

# Ergo, SMIRK is Safe: A Safety Case for a Machine Learning Component in a Pedestrian Automatic Emergency Brake System

Markus Borg<sup>1,2\*</sup>, Jens Henriksson<sup>3</sup>, Kasper Socha<sup>1,2</sup>, Olof Lennartsson<sup>4</sup>, Elias Sonnsjö Lönegren<sup>4</sup>, Thanh Bui<sup>1</sup>, Piotr Tomaszewski<sup>1</sup>, Sankar Raman Sathyamoorthy<sup>5</sup>, Sebastian Brink<sup>6</sup> and Mahshid Helali Moghadam<sup>1</sup>

<sup>1</sup>\*RISE Research Institutes of Sweden, Lund, Sweden.

<sup>2</sup>Dept. of Computer Science, Lund University, Lund, Sweden.

<sup>3</sup>Semcon AB, Gothenburg, Sweden.

<sup>4</sup>Infotiv AB, Gothenburg, Sweden.

<sup>5</sup>QRTECH AB, Gothenburg, Sweden.

<sup>6</sup>Combitech AB, Gothenburg, Sweden.

\*Corresponding author(s). E-mail(s): [markus.borg@ri.se](mailto:markus.borg@ri.se);

Contributing authors: [jens.henriksson@semcon.com](mailto:jens.henriksson@semcon.com);

[kasper.socha@ri.se](mailto:kasper.socha@ri.se); [olof.lennartsson@infotiv.se](mailto:olof.lennartsson@infotiv.se);

[elias.sonnsjo@infotiv.se](mailto:elias.sonnsjo@infotiv.se); [thanh.bui@ri.se](mailto:thanh.bui@ri.se);

[piotr.tomaszewski@ri.se](mailto:piotr.tomaszewski@ri.se); [sankar.sathyamoorthy@qrtech.se](mailto:sankar.sathyamoorthy@qrtech.se);

[sebastian.brink@combitech.com](mailto:sebastian.brink@combitech.com); [mahshid.helali.moghadam@ri.se](mailto:mahshid.helali.moghadam@ri.se);

## Abstract

Integration of Machine Learning (ML) components in critical applications introduces novel challenges for software certification and verification. New safety standards and technical guidelines are under development to support the safety of ML-based systems, e.g., ISO 21448 SOTIF for the automotive domain and the Assurance of Machine Learning for use in Autonomous Systems (AMLAS) framework. SOTIF and AMLAS provide high-level guidance but the details must be chiseled out for each specific case. We report results from an industry-academia collaboration on safety assurance of SMIRK, an ML-based pedestrian

automatic emergency braking demonstrator running in an industry-grade simulator. We present the outcome of applying AMLAS on SMIRK for a minimalistic operational design domain, i.e., a complete safety case for its integrated ML-based component. Finally, we report lessons learned and provide both SMIRK and the safety case under an open-source licence for the research community to reuse.

**Keywords:** machine learning safety, safety standards, safety case, automotive demonstrator

## 1 Introduction

Machine Learning (ML) is increasingly used in critical applications, e.g., supervised learning using Deep Neural Networks (DNN) to support automotive perception. Software systems developed for safety-critical applications must undergo assessments to demonstrate compliance with functional safety standards. However, as the conventional safety standards are not fully applicable for ML-enabled systems (Salay et al, 2018; Tambon et al, 2022), several domain-specific initiatives aim to complement them, e.g., organized by the EU Aviation Safety Agency, the ITU-WHO Focus Group on AI for Health, and the International Organization for Standardization.

In the automotive industry, several standardization initiatives are ongoing to allow safe use of machine learning in road vehicles. It is evident that the established functional safety as defined in ISO 26262 Functional Safety (FuSa) is no longer sufficient for the next generation of Advanced Driver-Assistance Systems (ADAS) and Autonomous Driving (AD). One complementary standard under development is ISO 21448 Safety of the Intended Functionality (SOTIF) (ISO, 2019). SOTIF is a standard that aims for absence of unreasonable risk due to hazards resulting from functional insufficiencies – also for systems that rely on ML.

Standards such as SOTIF mandate high-level requirements on what a development organization must provide in a safety case for an ML-based system. However, how to actually collect the evidence – and argue that it is sufficient – is up to the specific organization. Assurance of Machine Learning for use in Autonomous Systems (AMLAS) is one framework that supports the development of corresponding safety cases (Hawkins et al, 2021). Still, when applying AMLAS on a specific case, there are numerous details that must be analyzed, specified, and validated. The research community lacks demonstrator systems that can be used to explore such details.

We report results from an industry-academia collaboration on safety assurance of SMIRK, an ML-based ADAS that provides Pedestrian Automatic Emergency Braking (PAEB) in an industry-grade simulator. This work has

been completed as part of the Swedish research project SMILE III<sup>1</sup>. The contributions of this paper are two-fold. First, we introduce SMIRK, a demonstrator available on GitHub under an Open-Source Software (OSS) license – ready to be reused by others in the research community. Second, we report the outcome from a complete application of AMLAS on SMIRK. While parts of the framework have been demonstrated before [Gauerhof et al \(2020\)](#), we believe this is the first comprehensive use of the framework conducted independently from its authors. Moreover, we believe this paper constitutes a pioneering safety case for an ML-based component that is OSS and completely transparent. Thus, our contribution can be used as a starting point for studies on safety engineering aspects such as Operational Design Domain (ODD) extension, dynamic safety cases, and reuse of safety evidence.

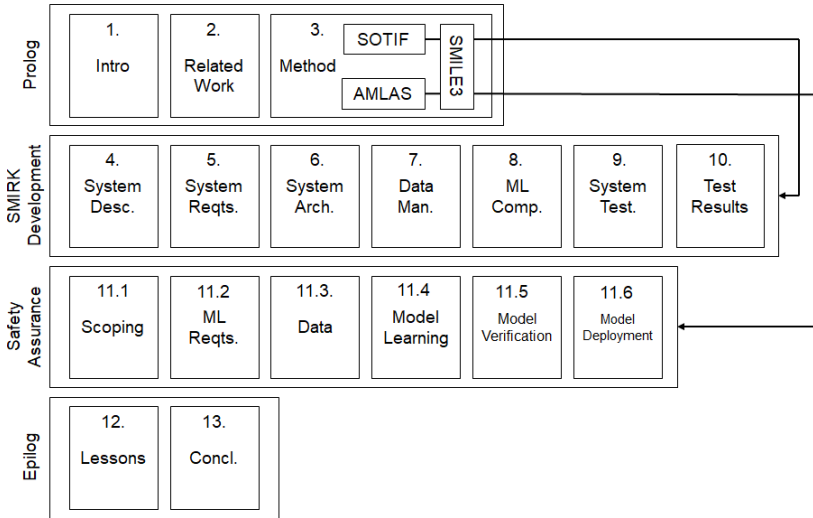
Our results show that even an ML component in an ADAS designed for a minimalistic ODD results in a large safety case. Furthermore, we consider three lessons learned to be particularly important for the community. First, using a simulator to create synthetic data sets for ML training particularly limits the validity of the negative examples. Second, evaluation of object detection is non-intuitive and necessitates internal training. Third, the fitness function used for model selection encodes essential tradeoff decisions, thus the project team must be aligned.

Figure 1 presents the structure of this long article. We are aware that we include a large piece of work in a single publication unit. Still, we find that a largely self-contained paper presenting a comprehensive safety case is missing in the research community. As opposed to the “salami publication” anti-pattern in academic publishing, we choose to present both SMIRK and its complete safety case in the same article. The figure shows the four main parts of this paper.

- **Prolog:** The first part of the paper consists of Sections 1–3, i.e., this introduction, a brief overview of related work, and a method section. The method section contains three parts, i.e., SOTIF, AMLAS, and how we applied them in the SMILE III project. As presented by arrows, SOTIF influenced the SMIRK development and we relied on AMLAS for the safety assurance.
- **SMIRK:** The second part describes the ML-based ADAS under study. We present an overall system description, system requirements, the system architecture, the data management strategy, the ML-based pedestrian recognition component, and the approach to testing in Sections 4–9, respectively. Finally, we present an overview of the SMIRK test results in Section 10.
- **Safety Assurance:** Section 11 presents how we apply the six stages of AMLAS to construct a safety case for SMIRK.
- **Epilog:** The final part of the paper presents lessons learned and limitations in Section 12 before we conclude the paper and outline directions for future work in Section 13.

---

<sup>1</sup><https://tinyurl.com/smileiii>

**Fig. 1** Paper organization.

## 2 Related Work

Many researchers argue that software and systems engineering practices must evolve as ML enters the picture, reflected by the organization of the first International Conference on AI Engineering in 2022<sup>2</sup>. Pioneering work in the dawn of AI engineering includes a research agenda by Bosch et al (2021), an analysis on novel best practices by Serban et al (2020), and the evolving book “Machine Learning in Production” by Kästner (2022). In this section, we focus on aspects of AI engineering for safety-critical ML-based automotive systems, i.e., safety argumentation in general and DNNs used for automotive perception in particular. Moreover, we stress that the remainder of the paper contains several references to related work as we discuss various design choices.

Many publications address the issue of safety argumentation for systems with ML-based components. A solid argumentation is required to enable safety certification, for example to demonstrate compliance with future standards such as SOTIF and ISO 8800 Road Vehicles – Safety and Artificial Intelligence. While there are several established safety patterns (e.g., simplicity, substitution, sanity check, condition monitoring, comparison, diverse redundancy, replication redundancy, repair, degradation, voting, override, barrier, and heartbeat Wu and Kelly (2004); Preschern et al (2015)), considerable research is now directed at understanding what is needed in the ML era. Our previous work provides an overview of verification and validation of DNN-based systems, including a challenge elicitation with the Swedish automotive industry (Borg et al, 2019).

We have found two review studies that focus on ML and safety certification. Schwalbe and Schels (2020) present an *ad hoc* overview on methods

<sup>2</sup><https://conf.researchr.org/track/cai-2022>

that support safety argumentation for ML-based systems, organized into the phases 1) requirements engineering, 2) development, 3) verification, and 4) validation. For each phase, the authors present example methods from the literature. [Tambon et al \(2022\)](#) present a systematic literature review covering 217 primary studies. The authors investigate fundamental topics such as robustness, uncertainty, explainability, and verification – and calls for deeper industry-academia collaborations. This paper responds to this call and explicitly targets the listed fundamental topics on an operational level. As the devil is in the detail, we recommend additional researches of this nature. By conducting hands-on development of an ADAS and its corresponding safety case, we have identified numerous design decisions that have not been discussed in prior work.

Not only AMLAS provides structured methods to support ML safety argumentation. [Picardi et al \(2020\)](#) present a set of patterns that can be used to develop assurance arguments for demonstrating the safety of the ML components. The argument patterns provide reusable templates for the types of claims that must be made in a compelling argument. [Kochanthara et al \(2021\)](#) propose a safety assessment method on the systems-of-systems level, i.e., for cooperative driving systems. While the method does not target ML specifically, it discusses the context of platooning with a manually driven lead vehicle is followed by autonomous vehicles – a solution that most likely requires ML-based perception. [Schwalbe et al \(2020\)](#) systematically establish and break down safety requirements to argue the sufficient absence of risk arising from such SOTIF-style functional insufficiencies. The authors stress the importance of diverse evidence for a safety argument involving DNNs. Moreover, they provide a generic approach and template to thoroughly respect DNN specifics within a safety argumentation structure. Finally, the authors show its applicability for an example use case based on pedestrian detection.

Just like [Schwalbe et al \(2020\)](#), several researchers choose pedestrian detection systems to illustrate different approaches to safety argumentation. [Wozniak et al \(2020\)](#) provide a safety case pattern for ML-based systems and showcase its applicability on a pedestrian avoidance system. The pattern is integrated within an overall encompassing approach for safety case generation. [Willers et al \(2020\)](#) discusses safety concerns for DNN-based automotive perception, including technical root causes and mitigation strategies. The authors argue that it remains an open question how to conclude whether a specific concern is sufficiently covered by the safety case – and stress that safety cannot be determined analytically through ML accuracy metrics. In our work on SMIRK, we provide safety evidence that goes beyond the level of the ML model. Related to pedestrian detection, we find that the work by [Gauerhof et al \(2020\)](#) is the closest to this study, and the reader will find that we repeatedly refer to it throughout the text.

In the current paper, not only do we present a holistic safety case building on previous work, we also present the demonstrator system under an OSS license. In contrast to most previous work that stop at pedestrian detection,

we present an ADAS that subsequently commences emergency braking in a simulated environment. This addition responds to calls for researchers to go from offline to online testing (Haq et al, 2021), as many safety violations identified by online testing could not be identified by offline testing. We hope that SMIRK can contribute to a shift in the testing community away from standalone image data sets.

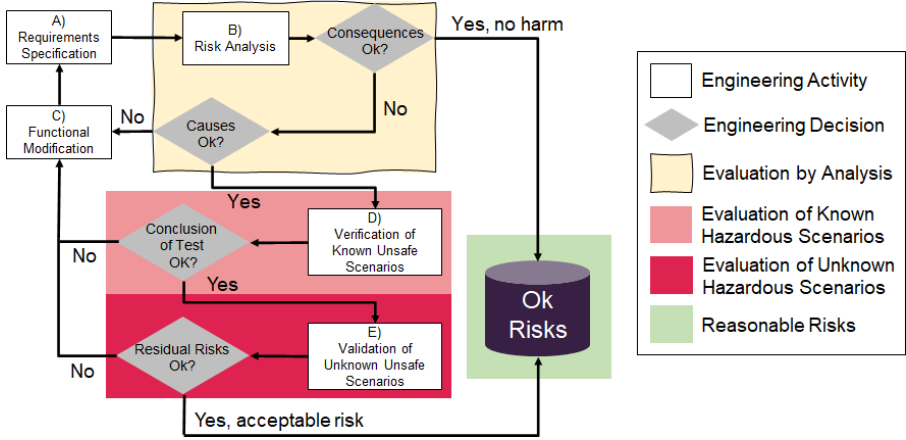
### 3 Method

The overall frame of our work is the *engineering research* standard as defined in the evolving community endeavor ACM SIGSOFT Empirical Standards (Ralph et al, 2020). Engineering research is an appropriate standard when evaluating technological artifacts, e.g., methods, systems, and tools – in our case SMIRK and its safety case. To support the validity of our research, we consulted the essential attributes of the corresponding checklist. While most attributes are clearly addressed in this manuscript, we provide three clarifications: 1) empirical evaluations of SMIRK are done using simulation in ESI Pro-SiVIC, 2) empirical evaluation of the safety case has been done through workshops and peer-review, and 3) we compare the SMIRK safety case against state-of-the-art implicitly by building on previous work. In this section, we first present our interpretations of SOTIF and AMLAS. Second, we present an overview of our ways of working in the development project.

#### 3.1 Systems Development Using SOTIF

The SMIRK development followed the process in ISO 21448 Safety of the Intended Functionality (SOTIF). SOTIF is a candidate standard under development to complement the established automotive standard ISO 26262 Functional Safety (FuSa). While FuSa covers hazards caused by *malfunctioning behavior*, SOTIF addresses hazardous behavior caused by the intended functionality. A system that meets FuSa can still be hazardous due to insufficient environmental perception or inadequate robustness within the ODD. The SOTIF process provides guidance on how to systematically ensure the absence of unreasonable risk due to *functional insufficiencies*. The goal of the SOTIF process is to perform a risk acceptance evaluation and then reduce the probability of 1) known and 2) unknown scenarios causing hazardous behavior.

Figure 2 shows a simplified version of the SOTIF process. The process starts in the upper left with A) Requirements specification. Based on the requirements, a B) Risk Analysis is done. For each identified risk, its potential *Consequences* are analyzed. If the risk of harm is reasonable, it is recorded as an acceptable risk. If not, the activity continues with an analysis of *Causes*, i.e., an identification and evaluation of triggering conditions. If the expected system response to triggering conditions is acceptable, the SOTIF process continues with V&V activities. If not, the remaining risk forces a C) Functional Modification with a corresponding requirements update.



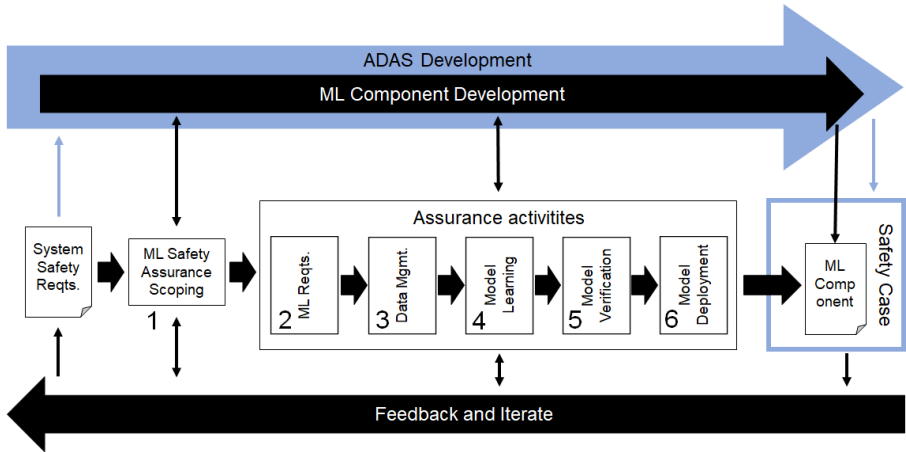
**Fig. 2** A simplified overview of the SOTIF process. Adapted from ISO (2019).

The lower part of Figure 2 shows the V&V activities in the SOTIF process, assuming that they are based on various levels of testing. For each risk, the development organization conducts D) Verification to ensure that the system satisfies the requirements for the known hazardous scenarios. If the F) Conclusion of Verification Tests are satisfactory, the V&V activities continues with validation. If not, the remaining risk requires a C) Functional Modification. In the E) Validation, the development organization explores the presence of unknown hazardous scenarios – if any are identified, they turn into known hazardous scenarios. The H) Conclusion of Validation Tests estimates the likelihood of encountering unknown scenarios that lead to hazardous behavior. If the residual risk is sufficiently small, it is recorded as an acceptable risk. If not, the remaining risk again necessitates a C) Functional Modification.

### 3.2 Safety Assurance Using the AMLAS Process

Our safety assurance work is guided by a methodology for the Assurance of Machine Learning for use in Autonomous Systems (AMLAS) developed by the Assuring Autonomy International Programme at the University of York. AMLAS provides an overall process and a set of safety case patterns for safety assurance of ML components. Figure 3 shows an overview of the six stages of AMLAS, which also provide an overall structure for this paper. Throughout this paper, the notation **[A]–[HH]**, in bold font, refers to 34 individual artifacts prescribed by AMLAS. Table 13 provides an overview of how those artifacts related to the stages of AMLAS and where in this paper they are described. Finally, in Section 11, the 34 artifacts are used to present a complete safety case for the ML component in SMIRK.

The upper part of Figure 3 stresses that the development of an ML component and its corresponding safety case is done in the context of larger systems context. In our case, the larger context is the development of the SMIRK ADAS, indicated by the gray arrow. The AMLAS process starts in the System



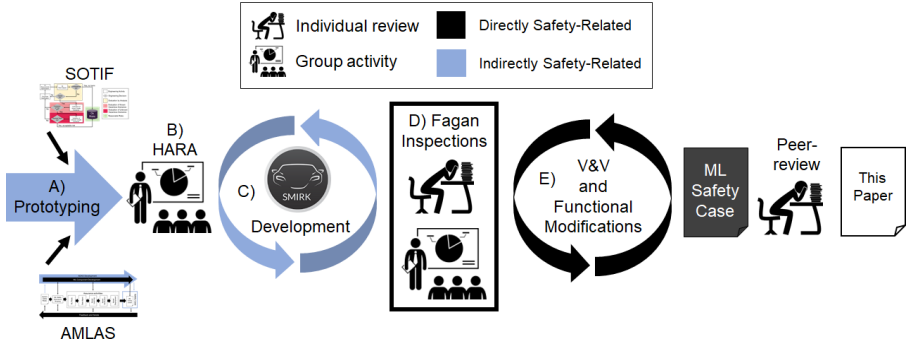
**Fig. 3** An overview of the AMLAS process, adapted from [Hawkins et al \(2021\)](#). Blue color denotes systems engineering, whereas black color relates specifically to the ML component. Numbers refer to the AMLAS stages.

Safety Requirements, which in our case come from following the SOTIF process. However, both SOTIF and AMLAS are iterative process, which means that their activities are performed in parallel and there are many interdependencies – for AMLAS, the iteration is highlighted by the black arrow in the bottom of Figure 3.

