The Cyber Security Aspects of Autonomy in Aircraft



Mick Warren discusses his thoughts on how autonomy will impact the cyber security aspects of safety within aviation. Including a brief background to the history of aviation safety, he discusses the challenges use of Artificial Intelligence (AI) bring. He outlines what industry is currently doing to address this, what he has done through research and application on an in-service aircraft and what the potential next steps should be.

Since the dawn of aviation accidents have occurred – from the crash of the first commercial flight in the UK in 1919 killing two to the latest, at the time of writing, of a Beechcraft B200 in Kenya. Safety is paramount as we have seen with recent problemd such as the Boeing 737 Max safety issues. This has cost 346 lives; resulted in the halting of production of the aircraft; closing of factories building the aircraft; grounding of all of the aircraft and reportedly has cost Boeing as much as $19 billion as of Jun 2020.

Historically, accidents tended to be as a result of mechanical failure. Over time, valuable lessons have been learned and the amount of mechanical failures causing accidents has reduced significantly. Accident attribution in more recent years has been the result of terrorism or crew issues. Processes and procedures have developed to mitigate such causes. Is cyber-security the next threat source for aviation accidents? We have already seen cyber-security attacks on aviation in an Airline Management Systems where 1,400 passengers were grounded and the flight plans of 10 aircraft went missing. Luckily, this was a mission impact and not a safety of flight one. However, the potential for cyber-attacks on aircraft was understood in the last century as military aircraft had already been subjected to cyber-attacks by then.

The outcomes of risks emerging from security issues are related to both mission and safety aspects. Mission issues are where the end impact of a successful cyber-security attack may cause loss of a mission; for example, an attack against the inventory management system may mean that an aircraft cannot be repaired and thus misses a scheduled flight. Alternatively, there could be a targeted attack on the inventory management system that affects flight records, potentially causing an incorrect part being fitted to an aircraft, thus resulting in a safety risk. It may also affect more than one aircraft at the same time. A potential safety risk is of particular concern in aviation – at worst case it could lead to a catastrophic accident with significant loss of life.

Cyber security threats are real and are happening within the aviation domain now

In the last 30 years aviation has moved from a traditionally isolated set of on-board and off-board systems into highly interconnected systems (such as Integrated Modular Avionics). This, when combined with more recent developments of highly integrated systems-of-systems (such as the Air Traffic Management network and Controller Pilot Data-Link Communications) and coupled with widespread internet access on-board presents significant potential security issues. Such integration provides potential significant additional threat paths – many of which are outside the control of the aviation domain – that have, hitherto, not been possible. Aviation is vulnerable to cyber security attacks. It is not necessarily that difficult to expose such issues. A security expert, Chris Roberts, tweeted that he could play with Engine Indicating and Crew Alerting System messages. This could have resulted in a significant safety impact from a cyber threat. Luckily, he didn't do it – he was making a point albeit he was arrested by the Federal Bureau of Investigation for his troubles.



Artificial Intelligence Challenges to Aviation safety

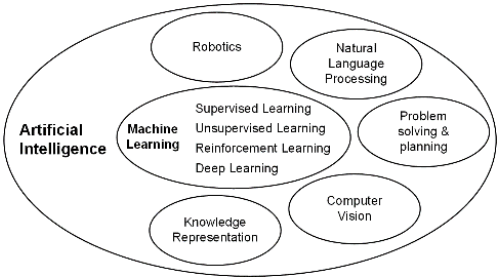
Aviation security & safety faces a new existential challenge – AI. The use of AI – and Machine Learning (ML) in particular – is on the rise, where tasks may be automated through the use of large and accurate input/output datasets. Such intelligent behaviour is already with us in day-to-day use, be it voice recognition on a smartphone, interaction with an automated call-centre, machine translation or search engines. Significant research and investments are being made in fields using ML such as autonomous cars (with $1 billion invested in a single week!), aviation, software fault prediction and electrical load forecasting. The growth is such that what may have been considered intelligent behaviour only 5 years ago is now considered hardly noteworthy.

The use of ML widens the potential cyber security threat landscape and the speed at which attacks can occur

With such growth comes the potential for malicious use – such use significantly impacts the potential threat landscape; moreover, it can significantly increase the speed at which attacks occur. In addition, current growth is such that it will invariably outstrip the ability to either provide certification (as the development of domain standards will invariably lag system development) or appropriate legislation (legislature works even slower than standards development).

What is AI? AI, helpfully, has neither an agreed general definition nor scope. It is also widely used and misused. Standard texts agree that it contains:

- Robotics
- Natural Language Processing (NLP)
- Problem solving and planning
- Computer vision
- Knowledge representation and
- Machine Learning (ML) (which has sub-sets as shown).



