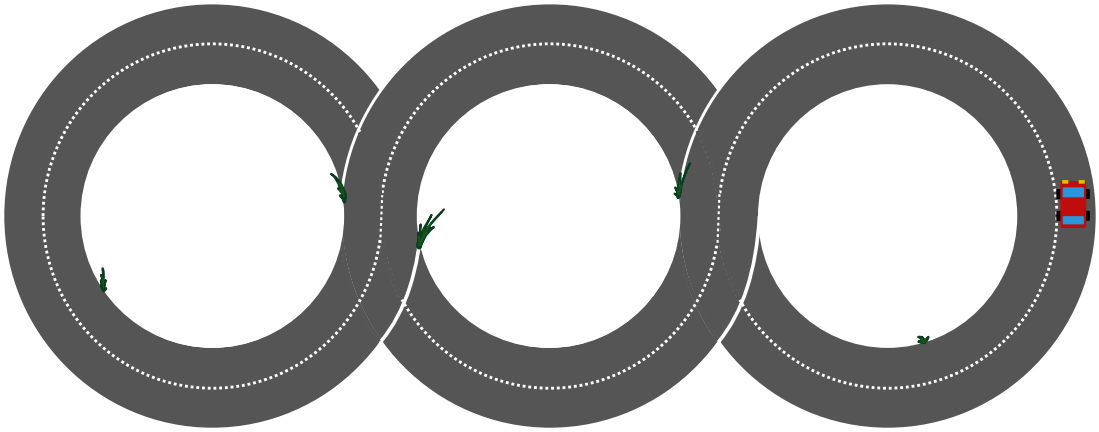

Safe AI Compliance in CI/CD for Software-Defined Vehicles

Researching Conformity to ISO/PAS 8800:2024

Master's Thesis in Computer Science and Engineering

Atosa Daneshvar-Minabi
Ann Heijkenskjöld

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG
Gothenburg, Sweden 2026

Master's Thesis 2026

Safe AI Compliance in CI/CD for Software-Defined Vehicles

Researching Conformity to ISO/PAS 8800:2024

Atosa Daneshvar-Minabi
Ann Heijkenskjöld

UNIVERSITY OF
GOTHENBURG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg
Gothenburg, Sweden 2026

Safe AI Compliance in CI/CD for Software-Defined Vehicles
Researching Conformity to ISO/PAS 8800:2024
Atosa Daneshvar-Minabi, Ann Heijkenskjöld

Supervisor: Beatriz Cabrero-Daniel, Department of Computer Science and Engineering
Supervisor: Mazen Mohamad, Department of Computer Science and Engineering
Industrial Supervisor: Peter Forsberg, IBM
Examiner: Christian Berger, Department of Computer Science and Engineering

Master's Thesis 2026
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in LaTeX
Gothenburg, Sweden 2026

Atosa Daneshvar-Minabi, Ann Heijkenskjöld
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg

Abstract

As software-defined vehicles (SDVs) are incorporating artificial intelligence (AI), ensuring AI safety presents significant challenges for development teams. This thesis addresses four research questions spanning the current application of *ISO/PAS 8800:2024 Road vehicles — Safety and artificial intelligence* in SDV development, the compliance challenges teams face, where in the Continuous Integration/Continuous Delivery (CI/CD) challenges arise, and how to mitigate them. A two-phase qualitative interview study was conducted with industry professionals, including domain experts and researchers. Through semi-structured interviews, recommendable actions were produced and evaluated in the form of artefacts and a set of guidelines. The results reveal that while some individual practitioners are aware of ISO/PAS 8800:2024, formalised organisational conformance frameworks are largely absent, and a variety of different conformance issues surface throughout the entire CI/CD process rather than at any single stage. Artefacts were developed for supporting conformance to ISO/PAS 8800:2024 within existing development practices. The first artefact guides teams on where human oversight is required, where AI could be helpful, and where AI automation is feasible across the AI safety lifecycle, including the trade-offs and acceptance of AI automation at each stage. The second artefact shows which selected clauses of the ISO/PAS 8800:2024 are most prominent at each phase of the MLOps lifecycle. Integrating ISO/PAS 8800:2024 into CI/CD requires conformance to be treated as an inherent part of development rather than an after-the-fact activity. Insights from this study could be valuable for organisations integrating ISO/PAS 8800:2024 into their development workflows.

Keywords: ISO/PAS 8800:2024, Software-Defined Vehicles, Safety Compliance, AI Safety, ISO 26262, DevOps, MLOps, CI/CD

Acknowledgements

While writing this master's thesis at Chalmers University of Technology in partnership with IBM, we reflect on how the completion of this study could not have been possible without the expertise surrounding us. Therefore, we would like to express our sincerest thanks to our wonderful supervisors Beatriz Cabrero-Daniel and Mazen Mohamad at Chalmers University of Technology, and Peter Forsberg at IBM Svenska AB, who not only have guided us throughout this process, but have also used their network to provide valuable connections for our study.

We moreover want to express our gratitude to our examiner, Christian Berger, for providing valuable feedback to us.

The results are built on insights gained from interviewees who voluntarily have taken the time of their day to participate in our study. We express a heartfelt thank you to each interviewee.

Lastly, we want to extend our thanks to all the colleagues at IBM for welcoming us so warmly and showing an eagerness to help us with both material and valuable contacts. Our time at IBM has truly been an enriching and fun experience. Thank you, Birgitta Krafft, for signing up to be our manager at the IBM Gothenburg office, and for showing up and staying in touch with us. We cherished every lunch break with you.

Atosa Daneshvar-Minabi, Ann Heijkenskjöld, Gothenburg, June 2026

Contents

List of Figures

List of Tables

1

Introduction

Technology is constantly evolving in the automotive industry, and the emergence of software-defined vehicles (SDVs) has become a significant area. SDVs are vehicles in which mainly software, as opposed to hardware in traditional vehicles, is responsible for how the vehicle functions, as well as how well it performs and how the user experiences it.

Safety has always been of a high priority in the automotive domain, however, SDVs introduce new challenges by relying on AI and complex software for core functions. This requires safety considerations to expand beyond traditional approaches as new technologies emerge. Consequently, organizations exist to standardise these technologies, such as the International Organization for Standardization (ISO). In addition to developing full International Standards, ISO also publishes other types of standardization deliverables, such as Publicly Available Specifications (ISO/PAS). These are produced faster than regular ISOs for quickly developing technologies such as artificial intelligence (AI). ISO/PAS documents still have to be purchased though, and are not freely available for anyone to access. ISO/PAS 8800 is one of them, which regards safety of AI in road vehicles [1].

Assuring conformance with such safety specifications, however, must be reconciled with how modern automotive software is actually built. A heavily adopted modern software development practice in the industry is DevOps, with high and elite performers representing two thirds of organizations globally [2, p. 11]. It extends the Agile software development methodology and supports faster and smaller software updates through automation and continuous feedback. It consists of a set of development and IT operation practices that promote collaboration between the two groups, development and operations, that traditionally work separately in distinct silos. These practices endorse automation, shared responsibility and continuous feedback across the software lifecycle. The DevOps method Continuous Integration and Continuous Delivery (CI/CD) automatises build, test, and delivery processes and is largely used as a method to integrate, test and deliver software changes continuously [3].

Yet the pace that CI/CD enables is at tension with how safety standards are conventionally applied. Safety standards are widely used and applied to the development sequentially, especially during the testing phase of verification and validation to

ensure conformance. This, in turn, can help achieve compliance with regulatory requirements. Safety risks can be reduced by running checks from the framework, however, analysing completed development will require time and effort to make modifications. Identifying software issues, such as non-compliance issues, late in the lifecycle process requires rework, with the cost of fixing a mistake increasing significantly the later in the process it is discovered [4, p. 40]. In an environment where DevOps and CI/CD are central and prioritised, this study aims to investigate how continuous development in SDVs can be integrated with conformance to ISO/-PAS 8800 as an inherent part of the development process rather than treated as an after-the-fact activity, and how AI tools can contribute throughout the process.

1.1 Problem Description

The growing popularity of DevOps and CI/CD has been fueled by the industry's need for faster development, quality assurance and reliable products. It allows teams to deploy incrementally while prioritising iterative and fast work and is therefore a modern software development practice in areas such as the automotive industry. Simultaneously, AI is being increasingly and widely adopted in SDVs, especially in the area of Advanced Driver Assistance Systems (ADAS) and Autonomous Driving (AD) where it can be used in, for instance, perception and decision-making with Deep Neural Networks [5].

ISO/PAS 8800 is a Publicly Available Specification (ISO/PAS). This means that it was developed faster than a full ISO standard to address the rapidly evolving field of AI safety in road vehicles, serving as an early framework in an area where formal standardization is still developing. It regards safety risks and inadequacies in AI technology usage for road vehicles containing one or more Electrical/Electronic systems (E/E systems). These are complex systems that are responsible for safety functions among other things. Conformance to this specification typically requires rigorous documentation through evidence. The issue is that in the automotive industry, retrofitting compliance is a weak point. DevOps is an iterative practice, however, in many cases, compliance is treated as an afterthought and often checked in retrospect [6]. Retrofitting requires correction through reconstruction and modification on finished work, thus resulting in an inefficient use of both resources and time.

Standards such as ISO 26262 (FuSa) and ISO 21448 (SOTIF) address the functional and systematic safety of deterministic software in road vehicles, but do not adequately cover the AI aspect that SDVs increasingly rely on because AI started being adopted in the industry after these standards were developed. This new development shifts vehicle behaviour from no longer being defined by explicit rules, to being defined by probabilistic models learned from data. Traditional standards assume a deterministic relationship between code and behaviour, which is incompatible with the black-box nature of Deep Learning (DL). SOTIF covers some performance limitations in AI as functional insufficiencies, however, it is not the main focus and does not provide a framework for the AI lifecycle (for instance, data quality assurance).

ISO/PAS 8800 specifically addresses the probabilistic and data-centric nature of AI, as well as its risks, in automotive vehicles.

Retrofitting ISO/PAS 8800 compliance into an SDV development process also entails AI-specific challenges. For example, safety in AI requires training and validation datasets. If reworked after failing to fulfil criteria, collection of new data as well as retraining and reverification of the model will require a large amount of time and resources in comparison to rewriting code.

One of the issues ISO/PAS 8800 addresses is output insufficiency, where AI produces dangerous incorrect outputs not because of a software bug, but because of limitations in how it was trained or specified. This is a priority when implementing safe AI in road vehicles and needs to be taken into consideration during implementation through safety metrics and thoroughly planned architecture. Developers cannot afford to take output insufficiency as an afterthought, as adding architectural safety mechanisms upon realization of output insufficiency is risky to the driver and other road users.

Furthermore, there is currently a knowledge gap in defining and exploring frameworks for practically integrating ISO/PAS 8800 into development processes. It was published in 2024 and is therefore considered novel as of writing this thesis, and empirical studies that demonstrate successful implementation into development pipelines are scarce. ISO/PAS 8800 also presents requirements, meaning "What needs to be achieved", and standards rarely dictate precisely how to implement the requirements in workflows.

1.2 Purpose of the Study

The purpose of this study is to investigate concurrently integrating conformance with ISO/PAS 8800 into the DevOps CI/CD practice of development within SDVs.

It would be more effective if conformance with safety regulations were integrated with development as an inherent part of the process, rather than treated as an after-the-fact activity. The results of the study take the form of recommendable actions, including two research artefacts which were subsequently evaluated through a second phase of industry interviews. The results are beneficial for both efficiency and in reducing time and cost resources when developing software for SDVs.

1.3 Significance of the Study

Since 2016, there have been documented injuries and fatalities involving ADAS/AD systems [7, 8, 9]. As AI capabilities advance, ADAS and autonomous driving systems are increasingly powered by AI-driven perception and decision-making. Real-world incidents have demonstrated the consequences of insufficiently assured AI systems, with notably a 2023 incident involving a robotaxi company in San Francisco where the technical investigation found no sources of hardware or software failure [10]. This

is the type of AI-specific functional insufficiency that ISO/PAS 8800 is intended to address, and which traditional functional safety frameworks such as FuSa were not originally designed to cover.

Findings from this study indicate that while individual practitioners are often familiar with the specification, organizations generally lack established frameworks for conformance. This may partly be because no legal mandate for compliance to ISO/PAS 8800 currently exists, but also because the standard itself is still maturing. This creates a gap where technical implementation is advancing faster than regulatory guidance. Addressing this gap is important because ISO/PAS specifications are often precursors to ISO standards. Organizations that begin establishing conformance practices early may therefore be better positioned if compliance becomes mandatory or industry-expected.

By producing practitioner-validated recommendable actions for integrating ISO/-PAS 8800 into CI/CD practices this study contributes to current safety practices and to the broader challenge of aligning fast-moving AI development with emerging regulatory expectations.

The study is also significant for the automotive industry because it proposes a method to reduce the rework costs associated with late-stage conformance checks. While the results of this study are grounded in the SDV domain, the methodology may be transferable to other safety-critical domains undergoing a similar transformation toward AI adoption, such as aviation or medical devices.

1.4 Research Questions (RQs)

The research questions (RQs) are divided into two parts: evaluating standard compliance issues (RQ1-3) and evaluating artefacts (RQ4).

Research Question 1 (RQ1): How is ISO/PAS 8800 currently applied in the software development of SDVs?
RQ1 justification: This RQ seeks to understand how industry professionals currently interpret and apply ISO/PAS 8800. As the specification is new, integration within organisations is likely not yet well established.

Research Question 2 (RQ2): What are the specific challenges development teams face regarding safety regulation compliance when developing software for SDVs?
RQ2 justification: Understanding precisely what the challenges are is a prerequisite for placing them on a timeline and addressing them, which is the responsibility of RQ3-4.

Research Question 3 (RQ3): Where in the DevOps CI/CD lifecycle are the aforementioned challenges?
RQ3 justification: Identifying where in the development lifecycle problems arise is necessary to understand where progress slows down and where solutions need to be implemented.

Research Question 4 (RQ4): How can development teams mitigate these challenges when integrating ISO/PAS 8800 compliance into CI/CD processes?
RQ4 justification: A solution to integrating ISO/PAS 8800 as part of the development process rather than as an after-the-fact consideration cannot be proposed without first understanding current practice, challenges, and their timing. Answers to RQ1-3 help conceptualise the answer to RQ4, resulting in the artefacts that are further evaluated.

1.5 Contributions

This study contributes to the knowledge of how ISO/PAS 8800 can be integrated into DevOps CI/CD workflows for software-defined vehicles. Specifically, the following contributions are made:

First, an understanding of how industry practitioners currently interpret and apply ISO/PAS 8800 in software-defined vehicle development is provided, addressing RQ1. As the standard is newly introduced in 2024, the findings in Chapter 5 show that there is some individual familiarity with the standard among practitioners, yet formalised organisational conformance frameworks are largely absent.

Second, an overview of the compliance challenges development teams face is also presented, along with an identification of where in the CI/CD lifecycle these challenges arise, addressing RQ2 and RQ3. The findings reveal that a variety of compliance issues are fragmented across the entire development process. This can be found in Chapter 5 and Appendix B.

Third, two artefacts are proposed and evaluated in response to RQ4. The first artefact overlays five prioritised clauses of ISO/PAS 8800 onto the AI safety lifecycle, evaluating the appropriate level of human oversight alongside the assessed degree of AI helpfulness and AI automation. The second artefact provides a complementary MLOps perspective, showing where each of the five clauses is prominent in the phases of the MLOps lifecycle. It is presented both as a numerical table intended as input for AI tools and as a visual representation for human readers. These can be found in Section 5.2.2 and in Appendix B.

Together, these contributions offer organisations concrete guidance for integrating ISO/PAS 8800 conformance continuously into their development process.

2

Background

To understand the basis of this study, background information on standard conformance relevant to AI safety in automotive development is presented. This chapter establishes the concepts that the study builds upon. The sections cover software-defined vehicles as the application domain, relevant development workflows such as DevOps and MLOps, and the relevant standards to provide regulatory context.

2.1 Software-Defined Vehicles (SDV)

Software-Defined Vehicles (SDVs) are vehicles in which the main functions and features are managed by software rather than hardware, hence the phrase "Software-defined". Traditionally, automotive vehicles are more hardware-oriented. They consist of many ECUs (Electronic Control Units) that control individual functions, for instance the brake light or airbag control system. New functionality in a traditional vehicle would introduce more ECUs, increasing complexity to the overall system.

SDVs rely on a centralised computing architecture and can be updated over-the-air (OTA), resulting in a significant shift in maintenance approach compared to its traditional counterpart [11]. Functions are integrated in central, zonal computers instead of several ECUs. This requires High-Performance Computers (HPC) which can handle cross-domain vehicle functions. The centralised computing architecture facilitates inter-system coordination and data processing [12].

Improvements and customisations can be updated remotely in the software as opposed to going to a physical dealership or purchasing a new car model. Traditional Original Equipment Manufacturers (OEMs) are also transitioning to software-first platforms, as SDVs can introduce new revenue streams in post-sale subscriptions [12].

SDV is an umbrella term and contains the vehicle subcategories autonomous vehicles (AVs) and connected vehicles [12]. AVs are a category of SDVs which operate by themselves with the help of technologies such as sensors, cameras, and, most importantly, AI. This is an illustration of how the automotive field is shifting more toward AI. Faulty decisions made by AI in AVs are critical to mitigate since they can cause serious physical harm to people, making it an important area where ISO/PAS

8800 can be applied. The AV industry is growing and there are companies that currently operate AVs for robo-taxi services such as Waymo in the US and Apollo Go in China [13, 14]. There are also companies such as Tesla that are developing increasingly autonomous features for private individuals [15].

2.2 DevOps

DevOps is the combination of Development and Operations. It may be seen as a way of working, where software development and operations are combined into one continuous process, whereas traditionally, development and operations teams worked separately. It can be described as a set of practices and a mindset that make software development and operations more efficient and reliable. The key principles are collaboration, automation, continuous improvement, and fast feedback when issues are detected. The DevOps lifecycle is iterative and consists of eight stages, as illustrated in Figure 2.1 [16].

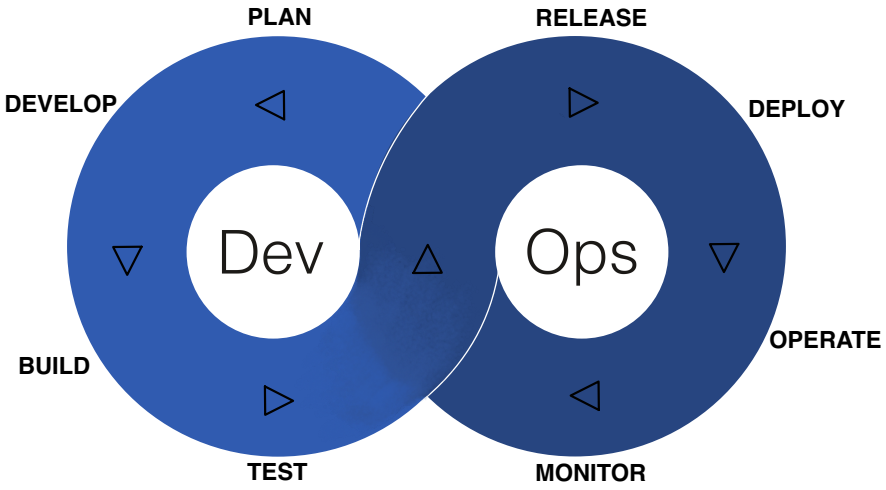

Figure 2.1: The DevOps lifecycle, consisting of the eight stages plan, develop, build, test, release, deploy, operate, and monitor.

2.2.1 MLOps

MLOps is a framework for building and running Machine Learning (ML) models. It describes how to apply DevOps procedures to ML practices, streamlining the automation of tasks and deployment of models [17, 18]. Additionally, MLOps accomplishes this by incorporating methodology from CI/CD to establish a production line for each stage of the development of an ML product. A simplified MLOps lifecycle with its 9 phases is illustrated in Figure 2.2 [16, 19, 20].

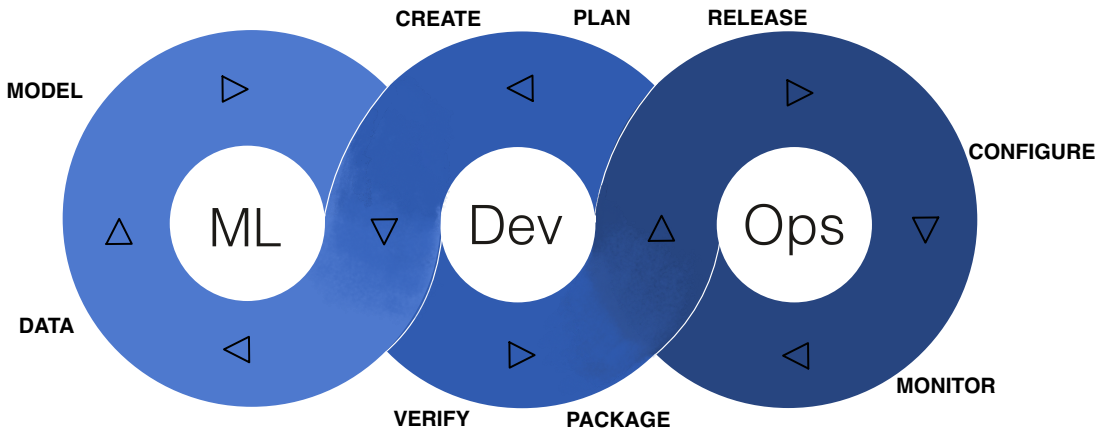

Figure 2.2: The MLOps lifecycle, consisting of the phases plan, create, data, model, verify, package, release, configure, and monitor.

2.2.2 Continuous Integration and Continuous Delivery (CI/CD)

CI/CD is a DevOps methodology consisting of Continuous Integration (CI) and Continuous Delivery (CD). Together, they aim to get software changes into production safely, quickly, and sustainably, by automating the process of integrating, testing, and deploying code. This process enables software developers and teams to produce software releases more frequently and more reliably [21].

CI refers to the process of continuously integrating new code with existing code by frequently merging code changes into a shared repository, where automated builds and tests are triggered to detect errors and conflicts early. This ensures that the existing codebase does not break by code changes. Version control systems such as Git are used to manage the shared repository, while CI platforms such as Jenkins automatically trigger builds and tests upon each code change [22, 23]. Testing frameworks such as JUnit are commonly used to validate the code throughout this process.

CD most commonly stands for Continuous Delivery, but sometimes also Continuous Deployment, depending on the context. They are related but distinct concepts. Continuous delivery means that code is automatically validated and kept in a releasable state, but it still requires manual approval before deploying to production. Continuous deployment refers to automatically deploying software updates. It is an extra step of automation. Tools such as Docker and Kubernetes are commonly used to facilitate this automated deployment process [24, 25]. Many automotive organisations have continuous delivery but do not adopt continuous deployment, as state of practice for dependable systems has been observed to remain more traditional than in non-safety-critical domains [26].

