



STEAM & MoSAFE: SOTIF Error-and-Failure Model & Analysis for AI-Enabled Driving Automation

Krzysztof Czarnecki University of Waterloo

Hiroshi Kuwajima DENSO CORPORATION

Citation: Czarnecki, K. and Kuwajima, H., "STEAM & MoSAFE: SOTIF Error-and-Failure Model & Analysis for AI-Enabled Driving Automation," *SAE Int. J. Advances & Curr. Prac. in Mobility* 7(1):168-193, 2025, doi:10.4271/2024-01-2643.

This article was presented at the WCX SAE World Congress Experience, April 16-18, 2024.

Received: 08 Nov 2023

Revised: 08 Jan 2024

Accepted: 20 Jan 2024

Abstract

Driving Automation Systems (DAS) are subject to complex road environments and vehicle behaviors and increasingly rely on sophisticated sensors and Artificial Intelligence (AI). These properties give rise to unique safety faults stemming from specification insufficiencies and technological performance limitations, where sensors and AI introduce errors that vary in magnitude and temporal patterns, posing potential safety risks. The Safety of the Intended Functionality (SOTIF) standard emerges as a promising framework for addressing these concerns, focusing on scenario-based analysis to identify hazardous behaviors and their causes. Although the current standard provides a basic cause-and-effect model and high-level process guidance, it lacks concepts required to identify and evaluate hazardous errors, especially within the context of AI.

This paper introduces two key contributions to bridge this gap. First, it defines the SOTIF Temporal Error and Failure Model (STEAM) as a refinement of the SOTIF cause-and-effect model, offering a comprehensive system-design perspective. STEAM refines error definitions, introduces error sequences, and classifies them as error sequence patterns, providing particular relevance to systems employing advanced sensors and AI. Second, this paper proposes the Model-based SOTIF Analysis of Failures and Errors (MoSAFE) method, which allows instantiating STEAM based on system-design models by deriving hazardous error sequence patterns at module level from hazardous behaviors at vehicle level via weakest precondition reasoning. Finally, the paper presents a case study centered on an automated speed-control feature, illustrating the practical applicability of the refined model and the MoSAFE method in addressing complex safety challenges in DAS.

Introduction

With the rapid proliferation of DAS in vehicles, such as ADASs (Advanced Driver Assistance Systems) and ADSs (Automated Driving Systems), assuring their safety becomes paramount. This paper addresses the challenges of safety assurance of DAS that are subject to complex road environments and vehicle behaviors, exacerbated by the growing reliance on sophisticated sensors and AI. Such challenges give rise to unique safety faults stemming from specification insufficiencies and technological performance limitations, where sensors and AI introduce errors that vary in magnitude and temporal patterns, posing potential safety risks. For example, an AI-based object detector may experience a False Negative (FN) detection in a sensor frame. While a singular FN may not cause any safety risk, this error repeating multiple times while approaching an obstacle may cause a collision. The SOTIF standard ISO 21448:2022[1] emerges as a promising framework for

addressing these concerns, focusing on scenario-based analysis to identify hazardous behaviors and their causes. Although the current standard provides a basic cause-and-effect model linking specification insufficiencies and technological performance limitations to hazardous behaviors, it lacks concepts required to adequately specify hazardous errors, especially within the context of AI, such as the sequences of FNs that may lead to a crash. Further, SOTIF suggests the analysis of system architecture to identify potential functional insufficiencies, but it does not provide any concrete method to do so. It also does not give detailed guidance on severity evaluation of the identified hazards.

This paper introduces two key contributions to bridge this gap. First, it defines the SOTIF Temporal Error and Failure Model (STEAM), offering a comprehensive system-design perspective. STEAM refines error definitions, introduces error sequences, and classifies them as error sequence patterns, providing particular relevance to systems employing advanced

sensors and AI. Furthermore, it categorizes scenario conditions based on their role in the causal chain, enabling a gradual refinement of scenario and system behavior models for safety analysis. Second, this paper proposes the Model-based SOTIF Analysis of Failures and Errors (MoSAFE) method, building upon the STEAM. MoSAFE leverages system design models, scenarios, and harmful events to derive scenario-specific causal error and failure models. This is achieved by adapting Fault Trees (FTs) to incorporate error sequence patterns as nodes. The approach enables the derivation of hazardous error sequence patterns from hazardous behaviors at vehicle level using weakest precondition analysis, allowing for probabilistic analysis of error occurrence, particularly in the context of sensors and AI. To enhance tractability, conservative approximation techniques are employed. Finally, the paper presents a case study centered on an automated braking feature, illustrating the practical applicability of the refined model and the MoSAFE method in addressing complex safety challenges in DAS.

In summary, this paper makes the following contributions:

1. STEAM refines the SOTIF cause-and-effect model by adding the concept of (i) hazardous error sequences to recognize the spatio-temporal nature of hazardous errors; and (ii) Hazardous Behavior Patterns (HBPs) at vehicle level and (iii) Hazardous Error Patterns (HEPs) at element level as a means to specify classes of hazardous behaviors and hazardous error sequences, respectively. It also categorizes scenario conditions based on their role in linking hazardous behaviors to harm and their effects on DAS inputs.
2. MoSAFE is a model-based method to identify HBPs and HEPs and evaluate their severity and likelihood as part of SOTIF analyses in Clause 6 and Clause 7 of ISO 21448:2022 [1]. MoSAFE relies on building scenario-specific models of the DAS and its road-and-vehicle environment, which are instrumented to inject deviations from the intended behavior. MoSAFE uses these models to identify HBPs and derive HEPs from HBPs as weakest preconditions (or their over-approximations). It additionally captures the causal links among the HBPs and HEPs in a novel form of an FT with temporal patterns as nodes. The FT can also express over-approximations using implication arrows. Finally, MoSAFE allows deriving safety requirements on the performance of AI-based components as upper bounds on HEP occurrence rates, and MoSAFE results can be used as evidence in a safety case.

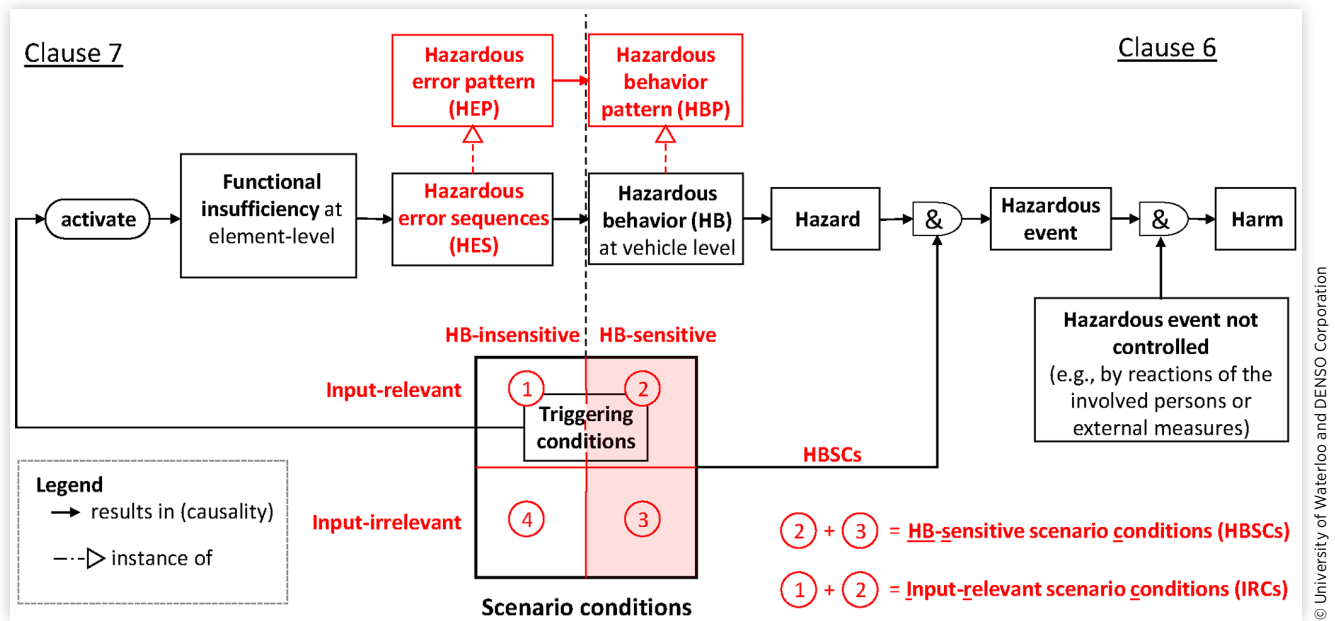
Background: SOTIF Cause-and-Effect Model and Assurance Process

ISO 21448:2022 [1] is an international standard providing guidance on assuring the safety of the intended functionality (SOTIF) of E/E (electrical and/or electronic) systems (including

software), especially emergency intervention systems and DAS at SAE levels 1 to 5 [2]. SOTIF is defined as the absence of unreasonable risk due to a hazard caused by *functional insufficiencies*, which are (i) *insufficiencies of specification* of the intended functionality or (ii) *performance insufficiencies*, both at the vehicle level or the level of the E/E elements implementing the DAS. SOTIF complements ISO 26262:2018 [3], which focuses on functional safety assurance (FuSA), that is, assuring the absence of unreasonable risk due to a hazard caused by deviating from the specified behavior. While assuring the absence of unreasonable risk due to a hazard caused by functional insufficiencies of components is the subject of both standards, functional insufficiencies of AI components are currently insufficiently covered by ISO 26262:2018 [4], but are explicitly in scope of SOTIF.

Fig. 1 summarizes the cause-and-effect model that underlies SOTIF at element level (ignore the red elements for now). SOTIF focuses on hazards that result from functional insufficiencies of the DAS, which consist of insufficiencies of specification and performance insufficiencies at the vehicle and element level, where “element” refers to one or more hardware parts and software units of the DAS. Fig. 1 explicitly shows functional insufficiencies at the element level, but functional insufficiencies at the vehicle level are also considered in our analysis, as will be explained. An example of an insufficiency of specification at the vehicle level is an incorrectly specified vehicle behavior to be implemented by the DAS, such as an inadequate braking level in a given scenario. An example of an insufficiency of specification at the element level is the detection range of the object detector that is inadequately selected for the target ODD (Operational Design Domain). An example of a performance insufficiency at the element level is an insufficient obstacle detection rate by the object detector. These insufficiencies can cause Hazardous Behavior (HB) of the Subject Vehicle (SV), such as unintended lack of braking, which may lead to *hazards*, defined as potential sources of harm, such as the potential of colliding with an object. The realization of the hazard and its potential severity depend on the operational scenario in which the HB transpires. In particular, the scenario may contain *scenario conditions* in which the HB can lead to harm, such as the presence of an obstacle blocking the lane ahead of the SV. The occurrence of an HB under such conditions is referred to as a *hazardous event*. An example is the unintended lack of braking when approaching a stopped vehicle such that the lack of braking can cause a collision. The harm from a hazardous event may be avoided by proper reactions of the involved persons, including the SV driver (for a driver assistance system) or the drivers of the other involved vehicles. For example, the collision may be avoided by the front vehicle accelerating or changing lane. The functional insufficiencies that lead to an HB do so when activated by specific scenario conditions, which are referred to as *triggering conditions*. In particular, a functional insufficiency on element level is activated by a triggering condition, which is a combination of scenario conditions that results in a hazardous erroneous output of the element, which then contributes to an HB on vehicle level (see Fig. 1). An example of an output

FIGURE 1 SOTIF cause-and-effect model (based on Figs. 3B and 4 in [1]), including new elements of SOTIF Temporal Error and Failure Model (STEAM) in red



© University of Waterloo and DENSO Corporation

error from an object detector is an FN. The triggering condition for an FN would include the object that is missed, but also other scenario conditions that contribute to the object not being detected, such as adverse weather conditions or an unusual appearance of the object. Another SOTIF concept is as a *reasonably foreseeable misuse*, which is a usage of the DAS in a way that was not intended by the manufacturer. It could itself be a triggering condition leading to an HB, or it could contribute to reduced controllability of an HB. Reasonably foreseeable misuse is outside the scope of this paper.

The SOTIF standard defines a multi-activity assurance process to identify and eliminate hazards or reduce risks related to SOTIF, spanning multiple document clauses. The process starts with the system specification and design, which among others specify the ODD, use cases, the driving policy and the system design (Clause 5). The next activity is the identification and evaluation of SOTIF hazards (Clause 6), which has three objectives: (1) identification of hazards; (2) evaluation of severity, exposure, and controllability, and (3) specification of acceptance criteria. The latter are used to determine whether the risk estimated in subsequent activities is reasonable. If this activity establishes that the severity and controllability of the identified hazard is above the lowest classes, respectively, S0 and C0, the potential SOTIF causes of the hazard need to be analyzed (Clause 7), otherwise the risk is deemed reasonable for the system to be deployed. The identification and evaluation of functional insufficiencies and triggering conditions (Clause 7) has two objectives: (1) identification of functional insufficiencies (i.e., insufficiencies of specification and performance insufficiencies) and triggering conditions; and (2) evaluation of system response to the identified functional insufficiencies and triggering conditions. The latter objective requires

estimating the likelihood of the hazards resulting from the identified functional insufficiencies and triggering conditions, so that their risk can be evaluated against the acceptance criteria from Clause 6. If the risk is deemed as reasonable, the system is subject to verification and validation (defined in Clause 9), which covers the known unsafe scenarios (Clause 10) and performs an exploration to uncover unknown unsafe scenarios (Clause 11). If at the end of any of these steps the risk does not meet the acceptance criteria, the system is modified to reduce the risk (Clause 8); otherwise, the residual risk is evaluated (Clause 12), and the system can be deployed. Additional SOTIF activities occur during operation in order to uncover any additional potential SOTIF issues (Clause 13), which then are cycled back into the SOTIF process.

The focus of this paper is on Clauses 6 and 7. Clause 6 focuses on identifying and analyzing the effects of HBs, which can be seen as hazardous vehicle-level failures, without the regard for their causes (the part of the cause-and-effect model on the right of the vertical dashed line in Fig. 1). Clause 7 focuses on identifying and analyzing the causes of HBs, especially the functional insufficiencies and the errors they cause at the element level (the part on the left of the vertical dashed line in Fig. 1).

SOTIF Temporal Error and Failure Model (STEAM)

This section introduces our proposed SOTIF Temporal Error and Failure Model (STEAM) as a refinement of the original SOTIF cause-and-effect model (see Fig. 1). STEAM makes three refinements.

On the vehicle-level (Clause 6), STEAM introduces the concept of a *hazardous behavior pattern* (HBP), which is a specification of the magnitude and temporal occurrence of the deviations of the SV behavior from its intended behavior that is hazardous under the given scenario conditions. Small deviations from the intended behavior may not be hazardous. For example, slightly reduced braking level when approaching a stopped vehicle may lead to a slight overshoot of the targeted stopping point, but it would not be hazardous if the SV still stops well before the stopped vehicle. Similarly, a temporary deviation from the intended braking level, even if large, may still not be hazardous if it can be compensated by subsequently harder braking.

On the element-level (Clause 7), STEAM introduces the concept of a *hazardous error sequence* (Hazardous Error Sequence (HES)), which is a temporal sequence of errors that is caused by a functional insufficiency of an element, and it causes an HB. An error sequence is a temporal refinement of the concept of error from ISO 26262:2018, defined as a “discrepancy between a computed, observed, or measured value or condition, and the true, specified or theoretically correct value or condition.” [3] The notion of the true or correct value depends on the DAS function. Typical functions in a DAS include perception, prediction, planning, and control [5]. For perception functions, error is defined wrt. ground truth. There is neither ground truth for prediction nor planning functions. Prediction should adequately reflect the distribution of possible futures. An example of a prediction error would be to miss a likely future, such as to exclude the possibility of another vehicle tuning when the turn is plausible. For planning, the planned actions can be assessed for their quality or level of deviation from a desired policy. For all three types of functions, a momentary error is not necessarily hazardous; it is typically a certain pattern of errors over time that becomes hazardous. For example, a single, momentary FN may cause a momentary lack of braking, but such lack of braking may be compensated by subsequent harder braking. However, a persistent FN or lack of braking for an obstacle ahead may become impossible to compensate and will cause a collision if not controlled for. A *hazardous error pattern* (HEP) is a specification of the pattern of errors in terms of their time and magnitude to cause an HB. As such, an HEP represents a set of concrete HESs (see Fig. 1), which may lead to a class of HBs; and the latter may itself be specified by an HBP. For example, an HEP for FNs could be specified as the total duration of missing detection of an obstacle during an approach scenario. This HEP would include error sequences where the missing detection is in one continuous period or spans multiple periods during the scenario, as long as the total duration meets the HEP specification. Whereas Fig. 1 shows a single HES (and a single HBP) causing an HB, there will normally be a chain of HESs through the system before they trigger an HB. For example, a perception HES may cause a prediction HES, and the latter may cause a planning HES leading to an HB. Thus, STEAM needs to be instantiated as error-and-failure causal chains across the DAS design, where HBs

are considered as hazardous DAS *failures*, for specific combinations of DAS design, hazards, and scenarios, as part of assurance.

As a third refinement, STEAM categorizes scenario conditions according to their role in the causal chain (see the four quadrants in Fig. 1). First, scenario conditions are categorized whether they influence the translation of an HB into harm or whether they do not (see the horizontal dimension of the four quadrants in Fig. 1): *HB-sensitive* scenario conditions (HBSC) are those in which the HB leads to harm; conversely, *HB-insensitive* scenario conditions are those that do not influence the translation of HB into harm. An example of an HBSC for the HB “lack of braking” is the presence of an obstacle ahead of the SV; however, the color of the obstacle is an HB-insensitive scenario condition. Second, scenario conditions are classified whether they affect the input into the DAS or not (see the vertical dimension): *input-relevant* scenario conditions (IRC) are those affecting the DAS input, including sensor inputs, vehicle-to-vehicle and vehicle-to-infrastructure messages, and pre-recorded maps; otherwise, they are *input-irrelevant*. For example, assuming a DAS with a camera sensor, both the presence and the color of an obstacle are *input-relevant*, as they influence the image produced by the camera. Thus, the presence of the obstacle falls into quadrant 2, and its color falls into quadrant 1. Road friction is an example of an HBSC that is not affecting the camera and thus is input-irrelevant and falls into quadrant 3. Triggering conditions are necessarily input-relevant, but not all input-relevant scenario conditions are triggering conditions. For example, the color of an obstacle may or may not be responsible for a particular FN.