Starting from the System Safety Requirements from the left, Stage 1 is **ML Safety Assurance Scoping**. This stage operates on a systems engineering level and defines the scope of the safety assurance process for the ML component as well as the scope of its corresponding safety case – the interplay with the non-ML safety engineering is fundamental. The next five stages of AMLAS all focus on assurance activities for different constituents of ML development and operations. Each of these stages conclude with an assurance argument that when combined, and complemented by evidence through artifacts [A]–[HH], compose the overall ML safety case.

- Stage 2 **ML Safety Requirements Assurance**. Requirements engineering is used to elicit, analyze, specify, and validate ML safety requirements ([Vogelsang and Borg, 2019](#)) in relation to the software architecture and the ODD.
- Stage 3 **Data Management Assurance**. Requirements engineering is first used to develop data requirements that match the ML safety requirements. Subsequently, data sets are generated (development data, internal test data, and verification data) accompanied by quality assurance activities.
- Stage 4 **Model Learning Assurance**. The ML model is trained using the development data. The fulfilment of the ML safety requirements is assessed using the internal test data.
- Stage 5 **Model Verification Assurance**. Different levels of testing or formal verification to assure that the ML model meets the ML safety requirements.





**Fig. 4** An overview of the SMIRK development in the SMILE III project.

Most importantly, the ML model shall be tested on verification data that has not influenced the training in any way.

**Stage 6 Model Deployment Assurance.** Integrate the ML model in the overall system and verify that the system safety requirements are satisfied. Conduct integration testing in the specified ODD.

The rightmost part of Figure 3 shows the overall safety case for the system under development with the argumentation for the ML component as an essential part, i.e., the target of the AMLAS process. The AMLAS argumentation patterns are presented using the graphical format Goal Structuring Notation (GSN) (Assurance Case Working Group, 2021). All semantics used in the figures in Section 11 are defined in this open standard.

### 3.3 SMIRK Development in the SMILE III Project

Figure 4 shows an overview of the two-year development project (SMILE III) that resulted in the SMIRK MVP (Minimum Viable Product) and the safety case for its ML component. Starting from the left, we relied on A) Prototyping to get an initial understanding of the problem and solution domain (Käpyaho and Kauppinen, 2015). As our pre-understanding during prototyping grew, SOTIF and AMLAS were introduced as fundamental development processes and we established a first System Requirements Specification (SRS).

Based on the SRS, we organized a B) Hazard Analysis and Risk Assessment (HARA) workshop (cf. ISO 262626) with all author affiliations represented. Then, the iterative C) SMIRK development phase commenced, encompassing both software development, ML development, and a substantial amount of documentation. When meeting our definition of done, i.e., an MVP implementation and stable requirements specifications, we conducted D) Fagan Inspections as described in Section 3.3.1. After corresponding updates, we baselined the SRS and the Data Management Specification (DMS). Note that due to the Covid-19 pandemic, all group activities were conducted in virtual settings.

Subsequently, the development project turned to E) V&V and Functional Modifications as limitations were identified. In line with the SOTIF process

(cf. Figure 2), also this phase of the project was iterative. The various V&V activities generated a significant part of the evidence that supports our safety argumentation. The rightmost part of Figure 4 depicts the safety case for the ML component in SMIRK, which is peer-reviewed as part of the submission process of this paper.

### 3.3.1 Fagan Inspections

We conducted two formal Fagan inspections (Fagan, 1976) during the SMILE III project with representatives from the organizations listed as co-authors of this paper. All reviewers are active in automotive R&D. The inspections targeted the Software Requirements Specification and the Data Management Specification, respectively. The two formal inspections constitute essential activities in the AMLAS safety assurance and result in ML Safety Requirements Validation Results [J] and a Data Requirements Justification Report [M]. A Fagan inspection consists of the steps 1) Planning, 2) Overview, 3) Preparation, 4) Inspection meeting, 5) Rework, and 6) Follow-up.

1. Planning: The authors prepared the document and invited the required reviewers to an inspection meeting.
2. Overview: During one of the regular project meetings, the lead authors explained the fundamental structure of the document to the reviewers, and introduced an inspection checklist, available on GitHub. Reviewers were assigned particular inspection perspectives based on their individual expertise. All information was repeated in an email, as not all reviewers were present at the meeting.
3. Preparation: All reviewers conducted an individual inspection of the document, noting any questions, issues, and required improvements.
4. Inspection meeting: Two weeks after the individual inspections were initiated, the lead authors and all reviewers met for a virtual meeting. The entire document was discussed, and the findings from the independent inspections were compared. All issues were compiled in inspection protocols that can be found on GitHub.
5. Rework: The lead authors updated the SRS according to the inspection protocol. The independent inspection results were used as input to capture-recapture techniques to estimate the remaining amount of work (Petersson et al, 2004). All changes are traceable through individual GitHub commits.
6. Follow-up: Selected reviewers verified that the previously found issues had been correctly resolved.

## 4 SMIRK System Description [C]

SMIRK is a pedestrian automatic emergency braking (PAEB) system that relies on machine learning (ML). As an example of an advanced driver-assistance system (ADAS), SMIRK is intended to act as one of several systems supporting the driver in the dynamic driving task, i.e., all the real-time operational and tactical functions required to operate a vehicle in on-road traffic.

SMIRK, including the accompanying safety case, is developed with full transparency under an open-source software (OSS) license. We develop SMIRK as a demonstrator in a simulated environment provided by ESI Pro-SiVIC – we stress that SMIRK shall never be used in a real vehicle and take the authors take no responsibility in any such endeavors.

The SMIRK product goal is to assist the driver on country roads in rural areas by performing emergency braking in the case of an imminent collision with a pedestrian. The level of automation offered by SMIRK corresponds to SAE Level 1 – Driver Assistance, i.e., “the driving mode-specific execution by a driver assistance system of either steering or acceleration/deceleration.” – in our case only braking. However, SMIRK is developed with evolvability in mind, thus future versions might include steering and thus comply with SAE Level 2. The first release of SMIRK is an MVP, i.e., an implementation limited to a highly restricted ODD.


Sections 4 and 5 presents the core parts of the SMIRK SRS. The SRS, as well as this section, largely follows the structure proposed in IEEE 830-1998 – IEEE Recommended Practice for Software Requirements Specifications (IEEE, 1998) and the template provided by Wiegers (2008). To support readability, this section presents a SMIRK overview whereas Section 5 specifies the system requirements.

## 4.1 Product Scope

SMIRK is designed to send a brake signal when a collision with a pedestrian is imminent. Figure 5 illustrates the overall function provided by SMIRK. SMIRK shall commence emergency braking if collision with a pedestrian is imminent. Pedestrian are expected to cross the road at arbitrary angles, including perpendicular movement and moving in the toward or away from the car. Furthermore, a stationary pedestrian on the road must also trigger emergency braking, i.e., a scenario known to be difficult for some pedestrian detection systems. Finally, Figure 5 stresses that SMIRK must be robust against false positives, also know as “braking for ghosts.” Trajectories are illustrated with blue arrows accompanied by a speed ( $v$ ) and possibly an angle ( $\theta$ ). In the superscript,  $c$  and  $p$  denote car and pedestrian, respectively, and 0 in the subscript indicates initial speed.

SMIRK is an ADAS that is intended to co-exist with other ADAS in a vehicle. We assume that sensors and actuators will be shared among different systems. SMIRK currently implements its own perception system based on radar and camera input. In future versions, it is likely that a central perception system operating on the vehicle will provide reliable input to SMIRK. This is not yet the case for the SMIRK MVP and this version of the SRS does not specify any requirements related to shared resources. The SMIRK scope is further explained through the context diagram in Section 2.1.

Figure 6 shows the SMIRK context diagram. The sole purpose of SMIRK is PAEB. The design of SMIRK assumes that it will be deployed in a vehicle



**Fig. 5** Example scenario with a pedestrian crossing the road at an arbitrary angle. A false positive is also presented, i.e., a ghost.

with complementary ADAS, e.g., large animal detection, lane keeping assistance, and various types of collision avoidance (cf. “Other ADAS 1 – N”). We also expect that sensors and actuators will be shared between ADAS. For the SMIRK MVP, however, we do not elaborate any further on ADAS co-existence and we do not adhere to any particular higher-level automotive architecture. In the same vein, we do not assume a central perception system that fuses various types of sensor input for individual ADAS such as SMIRK to use. SMIRK uses a standalone ML model trained for pedestrian detection and recognition. In the SMIRK terminology, to mitigate confusion, the radar detects objects and the ML-based pedestrian recognition component identifies potential pedestrians in the camera input. Solid lines in the figure show how SMIRK interacts with sensors and actuators in the ego car. Dashed lines indicate how other ADAS might use sensors and actuators.

Product development inevitably necessitates quality trade-offs. While we have not conducted a systematic quality requirements prioritization, such as an analytical hierarchy process workshop (Kassab and Kilicay-Ergin, 2015), this section shares our general aims with SMIRK. The software product quality model defined in the ISO/IEC 25010 standard consists of eight characteristics. Furthermore, as recommend in requirements engineering research (Horkoff, 2019), we add the two novel quality characteristics explainability and fairness. For each characteristic, we share how important it is considered during the development and assign it a low, medium or high priority. Our priorities influence architectural decisions in SMIRK and support elicitation of architecturally significant requirements (Chen et al, 2012).

- **Functional suitability.** No matter how functionally restricted the SMIRK MVP is, it must meet the stated and implied needs of a prototype ADAS. This quality characteristic is fundamentally important. [**High priority**]
- **Performance efficiency.** When deployed in the simulated environment, SMIRK must be able to process input, conduct ML inference, and possibly

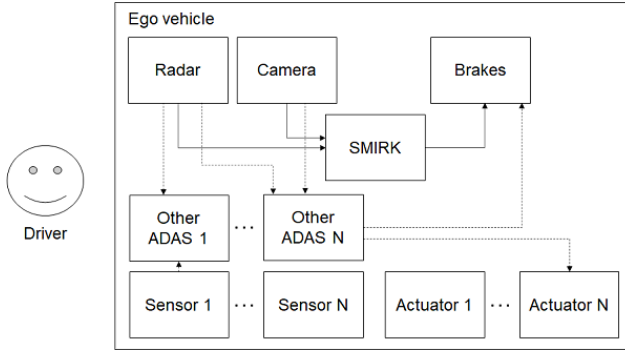
commence emergency braking in realistic driving scenarios. As a real-time system, SMIRK must be sufficiently fast and finding when performance efficiency reached excessive levels is vital in the requirements engineering process. [**Medium priority**]

- **Compatibility.** A product goal is to make SMIRK compatible with other ADAS. So far we have not explored this further, thus this is primarily an ambition beyond the MVP development. [**Low priority**]
- **Usability.** SMIRK is an ADAS that operates in the background and ideally never intervenes in the dynamic driving task. SMIRK does not have a user interface for the direct driver interaction. [**Low priority**]
- **Reliability.** A top priority in the SMIRK development that motivates the application of AMLAS. Note, however, that safety is not covered in the ISO/IEC 25010 product quality model but in its complementary quality-in-use model. [**High priority**]
- **Security.** Not prioritized in the SMIRK MVP. SOTIF is limited to “reasonably foreseeable misuse” but does not address antagonistic attacks. While safety and security shall be co-engineered, we leave this quality characteristic as future work. [**Low priority**]
- **Maintainability.** As mentioned in Section 1.1 Purpose, evolvability from the SMIRK MVP is a key concern. Consequently, maintainability is important, although not more important than functional suitability and reliability. [**Medium priority**]
- **Portability.** We aim to develop SMIRK in a manner that allows porting the ADAS to both other simulated environments and to physical demonstration platforms in future projects. We consider this quality characteristic during the SMIRK development, but it is not a primary concern. [**Low priority**]
- **Explainability.** Explainability is an important characteristic for any cyber-physical system, but the challenge grows with the introduction of DNNs. There is considerable research momentum on “Explainable AI” and we expect that new findings will be applicable to SMIRK. For the MVP development, however, our explainability focus is restricted to the auditability resulting in following AMLAS. [**Medium priority**]
- **Fairness.** Obviously a vital quality characteristic for a PAEB ADAS that primarily impacts the data requirements specified in the Data Management Specification. We have elaborated on SMIRK fairness in a previous study (Borg et al, 2021b). [**High priority**]

## 4.2 Product Functions

SMIRK comprises implementations of four algorithms and uses external vehicle functions. In line with SOTIF, we organize all constituents into the categories sensors, algorithms, and actuators.

- Sensors




**Fig. 6** SMIRK context diagram.

- Radar detection and tracking of objects in front of the vehicle (see Section 6.1).
- A forward-facing mono-camera (see Section 6.1).
- Algorithms
  - Time-to-collision (TTC) calculation for objects on collision course.
  - Pedestrian detection and recognition based on the camera input where the radar detected an object (see Section 8.1).
  - Out-Of-distribution (OOD) detection of never-seen-before input (part of the safety cage mechanism, see Section 8.3).
  - A braking module that commences emergency braking. In the MVP, maximum braking power is always used.
- Actuators
  - Brakes (provided by ESI Pro-SiVIC, not elaborated further).

Figure 7 illustrates detection of a pedestrian on a collision course, i.e., PAEB shall be commenced. The ML-based functionality of pedestrian detection and recognition, including the corresponding OOD detection, is embedded in the **Pedestrian Recognition Component** (defined in Section 6.1).

## 5 SMIRK System Requirements

This section specifies the SMIRK system requirements, organized into system safety requirements and ML safety requirements. ML safety requirements are further refined into performance requirements and robustness requirements. The requirements are largely re-purposed from the system for pedestrian detection at crossings described by [Gauerhof et al \(2020\)](#) to our PAEB ADAS, thus allowing for comparisons to previous work within the research community.



**Fig. 7** Illustrative example of pedestrian detection that shall trigger emergency braking.

### 5.1 System Safety Requirements [A]

- **SYS-SAF-REQ1** SMIRK shall commence automatic emergency braking if and only if collision with a pedestrian on collision course is imminent.

Rationale: This is the main purpose of SMIRK. If possible, ego car will stop and avoid a collision. If a collision is inevitable, ego car will reduce speed to decrease the impact severity. Hazards introduced from false positives, i.e., braking for ghosts, are mitigated under ML Safety Requirements.

### 5.2 Safety Requirements Allocated to ML Component [E]

Based on the HARA (see Section 3.3), two categories of hazards were identified. First, SMIRK might miss pedestrians and fail to commence emergency braking - we refer to this as a *missed pedestrian*. Second, SMIRK might commence emergency braking when it should not - we refer to this as an instance of *ghost braking*.

- Missed pedestrian hazard: The severity of the hazard is very high (high risk of fatality). Controllability is high since the driver can brake ego car.
- Ghost braking hazard: The severity of the hazard is high (can be fatal). Controllability is very low since the driver would have no chance to counteract the braking.

To conclude, we refine **SYS-SAF-REQ1** in the next section to specify requirements in relation to the missed pedestrian hazard. Furthermore, the ghost braking hazard necessitates the introduction of **SYS-ML-REQ2**.

### 5.3 Machine Learning Safety Requirements [H]

This section refines **SYS-SAF-REQ** into two separate requirements corresponding to missed pedestrians and ghost braking, respectively.

- **SYS-ML-REQ1.** The pedestrian recognition component shall identify pedestrians in all valid scenarios when the radar tracking component returns a  $TTC < 4s$  for the corresponding object.
- **SYS-ML-REQ2** The pedestrian recognition component shall reject false positive input that does not resemble the training data.

Rationale: **SYS-SAF-REQ1** is interpreted in light of missed pedestrians and ghost braking and then broken down into the separate ML safety requirements **SYS-ML-REQ1** and **SYS-ML-REQ2**. The former requirement deals with the “if” aspect of **SYS-SAF-REQ1** whereas its “and only if” aspect is targetted by **SYS-SAF-REQ2**. SMIRK follows the reference architecture from [Ben Abdesslem et al \(2016\)](#) and **SYS-ML-REQ1** uses the same TTC threshold (4 seconds, confirmed with the original authors). Moreover, we have confirmed that the TTC threshold is valid for SMIRK in its ODD based on calculating braking distances. **SYS-ML-REQ2** motivates the primary contribution of the SMILE III project, i.e., an OOD detection mechanism that we refer to as a *safety cage*.

### 5.3.1 Performance Requirements

The performance requirements are specified with a focus on quantitative targets for the pedestrian recognition component. All requirements below are restricted to pedestrians on or close to the road.

For objects detected by the radar tracking component with a  $TTC < 4s$ , the following requirements must be fulfilled:

- **SYS-PER-REQ1** The pedestrian recognition component shall identify pedestrians with an accuracy of 93% when they are within 80 m.
- **SYS-PER-REQ2** The false negative rate of the pedestrian recognition component shall not exceed 7% within 50 m.
- **SYS-PER-REQ3** The false positives per image of the pedestrian recognition component shall not exceed 0.1% within 80 m.
- **SYS-PER-REQ4** In 99% of sequences of 5 consecutive images from a 10 FPS video feed, no pedestrian within 80 m shall be missed in more than 20% of the frames.
- **SYS-PER-REQ5** For pedestrians within 80 m, the pedestrian recognition component shall determine the position of pedestrians within 50 cm of their actual position.
- **SYS-PER-REQ6** The pedestrian recognition component shall allow an inference speed of at least 10 FPS in the ESI Pro-SiVIC simulation.

Rationale: SMIRK adapts the performance requirements specified by [Gauerhof et al \(2020\)](#) for the SMIRK ODD. **SYS-PER-REQ1** reuses the accuracy threshold from Example 7 in AMLAS. **SYS-PER-REQ2** and **SYS-PER-REQ3** are two additional requirements inspired by [Henriksson et al \(2019\)](#). Note that **SYS-PER-REQ3** relies on the metric false positive per image rather than false positive rate as true negatives do not exist for object



detection (further explained in Section 9.1 and discussed in Section 12). **SYS-PER-REQ6** means that any further improvements to reaction time have a negligible impact on the total brake distance.

### 5.3.2 Robustness Requirements

Robustness requirements are specified to ensure that SMIRK performs adequately despite expected variations in input. For pedestrians present within 50 m of Ego, captured in the field of view of the camera:

- **SYS-ROB-REQ1** The pedestrian recognition component shall perform as required in all situations Ego may encounter within the defined ODD.
- **SYS-ROB-REQ2** The pedestrian recognition component shall identify pedestrians irrespective of their upright pose with respect to the camera.
- **SYS-ROB-REQ3** The pedestrian recognition component shall identify pedestrians irrespective of their size with respect to the camera.
- **SYS-ROB-REQ4** The pedestrian recognition component shall identify pedestrians irrespective of their appearance with respect to the camera.

Rationale: SMIRK reuses robustness requirements for pedestrian detection from previous work. **SYS-ROB-REQ1** is specified in [Gauerhof et al \(2020\)](#). **SYS-ROB-REQ2** is presented as Example 7 in AMLAS, which has been limited to upright poses, i.e., SMIRK is not designed to work for pedestrians sitting or lying on the road. **SYS-ROB-REQ3** and **SYS-ROB-REQ4** are additions identified during the Fagan inspection of the System Requirements Specification (see Section 3.3.1).

## 5.4 Operational Design Domain [B]

This section briefly describes the SMIRK ODD. As the complete ODD specification, based on the taxonomy developed by NHTSA ([Thorn et al, 2018](#)), is lengthy, we only present the fundamental aspects in this section. We refer interested readers to the GitHub repository. Note that we deliberately specified a minimalistic ODD, i.e., ideal conditions, to allow the development a complete safety case for the SMIRK MVP.

- **Physical infrastructure** Asphalt single-lane roadways with clear lane markings and a gravel shoulder. Rural settings with open green fields.
- **Operational constraints** Maximum speed of 70 km/h and no surrounding traffic.
- **Objects** No objects except 0-1 pedestrians, either stationary or moving with a constant speed ( $< 15$  km/h) and direction.
- **Environmental Conditions** Clear daytime weather with overhead sun. Headlights turned off.

## 6 SMIRK System Architecture

SMIRK is a pedestrian emergency braking ADAS that demonstrates safety-critical ML-based driving automation on SAE Level 1. The system uses input from two sensors (camera and radar/LiDAR), implements a deep neural network trained for pedestrian detection and recognition. If the radar detects an imminent collision between the ego car and an object, SMIRK will evaluate if the object is a pedestrian. If SMIRK is confident that the object is a pedestrian, it will apply emergency braking. To minimize hazardous false positives, SMIRK implements a SMILE safety cage to reject input that is OOD. To ensure industrial relevance, SMIRK builds on the reference architecture from PeVi, an ADAS studied in previous work [Ben Abdesslem et al \(2016\)](#).

