Universitatea
Transilvania
din Brașov

INTERDISCIPLINARY DOCTORAL SCHOOL

Faculty of Electrical Engineering and Computer Science

Obaida ALHOUSRYA

# IoT and IIoT Technologies Applied in Electrical Vehicle's Industry

SUMMARY

Scientific supervisor

Prof. Dr. Petru A. COTFAS

BRAȘOV, 2026

**Contents**

Abstract ........ 3

Chapter 1 - Introduction ........ 4

1.1 Topic, Field and Relevance of the Thesis ........ 4

1.2 Purpose, Research Gaps and Objectives of the Thesis ........ 6

1.3 Structure of the Thesis ........ 8

1.4 Research Methodology ........ 10

Chapter 2: State of the Art in IoT/IIoT for Electric Vehicles ........ 13

Chapter 3 - IoT/IIoT Architecture & System Implementation for Driver Health Monitoring ........ 15

Chapter 4: Hybrid Metaheuristic Optimization Under Uncertainty ........ 19

Chapter 5 - AI Fatigue/Health Prediction & Integrated Evaluation ........ 23

Chapter 6: Conclusions and Future Work ........ 28

6.1 Original Results, Contributions and Relevance ........ 28

6.2 General Conclusions and Future Work ........ 30

6.3 Valorization of the Research Results ........ 31

References ........ 33

## Abstract

This doctoral thesis addresses the need to treat driver health as an essential component of connected electric vehicle (EV) operation rather than as an isolated driver-monitoring add-on. The work is placed in the doctoral field of Electronics, Telecommunications and Informational Technologies and integrates several connected disciplines: Internet of Things (IoT), Industrial Internet of Things (IIoT), electric vehicles, cloud-edge computing, artificial intelligence, machine learning, multimodal sensing, driver health monitoring, data analytics, and uncertainty-aware optimization. The central idea is that an EV should not be understood only as a battery-driven vehicle, but also as a connected cyber-physical platform in which vehicle telemetry, driver-state information, communication networks, cloud services, and optimization algorithms interact in real time.

The purpose of the thesis is to develop and validate an integrated IoT/IIoT framework that can acquire driver-related signals, transmit them through a structured communication pipeline, process them through cloud-edge services, infer health and fatigue risks using AI models, and support decision-making through optimization under uncertainty. The research responds to the fragmentation observed in existing studies, where driver fatigue detection, EV connectivity, IoT/IIoT architecture, and EV routing or resource optimization are often developed separately. In contrast, the thesis connects these elements into a single engineering-oriented workflow, where driver health is treated as a first-class service inside connected EV ecosystems.

The original contribution of the thesis is developed at three levels. First, it proposes a layered IoT/IIoT architecture that links multimodal sensing, structured telemetry, MQTT/HTTP communication, edge preprocessing, Azure IoT Hub ingestion, Azure Stream Analytics, Cosmos DB storage, Power BI visualization, and alerting. Second, it develops a Hybrid Ant Colony Optimization (HACO) algorithm for shortest-path optimization with mixed fuzzy arc weights, providing an uncertainty-aware routing core. Third, it develops and evaluates multimodal AI pipelines for driver fatigue and health prediction based on ECG/HRV features, eye-state images, and respiration-related features. The results demonstrate near-real-time monitoring, strong HACO convergence behavior, and high predictive performance for ECG/HRV and eye modalities, while also clarifying the limitations of respiration-only detection.

# Chapter 1 - Introduction

## 1.1 Topic, Field and Relevance of the Thesis

The work is positioned at the intersection of connected mobility, embedded sensing, cloud-edge information systems, machine learning and intelligent optimization. The thesis does not treat IoT/IIoT only as a communication layer, and it does not treat fatigue detection only as a classification exercise. Instead, it uses the connected EV as a complete operational environment in which sensing, telemetry, data quality, inference, visualization, alerting and decision support must work together. This positioning is important because a driver-health warning has no practical value if the signal is not acquired correctly, if it is delayed in the network, if the dashboard does not display it clearly, or if the downstream routing and resource-allocation logic cannot use the risk signal.

The thesis also responds to the practical limitation that many advanced driver-monitoring studies remain outside the engineering constraints of deployment. High classification accuracy alone is not sufficient for connected EV applications. A model must be compatible with low-bandwidth telemetry, must tolerate intermittent communication, must support timestamped and traceable messages, and must be able to distinguish low-quality sensor windows from meaningful fatigue evidence. For this reason, the thesis combines algorithmic evaluation with system-level evaluation, so that the proposed results can be interpreted not only as laboratory results but also as building blocks for connected vehicle services.

Road transport is undergoing a major transition due to decarbonization targets, energy-security concerns, the growth of battery technologies, and the increasing digitalization of mobility. Electric vehicles are central to this transition, and recent market reports show continuous expansion in global EV adoption and a rising share of new vehicle sales. At the engineering level, modern EVs are no longer isolated electromechanical products. They integrate battery management systems, power electronics, drivetrain control, advanced driver-assistance systems, telematics, over-the-air updates, remote diagnostics, and cloud-connected services. These elements transform EVs into software-defined and data-driven platforms that continuously generate telemetry and interact with digital infrastructure [1].

The thesis topic, “IoT and IIoT Technologies Applied in Electrical Vehicle's Industry,” integrates directly into the doctoral field of Electronics, Telecommunications and Informational Technologies. Its electronic component is reflected in sensing and data acquisition, its telecommunications component in MQTT/HTTP telemetry and cloud ingestion, and its informational-technologies component in stream processing, storage, dashboards, AI inference, and decision support. At the same time, the research has an interdisciplinary character because it combines connected vehicle systems, cloud computing, embedded sensing, biomedical-signal processing, human-state monitoring, artificial intelligence, fuzzy optimization, and fleet-oriented IoT/IIoT operation, where interoperability, reliability, lifecycle management, and integration with back-end platforms are essential [2,3].

The practical value of IoT/IIoT appears when telemetry is converted into action. A connected EV ecosystem requires data acquisition from heterogeneous sources, reliable transport through communication protocols, edge processing for low-latency responses and bandwidth reduction, cloud services for storage and analytics, and decision mechanisms capable of acting on the acquired information. This is why edge-cloud architectures are increasingly presented as essential for mobility and cyber-physical systems that require both responsiveness and scalability. The thesis adopts this perspective and develops an architecture in which sensing, communication, cloud processing, visualization, inference, and decision support are treated as parts of the same pipeline.

Driver health and fatigue monitoring represent the safety-oriented core of the research. Fatigue, drowsiness, stress, and sudden health degradation can reduce reaction time and situational awareness, especially during long, monotonous, night, or cognitively demanding driving conditions. Existing reviews show that reliable detection benefits from combining multiple indicators, because a single signal can be affected by noise, environmental conditions, individual variability, or sensor failure. Therefore, driver health monitoring should not be limited to one algorithm or one sensor, but should rely on a multimodal approach combining physiological, behavioral, and contextual evidence [2,3].

The relevance of the thesis comes from the fact that connected EV operation depends on both vehicle state and driver state. Monitoring only the vehicle may miss fatigue or health events, while monitoring only the driver without considering vehicle and route constraints limits the ability to make safe and feasible recommendations. For example, a detected fatigue risk may require an alert, a rest recommendation, a change in route, or a decision-support action that considers traffic, energy, charging availability, and uncertainty in travel conditions. This motivates an integrated IoT/IIoT framework that jointly supports safety, monitoring, prediction, and operational decision-making in connected EV environments.

## 1.2 Purpose, Research Gaps and Objectives of the Thesis

The purpose of this thesis is to move driver health monitoring from an isolated detection task toward an operational IoT/IIoT service for connected electric vehicles. Therefore, the research does not stop at fatigue classification; it investigates how driver-state information can be acquired, transmitted, processed, visualized, and connected to decision support under uncertainty.

The main purpose of the thesis is to design and validate an integrated IoT/IIoT-based framework for real-time driver health monitoring in connected electric vehicles and to develop intelligent methods that transform multimodal observations into operational decision support. The research is guided by two main objectives. The first objective is system-level: to design and experimentally validate a cloud-edge IoT/IIoT framework for driver health monitoring using multimodal sensing, structured telemetry, scalable cloud services, dashboards, and alerts. The second objective is methods-level: to develop advanced hybrid metaheuristic and AI models for routing, resource allocation, and early detection of driver health or fatigue issues under uncertainty.

The thesis starts from four main research gaps. The first gap is the lack of end-to-end architectures in which driver health is treated as a first-class IoT/IIoT service. Many fatigue and drowsiness solutions are evaluated as local or offline classifiers, without explicit consideration of ingestion, streaming analytics, edge-cloud partitioning, latency, storage, dashboards, security, or integration with EV services. Conversely, EV-IIoT research often emphasizes connectivity, charging, predictive maintenance, fleet management, and energy optimization, but does not treat driver health monitoring as a central architectural requirement [1,2].

