5-1-2021

Resilience Engineering's Potential For Advanced Air Mobility (AAM)

Christopher Nemeth

Jon Holbrook

RESILIENCE ENGINEERING'S POTENTIAL FOR ADVANCED AIR MOBILITY (AAM)

Christopher Nemeth
Applied Research Associates, Inc.
Albuquerque, NM

Jon Holbrook
National Aeronautics and Space Administration
Hampton, VA

The national airspace (NAS) will rapidly evolve in the next ten to twenty years. Plans for Advanced Air Mobility (AAM) during that period envision highly automated airspace management systems and electrically powered vehicles. AAM concepts also anticipate limited human roles. The goal of limiting the human role is to minimize the potential for misadventures, yet how the human role is limited needs to be carefully considered in order to also preserve the potential for human successes. The field of resilience engineering (RE) focuses on how systems can change in order to seize an opportunity or withstand an unforeseen challenge. RE methods rely on the use of empirical data to optimize the ability of any system to adapt. RE studies have shown how individual and team initiatives ensure resilient system performance by creating safety through flexibility. Benefits of the RE approach include improved awareness of operational circumstances and how system elements depend on each other, and the ability to allocate limited resources and prepare for surprise. RE offers the ability to account for and incorporate the human role as an essential element in order to ensure NAS systems' resilient performance. Data on the human contribution to safe and resilient system performance, which is termed "work as done," are available but are not being considered as the NAS evolves. We present an approach that describes how use of RE can enable the evolving NAS to adapt, and perform, in a resilient manner.

Incremental advancements in computer software, sensors, energy storage, and electric propulsion are fueling the development of new air vehicles that promise to change the way that cargo and people are moved. Simplified electric vehicles capable of lower noise levels and vertical flight, with lower operating and maintenance costs than today's vehicles, could lead to a vast expansion of opportunities for tasks to be accomplished by flight that are currently accomplished using ground-based systems. Urban Air Mobility (UAM) represents one such opportunity, focused on moving people and goods within and around densely populated urban centers, with the eventual goal of providing the public with airborne personal transportation and cargo services. Services may be scheduled, on demand, or part of an intermodal transportation link, connecting passengers or goods to ground-based networks of road or rail systems. UAM vehicles with electrically powered vertical takeoff and landing capabilities will range from small drones that deliver packages to passenger-carrying aircraft that operate in and around metropolitan areas.

Opportunities to leverage these emerging aerial technologies also exist in non-urban areas, or in support of other missions including longer range regional transport of people and goods using electric vehicles that operate out of more conventional airstrips; air-ambulance and medical transfer services; search-and-rescue or disaster relief operations; power-line inspection or other visual surveillance operations, et cetera. This larger ecosystem, known as Advanced Air Mobility (AAM), would represent a highly complex system of individual vehicles, local fleet operations, and regional networks that must all work cooperatively, not only within the AAM ecosystem, but with connected ground-based systems and any adjacent conventional aviation airspace. As a result, AAM comprises a broad range of stakeholders, living and operating in a wide range of locations with different geographic features, using different classes and sizes of airborne vehicles to accomplish a diverse set of missions.

AAM Challenges and Barriers

Despite ongoing technological advances, the potential benefits of AAM do not come without challenges (National Academies of Sciences, Engineering, & Medicine, 2020). Implementing a versatile advanced aerial mobility system with multiple applications and users is a complex, multidisciplinary challenge. No entity within the US, however, has the mandate to promote the development, adoption, and commercialization of new aviation technologies and applications. Nor does any single entity currently have sufficient oversight or responsibility to effectively make advanced aerial mobility a reality, while maximizing the potential societal benefits. Commercialization of AAM will require clarity from regulators and a timely regulatory progress, to support new flight operation types or applications. Without regulatory certainty, advanced aerial mobility systems may develop in an ad-hoc manner, with private point-to-point systems instead of open many-to-many systems. Closing the AAM business case means lowering "cost-per-mile" to the point that perceived benefits to the consumer make the cost acceptable. Highly trained expert human operators represent a significant cost in today's air operations – a cost that will have to be addressed in AAM. Expanding use cases for aviation across the economy and increasing the scale of airspace activity by orders of magnitude are key components of the AAM vision. More vehicles operating in more densely packed airspace, will require increased data needs to schedule, track, and separate those vehicles. The pilot recruitment and training pipeline is not expected to be able to keep up with the anticipated massive expansion of vehicles and operations. Insufficient pilot supply and high pilot cost are driving demand for increased levels of automation and increased demands on automation capability and reliability.