Based on a stakeholder analysis in the SMILE III project, this architecture description considers the following stakeholders:

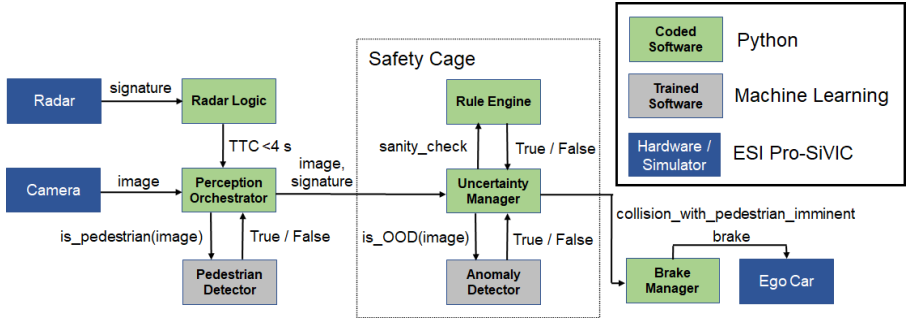
- Researchers who want to study the design of SMIRK.
- Safety assessors who want to investigate the general design in the light of the safety case.
- Software developers building or evolving SMIRK.
- ML developers designing and tuning the ML perception model.
- Hardware developers interested in the SMIRK sensors, incl. replacing them or adding sensor fusion.
- Simulator developers looking for ways to port SMIRK to their virtual prototyping environments.
- Testers developing test plans for SMIRK.
- System integrators who are about to include SMIRK in other systems, incl. co-existence with other ADAS.

Explicitly defined architecture viewpoints support effective communication of certain aspects and layers of a system architecture. The different viewpoints of the identified stakeholders are covered by the established 4+1 view of architecture by [Kruchten \(1995\)](#). The 4+1 view model supports documentation and communication of software-intensive systems. The model is a generic tool that does not restrict its users in terms of notations, tools or design methods. For SMIRK, we describe the logical view using a simple illustration with limited embedded semantics complemented by textual explanations. The process view is presented through a bulleted list, whereas the interested reader can find the remaining parts in the GitHub repository ([RISE Research Institutes of Sweden, 2022](#)). Scenarios are illustrated with figures and explanatory text.

### 6.1 Logical View

The SMIRK logical view is constituted by a description of the entities that realize the PAEB. Figure 8 provides a graphical depiction.

SMIRK interacts with three external resources, i.e., hardware sensors and actuators in ESI Pro-SiVIC: A) Mono Camera (752x480 (WVGA), sensor dimension 3.13 cm x 2.00 cm, focal length 3.73 cm, angle of view 45 degrees),



**Fig. 8** SMIRK logical view.

B) Radar unit (providing object tracking and relative lateral and longitudinal speeds), and C) Ego Car (An Audi A4 for which we are mostly concerned with the brake system). SMIRK consists of the following constituents. We refer to E), F), G), I), and J) as the **Pedestrian Recognition Component**, i.e., the ML-based component for which this study presents a safety case.

- Software components implemented in Python:
  - D) Radar Logic (calculating TTC based on relative speeds)
  - E) Perception Orchestrator (the overall perception logic)
  - F) Rule Engine (part of the safety cage, implementing heuristics such as pedestrians do not fly in the air)
  - G) Uncertainty Manager (main part of the safety cage, implementing logic to avoid false positives)
  - H) Brake Manager (calculating and sending brake signals to the ego car)
- Trained Machine Learning models:
  - I) Pedestrian Detector (a YOLOv5 model trained using PyTorch<sup>3</sup>)
  - J) Anomaly Detector (an autoencoder provided by Seldon<sup>4</sup>)

## 6.2 Process View

The process view deals with the dynamic aspects of SMIRK including an overview of the run time behavior of the system. The overall SMIRK flow is as follows:

1. The Radar detects an object and sends the signature to the Radar Logic class.
2. The Radar Logic class calculates the TTC. If a collision between the ego car and the object is imminent, i.e., TTC is less than 4 seconds assuming constant motion vectors, the Perception Orchestrator is notified.

<sup>3</sup><https://pytorch.org/>

<sup>4</sup><https://www.seldon.io/>

3. The Perception Orchestrator forwards the most recent image from the Camera to the Pedestrian Detector to evaluate if the detected object is a pedestrian.
4. The Pedestrian Detector performs a pedestrian detection in the image and returns the verdict (True/False) to the Pedestrian Orchestrator.
5. If there appears to be a pedestrian on a collision course, the Pedestrian Orchestrator forwards the image and the radar signature to the Uncertainty Manager in the safety cage.
6. The Uncertainty Manager sends the image to the Anomaly Detector and requests an analysis of whether the camera input is OOD or not.
7. The Anomaly Detector analyzes the image in the light of the training data and returns its verdict (True/False).
8. If there indeed appears to be an imminent collision with a pedestrian, the Uncertainty Manager all available information is forwarded to the Rule Engine for a sanity check.
9. The Rule Engine does a sanity check based on heuristics, e.g., in relation to laws of physics, and returns a verdict (True/False).
10. The Uncertainty Manager aggregates all information and, if the confidence is above a threshold, notifies the Brake Manager that collision with a pedestrian is imminent.
11. The Brake Manager calculates a safe brake level and sends the signal to Ego Car to commence PAEB.

## 7 SMIRK Data Management Specification

This section describes the overall approach to data management for SMIRK and the explicit data requirements. SMIRK is a demonstrator for a simulated environment. Thus, as an alternative to longitudinal traffic observations and consideration of accident statistics, we have analyzed the SMIRK ODD through the ESI Pro-SiVIC “Object Catalog.” We conclude that the demographics of pedestrians in the ODD is constituted of the following: adult males and females in either casual, business casual, or business casual clothes, young boys wearing jeans and a sweatshirt, and male road workers. As other traffic is not within the ODD (e.g., cars, motorcycles, and bicycles), we consider the following basic shapes from the object catalog to as examples of OOD objects (that still can appear in the ODD) for SMIRK to handle in operation: boxes, cones, pyramids, spheres, and cylinders.

### 7.1 Data Requirements [L]

This section specifies requirements on the data used to train and test the pedestrian recognition component. The data requirements are specified to comply with the ML Safety Requirements in the System Requirements Specification. All data requirements are organized according to the assurance-related desiderata proposed by [Ashmore et al \(2021\)](#), i.e., the key assurance requirements that ensure that the data set is relevant, complete, balanced, and accurate.

**Table 1** Requirements-Data traceability matrix.

		Data Requirements per Assurance Desiderata																			
		Relevant							Complete							Balanced		Accurate			
		DAT-REL-REQ1	DAT-REL-REQ2	DAT-REL-REQ3	DAT-REL-REQ4	DAT-REL-REQ5	DAT-REL-REQ6	DAT-REL-REQ7	DAT-COM-REQ1	DAT-COM-REQ2	DAT-COM-REQ3	DAT-COM-REQ4	DAT-COM-REQ5	DAT-COM-REQ6	DAT-COM-REQ7	DAT-BAL-REQ1	DAT-BAL-REQ2	DAT-ACC-REQ1	DAT-ACC-REQ2	DAT-ACC-REQ3	
Machine Learning Requirements	Safety	SYS-ML-REQ1	X	X	X	X	X	X		X	X	X	X	X	X	X	X	X	X	X	
		SYS-ML-REQ2							X	X							X	X			
	Performance	SYS-PER-REQ1										X	X	X	X	X					
		SYS-PER-REQ2									X										
		SYS-PER-REQ3							X	X		X	X	X	X	X					
		SYS-PER-REQ4	Not related to the data requirements																		
		SYS-PER-REQ5												X					X	X	X
		SYS-PER-REQ6	Not related to the data requirements																		
	Robustness	SYS-ROB-REQ1							X	X	X	X	X	X	X						
		SYS-ROB-REQ2							X	X	X	X	X	X	X						
SYS-ROB-REQ3								X	X	X	X	X	X	X							
SYS-ROB-REQ4								X	X	X	X	X	X	X							

Table 1 shows a requirements traceability matrix between ML safety requirements and data requirements. The matrix presents an overview of how individual data requirements contribute to the satisfaction of ML Safety Requirements. Entries in individual cells denote that the ML safety requirement is addressed, at least partly, by the corresponding data requirement. **SYS-PER-REQ4** and **SYS-PER-REQ6** are not related to the data requirements.

### 7.1.1 Desideratum: Relevant

This desideratum considers the intersection between the data set and the supported dynamic driving task in the ODD. The SMIRK training data will not cover operational environments that are outside of the ODD, e.g., images collected in heavy snowfall.

- **DAT-REL-REQ1** All data samples shall represent images of a road from the perspective of a vehicle.
- **DAT-REL-REQ2** The format of each data sample shall be representative of that which is captured using sensors deployed on the ego car.
- **DAT-REL-REQ3** Each data sample shall assume sensor positioning representative of the positioning used on the ego car.
- **DAT-REL-REQ4** All data samples shall represent images of a road environment that corresponds to the ODD.
- **DAT-REL-REQ5** All data samples containing pedestrians shall include one single pedestrian.
- **DAT-REL-REQ6** Pedestrians included in data samples shall be of a type that may appear in the ODD.

- **DAT-REL-REQ7** All data samples representing non-pedestrian OOD objects shall be of a type that may appear in the ODD.

Rationale: SMIRK adapts the requirements from the Relevant desiderata specified by [Gauerhof et al \(2020\)](#) for the SMIRK ODD. **DAT-REL-REQ5** is added based on the corresponding fundamental restriction of the ODD of the SMIRK MVP. **DAT-REL-REQ7** restricts data samples providing OOD examples for testing.

### 7.1.2 Desideratum: Complete

This desideratum considers the sampling strategy across the input domain and its subspaces. Suitable distributions and combinations of features are particularly important. [Ashmore et al \(2021\)](#) refer to this as the external perspective on the data.

- **DAT-COM-REQ1** The data samples shall include the complete range of environmental factors within the scope of the ODD.
- **DAT-COM-REQ2** The data samples shall include images representing all types of pedestrians according to the demographics of the ODD.
- **DAT-COM-REQ3** The data samples shall include images representing pedestrians paces from standing still up to running at  $15km/h$ .
- **DAT-COM-REQ4** The data samples shall include images representing all angles an upright pedestrian can be captured by the given sensors on the ego car.
- **DAT-COM-REQ5** The data samples shall include images representing all distances to crossing pedestrians from 10 m up to 100 m away from ego car.
- **DAT-COM-REQ6** The data samples shall include examples with different levels of occlusion giving partial views of pedestrians crossing the road.
- **DAT-COM-REQ7** The data samples shall include a range of examples reflecting the effects of identified system failure modes.

Rationale: SMIRK adapts the requirements from the Complete desiderata specified by [Gauerhof et al \(2020\)](#) for the SMIRK ODD. We deliberately replaced the original adjective “sufficient” to make the data requirements more specific. Furthermore, we add **DAT-COM-REQ3** to cover different poses related to the pace of the pedestrian and **DAT-COM-REQ4** to cover different observation angles.

### 7.1.3 Desideratum: Balanced

This desideratum considers the distribution of features in the data set, e.g., the balance between the number of samples in each class. [Ashmore et al \(2021\)](#) refer to this as an internal perspective on the data.

- **DAT-BAL-REQ1** The data set shall have a representation of samples for each relevant class and feature that ensures AI fairness with respect to gender.

- **DAT-BAL-REQ2** The data set shall have a representation of samples for each relevant class and feature that ensures AI fairness with respect to age.
- **DAT-BAL-REQ3** The data set shall contain both positive and negative examples.

Rationale: SMIRK adapts the requirements from the Balanced desiderata specified by [Gauerhof et al \(2020\)](#) for the SMIRK ODD. The concept of AI fairness is to be interpreted in the light of the Ethics guidelines for trustworthy AI published by the European Commission ([High-Level Expert Group on Artificial Intelligence, 2019](#)). Note that the number of ethical dimensions that can be explored in through the ESI Pro-SiVIC object catalog is limited to gender (**DAT-BAL-REQ1**) and age (**DAT-BAL-REQ2**). Moreover, the object catalog does only contain male road workers and all children are boys. Furthermore, **DAT-BAL-REQ3** is primarily included to align with [Gauerhof et al \(2020\)](#) and to preempt related questions by safety assessors. In practice, the concept of negative examples when training object detection models are typically satisfied implicitly as the parts of the images that do not belong to the annotated class are true negatives (further explained in [Section 9.1](#)).

#### 7.1.4 Desideratum: Accurate

This desideratum considers how measurement issues can affect the way that samples reflect the intended ODD, e.g., sensor accuracy and labeling errors.

- **DAT-ACC-REQ1:** All bounding boxes produced shall include the entirety of the pedestrian.
- **DAT-ACC-REQ2:** All bounding boxes produced shall be no more than 10% larger in any dimension than the minimum sized box capable of including the entirety of the pedestrian.
- **DAT-ACC-REQ3:** All pedestrians present in the data samples shall be correctly labeled.

Rationale: SMIRK reuses the requirements from the Accurate desiderata specified by [Gauerhof et al \(2020\)](#).

## 7.2 Data Generation Log [Q]

This section describes how the data used for training the ML model in the pedestrian recognition component was generated. Based on the data requirements, we generate data using ESI Pro-SiVIC. The data are split into three sets in accordance with AMLAS.

- Development data: Covering both training and validation data used by developers to create models during ML development.
- Internal test data: Used by developers to test the model.
- Verification data: Used in the independent test activity when the model is ready for release.

The SMIRK data collection campaign focuses on generation of annotated data in ESI Pro-SiVIC. All data generation is script-based and fully reproducible. The following two lists present the scripts used to play scenarios and capture the corresponding annotated data. The first section describes positive examples [PX], i.e., humans that shall be classified as pedestrians. The second section describes examples that represent OOD shapes [NX], i.e., objects that shall not initiate PAEB in case of an imminent collision. These images, referred to as OOD examples, shall either not be recognized as a pedestrian or be rejected by the safety cage (see Section 8.3).

For each listed item, there is a link to a YAML configuration file that is used by the Python script that generates the data in the ESI Pro-SiVIC output folder “Sensors.” Ego car is always stationary during data collection, and pedestrians and objects move according to specific configurations. Finally, images are sampled from the camera at 10 frames per second with a resolution of  $752 \times 480$  pixels. For each image, we add a separate image file containing the ground truth pixel-level annotation of the position of the pedestrian.

In total, we generate data representing  $8 \times 616 = 4,928$  execution scenarios with positive examples and  $5 \times 40 = 200$  execution scenarios with OOD examples. In total, the data collection campaign generates roughly 185 GB of image data, annotations, and meta-data (including bounding boxes).

### 7.2.1 Positive Examples

We generate positive examples from humans with eight visual appearances (see the upper part of Figure 9) available in the ESI Pro-SiVIC object catalog.

- P1** Casual female pedestrian
- P2** Casual male pedestrian
- P3** Business casual female pedestrian
- P4** Business casual male pedestrian
- P5** Business female pedestrian
- P6** Business male pedestrian
- P7** Child
- P8** Male construction worker

Each configuration file for positive examples specify the execution of 616 scenarios in ESI Pro-SiVIC. The configurations are organized into four groups (A-D). The pedestrians always follow rectilinear motion (a straight line) at a constant speed during scenario execution. Groups A and B describe pedestrians crossing the road, either from the left (Group A) or from the right (Group B). There are three variation points, i.e., 1) the speed of the pedestrian, 2) the angle at which the pedestrian crosses the road, and 3) the longitudinal distance between ego car and the pedestrian’s starting point. In all scenarios, the distance between the starting point of the pedestrian and the edge of the road is 5 m.

- A. Crossing the road from left to right (280 scenario configurations)



- Speed (m/s): [1, 2, 3, 4]
- Angle (degree): [30, 50, 70, 90, 110, 130, 150]
- Longitudinal distance (m): [10, 20, 30, 40, 50, 60, 70, 80, 90, 100]
- B. Crossing the road from right to left (280 scenario configurations)
  - Speed (m/s): [1, 2, 3, 4]
  - Angle (degree): [30, 50, 70, 90, 110, 130, 150]
  - Longitudinal distance (m): [10, 20, 30, 40, 50, 60, 70, 80, 90, 100]

Groups C and D describe pedestrians moving parallel to the road, either toward ego car (Group C) or away (Group D). There are two variation points, i.e., 1) the speed of the pedestrian and 2) an offset from the road center. The pedestrian always moves 90 m, with a longitudinal distance between ego car and the pedestrian's starting point of 100 m for Group C (towards) and 10 m for Group D (away).

- C. Movement parallel to the road toward ego car (28 scenario configurations)
  - Speed (m/s): [1, 2, 3, 4]
  - Lateral offset (m): [-3, -2, -1, 0, 1, 2, 3]
- D. Movement parallel to the road away from ego car (28 scenario configurations)
  - Speed (m/s): [1, 2, 3, 4]
  - Lateral offset (m): [-3, -2, -1, 0, 1, 2, 3]


### 7.2.2 Out-of-Distribution Examples

We generate OOD examples using five basic shapes available in the ESI Pro-SiVIC object catalog. The OOD examples, visualized in the lower part of Figure 9, are:

- N1** Sphere
- N2** Cube
- N3** Cone
- N4** Pyramid
- N5** Cylinder

All four configuration files for OOD examples specify the execution of 10 scenarios in ESI Pro-SiVIC. The configurations represent a basic shape crossing the road from the left or right at an angle perpendicular to the road. Since basic shapes are not animated, we fix the speed at 4 m/s. In all scenarios, the distance between the starting point of the basic shape and the edge of the road is 5 m. The only variation point is the longitudinal distance between ego car and the objects' starting point. The objects always follow rectilinear motion (a straight line) at a constant speed during scenario execution.

- Longitudinal distance (m): [10, 20, 30, 40, 50, 60, 70, 80, 90, 100]



**Fig. 9** Visual appearance of pedestrians (P1–P8) and basic shapes (N1–N5).

### 7.2.3 Preprocessing and Data Splitting

As the SMIRK data collection campaign relies on data generation in ESI Pro-SiVIC, the need for pre-processing differs from counterparts using naturalistic data. To follow convention, we refer to the data processing between data collection and model training as pre-processing – although post-processing would be a more accurate term for the SMIRK development. We have developed scripts that generate data sets representing the scenarios listed in Sections 7.2.1 and 7.2.2. The scripts ensure that the crossing pedestrians and objects appear at the right distance with specified conditions and with controlled levels of occlusion. All output images share the same characteristics, thus no normalization is needed.

SMIRK includes a script to generate bounding boxes for training the object detection model. ESI Pro-SiVIC generates ground truth image segmentation on a pixel-level. The script is used to convert the output to the appropriate input format for model training.

The development data contains images with no pedestrians, in line with the description of “background images” in the YOLOv5 training tips provided by Ultralytics<sup>5</sup>. Background images have no objects for the model to detect, and are added to reduce FPs. Ultralytics recommends 0-10% background images to help reduce FPs and reports that the fraction of background images in the well-known COCO data set is 1% (Lin et al, 2014). In our case, we add background images with cylinders (N5) to the development data. In total, the SMIRK development data contains 1.98% background images, i.e., 1.75% images without any objects and 0.23% with a cylinder.

The generated SMIRK data are used in sequestered data sets as follows:

<sup>5</sup><https://github.com/ultralytics/yolov5/wiki/Tips-for-Best-Training-Results>

- Development data: **P2**, **P3**, **P6**, and **N5**
- Internal test data: **P1**, **P4**, **N1**, and **N3**
- Verification data: **P5**, **P7**, **P8**, **N2**, and **N4**

Note that we deliberately avoid mixing pedestrian models from the ESI Pro-SiVIC object catalog in the data sets due to the limited diversity in the images within the ODD.

## 8 Machine Learning Component Specification

The pedestrian recognition component consists of, among other things, two ML-based constituents: a pedestrian detector and an anomaly detector (see Figure 8).


### 8.1 Pedestrian Recognition Component

The pedestrian detector uses the third-party OSS framework YOLOv5 by Ultralytics implemented using PyTorch. YOLO is an established real-time object detection algorithm that was originally released by [Redmon et al \(2016\)](#). The first version of YOLO introduced a novel object detection process that uses a single DNN to perform both prediction of bounding boxes around objects and classification at once. Compared to the alternatives, YOLO was heavily optimized for fast inference to support real-time applications. A fundamental concept of YOLO is that the algorithm considers each image only once, hence its name “You Only Look Once.” YOLO is referred to as a single-stage object detector. While there have been several versions of YOLO (and the original authors maintained them until v3), the fundamental ideas of YOLO remains the same across versions - including YOLOv5 used in SMIRK.

