Redundancy-Aware Multi-Sensor Fusion for Resilient Perception in Intelligent Vehicles

Tianqi Chu *

Department of mechanical engineering, Northeastern University, Beijing, China

* Corresponding Author Email: chu.tianq@northeastern.edu

Abstract. Reliable perception in intelligent vehicles rests on redundant sensing and fault-tolerant data fusion rather than any single modality. This study explains how redundancy is realized in practice across cameras, LiDAR, radar, and GNSS/IMU, and organizes fusion techniques from raw-level integration to feature- and decision-level reasoning. Emphasis is placed on reliability awareness—detecting, isolating, and mitigating sensor faults—so that the fused estimate degrades gracefully instead of catastrophically. Quantitative evidence from recent evaluations anchors the discussion: learning-based health monitoring on real vehicle bus signals achieves near-ceiling fault-detection accuracy with usable fault identification/isolation; fault-aware vision fusion consistently recovers mean intersection-over-union by roughly 0.08–0.22 under defocus and field-of-view truncation; and localization pipelines that adapt exclusion thresholds with a deep model reduce extreme position-error tails when GNSS measurements become unreliable. Persisting obstacles include uncertainty miscalibration under distribution shift, limited explainability of deep fusion, compute–latency constraints on embedded platforms, and the absence of fault-rich, standardized benchmarks for certification. The paper closes by outlining practical directions—self-auditing fusion that reports sensor contributions and uncertainty, context-aware reliability modeling, resource-conscious implementations, and datasets that "fail on purpose"—with the goal of enabling certifiable, fail-operational perception that maintains situational awareness in the adverse and ambiguous conditions where safety matters most.

Keywords: Redundant Sensing, Multi-Sensor Fusion, Fault Detection and Isolation, Robust Localization, Explainable Autonomy.

1. Introduction

In intelligent and automated driving, perception constitutes the operational core of the stack rather than a peripheral add-on. Whenever visibility deteriorates, objects are partially occluded, or transducers drift over time, single-modality pipelines tend to fail at precisely the moments when reliability is most needed. Redundancy and data fusion counter this fragility by arranging for alternate coverage, by cross-checking disagreeing measurements, and by enabling graceful degradation when one channel becomes unreliable. Recent empirical results give quantitative support to this claim. A unified vehicle-health monitor trained on real bus signals—accelerator-pedal position, steering-wheel angle, and brake pressure—reports an overall fault-detection accuracy close to 99.84%, while identification and isolation rates span roughly seventy-three to one hundred percent depending on fault class, with drift emerging as the hardest case to type [1]. In vision-centric setups, on-line fusion equipped with a diagnosis-and-avoidance gate consistently recovers mean intersection-over-union by around eight to twenty-two percentage points when camera streams are deliberately defocused or their fields of view are truncated [2]. Localization benefits as well: an α-Rényi-divergence residual whose threshold is adapted by a deep network preferentially removes harmful measurements and reduces extreme position-error bursts within a GNSS/INS/odometry stack [3]. Guided by these outcomes, this paper reviews how redundancy is realized in practice across cameras, LiDAR, radar, and GNSS/IMU; organizes data fusion techniques with attention to reliability; interprets the most instructive quantitative results; and sketches prospects for self-auditing fusion that can meet fail-operational requirements without exceeding embedded compute budgets. In addition to the empirical evidence, there is a conceptual reason to prefer fusion over isolated sensing. Perception is not a single variable but a coupled set of hypotheses about geometry, semantics and motion, and those hypotheses are constrained by physics. Combining modalities allows these constraints to be enforced across channels: inertial increments explain short-term motion even when visual texture disappears; radar Doppler validates whether a visually detected target is stationary or moving; LiDAR geometry anchors scale for monocular cameras. When the system reasons jointly, it can maintain continuity through brief dropouts and it can expose contradictions that are often the first symptom of a fault. The following sections therefore treat redundancy and fusion not as an added feature but as the organizing principle behind a robust perception stack.