The ultimate success of AAM will depend on providing benefits not only cost-effectively but also safely. The increased levels of automation in AAM systems, however, create challenges for traditional safety assurance methods. Testing and simulation alone will not be adequate to ensure safety in these complex software-intensive systems, which can fail very differently than the more hardware-based systems of the past for which legacy hazard analysis tools were developed. The demand for automated systems that can learn and adapt will require new methods of certification – methods that address automation capabilities that can change how they perform over time. Despite supply and cost demands to reduce human involvement in the control of AAM systems, to date, humans remain the most capable source of information ingestion, situation understanding, and real-time decision adaptation.

In today's systems, human operators participate directly in the control and safety management of the system. Human roles in AAM will share vehicle control and contingency management responsibilities with automation and will be expected to perform with less training. Although well-understood vehicle control tasks may be simplified for UAM vehicles, contingency responses will still be required when vehicle or infrastructure systems fail, environmental conditions are hazardous, passengers are disruptive, et cetera. Given the dramatic anticipated increase in the number of operations in AAM compared to today, the overall frequency of contingency operations is also likely to increase. One barrier to effective and timely safety management of a complex dynamic system on this scale is identifying, collecting, and analyzing the key system configuration, health status, and performance data from all of the entities that operate in or support the ecosystem. Most plans for enabling automated contingency management involve coding "well-established" contingency management procedures into the automation. However, many of these "well-established" procedures depend upon significant interpretation and adaptation by human operators to be successful. A second barrier is that these "well-established" procedures may not be as "well understood" as some may believe. A third, and largely unrecognized, barrier that applies to both data needs and to contingency management is a barrier that results from how we think about safety. Contingency management is not limited to responding and recovering from anomalies, but also routinely preparing for and preventing them from happening in the first place. Our safety thinking can limit the performance data we choose to collect and analyze (Holbrook, 2021).

Resilience Engineering

Safety and risk management thinking has often led to the assumption that human error was the cause for adverse outcomes, that counting "errors" is a way to limit adverse results, and removal of humans would mitigate this risk. AAM, just like the NAS, is a socio-technical system. In order to be effective, socio-technical systems must reflect intense attention to behavior of operators, users, and maintainers who work as participants in what can be considered a joint cognitive system (Woods and Hollnagel, 2006). Resilience Engineering (RE) has evolved from safety studies over the past 10 years to enable systems in high stakes sectors to anticipate and sustain operation when confronted by unforeseen threats (Hollnagel, Woods, & Leveson, 2006). Hollnagel (*in press*) more recently defined resilient performance as "the ability to succeed under varying conditions, so that the number of intended and acceptable outcomes (in other words, everyday activities) is as high as possible" both in the face of adversity as well as during normal conditions. RE studies collect empirical data on work as it is actually done (rather than as it is imagined), and what goes right, and why, at the system level. Results demonstrate how operators ensure resilient performance, making adaptation possible in the face of complexity. They reveal barriers to cognitive work operators perform and show actual (rather than assumed) system performance and interdependencies among system elements. Methods such as the Resilience Analysis Grid (RAG) (Hollnagel, 2010) can be used to identify opportunities to anticipate, monitor, respond, and learn during routine and exceptional conditions. The RE approach can help to understand the emergent, interdependent, irregular nuances and complexities that can be expected in AAM, because those traits and operator performance data already exist in the manned NAS. AAM requires a valid grasp of how operators create resilient performance. RE makes it possible to develop an understanding of what goes well in the NAS, and how to capitalize on that understanding as an asset for AAM.

