An HMI Personalization Concept for Increasing Safety of the Intended Functionality

1st Swagat Sukumar Sahoo
Scania
Södertälje, Sweden
swagat.sukumar.sahoo@scania.com

2nd Barbara Gallina
Mälardalen University
Västerås, Sweden
barbara.gallina@mdu.se

3rd Soheila Sheikh Bahaei
Scania
Södertälje, Sweden
soheila.sheikh.bahaei@scania.com

***Abstract*—Many safety-relevant vehicle functions, both ADAS (Advanced Driver Assistance Systems) and non-ADAS, depend on timely and correct driver responses to system information or warnings. When Human Machine Interface (HMI) assumptions are underspecified, Safety Of The Intended Functionality (SOTIF) may be undermined by delayed reactions, nuisance-driven, mistrust, or misuse. We use the Moving Off Information System (MOIS, UN R159) as a case study and we propose a guardrailed HMI personalisation concept that adapts presentation, not function, to driver variability across users and over time. We define a practical, measurable notion of non-intrusive optical information and map it to evaluation hooks and evidence artifacts. The concept is coupled with attention-aware safeguards and functional safety timing budgets consistent with ISO 26262:2018 and ISO 21448:2022. We outline evaluation hooks and evidence mapping that link SOTIF claims to measurable HMI outcomes without processing personal data. Although developed for MOIS, the concept is intended to extend to other driver-in-the-loop functions beyond ADAS.**

***Index Terms*—Safety Of The Intended Functionality (SOTIF), Human–Machine Interface (HMI), Moving Off Information System (MOIS), UN R159, Personalisation, Functional Safety.**

I. Introduction

Driver-in-the-loop vehicle functions covering both ADAS and non-ADAS deliver their intended safety benefit only if drivers can notice, interpret, and respond to system cues in time. When HMI assumptions are under-specified (e.g., attention demand, workload, shared-vehicle use, sensory variability, and in some cases temporary cognitive or mental conditions that affect attention or perception), Safety of the Intended Functionality (SOTIF) risks may arise even without component faults: drivers may react late, lose trust due to nuisance, or deactivate the feature. This is particularly critical in low-speed, close-proximity situations where response margins are small and the driver must manage multiple external cues.

In this paper, we propose a regulatory-guardrailed HMI personalisation concept for driver-in-the-loop functions: only presentation adapts (how information is shown), while regulated behaviour remains unchanged (when/where the function operates and what it must do). Using the UNECE Moving Off Information System (MOIS, UN R159 [1]) as a case study, we contribute (i) an explicit separation between fixed regulatory guardrails and tunable presentation parameters, (ii) a practical, measurable definition of "non-intrusive" optical information expressed through verification-friendly criteria, and (iii) an evidence strategy that couples attention-aware safeguards and timing considerations consistent with ISO 26262:2018 [2] and ISO 21448:2022 [3], mapping SOTIF-relevant claims to evaluation hooks and evidence artifacts without processing personal data. The rest of this paper is organized as follows. In Section II, we recall essential background and related work. In Section III, we propose our SOTIF-oriented HMI personalisation concept and its verification hooks. In Section IV, we discuss limitations and generalisation. In Section V, we conclude and outline future work.

II. Background and related work

A. MOIS context (UN R159): The Moving Off Information System (MOIS) addresses close-range conflicts with vulnerable road users when a heavy vehicle starts to move off or travels straight at very low speed. UN R159 specifies the operational constraints, the monitored zone, and the required driver signalling including availability/unavailability behaviour. In these manoeuvres the driver's attention is often split across mirrors, traffic lights, surrounding traffic, and immediate vehicle control, leaving limited time to perceive and interpret the information signal. As a result, the timing and placement of the driver-facing cue become safety-relevant even when the underlying detection function behaves as specified.

B. SOTIF perspective. SOTIF [3] emphasizes that hazardous scenarios can arise from functional insufficiencies (specification or performance insufficiencies). Even when no component is faulty, safety can be reduced if the intended functionality does not work safely for real users and real operating contexts. For driver-in-the-loop functions, this implies that assurance should cover not only sensing performance but also whether the driver interface mitigates forseeable misuse (unintended use as well as human factors (HF)) by supporting timely and correct interpretation and response under realistic conditions. As suggested by Gallina [4], considerations regarding mental harm are appropriate for the SOTIF of driver-in-the-loop functions.

