



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/231219/>

Version: Accepted Version

Proceedings Paper:

Shahbeigi Roudposhti, Sepeedeh, Prama, Nawshin Mannan, Hodge, Victoria J. et al. (2025) Robustness requirement coverage using a Situation Coverage Approach for Vision-based AI Systems. In: 33rd International Conference on Requirement Engineering (RE'2025): Requirement Engineering for Trustworthy Artificial Intelligence (RETRAI) Workshop. 33rd International Conference on Requirement Engineering (RE'2025), 01 Sep 2025 International Requirements Engineering Conference Workshops (REW). , ESP, pp. 504-507.

<https://doi.org/10.1109/REW66121.2025.00075>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Robustness Requirement Coverage using a Situation Coverage Approach for Vision-based AI Systems

1st Sepeedeh Shahbeigi
Department of Computer Science
University of York
York, United Kingdom
Sepeedeh.shahbeigi@york.ac.uk

2nd Nawshin Mannan Proma
Department of Computer Science
University of York
York, United Kingdom
nawshinmannan.proma@york.ac.uk

3rd Victoria Hodge
Department of Computer Science
University of York
York, United Kingdom
victoria.hodge@york.ac.uk

4th Richard Hawkins
Department of Computer Science
University of York
York, United Kingdom
richard.hawkins@york.ac.uk

5th Boda Li
WMG
University of Warwick
Coventry, United Kingdom
Boda.Li.4@warwick.ac.uk

6th Valentina Donzella
School of Engineering and Materials Science
Queen Mary University of London
London, United Kingdom
v.donzella@qmul.ac.uk

Abstract—AI-based robots and vehicles are expected to operate safely in complex and dynamic environments, even in the presence of component degradation. In such systems, perception relies on sensors such as cameras to capture environmental data, which is then processed by AI models to support decision-making. However, degradation in sensor performance directly impacts input data quality and can impair AI inference. Specifying safety requirements for all possible sensor degradation scenarios leads to unmanageable complexity and inevitable gaps. In this position paper, we present a novel framework that integrates camera noise factor identification with situation coverage analysis to systematically elicit robustness-related safety requirements for AI-based perception systems. We focus specifically on camera degradation in the automotive domain. Building on an existing framework for identifying degradation modes, we propose involving domain, sensor, and safety experts, and incorporating Operational Design Domain specifications to extend the degradation model by incorporating noise factors relevant to AI performance. Situation coverage analysis is then applied to identify representative operational contexts. This work marks an initial step toward integrating noise factor analysis and situational coverage to support principled formulation and completeness assessment of robustness requirements for camera-based AI perception.

Index Terms—Requirements elicitation, robustness requirements, noise factors identification, situation coverage analysis

I. INTRODUCTION

AI components are increasingly used in safety-critical systems such as autonomous vehicles and robots, which must operate reliably under diverse and potentially degraded conditions. However, AI-based perception, particularly when reliant on vision sensors, is highly sensitive to such variations [1]. This makes their safety assurance particularly challenging.

Traditional safety engineering mitigates risk by identifying hazards and deriving explicit safety requirements. In AI systems, especially those involving perception, this approach becomes intractable due to the open-ended nature of the input space and their data-driven performance under noise and environmental stressors [2].

To address this, we propose a structured method that constrains the infinite input space and systematically defines variations in the operational context (i.e., the set of environmental and system conditions under which the AI component operates). These contextual variations constitute the basis for robustness requirements on the AI component. Building on prior work that classifies camera noise factors using a P-Diagram-based framework [3], we extend this foundation to support systematic robustness derivation. We adopt the situation coverage framework [4], not to generate test cases, but to formally represent the space of degraded operating conditions and guide robustness requirements elicitation.

Our motivation is twofold: (1) situation coverage provides a principled and extensible way to model sensor degradation factors and their combinations; and (2) this structured representation enables systematic identification of the performance boundaries within which AI-based perception must remain safe. To our knowledge, this is the first approach to integrate noise factor identification with situation coverage to derive robustness safety requirements in perception under degradation.