The classification of scenario conditions enables a gradual refinement of scenario behavior models for safety analysis in Clause 6 and 7. The HBSCs (quadrants 2 and 3 in Fig. 1) and driving policy are relevant in the high-level scenario modeling that targets the severity assessment of hazardous behavior (Clause 6). IRCs (quadrants 1 and 2) are relevant—in addition to quadrant 3—in the detailed modeling and analysis of scenarios, including the DAS design, to uncover the causal chains through the system that trigger HBs (Clause 7). Quadrant 2 impacts both the DAS input and the translation of HBs into harm. Quadrant 4 can be ignored during modeling and analysis.

Model-Based SOTIF Analysis of Failures and Errors (MoSAFE)

We now describe our proposed Model-based SOTIF Analysis of Failures and Errors (MoSAFE) method, which allows instantiating STEAM based on scenarios; DAS, SV, and road environment models; and hazards. The method is divided into two activities, matching Clause 6 and 7, respectively. The first activity focuses on the identification and evaluation of Hazardous Behavior Patterns (HBPs), which corresponds to the right side of the STEAM in Fig. 1.

The second activity focuses on the identification and evaluation of Hazardous Error Patterns (HEPs), which corresponds to the left side of the STEAM in Fig. 1. Both activities rely on modeling the DAS and its environment, targeting a level of abstraction that is appropriate for the given activity’s objective. The first activity’s objective is to identify the HEPs and evaluate their severity. For this purpose, it uses a high-level model of DAS, being the intended driving policy, and an environment model capturing the HBSCs. The second activity’s objective is to identify the HEPs and evaluate their likelihood. Therefore, this activity uses a more detailed design model of the DAS and an environment model refined with the IRCs. Together, the two activities provide a risk evaluation of the identified hazards, consisting of their severity and likelihood.

Identification and Evaluation of Hazards (Clause 6)

Clause 6 focuses on the identification and evaluation of SOTIF hazards (the right side of Fig. 1), and the main idea of its refinement as part of MoSAFE is to specify them as HBPs and evaluate their severity using high-level behavior models of the DAS and its environment (Fig. 2).

Hazard Identification

SOTIF hazards are potential sources of harm caused by the HBs at the vehicle level [1], and thus their identification involves the identification of the HBs and the HBSCs (see the right side of Fig. 1). The identification of HBs involves both analyzing the safety of the specified behavior on vehicle level and the safety of deviations from the specified behavior. The deviations can be identified using Hazard and Operability Analysis (HAZOP) [6], by instantiating guide words such as “no”, “more”, and “less”. The HBSCs are identified by systematically eliciting the

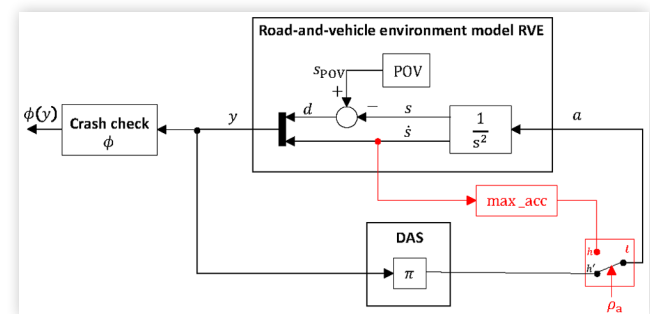
operational scenarios relevant to the ODD of the DAS (see [7] for additional guidance).

Table 1 lists two sample intended behaviors and the corresponding HBs and hazards for an automated speed control feature. The feature could be part of an ADAS, such as a full-speed range adaptive cruise control, or it could represent the longitudinal behavior aspect of an ADS. The sample intended “behaviors are braking for a stationary vehicle ahead” and “maintaining a safe distance when following a vehicle.” The second column lists HBs, which are hazardous deviations from the intended behavior, including Unintended Braking Interruption (UBI) (i.e., no braking when braking needed), Unintended Insufficient Braking (UIB) (i.e., less braking than needed), Unintended Hard Braking (i.e., more braking than needed), and Unintended Acceleration (UA) (i.e., more acceleration than needed). The third column lists the hazards resulting from the HBs: rear-ending another vehicle or being rear-ended by another vehicle. The hazards indicate the key HBSCs in which the given HB may lead to harm, such as the presence of a stationary or braking vehicle ahead. These initial HBSCs are then refined to cover the full range of operational scenarios and conditions within the ODD of the DAS, such as the full ranges of speed, road friction, road grade, and road curvature occurring within the ODD.

Hazard Evaluation and Acceptance Criteria Specification

Hazard identification is followed by the evaluation of the severity, exposure, and controllability of the identified hazards. As recommended by the SOTIF standard, this step leverages concepts and methods from Part 3 of ISO

FIGURE 2 High-level scenario model of the Road-and-Vehicle Environment (RVE), capturing the Hazardous-Behavior-Sensitive Scenario Condition (HBSC), and the DAS driving policy for “braking for a stationary vehicle ahead”, with the HB injection logic for Unintended Braking Interruption (UBI) marked in red



© University of Waterloo and DENSO Corporation

TABLE 1 Sample intended behaviors, HBs, and hazards related to longitudinal behavior of a DAS

Intended longitudinal behaviour	Longitudinal HB	Hazard
Braking for a stationary vehicle ahead	Unintended braking interruption (UBI)	Rear-end collision with the stationary vehicle
	Unintended insufficient braking (UIB)	Rear-end collision with the stationary vehicle
	Unintended hard braking (UHB)	Another vehicle hitting from behind
Maintaining a safe distance when following a vehicle	Unintended braking interruption (UBI)	Rear-end collision with the front vehicle when it brakes
	Unintended acceleration (UA)	Rear-end collision with the front vehicle
	Unintended hard braking (UHB)	Another vehicle hitting from behind

© University of Waterloo and DENSO Corporation

26262:2018 [3], including the classification of severity (S0-S3), exposure (E0-E4), and controllability (C0-C3). For collisions, severity can be estimated using the Delta-V method [8], which maps the collision configuration and the change in velocity resulting from the collision to a range of injury severity based on crash statistics. Exposure is the estimated likelihood of the crash-relevant HBSCs during the operation of the DAS. Finally, controllability depends on the type of DAS and its level of automation. Whereas ADAS rely on the driver to intervene, an ADS-equipped vehicle may be driverless, and the ability of other road users to control the hazardous event is often limited.

Acceptance criteria are qualitative or quantitative criteria representing the absence of an unreasonable level of risk. An example of quantitative acceptance criteria would be an upper bound on the crash rate for each severity class, possibly expressed as the mean distance travelled between crashes. Acceptance criteria may be allocated to different combinations of HBs and HBSCs. For example, for an HBSC set and its exposure $P(\text{HBSC})$, acceptance criteria would put an upper bound on $P(\text{HB} | \text{HBSC})$, the conditional probability of HB under HBSC, and the probability of the hazardous scenario would be given as follows:

$$P(\text{HB}, \text{HBSC}) = P(\text{HB} | \text{HBSC})P(\text{HBSC}) \quad (1)$$

Model-Based Severity Evaluation

The MoSAFE method leverages high-level behavior models to evaluate the severity of the different degrees of HB under the different HBSCs. Although Delta-V allows evaluating the severity of harm resulting from a collision, hazard evaluation also requires mapping different degrees of HB under different HBSCs to crash severity. For example, the duration and timing of a braking interruption or the level of insufficient braking when approaching a stationary or braking vehicle will influence the occurrence of a collision and its severity. As part of severity evaluation, MoSAFE uses a High-Level Scenario Model (HLSM), consisting of (i) the driving policy of the DAS as its specified intended behavior and (ii) a Road-and-Vehicle Environment (RVE) model that captures the HBSCs. HLSMs are specific to the intended behavior and the HBSCs being evaluated. In other words, “braking for a stationary vehicle ahead” and “maintaining a safe distance when following a vehicle” would each lead to a different HLSM. Given an HLSM, the safety of the specified behavior of the DAS under the given HBSCs is first evaluated. This is followed by the evaluation of HBs, which are injected into the HLSM. Varying levels of HBs are then linked to crashes of different severity.


We illustrate the key ideas of severity evaluation in MoSAFE for the intended behavior of “braking for a stationary vehicle ahead”, which we refer to as the Principle

Other Vehicle (POV), and the HB of unintended braking interruption (UBI) (see Fig. 3). The intended behavior for the SV is to brake at a comfortable level $a_{b,\min}$ to stop at a required standstill distance Δs_{stand} behind POV. An HB such as a UBI would cause the SV to approach the POV too fast, so that a higher level of braking would be necessary to stop at Δs_{stand} behind the POV. If the required braking to stop behind the POV without colliding with it exceeds the maximum braking capacity a_{\max} of the SV, the UBI becomes hazardous and will lead to a rear-end collision at a certain collision velocity v_{impact} , which then can be mapped to collision severity using the Delta-V method.

Establishing high-level scenario models to determine v_{impact} . The first step in the severity evaluation is to establish the HLSM of the DAS and its environment for the given intended behavior (see Fig. 2). The environment of the DAS is represented by the RVE model, which consists of the relevant road environment elements and the SV dynamics. It takes the control input from the DAS and provides the DAS with observations y . The RVE model captures the HBSCs for our example, including the presence of the stationary POV, modeled by its position s_{POV} ; and the SV kinematics, modeled by a double integrator $\frac{1}{s^2}$, with the SV’s acceleration a as control input and its speed \dot{s} and position s as output. The initial position of SV is $s_{\text{init}} = 0$, and its initial speed is $v_{\text{init}} = v_{\max}$, to allow for a full braking scenario from v_{\max} of the DAS. The POV position is such that the SV needs to brake with constant $a_{b,\min}$ to stop at $s_{\text{stop}} = s_{\text{POV}} - \Delta s_{\text{stand}}$ (see Fig. 4a); thus, we have:

$$s_{\text{POV}} = s_{\text{stop}} + \Delta s_{\text{stand}} \quad s_{\text{stop}} = \frac{v_{\text{init}}^2}{2a_{b,\min}} \quad (2)$$

FIGURE 3 The SV (left) experiences an unintended braking interruption (UBI) when braking for a POV (right) stopped ahead. The UBI transforms a safe situation (top) into an unsafe one (bottom).



The output y of the RVE model is the SV's speed \dot{s} and its distance $d = s_{POV} - s$ to the POV. The following time-continuous linear ODE gives the state-space representation of the model, where x is the system state and $x^{[0]}$ is the initial state at $t = 0$:

$$\begin{aligned} x &= \begin{bmatrix} s \\ \dot{s} \end{bmatrix} & y &= \begin{bmatrix} d \\ \dot{s} \end{bmatrix} \\ \dot{x} &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} a \\ y &= \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} x + \begin{bmatrix} s_{POV} \\ 0 \end{bmatrix} \\ x^{[0]} &= \begin{bmatrix} 0 \\ v_{init} \end{bmatrix} \end{aligned} \quad (3)$$

The intended behavior of the SV for our example is represented by a driving policy π that applies the required braking $a_{b,req}$ to stop at Δs_{stand} behind the POV, when $a_{b,req}$ is between $a_{b,min}$ and $a_{b,max}$, where

$$a_{b,req} = \frac{\dot{s}^2}{2(d - \Delta s_{stand})} \quad (4)$$

When $a_{b,req}$ is less than $a_{b,min}$, the SV is free to accelerate to v_{max} and continue at that speed. If $a_{b,req}$ reaches or exceeds $a_{b,max}$, the SV will apply $a_{b,max}$. Also, the SV will apply $a_{b,max}$ whenever it moves closer to the POV than Δs_{stand} . These cases are captured by the following policy function:

$$a = \pi(d, \dot{s}) = \begin{cases} 0, & \text{if } d > \Delta s_{stand} \wedge a_{b,req} < a_{b,min} \wedge \dot{s} \geq v_{max} \\ a_{max}, & \text{if } d > \Delta s_{stand} \wedge a_{b,req} < a_{b,min} \wedge \dot{s} < v_{max} \\ -a_{b,req}, & \text{if } d > \Delta s_{stand} \wedge a_{b,min} \leq a_{b,req} < a_{b,max} \\ -a_{b,max}, & \text{if } d > \Delta s_{stand} \wedge a_{b,req} > a_{b,max} \\ 0, & \text{if } d \leq \Delta s_{stand} \wedge \dot{s} = 0 \\ -a_{b,max}, & \text{if } d \leq \Delta s_{stand} \wedge \dot{s} > 0 \end{cases} \quad (5)$$

Note that this policy has discrete transitions in acceleration to simplify the analysis, but these transitions would lead to jerky driving. An actually implemented policy would have smooth transitions, but its braking level would need to be close to $a_{b,req}$ on average, and therefore the discrete policy is adequate for evaluating the safety of applying $a_{b,req}$ as an intended target and the safety of deviating from it. The adequacy of analyzing a smooth policy using a discrete approximation is confirmed by the model validation results using simulation testing on p. 14.

Evaluation of the intended behavior. The intended behavior represented by this policy under the modeled HBSCs is safe, as the vehicle is guaranteed to stop at x_{stop} , i.e., Δs_{stand} behind the POV (Fig. 4a). The combination of the RVE model (3) and the policy (5) (i.e., the HLSM in Fig. 2 where the red part is ignored) results in the following non-linear ODE during braking:

$$\ddot{s} = \frac{\dot{s}^2}{2(s_{stop} - s)} \quad (6)$$

Assuming the initial condition $s^{[0]} = 0$ and $\dot{s}^{[0]} = v_{init}$, this ODE can be shown to have a solution that corresponds to the application of constant acceleration $a_{b,min}$, resulting in the speed profile in Fig. 4a.

Figure 2 also includes a crash check $\phi(y)$, which has two components: $\phi_1(y)$ checks whether the SV has collided with the POV, and $\phi_2(y)$ keeps track of v_{impact} :

$$\begin{aligned} \phi_1(d, \dot{s}) &= I_{\exists t \in [0..T]: d^{[t]} \leq 0} \\ v_{impact} &= \phi_2(d, \dot{s}) = \dot{s}^{[t_c]}, t_c = \min t \in [0..T], d^{[t]} \leq 0 \end{aligned} \quad (7)$$

where I is the indicator function, and T is duration of the braking scenario. Note that safety properties such as $\phi_1(y)$ could be expressed using a temporal logic, such as Signal Temporal Logic (STL) [9], but we choose to use first-order logic for simplicity.

Evaluation of the HBs. Having evaluated the intended behavior as safe, we turn to evaluating the severity of the HBs, which are hazardous deviations for the intended behavior. This is achieved by injecting the HBs into the intended behavior of the HLSM. For UBI, the injection is accomplished by adding the switch ι in Fig. 2, which interrupts braking by injecting a_{max} when the SV's speed is below v_{max} or zero acceleration otherwise:

$$\max_acc(\dot{s}) = \begin{cases} a_{max}, & \text{if } \dot{s} < v_{max} \\ 0, & \text{if } \dot{s} \geq v_{max} \end{cases} \quad (8)$$

The switch ι operates in discrete time and is controlled by the sequence $\rho_a \subseteq \mathbb{N}$, which contains the time steps for which the switch should be in on-position, i.e., connecting to h rather than h' , and injecting a braking interruption. The current time step k is computed as the integer part, represented by the floor operator, of the current continuous time t divided by the time-step duration Δt . Formally, the switch is defined by the following function:

$$\iota^{[t]}(h, h', \rho_a) = \begin{cases} h, & \text{if } k(t) \in \rho_a \quad \rho_a \subseteq \mathbb{N} \\ h', & \text{if } k(t) \notin \rho_a \quad k(t) = \frac{t}{\Delta t} \end{cases} \quad (9)$$

The switch allows injecting an arbitrary sequence of braking interruptions, up to the time resolution Δt , which can be set as finely as needed. Figure 4b shows an example speed profile resulting from injecting $\rho_a = \{26.45, 66.87\}$ with $\Delta t = 0.1$ s. This sequence injects two braking interruption intervals, the first one with the duration $\tau_1 = 1.9$ s and the second one with $\tau_2 = 2.1$ s. The first interruption changes the approach situation of the SV, requiring it to apply $a_{b,req} = 2$ m/s², rather than the initial $a_{b,min} = 1$ m/s². The second interruption puts the SV on a maximum braking trajectory with $a_{b,max} = 8$ m/s² to crash into the POV with $v_{impact} = 6$ m/s.

Evaluating the severity of a multi-interval braking interruption is complex, as the number of the intervals, their start times, and their duration influence the resulting v_{impact} . This complexity likely arises for many other HBs, where a multitude of specific HB sequences need to be mapped to the resulting v_{impact} .

A general solution strategy in MoSAFE is to employ a conservative over-approximation, where a class of sequences $\mathbb{P} \subseteq 2^{\mathbb{N}}$, specified by an abstract pattern, is bounded by the maximum v_{impact} that any of the sequences in \mathbb{P} can result in. For UBI, we define the pattern $\mathbb{P}_{a, k_{\min}, k_{\max}}$ to denote all sequences ρ_a with the total duration of braking interruption being at least k_{\min} and at most k_{\max} time steps. The lower bound helps eliminate UBI sequences that are guaranteed to be safe. Without any loss of generality, these sequences are also limited by the

maximum duration T_{max} of an approach scenario. This maximum occurs in the nominal case of braking with $a_{b, \min}$ (Fig. 4a), since any UBI would lead to higher approach velocities and thus shorter approach time:

$$T_{\text{max}} = \frac{v_{\text{init}}}{a_{\text{min}}} \quad (10)$$

We also define the corresponding maximum duration n_{max} in discrete steps (the ceiling operator ensures that n_{max} covers T_{max} completely):

$$n_{\text{max}} = \frac{T_{\text{max}}}{\Delta t} \quad (11)$$

The UBI pattern $\mathbb{P}_{a, \tau_{\min}, \tau_{\max}}$ is formally specified as follows:

$$\mathbb{P}_{a, k_{\min}, k_{\max}} = \left\{ \rho_a \subseteq \mathbb{N}_{n_{\text{max}}} \mid k_{\min} \leq |\rho_a| \leq k_{\max} \right\} \quad (12)$$

where $\mathbb{N}_n = \{k \in \mathbb{N} \mid k \leq n\}$ and $|\rho_a|$ denotes the size of ρ_a and corresponds to the number of discrete time-steps where UBI is injected.

