



# IVORY

AI for Vision Zero in Road Safety



## D5.4 Methodologies for road user assistance and human-vehicle-environment interactions



This project has received funding from the European Union's Horizon Europe Research and Innovation programme under the Marie Skłodowska-Curie grant agreement No 101119590. The opinions expressed in this document reflect only the authors' view and in no way reflect the European Commission's opinions. The European Commission is not responsible for any use that may be made of the information it contains.



[www.ivory-network.eu](http://www.ivory-network.eu)

## Project identification

<b>Grant Agreement No</b>	101119590
<b>Acronym</b>	IVORY
<b>Project Title</b>	AI for Vision Zero in Road Safety
<b>Start Date</b>	01/11/2023
<b>End-Date</b>	30/10/2027
<b>Project URL</b>	<a href="http://www.ivory-network.eu">www.ivory-network.eu</a>

## Document summary

<b>Deliverable No</b>	5.4
<b>Deliverable Title</b>	Methodologies for road user assistance and human-vehicle-environment interactions
<b>Work Package</b>	5
<b>Contractual due date</b>	October 31, 2025
<b>Actual submission date</b>	October 30, 2025
<b>Nature</b>	Report
<b>Dissemination level</b>	PUB - Public
<b>Lead Beneficiary</b>	UH
<b>Responsible Author</b>	Tom Brijs (UH)
<b>Contributions from</b>	Tom Brijs, Shi Qiu, Ajay Iyer, Mohammad Pashae, Aristotelis Tsoutsanis, Aristotelis Styanidis

## Document referencing

Brijs T., Qiu S., Iyer A., Pashae M., Tsoutsanis A., Styanidis A. (2025), Methodologies for road user assistance and human-vehicle-environment interactions. Deliverable 5.4 of the MSCA IVORY Project.

## AI usage disclosure

Parts of this document have been created with assistance from AI tools. The content has been reviewed and edited by a human. For more information on the extent and nature of AI usage, please contact the authors.

# Table of Contents

<b>PROJECT IDENTIFICATION</b> .....	<b>2</b>
<b>DOCUMENT SUMMARY</b> .....	<b>2</b>
<b>DOCUMENT REFERENCING</b> .....	<b>2</b>
<b>AI USAGE DISCLOSURE</b> .....	<b>2</b>
<b>TABLE OF CONTENTS</b> .....	<b>3</b>
<b>LIST OF FIGURES</b> .....	<b>6</b>
<b>LIST OF TABLES</b> .....	<b>6</b>
<b>LIST OF ABBREVIATIONS</b> .....	<b>6</b>
<b>EXECUTIVE SUMMARY</b> .....	<b>9</b>
<b>1 INTRODUCTION</b> .....	<b>12</b>
1.1 ABOUT THE PROJECT.....	12
1.2 ABOUT THIS DELIVERABLE .....	12
1.3 OVERARCHING RESEARCH GOALS FOR WP5 .....	13
1.4 THEMATIC AREAS .....	14
1.5 REPORT OUTLINE .....	14
<b>2 STATE-OF-THE-ART REVIEW OF METHODOLOGIES</b> .....	<b>15</b>
2.1 METHODOLOGIES FOR MULTIMODAL DATA INTEGRATION AND FEATURE ENGINEERING .....	15
2.1.1 <i>Driver profiling using multimodal data</i> .....	15
2.1.2 <i>Preventing driver distraction and drowsiness: A multimodal data integration approach</i> .....	16
2.1.3 <i>Multimodal learning in drowsiness detection</i> .....	17
2.1.4 <i>Integrating telematics, traffic, and infrastructure data for context-aware safety assessments</i> ..	20
2.2 METHODOLOGIES FOR DRIVER STATE, BEHAVIOUR, AND SAFETY PREDICTION .....	21
2.2.1 <i>Predicting driver safety: Methodological approaches for driver state and behaviour assessment for distraction and drowsiness</i> .....	22
2.2.2 <i>Driver drowsiness detection</i> .....	22
2.2.3 <i>Behaviour modelling: from individual risky actions to resilient or optimal driving</i> .....	23
2.3 METHODOLOGIES FOR PROFILING, COACHING, AND DRIVER ASSISTANCE.....	24
2.3.1 <i>Personalized driver profiling and coaching</i> .....	24
2.3.2 <i>Smart driver development: Personalized profiling and driver assistance to prevent driver distraction and drowsiness</i> .....	25
2.3.3 <i>Personalised profiling using telematics, traffic and infrastructure data</i> .....	26
2.4 METHODOLOGIES FOR ETHICS, FAIRNESS, AND HUMAN-CENTRIC DESIGN IN ROAD SAFETY .....	27
2.4.1 <i>Ethical Road Safety: Human-Centric Design Principles for Driver Monitoring</i> .....	27
2.4.2 <i>Federated learning</i> .....	29

<b>3</b>	<b>OVERVIEW OF DATA SOURCES FOR IVORY DCS.....</b>	<b>30</b>
3.1	DATA SOURCE OVERVIEW .....	30
3.2	THE PREVENTION DATASET .....	30
3.2.1	<i>Description/origin</i> .....	30
3.2.2	<i>Availability/replicability</i> .....	31
3.3	THE I-DREAMS DATASET .....	31
3.3.1	<i>Description/origin</i> .....	31
3.3.2	<i>Availability/replicability</i> .....	31
3.4	UDRIVE DATASET .....	32
3.4.1	<i>Description/origin</i> .....	32
3.4.2	<i>Availability/replicability</i> .....	32
3.5	OSEVEN DATASET.....	32
3.5.1	<i>Description/origin</i> .....	32
3.5.2	<i>Availability/replicability</i> .....	32
3.6	CARDIOID/ISEL DATASET.....	32
3.6.1	<i>Description/origin</i> .....	32
3.6.2	<i>Availability/replicability</i> .....	33
3.7	SLEEPEYE DATASET .....	33
3.7.1	<i>Description/origin</i> .....	33
3.7.2	<i>Availability/replicability</i> .....	34
3.8	VALU3S DATASET.....	34
3.8.1	<i>Description/origin</i> .....	34
3.8.2	<i>Availability/replicability</i> .....	35
3.9	TNO DATASET .....	35
3.9.1	<i>Description/origin</i> .....	35
3.9.2	<i>Availability/replicability</i> .....	35
3.10	DATA ACCESS AND USAGE BY DCS .....	35
<b>4</b>	<b>INNOVATIVE METHODOLOGICAL APPROACHES BY IVORY DCS .....</b>	<b>37</b>
4.1	METHODOLOGIES ADOPTED BY DC 4.....	37
4.1.1	<i>Objectives of the research</i> .....	37
4.1.2	<i>Overview of methods</i> .....	38
4.1.3	<i>Research flowchart</i> .....	41
4.2	METHODOLOGIES ADOPTED BY DC 5.....	42
4.2.1	<i>Objectives of the research</i> .....	42

4.2.2	<i>Overview of methods</i> .....	43
4.2.3	<i>Research flowchart</i> .....	47
4.3	METHODOLOGIES ADOPTED BY DC 6 .....	48
4.3.1	<i>Objectives of the research</i> .....	48
4.3.2	<i>Overview of methods</i> .....	49
4.3.3	<i>Research flowchart</i> .....	49
4.4	METHODOLOGIES ADOPTED BY DC 7 .....	52
4.4.1	<i>Objectives of the research</i> .....	52
4.4.2	<i>Overview of methods</i> .....	53
4.4.3	<i>Research Flowchart</i> .....	54
4.5	METHODOLOGIES ADOPTED BY DC 14 .....	55
4.5.1	<i>Objectives of the research</i> .....	55
4.5.2	<i>Overview of methods</i> .....	56
4.5.3	<i>Research flowchart</i> .....	58
<b>5</b>	<b>WIDER METHODOLOGICAL CONTRIBUTIONS</b> .....	<b>59</b>
5.1	AI-SUPPORTED DETECTION OF ROAD USER RISKS .....	59
5.2	AI-SUPPORTED ANALYSIS OF ROAD USER RISKS .....	60
5.3	AI-SUPPORTED MITIGATION OF ROAD USER RISKS .....	62
<b>6</b>	<b>SUMMARY AND FUTURE STEPS</b> .....	<b>64</b>
6.1	SUMMARY.....	64
6.2	FUTURE STEPS .....	65
	<b>REFERENCES</b> .....	<b>67</b>

## List of figures

Figure 1: work package overview .....	14
Figure 2: DC4's research flowchart.....	41
Figure 3: flowchart of the research design to achieve research objectives .....	48
Figure 4: flowchart of DC6's research design .....	50
Figure 5: DC7 research area .....	54
Figure 6: DC14's conceptual doctoral research outline.....	58

## List of tables

Table 1: use of datasets by DCs .....	36
---------------------------------------	----

## List of abbreviations

Abbreviation	Description
AI	Artificial Intelligence
ADAS	Advanced Driver Assistance System
ANN	Artificial Neural Networks
ANOVA	Analysis of Variance
ApEn	Approximate Entropy
AUC	Area under the curve
CLT	Cognitive Load Theory
CNN	Convolutional Neural Network
DBN	Dynamic Bayesian Network
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DC	Doctoral Candidate
DDD	Driver Distraction and Drowsiness
DFA	Detrended Fluctuation Analysis
DL	Deep Learning
DMS	Driver Monitoring Systems
EC	European Commission
ECG	Electrocardiogram
ECTS	European Credit Transfer System

EEG	Electroencephalography
EOG	Electrooculogram
EQD	Equal Opportunity Difference
EU	European Union
EURO NCAP	The European New Car Assessment Programme
FCW	Forward Collision Warning
FL	Federate Learning
GDPR	General Data Protection Regulation
GNN	Graph Neural Networks
GPS	Global Positioning System
GSR	General Safety Regulations
HMI	Human-Machine Interface
HRV	Heart Rate Variability
IRL	Inverse Reinforcement Learning
ITS	Intelligent Transportation Systems
KPI	Key Performance Indicators
KSS	Karolinska Sleepiness Scale
LiDAR	Light Detection And Ranging
LSTM	Long Short-Term Memory
LLM	Large Language Model
MSCA	Marie Skłodowska-Curie Actions
ML	Machine Learning
MRT	Multiple Resource Theory
NDS	Naturalistic Driving Study
NTUA	National Technical University of Athens
OBD	On Board Diagnostics
PERCLOS	The percentage of time that the eyes are more than 80% closed
PET	Post-Encroachment Time
PRP	Predictive Rate Parity Difference

PSD	Proportion of Stopping Distance
RMSSD	Root Mean Square of Successive Differences
RNN	Recurrent Neural Network
RQ	Research Question
SAE	Society of Automotive Engineers
SampEn	Sample Entropy
SCE	Safety Critical Event
SDNN	Standard Deviation of Normal-to-Normal intervals
SEM	Structural Equation Modelling
SHAP	SHapley Additive exPlanations
SSL	Self Supervised Learning
SSM	Surrogate Safety Measures
STZ	Safety Tolerance Zone
SVM	Support Vector Machine
SWA	Steering Wheel Acceleration
STA	Subject to Availability
TGAT	Temporal Graph Neural Networks
TNO	The Netherlands Organisation for Applied Scientific Research
TTC	Time to collision
TU Delft	Delft University
UHASSELT	Hasselt University
VTI	The Swedish National Road and Transport Research Institute
V2X	Vehicle To Everything
WP	Work Package
XAI	Explainable AI

### Project Context and Objectives

The IVORY project — *AI for Vision Zero in Road Safety* — is a Marie Skłodowska-Curie Industrial Doctorate Network funded under the Horizon Europe programme (Grant Agreement No. 101119590). It aims to advance the integration of Artificial Intelligence (AI) into road safety research and to train a new generation of researchers in responsible and ethical AI for transport safety. The project contributes directly to the European Commission’s *Vision Zero* target and the UN Sustainable Development Goal 3.6 of halving road traffic deaths by 2030.

Within this framework, **Work Package 5 (WP5)**, *Road User Assistance*, develops innovative AI-based methodologies for understanding, predicting, and improving driver behaviour and human–vehicle–environment interactions. The work focuses on four overarching themes:

1. Multimodal data integration and feature engineering,
2. Driver state, behaviour, and safety prediction,
3. Personalised profiling, coaching, and driver assistance, and
4. Ethics, fairness, and human-centric design in AI for road safety.

Deliverable 5.4 presents methodologies emerging from five doctoral research projects (DC4, DC5, DC6, DC7, DC14), each contributing to novel data-driven safety assessment and intervention techniques. While still *work in progress*, these efforts collectively advance the project’s second research goal: *developing new ways of supporting road users and improving human–vehicle–environment interaction through AI*.

### Methodological Framework and Innovations

#### ***Multimodal Data Integration and Feature Engineering***

A key methodological pillar of WP5 is the integration of multimodal data from heterogeneous sources — in-vehicle sensors, telematics, video, physiological signals, and contextual traffic and infrastructure data — to derive meaningful indicators of driver state and safety. This integration requires rigorous preprocessing (synchronisation, harmonisation, cleaning) followed by advanced feature engineering. Derived features such as head pose, blink rate, steering wheel variability, heart rate variability (HRV), or time-to-collision (TTC) serve as crucial inputs for machine learning models. This holistic data fusion enables richer representations of driver behaviour and risk than any single modality could provide.

#### ***Driver State, Behaviour, and Safety Prediction***

Building on the engineered data, predictive methodologies combine classical statistical analysis with advanced AI models — including convolutional and recurrent neural networks (CNNs, LSTMs), transformers, and graph neural networks. These models detect and forecast risky behavioural states, particularly driver distraction and drowsiness (DDD), across different levels of vehicle automation. Predictive frameworks capture temporal dependencies among

behavioural, physiological, and vehicular signals, enabling early detection of precursors to safety-critical events. Complementary research also models *positive* or *resilient* driving, expanding safety research from a failure-prevention focus (Safety-I) to an adaptation and success-based perspective (Safety-II).

### ***Personalised Profiling, Coaching, and Assistance***

The transition from descriptive analytics to proactive, individualised support represents another core innovation. By combining telematics, video, and physiological data, AI systems can generate personalised driver profiles that evolve with context and time. More specifically, these profiles support tailored feedback and coaching, including intelligent recommendation of post-trip videos and behaviour summaries. Adaptive feedback mechanisms, implemented through connected apps or in-vehicle systems, can foster continuous driver improvement. Moreover, multimodal models allow integration into Advanced Driver Assistance Systems (ADAS), supporting real-time alerts or adaptive vehicle responses based on individual risk patterns.

### ***Ethics, Fairness, and Human-Centric AI***

Given the sensitive nature of behavioural and physiological data, WP5 embeds ethical and human-centric design principles across all methodological activities. The deliverable highlights strategies to ensure fairness, transparency, and privacy in AI-driven safety systems. Fairness-aware algorithms are employed to mitigate demographic biases (e.g., gender, age, cultural differences) in model predictions. Explainable AI (XAI) methods enhance transparency, providing interpretable insights into model decisions, while Federated Learning offers a privacy-preserving framework that allows distributed model training without centralising sensitive data. These practices align with the EU Artificial Intelligence Act's emphasis on trustworthy, human-centric AI.

## **Doctoral Research Contributions**

Each doctoral project contributes a specific methodological innovation:

- **DC4** develops AI methods for analysing dashcam videos to detect risky interactions such as cut-ins and tailgating, constructing dynamic driver profiles and personalised coaching systems.
- **DC5** explores multimodal AI models to detect and predict driver distraction and drowsiness across automation levels, addressing data fusion, precursor detection, and bias mitigation.
- **DC6** analyses the full spectrum of driver behaviour—from unsafe to optimal driving—to identify adaptive patterns and inform AI-driven behaviour modelling.
- **DC7** focuses on fusing telematics, infrastructure, and traffic data for holistic driver assistance and context-aware risk assessment.
- **DC14** investigates the use of ethically sound physiological measurements (EEG, ECG) for predicting road safety outcomes, advancing privacy-preserving federated learning models.

Together, these projects demonstrate a coherent methodological ecosystem combining advanced AI, behavioural science, and ethical governance to create scalable, equitable, and human-centred road safety solutions.

## Conclusions and Future Outlook

Deliverable 5.4 consolidates the methodological foundation for AI-based road user assistance within IVORY. The research contributes to:

- a **new conceptualisation of road safety as a behavioural spectrum**,
- **novel multimodal data fusion** and **privacy-preserving algorithms**,
- **AI software tools** for driver profiling and distraction/drowsiness detection, and
- **ethical frameworks** ensuring fairness and transparency in safety-critical applications.

Future deliverables (D5.5–D5.9) will expand upon these preliminary methodologies with empirical validations, model implementations, and field demonstrations. The expected outcomes include open-access code and datasets, validated AI models for driver assistance, and deployable prototypes contributing directly to Europe’s Vision Zero ambitions.

# 1 Introduction

## 1.1 About the project

IVORY - 'AI for Vision Zero in Road Safety' is an MSCA Industrial Doctorates Network aiming to develop a new framework for optimal integration of Artificial Intelligence (AI) in road safety research, and to train a new generation of leading researchers in the field. It addresses the UN Sustainable Development Goals target 3.6 and the EC Vision Zero strategy, of halving traffic fatalities by 2030 and eliminating them by 2050.

IVORY addresses the lack of common understanding of the challenges and opportunities of AI for road safety by means of 4 research goals, aiming to develop:

- (i) responsible, fair and impactful AI for road safety,
- (ii) new ways of road user support and human-vehicle-environment interaction,
- (iii) new scalable and equitable AI technologies for proactive infrastructure safety management,
- (iv) a sustainable knowledge sharing network on AI for road safety.

IVORY outputs will not only provide more robust user support through AI in vehicle automation, but will also allow to responsibly and proactively manage the persistent problems of existing conventional, low automation transport systems, so that new opportunities for global road safety impact can emerge. Moreover, IVORY adopts a design for values approach for AI in road safety, operationalising the ethical principles of justice and explainability, and providing efficient AI solutions also for disadvantaged groups (e.g. vulnerable road users, low-to-middle-income countries).

IVORY consists of 4 academic and 8 non-academic beneficiaries, and 8 associated partners, joining from engineering, data science and ethics of technology disciplines, from 11 countries. 15 young researchers will receive high-level doctoral education, industrial exposure, local training, and 8.5 ECTS of network-wide training on key advanced, core and transferable skills. IVORY will create an online learning & networking platform for AI in road safety, to be available after the end of the project for future researchers in this field.

## 1.2 About this deliverable

More specifically, the work of the following Doctoral Candidates (DCs) is included in this deliverable, the topics of which find common thematic ground in novel proactive road safety approaches:

- DC 4: Shi Qiu - Road user profiling using multimodal data of naturalistic driving databases
- DC 5: Ajay Iyer - AI to mitigate driver distraction and drowsiness at different levels of automation
- DC 6: Mohammad Pashae - Learning from the whole spectrum of driver behaviour: from unsafe to optimal driving
- DC 7: Aristotelis Tsoutsanis - Data fusion of traffic, behaviour & infrastructure for holistic driver assistance
- DC 14: Aristotelis Styanidis: Road safety prediction on the basis of ethically sound physiological measurements

It has to be emphasized that the contributions described in this document are based on the current status of the research carried out by the DCs (about 1 year progress in their PhD studies). Hence, the contributions in this deliverable have to be viewed as ‘work in progress’. In the future, when their PhD work evolves, updates and new insights will be added in later progress and methodological reports (D5.5, D5.6, D5.7, D5.8 and D5.9).

### 1.3 Overarching research goals for WP5

Work Package 5 (WP5), entitled ‘Road User Assistance’ will advance road user profiling and coaching tools, by developing new Artificial Intelligence (AI) models for data fusion and automatic creation and feedback of driver profiling, on the basis of advanced video and spatio-temporal data from naturalistic driving (DC4) and from driver telematics, traffic and street map data (DC7). It will study for the first time the impact of key human factors such as driver distraction and drowsiness at higher SAE levels of automation, by enhancing current detection algorithms with video and physiological data (DC5, DC14). At the same time, it will introduce a new paradigm in driver support by Machine Learning (ML) from the whole spectrum of driving behaviour analytics, including “successful” interactions, and not just event-based warnings data (i.e., based on occasions where a risk is detected (DC6) – in line with Research Goal 2 (RG2) of the project (To develop new ways of supporting road users and human-vehicle-environment interaction by means of AI).

The research in this WP aims to contribute to the following innovations and key exploitable results.

Innovations:

- New definition of road safety as a spectrum (from normal driving to high risk)
- Novel data mining techniques for road user profiling
- Driver behaviour assessment in different levels of automation
- Novel data fusion techniques & algorithms to preserve privacy

Key exploitable results:

- Open-access backend data and code
- Road safety-as-a-spectrum theoretical model
- In-vehicle gaze detection and prediction
- AI software for road user profiling
- Enhanced data fusion and protection

The work in WP5 fits within the larger framework of the project as shown in Figure 1, where WP4, WP5 and WP6 focus on the actual research within the IVORY project. WP4 addresses topics such as responsible, fair and impactful AI, whereas WP6 focuses on proactive infrastructure safety management.

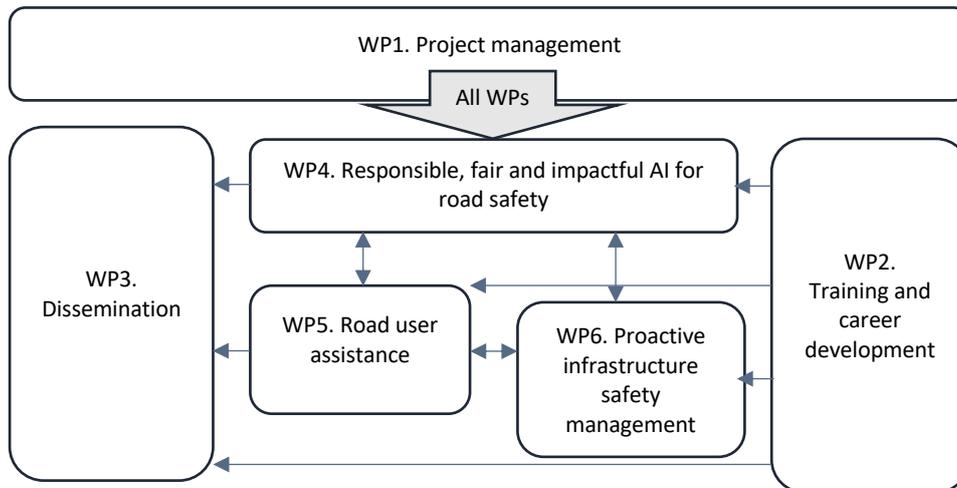


FIGURE 1: WORK PACKAGE OVERVIEW

## 1.4 Thematic areas

We have chosen to group the methodological work of the different DCs in this WP5 into the following overarching thematic areas. Each of the DCs contributes to one or more of these thematic areas, as indicated below. A more in-depth coverage of each of these themes is presented in section 2:

- Methodologies for multimodal data integration and feature engineering (DC4, DC5, DC7, DC14)
- Methodologies for driver state, behaviour, and safety prediction (DC5, DC14)
- Methodologies for profiling, coaching, and driver assistance (DC4, DC5, DC7)
- Methodologies for ethics, fairness, and human-centric design in road safety (DC5, DC14)

## 1.5 Report outline

In section 2, a state-of-the-art overview based on a literature synthesis will be presented of methodologies in each of the thematic areas presented in Section 1.4. Section 3 provides an overview of data sources that are used by the DCs to build their models and/or empirically validate their research. Next, section 4 presents how each of the DCs goes beyond this state-of-the-art by presenting the methodologies they work on in their research. Section 5 presents wider methodological contributions from the DCs' work in the IVORY project. Finally, section 6 presents a summary of this deliverable and provides an overview of future steps.

