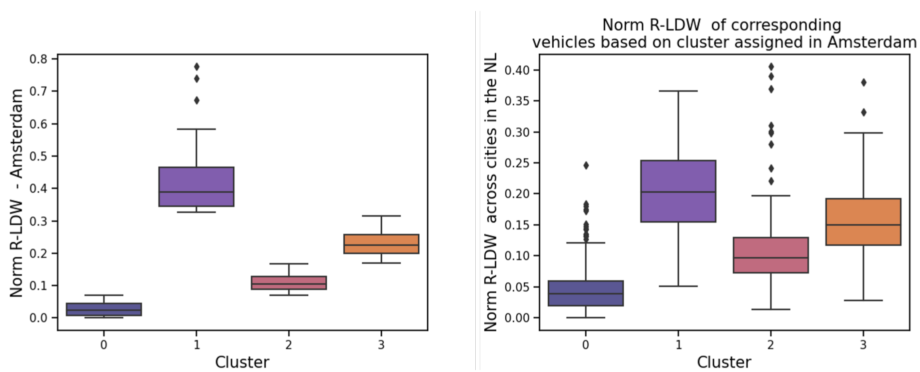

Fig. 21: Box plots representing clusters - Normalized Right Lane Departure Warnings In Amsterdam (left) and corresponding vehicles in the Netherlands (right)

TABLE VI: Summary of correlation coefficient in Urban Areas

	Amstedam and (Cities across the NL)	Rotterdam and (Cities across the NL)	Zwolle and (Cities across the NL)	Mean r-value	Description
	r-value	r-value	r-value		
Mean Point Speed	0.22	-0.04	0.28	0.15	**Very Weak Correlation**
Norm Braking Events	0.78	0.81	0.73	0.77	**Strong Correlation**
Norm HW-L(I)	0.63	0.62	0.63	0.62	
Norm HW-L(II)	0.61	0.71	0.70	0.67	**Moderate-Strong Correlation**
Norm HW-L(III)	0.57	0.50	0.52	0.53	
Norm R-LDW	0.62	0.68	0.68	0.66	**Moderate-Strong Correlation**
Norm L-LDW	0.53	0.65	0.50	0.56	

B. Spatial Stability - Motorways

Here, the stability of different features on Motorways is discussed. For this analysis, we will consider a highway between two locations in the Netherlands and compare it with the behaviour of corresponding trucks across the Netherlands (motorways). This is depicted in Figure 22.

Fig. 22: Trips between Utrecht and Leek (left) and corresponding trips of trucks across the Netherlands (motorways) (right)
Trip Detail: 6×10^4 (left) and 10^5 (right) data points

For example, if the mean point speed has high stability, it would mean that a truck overspeeding on highway A is likely to overspeed on highway B. We will analyze two different highways, i.e. Utrecht to Leek (M1 - 32 trucks) and Utrecht to Eindhoven (M2 - 112 trucks). These routes were chosen based on the density of trucks in a particular city. Secondly, the distance between the two locations should be at least around 100 km so that meaningful conclusions can be drawn about stability.

The distribution of features can be found in Figure 23.

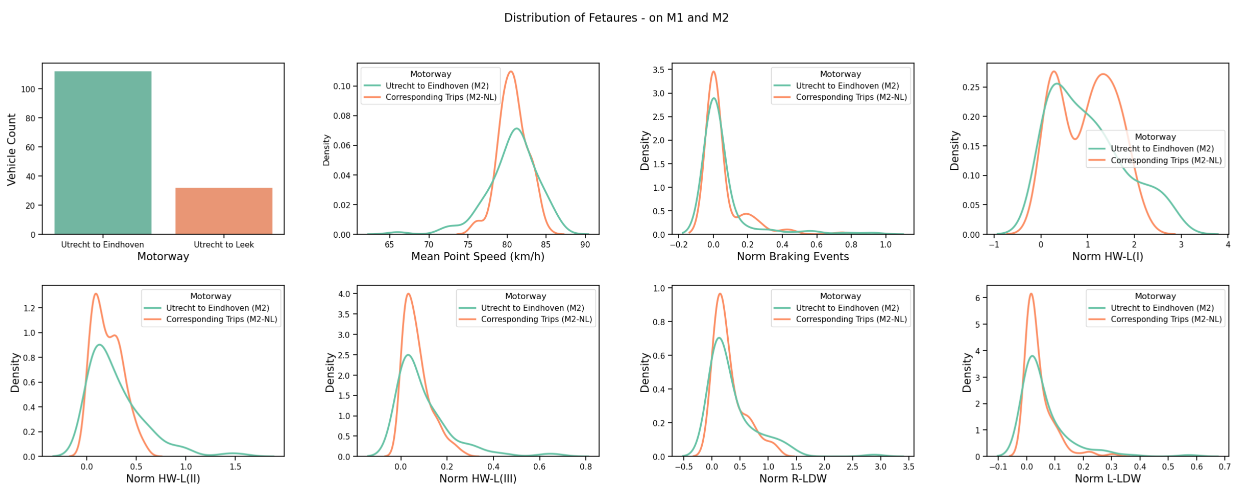

Fig. 23: Distribution of features on motorways
Utrecht-Leek (M1) and Utrecht-Eindhoven (M2)

1) Mean Point Speed: Compared to urban areas, mean point speed has **higher stability** for motorways. The plausible reason behind this could be the comparable environment and traffic conditions the truck encounters. Here it is observed that the range of speed of different trucks is significantly lower than in urban areas. A noteworthy difference exists between vehicles assigned the lowest and highest clusters in both locations. Despite the increase in stability, we cannot accurately predict the speed of a vehicle on any arbitrary highway solely on information from a single motorway (Fig. 24).

2) Normalized Braking Events: We can infer from Fig. 25 that normalized braking events have **higher stability** in the case of M2 than M1. The number of trucks is higher for M2 than for M1; therefore, the correlation is a more accurate representation compared to M1. Secondly, if we further break down the trucks into different clusters, we can notice that different clusters follow a similar trend (especially for clusters 0 and 1). Over 87.5% trucks are assigned the correct cluster.

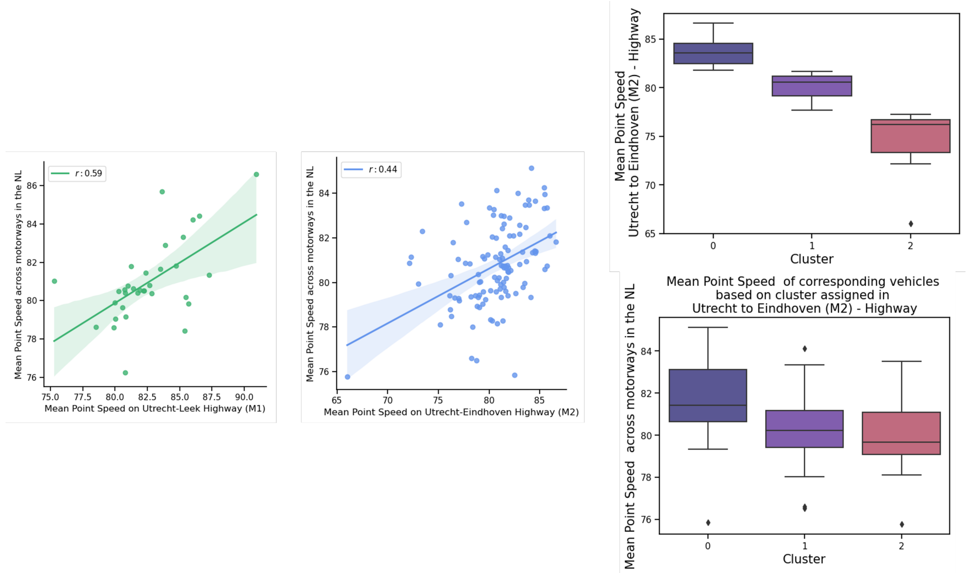

Fig. 24: Scatter Plots of Mean Point Speed on Motorways (Utrecht-Leek and Utrecht-Eindhoven) - Higher Stability (overall)

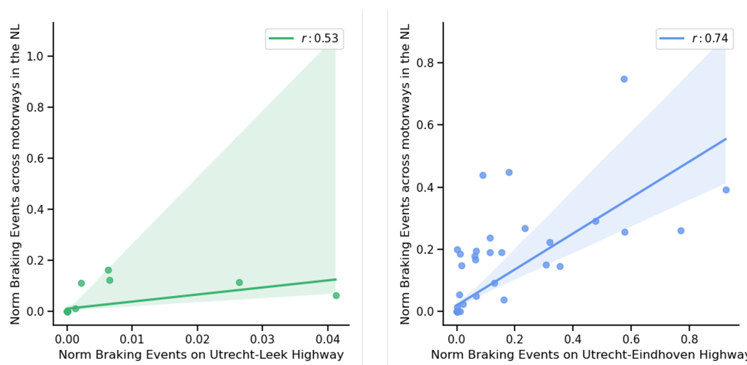

Fig. 25: Scatter Plot of Normalized Braking Events on Motorways
(Utrecht-Leek and Utrecht-Eindhoven) - Higher Stability (overall)

3) Normalized Headway Warnings: Normalized headway warnings have **weak to moderate** correlations across different levels. The stability is lower in the case of motorways than in urban areas. Since the correlation is weak (Fig. 26), we will not notice consistent behaviour among clusters. Of all the trucks analyzed for M2, about 39% of trucks are assigned to the same cluster, which is lower when compared to other metrics (with higher stability).

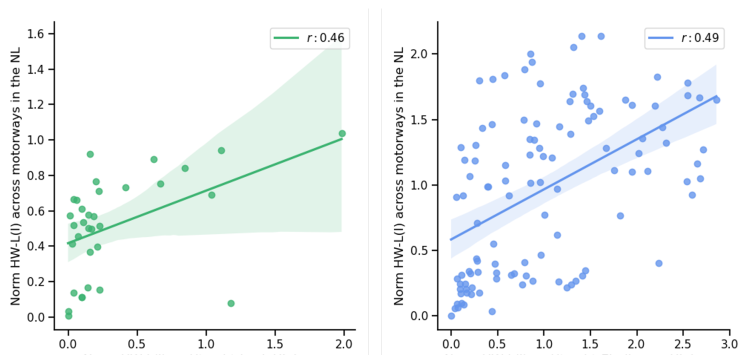

Fig. 26: Scatter Plot of Normalized Headway Warnings Level-I on Motorways
(Left to Right: Utrecht-Leek and Utrecht-Eindhoven) - Lower Stability (overall)

This can be visualized in Fig. 27. When clusters allocated to trucks (M2) are evaluated across the Netherlands, their behaviour does not follow a comparable pattern.

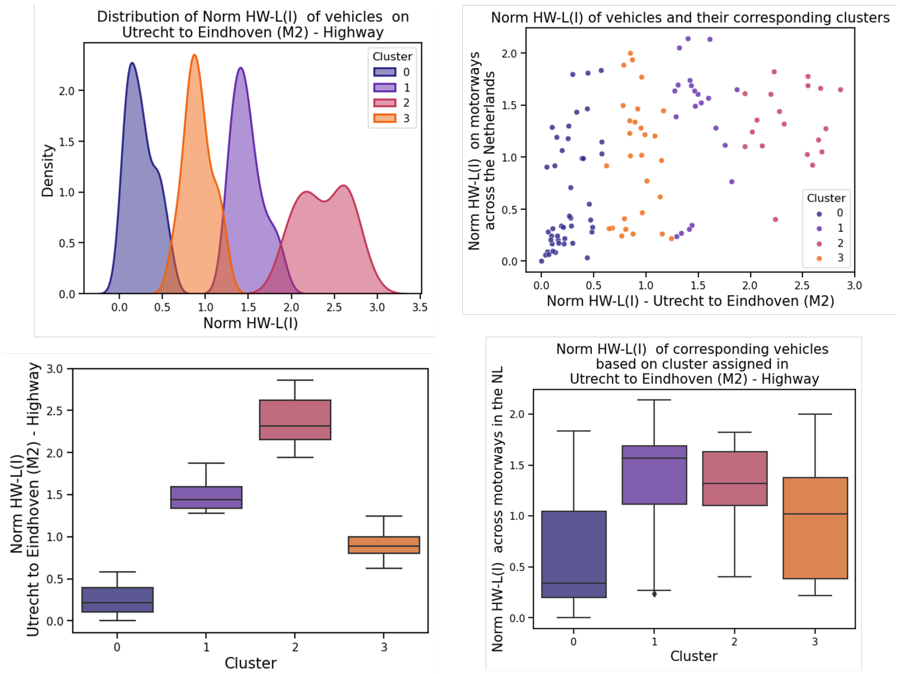

Fig. 27: Distribution and Box plots representing clusters - Normalized Headway Warning (Level-I)
Utrecht Eindhoven and corresponding vehicles in the Netherlands

4) Normalized Lane Departure Warnings - Right/Left: Normalized lane departure warnings exhibit a **strong** correlation consistent with their behaviour in urban areas. Over 59% trucks are assigned the same cluster, out of which 88% belong to cluster 0 (lowest values). This is consistent with the findings for urban regions.

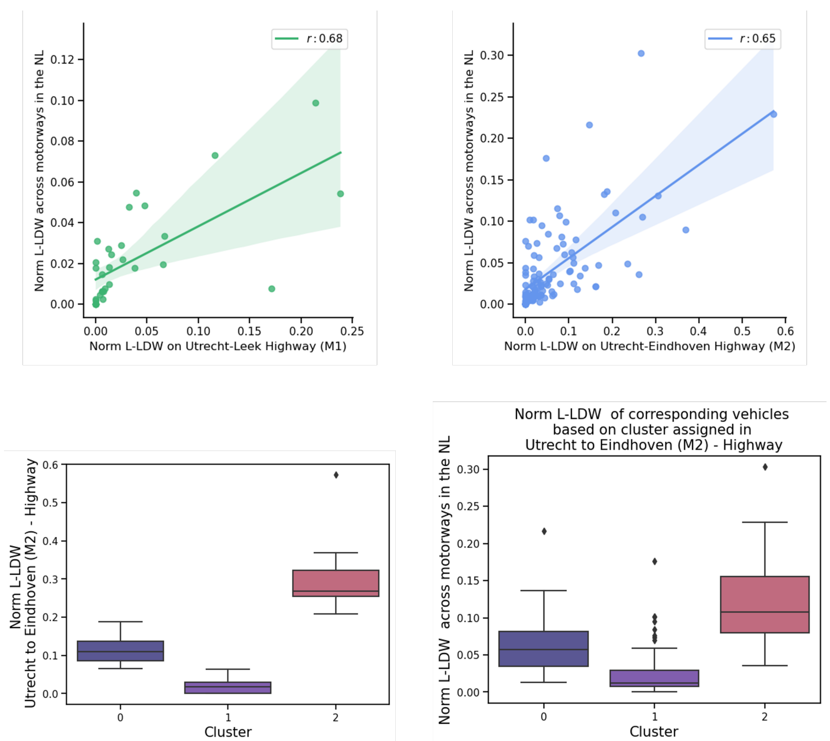

Fig. 28: Scatter Plot of Normalized Lane Departure Warnings- Left: on Motorways (Top) - Higher Stability (overall)
Box plots representing clusters - Normalized Lane Departure Warnings - (Bottom)

TABLE VII: Summary of correlation coefficient on motorways for different features

	Utrecht-Leek (M1) (and motorways across the NL)	Utrecht - Eindhoven (M2) (and motorways across the NL)	Mean (r-value)	Description
Mean Point Speed	0.59	0.44	0.51	Moderate Correlation
Norm Braking Events	0.53	0.74	0.63	Strong Correlation
Norm HW-L(I)	0.46	0.49	0.47	Weak-Moderate Correlation
Norm HW-L(II)	0.26	0.49	0.37	
Norm HW-L(III)	0.44	0.44	0.44	
Norm L-LDW	0.68	0.65	0.66	Strong Correlation
Norm R-LDW	0.81	0.52	0.67	

C. Temporal

For the analysis of temporal stability, the entire dataset has been used. So, these results are based on data from over 1700 trucks. In the following sections, temporal stability for different metrics has been discussed.

1) Mean Point Speed: We can infer from Fig 29 that, contrary to spatial stability, mean point speed exhibits high temporal stability over consecutive hours of the day, i.e. a strong correlation between (t and t+1) hour over the entire day. The stability is more pronounced between 6:00 to 16:00 (purple triangle). This is in line with our observations in Fig. 11, where a drop in speed can be seen in the same period and a similar range of speed.

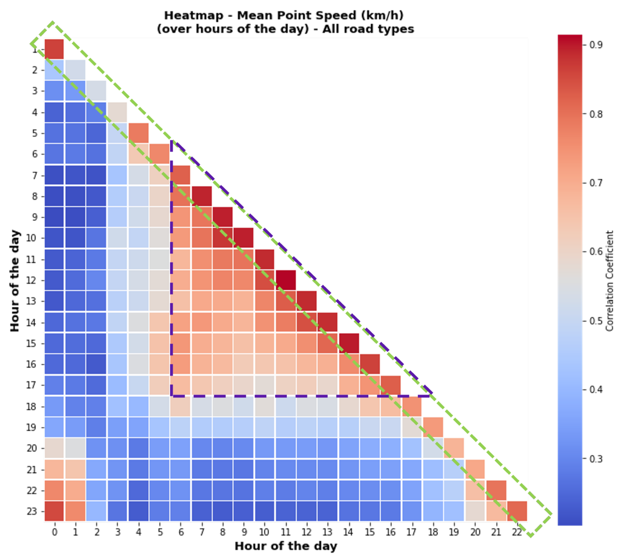

Fig. 29: Heatmap depicting correlation for Mean Point Speed over time (hours of the day)

2) AOS Events: AOS events follow a similar pattern and present strong stability over t and t+1 hours of the day. The number of data points recorded shows a similar pattern, where a maximum number of data points are recorded during the daytime (Figs. 30 and 31).