Human Performance and AAM

Efforts to accomplish goals in high hazard domains such as AAM include individual behavior as well as macrocognitive activities such as contingency planning that was mentioned previously. Cacciabue and Hollnagel (1995) describe macrocognitive activities as "the cognitive functions that are performed in natural (rather than artificial laboratory) decision-making settings." The macrocognitive view (Table 1) can be used to develop descriptive models of these activities, in order to identify and understand how they occur in the NAS and the AAM.

Table 1. *Macrocognitive Activities* (adapted from: Crandall, Klein, & Hoffman, 2006)

Activity	Description
Naturalistic decision making	Reliance on experience to identify a plausible course of action and use of mental simulation to evaluate it.
Sensemaking/situation assessment	Diagnosis of how a current state came about, anticipation of how it will develop.
Planning	Changing action in order to transform a current state into a desired state.
Adaptation/re-planning	Modification, adjustment, or replacement of implemented plan.
Problem detection	Ability to notice potential problems at an early stage.
Coordination	How team members sequence actions to perform a task.
Developing mental models	Mental imagery and event comprehension, based on abstract knowledge and domain concepts and principles.
Mental simulation and storyboarding	Use of mental models to consider the future, enact a series of events, and ponder them as they lead to possible futures.
Maintaining common ground	Ongoing maintenance and repair of a calibrated understanding among team members.
Managing uncertainty and risk	Coping with a state or feeling in which something is unknown or not understood.
Turning leverage points into courses of action	Ability to identify opportunities, turn into courses of action.
Managing attention	Use of perceptual filters to determine the information a person will seek and notice.

While automated systems can follow rules, human intervention is routinely required at the knowledge level (Rasmussen, 1983). Humans ensure resilient system performance in multiple high hazard settings, including the NAS, as data on manned NAS human performance demonstrate (Weick and Sutcliffe, 2007; Holbrook et al, 2019). Automated system inability to function beyond the rules level has shown how understanding human performance is essential to understand and manage risk in the NAS. Organizations must now consider the interplay of different types of risk. More automation reduces the risk of human errors, most of the time, as shown by aviation's excellent and improving safety record. But automation also leads to the subtle erosion of cognitive abilities that may only manifest themselves in extreme and unusual situations (Oliver, Calvard and Potocnik, 2017). Data that describe desired performance already exist. The research and development literature (e.g., job design, work procedures, standards) already describes aviation performance as it is intended and is routinely accomplished. Even reports of adverse outcomes included in the Aviation Safety Reporting System (ASRS) (Billings et al., 1976) include data on resilient pilot performance.

How RE Can Incorporate Human Performance into AAM

While increased automation is a key tenet of AAM, RE offers some direction in thinking about what kinds of functions we might want to automate. For example, how could automation be used to enhance a system's capability to anticipate unforeseen events? How could automation be used to enhance a system's capability to monitor the environment and its own performance? By recognizing that humans are a source of flexibility and resilience, and not just a source of errors and hazards, we can focus design on how the automation can support human performance, rather than on trying to replace the human or protect the system from the human. Or, in the event that we must replace the human, we can better recognize and understand the range of capabilities that we are attempting to replace. We can start by investing in how we think about safety, which informs the safety data we choose to collect and analyze. However, because data collection and analysis are typically triggered by failure outcomes, we rarely study how failure preparation, response, and recovery lead to successful outcomes. We typically wait for something to go wrong before we start to learn from "what happened." We diligently learn from our mistakes, but do we systematically learn from our successes? The answer, far too often, is "no."

Rethink Safety Policy. When we only analyze data from errors and failures, we are ignoring that vast majority of human impacts on system performance. Without understanding how safety is produced, claims about the predicted safety of autonomous machine capabilities that cannot account for this are inherently suspect. Plans to minimize or even remove the only demonstrated, reliable source of safety-producing behavior, without first understanding the capability being minimized or removed, introduces unknown, potentially unaccounted-for risk.

Collect and Analyze Human Performance Data. Fortunately, there are opportunities to address these risks. Data already exist or could be collected with minimal effort on successful preparation for, response to, and recovery from failure. While we don't have the many decades of experience and infrastructure for doing this like we do for human error, emerging approaches to safety and risk management, such as Resilience Engineering, offer a useful place to begin.