YOLO segments input images into smaller images. Each input image is split into a square grid of individual cells. Each cell predicts bounding boxes capturing potential objects and provides confidence scores for each box. Furthermore, YOLO does a class prediction for objects in the bounding boxes. Note that for the SMIRK MVP, the only class we predict is *pedestrian*. Relying on the Intersection over Union (IoU) method for evaluating bounding boxes, YOLO eliminates redundant bounding boxes. The final output from YOLO consists of unique bounding boxes with class predictions. Further details are available in the original paper by [Redmon et al \(2016\)](#).

The pedestrian recognition component in SMIRK uses the YOLOv5 architecture without any modifications. This paragraph presents a high-level description of the model architecture and the key technical details. We refer the interested reader to further details provided [Rajput \(2020\)](#) and the OSS repository on GitHub. YOLOv5 provides several alternative DNN architectures. To enable real-time performance for SMIRK, we select YOLOv5s with 191 layers and  $\approx 7.5$  million parameters.

Figure 10 shows the speed/accuracy tradeoffs for different YOLOv5 architectures with YOLOv5s depicted in orange. The results are provided by



**Fig. 10** Speed/accuracy tradeoffs for different YOLOv5 architectures. (Image source: Ultralytics under GPLv3)

Ultralytics including instructions for reproduction. On the y-axis, COCO AP val denotes the mAP@0.5:0.95 metric measured on the 5,000-image COCO val2017 data set over various inference sizes from 256 to 1,536. On the x-axis, GPU Speed measures average inference time per image on the COCO val2017 data set using an AWS p3.2xlarge V100 instance at batch-size 32. The curve EfficientDet illustrates results from Google AutoML at batch size 8.

As a single-stage object detector, YOLOv5s consists of three core parts: 1) the model backbone, 2) the model neck, and 3) the model head. The model backbone extracts important features from input images. The model neck generates so called “feature pyramids” using PANet (Liu et al, 2018) that support generalization to different sizes and scales. The model head performs the detection task, i.e., it generates the final output vectors with bounding boxes and class probabilities.

In SMIRK, we use the default configurations proposed in YOLOv5s regarding activation, optimization, and cost functions. As activation functions, YOLOv5s uses Leaky ReLU in the hidden layers and the sigmoid function in the final layer. We use the default optimization function in YOLOv5s, i.e., stochastic gradient descent. The default cost function in YOLOv5s is binary cross-entropy with logits loss as provided in PyTorch, which we also use.

## 8.2 Model Development Log [U]

This section describes how the YOLOv5s model has been trained for the SMIRK MVP. We followed the general process presented by Ultralytics for training on custom data.

First, we manually prepared two SMIRK data sets to match the input format of YOLOv5. In this step, we also divided the development data [N] into two parts. The first part containing approximately 80% of development data, was used for training. The second part, consisting of the remaining data, was

used for validation. Camera frames from the same video sequence were kept together in the same partition to avoid having almost identical images in the training and validation sets. Additionally, we kept the distribution of objects and scenario types consistent in both partitions. The internal test data [O] was used as a test set. We then prepared these three data sets, training, validation, and test, according to Ultralytic’s instructions. We created a `dataset.yaml` with the paths to the three data sets and specified that we train YOLOv5 for a single class, i.e., pedestrians. The data sets were already annotated using ESI Pro-SiVIC, thus we only needed to export the labels to the YOLO format with one `txt`-file per image. Finally, we organize the individual files (images and labels) according to the YOLOv5 instructions. More specifically, each label file contains the following information:

- One row per object
- Each row contains class, `x_center`, `y_center`, width, and height.
- Box coordinates are stored in normalized `xywh` format (from 0 – 1).
- Class numbers are zero-indexed, i.e., they start from 0.

Second, we trained a YOLO model using the YOLOv5s architecture with the development data without any pre-trained weights. The model was trained for 10 epochs with a batch-size of 8. The results from the validation subset (27,843 images in total) of the development data guide the selection of the confidence threshold for the ML model. We select a threshold to meet **SYS-*PER-REQ3*** with a safety margin for the development data, i.e., an FPPI of 0.1%. This yields a confidence threshold for the ML model to classify an object as a pedestrian that equals 0.448. The final pedestrian detection model, i.e., the ML model [V], has a size of  $\approx 14$  MB.


### 8.3 OOD Detection for the Safety Cage Architecture

SMIRK detects OOD input images as part of its safety cage architecture. The OOD detection relies on the OSS third-party library Alibi Detect<sup>6</sup> from Seldon. Alibi Detect is a Python library that provides several algorithms for outlier, adversarial, and drift detection for various types of data (Klaise et al, 2020). For SMIRK, we trained Alibi Detect’s autoencoder for outlier detection, with three convolutional and deconvolutional layers for the encoder and decoder respectively.

Figure 11 shows an overview of the DNN architecture of an autoencoder. An encoder and a decoder are trained jointly in two steps to minimize a *reconstruction error*. First, the autoencoder receives input data  $X$  and encodes it into a latent space of fewer dimensions. Second, the decoder tries to reconstruct the original data and produces output  $X'$ . An and Cho (2015) proposed using the reconstruction error from a autoencoder to identify input that differs from the training data. Intuitively, if *inlier* data is processed by the autoencoder, the difference between  $X$  and  $X'$  will be smaller than for *outlier* data. By carefully selecting a threshold, this approach can be used for OOD detection.

---

<sup>6</sup><https://github.com/SeldonIO/alibi-detect>



**Fig. 11** Overview architecture of an autoencoder. Adapted from WikiUser:EugenioTL (CC BY-SA 4.0)

For SMIRK, we trained Alibi Detect’s autoencoder for OOD detection on the training data subset of the development data. The encoder part is designed with three convolutional layers followed by a dense layer resulting in a bottleneck that compresses the input by 96.66%. The latent dimension is limited to 1,024 variables to limit requirements on processing VRAM of the GPU. The reconstruction error from the autoencoder is measured as the mean squared error between the input and the reconstructed instance. The mean squared error is used for OOD detection by computing the reconstruction error and considering an input image as an outlier if the error surpasses a threshold  $\theta$ . The threshold used for OOD detection in SMIRK is 0.004, roughly corresponding to the threshold that rejects a number of samples that equals the amount of outliers in the validation set. As explained in Section 10.4, the OOD detection is only active for objects at least 10 m away from ego car as the results for close-up images are highly unreliable. Furthermore, as the constrained SMIRK ODD ensures that only one single object appears in each scenario, the safety cage architecture applies the policy “once an anomaly, always an anomaly” – objects that get rejected once will remain anomalous no matter what subsequent frames might contain.

## 9 SMIRK System Test Specification

This section describes the overall SMIRK test strategy. The ML-based pedestrian recognition component is tested on multiple levels. We focus on four aspects of the ML testing scope facet proposed by [Song et al \(2022\)](#):

- **Data set testing:** This level refers to automatic checks that verify that specific properties of the data set are satisfied. As described in the ML Data Validation Results, the data validation, presented in Section 10.1, includes automated testing of the Balance desiderata. [Zhang et al \(2022\)](#) refer to data set testing as Input testing.
- **Model testing:** Testing that the ML model provides the expected output. This is the primary focus of academic research on ML testing, and includes white-box, black-box, and data-box access levels during testing ([Riccio et al,](#)

2020). SMIRK model testing is done independently from model development and results in ML Verification Results [Z] as described in Section 10.2.2.

- Unit testing: Conventional unit testing on the level of Python classes. A test suite developed for the pytest framework is maintained by the developers and not elaborated further in this paper.
- System testing: System-level testing of SMIRK based on a set of Operational Scenarios [EE]. All test cases are designed for execution in ESI Pro-SiVIC. The system testing targets the requirements in the System Requirements Specification. This level of testing results in Integration Testing Results [FF] presented in Section 10.3.

## 9.1 ML Model Testing [AA]

This section corresponds to the Verification Log [AA] in AMLAS Step 5, i.e., Model Verification Assurance. Here we explicitly document the ML Model testing strategy, i.e., the range of tests undertaken and bounds and test parameters motivated by the SMIRK system requirements.

The testing of the ML model is based on assessing the object detection accuracy for the sequestered verification data set. A fundamental aspect of the verification argument is that this data set was never used in any way during the development of the ML model. To further ensure the independence of the ML verification, engineers from Infotiv, part of the SMILE III research consortium, led the corresponding verification and validation work package and were not in any way involved in the development of the ML model. As described in the Machine Learning Component Specification (see Section 8), the ML development was led by Semcon with support from RISE Research Institutes of Sweden.

The ML model test cases provide results for both 1) the entire verification data set and 2) eight slices of the data set that are deemed particularly important. The selection of slices was motivated by either an analysis of the available technology or ethical considerations, especially from the perspective of AI fairness (Borg et al, 2021b). Consequently, we measure the performance for the following slices of data. Identifiers in parentheses show direct connections to requirements.

- S1 The entire verification data set
- S2 Pedestrians close to the ego car (longitudinal distance < 50 m) (**SYS-PER-REQ1**, **SYS-PER-REQ2**)
- S3 Pedestrians far from the ego car (longitudinal distance  $\geq$  50 m)
- S4 Running pedestrians (speed  $\geq$  3 m/s) (**SYS-ROB-REQ2**)
- S5 Walking pedestrians (speed > 0 m/s but < 3 m/s) (**SYS-ROB-REQ2**)
- S6 Occluded pedestrians (entering or leaving the field of view, defined as bounding box in contact with any edge of image) (**DAT-COM-REQ4**)
- S7 Male pedestrians (**DAT-COM-REQ2**)
- S8 Female pedestrians (**DAT-COM-REQ2**)
- S9 Children (**DAT-COM-REQ2**)

Evaluating the output from an object detection model in computer vision is non-trivial. We rely on the established IoU metric to evaluate the accuracy of the YOLOv5 model. After discussions in the development team, supported by visualizations<sup>7</sup>, we set the target at 0.5. We recognize that there are alternative measures tailored for pedestrian detection, such as the log-average miss rate (Dollar et al, 2011) but we find such metrics to be unnecessarily complex for the restricted SMIRK ODD with a single pedestrian. There are also entire toolboxes that can be used to assess object detection (Bolya et al, 2020). In our safety argumentation, however, we argue that the higher explainability of a simpler – but valid – evaluation metric outweighs the potential benefits of a customized metric customized for a more complex ODD.

Even using the standard IoU metric to assess how accurate SMIRK’s ML model is, the evaluation results are not necessarily intuitive to non-experts. Each image in the SMIRK data set either has a ground truth bounding box containing the pedestrian or no bounding box at all. Similarly, when performing inference on an image, the ML model will either predict a bounding box containing a potential pedestrian or no bounding box at all. IoU is the intersection over the union of the two bounding boxes. An IoU of 1 implies a perfect overlap. For the ML model in SMIRK, we evaluate pedestrian detection at IoU = 0.5, which for each image means:

TP True positive:  $\text{IoU} \geq 0.5$

FP False positive:  $\text{IoU} < 0.5$

FN False negative: There is a ground truth bounding box in the image, but no predicted bounding box.


TN True negative: All parts of the image with neither a ground truth nor a predicted bounding box. This output carries no meaning in our case.

Figure 12 shows predictions from the the ML model. The green rectangles show the ground truth and the red rectangles show ML model’s prediction of where a pedestrian is present. The left example is a FP since  $\text{IoU}=0.3$  with a predicted box substantially smaller than the ground truth. On the other hand, the ground truth is questionable, as there probably is only a single pixel containing the pedestrian below the visible arm that drastically increases the size of the green box. The center example is a TP with  $\text{IoU}=0.9$ , i.e., the overlap between the boxes is very large. The right example is another FP with  $\text{IoU}=0.4$  where the predicted box is much larger than the ground truth. These examples show that FPs during model testing do not directly translate to FPs on the system level as discussed in the HARA (Safety Requirements Allocated to ML Component [E]). If any of the objects within the red bounding boxes were on a collision course with the ego car, commencing PAEB would indeed be the right action for SMIRK and thus not violate **SYS-SAF-REQ1**. This observation corroborates the position by (Haq et al, 2021), i.e., system level testing that goes beyond model testing on single frames is critically needed.

---

<sup>7</sup><https://zapire.com/docs/visualizing-ml/iou.html>





**Fig. 12** Example predictions from the SMIRK ML model. The center image represents a TP, whereas the left and right examples are FPs with IoU scores of 0.3 and 0.4, respectively.

All results from running ML model testing, i.e., ML Verification Results [Z], are documented in the Protocols folder.

## 9.2 System Level Testing

System-level testing of SMIRK involves integrating the ML model into the pedestrian recognition component and the complete PAEB ADAS. We do this by defining a set of Operational Scenarios [EE] for which we assess the satisfaction of the ML Safety Requirements. The results from the system-level testing, i.e., the Integration Testing Results [FF], are presented in Section 10.3.

### 9.2.1 Operational Scenarios

SOTIF defines an operational scenario as “a description of an imagined sequence of events that includes the interaction of the product or service with its environment and users, as well as interaction among its product or service components” (ISO, 2019). Consequently, the set of operational scenarios used for testing SMIRK on the system level must represent the diversity of real scenarios that may be encountered when SMIRK is in operation. Furthermore, for testing purposes, it is vital that the set of defined scenarios are meaningful with respect to the verification of SMIRK’s safety requirements.

As SMIRK is designed to operate in ESI Pro-SiVIC, the difference between defining operational scenarios in text and implementation scripts to execute the same scenarios in the simulated environment is very small. We will not define any operational scenarios that cannot be scripted for execution in ESI Pro-SiVIC. To identify a meaningful set of operational scenarios, we use equivalence partitioning as proposed by Masuda (2017) as one approach to limit the number of test scenarios to execute in vehicle simulators. Originating in the equivalence classes, we use combinatorial testing to reduce the set of operational scenarios. Using combinatorial testing to create test cases for system testing of a PAEB testing in a vehicle simulator has previously been reported by Tao et al (2019). We create operational scenarios that provide complete

pair-wise testing of SMIRK considering the identified equivalence classes using the AllPairs test combinations generator<sup>8</sup>.

Based on an analysis of the ML Safety Requirements and the Data Requirements, we define operational scenarios addressing **SYS-ML-REQ1** and **SYS-ML-REQ2** separately. For each subset of operational scenarios, we identify key variation dimensions (i.e., parameters in the test scenario generation) and split dimensions into equivalence classes using explicit ranges. Note that ESI Pro-SiVIC enables limited configurability of basic shapes compared to pedestrians, thus the corresponding number of operational scenarios is lower.

#### Operational Scenarios for **SYS-ML-REQ1**:

- Pedestrian starting point (lateral offset from the road in meters): Left side of the road (-5 m), Center of the road (0 m), Right side of the road (5 m)
- Longitudinal distance from ego car (offset in meters): Close distance (< 25 m), Medium distance (25–50 m), Far away (> 50 m)
- Pedestrian appearance: Female casual, Male business casual, Male business, Female business, Child, Male worker
- Pedestrian speed (m/s): Stationary (0 m/s), Slow (1 m/s), Fast (3 m/s)
- Pedestrian crossing angle (degrees): Toward ego car (0), Diagonal toward (45), Perpendicular (90), Diagonal away (135), Away from car (180)
- Ego car speed (m/s): Slow (< 10 m/s), Medium (10–15 m/s), Fast (15–20 m/s)

The dimensions and ranges listed above result in 2,430 possible combinations. Using combinatorial testing, we create a set of 25 operational scenarios that provides pair-wise coverage of all equivalence classes.

#### Operational Scenarios for **SYS-ML-REQ2**:

- Object starting point (lateral offset from the road in meters): Left side of the road (-5 m), On the road (0 m), Right side of the road (5 m)
- Longitudinal distance from ego car (offset in meters): Close (< 25 m), Medium distance (25–50 m), Far away (> 50 m)
- Object appearance: Sphere, Cube, Cone, Pyramid
- Object speed (m/s): Stationary (0 m/s), Slow (1 m/s), Fast (3 m/s)
- Ego car speed (m/s): Slow (<10 m/s), Medium (10–15 m/s), Fast (15–20 m/s)

The dimensions and ranges listed above result in 324 possible combinations. Using combinatorial testing, we create a set of 13 operational scenarios that provides pair-wise coverage of all equivalence classes.

For each operational scenario, two test parameters represent ranges of values, i.e., the longitudinal distance between ego car and the pedestrian and the speed of ego car. For these two test parameters, we identify a combination

---

<sup>8</sup><https://github.com/thombashi/allpairspy>

**Table 2** Pairwise-testing of equivalence classes constituting 38 operational scenarios.

ID	Object	X	Y	Angle	Speed	Car_Speed
TC-OS-1	P7	close	left	diagonal.away	slow	slow
TC-OS-2	P7	medium	center	away	stationary	medium
TC-OS-3	P7	far	right	diagonal.towards	fast	fast
TC-OS-4	P5	medium	right	perpendicular	slow	medium
TC-OS-5	P5	close	center	towards	fast	fast
TC-OS-6	P1	far	left	diagonal.away	fast	fast
TC-OS-7	P6	medium	left	perpendicular	fast	medium
TC-OS-8	P6	far	center	away	slow	slow
TC-OS-9	P1	close	right	diagonal.towards	slow	slow
TC-OS-10	P4	left	right	diagonal.away	slow	slow
TC-OS-11	P8	close	left	perpendicular	fast	slow
TC-OS-12	P8	far	left	diagonal.towards	slow	medium
TC-OS-13	P4	far	center	perpendicular	stationary	fast
TC-OS-14	P5	close	left	diagonal.away	slow	slow
TC-OS-15	P6	close	right	diagonal.away	slow	slow
TC-OS-16	P4	medium	left	diagonal.towards	slow	slow
TC-OS-17	P8	close	right	diagonal.away	slow	slow
TC-OS-18	P1	close	center	diagonal.towards	stationary	medium
TC-OS-19	P1	medium	left	diagonal.towards	slow	slow
TC-OS-20	P6	medium	left	diagonal.towards	slow	slow
TC-OS-21	P1	medium	left	perpendicular	slow	medium
TC-OS-22	P5	medium	left	diagonal.towards	slow	slow
TC-OS-23	P4	medium	left	diagonal.away	slow	medium
TC-OS-24	P7	medium	left	perpendicular	slow	slow
TC-OS-25	P8	medium	left	diagonal.towards	slow	slow
TC-OS-26	N2	close	left	perpendicular	slow	slow
TC-OS-27	N3	medium	right	perpendicular	fast	medium
TC-OS-28	N4	far	right	perpendicular	slow	fast
TC-OS-29	N1	far	left	perpendicular	slow	medium
TC-OS-30	N1	medium	left	perpendicular	fast	fast
TC-OS-31	N1	close	right	perpendicular	fast	slow
TC-OS-32	N4	close	left	perpendicular	fast	fast
TC-OS-33	N3	medium	left	perpendicular	slow	slow
TC-OS-34	N2	far	right	perpendicular	slow	slow
TC-OS-35	N3	far	left	perpendicular	slow	fast
TC-OS-36	N2	medium	left	perpendicular	fast	medium
TC-OS-37	N3	close	left	perpendicular	slow	slow
TC-OS-38	N1	close	left	perpendicular	slow	medium

of values that result in a collision unless SMIRK initiates emergency braking. Table 2 shows an overview of the 38 operational scenarios, whereas all details are available as executable test scenarios in the GitHub repository.

### 9.2.2 System Test Cases

The system test cases are split into three categories. First, each operational scenario identified in Section 9.2.1 constitutes one system test case, i.e., Test Cases 1-38. Second, to increase the diversity of the test cases in the simulated environment, we complement the reproducible Test Cases 1-38 with test case counterparts adding random jitter to the parameters. For Test Cases 1-38, we

create analogous test cases that randomly add jitter in the range from -10% to +10% to all numerical values. Partial random testing has been proposed by Masuda (2017) in the context of test scenarios execution in vehicle simulators. Note that introducing random jitter to the test input does not lead to the test oracle problem (Barr et al, 2014), as we can automatically assess whether there is a collision between ego car and the pedestrian without emergency braking in ESI Pro-SiVIC or not (TC-RAND-[1-25]). Furthermore, for the test cases related to provoking ghost braking, we know that emergency braking shall not commence.

