Safety Expertise and the Perils of Novelty

Frank McCormick
Certification Services, Inc. (ret.)
islandtimekeeper@gmail.com

Mallory Graydon, Natasha Neogi, Paul Miner, and Jeffrey Maddalon
NASA Langley Research Center, Hampton, VA, USA
{m.s.graydon, natasha.a.neogi, p.s.miner, j.m.maddalon}@nasa.gov

***Abstract*—Emerging aviation markets such as urban air mobility are giving rise to new technologies and means of operation. However, novelty may hide 'unknown unknowns,' raising new hazards. This paper examines how expertise and safety techniques enable transformative technologies such as reduced crew operations, hybrid wing-borne and rotor-born flight, federated air traffic services, and urban operations. We explore how analysts use expertise to address common-cause failures, collect and interpret safety data, and perform exacting tradeoffs between dissimilarity, redundancy, independence, and diversity (human, process lifecycle, or otherwise) to ensure safety. When novelty is present, analysts might not possess the expertise needed to fully understand the implications of design decisions and tradeoffs being made, especially in early lifecycle phases, on emergent properties such as safety. Safety expertise must be carefully cultivated. The conflicting views of safety experts must be unpacked to identify the divergence in fundamental assumptions, models, means, and methods that may be causing them. Once systems venture beyond the basis of what safety expertise can reliably guarantee, projects take on risk that must be managed.**

The paper contains key takeaways and actionable recommendations for novel OEMs and regulators touching on topics such as robust monitoring; clear and transparent reporting; incremental approaches to fielding novel systems in hazard-rich, risk-tolerant environments; the cultivation of safety culture and expertise in an organization; and the use of scientific study to reduce epistemic uncertainty in novel operations with new technologies. Since excessive novelty in aviation can undermine the current foundation of safety, humility and incrementalism are necessary to enable emerging aviation markets safely.

***Index Terms*—safety, expertise, novelty, urban air mobility**

I. Introduction

Emerging aviation markets such as urban air mobility, autonomous cargo, increasingly autonomous air traffic management, and commercial space operations are giving rise to new technologies and novel means of operation. Advanced capabilities such as autonomy, distributed electric propulsion, vehicle-to-vehicle communication, virtual and augmented reality, cloud computing, new materials and manufacturing, structural health monitoring, and others will transform the national airspace system (NAS). However, 'unknown unknowns' may hide in novel system designs and operations raising both new hazards and old hazards in unexpected ways. Judicious, incremental application of novelty becomes necessary for the concentrated application of well-known and emerging safety techniques to manage risk carefully.

This paper explores the expertise that underpins all safety engineering, with a specific eye on software considerations, accident and incident investigation, and operational safety. This expertise is crucial to both performing hazard analyses and to making exacting tradeoffs between, and discriminating blends of, dissimilarity, redundancy, independence, and diversity (human, process lifecycle, or otherwise).

Once systems venture beyond the basis of what expertise can reliably guarantee, projects take on risk. While we focus here on safety risk, acting without the necessary expertise raise to other risks, including project risk [9]. Minimizing harm from insufficient safety expertise requires learning what is necessary as quickly as possible. It also requires a willingness to open the system and safety process up to scrutiny at all phases of the design lifecycle. This process must persist throughout the operational lifecycle with thorough investigation and comprehensive reporting of incidents and accidents to regulatory agencies, competitors, and the public. Potential competitors must learn from each other and recognize that their shared need for a safe reputation for the industry outweighs competition concerns.

The paper contains multiple key takeaways and actionable recommendations for novel OEMs and regulators touching on topics such as the use of robust monitoring programs; clear and transparent reporting practices; incremental approaches to fielding novel systems in hazard-rich, risk-tolerant environments; the cultivation of safety culture and expertise in an organization; and the use of scientific study to reduce epistemic uncertainty in novel operations with new technologies. Since excessive novelty in aviation can undermine the current foundation of safety, humility and incrementalism are necessary to enable emerging aviation markets safely.

II. Novelty in future aviation systems

Aviation is poised to expand to new markets and embrace new technologies. Safely embracing this novelty will require new expertise and a judicious approach.

A. Novel role and responsibility paradigms

There are multiple proposed novel role and responsibility paradigms for electric vertical take-off and landing (eVTOL) aircraft performing urban air mobility (UAM) operations. Simplified vehicle operations (SVO) [15], where the pilot no longer retains the final authority and responsibility for making some safety critical decisions, is often cited as a key enabler of scalable UAM due to the shortage of pilots and the cost associated with piloting. However, the feasibility and safety effects of this paradigm have not yet been fully investigated, and it is unclear what the implications of reducing the human pilot's understanding of aircraft flight will be.

Another proposed paradigm is 'm-to-N' operations, where m is a small number of operators remotely supervising a large number of N aircraft. To lower the risk of entry, this might most suitably be employed first for cargo-carrying operations. However the issues involved with pilot incapacitation remain a dominant feature of any safety discussion of this paradigm.

Both paradigms, like others, shift critical decisions from humans to automated agents. For instance, some inner-loop stabilization controls tasks cannot be performed by human pilots. However, per 14 CFR Part 91.3 [1], "The pilot in command of an aircraft is directly responsible for, and is the final authority as to, the operation of that aircraft." Limitations on the pilot's capability to intervene to assure the safety of the aircraft requires a paradigm shift in safety assessment.

Once automation becomes responsible for a critical decision, it must evince all contributions to safety that a human pilot made in the same role. This includes handling edge cases: the automation should be resilient and be able to handle unforeseen, unplanned, and unanticipated contingencies. It is currently difficult to respond to single failures in complex scenarios, let alone having a robust response to arbitrary combinations of multiple failures.

B. Novel operations

Scalable UAM operations over an urban environment would bring a dramatic expansion in the number of air vehicles over the densest parts of cities in prolonged, low overflight. In addition to posing challenges to existing noise regulations, land management issues in urban cores means that good alternative landing sites for distressed vehicles will be scarce. This poses challenges for emergency landings. But, unlike current operations where mitigating vehicle hull loss adequately address the majority of risk to both passengers and third parties on the ground not connected to the operation, there may be extensive risks posed to ground-based parties and infrastructure from dropped parts, loss of separation with structures by low-flying craft, etc. These risks will need to be accounted for in safety engineering and certification.