While existing frameworks [5]–[7] recognise the importance of robustness, they do not specify how to identify the full range of degraded conditions in a tractable, systematic way. Situation coverage addresses this gap by enabling comprehensive consideration of environmental variations and their interactions. Although previously applied to system-level testing [8], it has not yet been used for structured robustness requirement elicitation in AI-based perception systems.

II. AUTOMOTIVE CAMERA NOISE FACTORS IDENTIFICATION

Camera-based perception systems are highly sensitive to both internal and external sources of variation, which can degrade image quality and impair the performance of AI models. To support systematic identification and analysis of such effects, Li et al. [3] proposed a Parameter Diagram (P-Diagram)

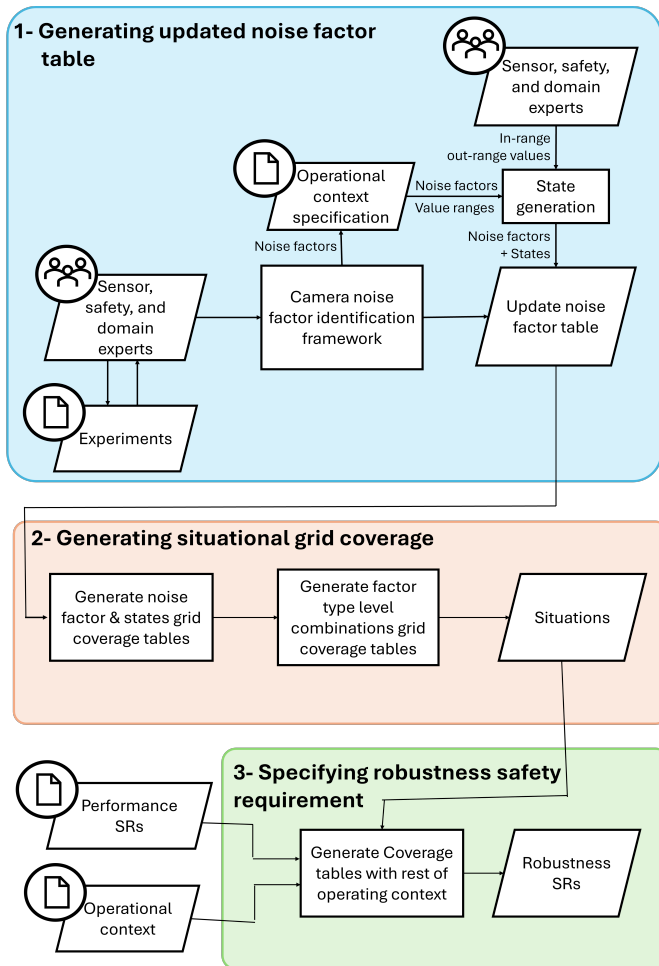


Fig. 1: Framework overview for noise-based requirement derivation.

approach to classify noise factors affecting automotive camera sensors. We adopt this method within our framework to derive traceable and structured safety requirements for perception under degradation.

The list of camera noise factors [3] was developed through practical experience and expert consultation. It captures a broad range of systematic or recurring sources of sensor degradation. Currently, it does not include corner cases such as surface reflections that result in object ghosting, which fall outside the scope of regular degradation patterns.

The P-Diagram approach categorises noise factors into five groups. The first category, **piece-to-piece** variation, refers to manufacturing or assembly inconsistencies, such as sensor–lens misalignment or variability in lens material properties. The second, **change over time**, includes degradation mechanisms like electronic ageing and mechanical loosening due to prolonged vibration. The third, **usage**, captures operational stressors such as scratches, sensor misplacement, or the accumulation of contaminants during vehicle operation or maintenance. The fourth, **environmental factors**, covers external conditions that can lead to camera sensor

degradation. Finally, **system interactions** refer to integration-related effects, such as electromagnetic interference, unstable power supply, or optical distortion introduced by the vehicle’s windshield curvature. Each identified factor is evaluated for its impact on key sensor outputs: frame rate (FR), pixel intensity (I_{RGB}), pixel position ($P_{(x,y)}$), and dropped frames (DF).