The second gap concerns the limited exploitation of multimodal and contextual data under realistic connected-system constraints. Multimodal sensing is widely recognized as beneficial, but many studies still rely on restricted datasets, single- model assumptions, or offline evaluation. They rarely assess sensor dropouts, latency, bandwidth limitations, heterogeneous sampling rates, and missing data, although these factors strongly affect IoT deployment. The thesis responds by designing data quality flags, synchronized telemetry, model-specific preprocessing, and quality-aware fusion logic to make driver-health inference suitable for operational IoT/IIoT pipelines [2,3].

The third gap is that EV routing and resource optimization are rarely driver-aware, especially under uncertainty. Many EV routing and charging studies consider travel time, energy consumption, station availability, time windows, and fuzzy or stochastic uncertainty. However, driver state is often ignored or represented only indirectly, and few approaches integrate health or fatigue indicators as constraints, penalties, or triggers for routing decisions. This creates an opportunity to connect driver-state prediction with uncertainty-aware optimization, so that the system can respond to risk rather than simply detect it.

The fourth gap is validation fragmentation. Some studies demonstrate hardware prototypes but do not provide flexible algorithmic validation, while others simulate optimization or edge-cloud behavior without considering driver health. The thesis therefore adopts software-centric and simulation-oriented validation, supported by public datasets, reproducible workflows, benchmark comparisons, and representative operational scenarios. This approach enables the architecture, optimization algorithm, and AI models to be evaluated together rather than as separate components.

Based on these gaps, the thesis defines five specific objectives. O1 is to review and analyze IoT/IIoT technologies, AI methods, driver health monitoring, and optimization in connected EV ecosystems. O2 is to design an end-to-end IoT/IIoT architecture for real-time driver health monitoring. O3 is to develop a hybrid metaheuristic method for uncertainty-aware routing or resource decision support. O4 is to develop AI models for early detection and prediction of driver health and fatigue using multimodal data. O5 is to validate the proposed architecture, optimization algorithm, and AI models through reproducible experiments, benchmarking, and integrated case studies.

### 1.3 Structure of the Thesis

The thesis is organized in six chapters, following a progressive logic from research motivation and state-of-the-art analysis to system architecture, optimization, AI-based prediction, integrated validation, conclusions, and valorization. Table 1 summarizes the role of each chapter and its result-oriented contribution.

*Table 1. Thesis chapter structure and result focus.*

| Chapter | Main content | Result-oriented role in the summary |
|---|---|---|
| Chapter 1 | Introduction, motivation, research gaps, objectives and scope | Defines the problem and research questions |
| Chapter 2 | State of the art in IoT/IIoT, EV systems, driver health monitoring and optimization | Identifies the scientific gaps and positions the thesis |
| Chapter 3 | IoT/IIoT architecture and implementation for driver health monitoring | Provides the cloud-edge system, telemetry logic and operational latency results |
| Chapter 4 | Hybrid metaheuristic optimization under uncertainty | Develops and benchmarks HACO for mixed fuzzy shortest paths |
| Chapter 5 | AI fatigue/health prediction and integrated evaluation | Reports ECG/HRV, eye, respiration and fusion results |
| Chapter 6 | Conclusions, limitations, future work and valorization | Synthesizes contributions and practical relevance |

Chapter 1 introduces the background and motivation, defines the research problem and gaps, states the objectives and research questions, presents the methodology, and summarizes the thesis contributions. It positions driver health monitoring as a necessary element of connected EV operation rather than a peripheral feature.

Chapter 2 presents the state of the art in IoT/IIoT for electric vehicles. It reviews IoT and IIoT architectures, connectivity patterns, protocol stacks, edge-cloud computing, digital twins, cybersecurity, lifecycle management, multimodal driver health monitoring, machine-learning and deep-learning models, IoT-based driver monitoring platforms, EV routing and resource optimization, and AI analytics for EV reliability and predictive maintenance. The chapter establishes the conceptual and technical basis for the thesis by showing that existing research provides useful components but not a complete integrated driver-health-aware EV framework.

Chapter 3 develops the proposed IoT/IIoT architecture and system implementation for driver health monitoring. It describes the design rationale, functional requirements, sensing and acquisition layer, communication protocols, message schemas, timing constraints, cloud ingestion, stream processing, storage, visualization, security controls, alerting, and feedback. The chapter is implementation-oriented and shows how multimodal driver data can be converted into structured telemetry and processed through a cloud-edge pipeline based on Azure IoT Hub, Azure Stream Analytics, Cosmos DB, Power BI, MQTT, HTTP, and TLS-protected communication.

Chapter 4 develops the hybrid metaheuristic optimization component. It formulates the shortest-path problem under mixed fuzzy arc weights, presents normal and trapezoidal fuzzy models, uses alpha-cut representation and distance-based ranking, and introduces the HACO method. The proposed algorithm combines Ant Colony Optimization with GA-inspired mutation and a Tabu-like no-repetition mechanism. The chapter compares HACO with GA, PSO, ABC, and standard ACO using simple, complex, and large benchmark graphs, evaluating convergence iterations, convergence time, total runtime, robustness, and statistical significance.

Chapter 5 develops and evaluates the AI fatigue and health prediction layer. It presents datasets, preprocessing, feature engineering, subject-wise splitting, leakage prevention, performance metrics, ECG/HRV-based Random Forest modeling, eye-based CNN modeling, respiration-based Random Forest modeling, cross- model comparison, ablation analysis, safety-oriented error analysis, quality-aware fusion, online inference orchestration, dashboards, and latency-aware operational evaluation. Chapter 5 is the central results chapter because it connects the AI models to the IoT/IIoT architecture and translates predictions into risk signals suitable for alerts and decision support.

Chapter 6 summarizes the conclusions, achievement of objectives, principal contributions, limitations, future work, final remarks, and valorization of the research results. It confirms that the thesis achieves its system-level, optimization-level, AI-level, and integrated-validation objectives. It also clarifies the boundaries of the work, including software-centric validation, use of prototype sensors as signal proxies, public-dataset heterogeneity, and the need for future real-vehicle and fleet-scale experiments.

## 1.4 Research Methodology

The methodology combines five connected phases: analytical review, architecture design, uncertainty-aware optimization, AI modelling, and integrated validation. This structure was necessary because the thesis problem cannot be solved by one method alone: it requires both system-level implementation and algorithmic evaluation.

*Table 2. Research methodology phases and outputs.*

| Phase | Main actions | Outputs |
|---|---|---|
| Analytical review | Systematic mapping and SLR-style analysis of IoT/IIoT, driver monitoring and optimization | Gaps, requirements and thesis objectives |
| Architecture design | Layered sensing, communication, cloud, visualization and alerting design | End-to-end IoT/IIoT framework and telemetry contract |
| Optimization development | Mixed fuzzy shortest-path formulation and HACO design | Uncertainty-aware routing core |
| AI modelling | ECG/HRV RF, eye CNN, respiration RF and multimodal fusion | Fatigue risk inference models |
| Integrated validation | Latency, responsiveness, convergence, runtime, accuracy, F1 and ablation analysis | Evidence for feasibility, performance and limitations |

The research methodology follows a structured progression with five major phases: analytical, architectural, algorithmic optimization, AI modeling, and experimental validation. The analytical phase consists of a systematic state-of-the-art review and gap analysis. It investigates IoT/IIoT systems used in EV ecosystems, driver health and fatigue monitoring, edge-cloud architectures, multimodal sensing, and optimization under uncertainty. The review uses database searches, inclusion and exclusion criteria, thematic classification, data extraction, and quality appraisal. This phase establishes the evidence base and ensures that the proposed system responds to real limitations in the literature [4-6].

The architectural phase translates the identified requirements into a layered IoT/IIoT framework. The methodology defines the functional role of each layer: sensing and acquisition, edge preprocessing, communication, cloud ingestion, stream processing, persistence, visualization, alerting, and decision support. The architecture is not treated as a high-level diagram only; it is mapped to implementable cloud services and communication mechanisms. MQTT is selected as the primary telemetry protocol due to its publish-subscribe pattern and suitability for lightweight continuous messaging, while HTTP is maintained as a secondary route for compatibility and request-response operations. Timing, message structure, payload fields, quality flags, and buffering are considered as design constraints rather than afterthoughts [7-12].

The optimization phase develops an uncertainty-aware routing core. The problem is formulated as a graph-based shortest-path model in which arc weights are represented as mixed fuzzy numbers. This is important because EV travel and decision costs can be uncertain due to traffic, energy use, charging availability, environmental conditions, and potentially driver-state risk. The thesis uses alpha-cut processing and fuzzy cost ranking to compare candidate paths. The HACO method is then designed to improve search performance by combining ACO path construction with mutation and no-repetition logic. The method is benchmarked against GA, PSO, ABC, and ACO under common settings and repeated runs [13-16].