C. Novel vehicle technologies

Novel vehicle technologies are also being proposed. For example, a notional "Lift plus Cruise" vehicle has been developed by NASA's Revolutionary Vertical Lift Technologies Project [19], [20], [34]. It operates as a fixed-wing aircraft whenever possible and uses rotors for lift only during takeoff and landing. The most distinctive attribute of this concept is that the rotors for lift are separate from those used for propulsion, and the lift rotors are stopped and aligned with flow during cruising flight. Characteristics of a lift plus cruise vehicle include: (1) multicopter-style redundancy that can be employed to deal with failures in the lifting system; (2) simple propellers with RPM control that can be used with all rotors and motors of the lifting system of the same design, adding to ease of safety analysis; and (3) a small wing, optimized for cruise, that can be used, since slow-speed flight is performed with lift augmentation from the propellers [20].

The power flow throughout the system is tightly coupled with the control in hover and conversion for the lift+cruise aircraft [19], [20]. When in vertical flight configuration, the pilot cannot directly control the rotor speed to control the aircraft direction. Because humans are incapable of this, reversion to direct pilot control of a 'dead-stick' landing is not a viable mitigation for total loss of power in this configuration. Instead, topologies with multiple batteries serving all or some of the motors are used to mitigate such failures. Careful attention to the control algorithms, transition between vertical and wing-borne flight, and exceptional flight conditions is needed to ensure sufficient control authority and predictable behavior in both nominal and off-nominal conditions. Additionally, because of the excessive heat generated by the battery, battery fires and/or battery explosions pose additional failure scenarios.

D. Novel air traffic systems

Many proposed UAM operations involve carrying passengers at lower altitudes in and around major metropolitan areas through the use of UAM corridors. These UAM operations will be incorporated into the NAS, where traditional air traffic control and UAS traffic managed operations will coexist. UAM vehicles, eVTOL or otherwise, will have required equipage and meet performance requirements to operate in the current NAS. Both NASA and FAA have developed an initial concepts of operations for UAM [13], [14], [17]. The FAA UAM ConOps v1.0 is focused on "nearer-term", crewed operations [13], while the NASA/Deloitte UAM Maturity Level (UML) 4 "Vision" ConOps is focused into future operations with increasing automation and increasing number and/or density of flights [17].

In these proposals, the UAM airline operator's and UAM vehicle operator's actions may be supported by service providers known as providers of services to UAM (PSUs). The exact nature of PSU services, of communications between UAM craft, PSUs, vertiports (i.e., takeoff and landing areas), and existing air-traffic service providers has yet to be specified. For example, there is an open question of whether tactical separation within UAM corridors will be allocated to the UAM airline operator, UAM vehicle operator and/or pilot-in-command, and PSU, or whether it will remain a function of traditional air-traffic control. UAM vehicles must register their operational intent with a PSU to enter, transit, or leave a UAM corridor. However, non-UAM vehicles can cross a UAM corridor in Class C, D, or E airspace without declaring their operational intent [13]. This could lead to a significant number of disruptions of UAM operations and should be studied further. But while details of how UAM air traffic will be controlled are still being developed, it seems clear that significant novelty will be involved.

E. Enablers of Urban Air Mobility operations

The major difference between the proposed UAM concepts and current helicopter transit services offered from operators like Blade and Uber hinges on the key enablers of (1) electric

propulsion and (2) autonomy. The proposed eVTOL vehicles aim to reduce noise (and potentially other) emissions via electric propulsion systems and multi-rotor concepts to address critical concerns from overflown communities. Additionally, the scale of the near-mature UAM operations is seen as being hundreds (if not thousands) of vehicles at any given time over an urban core, a tenfold or greater increase over the maximal state of today's system. These eVTOL vehicles will likely possess more onboard automation in order to enable this scale of near-mature UAM operations. At the moment, both current day operations and proposed UAM operations cite the use of corridors [13], but current helicopter corridors are not seen as adequate for the scale of UAM operations envisioned. Similarly, with increased automation in both the cockpit and traffic management, it is necessary to revisit and redesign processes, procedures, and training for crew and all relevant personnel in order to understand where safety critical decision making may be occurring and how roles, authority, and responsibility paradigms may have shifted. This novelty in vehicle type and operations requires a thorough study of all of the potential ways hazards may evince themselves as well as a probing examination of what new hazard might arise.

As with any technology that has the potential to harm human beings, three questions for the developers, maintainers, and operators arise: (1) *How can you hurt people?*, (2) *What are you going to do about it?*, and (3) *How do you know when you're done?* The answers to these questions and the quality of those answers will both turn, in part, on safety expertise. To explore how and why, we next turn to how safety expertise underpins the safety of current aviation operations.

III. The knowledge that underpins current aviation safety

Aviation safety processes reflect the history of aviation. Using them, transport-category aircraft have amassed an enviable safety record. But following those procedures is not, by itself, sufficient to ensure safety. Safety expertise is required.

A. Standards and lifecycles

Three key safety standards are among the many recognized means of compliance with airworthiness regulations for transport-category aircraft: the safety engineering and safety analysis standards SAE ARP4754A and ARP4761 and the software standard RTCA DO-178C [30]–[32]. But safety engineering is not limited to design time; it is accompanied by both accident and incident investigations and operational hazard mitigations.

1) Safety engineering and safety assessment: Civil aircraft are developed following safety engineering and safety analysis standards. For example, compliance with the SAE ARP4754A and SAE ARP4761 standards is accepted in the USA and elsewhere as a means of compliance with the 14 CFR 25.1309 reliability requirements for transport category aircraft [11], [31], [32]. ARP4754A defines guidelines for a development assurance process that uses a safety assessment process to identify failure conditions for the aircraft and its systems, defines safety requirements to mitigate these failure conditions, and uses further safety analysis to confirm satisfaction of safety requirements. ARP4761 gives guidance on conducting this safety analysis, defining processes for functional hazard assessment (FHA), fault tree analysis (FTA), and failure modes and effects analysis (FMEA), among others. But effective use of these standards requires safety expertise.

2) Software considerations in safety: When deployed to a processor that controls safety-relevant actuators aboard an aircraft, an error in its construction might cause a dangerous failure condition. In the USA (and elsewhere), conformance to RTCA DO-178C is accepted as a means of showing compliance with the applicable airworthiness regulations for the software aspects of airborne systems and equipment [10], [30]. DO-178C and its companion documents define guidance for the production of software for airborne systems and equipment. The standard's objectives for the software lifecycle process include, famously, objectives related to the requirements coverage and structural code coverage of software test suites.

But the aim of DO-178C's guidance is limited to what might be described as software correctness considerations. DO-178C leaves the consideration of the safety implications of software to a wider safety process (such as described by ARP4754A). Where software designers choose behavior that is not specified by the software requirements they are given, they must communicate these to the system safety team who will, if necessary, define additional or revised software requirements.

