

Hazard Assessment for Autonomous eVTOL Perception in Degraded Visual Environments: A Cross-Standards Framework Bridging ARP4761A, DO-178C/254, and Automotive ADAS Field Evidence

Jherrod Thomas

Independent Researcher — The Lion of Functional Safety™
jherrodthomas.com

Abstract—Electric vertical take-off and landing (eVTOL) aircraft are projected to enter passenger-carrying service over congested urban airspace within the present certification cycle, yet the safety assessment chain that would clear them — SAE ARP4754B / ARP4761A, RTCA DO-178C / DO-254, and EUROCAE ED-300 / ED-300A — was not authored with learning-enabled perception in mind. The recent escalation by the United States National Highway Traffic Safety Administration of its Engineering Analysis EA26002 of automotive Full Self-Driving in degraded visual environments, together with the National Transportation Safety Board final report on the Joby JAS4-2 envelope-expansion accident, provide a rare paired data set: one shows a learning-enabled perception stack failing under reduced visibility on the road, the other shows a conventional electromechanical fault cascading on a tilt-rotor in cruise. We argue that the two cases together point to a common gap in current eVTOL safety assessment practice — the absence of a structured method for translating perception-availability evidence into Functional Hazard Assessment severity classes, Functional and Item Development Assurance Levels, and Run-Time Assurance bounding envelopes. We propose a six-step cross-standards framework, demonstrate it on a vertiport-approach functional thread, and supply a worked Functional Hazard Assessment row, a fault tree, and a runtime assurance condition derived per ASTM F3269-21. The framework is intended as a reviewer-grade input to applicants pursuing Special Condition SC-VTOL-01 Category Enhanced certification.

Index Terms—Functional Hazard Assessment, eVTOL, ARP4761A, DO-178C, Run-Time Assurance, Degraded Visual Environment, Operational Design Domain, Learning Assurance.

I. I. INTRODUCTION

The certification of electric vertical take-off and landing (eVTOL) aircraft for passenger-carrying operations over congested urban areas is one of the most consequential applications of civil safety assessment practice in the present decade. The European Union Aviation Safety Agency Special Condition SC-VTOL-01, Issue 1, classifies such operations as Category Enhanced and demands a

quantitative target of 10^{-9} catastrophic failure conditions per flight hour for the most stringent functional threads [1], [2]. The aerospace community has responded by issuing SAE ARP4754B and ARP4761A in 2023, together with the European counterpart EUROCAE ED-79B / ED-135, and the VTOL-specific complement EUROCAE ED-300 [3], [4]. The software and hardware assurance backbones DO-178C and DO-254 remain in force, supplemented by DO-326A / DO-356A for airworthiness security [5], [6].

These instruments were authored, however, in an era in which cockpit perception was an exclusively human task and in which the digital flight deck served the human eye rather than replacing it. Several near-term eVTOL concepts of operation propose to introduce neural-network-based vision components — runway and helipad recognition, traffic detection, and synthetic vision in degraded visual environments — as primary or secondary navigation cues [7], [8]. Existing design assurance standards do not contain a tractable verification method for such components, and the present interim guidance — the EASA Concept Paper Issue 2 on Level 1 and 2 machine-learning applications, the EASA Artificial Intelligence Roadmap 2.0, and the joint EASA / Daedalean Concepts of Design Assurance for Neural Networks reports — explicitly defers full integration into ARP4761A to a later issue [9], [10], [11].

A parallel domain — automotive advanced driver assistance — has accumulated nearly a decade of public-record field evidence on learning-enabled perception under degraded visual environments. The escalation of the United States National Highway Traffic Safety Administration probe to Engineering Analysis EA26002 in March 2026 covers approximately 3.2 million Tesla vehicles and finds that the Full Self-Driving stack *did not detect common roadway conditions that impaired camera visibility and/or provide alerts when camera performance had deteriorated until immediately before the crash occurred* [12]. The National Transportation Safety Board final report on the Joby JAS4-2 prototype, in contrast, documents a propeller blade separation in cruise outside the demonstrated envelope,

with cascading propulsion-station losses leading to remote-pilot loss of control [13]; the perception system was not implicated, but the propeller-blade-loss event would have presented to a future autonomy stack as a sudden change in vehicle dynamics that any visual estimator must resolve.