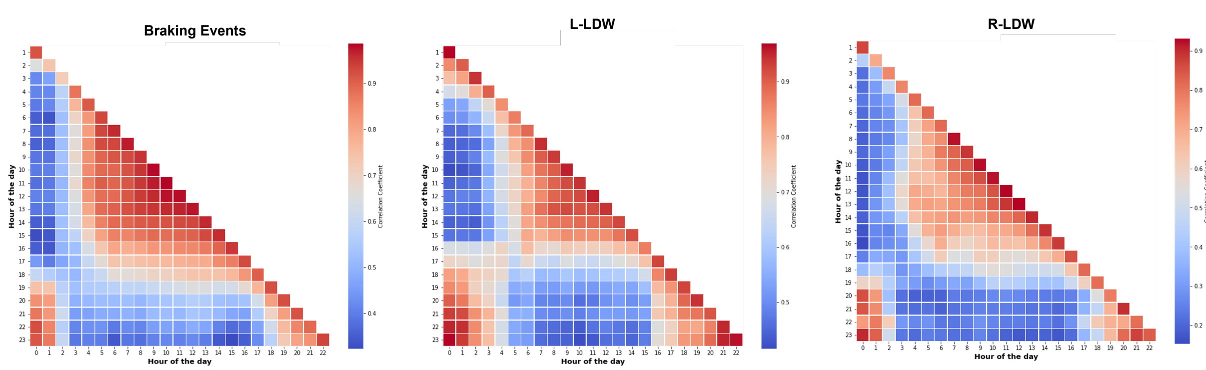

Fig. 30: Heatmap for different types of events recorded (over hours of the day)- All road types
(Left to Right: Braking events, Left Lane Departure Warnings and Right Lane Departure Warnings)

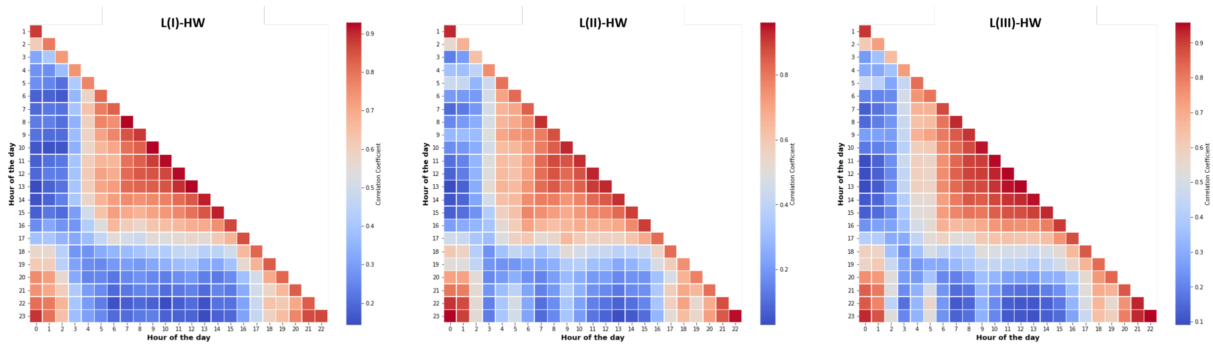

Fig. 31: Heatmap for different types of events recorded (over hours of the day)- All road types
(Left to Right: Normalized Headway Warnings Level-I, II, and III)

The following section will examine the relationship between location-specific features (urban) and vehicle characteristics.

V. Discussion

In the previous section, stability was analysed, and conclusions were drawn. Now, we will focus on understanding the validity of the (location-specific) features by exploring how they correlate with external criteria. This will be done with the help of vehicle characteristics.

Vehicle characteristics were obtained by automatically querying the Dutch vehicle registration database (45). Several third-party websites exist with more complete records. We used the following website: https://www.rdwdata.nl/ to find several missing records. In total, 1300 license plates of the vehicles could be retrieved out of 1727 license plates for the analysis. In this section, we will examine the impact of vehicle characteristics (mass of the vehicle (kg) and engine power (kW)) for different features.

A. Mean Point Speed

The mean point speed does not exhibit stability in either location (urban or motorway). However, it is more stable on motorways than in urban areas. This is likely because of the comparable environment drivers are exposed to on highways in contrast to cities (Figure 8).

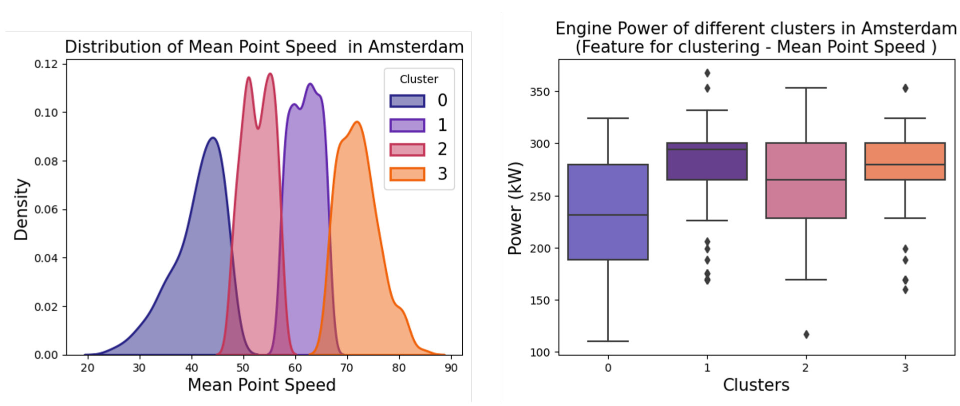

Fig. 32: Distribution of mean point speed and the corresponding variation in Engine Power (kW) for different clusters in Amsterdam

Figure 32 depicts the distribution of mean point speed and the corresponding variation in engine power for different clusters in Amsterdam. Vehicles assigned to cluster 0 have a lower median speed than the rest, and a similar trend can be observed for engine power. Secondly, the mass of vehicles assigned to cluster 0 is consistently higher than that of other clusters (refer Table VIII). Therefore, we can conclude that vehicles with higher mass and lower engine power tend to have lesser mean speed than the norm. This is corroborated by the fact that engine power directly translates to speed (46; 47).

However, for Rotterdam, we observe a deviation from this trend. The engine power of vehicles assigned to cluster 0 is higher than other clusters (Table VIII), possibly because larger trucks are assigned to go to the Rotterdam harbour. Since Rotterdam is a busy city, drivers with larger trucks might be more cautious while driving and maintain a lower speed compared to other trucks from different clusters.

TABLE VIII: Mass of Vehicle and Engine Power for different cities - Mean Point Speed

City	Mass of Vehicle (kg)		Engine Power (kW)	
	Cluster - 0	**Clusters -1, 2, and 3 (Mean)**	**Cluster - 0**	**Clusters -1, 2, and 3 (Mean)**
Amsterdam	7287	7150	188	277
Rotterdam	7372	7173	300	274
Zwolle	7745	7200	243	275

As discussed earlier, mean point speed is not a stable metric. A few plausible causes for the instability of mean speed include environment, traffic density and flow, type of work, or driver behaviour. Intriguingly, vehicles assigned to cluster 0 tend to have a lower speed than the norm across the Netherlands as well. This result cannot be extrapolated for the other clusters. A probable cause is that smaller trucks with lower engine power are more frequently assigned to deliver in cities (e.g., last-mile delivery).

A similar pattern is observed in the case of motorways as well. However, the number of data points is much lower, so it is difficult to draw a generalized conclusion.

B. Headway Warnings and Braking Events

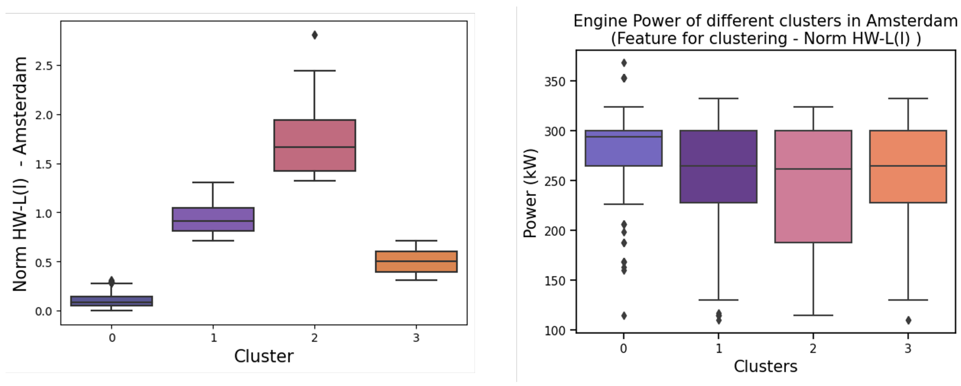

Fig. 33: Box-plots for different clusters in Amsterdam based on Normalized Headway Warnings Level-I followed by the corresponding variation in Engine Power (kW) for different clusters

TABLE IX: Median Mass of Vehicles and Engine Power for different clusters based on Normalized Headway Warning Level-I

Cluster	Mass of Vehicle (kg)	Engine Power (kW)
0 (Lowest Value Cluster)	7280	294
1	7150	265
2	7185	261.5
3	7100	265

Normalized braking events and different levels of headway warnings are more stable when compared to mean point speed. Figure 33 and 34 show the box plots of norm braking events and HW-L(I) and the corresponding variation in engine power for different clusters in Amsterdam. We observe a trend similar for both metrics. In both cases, clusters with the highest median value (braking events and headway warnings) are linked to low engine power. The same holds for all levels of headway warnings. Therefore we can conclude that a high number of braking events and headway warnings are associated with low engine power. The notion that trucks with higher engine capacity are associated with fewer harsh brakes is consistent with the work of other researchers (48).

TABLE X: Median Mass of Vehicles and Engine Power for different clusters based on Normalized Braking Events

Cluster	Mass of Vehicle (kg)	Engine Power (kW)
0 (Lowest Value Cluster)	7155	279
1	6960	258
2	8712	226

On the other hand, the trend w.r.t. mass of vehicles is different for both the metrics (summarised in Tables IX and X). We observe that vehicles assigned the lowest cluster, i.e. 0, based on Normalized Headway Warnings Level-I clusters, have

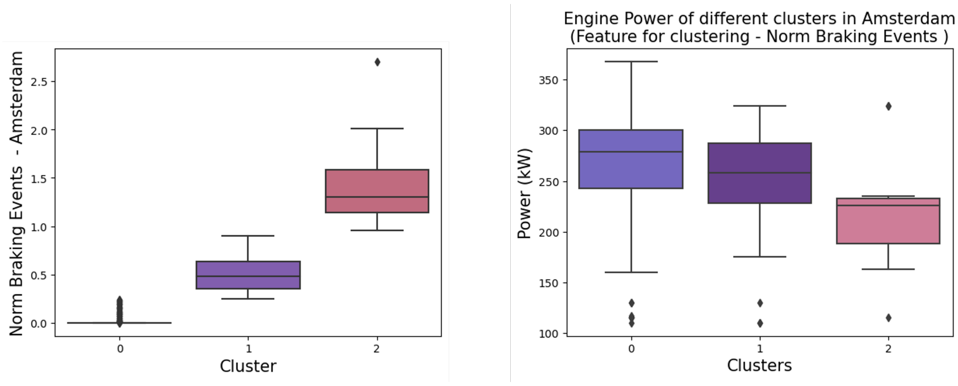

Fig. 34: Box-plots for different clusters in Amsterdam based on Normalized Braking Events (bottom) followed by the corresponding variation in Engine Power (kW) for different clusters

higher mass than other clusters. On the contrary, cluster 0, in the case of Normalized Braking Events, has a lower mass. The observations for the lowest cluster (for different metrics) have been summarised in Table XI.

TABLE XI: Summary of Vehicle and Engine Power Trend for different metrics (Lowest Value Cluster)

Metric	**Description (Cluster 0 - Lowest Value)**
Mean Point Speed	Higher Mass and Lower Power (compared to other clusters)
Normalized Headway Warnings	Higher Mass and Higher Power (compared to other clusters)
Normalized Braking Events	Lower Mass and Higher Power (compared to other clusters)

The plausible cause behind lower headway warnings for trucks with higher power and mass could be that larger trucks are driven more cautiously when compared to smaller trucks to avoid damage to the payload. For clusters with a higher number of normalized braking events, we observe that mass of vehicles is significantly higher than in other clusters. This might also hint that heavier vehicles are more difficult to control.

Since braking events are more stable, we observe a similar pattern for different clusters across the Netherlands (Figure 35).

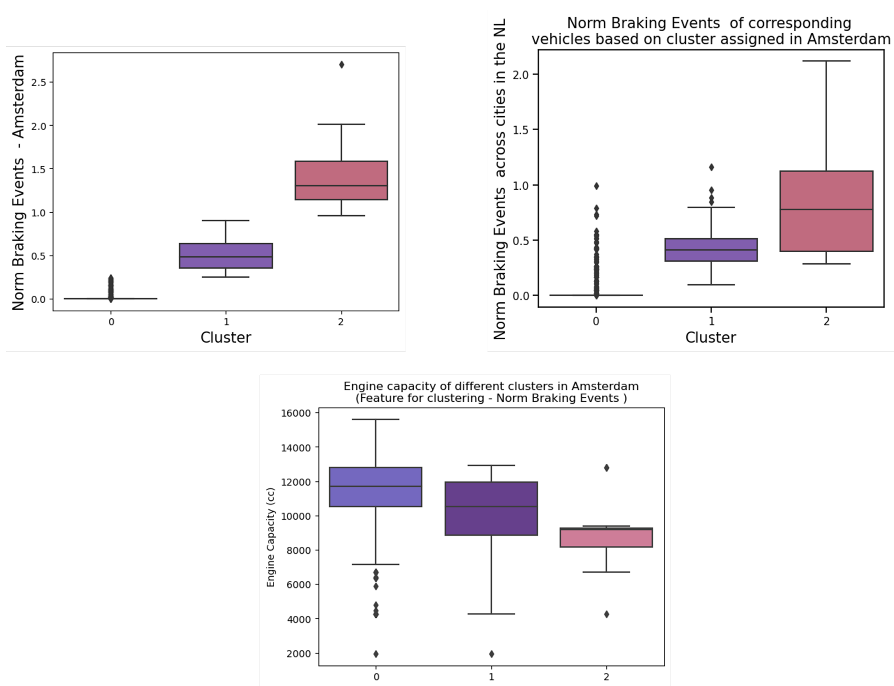

Fig. 35: Cylinder capacity (cc) for different clusters in Amsterdam and the NL

Finally, the results of the link between vehicle characteristics and lane departure warnings are inconclusive. The analysis did not yield a distinct relationship for any specific cluster or metric (R/L-LDW). We can conclude from Figure 37 that the environment substantially influences lane departure warnings. As we can infer from the plot, the usage of indicators is higher in urban areas compared to highways. Consequently, the number of lane departure warnings logged is higher on motorways.

Also, theories posit that a driver's goal to maintain an acceptable level of risk determines their adaptation to a road safety intervention system like lane departure warnings. Several psychological characteristics of an individual also influence this, for example, their trust in automation (49). In conclusion, the latent cause behind lane departure warnings could be driver behaviour and environment.

Realistically, lane departures are linked to driver inattention and drowsiness (50). It would be beneficial to include driver demographic, for example, age, gender, and type of work information, for more context and to corroborate the findings. It is also noteworthy that the type of work has an impact on the stability of different features. For example, drivers of a certain transport company tend to be more stable than others. This aspect needs to be further researched.

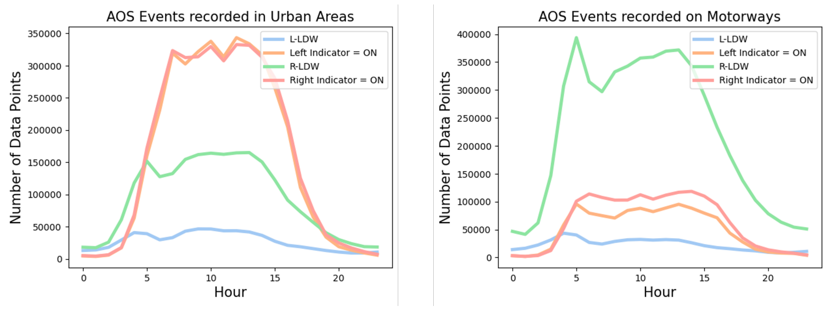

Fig. 36: AOS events - Number of data points recorded over hours of the day (Lane Departure Warnings and Indicators (Left and Right))

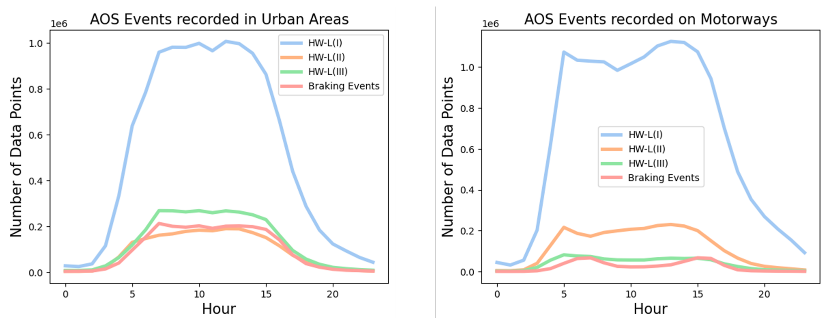

Fig. 37: AOS events - Number of data points recorded over hours of the day (Headway Warnings and Braking Events)

VI. Conclusion, Limitations and Future Work

A. Conclusion

Researchers have developed driver profiling models to predict driver behaviour. In this thesis, we aimed to explore the motivation behind including or excluding certain predictive features used to develop the models. This has been evaluated with naturalistic truck driving behaviour data, focusing on the stability of driving behaviour in time and location.

The current study provided insight into the stability of different truck driving-related metrics with respect to space and time. The underlying issue of predictability without examining the reliability or usefulness of features was raised. In order to address this issue, our analysis used a large-scale FOT, where different truck-driving-related metrics were examined. For temporal stability, it was concluded that virtually all metrics present strong stability over t and t+1 hours of the day.

In the case of spatial stability, we concluded that mean point speed exhibits low spatial stability in general. On the other hand, normalized braking events and lane departure warnings show high stability in both areas, whereas headway warnings only exhibit high stability in urban areas. Vehicle characteristics also affect stability. Vehicles with higher mass and lower engine power tend to have lower mean speeds when compared with other vehicles in a particular area. Lower engine power is also associated with high braking events and headway warnings. Therefore vehicle characteristics also contribute to determining vehicle stability along with route type and driver demographic and behaviour. In the next section, we will discuss the limitations of this study, followed by future work.

B. Limitations

Our analysis is limited by the fact that AOS-FOT only recorded vehicle data. Therefore this paper did not consider individual driving characteristics. This is a common issue when analysing truck driving behaviour, as, at times, drivers do not drive the exact vehicle but may switch between multiple trucks (51). Therefore, we cannot extrapolate our study to comment on driving behaviour. However, as we have concluded that vehicle characteristics and the environment significantly impact stability, we can extend this study to provide a basis for improvement in fleet management. For example, using trucks with a higher power is associated with lower headway warnings and braking events, which can be used as surrogates for safety.

C. Future Work

The information above should be considered in the future, and a more comprehensive dataset that accounts for driver information should be used for further research. By doing so, we can first examine driver behaviour's contribution towards stability. Consequently, we can compare the contribution of vehicles and drivers towards stability individually. Secondly, this paper has studied temporal and spatial stability in isolation. An intriguing direction could be examining these in combination, as the differences among drivers or vehicles might be more pronounced. The next step could be using these features to build generalised models to profile drivers and assess driving patterns. It can also be integrated with models developed for trajectory prediction in autonomous driving to supplement the understanding of driver intention (52).

