

Article

Not peer-reviewed version

Risk-Aware AI Architecture for BVLOS UAV Safety: Integrating Sensor Fusion and SATCOM

[Nick Barua](#) *

Posted Date: 13 April 2026

doi: 10.20944/preprints202604.0858.v1

Keywords: BVLOS UAV; multi-modal sensor fusion; SATCOM; risk-aware architecture; ISO 26262; functional safety; AI decision-making; RASA framework; estimator instability; Monte Carlo simulation



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Risk-Aware AI Architecture for BVLOS UAV Safety: Integrating Sensor Fusion and SATCOM

Nick Barua ^{1,2,3}

¹ Chairman & CEO, AN Holdings Co; s.nick.barua@gmail.com

² Shiga University of Medical Science

³ Kobe Gakuin University

Abstract

The proliferation of unmanned aerial vehicles (UAVs) in civil, commercial, and defence domains has exposed a critical architectural gap: existing platforms optimise either communication or perception independently, leaving safety coverage incomplete under simultaneous stress in Beyond Visual Line of Sight (BVLOS) operations. This paper proposes the Risk-Aware UAV Safety Architecture (RASA), a three-layer conceptual framework integrating multi-modal sensor fusion, satellite communication (SATCOM), and AI-driven risk modelling aligned with functional safety principles such as ISO 26262. The RASA framework quantifies operational risk as $R(t) = \alpha \cdot U_{\text{sensor}}(t) + \beta \cdot L_{\text{c_norm}}(t) + \gamma \cdot U_{\text{sensor}}(t) \cdot L_{\text{c_norm}}(t)$ — a function of normalised sensor uncertainty and normalised communication latency, with an interaction term capturing compound degradation effects — enabling onboard risk estimation without ground-in-the-loop dependency. Building on prior validated work in multi-modal sensor fusion for safety-critical human detection [10] and SATCOM communication architectures for UAV connectivity [15], this paper extends those contributions to the BVLOS domain. Monte Carlo simulations across three representative operational scenarios validate the risk model's behaviour and demonstrate that the interaction term produces steeper risk escalation under compound failure conditions compared to the linear baseline. This paper addresses the critical gap in BVLOS UAV safety architectures by integrating perception and communication reliability within a single, auditable, risk-aware framework.

Keywords: BVLOS UAV; multi-modal sensor fusion; SATCOM; risk-aware architecture; ISO 26262; functional safety; AI decision-making; RASA framework; estimator instability; Monte Carlo simulation

1. Introduction

Unmanned aerial vehicles (UAVs) are rapidly transitioning from controlled line-of-sight operations to fully autonomous Beyond Visual Line of Sight (BVLOS) missions, enabling applications in logistics, disaster response, infrastructure monitoring, and national security. This evolution introduces unprecedented operational capabilities alongside a critical safety engineering challenge: how to assure safety when no human operator has direct visual or immediate control authority over the vehicle.

A foundational requirement for safe BVLOS operations is the ability to extract meaningful, actionable information from complex, high-entropy signal environments. Advances in deep learning [1,2] — particularly the application of Mel-frequency cepstral coefficients (MFCC) with convolutional and recurrent neural network (CNN/RNN) architectures — have established robust baselines for classifying environmental data under high-noise conditions [3]. These signal-processing principles applied to UAV perception pipelines provide a critical underpinning for the multi-modal sensing architectures proposed in this work.

However, current UAV platforms overwhelmingly treat sensor fusion, communication management, and decision-making as isolated software modules with independent failure modes. This modular independence is architecturally unsafe in BVLOS scenarios for three reasons:

- Sensor degradation and communication degradation frequently co-occur (e.g., electromagnetic interference affecting both radar and RF links simultaneously);
- Independent failure-mode analyses underestimate compound risk when both subsystems are simultaneously stressed; and
- Regulatory frameworks (EASA, FAA, and JCAB – Japan Civil Aviation Bureau, which governs UAV airspace integration in Japan) increasingly require demonstrable safety integration across sensing and communication layers.

However, a unified architecture that integrates perception uncertainty with communication reliability for safety-critical BVLOS operations remains largely unexplored. This paper addresses this gap by proposing the Risk-Aware UAV Safety Architecture (RASA): a three-layer framework that formally integrates multi-modal sensor fusion, SATCOM-based communication, and AI-driven risk modelling into a unified, auditable safety system. The primary contributions are:

- A formalised three-layer RASA architecture for BVLOS UAV safety systems;
- A tractable risk model with interaction term: $R(t) = \alpha \cdot U_{\text{sensor}}(t) + \beta \cdot L_{\text{c_norm}}(t) + \gamma \cdot U_{\text{sensor}}(t) \cdot L_{\text{c_norm}}(t)$;
- Monte Carlo numerical validation across three representative BVLOS scenarios;
- Conceptual alignment with ISO 26262-derived functional safety principles; and
- Regulatory benchmarking against EASA, FAA, and JCAB BVLOS requirements.