2.2.3 Human-in-the-Loop (HITL) and Human-on-the-Loop (HOTL)

Human-in-the-loop (HITL) refers to the act of human participation in an automated system. There are several reasons for having a HITL. Usually, the human intervention acts as a safeguard that ensures accuracy, safety, accountability or ethics in AI workflows [27]. Automated tasks may be effective, but they can introduce insufficient output, ambiguity, bias or edge cases that deviate from the training data [27]. This limitation of automation in safety-critical domains can have dire consequences. By inserting human oversight into the AI workflow with feedback-based interaction between AI systems and humans, this risk can be minimised. Human-on-the-loop (HOTL) is another form of human AI governance. The distinction lies in the degree of involvement, where a HOTL would monitor and intervene if necessary. HITL and HOTL are especially relevant in the context of regulatory compliance and safety-critical systems. The EU AI Act's Article 14, for instance, mandates that

> *"High-risk AI systems shall be designed and developed in such a way, including with appropriate human-machine interface tools, that they can be effectively overseen by natural persons during the period in which they are in use."* [28]

Embedding human oversight into AI workflows mitigates the black-box effect in non-deterministic systems, supporting transparency and explainability commonly required by safety audits.

2.3 Compliance, Conformance and Standards

Compliance and conformance are terms that are sometimes used interchangeably, but they do not mean the same thing. Compliance refers to adhering to rules or regulations imposed by an external authority, such as a government body. Its focus is on meeting legally or formally enforced obligations in order to avoid consequences such as fines or legal action. Compliance is most often mandatory.

Conformance, on the other hand, refers to adhering to a specific set of guidelines, standards, or specifications. Unlike compliance, conformance is not always mandatory. Many specifications and also standards are voluntary and followed because they ensure best practices. That said, the two concepts can intersect. When a standard is adopted into law or regulation, conformance to that standard becomes a condition of compliance.

Standards are generally voluntary guidelines, however, as aforementioned, public authorities can refer to or mandate them through legislation, making it necessary for businesses to comply in order to function or sell their products [29]. Standard conformance helps to assure that the products are of a certain quality whilst also ensuring safety and reliability [30]. This drives consumer trust and can be a competitive advantage for the firms who choose to adhere to it. The involvement of globally established experts within their respective industry [30] reflects on the

standard itself, it contains the most current and agreed-upon technical knowledge and innovation, and makes for a state-of-the-art solution.

2.3.1 Normative Requirements

In specifications and standards, normative requirements are mandatory requirements that need to be fulfilled in order to claim conformity to a specification or compliance to a standard [1, 31]. All of the normative requirements together define the acceptable criteria that a product for instance must meet. In order to pass conformity assessments and audits they all need to be fulfilled. Failure to meet normative requirements means that conformance with the standard cannot be claimed. In standards documents, normative requirements are typically identified by the word *shall* [32].

2.3.2 Recommendations

Compared to normative requirements, recommendations provide guidance. They can add context, or examples to aid in the understanding and implementation of a standard, but they do not impose mandatory requirements. As such, recommendations are typically identified by the word *should* [32]. This indicates that it is recommended but not mandatory in order to pass conformity assessments or audits. While not required for compliance, informative content can be valuable in helping developers and engineers follow best practices.

2.4 ISO and ISO/PAS

A standard is a jointly agreed solution to a frequent problem spanning a broad range of industries [29]. ISO is an independent, non-governmental organization based in Switzerland that develops standards in many different domains [33]. The purpose of these is to ensure that products are developed by globally accepted best practices agreed upon by experts in their respective fields. As discussed in Section 2.3, these standards are generally voluntary but can be mandated by regulatory bodies [29]. An ISO/PAS is often a precursor to a formal ISO standard. An ISO/PAS is developed quickly for fast-developing markets and can evolve into a full ISO standard after iterations of feedback.

2.4.1 ISO 26262:2018 (FuSa)

ISO 26262:2018 (FuSa) [34] covers the functional safety aspect of electrical and/or electronic (E/E) systems in road vehicles. FuSa was originally published in 2011 and further revised in 2018. The purpose of this standard is to minimise possible hazards stemming from system malfunctions from for instance hardware or software bugs. It is also primarily intended to guide company-internal development processes. It is a voluntary standard, meaning there is no official legal requirement for compliance, though it is widely adopted as industry best practice in the automotive sector and often expected throughout the automotive supply chain.

The standard consists of twelve parts, ten of which are normative. They cover the entire development process using the V-Model, and help determine risk classes through a scheme called Automotive Safety Integrity Level (ASIL). Another important part of this standard is the Hazard Analysis and Risk Assessment (HARA) which is a method that defines safety goals and identifies potential hazards.

2.4.2 ISO 21448:2022 (SOTIF)

ISO 21448:2022 (SOTIF) [35], commonly referred to as SOTIF, revolves around safety of the intended functionality. Unlike ISO 26262, this standard aims to minimise the occurrence of hazards caused by functional insufficiencies, that is without a system failure, such as sensor failures. It covers performance and specification insufficiencies in both the intended functionality as well as the implementation of E/E elements in the system. It also handles "unknown unknowns", which are unforeseen hazardous scenarios formed by information gaps and lack of awareness. SOTIF also identifies triggering conditions that initiate a system reaction that can lead to unsafe behaviour. The process iterates through the three phases design, verification and validation.

2.4.3 ISO/PAS 8800:2024

ISO/PAS 8800:2024 [1] is a Publicly Available Specification published in 2024 regarding functional safety for AI. It addresses the safety risks arising from output insufficiencies and systematic errors in AI elements, as well as random hardware errors in the components running them. It additionally defines a framework for management through specific methods for data quality, model training and performance validation. The aim is to provide guidance on the usage of safety-related functions. ISO/PAS 8800 complements FuSa and SOTIF by extending their approaches with AI-specific requirements. What sets it apart is its definition of data quality-specific requirements and the techniques to verify AI robustness.

A clause is a numbered section in ISO/PAS 8800. It is structured into 15 clauses, of which clauses 7 to 15 contain the normative requirements. There are 49 normative requirements. At the end of these 9 clauses, there are work products that should be produced in order to be compliant to the specification.

Clauses 1-6 are not part of the AI Safety Lifecycle. They are about topics such as the document's scope and abbreviated terms. Clauses 7-15 are the technical clauses that make up the AI safety lifecycle. Each clause addresses a distinct phase or activity. One example of a technical clause is Clause 7 which is about AI safety management. It involves defining the AI safety lifecycle and its activities, as well as making a safety plan. A close replication of the AI safety lifecycle to the reference in ISO/PAS 8800 can be seen in Figure 2.3.

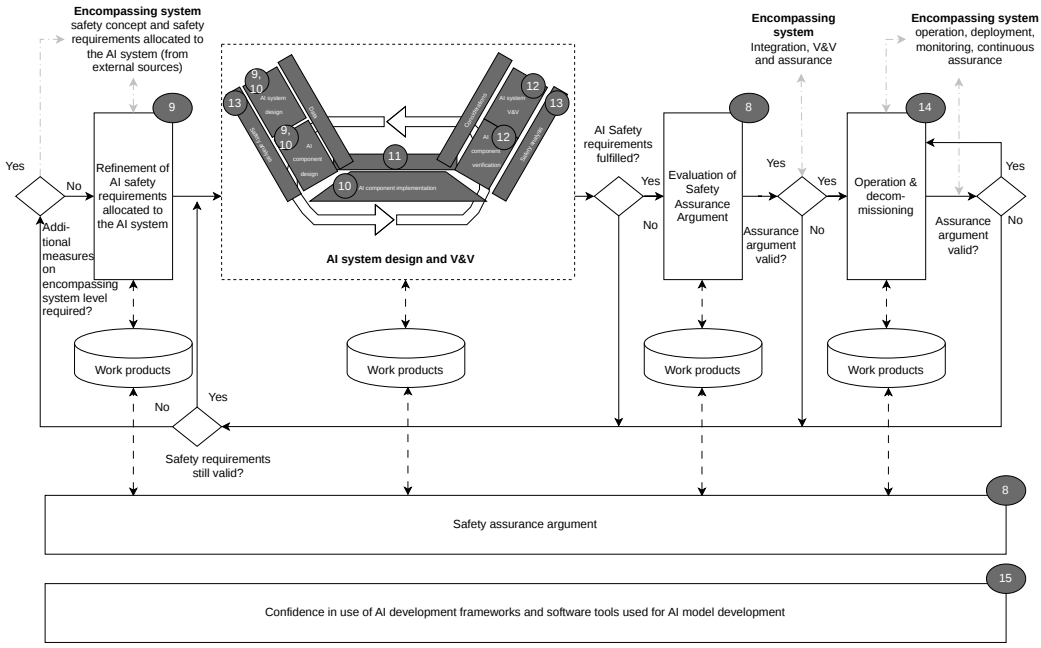

Figure 2.3: The AI safety lifecycle, adapted from the example contained in ISO/-PAS 8800, Clause 7 [1].

As Figure 2.3 shows, the lifecycle begins with safety requirements allocated from external sources, progressing through AI system design and V&V in Clauses 9–13, before concluding with operation and decommissioning in Clause 14. The safety assurance argument, covered in Clause 8, runs throughout the entire process and is evaluated at key transition points. Work products are produced at each stage as evidence of conformance, and feedback loops ensure that safety requirements are re-evaluated if the assurance argument is found to be invalid at any point.

3

Related Work

This chapter reviews prior research across three areas: adopting DevOps and continuous engineering in automotive development, AI safety assurance in automotive systems, and recent efforts to operationalise ISO/PAS 8800:2024 as a lifecycle framework. Research on investigating specifically conformance challenges for ISO/PAS 8800:2024 within CI/CD contexts is sparse.

3.1 DevOps in Automotive Development

Studies have investigated the challenges of applying DevOps and continuous engineering practices to safety-critical automotive software development. Nouri et al. (2022) report on challenges identified through interviews with domain experts, approaching this from an industrial practitioner perspective [36]. Their study identifies a variety of safety-related challenges as time-consuming and resource-intense, including impact analysis, conducting Hazard Analysis and Risk Assessment (HARA), safety argumentation, and verification and validation, which span all aspects of product development. Notably, the authors acknowledge a direct dependency between growing system complexity and the effort required to conduct these analyses, and explicitly call for more agile approaches during development and operations as a necessity. Simultaneously, they identify the tension that HARA is not easily compatible with agile processes. Building on these identified challenges, Nouri et al. (2025) present a systematic literature review on safe DevOps for autonomous driving, introducing DevSafeOps as a framework for integrating safety activities into DevOps iterative loops [6]. Beyond confirming challenges from the earlier work, the review identifies additional challenges such as the separation of safety and software workflows, and offers solutions to identified challenges. One of these solutions is AI-based development assistance for tasks such as safety assessment and HARA, however, human oversight remains essential.

Complementing the examination of these challenges, Sandgren and Antinyan address the challenge of conducting software safety analysis on a continuous basis within agile development [37]. Specifically, in the context of conformance to FuSa at Volvo Cars. They propose a two-phase software safety analysis method intended to be possible within agile development cycles, by combining automated checks with workshop-based architectural analysis. This hybrid method demonstrates that inte-

grating safety analysis into iterative development is possible. However, the approach assumes a Simulink-based development environment and relies on manual activities as well as automation. Their work addresses functional safety rather than AI safety, which leaves open the question of how comparable hybrid approaches could be made for the concerns of ISO/PAS 8800:2024.

Kugele and Broy (2022) argue that architecture spanning service, data, and technical aspects must serve as a stable backbone for agile and DevOps-oriented automotive development [38]. They propose that such an architectural backbone enables deployability, modifiability, testability, and monitorability as prerequisites for safe continuous delivery, offering a concrete architectural answer to the agility that Nouri et al. identify as necessary [6, 36]. While short development cycles are identified as essential for remaining competitive, the authors note that adopting agile practices requires a substantial cultural shift, which they identify as often the most challenging aspect of DevOps adoption in the automotive industry. The authors also acknowledge that homologation, the certification process confirming that a vehicle complies with local standards and regulations, remains an open issue when software changes occur within the vehicle.

Extending this line of argument to the vehicle level, Resing (2024) examines the practical challenges of implementing DevOps practices across the full development and operations chain of SDVs, from code authoring through to deployment [39]. Resing argues that a fundamental organisational transformation toward agile development is necessitated. Traditional automotive engineering is structured around siloed teams organised by vehicle domain or individual ECUs, each operating largely independently of one another. In this view, simply overlaying agile practices onto existing organisational structures is insufficient. Resing further cautions that the adoption of agile methods and organisational patterns in automotive contexts warrants careful investigation. Resing additionally identifies that the lack of homogeneous tooling across vehicle components leads to inefficient workarounds and the need for adapters to bridge incompatible systems. He also identifies that integrating Hardware-in-the-Loop (HiL) and Software-in-the-Loop (SiL) testing into a continuous pipeline is difficult. It requires manual steps that interrupt automation in deployment pipelines, causing bottlenecks.

Taken together, these works establish that the organisational, architectural, and tooling challenges of integrating continuous engineering practices into safety-critical automotive development are well recognised, yet remain difficult to resolve in practice. None of the works address how conformance to an AI-specific standard such as ISO/PAS 8800:2024 should be embedded into DevOps CI/CD practices, even though Nouri et al. did identify AI-based assistance for safety tasks as a potential solution [6]. This more specific question requires first understanding the conceptual foundations of AI safety assurance in automotive systems, to which the following section turns.

3.2 AI Safety in Automotive Systems

A foundational line of work on safety for machine learning (ML) in highly automated driving begins with Burton et al. (2017), who argue that adherence to ISO 26262 alone is insufficient when the functionality of a system is embedded in high-dimensional learned parameters rather than in explicit code [40]. To address this gap, the authors propose an assurance-case approach in which safety is argued from first principles by decomposing system-level goals into technical performance requirements. These requirements are structured around causes of functional insufficiencies such as inadequate training data coverage. This line of work would later inform the development of ISO/PAS 8800:2024, for which Burton chairs the ISO working group responsible for ISO/PAS 8800:2024 [41].

Building directly on this work, Gauerhof et al. (2018) apply the assurance-case approach to a case study of pedestrian detection [42]. The authors emphasise that industry consensus on a sufficient "weight of evidence" has not been reached. Adopting a broader perspective, Ashmore et al. (2021) survey assurance methods across the entire ML lifecycle, decomposed into the stages Data Management, Model Learning, Model Verification, and Model Deployment [43]. The authors conclude that assurance methods are unevenly distributed across the lifecycle, with the Model Deployment stage least mature and the continuous updating of deployed models identified as a particularly open challenge.

Taken together, these works establish the conceptual foundation for AI safety arguments in automotive systems, but all of them predate ISO/PAS 8800:2024. How the standard's requirements integrate with CI/CD workflows, and where AI can support this integration, remains the open question this thesis addresses.

3.3 ISO/PAS 8800:2024 Frameworks

Two recent studies are closely related to this study in that both propose lifecycle-level frameworks intended to operationalise ISO/PAS 8800:2024. Zafar et al. (2024) present a unified AI-product lifecycle that integrates ISO 26262, ISO 21448, and ISO/PAS 8800:2024 together within a V-cycle representation [44]. The authors argue that the absence of an integrated view leads to fragmented assurance, redundant work products, and slower time-to-market. Notably, the authors acknowledge that ISO/PAS 8800:2024 provides only high-level guidance and lacks specific, actionable activities.

Mohamed and Aslan (2025) offer a complementary perspective. The authors argue that ISO/PAS 8800:2024 offers an AI-specific safety foundation but that its practical application remains underdeveloped. They propose an integrated AI safety framework that decomposes assurance into five iterative stages: defining AI safety requirements, designing the system architecture and managing data, verifying and validating the AI system, compiling the safety assurance argument, and deploying, monitoring, and continuously improving the system [5]. However, the framework is

presented at a conceptual level, and the authors acknowledge that future work is needed to validate its effectiveness through industrial case studies.

Both works approach the standard primarily from a V-model or staged-lifecycle perspective and address neither how conformance activities should be integrated into DevOps CI/CD practices, nor where AI itself can support the conformance work.

3.4 Research Gap

Studies on DevOps in automotive development, reviewed in Section 3.1, establish that integrating continuous engineering practices into safety-critical development is both necessary and difficult, but their focus is not specifically on AI-specific standards such as ISO/PAS 8800:2024. The AI safety assurance works reviewed in Section 3.2 provide the conceptual foundations upon which ISO/PAS 8800:2024 was built, yet all predate the standard and do not consider how the assurance activities of ISO/PAS 8800:2024 can be embedded into DevOps CI/CD practices or how AI tooling could help in the process.

Taken together, the two studies in Section 3.3 confirm that the application of ISO/-PAS 8800:2024 is recognised as an open problem, and both approach it from a V-model perspective rather than a DevOps CI/CD one. This thesis addresses that specific gap by examining ISO/PAS 8800:2024 conformance within the DevOps CI/CD context of software-defined vehicles, and by proposing artefacts that map clauses onto the MLOps lifecycle alongside the appropriate degree of human oversight, AI helpfulness, and AI automation.

4

Methods

The study was conducted through qualitative interviews, namely semi-structured interviews with industry professionals. The interview study consisted of two phases, where the first derived a codebook found in Appendix C from twelve transcriptions to answer most importantly RQ1-3, but also RQ4. Empirical data collection was gathered through the interviews and analysed with thematic analysis. Inductive coding of the transcripts identified codes contained in codebooks, respectively for the first and second phase of interviews. These were used to identify themes to encapsulate interesting stories representative of the participants' experiences and insights. Two artefacts as part of the recommendable actions to answer RQ4 were created after the thematic analysis of the first phase. The purpose of the second phase was to evaluate the artefacts and further explore the themes from Phase 1 connected to RQ4. The artefacts were then refined and presented in Chapter 5. The research process can be seen in Figure 4.1.

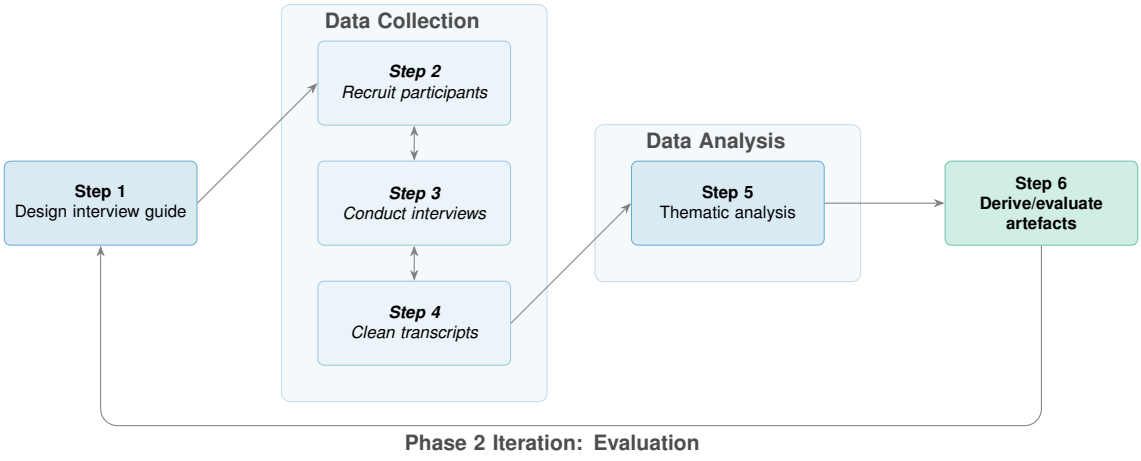

Figure 4.1: Overview of the research process. The interview guide (Step 1) initiates parallel recruitment, which then runs in parallel with conducting interviews and transcript cleaning (Steps 2 - 4), followed by thematic analysis (Step 5) and artefact derivation (Step 6), before starting a new evaluative iteration for Phase two.

4.1 Scope

This study focused primarily on ISO/PAS 8800. Related automotive safety standards, such as ISO 26262 and ISO 21448, were not the primary subject of investigation, but were drawn upon as contextual reference points and proxies given their close relationship to ISO/PAS 8800 and their established presence in industry practice. Methodologically, the study was limited to qualitative semi-structured interviews. The findings are therefore based on participants' reported experiences rather than direct observations of development practices.

4.2 Research Process

The methods and procedures are presented chronologically. The study was conducted in two phases, each following the same six steps with their respective adaptations. Phase 1 derived a codebook from twelve transcriptions to answer RQ1-4 and produced two artefacts as part of the recommendable actions. Phase 2 evaluated the two artefacts with six new participants, following the same process.

Step 1: An interview guide was designed. Semi-structured questions were constructed that contributed to answering RQ1-4.

Phase 1: The goal of the interviews was to elicit participants' experience with ISO/PAS 8800, find challenges related to safety regulation compliance, and where such challenges arise within the DevOps CI/CD lifecycle.
Phase 2: The guide aimed at evaluating the artefacts produced in Step 6 of Phase 1.

Step 2: Participants were recruited. They were intentionally selected based on their roles in the SDV domain or other relevant expertise. These individuals held positions such as system safety engineer, functional safety engineer, and AI researcher, see Table 4.1. This step was done in parallel with Steps 3, 4 and 5.

Phase 1: Twelve interviewees were recruited.
Phase 2: Six new industry practitioners were recruited.

Step 3: Qualitative semi-structured interviews were conducted based on the developed interview guide with the recruited participants on a Microsoft Teams video call. Although some interviews lasted longer, the interview had a set time limit of roughly 30 to 45 minutes. The conversations were transcribed using Microsoft Teams Transcription tool. This step was done in parallel with Steps 2, 4 and the coding in Step 5.

Phase 1: The data gathered from this step contributed to answering RQ1-3, and partly RQ4.
Phase 2: The data gathered from this step contributed to evaluating the artefacts from Step 6 in Phase 1. Phase 2 interviews were longer, at approximately one hour.

Step 4: Personal or identifiable information of the interviewee in the transcriptions was removed to make a clean copy. While most of the transcript was cleaned by hand, IBM's internal Microsoft Copilot was utilised to eliminate filler words and personal information. The cleaned version of the transcript was then sent to the interviewee with an offer to add further insights or clarifications. This step was done in parallel with Steps 2, 3, and the coding in Step 5.

Step 5: A thematic analysis of the interview data was conducted. Inductive coding of the transcripts identified codes contained in codebooks.

Phase 1: Recurrent themes pertaining to compliance issues and their recurring locations throughout the DevOps CI/CD lifecycle were identified for the main goal of the analysis. The data gathered in Step 5 led to answering RQ1-4. A codebook was derived from twelve transcriptions.
Phase 2: A new codebook was created, with an overlap of six old codes contained in the former codebook. Feedback from participants resulted in themes aimed at improving the recommendable actions.

Step 6: Based on the analysis of the interviews, recommendable actions were derived and evaluated. Steps 5 and 6 contributed to answering RQ1-4.