The AI modeling phase develops pipelines for multimodal driver health and fatigue inference. For ECG/HRV, the methodology includes window-level signal processing, R-peak detection, HRV-related feature extraction, heuristic labeling where ground truth is limited, Random Forest classification, subject-wise evaluation, and interpretation of feature importance. For the eye state, the methodology includes dataset preparation, image preprocessing, augmentation, lightweight CNN design, training, testing, and confusion-matrix analysis. For respiration, the methodology includes observation definition, preprocessing, feature selection, Random Forest modeling, grouped cross-validation, and safety-oriented evaluation. The models are evaluated not only by accuracy but also by precision, recall, F1-score, false-positive rate, false-negative rate, balanced accuracy, and cross-validation stability [17-25].

A central methodological decision is subject-wise splitting and leakage prevention. In driver monitoring, ordinary random splitting can overestimate performance when samples from the same driver appear in both training and testing. The thesis therefore emphasizes splitting strategies that reduce leakage and better reflect generalization to new drivers. This decision is especially important for physiological signals, where individual baseline differences can dominate the prediction task. By combining leakage-resistant evaluation with cross-dataset comparisons, the thesis provides a more realistic interpretation of model performance and deployment risk [17-19,25].

The experimental validation phase links the system, optimization, and AI metrics. At the system level, the thesis measures upload latency, alert responsiveness, detection accuracy, and robustness under representative scenarios such as calm driving, heavy traffic, night driving, and emergency braking. At the optimization level, it measures convergence iterations, convergence time, runtime, robustness, and statistical significance. At the AI level, it measures test performance, cross-validation stability, and safety-oriented errors. This multi-level validation is important because a model with high offline accuracy is insufficient if it cannot operate within the timing, data-quality, and communication constraints of a connected vehicle environment [12,16,25].

## Chapter 2: State of the Art in IoT/IIoT for Electric Vehicles

The state-of-the-art review is synthesized around three axes: IoT/IIoT architectures for EV ecosystems, multimodal driver health monitoring, and optimization under uncertainty. The review shows that these areas are individually mature but insufficiently connected, which justifies the integrated framework developed in the thesis [1-3,13-16].

The state-of-the-art review shows that IoT and IIoT technologies are increasingly important in EV ecosystems. IoT platforms support telemetry collection from battery systems, drivetrain components, sensors, charging behavior, location data, and driver-related variables. IIoT adds requirements linked to reliability, scalability, lifecycle management, interoperability, fleet supervision, predictive maintenance, and integration with energy infrastructure. However, the review also shows that EV-IIoT architectures often focus on vehicle-side and infrastructure-side variables while treating human state as secondary, even though the driver remains a critical safety component in many EV operating scenarios [1].

Connectivity and protocol selection are major themes in the reviewed literature. EV systems may use in-vehicle communication, vehicle-to-everything communication, cloud APIs, IoT messaging, and service-to-service integration. For continuous telemetry, publish-subscribe architectures such as MQTT are useful because they allow lightweight data flows between vehicle nodes, gateways, and cloud services. For compatibility, configuration, and discrete service calls, HTTP-based communication remains relevant. The thesis positions MQTT and HTTP as complementary engineering options rather than competing technologies, with MQTT used as the primary route for streaming driver-health telemetry and HTTP used as a secondary interface [7].

Edge-cloud computing is another key area. Edge processing can reduce latency, protect privacy, and reduce bandwidth, while cloud processing offers scalable storage, dashboards, model updating, fleet analysis, and integration with external services. In driver health monitoring, this split is important because some alerts may need low-latency local response, while longer-term analytics and fleet monitoring require cloud resources. The thesis uses this principle to define a system in which local acquisition and preliminary logic coexist with cloud ingestion, stream analytics, database persistence, visualization, and remote alerting [8,9]. The same connected architecture also requires cybersecurity governance, secure device identity, encrypted data movement and controlled access when driver-health information is processed across vehicle, edge and cloud services [10,11].

The review of driver health and fatigue monitoring shows that no single sensing model is sufficient in all conditions. Eye-based systems can capture blinking, eyelid closure, gaze, and visual signs of drowsiness, but may fail under poor illumination, occlusion, eyewear, or camera placement issues. Physiological signals such as ECG/HRV can provide deeper evidence of fatigue or cognitive state, but may be affected by motion artifacts, electrode/contact quality, and individual variability. Respiration can reflect stress and fatigue-related changes, but it often overlaps strongly between classes and may be unreliable as a standalone detector. These observations motivate the thesis focus on multimodal evidence and quality-aware fusion [2,3,17,22].

The review of machine-learning and deep-learning approaches shows a diverse landscape of Random Forest, SVM, kNN, CNN, recurrent models, transfer learning, and multimodal fusion strategies. Many works report high accuracy, but comparability is limited by different datasets, labels, evaluation splits, metrics, and controlled conditions. The thesis responds by reporting multiple metrics, using subject-wise evaluation where possible, comparing modalities, and interpreting safety-oriented indicators such as false-negative rate. This is important because missed fatigue events can be more dangerous than false alarms in safety-critical monitoring [3,23,24].

The review of optimization in EV systems shows that EV routing and resource allocation are commonly influenced by travel time, energy consumption, charging station availability, battery state, traffic, time windows, and uncertainty. Fuzzy formulations and metaheuristics are useful when exact deterministic formulations are insufficient or computationally expensive. The thesis builds on this background but adds an important perspective: routing and resource decisions should be capable of receiving driver health risk as an input. The HACO optimization core is therefore not only a benchmark exercise; it is positioned as a reusable service that can later accept risk penalties or constraints from the driver-monitoring layer [13-16].

## Chapter 3 - IoT/IIoT Architecture & System Implementation for Driver Health Monitoring

The proposed architecture is one of the central original results of the thesis. It transforms multimodal driver-health data into structured telemetry that can be processed through local logic, MQTT/HTTP communication, Azure cloud services, dashboards, and alerts. Figures 1-4 and Table 3 summarize the workflow, cloud implementation, operational response, and telemetry upload latency of the prototype [12,25].

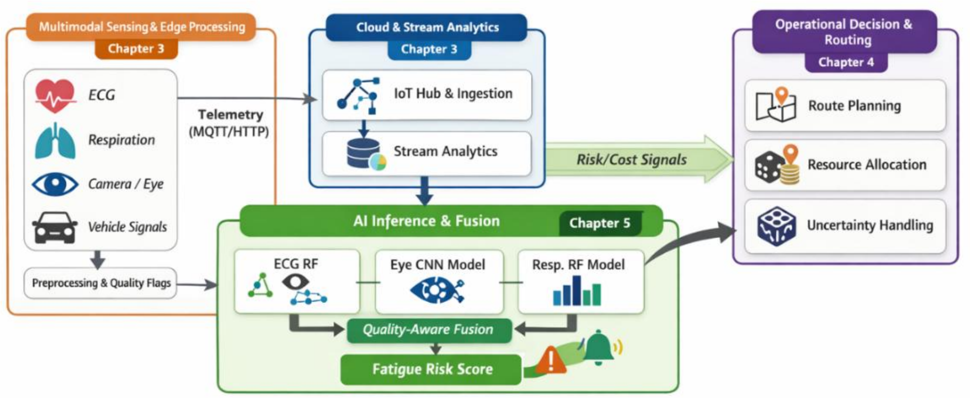


*Figure 1. Integrated thesis workflow linking sensing, cloud analytics, AI fusion and decision support.*

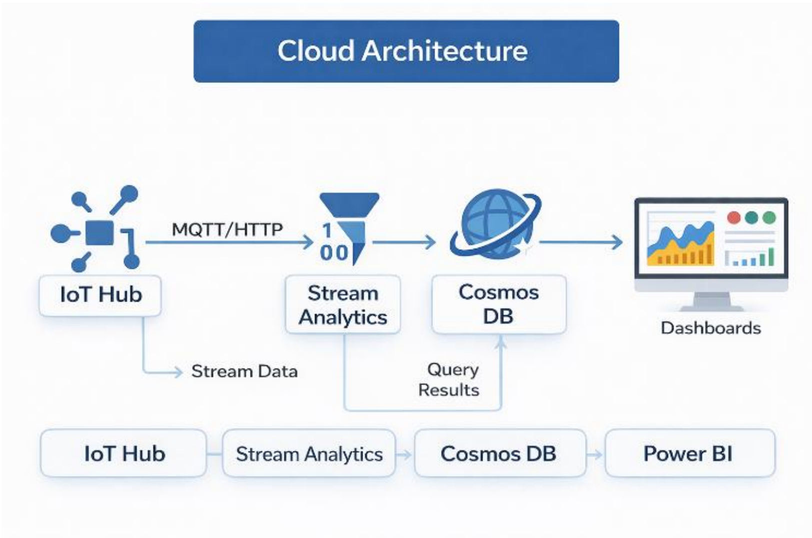


*Figure 2. Cloud architecture for IoT telemetry and dashboards.*

*Table 3. Operational evaluation of the IoT/cloud monitoring prototype.*

| Scenario | Upload latency (ms) | Detection accuracy (%) | Alert response (ms) | Observation |
|---|---|---|---|---|
| Calm driving | 180 | 95 | 250 | Stable performance |
| Heavy traffic | 250 | 92 | 300 | Slight delays under load |
| Night driving | 200 | 93 | 270 | Consistent detection |
| Emergency braking | 150 | 94 | 220 | Fast response to critical event |

