Thesis for The Degree of Licentiate of Engineering

Towards Leveraging Underutilized IoT Resources for Automotive Software: A Study on Resource Sharing for Connected Vehicles

Malsha Ashani Mahawatta Dona

Division of Interaction Design & Software Engineering
Department of Computer Science & Engineering
Chalmers University of Technology and Gothenburg University
Gothenburg, Sweden, 2025

"Any sufficiently advanced technology is
indistinguishable from ***magic****"*
- Arthur C. Clarke

Abstract

Background: Internet of Things (IoT) has found its way to day-to-day lives of people, making their lives convenient and connected. As a result of growing IoT usage, we are surrounded by potentially underutilized computing resources suggesting an opportunity to improve their utilization by sharing them with entities that may require from time to time more resources for complex computations. However, the resource allocation and resource utilization process is a rather complex task that involves heterogeneous and distributed entities. This concept can be particularly relevant in highly dynamic and safety-critical domains such as transportation and automotive systems, where edge devices such as vehicles operate in a resource-constrained environment.

Objective: This research focuses on to what extent underutilized resources in a highly dynamic heterogeneous environment can be utilized more efficiently. This thesis explores resource sharing in edge devices within close vicinity on the example of connected vehicles. We propose and evaluate this concept on a practical application scenario from the automotive domain, where one vehicle would benefit of being able to "look around the corner".

Method: A systematic mapping study was conducted to identify the key research areas and limitations of the automotive domain focusing Vehicular ad-hoc Networks (VANETs). Based on the results of the literature review, an explanatory study was conducted to present the proposed resource utilization framework exploring the novel research areas identified. A series of experimental-based evaluation studies following design science was conducted to explore the applicability of state-of-the-art large language models (LLMs) as dialogue interfaces within the proposed resource utilization framework. This line of studies includes identifying and mitigating the potential challenges of using LLMs as a tool to support resource utilization.

Findings: The results of the study revealed that resource utilization can be achieved through sharing underutilized computing resources of nearby IoT-enabled entities. Within the context of the selected practical application scenario, our experiments showed that LLMs can support pedestrian detection and localization. LLMs can initiate a dialogue between connected vehicles and process relevant multimodal data to contribute to improved decision-making in autonomous driving (AD). Further experiments evaluated novel techniques to assess the trustworthiness of such LLM-assisted systems.

Conclusion: The introduction of state-of-the-art artificial intelligence (AI) tools such as LLMs has the potential to positively impact advanced driver assistant systems (ADAS), establishing a new research dimension to the automotive context. This novel approach aims to enhance the adaptability and efficiency of the proposed framework for safety critical systems, demonstrated with an industrially relevant practical application scenario.

Keywords Internet of Things, Resource Utilization, Automotive, Large Language Models, Trustworthiness, Hallucination Detection and Mitigation

Acknowledgment

First I would like to extend my heartfelt gratitude to my supervisor Professor Christian Berger for his continuous guidance and unwavering support. His encouragement, patience, availability, and immense knowledge were invaluable to me throughout the past couple of years.

I am very thankful to my co-supervisor Assistant Professor Yinan Yu for providing me with insightful feedback and valuable advice throughout this journey. I am forever grateful for the kind and constant support she extended.

I would also like to thank Dr. Beatriz Cabrero-Daniel for her exceptional support and constant encouragement throughout the research study. I am grateful for the supportive and uplifting work environment she created within the research group which positively impacted my research journey.

Next, I would like to thank my examiner Professor Alejandro Russo for his kind support and constant feedback which guided me to keep the research aligned with the expected scope and the quality. I would also like to thank the CSE graduate school and administrative staff for all the assistance provided.

My special appreciation goes to all the wonderful friends I met, Krishna, Bea, Cristy, Ricardo, Ranim, Wardah, Habib, Razan, Mazen, Teodor, Hannah, Pierre, Sabina, Khaled, Hamdy, Amna, Sushant, and Tayssir. Thank you for being there with me sharing both happy and challenging moments in my life.

I would like to thank my parents Anoma Weerakoon and Leslie Mahawatta for their unwavering love and support, and my loving husband Damith Wijethunge, for supporting me, for believing in me, and for sharing my emotional baggage on a daily basis with zero complaints. I thank my aunt Sagarika and late uncle Sarath for loving me and being there for me as if I were their own child. I would also like to thank my siblings Chanaka, Thushani, Kasun, Samira, Vindya, and Akila, and my lovely niece Kenuki and nephew Dion for their unconditional love.

This research is supported by the Swedish Foundation for Strategic Research (SSF), Grant Number FUS21-0004 SAICOM, the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation and the Swedish Research Council (Diarienummer: 2023-03810).

List of Publications

Appended publications

This thesis is based on the following publications:

[A] M. A. Mahawatta Dona, B. Cabrero-Daniel, C. Berger, Y. Yu.
"Home Sharing for Internet-of-Vehicles - A Systematic Mapping Study"
Under submission to the journal ACM Transactions on Internet of Things, special issue on Autonomous Driving 2025.

[B] M. A. Mahawatta Dona, C. Berger, Y. Yu.
"AirDnD-Asynchronous In-Range Dynamic and Distributed Network Orchestration Framework"
Proceedings of the 43rd International Conference on Distributed Computing Systems (ICDCS) 2023.

[C] M. A. Mahawatta Dona, B. Cabrero-Daniel, Y. Yu, C. Berger.
"Tapping in a Remote Vehicle's onboard LLM to Complement the Ego Vehicle's Field-of-View"
Proceedings of the 50th Euromicro Conference on Software Engineering and Advanced Applications (SEAA - WiP), 2024.

[D] M. A. Mahawatta Dona, B. Cabrero-Daniel, Y. Yu, C. Berger.
"Evaluating and Enhancing Trustworthiness of LLMs in Perception Tasks"
Proceedings of the 27th IEEE International Conference on Intelligent Transportation Systems (ITSC) 2024.

[E] M. A. Mahawatta Dona, B. Cabrero-Daniel, Y. Yu, C. Berger.
"LLMs Can Check Their Own Results to Mitigate Hallucinations in Traffic Understanding Tasks"
Proceedings of the 36th International Conference on Testing Software and Systems (ICTSS) 2024 .

Research Contribution

The CRediT (Contribution Roles Taxonomy) model [1] was employed to determine the authors' contribution to the appended papers of this thesis. I bear the primary responsibility of planning, designing, and writing of the thesis. I contributed to co-conceptualizing, co-designing the methodology and experiments, co-execution of the experiments, data gathering, data analysis, and authoring the papers A, B, C, D, and E. Table 1 outlines my contributions to the appended papers.

Role	**Papers**				
	A	**B**	**C**	**D**	**E**
Conceptualization	✓	✓	✓	✓	✓
Methodology	✓		✓	✓	✓
Software	✓		✓	✓	✓
Validation	✓	✓	✓	✓	✓
Formal Analysis	✓	✓	✓	✓	✓
Investigation	✓	✓	✓	✓	✓
Resources			✓	✓	✓
Data curation	✓		✓	✓	✓
Writing – Original Draft	✓	✓	✓	✓	✓
Writing – Review & Editing					
Visualization	✓	✓	✓	✓	✓
Supervision					
Project Administration					
Funding acquisition					

Table 1: The author's contribution to the appended papers of the thesis.

List of Abbreviations

AD	Autonomous Driving
ADAS	Advanced Driver Assistant Systems
AI	Artificial Intelligence
AirDnD	Asynchronous in-Range Dynamic and Distributed Network
BO3	Best of Three
CCAM	Cooperative, connected and automated mobility
C-ITS	Cooperative Intelligent Transport Systems
CPS	Cyber-Physical Systems
CRediT	Contribution Roles Taxonomy
FoV	Field of View
GPS	Global Positioning System
HMI	Human Machine Interfaces
IoT	Internet of Things
IoV	Internet of Vehicles
ITS	Intelligent Transportation Systems
LiDAR	Light Detection and Ranging
LLMs	Large Language Models
NFV	Network Function Virtualization
NN	Neural Network
OEMs	Original Equipment Manufacturers
RAG	Retrieval Augmented Generation
RoIs	Regions of Interests
SDNs	Software Defined Networks
THV	Two historical votes in sequences
THV-2	Two historical votes in sequences 2
VANETs	Vehicular ad-hoc Networks
VLMs	Vision-Language-Models
V2I	Vehicle to Infrastructure
V2V	Vehicle to Vehicle
V2X	Vehicle to Everything

Contents

Chapter 1

Introduction

In today's society, the way how people interact with each other has drastically changed due to remarkable modern technological advancements. Communication has become more accessible, efficient, and faster upon the availability of advanced wireless communication infrastructure together with advanced computing nodes. This has increased the number of individually owned computational devices and has also affected the quality and resources of such devices. The rapid growth of the Internet of Things (IoT) has made a significant impact on most day to day chores of the general public [2]. Varying from smartphones, and smart home appliances to connected vehicles and wearable health monitoring devices, IoT devices have become an integral part of the day-to-day life of modern society.