3) Accident and incident investigation: Aviation accidents and serious incidents are investigated by professionals, such as, in the United States of America, staff from the National Transportation Safety Board. This process is noteworthy for its degree of openness and cooperation: aircraft manufacturers and air carriers understand that honest assistance to investigators is the right way to proceed. Official accident reports routinely publish details such as traces of parameter values from flight data recorders. Any corporate discomfort with public revelation of potentially proprietary information is buried deeply beneath the understanding that everyone, including the manufacturer of the incident aircraft, benefits from a comprehensive understanding of what went wrong and why. Without that understanding, similar accidents might recur, bringing with them loss of public confidence and patronage.

Accident and incident investigations benefit from forensic records that allow investigators to reconstruct what happened and why. In transport-category aircraft, this includes cockpit voice recorders that record radio transmissions, pilot utterances, aural alerts, and sounds from the machinery and environment; flight data recorders that capture parameters such as control inputs, system states, and air data parameters; and quick access recorders that record a variety of parameters for maintenance purposes. As aircraft have evolved, the quantity and scope of this record has expanded to ensure investigators have what they need.

4) Operational safety: The failures of many aircraft functions, such as passenger entertainment systems, pose little or no safety risk to a flight. Other failures are often addressed by mitigations outside the aircraft. For example, loss of all radio communications might sound disturbing, but procedures have been in place since before World War II to deal with silent aircraft. If voice contact is lost, flight crews know to execute certain standard responses such as adhering to previously agreed flight paths. Likewise, air traffic control has standard responses, which include clearing airspace around the silent craft. The problem is dealt with smoothly not by building more reliable radios or by installing redundant radios but instead by well-understood procedural mitigations.

Similar mitigations have been proposed and used to deal with unreliable automated or remotely piloted aircraft. For example, an agricultural spraying aircraft with no pilot aboard might safely treat a field if it can be established that it will remain within an area that is kept clear of personnel who might be injured if it crashes [16].

B. The underpinnings of aviation safety

From the Wright Brothers at Kitty Hawk to modern airliners, aviation's approaches to safety have changed with the times. Understanding the role of safety expertise and its implications for novelty requires understanding this history.

1) Selecting reliable components: In aviation's dim and distant beginnings—say, the time of Manfred von Richthofen and Eddie Rickenbacker—safety was treated as essentially a reliability problem: if engineers and mechanics could just make and maintain components that were durable and robust, then vehicles would generally remain airworthy and pilots would face mechanical emergencies only rarely, or at least rarely enough.

Later generations came to grips with a harsher reality: no component can be made wholly failure-proof. No physical component can be made eternally perfect. Things break. Things wear out. And, more insidiously, even in the absence of component failures, increasing complexity and sophistication of aircraft systems brought new dangers, new opportunities for misunderstanding, for mental paralysis of flight crews, and for confused mishandling in moments of cockpit crisis. For example, the addition of multiple automation modes raises the possibility of mode confusion [26].

2) Building in redundancy: If we can't make an airplane that never fails, some reasoned, perhaps we can make an airplane that continues to fly safely when important pieces of it do. Soon, the answer was shown to be yes, with more than one layer of reasoning behind it. Redundancy was an obvious layer. If engine failure is undesirable, then just provide two or more engines such that failure of any one powerplant does not doom the vehicle. Germany's Junkers 52 and America's Ford Trimotor, icons of the 1920s and '30s, are three-engine designs offering a decent chance of surviving an engine failure at a time when aircraft engines were, by today's standards, unreliable. Problem solved? Not quite.

3) Addressing common causes: Multiple engines were sensible, but there are assumptions hidden within all schemes based on redundancy—assumptions that have sometimes been later revealed to be unjustifiably optimistic—namely that the failures of interest are random and independent, which is to say uncorrelated. Distressingly, redundant resources such as multiple engines or replicated hydraulic controls can be rendered instantly useless if all are threatened at the same time for the same reasons. If one uncontained engine failure can disable all the parallel hydraulic systems on which flight control relies, it can disable flight control [28]. If the airplane's fuel is contaminated, it doesn't matter how many engines you enjoy [3]. If all four of an airliner's generator control devices are closely spaced when installed in the belly of the airplane and are flooded and ruined together when a galley sink above them overflows, then it was irrelevant that multiple (nominally separate and independent) generator control devices had been provided. All failed at the same time for the same reason [5]. It is precisely to address such possibilities that ARP4754A calls for analyses to identify and engineering to mitigate common causes of redundant system failures [31].

Redundant computing devices can fail simultaneously if they all use the same erroneous software [24]. Not even systems running independently developed software can be guaranteed to fail independently [21]. Careful attention is required to design fault tolerance systems, especially their voting and fault management systems. An error in such a system can itself produce a failure of supposedly independent computer systems [6], [7], [36].

4) The essential role of expertise: The modern approach to aviation safety—including attention to reliability concerns, built-in redundancy, mitigations for common-cause failures, and operational mitigations—has been successful. It achieved this success in large part by incorporating lessons from experience. Even safety processes conducted per the detailed guidance in ARP4754A and ARP4761 do not allege perfection but rather try to exploit the skills and knowledge of experienced practitioners who have spent their careers paying close attention to aviation accidents and incidents and their causes while reflecting deeply on the best ways to prevent future mishaps.

The aircraft industry's front-line defense in this regard has been an unassuming cordon of design engineers. No matter their technical specialty, the best designers typically share two personality traits. First, they are never not thinking about product safety. No answer or choice is ever offered except in the wake of the most comprehensive battery of intellectual what-ifs of which that expert is capable. Second, these people embody a native humility that recognizes fallibility, including their own, as hard-wired in humanity. No design element can be taken as divinely given, no axiom or assumption as unchallengeable. People make stupid choices and miss absurdly obvious blunders. They just do.

Good aeronautical engineers are thus always prepared, indeed eager, to uncover potential threats to safety, often subtle and seemingly implausible threats, and to chase them to bitter

ends, no matter where they might lead through the rest of the aircraft. These habits build up a safety expertise that is necessary for safety. Much of the art, science, and business of safety engineering lies in a judicious blend of redundancy, independence, and dissimilarity. Experienced designers can be superb in judging those blends.

Like engineers, operators and maintainers use and cultivate a specialized body of safety expertise. Crop dusters quickly learn of the specialized hazards of their profession, such as power lines, and put in place practices such as limitations on when they fly to address these. Dispatchers remain available to crews in flight to, among other things, serve as a single point of contact for technical questions that arise. Even the language used in conducting flight operations is scrutinized in the wake of accidents and incidents to identify where improvements could ward off dangerous miscommunication [8].