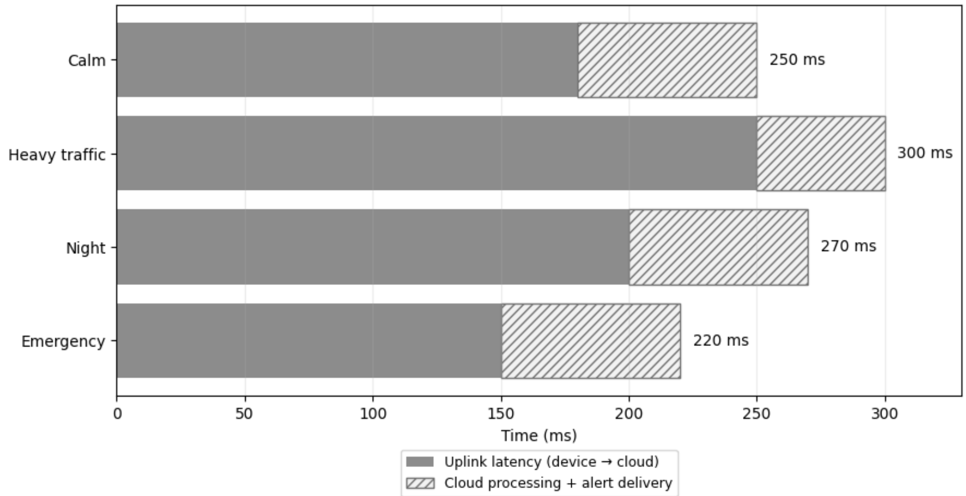


*Figure 3. End-to-end alert response across representative driving scenarios.*

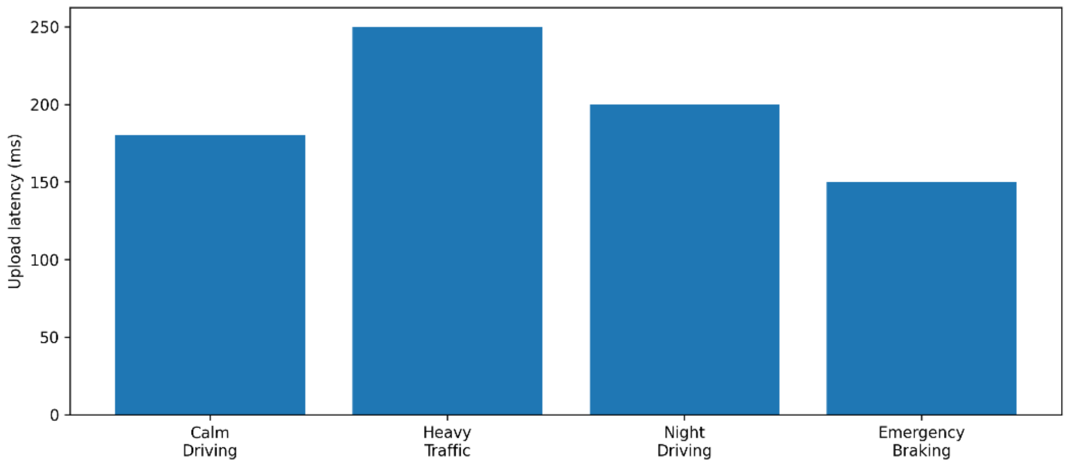


*Figure 4. Telemetry upload latency under representative driving scenarios.*

The telemetry upload latency remains within 150-250 ms, with the highest delay occurring under heavy traffic and the lowest delay during emergency braking, supporting the near-real-time feasibility of the IoT/cloud pipeline.

The system-level contribution of the thesis is an end-to-end IoT/IIoT architecture for driver health monitoring in connected EVs. The architecture is layered, modular, and deployment-oriented. Its main layers are multimodal sensing and acquisition, local preprocessing, edge/gateway communication, cloud ingestion, stream processing, persistent storage, visualization, alerting, and closed-loop feedback. The architecture treats driver health monitoring as an EV service that must satisfy engineering requirements such as interoperability, traceability, timing control, data quality, security, and scalability [12].

At the sensing layer, the thesis considers multimodal evidence related to heart activity, eye behavior, and respiration. In the prototype architecture, low-cost components are used as signal proxies to validate data acquisition, packet formatting, communication, and cloud integration. This is an important limitation and design choice: the prototype sensors are not presented as clinical instruments, but as engineering proxies for demonstrating the complete data path. In future deployment, these would be replaced or complemented by production-grade ECG, camera, respiration, and vehicle telemetry sensors. The value of the prototype is that it proves the architecture and data handling logic before expensive real-vehicle validation [12].

The acquisition layer defines a reproducible packet format and recommends feature update cadences for different modalities. The system uses timestamps, device identifiers, sequence numbers, status labels, and quality flags. These fields support traceability and synchronization, especially when data may be delayed, missing, duplicated, or affected by sensor noise. The thesis emphasizes that driver-health inference in connected vehicles requires more than raw measurements; it requires structured telemetry that allows the system to know when data were collected, where they came from, whether they are valid, and how they should be aligned with other modalities [12].

The communication design uses MQTT as the primary protocol and HTTP as the secondary protocol. MQTT is appropriate for continuous telemetry because it supports lightweight publish-subscribe data exchange and topic-based routing. HTTP is useful for compatibility, testing, configuration, and REST-style communication. The thesis defines topic schemes, payload fields, gateway variants, and timing constraints. It also considers latency, packet loss, jitter, buffering, and paced sending after outages. This approach reflects real IoT/IIoT constraints: a system that works only under perfect connectivity is not sufficient for connected mobility [7].

The cloud implementation uses Azure IoT Hub as the telemetry gateway, Azure Stream Analytics for parsing, windowing, and routing, Azure Cosmos DB for telemetry and event persistence, and Power BI for visualization. This service decomposition gives the architecture a practical cloud realization. IoT Hub handles device-to-cloud ingestion and device identity. Stream Analytics processes events and supports windowed logic. Cosmos DB stores telemetry and alerts in a scalable document model. Power BI provides dashboards for monitoring driver state, sensor quality, scenario behavior, and event status [12].

The system demonstrated near-real-time operational behavior across representative scenarios. Upload latency ranged from 150 to 250 ms, with calm driving at 180 ms, heavy traffic at 250 ms, night driving at 200 ms, and emergency braking at 150 ms. Alert responsiveness ranged from 220 to 300 ms, with the heaviest delay under heavy traffic and the fastest response under emergency braking. Detection or anomaly accuracy in the scenario table ranged from 92% to 95%. These results indicate that the architecture is suitable for low-bandwidth, cloud-integrated monitoring and can generate timely warnings under representative conditions [12].

## Chapter 4: Hybrid Metaheuristic Optimization Under Uncertainty

Chapter 4 provides the uncertainty-aware optimization component of the thesis. It formulates routing as a shortest-path problem with mixed fuzzy arc weights and proposes HACO as a fast hybrid metaheuristic solution. Fig. 5, 6, and 7 and Table 4 summarize the algorithmic logic, convergence behavior, and the main convergence/runtime results [13-16].