This work provides the foundation for targeted coaching to drivers, better fleet management, and developing models for insurance companies (Pay As You Drive - PAYD). Secondly, this dataset is quite extensive and provides insight into naturalistic truck driving behaviour at a large scale. Even though it is well-known among field-operational test (FOT) experts to date, it has yet to be investigated in the published academic literature in over fourteen years. As a first step to stimulate future researchers to explore this dataset further, the code used for analysing this dataset is accessible online. Usable research is vital to facilitate our understanding of truck driving in general and mitigate the impact trucks have on the environment while enhancing safety.

References

[1] J. C. de Winter and D. Dodou, "The driver behaviour questionnaire as a predictor of accidents: A meta-analysis," *Journal of safety research*, vol. 41, no. 6, pp. 463–470, 2010.

[2] Y. Ali, A. Sharma, M. M. Haque, Z. Zheng, and M. Saifuzzaman, "The impact of the connected environment on driving behavior and safety: A driving simulator study," *Accident Analysis & Prevention*, vol. 144, p. 105643, 2020.

[3] T. Melman, D. Abbink, X. Mouton, A. Tapus, and J. de Winter, "Multivariate and location-specific correlates of fuel consumption: A test track study," *Transportation Research Part D: Transport and Environment*, vol. 92, p. 102627, 2021.

[4] O. Carsten, K. Kircher, and S. Jamson, "Vehicle-based studies of driving in the real world: The hard truth?" *Accident Analysis & Prevention*, vol. 58, pp. 162–174, 2013.

[5] V. L. Neale, T. A. Dingus, S. G. Klauer, J. Sudweeks, and M. Goodman, "An overview of the 100-car naturalistic study and findings," *National Highway Traffic Safety Administration, Paper*, vol. 5, p. 0400, 2005.

[6] T. Johnson, R. Sherony, and H. C. Gabler, "Driver lane keeping behavior in normal driving using 100-car naturalistic driving study data," in *2016 IEEE Intelligent Vehicles Symposium (IV)*, 2016, pp. 227–232.

[7] R. Chen, K. D. Kusano, and H. C. Gabler, "Driver behavior during overtaking maneuvers from the 100-car naturalistic driving study," *Traffic injury prevention*, vol. 16, no. sup2, pp. S176–S181, 2015.

[8] J. Montgomery, K. D. Kusano, and H. C. Gabler, "Age and gender differences in time to collision at braking from the 100-car naturalistic driving study," *Traffic injury prevention*, vol. 15, no. sup1, pp. S15–S20, 2014.

[9] J. F. Antin, S. Lee, M. A. Perez, T. A. Dingus, J. M. Hankey, and A. Brach, "Second strategic highway research program naturalistic driving study methods," *Safety Science*, vol. 119, pp. 2–10, 2019.

[10] M. M. Ahmed and A. Ghasemzadeh, "The impacts of heavy rain on speed and headway behaviors: an investigation using the shrp2 naturalistic driving study data," *Transportation research part C: emerging technologies*, vol. 91, pp. 371–384, 2018.

[11] R. Eenink, Y. Barnard, M. Baumann, X. Augros, and F. Utesch, "Udrive: the european naturalistic driving study," in *Proceedings of Transport Research Arena*. IFSTTAR, 2014.

[12] M. A. Regan, A. Williamson, R. Grzebieta, J. Charlton, M. Lenne, B. Watson, N. Haworth, A. Rakotonirainy, J. Woolley, R. Anderson *et al.*, "The australian 400-car naturalistic driving study: Innovation in road safety research and policy," in *Proceedings of the 2013 Australasian road safety research, policing & education conference, Brisbane, Queensland*, 2013.

[13] L. Guyonvarch, T. Hermitte, F. Duvivier, C. Val, and A. Guillaume, "Driving style indicator using udrive nds data," *Traffic injury prevention*, vol. 19, no. sup1, pp. S189–S191, 2018.

[14] N. H. T. S. Administration, "National statistics : Vehicles involved in fatal crashes by vehicle type," 2021. [Online]. Available: https://www-fars.nhtsa.dot.gov/Vehicles/VehiclesAllVehicles.aspx

[15] "Traffic safety facts (nhtsa)," 2017. [Online]. Available: https://crashviewer.nhtsa.dot.gov/

[16] "Eu road safety observatory," 2021. [Online]. Available: https://road-safety.transport.ec.europa.eu/system/files/2022-03/FF_buses_hgv_20220209.pdf
[17] "Reducing co emissions from heavy-duty vehicles," 2021. [Online]. Available: https://climate.ec.europa.eu/eu-action/transport-emissions/road-transport-reducing-co2-emissions-vehicles/reducing-co2-emissions-heavy-duty-vehicles_en
[18] T. Zhou and J. Zhang, "Analysis of commercial truck drivers' potentially dangerous driving behaviours based on 11-month digital tachograph data and multilevel modelling approach," *Accident Analysis & Prevention*, vol. 132, p. 105256, 2019.
[19] M. Cai, M. A. A. Yazdi, A. Mehdizadeh, Q. Hu, A. Vinel, K. Davis, H. Xian, F. M. Megahed, and S. E. Rigdon, "The association between crashes and safety-critical events: Synthesized evidence from crash reports and naturalistic driving data among commercial truck drivers," *Transportation research part C: emerging technologies*, vol. 126, p. 103016, 2021.
[20] "Towards safe mobility for all: A data-driven approach," Sep 2020. [Online]. Available: https://www.nwo.nl/en/projects/40319243
[21] "Nextdriver - professional driver coching." [Online]. Available: https://www.nextdriver.nl/
[22] "Abw," Nov 2022. [Online]. Available: https://abw.nl/
[23] M. F. Carfora, F. Martinelli, F. Mercaldo, V. Nardone, A. Orlando, A. Santone, and G. Vaglini, "A "pay-how-you-drive" car insurance approach through cluster analysis," *Soft Computing*, vol. 23, no. 9, pp. 2863–2875, 2019.
[24] A. B. Ellison, "Evaluating changes in driver behaviour for road safety outcomes: a risk profiling approach," 2013.
[25] L. Lu, S. Xiong, and Y. Chen, "A bi-level distribution mixture framework for unsupervised driving performance evaluation from naturalistic truck driving data," *Engineering applications of artificial intelligence*, vol. 104, p. 104349, 2021.
[26] O. F. Ozgul, M. U. Cakir, M. Tan, M. F. Amasyali, and H. T. Hayvaci, "A fully unsupervised framework for scoring driving style," in *2018 International Conference on Intelligent Systems (IS)*. IEEE, 2018, pp. 228–234.
[27] G. P. Figueredo, U. Agrawal, J. M. Mase, M. Mesgarpour, C. Wagner, D. Soria, J. M. Garibaldi, P.-O. Siebers, and R. I. John, "Identifying heavy goods vehicle driving styles in the united kingdom," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 9, pp. 3324–3336, 2018.
[28] C. Wu, C. Sun, D. Chu, Z. Huang, J. Ma, and H. Li, "Clustering of several typical behavioral characteristics of commercial vehicle drivers based on gps data mining: Case study of highways in china," *Transportation Research Record*, vol. 2581, no. 1, pp. 154–163, 2016.
[29] U. Agrawal, J. M. Mase, G. P. Figueredo, C. Wagner, M. Mesgarpour, and R. I. John, "Towards real-time heavy goods vehicle driving behaviour classification in the united kingdom," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 2330–2336.
[30] C. Cabrall, "Human factors of monitoring driving automation eyes and scenes," *Delft University of Technology*, 2019.
[31] J. Desai, H. Li, J. K. Mathew, Y.-T. Cheng, A. Habib, and D. M. Bullock, "Correlating hard-braking activity with crash occurrences on interstate construction projects in indiana," *Journal of Big Data Analytics in Transportation*, vol. 3, no. 1, pp. 27–41, 2021.
[32] S. Niu and S. V. Ukkusuri, "Risk assessment of commercial dangerous-goods truck drivers using geo-location data: A case study in china," *Accident Analysis & Prevention*, vol. 137, p. 105427, 2020.
[33] O. Raddaoui, M. M. Ahmed, and S. M. Gaweesh, "Assessment of the effectiveness of connected vehicle weather and work zone warnings in improving truck driver safety," *IATSS research*, vol. 44, no. 3, pp. 230–237, 2020.
[34] H. J. Walnum and M. Simonsen, "Does driving behavior matter? an analysis of fuel consumption data from heavy-duty trucks," *Transportation research part D: transport and environment*, vol. 36, pp. 107–120, 2015.
[35] J. Lin and D. Niemeier, "Regional driving characteristics, regional driving cycles," *Transportation Research Part D: Transport and Environment*, vol. 8, no. 5, pp. 361–381, 2003.
[36] L. R. Zeitlin, "Estimates of driver mental workload: A long-term field trial of two subsidiary tasks," *Human Factors*, vol. 37, no. 3, pp. 611–621, 1995.
[37] I. R. Eenink, "Verkeersveiligheidseffecten van anti-ongevalsystemen."
[38] J. Hogema, "Anti-ongevalsystemen: data-analyse," p. 1–89, 2009.
[39] "Google scholar search result : Shrp2." [Online]. Available: https://bit.ly/3U85oR1
[40] "Aos for net wiki," Apr 2015. [Online]. Available: https://wiki.fot-net.eu/index.php/AOS
[41] Ernst and Young, "Aos project – data quality assessment," July, 2009.
[42] J. Hogema, "Anti ongeval systemen; onderzoeksdesign," January, 2009.
[43] O. T. SrL, "Aos project – interface control document," January, 2009.
[44] J. MacQueen, "Classification and analysis of multivariate observations," in *5th Berkeley Symp. Math. Statist. Probability*, 1967, pp. 281–297.
[45] "Rdw database." [Online]. Available: https://www.rdw.nl/
[46] A. J. Kean, R. A. Harley, and G. R. Kendall, "Effects of vehicle speed and engine load on motor vehicle emissions," *Environmental science & technology*, vol. 37, no. 17, pp. 3739–3746, 2003.
[47] M. Guillen, J. P. Nielsen, A. M. Pérez-Marín, and V. Elpidorou, "Can automobile insurance telematics predict the risk of near-miss events?" *North American Actuarial Journal*, vol. 24, no. 1, pp. 141–152, 2020.
[48] T. Driessen, D. Dodou, D. de Waard, and J. de Winter, "Predicting damages, fines, and fuel consumption from truck

driver data: A study from the netherlands." [Online]. Available: https://www.researchgate.net/publication/364322807_Predicting_damages_fines_and_fuel_consumption_from_truck_driver_data_A_study_from_the_Netherlands

[49] C. M. Rudin-Brown and Y. I. Noy, "Investigation of behavioral adaptation to lane departure warnings," *Transportation Research Record*, vol. 1803, no. 1, pp. 30–37, 2002.

[50] J. Navarro, J. Deniel, E. Yousfi, C. Jallais, M. Bueno, and A. Fort, "Influence of lane departure warnings onset and reliability on car drivers' behaviors," *Applied ergonomics*, vol. 59, pp. 123–131, 2017.

[51] Y. Li, L. Zhao, and L. R. Rilett, "Driving performances assessment based on speed variation using dedicated route truck gps data," *IEEE Access*, vol. 7, pp. 51 002–51 013, 2019.

[52] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conference on Computer Vision*. Springer, 2020, pp. 683–700.

[53] J. C. De Winter, S. D. Gosling, and J. Potter, "Comparing the pearson and spearman correlation coefficients across distributions and sample sizes: A tutorial using simulations and empirical data." *Psychological methods*, vol. 21, no. 3, p. 273, 2016.

[54] "Google scholar search result : "reliability is a necessary"." [Online]. Available: bit.ly/3GyiNyy

APPENDIX

This section provides more detailed information about the dataset structure and an overview of all the plots for different features.

A. Description

The dataset used in this study was recorded as part of a large-scale field operational test aimed at assessing accident prevention systems by the Dutch Ministry of Infrastructure and Water Management. "Connekt conducted a large-scale field operational test for trucks with active driver assistance systems, also known as accident prevention systems (APS). Five different accident prevention systems and a registration system were tested on Dutch highways over eight months. The test's purpose was to understand better the extent to which accident prevention systems can contribute to traffic safety and traffic flow on the Dutch road network. Until now, the contribution of these systems has only been examined to a limited extent (40)."

The data was collected from September 2008 – May 2009 in Europe. Anti-ongevalsystemen (AOS) dataset was recorded using two devices **Clear Box** and **Mobil Eye Terminal**.

The entire dataset was split into three groups based on the collected data type - Orderly Use data, Main AOS Data and Auxillary Data.

1) **Orderly Use Data**: This data was recorded using a clear box, and no real-time alert was provided to the driver.
   - Journey Data: Features recorded every 2 km, which include number plate, latitude, longitude, point speed, timestamp, distance travelled, and time elapsed. The speed limit of the road provides spatial information, the number of lanes, area, road type, class and form.
   - Crash Data: Speed and acceleration values were recorded a few seconds before and after the crash.
2) **Main AOS Data**: Recorded using Mobile Eye. Contrary to Orderly Use data, this data produced a warning when an event is triggered (the driver receives a real-time alert). The events include:
   - Braking Events
   - Headway Warnings: (Level I/II/III-HW)
   - Right and Left Indicator: (R/L-I)
   - Right and Left Lane Departure Warnings: (R/L-LDW)
3) **Auxiliary Data**: Consists of data points recorded corresponding to the Main AOS events file. This includes latitude, longitude, speed and acceleration (X and Y axis).

These groups were further divided based on the type of data recorded into three categories: **TRIP**, **AOS**, and **CRASH**. This study focuses on **TRIP** and **AOS** data.

B. Overview

1) TRIP: **TRIP** folder consists of **SUMMARY** and **DETAIL** files. The detailed folder contains the orderly use data described in Sec A. Fig. 4 depicts trips recorded across Europe during the field operational test. The data was recorded for over 96×10^6 kilometres from 1727 trucks.

2) AOS: AOS comprises SUMMARY and DETAIL files. SUMMARY file contains Main AOS Data discussed in Sec. A. Fig. 7. illustrates the different AOS events recorded in the Netherlands. These events are recorded when triggered and when the point speed is above 55 km/h. We can infer from the figure that certain areas, for example, the Randstad region is more prone to headway warnings. This is likely due to the higher traffic density in these areas. We observe a similar pattern in the case of braking events, which might indicate a correlation between the two. Over 120×10^6 events were recorded, including 41×10^6 Level I/II/II-HW, 3.5×10^6 braking events and over 9.7×10^6 R/L-LDW.

AOS study is well known among field-operational test (FOT) experts, but there have yet to be any scientific publications about its results. The only analysis that has been done divides different trucks based on ADAS equipped (for example, Group 1 is equipped with FCW/HWM system, Group 2 with LDWA system and Group 3 acts as the reference) [Note: not available online] (38). These trucks were analysed under different conditions to study the impact of speed based on the time of day, the speed limit of the road and the ADAS system. Significant results were found for the time of the day and the speed limit of the road, i.e. the higher the limit, the higher the average speed driven. However, the impact of different ADAS systems equipped in different trucks was insignificant. This also motivated us to explore the dataset in another light.

In Figure 38, we can see three plots. The top shows Trip Data recorded around the South Holland area. The plots on the bottom represent Headway Warnings (Level-I) recorded on two different days. We notice that trucks around urban areas (encircled) consistently have a higher density of headway warnings logged compared to another stretch (where we consistently see lower recordings logged). This shows that location and environment significantly impact truck driving behaviour.

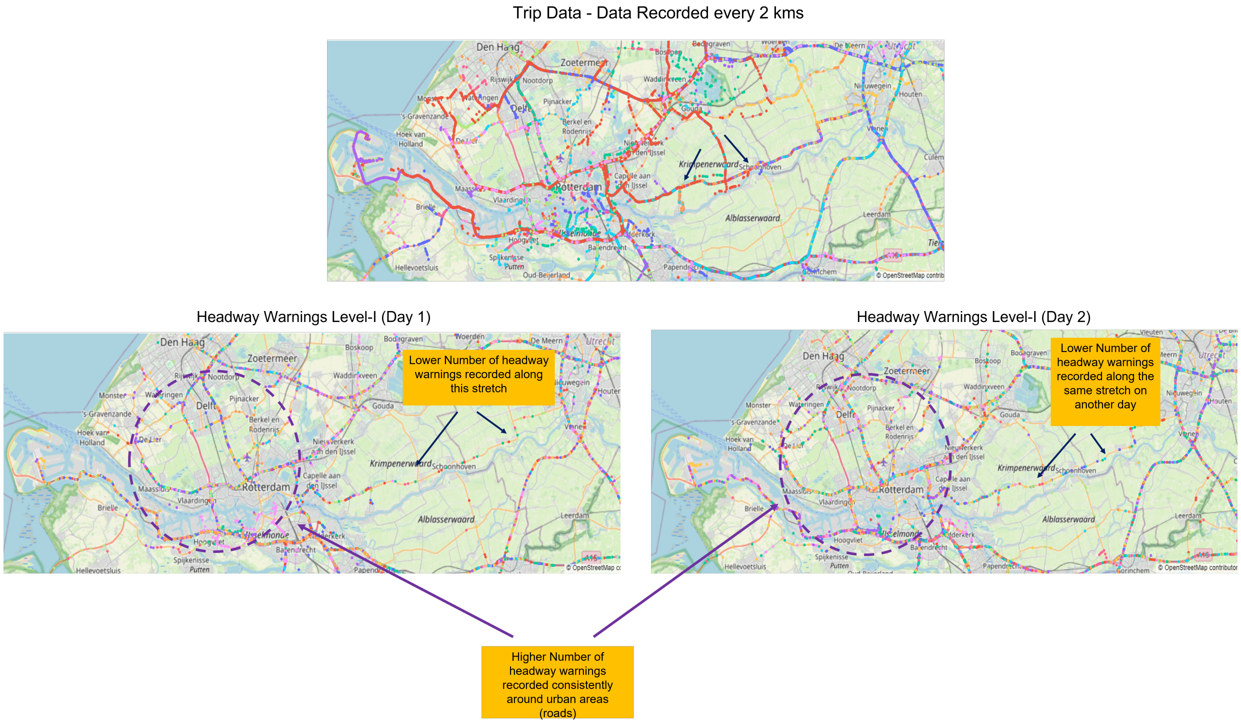

Fig. 38: Headway Warnings in Urban Areas vs Rural Areas

C. Overview Plots-Urban Areas

Table XII gives an overview of the number of data points used for analysis. The locations were filtered by creating polygons around the cities. Mean Speed and Total Distance Travelled were calculated using the Trip Detail folder. The AOS summary file was used to determine the count of different events, which include Braking Events, Headway Warnings (Level-I, II and III), and Lane Departure Warnings (Right/Left). Then all AOS events were normalized using the Total Distance Travelled within the reference location.