## 2 State-of-the-art review of methodologies

The state-of-the-art review of methodologies provides a structured overview of the diverse approaches currently shaping research and innovation in driver behaviour analysis and road safety. It highlights the technical, analytical, and ethical foundations required to advance this field, from data integration to human-centered design. First, methodologies for multimodal data integration and feature engineering address the challenge of processing and harmonising heterogeneous sources, enabling the extraction of meaningful indicators for behaviour and safety. Building on these inputs, methodologies for driver state, behaviour, and safety prediction explore AI and ML based techniques for detecting, modelling, and forecasting risky and optimal driving patterns. Next, methodologies for profiling, coaching, and driver assistance translate analytical insights into actionable feedback and support, enabling safer and more resilient driver practices through individualised interventions and real-time applications. Finally, the review considers methodologies for ethics, fairness, and human-centric design, emphasizing the need for transparency, inclusivity, and trust in developing and deploying AI systems in road safety. Together, these methodologies form a complementary framework where multimodal data integration provides the foundation for predictive modelling, which in turn informs personalized coaching and assistance, all guided by ethical and human-centric principles to ensure trustworthy and effective AI applications in road safety.

### 2.1 Methodologies for multimodal data integration and feature engineering

The integration of multimodal data streams and subsequent feature engineering has emerged as a foundational paradigm in transport safety research, predicated on the principle that a holistic understanding of human-vehicle-environment interactions cannot be achieved through any single data modality alone. This approach synthesizes heterogeneous data sources, including in-vehicle visual signals, physiological indicators, vehicle kinematics, and contextual data from infrastructure and real-time traffic feeds. A critical precursor to effective fusion is the meticulous pre-processing of these disparate datasets, which involves rigorous synchronization, harmonization, and cleaning to ensure the temporal alignment and feature comparability required for sophisticated modelling. Following data preparation, sophisticated feature engineering is employed to transform raw data into meaningful indicators of driver state and risk. This process encompasses a wide spectrum of techniques, from the calculation of established behavioural Key Performance Indicators (KPIs), Surrogate Safety Metrics (SSMs) and eye glance behaviour, to the extraction of complex physiological markers such as Electrocardiogram (ECG) and Electroencephalogram (EEG). These engineered features serve as the essential inputs for methodologies and models targeting driver state and behaviour in the context of safety prediction. The following sub-sections explore these core methodologies further.

#### 2.1.1 Driver profiling using multimodal data

Recent research emphasizes the integration of diverse sensory sources, such as onboard telematics (e.g., GPS, acceleration), cabin and road-facing video, and physiological signals, as the foundation for personalized driver behaviour profiling. Approaches like Multimodal Driver State Modeling decompose scenarios using techniques like Bayesian change-point detection

paired with Latent Dirichlet Allocation to unearth hidden driving states, such as harsh braking, abnormal heart rate, or gaze entropy (Tavakoli & Heydarian, 2022). Similarly, for multimodal input fusion, Interwoven Convolutional Neural Networks incorporate multiple video streams and optical flow for robust real-time driver behaviour classification, blending features across modalities through fusion layers (Zhang et al., 2020).

Personalization was advanced significantly via Inverse Reinforcement Learning (IRL), which infers latent driving preferences from real human behaviour traces. A maximum-entropy IRL framework was applied to Next Generation Simulation highway data, learning discrete intent-driven reward functions that capture individual trajectory choices, yielding better "human likeness" than traditional models (Huang et al., 2022). In another line of work, other methods, such as logistic-regression-based IRL, efficiently and reproducibly estimate a driver's latent style, enabling the reconstruction of personalized trajectories (Jiang et al., 2018).

Beyond behaviour classification, behavioural embeddings also enable driver identification. For example, Driver2vec was introduced to learn compact representations of driver-specific patterns via Temporal Convolutional Networks and triplet loss optimization, enabling identification from only 10 seconds of driving data with over 80% accuracy (Yang et al., 2021).

### **2.1.2 Preventing driver distraction and drowsiness: A multimodal data integration approach**

Considering the overarching goals of WP5 the key challenge in developing AI-based strategies for DDD mitigation is to harness the full potential of heterogeneous data sources available. This is essential for delivering on key project innovations and exploitable results, including 'Driver behaviour assessment in different levels of automation', 'Novel data fusion techniques' and 'AI software for road user profiling.' There have been significant advancements in data fusion techniques used in ITS, such as observation-level fusion, feature-level fusion, and decision-level fusion, but limited efforts have been made to employ these techniques in DDD analysis. Observation-level fusion is used in the early stages to combine raw data from multiple sensors (e.g. raw video frames, gyroscope data etc.) generally already used in ADAS systems. Feature-level fusion focusses on combining higher level engineered features from different features (e.g. head pose, vehicle speed, acceleration etc.). Finally decision level fusion combines outputs predictions of multiple models that maybe trained on different data-streams (e.g. One model using PERCLOS from video data and another using physiological data such as HRV to predict the probability of drowsiness) Furthermore, individual data modalities have been explored for detecting these dangerous driving states (Ramzan et al., 2019), whereas effectively combining these heterogeneous sources in the domain of DDD remains a critical research gap. As highlighted by (National Research Council, 2013) and (Durrant-Whyte and Henderson, 2008) that emphasize the importance of multi-sensor data fusion in robotics, a concept directly applicable to driver monitoring domain. However, the optimal fusion strategy for maximizing the accuracy of AI models in predicting driver distraction and drowsiness remains an open question. The current research approach does not fully utilize the existing DDD data and its underlying relationships using optimal data fusion strategies. Finding an optimal strategy would involve focusing on promising candidates

like DL-based feature fusion, which leverages neural networks to fuse video-based driver data kinematics and physiological sources or ensemble learning, a decision-level technique by combining the outputs of multiple specialized models. This optimal fusion approach will maximize the accuracy and reliability of AI-based mitigation systems across various driving scenarios and individual driver characteristics.

The proposed research methodologies are anchored in the fusion of diverse and heterogeneous data streams to create a comprehensive understanding of the driver-vehicle-environment nexus. This involves combining data from multiple sources, including in-vehicle video capturing visual signals such as facial expressions, eye movements, and head pose; vehicle kinematics data from telematics sensors measuring steering wheel inputs, acceleration, and braking patterns; and physiological indicators like PERCLOS (Percentage of eye closure) and blink frequency.

With a harmonised dataset, the methodological focus can shift to sophisticated feature engineering, such as extracting PERCLOS more than 80% for drowsiness, and head positions (yaw, pitch, etc.) which useful to demonstrate significant differences in safe and distracted driving, highlighted in the review paper by (Ounoughi and Ben Yahia, 2023). This process is central to identifying the systematic patterns of driver behaviour linked to distraction and drowsiness. This is achieved by using two distinct but related categories of metrics. The first one is KPI which is a metric that describes “what the driver is doing” during the distracted and drowsy state. Hence driver based KPIs like "Eyes off Road", “hands of wheel” directly quantify driver behavior and vehicle-based indicators such as Forward Collision Warning (FCW) which are not a direct measure of driver state but serves as an important indicator of a critical performance outcome during these states. The second metric is SSMs that quantify the objective risk of the traffic situation. An SSM like PET or TTC is calculated purely from the kinematics of the interacting road users providing an objective measure of "how close" the situation was to a crash. The holistic view of distraction and drowsiness emerges from quantifying the endogenous relationship between these two sets of metrics. For example, a demonstration that an increase in the "Eyes off Road" (KPI) which systematically precedes a decrease in the TTC (SSM), connects the driver's inattentive state directly to a quantifiable increase in objective driving risk. These engineered features—spanning behavioural, physiological, vehicular, and contextual domains—serve as the essential inputs for the advanced analytical models.

### **2.1.3 Multimodal learning in drowsiness detection**

Driver monitoring plays a critical role in ensuring road safety by continuously assessing a driver’s state and detecting conditions that may impair driving performance, such as distraction, fatigue, or drowsiness. Drowsiness has been identified as a major contributing factor in road accidents worldwide, as it degrades reaction time, attention, and decision-making. The goal of driver drowsiness detection is therefore to automatically identify early signs of reduced alertness and issue timely warnings or corrective interventions. Achieving reliable detection, however, remains a significant challenge due to the complex and dynamic nature of human fatigue and its subtle manifestations across different individuals and driving conditions.

Driver drowsiness is inherently a multifaceted phenomenon, with behavioural, physiological, and contextual signals all carrying complementary information. As a result, multimodal learning and fusion strategies have become a central research direction. By integrating facial expressions (Li et al., 2019), eye movements (Min et al., 2022), steering behaviour (Martensson et al., 2019), or biosignals like EEG (Shahbakhti et al., 2022; Guo et al., 2025), Electrooculogram (EOG) (Huo et al., 2016) and ECG (Du et al., 2023), these approaches aim to increase robustness and reduce false detections caused by relying on a single modality. Fusion can occur at different levels — feature-level, decision-level, or through deep neural architectures that jointly learn cross-modal representations.

Feature engineering and representation learning remain crucial for detection performance. Earlier approaches focused on manually extracted features (e.g., eye blink frequency, heart rate variability, steering entropy) combined with classic classifiers such as Support Vector Machines (SVMs) or Random Forests. More recent methods increasingly rely on Deep Learning (DL) architectures — convolutional neural networks (CNNs) (Krizhevsky et al., 2012), recurrent models such as Long Short-Term Memory (LSTMs) (Hochreiter & Schmidhuber, 1997), and transformers (Vaswani et al., 2017))— which learn hierarchical feature representations directly from raw data.

In video-based drowsiness detection, feature engineering predominantly targets ocular and facial behaviour, as these are reliable external manifestations of vigilance decline. Widely utilized metrics include blink frequency, eye closure duration, and PERCLOS (percentage of eye closure over time), which have been shown to correlate strongly with fatigue-related cognitive states. Additional indicators such as yawning frequency, head nodding, and gaze dispersion further enrich the feature space by capturing both micro- and macro-level behavioural cues. Beyond handcrafted features, modern approaches also leverage facial landmark tracking and DL-based embeddings to extract higher-order representations of head pose and facial expression dynamics, thereby enabling the system to capture subtle temporal changes associated with progressive drowsiness.

For physiological signals, electrocardiogram (ECG) and electroencephalogram (EEG) recordings provide direct insight into autonomic and cortical correlates of drowsiness. In ECG-based analysis, heart rate variability (HRV) is particularly informative, as it reflects the balance between sympathetic and parasympathetic nervous system activity (Shaffer & Ginsberg, 2017). Time-domain features such as the standard deviation of normal-to-normal (NN) intervals (SDNN) and root mean square of successive differences (RMSSD) reflect overall variability and short-term parasympathetic activity, respectively, offering a straightforward means of capturing fatigue-induced reductions in cardiac flexibility. Frequency-domain measures (e.g., low-frequency [LF: 0.04–0.15 Hz] and high-frequency [HF: 0.15–0.4 Hz] power components, and their ratio LF/HF) provide a spectral decomposition of HRV, enabling finer discrimination of autonomic shifts toward sympathetic dominance during fatigue. Beyond linear measures, non-linear HRV features—such as Poincaré plot indices (SD1, SD2), approximate entropy (ApEn), sample entropy (SampEn), detrended fluctuation analysis (DFA), and correlation dimension—capture the complexity and irregularity of cardiac dynamics,

offering enhanced sensitivity to the subtle physiological dysregulation that accompanies drowsiness.

In EEG-based analysis, spectral features derived from specific frequency bands are widely recognized as reliable neural markers of alertness. The theta band (4–7 Hz) is typically associated with drowsiness and cognitive slowing, while the alpha band (8–13 Hz) power increases during relaxed wakefulness and light fatigue. Conversely, beta activity (13–30 Hz), which reflects active engagement and cortical arousal, tends to decrease with reduced vigilance. Monitoring relative power shifts across these bands, or computing band ratios such as theta/alpha or theta/beta, provides a robust electrophysiological signature of drowsiness onset. Although EEG offers high sensitivity and provides direct access to cortical correlates of fatigue, its practical deployment remains limited due to sensitivity to motion artifacts, the need for stable electrode placement, and reduced user comfort in long-term use. Recent developments in dry-electrode EEG headbands and wearable systems, however, have significantly improved usability and signal stability, prompting renewed interest in EEG-based monitoring. Consequently, the research community increasingly explores EEG either as a standalone modality or within multimodal fusion frameworks, aiming to leverage its strong discriminative power while overcoming previous constraints related to comfort and robustness. In addition to ECG and EEG, EOG provides a robust means of quantifying ocular activity, which is strongly associated with drowsiness. By measuring the corneo-retinal potential generated by eye movements and blinks, EOG enables precise extraction of temporal and frequency-domain features such as blink duration, blink frequency, inter-blink intervals, and saccadic amplitude or velocity. Slow eye movements (SEMs), characterized by reduced saccadic velocity and increased drift, are particularly indicative of drowsiness. As EOG captures eye dynamics through electrical potentials rather than visual cues, it remains effective under poor lighting or partial occlusion, offering valuable complementarity to video-based systems in real-world conditions. However, its practical deployment is limited by the need for electrode placement near the eyes, potential discomfort during long-term use, and susceptibility to signal degradation from skin–electrode impedance changes. Consequently, EOG is increasingly regarded as a supportive modality that enhances the robustness of multimodal driver monitoring frameworks rather than a standalone solution.

Complementary to behavioural and physiological modalities, vehicle-based features, particularly those derived from steering wheel angle dynamics, capture the external manifestations of impaired motor control under fatigue. Drowsy driving is typically characterized by increased variability in steering behaviour, delayed corrective actions, and reduced smoothness in trajectory maintenance. Quantitative features such as the standard deviation of steering angle, steering entropy, and corrective steering frequency have been shown to reflect degraded motor coordination and attentional lapses. By incorporating steering wheel dynamics, systems can integrate contextual evidence of compromised driving performance, thereby enhancing robustness when combined with video and physiological features.

#### 2.1.4 Integrating telematics, traffic, and infrastructure data for context-aware safety assessments

Telematics forms the empirical backbone of modern driver behaviour and safety analysis, offering a continuous, sensor-based record of how vehicles interact with the road. It encompasses data captured from onboard units, smartphones, or connected-vehicle systems, typically including high-frequency measurements of position, speed, acceleration. These data streams allow researchers to reconstruct driving trajectories and infer behavioural patterns such as harsh braking, speeding, or abrupt cornering. However, while telematics provides granular insight into vehicle dynamics, it lacks the contextual information necessary to interpret whether such manoeuvres are genuinely risky or merely adaptive responses to surrounding conditions. Without situating these kinematic signals within their traffic and infrastructural context, assessments of safety risk remain partial and potentially misleading.

To bridge these gaps, researchers are increasingly turning to multimodal data integration, in which telematics, traffic, and infrastructure datasets are combined to produce a more holistic, context-aware view of road safety. This approach does not merely stack different data types but engineers meaningful features from their alignment, harmonisation, and transformation, producing indicators that can be operationalised in safety analysis and driver assistance applications.

The integration of telematics, traffic, and infrastructure data has become a cornerstone of modern road safety analytics, enabling fine-grained behavioural assessment and network-level risk profiling. Recent research has moved away from analysing each dataset in isolation towards building multimodal fusion pipelines in which heterogeneous data streams are synchronised, harmonised, and enriched to produce actionable safety indicators (Chen et al., 2022; Xu et al., 2025).

Telematics datasets, originating from smartphones, connected vehicles, or aftermarket devices, provide high-frequency kinematic measurements such as speed, acceleration, and device sensors, these raw data can be processed to produce event flags that indicate behaviours like harsh braking, rapid acceleration, speeding, or abrupt cornering (Nikolaou et al., 2025). Infrastructure datasets, often derived from open sources such as OpenStreetMap, contribute spatial context in the form of intersection geometry, lane configuration, curvature, control type, and posted speed limits (Newson & Krumm, 2009; Quddus et al., 2003). Traffic datasets, including floating-car data and traffic feeds provide dynamic exposure metrics including average speed, delay, travel time, and congestion levels (Wang et al., 2021).

State-of-the-art fusion strategies typically establish spatial correspondence between telematics records and the road network through map-matching (Luo et al., 2017), followed by spatial joins that link each kinematic observation to its underlying segment geometry and attributes. Temporal integration with traffic feeds situates behavioural observations within prevailing flow conditions, enabling researchers to differentiate, for example, between harsh braking in free-flow conditions and the same behaviour occurring under congestion (Chen et al., 2022).

A critical part of this process is data preprocessing and feature engineering. Synchronisation aligns heterogeneous streams to a common temporal resolution, often via resampling and interpolation (Newson & Krumm, 2009). Harmonisation corrects spatial noise through probabilistic map-matching methods such as Hidden Markov Models or Kalman filtering (Luo et al., 2017). Data cleaning mitigates artefacts including GPS spikes, dropouts, or duplicated points, and applies low-pass or median filtering to inertial channels to ensure physical plausibility.

These engineered features can be categorised into three main groups. SSMs quantify proxies of crash risk in the absence of reliable crash data and include measures such as the frequency and severity of harsh braking, rapid acceleration, abrupt cornering, or traffic flow instability (Lizarazo et al., 2023; Wang et al., 2021). Contextual indicators describe the environment in which an event occurred, including intersection type, curvature, grade, control devices, number of lanes, posted speed limits, and prevailing traffic conditions (Chen et al., 2022). Behavioural key performance indicators aggregate driver events over trips, drivers, or locations and normalise them by exposure such as distance travelled or time spent driving, producing measures such as ratios of harsh events per kilometre or compliance scores for speed limits (Nikolaou et al., 2025).

By combining high-resolution telematics with infrastructure topology and dynamic traffic conditions, contemporary approaches produce a multidimensional representation of driver behaviour that captures not only what occurred, but also where and under what circumstances (Abdel-Aty et al., 2023; Xu et al., 2025). This integrated view supports applications ranging from dynamic risk mapping and targeted infrastructure interventions to the development of personalised driver assistance and feedback systems.

## **2.2 Methodologies for driver state, behaviour, and safety prediction**

Building upon the foundations of multimodal data integration and feature engineering, this section focuses on the methodological frameworks employed to analyse, interpret, and predict driver state, behaviour, and safety outcomes. Whereas the previous section addressed how diverse data streams are collected and transformed into meaningful indicators, the present section examines how these engineered features are operationalised through analytical and ML techniques to detect, classify, and anticipate driver behavioural states such as distraction, drowsiness, and adaptive responses. The methodological landscape explored here encompasses both traditional statistical modelling and contemporary AI approaches, highlighting the progression from descriptive analyses to dynamic, predictive, and context-aware systems. Collectively, these methodologies form the analytical core of intelligent transportation research, enabling proactive safety assessment and the development of driver-centric decision-support systems.

### **2.2.1 Predicting driver safety: Methodological approaches for driver state and behaviour assessment for distraction and drowsiness**

The development of robust systems to mitigate risks associated with driver behavioural states linked to distraction and drowsiness hinges on sophisticated methodologies capable of processing and integrating diverse, multimodal data streams. A primary methodological focus involves the detection and prediction of driver states by leveraging a combination of behavioural, physiological, and vehicular measurements. Initial approaches employ descriptive and statistical analyses—such as ANOVA, regression modelling, and time-series analysis—to establish foundational relationships and identify how indicators DDD manifest and evolve across different levels of vehicle automation (Cooper et al., 2023). These initial insights inform the development of more complex AI models. For instance, Multimodal ANNs are proposed as a core technique to ingest and synchronize heterogeneous inputs, such as camera data capturing eye-gaze and head pose, with vehicle kinematics like steering wheel acceleration, to accurately classify a driver's state in real-time (Krizhevsky et al., 2012; LeCun et al., 2015).

Beyond simple state detection, a critical methodological tool is using predictive safety analytics through advanced behaviour modelling. This involves moving from identifying fully manifested DDD to detecting subtle, systematic patterns that serve as reliable precursors to these high-risk states. To capture the complex temporal dynamics inherent in driving data, methodologies such as RNNs. These techniques model the evolving interrelationships between various indicators over time—for example, linking a subtle increase in lane deviation with subsequent changes in blink frequency—to predict an impending drowsiness event before it becomes critical (Jackson et al., 2016). Similarly in the case of distraction an ‘established distracted state’ where the driver's attention is demonstrably diverted from the primary driving task (Regan et al., 2013) can be investigated verifying the increased time and frequency of short glances away from the road to the center console or phone using the time-series models (Noble et al., 2021). This predictive capability is further enhanced by creating representative driver profiles through clustering algorithms and quantifying risk using SSMs like TTC (Saleh et al., 2013). These integrated methodologies collectively enable a comprehensive understanding of human factors in transportation safety.

### **2.2.2 Driver drowsiness detection**

Driver drowsiness is a major contributor to traffic accidents worldwide, reducing vigilance, slowing reaction times, and impairing decision-making. Detecting drowsiness in real time is therefore a critical goal for intelligent transportation systems and advanced driver-assistance systems (ADAS). Research in this field spans behavioural monitoring (e.g., facial expressions, eye blinks, yawning), physiological sensing (e.g., EEG, ECG, EOG), and vehicle-based measures (e.g., steering patterns, lane-keeping) (Paul et al., 2017; Li et al., 2017). These modalities capture different aspects of drowsiness, but each has limitations in terms of reliability, intrusiveness, or susceptibility to external conditions, as discussed in Section 2.1.3.

Despite substantial progress, several challenges remain. A key limitation is the scarcity of annotated datasets, since labelling driver states requires time-consuming protocols,

subjective scales such as the Karolinska Sleepiness Scale (KSS) (Shahid, A., 2012), or labor-intensive video annotations, making it costly and impractical to collect at scale. Inter-individual variability further complicates generalization, as physiological and behavioural markers of drowsiness differ widely across drivers. Environmental factors such as lighting conditions, time of day, and road type also introduce difficulties. These challenges have motivated research into multimodal fusion (Guo et al., 2025) to exploit complementary signals, and self-supervised learning (Chen et al., 2020) to leverage large amounts of unlabeled driving data. Together, these directions represent the frontier of robust and scalable drowsiness detection.

### **2.2.3 Behaviour modelling: from individual risky actions to resilient or optimal driving**

Behaviour modelling seeks to represent how drivers perceive, decide, and act within dynamic traffic environments. It provides a structured way to explain how behaviour emerges from the continuous interaction between the driver, vehicle, and environment, and how these interactions shape safety outcomes. Recent developments in AI and ML have enabled new methodologies for understanding and predicting driver behaviour, with an increasing shift from detecting risky actions (such as those presented in sections 2.2.1 and 2.2.2) to modelling resilient and optimal driving. Traditional models largely follow the Safety-I paradigm, focusing on critical events and failure prevention, but newer approaches grounded in Safety-II theory emphasise successful adaptations and safe behaviour under varying conditions (Hollnagel, 2018; Papadimitriou et al., 2022). This shift supports a broader understanding of driver performance as a spectrum of behaviours, rather than binary safety states. Research frames driving as a multi-objective optimisation process, where drivers balance safety, efficiency, comfort, and personal goals (e.g., (Salvucci, 2006; Summala, 2007)). These insights are significant for defining “optimal” or “positive” driving behaviour, which has been variously referred to as “competent driving,” “roadmanship,” or “driving efficiency,” though no consensus has been reached (Papadimitriou et al., 2022; Tejada et al., 2023).

Within this framework, AI provides a powerful tool to advance behaviour modelling by detecting, classifying, and interpreting adaptation patterns in naturalistic driving data. AI models play a critical role in predictive safety analytics by analysing data from in-vehicle sensors to identify patterns of adaptation. These include supervised models (e.g., Support Vector Machine, decision trees, Convolutional Neural Networks, Long Short-Term Memory) and unsupervised methods (e.g., clustering, Dynamic Time Warping) applied to time-series data reflecting speed, steering, acceleration, braking, and driver physiological states. For example, LSTM-based models have shown promise in identifying context-sensitive driving patterns, such as sudden braking or lane changes, that vary depending on road type and traffic density (Garcia-Constantino et al., 2014; Lee et al., 2018; Sarker et al., 2021). Some models also integrate driver state monitoring (e.g., fatigue, arousal, comfort zone thresholds) to enhance prediction accuracy.