Phase 1: A small workshop was conducted with the research team and supervisors to select five prioritised clauses to investigate as part of the artefacts, the amount limitation was set to maintain a manageable scope. These specific clauses were chosen heuristically and prioritised according to the findings from Phase 1. These are Clauses 7, 8, 9, 11 and 12, which are described in Section 5.2.2 of the Results. The recommendable actions aimed at mitigating the challenges associated with integrating ISO/PAS 8800 compliance concurrently into continuous development and delivery practices in SDVs.
Phase 2: The purpose of the latter phase of the interview study was to evaluate the artefacts and further explore the themes from Phase 1 connected to RQ4.

4.3 Thematic Analysis

The interviews were analysed using Braun and Clarke's six-phase thematic analysis (TA) framework [45]. The TA was performed by the researchers first individually coding one transcription. Following a joint discussion of each code, a final set of codes from the transcription was chosen and a shared codebook was produced. A second transcript was then subjected to the same procedure.

The remaining transcripts were coded independently by each researcher. To establish intercoder agreement, two additional transcripts were coded separately and compared. Cohen's Kappa, a statistical measure of how consistently two coders independently assign the same codes to the same pieces of text, was then evaluated. Phase 1 received a Cohen's Kappa value of 64.2% on 15% of all transcripts, while Phase 2 received a value of 63.6% on 33% of all transcripts. Values between 61–80% are commonly classified as a substantial strength of agreement [46].

In total, twelve interview transcripts were collected and analysed in the first phase. Thematic saturation was considered to have been reached at this point, as new interviews were no longer generating substantially new codes or themes, and responses had become repetitive. This assessment is consistent with empirical research on saturation in inductive thematic analysis, where data saturation for the most part have occurred by the 12th interview [47]. The second phase comprised of six interviews, where saturation was reached. While this is a smaller number than Phase 1, the interviews often lasted approximately one hour, and the questions were more focused on structured artefact evaluation. This sample size is supported a study which demonstrated that six evaluators are sufficient to uncover approximately 80% of major usability problems in a system [48].

For Phase 1, 97 codes were coded at least once, some of these were subcodes, or subcodes of subcodes. Phase 2 resulted in 79 codes, making the total amount of codes 176 with 6 codes overlapping both codebooks. These codes were then grouped together, and themes and their meanings were extracted from them. All codes and the themes they belong to can be found in codebooks displayed in Appendix C. For Phase 1, RQ1-4 were directly connected to each theme. During the TA for Phase 2, codes were grouped to the relevant clause, which provided a cluster of feedback for each clause.

4.4 Interview Participant Roles and Demographics

The participants consisted of industry professionals, including both domain experts and researchers. Many participants were purposefully selected based on their expertise in automotive safety standards. For example, several participants held roles related to system safety or functional safety engineering. However, the interviews were not limited to safety-focused roles. The full distribution of roles can be seen in Table 4.1.

Table 4.1: Interview participant profession contributions, where * implies participation of the Phase 2.

Participant	**Profession contribution**
P1, P16*, P17*	System Safety
P2, P3, P7, P13*, P14*, P18*	Functional Safety
P4, P8	AI Research
P5, P6, P12	Industry Research
P9, P10	Automotive Business Development
P11	Automotive IT Architect
P15*	Standard Conformance

4.4.1 Years of Experience and Country of Work

The two phases differed notably in participant experience. Participants in Phase 1 reported an average of 21 years of industry experience with a median of 25 years. Participants in Phase 2 reported an average of 14 years and a median of 10 years. This reflects years of accumulated knowledge in the participant pool, many of whom have experience working through many relevant standards and legislations. Their answers contain a depth of contextual and professional knowledge which shapes how the findings of the research should be understood: the insights gathered are grounded in lived experience rather than being solely based on technical familiarity. Participant awareness of the research topics such as MLOps and ISO/PAS 8800 varied, with the most recurring baseline being a general familiarity with the concepts, but no particular practical experience. To provide sufficient information about the concepts, a brief theoretical presentation was given, which can be seen in Appendix A.

Across both phases, Sweden was the most represented country, accounting for two thirds of Phase 1 participants and half of Phase 2 participants. Germany was the second most represented country in both phases, with one quarter of Phase 1 and one sixth of Phase 2 participants. The remaining participants were more geographically distributed. Japan contributed one participant in Phase 1, while the United States and Italy each contributed one participant in Phase 2. Figure 4.2 presents the geographical distribution of country of work for the interview study.

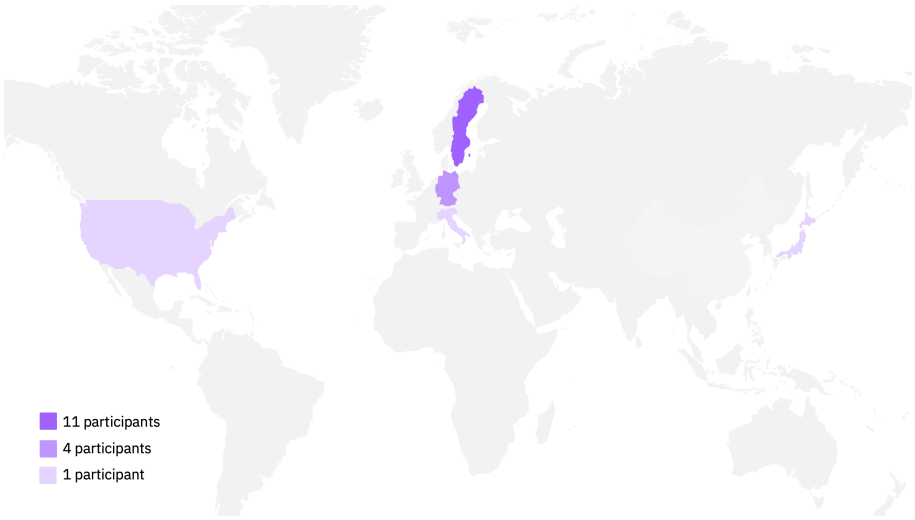

Figure 4.2: World map of participants.

4.5 Ethical Considerations

Ethical principles for software engineering interview studies were followed where applicable, throughout this two-phase study [49]. Before each interview began, participants were informed of the study's purpose and data handling practices, and explicit consent was obtained for recording the transcription, and the use of Mi-

crosoft Copilot to clean and anonymise transcripts. No participant or company names were disclosed in the cleaned transcripts. Transcripts were stored securely in IBM's OneDrive and were accessible only to the researchers. Participants were given the opportunity to review their cleaned and anonymized transcripts and to add, correct, or remove content.

5

Results

This chapter presents the results structured according to the four RQs introduced in Section 1.4. Phase 1 initially addressed RQ4, while Phase 2 refined and validated the response to RQ4. Phase 1 and Phase 2 consisted of twelve and six participants respectively. Through thematic analysis on the two interview phases, two codebooks were developed with 97 and 79 codes respectively. Six codes overlapped in the codebooks, resulting in 170 unique codes in total. Both codebooks and how the codes were grouped into themes can be found in Appendix C.

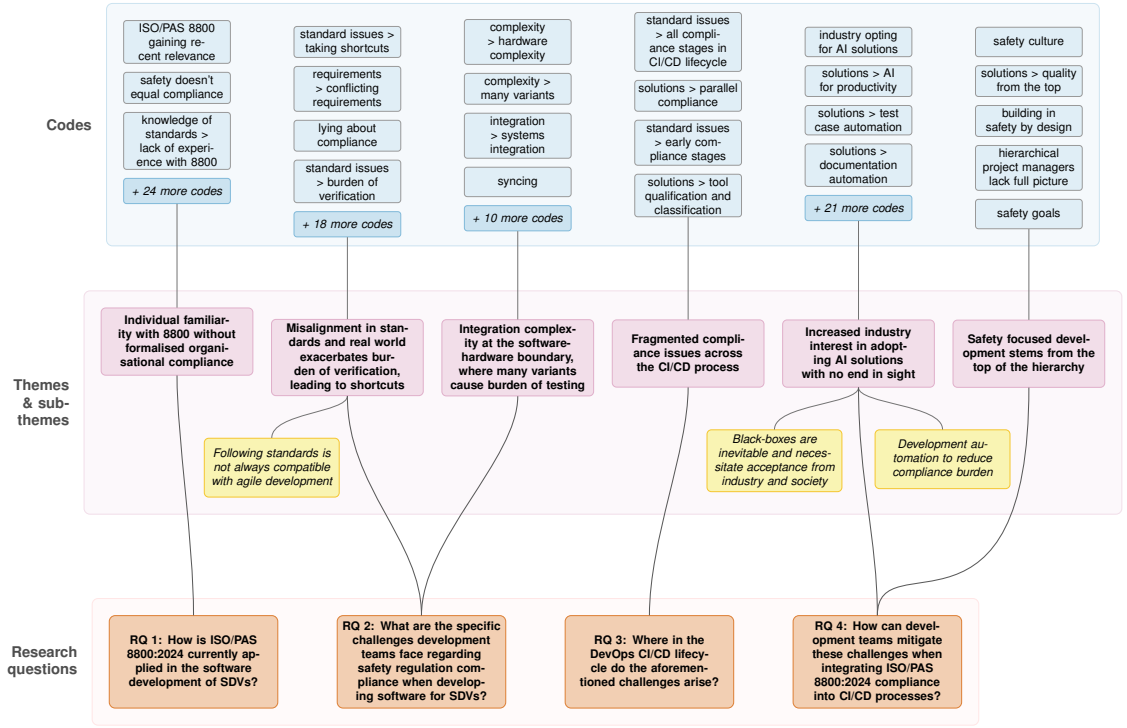

Figure 5.1: Themes emerged from the first phase of interviews. Codes and subcodes (indicated by ">") derived from the twelve interview transcripts, with representative codes shown, were grouped into six themes and three subthemes.

Figure 5.1 shows the results from Phase 1 and how the analysis was built up step by step. First, the data was broken down into codes, which were then grouped together based on similarities. From those groups, broader themes and subthemes were written. Note that some a small subset of codes did not belong to any theme.

Finally, the broader themes were linked back to the study's RQs, showing how the data from the interviews ultimately connects to the four RQs. From the first codebook, six themes and three subthemes were derived and connected to each RQ.

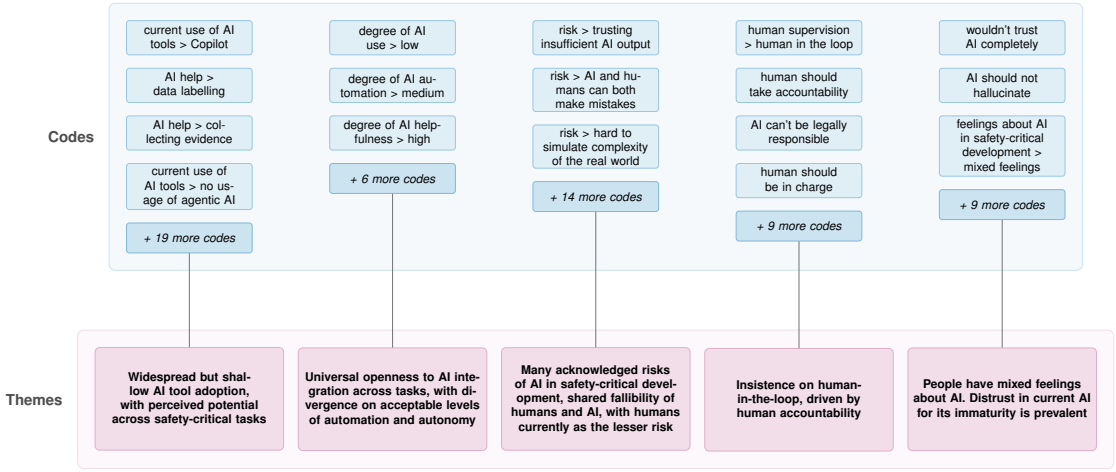

Figure 5.2: Themes emerged from the second phase of interviews, in which participants evaluated the artefacts in Section 5.2.2. Codes derived from the six interview transcripts, with representative codes shown, were grouped into five themes.

Figure 5.2 follows the same structure as Figure 5.1. However, it does not show the connection between the themes and the RQs. This is because Figure 5.2 relates to the second interview phase, where participants were asked to evaluate two artefacts rather than directly answer the RQs. The figure still shows how the codes were grouped and how themes emerged from those groups, but does not take that final step of linking the themes back to the RQs. From the codebook of the second phase, five themes were derived. Each theme is presented in the subsection corresponding to the RQ it addresses.

5.1 Findings Addressing RQ1-4

The findings effectively address all four RQs. Regarding RQ1, the findings indicate that ISO/PAS 8800 is not formally applied in practice, despite individual practitioners demonstrating some familiarity with the standard. The findings to RQ2 showed that development teams face a wide range of difficulties pertaining to safety regulation compliance. The findings to RQ3 showed that these difficulties are not concentrated in any one stage of the development lifecycle, but rather are dispersed throughout the whole lifecycle.

Additionally, two themes resulting from the analysis could be connected to RQ4. The two themes were: "Increased industry interest in adopting AI solutions with no end in sight" and "Safety focused development stems from the top of the hierarchy". The first theme also contained the subthemes "Black-boxes are inevitable and necessitate acceptance from industry and society" and "Development automation to reduce

compliance burden", as highlighted in yellow in Figure 5.1. The first subtheme, which relates to AI, is something this study can help address and that is how the two artefacts came about. The second subtheme, which is more closely related to broader organisational issues, is not something this study can provide direct practical support for. However, it is still valuable information for company leaders to consider when trying to tackle compliance issues related to safety-critical systems.

5.2 Recommendable Actions

After conducting Phase 1, the identified issues could be mapped onto the AI safety lifecycle. A few potential solutions had been identified as well, based on one of the themes connected to RQ4, "Increased industry interest in adopting AI solutions with no end in sight", as well as its two subthemes, a new focus was constructed. This reoriented the focus toward how AI could provide support, where its use is appropriate, and where human involvement remains necessary. These were not easy questions to answer. Therefore, two research artefacts were developed with the intention of visualising the integration of ISO/PAS 8800 into existing development practices. General guidelines were also created to answer the broader challenges identified through themes.

5.2.1 Guidelines

To tackle the hardware bottleneck challenge resulting from an abundance of variants, focus on fewer and better-defined variants. The automotive industry can reduce integration complexity by pushing out fewer variants. This is an organisational decision, and it requires leadership to prioritise the depth of verification rather than breadth of configurability. Fewer variants imply fewer integrations to verify, which in turn makes the burden of testing more manageable.

The structural misalignment between the V-model framing of ISO/PAS 8800 and iterative, agile development workflows represents a fundamental compatibility challenge. On the organisational side, consideration should be given to adopting more hierarchical engineering approaches where feasible, as these reduce the silo effect that arises from applying the V-model iteratively across sprints. On the standards side, future revisions of ISO/PAS 8800 could consider explicitly incorporating guidance for agile and CI/CD-oriented development contexts, rather than assuming a sequential lifecycle. Until such alignment exists, organisations should treat the misalignment itself as a known risk and establish governance mechanisms to mitigate risks, such as having designated conformance checkpoints in sprints.

5.2.2 Artefacts

The artefacts were built with the purpose of aiding the process of conformance to ISO/PAS 8800 in two ways. These can be seen in Figures 5.3 to 5.6b. The first artefact represented where the researchers believed AI, and more specifically AI automation, could provide support within the workflow, and to which degree. It also

identified where human involvement would still be necessary, either through a HITL or a HOTL approach. The second artefact was created to provide contextual support for the first artefact. It served as a complementary artefact intended to give AI systems context about which clause they could support in each phase of the MLOps lifecycle. The initial idea was to build the artefact around the traditional CI/CD model, however, it lacked in aspects of AI model development and extensions of CI/CD were therefore investigated. The MLOps lifecycle was consequently selected because it is specialised for the data-driven nature of AI model development. Evaluation interviews with industry professionals were then conducted to assess whether the artefacts were appropriate and where adjustments might be needed. The artefacts build on five technical clauses, where each clause addresses a phase or certain activities in the development. The selected clauses and their respective description can be seen in Table 5.1.

Table 5.1: Selected ISO/PAS 8800 clauses for artefacts and their descriptions

Clause Number	Description
Clause 7	Defining the AI safety lifecycle and its activities, which involves making a safety plan. A similar AI safety lifecycle to the reference in ISO/PAS 8800 can be seen in Figure 2.3.
Clause 8	Constructing a safety assurance argument: a structured, logical case demonstrating that the AI system meets its safety requirements. Evidence gathered across all development steps is arranged into a chain of reasoning showing why the system is considered safe.
Clause 9	Translating high-level safety goals into specific, measurable requirements allocated to the AI system. For instance, a goal such as "do not hit pedestrians" is decomposed into concrete technical criteria that can be verified and tested.
Clause 11	Managing the entire data lifecycle, including gathering, labelling, cleaning, verifying, validating, and maintaining datasets. A dataset safety analysis is performed to determine whether the data itself could introduce safety problems into the AI system.
Clause 12	Verifying and validating the AI system through structured and rigorous testing. Verification establishes whether the system meets its specified requirements, while validation establishes whether it performs safely under real-world conditions.

Figure 5.3 to 5.5 belong to the same first artefact. The artefact contains three different aspects of the AI safety lifecycle reference provided in Figure 2.3, two of which regard the helpfulness and automation that AI can provide in the workflow. There is a distinction between helpfulness and automation where the former does not regard the amount of AI-tooling that is being used. Helpfulness refers to the degree to which AI can meaningfully support human work within a given clause. Automation speaks to something different, namely how much of that work could, in principle, be handed off.

Figure 5.3 explains where human oversight and involvement should be placed on the AI safety lifecycle. All inspected clauses require a HITL, while Clause 11 also has potential for HOTL oversight.

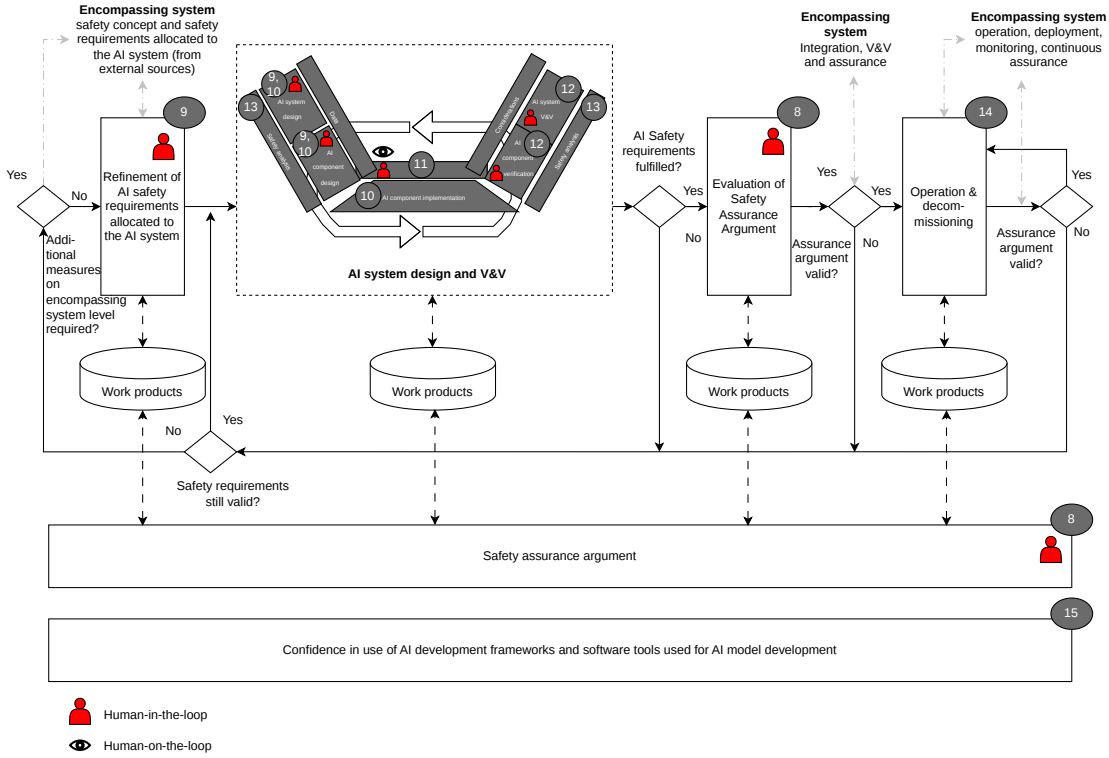

Figure 5.3: The final result of AI artefact with HITL/HOTL focus.

Figure 5.4 evaluates another aspect of the AI safety lifecycle: the helpfulness of using an AI tool, rated on a scale of low, medium, and high, where darker shading indicates a higher degree. Clauses 7, 8, and 11 received a high degree of perceived helpfulness of using an AI tool, while Clause 9 and 12 received a medium degree. No clause contained a low degree of AI helpfulness.

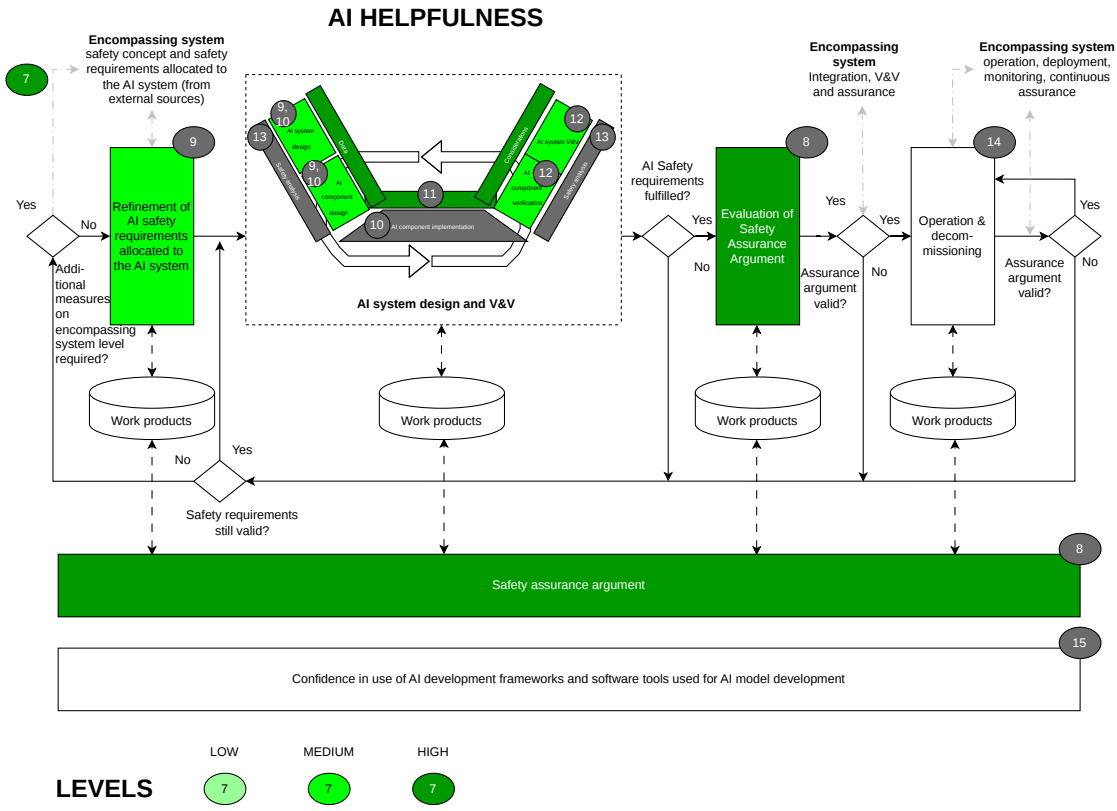

Figure 5.4: The final result of AI artefact with AI helpfulness focus.