1.1 Background

While IoT devices are widely used by people to keep their lives connected and convenient, it has been noted that such IoT devices are not fully utilized to their maximum capabilities at all times [3]. For instance, some devices may operate in isolation, leaving computing power, communication capabilities, and data idle. This situation allows for dynamic and efficient means for IoT resource allocation and utilization to improve performance and provide better quality service to the users [4]. However, such resource allocation and utilization processes are challenged by the heterogeneous and distributed nature of the IoT devices.

The underutilization of IoT resources can be both challenging and opportunistic, due to the new possibilities that can be unlocked by better coordination and sharing of resources. More adaptable connected systems can be created with less energy waste and improved resource utilization by introducing better collaboration between computing nodes. The underutilization of IoT resources could occur in many domains, irrespectively of the application scenario [5].

In safety-critical application domains such as transport and automotive, better resource utilization would generate a significant impact, given the limited amounts of available computing resources within connected vehicles. For instance, advanced multi-modal sensor systems may occasionally exceed the available computational capacities of the infrastructure, where resource utilization could help in allocating required resources [6].

1.1.1 Domain–Specific IoT Applications

The usage of IoT devices has spanned across many diverse application domains, whereas it showcases greater potential and a rather large usage in certain domains. Figure 1.1 shows a variety of domains where IoT applications are significantly used. The IoT devices are on the verge of transforming both industries and individuals, improving and constantly changing how they operate.

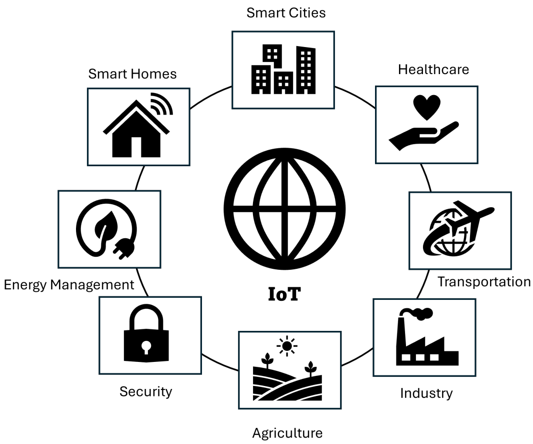

Figure 1.1: Applications of IoT in different domains (own diagram)

Domains such as healthcare have been rapidly changed for the better with IoT–powered wearable devices and monitoring devices that enable real-time tracking and monitoring of vital signs [7]. Such improvements led to medical marvels, including early diagnosis and improved personalized care. However, in healthcare, many IoT devices are used within controlled environments where reliable connections can be predicted together with centralized systems to offload and process data [8].

Even though real-time monitoring is vital for many critical applications within healthcare, most applications permit post-processing of the collected data. Also, such processes are supplemented with human oversight where manual interventions can take place. This allows flexibility and higher latencies for cloud-based computations to perform post-data analysis tasks [8]. Healthcare IoT devices also typically benefit from reliable networks, such as WiFi or cellular connections, often found within stationary environments such as hospitals and houses, and mobile environments if the device is a wearable.

The increased usage of IoT devices has a great impact not only on the healthcare domain but also on domains including agriculture and various other industries. In agriculture, IoT systems assist in improving farming by providing state-of-the-art monitoring systems for soil, weather, or crop status for example [9]. Similarly, initiatives such as smart cities also use IoT technologies to improve the living status of the people by providing services such as automated waste management, energy-efficient lighting systems, or smart heating systems. These systems are often powered by smart IoT devices, for instance, thermostats, cameras, switches, or lights, which constantly gather data [10, 11]. In domains like as manufacturing, supply chain management, and logistics, IoT has been an enabling driving force towards reduced downtime and operational costs. Predictive maintenance, real-time monitoring systems, fleet management, route optimization, inventory, and warehouse management have been greatly supported by IoT-based solutions [12, 13]. However, similar to the characteristics described for the healthcare domain, many of the above-mentioned domains benefit from similar properties that are not limited to reliable network connections, stationary environments, room for higher latencies and flexibility in computation-related tasks and analysis, and data offloading to centralized systems for further processing.

The automotive domain is also a fast-changing field that has benefited from many modern technological advances. The frequent use of IoT devices and their enhanced capabilities have positively impacted the future of the automotive industry, which envisions the evolution of advanced driver-assistant systems (ADAS) to autonomous driving (AD). For instance, IoT-enabled vehicle-to-everything (V2X) communication could significantly improve road safety and traffic management by providing real-time data sharing between vehicles and the infrastructure in close vicinity [14].

However, the automotive domain can be considered as a domain with critical requirements concerning IoT devices and their communication, given the highly dynamic environments where vehicles operate. The connectivity between IoT devices within the automotive domain may require reliable real-time adaptability despite the intermittent connectivity often found in such dynamic environments [15]. Sudden unpredictable conditions such as road hazards and traffic conditions tighten the need for seamless distributed computational frameworks within automotive systems.

The latency in processing the generated data and sharing it with connected devices should also be as low as possible within the automotive IoT systems, to prevent catastrophic failures [15]. For instance, the perception tasks that support ADAS systems require to process sensor data with low latencies. This does not only relate to the processing of data but also requires low latency connectivity within V2X devices to provide reliable driver assistance systems. Such systems operate autonomously in a real-time setup leaving minimal to no room for human intervention. This highlights the critical need for efficient allocation and utilization of resources including computational resources, storage and memory capacities, communication and bandwidth resources, software, and algorithmic resources, together with power and energy resources.

From all IoT application domains that have been covered here, the field of automotive stands out due to its requirements and opportunities for innovation. As mentioned above, automotive IoT applications work under highly dynamic environments where highly complex computational activities are often executed subject to resource constraints. Therefore, efficient utilization of underutilized resources by leveraging idle computational power and optimizing storage and memory usage is a relevant research area for the automotive field. An adaptive resource utilization based on dynamic connectivity could eventually contribute towards enhanced traffic safety.

1.1.2 Advancements in Automotive Computing and Connectivity

In the automotive field, many advanced technologies have been integrated into vehicles to improve the comfort and safety of passengers and other road users. With these modern technological advancements, vehicles are slowly turning into intelligent Cyber-Physical Systems (CPS) that could potentially enable AD with the support of V2X-enabled ADAS. Such modern vehicles aim for simplified system architectures [16] that could house powerful centralized processing units with adequate storage capacities to support future applications. These hardware-accelerated vehicles are not only capable of executing complex neural networks (NN) to support ADAS/AD features, but they are also capable of running Large Language Models (LLMs) locally, which have shown potential in supporting computer vision tasks that were earlier performed by specialized NN models (cf. [17, 18]). The successful adoption of such LLMs within vehicles is currently envisioned to improve in-car Human Machine Interfaces (HMI) [19] providing natural language-based conversations for its passengers [20, 21].

Such modern automotive applications show a wide range of objectives that vary from simple in-car passenger assistance systems to complex data processing tasks that support real-time decision-making activities, resulting in fluctuations in computational resource demands at a given time. Although some vehicles may perform relatively simple tasks, leaving portions of their computing resources underutilized, meaning that some of their processing powers remain idle or only partially used, there could be other vehicles in the vehicle's vicinity that may require additional computational power to complete demanding complex data processing tasks such as real-time object detection, localization, and decision making. The said imbalance of computational demand could present an opportunity to share resources by monitoring and disseminating such tasks and demands between different compute units in nearby vehicles to improve performance while utilizing excess computing power.

1.1.3 Resource Utilization for Automotive Industry

With the rise of IoT technologies, we are expected to be surrounded by millions of IoT devices within the next five years [2]. Such rapid increment of IoT devices will result in large amounts of potentially idling or partially used resources like processing power, storage, and network bandwidth in close vicinity. The systems formed by smartphones, smart appliances, and safety-critical components together with advanced communication infrastructure can be pictured as a "living organism" that dynamically adapts to changing demands. In such adaptive environments, resource utilization needs to cope with the allocation and use of idle or partially available resources to reduce waste of energy or computational resources, and enhance the performance of computing nodes such as edge devices that require higher amounts of computing resources for complex processing tasks from time to time.

A key aspect of this process is the orchestration of the computing tasks, which involves the management and distribution of tasks across multiple computing nodes. Efficient task orchestration would result in the dynamic allocation of computational tasks to available devices that would support real-time applications in automotive and other relevant domains. However, resource utilization and smart orchestration are intertwined together and are crucial aspects, especially for systems that consist of computing nodes such as edge devices and smart vehicles, where onboard computing resources are limited.

Strategically planned task orchestration and resource utilization of potential excess resources in vehicles could improve onboard processing capabilities and reduce modern vehicles' dependability on external infrastructure and advanced network connectivity. For instance, a vehicle that can leverage onboard computational resources can perform complex data processing to understand real-time traffic situations, which can not only be used by the same vehicle but even by a different vehicle in close vicinity that does not have sufficient computing resources or relevant data to provide its ADAS/AD service. This concept is commonly known as resource sharing, where one computing node allows the other computing nodes to use its own resources [22]. Resource sharing can be considered as one approach to resource utilization for the automotive domain.