Despite these advances, many studies continue to focus on classifying discrete events or a type of driving. This approach limits our ability to understand why drivers behave in specific ways and how safety is actively maintained through ongoing adjustments. To overcome this

limitation, the IVORY framework proposes an integrative approach that combines pattern recognition with behavioural theory, positioning AI as a diagnostic tool and a means to reveal underlying mechanisms of adaptation.

This gap highlights the need for methodologies that combine pattern recognition with behavioural theory, enabling a more meaningful interpretation of detected behaviours and their role in safety. Across the reviewed literature, a methodological trend is emerging that pairs large-scale data analysis with cognitive models of driver adaptation, providing a foundation for the research in adaptive and proactive road safety systems.

## **2.3 Methodologies for profiling, coaching, and driver assistance**

The evolution of intelligent transportation systems is increasingly moving from generic, one-size-fits-all solutions toward personalised, adaptive support for drivers. Central to this shift is the use of telematics, AI, and multimodal sensing to build individualized driver profiles that capture not only observable behaviour but also the underlying cognitive and contextual factors shaping driving performance. By analysing these profiles, emerging systems are able to move beyond descriptive analytics toward predictive and prescriptive functions; anticipating risky states, tailoring interventions, and delivering feedback that adapts to the unique characteristics of each driver. This personalised approach has significant implications across domains such as road safety, fleet management, and insurance, offering the potential to reduce accidents, improve driving habits, and foster long-term behavioural change. The following sections examine current advances in driver profiling, coaching, and assistance, highlighting the methods, applications, and challenges in operationalising personalised driver support.

### **2.3.1 Personalized driver profiling and coaching**

Recent research emphasizes the increasing integration of telematics, AI, and predictive analytics for driver behaviour. By leveraging data from onboard sensors, cameras, and Internet of Things devices, these systems generate context-aware driver profiles and provide timely feedback, with applications in insurance and fleet safety management (Shirole et al., 2025). Advances in ML, DL, and generative AI are enabling more personalized profiling, supported by adaptive techniques such as Bayesian filtering for estimating mental workload in real time (Caber et al., 2023) and Inverse Reinforcement Learning to model individual driving preferences (Huang et al., 2022). Telematics interventions that combine behavioural feedback with financial incentives have also been shown to reduce risky driving, although their effectiveness depends strongly on the design of the incentives (Chen et al., 2023; Malekpour et al., 2023). In fleet operations, real-time monitoring, gamified dashboards, and scorecard systems have demonstrated substantial reductions in harsh braking, speeding, and crash-related costs (Fitzharris et al., 2017; Pozueco et al., 2020).

Despite these advances, several challenges still remain in this research area. Naturalistic Driving Study (NDS) datasets, such as Brain4Cars, Drive&Act, and SHRP2, provide valuable real-world data but are often constrained by limited participant pools and driving contexts, reducing their generalizability for personalized profiling (Liao et al., 2025). This lack of

diversity and volume may force researchers to collect custom datasets tailored to their objectives, such as the i-DREAMS project. On the modeling side, DL methods achieve high predictive performance but face criticism for their lack of transparency, which undermines trust in safety-critical applications. Although approaches like attention-based networks and hybrid rule-based frameworks have been proposed to improve interpretability, their use remains limited (Mobini Seraji et al., 2025). Furthermore, these models often require significant computational resources, posing difficulties for real-time deployment in resource-constrained, in-vehicle environments.

Another key barrier is the scarcity of labeled data, particularly for rare but high-risk driving events. Manual annotation is labor-intensive, which limits the scalability of supervised approaches. Emerging semi-supervised and self-supervised methods, as well as anomaly detection techniques, offer promising alternatives by exploiting large volumes of unlabeled data, though their application in driver profiling remains at an early stage (Narayanan et al., 2018). Finally, data privacy remains a persistent concern, as NDS data frequently contains personally identifiable information, such as in-cabin video footage. While de-identification techniques, including Generative-Adversarial-Network-based face anonymization, are advancing, striking a balance between data utility and privacy protection continues to be an unresolved challenge (Thapa & Sarkar, 2025).

### **2.3.2 Smart driver development: Personalized profiling and driver assistance to prevent driver distraction and drowsiness**

The technological evolution of artificial intelligence has pushed the boundaries of one-size-fits-all approaches, embracing sophisticated AI-driven systems that adapt to individual driver characteristics, preferences, and behavioural patterns. This approach extends to the domain of mitigating risks associated with DDD using personalized driver analysis and support. The foundation of this approach lies in creating robust driver profiles that serve as a baseline for normal driving behaviour (Tselentis and Papadimitriou, 2023). This is accomplished by applying unsupervised ML techniques, such as K-means or DBSCAN clustering, to extensive naturalistic driving data. These algorithms group driving segments based on a multitude of features—encompassing vehicular dynamics, driver behaviour, and physiological indicators—to establish a representative profile of an individual's typical operational patterns (Mantouka et al., 2019). Subsequently, to quantify risk, these profiles are scored using established SSMs. This profiling and scoring framework can be further enhanced by sophisticated multimodal AI models, such as ANNs and LSTM networks, which classify complex driver states by fusing data from heterogeneous sources like video feeds, vehicle kinematics, and sensor inputs (Wang et al., 2022).

Building upon this capacity for detailed driver state classification, the research focuses on transitioning from reactive detection to proactive intervention through personalised feedback and assistance. The core objective is to identify systematic patterns and subtle indicators that act as reliable precursors to a fully manifested state of driver behaviour mainly linked to distraction or drowsiness. Advanced time-series analysis and graph-based models (Li et al., 2024) could be employed to examine the sequence of driving events and learn the

complex temporal relationships between different indicators (e.g., lane deviation, blink frequency, gaze concentration). By identifying these individualized precursor patterns, the system can provide timely, data-driven suggestions or alerts, forming the basis of a personalized coaching mechanism. The approach has already been employed in comparable parallel domains such as aviation safety (Ackley et al., 2020). The end-goal is the integration of these validated AI models into real-time in-vehicle driver assistance systems or telematics-based applications. By combining sophisticated AI algorithms with comprehensive data collection and real-time feedback mechanisms, these methodologies create an ecosystem of continuous improvement that enhances both individual driver performance and overall road safety outcomes.

### **2.3.3 Personalised profiling using telematics, traffic and infrastructure data**

Building upon the multimodal data integration framework, a natural next step is to operationalise the derived features for individualised assessment and support. By linking telematics-derived kinematic patterns with infrastructure context and prevailing traffic conditions, it becomes possible to move beyond aggregate safety analytics and deliver tailored feedback at the driver level. This shift reflects the growing recognition that risk is not solely embedded in the road environment but also in the behavioural heterogeneity across drivers, which calls for adaptive and personalised interventions (Eren et al., 2012; Sagberg et al., 2015).

Driver profiling and scoring systems leverage ML methods, such as clustering and classification, to identify recurrent patterns in driving style. Profiles are typically constructed from behavioural key performance indicators, including rates of harsh manoeuvres, speed compliance, and stability indices, normalised for exposure (Eren et al., 2012). These profiles can be used to stratify drivers into categories ranging from conservative to aggressive, or to derive continuous risk scores that reflect the likelihood of unsafe events. When infrastructure and traffic context are incorporated into the model, these scores gain additional interpretability: the same harsh-braking frequency may signal riskier tendencies on high-speed arterials than on congested urban streets. Such multimodally enriched profiling has found applications in usage-based insurance, fleet management, and road safety research, offering a data-driven foundation for incentivising safer behaviour (Toledo et al., 2008).

Personalised coaching extends these insights by providing actionable feedback to individual drivers (Meuleners et al., 2023). Feedback mechanisms range from simple post-trip summaries delivered via telematics-based applications to more advanced modalities such as annotated video replay, where critical manoeuvres are highlighted and contextualised. Data-driven suggestions can emphasise not only what behaviours occurred but also why they were problematic under the prevailing road and traffic conditions. For example, a driver might be informed that their late braking on a rural curve is both a sign of aggressive cornering and a mismatch with the geometric demands of the segment. Studies have shown that continuous, context-aware coaching can reduce risky behaviours more effectively than generic messaging, as it aligns recommendations with drivers' personal behavioural profiles and their specific environments (Toledo et al., 2008).

Finally, multimodal models can be embedded into real-time, telematics-based applications or ADAS. By monitoring driver inputs alongside contextual indicators in real time, such systems can issue timely alerts, ranging from collision warnings to speed compliance reminders, or adapt vehicle assistance thresholds to the driver's profile (Castignani et al., 2015). For instance, drivers classified as prone to abrupt lane changes may receive earlier stability control interventions, while conservative drivers may benefit from smoother adaptive cruise control responses. In connected-vehicle settings, the integration of traffic and infrastructure feeds further enables cooperative safety functions, such as advisory speed recommendations on approaches to intersections or proactive warnings in high-risk corridors (Sagberg et al., 2015). The fusion of behavioural profiling with real-time assistance in ADAS thus closes the loop between observation, interpretation, and intervention, forming the basis of a holistic driver support ecosystem.

## 2.4 Methodologies for ethics, fairness, and human-centric design in road Safety

The integration of AI into road safety introduces not only poses technical challenges but also presents profound ethical, societal, and human-centred considerations. As driver monitoring and feedback systems increasingly rely on behavioural, physiological, and biometric data, issues of fairness, transparency, and trust become central to their design and deployment. This section examines the ethical frameworks and human-centric principles necessary to ensure that AI-driven safety tools respect autonomy and privacy while equitably serving diverse populations. It explores how explainable and trustworthy AI methods can foster user acceptance, how fairness-aware strategies can mitigate demographic biases in predictive models, and how federated learning offers a privacy-preserving pathway for leveraging sensitive biometric and physiological data. Together, these discussions highlight the importance of embedding ethics, fairness, and human-centric design into AI systems to ensure that road safety innovations are not only technologically effective but also socially responsible and inclusive.

### 2.4.1 Ethical Road Safety: Human-Centric Design Principles for Driver Monitoring

The deployment of AI-powered mitigation strategies for driver assistance, especially in the domain of distraction and drowsiness, necessitates a comprehensive ethical framework that prioritizes human dignity, autonomy, and fairness. Existing research indicates that with the ongoing implementation of DMS in the auto-industry, there are significant gender differences in road safety attitudes, behaviors, and risk perception which have limited considerations within the algorithms that run these systems (Cordellieri et al., 2016). Driver profiles and age also play a significant role in driver distraction (Bock et al., 2021), risk perception (Zhang et al., 2019), and ADAS take-over performance (Wu et al., 2020). This raises questions about the fairness and transparency of AI models trained on datasets that may not adequately represent the diversity of the driving population.

Fairness and transparency constitute foundational ethical principles intrinsically linked to trust in AI systems. Fairness ensures equitable outcomes across demographic groups, preventing discriminatory treatment based on protected attributes (Foulds et al., 2020), while

transparency encompasses openness about data sources, algorithmic processes, and system limitations, enabling stakeholders to understand decision-making mechanisms (Felzmann et al., 2019). AI amplifies ethical concerns through automation at scale, where biased algorithms can rapidly perpetuate and magnify existing inequalities across populations, while opaque decision-making processes undermine accountability and hinder detection of systematic errors (Ferrara, 2024).

Current research demonstrates that human-centric approaches to AI development are essential for building trustworthy and socially acceptable systems. The European Union's emphasis on human-centric AI, as outlined in article 6 of the EU regulation on artificial intelligence (Regulation 2024/1689), provides a foundational framework that categorizes AI systems based on their inherent societal risks, distinguishing between unacceptable, high-risk, and minimal-risk applications. A human-centric approach is imperative, as the goal is not merely to deploy technology but to enhance the safety and well-being of all road users.

The research framework acknowledges this by planning to leverage behavioral, physiological, and vehicular measurements, using a rigorous ethical framework to govern data collection, privacy, and application. The acceptability of such systems hinges on driver trust, which can only be established if the AI's interventions are perceived as reliable, timely, and non-intrusive (Dunn et al., 2021). Consequently, a human-centered evaluation focusing on trust and transparency becomes a critical component of the research methodology. The proposed use of Explainable AI (XAI) techniques directly addresses the need for transparency (Kuznietsov et al., 2024). By providing insight into how a model arrives at a specific conclusion—for instance, why it flagged a particular sequence of behaviors as a precursor to a drowsiness event—researchers and developers can build more robust and trustworthy systems, fostering greater user acceptance. Beyond the ethical handling of data and system transparency, ensuring fairness and mitigating demographic bias in AI models is a cornerstone of responsible innovation in road safety. The research explicitly recognizes the documented differences in road safety attitudes and behaviours across demographic groups (e.g., gender, ethnicity, age) and addresses the critical gap concerning potential biases in predictive models. This involves a rigorous diagnostic phase to quantify the extent of bias within the developed DDD models (Adler et al., 2018). This would involve a formal evaluation using established fairness metrics, which examines if outcomes are proportional across groups (Zafar et al., 2017). Once biases are identified and measured, the research transitions to implementing targeted mitigation strategies. The proposed methods include both in-processing techniques like Fairness-Aware Learning, which integrates fairness constraints directly into the model training process by penalizing biased outcomes, and post-processing adjustments such as Threshold Moving, which fine-tunes classification thresholds for different demographic subgroups to equalize fairness metrics without necessitating a full model retrain (Putzel & Lee, 2022). By systematically identifying and correcting these biases, the research aims to ensure that the AI-driven safety benefits are distributed equitably across the entire driving population.

### 2.4.2 Federated learning

Federated Learning (FL) is an emerging ML paradigm designed to train models across distributed devices without requiring raw data to leave the local nodes. Unlike traditional centralized approaches, where sensitive information is aggregated on a central server, FL allows individual clients—such as vehicles or onboard sensors—to locally compute model updates, which are then aggregated to form a global model. This approach is particularly relevant in connected vehicle applications, where large volumes of driver data, including images, physiological signals, and behavioural metrics, are continuously generated. Directly transmitting these data for centralized training raises privacy concerns, regulatory challenges, and communication overhead. By enabling collaborative learning while preserving user privacy, FL provides a practical solution for applications like driver drowsiness detection, allowing models to leverage diverse, real-world data from multiple vehicles while minimizing risks associated with data sharing. Additionally, FL supports scalability and personalization, allowing models to adapt to individual driving behaviours without exposing sensitive personal information.

Most current studies have focused on image-based datasets, such as those used by Zafar et al. (2021), Zhao et al. (2021), Zhang et al. (2022), and Lindskog et al. (2024), where facial cues like blinking and yawning serve as primary indicators of fatigue. More recent work has turned to EEG and EOG biosignals, with Qin et al. (2023), Chen et al. (2024), and Bernal et al. (2024) exploring federated frameworks that address inter-subject variability, small dataset sizes, and multi-stakeholder collaboration. Yu et al. (2024) extended this line of research by proposing a multimodal FL framework that combines visual and physiological modalities, while Khoa et al. (2024) incorporated various DL models into FL using longitudinal physiological sensor data. Dautov et al. (2024) further emphasized the non-IID challenge of personal datasets, underscoring the need for solutions that balance personalization and generalization in real-world deployments.

Despite these advances, current FL-based drowsiness detection studies remain limited to images, EEG, EOG, or multimodal combinations thereof, with no work yet leveraging ECG signals. Given ECG's practicality as a non-invasive, widely adopted biosignal already integrated into wearable devices and vehicular monitoring systems, this represents a significant research gap. Exploring ECG within an FL framework could enable scalable, privacy-preserving, and physiologically grounded drowsiness detection, paving the way for more robust and deployable connected vehicle safety solutions.

## 3 Overview of data sources for IVORY DCs

This section provides an overview of the data sources used by the IVORY doctoral candidates (DCs), highlighting the diversity of datasets and their complementary roles in supporting methodological innovation. The data underpin both the development and calibration of models and the empirical evaluation of newly proposed approaches, ensuring methodological robustness and real-world relevance. The overview captures a broad spectrum of modalities, from video and telematics logs to physiological measurements, traffic and infrastructure data, and naturalistic driving records, reflecting the multimodal perspective emphasized in Section 2. These datasets enable the extraction of surrogate safety indicators, behavioural markers, and contextual variables.

### 3.1 Data source overview

The project relies on a diverse set of complementary data sources, each serving distinct roles in model calibration, testing, and validation, and together enabling a comprehensive exploration of driver behaviour and road safety. Naturalistic driving datasets such as **PREVENTION**, **i-DREAMS**, **UDrive**, **OSeven**, provide rich real-world information on traffic interactions, lane changes, cut-in events, and general driving patterns, forming the basis for behaviour modelling, risk assessment, and safety metric extraction. To deepen the analysis of driver states such as distraction and drowsiness, specialised physiological datasets including **CARDIOID/ISEL**, **SleepEYE**, and **Valu3s** offer detailed biometric signals that allow the detection and prediction of cognitive and physical states. In parallel, datasets focusing on professional drivers, such as the **TNO dataset**, contribute insights into optimal driving practices and adaptation strategies, serving as benchmarks for resilience and safe performance. Collectively, these heterogeneous sources provide the empirical foundation for advancing methodological approaches, supporting the development of predictive models and evaluating personalised, context-aware interventions.

For all datasets holds that they have been prepared with attention to ethical and legal standards, including GDPR compliance (if applicable). Prior to publication or distribution, all personal identifiable information was removed or anonymised to prevent re-identification of participants. By balancing open accessibility with rigorous anonymisation, the datasets enable broad research use while upholding data protection requirements and safeguarding participant privacy.

### 3.2 The PREVENTION dataset

#### 3.2.1 Description/origin

The PREVENTION dataset is the most widely used benchmark in the lane change detection research area. It comprises 356 minutes (corresponding to 540 km of driving) of data recorded in 2018 by three different drivers in both urban and highway settings, using a suite of sensors (LiDAR, radar, and cameras). The authors also provide extensive post-processed annotations for surrounding vehicles, including vehicle bounding box, unique vehicle IDs, trajectories, lane-marking delineations, and detailed lane-change behaviour labels.

Each behaviour label contains the target vehicle ID, the lane change direction, the start-frame, mid-frame, end-frame, and whether the turn indicator was activated. These behaviour labels are commonly utilized as prior information to prepare event-specific video clips for experiments in this research direction.

### **3.2.2 Availability/replicability**

The [PREVENTION dataset](#) is openly available to the public. Its open-access nature allows researchers and practitioners to freely download, explore, and analyse the data without the need for special agreements or approvals. The dataset is accompanied by thorough documentation on its structure, collection procedures, and variable definitions, ensuring that analyses can be replicated and extended across different research contexts. This transparency and accessibility support both methodological robustness and knowledge transfer within the wider research community.

## **3.3 The i-DREAMS dataset**

### **3.3.1 Description/origin**

The naturalistic driving data used in the research is collected from the EU Horizon 2020-funded i-DREAMS (Intelligent Driver and Road Environment Assessment and Monitoring System) project. It was developed to design, test, and validate a context-aware safety framework built around the concept of a Safety Tolerance Zone (STZ). The STZ continuously evaluates whether drivers operate within safe limits by considering background factors, real-time risk indicators, and task complexity. To improve safety, the platform delivers both in-vehicle interventions, such as real-time alerts to prevent unsafe situations, and post-trip feedback to promote long-term safety awareness.

The project collected extensive naturalistic driving data from 250 drivers across Belgium, Greece, the UK, Germany, and Portugal, spanning 49,651 trips and over 1.95 million minutes of driving. Vehicles were instrumented with On Board Diagnostics (OBD) devices, Mobileye and dash cameras, Cardio Gateway sensors, and smartphone-based telematics. These sensors captured multimodal information, including for example speed, acceleration, braking, lane departures, headway, phone use, event-related video data and physiological measures such as heart rate variability. The dataset also contains socio-demographic and attitudinal data from surveys.

### **3.3.2 Availability/replicability**

The [i-DREAMS dataset](#) is available to project partners and selected research initiatives under specific data-sharing agreements. Due to its richness and the inclusion of potentially sensitive information, access is managed with attention to legal and ethical requirements. For this project, a subset of the dataset has been made available for initial analysis, and further access is subject to approval in line with the data governance protocols of the i-DREAMS consortium. While the full dataset is not openly accessible to the public, its methodology and structure are well documented, supporting the replicability of analyses in comparable research contexts, provided similar data sources can be obtained.

## 3.4 UDrive Dataset

### 3.4.1 [Description/origin](#)

The UDRIVE dataset is a large scale European Naturalistic Driving study, providing unique behavioural data on car driving, truck driving and riding powered two-wheelers. In six European countries, cars, trucks and powered two wheelers were fitted with camera's, sensors (GPS/MobilEye + 7 internal/external cameras) that provide detailed information on driver behaviours, vehicle dynamics, and environmental factors across different levels of automation. The data consists of 186 drivers across 5 EU countries and approximately 45K hours of naturalistic driving.

### 3.4.2 [Availability/replicability](#)

This dataset which needs to be accessed for research is subjected to restrictions and specific availability based on appropriate authorizations. It is not available for public use due to the sensitive nature of the driving video content. It is only available to partners subject to approval by Stellantis LAB and its associated partner CEESAR. The data cannot be replicated outside the LAB and can only be accessed by a specific tool provided by CEESAR.

## 3.5 OSeven Dataset

### 3.5.1 [Description/origin](#)

The telematics dataset originates from OSeven Telematics and consists of anonymized time-series data collected from smartphone sensors at a frequency of 1 Hz. Each record includes GPS-derived location coordinates, speed, and heading (course over ground), as well as readings from embedded inertial sensors such as the accelerometer and gyroscope. The data is captured passively during driving, leveraging smartphones placed arbitrarily inside the vehicle without fixed orientation or mounting. As a result, the dataset reflects realistic deployment conditions, including variability in device orientation, sensor noise, and driving environments. The temporal resolution and the nature of the dataset make it suitable for modeling driving behaviour and detecting dynamic events such as harsh braking, harsh acceleration, and harsh cornering.

### 3.5.2 [Availability/replicability](#)

The dataset is exclusively available to DC7. External access is not permitted. This controlled access ensures the security of sensitive data and full compliance with applicable data protection regulations.

## 3.6 CardioID/ISEL dataset

### 3.6.1 [Description/origin](#)

The CardioID/ISEL simulator dataset will be collected by DC14 using a driving simulator equipped with a range of physiological sensors. Data acquisition will include signals from in-vehicle technologies such as CardioWheel, as well as additional physiological measurements gathered from wearable devices, including electroencephalogram (EEG), photoplethysmogram (PPG), and electrocardiogram (ECG). To enhance the multimodal nature of the dataset, a driver monitoring camera will also be used to capture the driver's

facial expressions and blinking rate, which may serve as supplementary inputs for future learning models. While the data collection protocol is still being finalized, it will include a mechanism for collecting subjective ground truth data, whereby participants will verbally report their sleepiness levels using the Karolinska Sleepiness Scale (KSS) at defined intervals throughout the experiment.

### **3.6.2 Availability/replicability**

The dataset to be acquired for this research is subject to specific availability and usage restrictions. Access to the data will be limited to collaborations involving CardioID, and it will not be publicly accessible for open use. Due to confidentiality and data protection considerations, the dataset will not be stored or processed on institutional partner storage; all data processing must be conducted on-site at CardioID premises. Furthermore, the dataset must be used exclusively within CardioID, and any form of external sharing or use beyond the defined scope of the collaboration will not be permitted. These restrictions impact the replicability of experiments by external parties but ensure compliance with the proprietary and ethical obligations associated with the data.