2. Redundant Sensing in Intelligent Vehicles

Redundancy in production-minded perception mixes hardware overlap with analytical cross-checking so that the system can keep operating even when one input stream degrades. Multiple cameras with overlapping fields of view, front and rear radars with partially shared coverage, and both roof-mounted and bumper-level LiDARs serve as physical insurance against blind spots and single-point failures. Analytical redundancy then reconciles heterogeneous evidence against motion priors and kinematic constraints. If a channel becomes unreliable, the fused estimate does not collapse; the system either down-weights the suspect signal or excludes it until self-diagnosis clears the fault. The approach reflects a simple engineering truth: no single modality dominates across the full operational design domain, hence resilience must emerge from informed combination rather than from perfect sensors. Cameras bring dense texture and semantic cues and remain indispensable for recognition and scene understanding, yet their dependence on illumination makes them fragile in glare, precipitation, or low-light conditions. LiDAR contributes precise three-dimensional geometry that stabilizes depth estimates and supports mapping, but rain and snow reduce effective range and airborne droplets scatter the signal; the sensor is also susceptible to occlusion by grime or by roof-rack accessories. Radar is far more tolerant to adverse weather and measures Doppler directly, making it valuable for tracking; however, multipath and ghost targets complicate association and its angular resolution lags behind the pixel-level detail available to cameras. GNSS anchors global pose and keeps the vehicle within a broader geospatial frame, yet it degrades in urban canyons and tunnels; IMUs offer short-term stability at the cost of drift, which must be corrected by absolute references. A recent survey canvassing more than one hundred publications maps this failure landscape in detail and converges on the same practical conclusion: redundancy is not a luxury but the default mechanism for bridging complementary strengths and weaknesses across modalities [4]. Because redundancy incurs hardware cost and power draw, the architecture also has to balance coverage against resources. Designers therefore exploit diversity rather than mere duplication. Spatial redundancy ensures that adjacent sensors can continue observing when part of the field is blocked; information redundancy, expressed through heterogeneous modalities measuring the same quantity, allows the system to compensate when a particular physics channel is unreliable; temporal redundancy uses history to validate current readings and to filter transient disturbances. Together these forms of redundancy support fail-operational behavior in which the vehicle continues to function safely even when faults are present. In practice, redundancy earns its keep only when combined with principled fusion, health monitoring, and clear policies for exclusion and recovery. The combination is what turns multiple sensors from parallel liabilities into a coordinated, robust observer that can tolerate realistic disturbance and still provide the estimates required by planning and control. Design practice has crystallized several principles for building such redundancy. Spatial placement is chosen so that overlapping coverage exists without creating perfect correlation; the goal is to ensure that an obstruction affecting one viewpoint is unlikely to affect all. Signal paths are routed to avoid shared failure points, and timing is disciplined by hardware clocks or software time bases so that the same physical event is observed in a comparable temporal window. Analytical redundancy demands careful calibration. Intrinsic and extrinsic camera parameters, LiDAR–camera rigid transforms, radar boresight and lever arms, and clock offsets must be estimated and monitored; small errors in calibration create large inconsistencies after projection, which fusion can mistakenly interpret as scene changes. Long-term operation adds further drift, so periodic self-calibration routines are scheduled when the vehicle is stationary or when landmarks with known geometry are visible. Functional safety standards such as ISO 26262 encourage this discipline by requiring that critical functions remain operational under single-point failures. In a perception context, that requirement translates into explicit fallbacks and into evidence that the system detects, announces and mitigates faults within bounded time. Redundancy helps satisfy the requirement, but it also introduces cost and complexity. Extra sensors consume power and generate heat; additional data increase bandwidth through the compute fabric; and added software creates more states to be tested. The architecture must therefore trade coverage against resources. Practical systems tend to privilege diversity: a wide camera may complement a telephoto unit; a short-range radar can be paired with a long-range module; a solid-state LiDAR can share the scene with a rotating unit. Diversity reduces the chance that a single disturbance knocks out every view and increases the probability that at least one channel remains informative. Finally, redundancy must be visible to the decision layers that follow perception. Planning and control reason over risk, not just point estimates. When the perception stack provides calibrated uncertainty and a clear account of which channels were trusted, the downstream modules can adapt their behaviour for instance by lowering speed or increasing following distance-so that the whole vehicle behaves consistently with the quality of its current observations.