C. HMI, attention, and misuse effects: Prior research on warning and information systems [5], [6] shows that underspecified HMI behavior can lead to misunderstanding, delayed response, or nuisance-driven disuse. When cues are perceived as unreliable or annoying, drivers may ignore them or deactivate the feature, undermining the intended safety

benefit. This motivates measurable HMI outcomes such as noticeability without excessive distraction, rapid understanding, and low nuisance, and the need to connect such outcomes to SOTIF reasoning.

D. Human variability and data constraints: In fleet operation, vehicles are commonly shared across multiple drivers, and driver capabilities may vary and change over time (e.g., sensory variability, workload, or temporary cognitive conditions affecting attention). If these differences are not handled explicitly, a compliant function can still become less effective due to delayed perception, misinterpretation, or avoidance behaviours. At the same time, any adaptation of the HMI must remain compatible with regulatory constraints and privacy expectations. A pragmatic direction is therefore to rely on non-identifying, coarse inputs and local decision logic, so that evaluation and assurance can be performed without processing personal identifiers.

E. Related work. In [7], the authors explore the evolution of automotive HMIs and the integration of advanced technologies such as AI (Artificial Intelligence), AR (Augmented Reality) and gesture-based controls. They highlight effects of proper HMI design in minimizing cognitive load considering safety and security challenges while maintaining functionality. They stress that personalization enhances the user experience to ensure both functional safety and user comfort.

In [8], the authors review and analyze adaptive human-machine interface as a solution for having an individual customer experience and not overwhelming the driver with growing amount of information. The paper consolidates findings from several studies on adaptive systems and inclusive design principles. The authors point out that personalization of HMI not only contributes to safety, but also enhances user experience and inclusivity in modern vehicles for all levels of automation. Investigation on standardized frameworks for implementing personalization in automotive systems are considered as future research to be planned.

In [9], the authors focus on user-adaptive in-vehicle Human-Machine Interfaces (HMIs) addressing diverse user requirements. It compares two HMI designs in automated vehicles one with permanent information presentation and the other offering context-adaptive information and evaluates perceived safety, trust and driver experience. The results indicate more positive ratings for context-adaptive HMI. In addition, it is highlighted that providing more information does not lead to better safety and user experience. Instead, presenting the right amount of information at the right time can enormously improve user confidence and safety and it is proposed to use adaptive HMIs and provide personalization for a better experience in automated driving systems.

In [10], authors introduce a methodology for safety analysis in human-automated vehicle interactions with focus on hazards arising from design and user behavior. The methodology consists of HMI specification and interaction sequence diagrams between human and Automated Driving Systems (ADS) and using a cause-consequence analysis in order to find the interaction faults leading to hazardous situations. Both ISO 26262:2018 safety standard and ISO 21448:2019 PAS – Safety of the Intended Functionality (SOTIF) are considered in this research and the aim is to take into account human factors expertise to have a safety analysis for the interaction between human and the ADS. The authors assert that using adaptive HMIs can improve safety by mitigating many identified risks and confusion during critical transitions can be reduced by personalization. It is further recommended to integrate safety analysis with personalization strategies while developing automated driving systems.

Our work goes beyond the summarised related works by providing a practical, auditable, privacy-preserving, SOTIF-oriented personalization concept for MOIS.

III. SOTIF PERSONALISATION CONCEPT

Here, we present a regulatory-guardrailed SOTIF personalisation concept for driver-in-the-loop functions, using MOIS (UN R159) as the case study. The guiding idea is that personalisation may adapt only presentation (how information is shown) to account for driver variability, while regulated MOIS behaviour remains unchanged (when/where the function operates and what it must do). The concept is structured around (i) fixed regulatory guardrails, (ii) a banded personalisation policy, and (iii) measurable verification hooks that link SOTIF-relevant claims to evidence without processing personal data.