The third category is requirements-based testing (RBT). RBT is used to gain confidence that the functionality specified in the ML Safety Requirements has been implemented correctly (Hauer et al, 2019). The top-level safety requirement **SYS-SAF-REQ1** will be verified by testing of all underlying requirements, i.e., its constituent detailed requirements. The test strategy relies on demonstrating that **SYS-ML-REQ1** and **SYS-ML-REQ2** are satisfied when executing TC-OS-[1-38] and TC-RAND-[1-38]. **SYS-PER-REQ1** – **SYS-PER-REQ5** and **SYS-ROB-REQ1** – **SYS-ROB-REQ4** are verified through the model testing described in Section 9.1. The remaining performance requirement **SYS-PER-REQ6** is verified by TC-REQ-3. Table 3 lists all system test cases, of all three categories, using the Given-When-Then structure as used in behavior-driven development (Tsilionis et al, 2021). For the test cases TC-RBT-[1-3], the “Given” condition is that all metrics have been collected during execution of TC-OS-[1-38] and TC-RAND-[1-38]. The set includes seven metrics:

1. **MinDist**: Minimum distance between ego car and the pedestrian during a scenario.
2. **TimeTrig**: Time when the radar tracking component first returned  $TTC < 4$  s for an object.
3. **DistTrig**: Distance between ego car and the object when the radar component first returned  $TTC < 4$  s for an object.
4. **TimeBrake**: Time when emergency braking was commenced.
5. **DistBrake**: Distance between ego car and the object when emergency braking commenced.
6. **Coll**: Whether a scenario involved a collision between ego car and a pedestrian.
7. **CollSpeed**: Speed of ego car at the time of collision.

## 10 SMIRK Test Results

This section presents the most important test results from three levels of ML testing, i.e., data testing, model testing, and system testing. Complete test reports are available in the protocols subfolder on GitHub<sup>9</sup>. Moreover, this section presents the Erroneous Behaviour Log.

---

<sup>9</sup><https://github.com/RI-SE/smirk/tree/main/docs/protocols>

**Table 3** SMIRK system test cases. VMC means valid metrics were collected during execution of the 38 preceding scenarios.

ID	Type	Given	When	Then
TC-OS-[1-25]	Op. Scenario	Scenario [1-25]	Pedestrian crosses the street and ego car is on collision course	SMIRK commences PAEB
TC-OS-[26-38]	Op. Scenario	Scenario [26-38]	Object crosses the street and ego car is on collision course	SMIRK does not commence PAEB
TC-RAND-[1-25]	Random Testing	TC-OS-[1-25]+jitter	Pedestrian crosses the street and ego car is on collision course	SMIRK commences PAEB
TC-RAND-[19-38]	Random Testing	TC-OS-[19-38]+jitter	Object crosses the street and ego car is on collision course	SMIRK does not commence PAEB
TC-RBT-1	RBT (SYS-ML-REQ1)	VMC	The radar tracking component returns a pedestrian with TTC < 4 s	The pedestrian recognition component identifies the pedestrian
TC-RBT-2	RBT (SYS-ML-REQ2)	VMC	The radar tracking component returns a basic shape with TTC < 4 s	The pedestrian recognition component does not identify a pedestrian
TC-RBT-3	RBT (SYS-PER-REQ6)	VMC	The radar tracking component returns a pedestrian with TTC < 4 s within 80 m	The inference speed of the pedestrian recognition component is at least 10 FPS

## 10.1 Results from Data Testing [S]

This section describes the results from testing the SMIRK data set. The data testing primarily involves a statistical analysis of its distribution and automated data validation using Great Expectations<sup>10</sup>. Together with the outcome of the Fagan inspection of the Data Management Specification (described in Section 3.3.1), this constitutes the ML Data Validation Results in AMLAS. As depicted later in Figure 22, the results entail evidence mapping to the four assurance-related desiderata, i.e., we report a validation of 1) data relevance, 2) data completeness, 3) data balance, and 4) data accuracy. Since we generate synthetic data using ESI Pro-SiVIC, data relevance has been validated through code reviews and data accuracy is implicit as the tool's ground truth is used. For both the relevance and accuracy desiderata, we have manually analyzed a sample of the generated data to verify requirements satisfaction.

We validate the ethical dimension of the data balance by analyzing the gender (**DAT-BAL-REQ1**) and age (**DAT-BAL-REQ2**) distributions of the pedestrians in the SMIRK data set. SMIRK evolves as a demonstrator in a Swedish research project, which provides a frame of reference for this analysis. Table 4 shows how the SMIRK data set compares to Swedish demographics from the perspective of age and gender. The demographics originate in a study on collisions between vehicles and pedestrians by the Swedish Civil Contingencies Agency (Schylander, 2014). We notice that 1) children are slightly over-represented in accidents but under-represented in deadly accidents, and

<sup>10</sup><https://greatexpectations.io/>

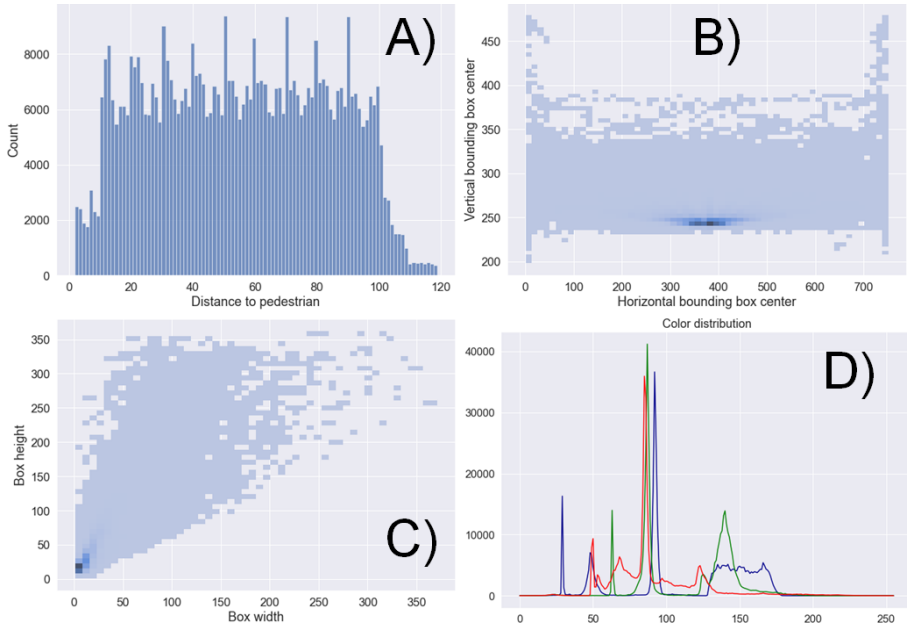
**Table 4** Distribution of pedestrian types in Sweden and in the SMIRK data set.

Pedestrian types	Population	Accidents	Deadly accidents	SMIRK
Children & young adults (0-19)	23%	27%	12%	12.5%
Adult males (20+)	39%	31%	57%	50%
Adult females (20+)	38%	42%	31%	37.5%

that 2) adult males account for over half of the deadly accidents in Sweden. The rightmost column shows the distribution of pedestrian types in the entire SMIRK data set. We designed the SMIRK data generation process to result in a data set that resembles the deadly accidents in Sweden, but, motivated by AI fairness, we increased the fraction of female pedestrians to mitigate a potential gender bias. Finally, as discussed in Section 7.2.3, code reviews confirmed that the development data contains roughly 2% “background images”.

Automated data testing is performed by defining conditions that shall be fulfilled by the data set. These conditions are checked against the existing data and any new data that is added. Some tests are fixed and simple, such as expecting the dimensions of input images to match the ones produced by the vehicle’s camera. Similarly, all bounding boxes are expected to be within the dimensions of the image. Other tests look at the distribution and ranges of values to assure the completeness, accuracy, and balance of the data set and catch human errors. This includes validating enough coverage of pedestrians at different positions of the image, coverage of varying range of pedestrian distances, and bounding box aspect ratios. For values that are hard to define rules for, a known good set of inputs can be used as a starting point and remaining and new inputs can be checked to against these reference inputs. As an example, this can be used to verify that the color distribution and pixel intensity are within expected ranges. This can be used to identify images that are too dark or dissimilar to existing images.

Figure 13 shows a selection of summary plots from the data testing that support our claims for data validity, in particular from the perspective of data completeness. Subplot A) presents the distance distribution between ego car and pedestrians, verifying that the data set contains pedestrians at distances 10–100 m (**DAT-COM-REQ5**). Subplot B) shows a heatmap of the bounding boxes’ centers captures by the 752x480 WVGA camera. We confirm that pedestrians appear from the sides of the field of view and a larger fraction of images contain a pedestrian just in front of ego car. The position distribution supports our claim that **DAT-COM-REQ4** is satisfied, i.e., the data samples represent different camera angles. Subplot C) shows a heatmap of bounding box dimensions, i.e., different aspect ratios. A variety of aspect ratios indicate that pedestrians move with a diversity of arm and leg movements – indicating walking and running – and thus support our claim that **DAT-COM-REQ3** is fulfilled. Finally, subplot D) shows the color histogram of the data set. In the automated data testing, we use this as a reference value when adding new images to ensure that they match the ODD. For example, a sample of nighttime images would have a substantially different color distribution.



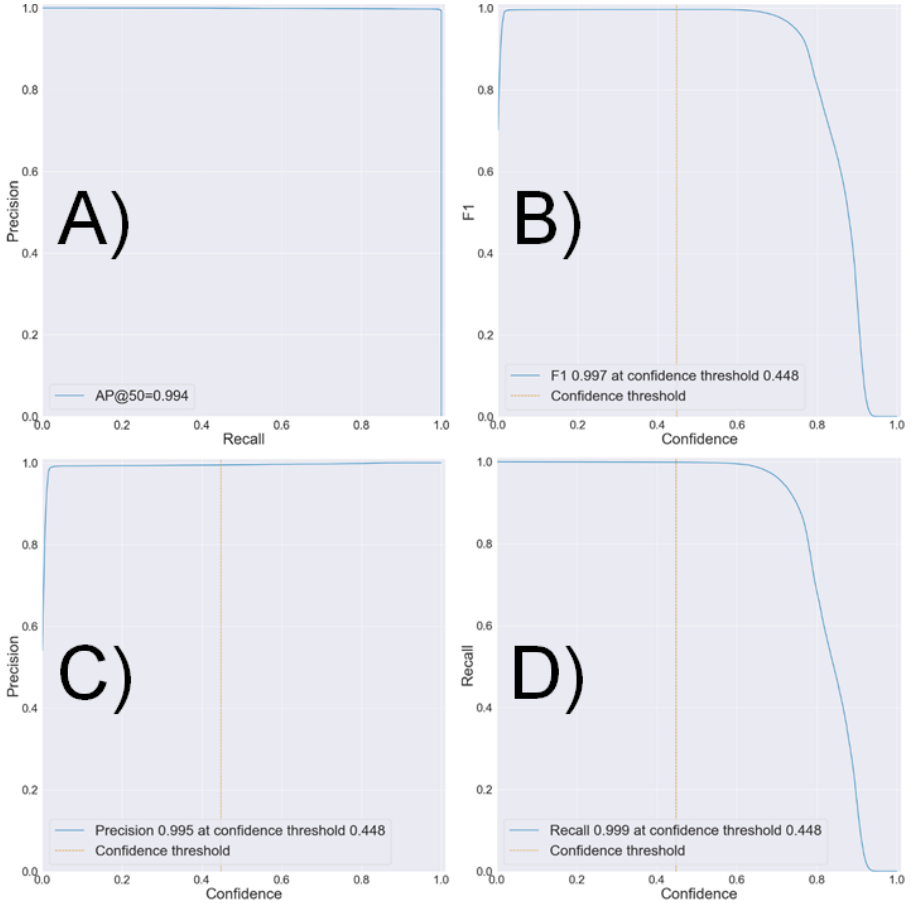
**Fig. 13** Four visualizations from the data testing. A) distance distributions between ego car and pedestrians (m), B) heatmap of center positions of the bounding boxes (WVGA pixel), C) heatmap of bounding box dimensions (pixels), and D) the color histogram of the data set.

## 10.2 Results from Model Testing

This section is split into results from testing done during development and the subsequent independent verification. Throughout this section, the following abbreviations are used for a set of standard evaluation metrics: Precision (P), Recall (R), F1-score (F1), Intersection over Union (IoU), True Positive (TP), False Positive (FP), FPs Per Image (FPPI), False Negative (FN), and Average Precision for IoU at 0.5 (AP@0.5), and Confidence (Conf).

### 10.2.1 Internal Test Results [X]

In this section, we present the most important results from the internal testing. These results provide evidence that the ML model satisfies the ML safety requirements (see Section 5.3) on the internal test data. The total number of images in the internal test data is 139,526 (135,139 pedestrians (96.9%) and 4,387 non-pedestrians (3.1%)). As described in Section 9.1, Figure 14 depicts four subplots representing  $\text{IoU} = 0.5$ : A) P vs R, B) F1 vs. Conf, C) P vs. Conf, and D) R vs. Conf. Subplot A) shows that the ML model is highly accurate, i.e., the unavoidable discrimination-efficiency tradeoff of object detection (Wu and Nevatia, 2008) is only visible in the upper right corner. Subplots B)–D) shows how P, R, and F1 vary with different Conf thresholds. Table 5 presents further details of the accuracy of the ML model for the selected Conf threshold, organized into 1) all distances from the ego car, 2) within 80 m, and 3) within



**Fig. 14** Evaluation of the ML model on the internal test data at IoU=0.5. A) P–R curve, B) F1 vs. Conf, C) P vs. Conf, and D) R vs. Conf for the internal test data.

50 m, respectively. The table also shows the effect of adding OOD detection using the autoencoder, i.e., a substantially reduced number of FPs.

Table 6 demonstrates how the ML model satisfies the performance requirements on the internal test data. First, the TP rate (95.9%) and the FN rate (0.31%) for the respective distances meet the requirements. The model’s FPPI (0.42%), on the other hand, is too high to meet **SYS-PER-REQ3** as we observed 444 FPs (cones outnumber spheres by 2:1). This observation reinforces the need to use a safety cage architecture, i.e., OOD detection that can reject input that does not resemble the training data. The rightmost column in Table 6 shows how the FPPI decreased to 0.012% with the autoencoder. All basic shapes were rejected, but 13 images with pedestrians led to FPs within the ODD due to too low IoU scores.

**SYS-PER-REQ4** is met as the fraction of rolling windows with more than a single FN is 0.24%, i.e.,  $\leq 3\%$ . Figure 15 shows the distribution of



**Table 5** ML model accuracy on the internal test data at the Conf threshold 0.448. The rows show results for all distances, within 80 m, and within 50 m, respectively. Every second row show results for the ML model followed by OOD detection using the autoencoder.

Distance	Total	TP	FP	FN	P	R	F1	AP@0.5
All	139,526	134,948	711	191	0.9948	0.9986	0.9967	0.9942
+OOD		134,927	20	212	0.9999	0.9984	0.9991	0.995
≤ 80 m	105,588	101,320	444	173	0.9956	0.9983	0.997	0.9948
+OOD		101,300	13	193	0.9999	0.9981	0.999	0.995
≤ 50 m	61,845	57,877	186	173	0.9968	0.9970	0.9969	0.9944
+OOD		57,857	13	193	0.9998	0.9967	0.9982	0.995

**Table 6** ML model satisfaction of the performance requirements on the internal test data at the Conf threshold 0.448. R1–R4 = SYS-PER-REQ1–4. The rightmost column show results for the YOLOv5 model followed by OOD detection using the autoencoder.


Req.	Expected	Observed (Model)	Observed (Model+OOD)
R1	TP rate $\geq 93\%$ for $\leq 80$ m	$\frac{101,320}{105,588} = 96.0\%$	$\frac{101,300}{105,588} = 95.9\%$
R2	FN rate $\leq 7\%$ for $\leq 50$ m	$\frac{173}{61,845} = 0.28\%$	$\frac{193}{61,845} = 0.31\%$
R3	FPPI $\leq 0.1\%$ for $\leq 80$ m	$\frac{444}{105,588} = 0.42\%$	$\frac{13}{105,588} = 0.012\%$
R4	$\leq 3\%$ of rolling windows contain $\geq 2$ misses in 5 frames for $\leq 80$ m	$\frac{216}{101,564} = 0.21\%$	$\frac{239}{101,564} = 0.24\%$

position errors in the object detection for pedestrians within 80 m of ego car, i.e., the difference between the object detection position and ESI Pro-SiVIC ground truth. The median error is 1.0 cm, the 99% percentile is 5.6 cm, and the largest observed error is 12.7 cm. Thus, we show that **SYS-PER-REQ5** is satisfied for the internal test data, i.e.,  $\leq 50$  cm position error for pedestrian detection within 80 m. Note that satisfaction of **SYS-PER-REQ6**, i.e., sufficient inference speed, is demonstrated as part of the system testing reported in Section 10.3. The complete test report is available on GitHub.

Table 7 presents the output of the ML model on the eight slices of internal test data defined in Section 9.1. Note that we saved the children in the ESI Pro-SiVIC object catalog for the verification data, i.e., S9 does not exist in the internal test data. Apart from the S6 slice with occlusion, the model accuracy is consistent across the slices which corroborates satisfaction of the robustness requirements on the internal test data, e.g., in relation to pose (**SYS-ROB-REQ2**), size (**SYS-ROB-REQ2**), and appearance (**SYS-ROB-REQ2**).

### 10.2.2 ML Verification Results [Z]

This section reports the key findings from conducting the independent ML model testing, i.e., the Verification Log in the AMLAS terminology. These results provide independent evidence that the ML model satisfies the ML safety requirements (see Section 5.3) on the verification data. The total number of images in the verification data is 208,884 (202,712 pedestrians (97.0%) and



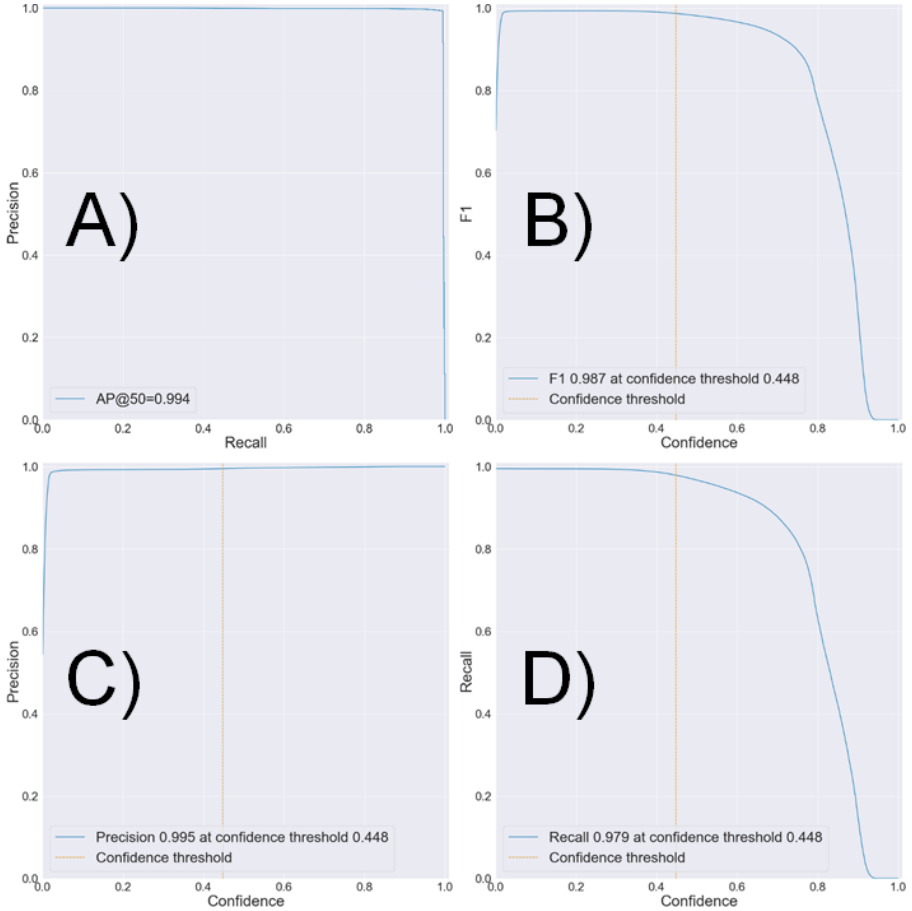
**Fig. 15** Distribution of position errors for the internal test data.

**Table 7** ML model accuracy on eight slices of the internal test data. S1=All data, S2=close distance, S3=far distance, S4=running pedestrians, S5=walking pedestrians, S6=occluded pedestrians, S7=males, and S8=females. Every second rows show results for the ML model followed by OOD detection using the autoencoder.