The onboard data processing tasks involved in the resource-sharing process may not heavily rely on distant cloud servers or network connections, given that the data are processed within the same areas where they were generated. The improved overall response time in onboard data processing tasks positively supports real-time decision-making for AD/ADAS. Furthermore, limiting data processing to where it needs to be processed, in our case onboard of the data producing vehicle to be shared with interested vehicles in its close vicinity, reduces the demand for cellular communication and remote data processing in the cloud. This does not only reduce processing it, it also helps maintain data privacy and sovereignty regarding the data collected by the onboard sensors of modern vehicles.

As modern vehicles become more complex with advanced heterogeneous subsystems, new avenues of opportunities will open for resource utilization within the automotive context. Seeing the possibility of adopting LLMs, which have already demonstrated positive influence in different domains including education, research, or healthcare [23, 24], within vehicles would enable a next generation of ADAS/AD systems including opportunities for improved resource utilization.

In the automotive industry, some automotive Original Equipment Manufacturers (OEMs) are already experimenting with incorporating LLMs within vehicles to provide better in-car experiences to their users by allowing natural language-based conversations between the passengers and the vehicle [20, 21]. However, the LLMs' capabilities not only lie within natural language-based content generation but also in specialized computer vision-related tasks such as image description and object detection, which could enhance perception systems for ADAS/AD. Such LLMs utilized for vision-based recognition tasks are commonly known as Vision-Language-Models (VLMs) [25] and are often investigated to find optimal integration within Intelligent Transportation Systems (ITS) [26]. Also, commercial level applications such as LINGO-2 [27] and research work such as "Dolphins" [28] accelerate the use of VLMs as driving assistants to support the navigation of vehicles in real-world scenarios by invoking autonomous driving functions.

The versatility of LLMs extends beyond the processing of multimodal data to support automotive perception-related tasks. We foresee LLMs as a novel instrument to serve as a unified protocol for communication, especially within heterogeneous systems that usually operate under different communication protocols. The subsystems such as LiDAR, radar, or control units that typically communicate with proprietary or domain-specific protocols can benefit from such a unified dialogue interface. LLMs can bridge such heterogeneous systems, providing task-agnostic dialogue interfaces that leverage prompt-based interaction mechanisms to assist scenario understanding and context-aware reasoning. As LLMs continue to evolve at rather a rapid pace, we can anticipate that their role as an essential instrument in automotive application scenarios will potentially increase, providing more intelligent and adaptive transportation systems in the near future.

1.2 Motivation and Problem Domain

Improving resource utilization in the automotive domain requires greater attention considering the technological advancements that challenge the existing and limited resource infrastructure within modern vehicle fleets. Orchestration of computing nodes that require additional resources to perform complex computational tasks, and the computing nodes that have excess resources is crucial to utilize the available resources within close proximity. Efficient resource utilization in automotive systems becomes more imperative due to the limited amount of onboard computing resources within vehicles. In this thesis, we outline the concept of resource sharing as a means of resource utilization on the example of the automotive domain, in particular, we focus on resource sharing as a process that can benefit from the latest technological advancements enabled by LLMs.

We consider an application scenario that enriches the vehicle's field of view in its close vicinity to support navigation in urban environments. This application scenario involves intersections and their traffic scenarios and can be referred to as the "looking around the corner" problem. The problem consists of multiple vehicles entering the same intersection that could support each other by initiating a V2X communication. In this scenario, we observe multiple safety-critical systems in the form of vehicles and many smart appliances such as mobile phones and wearables in a close vicinity that could benefit from the proposed concept of resource utilization. The underutilized computing resources of onboard computing systems can be shared between others in close vicinity to perform complex data processing tasks like understanding the traffic situation around all involved entities.

This application scenario is a challenging problem in the automotive context given the urban nature of intersections, where occlusions are frequently occurring due to the buildings, parked vehicles, and other stationary objects as depicted by Figure 1.2. Such occlusions would create limited viewing ranges and blind spots for the vehicles that intend to enter the intersection, hindering their timely decision-making functionalities: The approaching vehicle, which we denote by the term "ego vehicle" in Figure 1.2, lacks real-time critical traffic information such as oncoming vehicles or pedestrians crossing the road until they enter the visibility range. Such information is vital for the safe navigation of the ego vehicle, as timely available information would enable safer maneuvering.

However, in urban intersection situations, other vehicles that are approaching the same intersection could potentially be available as depicted in Figure 1.2 labeled "vehicle A and B". Such vehicles may have different visibility ranges from which the ego vehicle could benefit from. The oncoming vehicles may have different sensor data such as LiDAR, radar, and camera feeds that could provide a more comprehensive view of the surroundings including the precise positions of stationary objects and their foreseen trajectories.

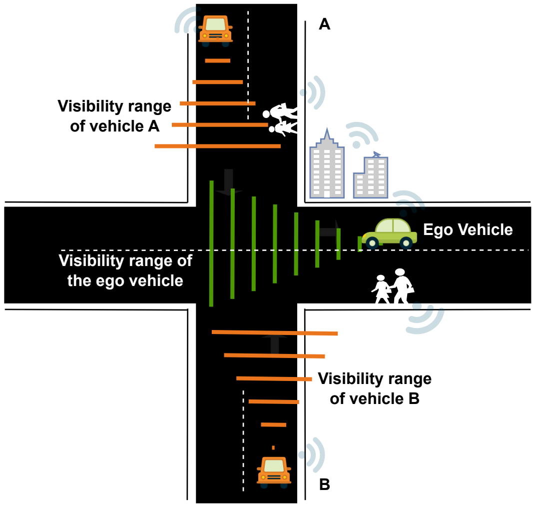

Figure 1.2: Graphical representation of a looking around the corner problem: The ego vehicle is approaching an intersection together with vehicles A and B. Each vehicle has its own field-of-view (FOV) represented by their respective colours, and that information could help each other understand the complete traffic situation at hand.

The collaborative perception between the ego vehicle and the oncoming vehicles in close vicinity can extend the field of view of the ego vehicle. However, in contrast to existing V2X approaches, we intend to not address this challenge by exchanging raw data and hence, consuming valuable resources from wireless communication infrastructure. Instead, we propose a resource utilization framework that enables the sharing of resources for processing sensor data using onboard computational resources to support collaborative sensing enabled by letting the involved vehicles having a dialogue about their respective information. Hence, addressing the practical and theoretical challenges that this scenario pose would contribute towards safer vehicle systems.

1.3 Scope of the thesis

This section provides an overview of the research boundaries outlining the parameters and limitations of the thesis. The specific areas covered within this thesis are highlighted to ensure clarity concerning the research direction and constraints.

We conducted a literature survey that aimed at understanding research gaps in the area of Vehicular Ad-Hoc Networks (VANETs), which facilitate the Internet of Vehicles (IoV) supporting the concept of resource sharing. VANETs provide vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication in localized surroundings [29]. Based on the results of the literature review, a resource orchestration framework was proposed to perform resource utilization in the field of automotive. This framework is proposed based on a relevant automotive application scenario called "looking around the corner", which is further described under Section 1.2.

The proposed resource utilization framework explores the use of LLMs within connected vehicles to support resource sharing between vehicles. This includes enabling dialogues between connected vehicles to query traffic situations in nearby environments to support perception-related tasks of onboard ADAS. However, evaluating the trustworthiness of such state-of-the-art technologies that support perception-related tasks in the field of automotive is essential in a safety-critical application setup. Trustworthiness plays a vital role in the field of automotive, especially within automated and connected vehicle environments, given that incorrect or unreliable perception could lead to fatal traffic safety risks. For instance, the use of LLMs often comes with a great possibility of generating hallucinated content; therefore, it is paramount to find suitable hallucination detection and mitigation techniques that can support LLM-powered perception-related tasks within connected vehicles. We propose hallucination detection and mitigation techniques that leverage multimodal data to address the safety-critical context of the proposed resource orchestration framework.

1.4 Research Design

We envision introducing resource utilization to enable a more efficient use of available computing resources within a selected vicinity. To contribute to this vision, we propose a framework entitled Asynchronous in-Range Dynamic and Distributed Network (AirDnD) for resource orchestration, which allows entities to engage in a problem-centric, solution-oriented dialogue. As modern ADAS already process substantial amounts of multimodal data from cameras, Global Positioning Systems (GPS), and inertial measurement platforms, simply streaming all such data between all interested participant entities is not an efficient use of available resources, including communication infrastructure such as VANETs and computing capabilities. We instead propose a concept to enable dialogues about the vehicles' environments instead of just sending what every entity sees on the raw data streaming level. We exploit onboard LLMs to initiate such dialogues between the vehicles to query and inform each other on higher levels about traffic situations focusing on how LLMs can support perception-related tasks such as pedestrian detection and localization.

However, LLMs may suffer from hallucination effects and therefore, we also need to consider strategies to detect and mitigate such risks. Hence, we propose suitable techniques for identifying and mitigating hallucinations by leveraging temporal information.

1.4.1 Research Goals and Questions

We derive the following RGs, RQs, and hypotheses to structure our research design:

RG1: To investigate the current state of research on resource sharing within VANETs for IoV, identifying key topics, trends, and open research challenges to inform future developments.