III. THE GRID COVERAGE-BASED FRAMEWORK

Our framework, illustrated in Fig. 1, comprises three high-level stages. Stage 1 constructs a noise factor table that includes relevant degradation factors, their discrete operational states, and inter-factor dependencies. This table focuses on factors that are relevant to the AI component within the specific application and task. For each factor, discrete states are defined based on the operating context specifications. Operating context refers to the environmental, spatial, and temporal conditions under which the system must operate safely.

Stage 2 generates situation coverage grids, enumerating all possible situations resulting from combinations of the noise factors defined in Stage 1. In this work, we adopt the term situation to refer specifically to aspects of the operating context related to camera degradation.

Stage 3 uses the situational grid table from Stage 2 to derive coverage tables that span the full set of operational context attributes. Each row in these coverage tables corresponds to a robustness safety requirement for the AI perception system.

A. Stage 1: Generating an Updated Noise Factor Table

Stage 1 focuses on identifying and organising the noise factors that affect camera sensors. We build on the taxonomy of thirty noise factors proposed by Li et al. [3], extending it to better suit safety-critical AI-based perception tasks. Our process begins with expert consultation to assess the relevance and severity of various degradation sources, taking into account the specific operational context and perception functions. As this is a proposed framework, expert input is used illustratively to highlight how this process would be conducted in practice. For instance, we assume consultation with sensor engineers and safety assessors to validate factor relevance. For assurance purposes, merely listing degradation sources is insufficient; it is essential to understand how each factor impacts the camera image data, as this directly affects AI performance. Therefore, this stage prioritises the analysis of image-level effects to identify degradation sources likely to impair perception.

To support this, we classify each noise factor based on its effect on image-level outputs, specifically whether it alters pixel intensity (I_{RGB}), pixel position ($P_{(x,y)}$), or both, following the methodology in [3]. Factors that only influence FR or DF are excluded, as they do not directly affect the spatial or visual content processed by the AI. In some cases, the effect of a factor may not be known *a priori* and may require targeted testing or simulation as illustrated by *experiments* block in Fig. 1.

We also capture dependencies between degradation factors in this stage. The dependencies might include relationships

TABLE I: Noise factor states and dependencies

Factor Type	ID / Noise Factor	State(s)	Dependency
Piece to Piece	1. Alignment	False (Normal), True (Misalignment)	—
	2. Fabrication Variability	False (Normal), Low, High	—
	3. Lens Shape, Purity	False (Normal), Low, High	—
	4. Dark Current Variability	False (Normal), True (High)	—
	5. Image Signal Processing (ISP)	False (Normal), True	—
Change over Time	1. Ageing of Electronics	False (Normal), True (High)	—
	2. Degradation of Lens	False (Normal), Low, High	—
	3. Vibration of Mounting	False (Normal), True	Co-occurrence with Misplacement
	4. Pollutant Ingress	False (Normal), Low, High	—
Usage	1. Misplacement of Sensor	False (Normal), True	—
	2. Vehicle Impact	False (Normal), True	Co-Occurrence with Misplacement
	3. Chemicals / Contaminants	False (Normal), True	—
	4. Obstructions	False (Normal), Low, High	—
	5. Lens Scratch	False (Normal), Low, High	Co-occurrence with Obstructions
	6. Vehicle Dynamic Settings	False (Normal), True	Co-occurrence with Misplacement
Environment	1. Sensor Saturation / Depletion	False (Normal), High	—
	2. Extreme Temperature	False (Normal), Low, High	—
	3. Low Illumination	False (Normal), High	—
System Interaction	1. Malicious Attacks	False (Normal), True	—
	2. Windshield Distortion	False (Normal), Low, High	—
	3. Power Supply	False (Normal), True	—
	4. Electromagnetic Interference (EMI)	False (Normal), True	—