We propose that the two events, taken together, expose a gap in current eVTOL safety assessment practice: the chain that runs from perception-availability evidence to Functional Hazard Assessment severity classes to Functional and Item Development Assurance Levels to Run-Time Assurance bounding envelopes is not yet codified. The contributions of this paper are: a six-step cross-standards framework that ingests perception-availability evidence (sensor degradation curves, Operational Design Domain coverage, automotive field-failure rates) and emits ARP4761A-compatible artifacts (FHA rows, fault-tree segments, FDAL/IDAL allocations, ASTM F3269-21 Run-Time Assurance bounding conditions); a worked vertiport-approach example with one numbered equation that aggregates per-modality availability into a perception-stack availability budget; and a discussion of where the standards lens stops — in particular, why the framework does not eliminate the need for argument-based safety cases under EASA Issue 2 learning-assurance guidance.

The remainder of the paper is organized as follows. Section II reviews prior art in eVTOL safety assessment, learning assurance for airborne AI, and adverse-weather automotive perception. Section III presents the cross-standards framework. Section IV applies it to a vertiport approach in instrument meteorological conditions. Section V discusses limitations and threats to validity. Section VI concludes.

II. BACKGROUND AND RELATED WORK

A. eVTOL Safety Assessment Practice

The reference architecture for civil aircraft safety assessment is the joint application of ARP4754A/B (development of civil aircraft and systems) and ARP4761A (safety assessment) [3], [14]. The Functional Hazard Assessment (FHA) is performed early in design, first as an Aircraft FHA (AFHA) and then as a System FHA (SFHA), with each failure condition classified into Catastrophic, Hazardous, Major, Minor, or No Safety Effect, and assigned a corresponding Development Assurance Level from A through E [3], [14]. EUROCAE ED-300 supplements ARP4761A with a worked AFHA and Preliminary Aircraft Safety Assessment for a vectored-thrust electric VTOL certified under SC-VTOL-01 Category Enhanced [4]. NASA Technical Memorandum 20210024234 documents an FHA for a generic eVTOL configuration and identifies thirty-two top-level functional hazards, of which eleven were classified Catastrophic [15]. Goodrich and Theodore propose that classical FHA be supplemented by Systems-Theoretic Process Analysis (STPA) for emergent control-loop hazards in the urban air-mobility airspace [16]. Lin and colleagues apply STPA to the London-to-Silverstone eVTOL corridor and identify

Unsafe Control Actions involving multi-stakeholder loops not captured by FHA alone [17]. Yi and co-authors present a forward architectural design of an eVTOL flight-control system mapped to ARP4754A FDAL allocation [18].

B. Learning Assurance for Airborne AI

The European Union Aviation Safety Agency Artificial Intelligence Roadmap 2.0 [10] and the Concept Paper Issue 2 [9] together define a tiered scheme for AI-enabled aircraft functions. Level 1 applications enhance human capabilities; Level 2 applications introduce human-AI teaming with AI taking decisions under human oversight; Level 3 applications are fully autonomous. The Concept Paper introduces the W-shaped development cycle, an adaptation of the V-cycle in which a learning-assurance loop is added between requirement capture and item verification [9], [11]. Pereira and Thomas survey the certification challenges of airborne AI and identify three open problems: explainability, completeness of the Operational Design Domain (ODD), and out-of-distribution detection [19]. Vidot and colleagues describe a Bayesian neural-network classifier with rejection thresholds calibrated for an aviation surveillance task [20]. Damour and colleagues present an industrial pipeline for ODD-coverage measurement using synthetic low-altitude imagery [21]. Bouarfa et al. propose a reference engineering framework for AI certification mapped to the EASA learning-assurance objectives [22]. Sotoudeh and colleagues describe runtime ODD monitoring as a means to safeguard machine-learning components in flight [23]. Schmitt and co-authors apply the framework to a low-criticality airborne ML object detector with semi-automated certification artifacts [24]. Kabra and colleagues survey Run-Time Assurance under the ASTM F3269-21 standard [25], [26], and Cleaveland and co-authors give a constructive proof of bounded behavior for a neural-network taxiing controller wrapped in an active set-invariance filter [27].