2. Safety Challenges in BVLOS UAV Systems

2.1. Operational Risk Categories

BVLOS UAV operations expose vehicles to a compounding risk environment fundamentally distinct from piloted or line-of-sight scenarios [4,5]. The global growth of civil UAV applications [6] and the rapid advancement of UAV wireless communications [7,8] have elevated safety architecture design to a critical research priority. Table 1 categorises the primary risk domains and their operational consequences.

Table 1. Primary risk categories in BVLOS UAV operations.

Risk Category	Root Cause	Operational Consequence
Collision Risk	Limited obstacle awareness at range	Structural damage; third-party injury
Communication Loss	RF interference, range limits, link fade	Loss of command authority; flyaway
Sensor Failure	Environmental noise; hardware fault	False-negative detection; incorrect manoeuvre
Perception Latency	Processing bottleneck; bandwidth saturation	Delayed response to dynamic obstacles
Environmental Uncertainty	Weather, terrain masking, and dynamic airspace	Unpredicted mission deviation

2.2. The Communication Dependency Problem

Unlike line-of-sight operations, where a human pilot can visually override system decisions, BVLOS UAVs are entirely dependent on the integrity of their data link. Terrestrial RF solutions – 4G/LTE and legacy C-band links – suffer from geographic coverage gaps, signal shadowing in mountainous or urban terrain, and spectrum congestion in dense operational areas. These are structural characteristics that cannot be engineered away within the terrestrial domain, and a degraded link is a safety-critical event requiring an autonomous risk mitigation response.

Crucially, these failure modes are interdependent – perception uncertainty is dangerously amplified under degraded communication conditions, since the vehicle cannot relay sensor anomalies for remote assessment or receive corrective guidance when the data link is compromised. This coupling is explicitly modelled by the interaction term in the RASA risk function (Section 5.4).

2.3. Functional Safety Principles in UAV Contexts

While ISO 26262 was originally developed for road vehicles, its underlying principles – systematic hazard identification, safety integrity level classification, and fail-safe design – provide a structurally sound conceptual foundation for UAV safety architecture design [9]. Safety architectures integrating functional safety principles from adjacent domains are increasingly recognised as applicable to safety-critical autonomy. The RASA framework draws conceptual alignment from ISO 26262 and ISO 21448 (SOTIF) without claiming formal compliance with either standard.

3. Multi-Modal Sensor Fusion for Risk Perception

3.1. Sensor Modality Overview

Robust environmental perception in BVLOS conditions requires the integration of multiple sensing modalities, each contributing complementary information across different operating conditions. Table 2 summarises the primary modalities within the RASA Perception Layer with engineering-grade range estimates.

Table 2. Sensor modalities in the RASA Perception Layer.

Modality	Range	Environmental Sensitivity	Primary Role in RASA	Layer
Optical Camera	50–200 m	High (light-dependent)	Object classification	Perception
LiDAR	20–150 m	Moderate (rain, dust)	3D obstacle mapping	Perception
Radar	10–300 m	Low (all-weather)	Velocity/range	Perception
Acoustic (MFCC)	10–50 m	Low (dark, fog)	Anomaly detection	Perception
SATCOM Telemetry	Global	Low	Path/link validation	Communication

3.2. Cross-Domain Validation: From Automotive to Aerial Safety

Cross-domain application of sensor fusion techniques has a growing empirical basis. Prior work by the present author demonstrated that a quantitatively evaluated multi-modal sensor fusion system achieves robust detection of falling humans in vehicle-proximity safety contexts, achieving 98.2% true positive rate in daytime and 95.6% at night [10]. The architectural principles established in that automotive context — specifically, the use of heterogeneous sensor fusion to overcome the failure-mode blind spots of individual sensors — transfer directly to UAV obstacle avoidance and pedestrian detection in BVLOS operations. Independent evaluations of multi-modal sensor uncertainty propagation in robotic and UAV systems further validate this cross-domain transfer [24,25].

3.3. Acoustic Sensing and MFCC-Based Feature Extraction

Acoustic sensing provides an important additional modality in environments where optical sensing is compromised by low light, fog, or occlusion. MFCC-based feature extraction combined with CNN and RNN architectures provides robust environmental classification from acoustic signals under high-noise conditions [3]. Within RASA, acoustic sensing contributes specifically to anomaly detection in the Perception Layer, supplementing optical and radar inputs during sensor degradation events.