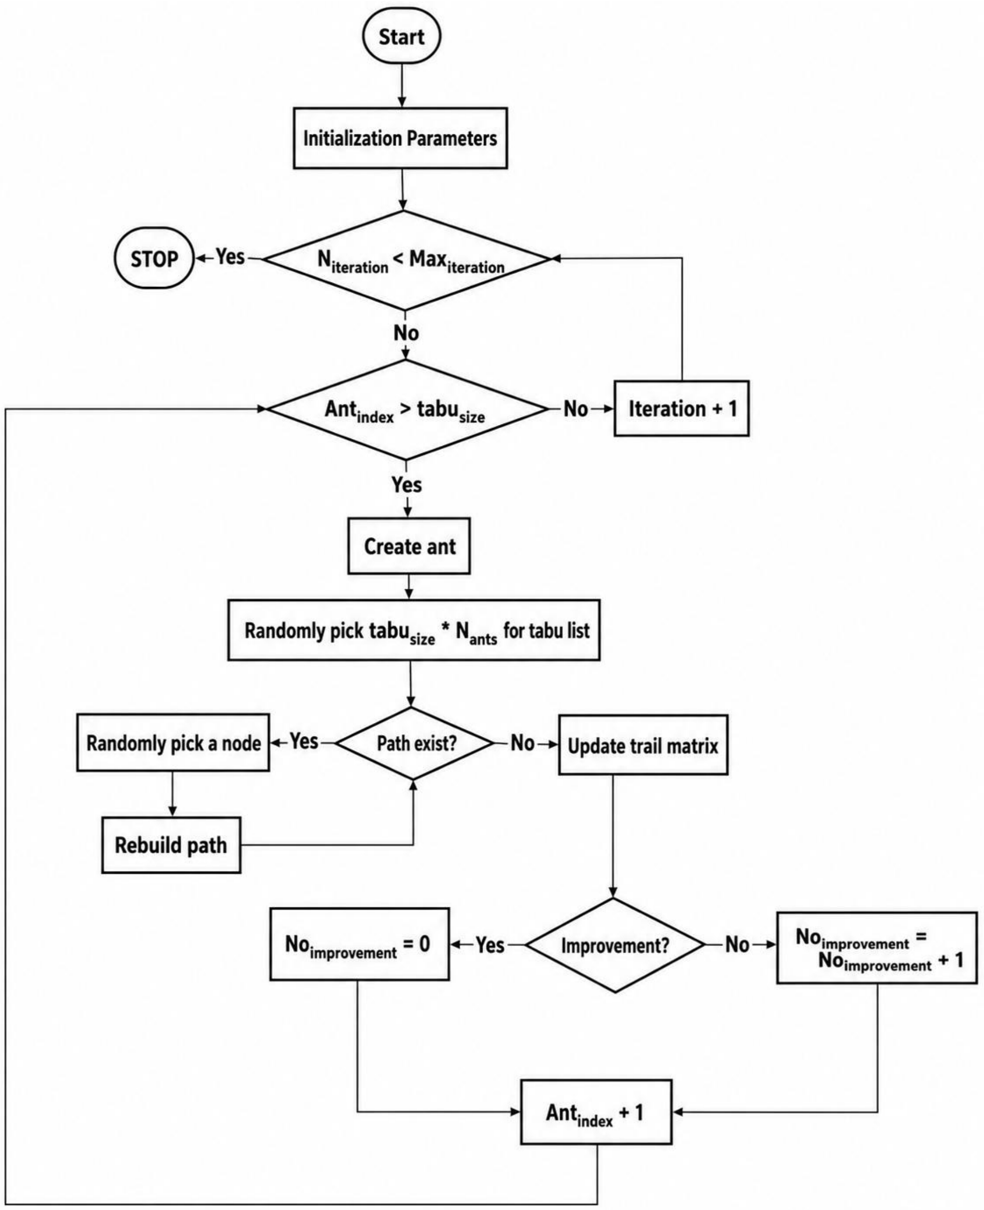


*Figure 5. HACO algorithm flowchart.*

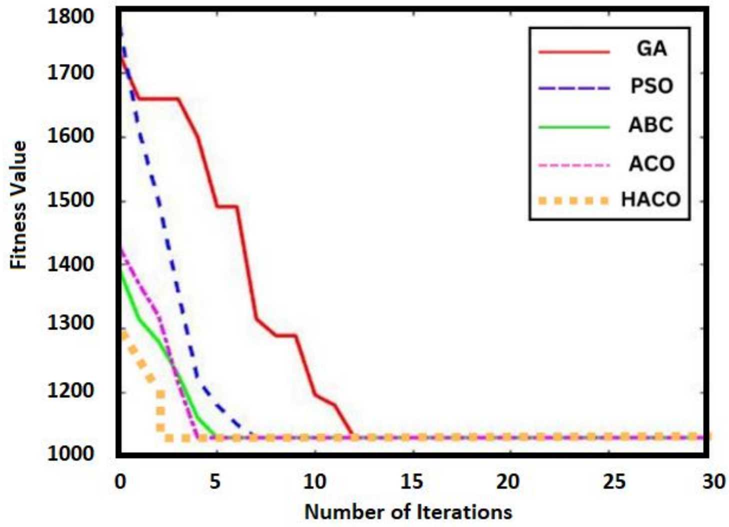


*Figure 6. Convergence behavior of HACO and benchmark algorithms for the simple graph.*

The convergence graph illustrates the rapid stabilization of HACO compared with GA, PSO, ABC and standard ACO, visually supporting the numerical reductions in convergence time and runtime reported in Table 4.

*Table 4. HACO convergence and runtime results across benchmark graphs.*

| Graph | Conv. iterations | Time to convergence (s) | Runtime (s) | Runtime reduction vs. ACO |
|---|---|---|---|---|
| Simple | 1.8 | 0.013 | 0.074 | 91.40% |
| Complex | 2.0 | 0.031 | 0.17 | 95.11% |
| Large | 7.0 | 0.083 | 0.18 | 96.73% |

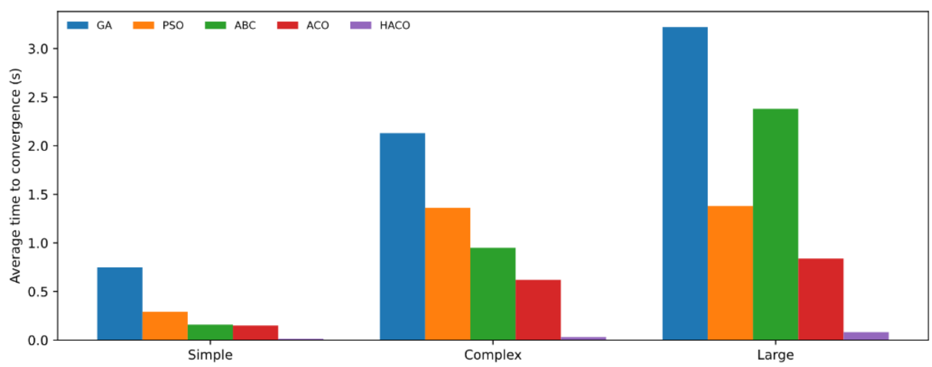


*Figure 7. Average time to convergence across benchmark graph sizes.*

The optimization part of the thesis addresses the need for decision support under uncertain EV operating conditions. Travel time, energy use, traffic congestion, charging availability, and driver-state risk can all be uncertain. To model this type of uncertainty, Chapter 4 formulates the shortest-path problem using mixed fuzzy arc weights. The graph model represents nodes and arcs, while arc weights may be expressed through normal and trapezoidal fuzzy numbers. Alpha-cut representation and aggregation allow fuzzy path costs to be computed, and distance-based ranking enables candidate paths to be compared [13,14,16].

The proposed solution is the Hybrid Ant Colony Optimization algorithm, HACO. Standard ACO is effective for graph problems because it uses pheromone-guided probabilistic construction of candidate paths. However, standard ACO can still suffer from premature convergence, repeated paths, and limited diversification. HACO improves the search by combining ACO path construction with a GA-inspired mutation operator and a Tabu-like no-repetition criterion. The mutation mechanism supports exploration by modifying repeated or stagnating paths, while the no-repetition rule discourages ineffective revisiting of the same solution patterns. Pheromone updating then reinforces good candidate paths, while pheromone evaporation reduces the influence of older or less effective paths and helps prevent premature convergence to suboptimal solutions [15,16].

The experimental design compares HACO with Genetic Algorithm, Particle Swarm Optimization, Artificial Bee Colony, and standard ACO. The algorithms are evaluated on three benchmark graphs: a simple graph with 11 nodes and 25 edges, a complex graph with 23 nodes and 40 edges, and a large graph with 30 nodes and 71 edges. For all algorithms, the benchmarking protocol uses common encoding, objective evaluation, parameter settings, population sizes, iteration budgets, repeated runs, and the same metrics: convergence iterations, convergence time, and total runtime [16].

The results show that HACO achieved the best convergence behavior and runtime across all benchmark sizes. In the simple graph, HACO reached an average of 1.8 convergence iterations, 0.013 s convergence time, and 0.074 s runtime. In the complex graph, it reached 2.0 iterations, 0.031 s convergence time, and 0.17 s runtime. In the large graph, it reached 7.0 iterations, 0.083 s convergence time, and 0.18 s runtime. These results are substantially better than the baseline algorithms, especially in time-to-convergence and total runtime [16].