Slice	Total	TP	FP	FN	P	R	F1	AP@0.5
S1	139,526	134,948	711	191	0.9948	0.9986	0.9967	0.9942
	+OOD	134,927	20	212	0.9999	0.9984	0.9991	0.995
S2	61,333	57,774	16	172	0.9997	0.997	0.9984	0.995
	+OOD	57,753	13	193	0.9998	0.9967	0.9982	0.995
S3	43,547	43,547	0	0	1	1	1	0.995
	+OOD	43,547	0	0	1	1	1	0.995
S4	38,786	37,804	9	48	0.9998	0.9987	0.9992	0.995
	+OOD	37,783	8	69	0.9998	0.9982	0.9990	0.995
S5	99,740	97,144	14	143	0.9999	0.9985	0.9992	0.995
	+OOD	97,144	12	143	0.9999	0.9985	0.9992	0.995
S6	778	609	16	169	0.9744	0.7823	0.8679	0.9211
	+OOD	593	13	185	0.9785	0.7618	0.8567	0.8899
S7	69,238	67,470	14	99	0.9998	0.9985	0.9992	0.995
	+OOD	67,460	11	109	0.9998	0.9984	0.9991	0.995
S8	69,288	67,479	9	91	0.9999	0.9987	0.9993	0.995
	+OOD	67,468	9	102	0.9999	0.9985	0.9992	0.995

6,172 non-pedestrians (3.0%)). Analogous to Section 10.2.1, Figure 16 depicts four subplots representing  $\text{IoU} = 0.5$ : A) P vs R, B) F1 vs. Conf, C) P vs. Conf, and D) R vs. Conf. We observe that the appearance of the four subplots closely resembles the corresponding plots for the internal test data (cf. Figure 14).

Table 8 shows the output from the ML model using the Conf threshold 0.448 on the verification data. The table is organized into 1) all distances from the ego car, 2) within 80 m, and 3) within 50 m, respectively. The table also shows the effect of adding OOD detection using the autoencoder, i.e., the number of FPs is decreased just as for the internal test data. Table 9 demonstrates how the ML model satisfies the performance requirements on the verification data. Similar to the results for the internal test data, the FPPI (0.21%) is



**Fig. 16** Evaluation of the ML model on the verification data at IoU=0.5. A) P-R curve, B) F1 vs. Conf, C) P vs. Conf, and D) R vs. Conf for the internal test data.

too high to satisfy **SYS-PER-REQ3** without additional OOD detection, i.e., we observed 330 FPs (roughly an equal share of pyramids and children). The rightmost column in Table 9 shows how the FPPI decreased to 0.015% with the autoencoder. All basic shapes were rejected, instead children at a long distance with too low IoU scores dominate the FPs. We acknowledge that it is hard for the YOLOv5 to achieve a high IoU for the few pixels representing a child almost 80 m away. However, commencing emergency braking in such cases is an appropriate action – a child detected with a low IoU is not an example of the ghost braking hazard described in Section 5.2. **SYS-PER-REQ4** is satisfied as the fraction of rolling windows with more than a single FN is 2.3%. Figure 17 shows the distribution of position errors. The median error is 1.0 cm, the 99% percentile is 5.4 cm, and the largest observed error is 12.8 cm. Consequently, we show that **SYS-PER-REQ5** is satisfied for the verification data.

**Table 8** ML model accuracy on the verification data at the Conf threshold 0.448. The three rows show results for all distances, within 80 m, and within 50 m, respectively. Every second row show results for the ML model followed by OOD detection using the autoencoder.

Distance	Total	TP	FP	FN	P	R	F1	AP@0.5
All	208,884	198,457	990	4,255	0.9950	0.9790	0.9870	0.9942
+OOD		195,695	533	7,017	0.9973	0.9654	0.9811	0.9878
$\leq 80$ m	158,066	151,976	330	210	0.9978	0.9986	0.9982	0.9945
+OOD		149,214	23	2,972	0.9998	0.9905	0.9901	0.988
$\leq 50$ m	92,592	86,847	178	165	0.9980	0.9981	0.9980	0.9949
+OOD		84,085	21	2,972	0.9998	0.9805	0.9901	0.988

**Table 9** ML model satisfaction of the performance requirements on the verification data at the Conf threshold 0.448. R1–R4 = SYS-PER-REQ1–4. The rightmost column show results for the YOLOv5 model followed by OOD detection using the autoencoder.

Req.	Expected	Observed (Model)	Observed (Model+OOD)
R1	TP rate $\geq 93\%$ for $\leq 80$ m	$\frac{151,976}{158,066} = 96.1\%$	$\frac{149,214}{158,066} = 94.4\%$
R2	FN rate $\leq 7\%$ for $\leq 50$ m	$\frac{165}{92,592} = 0.18\%$	$\frac{2,927}{92,592} = 3.2\%$
R3	FPPI $\leq 0.1\%$ for $\leq 80$ m	$\frac{330}{158,066} = 0.21\%$	$\frac{23}{158,066} = 0.015\%$
R4	$\leq 3\%$ of rolling windows contain $\geq 2$ misses in 5 frames for $\leq 80$ m	$\frac{201}{152,395} = 0.13\%$	$\frac{3,499}{152,090} = 2.3\%$


**Fig. 17** Distribution of position errors for the verification data.

Table 11 presents the output of the ML model on the nine slices of the verification data defined in Section 9.1. In relation to the robustness requirements, we notice that there the accuracy is slightly lower for S9 (children). This finding is related to the size requirement **SYS-ROB-REQ3**. Table 10 contains an in-depth analysis of children at different distances with OOD detection. We confirm that most FPs occur outside of the ODD, i.e., 507 out of 512 FPs occur for children more than 80 m from ego car. In extension, we show that

**Table 10** Detailed analysis for children. The rows show results results for the ML model followed by OOD detection using the autoencoder for four distance ranges. The bottom row in italic font is outside the ODD.

Distance	Total	TP	FP	FN	P	R	F1	AP@0.5
All	69,301	63,360	512	4,174	0.992	0.9382	0.9643	0.9877
$\leq 80$ m	50,696	50,402	5	294	0.9999	0.9942	0.9971	0.995
$\leq 50$ m	30,731	28,715	3	249	0.9999	0.9914	0.9956	0.995
$> 80$ m	<i>16,838</i>	<i>12,803</i>	<i>507</i>	<i>4,035</i>	<i>0.9619</i>	<i>0.7604</i>	<i>0.8493</i>	<i>0.942</i>

**Table 11** ML model accuracy on nine slices of the verification data. S1=All data, S2=close distance, S3=far distance, S4=running pedestrians, S5=walking pedestrians, S6=occluded pedestrians, S7=males, S8=females, and S9=children. Every second rows show results for the ML model followed by OOD detection using the autoencoder.

Slice	Total	TP	FP	FN	P	R	F1	AP@0.5
S1	208,884	198,457	990	4,255	0.995	0.979	0.987	0.9942
	+OOD	195,695	533	7,017	0.9998	0.9616	0.9803	0.9878
S2	92,028	86,691	22	165	0.9997	0.9981	0.9989	0.995
	+OOD	83,929	21	2,927	0.9997	0.9663	0.9827	0.9761
S3	65,330	65,285	2	45	1	0.9993	0.9996	0.995
	+OOD	65,285	2	45	1	0.9993	0.9996	0.995
S4	58,267	56,130	110	716	0.998	0.9874	0.9927	0.9949
	+OOD	54,964	110	1,882	0.998	0.9669	0.9822	0.9818
S5	149,617	142,328	424	3,538	0.997	0.9757	0.9863	0.9949
	+OOD	140,732	423	5,134	0.997	0.9648	0.9806	0.9882
S6	1,031	866	22	165	0.9752	0.84	0.9026	0.9289
	+OOD	805	21	226	0.9746	0.7808	0.867	0.8783
S7	69,292	67,555	15	54	0.9998	0.9992	0.9995	0.995
	+OOD	65,009	14	2,600	0.9998	0.9616	0.9803	0.9741
S8	69,291	67,495	7	74	0.9999	0.9989	0.9994	0.995
	+OOD	67,482	7	87	0.9999	0.9987	0.9993	0.995
S9	69,301	63,408	512	4,126	0.992	0.9389	0.9647	0.9879
	+OOD	63,205	512	4,329	0.992	0.9359	0.9631	0.9871


the performance requirements are still satisfied for the most troublesome slice of data as follows:

- TP rate children  $\leq 80m$ :  $\frac{50,402}{50,696} = 99.4\%$
- FN rate children  $\leq 50m$ :  $\frac{249}{30,731} = 0.81\%$
- FPPI children  $\leq 80m$ :  $\frac{5}{52,463} = 0.0099\%$

The independent verification concludes that all requirements are met, based on the same argumentation as for the internal test results. The complete verification report is available on GitHub.

### 10.3 Results from System Testing [FF]

This section presents an overview of the results from testing SMIRK in ESI Pro-SiVIC, which corresponds to the Integration Testing Results in AMLAS. As explained in Section 9.2.2, we measure seven metrics for each test case



**Fig. 18** Distribution of inference speeds during system testing.

execution, i.e., MinDist, TimeTrig, DistTrig, TimeBrake, DistBrake, Coll, and CollSpeed.

Table 12 presents the results from executing the test cases representing operational scenarios with pedestrians, i.e., TC-OS-[1–25]. From the left, the columns show 1) test case ID, 2) the minimum distance between ego car and the pedestrian during the scenario, 3) the difference between TimeTrig and TimeBrake, 4) the difference between DistTrig and DistBrake, 5) whether there was a collision, 6) the speed of ego car at the collision, and 7) the initial speed of ego car. We note that 2) and 3) are 0 for all 25 test cases, showing that the pedestrian is always detected at the first possible frame when  $TTC \leq 4s$ , which means that SMIRK commenced emergency braking in all cases. Moreover, we see that SMIRK successfully avoids collisions in all but two test cases. In TC-OS-5, the pedestrian starts 20 m from ego car and runs towards it while it drives at 16 m/s – SMIRK brakes but barely reduces the speed. In TC-OS-9, the pedestrian starts only 15 m from ego car but SMIRK significantly reduces the speed by emergency braking.

The remaining system test cases corresponding to non-pedestrian operational scenarios (TC-OS-[26–38]) and all test cases with jitter (TC-RAND-[1–38]) were also executed with successful test verdicts. All scenarios with basic shapes on collision course were rejected by the safety cage architecture, i.e., SMIRK did never commence any ghost braking. In a virtual conclusion of test meeting, the first three authors concluded that TC-RBT-1 and TC-RBT-2 had passed successfully. Finally, Figure 18 shows the distribution of inference speeds during the system testing. The median inference time is 22.0 ms and the longest inference time observed is 51.6 ms. Based on these results we conclude that TC-RBT-3 passed successfully and thus provide evidence that **SYS-PER-REQ6** is satisfied. The complete system test report is available on GitHub.

**Table 12** Test results of metrics collected during execution of TC-OS-[1-18] and TC-RAND-[1-18].

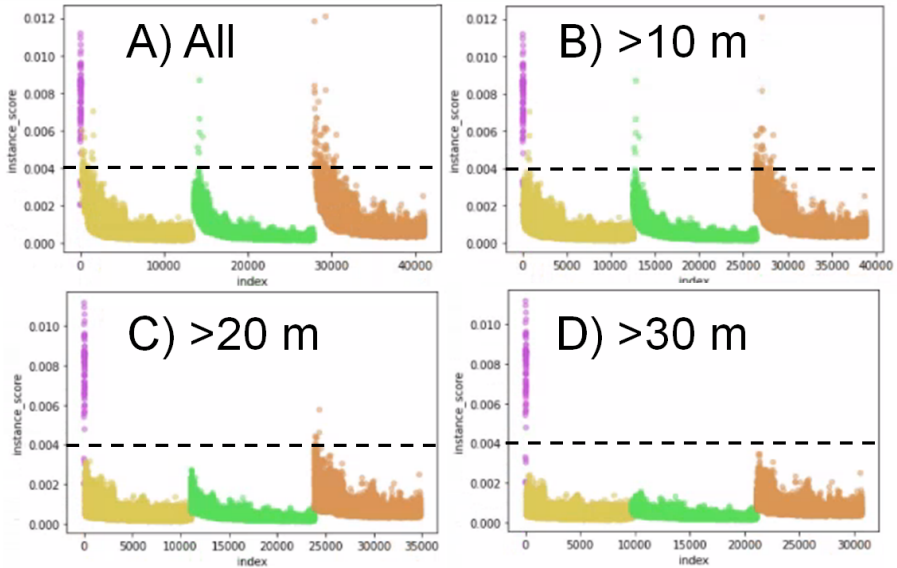
ID	MinDist(m)	$\Delta$ Time(s)	$\Delta$ Distance(m)	Collision	Collision Speed(m/s)	Initial
TC-OS-1	18	0.0	0.0	False	-	-
TC-OS-2	21	0.0	0.0	False	-	-
TC-OS-3	12	0.0	0.0	False	-	-
TC-OS-4	3	0.0	0.0	False	-	-
TC-OS-5	0	0.0	0.0	True	15.99	-
TC-OS-6	18	0.0	0.0	False	-	-
TC-OS-7	34	0.0	0.0	False	-	-
TC-OS-8	28	0.0	0.0	False	-	-
TC-OS-9	0	0.0	0.0	True	3.11	-
TC-OS-10	18	0.0	0.0	False	-	-
TC-OS-11	33	0.0	0.0	False	-	-
TC-OS-12	27	0.0	0.0	False	-	-
TC-OS-13	16	0.0	0.0	False	-	-
TC-OS-14	19	0.0	0.0	False	-	-
TC-OS-15	1	0.0	0.0	False	-	-
TC-OS-16	18	0.0	0.0	False	-	-
TC-OS-17	32	0.0	0.0	False	-	-
TC-OS-18	17	0.0	0.0	False	-	-
TC-OS-19	25	0.0	0.0	False	-	-
TC-OS-20	29	0.0	0.0	False	-	-
TC-OS-21	28	0.0	0.0	False	-	-
TC-OS-22	4	0.0	0.0	False	-	-
TC-OS-23	16	0.0	0.0	False	-	-
TC-OS-24	30	0.0	0.0	False	-	-
TC-OS-25	36	0.0	0.0	False	-	-

## 10.4 Erroneous Behaviour Log [DD]

As prescribed by AMLAS, the characteristics of erroneous outputs shall be predicted and documented. This section presents the key observations from internal testing of the ML model, independent verification activities, and system testing in ESI Pro-SiVIC. The findings can be used to design appropriate responses by other vehicular systems in the SMIRK context.

Tables 7 and 11 show that the AP@0.5 are lower for occluded pedestrians (S6). As occlusion is an acknowledged challenge for object detection, which we previously have studied for automotive pedestrian detection (Henriksson et al, 2021b), this is an expected result. Table 11 also reveals that the number of FPs and FNs for the S9 slice (children) is relatively high, resulting in slightly lower AP@0.5. Table 10 shows that the problem with children is primarily far away, explained by the few pixels available for the object detection at long distances. While the SMIRK fulfils the robustness requirements within the ODD, we recognize this perception issue in the erroneous behavior log.

During the iterative SMIRK development (cf. E) in Figure 4), it became evident that OOD detection using the autoencoder was inadequate at close range. Figure 19 shows reconstruction errors (on the y-axis) for all objects in the validation subset of the development data at A) all distances, B) > 10 m, C) > 20 m, and D) > 30 m. The visualization clearly shows that the



**Fig. 19** Reconstruction errors for different objects on the validation subset of the development data at different distances from ego car (magenta=cylinder, yellow=female business casual, green=male business, orange=male casual). The dashed lines show the threshold for rejecting objects. In SMIRK, we use alternative B) in the safety cage.

autoencoder cannot convincingly distinguish the cylinders from the pedestrians at all distances (in subplot A), different objects appear above the threshold), but the OOD detection is more accurate when objects at close distance are excluded (subplot D) displays high accuracy). Based on validation of the four distances, comparing the consequences of the trade-off between safety cage availability and accuracy, the design decision for SMIRK’s autoencoder is to only perform OOD detection for objects that are at least 10 m away. We explain the less accurate behaviour at close range by limited training data, a vast majority of images contain pedestrians at a larger distance – which is reasonable since the SMIRK OOD is limited to rural country roads.

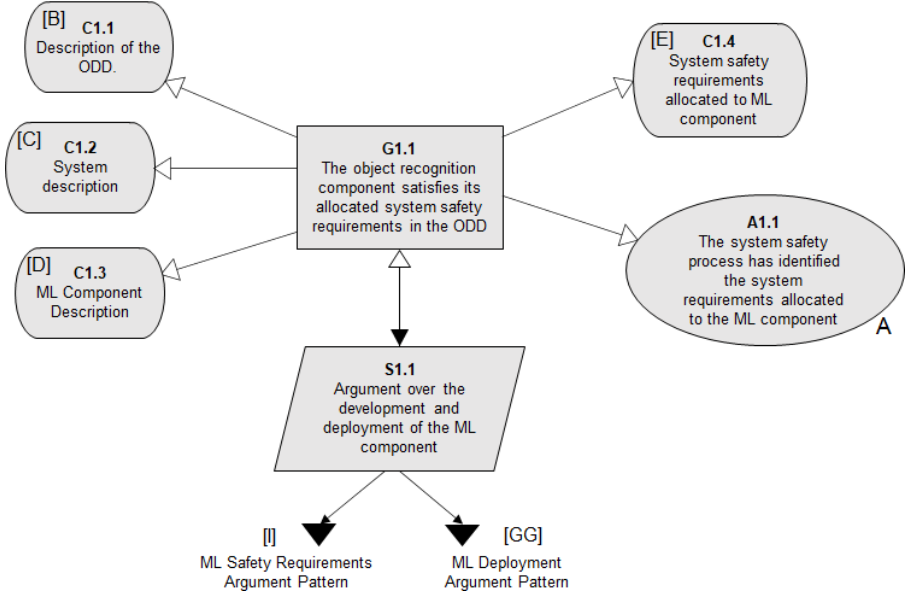
## 11 AMLAS Safety Assurance

This section describes the complete SMIRK safety argumentation organized by the six AMLAS stages. For each step, we present an argument pattern using GSN notation ([Assurance Case Working Group, 2021](#)) and present the final argument in a text box.

### 11.1 Phase 1: Machine Learning Assurance Scoping

Figure 20 shows the overall ML assurance scoping argument pattern for SMIRK. The pattern, as well as all subsequent patterns in this paper, follows



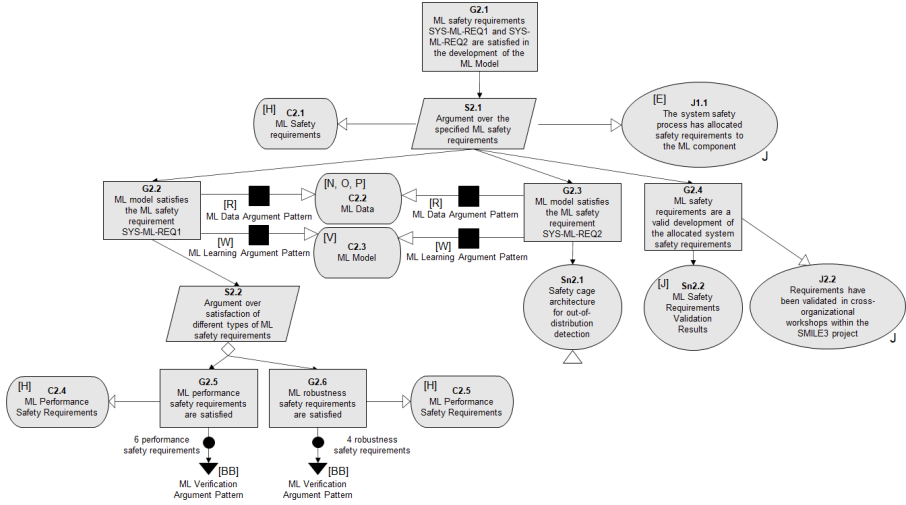


**Fig. 20** ML Assurance Scoping Argument Pattern [F].

the examples provided in AMLAS, but adapts it to the specific SMIRK case. Furthermore, we provide evidence that supports our arguments.

The top claim, i.e., the starting point for the safety argument for the ML-based component, is that the system safety requirements that have been allocated to the pedestrian recognition component are satisfied in the ODD (G1.1). The safety claim for the pedestrian recognition component is made within the context of the information that was used to establish the safety requirements allocation, i.e., the system description ([C]), the ODD ([B]), and the ML component description ([D]). The allocated system safety requirements ([E]) are provided as context. An explicit assumption is made that the allocated safety requirements have been correctly defined (A1.1), as this is part of the overall system safety process (FuSa and SOTIF) preceding AMLAS. Our claim to the validity of this assumption is presented in relation to the HARA described in [E]. As stated in AMLAS, “the primary aim of the ML Safety Assurance Scoping argument is to explain and justify the essential relationship between, on the one hand, the system-level safety requirements and associated hazards and risks, and on the other hand, the ML-specific safety requirements and associated ML performance and failure conditions.”

The ML safety claim is supported by an argument split into two parts. First, the development of the ML component is considered with an argument that starts with the elicitation of the ML safety requirements. Second, the deployment of the ML component is addressed with a corresponding argument.