Use RE to build new performance models. Start by broadening the data sources on both desired and undesired human performance. Use those data to distill criteria that will define resilient performance. Validate those criteria through the collection of empirical data using rigorous methods such as analysis of simulator runs. Develop requirements and use cases on human roles in AAM. Build means to learn from experience, and to grow the field of practice.

Summary

The UAM/AAM environment will evolve over the next 10+ years, and prior risk/safety models may not serve this new domain well. Organizations can systematically drive change, which can begin by using tools already at their disposal (e.g., policies, procedures, training, equipment) to effectively translate insights into action. Resilience Engineering offers new means to develop the AAM environment through deep insight into how humans ensure resilient performance. Data exist that RE methods can use to create an effective AAM.

Acknowledgments

This work was funded by NASA's System-Wide Safety Project, part of the Aeronautics Research Mission Directorate's Aviation Operations and Safety Program.

References

Billings et al. (1976, September). *NASA Aviation Safety Reporting System Quarterly Report* Number 76-1, April 15, 1976 through July 14, 1976. NASA TM X-3445. Ames Research Center, NASA, Moffett Field, California 94035. Available at: https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19760026757.pdf

Cacciabue, P.C. & Hollnagel, E. (1995). *Simulation of cognition: Applications.* Expertise and technology: Cognition and human-computer cooperation. In J.M. Hoc, P.C. Cacciabue, & E. Hollnagel (Eds.). Mahwah, NJ: Lawrence Erlbaum Associates. 55-73.

Crandall, B., Klein, G. & Hoffman, R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis.* Cambridge, MA: MIT Press.

Hollnagel, E. (2010). Epilogue: RAG. The Resilience Analysis Grid. In Hollnagel, E., Paries, J., Woods, D. & Wreathall, J. (Eds). *Resilience Engineering in Practice: A Guidebook.* Ashgate Publishing. Farnham, UK.

Holbrook, J. (2021). Exploring methods to collect and analyze data on human contributions to aviation safety. In Proceedings of the 2021 International Symposium on Aviation Psychology.

Holbrook, J. B., Stewart, M. J., Smith, B. E., Prinzel, L, J., Matthews, B. L., Avrekh, I., Cardoza, C. T., Ammann, O. C., Adduru, V., & Null, C. H. (2019). Human Performance Contributions to Safety in Commercial Aviation. NASA Technical Memorandum 2019-220417. Available at: https://ntrs.nasa.gov/api/citations/20190033462/downloads/20190033462.pdf

Hollnagel, E. (in press). From Resilience Engineering to resilient performance. In C. Nemeth and E. Hollnagel (Editors). *Advancing Resilient Performance*. New York: Springer Nature.

Hollnagel, E., Woods, D. D. & Leveson, N. (Eds.) (2006). *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate Publishing.

Klein, G., Ross, K. G., Moon, B. M., Klein, D. E., Hoffman, R. R., & Hollnagel, E. (2003). Macrocognition. IEEE Intelligent Systems, 18(3), 81-85.

National Academies of Sciences, Engineering, and Medicine 2020. *Advancing Aerial Mobility: A National Blueprint*. Washington, DC: The National Academies Press. https://doi.org/10.17226/25646.

Oliver, N., Calvard, T. & Potočnik, K. (2017, Sept 5) The Tragic Crash of Flight AF447 Shows the Unlikely but Catastrophic Consequences of Automation. *Harvard Business Review.* H03W8I-PDF-ENG.

Rasmussen, J. (1983) Skills, Rules and Knowledge: Signals, Signs and Symbols, and Other Distinctions in Human Performance Models. *IEEE Transactions on Systems, Man and Cybernetics.* Vol. SMC-13. No.3. May. 257-66.

Weick, K and Sutcliffe, K. (2007). *Managing the Unexpected: Resilient Performance in an Age of Uncertainty.* New York: John Wiley and Sons.

Woods, D. D. & Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering.* Boca Raton, FL: CRC Press.