TABLE XII: Number of Data Points used during stability analysis for Urban Areas

File	Location	Within the City	Across cities in the NL
Trip Detail	Amsterdam	1.8×10^6	59.2×10^6
	Rotterdam	8.69×10^6	74.4×10^6
	Zwolle	1.77×10^6	54.2×10^6
AOS Summary	Amsterdam	0.93×10^6	31.8×10^6
	Rotterdam	4.7×10^6	44.9×10^6
	Zwolle	0.99×10^6	29.4×10^6

The AOS Summary file was further filtered to extract specific data points corresponding to the events. (Note: The file (and table) also contains data points for other events not used for the analysis). The required data points were further extracted based on the event being analysed.

It is essential to mention here that we have used the Pearson correlation coefficient (r_p) to analyse stability. However, the underlying distribution of metrics is non-normal. We did not transform the dataset before correlation analysis to avoid diluting any meaningful association considering the majority of numbers were very small. Secondly, the Pearson correlation coefficient is not robust to outliers. It has been stated that the Spearman correlation coefficient (r_s) is more robust to such variations and is less variable when used for non-normal data (53). However, the Spearman correlation coefficient might present us with misleading results, considering it does not focus on linearity but on the strength of the relationship, which goes against the base of our argument. Hence, r_p is considered for analysis. As the analysis is based on a few assumptions, we can state that stability is a necessary but not sufficient condition for validity (54). Table XIII summarises the results of correlation analysis in cities.

TABLE XIII: Summary of correlation analysis - urban areas for different features

	Amstedam and (Cities across the NL)		Rotterdam and (Cities across the NL)		Zwolle and (Cities across the NL)		Mean r_p	Description
	r_p	r_s	r_p	r_s	r_p	r_s		
Mean Point Speed	0.22	0.18	-0.04	-0.05	0.28	0.24	**0.15**	Very Weak Correlation
Norm Braking Events	0.78	0.97	0.81	0.98	0.73	0.95	**0.77**	Strong Correlation
Norm HW-L(I)	0.63	0.74	0.62	0.77	0.63	0.71	**0.62**	
Norm HW-L(II)	0.61	0.72	0.71	0.83	0.70	0.74	**0.67**	Moderate-Strong Correlation
Norm HW-L(III)	0.57	0.68	0.50	0.70	0.52	0.66	**0.53**	
Norm R-LDW	0.62	0.77	0.68	0.81	0.68	0.79	**0.66**	Moderate-Strong Correlation
Norm L-LDW	0.53	0.64	0.65	0.73	0.50	0.57	**0.56**	

1) Mean Point Speed: This section presents the clustering analysis results for Mean Point Speed in three cities in the Netherlands and their corresponding trips across the Netherlands.

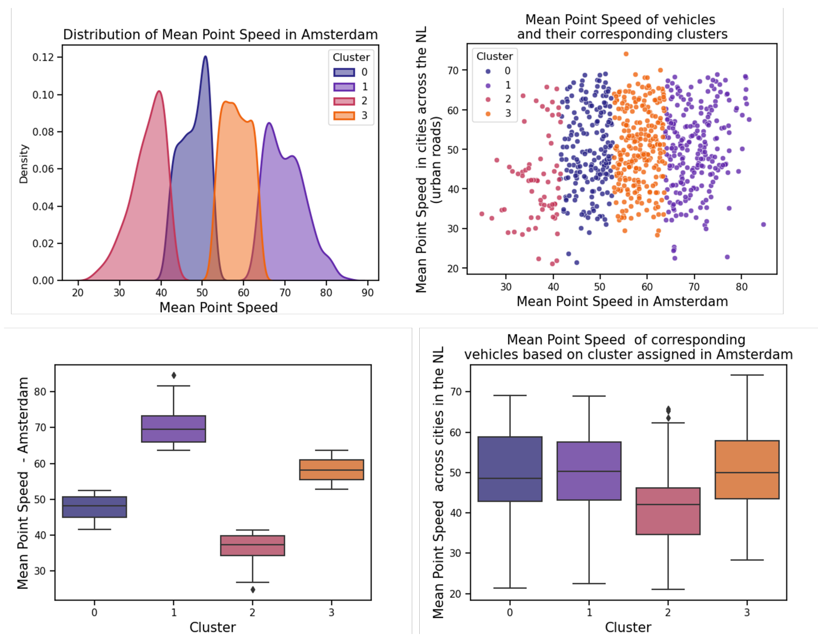

Fig. 39: Clustering Analysis - Mean Point Speed in Amsterdam and corresponding Trips across the Netherlands

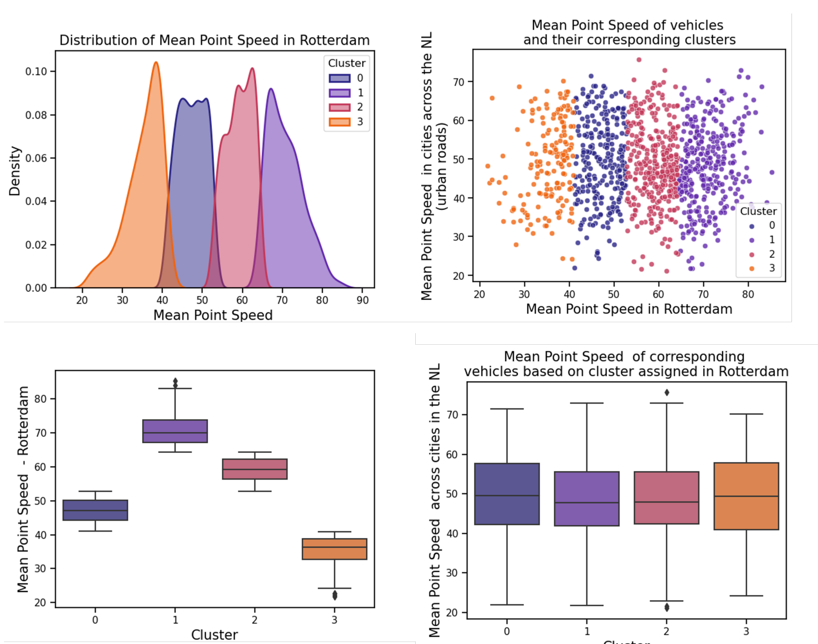

Fig. 40: Clustering Analysis - Mean Point Speed in Rotterdam and corresponding Trips across the Netherlands

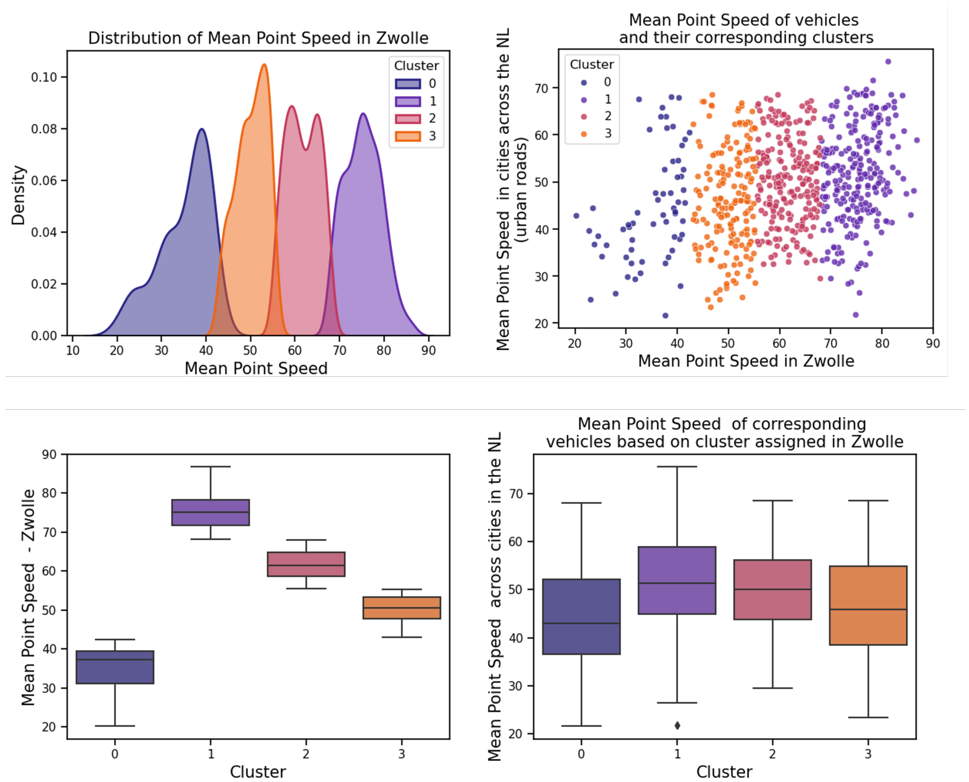

Fig. 41: Clustering Analysis - Mean Point Speed in Zwolle and corresponding Trips across the Netherlands

We can conclude from Table XIII that mean point speed has a very weak correlation. Therefore in Figures 39, 40, and 41, we consistently observe that distinct clusters in cities are indistinct in the Netherlands. Signifying that we cannot predict the speed of a truck predicated on location alone.

2) Normalized Braking Events: This section presents the clustering analysis results for Normalized Braking Events in three cities and their corresponding trips across the Netherlands.

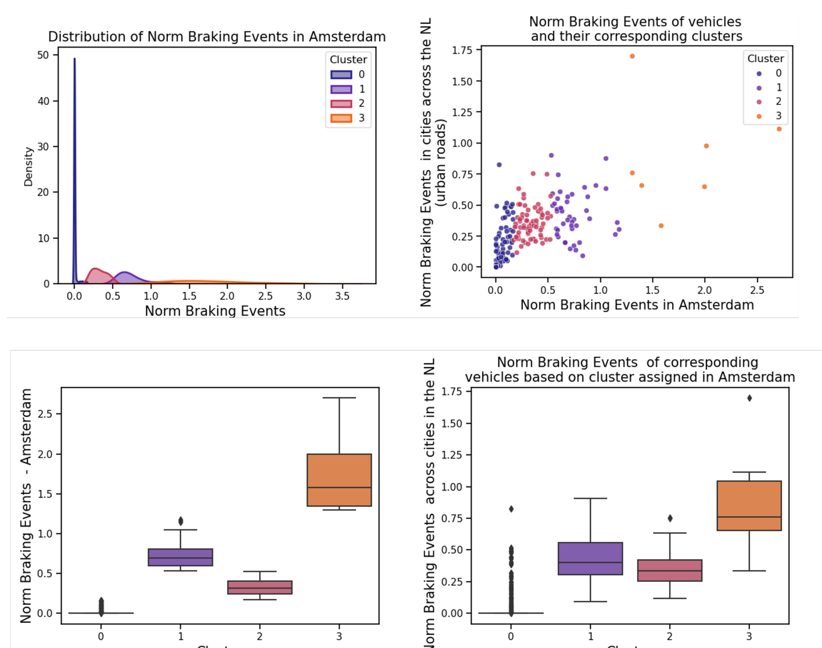

Fig. 42: Clustering Analysis - Normalized Braking Events in Amsterdam and corresponding Trips across the Netherlands

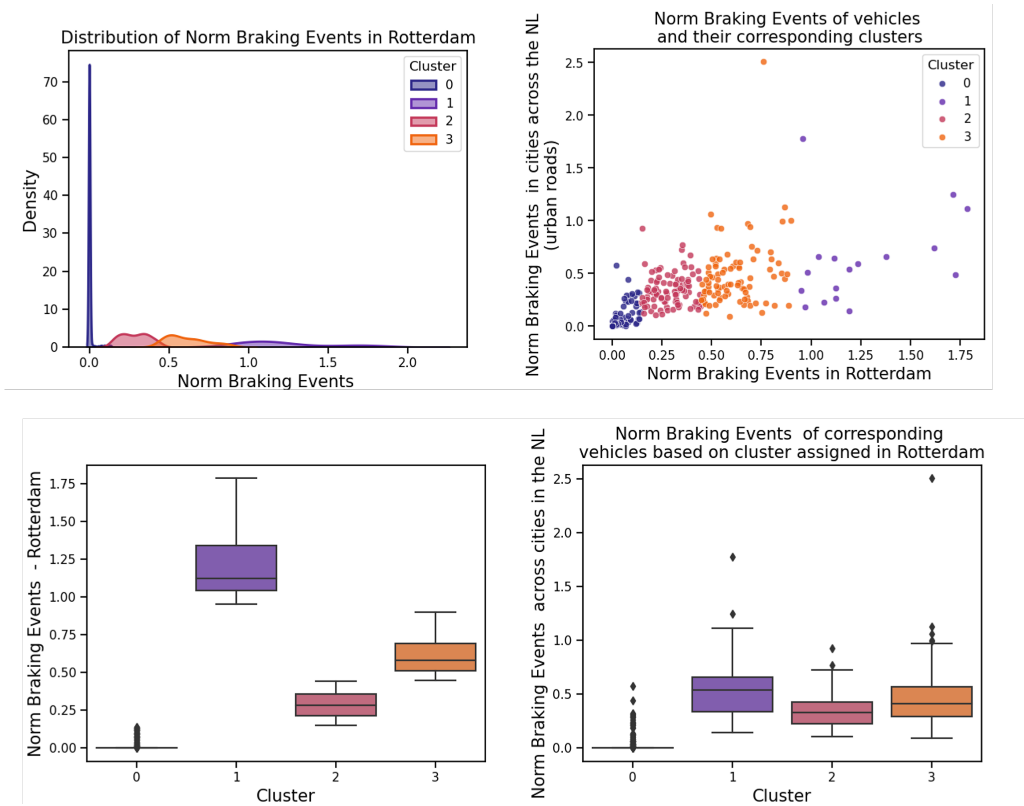

Fig. 43: Clustering Analysis - Normalized Braking Events in Rotterdam and corresponding Trips across the Netherlands

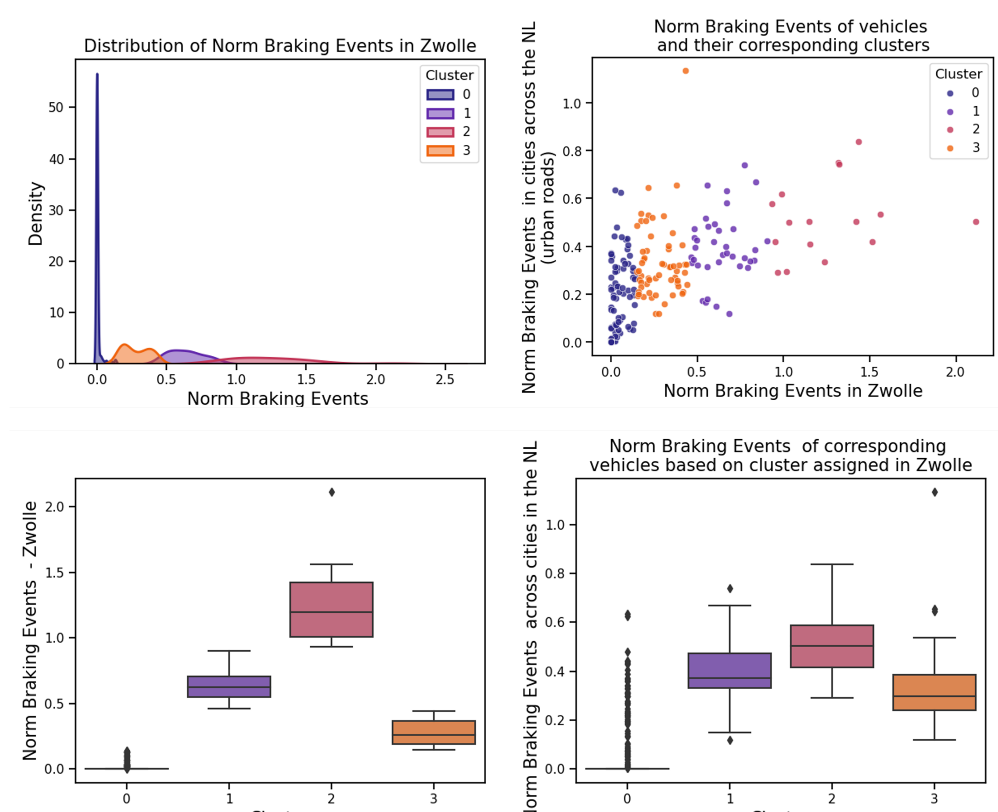

Fig. 44: Clustering Analysis - Normalized Braking Events in Zwolle and corresponding Trips across the Netherlands

We can conclude from Table XIII that normalized braking events have a strong correlation. Therefore in Figures 42, 43, and

44, we can observe that distinct clusters can be found in both cities and the Netherlands. The difference in the behaviour of the lowest-value clusters and the highest ones is most distinguishable.

3) Normalized Headway Warnings Level-I, II, and III: This section presents the clustering analysis results for Normalized Headway Warnings Level-I, II, and III in three cities and their corresponding trips across the Netherlands.