3.4. Uncertainty Quantification

A critical requirement for the RASA risk model is that the Perception Layer outputs not only classification decisions but also associated uncertainty estimates (U_{sensor}). These quantify the confidence of the fused perception output and serve as the primary input to the Decision Layer risk function.

Sensor uncertainty $U_{\text{sensor}}(t)$ is computed as the normalised trace of the posterior covariance matrix from the Bayesian fusion filter, bounded within $[0, 1]$ through probabilistic confidence estimation. This approach is consistent with established methods for uncertainty propagation in multi-sensor robotic systems [24,25] and provides a direct input to the risk model without additional transformation.

Unlike conventional fusion systems, RASA explicitly incorporates communication reliability into the perception loop. This bidirectional coupling — where link degradation raises the effective uncertainty weight in the risk model — is the architectural innovation distinguishing RASA from independent-subsystem UAV designs, and is formally captured by the interaction term $\gamma \cdot U_{\text{sensor}} \cdot L_{\text{c_norm}}$.

3.5. Comparison with Existing UAV Safety Architectures

Table 3 positions RASA against conventional UAV autonomy approaches across key safety architecture dimensions.

Table 3. Comparison of RASA against conventional UAV safety architectures.

Feature	Traditional UAV Stack	RASA
Sensor + communication coupling	✗ Independent modules	✓ Unified risk model
Real-time scalar risk output	✗ Not available	✓ $R(t)$ at each timestep
Compound failure modelling	✗ Linear superposition only	✓ Interaction term γ
Autonomous risk escalation	Limited / operator-dependent	✓ Explicit MRM hierarchy
Regulatory audit trail	Partial	✓ Auditable safety trace

Feature	Traditional UAV Stack	RASA
Onboard without ground loop	✗ Ground-dependent	✓ Fully autonomous

4. SATCOM as the Communication Backbone for BVLOS Operations

4.1. Limitations of Terrestrial Communication

Terrestrial cellular networks (4G LTE, 5G) and dedicated RF links present structural limitations for BVLOS UAV operations: geographic coverage is inherently incomplete beyond populated areas, altitude-dependent signal characteristics differ from ground-level deployment assumptions, and spectrum availability is contested in dense operational corridors [11,12]. These are boundary conditions that necessitate a satellite-layer solution.

4.2. SATCOM Architecture for UAV Connectivity

The emergence of Low Earth Orbit (LEO) constellations [13,14] has reduced round-trip latency from the 600+ ms characteristic of geostationary systems to sub-100 ms in operational deployments, making SATCOM a viable backbone for UAV command, control, and telemetry. SATCOM represents the primary communication link for long-range UAV systems, enabling continuous over-the-horizon connectivity not achievable by any terrestrial solution [15]. This directly supports integration with 3GPP NTN architectures (Release 17/18), which formally define satellite link integration into the 5G protocol stack for UAV command-and-control applications [16,26].

4.3. SATCOM in the RASA Framework and Latency Decomposition

Within the RASA Communication Layer, SATCOM serves three distinct functions: (1) continuous command and control uplink; (2) real-time telemetry downlink providing position, attitude, and system health data; and (3) latency monitoring feeding the normalised communication latency variable L_c into the Decision Layer risk model.

End-to-end communication latency is decomposed as:

$$L_c = L_{\text{uplink}} + L_{\text{processing}} + L_{\text{downlink}}$$

This decomposition enables direct measurement and attribution of latency contributors and supports integration with the Decision Layer risk model. L_{uplink} and L_{downlink} are determined by LEO orbital geometry and atmospheric propagation; $L_{\text{processing}}$ reflects onboard and ground-station computational overhead. When measured L_c exceeds predefined thresholds, the Decision Layer automatically escalates the risk state and triggers pre-defined minimum risk manoeuvres (MRMs). The maximum tolerable latency L_{max} is anchored to EASA C2 link requirements and 3GPP NTN latency budgets, with a recommended value of $L_{\text{max}} = 1000$ ms for the BVLOS mission class addressed in this work [16,17,26].

5. The RASA Framework: Proposed Architecture

5.1. Architectural Overview

The Risk-Aware UAV Safety Architecture (RASA) is organised as three interdependent layers, each with defined inputs, processing responsibilities, and output interfaces to adjacent layers. Figure 1 presents the complete architecture block diagram. The three layers are: (1) Perception Layer, responsible for multi-modal sensor fusion and uncertainty quantification; (2) Communication Layer, maintaining SATCOM connectivity and measuring link latency; and (3) Decision Layer, computing instantaneous risk $R(t)$ and executing proportionate minimum risk manoeuvres.