A. Fixed guardrails (UN R159 is preserved, not adaptable) We treat the UN R159 requirements as non-negotiable guardrails that remain identical in all modes. Personalisation is explicitly prohibited from affecting operational conditions, detection boundaries, or state semantics. Concretely, the following remain fixed: 1) Activation and Operational Design Domain: MOIS operates only under UN R159 conditions (e.g., moving-off/straight-ahead at very low speed), including the speed limit (10 km/h) and illumination gate (15 lx); 2) Detection zone: the regulated close-range zone geometry and target classes are unchanged; personalisation cannot shift zone boundaries; 3) Availability and fail-safe behaviour: self-check, contamination/failure handling, and required unavailability indication follow UN R159 with no silent degradation; 4) Telltale semantics: meaning of information, warning (if present), and unavailable/failure states is unchanged; personalisation cannot alter semantics. 5) Principle: adapt how information is presented, not when/where MOIS operates or what it does.

B. Personalisation policy (banded inputs, no personal data required) The personalisation policy aims to reduce misunderstanding, delayed response, and nuisance-driven mistrust while keeping regulated behaviour unchanged. To avoid processing personal data, it uses only coarse, non-identifying bands (optionally enabled) and always supports a safe default. Inputs are limited to simple bands such as: sensory band (default/vision-reduced/hearing-reduced/unknown), workload band (low/normal/high based on an OEM proxy or test condition), seat-to-display distance band (short/medium/long for legibility scaling), and optional cabin noise band (low/high to avoid ineffective modalities). The outputs are restricted to presentation tuning, such as icon size within legibility limits,

Criterion	Target (practical)	How checked	Evidence artifact
Contrast (day)	Meets a clearly readable contrast target	Photometric measurement/calibrated display test	Contrast report (day)
Contrast (night/glare)	Meets a stronger contrast target under night/glare	Glare setup + photometric check	Contrast report (night/glare)
Luminance window	Not too dim to miss; not too bright to distract	Luminance sweep vs ambient	Luminance compliance sheet
Legibility geometry	Icon/text remains legible at typical driver distance	Visual-angle check using distance jig	Geometry/legibility report
Temporal behavior	Steady icon; controlled dwell; no flashing for info	HMI event log + video	Timing + behavior log
Display allocation	No urgent cue on center display; primary cue stays on cluster	Requirements check + HMI inspection	HMI allocation checklist
Understandability	Drivers understand meaning quickly	Confusion/interpretation test	HF usability report
Non-startle	Cue does not cause startle or annoyance	Startle/annoyance rating	HF startle assessment

TABLE I
PRACTICAL NON-INTRUSIVE CRITERIA FOR MOIS OPTICAL INFORMATION SIGNAL

contrast theme (day/night), placement within the permitted cluster area, and controlled dwell/persistence.

The modality rule is strict: the MOIS information signal remains optical-only (steady icon; no flashing). If a separate collision warning exists in the system, an additional modality envelope (e.g., haptic or an audio ramp) may be applied to the warning only, while keeping information and warning semantics clearly distinct. If bands are unknown or inconsistent, the system falls back to a conservative default presentation and never suppresses required information.

C. Practical definition of "non-intrusive" optical information (testable, not re-defining regulation) UN R159 constrains what MOIS shall do, but it does not fully specify a measurable definition of "non-intrusive" for the optical information signal. To support SOTIF evidence, we define practical engineering targets that make the HMI testable without changing regulated behaviour. The targets are used as verification hooks rather than additional requirements. Example targets include: contrast ratio $\geq 4.5:1$ (day) and $\geq 7:1$ (night/glare), luminance bounded to avoid being too dim to miss or too bright to distract (e.g., day 400–1200 cd/m²; night 1–12 cd/m²), legibility geometry based on minimum visual angle at typical viewing distances (e.g., stroke 7 arc-min and icon height 20 arc-min at 60–90 cm), and temporal behaviour consistent with information signalling (steady icon, controlled dwell 150–350 ms, no flashing, and no urgent cues on the central information display). These targets provide a measurable operational meaning of "non-intrusive" for evaluation and evidence building.

D. Attention-aware safeguards (avoid distraction & misuse) Because personalisation can unintentionally increase distraction or nuisance, the concept includes attention safeguards that must remain non-inferior to a non-personalised baseline. We use glance behaviour metrics to cap long off-road glances and to ensure total off-road time does not increase beyond an acceptable margin (example gate: total off-road time within +5 percentage of baseline, and no single off-road glance exceeding a predefined limit). In addition, startle and confusion checks are applied to avoid abrupt warning escalation and misinterpretation; for example, avoid sudden audio level jumps (if warning audio exists) and keep tone/icon confusion below a defined threshold across cohorts (example gate:less than 5% confusion rate). The intent is simple: personalisation must not trade noticeability for distraction or mistrust.