The final aspect of the artefact is depicted in Figure 5.5: the degree to which AI automation can be implemented for each clause. Similarly to Figure 5.4, the degree is divided into sections of low, medium and high. Clause 11 had a high degree of automation, and so did Clause 12 partly. The reason a split has been done on Clause 12 in Figure 5.5 is because automation was favoured for the verification part of Clause 12, while validation could not reach the same amount of automation. Clause 7, 8, and partly 12 had medium degree of automation. Clause 9 had the lowest rated potential for automation.

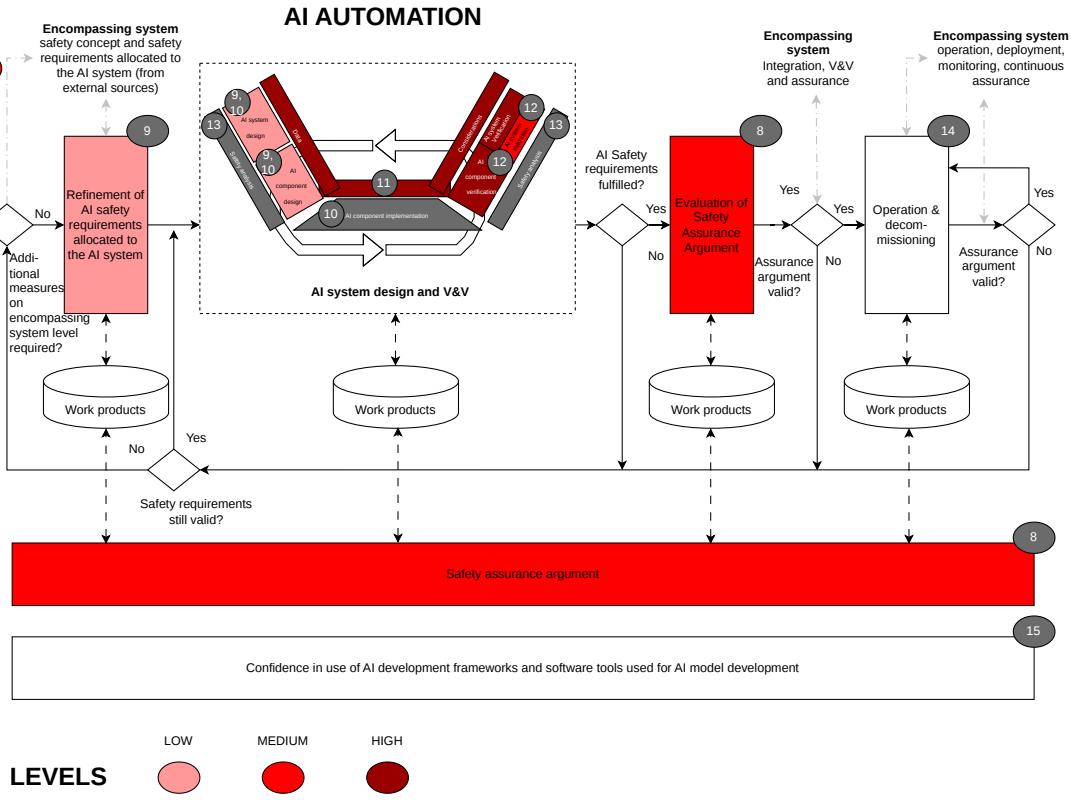

Figure 5.5: The final result of AI artefact with AI Automation focus.

Comparing the two Figures 5.4 and 5.5, both contain a high degree of helpfulness and automation in Clause 11. Further comparison between the two reveal that, although the degrees vary, the relationship between degrees tends to follow the same pattern. Darker colours (medium-to-high) are prominent in Clause 7, 8, 11 and 12.

During the analysis, a conclusion was reached that the traditional DevOps alone did not fully capture the challenges associated with AI-based systems. Since software-defined vehicles increasingly rely on machine learning components, an MLOps perspective was more appropriate. MLOps extends DevOps practices to include the machine learning lifecycle, making it better suited for representing AI development workflows. The artefact in Figure 5.6b has two components, one optimised for AI tool input and one for human readability.

(a) The MLOps artefact, developed for AI-tools. Clause prominence across MLOps loop phases (1-3 scale).

Phase	Clause 9	Clause 8	Clause 11	Clause 12	Clause 7
Plan	3	3	2	1	3
Create	3	2	2	2	2
Data	2	2	3	2	2
Model	2	2	3	2	2
Verify	2	2	2	3	2
Package	2	2	2	2	2
Release	2	2	2	1	2
Configure	2	2	2	2	2
Monitor	2	2	2	2	2

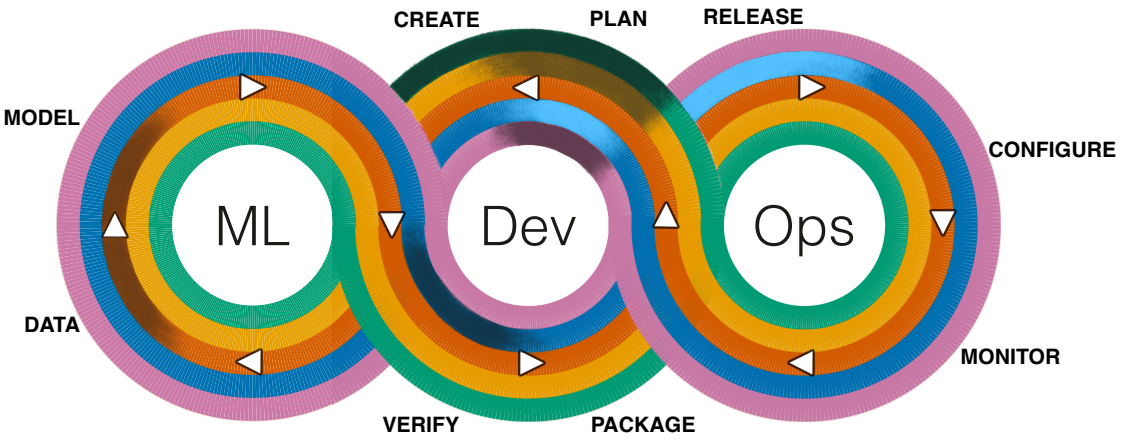

(b) Visual representation of the MLOps artefact, developed for the human eye. Lighter areas indicate that the clause is less prominent in that phase, whereas darker areas indicate the opposite.

Figure 5.6: The MLOps artefact: clause prominence across loop phases (a) presented numerically on a 1–3 scale, and (b) visualised for human readbility.

5.3 Current Application of ISO/PAS 8800

This section aims to answer *RQ1: How is ISO/PAS 8800 currently applied in the software development of SDVs?* Each subsection is a theme from the conducted thematic analysis providing an explanation to the RQ, beginning with the themes of the first phase. The themes can also be found in Figure 5.1 and 5.2. A selection of corresponding relevant codes and supporting quotes for each section can be seen in Table 5.2 - 5.5. Illustrative quotes for each theme are highlighted in blue. A table

containing the selected relevant codes and supporting quotes can be seen in full in Appendix D.

Table 5.2: Codes and supporting quotes from Phase 1 for answering RQ1

Code	Supporting Quotes
ISO/PAS 8800 gaining recent relevance	"[...] it was mentioned that they were thinking about ISO 8800 [...] But I think it's a big topic for [Company] and we should address it." *Participant 11*
compliance doesn't equal safety	"If we are talking about issues of complying with ISO 26262, it's the same issues. But if you are talking about making cars safe, that's a different thing." *Participant 6*

Individual Familiarity with ISO/PAS 8800 Without Formalised Organisational Compliance

This theme describes a familiarity with ISO/PAS 8800 from industry practitioners, who have either heard of it or read it themselves. Companies, however, lack the frameworks for compliance, since it is not a mandatory standard to adhere to through legislation. Being a PAS implicates the subject being a novel area in technology and thus no formal requirement of legal compliance exists in corporations. In comparison to its parent standard ISO 26262, ISO/PAS 8800 has yet to be a common standard to prove state-of-the-art safety-driven development in products for stakeholders, resulting in an environment where technical implementation and advancements precede regulatory guidance.

Furthermore, this lack of framework for compliance shows in the content of ISO/PAS 8800, where a large amount of the requirements, although normative, are in the early stages of development. A more settled standard consists of normative requirements which have been edited and iteratively improved through time, so that they are established to a higher degree. ISO/PAS 8800, in comparison, contains requirements that parties may have included that have yet to go through the development. Co-creators of ISO/PAS 8800 will update the requirements in new versions, until the sentences are agreed-upon by all parties involved. Thoroughly following a newly published PAS to state that the product is compliant will require an investment of time, and when the requirements are inevitably changed, the compliance, too, has to change with it. This might explain the reluctance towards compliance of ISO/PAS 8800.

> "[..] another advantage with publicly available specification is that even if you write the best of our standard with all companies [...], at the end, maybe even my colleague in my company think I was wrong. So they will provide comments [...] give the society some years. Maybe [...] a professor who is not part of that will read it and say that 'Oh, this is wrong. It should be written in another way'. And then [...] people publish papers on that [...] and then you publish the second version, the first version of the final version." *Participant 5*

The rapid growth of AI technology has resulted in a monumental demand for fast de-

velopment in multiple sectors, the automotive industry being one of these. Emerging technologies will require time to stabilise and become dependable. Because industry professionals produce standards and the industry depends on the standards' output, there is an underlying interdependency between industry and standards. It is difficult to adhere to the standard ISO/PAS 8800, because the standard regards immature technology that is undergoing development. Consequently, when ISO/PAS 8800 potentially becomes established in the industry, the industry will also have a more defined and mature technology. The synchronisation between technological maturity and standard refinement is a central challenge for the industry.

"Basically, there is a chicken and egg thing with these standards." *Participant 5*

Additionally, the theme encapsulates the perspective on safety that is not inherently addressed by compliance. The standard works as an advisory guideline to assure safety and minimise risks, and safety is a broad term that encompasses many aspects of being protected from danger or risk. Practitioners often view compliance as a secondary administrative layer rather than a primary safety driver. Companies also utilise alternate software-level activities to maintain safety; standard compliance does not equate to safety.

5.4 Identified Compliance Challenges

This section aims to answer *RQ2: What are the specific challenges development teams face regarding safety regulation compliance when developing software for SDVs?* Each subsection is a theme from the conducted thematic analysis providing an explanation to the RQ, the themes can also be found in Figures 5.1 and 5.2.

Table 5.3: Codes and supporting quotes from Phase 1 for answering RQ2

Code	Supporting Quotes
standard issues > burden of verification	"And you also need to make sure that it is reviewed by someone else, and that generally takes longer than quickly making a change without doing all these things, without being that thorough, which means it is often perceived as having to do a lot of extra documentation." *(Translated from original language) Participant 2*
common problem of hazard analysis in post development	"[...] the ISO standard requires that you actually deliver documents on time. But you cannot do that if the functionality does not exist. It is a fairly common problem that first they sit and program, and then the systems group is supposed to come in and design or finalise analyses. Doing hazard analysis after the fact is reprehensible." *(Translated from original language) Participant 3*
complexity > integration complexity	"With respect to a software defined vehicle you would then see a more horizontal approach. Maybe then a zonal architecture with four to five zones per vehicle and with a much higher degree of integration which requires let's say a kind of what we at [Company] tell a model-based systems engineering approach. Really integrating hardware and software development." *Participant 9*

Misalignment in Standards and Real World Exacerbates Burden of Verification, Leading to Shortcuts

One of the common challenges highlighted in the codes was the misalignment between standards and real world development. The terminology in standards can sometimes be inapplicable to software development practice, where the software ecosystem goes against the process of following a standard. For instance, it is common to provide evidence per product release, but product releases are continuous in the world of software. Problems therefore already arise in the stages of pure software development.

> "[...] it can say [...] some senseless requirement: Do software architectural analysis in the fashion of FMEA [...]. In software, it doesn't work that way. In software, you usually have millions of failure modes [...] You cannot really do a comprehensive analysis there and provide like FMEA for software." *Participant 6*

> "You can't really comply safety case for these thousands of updates. It just doesn't make sense. So, and in audit, people usually understand and are reasonable. So generally you show evidence that you are doing some job." *Participant 6*

The misalignment issue also encapsulates frustration of unrealistic expectations coming from the standards. Requirements can be mismatched with the development process due inadequacy or obsolescence, and therefore simply be wrong.

When combining standards to fulfil safety aspect throughout the whole system, conflicting requirements in different standards can arise. All of the aforementioned issues contribute to the burden of verification and leads many to taking shortcuts or even lying about compliance.

> "Nobody does that, even though everybody claims here, we are ISO 26262 compliant." *Participant 6*

Following Standards Is Not Always Compatible with Agile Development

This is a subtheme of the previous theme. The Agile methodology is widely used across many industries, including automotive development. It is iterative and flexible, with a strong emphasis on collaboration and adaptability. The V-model prescribed by both ISO 26262 and ISO/PAS 8800, however, is linear and sequential, placing considerable weight on upfront requirements and rigid documentation. This opposes Agile's adaptive nature, making compliance a challenge for teams working in an Agile environment. Agile methodology is misaligned with the standard workflows in safety-critical systems.

> "[...] another problem you have is that you try to introduce the standard into companies that don't work so hierarchically, while the entire standard is hierarchical [...] it states directly that the structure should be hierarchical." (Translated

from original language) *Participant 3*

"You can't work agile. You can't at a low level decide that 'we're going to have five wheels', and then there's someone else at a higher level who... I mean, a car is hierarchical. You should work agile towards the customer, but then you can't work agile." (Translated from original language) *Participant 3*

For instance, in requirements management, Agile provides flexibility in evolving or unfulfilled requirements. In safety-critical systems, every requirement is mandatory to validate system integrity. Changes in safety-critical systems pose a risk of introducing new hazards that require re-validation.

Rigorous documentation is also a strict priority for safety-critical systems conforming to safety standards. This contradicts the Agile Manifesto's emphasis on a working product rather than extensive documentation [50].

"Where I have worked, it has traditionally not been driven from an ISO 26262 perspective [...] a lot of focus has been placed on having 'working software' over documentation, which means they have perhaps chosen to do the documentation at the end [...] they have perhaps deprioritised writing requirements [...] until the end [...]. And when you try to correct that [...] ISO 26262 [...] is seen as an obstacle because it slows down the pace or efficiency." (Translated from original language) *Participant 2*

This conflict is seen as fundamental and contributes to a methodological dissonance between industry and standard compliance. The opposition between these two approaches suggests that resolving the conflict necessitates a re-evaluation of either the way of working or the standards themselves.

Integration Complexity at the Software-Hardware Boundary, Where Many Variants Cause Burden of Testing

When software and hardware meet, the amount of possible variants and configurations grow substantially. Each component contains internal complexity, and being dependent on other components in the ecosystem causes compounding complexity. The existing complexities will also increase when introducing new variants, which will require extensive testing.

"For instance, in hardware parts, hardware is a really complex thing. It's a material. And there are transistors [...] they are very, very hard to make. So it's a very advanced technology." *Participant 5*

"There are two big things affecting this, even making things even more difficult with this complexity of software[...]. The first thing that comes with it and makes things worse is having too many variants, because you cannot test that many variants. You have to take risks." *Participant 6*

The integration layer compounds this further, introducing new combinations that must all be accounted for. The result is a considerably greater testing burden, which adds to the difficulty of achieving and maintaining compliance with safety regulations.

> "After five years they come, they say this issue has gone bad, it doesn't work. You cannot really change that single ECU with its software. The software version has to work with the rest of the software in the other ECUs. So it has become too interconnected. So you have to somehow have a baseline of software that is compatible for all the ECUs [...]. It's very, very difficult to get that right [...] software complexity [...] makes the entire car or truck interconnected. Then on top of that, way too many variants and backwards compatibility [...]" *Participant 6*

5.5 Where Compliance Challenges Arise

This section aims to answer *RQ3: Where in the DevOps CI/CD lifecycle do the aforementioned challenges arise?* Each subsection is a theme from the conducted thematic analysis providing an explanation to the RQ, the themes can also be found in Figure 5.1 and 5.2.

Fragmented Compliance Issues Across the CI/CD Process

The compliance challenges identified by participants are not confined to any single phase of development. Instead, they surface throughout the entire lifecycle. Several participants expressed a need for compliance to run in parallel with development rather than being treated as a separate concern addressed after the fact.

> "What I would say is that it actually arises throughout the entire chain." (Translated from original language) *Participant 2*

> "It's always good to have it. I mean together with the lifecycle rather than having it at the end of the lifecycle. It should be there from the start of the lifecycle. That's how I feel." *Participant 2*

Problems introduced early in the process tend to carry forward, potentially triggering further compliance issues later in the development process. Addressing compliance continuously allows early-stage problems to be caught and resolved before they propagate, whereas retrospective checking will only surface them much later.

> "So early software development stages where you have pure software, most of the problems come up already there." *Participant 6*

5.6 Mitigation Approaches

This section aims to answer *RQ4: How can development teams mitigate these challenges when integrating ISO/PAS 8800 compliance into continuous development and delivery?* Each subsection is a theme from the conducted thematic analysis providing an explanation to the RQ, the themes can also be found in Figures 5.1 and 5.2.

Table 5.4: Codes and supporting quotes from Phase 1 for answering RQ4

Code	Supporting Quotes
solutions > AI for productivity	"Mostly, AI usage is discussed in relation to productivity improvements, but it is also being considered for innovation." *Participant 12*
likelihood estimation > society tolerance	"All technology always has an inherent risk. Then the question is what risk you can accept." *(Translated from original language) Participant 4*
requirements > limitations of blackboxes	"ISO 26262 is more deterministic and traceable, but when you come to ISO/PAS 8800, it is non-deterministic. You don't know where the failure can happen [...] the tools that are used in the CI/CD tool chain may introduce new errors." *Participant 1*

Table 5.5: Codes and supporting quotes from Phase 2 for answering RQ4

Code	Supporting Quotes
AI help > documentation	"I think it will definitely help expedite the creation of the artifacts." *Participant 15*
risk > hard to simulate complexity of the real world	"The question is how accurately the AI tool can mimic the real situation. That is one risk." *Participant 13*
trusting insufficient AI output	"[...] most of the time the tool simply generates some output, which should not be the case. If someone starts trusting that output, it can cause serious problems in the end." *Participant 13*
wouldn't trust AI completely > immature technology	"At the end of the day, it is a question of trust, and that is very common in safety-related topics. We are not at a point where we can fully trust AI yet. Maybe we will be in the future, but we are not there now." *Participant 14*

Increased Industry Interest in Adopting AI Solutions with No End in Sight

AI has become a pressing topic across many industries, and the conversation has shifted from *whether* to adopt it to a question of *when* and *how*. This momentum is reflected in the views of participants.

> "Still, I think companies are forced to use such tools in at least some areas, because you save so much money with them." *Participant 10*

> "[...] they are looking at actively using AI in SDV development processes. And then some areas they are looking at a big productivity improvement. And including regulatory compliance, like ISO 26262 or other ISO standards as well." *Par-*

ticipant 12

Simultaneously, competitive pressure is making the automotive industry turn towards more time-efficient solutions, as another participant expressed here:

"The problem is the long time that is needed for development [...] compare the development times of Chinese companies and European companies for a car. In Europe, the usual amount of time is three to four years, more four than three. In China, they need about two years. And you can ask yourself why they are so fast. [...] And regarding whether they use Gen AI [...]" *Participant 10*

Safety regulations are implemented to display the state-of-the-art development in safety-critical systems, while aiming to protect society from life-critical risks. At the same time, development and implementation time of AI in SDVs is surpassing that of European carmakers and OEMs.

Black-boxes are inevitable and necessitate acceptance from industry and society

This is a subtheme of the previous theme. Deep learning models are inherently opaque, and their limitations are not always well understood. This can create challenges when trying to write requirements for technologies such as LLMs.

"Still, we don't know what the limitations of LLMs are. We know that they hallucinate [...] So, writing requirements on those technologies [...] might be a bit challenging." *Participant 5*

The societal tolerance aspect of likelihood estimation is also voiced by several participants, where non-deterministic systems are often treated with scepticism by the public.

"I think we have to also consider that we as humans have [...] an extreme bias against automation in vehicles [...] we are much more [...] forgiving of essentially the same accident occurring if a person were [...] driving [...]. So I think we as a society need to somehow figure out what is good enough. If we basically say [...] probabilistic systems are the only way to go, then [...] the question becomes what is good enough versus [...] the only thing that is good enough is zero chance of these things happening." *Participant 8*

Prohibition of such technologies would be counterproductive given how embedded they are becoming. The more constructive path forward would be responsible adoption with safeguards and transparency about known limitations.

"Solutions, to me, [...] don't prohibit to use any of these technologies. Use them, although they are immature, but at the same time use them in a responsible way." *Participant 5*

Development Automation to Reduce Compliance Burden

This too is a subtheme. Automation in the development process has the potential to meaningfully reduce the overhead associated with compliance. Participants pointed in particular to documentation and verification tasks as areas where automation could have the greatest impact, especially when applied early in the development process before compliance debt accumulates.

> "I actually believe that AI could help formulate requirements in a better way, because from what I have seen with requirements writing at both the system level and software level, you tend to get a little lazy when writing your requirements." (Translated from original language) *Participant 2*

> "That's one of the goals of a proper CI/CD DevOps tool chain to really address these topics and automatically generate all the necessary reports you need for compliance [...] the goal is to automate all the stuff and reporting." *Participant 11*

Safety Focused Development Stems From the Top of the Hierarchy

Safety in complex development environments is not something that can be guaranteed through individual effort alone. It requires a culture that is embedded across the entire organisation, built into the design process from the outset and sustained by a commitment to quality at every level. The same applies to training the models from an architecture perspective, where safety can be built in by design.

> "I think it requires a lot of a return to fundamental work on sort of theory and applying that to training models and also figuring things out on the sort of like architecture side, so building in safety by design rather than as an add-on at the end" *Participant 8*

This mindset needs to be driven from the top. Without genuine buy-in from leaders and directors who make the company-wide decisions, a safety-first culture is unlikely to take meaningful hold.

> "And risks I have experienced. There is always a risk that it becomes acrobatically overdone with too much documentation. But at the same time, documentation is important. It is very difficult, and you need managers who understand and want to do this. That is what is called a Safety culture. It is very important to have. Otherwise it just becomes like showing that we follow this standard on paper." (Translated from original language) *Participant 3*

> "It is very difficult to introduce quality work from the bottom up. It almost always leads to burning people out. Quality must come from the top. And if you don't have management with you saying that this is important, then it is a swamp." (Translated from original language) *Participant 3*

Widespread but Shallow AI Tool Adoption, with Perceived Potential Across Safety-Critical Tasks

This and the following themes were created using the results of the thematic analysis of the second phase of interviews, they can be seen in Figure 5.2, and the corresponding relevant codes and supporting quotes can be seen in Table 5.5. Current use of AI tools among participants includes Copilot, Claude, and Gemini, used primarily for linguistic and research purposes. Development-specific use cases were mentioned less frequently, including test case generation, the code editor Cursor, and AI features embedded in standard tools.