3. Data Fusion Methods and Research Challenges

Fusion operates at several conceptual levels that differ in information content, synchronization demands, and certification posture. At the raw-data level, estimators combine unprocessed returns—LiDAR points with radar range–Doppler or GNSS positions with IMU increments—to produce precise state estimates when spatial and temporal alignment are tight. Kalman and Information filters, along with particle filters for non-Gaussian regimes, exemplify this layer; they provide elegant uncertainty propagation but become bandwidth and compute heavy on embedded platforms. Feature-level integration instead aligns intermediate representations, for instance camera feature pyramids with LiDAR bird's-eye-view voxels. This level has become dominant in three-dimensional detection because it strikes a reasonable accuracy–latency balance and is less brittle to slight miscalibration. Decision-level fusion aggregates per-modality hypotheses through probabilistic reasoning or evidential rules such as Dempster–Shafer, favoring modularity and facilitating safety arguments where traceability of decisions matters. Across levels, fault awareness is the differentiator between systems that merely concatenate data and systems that remain reliable under stress. In a detect–isolate–reconfigure pattern, anomalies are first flagged by residual analysis or learned detectors, then typed to the extent possible, and finally acted on by excluding or down-weighting the offending channel. Reliability-weighted integration assigns confidences per sensor and per frame so that healthy views dominate the consensus; context-adaptive thresholds, as demonstrated in localization, allow a learned policy to remove particularly harmful measurements rather than applying a static cutoff [1–3]. Quantitative results clarify the value of these patterns. In vehicle-health monitoring, learning-based detection on multi-sensor bus signals achieves accuracy near ninety-nine point eight four percent overall, while identification and isolation vary between roughly seventy-three and one hundred percent across classes; this mixture of high sensitivity with usable typing suggests the architecture can support both immediate mitigation and longer-horizon maintenance planning [1]. Perception under degraded sensing shows a similar trend. When camera streams are deliberately defocused or their fields of view are truncated, an on-line fusion network equipped with a diagnosis-and-avoidance gate consistently raises mean intersection-over-union by about eight to twenty-two percentage points; on a representative sequence, the IoU climbs from around 0.666 to 0.888 for truncation and from 0.801 to 0.882 for defocus, with comparable gains on additional videos [2]. Localization profits when fault handling is inside the loop. A deep network that adapts the threshold on an α-Rényi-divergence residual tends to discard measurements with substantially larger true errors—on average a little over eight meters with maxima near eighty—than those eliminated by a maximum-likelihood baseline whose discarded samples average closer to three and a half meters with maxima near fifteen; the asymmetry implies that the learned policy targets the harmful inputs and reduces tail risk in downstream position estimates [3]. Despite these advances, limitations remain. Uncertainty from learned fusion is often miscalibrated and therefore over-confident; distribution shift across weather, lens contamination, and rare targets remains under-represented in training and evaluation; black-box fusion complicates safety cases and public trust; and embedded compute budgets make some accurate raw-level approaches impractical at ASIL-D deadlines. Robust autonomy will require methods that expose their reasons, quantify and calibrate uncertainty, and meet hard timing without sacrificing coverage across the operational design domain. A closer look at canonical algorithms clarifies how uncertainty is represented and why certain choices interact well with redundancy. Extended and Unscented Kalman filters linearize nonlinear motion and observation models to maintain a Gaussian belief over state; the machinery exposes innovation residuals that double as health indicators, since large, sustained residuals often signal mis-calibration or sensor failure. Particle filters approximate arbitrary distributions at the cost of computational load and are therefore attractive when measurement likelihoods are multi-modal, as in urban canyons with GNSS multipath. For feature-level fusion, modern systems increasingly rely on shared intermediate spaces such as bird's-eye-view tensors onto which camera features are lifted via geometry and LiDAR points are voxelized. Transformer backbones and cross-attention then exchange context across modalities so that complementary cues reinforce one another. Decision-level fusion remains valuable wherever modularity and auditability dominate; information filters naturally combine estimates from independent sources, and Dempster–Shafer theory offers a way to represent ignorance separately from probability, which is useful when sensors provide conflicting evidence. Fault detection and isolation methods span classical and learned families. Parity-space checks and generalized likelihood ratio tests compare predicted and measured signals to flag anomalies; change-point detectors look for distributional shifts; autoencoders, one-class classifiers and contrastive models learn a compact description of normal behaviour and raise alarms when deviations exceed a bound. Isolation can be framed as a multi-class classification problem when labelled faults exist, or as a structured search where the system explains residuals by hypothesizing which channel failed. Reconfiguration closes the loop by excluding or down-weighting the suspect input and by reallocating computation to the remaining channels. The numerical results cited earlier provide concrete effect sizes to accompany these mechanisms. The bus-signal detector illustrates the value of modelling temporal dependence across correlated channels; its near-ceiling detection rate implies that the model leverages cross-channel consistency rather than relying on any single signal [1]. The vision experiments demonstrate that gains from fusion are largest exactly when a mechanism estimates health and acts on it. The change in IoU is not a small statistical nuisance; in practice it transforms borderline segmentation into stable object histories, which in turn stabilize planning. Localization with adaptive thresholds shows that reliability must be handled inside the estimator rather than as a post-hoc filter. Removing a few deceptive updates can reduce tail risk more than marginally improving average residuals, and that is precisely what certification cares about [3]. Against these benefits, several research challenges remain stubborn. Uncertainty calibration for deep fusion is an open problem: softmax scores tend to be over-confident, and naïve temperature scaling ignores multi-modal structure. Methods that tie uncertainty to physical consistency—such as enforcing cycle consistency between reprojections, or using self-supervised objectives that penalize disagreement across modalities—are promising yet not mature. Distribution shift raises parallel concerns. Weather, lens contamination and rare objects appear infrequently in training data; without targeted augmentation or explicit fault injection, learned systems overfit to clean conditions. Finally, explainability is not merely a social desideratum but a practical tool. Engineers must know why a fused estimate changed when a sensor degraded. Saliency on features, attribution to sensors, and causal probes that measure the effect of masking a channel can together form the basis of real-time audit trails that satisfy both debugging needs and regulatory expectations [5].