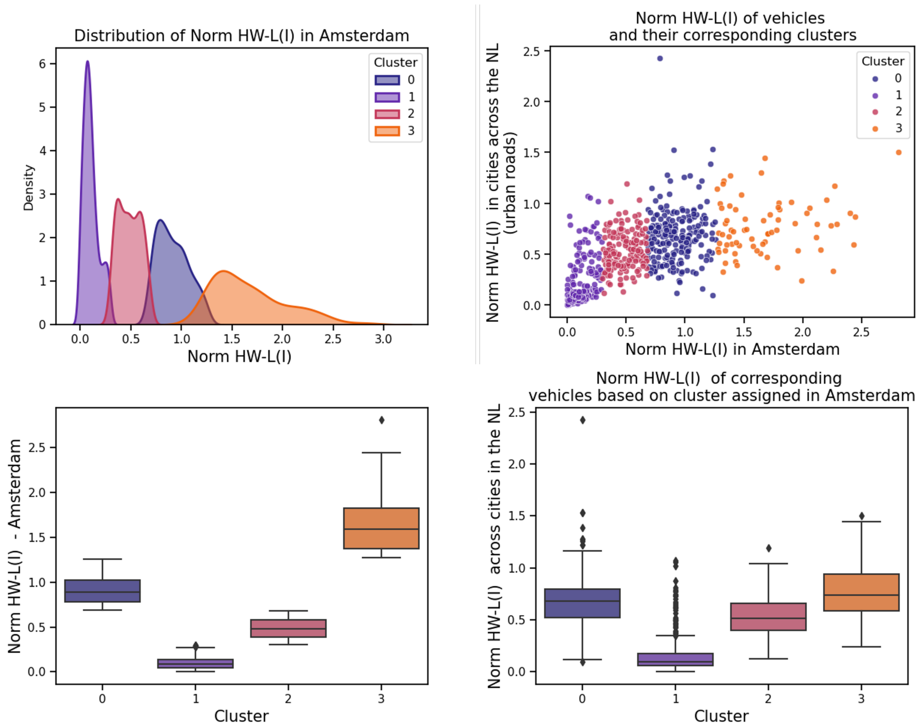

Fig. 45: Clustering Analysis - Normalized Headway Warnings Level-I in Amsterdam and corresponding Trips across the Netherlands

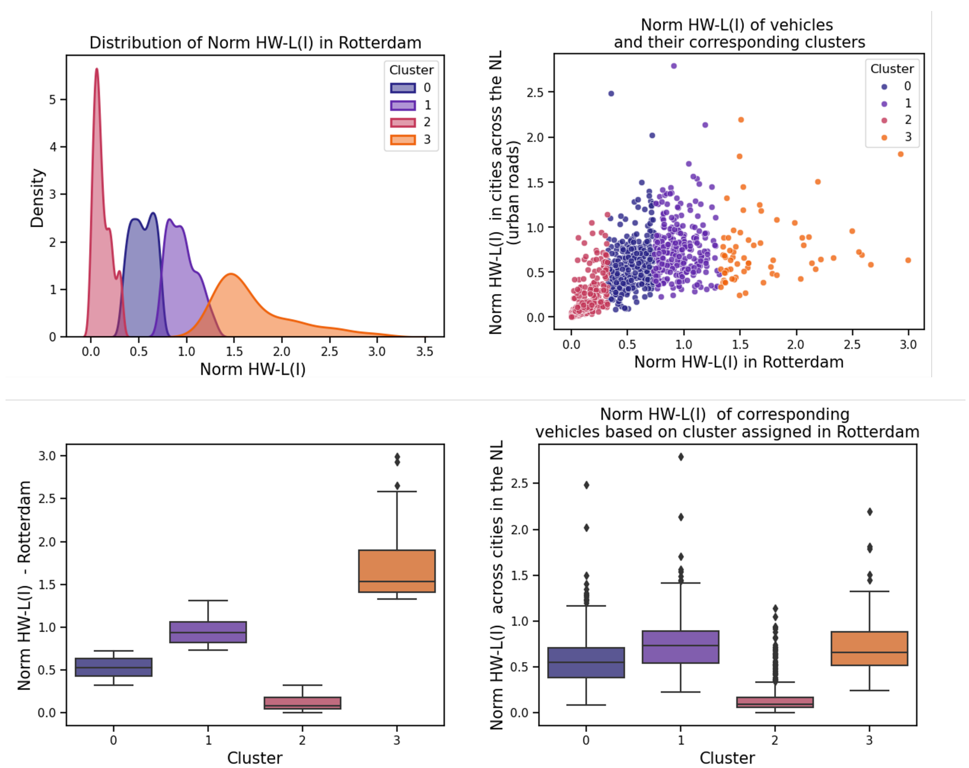

Fig. 46: Clustering Analysis - Normalized Headway Warnings Level-I in Rotterdam and corresponding Trips across the Netherlands

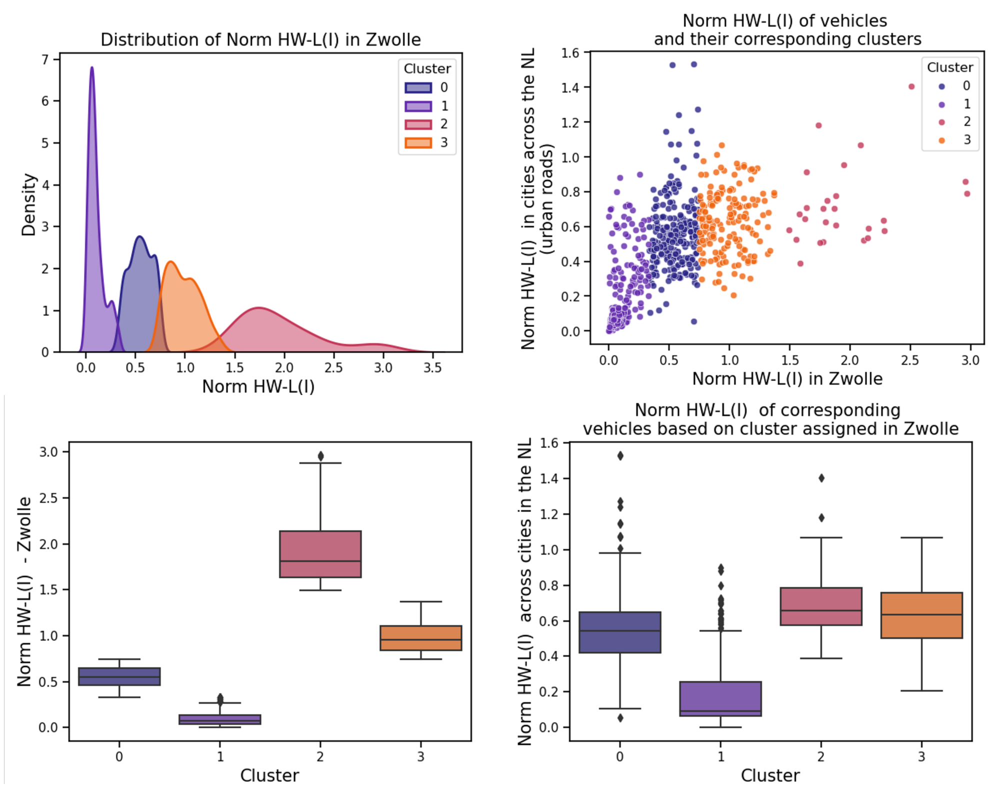

Fig. 47: Clustering Analysis - Normalized Headway Warnings Level-I in Zwolle and corresponding Trips across the Netherlands

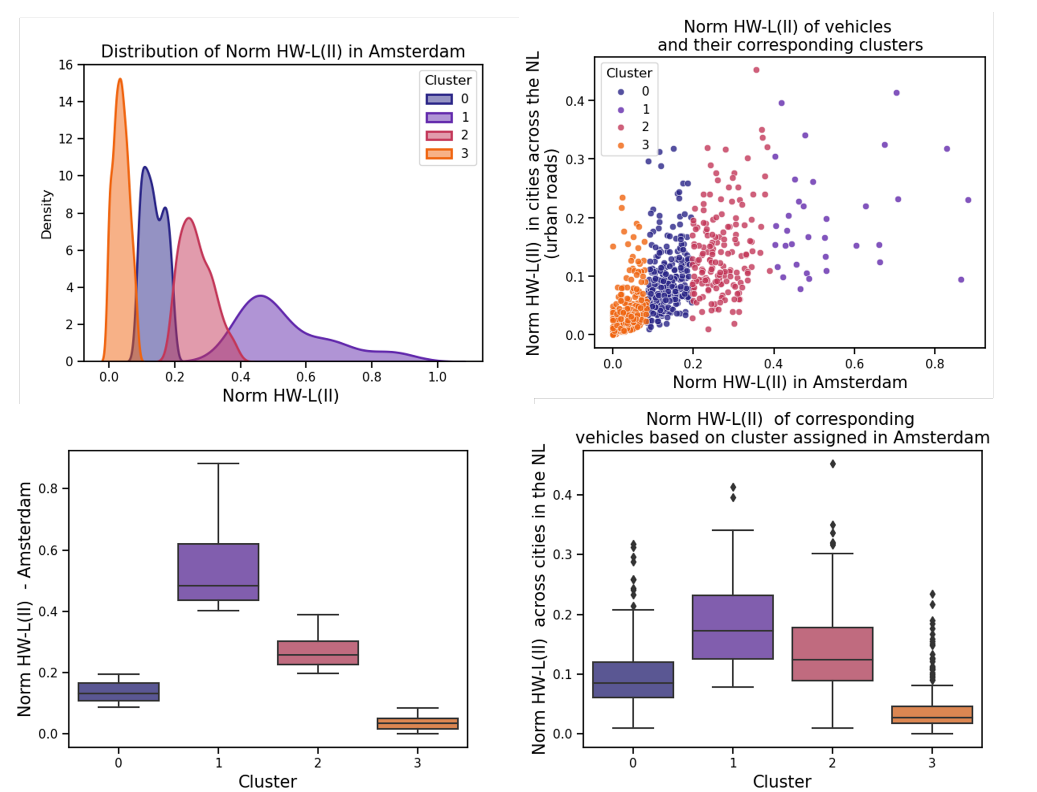

Fig. 48: Clustering Analysis - Normalized Headway Warnings Level-II in Amsterdam and corresponding Trips across the Netherlands

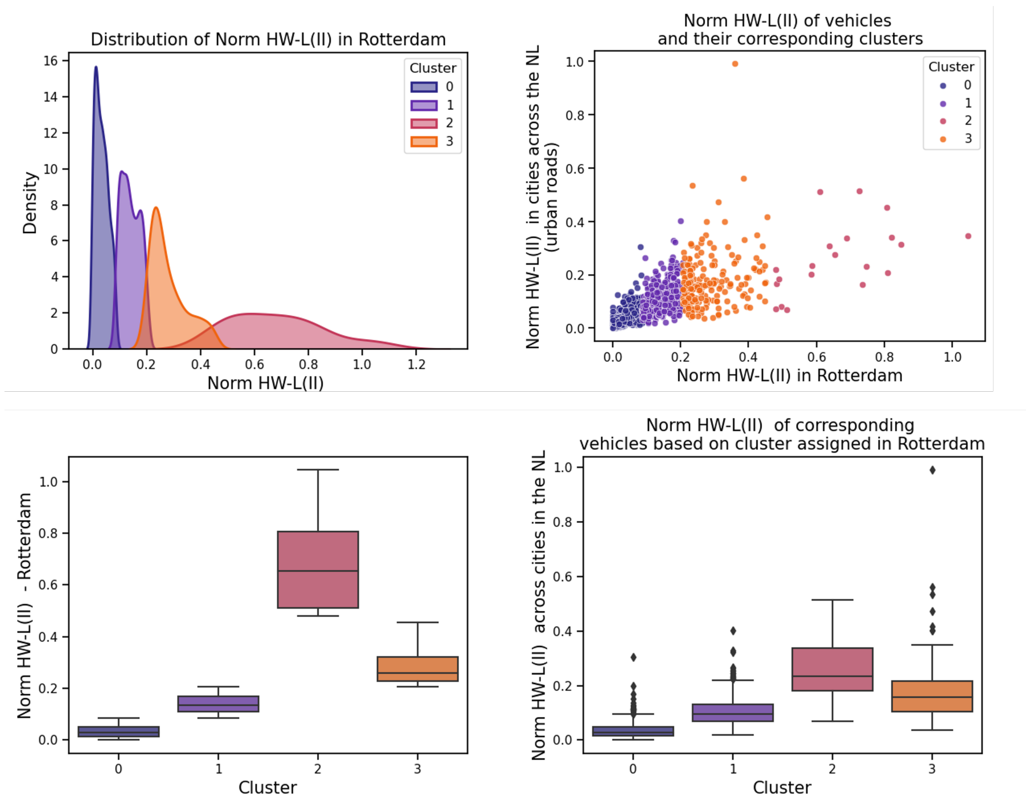

Fig. 49: Clustering Analysis - Normalized Headway Warnings Level-II in Rotterdam and corresponding Trips across the Netherlands

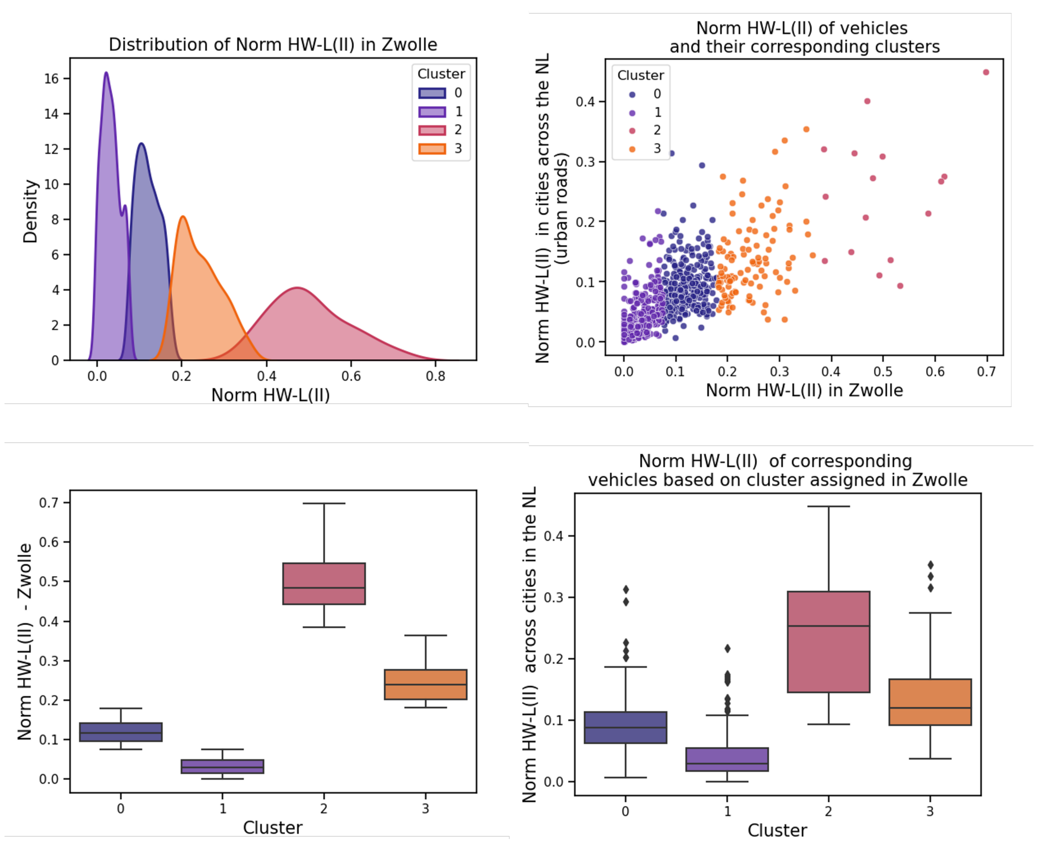

Fig. 50: Clustering Analysis - Normalized Headway Warnings Level-II in Zwolle and corresponding Trips across the Netherlands

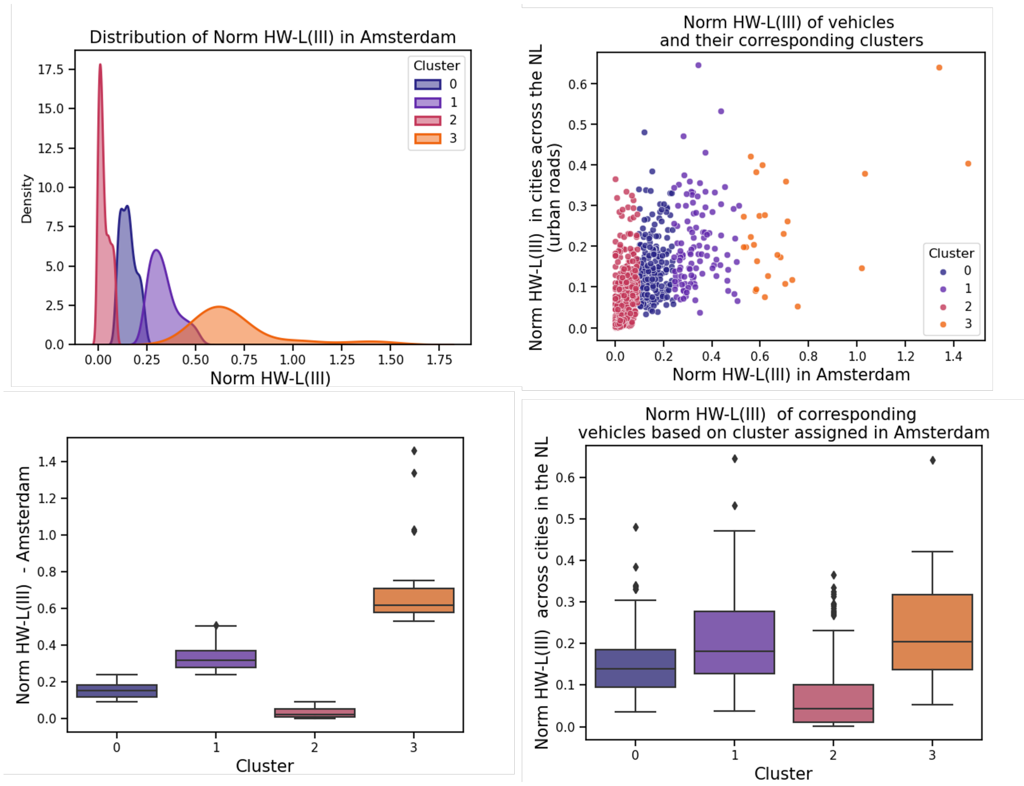

Fig. 51: Clustering Analysis - Normalized Headway Warnings Level-III in Amsterdam and corresponding Trips across the Netherlands

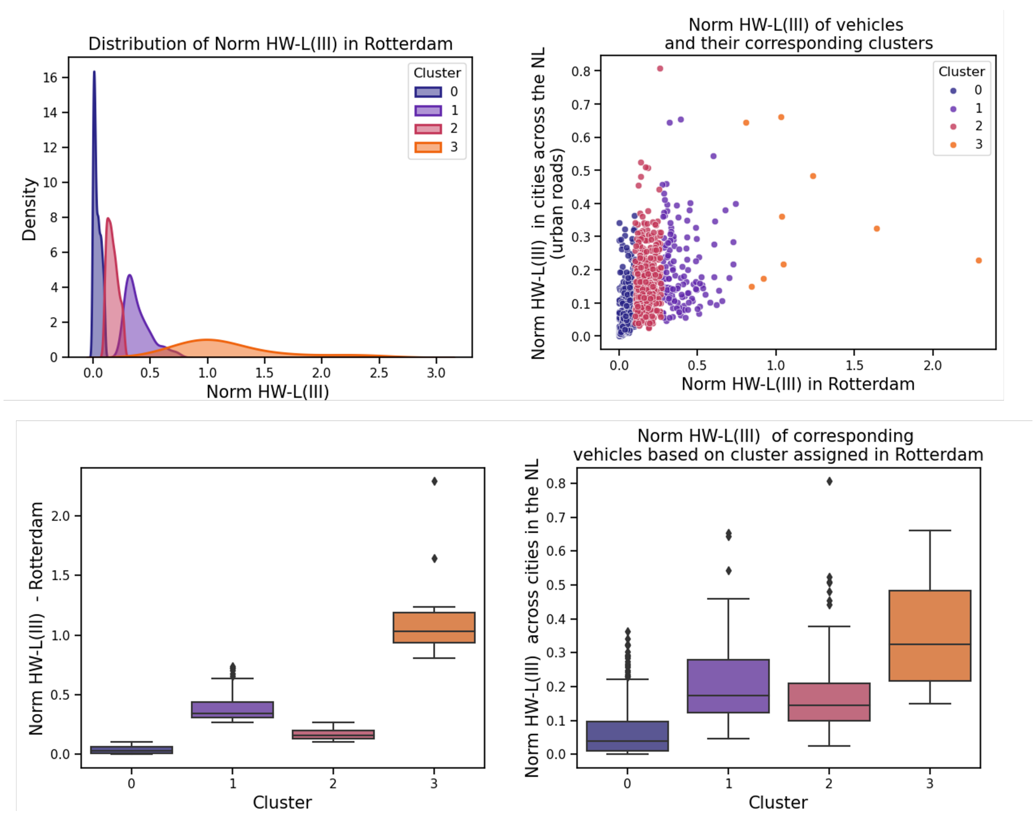

Fig. 52: Clustering Analysis - Normalized Headway Warnings Level-III in Rotterdam and corresponding Trips across the Netherlands