- **RQ 1.1** What are the topics in VANETs that have been addressed and what is the distribution of research articles across the different topics?
- **RQ 1.2** What novelty has been covered by the existing research articles and what topics are left for further research?

RG2: To design and evaluate a resource utilization framework for the automotive context, motivated by the "looking around the corner" application scenario by focusing on pedestrian detection and localization.

- **RQ 2.1** How to offload tasks among different computing nodes and what qualities and properties should be considered when selecting the nodes?
- **RQ 2.2** To what extent can a state-of-the-art LLM be used as a communication interface for automotive focusing on pedestrian detection and localization?

The following hypothesis was formulated under **RG2** to identify the approach to the utilization of resources for the automotive industry.

> **Hypothesis 2.1**: The offloading of computational tasks among different computing nodes in an automotive network can improve a vehicle's field-of-view in its surroundings for situations such as "looking around the corner". State-of-the-art LLMs can serve as an effective dialogue interface in such situations focusing on perception related traffic understanding tasks.

RG3: To evaluate and enhance the trustworthiness of the perception systems contribute to the concept of resource utilization.

- **RQ 3.1** How can hallucinations be characterized and how can hallucination detection strategies be enhanced in ADAS/AD perception and monitoring tasks?
- **RQ 3.2** What are the potential hallucination detection strategies suggested in the literature and to what extent can existing hallucination detection strategies be enhanced to support perception monitoring systems in ADAS/AD?

Within the scope of the **RG3**, the below hypotheses were formulated to investigate the trustworthiness of perception systems in the field of automotive.

Hypothesis 3.1 is based on the **RQ 3.1**. It is tested by subsequent experiments in our research design.

> **Hypothesis 3.1**: Characterizing hallucinations can enable the development of sophisticated hallucination detection mechanisms that could significantly enhance the safety and reliability of LLM-driven perception systems in ADAS/AD.

Hypothesis 3.2 states that existing hallucination detection and mitigation techniques can be adapted into the field of automotive to support perception-related tasks in ADAS/AD systems. This hypothesis is based on the **RQ 3.2** and later on, tested by conducting several experiments to evaluate the enhanced adaptations of selected hallucination detection strategies.

> **Hypothesis 3.2**: Hallucination detection strategies suggested in the literature, such as self-consistency checks and confidence scoring, can be adapted and enhanced to address the specific challenges of perception and monitoring tasks in ADAS/AD systems, improving reliability in dynamic automotive environments.

Table 1.1 represents how the research questions have been mapped to each appended paper and their local research questions.

1.4.2 Research Methodology

This section provides an overview of the research methodology followed in this thesis with respect to the appended papers. The first subsection outlines how the research goals, research questions, and hypotheses that are discussed in Section 1.4 are connected to each other, including a detailed description about how each paper has addressed these aspects. A detailed explanation of the research types followed in each appended paper and how they assisted in answering the relevant research questions are also included in the same section. The last subsection summarizes the overall research process followed in this thesis.

1.4.2.1 Research Methods

Figure 1.3 maps the research goals, research questions, and hypotheses to each of the appended papers.

The thesis consists of three main research goals that are spread across the five papers. RG1 consists of two research questions RQ 1.1 and RQ 1.2 that are answered by paper A. The paper A is a systematic mapping study that explores the literature on VANETs that facilitate the Internet of Vehicles (IoV).

Papers B and C address RG2, and each paper answers RQ 2.1 and RQ 2.2 respectively. Paper B proposes a framework for resource orchestration to achieve resource utilization within the field of automotive, using a practical application scenario. Paper C explores the state-of-the-art tools to be used as communication interfaces within such resource utilization frameworks focusing on pedestrian detection and localization in said practical application scenarios.

RQs	Paper ID	RQs addressed in the appended paper
RQ 1.1	A	**RQ 1:** What topics in VANETs have been addressed? **RQ 2:** What is the distribution of research articles across different topics?
RQ 1.2	A	**RQ 3:** What novelty has been covered by the existing research articles? **RQ 4:** What topics are left for further research?
RQ 2.1	B	**RQ 1:** What qualities and properties must be considered when selecting the computing nodes? **RQ 2:** How to offload the tasks? **RQ 3:** How to handle the computation? (Feasibility, privacy, integrity, and trust related to intellectual properties)
RQ 2.2	C	**RQ 1:** To what extent can a state-of-the-art LLM be used as a communication interface to qualitatively identify pedestrians? **RQ 2:** To what extent can a state-of-the-art LLM be used to reliably detect the location of pedestrians?
RQ 3.1	D	**RQ 1:** What are potential hallucination detection strategies suggested in literature? **RQ 2:** How can hallucinations be characterized when applying LLMs to pedestrian detection and localization for ADAS/AD? **RQ 3:** How can hallucination detection strategies be enhanced for use in ADAS/AD perception and monitoring systems?
RQ 3.2	E	**RQ 1:** To what extent can the SelfCheckGPT approach be adopted to spot potential hallucinations when using state-of-the-art LLMs for image captioning tasks for automotive usage scenarios? **RQ 2:** What is the performance of the adopted SelfCheckGPT approach on two state-of-the-art automotive datasets? **RQ 3:** To what extent is the performance of SelfCheckGPT affected by environmental conditions such as light or weather?

Table 1.1: Mapping the thesis research questions with the appended papers: This table illustrates how the overarching research questions of the thesis align with the research questions that are local to each appended paper.

RG3 is addressed by papers D and E where RQ 3.1 and RQ 3.2 are answered. Papers D and E present approaches for evaluating and enhancing the trustworthiness of the resource utilization approach presented in papers B and C.

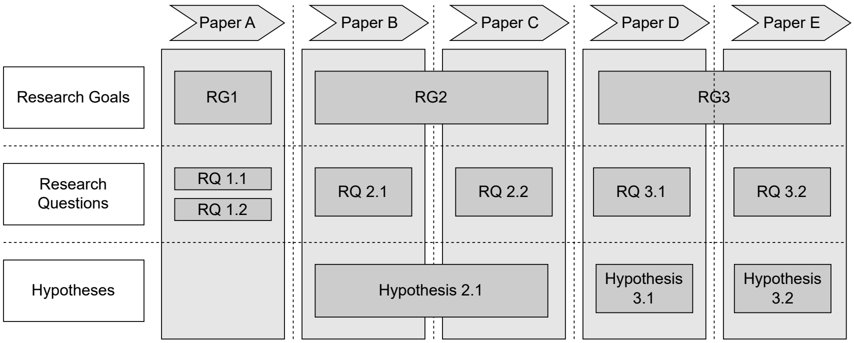

Figure 1.3: Overview of the research goals, research questions, and hypotheses of the thesis mapped to the appended papers

Figure 1.4 illustrates the detailed research design for this thesis concerning the appended papers. This figure follows the research decision-making structure proposed by Wohlin et al. [30]. This figure highlights that it is important to address multiple decision points when conducting research in both general and empirical software engineering. This figure mainly showcases three different phases of the research decision-making structure, which are depicted in three different colors as below.

Strategy Phase (Green) The strategy phase of the research design lays the foundation for the overall research to ensure that the research is carried out systematically. Once the research questions are formulated, the research outcome, research purpose and research paradigm are decided during this phase.

The first decision point of the strategy phase is the expected *research outcome*, which can be classified as *basic research* or *applied research* [31]. Basic research intends on generating knowledge whereas applied research focuses on finding solutions or improvements to a problem or an existing service. Based on these definitions, we can consider that paper A is basic research whereas the rest of papers B, C, D, and E are active research, where the research outcome is either a solution or an improvement.

The second decision point is the *research logic*, where the research direction is analyzed to classify whether the research logic is *deductive* or *inductive* [32]. The deductive research starts with a generic logic, which then evolves into a more specific logic. This is often considered as a top-down approach. Inductive research evolves from a specific point to a generic point, which can be considered the opposite of deductive research. The papers A, B, and C can be considered deductive research, where the study starts with a generic problem, which is then narrowed into a more specific form. However, papers D and E follow an inductive research logic, where the observations are further analyzed to detect patterns.