> "So I use Gemini both to gather information and to refine text." *Participant 14*

Participants identified considerable potential for AI assistance across various stages of conformance to ISO/PAS 8800, with the most prominent areas being documentation, initial drafting, and the repetitive work of data labelling. On the documentation side, collecting evidence was highlighted as a particularly promising application, as the current process can be tedious.

> "The evidence is spread across different places and is now collected manually and then assembled into a safety case. If this could be done by AI, it would be really, really helpful." *Participant 14*

> "Data labelling has been a human task for many years, but now it can be automated in very effective ways." *Participant 14*

Automating this process could save significant time and effort. While many potential areas of AI adoption were identified, implementation in the industry remains largely absent.

Universal Openness to AI Integration Across Tasks, with Divergence on Acceptable Levels of Automation and Autonomy

Participants expressed openness to AI integration across all discussed tasks, though the acceptable degree of automation and perceived helpfulness varied by clause.

> "Yes, yes, yes, very much because this one also is too much documentation and takes a lot of time and effort. [...] And it could easily be done using AI and there is less risk to have something missed in this one." *Participant 13*

> "So it's actually no harm to safety or humans here. So, I could accept a higher level of automation here." *Participant 16*

For Clause 9, perceived helpfulness was predominantly medium, and the acceptable degree of automation was low, though helpfulness was still considered meaningful. For Clause 8, helpfulness was rated mostly high, with a medium degree of automation considered acceptable. Clause 11 stood out with both high perceived helpfulness and a high acceptable degree of automation. Clause 12 received medium ratings

for both helpfulness and automation, with the latter ranging up to high for some participants. Clause 7 showed the strongest consensus, with every participant rating AI helpfulness as high, and automation rated medium to high. The potential for automation can be seen in Figure 5.5.

Many Acknowledged Risks of AI in Safety-Critical Development, Shared Fallibility of Humans and AI with Humans Currently as the Lesser Risk

Participants identified a range of risks associated with using AI tools in the development tasks covered by the selected clauses. The most commonly raised concerns were: AI missing important aspects, trusting insufficient AI output, difficulty in simulating the complexity of the real world, and the risk that developers stop thinking critically. These risks contribute to a general apprehension around adopting AI tools in safety-critical contexts.

> "The risk is that AI might miss or fail to understand the full scope of all the safety issues." *Participant 15*

> "Because I don't think AI retain that history and that knowledge and that depth of reasoning that humans do." *Participant 15*

> "For example, if you remember the Waymo cases, the systems were trained never to cross a white line because doing so could put them into oncoming traffic and cause a collision. Then people went out and drew white lines around the vehicle, which caused it to get stuck." *Participant 15*

At the same time, participants acknowledged that human developers are not free from error either, and in some cases may be more prone to mistakes than AI. This creates an interesting parallel: While the number of identified risks associated with AI was greater, the comparison with human fallibility complicates a rejection of AI involvement.

> "So what usually people do is copy-pasting, which is much more dangerous than telling AI a context [...]. It's much more effective than copy-pasting things and much less error-prone on my evaluation." *Participant 14*

Insistence on Human-in-the-Loop Oversight, Driven by Human Accountability

Across the majority of clauses discussed in Phase 2 of the interview study, participants expressed a strong preference for HITL oversight. A minority of participants found HOTL arrangements acceptable for certain tasks, but fully automated development without human involvement was considered unacceptable by all.

> "That is how the law works right now. We currently have no legal framework that defines how an artificial intelligence system can be held accountable." *Participant 14*

The primary reason given for this stance was accountability. Participants argued that AI cannot bear legal responsibility, and that when failures occur, the company or individual responsible for the development process will be held accountable. Given this, many felt that a human needs to remain actively involved rather than serving as a passive supervisor.

> "At the moment, we trust people, and we know how to make them accountable. We do not yet know how to make AI accountable [...]." *Participant 14*

> "You need someone to be accountable and also an organisation to be accountable." *Participant 13*

Several participants also framed the human role in more architectural terms: as AI takes on more operational tasks, the human becomes responsible for overseeing and directing the process rather than executing it. This shift was seen as reasonable only if the human retained genuine control and visibility over what the AI was doing.

People Have Mixed Feelings About AI. Distrust in Current AI for Its Immaturity Is Prevalent

When asked at the end of each interview how they felt about the use of AI tools in safety-critical development, participants gave a range of responses. While some expressed cautious optimism, others were more sceptical, and a general sense of caution ran through nearly all the interviews. A recurring theme was distrust rooted in the perceived immaturity of current AI technology, with participants uncertain whether the tools are reliable enough for the stakes involved.

> "Concerned. [...] But the risks associated with using it without thinking is high. [...] I can ask it to do a complete hazard analysis. And I will get something that looks like it has an analysis. But when you start reading it, every sentence is totally fine but it might be bogus. This is really concerning for me because it looks great, reads great at the fast glance. But the technical content and the technical correctness can sometimes be very wrong." *Participant 16*

> "It's an evolution just as part of our lives. I think we should embrace it. And I think that we should try to see how it can help us the most to camp. So I'm excited. I'm all for it." *Participant 15*

A code belonging to this theme was "AI doesn't have motives like humans". It regards a concern about AI reliability in the sense that it lacks motivation. This motivation is intrinsic to humans and drives initiative, it encourages effort that sometimes go beyond previously set limits. AI, however capable it may become at executing tasks, does not have the same motivation in the outcome.

> "[...] an AI does not want anything. I, as a human, want something. I want my company to be competitive on the market with the great products. So on top of things, there has to be a human in my mind." *Participant 16*

This perspective is another reason for the prevalent distrust in AI that goes beyond its technical immaturity.

6

Discussion

This chapter discusses the findings of the study in relation to the RQs, the existing literature and reflects on the findings broader implications. The sections follow the structure of seven thematic discussion sections, moving from how ISO/PAS 8800 is situated within current industry practice, through the challenges of conformance, to how development teams might mitigate them. The chapter closes with limitations, threats to validity, and possible directions for future work.

6.1 Situating ISO/PAS 8800 Within Safety Practices

ISO/PAS 8800 is a relatively recent specification published as an ISO/PAS because the field of AI safety in road vehicles is evolving faster than a full ISO standard process can accommodate. Therefore, its role within industry practice is still taking shape. The findings that speak most directly to RQ1 are captured in the theme "Individual Familiarity with ISO/PAS 8800 Without Formalised Organisational Compliance". The theme contained codes such as "compliance doesn't equal safety", "interdependency between standard and industry", and "interpretation of 8800 is okay". Together, these codes indicate that some individual practitioners have encountered the specification, yet their organisations lack any established framework for conformance. Both ISO 26262 and ISO/PAS 8800 are technically voluntary documents, but they occupy very different positions within the automotive industry. Through decades of iterative development and widespread industrial adoption, ISO 26262 has effectively become the de facto standard for automotive functional safety, and conformance to the standard is commonly treated as evidence that an organisation is applying recognised industry practices for managing functional safety throughout the development lifecycle [51].

The low organisational uptake of ISO/PAS 8800 may reflect the immaturity of the standard rather than a deliberate decision by companies to avoid it. At the same time, the finding that practitioners do not equate compliance with safety is important. This is not an individual opinion, but true as even when compliance with safety standards is achieved, products can still function unsafely, concluding that technical compliance alone is insufficient to ensure product safety [52]. ISO/PAS

8800 should be viewed as one tool among many for improving safety, rather than as proof that a vehicle is safe simply because the standard is followed. Conformance to the specification may indicate a stronger safety process, but it does not guarantee safety on its own. The current low adoption of ISO/PAS 8800 should therefore not be interpreted as evidence that the industry is managing well without it. The problem is not that companies think they are safe without it, it is that they do not yet know how to be safe with it and that is what needs to change.

6.2 Why Conformance Breaks on a System Level

An interesting finding from the results concerns why conformance to safety standards breaks down from a system-level perspective. It comes from the theme "Misalignment in standards and real world exacerbates burden of verification, leading to shortcuts" as well as its subtheme "Following standards is not always compatible with agile development", addressing RQ2 by identifying what the existing challenges are of standard conformance.

The subtheme presents the misalignment between ISO/PAS 8800's V-model structure, which demands a more sequential development process, and agile development which is iterative. The second point of the Agile Manifesto, "Working software over comprehensive documentation" demonstrates how the principles of Agile directly conflict with the standards' emphasis on documentation [50]. The waterfall process in safety standards has been identified across industry literature as a structural obstacle to agile adoption, with documentation burden being one of the most commonly stated issues [53].

Because of this paradox, organisations in the industry have to find a balance between meeting specification requirements and the rapid pace of software development. In practice, updating a safety case for a bug fix of a released product has been reported to take upwards of two weeks [54]. This exemplifies how conformance to safety standards demands documentation, which can disrupt the fast pace that iterative development aims to achieve. As the automotive industry continues shifting toward Agile and DevOps processes, especially in the development of SDVs, this misalignment between workflows requires careful investigation [39].

At the same time, the high complexity of hardware and software integration introduces new challenges of conformance, as told by the theme "Integration complexity at the software-hardware boundary, where many variants cause burden of testing". The presence of multiple hardware variants exacerbates the integration challenges, particularly in SDVs where software and hardware must function together across a myriad of configurations. The burden of verification consequently becomes more exacerbated, leading essentially to a misalignment between those responsible for development and those responsible for safety assurance. This has been identified as a structural challenge in safety-critical automotive development [54].

While they present as individual issues, these pressures should not be interpreted

in isolation. Together, the two challenges create conditions where shortcuts, incomplete conformance practices, or even inaccurate conformance reporting become more understandable responses to structural pressure. According to the results, auditors are typically reasonable and aware of this difficulty, which defeats the objective of audits and compliance checks. This finding suggests that the issue is one of system rather than one of individual level. The demands of the standard, the modern software development, and the hardware complexity constraints are fundamentally misaligned. Addressing any one of these issues in isolation is therefore unlikely to resolve the underlying pressure. Instead, solutions likely need to be pursued at an organisational level. Both guidelines in the recommendable actions reflect this conclusion, as there is no tooling-based solution alone can mitigate this challenge. A reduction in amount of variants requires product strategy decisions from leadership and hierarchical engineering requires a structural re-evaluation of the current development approach. Neither can be achieved through tooling, solutions to these challenges must rather be pursued at a systemic level.

6.3 Compliance Issues Are Fragmented Across the Lifecycle

The results from the first phase of interviews revealed that compliance challenges surface throughout the entire process rather than cluster in any single phase of the CI/CD lifecycle. When compliance obligations are distributed everywhere, they risk becoming everyone's concern in theory and no one's priority in practice.

This finding is captured in the theme "Fragmented compliance issues across the CI/CD process" connected to RQ3. It encompasses the two codes: "standard issues > all compliance stages in CI/CD lifecycle" and "standard issues > early compliance stages". Both codes indicate that compliance concerns are present both from the start and across all subsequent phases. When the corresponding challenges found to RQ2 were mapped onto the AI safety lifecycle figure during the first iterations of the artefacts (Appendix B.1), no pronounced concentration of issues in any particular place in the lifecycle was visible. However, this picture becomes more nuanced when examined alongside the MLOps artefact in Figure 5.6b. When the five chosen ISO/PAS 8800 clauses were mapped onto the MLOps lifecycle phases and their prominence was aggregated, the Plan phase received the highest total score (12), followed by the Create, Data, Model, and Verify phases (11 each). This can be seen in Table 6.1.

The Plan phase consistently scored at the highest end across Clauses 7, 8, and 9. Although these phases do not host a disproportionate number of challenges, they appear to carry the greatest overall workload, a finding consistent with prior literature on safety activity distribution in iterative development [6]. This is a noteworthy finding. While a single-point difference in isolation may appear marginal, it takes on greater significance in context. The Plan phase's score of 12 stands out against the uniform score of 11 shared by the four phases that follow it, and

Table 6.1: Final result of MLOps artefact, developed for AI-tools. Clause prominence across MLOps loop phases (1-3 scale) and the total value of the clause prominence values of each phase.

Phase	Clause 9	Clause 8	Clause 11	Clause 12	Clause 7	Total
Plan	3	3	2	1	3	12
Create	3	2	2	2	2	11
Data	2	2	3	2	2	11
Model	2	2	3	2	2	11
Verify	2	2	2	3	2	11
Package	2	2	2	2	2	10
Release	2	2	2	1	2	9
Configure	2	2	2	2	2	10
Monitor	2	2	2	2	2	10

the second half of the MLOps lifecycle, Package, Release, Configure, and Monitor, drops further still to 10 or 9. This cumulative front-loading of compliance weight aligns with the well-documented challenge of safety activities being isolated from development teams and executed retrospectively [55].

This suggests that the planning stage may still demand the most sustained compliance effort. This lends support to a shift-left approach, that is, deliberately anchoring compliance obligations earlier and more continuously in the development process as a potential strategy for reducing the burden of late-stage verification and preventing early-stage problems from propagating forward through the lifecycle. Another implication of this is that AI-based tooling and automation may offer the greatest potential value precisely in these beginning phases of the MLOps lifecycle.

6.4 Safety Culture Starts at the Top

The theme "Safety focused development stems from the top of the hierarchy" from Phase 1 reflects an aspiration expressed by practitioners. That is, the theme captures what practitioners believe should be the case rather than what commonly is. Some of the codes associated with this theme such as "hierarchical project managers lack full picture", "safety culture", and "solutions > quality from the top" collectively point in the same direction. They indicate that practitioners understand safety as something that requires ownership at a senior level. This finding connects directly to RQ4, which asks how development teams can mitigate the compliance challenges identified in RQ2. Crucially, when participants discussed potential mitigation strategies, they did not point solely to tools, automation, or process improvements. They also identified leadership commitment, organisational structure, and safety culture as necessary preconditions for technical solutions to have their intended effect.

Both Artefacts 1 and 2 in Section 5.2.2 are intended to support practitioners in

structuring their conformance activities more effectively. However, it is important to remember they operate within organisational constraints. That is, an artefact can indicate that a given clause requires human oversight or that the Plan phase carries the heaviest conformance workload, but it cannot instil the organisational commitment needed for those recommendations to be taken seriously and applied consistently. The finding that safety must be driven from the top is therefore not merely a practitioner preference, but a diagnosis of why technical solutions alone are insufficient to address conformance challenges. This implies that organisations seeking to integrate ISO/PAS 8800 into CI/CD should treat leadership engagement, safety culture and safety goal-setting as prerequisites to successfully deploying technical solutions to mitigate conformance challenges.

6.5 Everyone Wants AI Automation, Yet Nobody Trusts It

One interesting finding, that directly answers RQ4, is that two themes from the second phase of interviews are opposing each other. The themes are "Universal openness to AI integration across tasks, with divergence on acceptable levels of automation and autonomy" and "People have mixed feelings about AI. Distrust in current AI for its immaturity is prevalent". The first describes an openness towards utilising AI tools, while the latter presents a contradictory distrust in AI, stemming from its perceived current immaturity. For nearly every clause discussed, participants identified tasks where AI could be helpful, though the degree of helpfulness varied. As seen in Figure 5.4, the range varied between medium and high helpfulness, meaning that a low degree was not prevalent in any of the discussed clauses.

A previous theme in Phase 1 describes an "Increased industry interest in adopting AI solutions with no end in sight", so the industry is opting for AI. Practitioners have also acknowledged that humans in some instances can make more mistakes than AI, as seen in the code "risk > humans can sometimes make more mistakes than AI". The external competitive market is also putting pressure on the European automotive industry. As mentioned in the results under the theme "Increased industry interest in adopting AI solutions with no end in sight", shorter development times and less strict regulatory environments in other regions could be creating additional strain on European development practices. Safety conformance requirements could therefore become considerably more demanding. Trade measures currently in place are providing a degree of protection for European manufacturers [56], though how long this remains the case is uncertain. Standards set to ensure safety and prove state-of-the-art solutions can become a barrier in the competitive global race toward more autonomous vehicles when conformance challenges arise.

At the same time, the same participants expressed reservations that limited how much autonomy they were willing to grant AI. They insisted on HITL oversight in all clauses. When asked directly about their feelings toward AI in safety-critical development, responses varied, though caution was the most common thread. This

caution could be explained by the many identified risks of trusting insufficient AI output, hallucination, and the absence of legal accountability for AI decisions. This caution, combined with uncertainty about the maturity of current AI technology, reinforces the distrust that runs through a majority of the interviews.

Additionally, another reservation of a different character was expressed, one that is unlikely to diminish as AI matures. As noted in the codes for the theme "People have mixed feelings about AI. Distrust in current AI for its immaturity is prevalent", the lack of motivation in AI is a reason for distrust in it too. People conduct work with personal stakes: A desire to see the company succeed, to advance in their career, to avoid failure. These motivations drive initiative, accountability and sometimes even innovation. AI does not have stakes in the outcome. This aspect will not be resolved once AI reaches an acceptable level of maturity, it is rather a fundamental reason to keep humans actively involved regardless of capability of the AI. The conflict of wanting automation but not trusting it in a safety-critical environment may diminish as AI matures, but most participants still preferred to have a human involved in the tasks when asked if they would change their answer if it could do the work to the same standard as their best colleague.

The conflict of wanting AI while distrusting it indicates that the industry is in a transitional phase in which AI technology has outpaced the regulatory frameworks needed to govern it well and the organisational experience needed to calibrate trust in it appropriately. Until both catch up, adoption of AI tools is likely to continue being at a shallow level regardless of how capable they become. The artefacts proposed in this study are therefore timely, as they aim to give industry professionals a basis for deciding where AI involvement is appropriate.

6.6 Where Humans Must Remain and Where AI Can Operate Autonomously

The first artefact presents the AI safety lifecycle through three interconnected aspects, which when read together form an internal logic of where humans must remain and where AI can run free. HOTL oversight appears only once across the AI safety lifecycle, in Clause 11, where the repetitive aspects of data handling could be automated with less oversight. HITL is still required, but the possibility of having a HOTL is most feasible here. All clause activities still require human involvement rather than monitoring alone. AI helpfulness is rated medium or high across all clauses, indicating that AI has something to contribute at every stage. Furthermore, in Clause 12, a distinction between automation potential stands out, where verification scores high, while validation scores medium. Given the different nature of the two activities, automation is higher on verification because such tasks are typically more specifiable and deterministic. Verification asks if the system satisfies its defined requirements, which can be answered by comparing outputs with predefined criteria. Validation, by contrast, involves more contextual judgement, as it asks the question of whether the system behaves safely in the real world.

The overall logic after following the findings collectively is the indication that AI is best positioned to substitute for human effort where tasks are rule-based and verifiable, and least replacing where tasks require accountability or judgements on real-world matters. This directly addresses RQ4. The economic feasibility of safety conformance has been argued to depend on focusing on automation in tedious, repeating, and error-prone tasks [54]. Rather than opting for a blanket approach to human oversight or automation in safety conformance, the artefact provides an empirically grounded foundation for organisations in the industry on where to strategically allocate human oversight and AI automation.

6.7 Structural Misalignment and the Path Toward Responsible Conformance

The preceding discussion sections have focused on findings and their implications relating to each RQ. One condition permeates all of our findings: ISO/PAS 8800 built in one engineering archetype is applied to a product that is constructed in another.

What is more significant than any of the findings, however, is that practitioners were aware of the difficulties in standard conformance and development, they could also generally articulate potential solutions. This raises a question the findings cannot fully answer but make difficult to ignore: *Why do the conditions remain when practitioner opinions are broadly shared?* The answer is partly due to the difficult transition of comprehending an issue to organisationally committing to resolving it. Our artefacts might help in this process, but not entirely without the support of a present organisational safety mindset and culture. As ISO/PAS 8800 has yet to become a formally applied standard in the industry, this question is highly relevant to take into consideration for a sustainable adoption of ISO/PAS 8800.

The findings related to AI integration also have a meaning beyond what the challenges individually implied. The willingness to automate tasks with AI whilst also distrusting it due to its immaturity may change as the technology matures. Some sentiments and arguments against full AI automation that are not based on its maturity level will still challenge the open mindset. A future aspect of automation is that safety-critical engineering is enhanced through years of experience when accumulating exposure to the full context of a problem. When AI automates the difficult work and humans might have more of a reviewing role, that exposure to the deep understanding of a problem might gradually diminish. A developer who has faced problem-solving directly will likely bring a different quality of judgment to edge cases than a reviewer whose primary exposure to the task is through AI output. This points to that human oversight may be necessary but not sufficient to preserve the depth of understanding that the most consequential decisions in this domain depend on. If the domain knowledge that gives human oversight its value is not fully preserved and transmitted, it may find, some years from now, that its humans are in the loop without being fully equipped to manage it. That risk is not yet well

understood in the literature, and this study cannot resolve it. What it can do is name it clearly enough that it becomes part of how the industry reasons about AI adoption, rather than something discovered retrospectively.

6.8 Limitations

The first limitation revolves around the defined set of clauses selection contained in ISO/PAS 8800 for the second phase of interviews, namely Clause 7, 8, 9, 11 and 12. After conducting the thematic analysis of the first interview phase transcriptions, a small workshop was held where the researchers prioritised which clauses would be most interesting to investigate as part of the artefacts. The findings therefore do not cover all normative requirements of ISO/PAS 8800, and the artefacts are limited in scope. While this scope was necessary for an in-depth analysis of the selected clauses, it limits the generalisability of the research. The clauses were selected heuristically based on relevance for researching AI tooling implementation. The resulting findings are valid for the subset of clauses but cannot be assumed to apply across the entirety of ISO/PAS 8800.

The absence of direct observation of real CI/CD pipelines and the lack of audits of actual compliance documentation is a limitation of this study. The proposed artefacts have also not been implemented or tested in a real development pipeline. They were evaluated through interviews, which means that their practical utility is unverified. As the study was designed as a qualitative interview study, the collected data reflects reported practices rather than directly observed practices within safety-critical software development. Reported and actual practice do not always fully align, and this gap represents a limitation of the chosen methodology.

6.9 Threats to Validity

Threats to validity can be divided into threats to internal validity and threats to external validity. Internal validity concerns whether the study itself is valid, while external validity concerns whether the results are applicable to other contexts or geographical areas.

6.9.1 Threats to Internal Validity

One identified threat is self-censorship within the participant pool. Many of the interviewed safety engineers work professionally with arguing for compliance, which may have caused them to provide more idealistic answers when discussing challenges related to standard conformance and distrust in AI. Due to their professional roles, some participants may have found it difficult to openly discuss weaknesses or shortcomings in their own work practices. To mitigate this threat, participants were guaranteed anonymity and given the opportunity to review, edit, remove, or approve their interview transcriptions after the fact. This was intended to create a safer environment for honest expression of opinions and experiences.