## **3.7 SleepEYE dataset**

### **3.7.1 Description/origin**

The SleepEye dataset is provided by the Swedish National Road and Transport Research Institute (VTI). Data collection was carried out in Sweden and involved 20 participants (10 male, 10 female) aged between 30 and 60 years, all of whom were recruited from the Swedish register of vehicle owners. Participants were required to be in good health, of normal weight, non-shift workers, and not professional drivers. Individuals who wore glasses were excluded. Prior to the study, participants were instructed to maintain regular sleep schedules and to obtain at least seven hours of sleep per night during the three nights leading up to the experiment. They also completed pre-study sleep and wakefulness logs.

The study included four driving sessions per participant, consisting of both real-road and simulator drives under two different conditions: an alert condition conducted in the late afternoon and a sleep-deprived condition conducted after midnight. Each driving session lasted approximately 90 minutes. The driving simulator and real-road car were each equipped with Smart Eye camera systems: a 1-camera embedded system under evaluation, and a 3-camera Smart Eye Pro system (sampling at 60 Hz) used as the reference for ground truth. In the simulator, both alert and sleepy sessions were conducted in daylight due to environmental limitations, while in the real-road experiment, sessions were conducted in actual daylight or nighttime accordingly.

In both environments, vehicle data—including speed and lateral position—was recorded at 10 Hz, and physiological data was captured using a Vitaport 3 system (Temec Instruments BV, the Netherlands). The recorded physiological signals included electroencephalogram (EEG), electrooculogram (EOG), and electrocardiogram (ECG), sampled at 256 Hz for EEG and ECG, and 512 Hz for EOG. To assess subjective sleepiness, participants verbally reported their drowsiness levels during the sessions using the Karolinska Sleepiness Scale (KSS). Although the study primarily focused on sleepiness, a distraction event was introduced at the end of

each session, enabling the dataset to also support future research on driver distraction and simulator validation.

### **3.7.2 Availability/replicability**

The dataset used in this research is subject to specific availability and usage restrictions. Access to the data is limited to collaborations involving CardioID Ltd, and it is not publicly accessible for open use. Due to confidentiality and data protection considerations, the dataset cannot be stored or processed on institutional partner storage; all data processing must be conducted on-site at CardioID premises. Furthermore, the dataset must be used exclusively within CardioID, and any form of external sharing or use beyond the defined scope of the collaboration is not permitted. These restrictions impact the replicability of experiments by external parties but ensure compliance with the proprietary and ethical obligations associated with the data.

## **3.8 Valu3s dataset**

### **3.8.1 Description/origin**

The Valu3s dataset is provided by the Swedish National Road and Transport Research Institute (VTI) (Meyerson et al. 2024). Data collection was conducted in Sweden and involved 20 participants (10 male, 10 female) aged between 20 and 60 years. Participants were required to hold a valid driver's license for passenger cars and to be regular drivers. Individuals with a body mass index (BMI) exceeding 35 were excluded to mitigate potential confounding factors, such as undiagnosed sleep disorders and motion sickness.

Data collection was conducted using two driving simulators. Both simulators were equipped with physiological sensors to facilitate the collection of drowsiness-related data. The dataset includes recordings from CardioWheel, a steering wheel-embedded ECG sensor, as well as additional physiological signals relevant to drowsiness detection, such as electrooculogram (EOG) and electrocardiogram (ECG). The EOG (vertical channel, right eye) and the ECG (lead-II) were recorded with a Vitaport 3-system (Temec Instruments BV, the Netherlands). The electrodes used for the EOG and ECG were of the disposable Ag/AgCl type. CardioWheel data were sampled at a frequency of 1000 Hz, while EOG and ECG signals were sampled at 256 Hz.

Each participant completed four driving sessions: two during daytime and two at night after having been awake since early morning. Each session lasted approximately 60 minutes. For this study, only the ECG data acquired from the chest strap were used, as it provided a continuous signal without interruptions. In contrast, CardioWheel ECG data were subject to signal loss whenever participants removed their hands from the steering wheel, making it unsuitable for this analysis.

Throughout the driving sessions, participants were asked to verbally assess their sleepiness using the Karolinska Sleepiness Scale (KSS). The KSS consists of nine discrete levels, ranging from 1 (extremely alert) to 9 (very sleepy, making a great effort to stay awake). Sleepiness ratings were recorded as an average value every fifth minute during the driving sessions. These KSS ratings serve as ground truth labels for training classification models in subsequent analyses.

### **3.8.2 Availability/replicability**

The dataset used in this research is subject to specific availability and usage restrictions as outlined in the project's grant agreement. Access to the data is limited to collaborations involving CardioID, and it is not publicly accessible for open use. Due to confidentiality and data protection considerations, the dataset cannot be stored or processed on institutional partner storage; all data processing must be conducted on-site at CardioID premises. Furthermore, the dataset must be used exclusively within CardioID, and any form of external sharing or use beyond the defined scope of the collaboration is not permitted. These restrictions impact the replicability of experiments by external parties but ensure compliance with the proprietary and ethical obligations associated with the data.

## **3.9 TNO Dataset**

### **3.9.1 Description/origin**

The second data source for this research is being collected by TNO (DC6 Secondment Partner). This dataset is designed to support investigations into optimal driving behaviour and adaptation under varying conditions. The data collection involves a group of highly trained drivers who possess advanced driving skills, referred to as "Directiechauffeurs" in Dutch. These participants drive TNO's advanced instrumented vehicle platforms, the TNO CarLabs. Under real-world traffic conditions, the drives are performed along a fixed route between Eindhoven and Gilze.

The vehicle platforms are equipped with a wide array of sensors capable of capturing detailed information on driving dynamics, driver input, and vehicle–environment interaction. This dataset focuses on recording high-quality, high-resolution data under controlled yet realistic scenarios, making it particularly valuable for validating conceptual models related to optimal and resilient driving patterns.

### **3.9.2 Availability/replicability**

The dataset is being developed as part of an internal research initiative at TNO and is not publicly available. However, full access will be granted during the DC6's six-month secondment at TNO, starting in October 2025. Due to the high quality and specificity of the data collection process, the dataset provides a unique opportunity to explore controlled driving behaviour in naturalistic conditions. While replicability outside TNO may be limited due to proprietary tools and access to specialised drivers, the methods used in data collection are transparent and documented, supporting comparative studies if similar setups are available.

## **3.10 Data Access and Usage by DCs**

Table 1 below shows which datasets are currently in use, or planned to be used, in the respective PhD works of the different DCs. For some datasets the access has already been obtained, while for others the access request is still in progress or under evaluation.

**TABLE 1: USE OF DATASETS BY DCs**

Database Name	Current access	Current usage	Future access /usage interest
Prevention	DC4, DC5, DC6, DC7	DC4	-
i-DREAMS	DC4, DC5, DC6	DC4, DC5, DC6	DC7
Valu3s	DC14	DC14	-
SleepyEye	DC14	DC14	DC5
CardioID	DC14	DC14	DC5
TNO	DC6	DC6	DC6
OSeven	DC7	DC7	-
UDrive	DC5	-	-

## 4 Innovative methodological approaches by IVORY DCs

This section presents the innovative methodological approaches the IVORY doctoral candidates (DCs) developed, outlining how their individual research projects contribute to advancing the current state-of-the-art in driver behaviour analysis and road safety. Each contribution begins with a description of the specific research objectives, followed by an overview of the methodological frameworks and techniques applied, and is complemented by a research flowchart to illustrate the structure and progression of the work.

When reading the contributions from the different DCs below, it should be taken into account that each of them follows their own path towards achieving the research and innovation goals in their PhD. In some cases, the emphasis during the first year of PhD work is heavily on building theoretical foundations, conceptual frameworks and literature synthesis, while in other DCs the focus is more on concrete methodological innovations, data processing and algorithmic implementation. It explains why the sections below are not fully comparable in how the DCs present their progress until now.

### 4.1 Methodologies adopted by DC 4

#### 4.1.1 Objectives of the research

The research topic of DC4 is road user profiling using multimodal data of naturalistic driving databases. The project is academically supervised by UHasselt, complemented by industrial supervision from CardioID, and further enriched through a secondment at ISEL. The research objectives and expected outcomes are defined as follows.

#### RESEARCH OBJECTIVES:

- To create automated analysis techniques for extraction of important safety features and risk factors from road environment video data
- To automatically create road user profiles from naturalistic driving data
- To translate results of driving profiling into effective recommendation of videos and post-trip personalised coaching for improving safe driving behaviour

#### EXPECTED RESEARCH OUTCOMES:

- The creation of AI software for the extraction of safety features and risk factors from dashcam videos
- The creation of AI software for improved road user scoring and profiling
- The translation of driving profiles into intelligent recommendation of relevant videos and post-trip personalized coaching for improving road user behavior

Therefore, this research will prioritize the creation of AI-driven techniques to extract key safety features and risk factors from road environment video data, allowing for the development of dynamic, context-aware risk profiles. These profiles will serve as the basis for personalized post-trip coaching strategies, designed to improve driving safety through targeted feedback and behaviour adjustments.

By advancing AI software for driver scoring, profiling, and tailored safety interventions, this research aligns with the Vision Zero initiative, offering a deeper understanding of driver risk behaviours and delivering effective solutions to mitigate risky driving. Ultimately, this work aims to enhance road safety through innovative AI-powered tools that foster continuous driver improvement and risk reduction.

#### 4.1.2 Overview of methods

The research starts from the first objective above. Although the [i-DREAMS project](#) has collected and processed a large amount of data from various sensors, the potential of dashcam video data was not fully utilized for road safety risks analysis, e.g., tailgating events. Tailgating typically occurs when a driver maintains a headway time of less than two seconds behind the vehicle in front, which is a frequent cause of rear-end collisions (McCarty & Kim, 2024). In the i-DREAMS project, the time headway was measured to detect dangerous tailgating behaviour for ego vehicles. If such behaviour is detected, the driver's trip performance score will be deducted during post-trip safety interventions. However, research indicates that a surrounding vehicle's cut-in maneuver can also cause a sudden decrease in headway, potentially leading to false detections of tailgating. Therefore, to ensure fair driver performance assessments in post-trip safety interventions, it is crucial to differentiate whether the headway reduction results from the ego vehicle driver's tailgating or the surrounding vehicle's cut-in behaviour.

With the rise of DL, models such as CNNs, LSTM, Two-Stream-based networks, 3D video action recognition models, and Video Vision Transformer models have been gradually applied to leverage the scene's semantic information and visual features extracted from dashcam videos to detect dangerous lane change behaviours (Fernández-Llorca et al., 2020; Liang et al., 2023). However, when these state-of-the-art methods are evaluated from a practical application perspective, beyond just accuracy, several limitations become apparent, e.g., the need for prior information, the requirement to prepare event-specific video clips, limited algorithm testing scenarios, the inability to identify the vehicle performing the lane change event, insufficient analysis on the lane marking detection model's robustness and so on. In fact, these limitations mainly stem from the high computational demands of DL models. In other words, the cost of the inference process does not always match the value of its result, considering that most surrounding vehicles in most of the frames are irrelevant to lane change behavior. Those preprocessing steps, such as selecting relevant video segments for specifically identified target vehicles, are designed to enable the model to focus on meaningful targets and useful video periods, thereby producing valuable outputs with a relatively low computational cost. However, these steps introduce substantial manual labelling work and create significant barriers to practical application in real-world use.

Therefore, to develop practical algorithms for cut-in detection, several open challenges are identified and need to be addressed:

- Eliminate the requirement of target vehicle selection during data preprocessing.
- Eliminate the requirement of behaviours' start and end labels to generate a vehicle's corresponding event-specific video clips as input for the models. Actually, this challenge means that no matter the duration of the video, the algorithm can still process the video well.

- Robustness and accuracy, i.e., adaptation to different scenarios, e.g., different dash cameras, installation places, shooting angles, weather, light conditions, and so on.
- The identification of the lane-changing vehicle and also the determination of the start and end frames of each event, instead of just outputting a classification label for an input feature sequence.
- Highlighting the lane-changing vehicle in advance for warning drivers. This challenge means the algorithm can somehow detect the behaviour before the lane-changing vehicle completely enters the ego lane in real-time, instead of just summarizing the occurrence of the behaviour.

In addition, a novel algorithm which can meet these challenges is also proposed for cut-in detection and evaluated on a public dataset, i.e., the PREVENTION dataset (Izquierdo et al., 2019). Moreover, the videos of the i-DREAMS project are utilized to form a new dataset to make up for the PREVENTION dataset's shortcomings (e.g., inconvenient lane marking annotations, lack of direct connection with road safety questions, limited diversity in challenging conditions, and predominant focus on highway scenarios).

With respect to objective #2 above, driver profiling in road safety research is commonly defined as the process of modeling and categorizing drivers on the basis of their behavioral patterns and risk characteristics, in order to predict or mitigate unsafe events (Ellison et al., 2015). Although the term does not have a universally fixed meaning, most studies converge on the idea that profiling involves the systematic extraction of features from naturalistic driving data such as speed, acceleration, braking, and lane-keeping, and their translation into interpretable representations of driver style or safety risk (Singh & Kathuria, 2021). However, some authors emphasize long-term, habitual styles (e.g. conservative, moderate, aggressive), whereas others focus on dynamic risk states that evolve within and across trips (Al-Hussein et al., 2022).

Naturalistic driving data have become foundational because they deliver continuous real-road observations rather than behaviour in controlled or simulated settings. These datasets often include GPS, in-vehicle telematics (accelerations, decelerations, braking, throttle), sometimes camera or vision inputs (for attention, lane-keeping), and annotations of safety-critical or near-crash events (Abdelrahman et al., 2022). When paired with outcome labels (crashes or near-misses) or comparative measures, they allow the derivation of driver behaviour profiles that quantify risk or classify driving style (Chan et al., 2025).

ML and AI methods are widely used, including K-means clustering, neural networks, XGBoost, and ensemble learning (Oban et al., 2024). A growing trend is using sequential modelling to account for temporal dynamics. For instance, the hybrid CNN-BiLSTM + Hidden Markov Model structure was used for spatial/feature extraction and sequence modeling to identify distinct driving behavior states (straight driving, lane-changing, deceleration, turning) and to analyze temporal transitions between states (Sun et al., 2025).

In practical applications such as transportation fleet management, which involves large numbers of drivers and daily trips, these systems generate massive volumes of events, time-series data, and video recordings. Managers are then required to devote substantial effort to reviewing, analyzing, comparing, and summarizing each driver's profile and the risky events detected during each trip in order to determine the necessary actions for enhancing driving

safety. Relying on hierarchical drilldowns and manual per-trip reviews places heavy time and cognitive demands on fleet managers, often leading to information overload and alert fatigue.

Importantly, it is noted that most research has focused on developing sensors and improving detection accuracy, but far less on how drivers or managers interact with the resulting alerts and data (Ayas et al., 2024). False or frequent warnings can cause confusion, while user acceptance and system usability are rarely evaluated. For large-scale effectiveness, detection and scoring technologies must therefore be supported by interfaces that enable prioritization, summarization, and anomaly ranking.

Therefore, to address this issue, it is necessary to further extend the functionality of existing systems so that they not only detect and record unsafe driving behaviours, but also perform automated information filtering, prioritization, and comprehensive analysis at a higher level. The system should be able to automatically generate a list of high-risk drivers for each day or shift and provide a concise risk summary and recommended actions based on event type, frequency of occurrence, comparison with historical baselines, and potential consequences. In this way, managers no longer need to rely on large amounts of manual clicking and data retrieval, and can quickly identify the drivers and events that require attention. To this end, generative AI or AI agent techniques that utilize multiple AI models for cooperation may be utilized in future research to achieve these objectives.

The third objective is to translate driving profiling into effective video recommendation and post-trip personalized coaching. Graph-based driver profiles were proposed because dynamic graph representations can integrate contextual and behavioral features to enable deeper personalization of feedback (Zahraoui et al., 2022). Similarly, it was stressed that driver profiling methods are moving beyond static classification to adaptive and context-aware representations (Tselentis & Papadimitriou, 2023). In terms of interventions, Sangrar et al. (2022) demonstrate the potential of video-based coaching, reporting that their personalized ROADSkills program improved on-road performance and reduced risky behaviors in older drivers.

Despite these advances, several gaps remain. A systematic review of personalization in driving behavior concludes that there is no standardized methodology to evaluate the effectiveness of personalization in driving interventions (Liao et al., 2025). Moreover, while feedback reports can modestly improve driving safety in the short term, sustained long-term behavioral change is rarely achieved without ongoing intervention or incentives (Stevenson et al., 2021). Future research should focus on strategies to maintain improvements after feedback is discontinued. Another gap is causal attribution: without distinguishing whether risky driving events stem from driver error or external hazards, recommendations risk being overly generic and less effective. Research consistently shows that the majority of risky events and crashes are attributable to driver error, but external factors also play a significant role and require different intervention strategies (Khattak et al., 2021).

Therefore, future systems could combine causal event attribution, dynamic driver profiles, and automated micro-video generation into a scalable pipeline. Such systems would deliver short, personalized, video-based coaching modules with causal explanations (e.g., highlighting that a braking event was caused by speeding on a curve rather than external traffic). The novel AI techniques, such as multimodal Large Language Model (LLM) and AI agents, can be utilized for these purposes as explained in the following section.

### 4.1.3 Research flowchart

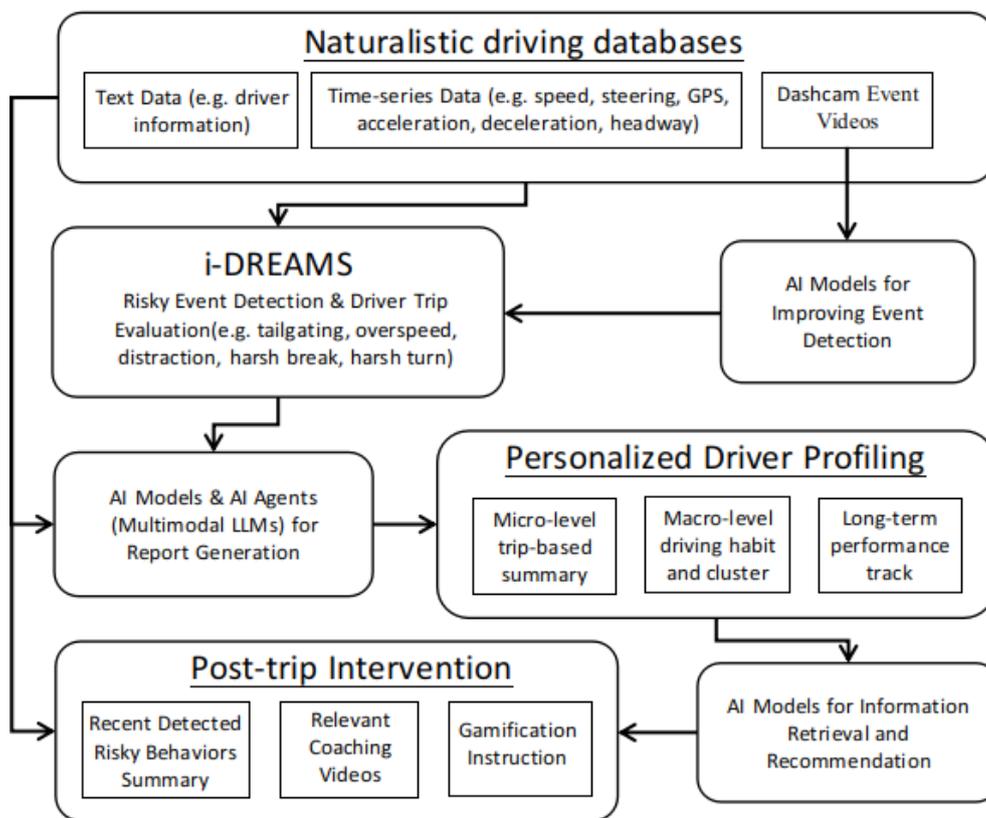


FIGURE 2: DC4'S RESEARCH FLOWCHART

As indicated in Figure 2 the research conducted by DC 4 highlights the application of naturalistic driving datasets in combination with AI techniques to detect safety- and risk-relevant patterns within multimodal data and to construct comprehensive driver profiles.

It begins with solving the tailgating detection error. As mentioned before, when a vehicle cuts into the driving lane, the detected time headway can quickly drop below the safety threshold, potentially triggering a false alert for low-headway risk. This could lead to the system unjustly issuing warnings and unfairly reducing the driver's score, prompting the driver to adjust their behaviour unnecessarily. Therefore, object detection, object tracking, and lane detection AI models are explored with dashcam video data to accurately identify these scenarios and enhance the reliability of the event detection system.

The second stage aims to assess drivers' road safety performance by analyzing their behaviour during previously recorded trips. In addition to evaluating specific driving patterns on a micro level within individual trips, the study also seeks to identify broader, personalized driving habits on a macro level and track their improvement over time after being coached by the developed behavioural intervention system. Ideally, we aim to present driver profiles in the form of readable text with explanations, rather than just simple performance scores across various aspects. To achieve this, the use of CNN models, multimodal DL, time-series AI models, and large language models is crucial for tasks such as feature extraction, risky event detection, summarizing driving behaviours during each trip, and tracking changes in driving behaviour over time.

The ultimate goal of this research is to convert both microscopic and macroscopic driver profiles into personalized recommendations for improving driving safety, as well as intelligent suggestions for relevant coaching videos. To achieve this, a road-safety-focused LLM with traffic knowledge can be developed to understand various risky driving events and offer specific improvement suggestions.

In summary, the expected real-life application pipeline resulting from this research would work as follows: after the necessary devices are installed in a driver's vehicle, the driver will receive an evaluation report on their driving safety performance through an app, based on their daily trips. The daily evaluation will include videos of detected risky events, text summaries of the microscopic driver profile explaining the risk factors during the trip, and personalized improvement suggestions. Additionally, models will be developed to intelligently recommend relevant coaching videos. An information retrieval and recommendation system will also be necessary. Monthly, the developed models will summarize the driver's progress in safety performance using the app. The overall pipeline will focus on specific driver behaviours, offering personalized feedback rather than generic, overly broad conclusions.

## 4.2 Methodologies adopted by DC 5

### 4.2.1 Objectives of the research

The research conducted by DC5 under guidance of the MCSA IVORY 'AI for Vision Zero in Road Safety' industrial doctoral network program, coordinated by TU Delft in partnership with the Stellantis/LAB (An Accidentology LAB, created as a collaborative effort of PSA and Renault), aims to gain further insights into interrelationships between distraction, drowsiness, and risk on the roads and contribute to the development of AI based systems that can intervene and prevent crashes caused by DDD and ultimately reduce road fatalities.

To achieve the above aim, DC5's research looks at establishing a framework for leveraging AI and mitigating DDD. The framework involves using behavioural, physiological, and vehicular measurements across different levels of automation, in an effective and ethical way. The initial purpose is to use AI models to understand how DDD events and their features change with automation levels, and second, to create AI models using data fusion that learn from these events to identify systematic patterns of DDD as possible precursors linked with risk of a crash or a traffic conflict. These possible precursors can enhance the future AI detection algorithms at higher automation levels. An additional purpose is also to ensure the AI models' demographic biases are managed in an inclusive way. Ultimately, this research seeks to develop and validate AI systems capable of preventing crashes caused by DDD, contributing to road safety goals. This leads to the following research objectives and expected outcomes.

#### RESEARCH OBJECTIVES:

- To understand factors contributing to DDD in different levels of automation.
- To develop AI-based models that identify, detect, and predict precursors or predictors of distraction and drowsiness among drivers.
- To understand the relationship between distraction/drowsiness and other risky driving behaviours.

## EXPECTED RESEARCH OUTCOMES:

- Analytical models that estimate the likelihood of DDD across different automation.
- Identification of a list of patterns and factors of DDD as precursors or predictors.
- Recommendations about DDD mitigation in the presence of these models and automation.

### 4.2.2 Overview of methods

To achieve the objectives mentioned in the previous section, the research is structured around four distinct but interconnected research axes. Each axis originates from a research gap identified in the scientific literature and is designed to systematically address one or more of the main research objectives. This structure allows for a focused investigation into key research objectives mentioned in the previous section while ensuring the work contributes to the overarching goals of the IVORY project.