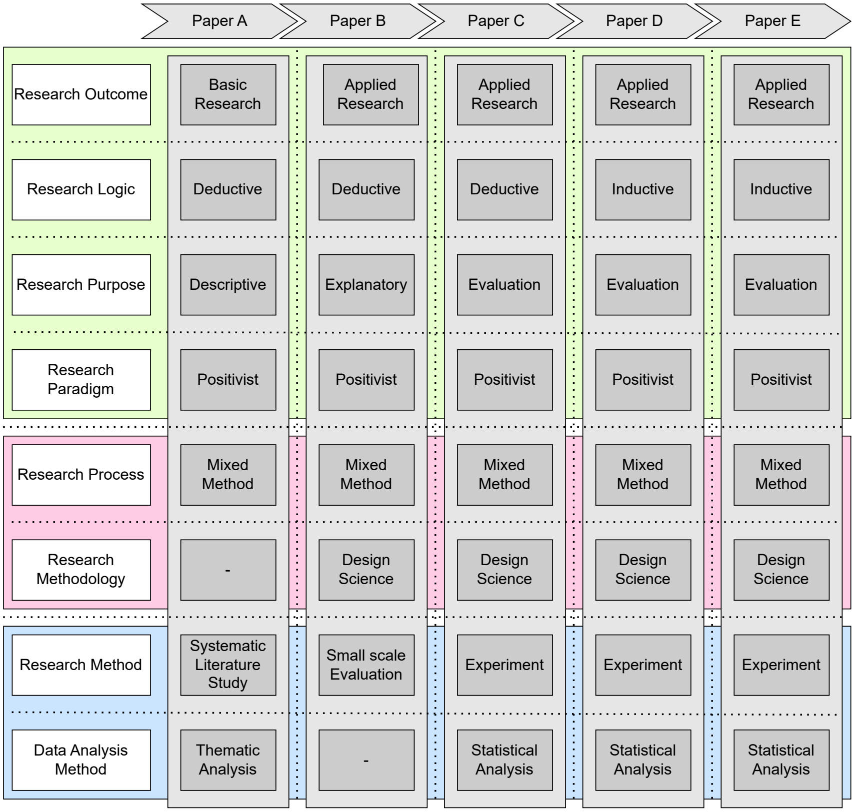

Figure 1.4: Research design of the appended papers. The first four categories indicated in green denote the research design's strategy phase, whereas the two categories separated by pink denote the tactical phase. The last category illustrated in blue demonstrates the operational phase of the research design.

The third discussion point focuses on the *purpose of the research*, which can be classified into four categories: *exploratory, descriptive, explanatory, and evaluation research* [31]. The papers that are appended in this thesis mainly fall under *descriptive, explanatory, and evaluation research* categories as illustrated by Figure 1.4. Paper A is a descriptive research paper, where the intention is to describe the existing literature and its gaps. Paper B is an explanatory research, where the paper focuses on addressing the problem of resource utilization. Paper C, D, and E are evaluation research studies, where the papers evaluate the proposed techniques to understanding their suitability in the automotive context.

Discussion point 4 is *research paradigm* and is classified into three main categories such as *Positivist, Constructivism, and Critical* [33]. The positivist research work is often based on the understanding that it is replicable due to the separation of concerns between the researcher and the reality. If another researcher repeats the research, the same conclusions should have been derived at the end regardless of who conducted the research. Constructivism, which is often referred to as interpretivism, is based on human activities and its subjectiveness. Usually, qualitative research methods such as interviews, ethnographies, etc. can be considered as interpretive research. The critical research targets on critically evaluating systems considering the impact of social and cultural variables. Given that papers A, B, C, D, and E are objective to the researcher and reality, they are designed to be replicable by a different set of researchers following the same procedures. Therefore, all appended papers can considered as positivist.

Tactical Phase (Pink) The tactical phase is linked to two main discussion points that prioritize on how to conduct the research to address the formulated research questions. The tactical phase enables the research to achieve its goals by following suitable processes and methodologies.

The first discussion point within this phase is the *research process* and it contains three main processes named *quantitative research, qualitative research, and mixed method research* [32]. Qualitative research consists of rich descriptions in natural language that is often used to understand and explain social phenomena. Qualitative research uses qualitative data that could potentially be collected from interviews, written documents, participant observations, or similar. Quantitative research uses quantitative data to describe the characteristics of the data population or predict relationships between the entities that participate in the experiments. Questionnaires, simulations, and experiments can be considered as some examples of quantitative research methods. Since we have used both qualitative and quantitative research methods in the appended papers, they can all be categorized under mixed method research as illustrated by Figure 1.4.

The next discussion point centers on *research methodology*, emphasizing three main approaches: *case study research, action research, and design science research* [30]. Case study research employs different data collection techniques to collect information from selected sources (cases) to explore and investigate natural phenomena [30]. Action research is similar to case studies; however, researchers may investigate further and develop a solution to address the problem that was studied using scientific approaches [30]. The design science methodology is more focused on building a knowledge base on a solution that addresses a specific problem at hand [30]. The solution could be an artifact that is both innovative and effective. In this thesis, we consider that papers B, C, D, and E follow design science, where the problems do not primarily come from organizations, yet building a knowledge base is crucial towards solving the problems.

Operational Phase (Blue) The operational phase includes decisions that are often related to the practical side of the research regarding its implementation. This is the last section of the overall research design decision-making structure that combines the research findings together with the research questions established before the strategic phase.

The first discussion point within the operational section is the *research method*, where the data collection techniques are decided. There could be various data collection techniques under this discussion point such as *interviews, observations, archival analysis, surveys, simulations, systematic literature studies, experiments, and small-scale evaluations* [34–36]. When we consider the data collection techniques of the papers of this thesis, Paper A falls under the systematic literature study category based on the systematic mapping study we conducted. Paper B proposes the framework AirDnD as a concept, therefore it follows the small-scale evaluation technique. The rest of the papers C, D, and E are experiments, where the objective is to test the formulated hypotheses.

The last discussion point in this stage is the *data analysis method* that targets how the collected data should be analyzed. Considering the papers of this thesis, we denote that paper A follows a thematic analysis to qualitatively analyze the collected data to report patterns and classifications. Since paper B follows a small-scale evaluation technique to conduct the research, it does not include any particular data analyzing step. However, papers C, D, and E use statistical analysis techniques to analyze the quantitative data.

1.4.2.2 Research Process

Figure 1.5 shows an overview of the research process followed in this thesis. The research process consists of a set of actions and steps that explain the flow of the research activities and how they are interconnected to each other with respect to the appended papers. The research activities and steps may not necessarily be mutually exclusive or not separate and distinct [37]. Therefore, it is important to note that some research activities may not entirely be completed within a single study and some connected steps may be repeated together as a cycle in multiple studies.

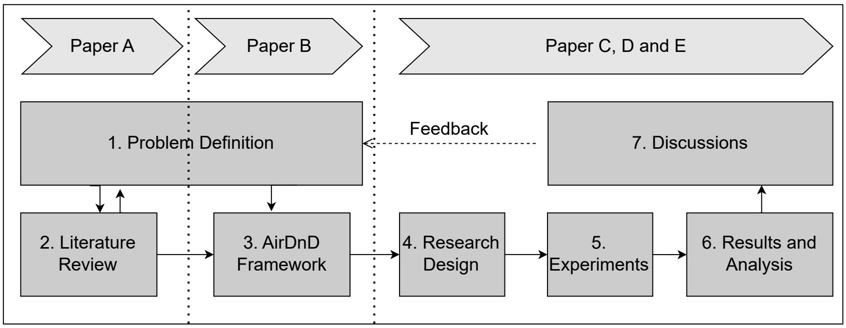

Figure 1.5: Research Process followed in the thesis.

The research process begins by defining the research problem and the problem domain. While this step is repeated in almost all appended papers, when we consider this thesis as a whole, the problem domain and the problem definition are formulated by papers A and B. The second step of the process is surveying the literature, which has been covered by paper A. As the next step, we propose AirDnD as a working hypothesis that answers the formulated research problem. The next research activities 4, 5, 6, and 7 are repeated in the papers C, D, and E as a cycle. The research design, conducting the experiments, analysis of the results, and discussion of the results are executed in that order in all three papers as illustrated in Figure 1.5.

1.4.2.3 Datasets and Data Pre-processing

During this research, we mainly used two recently collected, different automotive datasets with the intention of reducing non-controllable and potentially influential factors of a single given dataset. We used the Waymo Open Dataset - Perception dataset [38] and PREPER CITY [39] that were collected in 2021 by two different research groups.

The Waymo open dataset was collected by Google's subsidiary company Waymo in urban and suburban areas of the United States, supporting and facilitating research in the automotive field with a special focus on self driving technologies. The Waymo perception dataset contains 1150 driving scenes covering day and night scenarios. Each driving scene is approximately 20 seconds long and includes both Lidar and camera data with labels for predefined object classes. Three sample images from the Waymo open dataset that showcase different traffic patterns and different lighting conditions are shown in Figure 1.6. For this research study, the Waymo dataset was extracted with a sample rate of 10 Hz, focusing only on the front camera frames to capture the driver's perspective, which is relevant for perception related tasks such as pedestrian detection.

Figure 1.6: Three images from Waymo Open Dataset

The Waymo dataset is US-centric and therefore contains images that showcase vehicle types and driving styles that are specific to the US. To complement that, we used the PREPER CITY dataset, which was collected in 2021 in Gothenburg, Sweden. This dataset covers a metropolitan area, where different types of vehicles and different types of driving styles can be visible. This dataset includes 114 traffic segments and each of the segments is approximately 15 minutes long. This dataset also contains video feeds collected by multiple onboard cameras of the ego vehicle. Overall, the dataset contains more than 1.5 million video frames. Figure 1.7 shows three images taken from the PREPER CITY dataset that depict different city areas and different weather conditions.

We used the Waymo and PREPER CITY datasets for the experiments in papers C, D, and E. Even though the two datasets contain millions of images, given the limitations and constraints of the experimental setups, a limited number of images had to be selected depending on data curation criteria. Based on the research design of each of the papers, different traffic scenarios were selected, where different traffic agents are visible.