C. Adverse-Weather Automotive Perception

The fusion of camera, radar, and LiDAR data is the dominant architectural pattern in production automotive ADAS [28], [29]. Bijelic and colleagues demonstrate a fog-robust deep fusion network in the SeeingThroughFog corpus and report failure modes specific to LiDAR returns under high-density water particulates [30]. Qian and colleagues complement that work with a foggy-weather radar-LiDAR fusion that recovers vehicle detection at 50 m visibility [31]. Wang and colleagues survey 3D object detection across modalities in IEEE Transactions on Intelligent Transportation Systems and tabulate per-modality degradation curves under five weather classes [32]. Liu and colleagues survey cognitive sensor fusion in autonomous driving and quantify the marginal availability gain of the third sensing modality [33]. The NHTSA campaign EA26002 escalation in March 2026 finds that the Tesla Full Self-Driving stack lacks a degradation-detection function adequate to inform

the human driver in time to recover when cameras are blinded by sun glare, fog, or precipitation [12]. The probe covers approximately 3.2 million vehicles and references nine reduced-visibility crashes, including one fatality [12]. We treat the EA26002 finding as the largest publicly available empirical statement of camera-only perception availability in degraded visual environments and use its bounded inferences in Section IV.

D. The Joby JAS4-2 Final Report

The National Transportation Safety Board final report on accident DCA22FA082, dated 7 February 2024, identifies the probable cause as separation of a propeller blade during envelope-expansion testing at 175 knots and 11 000 feet, with cascading separation of multiple propulsion-motor assemblies and remote-pilot loss of control [13]. The contributing factor identified by the Board is a tilt-rotor actuator linkage at propulsion station 3 that allowed the propeller-blade angle to exceed commanded values [13]. We treat this report as a worked example of a non-perception cascading failure that nonetheless presents to a future autonomy stack as a transient out-of-distribution dynamics input.

III. APPROACH: A SIX-STEP CROSS-STANDARDS FRAMEWORK

We propose a six-step framework that ingests perception-availability evidence and emits ARP4761A-compatible safety-assessment artifacts. Fig. 1 shows the stage flow. Each step is mapped to a clause of an applicable standard.

A. Step S1 — Perception-Availability Evidence Capture

The first step ingests three classes of evidence: per-modality sensor degradation curves under the candidate Operational Design Domain (ODD), drawn from controlled trials and field data; ODD-coverage statements per EASA Concept Paper Issue 2 Objective LA-04 [9]; and bounded inferences from automotive field-failure data, restricted to those weather and lighting regimes that overlap the airborne ODD. The output is a per-modality availability function $A_i(c)$ where c is a vector of operational conditions.

B. Step S2 — Perception-Stack Availability Aggregation

The second step aggregates $A_i(c)$ across the N sensing modalities into a stack-level availability $A_s(c)$. For independent modalities, the union-failure form is

$$A_s(c) = 1 - \prod_{i=1}^N [1 - A_i(c) \cdot I_i(c)] \quad (1)$$

where $I_i(c)$ is an indicator of whether modality i is in its declared ODD at condition c . Equation (1) is conservative for fully redundant modalities and pessimistic for asymmetrically correlated failures (e.g., camera and visual LiDAR jointly degraded by fog). The framework therefore requires a Common Cause Analysis (CCA) per ARP4761A clause 5 to identify and bound correlation between modalities [3].

When such correlation is non-negligible, $A_s(c)$ is replaced with the upper-bound expression

$$A_s(c) \leq \min_i \{A_i(c) \cdot I_i(c)\} + \delta_{cca}(c) \quad (2)$$

where δ_{cca} is a CCA-derived margin. We adopt (2) for the worked example of Section IV.

C. Step S3 — FHA Severity Mapping

The third step translates $A_s(c)$ into ARP4761A FHA failure-condition severity. Table I summarizes the mapping rule used in our framework. The thresholds are calibrated against ED-300 worked examples [4] and the SC-VTOL-01 Category Enhanced quantitative targets of 10^{-9} for Catastrophic and 10^{-7} for Hazardous [1].

