Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference
Edited by Eirik Bjorheim Abrahamsen, Terje Aven, Frederic Bouder, Roger Flage, Marja Ylönen

doi: 10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P6532-cd

From rigid orderliness to barely controlled chaos: sociotechnical risk and AI in aviation

Jan Hayes

School of Property, Construction and Project Management, RMIT University, Australia. E-mail: Jan.hayes2@emit.edu.au

Martin Rasmussen Skogstad

NTNU Social Research, Norway. E-mail: martin.rasmussen@samforsk.no

Given the potentially catastrophic consequences of errors, faults and poor decisions in aviation, to date artificial intelligence (AI) applications are permitted only in non-safety related activities and tasks, and machine learning is banned during in flight operations. By restricting the possible adverse consequences of using AI technologies, this approach also severely restricts the possible benefits and so there are broad plans from regulatory bodies to allow further integration of AI into the sector. Based on interviews with aviation sector safety and AI experts, this study aims to understand the strengths and vulnerabilities in current aviation safety processes and how processes and practices may need to be adapted to address safety in AI. Drawing on Macrae's SOTEC (Structural, Organizational, Technological, Epistemic, and Cultural) framework for sociotechnical risk in autonomous and intelligent systems, we develop a preliminary set of risks posed by use of AI in aviation across these five domains. One of the significant challenges is the fact that different parts of the sector have different safety management approaches and so may be impacted by AI in different ways. Safety in aircraft manufacturing and flight operations is certification and compliance based. When it comes to safety in air traffic management, with multiple actors making judgment-based time pressured decisions, one interviewee described the environment as 'barely controlled chaos'. Uncertainty is high and risk-based processes prevail. This paper unpacks these issues and looks at the implications for identification and evaluation of novel risks linked to new AI applications.

Keywords: Uncertainty, Artificial Intelligence, Aviation, sociotechnical risk.

1. Introduction

The idea of artificial intelligence (AI) has been around for decades, but computing technology has now advanced to the point where what were once imaginative ideas can now become practical realities. '*Artificial intelligence (AI) refers to systems that display intelligent behavior by analyzing their environment and taking actions – with some degree of autonomy – to achieve specific goals' (HLEG 2019).* There is significant potential for AI in aviation to provide increased capacity and cost efficiency in aircraft operations and in air traffic management. Safety improvements are possible, but there is also potential for significant problems if this transition is not managed appropriately.

Drawing on the opinions of ten experts regarding aviation safety and AI, we address the following question:

What new sociotechnical risks are likely to be introduced to the aviation sector by the introduction of AI technologies?

2. Safety in Aviation

The early days of aviation are a story of daring and innovation with little regard for safety. Through the early part of the twentieth century, this exotic pastime of the wealthy gave way to the development of routine passenger transport services with a safety record that has largely gained the trust of the general public. The pace of change has slowed so much in the last 50 years that Downer described the key public safety strategy of commercial aviation as "innovative restraint" (2017). He further argues that aviation has achieved high reliability by two mechanisms – recursive practice (by which he means many flight hours) and design stability. This is not the conventional view, with other industry authors noting that the industry's strong ethos of compliance and certification combined with design redundancy makes the difference.

More recently, such rules-based approaches to safety which have been largely successful in preventing accidents have been augmented by risk-based processes driven by safety management system requirements introduced by ICAO Annex 19. Initially covering airlines and air traffic control, these requirements have recently been extended to cover aircraft design and manufacturing. The tools used in each of these areas, if not the overall safety paradigms, are set to be challenged by the introduction of AI.

3. Introduction of AI to Aviation

In both Europe and the US, AI use in aviation is currently restricted to systems that are not safety-critical and no machine learning is permitted during in flight operations. AI software deployed on an aircraft, for example, therefore comes with an integrated fixed database. More operational data may be gathered during each flight, but this can only be incorporated into software algorithms following development and recertification.