Figure 1.7: Three images from PREPER CITY Dataset

To answer RQ2.2 by appended paper C, we used a specifically curated dataset with two subsets of the Waymo open dataset images, wherein one subset, a single pedestrian on a crosswalk is visible towards the middle of the image. The other subset includes images, where there are no pedestrians visible, and it was used as a control group throughout the experiment. In paper C, the evaluations were done based on the bounding box coordinates of the pedestrians which were later referred to as the ground truth. To retrieve the bounding box coordinates of the pedestrians, the VGG Image Annotator [40] tool was used.

In paper D, a different data curation technique was used to pre-process the selected images before the experiments. To answer RQ3.1, according to the research design of paper D, driving scenarios were picked from both Waymo and PREPER CITY datasets, where each driving scenario contained at least four frames. This ensured that each selected driving segment contained temporal data on the video indicating the gradual movements of visible traffic objects. The traffic scenarios that captured a single pedestrian or a cohesive group of pedestrians were selected based on the research constraints during the data curation process. The selected driving scenarios were then pre-processed under two steps. The first step was to systematically remove the horizon of each image considering as only the ground-level details are important. The cropped images were then split into four regions of interest (ROIs) to identify and classify the left, right, far, and close regions of the ground plane. With the assumption that the camera's optical axis is parallel to the road surface given that the images are captured by the front camera from the ego vehicle, we considered that the ego vehicle's orientation provides a reference point that is consistent throughout the entire subset of images selected. These ROIs were later on used in the experiment to identify the semantic localization of the pedestrians. The last step of the data curation process was to collect the ground truth data on each ROI considering whether a pedestrian or a cohesive group of pedestrians was present. This process was completed by a human evaluator with the support of a simple Python tool.

To answer RQ 3.3, Paper E conducted an experiment using both datasets. In this experiment, the only constraints towards the data curation process were the processing time of each image and the cost that was spent on image caption generation when proprietary models were used. Considering these two constraints, a subset of images was picked ensuring a fair distribution of traffic objects within the selected subset of images. To understand the traffic object distribution, lighting conditions, and weather conditions in the images, the available object labels within both datasets were extracted.

	Paper C	**Paper D**	**Paper E**	**Total**
Waymo	264	209	920	1393
PREPER CITY	0	258	920	1178
Total	264	467	1840	**2571**

Table 1.2: The datasets curated for papers C, D, and E

The number of images used for the experiments conducted within the scope of paper C, D, and E are displayed in Table 1.2.

1.5 Summary of Findings

This section summarizes the findings of the appended papers with respect to the research goals and research questions formulated within the thesis. Subsections starting from Section 1.5.1 to Section 1.5.3 reflect the research goals of the thesis and the corresponding research questions. The findings of the relevant appended papers are discussed in detail under each research question.

1.5.1 Paper A - Exploring VANETs focusing on resource sharing for Internet of Vehicles

This goal is achieved by answering two research questions, which are mentioned below. The research study published in **paper A** was designed to answer the research questions.

Research Goal Paper A explores the existing literature published on VANETs focusing on resource sharing for IoV. A systematic mapping study was conducted to review and classify the literature surveys published under the relevant area with the intention of answering the below research questions.

Methodology Paper A provides a meta systematic mapping study on 66 literature reviews published on the topic of VANETs and LLMs to explore the research gaps and limitations that have not been addressed by the existing studies. A review protocol that consists of multiple steps was followed to gather and filter the existing related literature review papers. Four databases IEEE Xplore, ACM Digital Library, Google Scholar, and Springer Link were queried with a search string to curate the initial paper bank. Primarily, we identified 2336 papers through the database search, which was gradually reduced to 66 after applying the next steps of the exclusion criteria.

Results and Contributions The main findings of paper A are aggregated into multiple tables to showcase the contributions of each selected paper together with the results of a thematic analysis that highlights the clusters of domains each paper is focusing on.

RQ 1.1 - What are the topics in VANETs that have been addressed and what is the distribution of research articles across the different topics?

The results of Paper A showcase that a significant number of studies has focused on topics such as security, communication protocols, clustering algorithms, and routing protocols. Further analysis corroborates to the conclusion by showing that "Routing and Networking" and "Security and Trust" clusters often in the selected literature review papers. In addition to that, further investigation that includes the publication years of the selected studies with their citation scores reveals that a higher number of highly cited research work has been conducted in the above two categories in the past recent years.

RQ 1.2 - What novelty has been covered by the existing research articles and what topics are left for further research?

Paper A highlights fields where more research work has been conducted as well as the novel research areas and domains that are under-explored by the existing literature. The research areas such as real-world validations, robustness, and authentication have received limited attention from the research community compared to the research domains such as "Routing and Networking" and "Security and Trust". These areas can be considered as crucial and enabling segments towards the concept of "home sharing" in the context of IoV, where vehicles are enabled to tap into the nearby vehicles to retrieve relevant situational awareness related information. However these areas still remain as under-explored subject areas where insufficient research efforts have been made. Furthermore, paper A exposes that concepts such as task offloading, resource sharing, smart orchestration, and LLMs are also under-explored, but highly innovative novel areas within IoV where more research focus should be directed to unlock their potential to support VANET-based systems such as Cooperative, connected and automated mobility (CCAM) and Cooperative Intelligent Transport Systems (C-ITS).

1.5.2 Paper B and C - Resource utilization framework for automotive using "looking around the corner" application scenario focusing on pedestrian detection and localization

This research goal consists of two research questions that are answered by the two appended papers, **Paper B and Paper C**. The research work carried out under these two papers respectively answers the below research questions.

RQ 2.1 - How to offload tasks among different computing nodes and what qualities and properties should be considered when selecting the nodes?

Research Goal Paper B proposes a resource utilization framework for distributed edge computing resources by creating a dynamic mesh network capable of sharing computational workload. This proposed framework is called AirDnD and it assists geographically distributed edge devices to collaborate dynamically for task execution to minimize back-and-forth data transfer.

Methodology Paper B focuses on the literature published on VANETs, Software Defined Networks (SDNs), and Network Function Virtualization (NFV) to identify the gaps and limitations of the existing research studies, focusing on novel trends that can be applied to the proposed AirDnD framework. As a result, this study proposes to develop three connected models that facilitate the concept of task offloading through a dynamically formed mesh network within the edge devices. This framework is proposed to apply in the practical application scenario "looking around the corner" (cf. Section 1.2) for evaluation and validation purposes. This supports connected vehicles in a close vicinity to understand the traffic patterns with a wider range of view of the approaching vehicles in an intersection.

Results and Contributions Paper B proposes a novel vision to enable resource sharing among edge devices. The study presents the three-layered architecture for the three models, network description, task description, and data description. The models collectively facilitate task offloading to achieve resource sharing between the vehicles in the application scenario. This study proposes to generate and dissolve dynamic mesh networks between the vehicles for connectivity.

Paper B proposes the development and evaluation of the models that facilitate the task offloading process as mentioned above. The offloading of tasks is a process that needs to be carefully executed by considering different parameters and properties. The paper presents the layered architecture of the proposed three models, which explains how to offload the tasks between the nodes, i.e. connected vehicles. The properties and the qualities that should be considered are varied and mainly categorized under two types. The properties and the quality of the network should be considered when selecting the nodes and this is proposed to be formulated as the second research goal of paper B. The second type of properties and qualities are related to the data and the tasks. What type of data and the quality of the data expected for the offloading of computational tasks are investigated under the final research goal of paper B.

RQ 2.2 - To what extent can a state-of-the-art LLM be used as a dialogue interface for automotive focusing on pedestrian detection and localization?

Research Goal Paper C aims to evaluate how effective LLMs can serve as a dialogue interface between the connected vehicles, for instance, to complete complex tasks such as the detection and localization of traffic agents. This paper breaks down the perception-related tasks into pedestrian detection where pedestrians that are not in direct line of sight of the vehicle are focused more.

Methodology Paper C follows an experiment-based research method under the design science research methodology to achieve the above research goal. An initial preliminary set of experiments was conducted to test the feasibility of the idea and to formulate the hypothesis that is being tested within the line of research. A number of state-of-the-art LLMs was used to evaluate their image captioning capabilities. The main experiment consists of three sub-experiments that were conducted using an automotive dataset and two proprietary recently released version of LLMs. The first sub-experiment included binary pedestrian detection, where the LLM responded with yes or no answers, when asked whether there was a human pedestrian visible in the picture. The second experiment included evaluating how well LLMs can generate bounding boxes around the previously detected pedestrians, providing the localization of the pedestrian. The third experiment involved testing and evaluating different prompts to understand to what extent the prompt engineering could help enhancing the accuracy and recall of the pedestrian localization.