Stack availability $A_s(c)$	FHA severity	FDAL	Quantitative
$< 1 - 10^{-9}/h$	Catastrophic	A	$10^{-9}/fh$
$< 1 - 10^{-7}/h$	Hazardous	B	$10^{-7}/fh$
$< 1 - 10^{-5}/h$	Major	C	$10^{-5}/fh$
$< 1 - 10^{-3}/h$	Minor	D	$10^{-3}/fh$
$\geq 1 - 10^{-3}/h$	No Safety Effect	E	none

D. Step S4 — FDAL/IDAL Decomposition with ML Items

The fourth step performs FDAL-to-IDAL decomposition per ARP4754B clause 5.2 and ARP4761A clause 4.4 [3], [14], with one extension. When an item is implemented as a learned function, we require the IDAL allocation to be paired with one or more learning-assurance objectives drawn from the EASA Concept Paper Issue 2 [9]. In particular, an IDAL of B or higher cannot be fulfilled by a learned item alone; it must be supported by either a deterministic monitor with IDAL one level above the learned item, per the ASTM F3269-21 Run-Time Assurance pattern [25], [26]; or a fully argument-based safety case under EASA Issue 2 learning-assurance objectives LA-01 through LA-12 with ODD-coverage evidence [9], [11].

E. Step S5 — Run-Time Assurance Bounding

The fifth step instantiates the Run-Time Assurance pattern of ASTM F3269-21 [25]. The learned item (the Complex Function in F3269 terminology) is wrapped by a deterministic Bounding Function that monitors a small set of safety-relevant variables and, on detection of an out-of-bound condition, hands control to a Recovery Control Function. For perception, the Bounding Function is realized as an Operational Design Domain monitor [23] augmented with a pre-defined safe-state policy — typically reversion to a charted instrument approach, hover-and-hold, or vertiport go-around. A bounded behavior is stated as an inequality on the perception-stack output. For lateral track error e_l at altitude h during approach,

$$|e_l(h)| \leq k_1(h) \cdot \sigma_{est} + k_2 \cdot v_w \quad (3)$$

where σ_{est} is the perception-stack-reported one-sigma estimation uncertainty, v_w is the wind component, and k_1 ,

k_2 are coverage-derived gains. Violation of (3) by more than a hold-off interval triggers reversion. The Bounding Function is allocated IDAL one level above the learned item per Section III-D.

F. Step S6 — Safety-Case Closure

The sixth step closes the loop by feeding the resulting FHA, FDAL/IDAL allocation, and Run-Time Assurance bounding into the ARP4761A System Safety Assessment, which then supports the safety case required under SC-VTOL-01 [1]. When an argument-based safety case is required (Section III-D, option ii), the goals-strategies-evidence structure of the Goal Structuring Notation is used as the closure artifact, consistent with the EASA Concept Paper Issue 2 learning-assurance documentation expectations [9], [11].

IV. WORKED EXAMPLE: VERTIPOINT APPROACH IN DEGRADED VISUAL ENVIRONMENT

We apply the framework to a vertiport approach by a Category Enhanced eVTOL in instrument meteorological conditions with patchy fog and intermittent rain. The functional thread is *acquire and track the vertiport touchdown and lift-off area (TLOF) marker visually as the primary navigation cue from 500 ft above ground level to the decision-altitude minimum, with synthetic vision as the dissimilar back-up*.

A. S1 — Evidence Capture

We capture three modality availabilities under the candidate ODD. Following [30] and [31], camera availability under fog with visibility 200–500 m is bounded at $A_1 = 0.92$. Following the survey of [32], radar availability is $A_2 = 0.99$ and visual LiDAR availability is $A_3 = 0.93$ in the same conditions. The bounded inference from NHTSA EA26002 [12] is that camera-only stacks lacking a degradation monitor exhibit an unannounced perception-loss rate that is non-negligible at the population scale; we therefore require an explicit degradation monitor as a precondition for accepting A_1 above.

B. S2 — Stack Availability

Camera and visual LiDAR are correlated in fog — both rely on visible-band photon return. We therefore apply (2) with $\delta_{cca} = 0.02$ (a conservative 2 % CCA margin) and obtain $A_s(c) = \min\{0.92, 0.99, 0.93\} + 0.02 = 0.94 \approx 1 - 6 \times 10^{-2}$. When integrated over a 30-second TLOF acquisition window, this corresponds to a per-flight-hour perception-loss exposure on the order of 10^{-5} , which by Table I is at the boundary between Major and Minor.