Table 1. Human/AI interaction categories (EASA 2023).

Level 1 AI: Assistance to humans • 1A: Human augmentation • 1B: Human cognitive assistance in decision making and action selection.
Level 2 AI: Human – AI Teaming • 2A: Human and AI-based system 2B: cooperation • 2B: Human and AI-based system collaboration
Level 3 AI: Advanced Automation • 3A:The AI-based system performs decisions and actions that are overridable by the human. • 3B: The AI-based system performs non-overridable decisions and actions (e.g. to support safety upon loss of human oversight).

There is broad understanding that further modes of use of AI must be allowed to reap further benefits that the technology presents. Three classifications of AI applications are shown in Table 1 (EASA 2023). EASA further suggest that the likely timing of introduction of these various types of AI is starting now for Levels 1 and 2 with level 3 applications not likely until 2035 at the earliest. Issues to be resolved are not just technical with many questions of trust in the technology and ethics of its use remaining open.

4. Method

4.1. *Framework for Sociotechnical Risk*

To study the impact on risk of the introduction of AI to aviation, we have chosen to use a sociotechnical framework which focuses on the interactions between the technology and the people who are involved in every aspect of selecting, governing, designing, constructing, operating and maintaining technological systems (Reason 1990). Together, the people and the technology comprise what can be called a sociotechnical system. Macrae's SOTEC framework for describing and categorizing sociotechnical sources of risk integrates structural, organizational, technological, epistemic, and cultural sources of risk. This is not to say that these sources of risk are independent. As Macrae notes 'each of these five domains of sociotechnical risk is deeply interrelated to and constitutive of the others, with the patterns of risk identified here amplifying, reinforcing, interacting and overlapping with one another' (2022, pg 2013). The framework was developed based on a review of autonomous vehicles (Macrae 2022) and has since been applied in the context of AI use in healthcare (Macrae 2024) and robotics (Winter et al. 2024). This work applies the same framework to AI in aviation.

4.2. *Fieldwork Details*

This preliminary study is based on a review of aviation sector public domain publications regarding aviation and AI and interviews with ten aviation sector participants who are either experts in aviation safety or in AI (or both). They are drawn from across the sector (government, industry and academia) and come from a range of academic disciplines. Demographic details are shown in Table 2. The average working experience level of the interviewees is 31 years and the average experience in aviation is 22 years.

The interview guide opened with a conversation on key principles underlying current aviation safety practices and the reasons for their substantial success. Participants were then asked to reflect on how the introduction of AI technologies might require changes in aviation safety thinking and practices for all industry actors. Interviews were approximately 60 minutes in duration and were transcribed for analysis. Approval was obtained from relevant university and national research ethics committees.

Table 2. Interviewee Demographics.

Interviewees	Number	Percentage
Gender		
Male	7	70%
Female	3	30%
Discipline		
Avionics engineering	3	30%
Software engineering	3	30%
Human factors	4	40%
Employer		
Public sector	5	50%
Academic	2	20%
Industry	3	30%
Location		
Europe	8	80%
Other	2	20%

Interview analysis was conducted in NVivo. Transcripts were firstly coded descriptively focusing on participants' descriptions of current core conceptualizations of safety and their articulation of issues in safety raised by the introduction of AI. In a second pass, the material was reorganized into the five SOTEC categories.

5. Findings

5.1. *Structural Sources of Risk*

Structural sources of risk arise from interdependencies and interactions between different parts of the technical and social structures. The overall structure of the sector is well established and not changed in a major way by the introduction of AI. Our participants raised one new risk in this area.