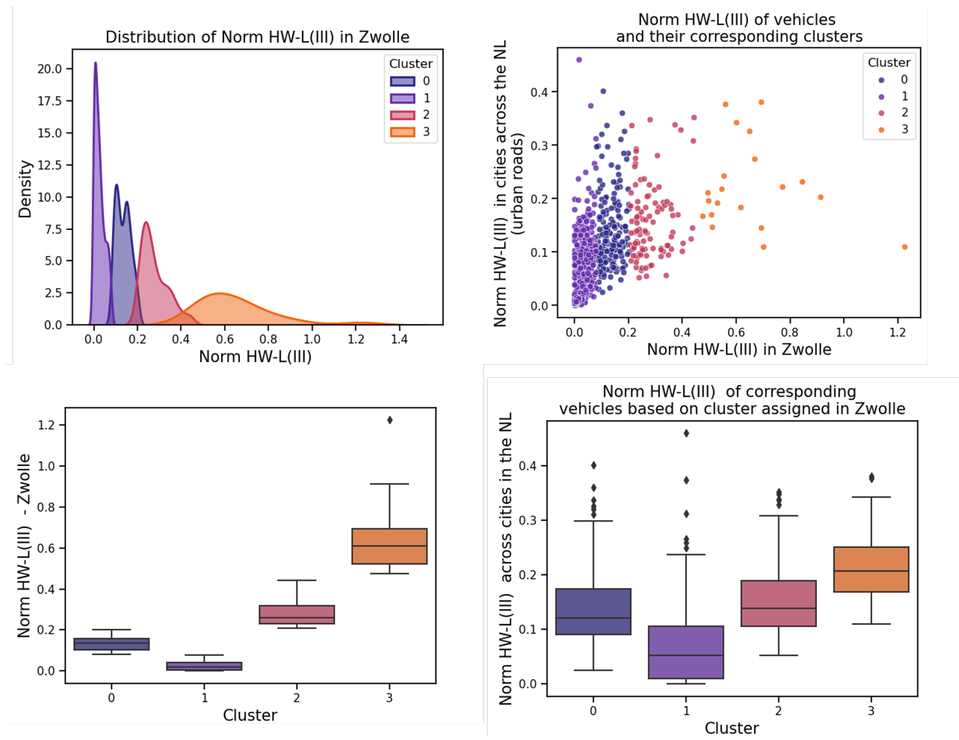

Fig. 53: Clustering Analysis - Normalized Headway Warnings Level-III in Zwolle and corresponding Trips across the Netherlands

We can conclude from Table XIII that normalized headway warnings have a moderate-strong correlation. From the figures in this Section, we can observe that distinct clusters can be found in both cities and the Netherlands. The difference in the behaviour of the lowest-value clusters and the highest ones is most distinguishable, the same as in the case of braking events, which also exhibited high stability.

4) Normalized Lane Departure Warnings-Left, and Right: This Section presents the clustering analysis results for Normalized Lane Departure Warnings-Left, and Right in three cities and their corresponding trips across the Netherlands.

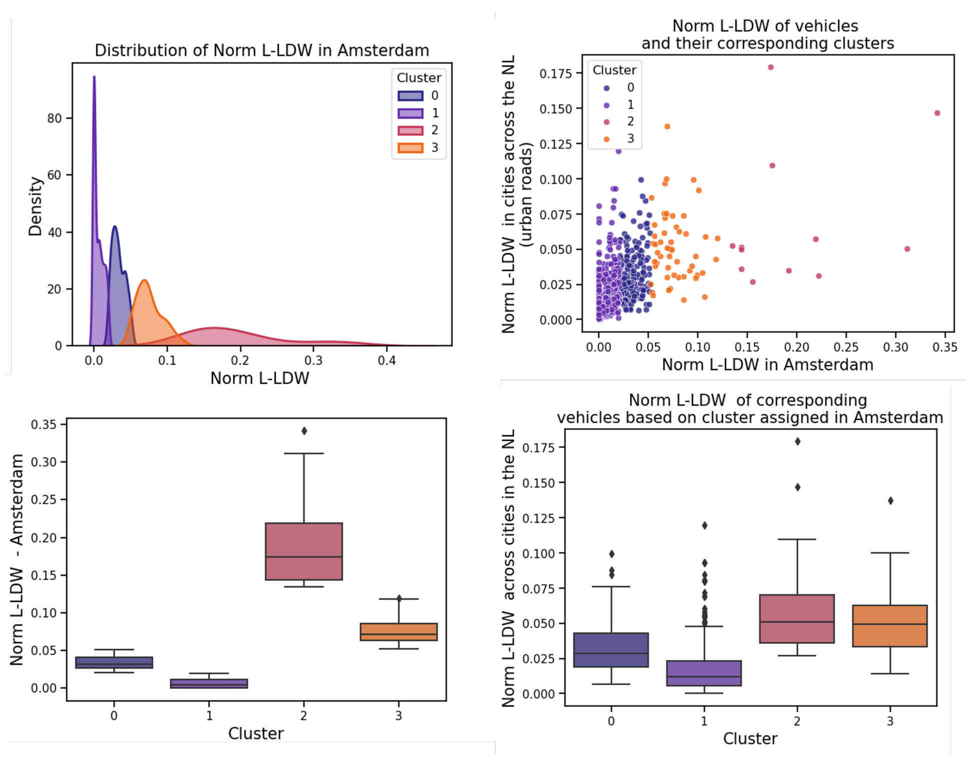

Fig. 54: Clustering Analysis - Normalized Left Lane Departure Warnings in Amsterdam and corresponding Trips across the Netherlands

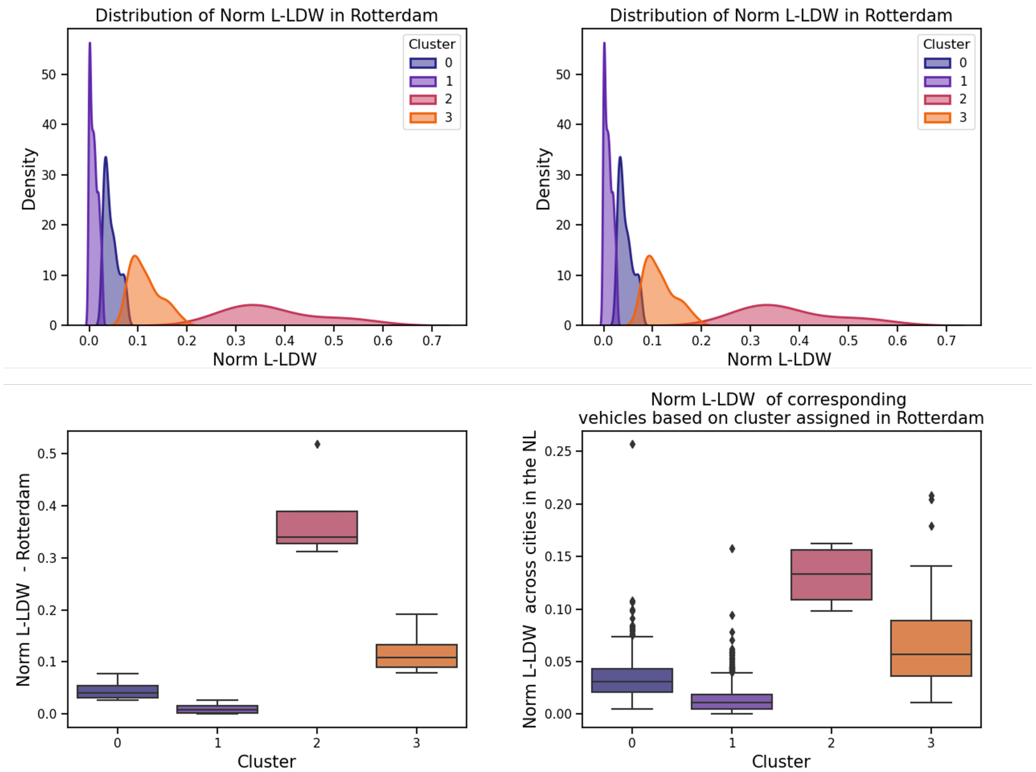

Fig. 55: Clustering Analysis - Normalized Left Lane Departure Warnings in Rotterdam and corresponding Trips across the Netherlands

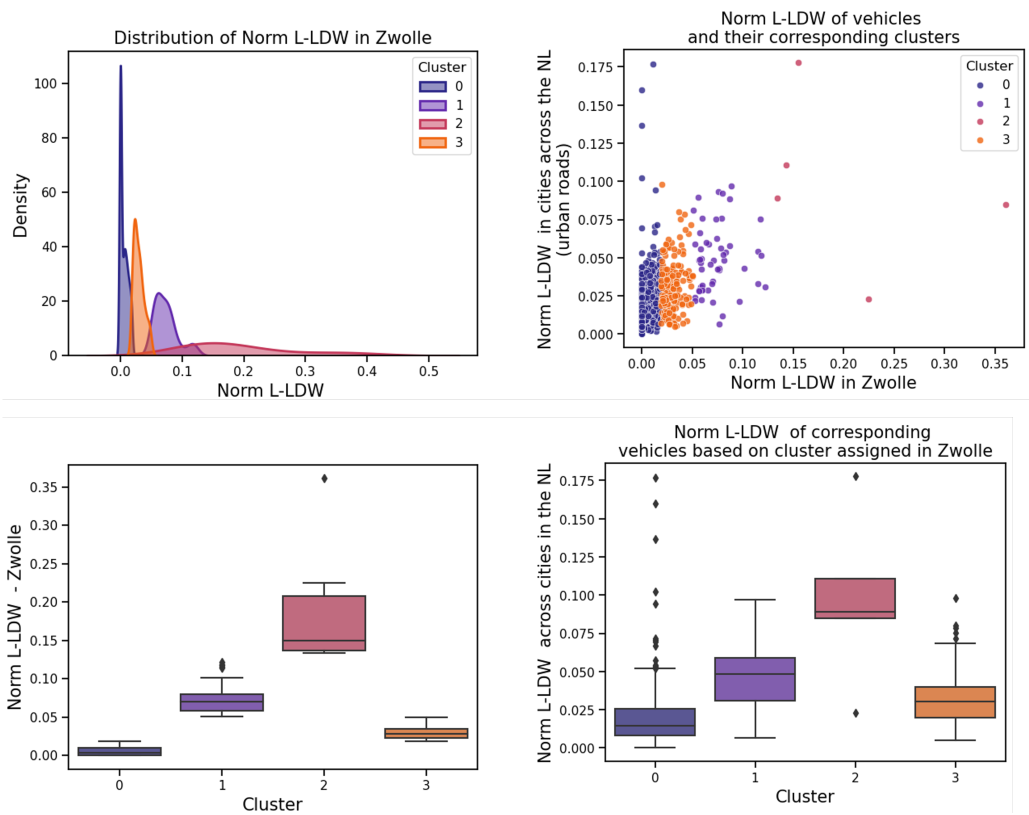

Fig. 56: Clustering Analysis - Normalized Left Lane Departure Warnings in Zwolle and corresponding Trips across the Netherlands

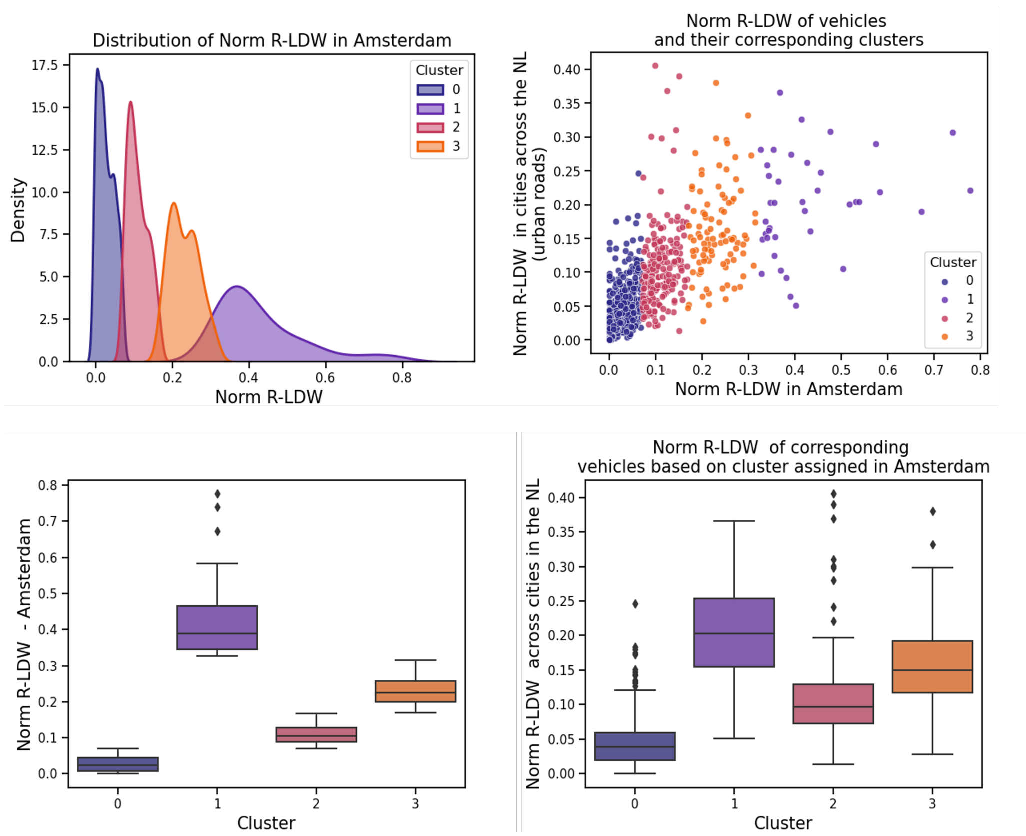

Fig. 57: Clustering Analysis - Normalized Right Lane Departure Warnings in Amsterdam and corresponding Trips across the Netherlands

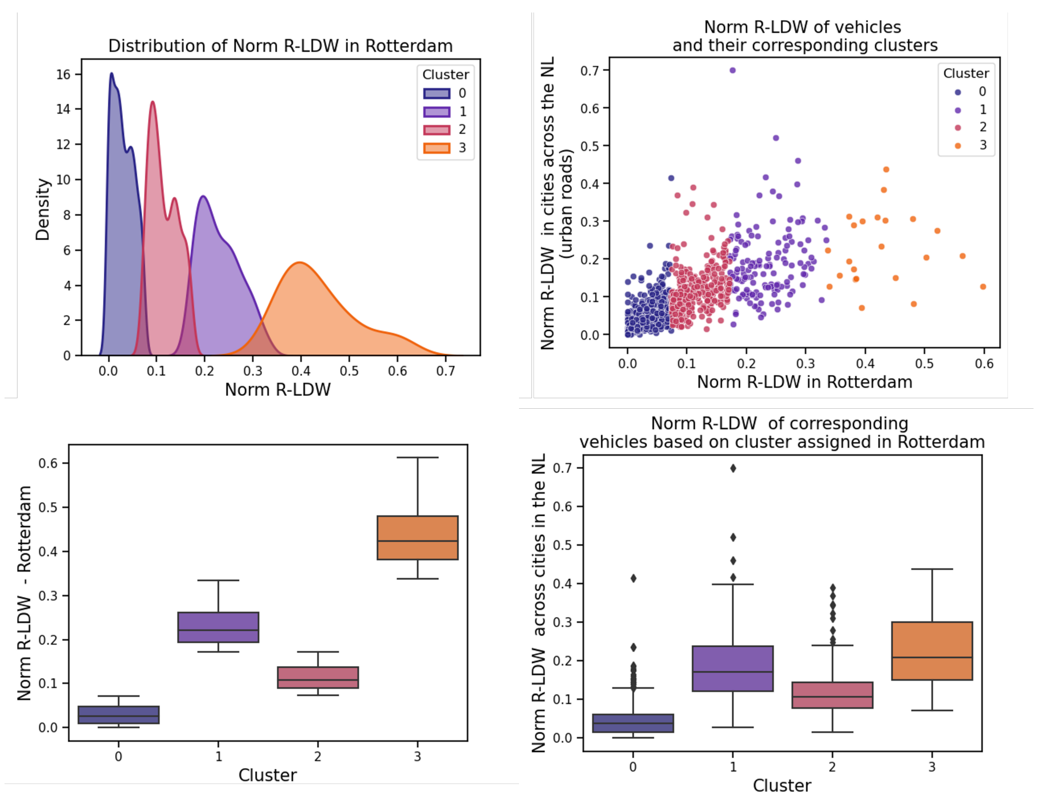

Fig. 58: Clustering Analysis - Normalized Right Lane Departure Warnings in Rotterdam and corresponding Trips across the Netherlands

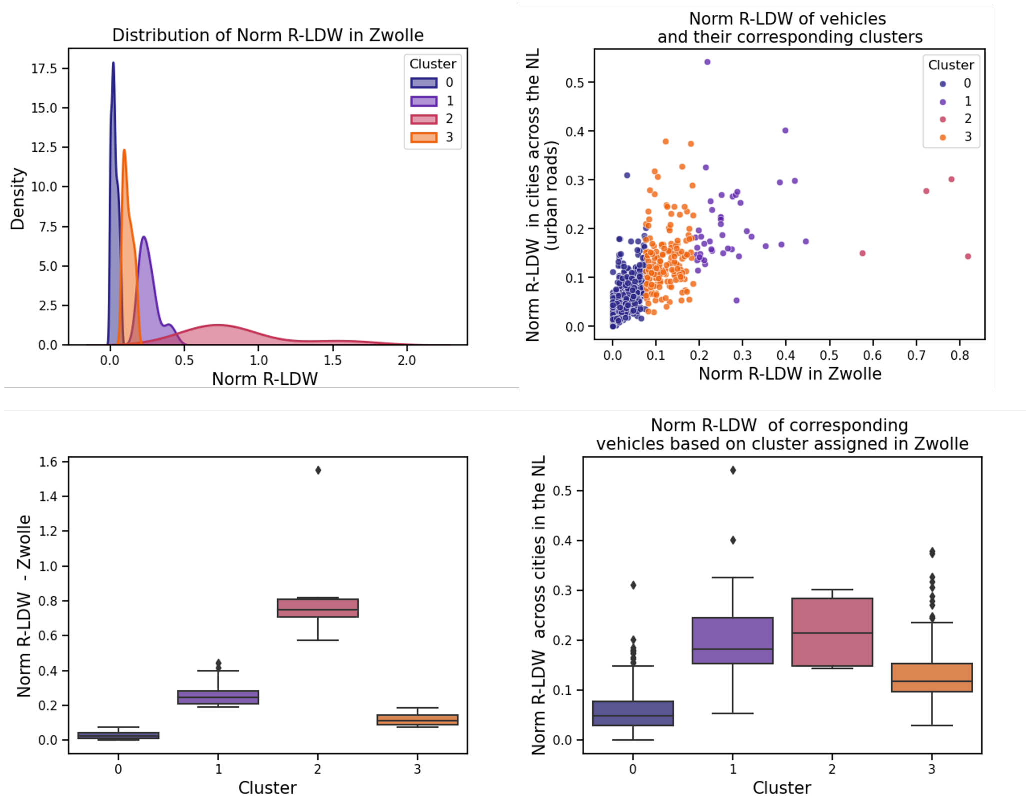

Fig. 59: Clustering Analysis - Normalized Right Lane Departure Warnings in Zwolle and corresponding Trips across the Netherlands

We can conclude from Table XIII that normalized lane departure warnings have a moderate-strong correlation. From the figures in this Section, we can observe that they exhibit higher stability in general and distinct patterns can be observed in both cities and the Netherlands.