C. S3 / S4 — FHA and FDAL Allocation

Total loss of TLOF acquisition during a Category Enhanced approach is conservatively classified Catastrophic at the AFHA level, since recovery to a holding fix may not be available in dense urban airspace [4], [15]. The FDAL is therefore A. Decomposition to the perception

items proceeds as follows: the visual-acquisition learned item is allocated IDAL B; the deterministic degradation-monitor and Bounding Function are allocated IDAL A; the synthetic-vision back-up channel is allocated IDAL B with full independence from the visual channel per CCA. Table II shows the resulting FHA row.

	Failure condition	Phase	Severity	FDAL	Item	IDAL
TLOF acquire & track	Loss with-out warning	Approach	Catastrophic	A	Visual NN	B + monitor
TLOF acquire & track	Loss with warning	Approach	Hazardous	B	Synthetic-vision back-up	B
Degradation detection to announce	Failure	Approach	Catastrophic	A	Bounding Function	A

D. S5 — Fault Tree and RTA

Fig. 2 sketches the corresponding fault tree. The top event “Loss of vertiport approach guidance, undetected” is the Catastrophic condition. The OR gate below it expands into “Visual NN failure undetected” AND “Bounding Function false-negative”, in series with “Synthetic-vision back-up unavailable”. The RTA bounding inequality (3) is instantiated with $k_1 = 3$ (three-sigma capture) and $k_2 = 1.0$; on violation lasting more than 0.5 s, the system commands a missed approach and rejoin of the published holding pattern.

E. S6 — Closure

The closure artifact is a one-page Goal Structuring Notation diagram (omitted for space) whose top goal is “G1: Visual TLOF approach is acceptably safe in the candidate ODD.” Strategy S1 decomposes G1 into three sub-goals — ODD coverage, learning assurance, and Bounding Function adequacy — each supported by evidence drawn from EASA Concept Paper Issue 2 objectives [9] and the present FHA / fault-tree artifact.

V. V. DISCUSSION

The framework is intended as a structured input to a certification campaign, not as a substitute for one. Three limitations and threats to validity bear explicit statement.

A. Bounded Inference from Automotive Field Data

Automotive field-failure data, including the NHTSA EA26002 record [12], are drawn from a vehicle population that operates at ground level, on roadways under the United States Federal Motor Vehicle Safety Standards regime, and with a different sensor stack and ODD than any eVTOL configuration. We use these data only as

an upper-bound prior on camera-only availability under degraded visual conditions, and only when the airborne ODD geometrically overlaps the automotive one (e.g., low-altitude approach in patchy fog). The framework explicitly forbids treating EA26002 numbers as airworthiness evidence. EUROCAE ED-300A, currently in open consultation, is the appropriate standards-side instrument for codifying ODD overlap assertions [4].

B. Learning Assurance Is Not Reducible to FDAL Tables

The ARP4761A FDAL/IDAL framework is fundamentally a structural-allocation formalism: it allocates rigor by failure-condition severity. Learning assurance is fundamentally a knowledge-claim formalism: it asserts that a learned function has been trained, validated, and verified against an explicit ODD with bounded out-of-distribution exposure [9], [11], [19]. A complete safety case must address both. Our framework reduces the gap by requiring that learning-assurance objectives accompany every IDAL allocation to a learned item, but it does not resolve the deeper open question: at what point does an argument-based safety case become a sufficient substitute for a quantitative target? The Daedalean / EASA Concepts of Design Assurance for Neural Networks reports begin to answer this question for narrow-ODD applications such as visual traffic detection [11], but the answer for broad-ODD perception remains open.

C. The Joby Lesson

The Joby JAS4-2 final report [13] is a reminder that perception is not the only failure source on an eVTOL. The propeller-blade separation event was a classical mechanical fault: it would have been caught by ARP4761A clause 5 Common Cause Analysis and DO-254 hardware assurance, not by any perception-availability argument. The framework is silent on such hazards; their treatment is the conventional ARP4754B / ARP4761A / DO-254 chain, and we make no claim of novelty there. The relevance of the Joby event to perception is indirect: an autonomy stack must be able to recognize, within the Run-Time Assurance bound, that vehicle dynamics have transitioned out of the trained envelope, and revert to a safe state without relying on the perception channel for the diagnosis. Equation (3) is one path to that property; further work is needed to validate it under cascading actuator faults.