Given an impact velocity v_{impact} , there exists the shortest total duration $\tau_{\min, v_{\text{impact}}}$ of UBI that results in a crash with v_{impact} , and this duration can be computed in closed form. The details of this statement are beyond the scope of this paper, but it relies on a theorem that multiple UBI intervals can always be replaced by a single one that results in a higher v_{impact} than the multiple ones. Thus, $\tau_{\min, v_{\text{impact}}}$ can be computed by minimizing the duration of a single UBI interval for a given v_{impact} . For example, Fig. 4c shows the shortest UBI interval $\tau_{\min, v_{\text{impact}} = 6 \text{ m/s}} = 2.39 \text{ s}$ that results in the same $v_{\text{impact}} = 6 \text{ m/s}$ as the multiple UBI intervals with $\tau_{\text{total}} = \tau_1 + \tau_2 = 4 \text{ s}$ in Fig. 4b.

Mapping collision configuration and v_{impact} to severity classes. The next step is to determine the range of v_{impact} for each of the S0-3 severity classes, which will allow us to determine the maximum severity class for a given total duration τ_{total} of UBI (Tab. 2). The HLSM gives us the crash configuration and v_{impact} due to the injected HB. For UBI, the configuration is a front-to-rear collision, and τ_{total} puts an upper bound on v_{impact} . The crash configuration and v_{impact} allow us to estimate the severity class based on statistical injury models according to the guidance in SAE J2980 [10]. As an example, Tab. 2 gives the severity class (first column) based on injury risk to belted front-row occupants of the SV during a front-to-rear collision at v_{impact}

FIGURE 4 Sample braking speed profiles for $v_{\text{init}} = 15 \text{ m/s}$ and $s_{\text{stop}} = 112.5 \text{ m}$

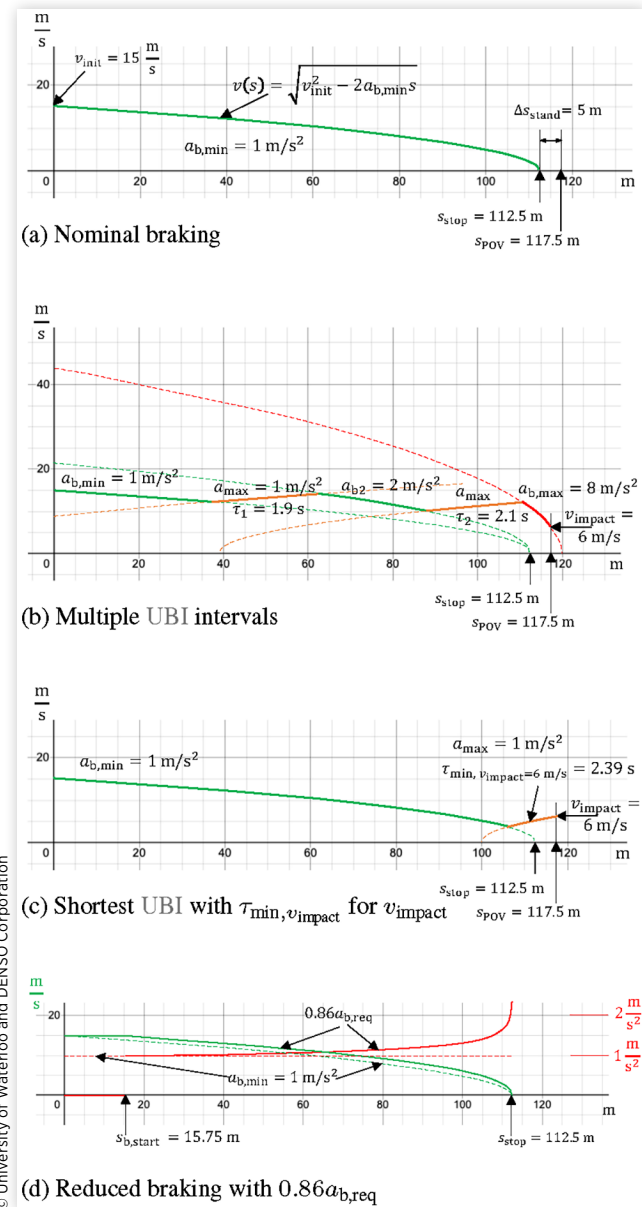


TABLE 2 Range of UBI duration τ_{total} that could result in a crash with a given maximum severity, assuming $a_{b, \min} = 1 \text{ m/s}^2$, $a_{b, \max} = 8 \text{ m/s}^2$, $a_{\text{max}} = 1 \text{ m/s}^2$, $v_{\text{init}} = 15 \text{ m/s}$, $s_{\text{POV}} = 117.5 \text{ m}$, and $\Delta s_{\text{stand}} = 5 \text{ m}$

Severity class	v_{impact} range (m/s)	τ_{total} range (s)
S0	$[v_{\text{min}}, v_{\text{S0}}] = [0.5..3]$	$[\tau_{\text{contact}}, \tau_{\text{S0}}] = [1.97..2.29]$
S1	$(v_{\text{S0}}, v_{\text{S1}}) = (5.3..7.8]$	$(\tau_{\text{S0}}, \tau_{\text{S1}}) = (2.29..2.76]$
S2	$(v_{\text{S1}}, v_{\text{S2}}) = (7.8..10.3]$	$(\tau_{\text{S1}}, \tau_{\text{S2}}) = (2.76..3.57]$
S3	$(v_{\text{S2}}, v_{\text{max}}) = (10.3..15]$	$(\tau_{\text{S2}}, \tau_{\text{max}}) = (3.57..7.83]$

that falls into the ranges specified in the second column. For example, when v_{impact} falls into the range $[v_{\text{min}}..v_{S0}] = [0..5.3]$ m/s the corresponding injury is estimated as severity class S0, where $v_{\text{min}} = 0$ is the minimum v_{impact} (when the SV just touches the POV) and $v_{S0} = 5.3$ m/s is the maximum v_{impact} that still results in S0. [Table 2](#) (in this paper) is based on an existing statistical model [11]; in particular, the severity classes are defined in [Table 2](#) there [11], and the delta V values for each corresponding traffic domain are specified in [Table 3](#) there [11]. Conservatively, delta V, which is the change of velocity of the bullet vehicle due to the collision, is equated with v_{impact} ; in reality, delta V is about half of v_{impact} if SV and POV have the same mass, and delta V approaches v_{impact} when SV strikes a heavy vehicle, such as a truck or a bus [12]. [Table 2](#) is also consistent with another model of severe injury [8], which estimates probability of MAIS (Maximum Abbreviated Injury Scale) 3+ injury as a function of v_{impact} (see [Fig. 4](#) in [8]). For example, S2 corresponds to more than 10% of MAIS 3+ (as defined in Table B.1 in [3], Part 3), and the model [8] predicts this risk at a critical impact speed of about 30 km/h (8.3 m/s).

Given the ranges of v_{impact} for each severity class, we can map them to the corresponding ranges of the shortest total duration $\tau_{\text{min}, v_{\text{impact}}}$ (third column in [Tab. 2](#)). For example, $[v_{\text{min}}..v_{S0}]$ in the first row is mapped to $[\tau_{\text{contact}}.. \tau_{S0}]$, where $\tau_{\text{contact}} = \tau_{\text{min}, v_{\text{impact}} = 0}$ represents the shortest duration of UBI for a crash at $v_{\text{impact}} = 0$, and $\tau_{S0} = \tau_{\text{min}, v_{\text{impact}} = v_{S0}}$ represents the shortest duration of UBI for a crash that still results in S0. As preciously described, the shortest duration $\tau_{\text{min}, v_{\text{impact}}}$ for a given v_{impact} can be computed as an optimization solution in closed form. As another example, the interval $(v_{S2}..v_{\text{max}}]$ of impact velocities that would likely result in S3 is mapped to $(\tau_{S2}.. \tau_{\text{max}}]$, where $\tau_{S2} = \tau_{\text{min}, v_{\text{impact}} = v_{S2}}$ represents the shortest duration of UBI for a crash that may result in a crash with maximum severity of S2, and τ_{max} is the maximum duration of UBI, which occurs when the SV continues at v_{max} without braking, i.e., $\tau_{\text{max}} = S_{\text{POV}}/v_{\text{max}} \approx 7.83$ s. Note that a bracket marks an interval bound that is included in the interval, and a parenthesis marks a bound that is excluded.

Specification of HBPs. The τ_{total} ranges in [Tab. 2](#) allow defining UBI patterns that are bounded by severity (see [Tab. 3](#)). The first pattern in [Tab. 3](#), $\mathbb{P}_{a, \text{nocrash}}$, contains all UBI sequences guaranteeing no crash, i.e., those with τ_{total} shorter than τ_{contact} . Thus, this UBI pattern summarizes all UBI sequences that are non-hazardous. The remaining patterns may lead to crashes and thus are hazardous, i.e., they are HBPs ([Fig. 1](#)). The first HBP (second row in [Tab. 2](#)), $\mathbb{P}_{a, S0..3}$, contains all hazardous UBI sequences, i.e., all UBI sequences that may lead to a crash. The next HBP, $\mathbb{P}_{a, S1..3}$, contains all UBI sequences that may lead to a crash at severity S1 or higher. The remaining two HBPs are defined analogously. Since we consider the probability occurrence of the pattern during the braking scenario and the scenario duration varies depending on the number of UBI time steps, we use the longest duration T_{max} and set k_{max} to n_{max} , instead of $k(\tau_{\text{max}})$.

TABLE 3 UBI patterns by severity crash they can cause; note that $k_{\text{contact}} = k(\tau_{\text{contact}})$, $k_{S0} = k(\tau_{S0})$, $k_{S1} = k(\tau_{S1})$, and $k_{S2} = k(\tau_{S2})$

UBI pattern	Definition	Description
$\mathbb{P}_{a, \text{nocrash}}$	$\mathbb{P}_{a, k_{\text{min}} = 0, k_{\text{max}} = k_{\text{contact}} - 1}$	All UBI sequences guaranteeing no crash
$\mathbb{P}_{a, S0..3}$	$\mathbb{P}_{a, k_{\text{min}} = k_{\text{contact}}, k_{\text{max}} = n_{\text{max}}}$	All UBI sequences that may lead to a crash
$\mathbb{P}_{a, S1..3}$	$\mathbb{P}_{a, k_{\text{min}} = k_{S0} + 1, k_{\text{max}} = n_{\text{max}}}$	All UBI sequences that may lead to a crash at severity S1 or higher
$\mathbb{P}_{a, S2..3}$	$\mathbb{P}_{a, k_{\text{min}} = k_{S1} + 1, k_{\text{max}} = n_{\text{max}}}$	All UBI sequences that may lead to a crash at severity S2 or higher
$\mathbb{P}_{a, S3}$	$\mathbb{P}_{a, k_{\text{min}} = k_{S2} + 1, k_{\text{max}} = n_{\text{max}}}$	All UBI sequences that may lead to a crash at severity S3

© University of Waterloo and DENSO Corporation

Note that each of the hazardous UBI patterns in [Tab. 3](#) contains all of the UBI sequences that lead to a crash within the pattern's severity range, but may also include UBI sequences of a lesser severity (see [Fig. 5](#)). In fact, they are the tightest over-approximations wrt. their specified severity range, given the pattern form in [eq. 12](#). For example, $\mathbb{P}_{a, S2..3}$ contains all UBI sequences that lead to a crash at severity S2 or higher and so does $\mathbb{P}_{a, S1..3}$. Further, while $\mathbb{P}_{a, S1..3}$ also includes all UBI sequences that lead to crashes at severity S1, some of these may also be in $\mathbb{P}_{a, S2..3}$, and both patterns may include UBI sequences of severity S0 or even non-hazardous ones. However, $\mathbb{P}_{a, S2..3}$ is a tighter over-approximation of all UBI sequences that lead to a crash at severity S2 or higher than $\mathbb{P}_{a, S1..3}$. In our sample scenario ([Fig. 4a](#)), assuming $\Delta t = 0.1$ s, injecting $\mathbb{P}_{a, S0..3}$ corresponds to the occurrence of $k_{\text{min}} = 19$ or more time steps with UBI within $n_{\text{max}} = 150$ time steps of the maximum scenario duration. This is because the duration T_{max} of the nominal scenario in [Fig. 4a](#) is 15 s, i.e., $n_{\text{max}} = 150$, and $k_{\text{min}} = k_{\text{contact}} = k(1.97 \text{ s}) = 19$ (from the first row in [Tab. 2](#) and the second row in [Tab. 3](#)). Thus, if UBI occurs in fewer than 19 out of 150 time steps, no collision will occur due to UBI. Similarly, $\mathbb{P}_{a, S3}$ tells us that no collision of severity S3 can occur in our scenario due to UBI if UBI occurs in fewer than $k_{S2} = 35 + 1 = 36$ out of 150 time steps.

Limiting the likelihood of occurrence of the hazardous UBI patterns in [Tab. 3](#) allows limiting the safety risk of crashes due to UBI. These likelihood limits, after being multiplied by exposure, would be assigned based on acceptance criteria.

Identification and Evaluation of Functional Insufficiencies & Triggering Conditions (Clause 7)

Clause 7 focuses on the identification and evaluation of functional insufficiencies and triggering conditions (the left side of [Fig. 1](#)). Functional insufficiencies include (i) insufficiencies of specification and (ii) performance insufficiencies at the vehicle level and element level. For our example,

functional insufficiencies at the vehicle level have already been addressed in the previous section: (i) the specification of the driving policy for our sample intended behavior has been shown to be safe and (ii) the target performance at the vehicle level is given by the acceptance criteria that assign an upper bound on the probability of a collision due to UBI. That probability would be further decomposed into upper bounds on the probability of UBI occurrence in different HBSCs (see [13] for more detail on this decomposition). Thus, this section will focus on the functional insufficiencies at the element level and their triggering conditions.

Model-based Identification and Evaluation of Functional Insufficiencies at the Element Level

The MoSAFE method helps identify and evaluate functional insufficiencies using a Detailed Scenario Model (DSM), which is a refinement of the HLSM from the previous activity. The DSM is used to (i) evaluate the intended behavior of the DAS components in the given scenario and (ii) identify, specify, and evaluate hazardous deviations from the intended behavior at the element level as HEPs.

Establishing detailed scenario models. The first step is to establish the detailed scenario model (DSM) for the intended behavior of the SV (see Fig. 6). To this end, the RVE model is refined to include all Input-Relevant Scenario Conditions (IRCs) (see quadrants 1 and 2 in Fig. 1), not just those IRCs that are also HBSCs (quadrant 2) and thus already included in the HLSM. In our example, the kinematic model from Fig. 2 is augmented with three sample IRCs being the POV's

appearance α , the remaining scenery Sc (such as the road appearance), and fog Fg . These IRCs may be represented at different levels of abstraction, such as using high-level attributes (typically based on an ODD ontology, e.g., [14]) or detailed shape and appearance information. The DAS model is refined based on the system design, and it is denoted as DAS'. In our example, the model includes the original driving policy π (eq. 5) as the intended behavior of the planning components, and four additional perception components: a sensor S , such as a camera or lidar, an object detector D , a tracker T , and an odometry component Odo . For simplicity, we do not further decompose π into planning, control, and actuation elements. All outputs from the refined RVE model RVE' are bundled in y' and input into the sensor S . They are all relevant to the sensor, potentially affecting its output frame F , and are thus IRCs. The output frame F , such as a camera image or lidar scan, is fed into the object detector D , which normally relies on deep neural networks to produce detections. For the purpose of the scenario, the intended behavior of

FIGURE 6 Detailed scenario model for “braking for a stationary vehicle ahead”, including the refined models RVE' and DAS' (refinements in red)

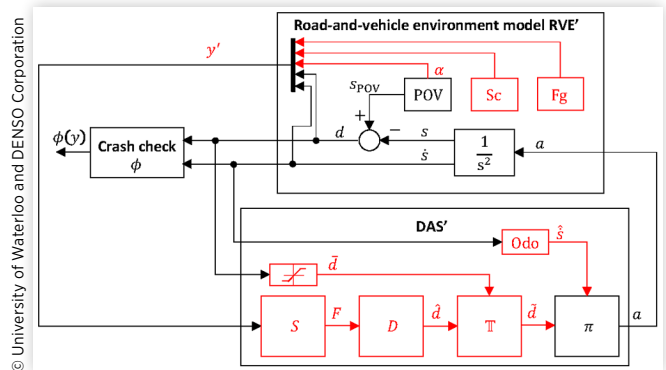
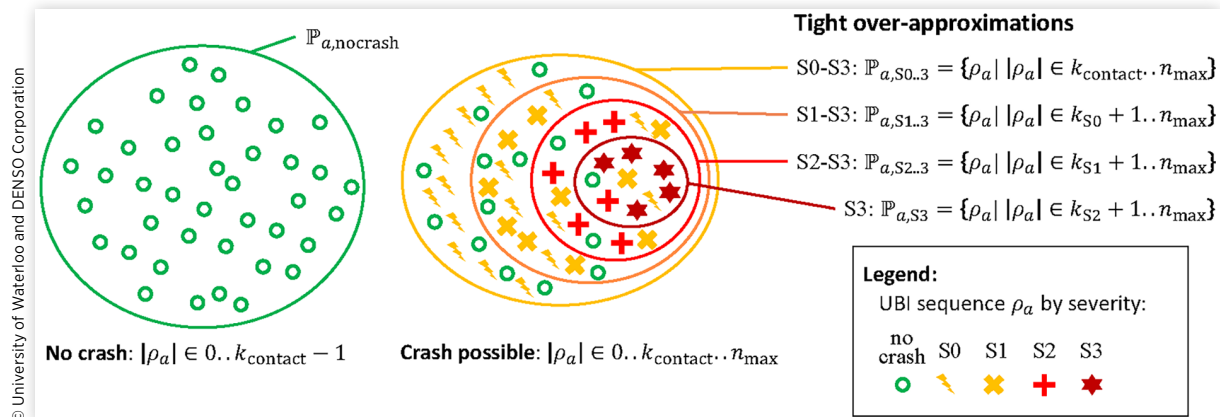