#### Relationship of Automation Levels and DDD

- **Gap:** Limited research explores how driver behaviour, specifically distraction and drowsiness, evolves across different levels of automation and how it links to risk factors (Hungund et al., 2021). This could ultimately result in a lack of comprehensive understanding of how different levels of automation, from basic driver assistance to conditional automation, specifically impact various indicators and types of DDD
- **Research Question 1 (RQ1):** How do specific indicators and patterns of DDD manifest as automation levels evolve?
- **Research Objective:** To understand the factors contributing to DDD by quantifying the relationship between specific indicators and patterns of DDD at different automation levels

#### Distracted and Drowsy Driving factors as Precursors Linked to Risk

- **Gap:** Studies mainly explore the influence of driver attributes, or reaction types during safety critical events on the occurrence of DDD (Soares et al., 2021), and existing systems often rely on detecting fully manifested DDD (Albadawi et al., 2023), and rarely look at the recurring patterns or precursors.
- **Research Question 2 (RQ2):** Are there specific indicators or systematic patterns of driver states that can be identified as precursors of DDD, while quantifying the precursor's endogenous relationship with risk?
- **Research Objective:** To develop AI-Based models identifying and predicting precursors of DDD that help understand the relationship between these emergent patterns and risky driving behaviours.

#### Improve AI models with Optimal Fusion of Diverse Data Sources

- **Gap:** The abundance of diverse data sources necessitates effective data fusion we maximize the accuracy of analytical models
- **Research Question 3 (RQ3):** How can we maximize the accuracy of analytical models for predicting DDD using the optimal data fusion strategy for combining heterogeneous data sources?

- **Research Objective:** To enhance the accuracy and robustness of AI –based models that detect and predict DDD by investigating data fusion techniques and identifying the optimal approach.

#### Bias in Models Related to Demographics

- **Gap:** Existing research highlights differences in road safety attitudes and behaviours, raising concerns about potential biases in models related to gender, ethnicity, and age (Cordellieri et al., 2016).
- **Research Question 4 (RQ4):** How can we ensure fairness and mitigate bias in AI models predicting driver distraction and drowsiness, considering demographic factors?
- **Research Objective:** To ensure equitable development of AI-based models by evaluating fairness-centred techniques to mitigate demographic bias in the prediction of DDD.

The methods explained below are chosen to focus on the above four axes or research gaps.

#### RQ1: Potential Methods and Techniques for DDD Detection across Automation Levels:

The proposed methods to target this area will largely focus on hybrid approaches combined with data fusion and multimodal ML techniques. The reason to select a hybrid modality approach involving combination of psychological or visual signals is based on evidence from literature reviews of higher success in hybrid methods (Albadawi et al., 2022; Koay et al., 2022), effectively mitigating some of the shortcomings of single-modality methods, and opening a path to build on this research as AI models become more efficient in ingesting and predicting DDD using multimodal data.

Initial steps for analysing the potential techniques would be *Descriptive analysis* using data visualization techniques that show how often, how long, and what types of distraction and drowsiness events occur at different levels of automation (Van Nes et al., 2019). This will be followed by *statistical tests* such as ANOVA, T-Tests, Time-series tests such as Durban Watson etc. to examine specific comparisons, like the frequency of phone use between drivers using different levels of automation (Cooper et al., 2023), and reveal associations between the type of distraction and the level of automation (Dunn et al., 2021). *Regression modelling* (linear, logistic, or mixed effect), will be used to analyse changes in likelihood of a drowsiness event based on the level of automation (Gershon et al., 2023). Finally, Time-series models will be used to analyse frequency and patterns of driver behaviour variables such as mobile phone use, head movements with automation features such as lane departures.

This initial analysis gives insights into the most promising features contributing to detection of DDD across automation levels. *Multimodal ANNs* is the core analytical technique or modelling strategy that would be implemented on these significant features after data transformation and standardization. It is a highly relevant hybrid model that can be used at different levels of automation to detect and assess the evolution of DDD. The intent here is to train the multimodal models for each automation level and identify the DDD evolution across SAE levels (L0-L2) with statistical methods such as mixed effect models that analyse probability distribution metrics of DDD states across automation.

This can be combined with explainable AI methods such as SHapley Additive exPlanations (SHAP) analysis to measure how the feature importance coefficients evolve during automation engagements (Ebel et al., 2023; Tahir et al., 2024).

- **Multimodal ANN Model for Distraction:** A single multimodal model can be designed to ingest data from multiple sources to detect distraction states. For instance, this model could ingest images from cameras capturing facial expressions, and eye movement while simultaneously processing input data from steering wheel sensors and accelerator/brake pedals at different automation levels and can be trained to detect the amount of time drivers spend looking away from the road in seconds (Eyes off Road), the amount of time the drivers' hands are not on the wheel (Hands off Wheel) and can analyse eye glances and SWA at each automation level.
- **Multimodal ANN Model for Drowsiness:** Similar to the distraction model, a multimodal ANN is effective for drowsiness detection as it can integrate diverse indicators of drowsiness, capturing both physiological and contextual factors such as PERCLOS, blink frequency, and head pose variations, this model can prioritize PERCLOS and blink frequency as indicators of homeostatic sleep pressure and circadian influences. Additionally, parameters such as time of day, average speed, collision warnings and harsh acceleration events can be incorporated from the NDS databases to predict their influence on driver drowsiness at different automation levels

#### RQ2: Potential Methods for Detecting DDD Precursors and their link to driving risk

Having analysed broader DDD patterns and evolution across automation, the focus of these methods is the early warning signs, predictors, or precursors, that signal an impending distraction or drowsiness event. Once precursors of DDD are identified an attempt to establish and quantify its relationship with risk of traffic conflicts or crash can be made. Several key techniques are promising and can be employed for precursor detection. Given the nature of naturalistic driving data, the most relevant are Time Series Analysis, GNNs and RNNs (mainly LSTMs) that allow the examination of driving data as a sequence of events.

The most promising method that would be the focus of this research is *Temporal Graph Neural Networks* (TGATs) that are used to modelling the complex interactions of driver behaviours over time, inspired by recent advancements in graph learning for risk prediction (Li et al., 2024). TGATs can represent DDD as dynamic graphs, where nodes represent different DDD indicators (e.g., eye closure duration, steering entropy, lane deviation) and edges capture their temporal relationships. By incorporating attention mechanisms, TGATs can learn to weigh the importance of different indicators and their interactions at various time points, effectively identifying precursor patterns. For example, a TGAT model could detect a precursor pattern where a subtle increase in lane deviation, followed by a decrease in blink frequency and a shift in gaze concentration away from the road, collectively indicates an impending drowsiness event.

Finally, to quantify the risk for drivers facing distraction and drowsiness the focus is on using the concept of *Driver Profiling* (Tselentis and Papadimitriou, 2023). The initial step would be to create a representative profile of driver's "normal" driving behaviour and then look for profiles which are riskier using the approaches mentioned below:

- **Clustering:** This promising method groups similar driving segments together based on extracted features with clustering algorithms (K-means, DBSCAN, Gaussian Mixture etc), applied to naturalistic driving data.
- **SSM selection and Anomaly Detection:** Relevant indicators (SSMs), e.g., TTC, PET, PSD etc for each driving segment in the datasets can be calculated. Thresholds for risky driving could be established based on the distribution of these SSM values.

### RQ3: Potential Methods for Data fusion and AI model performance

The fusion of multimodal data to enhance the prediction accuracy is an important strategy in this research. Different sensors offer complementary information; for instance, physiological sensors may excel at detecting the gradual onset of drowsiness, while behavioural and vehicle kinematic sensors might be more effective at identifying the abrupt signs of distraction (Fu et al., 2024). By intelligently combining these heterogeneous data sources, we can create a more complete picture of the driver's state, leading to more accurate and reliable DDD detection.

Based on the initial data analysis, various data fusion improvement techniques can be considered, including early fusion (combining raw data), feature-level fusion (combining extracted features), and decision-level fusion (combining the outputs of different models). The focus in this research would be on feature-level and decision-level fusion since they are more relevant to higher order processed which well-suited for combining the types of data available. Two primary methodologies will be prioritized:

- *Ensemble Learning* (Decision-Level Fusion): Ensemble methods operate on the principle of combining the predictions of multiple, independent models ("weak learners") to produce a single, more accurate and robust prediction ("strong learner"). In the context of this research, we can train distinct models, each specialized for a different data modality. For instance, a Random Forest model could be trained on vehicle dynamics to detect erratic driving patterns indicative of drowsiness, while a SVM is trained on physiological signals (e.g., PERCLOS) to identify markers of fatigue. The final prediction is then determined by combining the outputs of these individual models using techniques like weighted averaging or majority voting (Gwak et al., 2020). This approach is particularly valuable for its ability to reduce variance and bias, making it highly resilient to the noisy and heterogeneous data common in real-world driving scenarios.
- *DL-Based Feature Fusion:* This method leverages deep neural networks to automatically learn and combine the most salient features from multimodal data into a unified, high-level representation. This methodology optimizes and improves the Multimodal ANN models proposed in RQ1. For example, a CNN branch can process driver-facing video to extract facial features and eye-gaze patterns, while a RNN branch analyses time-series vehicle kinematics (e.g., steering angle, acceleration) to capture driving dynamics. The feature vectors extracted by these separate branches are then concatenated and fed into a final classification layer to predict DDD (Yang et al., 2020).

The optimal fusion strategy will be determined through a comparative evaluation using standard metrics such as accuracy, precision, recall, F1-score, and Area Under the Curve (AUC). By identifying and implementing the most effective fusion technique, this research axis is critical for improving the overall performance of the models developed under RQ1 and RQ2.

#### RQ4: Potential Methods for Fairness-Centred Bias Mitigation

To effectively address bias, the first crucial step is to evaluate the extent of demographic bias present in the generated DDD models. A promising method that can be employed for this purpose is using Fairness Metrics:

- *Fairness Metrics*: Beyond standard performance metrics, specific fairness metrics can be used to quantify bias more formally. Metrics such as *Demographic Parity*, which assesses whether the proportion of positive outcomes is equal across groups, *Equal Opportunity*, which checks if true positive rates are equal, and *Predictive Rate Parity*, which examines if positive predictive values are equal, can be employed (Zafar et al., 2017).

Once biases are identified and quantified, appropriate mitigation strategies can be implemented.

- *Fairness-Aware Learning (In-processing Mitigation)*: Fairness-aware learning is interesting because it directly integrates fairness into the model training process. By adding a penalty term to the loss function based on fairness metrics like demographic parity or equalized odds, the AI models can be encouraged to learn fair representations and make equitable predictions. This method is particularly suitable for the DL models being developed for multimodal DDD detection
- *Threshold Moving (Post-processing Mitigation)*: Threshold moving is a promising post-processing method due to its computational efficiency in adjusting fairness metrics without retraining models. After training the DDD models, threshold moving can be applied to fine-tune the classification thresholds based on demographic groups, aiming to equalize fairness metrics like EQD or PRP.

#### 4.2.3 Research flowchart

The proposed methodology is based on research questions and above methods aiming to understand the challenge of DDD as automation increases. First, examining the fundamental relationship between automation levels and DDD, secondly, investigating subtle cues that precede distraction and drowsiness, aiming to identify reliable DDD precursors. Third, recognizing that data is essential for AI, the focus is on optimizing the way information can be combined from diverse sources – vehicle dynamics, driver behaviour, physiological signals – to enhance the accuracy of the DDD detection models. Finally, to address the critical issue of DDD Model fairness across all demographic groups, the methodology focuses on theoretical concepts of bias and fairness in AI models. Given below in Figure 3 is a general flow diagram of the methodology for the research. The methodology is structured such that the outcomes of RQ1 investigating the relationship between automation levels and DDD provide crucial input for RQ2 identifying precursors to distraction and drowsiness and their relationship to risk and vice-versa. Research questions RQ3 and RQ4, focusing on data fusion and bias

mitigation, serve as overarching strategies to enhance the accuracy, and fairness of the AI/ML models.

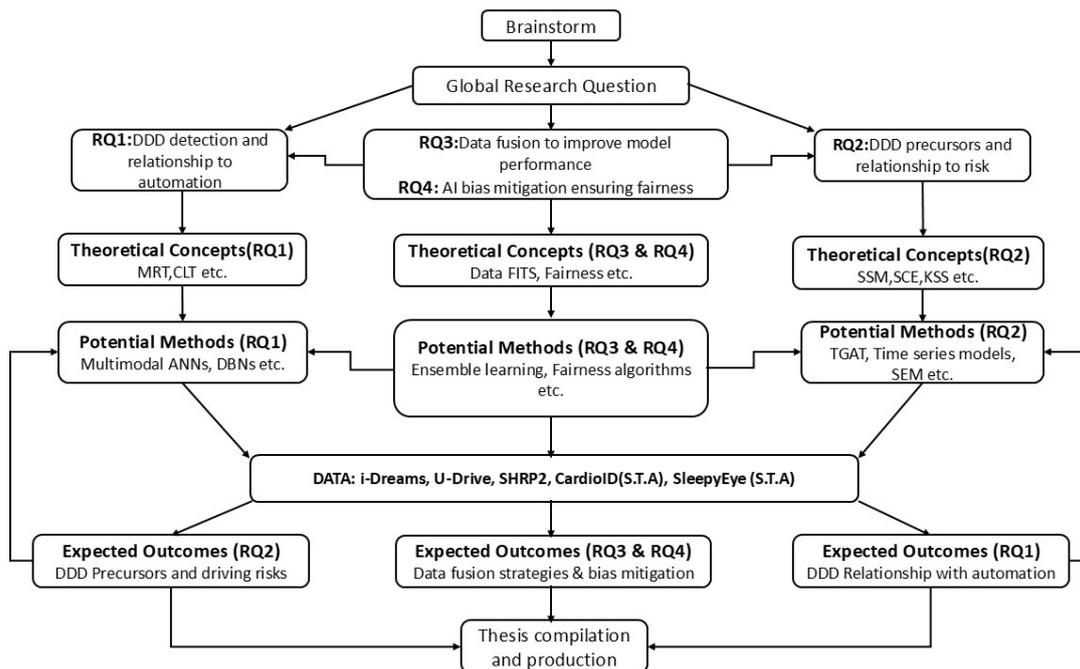


FIGURE 3: FLOWCHART OF THE RESEARCH DESIGN TO ACHIEVE RESEARCH OBJECTIVES

## 4.3 Methodologies adopted by DC 6

### 4.3.1 Objectives of the research

The research of DC6 aims to advance our understanding of road safety by moving beyond the traditional view of driving as “safe” or “unsafe”. Instead, the project adopts the Safety-II perspective, which conceptualizes safety as a spectrum of behaviours ranging from a highly unsafe to an optimally safe performance. The research systematically characterizes how drivers adapt and respond across varying contexts, such as weather, traffic conditions, and road environments within this framework. A key focus is identifying and describing what constitutes ‘optimal’ driving behaviour, considering diverse types of road users and situational factors. In addition, the project aims to analyse the factors that contribute to different driving outcomes, including behavioural adaptation, resilience, and their collective impact on safety. By integrating these insights, the research will support the development of proactive strategies and AI tools that can promote and sustain safer driving practices in complex real-world conditions. Specifically, the objectives are:

- To understand safety as a spectrum, rather than a dual state (safe vs unsafe), based on the Safety-II concept.
- To learn from drivers’ adaptation behaviours.
- To identify ‘optimal’ driving behaviours with respect to different road environments, road users, and driving contexts such as weather conditions and traffic.
- To understand the contributing factors and collective impact on the safety of different driver behaviours, including behavioural adaptation.

### 4.3.2 Overview of methods

The methodological approach in this research is grounded in the Safety-II paradigm, which frames safety not merely as the absence of accidents, but as the ability to succeed under varying conditions. This shift in perspective requires new ways of capturing and interpreting driver behaviour by focusing on incidents or violations, and how drivers successfully manage everyday driving tasks. To that end, the research applies a combination of *data-driven analysis, behavioural modelling, and pattern recognition techniques* to naturalistic driving data.

The method's core involves extracting *temporal and contextual driving patterns* from high-resolution sensor data, using *time-series analysis, clustering, and feature engineering techniques*. These methods help identify recurring behavioural profiles reflecting adaptive and resilient responses across driving environments. The goal is to understand how drivers transition between safety layers, particularly how they adapt and recover when moving toward or away from less safe states. These safety layers represent different driver performance levels and safety margins.

To better understand the underlying decision-making processes behind these adaptations, we conceptualise driver behaviour as an *optimisation problem*, where the driver continuously balances multiple goals under real-world constraints, such as safety, efficiency, and comfort. This methodological integration offers a new contribution to the field by operationalising Safety-II concepts in a scalable, data-driven framework. It complements existing research that typically focuses on risk prediction or incident detection by highlighting the conditions under which driving remains safe, flexible, and adaptive.

### 4.3.3 Research flowchart

The DC6's research follows a structured, multi-stage methodology that integrates data-driven analysis, pattern recognition, and behavioural modelling to study how drivers adapt to varying conditions and transition across different safety layers. The methodological approach builds on the Safety-II paradigm, emphasizing learning from successful and adaptive behaviours rather than focusing solely on errors or violations. Each flow stage feeds into the next, ultimately supporting the development of personalised safety systems and evidence-based definitions of optimal driving. The flowchart shown in Figure 4 supports the overarching research objectives.

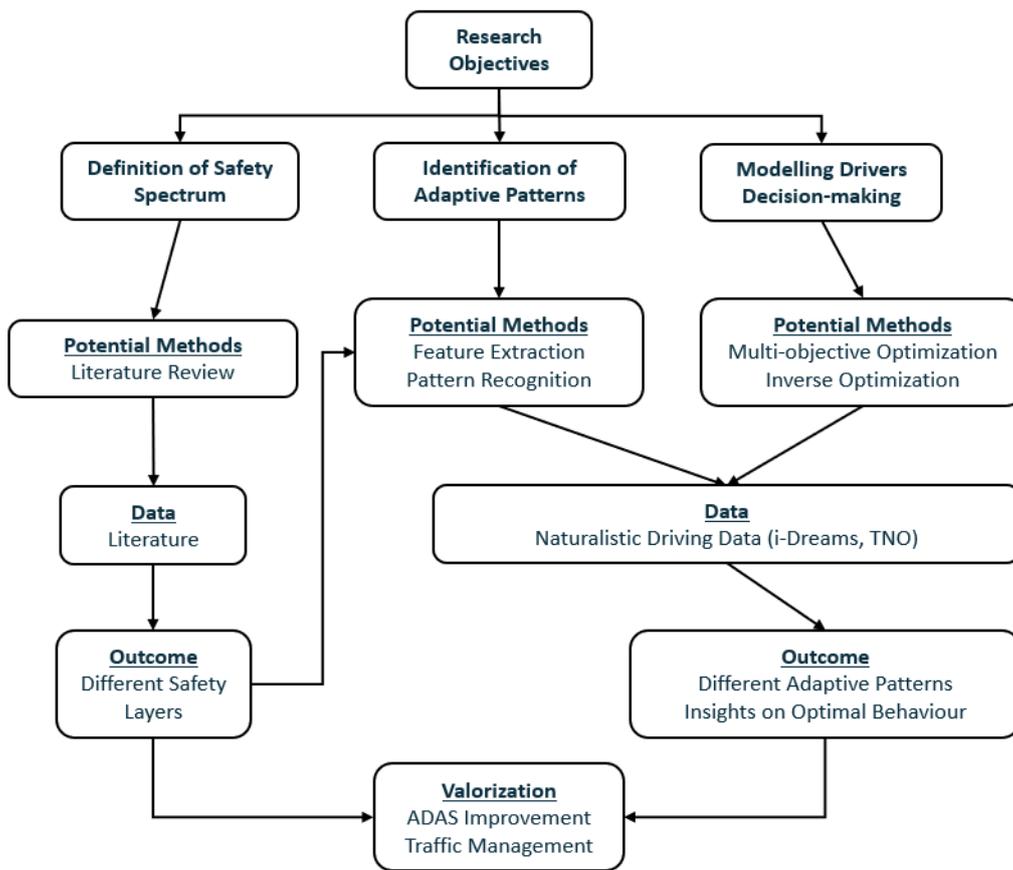


FIGURE 4: FLOWCHART OF DC6'S RESEARCH DESIGN

#### 4.3.3.1 Defining Safety Layers

The starting point is to establish a structured understanding of what constitutes varying levels of driving safety. This involves defining *safety layers* and conceptual zones of driver performance based on factors such as time-to-collision, headway, and contextual risk indicators. These layers reflect transitions from normal to elevated-risk conditions and serve as the basis for recognising meaningful behavioural adaptation.

The research draws on a thematic literature synthesis to define these layers, particularly studies that quantify conflict thresholds. Within this stage, undisturbed interaction, a key concept in the Safety-II framework, is formalised. The outcome of this step is a context-aware definition of safety layers to guide subsequent data analysis. With the safety layers established, we can now identify how and when drivers move between these layers in real-world scenarios.

#### 4.3.3.2 Scenario Selection and Data Preparation

This stage initiates the data-driven phase of the research by selecting and preparing real-world driving scenarios from the i-DREAMS and TNO datasets. The focus is on identifying *adaptive episodes*, such as car-following, speeding, or cut-in events, rich in behavioural variation and revealing how drivers respond to different contextual demands.

*Feature extraction and time-series processing* techniques convert raw sensor data into structured sequences representing variables like speed, acceleration, headway, and steering angle. The data is segmented around *moments of adaptation*, allowing for precise alignment of behaviour with surrounding environmental and contextual changes. These time-series segments form the foundation for extracting patterns of adaptation using ML and other pattern recognition methods.

#### 4.3.3.3 Pattern Recognition and Adaptation Detection

This phase applies *pattern recognition techniques* to uncover recurring behavioural strategies across drivers and contexts. The research identifies typical sequences of driver responses that reflect different adaptation styles using methods such as clustering, dimensionality reduction, and DL models (e.g., RNNs, LSTMs).

Different patterns will be classified, such as drivers transitioning between safety layers, some maintaining smooth control, and others showing late or abrupt reactions before recovery. These patterns are not interpreted as isolated events, but as expressions of a driver's adaptive capacity in response to varying demands. The labelled adaptation patterns are used to model the decision-making process behind driver responses and to identify the underlying structure of optimal behaviour.

#### 4.3.3.4 Driver Decision-making as Dynamic Optimisation

The next stage frames driver adaptation as a *multi-objective decision-making problem*, where drivers constantly optimize trade-offs between safety, comfort, and travel time. A *dynamic decision field* maps driver behaviour regarding lateral and longitudinal control, with each moment placed within a defined safety layer.

To explain why specific adaptation patterns emerge, the model applies *forward modelling techniques* (e.g., multi-objective optimisation) and *inverse optimisation* methods, which aim to infer the implicit cost functions drivers may be using. This stage provides insight into drivers' goal prioritisation under real-world constraints. These behavioural models help define what "optimal" means for different driver types and create the basis for personalising safety interventions.

#### 4.3.3.5 Application: Personalisation and Valorisation

The final stage translates the behavioural modelling results into practical applications. Personalised Advanced Driver Assistance Systems (ADAS) are one key outcome, designed to respond to each driver's specific adaptation patterns and decision-making style. Rather than relying on fixed thresholds, these systems can adjust intervention timing and intensity based on individual profiles.

Beyond ADAS, this research has implications for traffic management and training, offering new tools to identify resilient drivers, improve hazard anticipation, and promote adaptable driving strategies. The concept of "optimal driving" is thus reframed, not as a rigid target, but as an individual's ability to navigate trade-offs effectively across safety layers in diverse contexts. These applications feed directly into advancing road safety by learning from success rather than failure.