C. Process is not enough

The need for relevant safety expertise is not obviated by safety or hazard analysis process. The development assurance and safety assessment processes defined by ARP4754A and ARP4761 and the software standard DO-178C are powerful tools and deserve some of the credit for the excellent safety record of modern air travel. Practitioners use these standards to address complex threats to safety, including common-cause failures. But the efficacy of the practices these standards describe depends on the sufficiency of the hazard analysis processes and the proposal and judgment of the safety impacts of designs, which in turn depend on the safety expertise of practitioners.

1) Hazard analyses processes depend on analysts' safety expertise: There are many hazard analysis processes one might use to identify unsafe states of an aircraft and how the aircraft might enter such states. These include the tailored-for-aviation versions of analyses such as Functional Hazard Assessment (FHA), Zonal Hazard Analysis (ZHA), Failure Modes and Effects Analysis (FMEA), and Fault-Tree Analysis (FTA) described in ARP4761 [32]. These also include general techniques such as the System Theoretic Process Analysis (STPA) [23], [38]. But these techniques do not magically create safety insight where none existed. Analysts' ability to identify failure conditions and contributions to them in response to these questions depends on the analysts' safety expertise, e.g., their knowledge of what might go wrong or how it might go wrong.

Consider how analyses such as FHA, ZHA, FMEA, and STPA work [38]. The analysts gather a representation of the aircraft or system to be analyzed and scope safety questions of interest. They then systematically analyze that representation of the focus of analysis, a bit at a time, asking what-if questions at every stage. Where the answers are safety relevant, analysts document the resulting insight, e.g., the hazardous aircraft or system state or the design's potential contribution to entering it. But what-if questions can't elicit potential outcomes that the analysts can't conceive of. One can systematically examine a compartment in an aircraft and note both the potential of a fuel coupling to leak and the presence of a cross-feed duct that might get very hot. But if the analysts were somehow unaware of the possibility of fire where fuel, oxygen, and heat are gathered, their answer to the what-if question about spilled fuel might not contain the desired insight.

The idea of an analyst tasked with doing hazard analyses being unaware of fire is farfetched. But we have seen incidents and accidents where the failure mechanisms involved were not well understood beforehand, and thus could not be reliably identified by safety analysis. For example, the way hail behaves differently when ingested into turbofan engines only became clear after a dual engine failure incident [12]. The existence of a 'sticky range' of temperatures in which water in aviation fuel causes serious problems was revealed by an accident in which this phenomenon led to an aircraft landing 30 meters short of the runway [2].

Not even automated hazard analysis, if such a thing could be made effective, could solve the problem of novelty potentially requiring safety expertise that does not yet exist. Any kind of computer-based analysis would be a systematic examination of digitized models. But those models can only be guaranteed to accurately represent the properties the tool's designers thought relevant. And the tools' analysis of them can only identify as harmful the kinds of combinations of model features its creators programmed or trained it to identify.

Systematically guiding analysts to think about specific aspects of specific portions of an aircraft or system might be a great way to ensure that humans have the opportunity to notice all manner of potential dangers. But until those humans have the safety expertise needed to understand what they are seeing, they may not see it.

2) Design and safety assessment processes are not enough: Once hazard analyses have identified unsafe states of an aircraft or system, designers must mitigate these and confirm satisfaction of safety objectives with a safety assessment. But as with hazard analyses, neither the design process nor the safety assessment are sufficient on their own: both rely on the safety expertise of designers and safety engineers.

In practice, many of the hazardous states and failure modes for traditional fixed-wing transport-category and general-aviation aircraft, their systems, and components are well understood and have well-established mitigations. These have effectively have emerged over decades and have been formalized and standardized in ways that reliably boost confidence and enable repeatable review and approval by designers and regulatory authorities. Nevertheless, such strategies can be little more than cumulative scar tissue, accumulated in response to problematic real-life experiences. This approach does not allege perfection but rather tries to exploit the skills and knowledge of the most experienced practitioners available, those who have spent their careers paying close attention to aviation accidents and incidents and their causes, while reflecting deeply on the best ways to prevent future mishaps.

This experience is particularly crucial in cases where designers must make tradeoffs in hazard mitigation. For example,

there is heated debate about the need to mitigate potential design errors. Development assurance practices aim to ensure that engineered systems will do as their requirements dictate, with no surprises. But they are necessarily imperfect: Developers err. Testers err. The software used in verification and validation is itself imperfect. While this imperfect reality is widely acknowledged, experts differ on what should be done about it. Some insist that, when the consequence of failure might be catastrophic, good development assurance alone is insufficient and mitigations such as dissimilar redundancy must be deployed. If we can't guarantee that software is perfect, maybe it is better to have two independently written software components and compare their outputs. Others insist that such approaches are expensive, offer uncertain benefits, and may introduce more failure points. The cases for each position turn on interpretation of a diverse collection of evidence: Experimental results show that independently developed software may not fail in statistically independent ways [21]. History shows that when software is implicated in an incident or accident, it is often the very fault management logic on which redundancy-based mitigations depend that is at fault [6], [7], [36]. And developers weigh their own experience—often proprietary and unpublished—including a history of recorded faults and external conditions.

3) Software quality assurance is not enough: As noted in Section III-A2, the DO-178C software lifecycle does not directly address safety; it aims at satisfaction of requirements derived from systems engineering and system safety processes, with requirements emerging from design fed back to the safety process to check safety implications. But that process is again not sufficient by itself: it reflects safety expertise relevant to software like that its authors had experience with, and its efficacy for other kinds of software remains unknown.

The airborne software being created when the standard was written took the form of feedback control: typical aviation software gathers, checks, and processes input from sensors and controls; performs mode switches in response; and computes and exports outputs for actuators and displays. Such software is quite unlike some other kinds of software such as that using machine learning, and it is unclear how the standard's objectives could even be applied to such software, let alone how effective it would be. Achieving safety of complex autonomous systems based on machine learning can require addressing a wide variety of considerations, from bug guts on camera lenses to the suitability of the learning algorithm used [37]. An SAE committee is looking into how artificial intelligence might safety be deployed in an aviation context [33].

D. Talented people are not enough

Every aerospace development project needs capable engineers to complete its objectives. Different companies choose different strategies for hiring these engineers: some choose to hire engineers at the average salary with small adjustments for particularly experienced or inexperienced personnel, others choose an approach to hire the best at premium salaries and, of course, other strategies are also possible. The motivation for paying premium salaries rests on the idea that more difficult projects will be finished faster by hiring great employees, perhaps with punishing overtime schedules. Due to faster time to market with accompanying market dominance, the expectation is that this "premium hiring" approach will save the company money compared to other approaches.