4. Results and Indicative Evidence

The quantitative record just summarized can be interpreted through the lens of reliability economics: redundancy pays only when the fusion stack measures trust and acts on it. In the bus-signal detector, the near-ceiling overall accuracy is meaningful partly because the system also provides class-specific identification and isolation rates that range from the low seventies to one hundred percent. The pattern indicates that subtle, slowly accumulating phenomena such as drift remain challenging to type, whereas abrupt defects are easier to spot and classify. That division suggests a practical split of duties in an engineering deployment. A conservative front-end can maintain high sensitivity and trigger safe fallback modes, while a slower background process performs richer diagnosis and schedules maintenance when evidence accumulates. The two-speed structure uses redundancy not only to detect faults but also to plan around them [6].

Table 1. Representative quantitative evidence for redundancy-aware fusion in intelligent vehicles

Domain/Task	Signals / Data	Metric	Value (and gain)	Key condition	Ref.
Vehicle health monitoring (FDI/FDII/HF)	Vehicle bus signals (accelerator pedal, steering angle, brake pressure)	Fault detection; ID/Isolation	Detection ≈ 99.84%; ID/Isolation 73–100% (drift hardest)	Learned multi-sensor detector with prognostics	1
Perception under injected camera faults	KITTI-style sequences (defocus; FoV truncation)	Mean IoU (segmentation)	0.801→0.882 (defocus); 0.666→0.888 (truncation); gains ≈ +0.08 / +0.22	Fusion w/ fault diagnosis & avoidance (FDA) gate	2
Localization with adaptive FDI	GNSS/INS/odometry	True error of removed measurements	DNN-only removed: mean 8.247 m, max 80.141 m vs. MLE-only removed: mean 3.353 m, max 15.347 m	DNN-adapted α-Rényi residual threshold	3

In **Table 1** we summarize three representative evaluations spanning detection, perception, and localization. The figures show that redundancy only yields safety margins when the fusion stack is reliability-aware—faults are detected, isolated, and reflected in weights or exclusions