D. Clustering Analysis-Motorways

Table XIV gives an overview of the number of data points used to analyse motorways. The trips between two locations were extracted based on Number plate and date. The code for extraction can be found on Github. Trip and AOS data were extracted using this method. Then trips across different motorways in the Netherlands were extracted based on truck number plates. Only data within the Netherlands and some surrounding areas were considered for the analysis. All data points outside this area were filtered by creating polygons. Mean Speed and Total Distance Travelled were calculated using the Trip Detail folder. The AOS summary file was used to determine the count of different events, which include Braking Events, Headway Warnings (Level-I, II and III), and Lane Departure Warnings (Right/Left). Then all AOS events were normalized using the Total Distance Travelled within the reference location.

TABLE XIV: Number of Data Points used during stability analysis for Motorways

File	Location	Motorway	Motorways across the NL
Trip Detail	Utrecht-Leek	6.02×10^4	3.4×10^6
	Utrecht-Eindhoven	4.11×10^5	7.62×10^6
AOS Summary	Utrecht-Leek	1.4×10^4	1.91×10^6
	Utrecht-Eindhoven	1.1×10^5	8.27×10^6

The AOS Summary file was further filtered to extract specific data points corresponding to the events. (Note: The file also contains data points for other events which were not used for the analysis).

This section will discuss correlation analysis to determine spatial stability based on two different motorways in the Netherlands. The first motorway is between Utrecht-Leek (M1) and Utrecht-Eindhoven (M2). Table XV provides an overview of the results of correlation analysis.

TABLE XV: Summary of correlation analysis - motorways for different features

	Utrecht-Leek (M1) (and motorways across the NL)		**Utrecht - Eindhoven (M2) (and motorways across the NL)**		**Mean** r_p	**Description**
	r_p	r_s	r_p	r_s		
Mean Point Speed	0.59	0.54	0.44	0.45	0.51	Moderate Correlation
Norm Braking Events	0.53	0.98	0.74	0.95	0.63	Strong Correlation
Norm HW-L(I)	0.46	0.51	0.49	0.54	0.475	Weak-Moderate Correlation
Norm HW-L(II)	0.26	0.61	0.49	0.61	0.375	
Norm HW-L(III)	0.44	0.43	0.44	0.60	0.44	
Norm L-LDW	0.68	0.76	0.65	0.66	0.66	Strong Correlation
Norm R-LDW	0.81	0.76	0.52	0.53	0.67	

1) Mean Point Speed: This section presents the results of the clustering analysis for Mean Point Speed on two different motorways and their corresponding trips across motorways in the Netherlands.

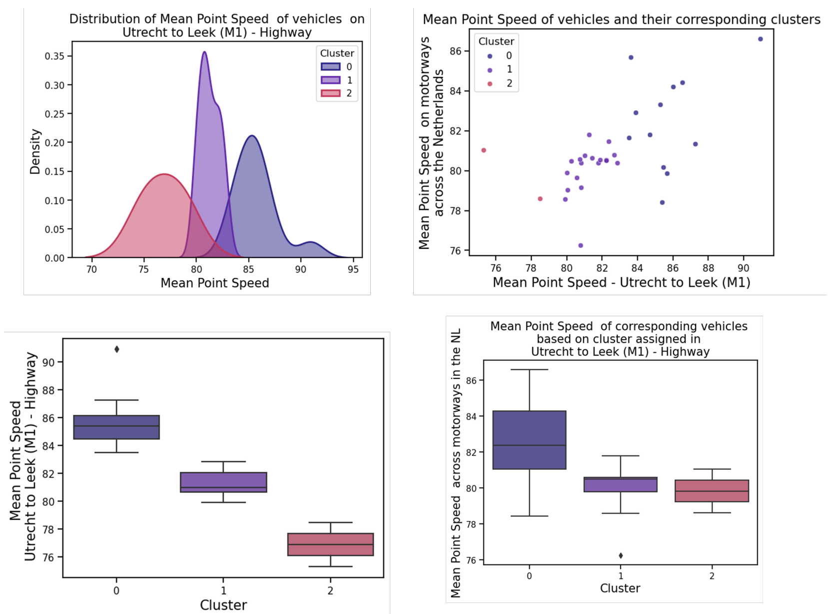

Fig. 60: Clustering Analysis - Mean Point Speed on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

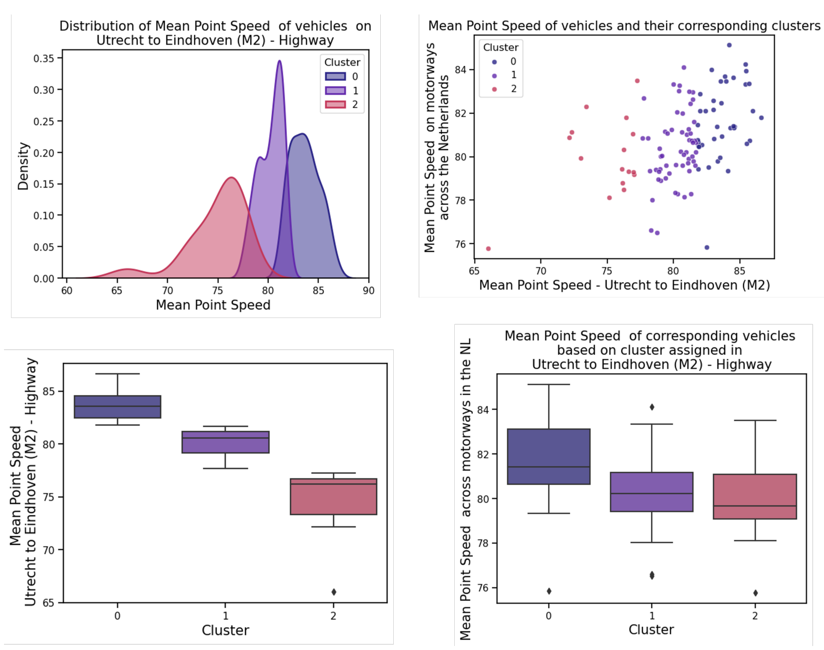

Fig. 61: Clustering Analysis - Mean Point Speed on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

Contrary to mean point speed in urban areas, which exhibited a very weak correlation, on motorways, it exhibits a moderate correlation. This is likely due to the comparable environment of the motorway. We can observe distinguishable clusters (for lowest and highest value).

2) Normalized Braking Events: This section presents the clustering analysis results for Normalized Braking Events on two motorways and their corresponding trips across motorways in the Netherlands.

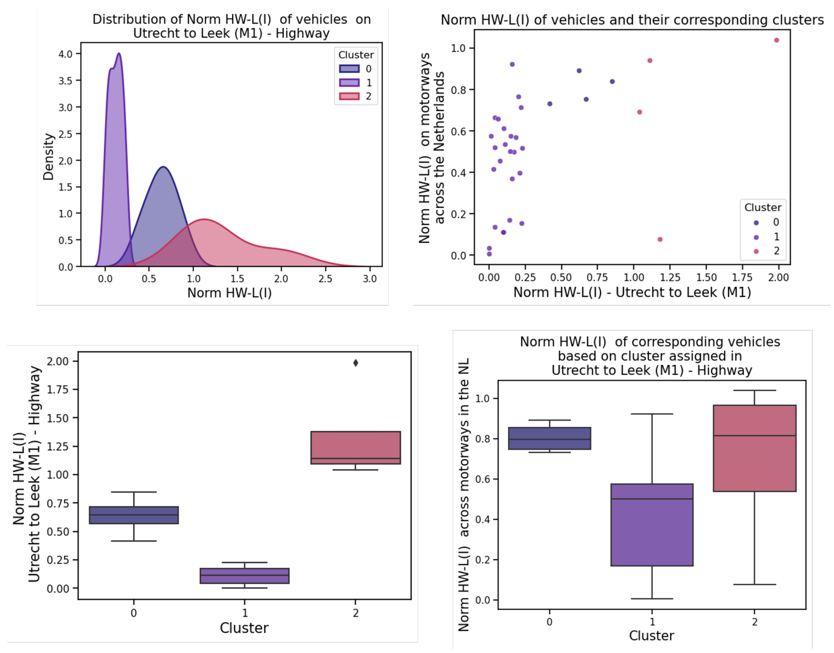

Fig. 62: Clustering Analysis - Normalized Braking Events on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

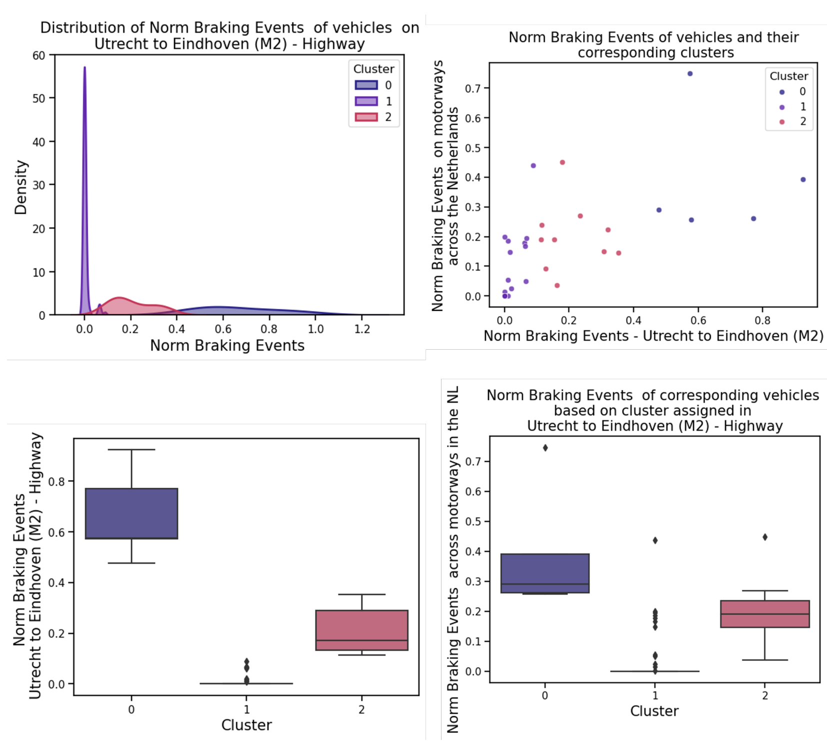

Fig. 63: Clustering Analysis - Normalized Braking Events on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

3) Normalized Headway Warnings Level-I, II, and III: This section presents the results of the clustering analysis for Normalized Headway Warnings Level-I, II, and III on two different motorways and their corresponding trips across motorways in the Netherlands.

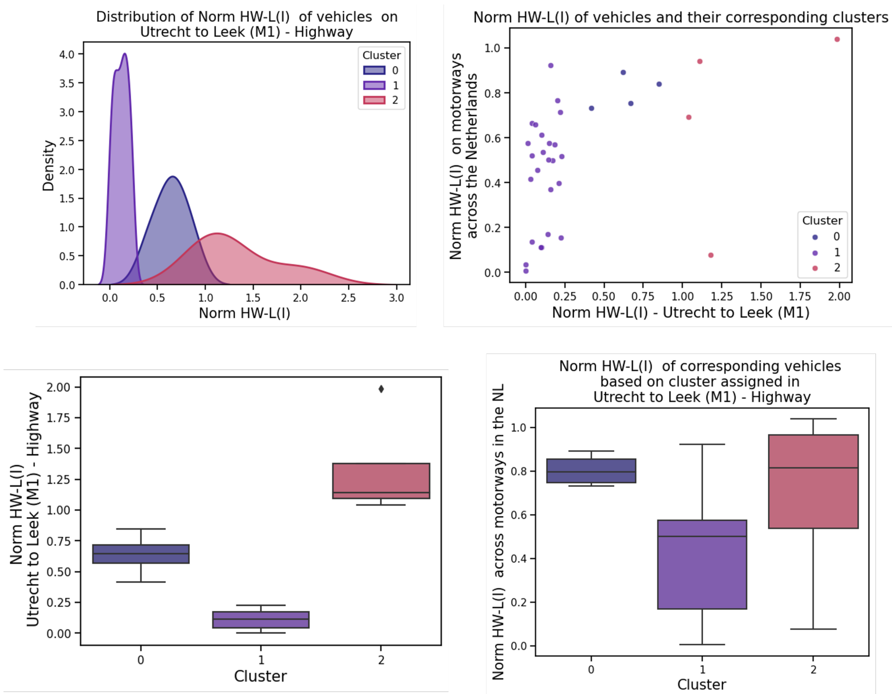

Fig. 64: Clustering Analysis - Normalized Headway Warnings Level-I on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

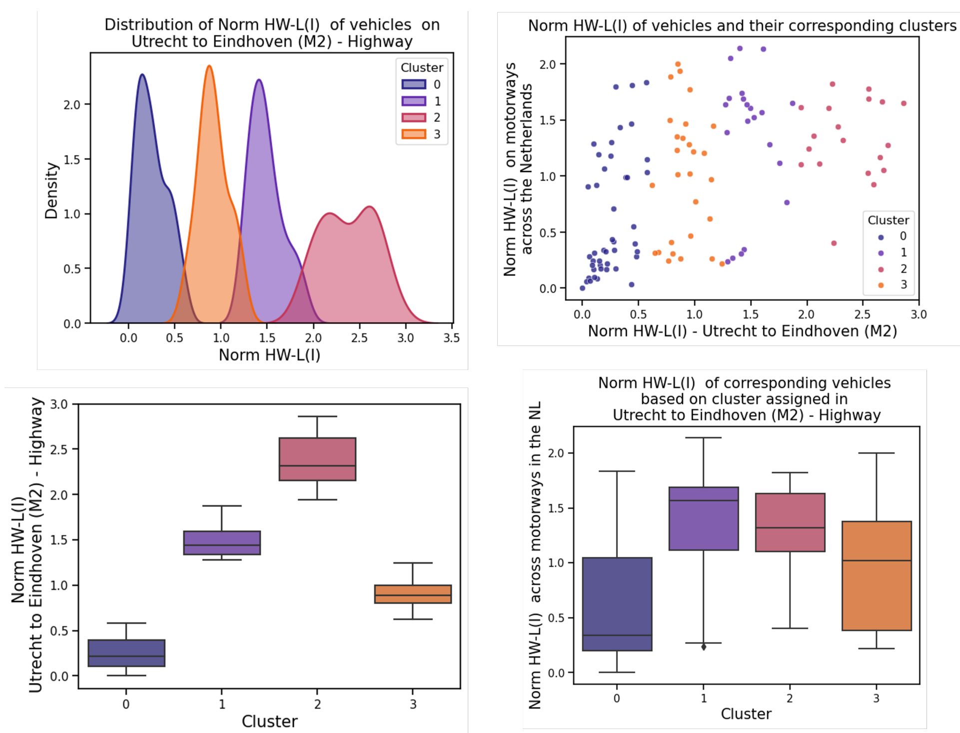

Fig. 65: Clustering Analysis - Normalized Headway Warnings Level-I on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

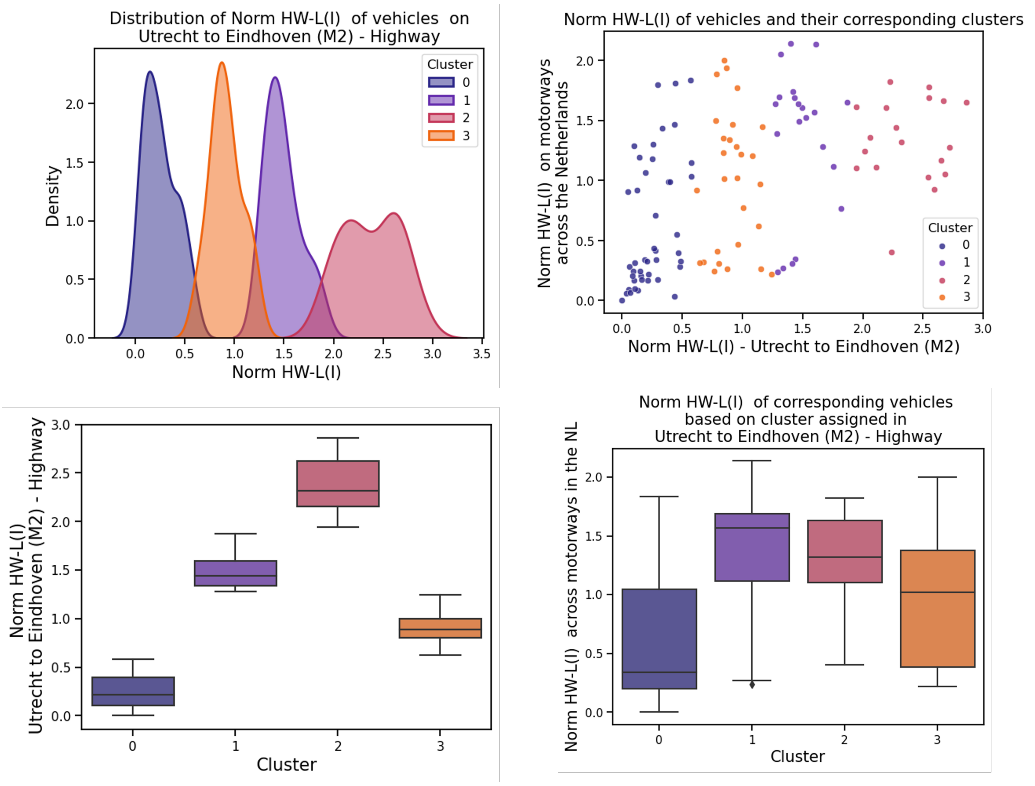

Fig. 66: Clustering Analysis - Normalized Headway Warnings Level-II on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