FIGURE 5 Illustration of the UBI patterns from Tab. 3 as sets of UBI sequences. Each glyph in the Venn diagram represents a particular UBI sequence ρ_a and its color indicates the severity of a crash that it would cause in the analyzed scenario. Note that the hazardous patterns are tight over-approximations of their corresponding severity range; e.g., $\mathbb{P}_{a,S3}$ contains all S3 sequences, but it also contains sequences of all other severities. The patterns are tight given their simple form as in eq. 12.



the composition of S and D is to estimate the distance d to the POV, limited by the maximum detection range r_{\max} of the sensor-detector combination:

$$\hat{d} = D(S(y')) \approx \bar{d}, \quad 0 \leq \hat{d} \leq r_{\max} \quad (13)$$

where \hat{d} is an estimate of the the ground-truth range-limited distance \bar{d} , defined as

$$\bar{d} = \begin{cases} d, & \text{if } 0 \leq d < r_{\max} \\ r_{\max}, & \text{otherwise} \end{cases} \quad (14)$$

Under these definitions, $\hat{d} < r_{\max}$ means that an object is detected, and $\hat{d} = r_{\max}$ means that no object is detected. If $\hat{d} = r_{\max}$ even though $\bar{d} < r_{\max}$, we have a case of a False Negative (FN) detection. Conversely, if $\hat{d} < r_{\max}$ even though $\bar{d} = r_{\max}$, we have a False Positive (FP) detection. Otherwise, we have a True Positive (TP) detection, with the estimate \hat{d} corrupted by a distance error $\epsilon_d = \hat{d} - \bar{d}$. The sensor-detector combination is complex and has a stochastic nature. Even though faults in the detector model are systematic, its input F is stochastic and conditioned on (i) the sensor characteristics, including random hardware noise, such as camera shot noise, and (ii) the IRCs, many of which are specified with uncertainty (because of their high-dimensional nature, such as appearance and weather). Thus, the sensor-detector combination is modeled probabilistically, and its error distribution is estimated through testing:

$$\hat{d} \sim p(\hat{d}|y') \quad (15)$$

The tracker \mathbb{T} takes a sequence of c latest distance estimates $\hat{d}^{[k(t)]}, \dots, \hat{d}^{[k(t)-c]}$ and produces the best estimate for the current time t . While an actual implementation would use an estimator such as a Kalman filter, we abstract the tracker to focus on its track management logic, which affects how detection errors propagate through the tracker. For our example, the tracker model \mathbb{T} expresses the typical “keep alive” logic, where a track is terminated (i.e., returning r_{\max}) when there is no detection for c consecutive time steps (i.e., each of the latest c distance estimates is r_{\max}); otherwise, it returns the ground-truth range-limited distance \bar{d} (for simplicity, our example ignores distance estimate errors):

$$\tilde{d} = \mathbb{T}(\hat{d}^{[k(t)]}, \dots, \hat{d}^{[k(t)-c]}, \bar{d}) = \begin{cases} r_{\max}, & \text{if } \hat{d}^{[k(t)]} = \dots = \hat{d}^{[k(t)-c]} = r_{\max} \\ \bar{d}, & \text{otherwise} \end{cases} \quad (16)$$

Note that while \mathbb{T} samples \hat{d} for c discrete time steps, it outputs \tilde{d} in continuous time, as expected by the DSM. Also, the range limiter block in Fig. 6 implements eq. 14. Further, we assume conservatively $\hat{d}^{[k]} = r_{\max}$ for $k < 0$.

Finally, the odometry Odo measures the SV speed. Internally, it consists of a sensor (e.g., a wheel encoder)

and an estimator (e.g., a Kalman filter). Its intended behavior is an identity function, but, similarly to object detection, its actual behavior is modeled probabilistically as $p(\hat{s}|\dot{s})$.

It is easy to show that the intended behavior of the DAS in this refined model is safe. Assuming a sufficient detection range $r_{\max} > s_{POV}$ and the intended behavior of each block as specified, the policy π receives the ground-truth d and \dot{s} , and thus the resulting SV braking behavior is same as for the HLSM in Figs. 2 and 4a.

Identification and Evaluation of HEPs. In the next step, the MoSAFE method helps to identify, specify, and evaluate hazardous deviations from the intended behavior as HEPs. The key idea is to determine error patterns on the input of each component in the DAS' that could lead to a given HBP identified in the previous activities (Clause 6). The analysis is performed in a backward direction starting from the HB at the output of the DAS' and identifying HEPs incrementally component by component towards the inputs y' of the DAS'. In our example (Fig. 6), the first step is to determine HEPs on the inputs into the policy π that would cause the policy to produce a UBI pattern of a given severity (Tab. 3). The next step is to determine HEPs on the inputs of the components that feed into the policy, e.g., \mathbb{T} , that would cause the HEPs on the policy inputs, and so on.

In order to support this analysis, the DSM is instrumented to allow injecting intended behavior and error sequences at specific component inputs, resulting in the instrumented DSM' and DAS" in Fig. 7. Intended behavior or error sequences are injected at a specific component input, e.g., \tilde{d} , by inputting specific sequences into the corresponding switches, i.e., $\gamma_{\tilde{d}}$ and $\rho_{\tilde{d}}$. For example, no injection at \tilde{d} occurs when $\gamma_{\tilde{d}} = \emptyset$; alternatively, the intended behavior for \tilde{d} is injected when $\gamma_{\tilde{d}} = \mathbb{N}_{r_{\max}} \wedge \rho_{\tilde{d}} = \emptyset$; and an error sequence from the pattern $\mathbb{P}_{\tilde{d}}$ is injected at \tilde{d} when $\gamma_{\tilde{d}} = \mathbb{N}_{r_{\max}} \wedge \rho_{\tilde{d}} \in \mathbb{P}_{\tilde{d}}$.

Table 4 defines four configurations of the instrumented DSM' that are relevant to the identification of HEP. Configuring the instrumented DSM' with the first configuration C_{DSM} , which turns all injection switches off, results in the original DSM from Fig. 6. The remaining three configurations allow injecting deviations from the intended behavior specified by patterns. The first of the three, $C_{\mathbb{P}_a}$, allows injecting UBI according to \mathbb{P}_a , while supplying π with ground-truth \dot{s} and \bar{d} ; this configuration is equivalent to the HLSM with UBI injection in Fig. 2. The remaining two configurations allow injecting r_{\max} (see Fig. 7), which corresponds to an FN error, at the output of the tracker \tilde{d} and detector \hat{d} , respectively. These three configurations help investigate how FNs cause UBIs and ultimately define HEPs for the detector and the tracker as FN patterns.

The first step in the HEP identification is to identify a tracker FN pattern $\mathbb{P}_{\tilde{d}}$ of all the FN sequences that would cause a particular UBI pattern \mathbb{P}_a . This can be phrased as determining the *weakest precondition* [15] on a component input to observe a particular output behavior. More

FIGURE 7 Detailed scenario model (DSM') for "braking for a stationary vehicle ahead" with element-level intended behavior and error injection (in red)

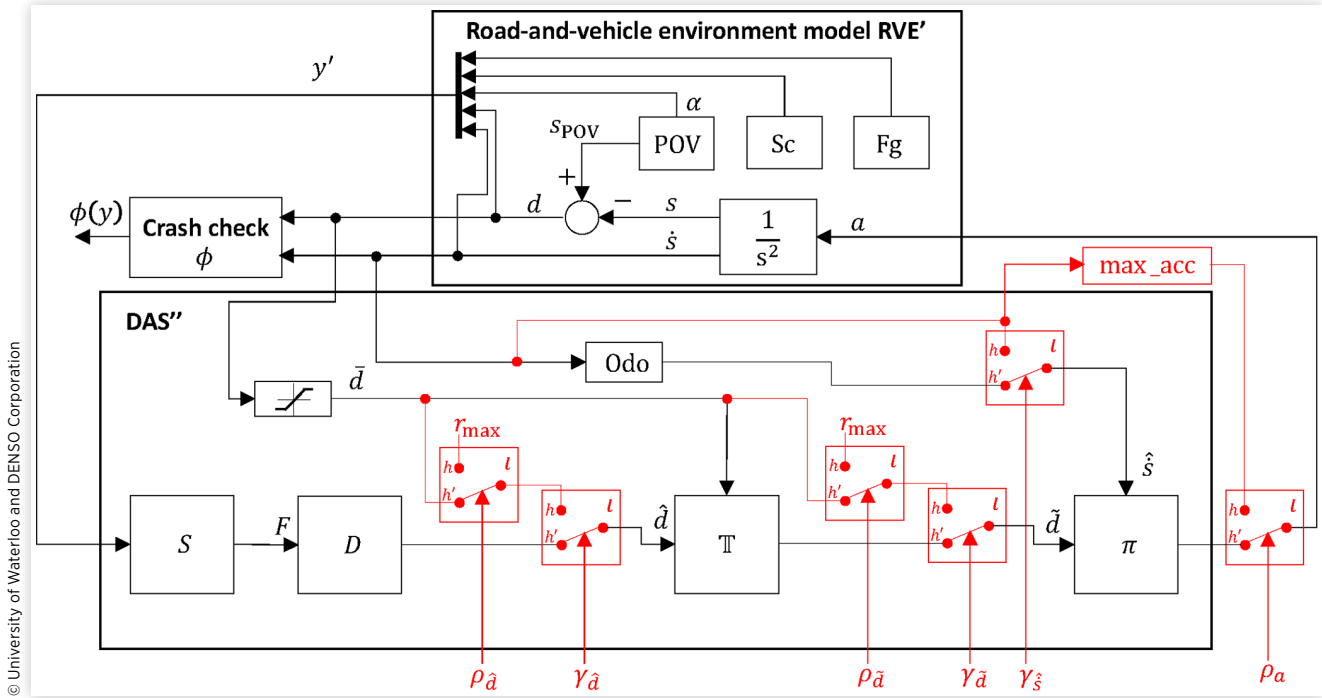


TABLE 4 Configurations of the instrumented DSM' from Fig. 7

Name	Definition	Description
C_{DSM}	$\gamma_{\dot{s}} = \gamma_{\dot{d}} = \gamma_{\dot{a}} = \rho_a = \emptyset$	DSM (as in Fig. 6)
C_{P_a}	$\gamma_{\dot{s}} = \gamma_{\dot{d}} = N_{n_{max}} \wedge \rho_{\dot{d}} = \emptyset \wedge \rho_a \in \mathbb{P}_a$	DSM with injected P_a and intended behavior for \dot{s} and \dot{d} (equivalent to HLSM with injected P_a)
$C_{\mathbb{P}_{\dot{d}}}$	$\gamma_{\dot{s}} = \gamma_{\dot{d}} = N_{n_{max}} \wedge \rho_a = \emptyset \wedge \rho_{\dot{d}} \in \mathbb{P}_{\dot{d}}$	DSM with injected $\mathbb{P}_{\dot{d}}$ and intended behavior for \dot{s}
$C_{\mathbb{P}_{\dot{d}}}$	$\gamma_{\dot{s}} = \gamma_{\dot{d}} = N_{n_{max}} \wedge \gamma_{\dot{d}} = \rho_a = \emptyset \wedge \rho_{\dot{d}} \in \mathbb{P}_{\dot{d}}$	DSM with injected $\mathbb{P}_{\dot{d}}$ and intended behavior for \dot{s}

precisely, the step is to determine $\mathbb{P}_{\dot{d}}$ such that the set of all behaviors of \dot{d} under $C_{\mathbb{P}_{\dot{d}}}$ is the largest for which the output a of π behaves like under C_{P_a} . In other words, we want to find all FN sequences $\rho_{\dot{d}} \in \mathbb{P}_{\dot{d}}$ that if injected at \dot{d} using configuration $C_{\mathbb{P}_{\dot{d}}}$ would result in the output a of π to behave as if UBI sequences \mathbb{P}_a were injected at a using C_{P_a} . Formally, we will denote these sequences by the Weakest Precondition Pattern (WPP) on \dot{d} for \mathbb{P}_a , defined as follows:

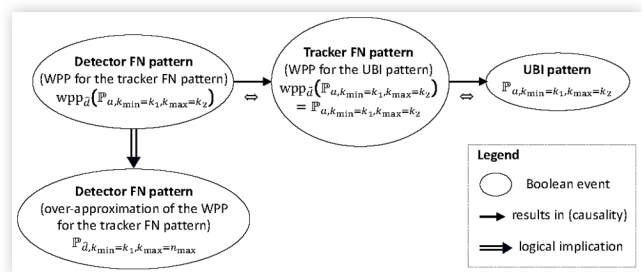
$$wpp_{\dot{d}}(\mathbb{P}_a) = \left\{ \rho_{\dot{d}} \in N_{n_{max}} \left(a \Big|_{C_{\mathbb{P}_{\dot{d}}=\{\rho_{\dot{d}}\}}} \subseteq a \Big|_{C_{P_a}} \right) \right\} \quad (17)$$

where $a|_{C_{P_a}}$ represents the set of a behaviors under configuration C_{P_a} , and $a|_{C_{\mathbb{P}_{\dot{d}}=\{\rho_{\dot{d}}\}}}$ represents a behaviors under configuration $C_{\mathbb{P}_{\dot{d}}}$ where $\mathbb{P}_{\dot{d}} = \{\rho_{\dot{d}}\}$.

The relationship between $wpp_{\dot{d}}(\mathbb{P}_a)$ and \mathbb{P}_a can be seen as causal (see the right arrow in Fig. 8): the occurrence of an FN tracking error sequence from $wpp_{\dot{d}}(\mathbb{P}_a)$ causes the occurrence of a UBI sequence from \mathbb{P}_a , and there are no other FN sequences not captured in $wpp_{\dot{d}}(\mathbb{P}_a)$ that could cause a UBI sequence in \mathbb{P}_a . Given a UBI pattern $\mathbb{P}_a, k_{min} = k_1, k_{max} = k_2$, specified by bounding the number of time steps when UBI occurs (eq. 12), its weakest

FIGURE 8 Causal model for UBI pattern $\mathbb{P}_a, k_{min} = k_1, k_{max} = k_2$.

Single-line arrows represent causality; double-line arrows represent logical implication (and no causality) and are used to model over-approximation. The bi-implication symbol \Leftrightarrow annotating each causal arrow indicates, by a slight abuse of notation, an identity function associated with each of them, meaning that the pattern on the right of a causal arrow annotated with the bi-implication occurs if and only if the pattern on the left of the arrow occurs. Because of the direction of causality, the pattern on the right of a causal arrow will occur with some delay relative to the occurrence of the pattern on the left.



precondition at \tilde{d} is the same pattern, i.e., $\mathbb{P}_{\tilde{d}, k_{\min}=k_1, k_{\max}=k_2} = \mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2} = \text{wpp}_{\tilde{d}}(\mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2})$.

This is easy to see from the policy definition (eq. 5), since injecting $r_{\max} > s_{\text{POV}}$ at \tilde{d} (which is used as the first argument into $\pi(d, \hat{s})$) restricts the policy to its first two cases (i.e., $a = 0$ or a_{\max}), depending on speed; this is because $d > \Delta s_{\text{stand}} \wedge a_{b, \text{req}} < a_{b, \text{min}}$ is true under $d = r_{\max} > s_{\text{POV}}$ and these are equivalent to injecting a UBI at a . In other words, injecting an FN error at \tilde{d} at step k causes a UBI at a at that step. Note that, for simplicity, this analysis assumes perfect distance \tilde{d} and speed \hat{s} estimation, as modeled in DSM' under $C_{\mathbb{P}_{\tilde{d}}}$. The first one is based on the tracker taking d as input (see eq. 16), and the latter, which is the intended behavior of *Odo*, is injected with $\gamma_{\hat{s}} = \mathbb{N}_{n_{\max}}$. We also ignore any perception-reaction delays. We will discuss how to relax these assumptions in the next section.

The next step is to establish the cause of the FN tracking error pattern $\mathbb{P}_{\tilde{d}, k_{\min}=k_1, k_{\max}=k_2}$, i.e., $\text{wpp}_{\tilde{d}}(\mathbb{P}_{\tilde{d}, k_{\min}=k_1, k_{\max}=k_2})$, which is equivalent to $\text{wpp}_{\tilde{d}}(\mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2})$ (see the left causality arrow in Fig. 8). Computing it requires analyzing the tracker logic (eq. 16), which compensates for c consecutive FN detection errors. Specifying $\text{wpp}_{\tilde{d}}(\mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2})$ is difficult, but we can easily provide a conservative over-approximation by observing that the tracker may reduce the number of FN detection errors by its compensation logic but will never introduce additional ones. Also, if all FN detection errors occur as a continuous sequence starting at $t = 0$, there will be the same number of FN tracking errors in the output. Thus, $\text{wpp}_{\tilde{d}}(\mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2})$ contains sequences each with at least k_1 FN detection errors. Also, some sequences may have more than k_2 FN detection errors, because of possible compensation by the tracker. Thus, we can over-approximate $\text{wpp}_{\tilde{d}}(\mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2})$ by $\mathbb{P}_{\tilde{d}, k_{\min}=k_1, k_{\max}=n_{\max}}$, which is represented by the bottom node in Fig. 8.

It is worth noting that the causal chain of the three events at the top of Fig. 8 can be understood as a Structural Causal Model (SCM) [16]. An SCM is a directed acyclic graph with nodes representing random variables, and arrows representing causal influence; the arrows incoming to a node are also associated with a function that maps the variables at the other end of the arrows to the node. In our case, the variables are Boolean and represent the occurrence of an error pattern, and the functions associated with the causal arrows are identity. The SCM can also be understood as an FT [17], which is a causal model with nodes as Boolean events; causal arrows associated with Boolean functions; the sink node (aka top event) representing a system-level failure; and its ancestors (in the causal direction) being errors or faults or both. In our case, the top event is an UBI pattern and the other events represent error patterns. The fourth node at the bottom of the figure and the logical

implication (which does not express causality) are not part of the SCM concept, but allow us to express over-approximation, which is useful when the exact cause is difficult to specify.