#### 4.3.3.6 Methodological Contributions and Novelty

This research introduces a novel, integrative framework that combines data-driven analysis, pattern recognition, behavioural modelling, and dynamic optimisation to operationalise the Safety-II perspective in road safety. By focusing on everyday adaptations rather than failures, it enables the understanding of how safety is achieved in practice. While most road safety research continues to rely on examining crashes, near misses, or other critical events, this work shifts the focus toward the undisturbed interactions that make up the majority of everyday driving. These interactions are a valuable source of information on how drivers continuously adapt to maintain safe performance.

The methodological innovation lies in its *multi-layered structure*: defining safety layers conceptually; detecting behavioural patterns empirically; modelling decision-making as an optimisation process; and translating insights into personalised and scalable interventions. Unlike existing studies that analyse behaviour only before or after specific events, this framework captures adaptation as a continuous process unfolding throughout routine driving. Building on the detection of driver adaptation patterns during undisturbed interactions offers a new way to study how safety emerges from normal variability in behaviour. This approach contributes new tools and perspectives to the field, supporting more intelligent safety systems.

### 4.4 Methodologies adopted by DC 7

#### 4.4.1 Objectives of the research

The following research is conducted under the supervision of OSeven Telematics and the Department of Transportation Planning and Engineering at NTUA and includes a secondment at Haskoning.

The central aim of this research is to develop a holistic artificial intelligence framework for road safety-related driver evaluation, achieved through the integration of multi-parametric datasets encompassing driver behaviour, infrastructure, and traffic conditions. The intention is to move beyond single-source or isolated-parameter analyses and instead exploit the synergistic value of heterogeneous datasets when harmonised into a unified analytical structure. This will allow the development of models capable of evaluating driving behaviour in context, providing a foundation for personalised driver assistance, targeted safety interventions, and broader transport policy support.

The specific objectives of the research are to:

- Determine the most relevant behavioural, infrastructure, and traffic parameters, along with key performance indicators (KPIs), that contribute to a holistic understanding of driver behaviour.
- Investigate the integration of supplementary modalities, such as computer vision-based driver monitoring, roadside sensing, or V2X communications, to enhance situational awareness and contextual understanding.
- Design and implement a novel AI-based fusion framework that integrates multimodal parameters to enable context-aware driver profiling and evaluation.
- Embed the developed models into telematics-based applications to deliver real-time assistance and personalised feedback to drivers.

The research is expected to deliver:

- A harmonised multimodal dataset that combines behavioural, traffic, and infrastructure information for advanced driver analysis.
- A novel AI framework that moves beyond isolated telematics indicators to achieve context-aware, holistic driver evaluation.
- Application capable of providing real-time, personalized feedback to support driver coaching and enhance safety, leveraging telematics and other multimodal data sources.

#### 4.4.2 Overview of methods

The methodological approach will progress in stages, beginning with the development of surrogate safety measures that can be reliably extracted from telematics data and then expanding to a multimodal fusion framework. Initial work has focused on harsh cornering detection, chosen because it is underrepresented in current telematics literature and is known to be associated with crash risk, especially in urban settings with frequent turns. Two complementary detection pipelines are envisaged: a supervised model where sensor placement is fixed and controlled, leveraging time-series DL architectures such as *Long Short-Term Memory (LSTM)* networks and *Transformer-based* sequence models; and an orientation-agnostic, unsupervised model for uncontrolled, real-world scenarios, for example using *Principal Component Analysis (PCA)* combined with *Isolation Forest* or *autoencoder-based anomaly detection* to identify harsh events as statistical outliers. This dual-track approach will inform strategies for other surrogate measures, including harsh braking, harsh acceleration, and mobile phone usage while driving.

The long-term plan is to extend these detection methods to a broader set of events, harmonising them into a unified behavioural profile for each driver. By combining telematics signals with infrastructure context from sources such as OpenStreetMap and traffic data from APIs, the framework will enable each detected event to be interpreted considering the road environment and prevailing traffic conditions. For example, an abrupt braking event can be distinguished as either an isolated risky manoeuvre or a defensive reaction to congestion, depending on the surrounding context.

Beyond telematics, the methodology will explore incorporating additional modalities such as V2X datasets, which can provide vehicle-to-vehicle and vehicle-to-infrastructure messages on speed harmonisation, signal phase, and roadside hazard alerts, and vision-based sources such as dashcams or roadside video analytics, which can detect lane deviations, pedestrian presence, and near-miss events. These additional data streams can act both as independent sources of safety-relevant information and as validation layers for telematics-derived detections.

The core of the framework will involve advanced AI-based fusion techniques capable of synchronising, harmonising, and jointly analysing these heterogeneous datasets. Temporal alignment will ensure that events detected in telematics data can be matched to concurrent traffic and infrastructure conditions. Spatial harmonisation through map-matching will link GPS points to the correct road segments, enabling the retrieval of geometric and regulatory attributes. Data cleaning and outlier removal will be critical to eliminate artefacts from GPS drift, sensor noise, or API anomalies.

Feature engineering will transform raw measurements into surrogate safety metrics, contextual indicators, and aggregated key performance indicators, normalised by exposure measures such as trip length or time driven.

The fused dataset will then be used to develop models for both *event-level risk assessment* and *aggregated safety scoring*. Event-level models will characterize the severity and situational context of each detected maneuver, while aggregate models will generate driver profiles and segment-level safety maps. The latter will be achieved through composite measures such as the Harsh Event Ratio and Adjusted Crash Score, which combine behavioural data with historical crash records to prioritize high-risk road segments. By linking micro-level behaviour to macro-level safety outcomes, the framework will provide a scalable methodology applicable to both individual driver feedback and city-wide safety planning.

#### 4.4.3 Research Flowchart

As shown in Figure 5, the high-level research pipeline begins with the acquisition of multisource datasets, including telematics, infrastructure, traffic, and, where possible, V2X data. These datasets will be pre-processed to achieve temporal and spatial synchronisation, with sensor calibration applied when orientation consistency is possible. Map-matching will link GPS traces to the road network, enabling the retrieval of geometric and regulatory attributes, while traffic data will contribute dynamic traffic context. Data harmonisation will address inconsistencies in sampling rates, coordinate systems, and measurement precision, producing a unified, high-quality dataset ready for feature engineering.

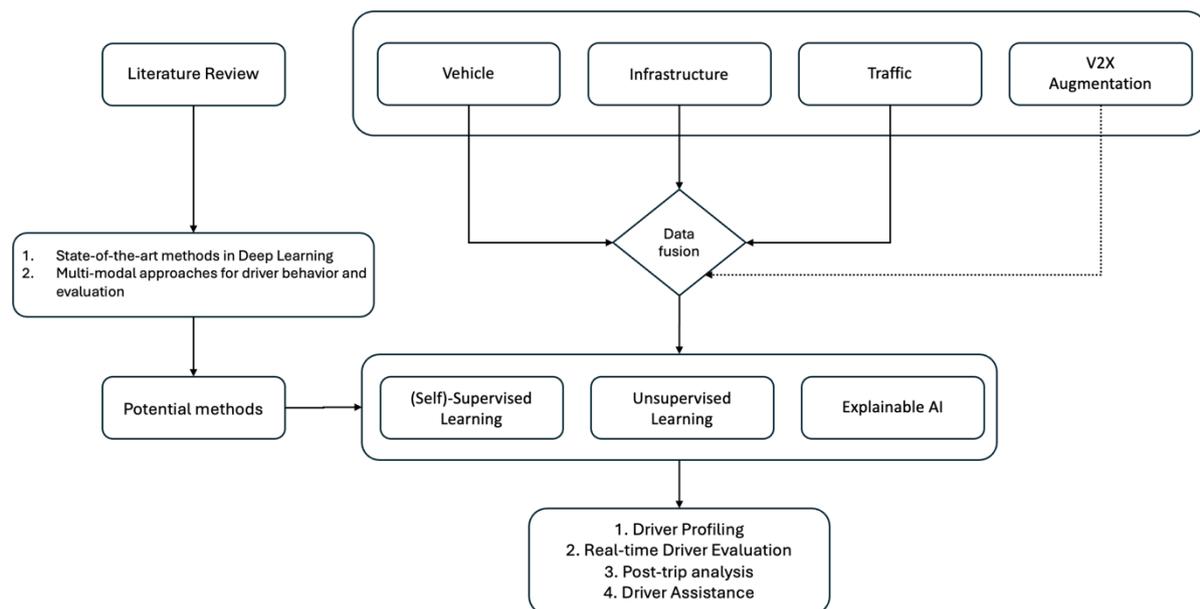


FIGURE 5: DC7 RESEARCH AREA

From this harmonised dataset, a set of surrogate safety measures will be extracted, initially focusing on harsh cornering, harsh braking, harsh acceleration, and mobile phone usage, with each event tagged to its road segment and accompanied by infrastructure and traffic context. These events will feed into both detection models and aggregated profiling processes, enabling the development of context-aware behavioural risk scores for individual drivers and

composite safety indicators for road segments. Additional modalities such as V2X and vision-based detections will be integrated either as standalone features or as verification layers, enhancing both accuracy and contextual richness.

The final stage of the pipeline will translate the analytical outputs into applications for driver assistance and support. In a telematics-based app environment, the system will be capable of delivering real-time alerts and post-trip feedback, with recommendations tailored to the driver's route, behavioural history, and environmental context. At the network level, aggregated results will support infrastructure improvements, enforcement planning, and targeted interventions in high-risk areas. By combining behavioural, infrastructural, and traffic datasets — and by designing the framework to accommodate emerging data sources such as V2X and vision analytics — this research will lay the groundwork for a next-generation, multimodal road safety evaluation system that is both scalable and transferable across regions.

The methodological novelty of this research lies in its integration of multiple underexplored elements into a single, coherent AI-driven framework for driver evaluation and assistance. While many existing studies focus on isolated behaviours or single-source data, this work combines diverse modalities; including telematics, infrastructure, traffic, and other modalities, within a harmonised pipeline. The framework is designed to operate in both controlled and uncontrolled conditions. The approach extends surrogate safety analysis beyond common measures such as harsh braking to include less-studied behaviours like harsh cornering and, in future stages, mobile phone distraction, thereby enriching the behavioural feature space. By explicitly linking micro-level behavioural events to macro-level network risk through composite indices such as the Harsh Event Ratio and Adjusted Crash Score, the framework bridges the gap between individual driver feedback and systemic road safety interventions. Furthermore, its modular structure ensures adaptability to emerging data sources and transferability across different geographic regions, making it a forward-looking methodology capable of supporting next-generation, context-aware driver assistance systems.

## 4.5 Methodologies adopted by DC 14

### 4.5.1 Objectives of the research

The original research objectives of this doctoral research aimed to exploit physiological measures obtained from naturalistic driving (i.e. tactile engagement of steering wheel, electrocardiogram (ECG), photoplethysmography (PPG), blood pressure, other physical activity KPIs) to create accurate and reliable real-time road safety models. These objectives also included to investigate scenarios, involving (i) individual driving scenarios (e.g. circumventing a fixed obstacle, distraction, reverse manoeuvring) and (ii) driver interaction scenarios (e.g. overtaking, lane-changing, right-of-way negotiating), as well as exploring the ethical dimensions of driver physiological measurements in road safety assessments, the type of biases that may arise, and how these can be eliminating for more objective and fair traffic safety assessments. However, over time, and through a comprehensive literature review and regular discussions with academic supervisors and industrial partners, it became evident that these initial objectives were too broad and abstract to effectively guide the research.

As a result, they were refined into a set of more specific, actionable research questions that better align with the scientific and technological challenges of the field. To refine these objectives, a detailed review of the current state of the art in driver monitoring systems and drowsiness detection was conducted, which revealed several emerging trends as well as notable gaps in existing research. Against this backdrop, the overarching objective of the research conducted by DC14 can be formally stated as follows: *to advance ML methods for analysing physiological data collected in both naturalistic and simulated driving environments, with the ultimate goal of improving drowsiness detection systems within the broader context of road safety*. This objective is pursued by first evaluating the capabilities and limitations of existing approaches on driver monitoring datasets, and then developing novel models that either outperform current benchmarks or provide new insights into physiological signal analysis.

The proposed methods build upon and extend key paradigms in ML, including self-supervised learning, multimodal learning, federated learning, and on-device sensing. In particular, this doctoral research aims to address the following four research questions:

- How can large amounts of unlabelled physiological data be leveraged to learn generalizable representations for drowsiness detection using self-supervision?
- How can DL be used to integrate cardiac signals with complementary physiological and behavioural signals to improve the robustness of detection models?
- How can high-performing, personalized models be developed for driver monitoring using physiological data distributed across multiple sources?
- How can lightweight physiological sensing models be deployed on edge devices for real-time drowsiness detection while ensuring both performance and privacy?

These revised research questions reflect a shift toward more focused, technically grounded problems that not only offer scientific contributions but are also practically relevant to the ongoing development of intelligent driver monitoring systems.

#### **4.5.2 Overview of methods**

To address these research questions, this doctoral research is structured around four complementary frameworks that collectively advance the analysis of physiological data for driver drowsiness detection. Each framework corresponds to a specific research focus: a self-supervised learning framework that leverages large-scale unlabelled physiological data to learn generalizable representations; a multimodal framework that integrates physiological and behavioural modalities to improve model generalization; a federated learning framework that enables the development of personalized models from distributed, imbalanced datasets while preserving privacy; and a lightweight on-device framework that ensures efficient, real-time deployment of detection models without compromising performance or user confidentiality. Together, these frameworks provide a coherent methodological pathway that links the research questions to practical and scalable advances in intelligent driver monitoring systems.

#### 4.5.2.1 A self-supervised learning framework using large-scale unlabelled physiological data for driver drowsiness detection

Electrocardiogram (ECG) signals are widely available from wearable and in-vehicle sensors, but labelled datasets for drowsiness detection remain scarce due to the cost and subjectivity of annotation. Self-supervised learning (SSL) provides a scalable alternative by leveraging pretext tasks that require no manual labelling. Recent advances in SSL for time-series data, such as contrastive learning and masked reconstruction, have demonstrated success in domains like speech and activity recognition. Applying these techniques to ECG could enable the extraction of robust, general-purpose physiological representations, which can be fine-tuned for drowsiness detection with minimal labelled data. This research question investigates whether and how SSL can effectively capture the latent structure of ECG for downstream drowsiness monitoring.

#### 4.5.2.2 A multimodal drowsiness detection framework combining physiological signals with other modalities

Drowsiness manifests through a complex interplay of physiological and behavioural cues, often varying between individuals and contexts. While ECG provides insight into autonomic nervous system responses, it may not capture the full spectrum of drowsiness indicators. Multimodal learning integrates diverse signals—such as EEG for brain activity, PPG for blood volume changes, or steering wheel angle for motor behaviour—to provide a more holistic view. The challenge lies in aligning heterogeneous data streams, dealing with missing modalities, and determining which fusion strategies (early, late, or hybrid) yield the most reliable predictions. This research question seeks to quantify the contribution of each modality and evaluate the trade-offs of multimodal architectures for drowsiness detection.

#### 4.5.2.3 A federated learning approach to develop personalized models using imbalanced physiological data distributed across in-vehicle devices

Driver drowsiness is inherently personalized, influenced by individual physiology, driving habits, and sleep patterns. Traditional centralized learning approaches are limited by privacy concerns and may fail to generalize across diverse populations. Federated learning (FL) offers a promising paradigm by training models locally on-device while sharing only model updates, not raw data. This facilitates privacy-preserving collaboration across users while enabling model personalization. However, applying FL to time-series physiological data introduces challenges in model heterogeneity, communication efficiency, and convergence under non-IID data. This question explores the design of federated architectures and strategies (e.g., personalization layers, client clustering) to build adaptive, secure drowsiness detection systems.

#### 4.5.2.4 A lightweight, quantized on device approach, achieving a balance between computational efficiency, inference speed, and model performance without compromising user privacy

Deploying ML models directly on edge devices—such as smartphones or in-vehicle embedded systems—eliminates the need for continuous data transmission, enhancing both privacy and latency. However, real-time inference on resource-constrained hardware requires model optimization techniques such as quantization, pruning, and knowledge distillation. These optimizations often trade off accuracy for speed and efficiency.

The question investigates which compression strategies are best suited for physiological time-series models (e.g., ECG-based), how to maintain performance under real-world conditions, and what design trade-offs are acceptable for practical, privacy-preserving deployment of drowsiness detection systems in cars or wearables.

### 4.5.3 Research flowchart

A preliminary conceptual structure of this doctoral research is illustrated in Figure 6. In terms of methods, this work focuses on adapting DL models to suit the properties of physiological data. From an application perspective, it addresses the complex challenges of drowsiness detection.

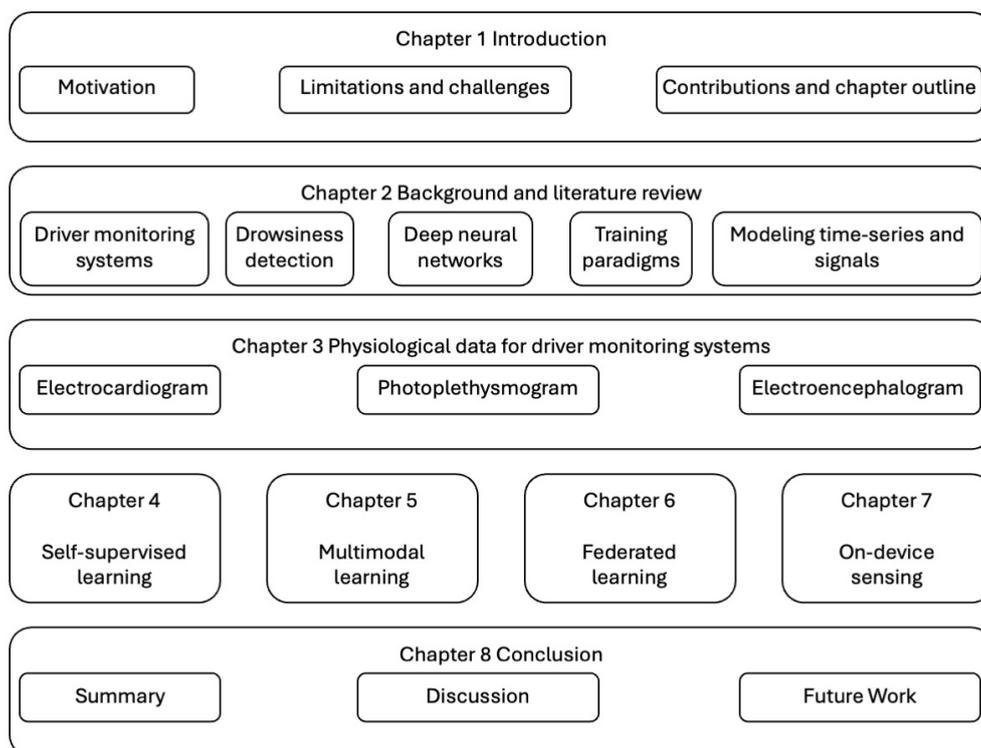


FIGURE 6: DC14'S CONCEPTUAL DOCTORAL RESEARCH OUTLINE

The doctoral research begins with an introduction that presents the motivation behind this research, outlines the limitations of current driver monitoring systems, and discusses the challenges associated with physiological data. This chapter concludes with a summary of the main contributions and the chapter outline. Next, the background chapter reviews relevant literature on driver monitoring systems and drowsiness detection, as well as foundational concepts in DL. This includes discussions on neural network architectures, training paradigms (e.g., transfer learning, self-supervised learning), and techniques for modelling time-series and signals. Following this, we analyse the core properties of the physiological data used throughout the dissertation and elaborate on the four primary research contributions in detail.

## 5 Wider methodological contributions

This section outlines the wider methodological contributions emerging from the research, focusing on how AI can enhance the detection, analysis, and mitigation of risks in road traffic environments. Initially, AI-supported detection of road user risks addresses methods for identifying early indicators of unsafe states, behaviours, or interactions across diverse road users. Building on this, AI-supported analysis of road user risks contains techniques for analyzing and modelling the underlying patterns, contexts, and contributing factors that shape risk levels in complex traffic systems. Finally, AI-supported mitigation of road user risks explores approaches that translate analytical insights into proactive interventions, ranging from personalised feedback and driver assistance to broader system-level safety measures. Together, these sections demonstrate the potential of AI methodologies to capture and understand risk and support evidence-based strategies that improve safety outcomes in real-world mobility contexts.

### 5.1 AI-supported detection of road user risks

An analysis of crash data from 2010 to 2017 reveals that abrupt lane changes were responsible for approximately 17.0% of severe accidents, while tailgating contributed 11.2% (Shawky, 2020). Tailgating typically occurs when a driver maintains a headway time of less than two seconds behind the vehicle in front, which is a cause of rear-end collisions (Michael et al., 2000). To address these risks, the lane change detection framework developed in DC 4 demonstrates how AI-based visual computing can automatically identify abrupt lane change maneuvers of surrounding vehicles without the need for prior information, manual labeling, or preselected video segments. Such automation enhances practical applicability by enabling real-time detection of dangerous cut-in behaviours. Moreover, the proposed cut-in detection algorithm refines risk estimation by distinguishing real tailgating events from false positives, which may occur when a vehicle's time headway drops below the safety threshold due to another car's sudden cut-in. This approach highlights the potential of AI-driven models to accurately capture and classify high-risk driving behaviours, thereby improving the detection and understanding of road user risks.

The ongoing research by DC5 outlines a multi-faceted methodological framework that aims to significantly advance the scientific understanding and mitigation of Driver Distraction and Drowsiness (DDD) within the context of increasing vehicle automation. This approach bridges critical gaps in transport safety research by establishing systematic methodologies that address the complex interrelationships between human factors, vehicle technology, and environmental contexts in automated driving systems. The framework is constructed upon three core pillars of methodological innovation: a shift from reactive detection to proactive precursor analysis, the integration of context-aware modeling through advanced data fusion and automation-level differentiation, and the foundational inclusion of fairness-centered bias mitigation. The methodology makes key contributions to the broader domain of AI-supported detection of road user risks. First, the hybrid multimodal approach combining physiological, behavioural, and vehicular data sources represents a shift from single-modality detection systems. It acknowledges that effective DDD detection requires comprehensive understanding of multiple interacting factors, such as cardiovascular responses, eye movement patterns, steering behaviours, and environmental conditions.

The integration of precursor detection introduces sophisticated temporal modeling capabilities that could identify subtle patterns preceding Driver distraction and drowsiness states. Finally, emphasis on fairness-centered bias mitigation techniques addresses a critical gap in AI-supported transportation safety systems. By incorporating demographic considerations and developing bias-aware algorithms, this approach ensures that AI-powered detection systems provide equitable safety benefits across diverse user populations

The research conducted in DC7, contributes to the detection of road user risks by developing AI-supported methods that fuse smartphone telematics, infrastructure topology, and traffic dynamics to identify hazardous manoeuvres such as harsh cornering. Using inertial measurements (e.g., acceleration, yaw rate) in combination with GPS trajectories, the system applies feature extraction and map-matching to align behavioural events with the underlying road network. This harmonisation enables risks to be interpreted relative to environmental context, for example distinguishing between abrupt cornering on a sharp rural bend versus the same action on a wide arterial curve. By further integrating traffic state indicators such as congestion levels and flow variability, the models can differentiate between manoeuvres that are appropriate reactions to traffic and those that reflect aggressive or unsafe driving styles. This multimodal AI pipeline thus advances proactive detection of road user risks by coupling individual driver behaviour with its situational context, supporting applications in real-time driver assistance, safety mapping, and personalised coaching.