**Fig. 21** ML Safety Requirements Argument Pattern [I].

### ML Safety Assurance Scoping Argument [G]

SMIRK instantiates the ML safety assurance scoping argument through the artifacts listed in the Table 13. The set of artifacts constitutes the safety case for SMIRK's ML-based pedestrian recognition component.

## 11.2 Phase 2: Machine Learning Requirements Assurance

Figure 21 shows the ML Safety Requirements Argument Pattern [I]. The top claim is that system safety requirements that have been allocated to the ML component are satisfied by the model that is developed (G2.1). The argument approach is a refinement strategy translating the allocated safety requirements into two concrete ML safety requirements (S2.1) provided as context (C2.1). Justification J2.1 explains how we allocated safety requirements to the ML component as part of the system safety process, including the HARA.

Strategy S2.1 is refined into two subclaims about the validity of the ML safety requirements corresponding to missed pedestrians and ghost braking, respectively. Furthermore, a third subclaim concerns the satisfaction of those requirements. G2.2 focuses on the ML safety requirement **SYS-ML-REQ1**, i.e., that the nominal functionality of the pedestrian recognition component shall be satisfactory. G2.2 is considered in the context of the ML data (C2.2) and the ML model (C2.3), which in turn are supported by the ML Data Argument Pattern [R] and the ML Learning Argument Pattern [W]. The argumentation strategy (S2.2) builds on two subclaims related to two types of safety requirements with respect to safety-related outputs, i.e., performance requirements (G2.5 in context of C2.4) and robustness requirements (G2.6 in

**Table 13** SMIRK Safety Assurance Table. Numbers in the Input/Output columns refer to the AMLAS stages in Figure 3. (TBD = AI Sweden has agreed to host our annotated data (185 GB) and provide it under a CC-BY-NC 4.0 licence, but the permalink is not ready yet.)

ID	AMLAS Artifact Title	Input to	Output from	Where?
[A]	System Safety Requirements	1, 6		Sec. 5.1
[B]	Description of Operating Environment of System	1, 6		Sec. 5.4
[C]	System Description	1, 6		Sec. 4
[D]	ML Component Description	1		Sec. 8
[E]	Safety Requirements Allocated to ML Component	2	1	Sec. 5.2
[F]	ML Assurance Scoping Argument Pattern	1		Fig. 20
[G]	ML Safety Assurance Scoping Argument		1	Sec. 11.1
[H]	ML Safety Requirements	3, 4, 5	2	Sec. 5.3
[I]	ML Safety Requirements Argument Pattern	2		Fig. 21
[J]	ML Safety Requirements Validation Results		2	Sec. 3.3.1
[K]	ML Safety Requirements Argument		2	Sec. 11.2
[L]	Data Requirements		3	Sec. 7.1
[M]	Data Requirements Justification Report		3	Sec. 3.3.1
[N]	Development Data		3	TBD
[O]	Internal Test Data		3	TBD
[P]	Verification Data		3	TBD
[Q]	Data Generation Log		3	Sec. 7.2
[R]	ML Data Argument Pattern	3		Fig. 22
[S]	ML Data Validation Results		3	Sec. 10.1
[T]	ML Data Argument		3	Sec. 11.3
[U]	Model Development Log		4	Sec. 8.2
[V]	ML Model	5, 6	4	GitHub repo
[W]	ML Learning Argument Pattern	4		Fig. 23
[X]	Internal Test Results		4	Sec. 10.2.1
[Y]	ML Learning Argument		4	Sec. 11.4
[Z]	ML Verification Results		5	Sec. 10.2.2
[AA]	Verification Log		5	Sec. 9.1
[BB]	ML Verification Argument Pattern	5		Fig. 24
[CC]	ML Verification Argument		5	Sec. 11.5
[DD]	Erroneous Behaviour Log		6	Sec. 10.4
[EE]	Operational scenarios	6		Sec. 9.2.1
[FF]	Integration Testing Results		6	Sec. 10.3
[GG]	ML Deployment Argument Pattern	6		Fig. 25
[HH]	ML Deployment Argument		6	Sec. 11.6

context of C2.5). The satisfaction of both G2.5 and G2.6 are addressed by the ML Verification Argument Pattern [BB].

Subclaim G2.3 focuses on the ML safety requirement **SYS-ML-REQ2**, i.e., that the pedestrian recognition component shall reject input that does not resemble the training data to avoid ghost braking. G2.3 is again considered in the context of the ML data (C2.2) and the ML model (C2.3). For SMIRK, the solution is the safety cage architecture (Sn2.1) developed in the SMILE research program (Henriksson et al, 2021a) (described in Section 8.3).

Subclaim G2.4 states that the ML safety requirements are a valid development of the allocated system safety requirements. The justification (J2.2) is that the requirements have been validated in cross-organizational workshops within the SMILE III research project. We provide evidence through ML Safety Requirements Validation Results [J] originating in a Fagan inspection (Sn2.2).

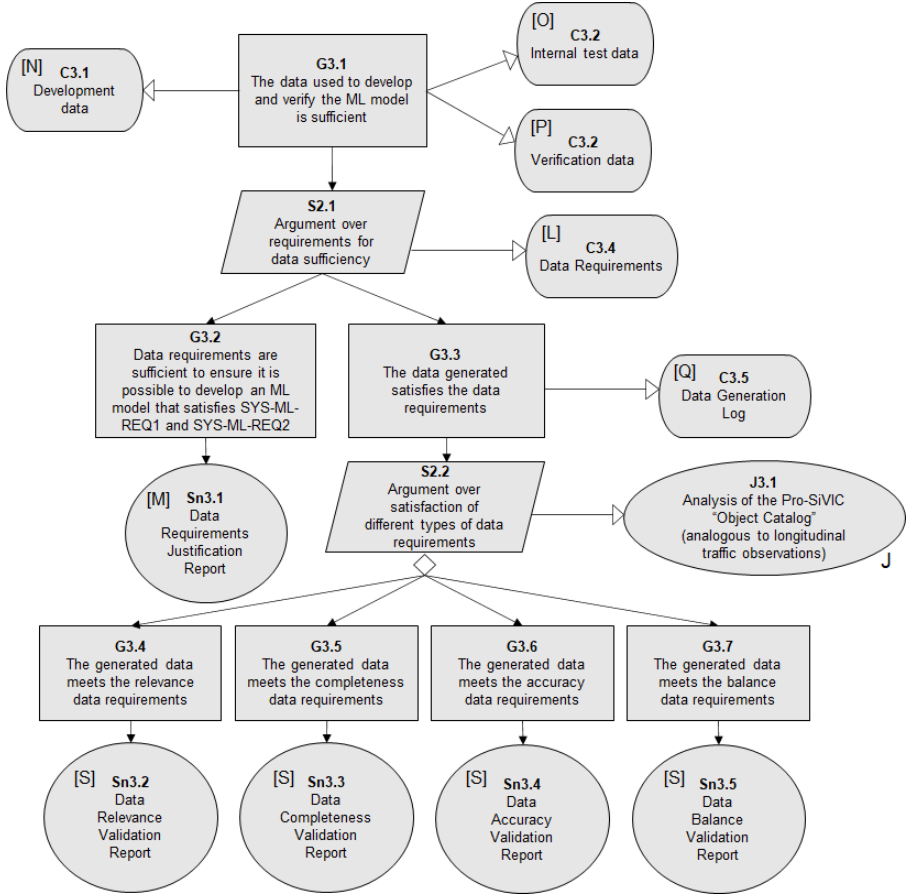
#### **ML Safety Requirements Argument [K]**

SMIRK instantiates the ML safety requirements argument through a subset of the artifacts listed in Table 13, i.e., ML Safety Requirements Argument Pattern [I], as well as: Safety Requirements Allocated to ML Component [E], ML Safety Requirements [H], and ML Safety Requirements Validation Results [J].

### **11.3 Phase 3: Data Management Assurance**

Figure 22 shows the ML Data Argument Pattern [R]. The top claim is that the data used during the development and verification of the ML model is sufficient (G3.1). This claim is made for all three data sets: development data [N], internal test data [O], and verification data [P]. The argumentation strategy (S2.1) involves how the sufficiency of these data sets is demonstrated given the Data Requirements [L]. The strategy is supported by arguing over subclaims demonstrating sufficiency of the Data Requirements (G3.2) and that the Data Requirements are satisfied (G3.3). Claim G3.2 is supported by evidence in the form of a data requirements justification report [M]. As stated in AMLAS, “It is not possible to claim that the data alone can guarantee that the ML safety requirements will be satisfied, however the data used must be sufficient to enable the model that is developed to do so.”

Claim G3.3 states that the generated data satisfies the data requirements in context of the decisions made during data collection. The details of the data collection, along with rationales, are recorded in the Data Generation Log [Q]. The argumentation strategy (S2.2) uses refinement mapping to the assurance-related desiderata of the data requirements. The refinement of the desiderata into concrete data requirements for the pedestrian recognition component of SMIRK, given the ODD, is justified by an analysis of the expected traffic agents and objects that can appear in ESI Pro-SiVIC. For each subclaim corresponding to a desideratum, i.e., relevance (G3.4), completeness (G3.5),



**Fig. 22** ML Data Argument Pattern [R].

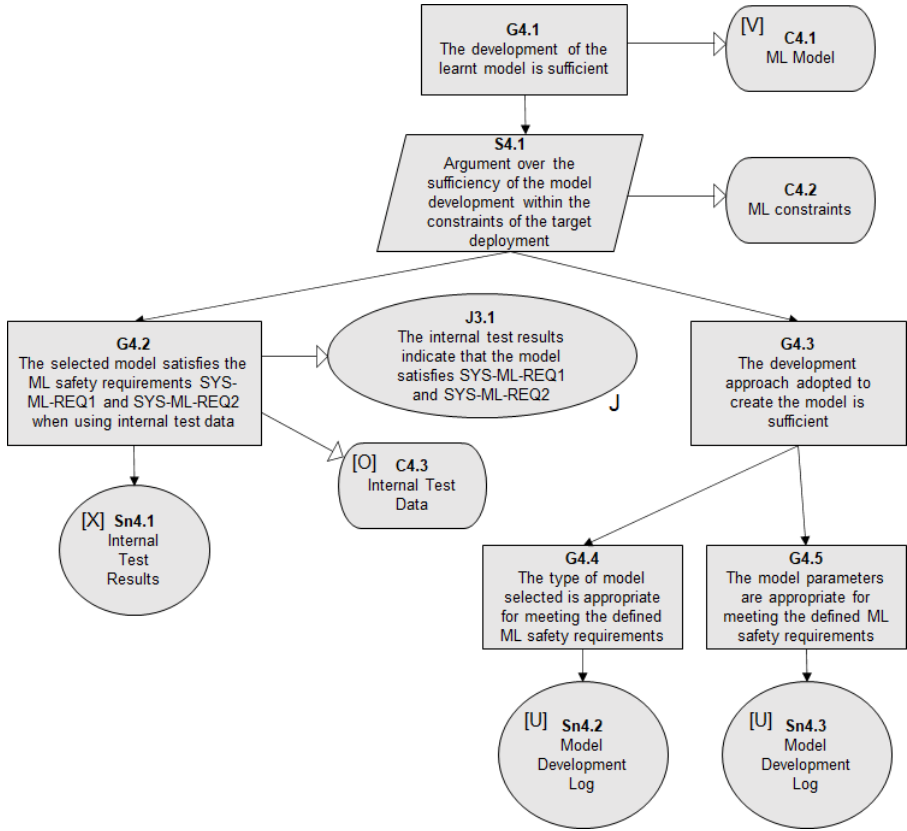
accuracy (G3.6), and balance (G3.7), there is evidence in a matching section in the ML Data Validation Report [S].

### ML Data Argument [T]

SMIRK instantiates the ML Data Argument through a subset of the artifacts listed in Table 13, i.e., the ML Data Argument Pattern [R], as well as: ML Safety Requirements [H], Data Requirements [L], Data Requirements Justification Report [M], Development Data [N], Internal Test Data [O], Verification Data [P], Data Generation Log [Q], and ML Data Validation Results [S].

## 11.4 Phase 4: Model Learning Assurance

Figure 23 shows the ML Learning Argument Pattern [W]. The top claim (G4.1) is that the development of the learned model [V] is sufficient. The



**Fig. 23** ML Learning Argument Pattern [W].

strategy is to argue over the internal testing of the model and that the ML development was appropriate (S4.1) in context of creating a valid model that meets practical constraints such as real-time performance and cost (C4.2). Subclaim (G4.2) is that the ML model satisfies the ML safety requirements when using the internal test data [O]. We justify that the internal test results indicate that the ML model satisfies the ML safety requirements (J3.1) by presenting evidence from the internal test results [X].

Subclaim G4.3 addresses the approach that was used when developing the model. This claim is in turn supported by two additional subclaims regarding the type of model selected and the model parameters selected, respectively. First, G4.4 claims that the type of model is appropriate for the specified ML safety requirements and the other model constraints. ML development processes, including transfer learning, are highly iterative thus rationales for development decisions must be recorded. Second, G4.5 claims that the parameters of the ML model are appropriately selected to tune performance toward the object detection task within the specified ODD. Rationales for all relevant decisions in G4.4 and G4.5 are recorded in the model development log [U].

**ML Learning Argument [T]**

SMIRK instantiates the ML Learning Argument through a subset of the artifacts listed in Table 13, i.e., the ML Learning Argument Pattern [W], as well as: ML Safety Requirements [H], Development Data [N], and Internal Test Data [O].

**11.5 Phase 5: Model Verification Assurance**

Figure 24 shows the ML Verification Argument Pattern [BB]. The top claim (G5.1) corresponds to the bottom claim in the safety requirements argument pattern [I], i.e., that all ML safety requirements are satisfied. The argumentation builds on a subclaim and an argumentation strategy. First, subclaim G5.2 is that the verification of the ML model is independent of its development. The verification log [AA] specifies how this has been achieved for SMIRK (Sn5.1). Second, the strategy S5.1 argues that test-based verification is an appropriate approach to generate evidence that the ML safety requirements are met. The justification (J5.1) is that the SMIRK test strategy follows the proposed organization in peer-reviewed literature on ML testing, which is a better fit than using less mature formal methods for ML models as complex as YOLOv5.

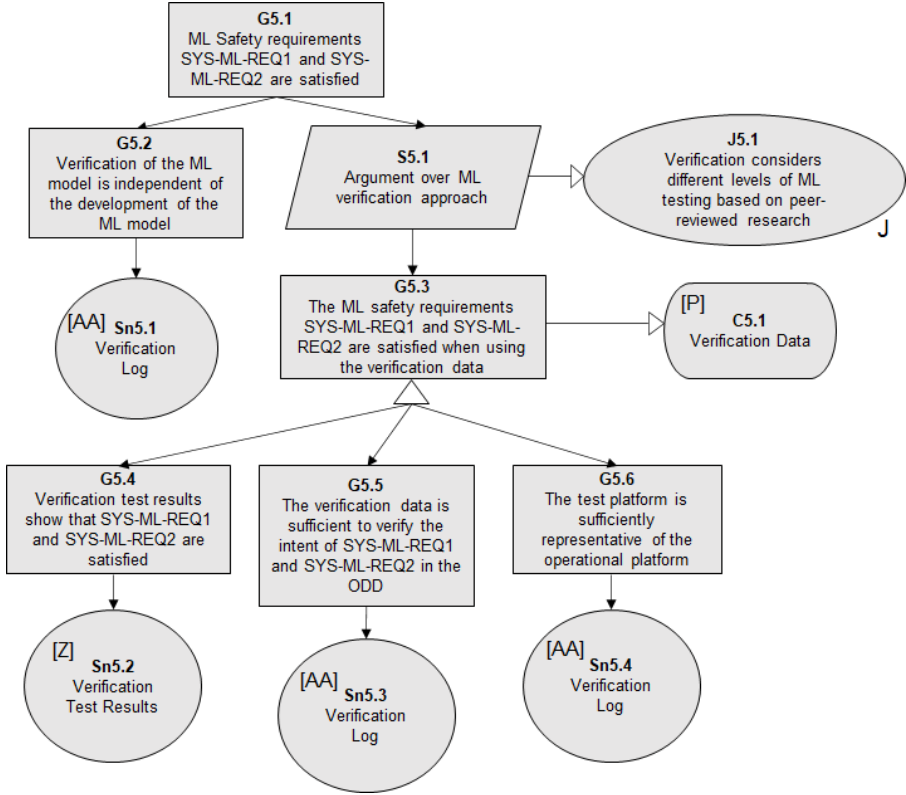
Following the test-based verification approach, the subclaim G5.3 argues that the ML model satisfies the ML safety requirements when the verification data (C5.1) is applied. The testing claim is supported by three subclaims. First, G5.4 argues that the test results demonstrate that the ML safety requirements are satisfied, for which Verification Test Results [Z] are presented as evidence. Second, G5.5 argues that the Verification Data [P] is sufficient to verify the intent of the ML safety requirements in the ODD. Third, G5.6 argues that the test platform is representative of the operational platform. Evidence for both G5.5 and G5.6 is presented in the Verification Log [AA].

**ML Verification Argument [CC]**

SMIRK instantiates the ML Verification Argument through a subset of the artifacts listed in Table 13, i.e., the ML Verification Argument Pattern [W], as well as: ML Safety Requirements [H], Verification Data [P], and the ML Model [V].

**11.6 Phase 6: Model Deployment Assurance**

Figure 25 shows the ML Verification Argument Pattern [GG]. The top claim (G6.1) is that the ML safety requirements **SYS-ML-REQ1** and **SYS-ML-REQ2** are satisfied when deployed to the ego car in which SMIRK operates. The argumentation strategy S6.1 is two-fold. First, subclaim G6.2 is that the ML safety requirements are satisfied under all defined operating scenarios when the ML component is integrated into SMIRK in the context (C6.1) of the specified operational scenarios [EE]. Justification J6.1 explains that the

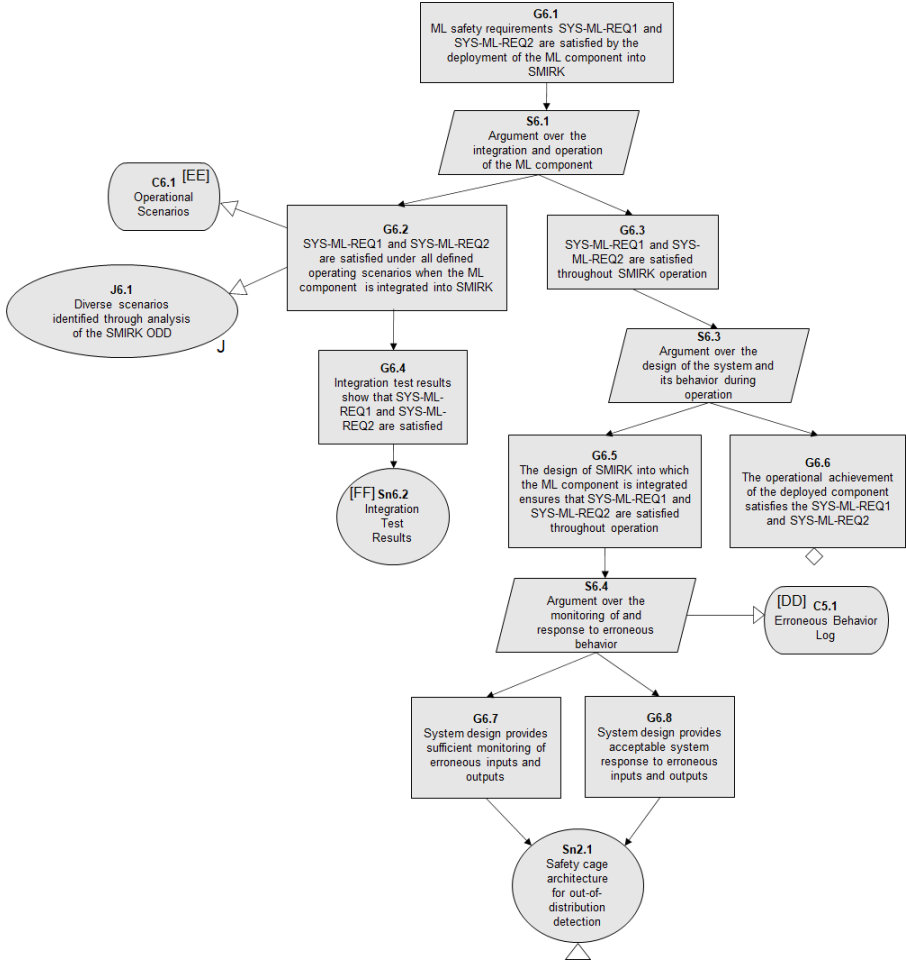


**Fig. 24** ML Verification Argument Pattern [BB].

scenarios were identified through an analysis of the SMIRK ODD. G6.2 has another subclaim (G6.4), arguing that the integration test results [FF] show that **SYS-ML-REQ1** and **SYS-ML-REQ2** are satisfied.