A further threat concerns interviewer bias and researcher influence. Semi-structured interviews are interactive, both the phrasing of questions and the choice of follow-up questions may unintentionally introduce bias. This is particularly relevant in Phase 2, where participants evaluated artefacts created by the researchers themselves. Participants may have been more inclined to validate the artefacts rather than critique them, as criticising someone else's work can feel socially uncomfortable. To reduce this effect, leading questions were avoided. Questions were framed in a more open-ended manner, which made feedback and suggestions for improvement feel more welcome. While this type of bias is difficult to eliminate entirely when evaluating one's own artefacts, awareness of it remained throughout the process. During Phase 1, no stance was taken by the researchers regarding conformance challenges, due to the previous unawareness of the domain.

A third internal validity threat concerns intercoder reliability. Cohen's Kappa scores were 64.2% for Phase 1 and 63.6% for Phase 2, both falling within the range of moderate agreement. Although these values are acceptable, they indicate that some coding decisions differed between researchers when transcripts were coded independently. It should be noted, however, that the majority of Phase 2 transcripts were coded collaboratively, and several Phase 1 transcripts were also coded jointly. Independent coding was applied, primarily where it was considered both beneficial and time-efficient.

The use of ISO 26262 as a proxy for studying research questions pertaining to ISO/PAS 8800 can also be seen as a threat. The AI-specific concerns addressed in ISO/PAS 8800 do not translate identically onto a standard originally designed for deterministic software systems. As a result, some insights derived through this proxy may not transfer directly to the ISO/PAS 8800 context. However, this approach was considered the most viable option available, given that ISO/PAS 8800 has not yet been widely adopted within industry, which became evident during the interviews. Furthermore, ISO/PAS 8800 builds substantially upon the material and structure already present in ISO 26262, with the primary distinction being its handling of non-deterministic rather than deterministic software systems.

The final threat to internal validity concerns the sample size of the Phase 2 interview round, which focused on artefact evaluation. The second phase included six participants, however, these interviews were considerably longer than those in Phase 1, often lasting approximately one hour or more. The discussions were also more specialised and questions more in-depth. On this basis, data saturation was considered to have been reached by the end of the second interview phase.

6.9.2 Threats to External Validity

A first threat to external validity is the geographical representation. The majority of participants in the study were based in Sweden, though participants from Germany, Italy, the United States, and Japan were also included. Given that Swedish participants formed the largest portion of the sample, it is important to acknowledge that the automotive safety culture in Sweden is generally considered mature and

safety-conscious. This may render the findings more representative of the Swedish automotive industry than of the global automotive industry more broadly.

A second threat regards the seniority of the participant pool. In Phase 1, the median level of industry experience among participants was 25 years. This notably high level of seniority may shape perspectives on AI automation, human accountability, and modern development workflows in ways that differ from less experienced practitioners. Junior professionals, who represent a substantial portion of the workforce in the safety-critical automotive industry, may hold considerably different views from those reflected in this study. As a result, the findings may not fully capture the range of perspectives held across the industry. The lowest amount of years of experience in the participant range was however four, six and seven years, meaning that representation of more junior industry practitioners is present in the data.

A third threat concerns the evolving nature of ISO/PAS 8800. As the specification continues to mature and potentially become a standard in the future, industry practices and interpretations of its requirements are likely to change from the version released in 2024. The long-term applicability of the study's findings is therefore uncertain, and conclusions drawn from the current state of adoption should be understood within that temporal context.

Finally, the study is constrained by its domain-specific sample and its focus on a particular set of standards. All participants worked within automotive software-defined vehicle contexts, meaning that the findings cannot readily be generalised to other safety-critical domains undergoing AI adoption, such as aviation or medical devices. These industries operate under distinct regulatory frameworks, follow different development processes, and exhibit varying levels of AI maturity, all of which may give rise to substantially different challenges and practices from those identified here.

6.9.3 Reliability

A study should be repeatable by other researchers and have consistent results to assure the reliability of a study. Although the interview settings cannot be repeated, the materials are documented throughout the study, especially in Appendix A, and may be utilised for replication.

7

Conclusion

This study investigated how ISO/PAS 8800 can be integrated into CI/CD development of software-defined vehicles, and how AI tooling can support that process through a two-phase interview study with industry practitioners. It was also found that organisational adoption of the standard remains low. Compliance challenges are fragmented across the entire development lifecycle, with the planning phase carrying the greatest overall burden, even though this disparity is marginal, it matters.

It can be concluded that there is a clear interest from the industry in using AI to support conformance-related work. This was consistently reflected throughout the findings. The next question consequently became: "Where should AI be used? How? Where does human involvement remain necessary and in which way?".

These were important questions to address, especially since AI-assisted compliance support, including agentic AI solutions, is already emerging on the market. In response, this study produced a set of recommendable actions, consisting of a set of guidelines and two artefacts which were evaluated, intended to give organisations concrete guidance on how to approach these questions. The artefacts provide guidance on where AI support may be appropriate and to which degree, where human oversight is needed, and how these considerations relate specifically to conformance to five prioritised clauses in ISO/PAS 8800. The accompanying guidelines address the broader organisational and structural conditions that must be in place for such guidance to take effect.

The main takeaway is that AI has a genuine role to play in conformance support, but human involvement remains non-negotiable particularly because of the need of accountability and distrust of current AI. The two artefacts produced in this study offer practitioners a principled starting point for deciding where AI can appropriately assist and where humans must remain in the loop. Ultimately, the question is not whether to involve AI in compliance work, but how to do so responsibly. As AI technology and specifications or standards continue to mature alongside each other, getting that balance right is not a one-time decision but an ongoing responsibility for the industry.

7.1 Future Work

For future work, the next step could be to implement some of the recommendable actions in a real industrial pipeline, measuring both their practical usability and their effect on the burden of conformance work.

Furthermore, the clause coverage in the artefacts could be extended beyond the five prioritised clauses and examined in the same way, namely Clause 10, 13, 14 and 15.

Finally, as this study relied on standards such as ISO 26262 as a partial proxy for studying ISO/PAS 8800, future studies should revisit the findings with participants who have direct and formalised experience with the standard if industry adoption of the standard grows. Such studies may reveal challenges and mitigation strategies not yet visible at the time of this study.

Bibliography

[1] International Organization for Standardization, "Road vehicles — safety and artificial intelligence," December 2024.

[2] D. Smith, D. Villalba, M. Irvine, D. Stanke, and N. Harvey, "Accelerate state of DevOps report 2021," research report, DORA (DevOps Research and Assessment), Google Cloud, 2021. Accessed: 12-May-2026.

[3] G. Lindemulder and M. Kosinski, "What is DevOps?." IBM. [Online]. Available: `https://www.ibm.com/think/topics/devops`. Accessed: 17-Dec-2025.

[4] B. W. Boehm, *Software Engineering Economics*. Englewood Cliffs, NJ: Prentice-Hall, 1981.

[5] A. A. Mohamed and H. K. Aslan, "Toward safe AI in autonomous vehicles: Challenges, standards, and an integrated safety engineering framework," *Procedia Computer Science*, vol. 272, pp. 706–711, 2025.

[6] A. Nouri *et al.*, "The DevSafeOps dilemma: A systematic literature review on rapidity in safe autonomous driving development and operation," *Journal of Systems and Software*, vol. 230, December 2025.

[7] S. Edwards, "Self driving car accidents trend chart (2026)." *ConsumerShield*, May 2026. Accessed: 19-May-2026.

[8] S. Levin and N. Woolf, "Tesla driver killed while using autopilot was watching Harry Potter, witness says." *The Guardian*, Jul 2016. Accessed: 19-May-2026.

[9] AI Incident Database, "Incident 231: A Tesla Crashed into and Killed a Road Sweeper on a Highway in China." AI Incident Database, Jan 2016. Accessed: 19-May-2026.

[10] P. Koopman, "Anatomy of a robotaxi crash: Lessons from the Cruise pedestrian dragging mishap," *arXiv preprint arXiv:2402.06046*, Feb 2024.

[11] R. Jolak, V. Antinyan, A. Åström, D. Durisic, O. Kopp, S. Kriebel, D. Krippner, M. Mohamad, J. Pfeiffer, C. Seiler, P. Svenson, A. Wortmann, and J. Bosch, "Navigating the future: Essential considerations for the engineering of software-

defined vehicles," *Computer*, vol. 59, no. 1, pp. 74–84, 2026.

[12] IBM, "What is a software-defined vehicle?," 2025. Accessed: 10-Mar-2026.

[13] Waymo LLC, "Waymo: Self-driving cars," 2026. Accessed: 10-Mar-2026.

[14] Apollo Go, "Apollo go: Autonomous ride-hailing service," 2026. Accessed: 10-Mar-2026.

[15] Tesla, Inc., "Tesla: Electric cars, solar and clean energy," 2026. Accessed: 10-Mar-2026.

[16] S. Islam, "MLOps vs. DevOps: What is the difference?." phData Blog, February 2022. Accessed: 18-May-2026.

[17] K. Salama, J. Kazmierczak, and D. Schut, "Practitioners guide to MLOps: A framework for continuous delivery and automation of machine learning," white paper, Google Cloud, May 2021.

[18] D. Kreuzberger, N. Kühl, and S. Hirschl, "Machine learning operations (MLOps): Overview, definition, and architecture," *IEEE Access*, vol. 11, 2023.

[19] R. Merritt, "What is MLOps?." NVIDIA Blog, September 2023. Accessed: 18-May-2026.

[20] G. Schouten, M. Sangiovanni, and W.-J. van den Heuvel, *IoT Beehives and Open Data to Gauge Urban Biodiversity*, pp. 225–235. Advances in Intelligent Systems and Computing, Springer, April 2021.

[21] N. Forsgren, J. Humble, and G. Kim, *Accelerate: The Science of Lean Software and DevOps*. Portland, OR: IT Revolution Press, 2018.

[22] Software Freedom Conservancy, "Git." `https://git-scm.com/`, 2026. Accessed: 13-Mar-2026.

[23] Jenkins Contributors, "Jenkins." `https://www.jenkins.io/`, 2026. Accessed: 13-Mar-2026.

[24] The Kubernetes Authors, "Kubernetes." `https://kubernetes.io/`, 2026. Accessed: 13-Mar-2026.

[25] Docker Inc., "Docker." `https://www.docker.com/`, 2026. Accessed: 13-Mar-2026.

[26] F. Warg, H. Blom, J. Borg, and R. Johansson, "Continuous deployment for dependable systems with continuous assurance cases," in *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*, (Berlin, Germany), pp. 318–325, 2019.

[27] IBM, "What is human-in-the-loop?," 2025. Accessed: 14-May-2026.

[28] European Parliament and Council of the European Union, "Article 14: Human oversight - artificial intelligence act," 2024. Official text of the EU AI Act.

[29] Svenska institutet för standarder, SIS, "What is a standard?." Accessed: 11-Mar-2026.

[30] International Organization for Standardization, "Iso - benefits of iso standards," 2022. Accessed: 11-Mar-2026.

[31] International Organization for Standardization, "Road vehicles — safety of the intended functionality," June 2022.

[32] International Organization for Standardization, "Foreword - supplementary information," n.d. Accessed: 18-Mar-2026.

[33] International Organization for Standardization, "Iso - international organization for standardization," 2026. Accessed: 10-Mar-2026.

[34] Svenska institutet för standarder, SIS, "Road vehicles - functional safety - part 1: Vocabulary (iso 26262-1, idt)," 2018. Accessed: 11-Mar-2026.

[35] International Organization for Standardization, "Road vehicles — safety of the intended functionality," 2022. Accessed: 11-Mar-2026.

[36] A. Nouri, C. Berger, and F. Törner, "An industrial experience report about challenges from continuous monitoring, improvement, and deployment for autonomous driving features," in *2022 48th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, (Gran Canaria, Spain), pp. 356–363, IEEE, 2022.

[37] H. Sandgren and V. Antinyan, "Software safety analysis to support iso 26262-6 compliance in agile development," *IEEE Software*, vol. 38, no. 3, 2021.

[38] S. Kugele and M. Broy, "Architecture as a backbone for safe DevOps in automotive systems," in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 4145–4150, IEEE, 2022.

[39] G. Resing, "End-to-end devops for the software-defined vehicle," white paper, IBM Corporation, 2024. IBM Consulting.

[40] S. Burton, L. Gauerhof, and C. Heinzemann, "Making the case for safety of machine learning in highly automated driving," in *Computer Safety, Reliability, and Security – SAFECOMP 2017 Workshops* (S. Tonetta, E. Schoitsch, and F. Bitsch, eds.), vol. 10489 of *Lecture Notes in Computer Science*, pp. 5–16, 2017.

[41] S. Burton, "Safety under uncertainty: Automotive standards for AI safety and

research perspectives." Invited talk at NASA Formal Methods (NFM 2024), Moffett Field, CA, USA, 2024.

[42] L. Gauerhof, P. Munk, and S. Burton, "Structuring validation targets of a machine learning function applied to automated driving," in *Computer Safety, Reliability, and Security – SAFECOMP 2018*, vol. 11093 of *Lecture Notes in Computer Science*, pp. 45–58, 2018.

[43] R. Ashmore, R. Calinescu, and C. Paterson, "Assuring the machine learning lifecycle: Desiderata, methods, and challenges," *ACM Computing Surveys*, vol. 54, no. 5, pp. 1–39, 2021.

[44] S. A. Zafar, J. Kelly, and N. Mata, "Unified AI-product lifecycle based on road-vehicle safety standards," in *2025 20th European Dependable Computing Conference Companion Proceedings (EDCC-C)*, pp. 189–194, 2025.

[45] S. K. Ahmed, R. A. Mohammed, A. J. Nashwan, R. H. Ibrahim, A. Q. Abdalla, B. M. M. Ameen, and R. M. Khdhir, "Using thematic analysis in qualitative research," *Journal of Medicine, Surgery, and Public Health*, vol. 6, p. 100198, Aug. 2025.

[46] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, no. 1, pp. 159–174, 1977.

[47] G. Guest, A. Bunce, and L. Johnson, "How many interviews are enough?: An experiment with data saturation and variability," *Field Methods*, vol. 18, no. 1, pp. 59–82, 2006.

[48] J. Nielsen and T. K. Landauer, "A mathematical model of the finding of usability problems," in *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, CHI '93, pp. 206–213, ACM, 1993.

[49] P. E. Strandberg, "Ethical interviews in software engineering," in *Proceedings of the 13th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*, (Porto de Galinhas, Brazil), 2019.

[50] K. Beck, M. Beedle, A. van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, J. Kern, B. Marick, R. C. Martin, S. Mellor, K. Schwaber, J. Sutherland, and D. Thomas, "Manifesto for agile software development." `https://agilemanifesto.org/`, 2001. Accessed: 17-May-2026.

[51] R. Debouk, "Overview of the second edition of iso 26262: Functional safety–road vehicles," *Journal of System Safety*, vol. 55, pp. 13–21, Mar. 2019.

[52] T. Maier, S. Aymans, and M. Mühlfelder, "Product safety culture as a concept for the development of safety–critical products in regulated industries," *Safety Science*, vol. 169, 2024.

[53] R. Kasauli, E. Knauss, B. Kanagwa, A. Nilsson, and G. Calikli, "Safety-critical systems and agile development: A mapping study," in *Proceedings of the 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pp. 470–477, IEEE, 2018.

[54] P. Munk and M. Schweizer, "DevOps and Safety? SafeOps! Towards Ensuring Safety in Feature-Driven Development with Frequent Releases," in *Computer Safety, Reliability, and Security. SAFECOMP 2022 Workshops*, vol. 13415 of *Lecture Notes in Computer Science*, pp. 145–157, 2022.

[55] U. Siddique, "SafetyOps," 2020. Bosch Research and Technology Center, Sunnyvale, California, USA.

[56] European Commission, "Questions and answers on the commission unveiling the action plan to drive innovation, sustainability, and competitiveness in the automotive sector." `https://ec.europa.eu/commission/presscorner/detail/en/qanda_25_636`, 2025. Press Corner, QANDA/25/636. Accessed: 21-May-2026.

A

Material for Interviews

A.1 Before the Interview

We plan to write our master's thesis and potentially an IBM Whitepaper based on the findings of this study. To support our analysis, we would like to transcribe the interview. As an individual interviewee, you are guaranteed anonymity, and no company names will be published in the resulting article. Only the researchers will know which interviewee provided what specific information. Only codes used in our thematic analysis and possibly selected direct quotes will be shared publicly in the article. There will be no traceability back to individual interviewees. We will store the transcripts securely in DropBox along with a file containing your contact information in case we need to validate any findings at a later stage. The content will be accessible only by us and our supervisors and stored for one year after the publication of the research article. At the beginning of the interview, we will explicitly ask you for your consent to store your contact information for the purpose described. You will participate in the interview voluntarily. You are free to skip answering any question and can choose to stop the interview at any time. Furthermore, you can request us to turn off the transcription for specific parts of the interview. Please interrupt us at any time if you need us to clarify terms or definitions. We will enable the automatic transcription service provided by Microsoft Teams during the interviews. After the interview, you will get an edited transcript per email. You will then have the chance to validate the content, remove any specific information if needed, or further clarify yourself.

If the Internet connection is lost or the meeting fails, please reconnect as soon as possible. If it is not possible to recover within a few minutes, we will reschedule the interview via email.

Thank you!

A.2 Interview Guide Phase 1

Short Introduction

- Introduce ourselves (Name, age, university and program)
- What this interview will be used for (SE Master's Thesis with Chalmers and IBM)
  - Our research focuses on the challenges of meeting AI safety standards within CI/CD workflows for Software-Defined Vehicles (SDV).
  - The purpose of the study is to investigate concurrently integrating conformance with the standard ISO/PAS 8800:2024 into the DevOps CI/CD practice of development within SDVs.
  - The aim of the interview is to identify challenges in conforming to standards in a CI/CD DevOps pipeline. And the results of the interview will be used for a Software Engineering master's thesis with Chalmers and IBM.
- Present the relevant standards shortly
  - ISO/PAS 8800:2024 provides a framework for safely integrating AI into vehicle systems such as autonomous cars
  - ISO 26262 is the bigger standard that revolves around functional safety of the development of a vehicle system.

Information and Consent

We will be using the Teams transcription tool to transcribe the interview. The full transcription is just for us and will not be publicised. Your data, however, will be collected and analysed with thematic analysis anonymously.

- Do we have your consent for conducting this interview?
- Do we have your consent for transcribing this interview using the Teams transcription tool?
- Do we have your consent for using IBM's internal Microsoft Copilot as a tool to clean the transcription and remove any personal information?

Start the transcript and warn them that they will need to unmute themselves when the transcription starts.

Collect Information About the Interviewee

"Let's begin, could you please introduce yourself, your job etc."

Things to look out for, can ask as a follow up question:

- What is your name?
- What are your years of experience?
- What is your company and your title?
- What are your daily responsibilities within this role?

Leading Questions to Ask
RQ1:

- Have you ever heard of the ISO/PAS 8800 which our thesis is about, or any other standard regarding vehicles and safety such as ISO 26262 or ISO 21448(SOTIF)?
- Do you use it in your work for compliance regulation?
  - If so, how?
  - Why do you apply it in this way?
  - Do you see any pain points?
  - If not, how do you ensure that your work is up to standard?
- Do you or anyone you work with continuously integrate and deploy (CI/CD) code in your organization or project? If so, please walk us through the process and the technologies used if you can.

RQ2:

- Do you find any challenges regarding safety regulation compliance?
  - What are they?
  - Where do they come up in the DevOps CI/CD lifecycle?

RQ3:

- Do you see any solutions that can mitigate these challenges?
- Do you think that it is possible to integrate the standard with the CI/CD lifecycle concurrently rather than sequentially?
- Do you know if there is anyone (else) at your organization that is responsible

that the software products are up to standards?

Further Insights and Thanks

"Can you come up with any further insights that we might have missed?"
Stop the transcript.
Tell them about what happens with the transcript.
"If you have further insights after reading the transcript, you can also add or complement".
Thank the interviewee for taking the time to answer our questions.

A.3 Interview Guide Phase 2

Short Introduction

- Introduce ourselves (Name, age, university and program)
- What this interview will be used for (SE Master's Thesis with Chalmers and IBM)
  - Our research focuses on the challenges of meeting AI safety standards within CI/CD workflows for Software Defined Vehicles (SDV).
  - The purpose of the study is to investigate concurrently integrating conformance with the specification ISO/PAS 8800:2024 into DevOps CI/CD practices within SDVs.
  - The aim of this interview is to evaluate where AI can and should play a role in the ISO/PAS 8800 AI Safety Lifecycle.
- Present the relevant standards shortly
  - ISO 26262 is the overarching standard for functional safety in vehicle system development.
  - ISO/PAS 8800:2024 builds on this by providing a framework for safely integrating AI into vehicle systems such as autonomous cars.
  - You do not need to know ISO/PAS 8800 in detail to answer the questions; we are interested in your professional judgment.

Information and Consent

We will be using the Teams transcription tool to transcribe the interview. The full transcription is only used by the researchers and will not be publicised. Your data will be analysed anonymously using thematic analysis.

- Do we have your consent for conducting this interview?
- Do we have your consent for transcribing this interview using the Teams transcription tool?
- Do we have your consent for using IBM's internal Microsoft Copilot to clean the transcription and remove personal information?

Start the transcript and warn the interviewee that they will need to unmute when transcription starts.

Collect Information About the Interviewee

"Let's begin, could you please introduce yourself, your job etc."

Things to look out for, can ask as follow-up questions:

- What are your years of experience?
- What is your company and your title?
- What are your daily responsibilities within this role?
- Do you or have you used any AI tools (e.g. ChatGPT, Copilot, internal tools) in your work? If so, for what purposes?

Clause-Based Interview Structure

Explain briefly:

"ISO/PAS 8800 is divided into clauses. Each clause contains work products and requirements that must be fulfilled to achieve conformance. We will briefly describe a clause and then ask questions about it. We will do this for several clauses. Does that sound alright?"

Clauses Introduced

- **Clause 7:** Establishing the AI safety lifecycle and creating a safety plan.
- **Clause 9:** Translating high-level safety goals into specific, measurable requirements.
- **Clause 8:** Building a safety assurance argument by collecting and structuring safety evidence.
- **Clause 11:** Managing the data lifecycle, including dataset safety analysis.
- **Clause 12:** Verification and validation through structured and rigorous testing.

Leading Questions to Ask for Each Clause

"I will ask the same set of questions for each clause and keep track of time."

- Assume AI could support you in performing this task. Would you want that?
  - What would you gain or lose?
- Should AI be in charge of performing this task, or should a human?
  - Why?
- What do you think the risks are?

  - What level of risk would you be willing to accept if AI performs this instead of a human?

- Where do you see the role of a human in this step?
  - If AI produces a first draft and a human reviews it, would that save time?
  - When must a human intervene?
  - Do you see any blockers?
- Would you trust the AI output enough to sign off on it?