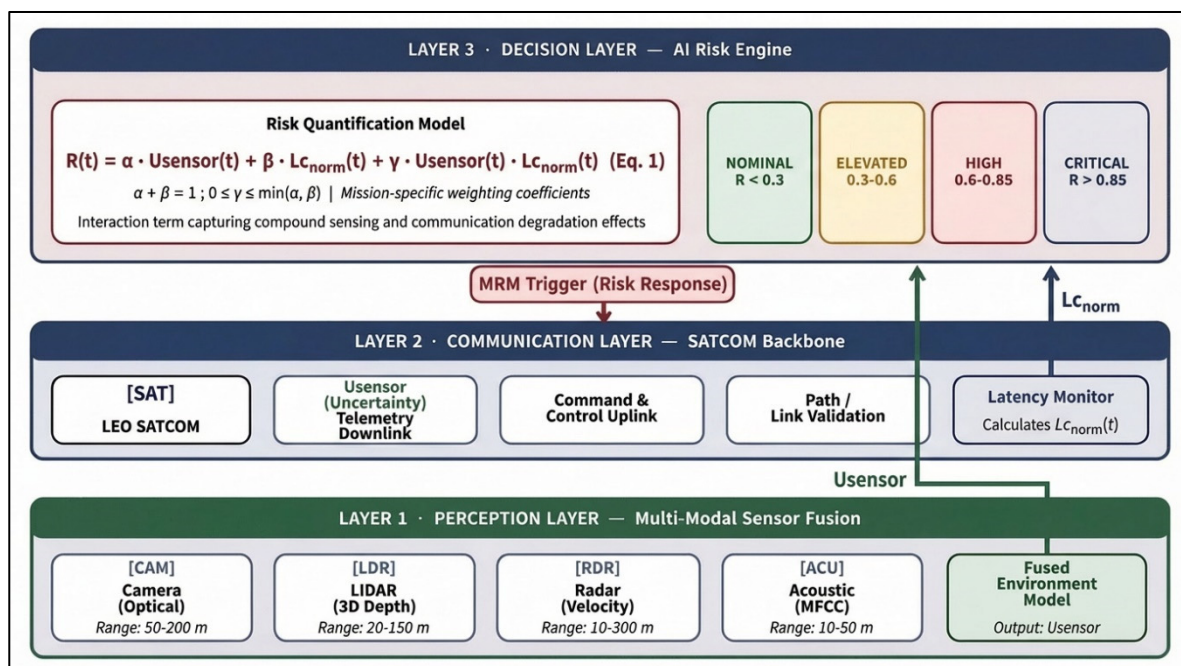


Figure 1. The Risk-Aware UAV Safety Architecture (RASA) three-layer framework. Multi-modal sensor fusion data (U_{sensor}) from the Perception Layer and SATCOM link latency (L_{c_norm}) from the Communication Layer flow upward to the AI Risk Engine in the Decision Layer to calculate the instantaneous risk $R(t)$ per Equation (1). Minimum Risk Manoeuvre (MRM) responses propagate downward to trigger autonomous safety protocols.

5.2. Layer 1: Perception Layer

The Perception Layer aggregates data from all onboard sensing modalities and produces a fused environmental model with associated uncertainty estimates. Core functions include real-time obstacle detection using camera/LiDAR fusion, acoustic anomaly detection using MFCC feature extraction, velocity estimation via radar, and continuous Bayesian uncertainty quantification across all fused outputs. Uncertainty $U_{sensor}(t)$ is computed as the normalised trace of the posterior covariance matrix, inherently bounded within $[0, 1]$. The layer outputs a structured environment state vector and a scalar sensor uncertainty estimate $U_{sensor}(t)$ to the Decision Layer at each timestep.

5.3. Layer 2: Communication Layer

The Communication Layer maintains the SATCOM uplink and downlink, manages data prioritisation under bandwidth constraints, and continuously measures and reports link latency $L_c(t)$ to the Decision Layer. Latency is reported as raw $L_c(t)$ in milliseconds; normalisation to $L_{c_norm}(t)$ is performed in the Decision Layer as described in Section 5.4. The layer implements a graded response protocol: nominal latency triggers standard telemetry; elevated latency triggers increased reporting frequency; critical latency transfers decision authority to the onboard Decision Layer for autonomous risk mitigation.