Second, subclaim G6.3 argues that **SYS-ML-REQ1** and **SYS-ML-REQ2** continue to be satisfied during the operation of SMIRK. The supporting argumentation strategy (S6.3) relates to the design of SMIRK and is again two-fold. First, subclaim G6.6 argues that the operational achievement of the deployed component satisfies the ML safety requirements. Second, subclaim G6.5 argues that the design of SMIRK into which the ML component is integrated ensures that **SYS-ML-REQ1** and **SYS-ML-REQ2** are satisfied throughout operation. The corresponding argumentation strategy (S6.4) is based on demonstrating that the design is robust by taking into account identified erroneous behavior in the context (C5.1) of the Erroneous Behavior Log [DD]. More specifically, the argumentation entails that predicted erroneous behavior will not result in the violation of the ML safety requirements. This is supported by two subclaims, i.e., that the system design provides sufficient monitoring of erroneous inputs and outputs (G6.7) and that the system design provides acceptable response to erroneous inputs and outputs (G6.8).





**Fig. 25** ML Deployment Argument Pattern [GG].

Both G6.7 and G6.8 are addressed by the safety cage architecture that monitors input through OOD detection using an autoencoder that rejects anomalies accordingly. The acceptable system response is to avoid emergency braking and instead let the human driver control ego car.

### ML Verification Argument [HH]

SMIRK instantiates the ML Deployment Argument through a subset of the artifacts listed in Table 13, i.e., the ML Deployment Argument Pattern [GG], as well as: System Safety Requirements [A], Environment Description [B], System Description [C], ML Model [V], Erroneous Behaviour Log [DD], Operational Scenarios [EE], and Integration Testing Results [FF].

## 12 Lessons Learned and Limitations

Long development projects lead to ample experience and many lessons learned. In this section, we share the lessons learned we believe are the most valuable for external readers. Furthermore, we discuss the primary limitations of our work and the most important threats to validity.

Using a simulator to create data sets limits the validity of the negative examples. On one hand, our data generation scripts enable substantial freedom and cheap access to data. On the other hand, there is barely any variation in the scenarios (apart from clouds moving on the skydome) as would be the case for naturalistic data. As anything that is not a pedestrian in our data is a *de facto* negative example (see rationale for **DAT-BAL-REQ3**), and nothing ever appears in our simulated scenarios unless we add it in our scripts, the diversity of our negative examples is very limited. Our approach to negative examples in the development data, referred to as “background images” in Section 7.2.3, involved including the outlier example Cylinder [N5]. From experiments on the validation subset of the development data, we found that adding frames with cylinders representing negative examples was essential to let the model distinguish between pedestrians and basic shapes. For ML components designed for use in the real world, trained on outcomes from real data collection campaigns, the natural variation of the negative examples would be completely different. When working with synthetic data from simulators, how to specify data requirements on negative examples remains an open question.

Evaluation of object detection models is non-trivial. We spent substantial time to align the understanding within the project and we believe other development endeavors will need to do the same. In particular, we observed that the definition of TP, FP, TN, and FN based on IoU (explained in Section 9.1) is difficult to grasp for novices. The fact that FPs appear due to low IoU scores despite parts of a pedestrian indeed is detected is often counter-intuitive, i.e., “how can a detected pedestrian ever be a FP?” To align the development team, organizations should ensure that the true meaning of those KPIs are communicated as part of internal training. In the same vein, FP rate is not a valid metric (as TPs do not exist) whereas FN rate is used in **SYS-PER-REQ2** – again internal training is important to align the understanding. What intuitively is perceived as a FP on the system level is not the same as a FP on the ML model level. To make the distinction clear, we restrict the use of FPs to the model level and refer to incorrect braking on the system level as “ghost braking.”

ML model selection post learning involves fundamental decisions. Model selection is an essential activity in ML. When training ML models over several epochs, the best performing model given some criterion shall be kept. Also, when training alternative models with alternative architectures or hyperparameter settings, there must be a way to select the best candidate. How to tailor a fitness function to quantitatively measure what “best” involves is a delicate engineering effort with inevitable tradeoffs. The default fitness function in YOLOv5 puts 10% of the weight at AP@0.5 and 90% at Mean AP for

a range of ten IoU values between 0.5 to 0.95. It would be possible to further customize the fitness function to also cover fairness aspects, i.e., to already during model selection value models that fulfill various quality aspects. There is no upper limit to the possible complexity, as this could encompass gender, size, ODD aspects etc. For SMIRK, however, we decided to do the opposite, i.e., to prioritize simplicity to gain explainability. As explained in Section 9.1, our fitness function solely uses AP@0.5.

OOD scores can be measured for different parts of an image. What pixels to send to the autoencoder is another important design decision. Initially, we used the entire image as input to the autoencoder, which showed promising results in detecting major changes in the environmental conditions, e.g., leaving the ODD due to nightfall or heavy fog. However, it quickly became evident that this input generated too small differences in the autoencoder’s reconstruction error between inliers and outliers, i.e., it was not a feasible approach to reject basic shapes. We find this to be in line with how the “curse of dimensionality” affects unsupervised anomaly detection in general (Zimek et al, 2012) – the anomalies we try to find are dwarfed by the background information. Instead, we decided to focus on squares (a good shape for the autoencoder) containing pixels close to the bounding box of the detected object, and tried three solutions: 1) extracting a square centered on the middle pixel, 2) extracting the entire bounding box and padding with gray pixels to make it a square, and 3) stretching the contents of the bounding box to fit a rectangle matching the average aspect ratio of pedestrians in the development set. The third approach was the most successful in our study, and is now used in SMIRK. Future OOD architectures will likely combine different selection of the input images.

The fidelity of the radar signatures in the simulator matters. While it is easy for a human to tell how realistic the visual appearance of objects are in ESI Pro-SiVIC, assessing the appropriateness of its radar signature model requires a sensor expert. In SMIRK, we attached the same radar signature to all pedestrians, i.e., the one provided for human bodies in the object catalog. For all basic shapes, on the other hand, we attach the same simplistic spherical radar signature. Designing customized signatures is beyond the scope of our project, thus we acknowledge this limitation as a threat to validity. It is possible that system testing results would have been different if more elaborate radar signatures were used.

Python is not a valid choice of programming language for safety-critical development. We are well aware that Python is not an ideal choice for development of safety-critical applications. Python is dynamically typed and might throw type errors at runtime. For our SMIRK demonstrator, the reasons we chose Python is two-fold. First, using Python enables easy access to numerous state-of-the-art ML libraries. Second, Python is the dominating language in the research community and thus others can more easily build on our work. A real-world in-vehicle implementation would lead to another language choice, e.g., adhering to MISRA C (Motor Industry Software Reliability Association

et al, 2012), a widely accepted set of software development guidelines for using the C programming language in safety-critical systems.

## 13 Conclusion and Future Work

Safe ML is going to be fundamental when increasing the level of vehicle automation. Several automotive standardization initiatives are ongoing to allow safety certification for ML in road vehicles, e.g., ISO 21448 SOTIF. However, standards provide high-level requirements that must be operationalized in each development context. Unfortunately, there is a lack of publicly available ML-based automotive demonstrator systems that can be used to study safety case development.

We present SMIRK, a PAEB designed for operation in the industry-grade simulator ESI Pro-SiVIC, available on GitHub under an OSS license. SMIRK uses a radar sensor for object detection and an ML-based component relying on a DNN for pedestrian recognition. Originating in SMIRK’s minimalistic ODD, we present a complete safety case for its ML-based component by following the AMLAS framework. To the best of our knowledge, this work constitutes the first complete application of AMLAS independent from its authors. We conclude that even for a very restricted ODD, the size of the ML safety case is considerable, i.e., there are many aspects of the AI engineering that must be clearly explained.

Based on this engineering research project, representing industry-academia collaboration in Sweden, we report several lessons learned. First, using a simulator to create synthetic data sets for ML training particularly limits the validity of the negative examples. Second, the complexity of object detection evaluations necessitates internal training within the project team. Third, composing the fitness function used for model selection is a delicate engineering activity that forces explicit tradeoff decisions. Fourth, what parts of an image to send to an autoencoder for OOD detection is an open question – for SMIRK, we stretch the content of bounding boxes to a larger square.

Thanks to the complete safety case, SMIRK can be used as a starting point for several avenues of future research. First, the SMIRK MVP enables studies on efficient and effective approaches to conduct safety assurance for ODD extension (Weissensteiner et al, 2021). In this context, SMIRK could be used as a platform to study dynamic safety cases (Denney et al, 2015), i.e., updating the safety case as the system evolves, and reuse of safety evidence for new operational contexts (de la Vara et al, 2019). Second, SMIRK could be used as a realistic test benchmark for automotive ML testing. The testing community has largely worked on offline testing of single frames, but we know that this is insufficient (Haq et al, 2021). Third, we recommend the community to port SMIRK to other simulators beyond ESI Pro-SiVIC. As we investigated in previous work, highly running highly similar test scenarios in different simulators can lead to considerably different results (Borg et al, 2021a) – further exploring this phenomenon using SMIRK would be a valuable research direction.

## Acknowledgment

Thanks go to ESI Group for supporting us with technical details along the way, especially Erik Abenius and François-Xavier Jegeden. This work was carried out within the SMILE III project financed by Vinnova, FFI, Fordonsstrategisk forskning och innovation under the grant number 2019-05871 and partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

## References

- An J, Cho S (2015) Variational autoencoder based anomaly detection using reconstruction probability. Special Lecture on IE 2(1):1–18
- Ashmore R, Calinescu R, Paterson C (2021) Assuring the machine learning lifecycle: Desiderata, methods, and challenges. *ACM Computing Surveys (CSUR)* 54(5):1–39
- Assurance Case Working Group (2021) Goal Structuring Notation Community Standard (Version 3). Tech. Rep. SCSC-141C, Safety-Critical Systems Club, UK
- Barr ET, Harman M, McMinn P, et al (2014) The oracle problem in software testing: A survey. *IEEE Transactions on Software Engineering* 41(5):507–525
- Ben Abdesslem R, Nejati S, Briand LC, et al (2016) Testing advanced driver assistance systems using multi-objective search and neural networks. In: *Proc. of the 31st IEEE/ACM International Conference on Automated Software Engineering*, pp 63–74, <https://doi.org/10.1145/2970276.2970311>
- Bolya D, Foley S, Hays J, et al (2020) Tide: A general toolbox for identifying object detection errors. In: *European Conference on Computer Vision*, Springer, pp 558–573
- Borg M, Englund C, Wnuk K, et al (2019) Safely entering the deep: A review of verification and validation for machine learning and a challenge elicitation in the automotive industry. *Journal of Automotive Software Engineering* 1(1):1–19
- Borg M, Abdesslem RB, Nejati S, et al (2021a) Digital twins are not monozygotic: Cross-replicating ADAS testing in two industry-grade automotive simulators. In: *2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST)*, IEEE, pp 383–393
- Borg M, Bronson J, Christensson L, et al (2021b) Exploring the assessment list for trustworthy AI in the context of advanced driver-assistance systems. In: *2021 IEEE/ACM 2nd International Workshop on Ethics in Software*

Engineering Research and Practice (SEthics), IEEE, pp 5–12

Bosch J, Olsson HH, Crnkovic I (2021) Engineering AI systems: A research agenda. In: Artificial Intelligence Paradigms for Smart Cyber-Physical Systems. IGI global, p 1–19

Chen L, Babar MA, Nuseibeh B (2012) Characterizing architecturally significant requirements. IEEE Software 30(2):38–45

Denney E, Pai G, Habli I (2015) Dynamic safety cases for through-life safety assurance. In: 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering, IEEE, pp 587–590

Dollar P, Wojek C, Schiele B, et al (2011) Pedestrian detection: An evaluation of the state of the art. IEEE Transactions on Pattern Analysis and Machine Intelligence 34(4):743–761

Fagan M (1976) Design and code inspections to reduce errors in program development. IBM Systems Journal 15(3):182–211

Gauerhof L, Hawkins R, Picardi C, et al (2020) Assuring the safety of machine learning for pedestrian detection at crossings. In: International Conference on Computer Safety, Reliability, and Security, Springer, pp 197–212

Haq FU, Shin D, Nejati S, et al (2021) Can offline testing of deep neural networks replace their online testing? Empirical Software Engineering 26(5):1–30

Hauer F, Schmidt T, Holzmüller B, et al (2019) Did we test all scenarios for automated and autonomous driving systems? In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), IEEE, pp 2950–2955

Hawkins R, Paterson C, Picardi C, et al (2021) Guidance on the assurance of machine learning in autonomous systems (amlas). Tech. Rep. Version 1.1, Assuring Autonomy International Programme (AAIP), University of York

Henriksson J, Berger C, Borg M, et al (2019) Towards structured evaluation of deep neural network supervisors. In: 2019 IEEE International Conference On Artificial Intelligence Testing (AITest), IEEE, pp 27–34

Henriksson J, Berger C, Borg M, et al (2021a) Performance analysis of out-of-distribution detection on trained neural networks. Information and Software Technology 130:106,409

Henriksson J, Berger C, Ursing S (2021b) Understanding the impact of edge cases from occluded pedestrians for ML systems. In: 2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA),

IEEE, pp 316–325

High-Level Expert Group on Artificial Intelligence (2019) Ethics guidelines for trustworthy AI. Tech. rep., Directorate-General for Communications Networks, Content and Technology, European Commission

Horkoff J (2019) Non-functional requirements for machine learning: Challenges and new directions. In: 2019 IEEE 27th International Requirements Engineering Conference (RE), IEEE, pp 386–391

IEEE (1998) IEEE recommended practice for software requirements specifications. Tech. Rep. IEEE 830-1998, Institute of Electrical and Electronics Engineers

ISO (2019) Road Vehicles - Safety of the Intended Functionality. Tech. Rep. ISO/PAS 21448:2019, International Organization for Standardization

Käpyaho M, Kauppinen M (2015) Agile requirements engineering with prototyping: A case study. In: 2015 IEEE 23rd International requirements engineering conference (RE), IEEE, pp 334–343

Kassab M, Kilicay-Ergin N (2015) Applying analytical hierarchy process to system quality requirements prioritization. *Innovations in Systems and Software Engineering* 11(4):303–312

Kästner C (2022) Machine Learning in Production. Self published under a Creative Commons 4.0 BY-SA license, <https://ckaestne.medium.com/machine-learning-in-production-book-overview-63be62393581>

Klaise J, Van Looveren A, Cox C, et al (2020) Monitoring and explainability of models in production. In: Proc. of the ICML Workshop on Challenges in Deploying and Monitoring Machine Learning Systems

Kochanthara S, Rood N, Saberi AK, et al (2021) A functional safety assessment method for cooperative automotive architecture. *Journal of Systems and Software* 179:110,991

Kruchten PB (1995) The 4+1 view model of architecture. *IEEE Software* 12(6):42–50

Lin TY, Maire M, Belongie S, et al (2014) Microsoft COCO: Common objects in context. In: European Conference on Computer Vision, Springer, pp 740–755

Liu S, Qi L, Qin H, et al (2018) Path aggregation network for instance segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp 8759–8768

- Masuda S (2017) Software testing design techniques used in automated vehicle simulations. In: 2017 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW), IEEE, pp 300–303
- Motor Industry Software Reliability Association, et al (2012) MISRA-C guidelines for the use of the C language in critical systems
- Petersson H, Thelin T, Runeson P, et al (2004) Capture-recapture in software inspections after 10 years research: Theory, evaluation and application. *Journal of Systems and Software* 72(2):249–264
- Picardi C, Paterson C, Hawkins RD, et al (2020) Assurance argument patterns and processes for machine learning in safety-related systems. In: *Proceedings of the Workshop on Artificial Intelligence Safety (SafeAI 2020)*, CEUR Workshop Proceedings, pp 23–30
- Preschern C, Kajtazovic N, Kreiner C (2015) Building a safety architecture pattern system. In: *Proceedings of the 18th European Conference on Pattern Languages of Program*, pp 1–55
- Rajput M (2020) YOLO V5 – Explained and Demystified. <https://towardsai.net/p/computer-vision/yolo-v5%E2%80%8A-%E2%80%8A8Aexplained-and-demystified>
- Ralph P, Ali Nb, Baltes S, et al (2020) Empirical standards for software engineering research. *arXiv preprint arXiv:201003525*
- Redmon J, Divvala S, Girshick R, et al (2016) You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp 779–788
- Riccio V, Jahangirova G, Stocco A, et al (2020) Testing machine learning based systems: A systematic mapping. *Empirical Software Engineering* 25(6):5193–5254
- RISE Research Institutes of Sweden (2022) SMIRK GitHub repository. URL <https://github.com/RI-SE/smirk/>
- Salay R, Queiroz R, Czarnecki K (2018) An analysis of ISO 26262: Machine learning and safety in automotive software
- Schwalbe G, Schels M (2020) A survey on methods for the safety assurance of machine learning based systems. In: *10th European Congress on Embedded Real Time Software and Systems (ERTS 2020)*
- Schwalbe G, Knie B, Sämann T, et al (2020) Structuring the safety argumentation for deep neural network based perception in automotive applications.



- In: International Conference on Computer Safety, Reliability, and Security, Springer, pp 383–394
- Schyllander J (2014) Fotgångarolyckor - statistik och analys. Tech. Rep. MSB744, Swedish Civil Contingencies Agency, URL <https://rib.msb.se/filer/pdf/27438.pdf>
- Serban A, van der Blom K, Hoos H, et al (2020) Adoption and effects of software engineering best practices in machine learning. In: Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM), pp 1–12
- Song Q, Borg M, Engström E, et al (2022) Exploring ML testing in practice: Lessons learned from an interactive rapid review with axis communications. In: Proceedings of the 1st International Conference on AI Engineering - Software Engineering for AI
- Tambon F, Laberge G, An L, et al (2022) How to certify machine learning based safety-critical systems? A systematic literature review. *Automated Software Engineering* 29(38)
- Tao J, Li Y, Wotawa F, et al (2019) On the industrial application of combinatorial testing for autonomous driving functions. In: 2019 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW), IEEE, pp 234–240
- Thorn E, Kimmel SC, Chaka M, et al (2018) A framework for automated driving system testable cases and scenarios. Tech. rep., Department of Transportation. National Highway Traffic Safety Administration. United States.
- Tsilionis K, Wautelet Y, Faut C, et al (2021) Unifying behavior driven development templates. In: 2021 IEEE 29th International Requirements Engineering Conference (RE), IEEE, pp 454–455
- de la Vara JL, Ruiz A, Gallina B, et al (2019) The AMASS approach for assurance and certification of critical systems. In: *Embedded World 2019*
- Vogelsang A, Borg M (2019) Requirements engineering for machine learning: Perspectives from data scientists. In: 2019 IEEE 27th International Requirements Engineering Conference Workshops (REW), IEEE, pp 245–251
- Weissensteiner P, Stettinger G, Rumetshofer J, et al (2021) Virtual validation of an automated lane-keeping system with an extended operational design domain. *Electronics* 11(1):72

- Wiegers K (2008) Karl Wiegers' software requirements specification (SRS) template. Tech. rep., Process Impact, URL <https://www.modernanalyst.com/Resources/Templates/tabid/146/ID/497/Karl-Wiegers-Software-Requirements-Specification-SRS-Template.aspx>
- Willers O, Sudholt S, Raafatnia S, et al (2020) Safety concerns and mitigation approaches regarding the use of deep learning in safety-critical perception tasks. In: International Conference on Computer Safety, Reliability, and Security, Springer, pp 336–350
- Wozniak E, Cârlan C, Acar-Celik E, et al (2020) A safety case pattern for systems with machine learning components. In: International Conference on Computer Safety, Reliability, and Security, Springer, pp 370–382
- Wu B, Nevatia R (2008) Optimizing discrimination-efficiency tradeoff in integrating heterogeneous local features for object detection. In: 2008 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, pp 1–8
- Wu W, Kelly T (2004) Safety tactics for software architecture design. In: Proceedings of the 28th Annual International Computer Software and Applications Conference, 2004. COMPSAC 2004., IEEE, pp 368–375
- Zhang JM, Harman M, Ma L, et al (2022) Machine learning testing: Survey, landscapes and horizons. *IEEE Transactions on Software Engineering* 48(1):1–36
- Zimek A, Schubert E, Kriegel HP (2012) A survey on unsupervised outlier detection in high-dimensional numerical data. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 5(5):363–387