Vision experiments with injected faults invite a different reading. Mean IoU gains in the range of eight to twenty-two percentage points are large enough to matter for downstream planning, because they convert borderline detections into stable tracks and reduce spurious obstacles that would otherwise lead to unnecessary braking. The observation that improvements persist across multiple sequences implies that the benefit is not tied to a particular scene but arises from reliability-weighted reallocation: when one view degrades, the network shifts trust to healthier channels instead of averaging in corrupted features. In practical systems, the same logic extends to cross-modal combinations, where radar or LiDAR can carry the scene when cameras fail and cameras can sharpen classification when geometry-centric sensors are ambiguous [7]. Localization tells yet another story. By adapting the exclusion threshold to context, the deep network effectively changes which parts of the measurement stream are believed at any given time. Measurements removed only by the learned threshold are associated with larger true position errors than those removed only by the maximum-likelihood baseline. The asymmetry is not accidental; it reveals that the model has learned to prefer discarding risky updates even if that choice slightly increases nominal variance, because the alternative is to admit a small number of disastrous outliers that dominate tail risk. In certification discussions this distinction matters, since it shows a concrete pathway by which learned policies can reduce worst-case behaviour without masking errors [8]. Taken together, the three strands of evidence justify a design stance. Redundancy should be coupled to detectors that estimate health, to fusion rules that express those estimates as weights or exclusions, and to monitoring that verifies tail behaviour rather than average metrics alone. Where such couplings exist, the data show gains in detection, perception, and localization that are large enough to translate into planning and control benefits. Where they do not, additional sensors merely increase bandwidth and cost without improving safety. These interpretations invite a few caveats. Reported numbers depend on datasets and experimental protocols, and care must be taken when extrapolating. The bus-signal study uses specific vehicle platforms and signal selections; different buses with different sampling rates or noise characteristics may alter absolute values. What appears robust is the relative result: models that combine signals and exploit their temporal structure detect more anomalies than single-signal baselines. In the vision setting, injected faults are a stand-in for real degradations. They approximate blur, occlusion and truncation but cannot capture every nuisance such as droplets or dirt, which produce structured artefacts. Nevertheless, the persistence of gains across several sequences reduces the likelihood that improvements arise from chance. For localization, the learned thresholding policy must be monitored so that it does not over-reject measurements and silently degrade availability; here audit trails that log reasons for exclusion are essential companions to the estimator. A second consideration concerns statistical power and effect sizes. Improvements in mean IoU of ten to twenty points often change the qualitative behaviour of a downstream planner because they move detections across operational thresholds. By contrast, small average changes in position error can still be safety-relevant if they reflect substantial mitigation of rare but large outliers. Interpreting results through the lens of decision thresholds clarifies which gains matter most to a vehicle on the road. A third consideration is transfer. Systems trained on one city or sensor set may not transfer cleanly to another. Redundancy can ease transfer by allowing the stack to rely on the modalities that remain stable across environments while discounting those that change. The practical value of reliability estimates becomes obvious in this setting: when confidence reflects scene conditions, the same fusion logic can operate sensibly even as absolute inputs shift. With these caveats in view, the overall message stands. Redundancy provides the raw material for resilience; reliability-aware fusion turns that material into gains that are visible in the metrics that matter for safety. Where the pipeline estimates health and takes decisions accordingly, the evidence shows robust improvements in detection, perception and localization; where it does not, additional sensors contribute little beyond cost and complexity.