Some propose to apply this approach to novel safety-critical aerospace development programs. The first problem with this approach is determining the kind of worker needed for aerospace. Is a smart worker all that is needed? Perhaps not: no evidence suggests that workers who did well on academic tests are the best workers for safety critical aerospace systems. Typically, students are taught how systems behave in the nominal case, they are not taught system misbehavior or how far from nominal conditions the system will operate over millions of hours of operation. In fact, academic programs rarely include courses that provide an academic understanding of safety expertise, let alone the practical application and judgement required to implement safety critical systems in the real world. The primary means to develop safety expertise in a workforce is on-the-job training, with the oversight of experienced workers. So, perhaps the talented workers are not primarily identified as smart, but instead are identified by experience. Unfortunately, novelty now enters the discussion. If the system is novel, then no known experience exists that will guarantee a successful completion of the project. An experienced worker may have learned lessons that are invalid for this novel system. An oft-told anecdote is that to increase the safety of physical systems (such as bridges) one often needs to add weight; however, the safety of software systems often comes by taking out code.

Talented people are not infallible, while safety critical systems must be. Smart people may be able to process information quickly or foresee issues that no one else has seen. But it is unjustified to expect that such people will be able to foresee all eventualities. Alternatively, the experienced safety engineer may know where the danger lies in a conventional design but not in all the places of a novel design. Of course, none of this is to say that smart or experienced workers should be avoided, only that such workers will not provide the same kinds of benefits that they do on projects without novelty.

IV. Novelty examples

To wit, let's have a quick glance at three novel technologies that have either been recently introduced or have been proposed: battery technologies, UAM primary flight controls, and artificial intelligence. This list is far from exhaustive. It is, at best, an illustration of how we might approach a thought process for the new bits we will be dealing with.

Unexplained ignition and self-combustion of rechargeable batteries in common use today is not a wildly rare phenomenon. Battery fires have occurred spontaneously in consumer electronics, in cargo shipments of lithium-ion batteries, and elsewhere, often with no apparent cause. Indeed, Boeing's 787 Dreamliner, shortly after its introduction in revenue service, suffered a pair of surprise battery fires in its then-new

high-energy rechargeable electrics, resulting in a grounded fleet, design rework, and financial losses [18], [27]. Such events speak to insufficient safety expertise: had the people designing this battery system or approving its design truly understood how the internals of such batteries work, we would not have introduced their associated hazards into an aircraft. The good news is that battery issues are grounded in physics and thus presumably eligible for proper investigation and understanding, leading to demonstrably effective mitigations.

The kinds of primary flight control systems proposed for UAM vehicles might be more problematic. The kinds of flight control systems used for transport-category fixed-wing aircraft are well understood, as are mitigations for their failures: redundant flight control computers, reversionary control modes, etc. In flight critical systems and equipment for commercial aviation, we have long relied on various categories of design choices to help aircraft and their crews deal with the unexpected: random component failures, unfavorable operating conditions, and so on. If one engine stops working in flight, you continue flying on the other engine, provided the cause of the first engine failure is not shared with the second engine.

But accidents happen for reasons other than simple, uncorrelated failure of batteries, motor controllers, and motors. We've seen physical damage from one failed system take out redundant components from other systems, both when those other systems were located together [5] and when they were not [28]. We've seen failures due to unplanned-for component failure modes defeat the very failure management strategies needed to leverage redundancy [6]. We've even seen accidents where control laws themselves produced behavior that pilots didn't understand, frustrating recovery from an upset condition [4]. Contingency maneuvers, flight control laws, and mode switching logic can be tricky to get right and have contributed to accidents [29]. Simple redundancy can't account for common causes of failure, and safety engineering can't account for causes of failure—common or otherwise—that aren't well-understood by designers and safety analysts.

The point is that weird things happen. They just do, and we can't anticipate all future possibilities. In revenue service, UAM flights will encounter weird and unexpected challenges. An eVTOL craft might be designed to rely on digital flight controls with no redundant/independent/dissimilar backup capable of backstopping flight-critical functions whose operation has become anomalous for reasons not anticipated during design. To the extent that a multicopter eVTOL design lacks such backup, the firm operating that vehicle is storing trouble for its future. Much of civil aviation's history has been a parade of eliminating Achilles Heels. Total reliance on the eternal perfection of a single solution—flight controls implemented solely in digital form—would be an excellent Achilles Heel to address before widespread deployment of passenger-carrying UAMs becomes a reality.

And, finally, we come to artificial intelligence. If treated as a universal backstop—whatever we can't figure out, we'll just dump on some machine learning gizmo—the AI could well be the most dangerous technology ever introduced into the national airspace. It is so new, so powerful, so tempting, and so unknown, we simply have no basis for making any meaningful predictions, let alone what to do with it or about it. Some features ("learning") fundamentally violate all previous notions of safety analysis namely that the system will behave predictably, even in an unpredictable world. Extreme humility and slow, careful, incremental exploration are advised.

V. Safety expertise and its implications

The design, construction, operation, and maintenance of aircraft requires safety expertise that accumulates through both experience and diligent thought. If this expertise is to be available when it is needed, it must be carefully cultivated at all phases of the aircraft design and operational life cycle.

A. Openness, transparency, and humility are essential

The safety expertise that is so critical to aviation safety is developed by the practitioners who employ it. Key to its development is the way these experts approach their work. Safety experts think deeply about safety in all aspects of their work. They offer no answers or choices until they've asked and answered the most comprehensive battery of intellectual what-if questions they can muster. They embody a native humility that recognizes fallibility, including their own, as hardwired in humanity. They take no design element, axiom, or assumption as beyond challenge. They understand that people make stupid choices and miss absurdly obvious blunders. They are thus prepared, or even eager, to uncover potential threats to safety—often subtle and seemingly implausible threats—and to chase them to bitter ends, no matter where they might lead through the aircraft. And this skeptical humility persists through the operational lifecycle, with thorough investigation and fulsome reporting of incidents and accidents [25, §6.4–5].

Takeaway 1: A robust monitoring program is essential to build safety knowledge. The civil aviation sector—particularly passenger transport operations—have long had a robust practice of accident and incident investigation. Airframers and engine manufacturers work with investigators and regulators to identify the causes of mishaps; aircrews and maintainers file reports in repositories such as the Aviation Safety Reporting System, and airframers and engine manufacturers monitor the operating performance of their products. This relentless re-asking, *But how safe is it really?* is critical to the gaining of safety expertise that could have further improved designs and can improve operation and maintenance. It is even more critical where there is novelty.