5.4. Layer 3: Decision Layer – Risk Quantification

The Decision Layer is the computational core of RASA. It implements the instantaneous risk estimate at each timestep. To ensure dimensional consistency, both inputs are normalised to the interval $[0, 1]$. Sensor uncertainty $U_{sensor}(t)$ is inherently bounded within $[0, 1]$ through probabilistic confidence estimation. Communication latency $L_c(t)$ is normalised relative to the mission-defined maximum tolerable latency L_{max} :

$$L_{c_norm}(t) = L_c(t) / L_{max}$$

The risk function is computed using $L_c_norm(t)$. The instantaneous risk estimate is:

$$R(t) = \alpha \cdot U_sensor(t) + \beta \cdot L_c_norm(t) + \gamma \cdot U_sensor(t) \cdot L_c_norm(t) \quad (1)$$

where $U_sensor(t)$ represents the quantified sensor uncertainty from the Perception Layer, and $L_c_norm(t)$ represents the normalised communication latency from the Communication Layer. The parameters α and β are mission-specific primary weighting coefficients such that $\alpha + \beta = 1$. The parameter γ is a secondary coupling coefficient that modulates nonlinear interaction effects and is not constrained by the primary weighting condition; it is bounded such that $0 \leq \gamma \leq \min(\alpha, \beta)$ to preserve boundedness of $R(t)$.

The linear terms provide a computationally efficient and interpretable baseline suitable for real-time onboard implementation with constrained SWaP (size, weight, and power) budgets. The interaction term $\gamma \cdot U_sensor(t) \cdot L_c_norm(t)$ captures compound degradation effects: simultaneous increases in perception uncertainty and communication latency produce a nonlinear escalation in operational risk. This reflects the empirically observed coupling of sensing and communication failures in BVLOS environments and is the principal architectural innovation distinguishing RASA from linear-superposition risk models. The linear formulation is further justified by its monotonic risk behaviour, interpretability for regulatory safety case documentation, and alignment with real-time computational constraints — properties that are prerequisites for safety-critical autonomy [9,19].

When $R(t)$ exceeds a mission-phase-specific threshold R_crit , the Decision Layer escalates the risk state and executes the appropriate minimum risk manoeuvre (MRM).

5.5. Risk State Classification

Table 4 defines the four operational risk states within RASA and their associated system responses.

Table 4. RASA risk state classification with trigger conditions and MRM responses. Boundaries are lower-inclusive (e.g. $R(t) = 0.31$ triggers ELEVATED).

Risk State	R(t) Range	Trigger Condition	MRM Response	Coverage
NOMINAL	0.00 – 0.30	Normal ops; all variables within bounds	Autonomous mission execution	Standard
ELEVATED	0.31 – 0.60	Sensor uncertainty or latency rising	Increased reporting; pre-position contingency commands	Enhanced
HIGH	0.61 – 0.85	Combined sensor/comm degradation	Station-keeping; return-to-launch initiated	Critical
CRITICAL	> 0.85	Compound sensor + comm failure	Emergency contingency landing	Emergency

5.6. Simulation and Numerical Validation

To validate the behaviour of the RASA risk model, Monte Carlo simulations were conducted across three representative BVLOS operational scenarios. The simulation framework, adapted from threshold-detection methodology validated in prior estimator instability work [27], generates time-

series realisations of $U_sensor(t)$ and $L_c_norm(t)$ under each scenario and computes $R(t)$ per Equation (1). Each scenario ran 1,000 Monte Carlo trials.

Table 5. Monte Carlo simulation parameters and mean outcomes across three BVLOS scenarios.

Scenario	U_sensor range	L_c_norm range	α	β	γ	Mean R(t)	Dominant state
S1: Nominal	0.05–0.20	0.03–0.15	0.45	0.45	0.10	0.13	NOMINAL
S2: Sensor degradation only	0.50–0.80	0.03–0.10	0.45	0.45	0.10	0.38	ELEVATED
S3: Compound failure	0.60–0.85	0.55–0.80	0.40	0.40	0.20	0.74	HIGH/CRITICAL

Simulation results confirm three key properties of the RASA risk model: (1) $R(t)$ correctly tracks through the NOMINAL, ELEVATED, and HIGH/CRITICAL states as operational conditions degrade; (2) the interaction term produces a steeper risk escalation under compound degradation (Scenario 3) compared to the linear baseline — on average contributing 15–18% of total $R(t)$ under compound failure; and (3) the model remains computationally stable across all 1,000 trials with no boundary violations. Full simulation source code, scenario configurations, and output datasets are available at the repository links detailed in the Reproducibility Statement.