such as co-occurrence or mutual exclusivity, as shown in the dependency column of Table I. These usually depend on domain assumptions (e.g., obstructions and lens scratches often co-occur) and are determined in close consultation with domain experts. Once relevant noise factors and their dependencies are identified, we define discrete operational states for each factor. These may be binary (e.g., *normal* vs. *degraded*) or categorical (e.g., *low*, *normal*, *high*), depending on the nature of the factor. For instance, lens scratches may be classified as: *normal* (visually insignificant), *low* (tolerable for AI perception), or *high* (impairing object detection). These classifications are informed by expert input, empirical data, and the operational context specifications. Capturing factor dependencies and discretised states at this stage is essential to constrain combinatorial complexity in the situational coverage analysis conducted in Stage 2.

B. Stage 2: Generating Situational Grid Coverage Table

Situational grid coverage [9] offers a structured method for evaluating autonomous systems by mapping real-world conditions, such as lighting, obstacles, and human presence, into a multi-dimensional grid. Each cell in the grid represents

a distinct subset of operating context, supporting systematic testing across a variety of environmental contexts to assess system safety and robustness. Unlike traditional coverage metrics, which focus on internal aspects such as code execution paths, situational coverage highlights external influences that affect system behaviour, particularly under complex or dynamic conditions [4], [10].

Stage 2 of our framework constructs situational coverage grids based on the noise factors, their discrete states, and known dependencies identified in Stage 1. The noise factor table (Table I) is organised hierarchically across three levels: *Type*, *Noise Factor*, and *State*. Situational coverage is applied at the most granular level by generating combinations of states within each type, while preserving logical consistency through dependency constraints. Incorporating consistencies helps to prune infeasible combinations, thereby reducing final grid sizes. For the five noise factor types, (1) piece-to-piece variation, (2) change over time, (3) usage, (4) environment, and (5) system interactions, we obtain pruned grid sizes of $n_1 = 72$, $n_2 = 36$, $n_3 = 108$, $n_4 = 12$, and $n_5 = 24$, respectively. These values exclude invalid or contradictory

state combinations, as reflected in Table I.

As an illustrative example, Table II presents the situational grid for the **Usage** type. This approach enables structured reasoning over a reduced but representative set of feasible situations, despite the theoretical infinitude of real-world scenarios. The resulting situational grid defines part of the operational contexts, associate with sensor degradation, in which robustness safety requirements will be derived, as described in Stage 3.

C. Stage 3: Specifying Robustness Safety Requirement

Assurance of Machine Learning (ML) for Autonomous Systems (AMLAS) [5] is a six-stage guidance framework for constructing a safety case for ML components. Stage two of AMLAS focuses on the development of ML safety requirements, which address both performance and robustness. Performance requirements define the expected functional behaviour of the ML model, while robustness requirements determine the range of operational contexts, including camera sensor degradation, under which these performance expectations must continue to hold.

We argue that performance and robustness aspects are inherently coupled and should not be treated separately. Instead, ML safety requirements should be formulated as performance requirements qualified by explicitly defined operational contexts. These operational contexts express the robustness dimension and are used to specify the range of conditions (e.g., sensor noise factors) within which the ML component must maintain safe performance. For instance, for a pedestrian detection task, we extend the performance safety requirement (RQ) in [11] to include an operational constraint as follows:

RQ: When the ego vehicle is 50 metres from the crossing, the object detection component shall identify pedestrians that are on or near the crossing in their correct position under all conditions defined in [PODs#1].

In this formulation, [POD#1] denotes a bounded operational context subset corresponding to the first row in the coverage table II, where the performance requirement must hold. For example, this represents an operational context without camera degradation.

We propose that robustness requirements are best conceptualised as these POD specifications, which collectively define the valid operating envelope for a given set of performance requirements. The union of all PODs should span the complete intended operational context.

Expressing safety requirements (SR) in the form *performance SRs + [PODs]*, where the performance requirements are explicitly qualified by the robustness context, enables targeted evaluation of ML models across well-defined slices of the Operational context. This approach enhances both the traceability and modularity of the resulting safety case, while supporting systematic robustness analysis.