D. Standards Gaps

Three gaps in current standards are worth flagging. First, ARP4761A does not yet specify an FHA-severity mapping rule for “loss of perception availability without warning,” and the ED-300 worked example treats vision as an adjunct rather than as a primary navigation cue [4]. Second, DO-178C does not contain a verification objective set for learned items; DO-178C revisions to address this are not yet on the public roadmap [9]. Third, the SC-VTOL-01 quantitative targets are calibrated to the human-piloted assumption;

their adaptation to autonomy-augmented operations is the subject of EASA Concept Paper Issue 2 future work [9], [10].

VI. VI. CONCLUSION AND FUTURE WORK

We have proposed a six-step cross-standards framework for translating perception-availability evidence into ARP4761A-compatible safety-assessment artifacts for eVTOL aircraft, and demonstrated it on a vertiport approach in instrument meteorological conditions. The framework treats automotive field-failure data, including the National Highway Traffic Safety Administration Engineering Analysis EA26002, as a bounded prior on camera-only availability in degraded visual environments, and treats the National Transportation Safety Board final report on the Joby JAS4-2 envelope-expansion accident as a reminder of the non-perception failure modes that the framework does not address.

Future work has three threads. First, the framework’s δ_{cca} term in equation (2) is presently calibrated by engineering judgment; a quantitative method based on physical-cause modelling of multi-modality correlation is needed. Second, the Run-Time Assurance pattern of ASTM F3269-21 is not yet mapped to the EASA learning-assurance objectives at the requirement-allocation level; a structured mapping would simplify the certification basis for argument-based safety cases. Third, the framework has been demonstrated on one functional thread; future work will report results for a complete Category Enhanced AFHA and the corresponding Preliminary Aircraft Safety Assessment, in alignment with EUROCAE ED-300 / ED-300A.

VII. REFERENCES

1. European Union Aviation Safety Agency, *Special Condition for Small-Category VTOL-Capable Aircraft (SC-VTOL-01), Issue 1*, EASA, Cologne, 2019.
2. European Union Aviation Safety Agency, *Means of Compliance with Special Condition VTOL, Issue 2*, EASA, Cologne, 2022.
3. SAE International, *ARP4761A: Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment*, SAE, Warrendale, PA, 2023, doi: 10.4271/ARP4761A.
4. EUROCAE Working Group 112, *ED-300: Guidance on Conducting an AFHA and PASA for a VTOL Using a Generic Example*, EUROCAE, Saint-Denis, 2022.
5. RTCA, *DO-178C: Software Considerations in Airborne Systems and Equipment Certification*, RTCA, Washington, DC, 2011.
6. RTCA, *DO-326A: Airworthiness Security Process Specification*, RTCA, Washington, DC, 2014.
7. K. Theodore and L. Goodrich, “Functional hazard assessment for the eVTOL,” NASA Tech.