The causal model in Fig. 8 allows us to limit the occurrence probability of a UBI of at least a given severity by limiting the occurrence of the FN detection error pattern that represents the over-approximation of the cause of the UBI pattern. This is because $P(\mathbb{P}_{\tilde{d}, k_{\min}=k_1, k_{\max}=n_{\max}}) \geq \mathbb{P}_{a, k_{\min}=k_1, k_{\max}=k_2}$ due to the over-approximation, and we can also set $k_2 = n_{\max}$ without affecting this inequality. Table 5 shows such HEPs at \tilde{d} and \hat{d} that cause or may cause, respectively, an UBI pattern of the corresponding severity range.

In our sample scenario (Fig. 4a), injecting $\mathbb{P}_{\tilde{d}, S0.3}$ corresponds to the occurrence of $k_{\min}=19$ or more FN detection errors within 150 frames, which may cause 19 or more time steps with UBI within the 150 time steps, and thus a potential collision. Conversely, no collision due to UBI caused by FN detection errors is possible in our scenario if fewer than 19 FNs occur within 150 frames. Similarly, no collision of severity S3 due to UBI caused by FN detection errors is possible in our scenario if fewer than 35 FNs occur within 150 frames.

Limiting the occurrence probability of the patterns in the third column in Tab. 5 allows us to limit the occurrence probability of the UBI patterns with the corresponding severity in the first column, and thus limit the risk of harm due to UBI.

IRC Identification and HEP Likelihood Estimation for AI-Based Components

The analysis in the previous subsection involved computing WPPs for the conventional components in our example, i.e., the driving policy and the tracker, whose scenario-specific behavior can be modeled precisely and whose implementation in conventional software can be verified against the model by the methods prescribed in ISO 26262:2018 [3], including inspection. It is typically infeasible

TABLE 5 Tracker FN patterns that cause and detector FN patterns that may cause UBI patterns of a given severity

UBI pattern	Tracker FN pattern (at \tilde{d}) being a WPP of UBI	Detector FN pattern (at \hat{d}) being an over- approximation of WPP of UBI
$\mathbb{P}_{a, S0.3}$	$\mathbb{P}_{\tilde{d}, S0.3} = \mathbb{P}_{a, S0.3}$	$\mathbb{P}_{\tilde{d}, S0.3} = \mathbb{P}_{\tilde{d}, k_{\min}=k_{\text{contact}}, k_{\max}=n_{\max}}$
$\mathbb{P}_{a, S1.3}$	$\mathbb{P}_{\tilde{d}, S1.3} = \mathbb{P}_{a, S1.3}$	$\mathbb{P}_{\tilde{d}, S1.3} = \mathbb{P}_{\tilde{d}, k_{\min}=k_{S0}, k_{\max}=n_{\max}}$
$\mathbb{P}_{a, S2.3}$	$\mathbb{P}_{\tilde{d}, S2.3} = \mathbb{P}_{a, S2.3}$	$\mathbb{P}_{\tilde{d}, S2.3} = \mathbb{P}_{\tilde{d}, k_{\min}=k_{S1}, k_{\max}=n_{\max}}$
$\mathbb{P}_{a, S3}$	$\mathbb{P}_{\tilde{d}, S3} = \mathbb{P}_{a, S3}$	$\mathbb{P}_{\tilde{d}, S3} = \mathbb{P}_{\tilde{d}, k_{\min}=k_{S2}, k_{\max}=n_{\max}}$

to produce such behavior specifications for AI-based components that rely on deep neural networks, such as the object detector, because they often implement very complex functions over highly dimensional inputs, such as images. Additionally, the logic implemented by neural networks cannot be inspected the way conventional programs can (for a comprehensive discussion of these issues in the context of safety assurance, see [4]). Therefore, AI-based components require a different approach.

Rather than attempting to determine WPPs over inputs like images that would cause specific HEPs on the output of an AI-based component, the combination of sensor and detector is modeled probabilistically, that is, $p(\hat{d}|y')$ in our example, and the probability of HEPs is estimated via testing. Testing requires test data that adequately covers the input conditions fed into an AI-based component. For a perception component, these are the IRCs reflecting y' in $p(\hat{d}|y')$. While IRCs like object and scenery appearance and weather conditions are multidimensional and complex, existing road ontologies (e.g., [14, 18]) can be used to express IRCs and partition the range of IRCs expected in operation. The test data, such as sensor recordings from drives in the target ODD, potentially augmented with synthetic data, would then be used to estimate the HEP probability in each partition. In our example, the IRCs would include different types of vehicles as POV, different poses, road configurations, and weather conditions. The test data would be used to estimate the probability of detector HEPs (from Tab. 5) under different IRCs. In our example, the modeled scenario would represent one of such partitions, and we would estimate $P(\mathbb{P}_{a,s1.3}|y')$. The HEP probabilities then need to be aggregated over the corresponding IRCs and HBSCs and their occurrence rates in operation to estimate the final safety risk due to UBI in operation. In our example, the occurrence rate of $\mathbb{P}_{a,s1.3}$ due to the analyzed scenario conditions in operation would be $\lambda(\mathbb{P}_{a,s1.3}, y) = P(\mathbb{P}_{a,s1.3}|y')\lambda(y)$, where y' covers the combination of IRCs and HBSCs, and $\lambda(y')$ denotes their occurrence rate (such as per kilometer driven; see [19, 10] for guidance on calculating exposure). The pattern occurrence rates would be then aggregated over the remaining scenario partitions (i.e., combinations of IRCs and HBSCs). This partitioning of HBSCs and IRCs and aggregation of probabilities and occurrence rates can be captured in a safety case, as described later.

In contrast to the approach outlined above, the SOTIF standard asks for the identification of triggering conditions, which may neither be feasible nor necessary to assure safety. The standard defines triggering conditions as those scenario conditions that cause HB; thus, they are the subset of IRCs that cause HEPs. Whereas knowing the causal link between IRCs and HEPs may benefit addressing the corresponding functional insufficiency and creating stronger safety cases, establishing this link may be difficult or even infeasible. For example, a misdetection of an object in an image may be triggered by the context of the object rather than the object itself; it may also

be due to particular noise or general appearance of the scene that cannot be easily specified or described in words. Further, creating counterfactual images that change just the specific triggering condition while leaving all other conditions unchanged, which may be necessary to establish causality, may be infeasible. On the other hand, statistical causal influences may be more practical to establish, e.g., by injecting specific weather conditions into a data set and relating it to HEP rate, but this approach may not be feasible for all types of IRCs. It is worth noting that identifying triggering conditions for specific decision of an AI-based component is the subject of Explainable AI [20], but existing research still lacks explanation methods that would benefit safety assurance of DAS in practice. Most importantly, establishing causal links between IRCs and HEPs does not seem required for assuring safety. Partitioning IRCs and estimating HEP rates in each partition can provide a statistical assurance [13] without the need for establishing causality. Lists of known common triggering conditions may still be used to inform the partitioning of IRCs and test data selection.

Guidance on assuring that AI-based components meet their safety-related performance requirements, such as HEP rates derived using MoSAFE, is subject to ISO/PAS 8800 [21]. The standard covers safety-related properties and risk factors impacting the insufficient performance and malfunctioning behavior of AI systems and components. Most related to IRC identification and HEP rate estimation, it provides guidance on the derivation of suitable safety requirements as well as data quality and completeness requirements. It also provides guidance on measures for the control and mitigation of failures in AI systems and components.

When the estimated HEP occurrence rates for an AI-based component exceed the upper limits imposed by the acceptance criteria, the DAS needs to be modified to reduce the risk (Clause 8). The possible modifications include improving the performance of the AI-based component to reduce the HEP occurrence rates, robustifying the downstream components to be less sensitive to the error sequences in the HEPs, such as by increasing safety margins (e.g., see [22, 23]), and restricting the ODD to exclude the IRCs that cause the high HEP occurrence rates.

Additional Considerations and Future Work

The previous sections demonstrated MoSAFE on a necessarily simplified example. This section discusses how to handle more complex scenarios and models, including multiple types of HBs and multi-input errors, relaxing some of the assumptions made earlier. It also illustrates how to use MoSAFE results as part of a safety case. Finally, it discusses challenges, limitations, and suggestions for future work.

Designing Adequate Environment and System Models

MoSAFE relies on models of the DAS and the road and vehicle environment to identify and evaluate HBPs (Clause 6) and HEPs (Clause 7), and the adequacy of these models determines the validity of the risk analysis results. The main choices relate to the level of abstraction, such as selection of the real-world phenomena and their approximations to be included. The HLSMs represent the HBSCs and driving policies thereunder and can be typically represented using simple kinematic models and rule-based policies, as in our example of braking for a stationary vehicle ahead. Existing safety frameworks provide examples of such kinematic and rule-based models for a wide range of scenarios [24, 25]. The selection of HBSCs can rely on existing road environment ontologies [14, 18]. The DSMs need to refine the input space into the DAS by adding IRCs. The refinement of this input space can again rely on the existing road environment ontologies, which provide standard classifications of road users and other objects, road structures, and weather and visibility conditions. The refined DAS model includes the decision logic relevant to the scenario, which can be extracted from the detailed DAS design. While our DAS model ignored time delays, perception and reaction time delays should be included. In our example, this would be done by starting the scenario before the braking needs to occur, extending the UBIs by reaction delay, and also including velocity and distance prediction logic in the model. The models should be verified against the DAS, for example, using simulation testing.

Misbehavior Injection

The UBI used in our example is applicable to a large range of scenarios that require braking as intended behavior. To evaluate a wider applicability of our approach, we have modeled UBIs occurring in braking for a stationary object ahead, braking for a slower vehicle merging in front, and braking for a hard-stopping front vehicle. We have also modeled Unintended Acceleration (UA) when waiting behind a stopped vehicle and steady following of a front vehicle, and Failure to Yield (FTY) at a stop-controlled intersection. We were able to model these scenarios and HBs using the same approach as in the presented example. In these scenarios, in addition to HBPs of the form $k_{\min}..k_{\max}$ HB occurrences within n_{\max} time steps, we have also identified two more useful forms: $k_{\min}..k_{\max}$ consecutive HB occurrences within n_{\max} time steps, and $k_{\min}..k_{\max}$ consecutive HB occurrences at the beginning of a scenario. Future research should explore HBs representing deviations from intended lateral behavior.

Another type of HB applicable to our example is unintended insufficient braking (UIB), which is braking with a lower level than intended (Tab. 1). In our example, we can define it as a reduction of the required braking $a_{b,req}$ by the factor $(1 - \eta_{a_{b,req}})$, where $\eta_{a_{b,req}} : [0..T_{\max}] \rightarrow [0..1]$ is an error signal representing the relative reduction of $a_{b,req}$ over time. Note that an error signal is the continuous analogue of an error sequence. Similar as for UBI, we can define a UIB pattern by bounding the magnitude of the reduction by some maximum $0 < \hat{\eta} \leq 1$:

$$\mathbb{P}_{a_{b,req}, \hat{\eta}} = \left\{ \eta_{a_{b,req}} : [0..T_{\max}] \rightarrow [0..\hat{\eta}] \right\} \quad (18)$$

Given a signal error $\eta_{a_{b,req}} \in \mathbb{P}_{a_{b,req}, \hat{\eta}}$, it can be injected into HLSM and DSM by multiplying the the right-hand side of eq. 4 by $(1 - \eta_{a_{b,req}})$.

We can analyze the worst-case effect of $\eta_{a_{b,req}} \in \mathbb{P}_{a_{b,req}, \hat{\eta}}$ with the constant error signal with value $\hat{\eta}$. For example, Fig. 4d shows the speed and acceleration profile assuming a reduction $\eta = 0.14$ of the required braking $a_{b,req}$ throughout the scenario. For comparison, the dashed curves show the intended behavior of braking with $a_{b,req}$. As a result of the reduced braking, the SV does not brake until $0.86a_{b,req}$ reaches $a_{b,min}$ at $s_{b,start}$; subsequently, the driving policy applying $0.86a_{b,req}$ compensates the initial lack of braking in order to stop for s_{stop} , requiring a high-level of deceleration before reaching it. For example, $0.86a_{b,req}$ reaches the maximum braking of $8m/s^2$ about 2 ms before stopping, but its speed is just 0.2m/s. As a result, the SV will overshoot s_{stop} , but only by less than 2.5 cm, which is safe. Even a 50% reduction of $a_{b,req}$ would lead to less than 2m overshoot, which would be safe assuming $\Delta s_{stand} = 5$ m.

The approach to represent and inject misbehaviors using binary or continuous patterns is quite general and flexible. Although our examples inject only constants and simple functions of the ground truth, more complex misbehaviors can also be injected.

Fault Tree Derivation under Multi-Input Failures

The example in Fig. 7 considered a single input into an element being erroneous at a time, such as a FN track detection at \hat{d} into the policy and the FN object detection at \hat{d} into the tracker. This resulted in the fault tree at the top of Fig. 8 being a sequence of three nodes. In practice, an element may receive more than one erroneous input at the same time. For example, the policy π could receive both a pattern of FN tracking errors at \hat{d} and a pattern of erroneous speed estimates \hat{s} at the same time. One way to deal with this more general case is to partition the input error cases into *single-input failure* and *multi-input failure*, which will cause the fault tree to be a general tree rather than a sequence, and even more generally a directed acyclic graph if inputs are shared among multiple elements.

As an example, consider Hazardous Braking (HBB) as a combination of hazardous UBI and UIB, ignoring Unintended Hard Braking (UHB) for simplicity. HBB constitutes the top event of the fault tree in Fig. 9. This event is partitioned into three cases:

1. $UBI \wedge \neg UIB$: a hazardous UBI with an otherwise nominal required braking level (UIB), i.e., one that is not reduced by more than some nominal maximum $\hat{\eta}_{a,max}$ when braking is not interrupted; this case is formalized as a combination of $\mathbb{P}_{a_b,req}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$ representing $\neg UIB$ (eq. 18) and $\mathbb{P}_{a,S0.3}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$ representing UBI, with the latter defined same as $\mathbb{P}_{a,S0.3}$ in Tab. 3, but with $k_{contact}$ determined under $a_{b,req}$ reduced according to $\mathbb{P}_{a_b,req}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$. Thus, the severity bounds for τ_{total} in Tab. 2 would need to be recomputed assuming the worst-case but still nominal reduction of $a_{b,req}$ by $\hat{\eta}_{a,max}$; this corresponds to UBIs on a braking profile like in Fig. 4d.
2. $\neg UIB \wedge UIB$: a (potentially) hazardous UIB that exceeds the nominal level $\hat{\eta}_{a,max}$, denoted by $\mathbb{P}_{a_b,req}^{\hat{\eta} > \hat{\eta}_{a,max}}$, while UBI is bounded to a level that would be safe under nominal braking reduction levels $\mathbb{P}_{a_b,req}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$, denoted by $\mathbb{P}_{a,nocrash}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$;
3. $UBI \wedge UIB$: a combination of hazardous UBI and UIB, i.e., $\mathbb{P}_{a,S0.3}^{\hat{\eta} \leq \hat{\eta}_{a,max}} \times \mathbb{P}_{a_b,req}^{\hat{\eta} > \hat{\eta}_{a,max}}$.

Formally, each of the three composite patterns is a cross-product.

Looking at eq. 4, a UIB can be caused by an underestimate of \hat{s} or an overestimate of d or both. For simplicity, we only consider underestimates of \hat{s} . Based on eq. 4, injecting $\mathbb{P}_{\hat{s},\hat{\eta} \leq \hat{\eta}_{s,max}}^{\hat{s}}$ with $\hat{\eta}_{s,max} = 1 - \sqrt{1 - \hat{\eta}_{a,max}}$ is equivalent to injecting $\mathbb{P}_{a_b,req}^{\hat{\eta} \leq \hat{\eta}_{a,max}}$; similarly, injecting $\mathbb{P}_{\hat{s},\hat{\eta} > \hat{\eta}_{s,max}}^{\hat{s}}$ is equivalent to injecting $\mathbb{P}_{a_b,req}^{\hat{\eta} > \hat{\eta}_{a,max}}$. The patterns for \hat{s} are defined similarly to eq. 18 and injected by multiplying \hat{s} in eq. 4 by $(1 - \eta_{\hat{s}})$.

Each of the three cases under the top event are then decomposed using *and*-gates (Fig. 9). The patterns on \tilde{d} are defined as in the second column of Tab. 3, but with severity bounds determined by assuming nominal speed underestimates $\mathbb{P}_{\hat{s},\hat{\eta} \leq \hat{\eta}_{s,max}}^{\hat{s}}$. The first two cases are examples of single-input failures, where one of the two inputs \tilde{d} and \hat{s} exceeds its threshold while the other does not, and the third case is a multi-input failure where both inputs exceed their thresholds.