5.1.1. *Increased pace of change*

As already described, the pace of technological change in aviation has historically been slow and incremental and this has been a key factor in the robust safety record of the sector. Several of our interviewees noted that introduction of AI is set to change this longstanding feature of the sector as a result of the introduction of new actors. One interviewee described current practice like this:

[When] the manufacturers ... build a component they spent five years to research, to test it, to qualify. One day it's certified. It goes on the airplane. ... they don't want to touch it unless something really wrong happens ... they're not into a release 5.14 and then 5.15 the next day. (I01)

In contrast, something different is occurring with AI developments:

We expect the pace of innovation to move very fast in this area, just because of the advancement in computer technology. And at the same time people will see the advantage and they start to figure out how to use it in different areas ... we see a faster design cycle for what is currently designed through traditional engineering. (I08)

All the regulators we interviewed spoke of the pressure this brings to their decision making:

But we know also that it's going to have a very long learning curve ... and what I'm ... worried about is that maybe the learning curve will be not in the same time frame as the inventions. They will always bring new versions of the apps ... and we don't even have the time to look back and say what is the assessment in real life of this version. Then you go to the next. (I01)

The risk that we see is the fast advancement of the technology, and sometimes if we don't slow down to make sure it's safe, and if we're facing too much pressure from the market, that might be a risk that we need to learn how to handle. (I08)

The Boeing 737 MAX losses demonstrate what can occur when the scale of change is underestimated and not appropriately reviewed (CTI 2020).

5.2. *Organizational Sources of Risk*

Organizational sources of risk arise from the social processes, organizing activities, and human and contextual factors that underpin new technologies. Organizationally based risk controls are key tools in the aviation sector. Interviewees highlighted two areas in which new organizational risks may arise due to AI.

5.2.1. *Holding Someone Responsible*

When it comes to liability for decisions in aircraft operations, US law states: 'The pilot in command of an aircraft is directly responsible for, and is the final authority as to, the operation of that aircraft' (US CFR Title 14 Part 91.3). In the short term, this will not change with AI applications only providing advice and assistance to decrease pilot workload and provide optimizing suggestions.

In the longer term, some interviewees were of the view that responsibility will shift. In considering where responsibility will lie, it is important to remember that AI is not a person but an engineering component (see also Section 5.3.1) with some unusual features. This means, in the view of some interviewees, that we should *'hold the designer of that engineering component responsible for it. If anybody thinks of AI as a person like a pilot, then it's unclear who is responsible for the extreme case when something is happening to the aircraft.'* (I08)

Other interviewees felt that this was not a reasonable way forward and that *'the human still needs to have safety agency. ... so even if the pilot or the controller ends up in a slightly more supervisory role, they've still got to be in that loop and still driving things.*' (I10)

This has epistemic implications as discussed below.

5.2.2. *Changing Operational Roles*

The role of pilots (and controllers) has been in flux for several decades. As one senior pilot wrote, pilots who were once hired for their judgment are now being evaluated on their compliance (Sullenberger 2009). In some ways, moves towards AI-assisted flight operations are continuing trends introduced by increasing flight automation. As one interviewee explained with increasing automation and ultimately use of AI in cockpits, *'you want people who are not doing much for a long, long time to then know the system well enough to actually go "right, OK, this is the root cause. This is what we need to do"'.* (I05) In response to this, early AI projects are *'trying to maintain real jobs'* (I10) with a focus on ethics and professional dignity although the extent to which this will happen in the longer term is moot.

Several interviewees expressed concerns about how skills will be developed and maintained in this very different role which is based around monitoring and abnormal operations. Not only will it be more difficult to maintain skills, but major changes will be needed in the way pilots come up through the ranks and learn about what it means to have that role. As one interviewee explained in the context of single pilot operations, *'currently, as a new pilot you become a first officer working with a captain. So there's that ongoing training happening. And, of course, once you've got one pilot in the cockpit, then a new pilot has to be at the level of captain before they can sit there'*. (I10)

5.3. *Technical Sources of Risk*

Technological sources of risk arise from the capabilities, affordances, and constraints inscribed into and produced by new material technologies.

Most of the AI applications being proposed in aviation in the short to medium term are forms of machine learning. This introduces both benefits to the industry and to society but also risks.