Takeaway 2: The entire industry sector must be open and transparent. Firms in most industries have an interest in protecting company trade secrets and proprietary data. There is a potential for short-sighted employers to punish or sideline staff who report incidents or insights that threaten a profitable status quo. But the aviation industry has long had a remarkable culture of openness and transparency. There are anonymous reporting systems. Vendors work with accident investigators, providing information that appears in detailed, publicly accessible reports. This openness is recognized as

enabling the same continuous, vigilant monitoring process that keeps aviation historically safe, which is good for everyone. Given the even higher potential for unknown unknowns that comes with novelty, new entrants to the aviation market should continue this history of openness and transparency.

B. Novelty must be approached cautiously

New contexts, techniques, and technologies bring new unknowns, both known and unknown. There might be new failure modes we do not yet understand. Undesired events, both in our systems and our environment, might occur more frequently than we anticipated. The balance between our concerns might change. And because safety expertise is built up through experience, no amount of safety analysis is entirely sufficient to de-risk novelty. Before TACA 110 suffered loss of thrust in both engines due to ice ingestion, we did not know that ice and water needed separate considerations in engine certification [12]. Before an all-engines-out incident in 1982, the effects of volcanic ash on aircraft were not well appreciated [35]. Design and procedures cannot be guaranteed to address concerns that are not known (or understood) in advance.

Takeaway 3: Don't embrace novelty for novelty's sake. Venturing into novelty brings risk and cost. Even when the risks do not manifest in lives lost, belated discovery of how novelty affects safety can give rise to expensive remediation. For example, the 787 Dreamliner fleet was grounded after incidents demonstrated unappreciated safety considerations of its novel lithium-ion batteries [18], [27]. Responsibly introducing novelty requires careful monitoring to detect how and where our incomplete safety expertise left a design riskier than it should be. When those factors are fully appreciated, airframers might find it worthwhile to stick to the familiar, especially in cases where the potential benefits novelty might enable are modest.

Takeaway 4: Try out novelty in safer applications. One answer to the unknown risk raised by what we don't know about novel things is to field them first in applications where failure can be made either harmless or less harmful. For example, one might deploy an autonomous or remotely piloted agricultural spraying aircraft over a field that is being kept clear of humans, minimizing many of the potential risks that might accompany an aircraft that has not been verified to meet standard structural airworthiness criteria [16]. One might also deploy novel technology in situations where some harm remains but where the potential good that only an aircraft with no humans aboard might do is so large as to create tolerance for unknown risk [22]. Just as early experience with turbine engined aircraft built up the experience base that led to regulatory permission to fly long distances away from suitable landing facilities, such experiences might build the safety expertise that might justify wider future use of once-novel technologies.

C. Capture the debt

The behavior of novel systems contains a significant dimension of "unknown unknowns." The safety assurance tasks for conventional systems heavily rely on common knowledge possessed by experienced engineers on the subject of these systems. Some of the most important common knowledge concerns the well-known interfaces between subsystems. For these conventional systems, only deviations from the norm needed to be addressed through special assurance procedures. However, with novel systems, there is no well accepted common knowledge. Thus, how to gather this knowledge is an open question. The practice of each engineer having a constant activity of considering what assumptions or shortcuts they have made is crucial to establishing a common knowledge base.

These assumptions and shortcuts should include considering why this system works, what information (and its characteristics) is being received from outside the system, how this system can fail, what are the effects of this failures, what mitigations could be used, what must be monitored and measured in service in order to judge authoritatively whether mitigations are working as intended, do these choices on monitoring affect vehicle design or flight operations, will this design feature induce new operational or maintenance activities, and when existing rules and guidance are inadequate or silent on the relevant issues, what are the thresholds for regulatory involvement and emergency intervention such as Airworthiness Directives or operational restrictions.

As an example, a software engineer might skip an array-bounds check or a memory reclamation step if he or she will be the only participant affected by execution of that code. The engineer would, however, be expected to document this "technical debt" in a canonical program-level repository to ensure that the expedient shortcoming is not forgotten later, and not allowed to remain in production. Capturing such items in a "technical debt" log allows the eventual evaluation of each assumption, in a team setting, to ensure the debt is eventually paid. This activity will provide some level of identifying the "known unknowns." This is a good first step towards positioning yourself to discover the "unknown unknowns."

Takeaway 5: Keep a log of "technical debt" that is identified in the development of a novel system. Pay this "debt" at some point in the development.

D. Test a lot

Testing is expensive in both money and time. For this reason, testing is often viewed as a place to look for cost cutting, or at least as a place where special approaches are taken to perform the absolute minimum number of tests. Any attempt to develop a minimal set of tests for a novel system is unachievable and ultimately unhelpful. The primary way to deal with "unknown unknows" of novel systems is to gain knowledge and this knowledge will come through testing. Explore the behavior of the system in as many environmental conditions and operational modes as you can. In particular, explore combinations of these conditions, even the ones where

you expect to get the same answer as previously tested sets of conditions. When you get answers that are not what was expected, try to imagine why and refine the tests conditions to explore the matter further. As always, theoretical models of system behavior will guide the testing, however such models have particularly low value for novel systems. Such models will rely on (often unstated) assumptions. The validity of these assumptions in real-world conditions is unknown and it should be assumed that they are invalid. Finally, testing should be viewed as a learning activity unto itself. The goal of testing in novel systems is not merely to complete the test matrix, but as a guide to discover where the system's behavior was particularly unpredictable. And these results point to areas where finer grain testing is to be conducted.

Takeaway 6: Test maximally, not minimally. Do not expect the testing to have a clear up-front schedule. Use testing to learn about the system's behavior and as a means to determine where more testing should be conducted.

E. Safety expertise must be cultivated

The lifetime's work of a safety expert is where aviation safety expertise is both developed and practiced. The minds that do this work hold the detailed and expansive repository of safety expertise compared to which any standard or textbook or policy document is but a distilled digest. But such minds can retire. They can quit. They can die. It is thus essential for any airframer to cultivate safety expertise in its workforce.

Takeaway 7: Safety expertise must be cultivated in employees. It is not enough to simply hope that young engineers will absorb the examples set by experts despite temptations to cut corners or rationalize their own suboptimal, naïve solutions. Nor does every promotion process ensure that the people who step up into key technical roles possess the safety expertise they will need. Ensuring that expertise is fostered, passed on, and selected for is a serious institutional challenge that demands deliberate attention.