The research conducted by DC14 contributes to the detection of road user risks by developing AI-supported methods that monitor physiological signals (such as ECG, PPG, EEG) to detect signs of driver drowsiness, a major contributor to road accidents. The methodology involves the use of both wearable and in-vehicle sensors, possibly along with video-based data (e.g., facial expressions and eye-blinking), to capture real-time indicators of cognitive and physiological fatigue. Through experiments conducted in simulated and naturalistic driving conditions, the data supports the development of robust models that can detect early signs of fatigue and inattention before they translate into hazardous behaviour. This approach extends traditional vehicle-based risk detection (like lane departure or speed monitoring) by integrating driver-centric sensing, enabling a more comprehensive understanding of risk from the human operator perspective.

## 5.2 AI-supported analysis of road user risks

The research of DC4 contributes to AI-supported risk analysis by leveraging multimodal deep learning and time-series models to systematically identify and estimate driver-related road safety risks. The methodology focuses on extracting high-resolution features from naturalistic driving data streams, such as vehicle dynamics, driver inputs, and contextual traffic conditions, using CNN models and other representation learning techniques. These features are then analyzed to detect transitions between distinct driving states (e.g., safe, risky, distracted, or aggressive), as well as to quantify the frequency and severity of risky behaviours. Furthermore, large language models are applied to summarize identified risk events and contextualize them within interpretable driver profiles, thereby bridging the gap between raw sensor data and actionable insights. This dual-layered analytical approach enables both microscopic assessments of individual trips and macroscopic evaluations of

long-term driving trends, providing a comprehensive framework for understanding and quantifying road user risk.

The research conducted by DC5 generates methodological contributions domain of AI-supported analysis of road user risks associated with Driver Distraction and Drowsiness (DDD). The core contribution lies in shifting the analytical paradigm from reactive detection of fully manifested DDD states to a proactive, predictive framework focused on identifying and quantifying risk through precursor analysis. The proposed methodologies aim to establish a sophisticated understanding of the endogenous relationship between driver state and crash risk, particularly within the complex and evolving context of vehicle automation. A core strategy is the application of multimodal ML and data-fusion techniques to capture the complex interplay of indicators, analysing of how DDD manifests at distinct SAE automation levels. Initial descriptive and statistical analyses—such as ANOVA, t-tests, and time-series modeling—map out differences in distraction duration, frequency of phone use, head-movement patterns, and lane-keeping variability etc. A key innovation is the focus on identifying systematic, recurring patterns in driver behaviour that precede distraction or drowsiness events. This precursor-based approach moves beyond traditional systems that trigger once clear signs of impairment are realised, aiming instead to intervene based on early, subtle indicators. This precursor-based analysis is intended to be achieved using graph-based neural networks for modeling the complex, time-series nature of driving data. By representing various DDD indicators as nodes in a dynamic graph, they can capture the temporal relationships and interdependencies between them. The integrated attention mechanism allows the model to learn and weigh the importance of different indicators at specific time points, making it highly effective at identifying subtle, collective patterns that signal an impending high-risk state.

The contribution of DC6 research is to apply ML and pattern recognition techniques to detect and interpret behavioural transitions across defined safety layers. Using high-resolution time-series data from naturalistic driving datasets, the methodology identifies adaptive driving patterns, specifically, how drivers respond to changing conditions and recover from them. Recurrent neural networks (RNNs), particularly LSTM architectures, are a possible solution to model temporal dependencies in driver behaviour, enabling the system to track shifts in performance and classify trajectories. By recognising these behavioural transitions, the approach provides a dynamic view of road user risk, grounded not in isolated events but in the continuum of everyday adaptation. This insight contributes to a more subtle and scalable method for analysing road safety, supporting the development of AI systems that proactively monitor and respond to risk before incidents occur.

The research of DC7 advances the analysis of road user risks by applying ML techniques to identify and estimate unsafe driving behaviours from multimodal time-series data. By using supervised, self-supervised and unsupervised techniques, the system detects unsafe events such as harsh braking, harsh acceleration and harsh cornering. These time-series based methods transform raw kinematic inputs into surrogate safety metrics that quantify the severity of hazardous events and link them with contextual indicators such as infrastructure and traffic data. This methodology is a core framework that can be integrated in different road safety applications such as modelling driver behaviour, driver profiling and driver assistance applications.

Beyond individual applications, the framework also supports systemic safety monitoring by enabling the generation of risk maps and network-level performance indicators, thereby bridging the gap between driver-level behaviour and broader traffic safety management.

The core methodological contribution of the research conducted by DC14 lies in the AI-supported analysis of drowsiness as a road user risk. The work advances the field through the development of ML models tailored to physiological time-series data. These models include techniques in self-supervised learning for unlabelled data, multimodal learning to integrate different biosignals and behavioural data, and federated learning for privacy-preserving personalization. The aim is to understand how physiological signals correlate with cognitive states such as fatigue or inattention, and how these states evolve over time. By doing so, the research enables real-time risk estimation based on internal human states, complementing external risk indicators like road conditions or vehicle dynamics. Additionally, the use of personalized models acknowledges inter-individual variability, ensuring more accurate risk assessments across diverse user populations.

### 5.3 AI-supported mitigation of road user risks

In terms of mitigation, the study of DC 4 introduces AI-driven intervention strategies that aim to reduce driver risk exposure through personalized and adaptive feedback mechanisms. The proposed system integrates post-trip evaluations with real-time monitoring, enabling both immediate and long-term interventions. Detected risky events are translated into comprehensible explanations and improvement suggestions via a safety-focused LLM, which further supports the recommendation of tailored coaching videos through an intelligent retrieval system. Over time, the models track behavioural changes, highlight improvements, and identify persistent risk patterns, thereby enabling targeted coaching interventions. By combining continuous monitoring, event-specific feedback, and longitudinal performance summaries, the methodology advances beyond generic safety assessments to deliver a proactive, personalized driver coaching framework. This AI-supported mitigation pipeline therefore not only identifies unsafe behaviours but actively contributes to their reduction, ultimately enhancing overall road safety performance.

Within the broader domain of methodologies for road user assistance and human-vehicle-environment interactions, this research of DC5 focusses on a framework that aims to advance the AI-supported mitigation of risks associated with DDD. The core methodological contribution lies in its deliberate shift from conventional, reactive detection of fully manifested driver impairment to a proactive, predictive approach aimed at preventing safety-critical events before they occur. Additionally, the explainable AI techniques planned to be used in the research would help elucidate evolving feature importance as vehicles transition from manual to assisted modes, guiding system designers toward the most salient predictors for real-time monitoring and intervention. A crucial focus in the research of employing driver profiling techniques, clustering algorithms (e.g., K-means, DBSCAN), and the calculation of SSMS such as TTC and PET helps to generate a pathway from identifying early warning signs to assessing the imminent risk of a traffic conflict or crash. This two-stage method contributes to mitigation of road user risk by enabling timely interventions based on predictive analytics rather than event detection.

The research of DC7 contributes to the mitigation of road user risks by exploring methods that translate behavioural and contextual insights into proactive interventions. Predictive models can anticipate high-risk situations and enable systems to trigger warnings, adapt driver assistance thresholds, or suggest safer driving strategies. By embedding predictive modelling and adaptive feedback into connected vehicle platforms and telematics services, this approach transforms risk analysis into tangible safety improvements, reducing the likelihood of hazardous events and supporting the development of intelligent, context-aware driver assistance solutions.

While the primary focus of the research conducted by DC14 is on detection and analysis, it also lays important groundwork for risk mitigation strategies. The deployment of lightweight, on-device drowsiness detection models enables timely interventions in real-world applications, such as alerting systems. By making real-time monitoring feasible on edge devices, the methodology supports proactive mitigation — alerting drivers before they reach dangerous levels of fatigue. In future extensions, such systems could be integrated into Advanced Driver Assistance Systems (ADAS) or semi-autonomous driving platforms, allowing for dynamic handover decisions or adaptive automation behaviour based on the driver's cognitive state.

## 6 Summary and future steps

### 6.1 Summary

DC4's research advances driver profiling by developing AI-driven methods to extract safety-relevant features from naturalistic driving data and dashcam videos. The work addresses challenges such as false tailgating detections caused by cut-in maneuvers, using object detection, tracking, and lane detection models to improve event accuracy. Beyond event-level analysis, the research builds comprehensive driver profiles at both trip and long-term scales, leveraging DL and large language models to provide interpretable, text-based feedback rather than simple scores. These profiles form the basis for personalized post-trip coaching, including tailored recommendations, coaching video suggestions, and progress tracking, with the overall goal of delivering scalable, intelligent tools that improve road safety and support the Vision Zero initiative.

The research on AI powered DDD mitigation strategies by DC5 provides a framework by integrating behavioural, physiological, and vehicular measurements systematically investigates how DDD indicators evolve across automation levels. Firstly, it seeks to elucidate the relationship between automation levels and the manifestation of DDD, quantifying how driver behaviour patterns evolve as vehicle control shifts between the human and the system. To achieve this, a hybrid approach combining multimodal ANNs with robust statistical analyses (e.g., ANOVA, mixed-effect models) would be employed. Secondly, the research moves beyond reactive detection to proactively identify systematic patterns or precursors of DDD using graph neural networks that analyze sequences of driving behaviour and identify early warning signs. The endogenous relationship between these precursors and driving risk will be quantified using driver profiling methodologies, clustering algorithms, and the calculation of SSMS. Finally, it embeds a critical focus on ethical AI by directly addressing demographic bias. It proposes a two-stage process of first quantifying bias using formal fairness metrics and subsequently implementing mitigation techniques such as in-processing fairness-Aware Learning and post-processing threshold moving to ensure equitable model performance across diverse driver populations.

The DC6's research focuses on understanding driver behaviour as a process of continuous adaptation in everyday traffic. Analyzing naturalistic driving data explores how drivers manage competing goals such as safety, efficiency, and comfort, and how the decisions can be represented through observable driving patterns. The conceptual model combines traffic psychology, behavioural theories, and data-driven approaches to capture adaptation as a form of multi-objective optimization. The aim is to provide insights into how safety is maintained in rare high-risk events and in most undisturbed interactions, ultimately contributing to the development of more personalized and supportive driver assistance systems.

The DC7's research focuses on advancing the detection, analysis, and mitigation of road user risks through multimodal data fusion and ML modelling. By integrating telematics signals from smartphones with infrastructure attributes and traffic conditions, the work develops a core framework capable of harmonizing different data sources and extracting meaningful behavioural indicators. Leveraging supervised, self-supervised. And unsupervised learning techniques, the models transform raw sensor inputs into surrogate safety metrics and

contextualized risk measures. Ultimately, the contribution lies in bridging behavioural modelling with proactive safety interventions, supporting the development of intelligent mobility systems that improve both driver safety and broader traffic system resilience.

The research conducted by DC14 is centred on enhancing driver monitoring systems through the use of physiological signals for drowsiness detection. This work aims to develop advanced ML methods capable of analysing data collected in both naturalistic and simulated driving environments, ultimately contributing to improved road safety. The research leverages cutting-edge paradigms in ML, including self-supervised learning for representation learning from unlabelled data, multimodal learning for integrating diverse physiological and behavioural signals, federated learning for privacy-preserving model development across distributed sources, and on-device sensing for real-time implementation in embedded automotive systems.

## 6.2 Future steps

The first research step of DC4 addresses the limitations of current tailgating detection methods, where false alerts can occur due to surrounding vehicles' cut-in maneuvers. Following work, DC4 will start to build methods to construct comprehensive driver profiles at both microscopic (trip-level) and macroscopic (long-term) scales. This involves the use of CNNs, multimodal DL architectures, time-series models, and large language models to extract features, detect risky events, and generate meaningful summaries of driving behaviours. The profiles are designed to go beyond numerical safety scores, offering interpretable, text-based feedback and actionable insights tailored to each driver. In addition, these profiles will then be integrated into personalized post-trip coaching systems for practical application, including generating context-specific recommendations, suggesting relevant coaching videos, and tracking driver improvement over time. By combining multimodal AI techniques with naturalistic driving datasets, DC4's work contributes novel methods for robust risk event detection and intelligent, interpretable driver profiling.

The research on AI powered DDD mitigation strategies by DC5 provides foundational methodological framework to reduce DDD and enhance driver assistance with responsible human-vehicle-environment interactions. As drivers transition into a supervisory role, the nature of distraction and the physiological onset of drowsiness fundamentally change, necessitating deeper research into monitoring disengagement, out-of-the-loop unfamiliarity, and takeover readiness. Thus, the future steps related to this research would be undertaking collaborative projects with industrial partner as part of the industrial IVORY PHD program. This would involve analysis of extensive naturalistic driving datasets to empirically model the evolution DDD across various automation levels. This effort will directly capitalize the partners' existing projects and curated data, such as identified distraction related SCE scenarios and documented drowsiness episodes, allowing the validation and detection of DDD precursors and their link to risk. It will also enable to leverage the partners' invaluable domain expertise in driver behaviour and compare findings with regulations such as EU GSR2 and Euro NCAP protocols.

Future DC6's research could leverage ML and optimization advances to improve the recognition of adaptive patterns and refine the modeling of decision-making processes.

From an application perspective, integrating these insights into driver assistance systems has the potential to move beyond rule-based interventions toward context-sensitive support that aligns with natural driver behaviour. In the longer term, collaborations with industry and policymakers could accelerate the translation of these ideas into real-world systems, enhancing safety and personalization in future mobility.

Future steps for DC7's research will focus on advancing the multimodal data fusion framework by incorporating additional perception-oriented modalities, particularly vision-based sources such as onboard camera data. This will enable a more holistic understanding of the driving context, capturing not only how the vehicle moves, but also what the driver sees and how the surrounding environment evolves. This fusion framework aims to support context-aware driver assistance and safety assessment in real-world scenarios. Finally, to ensure ethical scalability, privacy-preserving learning paradigms such as federated learning will be explored to facilitate collaborative model development across distributed data sources without exposing sensitive personal information.

Looking ahead, future research could extend the work of DC14 by exploring larger-scale studies to better capture inter-individual differences and adapt drowsiness detection systems to diverse driver populations. From a methodological perspective, advancing interpretable AI approaches would not only strengthen the scientific contribution to the field of driver state monitoring but also build trust in real-world applications. In terms of deployment, collaboration with automotive manufacturers and regulatory bodies could accelerate the translation of research into commercial in-vehicle systems, enhancing drowsiness detection and consequently road safety at scale. Furthermore, the techniques developed, such as privacy-preserving federated learning and efficient on-device sensing, could be applied beyond transportation to domains like healthcare monitoring, broadening the societal impact of this work.

## References

- A. Meyerson, J. Eklind, F. Fischer, M. Aramrattana, I. Fredriksson, and C. Ahlström, "Effects of daylight and darkness at daytime versus night-time on driver sleepiness: A driving simulator study," *Transportation Research Interdisciplinary Perspectives*, vol. 24, p. 101087, 2024.
- Abdel-Aty, M., Wang, Z., Zheng, O., & Abdelraouf, A. (2023). Advances and Applications of Computer Vision Techniques in Vehicle Trajectory Generation and Surrogate Traffic Safety Indicators. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4398932>
- Abdelrahman, A. E., Hassanein, H. S., & Abu-Ali, N. (2022). Robust Data-Driven Framework for Driver Behavior Profiling Using Supervised Machine Learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(4), 3336–3350. <https://doi.org/10.1109/TITS.2020.3035700>
- Ackley, J. L., Puranik, T. G., & Mavris, D. (2020, June 15). A Supervised Learning Approach for Safety Event Precursor Identification in Commercial Aviation. *AIAA AVIATION 2020 FORUM*. AIAA AVIATION 2020 FORUM, VIRTUAL EVENT. <https://doi.org/10.2514/6.2020-2880>
- Adler, P., Falk, C., Friedler, S. A., Nix, T., Rybeck, G., Scheidegger, C., Smith, B., & Venkatasubramanian, S. (2018). Auditing black-box models for indirect influence. *Knowledge and Information Systems*, 54(1), 95–122. <https://doi.org/10.1007/s10115-017-1116-3>
- Afghari, A. P., Washington, S., Haque, Md. M., & Li, Z. (2018). A comprehensive joint econometric model of motor vehicle crashes arising from multiple sources of risk. *Analytic Methods in Accident Research*, 18, 1–14. <https://doi.org/10.1016/j.amar.2018.03.002>
- Albadawi, Y., AlRedhaei, A., & Takruri, M. (2023). Real-time machine learning-based driver drowsiness detection using visual features. *Journal of Imaging*, 9(5), Article 91. <https://doi.org/10.3390/jimaging9050091>
- Albadawi, Y., Takruri, M., & Awad, M. (2022). A review of recent developments in driver drowsiness detection systems. *Sensors*, 22(5), Article 2069. <https://doi.org/10.3390/s22052069>
- Al-Hussein, W. A., Por, L. Y., Kiah, M. L. M., & Zaidan, B. B. (2022). Driver Behavior Profiling and Recognition Using Deep-Learning Methods: In Accordance with Traffic Regulations and Experts Guidelines. *International Journal of Environmental Research and Public Health*, 19(3), 1470. <https://doi.org/10.3390/ijerph19031470>
- Ayas, S., Donmez, B., & Tang, X. (2024). Drowsiness Mitigation Through Driver State Monitoring Systems: A Scoping Review. *Human Factors*, 66(9), 2218–2243. <https://doi.org/10.1177/00187208231208523>
- Bakker, B., et al. (2022). A multi-stage, multi-feature machine learning approach to detect driver sleepiness in naturalistic road driving conditions. *IEEE Transactions on Intelligent Transportation Systems*, 23(5), 4791–4800. <https://doi.org/10.1109/TITS.2021.3090272>

Barua, S., Ahmed, M. U., Ahlström, C., & Begum, S. (2019). Automatic driver sleepiness detection using EEG, EOG and contextual information. *Expert Systems with Applications*, 115, 121–135. <https://doi.org/10.1016/j.eswa.2018.07.054>

Bin Mohamad Azmi, M. M., & Kamaru Zaman, F. H. (2024). Driver drowsiness detection using vision transformer. 2024 IEEE 14th Symposium on Computer Applications & Industrial Electronics (ISCAIE), 329–336. <https://doi.org/10.1109/ISCAIE61308.2024.10576317>

Bock, O., Stojan, R., Wechsler, K., Mack, M., Voelcker-Rehage, C., 2021. Distracting tasks have persisting effects on young and older drivers' braking performance. *Accid. Anal. Prev.* 161, 106363. <https://doi.org/10.1016/j.aap.2021.106363>

Caber, N., Ahmad, B. I., Liang, J., Godsill, S., Bremers, A., Thomas, P., Oxtoby, D., & Skrypchuk, L. (2023). Driver profiling and Bayesian workload estimation using naturalistic peripheral detection study data [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2303.14720>

Castignani, G., Derrmann, T., Frank, R., & Engel, T. (2015). Driver behaviour profiling using smartphones: A low-cost platform for driver monitoring. *IEEE Intelligent Transportation Systems Magazine*, 7(1), 91–102. <https://doi.org/10.1109/MITS.2014.2328673>

Chan, I. W., Badescu, A. L., & Lin, X. S. (2025). Assessing Driving Risk Through Unsupervised Detection of Anomalies in Telematics Time Series Data. *ASTIN Bulletin*, 55(2), 205–241. <https://doi.org/10.1017/asb.2025.14>

Chen, H.-Y. W., Xie, J. Y., & Donmez, B. (2023). Gamification of driver distraction feedback: A simulator study with younger drivers. *IEEE Transactions on Human-Machine Systems*, 53(5), 813–822. <https://doi.org/10.1109/THMS.2023.3298309>

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. E. (2020). A simple framework for contrastive learning of visual representations [Preprint]. arXiv. <https://arxiv.org/abs/2002.05709>

Chen, X., Wang, H., Razi, A., Russo, B., Pacheco, J., Roberts, J., Wishart, J., Head, L., & Baca, A. G. (2022). Network-level safety metrics for overall traffic safety assessment: A case study. *IEEE Access*, 11, 17755–17778.

Cooper, J. M., Crabtree, K. W., McDonnell, A. S., May, D., Strayer, S. C., Tsogetbaatar, T., Cook, D. R., Alexander, P. A., Sanbonmatsu, D. M., & Strayer, D. L. (2023). Driver behaviour while using Level 2 vehicle automation: A hybrid naturalistic study. *Cognitive Research: Principles and Implications*, 8(1), Article 71. <https://doi.org/10.1186/s41235-023-00527-5>

Cordellieri, P., Baralla, F., Ferlazzo, F., Sgalla, R., Piccardi, L., & Giannini, A. M. (2016). Gender effects in young road users on road safety attitudes, behaviours and risk perception. *Frontiers in Psychology*, 7, Article 1412. <https://doi.org/10.3389/fpsyg.2016.01412>

Du, G., Long, S., Li, C., Wang, Z., & Liu, P. X. (2023). A product fuzzy convolutional network for detecting driving fatigue. *IEEE Transactions on Cybernetics*, 53(7), 4175–4188.