E. Timing budget and fail-safe behaviour (end-to-end) The concept makes timing and availability explicit to support driver-in-the-loop reaction. We define an information-path latency budget from first valid detection to visible information (example target: 200 ms median, 350 ms P95). We also require fail-safe unavailability indication to be asserted promptly when operation cannot be guaranteed (example target: within 1 s from degradation detection due to contamination, failure, or ODD gating), with no silent disable.

F. Evaluation hooks and evidence mapping (audit-ready, privacy-preserving) To avoid informal ("hand-wavy") arguments, each SOTIF-relevant claim is linked to (i) a measurable evaluation hook and (ii) a concrete evidence artifact. Evaluation combines simulator studies (e.g., age-stratified cohorts and dual-task conditions) with vehicle tests addressing zone coverage, illumination boundary conditions, and contamination-triggered unavailability. Measurements include photometric and geometry checks for non-intrusiveness and attention metrics via eye-tracking or equivalent instrumentation. Evidence artifacts are kept lightweight and auditable (e.g., measurement reports, log extracts, and test records). To preserve privacy, only band-level settings and aggregated outcomes are logged; no personal identifiers are stored.

IV. DISCUSSION

In this section, we discuss limitations and generalization of the proposed guard-railed HMI personalization concept. The key boundary condition is that personalization is limited to presentation parameters (e.g., icon size, contrast theme, dwell/persistence, and placement within the permitted display area; and, if a separate collision warning exists, its warning modality envelope). All UN R159 functional behavior remains unchanged. Personalization does not modify perception performance, object classification, decision logic, detection-zone geometry, activation conditions, or the Operational Design Domain (ODD). Thus, regulated guardrails such as zone geometry, speed and illumination thresholds, availability/unavailability behavior, contamination/self-check handling, and tell-tale semantics remain fixed; only how the information is shown may vary.

The concept assumes a driver-in-the-loop setting in which the system provides correct and timely information and the driver can perceive and respond as intended. It does not infer medical fitness to drive and does not require personal identifiers. When banding is used, it relies on coarse, non-identifying categories, a safe default presentation, and user opt-out. These assumptions help keep the assurance argument compatible with privacy expectations while still enabling measurable evaluation of HMI outcomes.

TABLE II
CLAIM → METRIC → ACCEPTANCE → EVIDENCE MAPPING

SOTIF claim	What we measure	Acceptance gate (example)	Evidence source
Information is noticed without distraction	Eye glance behavior	No single off-road glance exceeds limit; total off-road time non-inferior to baseline	Eye-tracking report
Information is understood	Interpretation / confusion	Confusion rate below threshold across cohorts	HF comprehension study
Information is non-intrusive	Startle/annoyance + subjective workload	Startle score below threshold; workload not increased	HF startle + workload report
No false trust when unavailable	Unavailable indicator behavior	Unavailable is visible and consistent when operation is not guaranteed	Test report + HMI log
Timing supports meaningful reaction	End-to-end latency distribution	Latency meets defined percentile targets	Latency histogram report
Personalisation does not change function	Guardrail compliance check	Zone/speed/availability behavior unchanged across modes	Requirements + test evidence
Privacy preserved	Data handling	Only banded profiles; no identifiers stored	Privacy design note

There are three main limitations. First, coarse banding is intentionally simple and cannot represent all individual differences, so conservative defaults, stability rules, and fallbacks are needed when band information is missing or contradictory. Second, simulator evidence may not fully transfer to on-road operation, where workload, stress, and risk perception differ; therefore, controlled field validation (e.g., track trials or limited fleet pilots) is needed. Third, personalization can unintentionally create nuisance or inconsistency; repeated non-helpful cues may reduce trust and increase delayed reactions or manual deactivation. Therefore, the evidence strategy should include attention-related metrics and acceptance indicators, not only technical compliance checks.

Although MOIS is the case study, the same structure can generalize to other driver-in-the-loop functions in both ADAS and non-ADAS domains. The key reusable element is the separation of fixed guardrails (regulatory/functional constraints) from tunable presentation, supported by measurable evidence that the presentation remains non-intrusive, understandable, attention-compatible, and timely.