F. Safety expertise must be brought in where needed

Some newly-formed organizations may lack an existing old guard of aviation safety experts. But even established organizations may need to bring into the design process particular expertise that design staff lacks. Neither safety nor safety expertise is limited to vehicle design. Safety depends on maintenance. It depends on training. It includes organizational cultures and operator priorities. And it emphatically includes procedural and operational restrictions. Crop dusters have a compelling interest in being able to see powerlines and other obstacles in the fields they are spraying. Thus, most crop dusters fly only in daylight and avoid working in nighttime darkness or dense fog. Relevant safety expertise exists in folks whose roles span the aircraft lifecycle, and many of them are not directly employed by an airframer.

Takeaway 8: Safety expertise must be brought in where in-house expertise is insufficient. For organizations new to aviation, this might include hiring experienced aviation safety practitioners. It might be tempting to dismiss such folks' recommendations as the old way of doing things, to be broken by fast-moving new organizations, but their knowledge is key to ensuring safety even in novel designs. For all organizations, this might include bringing in people with relevant knowledge about later phases of the aircraft lifecycle, e.g., pilots and maintenance staff. Such folks should be engaged early in the aircraft design process when their input has the greatest chance of warding off the design blunders that will lead to safety compromises.

G. Scientific study might help . . . to a degree

One recommendation that might be made is for scientific study of how novel technologies impact safety. Such study can be challenging. For example, controlled experiments in which a significant number of teams use either traditional or new techniques to conduct full-scale development activities would be expensive. And it is difficult to persuade private businesses to part with data that might be considered a trade secret.

VI. Conclusion

The safety processes behind the enviable safety record of commercial aviation depends on the safety expertise of the engineers, operators, and maintainers of those aircraft. As aviation embraces new technologies and enters new markets, such as advanced aerial mobility, new expertise will be needed to maintain safety. Ensuring safety in the face of novelty requires carefully cultivating the necessary expertise. This requires a robust monitoring program, including data collection, incident and accident investigation, and anonymous reporting of concerns. It requires openness and transparency. It requires embracing novelty slowly and cautiously, and only as needed, possibly in applications where the benefits justify the risk. It requires acknowledging the ignorance associated with novelty as a technical debt and tracking it appropriately. It requires testing to thoroughly explore the system in all its familiar and unfamiliar aspects. And it requires both cultivating safety expertise in staff and bringing in expertise where needed.

References

[1] 14 CFR Part 91. *Title 14 Code of Federal Regulations, Electronic Code of Federal Regulation, Part 91—General Operating and Flight Rules*. United States Government, 2023. URL: https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-91.

[2] Air Accidents Investigation Branch. Report on the accident to Boeing 777-236ER, G-YMMM, at London Heathrow Airport on 17 January 2008. Aircraft Accident Report 1/2010, Department for Transport, Aldershot, UK, February 2010. URL: https://assets.publishing.service.gov.uk/media/5422f3dbe5274a1314000495/1-2010_G-YMMM.pdf.

[3] Air Accidents Investigation Branch. Report on the serious incident to Airbus A321-211, G-POWN, at London Gatwick Airport on 26 February 2020. Aircraft Accident Report 1/2021, Civil Aviation Authority, London, UK, May 2021. URL: https://assets.publishing.service.gov.uk/media/6087c670e90e076ab1e3492c/1-2021_Airbus_A321-211_G-POWN.pdf.

[4] Aircraft Accident Investigation Board. F-35A, T/N 12-005053, 58th Fighter Squadron, 33rd Fighter Wing, Eglin AFB, Florida, 19 May 2020. Report, United States Air Force, Eglin Air Force Base, Florida, USA, September 2020. URL: https://www.afjag.af.mil/Portals/77/AIB-Reports/2020/May/Eglin%20AFB%20F35A%20AIB%20Report_Signed.pdf.

[5] Australian Transport Safety Bureau. Electrical system event: Boeing 747-438, VH-OJM, 25 km NNW of Bangkok International Airport, Thailand, 7 January 2008. Final report AO-2008-003, ATSB, December 2010. URL: https://www.atsb.gov.au/publications/investigation_reports/2008/aair/ao-2008-003.

[6] Australian Transport Safety Bureau. In-flight upset 154 km west of Learmonth, WA, 7 October 2008, VH-QPA, Airbus A330-303. Transport Safety Report AO-2008-070, ATSB, December 2011. URL: http://www.atsb.gov.au/publications/investigation_reports/2008/aair/ao-2008-070.aspx.

[7] Australian Transport Safety Bureau. In-flight upset event, 240 km north-west of Perth, WA, Boeing Company 777-200, 9M-MRG, 1 August 2005. Aviation Occurrence Report 200503722, ATSB, March 2007. URL: http://www.atsb.gov.au/publications/investigation_reports/2005/AAIR/aair200503722.aspx.

[8] Bureau d'Enqueêtes et d'Analyses. Serious incident between the Boeing 787-10 registered N16009 and the Airbus A320-214 registered OE-IJF. Technical report, BEA, France, September 2021. URL: https://bea.aero/fileadmin/user_upload/BEA2020-0289.en.pdf.

[9] Audrey Decker. Tardy training jet reveals limits of digital design, Air Force says. *Defense One*, 23 May 2023. URL: https://www.defenseone.com/defense-systems/2023/05/tardy-training-jet-reveals-limits-digital-design-air-force-says/386687/.

[10] Federal Aviation Administration. Airborne software assurance. Advisory Circular 20-115C, FAA, Washington, DC, USA, July 2013. URL: http://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_20-115C.pdf.

[11] Federal Aviation Administration. Development of civil aircraft and systems. Advisory Circular 20-174, FAA, Washington, DC, USA, September 2011. URL: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_20-174.pdf.

[12] Federal Aviation Administration. Lessons learned: TACA International Airlines 737 near New Orleans. Web page, FAA, January 2015. URL: http://lessonslearned.faa.gov/ll_main.cfm?TabID=1&LLID=40&LLTypeID=2.

[13] Federal Aviation Administration. Urban Air Mobility (UAM) concept of operations. Technical report, FAA, April 2023. Version 2.0. URL: https://www.faa.gov/sites/faa.gov/files/Urban%20Air%20Mobility%20%28UAM%29%20Concept%20of%20Operations%202.0_0.pdf.

[14] Federal Aviation Administration. Urban Air Mobility (UAM) use case document. Technical report, FAA, February 2021. Version 1.2.

[15] General Aviation Manufacturers Association (GAMA). Transitioning to electric vertical takeoff and landing (eVTOL) and other aircraft equipped for simplified vehicle operations (SVO). Whitepaper, GAMA, Washington, DC, USA, February 2020. URL: https://gama.aero/facts-and-statistics/consensus-standards/publications/gama-and-industry-technical-publications-and-specifications.