Compared with standard ACO, HACO reduced convergence iterations by 40.0% in the simple graph, 20.0% in the complex graph, and 23.6% in the large graph. It reduced time to convergence by 91.33%, 95.00%, and 90.24% in the three graphs, respectively, and reduced runtime by 91.40%, 95.11%, and 96.73%. These improvements show that the hybridization strategy provides a practical benefit, not only a theoretical modification. The mutation and no-repetition mechanisms help HACO maintain exploration while preserving the path-construction strengths of ACO [16].

Robustness analysis also supports the conclusion. Across the benchmark outcomes, HACO showed low standard deviation for convergence iterations, time to convergence, and runtime. Statistical testing confirmed that the differences among algorithms were significant. The Kruskal-Wallis test returned $H = 15.67$ with $p < 0.01$, and Dunn post-hoc comparisons showed significant pairwise differences between HACO and GA, PSO, ABC, and ACO for both time and iterations. This confirms that HACO's superiority in the reported benchmark is statistically meaningful [16].

The relevance of HACO within the thesis is broader than solving an abstract shortest-path problem. It establishes an uncertainty-aware routing core that can be connected to driver health monitoring. Once the AI layer estimates fatigue or health risk, the optimization layer can use that risk as an additional penalty, constraint, or rerouting trigger. This creates a closed-loop relation between sensing and action: the vehicle or fleet system does not simply detect a driver condition, but can use it to support safer decisions under uncertain traffic and operational conditions [16,25].

## Chapter 5 - AI Fatigue/Health Prediction & Integrated Evaluation

Chapter 5 develops and evaluates the AI inference layer of the thesis using ECG/HRV, eye-based, and respiration-based fatigue indicators. The aim is to detect or predict driver fatigue and health-related risk from multimodal data, determine which modalities are reliable enough for deployment, and operationalize their outputs inside the IoT/IIoT monitoring pipeline. The evaluation is deployment-oriented, using windowed observations, subject-wise splitting where possible, quality-aware reasoning, confusion matrices, and multiple metrics that reflect both predictive accuracy and safety relevance [17-25].

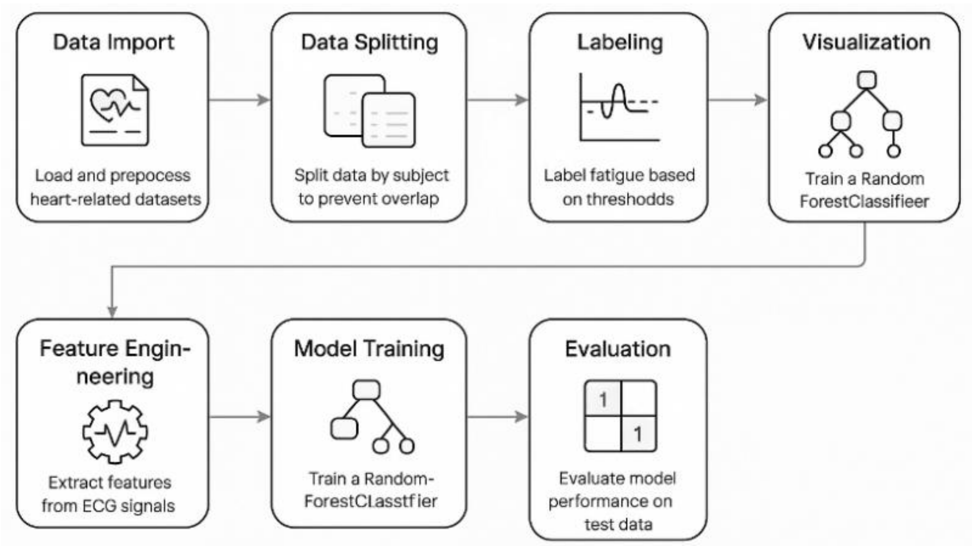


*Figure 8. ECG/HRV fatigue-detection workflow adapted from the thesis workflow.*

*Table 5. Multimodal AI performance and safety interpretation.*

| Model | Dataset/setting | Accuracy | F1-score | Interpretation |
|---|---|---|---|---|
| ECG/HRV (RF) | OpenDriver | 0.964 | 0.958 | Strong primary physiological channel |
| ECG/HRV (RF) | MAUS | 0.8095 | 0.826 | Useful but context-sensitive |
| Eye (CNN) | YawDD | 0.93 | 0.93 | Strong primary behavioral channel |
| Eye (CNN) | TinyML Eye | 0.92 | 0.92 | Stable and balanced performance |
| Respiration (RF) | Respiration dataset | 0.46 | 0.43-0.51 | Auxiliary only; high false-negative risk |

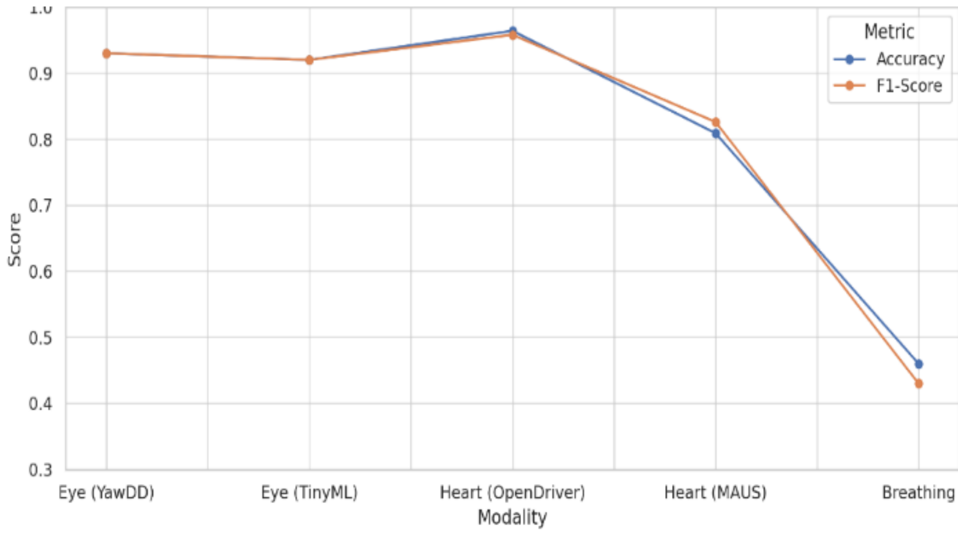


*Figure 9. Accuracy and F1-score comparison across sensing modalities.*

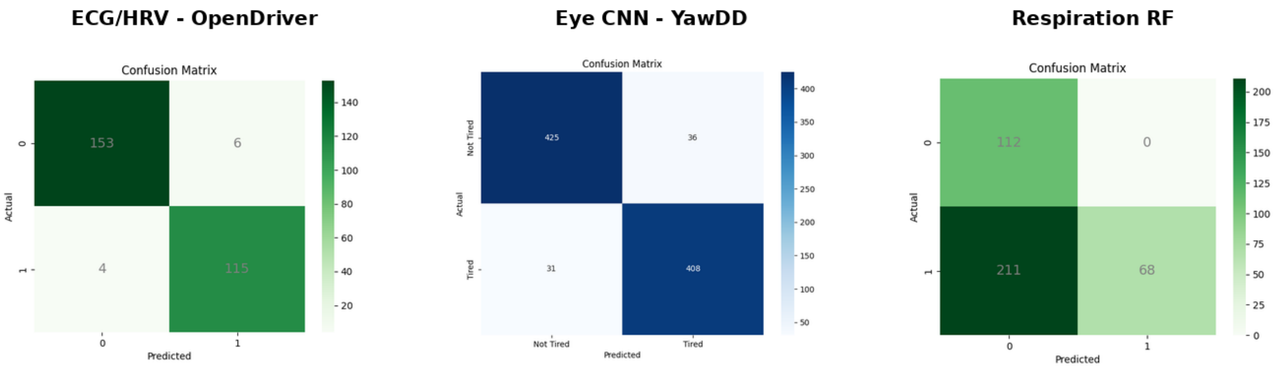


*Figure 10. Representative confusion matrices for ECG/HRV, eye-based and respiration-based fatigue detection.*

The confusion matrices make the safety interpretation clearer: ECG/HRV and eye-based models show balanced classification behavior, while the respiration-only model misses many fatigue cases, confirming why respiration is treated as auxiliary evidence rather than as a primary detector.

The ECG/HRV pipeline uses physiological features extracted from electrocardiogram or heart-rate variability data. The workflow includes preprocessing, R-peak detection, feature engineering, heuristic fatigue labeling under missing ground truth, Random Forest classification, and cross-dataset analysis. ECG/HRV is treated as a primary physiological channel because it can capture changes in autonomic regulation and fatigue-related variability. At the same time, the thesis recognizes that ECG/HRV is sensitive to contact quality, motion artifacts, and individual baseline differences. This is why subject-wise evaluation and quality flags are important [17,23,25].