AI is important to the Aviation sector – it may: reduce cost (through quicker design); improve performance (through improved engine fuel burn); introduce new markets (such as local Unmanned Air Vehicle (UAV) deliveries); increase payload (by reducing the need for multiple pilots) and, eventually, provide fully automated passenger aircraft (by removing the pilot).

The use of AI may present particular challenges – requirements being one of them. DO-178C (RTCA, 2011) – Software Considerations in Airborne Systems and Equipment Certification – relies on the system safety and hazard analysis to determine system failure conditions and then applies a rigour to the software development process commensurate with their perceived effect(s). Is this feasible for AI? For a start it assumes that requirements can be traced through decomposition to design to determination and evidence. This would be a challenge for AI as not only are requirements definition, decomposition and traceability a significant challenge, there is also the potential for complex unexpected behaviour – this is often a challenge for 'normal' software. Thus, AI presents particular challenges when demonstrating its acceptable safety. Given the potential breadth of AI, a decision was made to focus on the currently most widely used sub-set of AI – that of ML.

What has the Aviation Industry done about it?

The Aviation industry is still struggling with the introduction of multicore processing into safety critical systems never mind AI. Thus far, they have concentrated on the potential safety risks associated with cyber-security threats. DO-326A / ED-202A (RTCA, 2014) – Airworthiness Security Process Specification – was issued on 6 Aug 2014. It has only recently been included within Amendment 25 of European Union Aviation Safety Agency (EASA) Certification Specifications (CS) under CS25.1319. It was previously a Special Condition for Type Certificate (TC) applicants and Supplemental Type Certificate (STC) applicants. It is incumbent on TC and STC applicants to demonstrate that security risks have been identified, assessed and mitigated appropriately. This is managed as a certification activity although the potential severity of the impact of a cyber-security threat to a failure condition comes from the system safety analysis, as enacted through the use of industry standard approaches

such as Aerospace Recommended Practise (ARP) 4754A (SAE International, 2010) and ARP4761 (SAE International, 1996). In simple terms, it has its origins in safety but is then dealt with from a certification perspective. As we move from our current paradigm into a world where ML becomes more prevalent, there is a need to ensure that the potential threats that ML may pose are recognized and appropriately mitigated.

Current aviation security standards do not consider potential impact of Machine Learning

DO-326A and supporting standards do not even mentioned ML. Arguably, if the threat resultant from ML is identified then application of DO-326A may address potential threats. It is identifying the potential threats resulting from ML and the provision of assurance on the associated risk that are the challenges.

What have I done about it?

In order to better understand where DO-326A may be adapted to explicitly consider ML and the challenges it presents, I took the following approach to help facilitate the development of an assurance case:

- Decomposed the standard and its supporting standard DO-355 – Airworthiness Security Methods and Considerations – into a Goal Structuring Notation (GSN) argument
- Rebuilt the GSN Goals and Solutions using a well-known assurance approach – that of the 4+1 Principles (R. Hawkins, I. Habli and T. Kelly, 2013). This was undertaken from three differing viewpoints: safety, cyber-security and ML. This was extended with the 4+2 Principle as developed by Dstl (R. Ashmore and E. Lennon, 2017) (which is due for inclusion in the latest issue of Def Stan 00-055). This resulted in an adapted 4+2 set of principles from a cyber security perspective
- Reviewed literature to develop 'particular challenges' that ML may introduce, convert them into GSN and adding them where appropriate in the new 4+2 structure
- Took the developed assurance case in GSN, validated it, and applied it on an actual in-service aircraft.

This approach readily allows the assurance case structure provided through application of the Principles to create a link between the potential cyber security aspects of ML, which, with the extension mentioned above, allows ML specific cyber security issues to also be fed back into the safety domain within the context of a safety-critical system. This provides a link which is presently not there.

A number of 'particular challenges' were identified which would need to be considered specifically for ML as part of any assurance argument. It is worth a brief exploration into some examples of the 'particular challenges' discovered during the work. These are:

- **Aircraft specific ML threats affecting safety of flight have been appropriately considered.** Arguably, in an assurance case all potential threats need to have been considered. As the use of ML is emergent then specific attention was deemed necessary.
- **Aircraft specific ML threats which may be architecturally mitigated have been appropriately considered.** One would expect them to have been mitigated as part of any design. However, there are two potential issues here. Most designs are adaptations/evolutions of in-service equipment and the failure, from a safety viewpoint, may

have been mitigated from an unintended perspective only; questions may arise as to the frequency of intentional attacks and the ability of current architecture mitigations to address potentially significantly increased probabilities;
- **Maintenance specific ML threats have been appropriately considered**. On-aircraft systems and equipment tend to be tightly regulated; the off-aircraft systems – such as maintenance laptops – that may connect to the aircraft less so. These form part of the aircraft security environment which may readily introduce significant threats to on-aircraft systems. Discussion later on shows that this is of concern in a real-life application.