IV. CONCLUSION AND FUTURE WORK

This paper proposes a structured framework for deriving robustness safety requirements for AI-based perception systems

TABLE II: Coverage grid for usage type factor

ID	Misplacement	Impact	Contaminants	Obstructions	Lens Scratch	Dynamics
1	False	False	False	False	False	False
2	False	False	False	False	False	True
3	False	False	False	False	Low	False
...
108	True	True	True	High	High	True

affected by camera degradation in automotive settings. By combining noise factor identification with situation coverage analysis, the approach produces performance requirements tied to bounded operational contexts (*SRs + [PODs]*), improving traceability and evaluation. Future work will address the scalability challenge by developing methods to prioritise the most critical PODs for efficient testing and validation.

ACKNOWLEDGMENT

This work was supported by the Centre for Assuring Autonomy, a partnership between Lloyd’s Register Foundation and the University of York (<https://www.york.ac.uk/assuring-autonomy/>)

REFERENCES

- [1] P. H. Chan, G. Dhadyalla, and V. Donzella, “A framework to analyze noise factors of automotive perception sensors,” *IEEE Sensors Letters*, vol. 4, no. 6, pp. 1–4, 2020.
- [2] J. Molloy, S. Shahbeigi, and J. A. McDermid, “Hazard and safety analysis of machine-learning-based perception capabilities in autonomous vehicles,” *Computer*, vol. 57, no. 11, pp. 60–70, 2024.
- [3] B. Li, P. H. Chan, G. Baris, M. D. Higgins, and V. Donzella, “Analysis of automotive camera sensor noise factors and impact on object detection,” *IEEE Sensors Journal*, 2022.
- [4] N. M. Proma, V. J. Hodge, and R. Alexander, “SCALOFT: An Initial Approach for Situation Coverage-Based Safety Analysis of an Autonomous Aerial Drone in a Mine Environment,” in *Accepted for, 44th International Conference on Computer Safety, Reliability and Security (safecomp 2025)*, 2025. <https://arxiv.org/abs/2505.20969>.
- [5] R. Hawkins, C. Paterson, C. Picardi, Y. Jia, R. Calinescu, and I. Habli, “Guidance on the assurance of machine learning in autonomous systems (amlas),” *arXiv preprint arXiv:2102.01564*, 2021.
- [6] International Organization for Standardization, “ISO 21448:2022 Road vehicles – Safety of the intended functionality.” <https://www.iso.org/standard/70939.html>, 2022. ISO 21448:2022.
- [7] S. Burton, L. Gauerhof, and C. Heinzemann, “Making the case for safety of machine learning in highly automated driving,” in *Computer Safety, Reliability, and Security, Trento, Italy, September 12, 2017, Proceedings 36*, pp. 5–16, Springer, 2017.
- [8] N. Proma, V. Hodge, and R. Alexander, “Situation coverage based safety analysis of an autonomous aerial drone in a mine environment,” in *The Yorkshire Innovation in Science and Engineering Conference (YISEC)*, York UK, June 2024.
- [9] R. Alexander, H. R. Hawkins, and A. J. Rae, “Situation coverage—a coverage criterion for testing autonomous robots,” tech. rep., Department of Computer Science, University of York, 2015.
- [10] N. M. Proma and R. Alexander, “Systematic situation coverage versus random situation coverage for safety testing in an autonomous car simulation,” in *Procs of the 12th Latin-American Symposium on Dependable and Secure Computing, LADC ’23*, p. 208–213, 2023.
- [11] L. Gauerhof, R. Hawkins, C. Picardi, C. Paterson, Y. Hagiwara, and I. Habli, “Assuring the safety of machine learning for pedestrian detection at crossings,” in *Computer Safety, Reliability, and Security: 39th International Conference, SAFECOMP 2020, Lisbon, Portugal, September 16–18, 2020, Proceedings 39*, pp. 197–212, Springer, 2020.