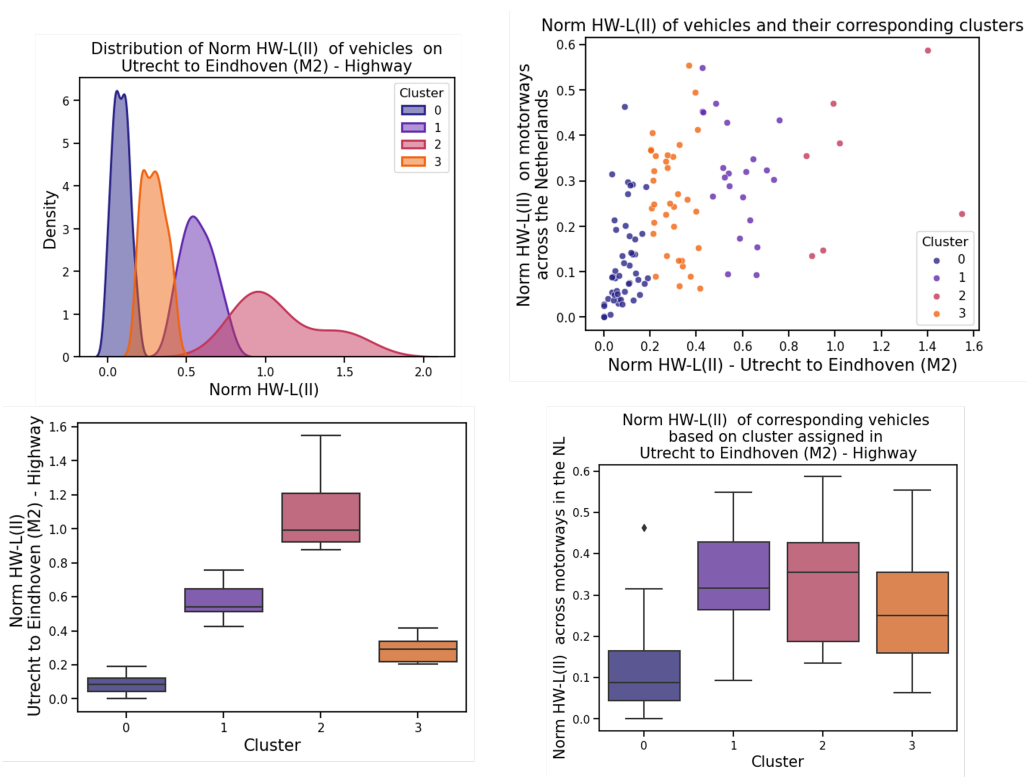

Fig. 67: Clustering Analysis - Normalized Headway Warnings Level-II on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

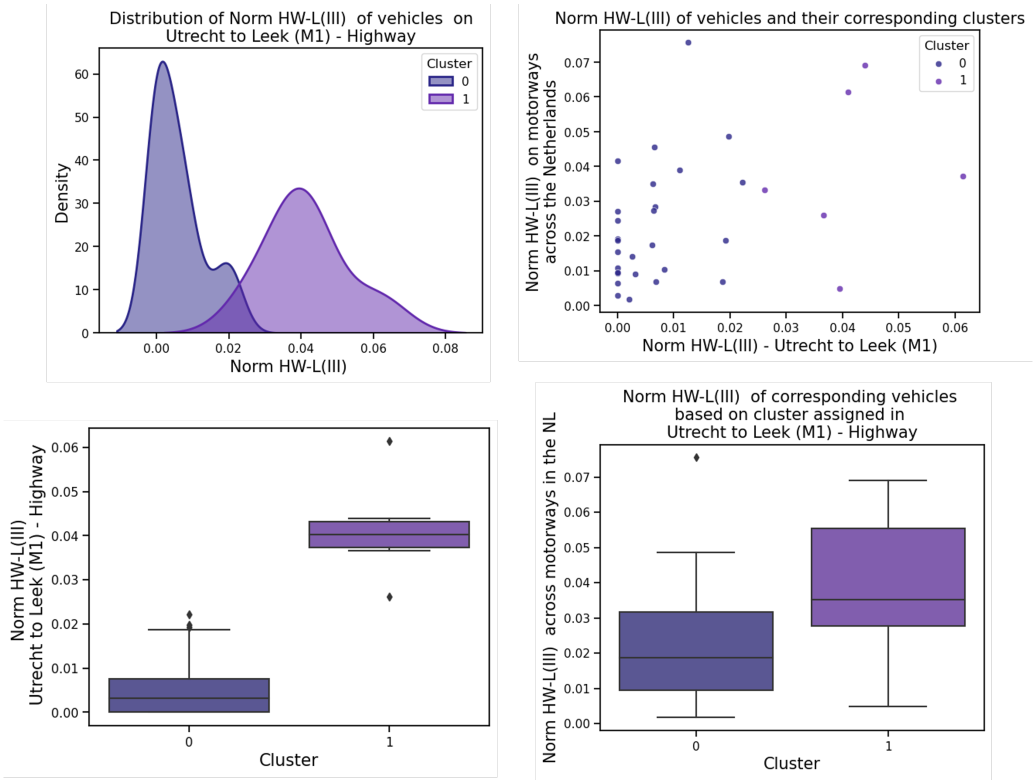

Fig. 68: Clustering Analysis - Normalized Headway Warnings Level-III on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

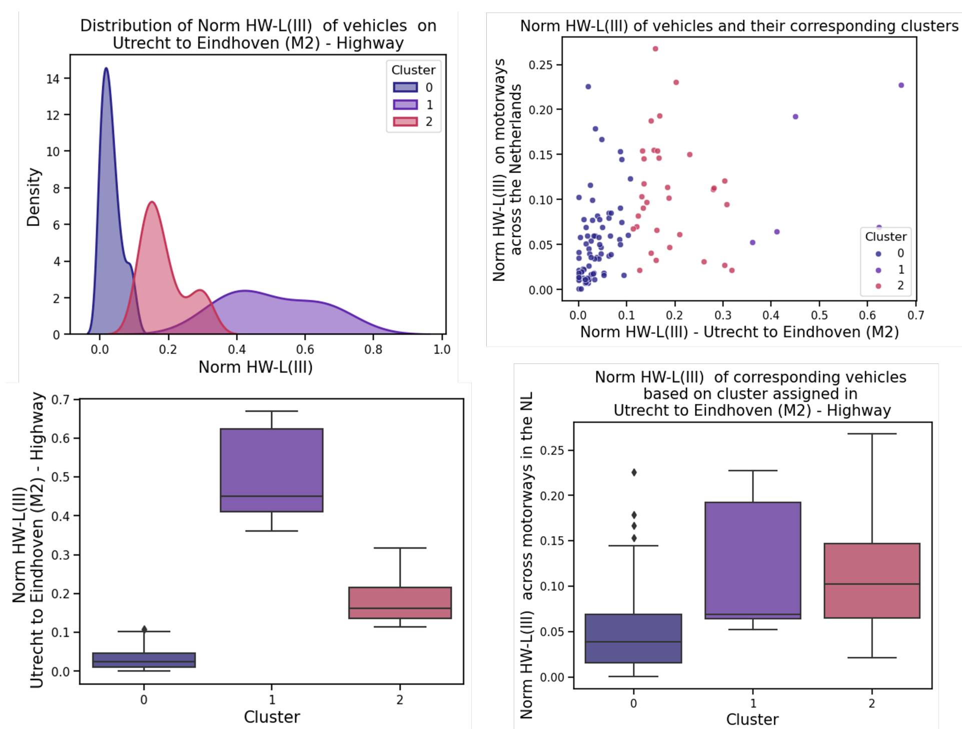

Fig. 69: Clustering Analysis - Normalized Headway Warnings Level-III on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

Contrary to headway warnings in urban areas, stability is weaker on motorways. Trucks with high engine power are consistently associated with low headway warnings in the case of urban areas.

4) Normalized Lane Departure Warnings-Left and Right: This section presents the clustering analysis results for Normalized Lane Departure Warnings-Left and Right on two different motorways and their corresponding trips across motorways in the Netherlands.

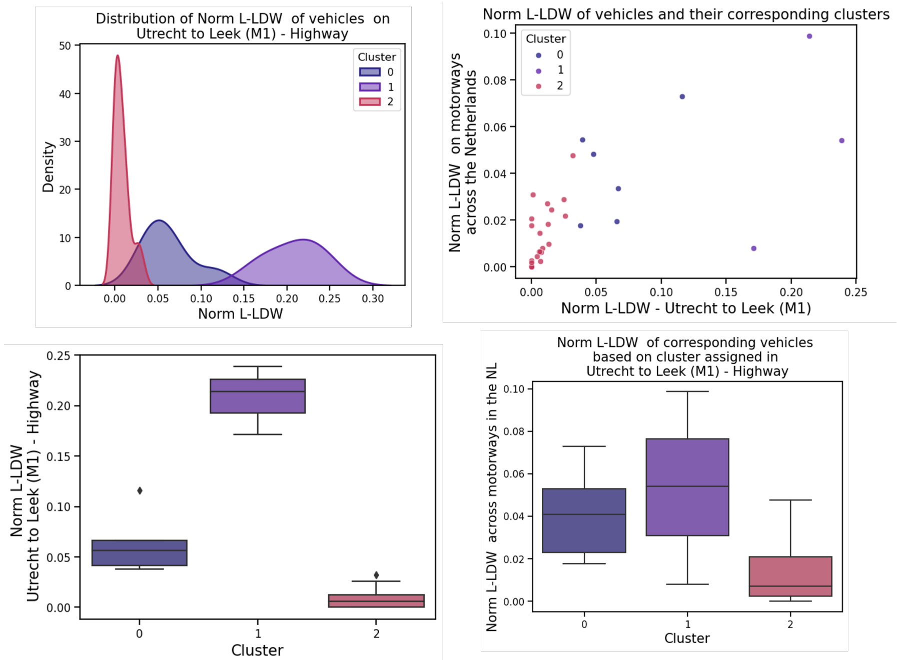

Fig. 70: Clustering Analysis - Normalized Lane Departure Warnings-Left on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

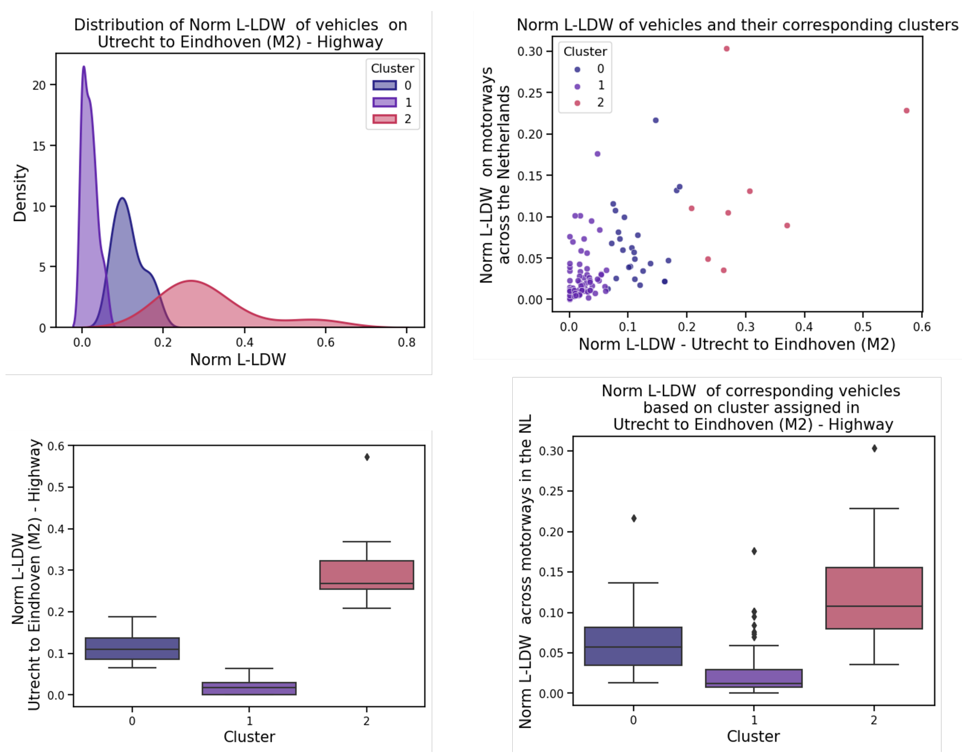

Fig. 71: Clustering Analysis - Normalized Lane Departure Warnings-Left on Utrecht-Eindhoven Motorway and corresponding trips across motorways the Netherlands

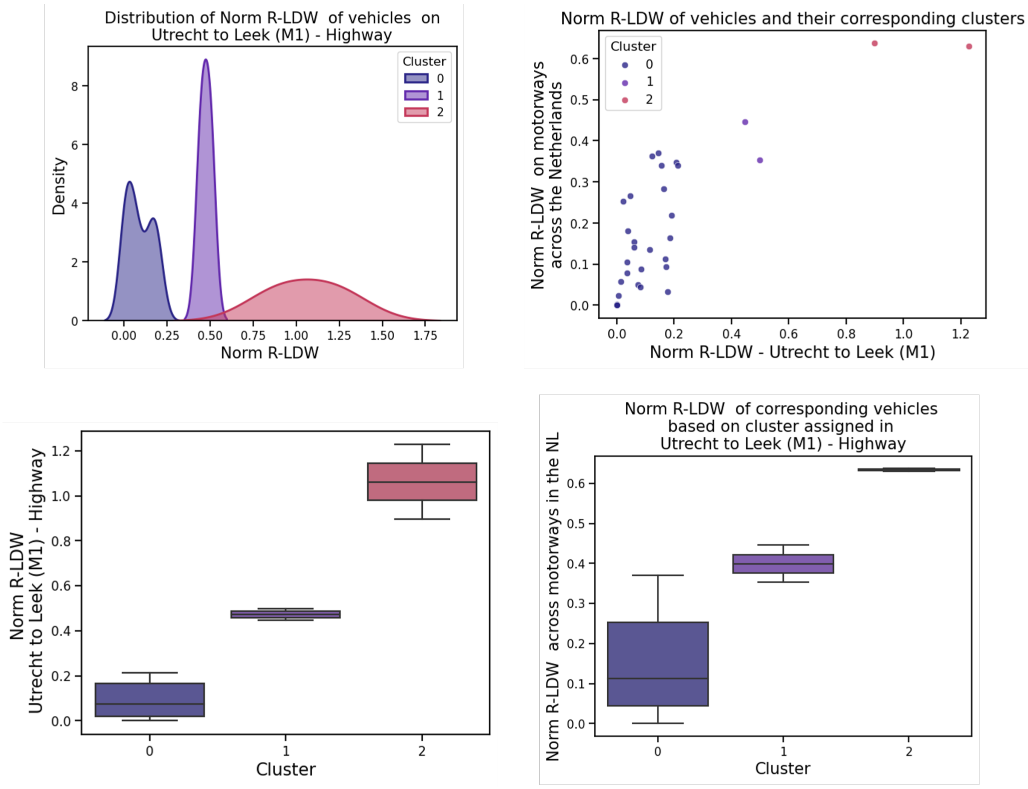

Fig. 72: Clustering Analysis - Normalized Lane Departure Warnings-Right on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

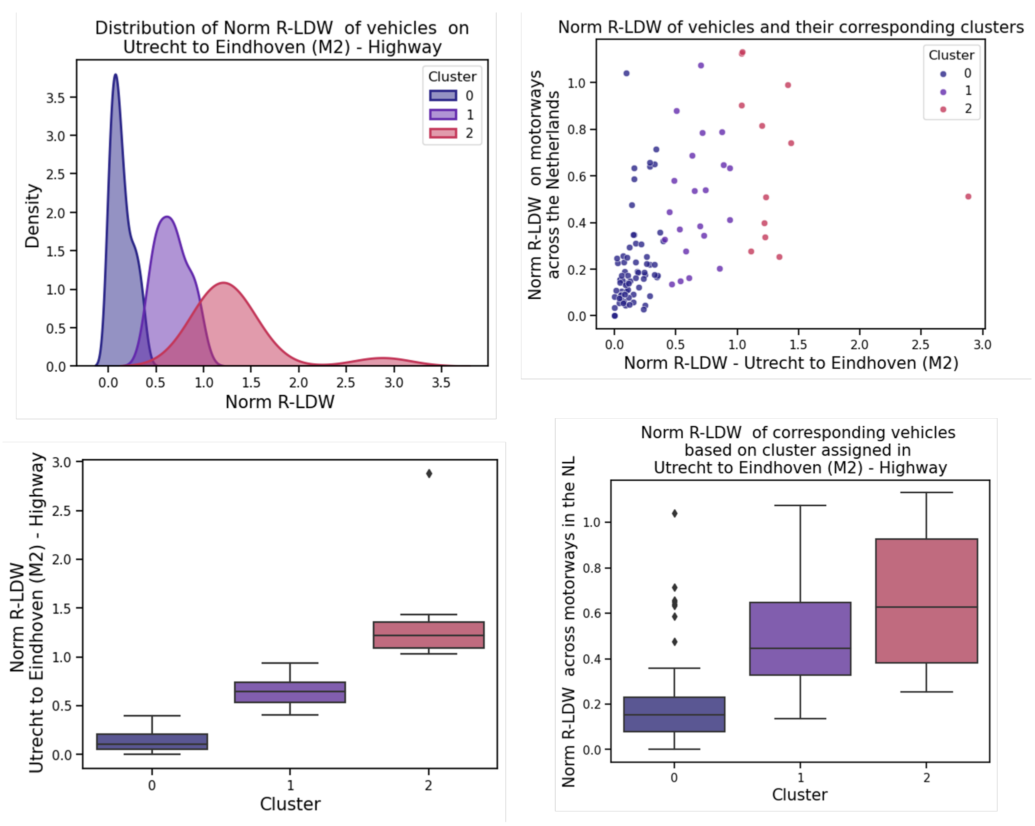

Fig. 73: Clustering Analysis - Normalized Lane Departure Warnings-Right on Utrecht-Leek Motorway and corresponding trips across motorways the Netherlands

Lane departure warnings have higher stability on motorways compared to urban areas.