Early practitioner feedback from commercial vehicle drivers (truck/bus) indicates that HMI effectiveness depends not only on correct function, but also on how information is presented under real workload. Drivers value clear and timely cues, but repeated non-helpful or overly prominent cues can reduce trust and increase annoyance. This supports the need for guardrailed personalization of presentation while keeping regulated behavior unchanged.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented a regulatory-guardrailed HMI personalisation concept for MOIS, in which only presentation adapts while all regulated behaviour remains unchanged. First, we made explicit which elements are fixed as regulatory guardrails and which presentation parameters may be tuned to account for driver variability. Second, we provided a measurement-oriented verification approach that links SOTIF-relevant claims to observable HMI outcomes and concrete evidence artifacts without processing personal data. Our concept is complemented by attention-aware safeguards and end-to-end timing considerations aligned with ISO 26262:2018 and ISO 21448:2022, to ensure that increased noticeability does not come at the cost of distraction, nuisance, or misuse.

As future work, we will refine and validate the evidence strategy through structured experimental studies and practical demonstrations. This includes confirming the proposed presentation bounds and attention safeguards with representative driver cohorts, and complementing simulator-based results with limited field validation to assess transferability under real operating conditions. In addition, we plan to investigate the exploitation of an ontology-based representation of claims, measures, and artifacts to strengthen traceability and reuse across iterations [11], [12]. Finally, the approach will be applied to additional driver-in-the-loop functions beyond MOIS to evaluate generalisability under different regulatory and operational constraints.

REFERENCES

[1] UNECE, "*UN Regulation No 159 – uniform provisions concerning the approval of motor vehicles with regard to the Moving Off Information System for the Detection of Pedestrians and Cyclists [2021/829]*," 2021.

[2] International Organization for Standardization, "ISO 26262:2018 Road vehicles – Functional safety," 2018.

[3] International Organization for Standardization, "ISO 21448:2022 Road Vehicles – Safety of the intended functionality," 2022.

[4] B. Gallina, "Safety of the Intended Functionality: What about Mental Harm?," in *CARS@EDCC2024 Workshop - Critical Automotive applications: Robustness & Safety, Leuven, Belgium. ⟨hal-04558509⟩*, 2024.

[5] H. Eom and S. H. Lee, "Mode confusion of human–machine interfaces for automated vehicles," *Journal of Computational Design and Engineering*, vol. 9, no. 5, pp. 1995–2009, 2022.

[6] J. Li, Y. He, S. Yin, and L. Liu, "Effects of automation transparency on trust: Evaluating HMI in the context of fully autonomous driving," in *Proceedings of the 15th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '23)*, pp. 1–16, ACM, 2023.

[7] I. Grobelna, D. Mailland, and M. Horwat, "Design of Automotive HMI: New Challenges in Enhancing User Experience, Safety, and Security," *Applied Sciences*, vol. 15, no. 10, p. 5572, 2025.

[8] L. Wirtz, E. Sever, and L. Eckstein, "Adaptive Human-Machine Interfaces and Inclusivity in the Automotive Field: A Review," *Accessibility, Assistive Technology and Digital Environments*, vol. 121, 2024.

[9] F. Hartwich, C. Hollander, D. Johannmeyer, and J. F. Krems, "Improving Passenger Experience and Trust in Automated Vehicles Through User-Adaptive HMIs: "The More the Better" Does Not Apply to Everyone," *Frontiers in Human Dynamics*, vol. 3, 2021.

[10] F. Warg, S. Ursing, M. Kaalhus, and R. Wiik, "Towards Safety Analysis of Interactions Between Human Users and Automated Driving Systems," in *Proceedings on Safety Analysis Methods for Automated Driving Systems*, RISE Research Institutes of Sweden, 2020.

[11] B. Gallina, M. Schweizer, and H. Dibowski, "Regulatory compliance-aware system change management via an ontology-based approach," in *Systems, Software and Services Process Improvement*, (Cham), pp. 245–259, Springer Nature Switzerland, 2026.

[12] T. B. Momcilovic, B. Gallina, I. Kessler, J. Hendricks, and D. Balta, "Ontogsn: An ontology-based framework for semantic management and extension of assurance cases," 2025.