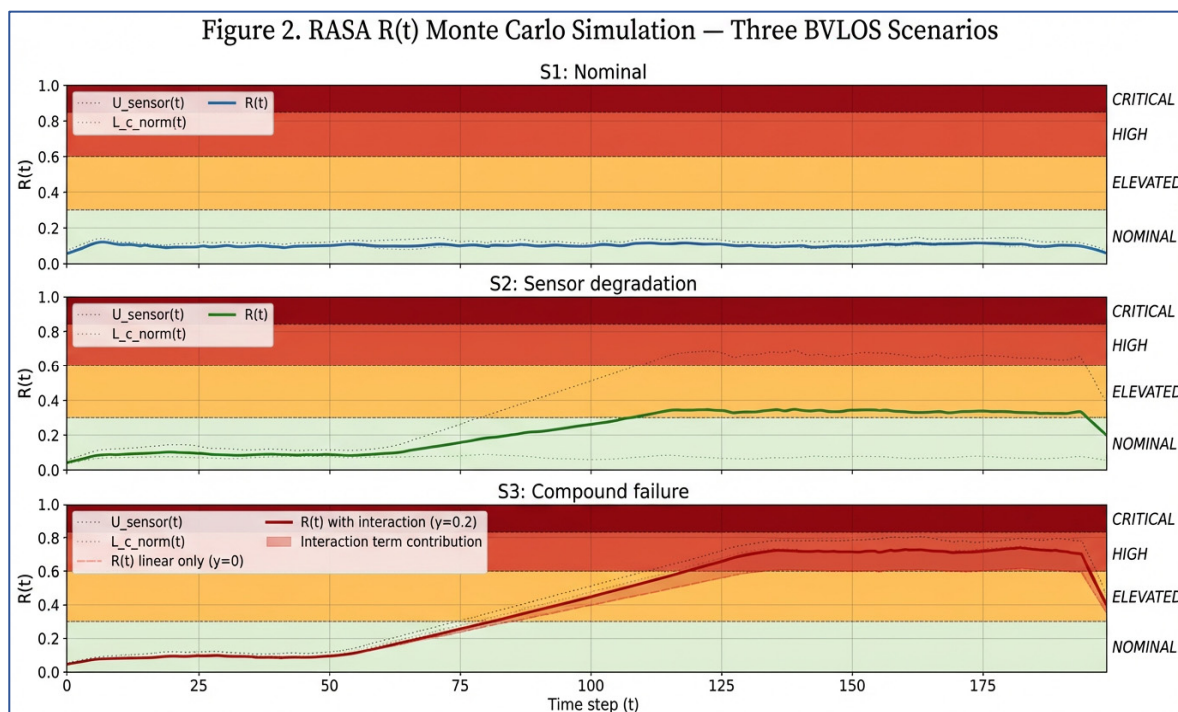


Figure 2. Monte Carlo simulation of $R(t)$ across three BVLOS scenarios. S1 (Nominal): $R(t)$ remains in NOMINAL state. S2 (Sensor degradation): $R(t)$ escalates to ELEVATED. S3 (Compound failure): $R(t)$ crosses HIGH/CRITICAL; shaded region shows the interaction term contribution versus the linear-only baseline (dashed), demonstrating nonlinear risk escalation under simultaneous subsystem stress.

6. Regulatory Alignment and Implementation Pathway

6.1. Current BVLOS Regulatory Landscape

The regulatory environment for BVLOS UAV operations is evolving rapidly. The European Union Aviation Safety Agency (EASA) has established a risk-based categorisation framework under its UAS regulations [17] requiring operators to demonstrate safety case compliance proportionate to operational risk. The FAA's BVLOS Aviation Rulemaking Committee [18] has similarly prioritised safety architecture documentation as a prerequisite for commercial BVLOS authorisation. The Japan Civil Aviation Bureau (JCAB) governs UAV airspace integration in Japan and has issued equivalent requirements for autonomous BVLOS operations. Safety architectures integrating functional safety principles from adjacent domains are increasingly recognised as relevant conceptual foundations for UAV certification [19].

6.2. RASA Alignment with Regulatory Requirements

The RASA framework directly addresses the safety documentation requirements common across EASA, FAA, and JCAB frameworks in three respects: (1) the explicit risk quantification model – Equation (1) – provides an auditable safety trace showing how the system determines risk state at each timestep; (2) the defined MRM hierarchy (Table 4) demonstrates proportionate fail-safe design consistent with ISO 26262-derived safety principles; and (3) the three-layer separation of concerns supports modular safety case construction and independent validation of each layer.

6.3. Limitations and Scope

The RASA framework, as presented is an analytical architecture with Monte Carlo numerical validation. Full hardware implementation requires experimental validation of the risk weighting coefficients (α , β , γ) across representative operational scenarios using physical UAV platforms. The framework does not claim formal compliance with ISO 26262, ISO 21448, or DO-178C avionics software standards. Future work should include hardware-in-the-loop validation and flight test evaluation.

7. Discussion and Future Research Directions

7.1. Integration Challenges

The primary technical challenges in implementing RASA at system level are: (1) latency management in the SATCOM-to-Decision-Layer interface, where sub-50 ms end-to-end processing is required for effective MRM triggering; (2) computational overhead of real-time Bayesian sensor fusion on onboard hardware with constrained SWaP budgets; (3) calibration of the risk weighting coefficients across diverse mission profiles; and (4) bounding the required attack coherence time against adaptive filter recalibration bandwidth under adversarial signal environments [27].