The fault tree (Fig. 9) allows computing the probability of the top event given the probability of the leaves (we use the descriptive names of the nodes rather than the patterns for readability):

$$P(\text{HBB}) = P(\text{Hazardous-tracking-FNs} \wedge \text{Nominal-speed-estimate}) + P(\text{Safe-tracking-FNs} \wedge \text{Off-nominal-speed-estimate}) + P(\text{Hazardous-tracking-FNs} \wedge \text{Off-nominal-speed-estimate}) \quad (19)$$

Since FNs and SV speed estimation errors are independent in our system, we have

$$P(\text{HBB}) = P(\text{Hazardous-tracking-FNs}) + P(\text{Nominal-speed-estimate}) + P(\text{Safe-tracking-FNs}) + P(\text{Off-nominal-speed-estimate}) + P(\text{Hazardous-tracking-FNs}) + P(\text{Off-nominal-speed-estimate}) \quad (20)$$

Finally, since both $P(\text{Hazardous-tracking-FNs})$ and $P(\text{Off-nominal-speed-estimate})$ are small, and thus the probabilities of their complements, $P(\text{Safe-tracking-FNs})$ and $P(\text{Nominal-speed-estimate})$, respectively, are close to one, we have

$$P(\text{HBB}) \approx P(\text{Hazardous-tracking-FNs}) + P(\text{Off-nominal-speed-estimate}) \quad (21)$$

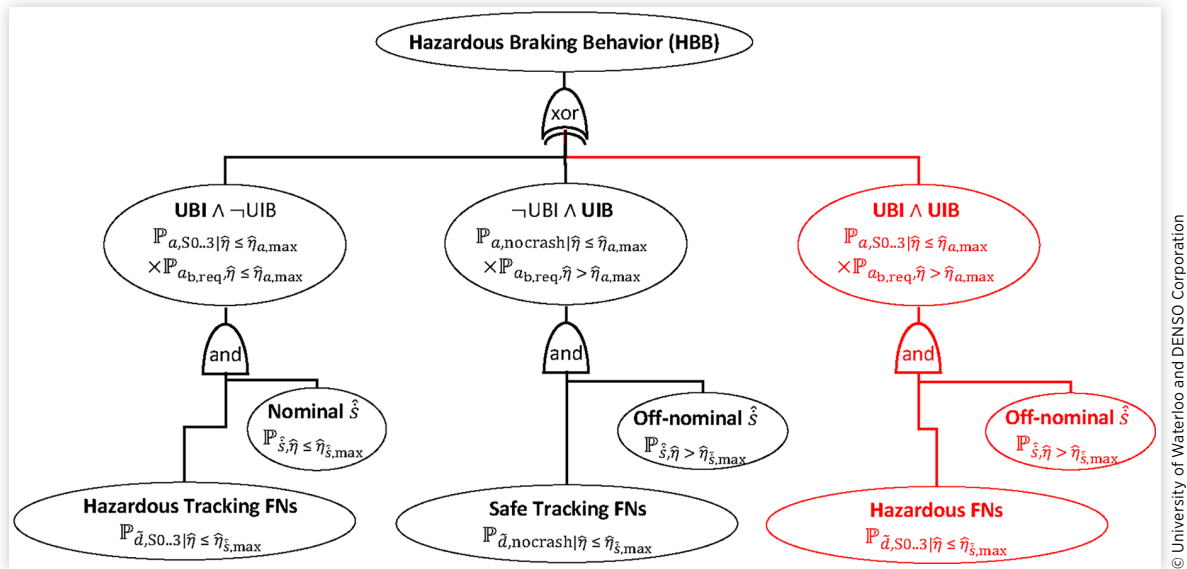
This modeling and probability-calculation approach, allowing to eliminate the multiple-failure branch marked in red in Fig. 9, is standard in Fault Tree Analysis (FTA) [17, 26]. Further, the approach of combining error patterns on multiple inputs can be easily extended to more than two erroneous inputs, which would allow us to also include a hazardous over-estimate of \tilde{d} , $\mathbb{P}_{\tilde{d},\hat{\eta} > \hat{\eta}_{\tilde{d},max}}^{\tilde{d}}$, as another cause of UIB, and a “nominal” over-estimate of \tilde{d} , $\mathbb{P}_{\tilde{d},\hat{\eta} \leq \hat{\eta}_{\tilde{d},max}}^{\tilde{d}}$, for $\neg UIB$.

In some cases, such as when modeling fault-tolerance of redundant elements, we want to define an HEP over two or more inputs rather than freely composing HEPs over individual inputs. Such HEPs over multiple inputs would be useful to model simultaneous FN detection errors in multiple sensor modalities, e.g., camera and lidar.

WPP Derivation

The computation of WPPs with FN tracking error sequences causing hazardous UBI patterns in our example was particularly simple due to the simple driving policy logic, but even the relatively simple tracker logic required us to resort to an over-approximation to express WPPs with FN detection error sequences that cause hazardous UBI patterns. The scenario-specific decision logic of an element captured by the element model, although simplified compared to the full decision logic of the element, may still cause the WPP computation to be challenging. The computation should be supported by reasoning tools, which should be explored in future work. In particular, the

FIGURE 9 Sample fault tree considering multiple inputs into an element being erroneous: hazardous braking behavior caused by FN tracking errors or SV speed underestimates or both



tooling should allow engineers to strike a balance between the complexity of a WPP and its precision.

A causal model may contain multiple over-approximation nodes and implication arrows. WPP computation may be applied to an over-approximation, and the WPP itself might need to be over-approximated again. Thus, the ability to mix causal arrows and implications in causal models is important.

Model Validation Using Simulation Testing

Testing is also necessary to validate the DAS and SV models used in MoSAFE, that is, assuring that they reflect the actual behavior of the DAS and the SV with sufficient accuracy. This approach of targeted testing to validate assumptions follows the layered residual-risk framework to ADS testing [27]. One of the basic ideas of this framework is to “use higher-fidelity simulation and tests to reduce residual risks due to assumptions and gaps in lower fidelity simulations and tests.” [27] Similarly, closed-course tests can be used to reduce residual risks due to assumptions and gaps in simulation testing. Finally, field testing is used primarily to discover new scenarios and requirements, including gaps in scenario catalogs, and to collect field data for testing AI-based components and for estimating exposure, i.e., occurrence rates of scenario conditions.

For our sample braking scenario, we have compared the analytical results of MoSAFE from the HLSM and DSM with the test results of injecting HEPs into WISE-ADS, a real ADS software stack running in a high-fidelity simulator [28]. The ADS uses a lidar-based object detector using a deep neural network, a multi-object

tracker based on a Kalman filter, a rule-based behavior planner, an optimization-based trajectory planner, a PID speed controller and a geometric steering controller, implemented on top of the Robot Operating System. In particular, the testing involved injecting a single sequence of consecutive FN detection errors of various lengths and at different points in time during the braking scenario by intercepting the communication between the object detector and the tracker. Based on the dynamics represented by the HLSM, we know that only consecutive FNs provide the largest risk effect, i.e., the smallest slowdown, for a given number of FNs, and thus patterns of intermittent FNs do not need to be simulated for this scenario. The simulation environment includes a high-fidelity 14-DOF vehicle dynamics model of the physical SV [29, 30]. The braking behavior of the ADS and the resulting velocity profiles in simulation were experimentally confirmed to closely resemble the behavior of the SV, 2015 Lincoln MKZ, controlled by the ADS on a test track [28].

The DAS model in Fig. 6 reflects the main components of the ADS that are involved in the braking scenario. The model was updated to account for perception-reaction time, including the transition from maximum acceleration the maximum braking (such as at the end of the second UBI interval in Fig. 4b), effectively extending the duration of the UBI with a_{\max} before $a_{b, \max}$ is applied. The DSM parameters (i.e., $a_{b, \min}$, $a_{b, \max}$, a_{\max} , v_{\max} , Δs_{stand} , Δt , and c in Fig. 6 and the perception-reaction delays) were calibrated with parameter values matching those of the ADS and SV combination, based on simulation experiments. Given these parameter values, WPP analysis as presented earlier predicts that injecting 7 or fewer frames with an FN detection error for the stationary POV during the scenario should be safe (this injection corresponds to $C_{\mathbb{P}_g}$ in the last row in Tab. 4). If the POV has been detected at the beginning of the scenario, then this number

increases by $c = 9$ frames, which are compensated by the tracker. Since our DSM predicts that injecting 7 or fewer frames with an FN detection error for the POV at the beginning of the scenario is safe, we determine that injecting $k=7+9=16$ or fewer frames with an FN detection error for the POV is safe.

To assess the result from MoSAFE, we run a simulation experiment of injecting $k=23, 24, 25, 26, 27$ consecutive frames with FN detection error for the POV into the tracker in the actual ADS running in the high-fidelity simulator, with the injection occurring at the optimal point during the scenario as determined by the calibrated DSM, and repeating the simulation 100 times for each k . We observed that $k=23$ or less was safe: $k=23$ resulted in 0 collisions, whereas $k=24$ resulted in 17% of the runs, and $k=27$ resulted in 100% of the runs experiencing a collision. The estimate from the DSM of $k=16$ or less being safe is lower than $k=23$ from simulation. This is mainly because the DSM makes the conservative assumption that the SV accelerates with a_{\max} during the UBI, but the SV accelerated at lower rates than the maximum in the actual simulation runs. Given that the SV with no UBI stops within $n_{\max}=56$ frames, the crash risk from UBI in this scenario can be limited by limiting the probability of $k=24$ or more frames with FN detection error for the POV occurring anywhere (including intermittently) within $n_{\max}=56$ consecutive frames.

Whereas simulation testing is still required to validate the DSM, the validation is much less costly than exploring the effect of injecting error sequences without a model. The analytical model states that multi-interval UBI can always be replaced by a single UBI interval with the same duration and a more severe effect, and thus the main uncertainty to be addressed by targeted simulation is about the system transition dynamics, the actual acceleration and deceleration profiles, and the SV-speed and distance-to-POV estimation errors. Running the 500 simulations in our validation experiment (i.e., 100 runs for each of the five values of k) took 7.5 h on a single modern desktop computer. Exploration of the effects of injecting FN detection error sequences would have likely required orders of magnitude more simulation runs with random error sequences injected to arrive at similar results, which would have been prohibitively expensive.

Trade-offs Between Model-Based Verification and Simulation Testing

MoSAFE can be viewed as model-based, modular verification, and an alternative to this approach is to perform *automated black-box testing* in simulation [31]. Such testing samples inputs automatically and executes the system under test as a black-box in a simulation environment. To be effective, it needs to use some form of optimization during input sampling in order to focus on finding and

exploring hazardous inputs efficiently. These automated approaches can be divided into three categories: (i) falsification, which finds inputs that violate a safety property, i.e., cause an HB; (ii) most-likely failure analysis, which tries to find maximum-likelihood failures, i.e., most-likely HB occurrences; and (iii) failure probability estimation, which estimates the probability of HB occurrence [31]. Approaches in the third category are most relevant to risk estimation.

Black-box testing can be applied to a whole DAS or parts of it. For example, it could be used to test and evaluate risk by generating perception inputs into a DAS, either by generating them using computer graphics or synthesizing them using neural rendering (e.g., neural radial fields [32]), or synthetically perturbing real or synthetic inputs, e.g., by adding weather effects. Alternatively, it can be applied to the prediction and planning portion by using simulated perception results as input. Black-box testing can also be applied in a gray-box setting by injecting error sequences into internal interfaces of the DAS, for example, to understand how prediction errors might propagate through the system (e.g., [33]), and to leverage the test optimization capability to find component inputs that cause HBs.

There are important trade-offs between model-based verification and simulation-based black-box testing. The key advantage of black-box testing methods is that it can be performed on the actual system software implementation and does not require a model of the system. On the other hand, while model-based verification requires the additional effort to create models, it can provide stronger guarantees than black-box testing, because it performs an exhaustive analysis of the models. In particular, an HLSM represents an exponentially large class of concrete scenarios, because it is parameterized, and MoSAFE analyzes it exhaustively. Further, the assume-guarantee reasoning over a DSM, as in the WPP derivation, establishes the full causal links between errors and system failures, which can be leveraged in a safety case. Black-box testing may require large number of samples to generate insights comparable to verification, and it does not provide guarantees. Also, having these WPPs as part of interface contracts allows modular verification, including modular testing. For example, AI components can be subjected to unit testing against an interface specifying HEP occurrence rate limits, which enables independent development, as in the case of automotive supply chains, and it also can improve test depth given the same amount of test budget [34].

Ultimately, it may be best to combine both approaches in practice. At the vehicle level, kinematic models of the road-and-vehicle environment and simple rule-based driving policies often allow for closed-form specification of HB patterns, as in our example. More complex scenarios and policies may require black-box test simulation, with a potentially significant computational cost and loss of guarantees. Similarly, element-level models of DAS may include complex logic necessitating black-box testing, possibly by injecting error sequences at element level. One may also combine black-box testing with

specification mining to approximate WPPs [35]. Finally, simulation testing is required to validate DAS and RVE models, as already discussed.

Using MoSAFE Results in a Safety Case

MoSAFE results can be used as evidence in a DAS safety case. A DAS safety case is a claim supported by a structured argument founded on evidence that justifies the DAS is considered to be safe in its operation [36]. We illustrate the use of MoSAFE results in a safety case template for an ADS. A safety case template provides a reusable argument structure, and our example is based on ISCaP [13]. Whereas ISCaP focuses on assuring perception components, the template in this paper generalizes to any type of components and it also includes vehicle and system-level arguments.

The safety case template is organized into modules, reflecting the different levels of system decomposition (Fig. 10). The top module, **M-Vehicle**, represents the vehicle level. The subsequent level focuses on the vehicle systems; although Fig. 10 only displays **M-ADS-System** for the ADS, other modules at this level would encompass the remaining vehicle systems, including power train, braking, and electrical-power supply. The following level addresses the integration of components into their respective systems, exemplified by **M-D-integration**, which pertains to the integration of the object detector *D* into the ADS. Lastly, the component-level deals with the safety requirements at the component level, such as **M-D-component** for the object detector. The vehicle level is governed by the SOTIF Clause 6 and the HLSM analysis in MoSAFE. The system and integration levels are covered by the SOTIF Clause 7 and the DSM analysis in MoSAFE. Finally, the component-level for AI components is addressed by ISO/PAS 8800 [21].

The contents of the modules in the scope of MoSAFE are detailed in Figs. 11-13. This template adopts a deductive reasoning approach for deconstructing the top-level claim, while confining any inductive reasoning processes, such as extrapolating from test results, to terminal claims that are substantiated directly by evidence [37]. The structure of the argument is represented using Goal Structuring Notation (GSN) [38] (see the legend in Fig. 11 for a

summary). In this notation, each goal node symbolizes a claim, and the strategy nodes elucidate how a higher-level goal is derived from its underlying supportive goals. Assumptions and justifications are articulated through oval-shaped nodes. Context nodes supply relevant contextual information, which includes additional documents and specifications necessary for interpreting the goals or strategies. Lastly, solution nodes cite references to the supporting evidence.

The content of the vehicle-level module, as illustrated in Fig. 11, centers on the top-level goal (**G-V-vehicle**), which claims that the overall crash rate attributable to the SV behavior within the ODD adheres to the upper limit λ_v as a SOTIF acceptance criterion. This criterion, λ_v is established using principles such as GAMAB, ALARP, and MEM as per [1] and [39]. It is then distributed across the individual HBSCs constituting the ODD: $\lambda_{\text{HBSC, res}} + \sum_i \lambda_i \leq \lambda_v$. The black dot on the 'is-supported-by' arrow denotes a multiplicity of HBSC_{*i*}. Assuming these HBSC_{*i*} are mutually exclusive, this strategy can be understood mathematically as an application of the law of total probability: $P(\text{Crash}) = P(\text{Crash} \wedge \neg v_i \text{HBSC}_i) + \sum_i P(\text{Crash} \wedge \text{HBSC}_i)$. The residual risk, $\lambda_{\text{HBSC, res}}$, accounts for any overlooked crash-relevant HBSCs, and established scenario catalogs (e.g., [40, 41]) can be used to estimate this residual risk, as indicated by the solution node **Sn-HBSC-Res**. The crash rate for a specific HBSC_{*i*}, i.e., $P(\text{Crash} \wedge \text{HBSC}_i)$, is then estimated via its corresponding HBP_{*i*}, which fully encapsulates the hazardous behavior of the SV under HBSC_{*i*}. Thus, we require that, under HBSC_{*i*}, HBP_{*i*} is necessary for a crash:

$\text{HBSC}_i \Rightarrow (\text{Crash} \Rightarrow \text{HBP}_i)$. This necessity claim, **G-HBP_{*i*}-N**, can be established using the HLSM analysis in MoSAFE, indicated by **Sn-HBP_{*i*}-N**. In our running example, the RVE in Fig. 2 and $\mathbb{P}_{a, 50..3}$ in Table 3 represent sample HBSC_{*i*} and HBP_{*i*}, respectively, assuming UBI as the only HB for this scenario. This necessity claim can be equivalently stated as $P(\text{Crash} \wedge \text{HBSC}_i) = P(\text{Crash} \wedge \text{HBP}_i \wedge \text{HBSC}_i)$. Consequently, the goal **G-HBSC_{*i*}** can be restated as **G-HS_{*i*}** (see Fig. 11), where HS_{*i*} refers to the hazardous scenario combining HBP_{*i*} and HBSC_{*i*}: $P(\text{Crash} \wedge \text{HBP}_i \wedge \text{HBSC}_i) \leq \lambda_i$. This joint probability is then decomposed using the chain rule: $P(\text{Crash} \wedge \text{HBP}_i \wedge \text{HBSC}_i) = P(\text{Crash} | \text{HBP}_i, \text{HBSC}_i) P(\text{HBP}_i | \text{HBSC}_i) P(\text{HBSC}_i)$. Establishing a tight bound on $P(\text{Crash} | \text{HBP}_i, \text{HBSC}_i)$ is difficult, but it can be conservatively bounded by $\lambda_{\text{CR}, i} = 1$. The occurrence rate of HBSC_{*i*}, i.e., $P(\text{HBSC}_i)$, can be estimated from field data. Lastly, the

FIGURE 10 Safety case modules structured by the system decomposition. MoSAFE addresses vehicle-level, system-level, and integration safety arguments.