[16] K. J. Hayhurst, J. M. Maddalon, N. A. Neogi, H. A. Verstynen, B. Buelow, and G. F. McCormick. Mock certification basis for an unmanned rotorcraft for precision agricultural spraying. Technical Memorandum NASA/TM-2015-218979, National Aeronautics and Space Administration, Hampton, VA, USA, November 20215. URL: https://ntrs.nasa.gov/citations/20160000766.

[17] B. Hill, D. DeCarme, and M. Patterson. National Aeronautics and Space Administration (NASA) UAM vision concept of operations (ConOps) UAM maturity level (UML) 4. Technical report, Deloitte, December 2020. URL: https://ntrs.nasa.gov/api/citations/20205011091/downloads/UAM%20Concept%20of%20Operations%2018DEC2020.pdf.

[18] U. Irfan. How lithium ion batteries grounded the Dreamliner. *E&E News*, 18 December 2014. URL: https://www.scientificamerican.com/article/how-lithium-ion-batteries-grounded-the-dreamliner/.

[19] W. Johnson and C. Cilva. NASA concept vehicles and the engineering of advanced air mobility aircraft. *The Aeronautical Journal*, 126(1295):59–91, October 2021. doi:10.1017/aer.2021.92.

[20] L. W. Kohlman, M. D. Patterson, and B. E. Raabe. Urban air mobility network and vehicle types: Modeling and assessment. Technical Memorandum NASA/TM-2019-220072, National Aeronautics and Space Administration, Hampton, VA, USA, February 2019. URL: https://ntrs.nasa.gov/api/citations/20190001282/downloads/20190001282.pdf.

[21] J. C. Knight and N. G. Leveson. An experimental evaluation of the assumption of independence in multi-version programming. *IEEE Transactions on Software Engineering (TSE)*, 12(1):96–109, January 1986. doi:10.1109/TSE.1986.6312924.

[22] S. M. Lehman, J. T. Slagel, S. Adrade, H. Walsh, A. Goodloe, S. Brandt, and N. Neogi. NASA System-Wide Safety wildland firefighting operations workshop report. Technical Memorandum NASA/TM-20220014721, National Aeronautics and Space Administration, Hampton, VA, USA, September 2020. URL: https://ntrs.nasa.gov/citations/20220014721.

[23] N. G. Leveson and J. P. Thomas. STPA handbook. Electronic document, March 2018. URL: http://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf.

[24] Ariane 501 Inquiry Board. *Ariane 5 Flight 501 Failure: Report by the Inquiry Board*. Paris, France, July 1996. URL: https://esamultimedia.esa.int/docs/esa-x-1819eng.pdf.

[25] N. Neogi, M. Graydon, and G. F. McCormick. Partial preliminary hazard assessment for an eVTOL aircraft supporting urban air mobility (UAM) operations. Technical memorandum, National Aeronautics and Space Administration, 2023. In Press.

[26] National Transportation Safety Board. Descent below visual glidepath and impact with seawall: Asiana Airlines flight 214, Boeing 777-200ER, HL7742, San Francisco, California, July 6, 2013. Accident Report NTSB/AAR-14/01, NTSB, Washington, DC, USA, June 2014. URL: http://www.ntsb.gov/doclib/reports/2014/AAR1401.pdf.

[27] National Transportation Safety Board. Auxiliary power unit battery fire Japan Airlines Boeing 787-8, JA829J, Boston, Massachusetts, January 7, 2013. Aviation Incident Report NTSB/AIR-14/01, NTSB, Washington, DC, USA, November 2014. URL: http://www.ntsb.gov/investigations/AccidentReports/Reports/AIR1401.pdf.

[28] National Transportation Safety Board. United Airlines flight 232, McDonnell Douglas DC-10-10, Sioux Gateway Airport, Sioux City, Iowa, July 19, 1989. Accident Report NTSB/AAR-90/06, NTSB, Washington, DC, November 1990. URL: https://www.ntsb.gov/investigations/AccidentReports/Reports/AAR-90-06.pdf.

[29] National Transportation Safety Board. Kansas State University Meridian UAS hard landing, McMurdo, Antarctica, December 20, 2011. Aviation Investigation Docket DCA12CA023, NTSB, Washington, DC, April 2013. URL: https://data.ntsb.gov/Docket?ProjectID=82580.

[30] RTCA DO-178C. *Software Considerations in Airborne Systems and Equipment Certification*. RTCA, Inc., Washington, DC, USA, December 2011. URL: https://my.rtca.org/NC__Product?id=a1B36000001IcmqEAC.

[31] SAE ARP4754A. *Guidelines for Development of Civil Aircraft and Systems*. , December 2010. URL: https://www.sae.org/standards/r/arp4754a/.

[32] SAE ARP4761. *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment*. SAE International, December 1996. URL: https://www.sae.org/standards/content/arp4761/.

[33] SAE International. G-34 Artificial Intelligence in Aviation. Standards committee, 2023. URL: https://standardsworks.sae.org/standards-committees/g-34-artificial-intelligence-aviation.

[34] C. Silva, W. Johnson, K. R. Antcliff, and M. D. Patterson. VTOL urban air mobility concept vehicles for technology development. In *Proceedings of the AIAA Aviation Technology, Integration, and Operations Conference*, Atlanta, GA, USA, June 2018. URL: doi:10.2514/6.2018-3847.

[35] SKYbrary. B742, en-route, south southeast of Jakarta Indonesia, 1982. Web page. URL: https://www.skybrary.aero/accidents-and-incidents/b742-en-route-south-southeast-jakarta-indonesia-1982.

[36] Taiwan Transportation Safety Board. June 14, 2020, China Airlines Flight CI202, Airbus A330-302, Registration Number B-18302, The aircraft experienced multiple systems failures during landing at Songshan Airport. Major Transportation Occurance, Final Report TTSB-AOR-21-09-001, TTSB, September 2021. URL: https://www.ttsb.gov.tw/media/4936/ci-202-final-report_english.pdf.

[37] UL 4600. *Evaluation of Autonomous Products*. UL Solutions, 17 March 2023. URL: https://www.shopulstandards.com/ProductDetail.aspx?productid=UL4600.

[38] K. Wasson, N. Neogi, M. Graydon, J. Maddalon, P. Miner, and G. F. McCormick. Functional hazard assessment for the eVTOL aircraft supporting urban air mobility (UAM) applications: Exploratory demonstrations. Technical memorandum, National Aeronautics and Space Administration, 2022. *In press. URL: https://ntrs.nasa.gov/api/citations/20210024234/downloads/NASA-TM-20210024234.pdf*.