5.3.1. *Difficulty in Certification*

Many interviewees were of the view that current systems of certification would need radical review if AI applications that include machine learning are allowed in live systems. This interviewee articulated the problem:

'the difficulty I see is that when you when you are working with specification as we are working today it's fixed. The logic is such that what you are certifying today will be the same in 10 years. When you're talking about AI and machine learning what you are certifying is maybe different five years after.' (I03)

In the new world of AI applications:

'Every time you train and fine-tune you've got the potential to create hazards that weren't available, or weren't considered, in the earlier versions.' (I04)

Most interviewees though that some form of certification would still be possible after all, *'just like when we put in a light bulb, there's always a mean time between failure. When we put in a screw, there is certain engineering aspect that we have to know. So, with any machine learning component we expect to know certain characteristics. Whether that characteristic is meeting expectations, that is the job that the person who designed the aircraft [has to decide]. And if he decided to go with some component, then he need to figure out whether that performance can comply with the safety* *requirement that we have'*. (I08) The solution might therefore be to *'constrain the outcomes of the AI so that it fits within the expected parameters of the performance'*. (I01)

On the other hand, one interviewee was of the view that the industry is looking at this issue in the wrong way. They said, *'there's a huge gap here that again is just thought in terms of, "how will we certify AI?" and so on. But to me, the question is, "does it makes sense to try and certify AI?" Isn't there a deep paradox here to try and certify a priory a system that has an evolving behavior by design?'* (I09)

5.3.2. *AI is not Human*

Another technical factor raised by interviewees is the very nature of AI technology. As one interviewee said, AI can be defined as *'mimicking human intelligence at the computational machine level'* (I08) People are frightened of AI but it is not human, it's just a piece of software so regulators *'should be able to verify the performance and certify it'*. The problem *'is how to explain the functionality that cannot be explained in the traditional manner ... we don't know what the output is, but we know that the output will have certain performance characteristics.'* (I08)

One interviewee noted the current paradox that by seeking to make AI function ethically we are attributing human characteristics (having values) to technology. At the same time, there are ever increasing moves to codify and constrain the actions of individuals to remove any component of judgment from their actions, thereby making people more like machines. This view of the relationship between technology and people does not seem to be making the best of the characteristics of either and has been raised previously in the sector as problematic (Prahl, Leung, and Chua 2022).

5.4. *Epistemic Sources of Risk*

Epistemic sources of risk arise from the ways that knowledge and ignorance are constructed in relation to, and within, the new technology. Our participants had a lot to say epistemic risks.

5.4.1 *AI Learning About Rare Cases*

Several interviewees raised concerns about the ability of AI technology to initiate appropriate action or provide appropriate advice in rare but potentially catastrophic situations.

Firstly, most of the data collected in flight operations or air traffic management is normal operations. Several interviewees were concerned about the lack of abnormal operations data to learn from, as one interviewee phrased it when discussing abnormal events *'I don't know how an AI agent will be supporting you because it's never seen one itself'*. (I01)

A second interviewee thought the data problem was even greater as abnormal cases will be dismissed by machine learning algorithms. *'So machine learning is completely different from human learning. ... When we learn a language; we learn how to conjugate verbs. So there are regular verbs that follow a rule that we can generalize, and then there are some special case that we have to memorize. But with machine learning they don't remember these special cases. They just generalize everything into rules. ... It means that machine learning will not learn something new [from an] unexpected function. It will generalize it into one unique function.'* (I08)

Another interviewee raised the concern that *'there are lots of things going on in emergencies that ... you have no sensor to get data from. And so again you would address emergencies through the lenses of the data available.'* (I09) This interviewee implies that embodied knowledge i.e. forms of knowledge that are bodily in character, meaning they relate to a particular pattern of movement or way of perceiving (Maslen 2025), are used for decision making now and would be unavailable to AI.