FIGURE 11 Vehicle-level safety argument template (i.e., content of the module **M-Vehicle** from Figure 10) expressed in GSN

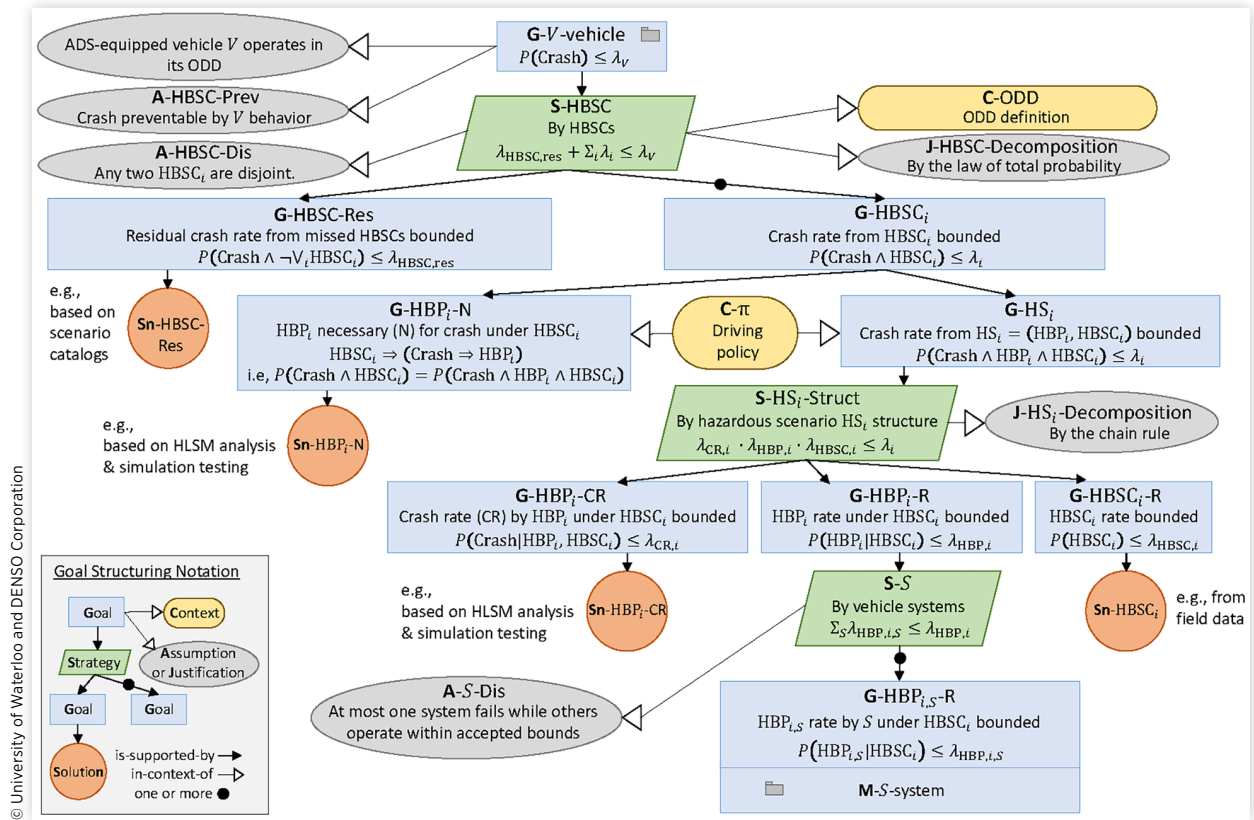
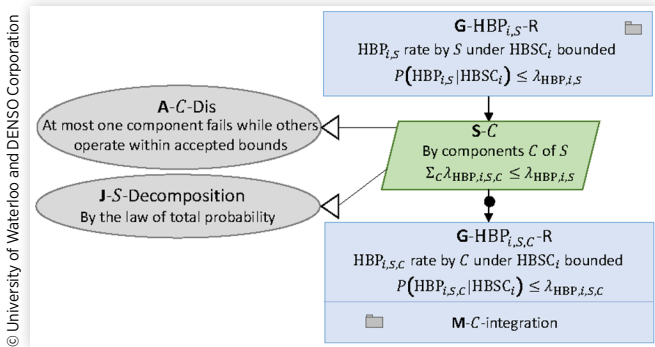


FIGURE 12 System-level safety argument template for system *S* (i.e., content of the module **M-ADS-system** from Figure 10, where *S*=ADS) expressed in GSN



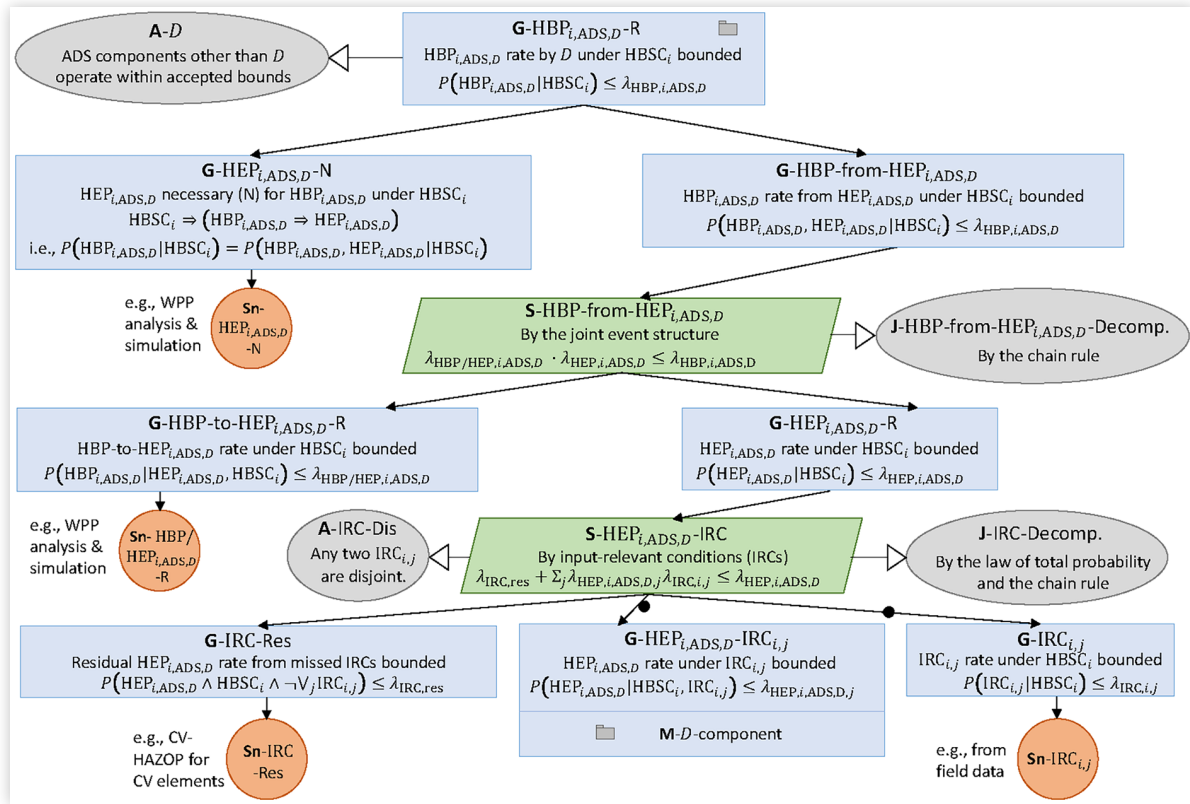
HBP_{*i*} occurrence rate under HBSC_{*i*} is distributed over the vehicle systems, resulting in a set of λ_{HBP_{*i*},*S*} in **G-HBP_{*i*},*S*-R**, with each λ_{HBP_{*i*},*S*} defining an HS_{*i*}-specific acceptance criterion for a given system *S*. The latter goal is then further decomposed in the module **M-S-system** over the components *C* of *S*, as illustrated in Fig. 12.

Figure 13 shows the content of the integration module for the object detector *D* as a component of the ADS. As a first step, $P(HBP_{i,ADS,D}|HBSC_i)$, which is the HBP_{*i*} rate due to *D* as part of the ADS under HBSC_{*i*}, is over-approximated by $P(HEP_{i,ADS,D}|HBSC_i)$, i.e., the rate of HEP_{*i*},_{ADS,*D*} under HBSC_{*i*}. This is done by showing HEP_{*i*},_{ADS,*D*} to be necessary for the HBP_{*i*},_{ADS,*D*} under HBSC_{*i*} using the WPP analysis of

the DSM for the HS_{*i*} and assuming λ_{HBP/HEP_{*i*},_{ADS,*D*} = 1. In our running example, the HEP_{*i*},_{ADS,*D*} for causing P_{*a*},_{50,3} as HBP_{*i*},_{ADS,*D*} under our sample scenario “approach to a stationary POV” is $\mathbb{P}_{a,50,3}$ from Tab. 5. Thus, the goal of limiting $P(HBP_{i,ADS,D}|HBSC_i)$ is realized by limiting $P(HEP_{i,ADS,D}|HBSC_i)$. This is supported by the WPP analysis of the DSM in Figs. 6 and 8 and the validation test results in simulation, as described earlier, and indicated by **Sn-HEP_{*i*},_{ADS,*D*}-N**. For our sample approach scenario and the experimental SV equipped with WISE-ADS, the analysis and simulations discussed earlier show that bounding the probability of $k_{contact} = 24$ or more FNs occurring within $n_{max} = 56$ consecutive frames will bound the probability of a collisions due to UBI. This probability $P(HEP_{i,ADS,D}|HBSC_i)$ is then distributed in **S-HEP_{*i*},_{ADS,*D*}-IRC** over IRC_{*ij*}, which are the IRCs identified based on sensor-relevant condition catalogs, such as CV-HAZOP [42]. The IRC_{*ij*} occurrence rate under a given HBSC_{*i*} is estimated from field data. Thus, the leaf goal for *D*, **G-HEP_{*i*},_{ADS,*D*}-IRC_{*ij*}**, is to limit the HEP_{*i*},_{ADS,*D*} rate in each HBSC_{*i*} and IRC_{*ij*}, which could be achieved using appropriate test data and other AI-component-level measures, as governed by ISO/PAS 8800 [21].}

The presented template makes several simplifying assumptions, which can be relaxed in a practical application. First, the top-level acceptance criterion does not consider crash severity; this can be addressed by repeating the decomposition for different severities, e.g., λ_{v,S1,3}, λ_{v,S2,3}, and λ_{v,S3}. Further, the template assumes operation within the ODD; however, an ADS needs to properly

FIGURE 13 Integration safety argument template for the object detector D (i.e., content of the module **M-D-integration** from Fig. 10) expressed in GSN



recognize and react to out-of-ODD situations, such as through a fallback. A complete safety case would also analyse these scenarios. Next, the template assumes HBSC $_i$ being mutually exclusive. This is not a limitation, since HBSCs can always be made disjoint, simply by designating overlaps as separate HBSCs. Further, the template assumes a complete HBP $_i$ for each HBSC $_i$; however, in practice, there would be a decomposition of HBPs per HBSC $_i$, such as UBI and UIB for the sample scenario “approach to a stationary POV.” This decomposition would be based on scenario catalogs, and it would also require a residual risk estimate due to potentially missed HBSCs. Further, the template states occurrence rates of HBSCs and IRCS as probabilities without specifying the meaning of a “trial.” In practice, these rates would be given as (i) frequencies per hour or kilometer driven for event-based conditions, such as approaches to a stopped vehicle or a stop sign; and (ii) percentages of hour or kilometer driven for conditions that can last arbitrarily long, such as vehicle following [10]. Ultimately, crash rates are given as frequencies per hour or kilometer driven or, equivalently, as a mean number of hours or kilometers driven between crashes. Additionally, the system decomposition in the vehicle- and system-level modules of the template assume mutually exclusive safety-relevant failures of systems and components, respectively, which is sufficient if their rate is kept low to satisfy the acceptance criteria. However, if the criteria satisfaction relies on redundant systems or components, their simultaneous failures need

to be modeled explicitly in the decomposition. Lastly, the template is designed for an ADS and thus does not deal with human driver behavior. For an ADAS, the interaction between the human driver of the SV and the system would need to be also considered.

MoSAFE Limitations

The HLSM analysis to derive HBPs at the vehicle level (Clause 6) is applicable to any DAS architecture and technology, but the WPP analysis of the DAS’ in the DSM at the element level (Clause 7) mainly targets systems that mix conventional software components with AI-based components. In particular, the analysis progresses backwards starting from an HBP at the DAS output and through the models of the conventional components up to the AI-based components. It stops at these components, deriving safety-related performance specifications on them as HEP occurrence rate limits. The analysis is still applicable to end-to-end optimizable systems that mix neural networks and differentiable variants of classical algorithms. However, the assume-guarantee reasoning espoused by the WPP derivation is not applicable to systems that contain mainly neural networks, even recently proposed architectures that strive to achieve modularity [43]. This is because some or all of the interfaces in such architectures are latent

representations, which are not human-interpretable, and thus it is difficult to impose specifications on them. Whereas some interpretable information can be extracted from them using suitable decoders, this decoding is incomplete and thus the flow of information and causality among the modules cannot be fully modeled. Future works should explore the development of effective modular reasoning approaches for such end-to-end AI systems.

Finally, MoSAFE analysis may face limitations in the context of complex scenarios and components, as already discussed. Both in our work and in related works, such as on RSS [24, 25], simple kinematic models and rule-based policies seem to be sufficient to produce adequate HLSMs for a wide variety of driving scenarios; however, complex road-user behaviors and complex vehicle dynamics, especially for high-performance collision avoidance, are likely to defy closed-form solutions and thus require sampling-based approaches in simulation. Similarly, complex component logic may make WPP analysis over a DSM intractable, in which case the assurance may need to be supplemented by automated black-box testing. Again, our work has demonstrated the feasibility of WPP analysis for a sample real-world system, but the limits of such analyses need to be explored in future work.

Related Work

Probably the most related work targeting SOTIF is by Vaicenavicius et al. [44], who present an analysis of an automated emergency braking scenario. Although the sample scenario is similar to ours, their work focuses on the statistical analysis of object detection errors that might cause a crash. However, it does not derive temporal specifications of error sequences, and it also acknowledges the toy nature of their illustrative example. In contrast, we have applied MoSAFE to four types of HBs (i.e., UBI, UIB, UA, and Failure to Yield (FTY)) and six different scenarios, and also validated our UBI model against the behavior of a real ADS in high-fidelity simulator.

Several behavioral safety frameworks, such as Responsibility-Sensitive Safety (RSS) [24], goal-oriented RSS [25], and Safety-Force Field [45], define safe behavior of an SV while formalizing reasonable behavior of other road users that capture common traffic rules as behavioral contracts and use them as assumptions. In particular, goal-oriented RSS formalizes intended SV behaviors (aka “proper responses”) in a range of traffic situations using Hoare quadruples [46] to allow sequential composition of behavioral contracts in traffic. These quadruples consist of behavior models with pre- and post-condition and invariant specifications, expressed using dFHL (differential Floyd-Hoare logic). MoSAFE has a different objective from these frameworks, namely to identify and evaluate HBPs, i.e., specifications of hazardous deviations from intended behavior (i.e., hazardous deviations from

proper response), and then use the identified HBPs to identify the corresponding HEPs in the DAS design. However, the frameworks provide the intended behavior as a starting point for HBP identification. Further, dFHL could potentially be used to represent the HLSMs, DSMs, HBPs, and HEPs and provide a formal basis to develop WPP derivation tools based on these models and specifications.

A related family of works applies automated black-box testing in simulation to identify hazardous behaviors of DAS (see [31] for a survey). These approaches use optimization methods, such as genetic algorithms [47], importance sampling [48, 49], and reinforcement learning [50], to find DAS inputs that cause HBs, or even estimate their probability (e.g., [48, 49]). As an example, Dreossi et al. [51] proposed to analyze the AI-based components and the conventional ones separately. They run the conventional part of the system with AI-based perception replaced by their intended behavior to determine range of low-dimensional inputs y for a given scenario, and then generate images that are consistent with y but cause misperceptions. Finally, they run a falsification tool to search among these images for a sequence to cause a crash. MoSAFE is fundamentally different from all these methods. It is model-based, rather than executing the DAS as a black-box. It focuses on establishing explicit human-interpretable specifications of hazardous error sequences on a module by module basis. Further, it aligns with the SOTIF standard requirements by separating vehicle-level analyses from element-level ones, to allow for reusable HBPs that are unaffected by DAS implementation. Finally, it establishes a detailed causal model of how hazardous error sequences propagate through the DAS in the form of a fault tree with temporal specifications as nodes. The use of models and specifications aids engineers in a deeper understanding of the system that what would be possible with a black-box technique.

The derivation of FTs from software using weakest-precondition reasoning has been explored before [52]. This related work proposes the use of a weakest-precondition calculus for the programming language at hand to derive software fault trees [53], where the faults could be defective program lines or random hardware errors corrupting program memory. In contrast to the usual approach of using weakest preconditions to characterize safe states or inputs, they are used to specify states or inputs that will result in a given fault. While this related work considers only singular faults as FT nodes, our FTs use temporal specifications as nodes. There are several temporal extensions of fault trees to capture temporal dependencies among events (see [17] for a comprehensive survey). They include dynamic [54] and temporal fault trees [55], which add different types of temporal gates. None of them consider nodes as temporal specifications, however. We are also not aware of the prior use of implications to represent over-approximation in fault trees.

Finally, in our prior work [13], we introduced the concept of hazardous misperception patterns, which can

be viewed as HEPs applied to perceptual components, and used them along with HBSCs and perception-only (PO) conditions (which correspond to quadrant 1 in Fig. 1) to propose a safety case template for assuring AI-based perception as part of a DAS. The template, expressed in the goal structuring notation, targets the development of an Integration Safety Case for Perception (ISCaP), which focuses on integrating safety requirements at the system level with perception-component performance requirements at the unit level. STEAM generalizes the concepts of misperceptions to AI errors and hazardous misperception patterns to hazardous error patterns. Further, MoSAFE is complementary to ISCaP. In particular, a template similar to ISCaP could be used to organize the results of MoSAFE into an integration safety argument.

Conclusion

The paper presented STEAM, a refinement of the SOTIF cause-and-effect model. STEAM adds the concept of hazardous error sequences, recognizing the fact that singular errors are often safe, and also adds HBPs and HEPs as a means to specify classes of hazardous behaviors and hazardous error sequences, respectively. Further, STEAM classifies scenario condition as HB-sensitive or insensitive and input-relevant or irrelevant, which aids the gradual refinement of scenario models for safety analysis, from HLSMs to DSMs.

Leveraging STEAM, MoSAFE helps identify HBPs and HEPs and evaluate their severity and likelihood. As part of Clause 6 analysis, an HLSM, which captures the HBSCs and the intended driving policy for the modeled scenario is developed and instrumented to inject HB sequences. The instrumented HLSM is then used to identify HBPs of different severities. The SOTIF acceptance targets can then be expressed as upper bounds on the occurrence rate of HBPs of different severities.

As part of Clause 7 analysis, a DSM and the HBPs are used to identify and evaluate HEPs. The DSM is developed by refining the HLSM with IRCs and the scenario-relevant DAS design and instrumenting it for injecting error sequences. The instrumented DSM is then used to identify HEPs that cause HBPs. The HEPs are derived from the HBPs as WPPs, and the causal links among them are captured in an FT with patterns as nodes. Conservative over-approximations of WPPs are applied as needed to simplify the analysis and captured in the FT as implication arrows. Using the FT, the SOTIF acceptance targets are then translated into upper bounds on the occurrence rate of the corresponding HBPs. This way, safety requirements on the performance of AI-based components are established as the upper bounds on the HEP occurrence rates.

The key benefit of MoSAFE is its modular and rigorous nature, establishing HEPs for each scenario-relevant element in the design and their causal links to HBPs in a systematic way. The WPP derivation is an instance of a

formal assume-guarantee reasoning and thus provides guarantees—under the assumptions made—that are not achievable with simulation testing. Although simulation testing is still needed to validate the models used in WPP derivation, it can be more targeted and thus more efficient than without model guidance. Safety engineers are already comfortable with the use of FTs, and the FTs capturing the causal links among the HBPs and HEPs can be used as part of a safety case. Finally, the use of models and specifications aids the engineers to acquire a deeper understanding of the system and its potentially hazardous behaviors than what would be possible with black-box testing in simulation.