5. Future Directions and Prospects

Looking forward, progress toward fail-operational perception will hinge on several complementary strands that can be advanced in parallel. The first concerns self-auditing fusion, in which each combined estimate carries a compact account of which sensors dominated the decision, which inputs were excluded and why, and how uncertainty aggregated. Real-time explainability of this kind is more than a research nicety; it creates artifacts that can be inspected during incident analysis and formal safety cases, and it allows engineers to diagnose failure modes without guessing [9]. The second strand is context-aware reliability modeling. Sensor confidences should depend on weather, illumination and traffic density, not just on instantaneous residual magnitudes. The success of adaptive thresholds within localization suggests a general recipe for coupling fault detection to environmental context so that exclusion becomes a principled response to measured risk rather than a static rule [10]. A third strand emphasizes resource awareness. Mid-level feature sharing, dynamic sensor selection, and adaptive numerical precision can meet hard timing requirements while preserving safety margins; without such co-design, otherwise accurate fusion models may remain too heavy for embedded deployment. A fourth strand is the construction of fault-rich benchmarks. Public datasets rarely fail on purpose, yet certification requires repeatable degradations—blur, occluders, interference, truncation and multipath—with ground-truthed exclusions so that fail-operational behaviour can be measured rather than asserted. Creating such suites will make research results comparable and will help close the gap between laboratory conditions and real roads. These technical programs should be accompanied by standardization of testing protocols and by ethical governance that clarifies accountability and transparency. Surveys from 2024–2025 argue that the lack of explainability in deep fusion remains a primary barrier to acceptance; attaching reasons to fused outputs and exposing calibrated uncertainty would directly address that concern [5]. Equally important is the integration of predictive maintenance and health forecasting so that the perception stack not only reacts to faults but also anticipates them. Evidence from bus-signal studies shows that learning-based prognostics can support maintenance scheduling in ways that reduce spurious alarms and avoid unplanned downtime [1]. If these strands advance together—self-auditing fusion, context-aware reliability, resource-conscious implementation, fault-rich evaluation, and predictive maintenance—the result will be perception systems that satisfy both engineering pragmatics and regulatory demands. Cooperative perception amplifies these themes. Vehicle-to-everything communication allows one platform to borrow sight from another and extends redundancy beyond the physical footprint of a single car. Sharing detections, intentions and environmental context can resolve occlusions and blind spots that no amount of on-board hardware can overcome. The same reliability questions reappear at network scale: messages must advertise uncertainty and provenance; fusion must discount stale or untrusted data; and cybersecurity must be treated as part of perception rather than a separate concern. Edge computing provides the architectural counterpart by pushing light-weight fusion to sensors and microcontrollers while reserving global aggregation for more capable units. Hybrid designs reduce bandwidth without abandoning the benefits of a consolidated world model. Another near-term opportunity lies in deliberate data generation. Fault-rich benchmarks should encode repeatable degradations with ground-truth exclusions so that fail-operational behaviour can be measured. Synthetic pipelines can create controlled perturbations—fog, rain, turbidity, occluders, multipath and spoofing—and can blend them with real sequences to approximate rare events. When evaluation protocols reward tail-risk reduction rather than only average accuracy, research will naturally move toward the behaviours that most affect safety. Finally, deployment will benefit from simple contracts between perception and the rest of the stack. If a fused estimate is accompanied by calibrated uncertainty, by a compact explanation of sensor contributions, and by flags that summarize recent health trends, planners can adapt aggressiveness and speed profiles to current observability. Such contracts turn perception from a brittle source of numbers into a negotiated interface that expresses both content and quality. Over time, the result should be vehicles that drive not only accurately but also cautiously when the world becomes difficult. A related research lane pursues metrics and tooling for explainability that are specific to fusion. Generic saliency often highlights textures rather than the modal contributions that matter for safety. What is needed are per-sensor attributions and compact rationales that tell an engineer, and eventually a regulator, why radar outweighed camera on a particular frame, or why GNSS updates were excluded during a localization spike. Lightweight Shapley-style proxies and counterfactual probes—masking a channel and measuring the delta in the fused estimate—can supply such evidence without halting real-time execution. On the uncertainty front, calibration should be evaluated at the object, trajectory and scene levels rather than only per-pixel or per-detection. Metrics that penalize over-confidence in the tails encourage models that align stated confidence with empirical error, which is the quantity planners actually consume. Validation infrastructure will need to grow accordingly. Hardware-in-the-loop testbeds that replay synchronized multi-modal logs through real sensors and compute stacks can expose timing issues and bus saturation that do not appear in offline experiments. Scenario libraries that combine naturalistic traffic with scripted degradations make it possible to reproduce edge cases, and coverage tools can track which combinations of weather, lighting, speed and sensor health have been tested. Finally, governance must keep pace with engineering. Clear responsibilities for audit trails, retention of fusion rationales, and disclosure of known failure modes will accelerate acceptance and allow independent assessment. In short, the next phase is as much about making reliable systems inspectable as it is about making them accurate.