MLOps Lifecycle Discussion

- Are you familiar with the MLOps lifecycle shown here?
- We are trying to place the ISO/PAS 8800 clauses onto an MLOps lifecycle.
- This is our proposed placement of the clause in the MLOps cycle. Do you agree, yes or no?

Additional Clause Descriptions (If Needed)

- **Clause 13:** Analysis of AI-specific failure modes (e.g., misclassification, hallucinations) and safety measures to prevent or control them.
- **Clause 14:** Post-deployment activities such as monitoring performance, collecting real-world data, re-training, re-validating, and updating the safety assurance argument.
- **Clause 10:** Choosing AI technologies, system architecture, and implementing the AI component.
- **Clause 15:** Trustworthiness and qualification of tools used to develop, train, or test AI systems.

Extra Questions

- Across all the clauses, where do you think AI can add the most value?
- Where would you absolutely draw the line, where should AI never be in charge?
- If you could automate one thing in this entire process tomorrow, what would it be?
- Is there a step where humans are the main bottleneck and AI could do better?
- How does accountability factor into your thinking?

- If AI produces a safety artifact and something goes wrong, who is responsible?
- Any final thoughts on the future of AI-assisted safety engineering?
- Are you worried about AI automation? Why or why not?

Further Insights and Thanks

"In general, how do you feel about AI being involved in safety-critical development? Excited? Cautious? Skeptical?"

Stop the transcript.
Explain what happens with the transcription.
"If you have further insights after reading the transcript, you can also add or complement them."
Thank the interviewee for taking the time to participate.

A.4 After the Interview

Dear [Name],

Thank you for participating in our interview, we truly appreciate it!

Attached is the cleaned-up version of the transcript where personal information has been removed. If you have further insights, you can add them through email. After reviewing it, you are also very welcome to clarify, adjust, or remove any parts you do not wish to be included, or simply leave as is.

Thank you again for your valuable contribution!

Best regards,
Ann & Atosa

B

Iterations of Artefacts

B.1 First Iteration

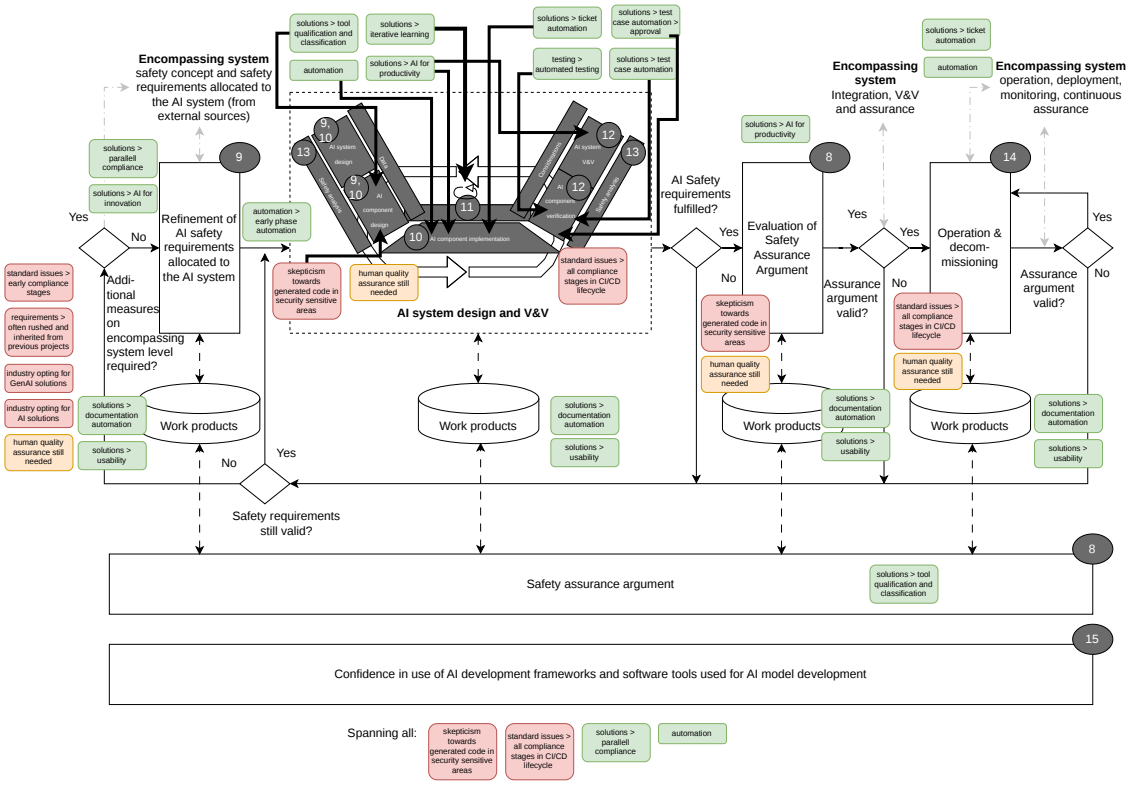

Figure B.1: The unevaluated first iteration of the AI artefact. Using the themes, codes, and interview findings from the thematic analysis, identified challenges were mapped out onto different stages of the lifecycle. The purpose was to understand where in the process these issues were most likely to emerge.

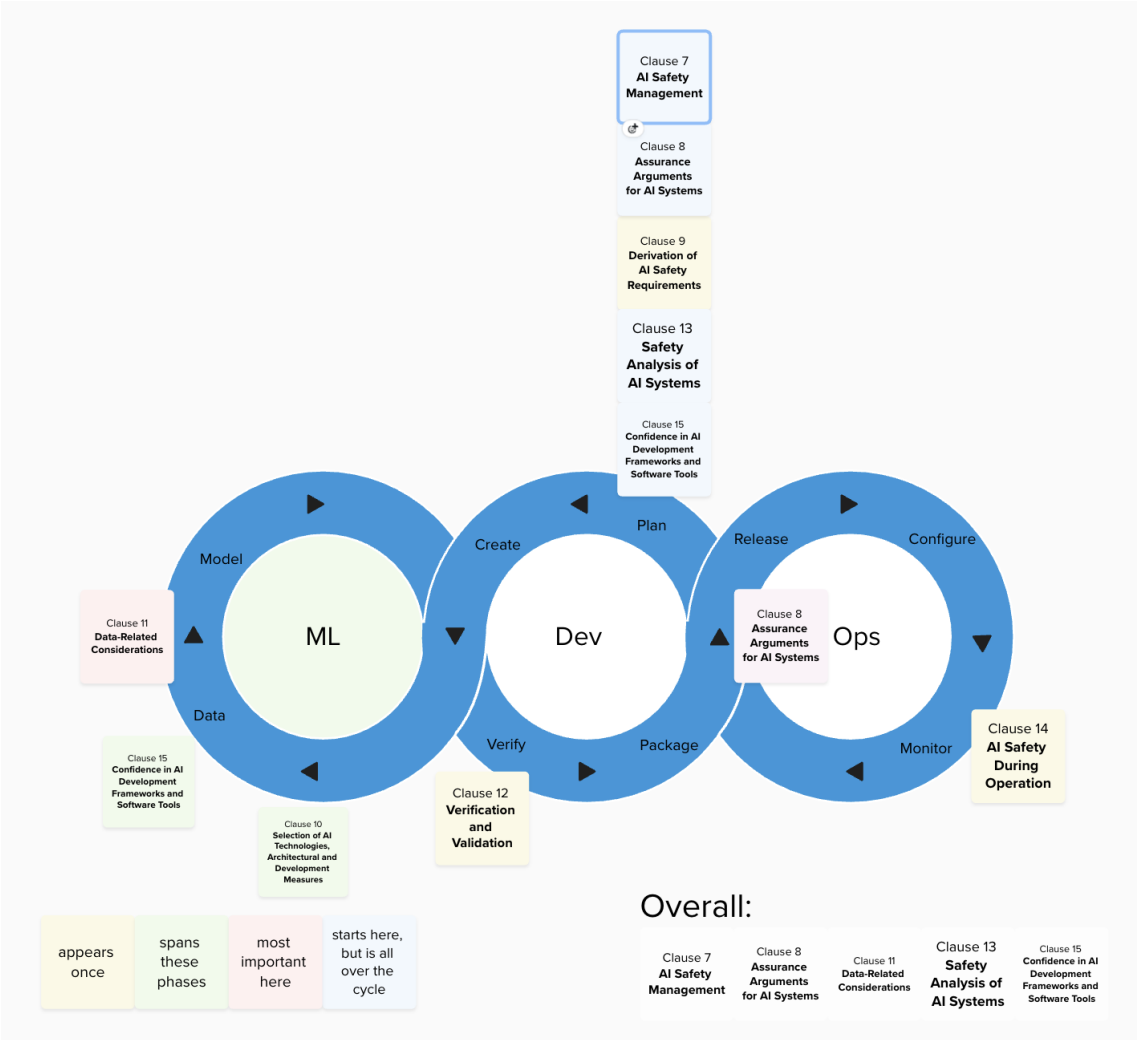

Figure B.2: The unevaluated first iteration of the MLOps artefacts based on the researchers opinions.

B.2 Second Iteration

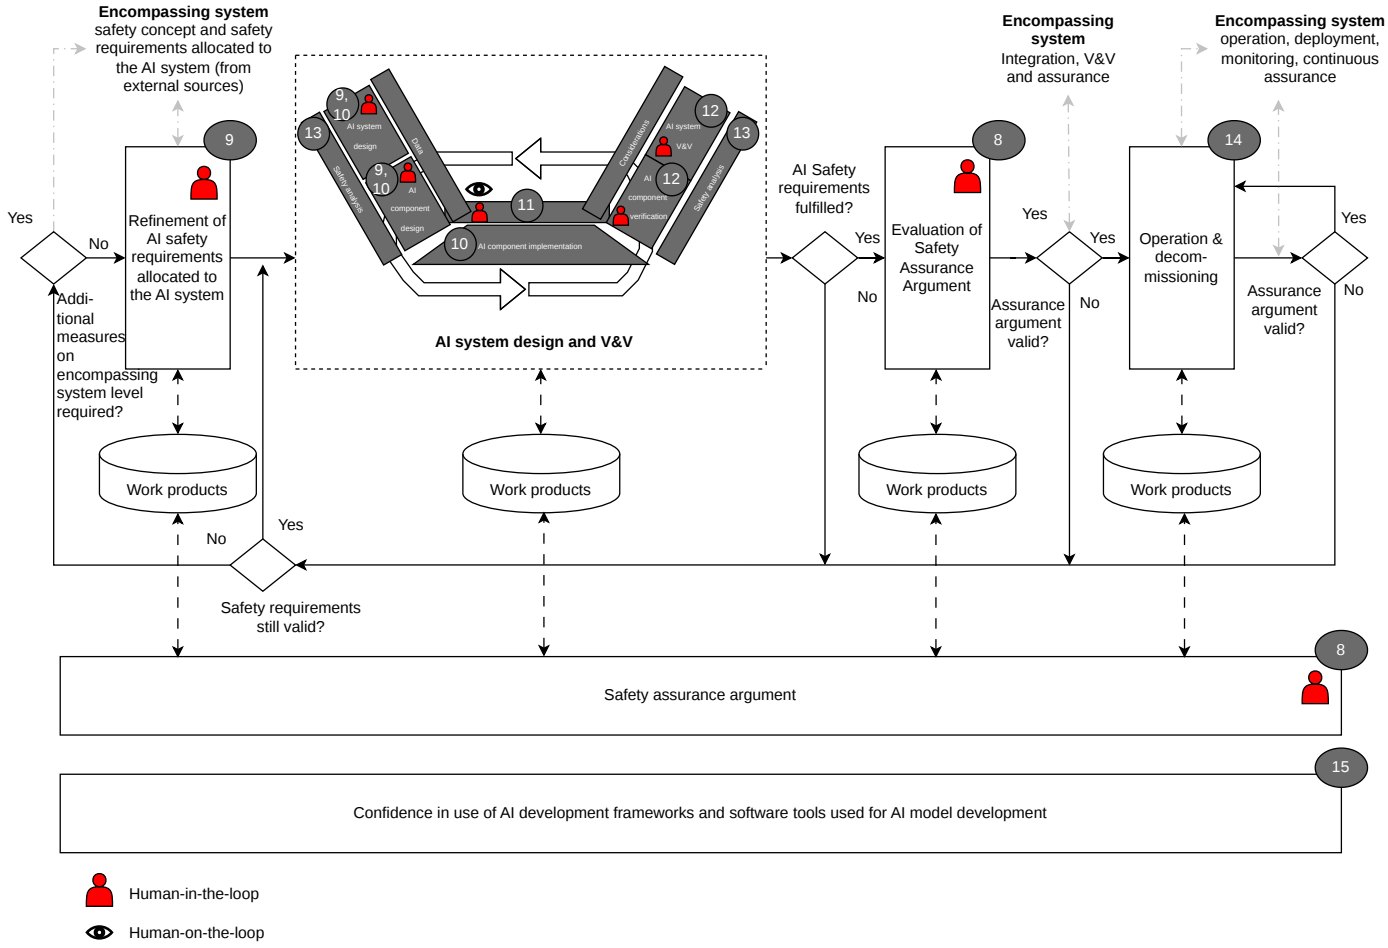

Figure B.3: The final result of AI artefact with HITL/HOTL focus. In landscape orientation for detailed view.

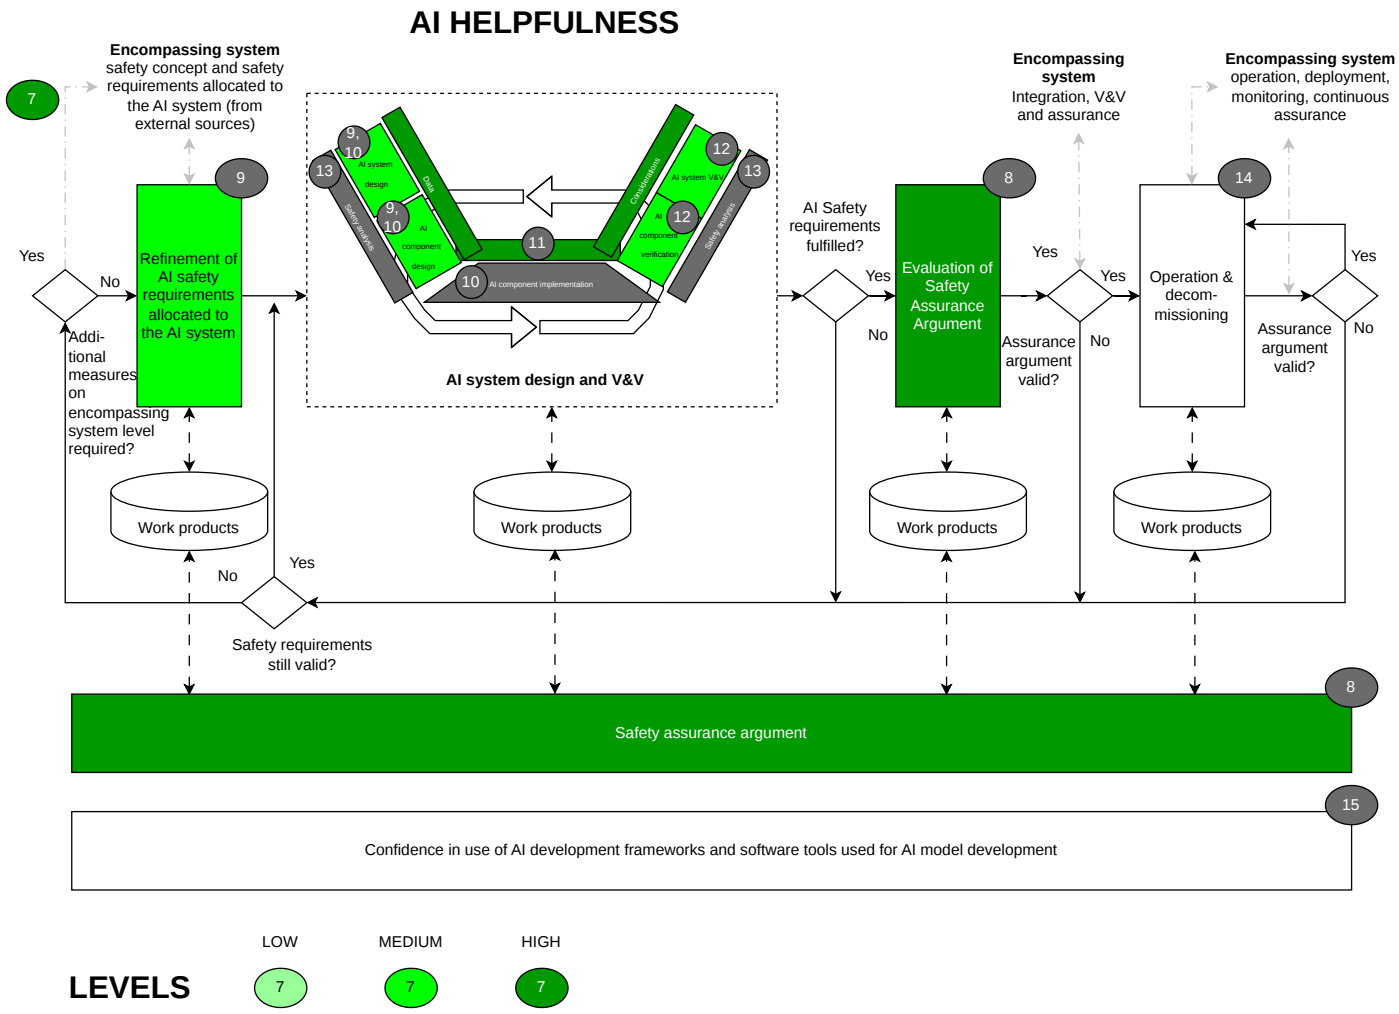

Figure B.4: The final result of AI artefact with AI helpfulness focus. In landscape orientation for detailed view.

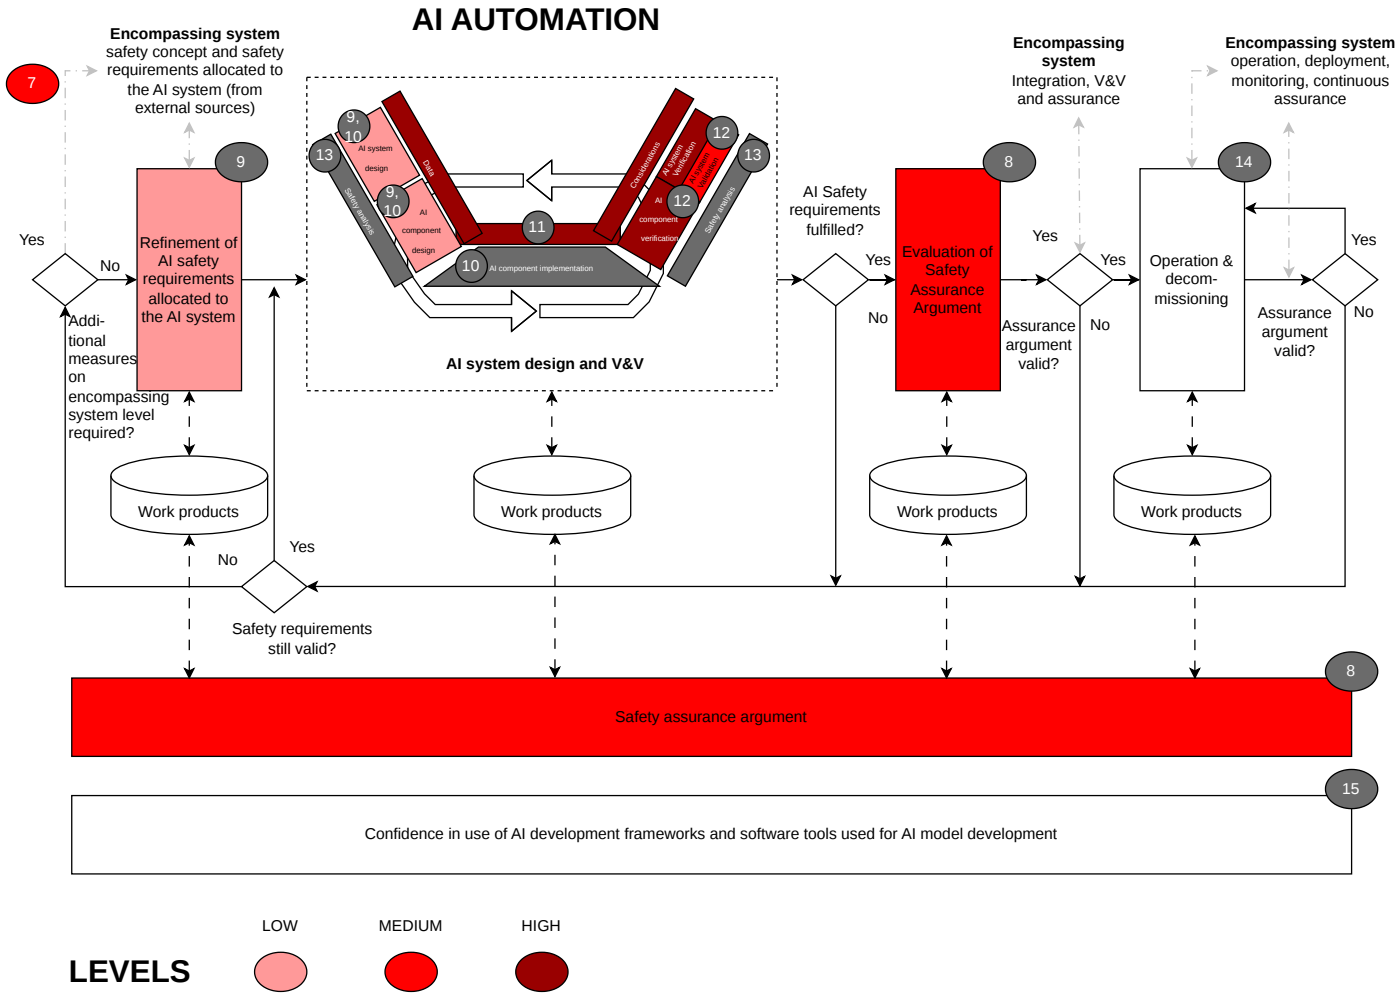

Figure B.5: The final result of AI artefact with AI Automation focus. In landscape orientation for detailed view.