Dunn, N. J., Dingus, T. A., Soccolich, S., & Horrey, W. J. (2021). Investigating the impact of driving automation systems on distracted driving behaviours. *Accident Analysis and Prevention*, 156, Article 106152. <https://doi.org/10.1016/j.aap.2021.106152>

Durrant-Whyte, H., & Henderson, T. C. (2008). Multisensor data fusion. In B. Siciliano & O. Khatib (Eds.), *Springer handbook of robotics* (pp. [585-610]). Springer. [https://doi.org/10.1007/978-3-540-30301-5\\_26](https://doi.org/10.1007/978-3-540-30301-5_26)

Ebel, P., Lingenfelder, C., & Vogelsang, A. (2023). On the forces of driver distraction: Explainable predictions for the visual demand of in-vehicle touchscreen interactions. *Accident Analysis and Prevention*, 183, Article 106956. <https://doi.org/10.1016/j.aap.2023.106956>

Ellison, A. B., Greaves, S. P., & Bliemer, M. C. J. (2015). Driver behaviour profiles for road safety analysis. *Accident Analysis & Prevention*, 76, 118–132. <https://doi.org/10.1016/j.aap.2015.01.009>  
Eren, H., Makinist, S., Akin, E., & Yilmaz, A. (2012). Estimating driving behaviour by a smartphone. 2012 IEEE Intelligent Vehicles Symposium, 234–239. <https://doi.org/10.1109/IVS.2012.6232298>

Felzmann, H., Villaronga, E. F., Lutz, C., & Tamò-Larrioux, A. (2019). Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society*, 6(1), 2053951719860542. <https://doi.org/10.1177/2053951719860542>

Fernández-Llorca, D., Biparva, M., Izquierdo-Gonzalo, R., & Tsotsos, J. K. (2020). Two-stream networks for lane-change prediction of surrounding vehicles. 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), 1–6. <https://doi.org/10.1109/ITSC45102.2020.9294326>

Ferrara, E. (2024). The Butterfly Effect in Artificial Intelligence Systems: Implications for AI Bias and Fairness. *Machine Learning with Applications*, 15, 100525. <https://doi.org/10.1016/j.mlwa.2024.100525>

Fitzharris, M., Liu, S., Stephens, A. N., & Lenné, M. G. (2017). The relative importance of real-time in-cab and external feedback in managing fatigue in real-world commercial transport operations. *Traffic Injury Prevention*, 18(sup1), S71–S78. <https://doi.org/10.1080/15389588.2017.1306855>

Foulds, J. R., Islam, R., Keya, K. N., & Pan, S. (2020). An Intersectional Definition of Fairness. 2020 IEEE 36th International Conference on Data Engineering (ICDE), 1918–1921. <https://doi.org/10.1109/ICDE48307.2020.00203>

Fors, C., Ahlström, C., Sörner, P., Kovaceva, J., Hasselberg, E., Krantz, M., Grönvall, J., Kircher, K., & Anund, A. (2011). Camera-based sleepiness detection: Final report of the project SleepEYE

Fredriksson, R., Lenné, M. G., Van Montfort, S., & Grover, C. (2021). European NCAP program developments to address driver distraction, drowsiness and sudden sickness. *Frontiers in Neuroergonomics*, 2, Article 786674. <https://doi.org/10.3389/fnrgo.2021.786674>

Frontiers in Massive Data Analysis. (2013). National Academies Press, Washington, D.C. <https://doi.org/10.17226/18374>

Garcia-Constantino, M., Missier, P., Blythe, P., & Guo, A. W. (2014). Measuring the impact of cognitive distractions on driving performance using time series analysis. 2014 IEEE 17th International Conference on Intelligent Transportation Systems (ITSC), 2047–2053. <https://doi.org/10.1109/ITSC.2014.6958005>

Gershon, P., Mehler, B., & Reimer, B. (2023). Driver response and recovery following automation initiated disengagement in real-world hands-free driving. *Traffic Injury Prevention*, 24(4), 356–361. <https://doi.org/10.1080/15389588.2023.2189990>

Guo, Y., Yang, K., & Wu, Y. (2025). A multi-modality attention network for driver fatigue detection based on frontal EEG, EDA and PPG signals. *IEEE Journal of Biomedical and Health Informatics*, 29(6), 4009–4022. <https://doi.org/10.1109/JBHI.2025.3527964>

Gwak, J., Hirao, A., & Shino, M. (2020). An investigation of early detection of driver drowsiness using ensemble machine learning based on hybrid sensing. *Applied Sciences*, 10(8), Article 2890. <https://doi.org/10.3390/app10082890>

Hassan, O. F., Ibrahim, A. F., Gomaa, A., Makhoul, M. A., & Hafiz, B. (2025). Real-time driver drowsiness detection using transformer architectures: A novel deep learning approach. *Scientific Reports*, 15(1), Article 17493. <https://doi.org/10.1038/s41598-025-02111-x>

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>

Hollnagel, E. (2018). *Safety-I and Safety-II: The past and future of safety management*. CRC Press. <https://doi.org/10.1201/9781315607511>

Huang, Y., Liu, C., Chang, F., & Lu, Y. (2024). Self-supervised multi-granularity graph attention network for vision-based driver fatigue detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 8(4), 3067–3080. <https://doi.org/10.1109/TETCI.2024.3369937>

Huang, Z., Wu, J., & Lv, C. (2022). Driving behaviour modeling using naturalistic human driving data with inverse reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 10239–10251. <https://doi.org/10.1109/TITS.2021.3088935>

Hultman, M., Johansson, I., Lindqvist, F., & Ahlström, C. (2021). Driver sleepiness detection with deep neural networks using electrophysiological data. *Physiological Measurement*, 42(3), Article 034001. <https://doi.org/10.1088/1361-6579/abe91e>

Hungund, A. P., Pai, G., & Pradhan, A. K. (2021). Systematic review of research on driver distraction in the context of advanced driver assistance systems. *Transportation Research*

Record: Journal of the Transportation Research Board, 2675(9), 756–765.  
<https://doi.org/10.1177/03611981211004129>

Huo, X.-Q., Zheng, W.-L., & Lu, B.-L. (2016). Driving fatigue detection with fusion of EEG and forehead EOG. 2016 International Joint Conference on Neural Networks (IJCNN), 897–904.  
<https://doi.org/10.1109/IJCNN.2016.7727294>

Izquierdo, R., Quintanar, A., Parra, I., Fernández-Llorca, D., & Sotelo, M. A. (2019). The PREVENTION dataset: A novel benchmark for PREdiction of VEHicles iNTentIONS. 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 3114–3121.  
<https://doi.org/10.1109/ITSC.2019.8917433>

Jackson, M.L., Raj, S., Croft, R.J., Hayley, A.C., Downey, L.A., Kennedy, G.A., Howard, M.E., 2016b. Slow eyelid closure as a measure of driver drowsiness and its relationship to performance. *Traffic Inj. Prev.* 17, 251–257.  
<https://doi.org/10.1080/15389588.2015.1055327>

Jarndal, A., Tawfik, H., Siam, A. I., Alsyof, I., & Cheaitou, A. (2024). A real-time vision transformers-based system for enhanced driver drowsiness detection and vehicle safety. *IEEE Access*, 13, 1790–1803. <https://doi.org/10.1109/ACCESS.2024.3522111>

Jiang, Y., Deng, W., Wang, J., & Zhu, B. (2018). Studies on drivers' driving styles based on inverse reinforcement learning (SAE Technical Paper 2018-01-0612). SAE International.  
<https://doi.org/10.4271/2018-01-0612>

Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. *Information Fusion*, 14(1), 28–44.  
<https://doi.org/10.1016/j.inffus.2011.08.001>

Khattak, A. J., Ahmad, N., Wali, B., & Dumbaugh, E. (2021). A taxonomy of driving errors and violations: Evidence from the naturalistic driving study. *Accident Analysis & Prevention*, 151, 105873. <https://doi.org/10.1016/j.aap.2020.105873>

Ko, W., Choe, J., & Kang, J. (2025). SLEEP-SAFE: Self-supervised learning for estimating electroencephalogram patterns with structural analysis of fatigue evidence. *IEEE Access*, 13, 35805–35817. <https://doi.org/10.1109/ACCESS.2025.3545094>

Koay, H. V., Chuah, J. H., Chow, C.-O., & Chang, Y.-L. (2022). Detecting and recognizing driver distraction through various data modality using machine learning: A review, recent advances, simplified framework and open challenges (2014–2021). *Engineering Applications of Artificial Intelligence*, 115, 105309. <https://doi.org/10.1016/j.engappai.2022.105309>

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6), 84–90.  
<https://doi.org/10.1145/3065386>

Kuznietsov, A., Gyevnar, B., Wang, C., Peters, S., & Albrecht, S. V. (2024). Explainable AI for Safe and Trustworthy Autonomous Driving: A Systematic Review. *IEEE Transactions on*

Intelligent Transportation Systems, 25(12), 19342–19364.  
<https://doi.org/10.1109/TITS.2024.3474469>

LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521, 436–444.  
<https://doi.org/10.1038/nature14539>

Lee, M. S., Park, J., Jang, Y. J., Kim, W., & Yoon, D. (2018). Individual stable driving pattern analysis for evaluating driver readiness at autonomous driving levels 2 and 3. 2018 International Conference on Information and Communication Technology Convergence (ICTC), 315–319. <https://doi.org/10.1109/ICTC.2018.8539722>

Li, S., Qian, D., Li, P., Yuan, X., & Fang, Q. (2024). Risk prediction model for distracted driving: Characterizing interactions of eye glances and manual sequences. *Travel Behaviour and Society*, 37, 100851. <https://doi.org/10.1016/j.tbs.2024.100851>

Li, X., Hong, L., Wang, J.-C., & Liu, X. (2019). Fatigue driving detection model based on multi-feature fusion and semi-supervised active learning. *IET Intelligent Transport Systems*, 13(9), 1401–1409. (Note: Issue number was inferred from standard formats, as it was missing). <https://doi.org/10.1049/iet-its.2018.5590>

Li, Z., Li, S. E., Li, R., Cheng, B., & Shi, J. (2017). Online detection of driver fatigue using steering wheel angles for real driving conditions. *Sensors*, 17(3), Article 495. <https://doi.org/10.3390/s17030495>

Liang, K., Wang, J., & Bhalerao, A. (2023). Lane Change Classification and Prediction with Action Recognition Networks. In L. Karlinsky, T. Michaeli, & K. Nishino (Eds.), *Computer Vision – ECCV 2022 Workshops* (pp. 617–632). Springer Nature Switzerland. [https://doi.org/10.1007/978-3-031-25056-9\\_39](https://doi.org/10.1007/978-3-031-25056-9_39)

Liao, X., Zhao, Z., Barth, M. J., Abdelraouf, A., Gupta, R., Han, K., Ma, J., & Wu, G. (2025). A review of personalization in driving behaviour: Dataset, modeling, and validation. *IEEE Transactions on Intelligent Vehicles*, 10(2), 1241–1262. <https://doi.org/10.1109/TIV.2024.3425647>

Lizarazo, C., Panasci, A., Madden, S., & Tebbens, R. (2023). What do surrogate safety metrics measure? Understanding driving safety as a continuum. *Accident Analysis and Prevention*, 195, Article 107245. <https://doi.org/10.1016/j.aap.2023.107245>

Luo, A., Chen, S., & Xv, B. (2017). Enhanced map-matching algorithm with a hidden Markov model for mobile phone positioning. *ISPRS International Journal of Geo-Information*, 6(11), Article 327. <https://doi.org/10.3390/ijgi6110327>

Malekpour, M.-R., Ghamari, S.-H., Ghasemi, E., Hejaziyeganeh, S., Abbasi-Kangevari, M., Bhalla, K., Rezaei, N., Shahraz, S., Dilmaghani-Marand, A., Taghi Heydari, S., Rezaei, N., Lankarani, K. B., & Farzadfar, F. (2023). The effect of real-time feedback and incentives on speeding behaviours using telematics: A randomized controlled trial. *Accident Analysis and Prevention*, 191, Article 107216. <https://doi.org/10.1016/j.aap.2023.107216>

Mantouka, E.G., Barmounakis, E.N., Vlahogianni, E.I., 2019. Identifying driving safety profiles from smartphone data using unsupervised learning. *Saf. Sci.* 119, 84–90. <https://doi.org/10.1016/j.ssci.2019.01.025>

Mårtensson, H., Keelan, O., & Ahlström, C. (2019). Driver sleepiness classification based on physiological data and driving performance from real road driving. *IEEE Transactions on Intelligent Transportation Systems*, 20(2), 421–430. <https://doi.org/10.1109/TITS.2018.2814207>

McCarty, D., & Kim, H. W. (2024). Risky behaviours and road safety: An exploration of age and gender influences on road accident rates. *PLOS ONE*, 19(1), e0296663. <https://doi.org/10.1371/journal.pone.0296663>

Meuleners, L., Fraser, M., Stevenson, M., & Roberts, P. (2023). Personalized driving safety: Using telematics to reduce risky driving behaviour among young drivers. *Journal of Safety Research*, 86, 164–173. <https://doi.org/10.1016/j.jsr.2023.05.007>

Michael, P. G., Leeming, F. C., & Dwyer, W. O. (2000). Headway on urban streets: Observational data and an intervention to decrease tailgating. *Transportation Research Part F: Traffic Psychology and Behaviour*, 3(2), 55–64. [https://doi.org/10.1016/S1369-8478\(00\)00015-2](https://doi.org/10.1016/S1369-8478(00)00015-2)

Min, J., Cai, M., Gou, C., et al. (2023). Fusion of forehead EEG with machine vision for real-time fatigue detection in an automatic processing pipeline. *Neural Computing and Applications*, 35, 8859–8872. <https://doi.org/10.1007/s00521-022-07466-0>

Mobini Seraji, M. H., Shaffiee Haghshenas, S., Shaffiee Haghshenas, S., Simic, V., Pamucar, D., Guido, G., & Astarita, V. (2025). A state-of-the-art review on machine learning techniques for driving behaviour analysis: Clustering and classification approaches. *Complex and Intelligent Systems*, 11(9), Article 386. <https://doi.org/10.1007/s40747-025-01988-5>

Mou, L., et al. (2023). Isotropic self-supervised learning for driver drowsiness detection with attention-based multimodal fusion. *IEEE Transactions on Multimedia*, 25, 529–542. <https://doi.org/10.1109/TMM.2021.3128738>

Narayanan, A., Chen, Y.-T., & Malla, S. (2018). Semi-supervised learning: Fusion of self-supervised, supervised learning, and multimodal cues for tactical driver behaviour detection [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.1807.00864>

Newson, P., & Krumm, J. (2009). Hidden Markov map matching through noise and sparseness. *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 336–343. <https://doi.org/10.1145/1653771.1653818>

Nikolaou, D., Ziakopoulos, A., Kontaxi, A., Theofilatos, A. A., & Yannis, G. (2025). Spatial analysis of telematics-based surrogate safety measures. *Journal of Safety Research*, 92, 98–108. <https://doi.org/10.1016/j.jsr.2024.09.012>

Noble, A.M., Miles, M., Perez, M.A., Guo, F., Klauer, S.G., 2021. Evaluating driver eye glance behavior and secondary task engagement while using driving automation systems. *Accid. Anal. Prev.* 151, 105959. <https://doi.org/10.1016/j.aap.2020.105959>

Oban, V., Kaya, M., Safi, G., Çimen, İ. N., Çatak, T., Karadağ, B., Görmüş, G., Çandır, A., & Alagöz, F. (2024). Profiling Driver Behaviors Using AI-Based Methods and Deep Learning Techniques for Improving Road Safety: A Comparative Study of Algorithms. 2024 9th International Conference on Computer Science and Engineering (UBMK), 955–960. <https://doi.org/10.1109/UBMK63289.2024.10773587>

Ounoughi, C., & Ben Yahia, S. (2023). Data fusion for ITS: A systematic literature review. *Information Fusion*, 89, 267–291. <https://doi.org/10.1016/j.inffus.2022.08.016>

Papadimitriou, E., Pooyan Afghari, A., Tselentis, D., & Van Gelder, P. (2022). Road-safety-II: Opportunities and barriers for an enhanced road safety vision. *Accident Analysis and Prevention*, 174, Article 106723. <https://doi.org/10.1016/j.aap.2022.106723>

Paul, A., Boyle, L., Boer, E., Tippin, J., & Rizzo, M. (2017). Steering entropy changes as a function of microsleeps. *Proceedings of the Ninth International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design*, 441–447. <https://doi.org/10.17077/drivingassessment.1196>

Pozueco, L., Gupta, N., Pañeda, X. G., García, R., Tuero, A. G., Melendi, D., Rionda, A., & Corcoba, V. (2020). Analysis of driving patterns and on-board feedback-based training for proactive road safety monitoring. *IEEE Transactions on Human-Machine Systems*, 50(6), 529–537. <https://doi.org/10.1109/THMS.2020.3027525>

Putzel, P., & Lee, S. (2022). Blackbox Post-Processing for Multiclass Fairness (No. arXiv:2201.04461). arXiv. <https://doi.org/10.48550/arXiv.2201.04461>

Qian, K., Koike, T., Nakamura, T., Schuller, B. W., & Yamamoto, Y. (2022). Learning multimodal representations for drowsiness detection. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 11539–11548. <https://doi.org/10.1109/TITS.2021.3105326>

Quddus, M., Ochieng, W., Zhao, L., & Noland, R. (2003). A general map matching algorithm for transport telematics applications. *GPS Solutions*, 7(3), 157–167. <https://doi.org/10.1007/s10291-003-0069-z>

Ramzan, M., Khan, H.U., Awan, S.M., Ismail, A., Ilyas, M., Mahmood, A., 2019. A Survey on State-of-the-Art Drowsiness Detection Techniques. *IEEE Access* 7, 61904–61919. <https://doi.org/10.1109/ACCESS.2019.2914373>

Regan, M.A., Hallett, C., Gordon, C.P., 2011. Driver distraction and driver inattention: Definition, relationship and taxonomy. *Accid. Anal. Prev.* 43, 1771–1781. <https://doi.org/10.1016/j.aap.2011.04.008>

Sagberg, F., Selpi, S., Bianchi Piccinini, G., & Engström, J. (2015). A review of research on driving styles and road safety. *Human Factors*, 57(8), 1221–1254. <https://doi.org/10.1177/0018720815591313>

Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors*, 48(2), 362–380. <https://doi.org/10.1518/00187200677724417>

Sangrar, R., Griffith, L. E., Letts, L., Ma, J., Porter, M. M., & Vrkljan, B. (2022). Refreshing Older Adults' Driving Skills (ROADSkills): A randomized controlled trial examining the effect of video feedback. *Transportation Research Part F: Traffic Psychology and Behaviour*, 90, 35–46. <https://doi.org/10.1016/j.trf.2022.07.019>

Santos, M., Coelho, P. J., Pires, I. M., Gonçalves, P., & Dias, G. P. (2024). An overview of machine learning algorithms to reduce driver fatigue and distraction-related traffic accidents. *Procedia Computer Science*, 238, 97–102. <https://doi.org/10.1016/j.procs.2024.06.003>

Sarker, S., Haque, M. M., & Dewan, M. A. A. (2021). Driving maneuver classification using domain specific knowledge and transfer learning. *IEEE Access*, 9, 86590–86606. <https://doi.org/10.1109/ACCESS.2021.3089660>

Shaffer, F., & Ginsberg, J. P. (2017). An Overview of Heart Rate Variability Metrics and Norms. *Frontiers in public health*, 5, 258. <https://doi.org/10.3389/fpubh.2017.00258>

Shahbakhti, M., Beiramvand, M., Rejer, I., et al. (2022). Simultaneous eye blink characterization and elimination from low-channel prefrontal EEG signals enhances driver drowsiness detection. *IEEE Journal of Biomedical and Health Informatics*, 26(3), 1001–1012. <https://doi.org/10.1109/JBHI.2021.3096984>

Shahid, A., Shen, J., & Shapiro, C. M. (2012). STOP, that and one hundred other sleep scales. In R. B. Berry, R. Brooks, C. E. Gamaldo, S. M. Harding, C. L. Marcus, & B. V. Vaughn (Eds.), *The AASM manual for the scoring of sleep and associated events: Rules, terminology and technical specifications* (pp. 209–210). American Academy of Sleep Medicine. <https://doi.org/10.1007/978-1-4419-9893-4>

Shawky, M. (2020). Factors affecting lane change crashes. *IATSS Research*, 44(2), 155–161. <https://doi.org/10.1016/j.iatssr.2019.12.002>

Shirole, V., Shahade, A. K., & Deshmukh, P. V. (2025). A comprehensive review on data-driven driver behaviour scoring in vehicles: Technologies, challenges and future directions. *Discover Artificial Intelligence*, 5(1), 26. <https://doi.org/10.1007/s44163-025-00244-6>

Singh, H., & Kathuria, A. (2021). Profiling drivers to assess safe and eco-driving behavior—A systematic review of naturalistic driving studies. *Accident; Analysis and Prevention*, 161, 106349. <https://doi.org/10.1016/j.aap.2021.106349>

Soares, S., Campos, C., Leitão, J. M., Lobo, A., Couto, A., & Ferreira, S. (2021). Distractive tasks and the influence of driver attributes. *Sustainability*, 13(9), Article 5094. <https://doi.org/10.3390/su13095094>

Stevenson, M., Harris, A., Wijnands, J. S., & Mortimer, D. (2021). The effect of telematic based feedback and financial incentives on driving behaviour: A randomised trial. *Accident Analysis & Prevention*, 159, 106278. <https://doi.org/10.1016/j.aap.2021.106278>

Summala, H. (2007). Towards understanding motivational and emotional factors in driver behaviour: Comfort through satisficing. In P. C. Cacciabue (Ed.), *Modelling driver behaviour in automotive environments: Critical issues in driver interactions with intelligent transport systems* (pp. 189–207). Springer. [https://doi.org/10.1007/978-1-84628-618-6\\_11](https://doi.org/10.1007/978-1-84628-618-6_11)

Sun, G., Zhang, H., Zhong, L., & Li, Q. (2025). Identification of driving behavior in continuous diverging sections of expressway system interchange based on CNN-BiLSTM. *Scientific Reports*, 15(1), 10631. <https://doi.org/10.1038/s41598-025-94000-6>

Tahir, H. A., Alayed, W., Hassan, W. U., & Haider, A. (2024). A novel hybrid XAI solution for autonomous vehicles: Real-time interpretability through LIME–SHAP integration. *Sensors*, 24(21), Article 6776. <https://doi.org/10.3390/s24216776>

Tavakoli, A., & Heydarian, A. (2022). Multimodal driver state modeling through unsupervised learning. *Accident Analysis and Prevention*, 170, Article 106640. <https://doi.org/10.1016/j.aap.2022.106640>

Tejada, A., Legius, M. J. E., Kalose, A. R., Oliveira, P. F. V., van Dam, E., & Hogema, J. H. (2023). Towards a practical methodology for defining competent driving. 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), 569–575. <https://doi.org/10.1109/ITSC57777.2023.10422458>

Thapa, S., & Sarkar, A. (2025). A deep dive into enhancing sharing of naturalistic driving data through face deidentification. *The Visual Computer*, 41(4), 2563–2594. <https://doi.org/10.1007/s00371-024-03552-7>

Toledo, T., Musicant, O., & Lotan, T. (2008). In-vehicle data recorders for monitoring and feedback on drivers' behaviour. *Transportation Research Part C: Emerging Technologies*, 16(3), 320–331. <https://doi.org/10.1016/j.trc.2008.01.001>

Tselentis, D. I., & Papadimitriou, E. (2023). Driver profile and driving pattern recognition for road safety assessment: Main challenges and future directions. *IEEE Open Journal of Intelligent Transportation Systems*, 4, 83–100. <https://doi.org/10.1109/OJITS.2023.3237177>

Van Nes, N., Bärghman, J., Christoph, M., & Van Schagen, I. (2019). The potential of naturalistic driving for in-depth understanding of driver behaviour: UDRIVE results and beyond. *Safety Science*, 119, 11–20. <https://doi.org/10.1016/j.ssci.2018.12.029>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need [Preprint]. arXiv. <https://arxiv.org/abs/1706.03762>

Wang, C., Xie, Y., Huang, H., & Liu, P. (2021). A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling. *Accident Analysis and Prevention*, 157, Article 106157. <https://doi.org/10.1016/j.aap.2021.106157>

Wang, X., Xu, R., Zhang, S., Zhuang, Y., & Wang, Y. (2022). Driver distraction detection based on vehicle dynamics using naturalistic driving data. *Transportation Research Part C: Emerging Technologies*, 136, 103561. <https://doi.org/10.1016/j.trc.2022.103561>

World Health Organization. (2019, January 1). European regional status report on road safety. <https://www.who.int/europe/publications/i/item/9789289054980>

Wu, Y., Kihara, K., Hasegawa, K., Takeda, Y., Sato, T., Akamatsu, M., Kitazaki, S., 2020. Age-related differences in effects of non-driving related tasks on takeover performance in automated driving. *J. Safety Res.* 72, 231–238. <https://doi.org/10.1016/j.jsr.2019.12.019>

Xu, G., Chen, J., Wang, Z., Zhou, A., Schrader, M., Bittle, J., & Shao, Y. (2025). Enhancing traffic safety analysis with digital twin technology: Integrating vehicle dynamics and environmental factors into microscopic traffic simulation [Preprint]. arXiv. <https://arxiv.org/abs/2502.09561>

Yang, J., Zhao, R., Zhu, M., Hallac, D., Sodnik, J., & Leskovec, J. (2021). Driver2vec: Driver identification from automotive data [Preprint]. arXiv. <https://arxiv.org/abs/2103.01822>

Zahraoui, Y., Errajraji, K., Ramah, S., Bouhoute, A., & Berrada, I. (2022). Driver profiling: The pathway to deeper personalization. *Journal of King Saud University - Computer and Information Sciences*, 34(10, Part B), 9088–9101. <https://doi.org/10.1016/j.jksuci.2022.08.034>

Zafar, M. B., Valera, I., Rogriguez, M. G., & Gummadi, K. P. (2017). Fairness constraints: Mechanisms for fair classification. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 962–970. <https://proceedings.mlr.press/v54/zafar17a.html>

Zhang, C., Li, R., Kim, W., Yoon, D., & Patras, P. (2020). Driver behaviour recognition via interwoven deep convolutional neural nets with multi-stream inputs. *IEEE Access*, 8, 191138–191151. <https://doi.org/10.1109/ACCESS.2020.3032344>

Zhang, J., Guo, G., Wang, P., Tang, Q., Zhang, Z., 2019. Research on the Differences of Risk Perception Ability Between Novice and Experienced Drivers, in: Stephanidis, C. (Ed.), *HCI International 2019 – Late Breaking Papers, Lecture Notes in Computer Science*. Springer International Publishing, Cham, pp. 433–447. [https://doi.org/10.1007/978-3-030-30033-3\\_34](https://doi.org/10.1007/978-3-030-30033-3_34)