MoSAFE is subject to several challenges and limitations, which point to exciting directions for future research. First, MoSAFE is subject to the same challenges of model creation and maintenance as any other model-based approach. There is an additional effort required to create and maintain models and the quality of the MoSAFE results is limited by the adequacy of models. This challenge creates opportunities to research and develop abstraction and slicing tools to help derive scenario-specific models from complete system designs or implementations or both. Additionally, the models need to be validated against implementations, which can be aided by automated black-box testing in simulation, including falsification tools to find counterexamples. Such tools, in combination with adversarial testing of AI-based components, could also be used to reduce the residual risk as part of SOTIF Clause 11. Second, WPP derivation can be challenging in the face of complex decision logic, and the WPPs themselves may become complex. Again, this creates an opportunity to research and develop tools to support WPP derivation. These could be adaptations of theorem provers as used in proving program correctness, especially for hybrid systems specifications, such as dFHL [25]. Another direction is to combine automated black-box testing, such as falsification, with temporal specification mining [35]. The WPP derivation should support over-approximation to balance complexity and precision. Furthermore, the tooling should also support WPP-based derivation of FTs. In particular, multi-failure analyses can be complex and will require adequate tool support. Further refinements of temporal FT notations may also be needed. Another opportunities for future research is to model and analyze a wider set of scenarios and HBPs at the vehicle level, potentially building on other safety frameworks such as goal-oriented RSS [25], which could lead to standard HLSMs and HBPs that are reusable across the DAS industry.

Finally, whereas MoSAFE currently targets DAS designs that mix conventional and AI-based components, future work should address the emerging “modular” end-to-end AI architectures (e.g., [43]). This means both improving the modularity of these architectures and developing effective assume-guarantee reasoning techniques for them. These techniques will likely be probabilistic and may include causal model learning.

References

- International Organization for Standardization, "ISO 21448: Road Vehicles – Safety of the Intended Functionality," 2021.
- SAE, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," J3016 202104, 2021.
- International Organization for Standardization, "ISO 26262: Road Vehicles – Functional Safety, 2nd edition, 2018.
- Salay, R., and Czarnecki, K., "Using Machine Learning Safely in Automotive Software: An Assessment and Adaptation of Software Process Requirements In ISO 26262," arXiv preprint arXiv:1808.01614, 2018.
- SAE, "Definitions for Terms Related to Automated Driving Systems Reference Architecture," J3131 202203, 2022.
- Ministry of Defence, "Defence Standard 00-58 HAZOP Studies on Items Containing Programmable Electronics," 2000.
- International Organization for Standardization, "ISO 34502: Road Vehicles – Engineering Framework and Process of Scenario-Based Safety Evaluation," 2022.
- Jurewicz, C., Sobhani, A., Woolley, J., Dutschke, J. et al., "Exploration of Vehicle Impact Speed—Injury Severity Relationships for Application in Safer Road Design," *Transportation Research Procedia* 14 (2016): 4247-4256.
- Maler, O., and Nickovic, D., "Monitoring Temporal Properties of Continuous Signals," in Lakhnech, Y. and Yovine, S. (eds.), *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems* (Berlin, Heidelberg: Springer), 152-166, 2004.
- SAE, "Considerations for ISO 26262 ASIL Hazard Classification," J2980 201804, 2018.
- Krampe, J. and Junge, M., "Injury Severity for Hazard & Risk Analyses: Calculation of ISO 26262 S-Parameter Values from Real-World Crash Data," *Accident Analysis & Prevention* 138 (2020).
- Bonnett, G.M., "Stiffness Coefficients—Energy and Damage," Tech. Rep., REC-TEC LLC, 2001, <http://www.rec-tec.com/Energy%20and%20Damage.html>.
- Salay, R., Czarnecki, K., Kuwajima, H., Yasuoka, H. et al., "The Missing Link: Developing a Safety Case for Perception Components in Automated Driving," *SAE International Journal of Advances and Current Practices in Mobility* 5 (2022): 567-579.
- International Organization for Standardization, "ISO 34503: Road Vehicles – Test Scenarios for Automated Driving Systems – Specification for Operational Design Domain," 2023.
- Dijkstra, E.W., "Guarded Commands, Nondeterminacy and Formal Derivation of Programs," *Commun. ACM* 18 (1975): 453-457.
- Pearl, J., *Causality: Models, Reasoning and Inference*, 2nd ed. (USA: Cambridge University Press, 2009).
- Ruijters, E. and Stoelinga, M., "Fault Tree Analysis: A Survey of the State-of-the-Art in Modeling, Analysis and Tools," *Computer Science Review* 15-16 (2015): 29-62.
- Czarnecki, K., "Operational World Model Ontology for Automated Driving Systems (Part 1 and 2)," 2018. Waterloo Intelligent Systems Engineering (WISE) Lab, Part 1: <http://dx.doi.org/10.13140/RG.2.2.15521.30568>, Part 2: <http://dx.doi.org/10.13140/RG.2.2.11327.00165>.
- de Gelder, E., Saberi, A.K., and Elrofai, H., "A Method for Scenario Risk Quantification for Automated Driving Systems," in *26th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, 2019.
- Atakishiyev, S., Salameh, M., Yao, H., and Goebel, R., "Explainable Artificial Intelligence for Autonomous Driving: A Comprehensive Overview and Field Guide for Future Research Directions," *CoRR* abs/2112.11561 (2021).
- International Organization for Standardization, "ISO/AWI PAS 8800: Road Vehicles – Safety and Artificial Intelligence," 2024.
- Salay, R., Czarnecki, K., Alvarez, I., Elli, M.S. et al., "PURSS: Towards Perceptual Uncertainty Aware Responsibility Sensitive Safety with ML," in *AAAI Workshop on Artificial Intelligence Safety (SafeAI)*, New York. CEUR, 2020.
- Kobayashi, T., Salay, R., Hasuo, I., Czarnecki, K. et al., "Robustifying Controller Specifications of Cyber-Physical Systems against Perceptual Uncertainty," in *NASA Formal Methods: 13th International Symposium, NFM 2021, Virtual Event, May 24–28, 2021, Proceedings* (Berlin, Heidelberg: Springer-Verlag, 2021), 198-213.
- Shalev-Shwartz, S., Shammah, S., and Shashua, A., "On a Formal Model of Safe and Scalable Self-Driving Cars," arXiv preprint: 1708.06374, 2018.
- Hasuo, I., Eberhart, C., Haydon, J., Dubut, J. et al., "Goal-Aware RSS for Complex Scenarios via Program Logic," *IEEE Transactions on Intelligent Vehicles* 8, no. 4 (2023): 3040-3072.
- NASA, "Fault Tree Handbook with Aerospace Fault Tree Handbook with Aerospace Applications Applications," version 1.1, 2002.
- Koopman, P., and Wagner, M., "Toward a Framework for Highly Automated Vehicle Safety Validation," 2018.
- Antkiewicz, M., Kahn, M., Ala, M., Czarnecki, K. et al., "Modes of Automated Driving System Scenario Testing: Experience Report and Recommendations," *SAE Int. J. Adv. & Curr. Prac. in Mobility* 2, no. 4 (2020): 2248-2266, <https://doi.org/10.4271/2020-01-1204>.
- Van Gennip, M., "Vehicle Dynamic Modelling and Parameter Identification for an Autonomous Vehicle," Master's thesis, University of Waterloo, 2018, <http://hdl.handle.net/10012/14260>.
- Hosking, B.A., "Modelling and Model Predictive Control of Power-Split Hybrid Powertrains for Self-Driving Vehicles," Master's thesis, University of Waterloo, 2018, <http://hdl.handle.net/10012/14094>.
- Corso, A., Moss, R., Koren, M., Lee, R. et al., "A Survey of Algorithms for Black-Box Safety Validation of Cyber-Physical Systems," *J. Artif. Int. Res.* 72 (2022): 377-428.
- Ost, J., Mannan, F., Thuerey, N., Knodt, J. et al., "Neural Scene Graphs for Dynamic Scenes," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2855-2864, 2021.

33. Jha, S., Banerjee, S., Tsai, T., Hari, S.K.S. et al., "ML-Based Fault Injection for Autonomous Vehicles: A Case for Bayesian Fault Injection," in *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 112-124, 2019.
34. Shalev-Shwartz, S., and Shashua, A., "On the Sample Complexity of End-to-end training vs. Semantic abstraction training," arXiv preprint arXiv:1604.06915, 2016.
35. Bartocci, E., Mateis, C., Nesterini, E., and Nickovic, D., "Survey on Mining Signal Temporal Logic Specifications," *Information and Computation* 289 (2022): 104957.
36. International Organization for Standardization, "ISO/AWI TS 5083: Road Vehicles – Safety for Automated Driving Systems – Design, Verification and Validation," 2024.
37. Rushby, J., "The Interpretation and Evaluation of Assurance Cases," Comp. Science Laboratory, SRI International, Tech. Rep. SRI-CSL-15-01, 2015.
38. SCSC Assurance Case Working Group, "Goal structuring notation community standard version 3," GSN Community Standard SCSC-141C, Safety-Critical Systems Club, CA, USA, 2021.
39. European Committee for Electrotechnical Standardization (CENELEC), "EN 50126: Railway Applications. The Specification and Demonstration of Reliability, Availability, Maintainability and Safety (RAMS) Generic RAMS Process, Part 1," 2017.
40. Najm, W.G., Smith, J.D., Yanagisawa, M., et al., "Pre-Crash Scenario Typology for Crash Avoidance Research," Tech. Rep., United States. National Highway Traffic Safety Administration, 2007.
41. Deepen AI and WMG University of Warwick, "Safety Pool," accessed January 6, 2024, <https://www.safetypool.ai>.
42. Zende, O., Murschitz, M., Humenberger, M., and Herzner, W., "CV-HAZOP: Introducing Test Data Validation for Computer Vision," in *Proceedings of the IEEE International Conference on Computer Vision*, 2066-2074, 2015.
43. Hu, Y., Yang, J., Chen, L., Li, K. et al., "Planning-Oriented Autonomous Driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
44. Vaicenavicius, J., Wiklund, T., Grigaitė, A., Kalkauskas, A. et al., "Self-Driving Car Safety Quantification via Component-Level Analysis," *SAE Intl. J CAV* 4 (2021): 35-45, <https://doi.org/10.4271/12-04-01-0004>.
45. Nistér, D., Lee, H.-L., Ng, J., and Wang, Y., "The Safety Force Field," Tech. Rep., NVIDIA, 2019.
46. de Boer, F.S., Hannemann, U., and de Roeper, W.P., "Hoare-Style Compositional Proof Systems for Reactive Shared Variable Concurrency," in Ramesh, S. and Sivakumar, G. (eds.), *Foundations of Software Technology and Theoretical Computer Science* (Berlin, Heidelberg: Springer), 267-283, 1997.
47. Ben Abdesslem, R., Nejati, S., Briand, L.C., and Stifter, T., "Testing Vision-Based Control Systems Using Learnable Evolutionary Algorithms," in *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*, 1016-1026, 2018.
48. Zhao, D., Lam, H., Peng, H., Bao, S. et al., "Accelerated Evaluation of Automated Vehicles Safety in Lane-Change Scenarios Based on Importance Sampling Techniques," *IEEE Transactions on Intelligent Transportation Systems* 18, no. 3 (2017): 595-607.
49. Sarkar, A., and Czamecki, K., "A Behavior Driven Approach for Sampling Rare Event Situations for Autonomous Vehicles," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 6407-6414, 2019.
50. Koren, M., Alsaif, S., Lee, R., and Kochenderfer, M.J., "Adaptive Stress Testing for Autonomous Vehicles," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, 1-7, 2018.
51. Dreossi, T., Donzé, A., and Seshia, S.A., "Compositional Falsification of Cyber-Physical Systems with Machine Learning Components," *J. Autom. Reason.* 63, no. 4 (2019): 1031-1053.
52. Clarke, S.J. and McDermid, J.A., "Software Fault-Trees and Weakest Preconditions - A Comparison and Analysis," *Software Engineering Journal* 8 (1993): 225-236.
53. Leveson, N.G. and Harvey, P.R., "Software Fault Tree Analysis," *Journal of Systems and Software* 3, no. 2 (1983): 173-181.
54. Dugan, J., Bavuso, S., and Boyd, M., "Dynamic Fault-Tree Models for Fault-Tolerant Computer Systems," *IEEE Transactions on Reliability* 41, no. 3 (1992): 363-377.
55. Palshikar, G.K., "Temporal Fault Trees," *Information and Software Technology* 44, no. 3 (2002): 137-150.

Acknowledgments

We acknowledge Dr. Rick Salay for his valuable discussions and fruitful feedback during the research development phase. Our thanks go to Maximilian Khan and Glen Tipold for their critical role in conducting the simulation experiments that validated the UBI model for the WISE ADS. We also extend our gratitude to Dr. Andrzej Wasowski and the anonymous reviewers for their constructive comments on an earlier draft of this manuscript.

Acronyms

- ADAS** - Advanced Driver Assistance System. 1, 4, 5, 17
- ADS** - Automated Driving System. 1, 4, 5, 14-16, 18
- AI** - Artificial Intelligence. 1, 2, 11, 12, 14-16, 18-20
- ALARP** - As Low as Reasonably Practicable. 16
- DAS** - Driving Automation System. 1-5, 8, 9, 12, 14-16, 18-20, 23
- dFHL** - differential Floyd-Hoare logic. 18-20
- DOF** - Degrees of Freedom. 15
- DSM** - Detailed Scenario Model. 8-10, 12-16, 18, 19, *Glossary: DSM*
- E/E** - electrical and/or electronic. 2
- FN** - False Negative. 1-4, 9-11, 13-16, *Glossary: FN*
- FP** - False Positive. 9, *Glossary: FP*

FT - Fault Tree. 1, 2, 11, 19, 20
FTA - Fault Tree Analysis. 14
FTY - Failure to Yield. 13, 18, *Glossary*: FTY
GAMAB - Globalement Au Moins Aussi Bon. 16
GSN - Goal Structuring Notation. 16–18
HB - Hazardous Behavior. 2–7, 9, 12, 13, 15–20, 23, *Glossary*: HB
HBB - Hazardous Braking. 13, 14, *Glossary*: HBB
HBP - Hazardous Behavior Pattern. 2–4, 8, 9, 12, 13, 17–20, 23, *Glossary*: HBP
HBSC - Hazardous-Behavior-Sensitive Scenario Condition. 4–6, 8, 12, 16, 17, 19, 23, *Glossary*: HBSC
HEP - Hazardous Error Pattern. 2–4, 8–12, 14, 15, 18, 19, 23, *Glossary*: HEP
HES - Hazardous Error Sequence. 3, 4, 23, *Glossary*: HES
HLSM - High-Level Scenario Model. 5–10, 12–16, 18–20, 23, *Glossary*: HLSM
IRC - Input-Relevant Scenario Condition. 4, 8, 9, 11, 12, 16, 17, 19, 23, *Glossary*: IRC
ISCaP - Integration Safety Case for Perception. 16, 19
MAIS - Maximum Abbreviated Injury Scale. 7
MEM - Minimal Endogenous Mortality. 16
MoSAFE - Model-based SOTIF Analysis of Failures and Errors. 1, 2, 4–6, 8, 9, 12, 14–20, *Glossary*: MoSAFE
ODD - Operational Design Domain. 2, 4, 5, 8, 12, 16
ODE - Ordinary Differential Equation. 6
POV - Principle Other Vehicle. 5–8, 12, 15–17, *Glossary*: POV
RSS - Responsibility-Sensitive Safety. 18, 20
RVE - Road-and-Vehicle Environment. 4–6, 8, 16, 23, *Glossary*: RVE
SCM - Structural Causal Model. 11
SOTIF - Safety of the Intended Functionality. 1–5, 12, 16, 18–20, 22, 23
STEAM - SOTIF Temporal Error and Failure Model. 1–4, 19, *Glossary*: STEAM
SV - Subject Vehicle. 2–9, 13–18, 23, *Glossary*: SV
TP - True Positive. 9, *Glossary*: TP
UA - Unintended Acceleration. 4, 5, 13, 18, *Glossary*: UA
UBI - Unintended Braking Interruption. 4–18, 23, *Glossary*: UBI
UHB - Unintended Hard Braking. 5, 13, 23, *Glossary*: UHB

UIB - Unintended Insufficient Braking. 4, 5, 13, 14, 17, 18, 23, *Glossary*: UIB
WPP - Weakest Precondition Pattern. 10–12, 14–16, 18–20, *Glossary*: WPP

Glossary

DSM - refinement of the HLSM to include IRCs and DAS elements for HEP identification. 8
FN - a foreground object misclassified as background. 1, 9
FP - background misclassified as a foreground object. 9
FTY - an HB representing failure to yield when required otherwise. 13, 18
HB - DAS behavior at the vehicle-level that may lead to harm. 2
HBB - UBI or UIB or UHB. 13
HBP - a class of HBs. 2, 4
HBSC - a scenario condition that links HB to harm. 4
HEP - a class of hazardous error sequences (or its over-approximation). 2, 4
HES - a sequence of errors that causes an HB. 3
HLSM - model of the driving policy and the RVE with HBSCs for HBP identification and evaluation. 5
IRC - scenario conditions that affect DAS inputs. 8
MoSAFE - a model-based method for identification and evaluation of HBPs and HEPs. 1
POV - the main other vehicle interacting with the SV in a scenario. 5
RVE - environment of the DAS, consisting of the SV and its road environment. 4, 5
STEAM - a refinement of SOTIF cause-and-effect model with error patterns and scenario-condition categories. 1
SV - a vehicle controlled by the DAS under analysis. 2
TP - correct detection. 9
UA - acceleration by the SV when the intended behavior is to maintain speed. 4, 13
UBI - lack of braking when the intended behavior is to brake. 4
UHB - hard braking when the intended behavior is no braking or braking at lower intensity. 13
UIB - braking at lower intensity than intended. 4
WPP - HEP consisting of all HESs causing another HEP or HBP. 10