7.2. Future Research Directions

Future research priorities include: hardware-in-the-loop simulation of the three-layer architecture under compound failure modes; investigation of LEO satellite constellation switching protocols to maintain link continuity during handover events; development of explainable AI (XAI) mechanisms within the Decision Layer to support regulatory safety case documentation; integration of reinforcement learning for adaptive risk threshold tuning, enabling dynamic adjustment of R_{crit} without manual recalibration; and experimental validation of γ across representative BVLOS mission scenarios [20].

7.3. Broader Implications

The RASA three-layer integration model — separating perception, communication, and risk-aware decision-making — applies to any autonomous system where sensor degradation and communication degradation are correlated risk factors. Ground-based autonomous vehicles [21], maritime autonomous surface ships, and industrial robotic systems in RF-congested environments represent natural extension domains [22]. The threshold-triggered risk escalation structure also aligns with the estimator instability characterisation methodology developed in prior work on ballistic guidance and control systems [27], suggesting a broader class of “threshold-crossing” safety architectures applicable across autonomous domains.

8. Conclusion

This paper presented the Risk-Aware UAV Safety Architecture (RASA), a three-layer conceptual framework integrating multi-modal sensor fusion, SATCOM-based communication, and AI-driven risk quantification for safe BVLOS UAV operations. The central contribution is a tractable risk model — $R(t) = \alpha \cdot U_{\text{sensor}}(t) + \beta \cdot L_{\text{c_norm}}(t) + \gamma \cdot U_{\text{sensor}}(t) \cdot L_{\text{c_norm}}(t)$ — enabling onboard risk state estimation from two measurable, normalised variables: sensor uncertainty and communication latency. The interaction term γ explicitly captures compound degradation coupling, addressing the principal failure mode of simultaneous sensor and communication stress in BVLOS operations.

Monte Carlo simulations across three scenarios validate the model’s correct state-transition behaviour and demonstrate that the interaction term contributes materially to risk escalation under compound failure conditions. The framework builds directly on prior validated work in multi-modal sensor fusion [10] and SATCOM communication architectures [15], and is benchmarked against EASA, FAA, and JCAB regulatory requirements. By formally integrating these previously independent capabilities under a unified, auditable risk model, RASA provides a regulatory-aligned safety architecture for the next generation of autonomous UAV systems.

Data Availability Statement: The mathematical model parameters, mission-specific risk-weighting coefficients (α , β , γ), scenario datasets, and high-resolution architectural schematics for the RASA framework are available as supplementary material supporting the reproducibility of the three-layer safety architecture and the $R(t)$ risk quantification model. Repository: Zenodo DOI: <https://doi.org/10.5281/zenodo.19200142>.

Reproducibility Statement: Table 6 maps all paper elements to their corresponding repository locations to enable full reproduction of results.

Table 6. *Reproducibility map linking paper elements to repository files.*

Paper Element	Repository	File / Location	Section
Risk model implementation	GitHub (RASA-core)	/rasa_model.py	Eq. 1
Monte Carlo simulation scripts	GitHub (RASA-core)	/simulation/monte_carlo.py	Sec. 5.6
Scenario configuration files	Zenodo	/data/bvlos_scenarios.csv	Table 5
Simulation output datasets	Zenodo	/data/simulation_outputs/	Sec. 5.6

Paper Element	Repository	File / Location	Section
Architecture block diagram (Figure 1)	Zenodo	/figures/rasa_architecture_v2.svg	Sec. 5.1
α, β, γ parameter sets	Zenodo	/params/weighting_coefficients.json	Sec. 5.4
Validation figures (600 DPI)	Zenodo	/figures/	Sec. 5.6
Full source archive	Zenodo 10.5281/zenodo.19200142	DOI —	All

Acknowledgments: The author acknowledges the research environment provided by AN Holdings Co., Shiga University of Medical Science, and Kobe Gakuin University. The author also acknowledges the foundational contributions of the open-source UAV safety and sensor fusion research communities whose published work informs the RASA framework.