Application on in-service aircraft

At the time of development of the assurance case approach there were no known instances of systems on-board aircraft that are labelled as using ML. There is significant work underway on how systems may use ML. This includes: data analysis on Health and Usage Monitoring Systems; Category IIIc Instrument Landing Systems where the aircraft lands autonomously with pilot supervision and Air Traffic Management systems.

A system on-board a typical aircraft that may be subjected to ML-like off-board analysis is the Engine Health Monitoring (EHM) Systems. EHM uses sensors on-board the engine and its ancillaries to monitor and then record the status of engine systems and components. Its aim is to allow the early detection or indications of defects before they cause safety issues. The information is stored and downloaded – on some systems it is actually transmitted, real-time, back to the manufacturer – thus introducing potentially significant additional threat paths. Off aircraft analysis takes the significant amount of data produced and performs complex analysis on it. The analysis can help predict early onset failure of components. This is not strictly speaking ML, but the complexity of the trending and monitoring on the data has ML like tendencies. Indeed, work is ongoing in the use of ML to perform such data analysis.

The totality of the assurance case approach was applied on an in service aircraft. This included an assessment against the EHM system. The main conclusions discovered from the application were:

- The argument structure provided a robust methodology to develop an assurance case to include the 'particular challenges' that an aircraft may face from ML. It did not solve the fundamental problem of how certain evidence may be developed – that was beyond its remit;
- There are weaknesses in DO-326A and DO-355 which need to be considered when applying it. For example, it was noted that the standard focussed on the word 'adequate' for test procedures. Whilst not necessarily an issue – as it is an aim to be demonstrated – the use of ML potentially challenges this – what is 'adequate'?

Next Steps

I have taken the work above and developed templates and guides on how to implement DO-326A, DO-355A and DO-356A. These are presently being trialled on differing projects to refine further. These will then act as the basis for a number of aircraft in order to demonstrate compliance with the appropriate certification specifications – these specifications are the same as those used in civil industry. This work has resulted in an understanding that more effort needs to be undertaken to ensure that the dynamic nature of cyber-security and the static nature of safety analysis is further considered. More importantly, the potential use of ML may challenge the cyber security aspects of the air domain from a developmental, assurance, integration, testing and safety assurance time perspective – it needs to be done significantly quicker than current process facilitates as it is sensitive to real-time changes to the threat landscape.

It is possible to develop an assurance case that considers the use of ML

References

[1] RTCA, DO-178C - Software Considerations in Airborne Systems and Equipment Certification, Washington, DC, USA: RTCA, 2011.

[2] RTCA, DO-326A / ED-202A - Airworthiness Security Process Specification, Washington, DC, USA: RTCA, 2014.

[3] SAE International, ARP 4754A - Guidelines for Development of Civil Aircraft and Systems, Warrendale, PA, USA: SAE International, 2010.

[4] SAE International, ARP 4761 - Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipments, Warrendale, PA, USA: SAE International, 1996.

[5] R. Hawkins, I. Habli and T. Kelly, "The principles of software safety assurance," in 31st International System Safety Conference, Boston, MA, USA, 2013.

[6] R. Ashmore and E. Lennon, "Progress Towards the Assurance of Non-Traditional Software," in Twenty-fifth Safety-Critical Systems Symposium, Bristol, UK, 2017.

Image Attribution
lead image: Crown Copyright 2021 (https://www.defenceimagery.mod.uk/fotoweb/archives/5042-Downloadable%20Stock%20Images/Archive/RAF/2021/March/45167916.jpg)
Visualization of interlinked systems in the civil aviation domain showing several potential vectors for cyber-attacks: Gontar, Patrick and Homans, Hendrik and Rostalski, Michelle and Behrend, Julia and Dehais, Frédéric and Bengler, Klaus Are pilots prepared for a cyber-attack? A human factors approach to the experimental evaluation of pilots' behavior. (2018) Journal of Air Transport Management, vol. 69. pp. 26-37. ISSN 0969-6997

Mick Warren, Frazer-Nash Consultancy Ltd

Mick Warren is a highly experienced and qualified systems safety professional who has undertaken aircraft, submarine and ATM safety work. He also specialises in software assurance and, more recently, has returned to his 'roots' undertaking cyber security work having completed a research-based MSc in Cyber Security focused on DO-326A. He has been the lead in developing a Cyber Security for Airworthiness approach (against DO-326A) where he now leads a team applying this on several aircraft to determine the potential safety impact of cyber security threats. All opinions expressed in this article are entirely his own and do not reflect the views of his employer.

The author retains copyright of this article.