6. Conclusion

Redundancy without a notion of reliability spreads errors across channels, while reliability without fusion leaves large blind spots. The quantitative record now supports a middle course in which redundancy-aware fusion detects, isolates and responds to faults, yielding measurable safety margins in detection, perception and localization. The examples considered here show how accuracy improves when camera faults are injected yet handled by diagnosis-and-avoidance gates, and how position-error tails contract when adaptive thresholds remove harmful GNSS measurements. Outstanding issues-uncertainty calibration, explainability, and compute ceilings on embedded platforms—should be treated as design constraints rather than afterthoughts. By pairing self-auditing fusion with context-aware reliability models and resource-conscious implementations, and by evaluating these systems on benchmarks that encode realistic failures, the field can move toward certifiable, fail-operational perception suitable for complex, dynamic environments. The practical message is straightforward: additional sensors add value only when their reliability is explicitly measured and expressed in the fusion logic that informs planning and control. The review also exposes limits that should temper expectations. Fusion cannot create information that was never observed; it can only recombine and qualify what the sensors provide. When all channels degrade simultaneously-for instance in heavy snow at night-no amount of cleverness removes the need for graceful fallback policies that reduce speed or suspend automation. Likewise, explainability and calibration add overhead, and there will always be pressure to trade them away for nominal accuracy. The case made here is that the trade is false: systems that explain themselves and that quantify their uncertainty are precisely those that avoid catastrophic mistakes and that earn public trust. The practical path forward is incremental. Add reliability estimates where they are missing; attach concise reasons to major fusion decisions; and test on datasets that fail on purpose. Even modest progress on these fronts will compound into perception stacks that remain effective when conditions are worst, which is when safety matters most. In bringing these strands together, the paper argues for treating redundancy-aware fusion as the architectural center of the perception stack and for equipping it with the diagnostics, explanations and calibrated uncertainties that downstream modules require. The quantitative examples—high detection and usable isolation on vehicle buses, large perception gains under camera faults, and localization improvements from adaptive exclusion—serve as proof points that such a stance is both technically feasible and practically meaningful. What remains is sustained engineering to turn these patterns into defaults across platforms. If that happens, intelligent vehicles will not simply perceive more accurately on average; they will behave more predictably when the environment is at its worst, which is the standard by which safety-critical systems are ultimately judged.

References

[1] Safavi, S., Hamid, H., Fallah, S. Multi-Sensor Fault Detection, Identification, Isolation and Health Forecasting for Autonomous Vehicles. Sensors, 21 (7):2547, 2021.

[2] Pan, H., Sun, W., Gao, H., et al. Deep Learning Based Data Fusion for Sensor Fault Diagnosis and Tolerance in Autonomous Vehicles. Chinese Journal of Mechanical Engineering, 34 (6):115, 2021.

[3] Harbaoui, N., et al. Context-Adaptive Fault-Tolerant Multi-Sensor Fusion for Vehicle Localization. Journal of Intelligent & Robotic Systems, 2024.

[4] Matos, F., Bernardino, J., Durães, J., Cunha, J. A Survey on Sensor Failures in Autonomous Vehicles: Challenges and Solutions. Sensors, 24 (16):5108, 2024.

[5] Panduru, K., Walsh, J. Exploring the Unseen: A Survey of Multi-Sensor Fusion and the Role of XAI in Autonomous Vehicles. Sensors, 25 (3), 2025.

[6] Zhang, Y., Tu, C., Gao, K., Wang, L. Multisensor Information Fusion: Future of Environmental Perception in Intelligent Vehicles. Journal of Intelligent and Connected Vehicles, 2024.

[7] Omoseebi, A., Ali, A., Muthanna, A., et al. Predictive Maintenance of Autonomous Vehicles Based on Sensor Fault Detection. ResearchGate preprint, 2025.

[8] Romanelli, F. Multi-Sensor Fusion for Autonomous Resilient Perception: Classical vs. Deep Learning Techniques. Thesis, 2025.

[9] Nemra, A., et al. Fault Diagnosis and Reconfiguration for Mobile Robot Localization Based on Multi-Sensor Data Fusion. Unmanned Systems, 2022.

[10] Rodger, J. A. Reducing Failure Risk in Integrated Vehicle Health Maintenance: A Fuzzy Multi-Sensor Kalman Filter Approach. Expert Systems with Applications, 39 (10):9821–9836, 2012.