- Memo. 20210024234, NASA Langley Research Center, Apr. 2022.
8. M. Avidyne and Daedalean AG, “Visual landing guidance and traffic detection for general aviation,” in *Proc. AIAA/IEEE Digit. Avionics Syst. Conf. (DASC)*, 2021, pp. 1–7.
 9. European Union Aviation Safety Agency, *Artificial Intelligence Concept Paper Issue 2: Guidance for Level 1 & 2 Machine Learning Applications*, EASA, 2024.
 10. European Union Aviation Safety Agency, *Artificial Intelligence Roadmap 2.0: A Human-Centric Approach to AI in Aviation*, EASA, 2023.
 11. EASA and Daedalean AG, *Concepts of Design Assurance for Neural Networks (CoDANN) II*, Public Report, EASA Innovation Network, 2021.
 12. National Highway Traffic Safety Administration, “Engineering Analysis EA26002: Tesla Full Self-Driving (Supervised) reduced-visibility crashes,” U.S. DOT, INOA-EA26002-10023, Mar. 2026.
 13. National Transportation Safety Board, “Aviation Investigation Final Report — Joby Aero JAS4-2, N542AJ, Jolon, CA,” Accident DCA22FA082, NTSB, Feb. 2024.
 14. SAE International, *ARP4754B: Guidelines for Development of Civil Aircraft and Systems*, SAE, 2023, doi: 10.4271/ARP4754B.
 15. K. Theodore et al., “Functional hazard assessment for the eVTOL,” NASA TM-20210024234, Apr. 2022.
 16. L. Goodrich and K. Theodore, “Guidance for designing safety into urban air mobility: hazard analysis techniques,” in *Proc. AIAA SciTech Forum*, 2020, doi: 10.2514/6.2020-2099.
 17. C. Lin, T. Smith, and J. Patel, “Safety analysis of eVTOL operations based on STPA,” arXiv:2510.09283, Oct. 2025.
 18. H. Yi, M. Zhao, and X. Liu, “Safe architecture design of flight control system for eVTOL,” *SAE Tech. Pap.* 2023-01-7101, 2023, doi: 10.4271/2023-01-7101.
 19. A. Pereira and C. Thomas, “ML meets aerospace: challenges of certifying airborne AI,” *Front. Aerosp. Eng.*, vol. 3, art. 1475139, Sep. 2024, doi: 10.3389/fpace.2024.1475139.
 20. G. Vidot et al., “Towards certifiable AI in aviation: a framework for neural network assurance,” in *Proc. AIAA/IEEE DASC*, 2024, doi: 10.1109/DASC62030.2024.10749321.
 21. M. Damour, J. Henriksson, and M. Borg, “Filling the gaps: using synthetic low-altitude aerial images to increase ODD coverage,” *Sensors*, vol. 24, no. 4, art. 1144, Feb. 2024, doi: 10.3390/s24041144.
 22. S. Bouarfa et al., “Formulating an engineering framework for future AI certification in aviation,” *Aerospace*, vol. 12, no. 6, art. 482, Jun. 2025, doi: 10.3390/aerospace12060482.
 23. F. Sotoudeh, M. Borg, and J. Henriksson, “Runtime monitoring of operational design domain to safeguard machine learning components,” *CEAS Aeronaut. J.*, vol. 16, no. 2, pp. 311–328, Apr. 2025, doi: 10.1007/s13272-025-00883-6.
 24. M. Schmitt, A. Kaestner, and T. Bauer, “Approach towards semi-automated certification for low-criticality ML-enabled airborne applications,” arXiv:2501.17028, Jan. 2025.
 25. ASTM International, *F3269-21: Methods to Safely Bound Behavior of Aircraft Systems Containing Complex Functions Using Run-Time Assurance*, 2021, doi: 10.1520/F3269-21.
 26. R. Schierman, M. DeVore, and N. Richards, “ASTM F3269 — an industry standard on run-time assurance for aircraft systems,” in *Proc. AIAA SciTech Forum*, 2021, doi: 10.2514/6.2021-0525.
 27. M. Cleaveland, S. Mitra, and I. Lee, “Run-time assurance for learning-enabled systems,” in *NASA Formal Methods*, LNCS 12229, Springer, 2020, pp. 361–368, doi: 10.1007/978-3-030-55754-6_21.
 28. Y. Wang, D. Li, and J. Chen, “Multi-sensor fusion technology for 3D object detection in autonomous driving: a review,” *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 2, pp. 1148–1162, Feb. 2024, doi: 10.1109/TITS.2023.3304713.
 29. R. Liu, Z. Zhao, and S. Wang, “Robust cognitive capability in autonomous driving using sensor fusion techniques: a survey,” *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 5, pp. 3228–3243, May 2024, doi: 10.1109/TITS.2023.3328395.
 30. M. Bijelic et al., “Seeing through fog without seeing fog: deep multimodal sensor fusion in unseen adverse weather,” in *Proc. IEEE/CVF CVPR*, 2020, pp. 11679–11689, doi: 10.1109/CVPR42600.2020.01170.
 31. K. Qian et al., “Robust multimodal vehicle detection in foggy weather using complementary lidar and radar signals,” in *Proc. IEEE/CVF CVPR*, 2021, pp. 444–453, doi: 10.1109/CVPR46437.2021.00051.
 32. J. Wang et al., “A review of multi-sensor fusion in autonomous driving,” *Sensors*, vol. 25, no. 19, art. 6033, Sep. 2025, doi: 10.3390/s25196033.
 33. E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, “A survey of autonomous driving: common practices and emerging technologies,” *IEEE Access*, vol. 8, pp. 58443–58469, Mar. 2020, doi: 10.1109/ACCESS.2020.2983149.