5.4.2. *AI Specialists Learning About Aviation*

Reflecting the range of views about AI and safety expressed by the expert interviewees themselves, some interviewees noted the wide range of background and experience levels of people coming into the aviation industry as AI specialists. Some of them do not have strong backgrounds in aviation and so have little background knowledge about aviation safety and the processes that currently keep the flying public safe.

One example given is that some AI designers from other sectors don't understand that rare cases are very important in aviation and hence can fail to address this in system design until this is raised by regulators. Consistent with this, other researchers (Rismani et al. 2023) have noted the wide variation in understanding and application of risk principles, safety engineering and professional ethics among AI professionals.

5.4.3. *Changes in System Knowledge*

Another epistemic risk raised by some interviewees is how on the job knowledge and professional learning may be changed by use of AI. There were two examples given of this.

First was the use of AI generated data as a decision making input by designers, controllers and pilots with those users of AI outputs not necessarily understanding the strength of knowledge held by particular pieces of AI-generated advice. As one interviewee described this: *'the controller will still be presented with the same information it's just that it's source will be unknown and it's that variability in the source that I think could in theory increase uncertainty'*. (I01). Similar issues arose in a recent study of automation in the maritime sector and the impact on seafarers (Aalberg et al. 2024).

A similar issue was raised in aircraft design and manufacturing. In this context, questions that were formerly resolved by direct access to an aircraft in the factory may now rely on an AI generated model with an associated loss of competency linked to embodied knowledge in engineering (Maslen and Hayes 2022).

5.4.4. *Learning from AI-related Failures*

Other interviewees raised epistemic risks linked to incident investigation and the perceived inability to determine causation in the way we currently understand it in cases where AI is involved in accidents or dangerous occurrences. As one interviewee described, *'In AI, we know it's almost impossible to backtrack and to understand back what should be changed in the system or in the algorithm, or in the way the system is trained and the machine learning process to eliminate an undesirable outcome'*. (I02)

This could impact the ability to learn from small failures which is critical in safe operations (Weick, Sutcliffe, and Obstfeld 1999).

5.5. *Cultural Sources of Risk*

In the view of our interviewees, aircraft manufacturing is highly controlled bottom-up process of certification against standards for managing risk and reliability so that uncertainty is effectively eliminated. Similarly, safe and efficient aircraft operations is substantially a matter of procedural compliance on the part of flight crew and flight engineers. Air traffic control is acknowledged to be different, described as either 'a carefully choreographed dance' (Sullenberger 2009) or 'barely controlled chaos' (I01).

AI-based technologies will be introduced to an aviation industry with these well-defined cultural norms that will be challenged by AI in at least two ways.

5.5.1 *Work as done?*

Previous research has noted the critical difference between 'work as imagined' and 'work as done' (Hollnagel 2004). We have already noted that currently AI is not permitted to be used in safety-critical aircraft systems and yet one interviewee highlighted that the distinction between safety critical systems and others might not be as hard as assumed. One interviewee described this distinction as naïve: *'You have some systems that are certified, because they're considered to be critical. And they live next to systems that are not certified, for example, pilots are asked not to use [them] as ... primary sources of information .... But obviously since these systems are much more sexy than the others, the practices in the field are completely different. They use sexy systems as a primary source. And even though these sources are not certified according to the whole process. ... that's what will happen.'* (09)

This distinction becomes important when using assumptions about current work arrangements as the foundation for AI development or management of AI rollout.

5.5.2. *A Priori Approaches to Safety*

While some interviewees talk about culture and professional judgment as important factors in aviation safety, by far the majority of interviewees and the majority of discussion centers on attitudes to safety assurance that could be described as positivist, deterministic or a priori approaches. Compliance, certification and risk-based management systems are safety assurance strategies in this tradition. Key largely uncontested assumptions are that the past is a good predictor of the future and that all possible causal chains can be predicted and addressed in comprehensive procedures that eliminate uncertainty and the need for judgement.