Results and Contributions The preliminary experiments showcased that LLMs can accurately describe the images to an exceptional level. Furthermore, it was revealed that LLMs were excelling at object detection in a zero-shot setting [41]. These results paved the way for the main experiments, where pedestrian detection and localization were tested using LLMs. The binary pedestrian detection was successful with near-perfect recall values for the included LLMs. Even though the higher recall values represent a highly reliable pedestrian detection mechanism that is capable of identifying pedestrians accurately, this comes with false positives, where the system tends to detect pedestrians in images, where actual pedestrians are not visible. The second experiment shows that general-purpose, commercially available LLMs can also localize pedestrians with bounding boxes. The LLM-generated bounding boxes covered almost half of the ground truth bounding boxes on average. The third experiment showcased that the results were arbitrary from the prompt patterns, however alternative prompt formulations should be studied further to provide a concrete conclusion on that front.

Paper C discovers that state-of-the-art LLMs are capable of detecting pedestrians with near-perfect recall values. Higher recall values are critical in the domain of automotive, specially when focused on perception-related tasks such as pedestrian detection, where not detecting an actual pedestrian could lead to potential fatal collisions. The above exemplary results indicate that LLMs can be used within vehicles as dialogue interfaces to perform complex tasks such as pedestrian detection. Even though the results of pedestrian localization were not satisfactory to be applied in the automotive domain given its safety critical requirements, this study still showcased that the idea of using an LLM as a dialogue interface for pedestrian localization is possible. Future updates of these state-of-the-models and new findings of fields such as prompt engineering and retrieval-augmented generation (RAG) [42] could potentially enhance the object classification and localization skills of LLMs.

Collectively, papers B and C support **hypothesis 2.1** by showing that offloading computational tasks among several other nodes in an automotive network would improve a vehicle's FoV. This allows vehicles to better utilize available computational resources for understanding better traffic situations such as looking around the corner, where LLMs can serve as dialogue interfaces.

1.5.3 Paper D and E - Evaluating and enhancing the trustworthiness of the perception systems allows resource utilization

This research goal was addressed by **paper D and paper E**. The studies conducted under these papers respectively answered the two research questions below.

RQ 3.1 - How can hallucinations be characterized and how can the hallucination detection strategies be enhanced in ADAS/AD perception and monitoring tasks?

Research Goal The research goal of paper D is to evaluate how trustworthy LLMs are when they are used to assist perception-related tasks such as pedestrian detection in urban environments to support safety critical systems for instance ADAS/AD systems.

Methodology Paper D first identifies the types of hallucinations that may occur within the context of ADAS/AD perception and monitoring tasks that particularly focus on pedestrian detection. Existing hallucination techniques are reviewed to identify which techniques are suitable to apply within the context of the automotive domain when LLMs are used to detect pedestrians using multimodal data. The main experiment was conducted to assess hallucination detection techniques extracted from recent publications. The experiment involved curating an image dataset using several selected driving scenarios based on pre-defined inclusion criteria and feeding them to the LLM with a prompt that checks whether a human pedestrian is visible in the image. This experiment was followed under the "best of three (BO3)" strategy proposed in an existing publication. The data collected by the main experiment was analyzed to propose a new hallucination technique that could assist in mitigating hallucinations in an LLM-supported perception and monitoring system.

Results and Contributions Paper D identifies three main types of hallucinations that may occur in perception-related tasks within the automotive domain. The second data analysis activity reported that the hallucination detection technique BO3 was able to increase the performances for automotive-related perception tasks that are performed using proprietary LLMs, whereas poor overall performance was recorded for the pedestrian detection tasks conducted using open source models despite which hallucination detection technique was applied. The use of RoIs instead of full pictures demonstrated slightly deviating results for each technique, however, in general, the processing of RoIs supported in localizing the pedestrians. The proposed hallucination detection techniques Two historical votes in sequences (THV) and Two historical votes in sequences 2 (THV-2) use temporal data to reduce hallucinations. Even though both techniques resulted in lower results compared to full-frame image analysis, THV-2 increased the performance of the open-source model by a considerable amount. The proposed physical plausibility check technique increased the recall values for the proprietary model indicating that in general the hallucination detection techniques that use temporal data could be promising and suitable for the applications in the ADAS/AD domain.

Paper D answers the research question by explaining how hallucinations can be characterized in the automotive domain focusing on pedestrian detection. This paper elaborates why false negative scenarios (LLMs being unable to detect the actual pedestrian) are more critical for automotive use cases when LLMs are applied to support perception and monitoring tasks within safety critical systems. The paper improves and evaluates the BO3 method to be used within ADAS/AD perception systems. In addition to that, as a result of the statistical analysis of the experiments, this paper presents a new hallucination detection approach that exploits temporal data to detect hallucinations in pedestrian detection.

The overall findings of the paper support the **hypothesis 3.1**. Characterization of hallucinations enables the identification of failure patterns in LLM output, which allows the development of sophisticated hallucination detection mechanisms to recognize and flag or even correct failures. This resulted in proposing new hallucination detection techniques that enhance the LLM-assisted perception monitoring systems in ADAS/AD.

RQ 3.2 - What are the potential hallucination detection strategies suggested in the literature and to what extent existing hallucination detection strategies can be enhanced to support perception monitoring systems in ADAS/AD?

Research Goal The main research goal of paper E was to adopt the existing hallucination detection strategy SelfCheckGPT [43] for multimodal data retrieved from the automotive domain to support perception tasks for ADAS/AD and assess the performance of the adopted technique using state-of-the-art LLMs.

Methodology This study designs an experimental setup that adopts SelfCheckGPT for multi modal automotive data that are selected from two different datasets, collected in 2021 in two countries. The experimental setup starts with multiple VLMs describing the traffic objects in each image. This process is repeated several times to record multiple responses for each image. The first response that is recorded for each image is then broken down into sentences and processed further by an LLM in a Retrieval-augmented generation (RAG) [44] based approach, providing the latter responses as context to check the consistency. After the consistency level analysis, this technique is capable of eliminating the sentences with a lower level of consistency that could potentially be hallucinations. SelfCheckGPT is tested and evaluated with different LLMs to understand the feasibility of executing the hallucination detection and mitigation process completely offline, making it more applicable within safety-critical systems. The collected results were also further analyzed to see how different datasets and times of the day pictures were captured would impact the LLMs in causing hallucinations.

Results and Contributions This study identifies two main definitions of correctness in the recorded responses based on the traffic objects that are described. Considering the definition of hallucinations, the first definition of correctness depends on not having nonsensical traffic agents in the generated response. However, this definition does not necessarily check whether all traffic agents present in the image are described by the response, even though it is capable of identifying hallucinated agents. As long as the traffic agents identified by the response are all present in the image, the response is considered as correct. The second definition of correctness was introduced by considering the possibility of overlooking traffic agents. The LLM could have missed identifying some traffic agents even though it did not generate any nonsensical objects. This scenario can also be considered a hallucination since it provides unfaithful information. Therefore the second definition relates to the correctness as not overlooking traffic agents. This definition is crucial for the automotive context since missing information about traffic agent presence could lead to fatal collisions when automated decision-making processes are followed. Based on the above two definitions, results are reported separately to identify the impact made by each dataset on hallucination detection and how the time of the day affected the hallucination detection process. In addition, results are reported for both original and fixed responses after applying the sentence-level filtering technique. The results show that the proprietary model generated captions provide better results when SelfCheckGPT is applied whereas the open-source model showcased comparatively lower results. The results showcase that SelfCheckGPT is better at identifying non-hallucinated content in the automotive domain, however, it comes with the price of missing some hallucinations. No significant deviations of performances were recorded between the two datasets used in the experiments indicating that the differences in traffic patterns in respective geographical areas do not pose any impact on hallucination detection. The next step of the analysis inspected how light conditions affected the hallucination detection process. The results indicated that daytime captured images show better results compared to the images captured during nighttime and dawn and dusk time.

Paper E presents an overview of potential hallucination detection techniques that use SelfCheckGPT as the baseline. In addition to that, this paper demonstrates how well the text-based SelfCheckGPT technique can be enhanced to support multi-modal, automotive data indicating their applicability in perception and monitoring systems in ADAS/AD supporting the **hypothesis 3.2**. The detailed analysis of results emphasizes how different external factors may affect the performances of the adopted hallucination detection technique.

1.6 Threats to Validity

Even though the thesis and its appended papers present many insights related to resource utilization for the automotive domain, it is crucial to acknowledge and address potential threats to the validity of the findings. The possible threats to validity could be identified after conducting a careful validity analysis which can be later addressed by following a mitigation strategy [45]. We followed the study conducted by Feldt and Magazinius (2010) [45] and the guidelines presented by Wohlin et al. [46] to assess the potential threats. This section is based on four main types of validity threats presented by Wohlin et al. [46]: conclusion, internal, construct, and external.

Conclusion Validity refers to the reliability of the observed relationship between research treatments used within the study and the actual outcome. The factors and potential issues that may hinder the ability to draw accurate conclusions on connections between variables in the study are fallen under this category. In this thesis, the appended paper A, the systematic mapping study, inherits a conclusion validity threat given that its findings are dependent on the selected inclusion and exclusion criteria. A diverse, representative set of included papers could affect the robustness of the conclusions of the study. This effect has been minimized within the study by adopting a rigorous review protocol. In addition to that, the experimental studies conducted under papers B and C can be dependent on the datasets and the prompts that were utilized. To mitigate these threats, the experiments were designed and conducted with multiple LLMs and datasets.