The strongest ECG/HRV result was obtained on the OpenDriver dataset. The Random Forest model achieved 97.6% mean cross-validation accuracy, 96.4% test accuracy, and 95.8% F1-score. These results indicate that ECG/HRV features provide strong separability in the evaluated naturalistic driving setting. On MAUS, the same model produced lower but still meaningful performance, with 80.95% test accuracy and 82.61% F1-score. This reduction is important because it shows context sensitivity: fatigue type, task conditions, labels, and dataset construction strongly affect model performance. On DriveDB, subject-wise performance showed high average results for selected subjects, but the thesis treats the small number of subjects carefully and avoids overstating generalization [18,19,23,25].

The eye-based modeling pipeline uses a lightweight convolutional neural network to detect fatigue-related visual patterns from eye-state images. Eye-based evidence is operationally attractive because cameras are increasingly available in modern vehicles and can capture behavioral indicators such as eye closure and drowsiness. The thesis evaluates the model on YawDD and TinyML Eye datasets. The CNN achieved 92.56% accuracy, 91.89% precision, 92.94% recall, and 92.41% F1-score on YawDD. On TinyML Eye, it achieved 92.18% accuracy, 92.36% precision, 91.94% recall, and 92.15% F1-score. These values show stable and balanced performance across two eye-related datasets [20,21,24,25].

The derived error metrics strengthen this interpretation. For YawDD, specificity was 0.922, negative predictive value was 0.932, balanced accuracy was 0.926, Matthews correlation coefficient was 0.851, false-positive rate was 0.078, and false-negative rate was 0.071. For TinyML Eye, specificity was 0.924, negative predictive value was 0.920, balanced accuracy was 0.922, MCC was 0.844, FPR was 0.076, and FNR was 0.081. These results show that the eye model is not only accurate but also balanced between false alarms and missed fatigue events [20,21,24,25].

The respiration-based model produced the weakest standalone results. The grouped cross-validation mean accuracy was 0.708 with standard deviation 0.187, indicating unstable generalization across folds. On the held-out test setting, the Random Forest model achieved only 46% accuracy. For the fatigue class, precision was 1.00 but recall was only 0.24, with F1-score 0.39. The confusion-matrix-derived indicators showed specificity of 1.000, sensitivity of 0.244, false-negative rate of 0.756, negative predictive value of 0.347, balanced accuracy of 0.622, and MCC of 0.291. This means that the model frequently missed true fatigue events, which is unacceptable if used as a primary safety detector [22,23,25].

The cross-model comparison gives one of the most important results of the thesis. ECG/HRV and eye modalities are both strong primary channels, but they provide different types of evidence. ECG/HRV offers interpretable physiological features such as RR variability and statistical descriptors, while eye CNNs provide behavioral evidence from learned visual patterns. Respiration has medium interpretability but weak standalone separability in the evaluated dataset. Therefore, the thesis recommends a quality-aware fusion strategy in which ECG and eye are the main channels, while respiration is used only as auxiliary evidence when signal quality and confidence are sufficient [3,17,20-22,25].

The ablation analysis confirms this recommendation. ECG alone reached 0.964 accuracy and 0.958 F1-score. Eye alone reached 0.930 accuracy and 0.930 F1-score. Respiration alone produced only 0.460 accuracy and class-dependent F1-score around 0.430-0.510. Combining ECG and eye improved performance to 0.978 accuracy and 0.969 F1-score. Adding respiration produced only a marginal gain, reaching 0.981 accuracy and 0.972 F1-score. The result is scientifically useful because it does not simply report that "more sensors are better"; it shows that only high-quality and complementary sensors should dominate the fusion decision [25].

The safety-oriented error analysis is also central. ECG-OpenDriver had approximately 0.038 false-positive rate, 0.034 false-negative rate, and 0.964 balanced accuracy. ECG-MAUS had approximately 0.167 FPR, 0.208 FNR, and 0.812 balanced accuracy. Eye-YawDD had approximately 0.078 FPR, 0.071 FNR, and 0.926 balanced accuracy. Eye-TinyML had approximately 0.076 FPR, 0.081 FNR, and 0.922 balanced accuracy. Respiration had 0.000 FPR but approximately 0.756 FNR and 0.622 balanced accuracy. Since false negatives represent missed fatigue events, these results clearly justify why respiration should not be used as the dominant detector in the current setup [25].

The final part of Chapter 5 integrates the AI results into the real-time IoT/cloud system. Model outputs are treated as events that can be ingested, stored, visualized, and used for alerts. The operational evaluation uses calm driving, heavy traffic, night driving, and emergency braking scenarios. Upload latency ranged from 150 to 250 ms, detection accuracy ranged from 92% to 95%, and alert response ranged from 220 to 300 ms. These values confirm that the AI layer is compatible with the monitoring pipeline developed in Chapter 3. The thesis therefore validates not only offline classification, but also end-to-end operational feasibility [12,25].

# Chapter 6: Conclusions and Future Work

## 6.1 Original Results, Contributions and Relevance

The original contributions of the thesis can be grouped into four main results: the IoT/IIoT driver-health architecture, the HACO optimization method, the multimodal AI evaluation and fusion rule, and the integrated validation workflow. Table 6 summarizes these contributions before they are discussed in detail [12,16,25].

*Table 6. Original contributions and thesis-level relevance.*

| Contribution | Original result | Relevance |
|---|---|---|
| IoT/IIoT architecture | Layered cloud-edge driver-health monitoring framework with telemetry contract and dashboards | Transforms driver health into a connected EV service |
| HACO optimization | Hybrid ACO with GA-inspired mutation and Tabu-like no-repetition for mixed fuzzy shortest paths | Enables fast uncertainty-aware routing support |
| Multimodal AI models | ECG/HRV RF, eye CNN, respiration RF and quality-aware fusion | Provides practical fatigue-risk inference with model-specific interpretation |
| Integrated validation | Latency, alert response, convergence, runtime, accuracy, F1 and ablation evidence | Links system feasibility with algorithmic performance |

The first original result is the development of a deployable IoT/IIoT architecture in which driver health monitoring is treated as a core EV service. The architecture defines the complete flow from sensing to action: multimodal acquisition, local preprocessing, structured telemetry, MQTT/HTTP communication, cloud ingestion, stream processing, storage, visualization, alerts, and feedback. It also specifies message schemas, timing constraints, device identity, transport security, data quality flags, and cloud-service decomposition. This gives the thesis a strong engineering contribution, because it moves beyond a conceptual block diagram and presents an implementable architecture [12].

The second original result is the HACO algorithm for mixed fuzzy shortest-path optimization. HACO combines ACO search, GA-inspired mutation, and a Tabu-like no-repetition mechanism to improve exploration and avoid premature convergence. Its results on simple, complex, and large benchmark graphs show major reductions in convergence time and runtime compared with standard ACO and other metaheuristics. The algorithm therefore contributes to uncertainty-aware EV decision support and can act as a routing primitive for future driver-health-aware optimization [16].

The third original result is the multimodal AI evaluation for driver fatigue and health prediction. The thesis provides a comparative analysis of ECG/HRV, eye, and respiration modalities using multiple datasets and metrics. It shows that ECG/HRV and eye evidence form the preferred primary combination, while respiration should be treated as auxiliary in the current setup. This conclusion is important because it is based on both predictive performance and safety-oriented error analysis. The thesis does not overstate the weakest model; instead, it translates empirical limitations into a practical fusion rule [17-25].

The fourth original result is integrated validation. The thesis connects architecture, optimization, and AI models into a single workflow. The cloud-edge architecture provides the operational channel, HACO provides the optimization core, and the AI models provide risk signals. This integration supports a closed-loop concept in which driver-state observations can influence alerts, dashboards, and routing or decision constraints. The resulting contribution is not only a collection of methods, but a framework for health-aware connected EV operation [12,16,25].

The thesis contribution to the field is therefore both system-level and method-level. At the system level, it offers an architecture for real-time driver health monitoring in connected EV ecosystems. At the method level, it offers a high-performance hybrid metaheuristic and validated AI models for fatigue inference. At the integration level, it demonstrates how risk predictions can become actionable events inside a cloud-edge IoT/IIoT pipeline. This is relevant for EV safety, fleet monitoring, intelligent transportation systems, cloud-connected mobility services, and future vehicle platforms that combine human-state awareness with operational decision-making [1,12,16,25].

The research also contributes by clarifying limitations. The prototype sensors are treated as signal proxies, not medical-grade instruments. The AI results are based on public datasets with heterogeneous conditions and, in some cases, proxy labels. Respiration-only modeling is shown to be weak and false-negative-heavy. Security and privacy are addressed architecturally through encryption, identity, and controlled access, but not through new cryptographic methods or full compliance workflows. These limitations are important because they make the thesis conclusions realistic and define the next steps needed before full in-vehicle or fleet deployment [10-12,17-25].