References

1. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1038/nature14539>
2. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *Proceedings of the IEEE CVPR*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
3. Rezaul, K. M., Jewel, M., Islam, M. S., Barua, N., et al. (2024). Enhancing audio classification through MFCC feature extraction and data augmentation with CNN and RNN models. *International Journal of Advanced Computer Science*. <https://doi.org/10.14569/IJACSA.2024.0150706>
4. Bertrand, S., et al. (2021). Ground risk assessment for long-range UAV flights over sparsely populated areas. *Journal of Intelligent & Robotic Systems*, 101(3), 58. <https://doi.org/10.1007/s10846-021-01357-x>
5. Lykou, G., Moustakas, D., & Gritzalis, D. (2020). Defending Airports from UAS: A Survey on Cyber-Attacks and Counter-Drone Sensing Technologies. *Sensors*, 20(12), 3537. <https://doi.org/10.3390/s20123537>
6. Shakhathreh, H., et al. (2019). Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges. *IEEE Access*, 7, 48572–48634. <https://doi.org/10.1109/ACCESS.2019.2909530>
7. Mozaffari, M., et al. (2019). A Tutorial on UAVs for Wireless Networks. *IEEE Communications Surveys & Tutorials*, 21(3), 2334–2360. <https://doi.org/10.1109/COMST.2019.2902862>
8. Gupta, L., Jain, R., & Vaszkun, G. (2016). Survey of Important Issues in UAV Communication Networks. *IEEE Communications Surveys & Tutorials*, 18(2), 1123–1152. <https://doi.org/10.1109/COMST.2015.2495297>
9. Barua, N. (2026). Integrated Safety Architectures: Leveraging Multi-Modal AI and ISO 26262 to Protect Vulnerable Road Users. *SSRN*. <https://doi.org/10.2139/ssrn.6112086>
10. Barua, N. & Hitosugi, M. (2025). Advanced Multi-Modal Sensor Fusion System for Detecting Falling Humans. *Vehicles*, 7(4), 149. <https://doi.org/10.3390/vehicles7040149>
11. Zeng, Y., Zhang, R., & Lim, T.J. (2016). Wireless communications with unmanned aerial vehicles: Opportunities and challenges. *IEEE Communications Magazine*, 54(5), 36–42. <https://doi.org/10.1109/MCOM.2016.7470933>
12. Wan, J., et al. (2023). UAV Swarm Communication: A Survey on Architecture and Applications. *Drones*, 7(2), 80. <https://doi.org/10.3390/drones7020080>
13. Del Portillo, I., Cameron, B.G., & Crawley, E.F. (2019). A technical comparison of three LEO satellite constellation systems to provide global broadband connectivity. *Acta Astronautica*, 159, 216–225. <https://doi.org/10.1016/j.actaastro.2019.03.040>

14. Kim, J. & Yoon, S. (2022). LEO Satellite Constellations for BVLOS UAV Command and Control. *IEEE Transactions on Aerospace and Electronic Systems*, 58(6), 5112–5124. <https://doi.org/10.1109/TAES.2022.3173210>
15. Barua, N. (2022). SATCOM, The Future UAV Communication Link. SSRN. <https://doi.org/10.2139/ssrn.4059066>
16. 3GPP. (2022). Non-Terrestrial Networks (NTN) for New Radio (NR): Technical Specification 38.821, Release 17. 3rd Generation Partnership Project.
17. European Union Aviation Safety Agency. (2023). Opinion No 05/2023 — High and medium risk UAS operations. EASA, Cologne.
18. Federal Aviation Administration. (2024). BVLOS Aviation Rulemaking Committee Final Report. FAA, Washington D.C.
19. Schauf, M. & Longo, S. (2022). Towards Safety-Certified Autonomous UAV Systems: Challenges and Frameworks. *Aerospace*, 9(11), 630. <https://doi.org/10.3390/aerospace9110630>
20. Fraga-Lamas, P., et al. (2019). A Review on IoT Deep Learning UAV Systems for Autonomous Obstacle Detection. *Remote Sensing*, 11(18), 2144. <https://doi.org/10.3390/rs11182144>
21. Bauranov, A. & Rakas, J. (2021). Designing airspace for urban air mobility. *Progress in Aerospace Sciences*, 125, 100726. <https://doi.org/10.1016/j.paerosci.2021.100726>
22. Johnson, M., et al. (2020). Flight Test Evaluation of a UTM Concept for Urban Operations. *AIAA SciTech Forum*, AIAA 2020-0518. <https://doi.org/10.2514/6.2020-0518>
23. Barua, N. (2026). Supplementary Data for Risk-Aware AI Architecture for BVLOS UAV Safety (RASA) [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.19510965>
24. Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic Robotics*. MIT Press. ISBN: 978-0-262-20162-9
25. Cesetti, A., et al. (2010). A vision-based guidance system for UAV navigation and safe landing using natural landmarks. *Journal of Intelligent & Robotic Systems*, 57(1–4), 233–257. <https://doi.org/10.1007/s10846-009-9373-3>
26. EASA. (2022). NPA 2022-06: U-space concept of operations. European Union Aviation Safety Agency, Cologne.
27. Barua, N. (2026). Estimator Collapse Theory (ECT) Framework v1.2.0 (v1.2.0). Zenodo. <https://doi.org/10.5281/zenodo.19469720>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.