One interviewee explained the current situation: *'We're in the same kind of impossible* *situation that was socially accepted for a very long time. Everybody's convinced that if the aircraft is certified, and if you comply with the procedures, it's safe. But if you're a safety management expert, you know that certification only considers a limited number of situations, and you will also know that it's the same for procedures. They are adapted to a limited number of situations. But it's a kind of social hypocrisy that is accepted.'* (I09)

Our interviewees expect that the sector will continue to focus on deterministic safety assurance strategies and try to modify them to suit AI. The extent to which this is feasible and effective, especially at higher levels of automation, remains untested.

6. Discussion

In the views of expert participants in this research, introduction of AI applications poses a series of sociotechnical risks. These are:

- The increased pace of technology change and the need for safety assurance practices to keep up.
- Uncertainty over responsibility for bad decision outcomes.
- Changing operational roles and the associated need for new modes of learning.
- How to certify AI-based applications.
- Confused expectations over the apparent humanity of AI technology.
- How to teach AI about rare cases.
- Bringing AI professionals into the safety culture of aviation.
- Dealing with loss of embodied knowledge in current work practices.
- How to investigate the causality of AI-related incidents.
- Ensuring that AI development is grounded in current work as done.
- Challenges to the current a priori approaches to aviation safety.

There is significant overlap between this list and EASA's articulation of common AI challenges in aviation (EASA 2023), but the sociotechnical perspective provided by the SOTEC framework (Macrae 2022) in conjunction with the views of research participants has highlighted some additional important issues. Introduction of AI may pose significant risks due to interactions with existing industry features such as cultural assumptions about safety assurance, cultural attitudes towards the speed of change, tacit noncompliance built into existing work practices and current uses of embodied knowledge.

While these issues were all raised by safety and AI experts in the industry, and so in one sense are known, it is notable that the identified risks that fall outside EASA's list of AI challenges are all issues that are less visible to the deterministic processes of certification and risk management currently used for safety assurance. This is consistent with Hardy and Maguire's (2020) finding that organizations may struggle to identify and act on novel risks characterized by uncertainty and unfamiliarity because prevailing processes direct thinking into more familiar areas where risks have traditionally been seen to lie.

Effective management of AI risks requires new and more creative ways to find potential dangers before the arise. These new issues are identified by interviewees thinking about an imagined future standing in the shoes of those who will be going about their work of designing aircraft, flying a plane and certifying designs. In taking this creative approach, they have identified cultural pressures which could impact decision making. For successful risk management, existing safety assurance processes will need to be augmented by new forward thinking and imaginative approaches.

There are some significant limitations to this study including the small sample of ten experts interviewed which prevents examination of how views on AI risk might vary between specialist disciplines and across different jurisdictions. Our reported results merge together issues from across different parts of the sector as a result of lack of data but also lack of space. In such a short paper, we have necessarily focused primarily on frontline uses of AI but there are many backroom applications that deserve attention, too.

The potentially catastrophic results of software problems are not unique to AI applications. Algorithmic errors have played a critical role in some previous aviation sector accidents such as the Boeing 737 MAX cases (STI 2020). A comparison between AI risks and previous automation challenges would also provide valuable context for assessing future risks.

7. Conclusion

This brief review of the sociotechnical sources of safety risk introduced into the aviation sector as a result of increasing use of AI-based applications has highlighted eleven risks that need attention as moves to integrate more AI accelerate. Macrae's SOTEC framework (2022) has provided a useful method of analysis, showing that risks exist across structural, organizational, technical, epistemic and cultural domains. Some are well known and have been flagged in EASA's AI Roadmap (2023) but others are new and yet worthy of significant attention to ensure the industry record for safety is maintained in the transition. The nature of the newly identified risks strengthens the case for new thinking in safety assurance for AI.