### 6.2 General Conclusions and Future Work

The conclusions synthesize the thesis at three levels: system implementation, optimization performance, and AI-based fatigue prediction. The following paragraphs explain how the objectives were achieved and what future work is required before full in-vehicle or fleet deployment.

The main conclusion of the thesis is that driver health can be embedded as a first-class IoT/IIoT service in connected electric vehicle ecosystems. The proposed framework links sensing, communication, cloud processing, AI inference, visualization, alerting, and uncertainty-aware optimization. It shows that driver monitoring should not be an isolated classifier but part of a broader operational loop in which detected or predicted risk can support decisions. This conclusion directly addresses the main problem statement and confirms the value of integrating vehicle connectivity, driver-state sensing, cloud-edge analytics, and optimization [1,12,16,25].

The system-level objective was achieved through the implementation and validation of the IoT/IIoT monitoring pipeline. The architecture showed near-real-time behavior, with upload latency between 150 and 250 ms and alert responsiveness between 220 and 300 ms. These values indicate practical feasibility for cloud-integrated monitoring, especially when combined with local logic for critical events. The thesis also demonstrated that structured telemetry, quality flags, and service decomposition are necessary for reliable deployment [12].

The optimization objective was achieved through the development and benchmarking of HACO. The algorithm achieved the lowest convergence time and runtime among the compared algorithms and showed statistically significant improvements. This confirms that hybrid metaheuristic design can improve uncertainty-aware routing performance. More importantly, the optimization layer is compatible with the thesis-level idea of driver-health-aware decision support, because risk scores from the AI layer can be added as penalties, constraints, or route-selection modifiers [16].

The AI objective was achieved through multimodal fatigue and health prediction models. ECG/HRV and eye modalities showed strong predictive performance, while respiration was shown to be unsuitable as a primary detector in the evaluated configuration. The best fusion results were obtained by combining ECG and eye, with all modalities producing only a marginal additional improvement. This supports the conclusion that multimodal systems should be selective and quality-aware rather than simply combining every available signal [17-25].

Future work should move the architecture from software-centric validation to real in-vehicle and fleet deployment. This requires production-grade sensors, long-duration recordings, varied driver populations, different vehicle types, and field conditions that include vibration, illumination changes, connectivity disruptions, and driver diversity. It should also include more advanced multimodal AI, calibrated uncertainty estimates, personalized baselines, explainability methods, edge-efficient inference, privacy-preserving learning, and stronger coupling between predicted driver risk and routing/resource policies [10-12,17-25].

Further research should also extend HACO and the decision-support layer beyond benchmark shortest-path problems. Future formulations can integrate real EV energy consumption, charging station availability, route risk, driver fatigue evolution, rest requirements, traffic uncertainty, and fleet-level constraints. In this way, the optimization component can evolve from an uncertainty-aware routing primitive into a complete health-aware EV planning tool [13-16].

## 6.3 Valorization of the Research Results

The valorization of the research results confirms that the thesis produced both theoretical and applied scientific outputs. These outputs support the review foundation, the IoT/cloud prototype, the HACO optimization contribution, and the final multimodal AI driver-fatigue framework [1,12,16,25].

The thesis research was valorized through several direct and complementary scientific outputs. The review-oriented foundation is supported by the publication "The Role of the Industrial IoT in Advancing Electric Vehicle Technology: A Review," which contributes to the state-of-the-art synthesis and supports the justification for treating driver health as a first-class service in EV ecosystems. The system architecture is supported by the paper "IoT-Enabled Driver Health Monitoring System Using Multi-Modal Sensors and Cloud Integration," which contributes directly to Chapter 3 through the low-cost multimodal monitoring architecture, ESP32-based sensing, MQTT/HTTP communication, Azure IoT Hub connectivity, visualization, and cloud-side anomaly detection [1,12].

The optimization contribution is supported by the Alexandria Engineering Journal paper "A novel Hybrid ant colony algorithm for solving the shortest path problems with mixed fuzzy arc weights," which forms the core of Chapter 4 and provides the HACO formulation, benchmark comparisons, and performance analysis. The AI and integrated monitoring contribution is supported by "A Hybrid Machine Learning and IoT System for Driver Fatigue Monitoring in Connected Electric Vehicles," which forms the main basis of Chapter 5 and contributes ECG/HRV modeling, eye-state analysis, respiration evaluation, fusion analysis, and cloud-connected deployment logic [16,25].

Additional publications contributed to the broader scientific development of the doctoral work, especially in AI, metaheuristic optimization, and hybrid predictive modeling. These complementary works strengthened the methodological maturity needed for model construction, benchmarking, performance interpretation, and interdisciplinary engineering integration. Overall, the valorization demonstrates a coherent correspondence between published outputs and thesis chapters: Chapter 2 is supported by the IIoT/EV review, Chapter 3 by the IoT-enabled monitoring system, Chapter 4 by the HACO optimization paper, and Chapter 5 by the hybrid machine learning and IoT fatigue-monitoring paper. Together, these outputs confirm the originality, relevance, and dissemination strength of the research [1,12,16,25].

## References

1. AlHousrya, O., et al., The Role of the Industrial IoT in Advancing Electric Vehicle Technology: A Review. Applied Sciences, 2025. 15(17): 9290.

2. Albadawi, Y., Takruri, M., and Awad, M., A review of recent developments in driver drowsiness detection systems. Sensors, 2022. 22(5): 2069.

3. Baltrušaitis, T., Ahuja, C., and Morency, L.-P., Multimodal machine learning: A survey and taxonomy. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2018. 41(2): 423-443.

4. Page, M.J., et al., The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ, 2021. 372.

5. Keele, S., Guidelines for performing systematic literature reviews in software engineering. Technical Report, 2007.

6. Petersen, K., et al., Systematic mapping studies in software engineering. 12th International Conference on Evaluation and Assessment in Software Engineering, 2008.

7. OASIS Standard, MQTT Version 5.0. 2019.

8. Satyanarayanan, M., The emergence of edge computing. Computer, 2017. 50(1): 30-39.

9. Shi, W., et al., Edge computing: Vision and challenges. IEEE Internet of Things Journal, 2016. 3(5): 637-646.

10. UNECE, Regulation No. 155 - Cyber security and cyber security management system. United Nations Economic Commission for Europe, 2021.

11. ISO, ISO/SAE 21434:2021 - Road vehicles: Cybersecurity engineering. International Organization for Standardization, 2021.

12. AlHousrya, O., et al., IoT-Enabled Driver Health Monitoring System Using Multi-Modal Sensors and Cloud Integration. IEEE SIITME, 2025.

13. Deng, Y., et al., Fuzzy Dijkstra algorithm for shortest path problem under uncertain environment. Applied Soft Computing, 2012. 12(3): 1231-1237.

14. Tajdin, A., et al., Computing a fuzzy shortest path in a network with mixed fuzzy arc lengths using alpha-cuts. Computers & Mathematics with Applications, 2010. 60(4): 989-1002.

15. Di Caprio, D., et al., A novel ant colony algorithm for solving shortest path problems with fuzzy arc weights. Alexandria Engineering Journal, 2022. 61(5): 3403-3415.

16. AlHousrya, O., et al., A novel Hybrid ant colony algorithm for solving the shortest path problems with mixed fuzzy arc weights. Alexandria Engineering Journal, 2024. 109: 841-855.

17. Lu, K., et al., Detecting driver fatigue using heart rate variability: A systematic review. Accident Analysis & Prevention, 2022. 178: 106830.

18. Liu, D., et al., OpenDriver: An open-road driver state detection benchmark. Journal of Network and Computer Applications, 2025. 242: 104279.

19. Meteier, Q., et al., A dataset on the physiological state and behavior of drivers in conditionally automated driving. Data in Brief, 2023. 47: 109027.

20. Alajlan, N.N., and Ibrahim, D.M., DDD TinyML: A TinyML-based driver drowsiness detection model using deep learning. Sensors, 2023. 23(12): 5696.

21. Abtahi, S., et al., YawDD: A yawning detection dataset. Proceedings of the 5th ACM Multimedia Systems Conference, 2014.

22. Guede-Fernandez, F., et al., Driver drowsiness detection based on respiratory signal analysis. IEEE Access, 2019. 7: 81826-81838.

23. Breiman, L., Random forests. Machine Learning, 2001. 45(1): 5-32.

24. Alzubaidi, L., et al., Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. Journal of Big Data, 2021. 8(1): 53.

25. AlHousrya, O., et al., A Hybrid Machine Learning and IoT System for Driver Fatigue Monitoring in Connected Electric Vehicles. Machine Learning with Applications, 2026: 100845.