Internal Validity focuses on the relationship that could exist between the proposed research methodologies and the outcomes of the studies. The effectiveness and reliability of the resource utilization framework proposed in paper B could be potentially influenced by the assumptions related to edge device capabilities and network stability. Also in the latter experimental studies, the authors have carefully designed the experiments including very diverse scenarios which may not generalize the results for all ADAS/AD context. Furthermore, the LLMs-based experiments conducted within this study often use heuristically designed prompts depending on the operational setup. This may pose limitations on the LLMs' answering capabilities. In addition to that, for almost all experimental studies conducted under this thesis, the object annotations that were available within the datasets were used to continue the evaluation processes. The quality of the available annotations could directly impose an internal threat to validity where poor quality annotations could lead to incorrect results which lower the overall performance of the experiments.

Construct Validity checks on the relation between the concepts behind the research and the observations recorded. For instance, the LLM performance on pedestrian detection and localization in paper C is dependent on the quality of the prompts and the choice of performance metrics used to report results. In cases such as pedestrian localization, metrics such as recall and accuracy may not impact much, whereas, metrics such as intersection over union would add more value to the recorded results by helping readers well understand the LLMs' capability in localizing the focused object. While these threats were reduced by selecting appropriate metrics for the analysis and by following prompt engineering techniques that optimized the selected prompt to yield better results, the mitigation strategies followed could be limited to the authors' knowledge. In addition to that, all LLM-related experiments including the types of potential hallucinations could have been impacted by edge cases and coarse granularity. The iterative experiment design processes and comprehensive validations have been applied to mitigate these threats.

External Validity underlines to what extent the results can be generalized outside the scope of the respective study. Paper A may hold an external threat due to the limitations and constraints considered while selecting databases and formulating the search queries. During the LLM-related experiments conducted under this thesis, we used different LLMs including proprietary models, to which the authors had no control. These industry-owned models may have undergone uncontrollable model updates unnoticeably during the experimentation and that could potentially affect the robustness and generalization of the results. In addition to that, ethical concerns, newly imposed rules and regulations related to AI [47], and frequent technical advancements in the field of LLMs could limit the generalizability of some experiments and their results. Furthermore, the datasets used in each paper have been selectively curated focusing on the problem domain and the motivation. Based on the experimental setup, some datasets have undergone image pre-processing steps which may have eliminated some parts from the images. Such curated datasets would affect the generalizability of the study compared to the use of unmodified images. Also when dealing with temporal data, the frequency of extracting images from the video data would have affected the results. Therefore further studies are required to understand the impact of such dataset-related characteristics. Moreover, the LLM capabilities in pedestrian detection and localization could change from other perception-related tasks or with different model architectures. To mitigate this threat, we have experimented with different LLM combinations with the intention of covering different model architectures. However future work in real-world validations is crucial for continuous adaptation in a rapidly advancing domain such as automotive to maintain the practicality of the research, even though the generalizability of the findings is limited.

By having addressed the identified threats to validity, this thesis presents a reliable and robust contribution to the automotive field, targeting connected and autonomous vehicles. The use of state-of-the-art AI tools in perception and monitoring tasks brings a new dimension to the research area while enhancing adaptability and efficiency in dynamic and complex environments. This opens new avenues of research that solve challenging problem domains such as resource utilization for resource-constrained, safety-critical systems.

1.7 Discussion

This section provides a discussion of the thesis findings concerning the research goals and research questions formulated at the beginning of the studies. This thesis contributes to a broader understanding of resource utilization in the automotive context by focusing on application scenarios related to perception systems. This thesis provides insights and novel methods for utilizing resources within advanced safety critical systems where state-of-the-art tools such as LLMs are proposed to be used as dialogue interfaces.

The findings of the systematic mapping study identify the key trends and research gaps in the domain of VANETs focusing on how vehicular networks can contribute towards the concept of resource sharing. The study reveals that topics related to networking, routing protocols, and security are extensively researched across the considered publication years, whereas real-world validation, task offloading, and authentication are some under-valued areas within the domain of VANETs. The findings of this study emphasize the significance of practical implementation-related research areas that will complement theoretical frameworks and real-world applications.

The research gaps and key trends identified within VANETs laid the foundation for the vision of the next explanatory study where the AirDnD framework was introduced for resource utilization. This innovative framework demonstrates how the proposed three-layered architecture supports the generation of a dynamic mesh network for task offloading among edge devices. This proposed concept helps in utilizing the network bandwidth to provide an improved, flexible model to fulfill the task offloading while considering data and network-specific properties. The subsequent experimental study explores the use of state-of-the-art LLMs as a tool that plays the role of dialogue interface within the resource utilization framework. This study infers that given the rapid growth of modern technological advancements, smart vehicles with relatively larger computing resources will be readily available within our society. As a result of that, it will only be a matter of time until OEMs incorporate foundational models within vehicles to provide better services to its passengers [20, 21]. Therefore, this experimental study demonstrates that the use of LLMs as a communication tool for resource utilization is a viable approach. To evaluate the performance of the proposed framework, perception and monitoring-related tasks such as pedestrian detection and localization were considered.

Along with the research avenue laid by the previous studies, a new set of challenges was introduced related to the trustworthiness of perception and monitoring systems. Therefore, the focus of the thesis was shifted to evaluating and enhancing the trustworthiness of LLM-supported perception-related systems. The following experimental studies identify the potential types of hallucinations that could occur in the automotive context focusing on perception tasks such as pedestrian detection. The related literature is thoroughly surveyed to understand existing hallucination detection and mitigation techniques that can be adopted into the automotive context following multimodal data. One study proposes a novel approach that exploits temporal data to detect hallucinations whereas the next study adopts a widely recognized text-based hallucination detection technique in the automotive context where it starts processing camera images recorded by front cameras. These studies immensely support ensuring the trustworthiness of LLM-supported perception systems and thereby contribute to reliable and safe advanced driver assistant systems.

1.8 Conclusions and Future Work

This research explores resource utilization in the automotive domain by examining VANETs and the application of LLMS within connected vehicles as dialogue interfaces to support resource sharing. A novel framework was introduced, demonstrating how connected vehicles can support resource utilization by sharing a dialogue between them that offloads computational tasks to nearby vehicles instead of transmitting raw data. To evaluate the proposed approach, a practical application scenario was referred to, focusing on perception and monitoring tasks including pedestrian detection and localization. The use of LLMs in the said resource utilization framework was validated through further research to enable communication between connected vehicles. The subsequent research not only showcased the potential of LLMs in supporting perception and monitoring tasks in a connected vehicular setup but also discovered challenges and limitations related to the trustworthiness of LLMs. To address the challenges, the research proposed hallucination detection and mitigation techniques that enhance the reliability of LLM-assisted perception-related tasks.

The proposed methodologies for resource utilization and LLM trustworthiness evaluation demonstrate where advanced AI-assisted technologies could support resource utilization in the automotive context focusing on connected vehicles. While the proposed novel approach could facilitate further significant improvements, it also opens up many other opportunities for future research. Based on this thesis' findings and limitations, several directions have been identified as future work.

Based on the succeeding step of the research work covered in this thesis, improving the LLM-assisted dialogue will be considered with the aim of ensuring seamless communication between heterogeneous edge devices. Future work would investigate and design unified protocols to be followed when prompts are shared between foundation models deployed within different edge devices. As a part of this research direction, different prompt patterns following novel prompt engineering techniques and RAG will be further investigated. Exploring lightweight, open-source model architectures that can be fine-tuned for multimodal automotive data would be the pinnacle of this proposed direction with the motivation of providing efficient, scalable, and reliable communication interfaces powered by state-of-the-art AI tools.

Furthermore, the continuous evolution of technologies will be tackled by actively participating in research that adapts to newer technological advancements such as emerging datasets, new model architectures, or training techniques. The main motivation is to explore the integration of next-generation AI models into the proposed AirDnD framework to leverage their capabilities and deliver improved performances. Adapting to new technologies and using state-of-the-art tools always comes with the responsibility of adhering to critical regulatory compliances. Therefore, further research on ethical and regulatory challenges is important to ensure fairness, accountability, and transparency in such AI-assisted systems.

The time-series information in the automotive context as introduced in paper D can be investigated to explore enhanced perception techniques that support solving the "looking around the corner" problem. Exploiting temporal data on different views of the same intersection with the aid of AI tools would help in identifying traffic patterns and environmental changes with an increased level of detail. This could potentially assist autonomous and connected vehicles to generate better context-aware decisions which ultimately helps in utilizing the resources. To ensure the trustworthiness of such systems, we can further extend SelfCheckGPT and apply it to detect and mitigate hallucinations.

Finally, strengthening security and privacy in LLM-assisted systems would be another vital and interesting research area that can be considered as future work of this thesis. While this could be a very broad and complex research area, the subdomains such as preserving data privacy, securing communication interfaces, ensuring ethical AI techniques, securing data sharing, and securing edge devices would be closely connected to the proposed resource utilization framework in the automotive context. Future work on these subdomains would increase the resilience of LLM-based systems against potential threats ensuring safe and reliable real-world adaptations.