Future research could consider expanding the study to a larger group of experts, evaluating potential legal frameworks for AI in aviation, further investigation of existing cultural factors including likely interactions with AI and drawing lessons from historical automation failures.

Acknowledgement

The authors would like to acknowledge The Research Council of Norway for funding this research through the Consequences of Fundamental Changes in Risk Regulation (RISKY; project number 315302). Thanks to Martin Inge Standal for support in interviewing. We would also like to acknowledge our interviewees and thank them for sharing their views and experience.

References

Aalberg, A. L., S.M. Holen, T. Kongsvik, and A.M Wahl. (2024). Does it do the same as we would? How trust in automated shipboard systems relates to seafarers' professional identity. *Safety Science 172*:106426. doi: https://doi.org/10.1016/j.ssci.2024.106426.

CTI. (2020). Final Committee Report: The Design, Development and Certification of the Boeing 737 MAX. The House Committee on Transportation and Infratructure.

Downer, J. (2017). The Aviation Paradox: Why We Can 'Know' Jetliners But Not Reactors. *Minerva 55 (2),* 229-248.

EASA. (2023). Artificial Ontelligence Roadmap 2.0: Human-centric approach to AI in aviation May 2023 Version 2.0. European Union Aviation Safety Agency.

Hardy, C., and S Maguire. (2020). Organizations, risk translation, and the ecology of risks: the discursive construction of a novel risk. *Academy of Management Journal 63 (3)*, 685–716.

HLEG. (2019). A Definition of AI: Main Capabilities and Disciplines. Brussels: High-Level Expert Group on Artificial Intelligence set up by the European Commission.

Hollnagel, E. (2004). *Barriers and Accident Prevention*. Ashgate.

Macrae, C. (2022). Learning from the Failure of Autonomous and Intelligent Systems: Accidents, Safety, and Sociotechnical Sources of Risk. *Risk Analysis 42 (9*), 1999-2025. doi: 10.1111/risa.13850.

Macrae, C. (2024.) Managing risk and resilience in autonomous and intelligent systems: Exploring safety in the development, deployment, and use of artificial intelligence in healthcare. *Risk Analysis*. doi: 10.1111/risa.14273.

Maslen, S. (2025). *Learning to Hear: The Auditory Bases of Excellence in Practicing Medicine, Climbing Mountains, Making Music, and Communicating in Morse Code*: Columbia University Press.

Maslen, S., and J. Hayes. (2022). "It's the Seeing and Feeling": How Embodied and Conceptual Knowledges Relate in Pipeline Engineering Work. *Qualitative Sociology 45 (4*), 593-616. doi: 10.1007/s11133-022-09520-8.

Prahl, A., R.K.H. Leung, and A.N.S. Chua. (2022). Fight for Flight: The Narratives of Human Versus Machine Following Two Aviation Tragedies. *Human-Machine Communication 4*, 27-44.

Reason, J. 1990. *Human Error*. Cambridge University Press.

Rismani, S., R. Shelby, A. Smart, E. Jatho, J. Kroll, A. Moon, and N. Rostamzadeh. (2023). From Plane Crashes to Algorithmic Harm: Applicability of Safety Engineering Frameworks for Responsible ML. *CHI '23: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems 2*, 1-18.

Sullenberger, C. B. (2009). *Sully*. Harper Collins.

Weick, K. E, K. M. Sutcliffe, and D. Obstfeld. (1999). Organizing for High Reliability: Processes of Collective Mindfulness. In R. I. Sutton and B. M. Staw (Eds) *Research in Organizational Behavior*. JAI Press Inc.

Winter, P., J. Downer, J. Wilson, D. B. Abeywickrama, S. Lee, S. Hauert, and S. Windsor. (2024). Applying the "SOTEC" framework of sociotechnical risk analysis to the development of an autonomous robot swarm for a public cloakroom. *Risk Analysis*. doi: 10.1111/risa.17632.