C

Codebooks and Themes

C.1 Codebook 1

Table C.1: Codebook 1

Code	Subcodes	Sub-subcodes
advantages to ISO/PAS		
agile not built for safety-related systems		
automation	• early phase automation • semi-automated ECU phase	
building in safety by design		
common problem of hazard analysis in post development		
comparing running old and new models in parallel		
complexity	• hardware complexity • high complexity of the real world • integration complexity • many variants	
dependency *(parent code only, not applied independently)*	• backwards compatibility	
development time in comparison to Chinese companies		
documentation		
final stage alignment		
hardware bottleneck		
hazard analysis or risk assessment		
hierarchical project managers lack full picture		
human quality assurance still needed		
immature technologies		
industry opting for AI solutions		
industry opting for GenAI solutions		
integration *(parent code only, not applied independently)*	• systems integration	
interdependency between standard and industry		

Continued on next page

Table C.1 - continued from previous page

Code	Subcodes	Sub-subcodes
interpretation of 8800 is okay		
interviewee characteristics		
ISO/PAS 8800 gaining recent relevance		
knowledge of standards *(parent code only, not applied independently)*	• experience with 8800 • expertise area related standards background • lack of experience with 8800 • related standards background	
legal compliance	• audits • internal audits	
likelihood estimation	• society tolerance • who is responsible for deaths	
lying about compliance		
not enough knowledge about standards		
required effort in tailoring to software development		
requirements	• conflicting requirements • early adaptation of recommendations • limitations of black boxes • normative requirements • often rushed and inherited from previous projects • standard building • traceability is important	
role responsibilities		
SAFe		
safety analysis		
safety culture		
compliance doesn't equal safety		
safety goals		
scenario-based testing *(parent code only, not applied independently)*	• alternatives to standard	
scepticism towards generated code in security-sensitive areas		
software level activities to ensure safety		
solutions *(parent code only, not applied independently)*	• AI for innovation • AI for productivity • avoid prohibiting the use of new technologies • documentation automation • high quality data • iterative learning • parallel compliance • quality from the top • return to V-model work • senior software engineer perspective • test case automation • ticket automation • tool qualification and classification • usability	• approval
standard courses		

Continued on next page

Table C.1 - continued from previous page

Code	Subcodes	Sub-subcodes
standard issues	• all compliance stages in CI/CD lifecycle • burden of verification • covers only faults in the system • early compliance stages • hierarchical in agile dev • inadequate • limitations for software dev • misalignment • obsolete • taking shortcuts • unrealistic expectations • vague or non-operationalisable • wrong requirements	• terminology
syncing	• automotive specific challenges	
testing *(parent code only, not applied independently)*	• automated testing • scenario-based testing	
unknown unknowns are problematic		

C.2 Themes and Grouped Codes Phase 1

Table C.2: Themes and Grouped Codes Phase 1

Code	Subcodes	Sub-subcodes
Individual familiarity with 8800 without formalised organisational compliance		
advantages to ISO/PAS		
immature technologies		
interdependency between standard and industry		
interpretation of 8800 is okay		
interviewee characteristics		
ISO/PAS 8800 gaining recent relevance		
knowledge of standards *(parent code only, not applied independently)*	• experience with 8800 • expertise area related standards background • lack of experience with 8800 • related standards background	
legal compliance	• audits • internal audits	
not enough knowledge about standards		
requirements	• early adaptation of recommendations • normative requirements • recommendations • standard building • traceability is important	
role responsibilities		
compliance doesn't equal safety		
scenario-based testing *(parent code only, not applied independently)*	• alternatives to standard	

Continued on next page

Table C.2 - continued from previous page

Code	Subcodes	Sub-subcodes
software level activities to ensure safety		
solutions *(parent code only, not applied independently)*	• high quality data • senior software engineer perspective	
standard courses		
Misalignment in standards and real world exacerbates burden of verification, leading to shortcuts		
documentation		
lying about compliance		
required effort in tailoring to software development		
requirements *(parent code only, not applied independently)*	• conflicting requirements • often rushed and inherited from previous projects	
standard issues	• burden of verification • covers only faults in the system • inadequate • limitations for software dev • misalignment • obsolete • taking shortcuts • unrealistic expectations • vague or non-operationalisable • wrong requirements	• terminology
Subtheme: Following standards is not always compatible with agile development		
agile not built for safety-related systems		
agile SAFe		
common problem of hazard analysis in post development		
solutions *(parent code only, not applied independently)*	• return to V-model work	
standard issues *(parent code only, not applied independently)*	• hierarchical in agile dev	
Integration complexity at the software-hardware boundary, where many variants cause burden of testing		
automation *(parent code only, not applied independently)*	• semi-automated ECU phase	
complexity	• hardware complexity • high complexity of the real world • integration complexity • many variants	
dependency *(parent code only, not applied independently)*	• backwards compatibility	
hardware bottleneck		
hazard analysis or risk assessment		
integration *(parent code only, not applied independently)*	• systems integration	
safety analysis		
syncing	• automotive specific challenges	
testing *(parent code only, not applied independently)*	• scenario-based testing	

Continued on next page

Table C.2 - continued from previous page

Code	Subcodes	Sub-subcodes
Fragmented compliance issues across the CI/CD process		
solutions *(parent code only, not applied independently)*	• parallel compliance • tool qualification and classification	
standard issues *(parent code only, not applied independently)*	• all compliance stages in CI/CD lifecycle • early compliance stages	
Increased industry interest in adopting AI solutions with no end in sight		
comparing running old and new models in parallel		
development time in comparison to Chinese companies		
final stage alignment		
human quality assurance still needed		
industry opting for AI solutions		
industry opting for GenAI solutions		
scepticism towards generated code in security-sensitive areas		
solutions *(parent code only, not applied independently)*	• AI for innovation • AI for productivity	
Subtheme: Black-boxes are inevitable and necessitate acceptance from industry and society		
likelihood estimation	• society tolerance • who is responsible for deaths?	
requirements *(parent code only, not applied independently)*	• limitations of blackboxes	
solutions *(parent code only, not applied independently)*	• avoid prohibiting the use of new technologies • hybrid modelling • iterative learning	
unknown unknowns are problematic		
Subtheme: Development automation to reduce compliance burden		
automation	• early phase automation	
solutions *(parent code only, not applied independently)*	• documentation automation • test case automation • ticket automation • usability	• approval
testing *(parent code only, not applied independently)*	• automated testing	
Safety focused development stems from the top of the hierarchy		
building in safety by design		
hierarchical project managers lack full picture		
safety culture		
safety goals		
solutions *(parent code only, not applied independently)*	• quality from the top	

C.3 Codebook 2

Table C.3: Codebook 2

Code	Subcodes
AI can't be legally responsible	
AI doesn't have motives like humans	
AI help *(parent code only, not applied independently)*	• broaden perspective • collecting evidence • data labelling • documentation • first draft • impact analysis • logic check at each step • low to high level safety goal • MCP server • quality perspective • regression test
AI should be in charge	
AI should never be completely in charge	
AI should not hallucinate	
challenge *(parent code only, not applied independently)*	• tool-chain confidence report
clause in the right place	• continuous • map work products/requirements to stages
clause should be in additional phase or phases	
current use of AI tools *(parent code only, not applied independently)*	• AI in standard tools • Claude • Copilot • Cursor • gather information • generate test cases • Gemini • linguistic usage • no usage of agentic AI
degree of AI automation *(parent code only, not applied independently)*	• high • low • medium
degree of AI helpfulness *(parent code only, not applied independently)*	• high • low • medium
degree of AI use *(parent code only, not applied independently)*	• high • low • medium
feelings about AI in safety-critical development *(parent code only, not applied independently)*	• cautious • concerned • excited • mixed feelings • worried of losing job
human involvement *(parent code only, not applied independently)*	• control over algorithm • control over process
human should be in charge	
human should take accountability	
human supervision *(parent code only, not applied independently)*	• human in the loop • human on the loop

Continued on next page

Table C.3 - continued from previous page

Code	Subcodes
input	• manual
interviewee characteristics	
knowledge of standards *(parent code only, not applied independently)*	• experience with 8800 • expertise area related standards background
legislation on AI	
likelihood estimation	• society tolerance
no opinion about clause placement	
people opting for automation over safety	
risk *(parent code only, not applied independently)*	• AI and humans can both make mistakes • AI missing important aspects • bias • falling behind in AI adoption • hard to simulate complexity of the real world • hardware-in-the-loop testing • humans can sometimes make more mistakes than AI • incorrect labelling • input • people stop thinking critically • security of data • time estimation can be difficult for AI • trusting insufficient AI output
role responsibilities	
separation of human and AI generated content	
system architecture	• AI tool for initial manual work
wouldn't trust AI completely	• immature technology

C.4 Themes and Grouped Codes Phase 2

Table C.4: Themes and Grouped Codes Phase 2

Code	Subcodes
Widespread but shallow AI tool adoption, with perceived potential across safety-critical tasks	
AI help *(parent code only, not applied independently)*	• broaden perspective • collecting evidence • data labelling • documentation • first draft • impact analysis • logic check at each step • low to high level safety goal • MCP server • quality perspective • regression test
challenge *(parent code only, not applied independently)*	• tool-chain confidence report

Continued on next page

Table C.4 - continued from previous page

Code	Subcodes
current use of AI tools *(parent code only, not applied independently)*	• AI in standard tools • Claude • Copilot • Cursor • gather information • Gemini • generate test cases • linguistic usage • no usage of agentic AI
system architecture	• AI tool for initial manual work
Universal openness to AI integration across tasks, with divergence on acceptable levels of automation and autonomy	
degree of AI automation *(parent code only, not applied independently)*	• high • low • medium
degree of AI helpfulness *(parent code only, not applied independently)*	• high • low • medium
degree of AI use *(parent code only, not applied independently)*	• high • low • medium
Many acknowledged risks of AI in safety-critical development, shared fallibility of humans and AI, with humans currently as the lesser risk	
likelihood estimation	• society tolerance
risk	• AI and humans can both make mistakes • AI missing important aspects • bias • falling behind in AI adoption • hard to simulate complexity of the real world • hardware-in-the-loop testing • humans can sometimes make more mistakes than AI • incorrect labelling • input • people stop thinking critically • security of data • time estimation can be difficult for AI • trusting insufficient AI output
input	• manual
Insistence on human-in-the-loop, driven by human accountability	
AI can't be legally responsible	
AI should be in charge	
human involvement *(parent code only, not applied independently)*	• control over algorithm • control over process
human should be in charge	
human should take accountability	
human supervision *(parent code only, not applied independently)*	• human in the loop • human on the loop
interviewee characteristics	
role responsibilities	
knowledge of standards *(parent code only, not applied independently)*	• experience with 8800 • expertise area related standards background
legislation on AI	

Continued on next page

Table C.4 - continued from previous page

Code	Subcodes
People have mixed feelings about AI. Distrust in current AI for its immaturity is prevalent.	
AI doesn't have motives like humans	
AI should never be completely in charge	
AI should not hallucinate	
feelings about AI in safety-critical development *(parent code only, not applied independently)*	• cautious • concerned • excited • mixed feelings • worried of losing job
people opting for automation over safety	
separation of human and AI generated content	
wouldn't trust AI completely	• immature technology
Codes not in any theme	
clause in the right place	• continuous • map work products/requirements to stages
clause should be in additional phase or phases	
no opinion about clause placement	

D

Codes and Supporting Quotes

D.1 Codes and Supporting Quotes for Phase 1

Table D.1: Quotes: Individual familiarity with 8800 without formalised organisational compliance

Code	Supporting Quotes
ISO/PAS 8800 gaining recent relevance	"[...] it was mentioned that they were thinking about ISO 8800 [...] But I think it's a big topic for [Company] and we should address it." *Participant 11*
compliance doesn't equal safety	"If we are talking about issues of complying with ISO 26262, it's the same issues. But if you are talking about making cars safe, that's a different thing." *Participant 6*
interdependency between standard and industry	"Basically, there is a chicken and egg thing with these standards." *Participant 5*
advantages to ISO/PAS	"[...] another advantage with publicly available specification is that even if you write the best of our standard with all companies [...], at the end, maybe even my colleague in my company think I was wrong. So they will provide comments [...] give the society some years. Maybe [...] a professor who is not part of that will read it and say that 'Oh, this is wrong. It should be written in another way'. And then [...] people publish papers on that [...] and then you publish the second version, the first version of the final version." *Participant 5*

Table D.2: Quotes: Misalignment in standards and real world exacerbates burden of verification, leading to shortcuts

Code	Supporting Quotes
standard issues > misalignment	"[...] it can say [...] some senseless requirement: Do software architectural analysis in the fashion of FMEA [...]. In software, it doesn't work that way. In software, you usually have millions of failure modes [...] You cannot really do a comprehensive analysis there and provide like FMEA for software." *Participant 6*
standard issues > burden of verification	"You can't really comply safety case for these thousands of updates. It just doesn't make sense. So, and in audit, people usually understand and are reasonable. So generally you show evidence that you are doing some job." *Participant 6*

Continued on next page

Table D.2 - continued from previous page

Code	Supporting Quotes
	"And you also need to make sure that it is reviewed by someone else, and that generally takes longer than quickly making a change without doing all these things, without being that thorough, which means it is often perceived as having to do a lot of extra documentation." *(Translated from original language) Participant 2*
lying about compliance	"Nobody does that, even though everybody claims here, we are ISO 26262 compliant." *Participant 6*

Table D.3: Quotes: Following standards is not always compatible with agile development

Code	Supporting Quotes
agile not built for safety-related systems	"You can't work agile. You can't at a low level decide that 'we're going to have five wheels', and then there's someone else at a higher level who... I mean, a car is hierarchical. You should work agile towards the customer, but then you can't work agile." *(Translated from original language) Participant 3*
standard issues > hierarchical in agile dev	"[...] another problem you have is that you try to introduce the standard into companies that don't work so hierarchically, while the entire standard is hierarchical [...] it states directly that the structure should be hierarchical." *(Translated from original language) Participant 3*
	"Where I have worked, it has traditionally not been driven from an ISO 26262 perspective [...] a lot of focus has been placed on having 'working software' over documentation, which means they have perhaps chosen to do the documentation at the end [...] they have perhaps deprioritised writing requirements [...] until the end [...]. And when you try to correct that [...] ISO 26262 [...] is seen as an obstacle because it slows down the pace or efficiency." *(Translated from original language) Participant 2*
common problem of hazard analysis in post development	"[...] the ISO standard requires that you actually deliver documents on time. But you cannot do that if the functionality does not exist. It is a fairly common problem that first they sit and program, and then the systems group is supposed to come in and design or finalise analyses. Doing hazard analysis after the fact is reprehensible." *(Translated from original language) Participant 3*

Table D.4: Quotes: Integration complexity at the software-hardware boundary, where many variants cause burden of testing

Code	Supporting Quotes
complexity > hardware complexity	"For instance, in hardware parts, hardware is a really complex thing. It's a material. And there are transistors [...] they are very, very hard to make. So it's a very advanced technology." *Participant 5*
complexity > integration complexity	"With respect to a software defined vehicle you would then see a more horizontal approach. Maybe then a zonal architecture with four to five zones per vehicle and with a much higher degree of integration which requires let's say a kind of what we at [Company] tell a model-based systems engineering approach. Really integrating hardware and software development." *Participant 9*

Continued on next page

Table D.4 - continued from previous page

Code	Supporting Quotes
complexity > many variants	"There are two big things affecting this, even making things even more difficult with this complexity of software[...]. The first thing that comes with it and makes things worse is having too many variants, because you cannot test that many variants. You have to take risks." *Participant 6*
dependency > backwards compatibility	"After five years they come, they say this issue has gone bad, it doesn't work. You cannot really change that single ECU with its software. The software version has to work with the rest of the software in the other ECUs. So it has become too interconnected. So you have to somehow have a baseline of software that is compatible for all the ECUs [...]. It's very, very difficult to get that right [...] software complexity [...] makes the entire car or truck interconnected. Then on top of that, way too many variants and backwards compatibility [...]." *Participant 6*

Table D.5: Quotes: Fragmented compliance issues across the CI/CD process

Code	Supporting Quotes
standard issues > all compliance stages in CI/CD lifecycle	"What I would say is that it actually arises throughout the entire chain." *(Translated from original language) Participant 2*
standard issues > early compliance stages	"So early software development stages where you have pure software, most of the problems come up already there." *Participant 6*
solutions > parallel compliance	"It's always good to have it. I mean together with the lifecycle rather than having it at the end of the lifecycle. It should be there from the start of the lifecycle. That's how I feel." *Participant 1*

Table D.6: Quotes: Increased industry interest in adopting AI solutions with no end in sight

Code	Supporting Quotes
industry opting for GenAI solutions	"Still, I think companies are forced to use such tools in at least some areas, because you save so much money with them." *Participant 10*
industry opting for AI solutions	"[...] they are looking at actively using AI in SDV development processes. And then some areas they are looking at a big productivity improvement. And including regulatory compliance, like ISO 26262 or other ISO standards as well." *Participant 12*
solutions > AI for productivity	"Mostly, AI usage is discussed in relation to productivity improvements, but it is also being considered for innovation." *Participant 12*
development time in comparison to Chinese companies	"The problem is the long time that is needed for development [...] compare the development times of Chinese companies and European companies for a car. In Europe, the usual amount of time is three to four years, more four than three. In China, they need about two years. And you can ask yourself why they are so fast. [...] And regarding whether they use Gen AI [...]." *Participant 10*

Table D.7: Quotes: Black-boxes are inevitable and necessitate acceptance from industry and society

Code	Supporting Quotes
likelihood estimation > society tolerance	"I think we have to also consider that we as humans have [...] an extreme bias against automation in vehicles [...] we are much more [...] forgiving of essentially the same accident occurring if a person were [...] driving [...]. So I think we as a society need to somehow figure out what is good enough. If we basically say [...] probabilistic systems are the only way to go, then [...] the question becomes what is good enough versus [...] the only thing that is good enough is zero chance of these things happening." *Participant 8*
	"All technology always has an inherent risk. Then the question is what risk you can accept." *(Translated from original language) Participant 4*
requirements > limitations of blackboxes	"ISO 26262 is more deterministic and traceable, but when you come to ISO/PAS 8800, it is non-deterministic. You don't know where the failure can happen [...] the tools that are used in the CI/CD tool chain may introduce new errors." *Participant 1*
	"Still, we don't know what the limitations of LLMs are. We know that they hallucinate [...] So, writing requirements on those technologies [...] might be a bit challenging." *Participant 5*
solutions > avoid prohibiting the use of new technologies	"Solutions, to me, the main solution is that you for sure don't prohibit to use any of these technologies. Use them, although they are immature, but at the same time use them in a responsible way." *Participant 5*

Table D.8: Quotes: Development automation to reduce compliance burden

Code	Supporting Quotes
solutions > documentation automation	"I actually believe that AI could help formulate requirements in a better way, because from what I have seen with requirements writing at both the system level and software level, you tend to get a little lazy when writing your requirements." *(Translated from original language) Participant 2*
	"That's one of the goals of a proper CI/CD DevOps tool chain to really address these topics and automatically generate all the necessary reports you need for compliance [...] the goal is to automate all the stuff and reporting." *Participant 11*

Table D.9: Quotes: Safety focused development stems from the top of the hierarchy

Code	Supporting Quotes
building in safety by design	"I think it requires a lot of a return to fundamental work on sort of theory and applying that to training models and also figuring things out on the sort of like architecture side, so building in safety by design rather than as an add-on at the end." *Participant 8*
safety culture	"And risks I have experienced. There is always a risk that it becomes acrobatically overdone with too much documentation. But at the same time, documentation is important. It is very difficult, and you need managers who understand and want to do this. That is what is called a Safety culture. It is very important to have. Otherwise it just becomes like showing that we follow this standard on paper." *(Translated from original language) Participant 3*

Continued on next page

Table D.9 - continued from previous page

Code	Supporting Quotes
solutions > quality from the top	"It is very difficult to introduce quality work from the bottom up. It almost always leads to burning people out. Quality must come from the top. And if you don't have management with you saying that this is important, then it is a swamp." *(Translated from original language) Participant 3*

D.2 Codes and Supporting Quotes for Phase 2

Table D.10: Quotes: Widespread but shallow AI tool adoption, with perceived potential across safety-critical tasks

Code	Supporting Quotes
AI help > documentation	"I think it will definitely help expedite the creation of the artifacts." *Participant 15*
AI help > data labelling	"Data labelling has been a human task for many years, but now it can be automated in very effective ways." *Participant 14*
current use of AI tools > gather information	"So I use Gemini both to gather information and to refine text." *Participant 14*

Table D.11: Quotes: Universal openness to AI integration across tasks, with divergence on acceptable levels of automation and autonomy

Code	Supporting Quotes
degree of AI helpfulness > high	"Yes, yes, yes, very much because this one also is too much documentation and takes a lot of time and effort. [...] And it could easily be done using AI and there is less risk to have something missed in this one, yeah." *Participant 13*
degree of AI automation > high	"So it's actually no harm to safety or humans here. So, I could accept a higher level of automation here." *Participant 16*

Table D.12: Quotes: Many acknowledged risks of AI in safety-critical development, shared fallibility of humans and AI, with humans currently as the lesser risk

Code	Supporting Quotes
risk > AI missing important aspects	"The risk is that AI might miss or fail to understand the full scope of all the safety issues." *Participant 15*
	"Because I don't think AI retain that history and that knowledge and that depth of reasoning that humans do." *Participant 15*
risk > hard to simulate complexity of the real world	"For example, if you remember the Waymo cases, the systems were trained never to cross a white line because doing so could put them into oncoming traffic and cause a collision. Then people went out and drew white lines around the vehicle, which caused it to get stuck." *Participant 15*
	"The question is how accurately the AI tool can mimic the real situation. That is one risk." *Participant 13*

Continued on next page

Table D.12 - continued from previous page

Code	Supporting Quotes
risk > humans can sometimes make more mistakes than AI	"So what usually people do is copy-pasting, which is much more dangerous than telling AI a context [...]. It's much more effective than copy-pasting things and much less error-prone on my evaluation." *Participant 14*
trusting insufficient AI output	"[...] most of the time the tool simply generates some output, which should not be the case. If someone starts trusting that output, it can cause serious problems in the end." *Participant 13*

Table D.13: Quotes: Insistence on human-in-the-loop, driven by human accountability

Code	Supporting Quotes
legislation on AI	"That is how the law works right now. We currently have no legal framework that defines how an artificial intelligence system can be held accountable." *Participant 14*
human should take accountability	"At the moment, we trust people, and we know how to make them accountable. We do not yet know how to make AI accountable [...]." *Participant 14*
	"You need someone to be accountable and also an organization to be accountable." *Participant 13*

Table D.14: Quotes: People have mixed feelings about AI. Distrust in current AI for its immaturity is prevalent

Code	Supporting Quotes
AI doesn't have motives like humans	"[...] an AI does not want anything. I, as a human, want something. I want my company to be competitive on the market with the great products. So on top of things, there has to be a human in my mind." *Participant 16*
feelings about AI in safety-critical development > concerned	"Concerned. [...] But the risks associated with using it without thinking is high. [...] I can ask it to do a complete hazard analysis. And I will get something that looks like it has an analysis. But when you start reading it, every sentence is totally fine but it might be bogus. This is really concerning for me because it looks looks great, reads great at the fast glance. But the technical content and the technical correctness can sometimes be very wrong." *Participant 16*
feelings about AI in safety-critical development > excited	"It's an evolution just as part of our lives. I think we should embrace it. And I think that we should try to see how it can help us the most to camp. So I'm excited. I'm all for it." *Participant 15*
wouldn't trust AI completely > immature technology	"At the end of the day, it is a question of trust, and that is very common in safety-related topics. We are not at a point where we can fully trust AI yet. Maybe we will be in the future, but we are not there now." *Participant 14*