The Autopilot Problem

By Amanda Draeger

April 25, 2023

Amanda Draeger will be speaking more about talent development at HammerCon 2023 on Thursday, May 18, 2023 - get your tickets today!

In 2009, AirFrance Flight 447 disappeared from the skies over the Atlantic Ocean. Like many aviation disasters, there was no one single cause for the crash; the causes included the functioning of the aircraft and the reaction of the crew. However, this was a groundbreaking incident because it highlighted the impact of the interaction between the crew and the automation technology that helps them to operate the aircraft.

During the flight, the pitot tubes, which measure airspeed, on the aircraft iced up. As a result of this, the aircraft transitioned from “normal law” to “alternate law”. The pilots

correctly recognized that, as part of this transition, autopilot disengaged, and they resumed manual operation of the aircraft. What the pilots did not know, because it was not part of their certification to operate the aircraft, was that some other systems, such as the one that would work to prevent stall, also disengaged.

This lack of awareness of what automation systems were currently engaged, coupled with the startle effect, number of alarms in the cockpit, and unreliability of cockpit indicators, led the pilots to make poor decisions. They operated the controls as if those other systems were still in place. This led to the aircraft entering stall, which the pilots failed to recover from, resulting in the deaths of everyone on board.

The BEA, the French Civil Aviation Investigation Authority, investigated the incident and issued several interim reports before its final report. One of its recommendations regarded the knowledge of pilots on not just how to operate their aircraft, but how flight works:

In the absence of any reliable speed indications, understanding of the overall physics of flight at high altitude could have considerably helped the pilots to anticipate the rapid degradation of the situation. The same applies to the overspeed phenomena that have evolved with modern aeroplanes.

Consequently, the BEA recommends that:

EASA define recurrent training programme requirements to make sure, through practical exercises, that the theoretical knowledge, particularly on flight mechanics, is well understood. [Recommendation FRAN-2012-041] [1]

Why Study Aviation?

You may be wondering what a 14-year-old aviation disaster has to do with computer security. The simple answer is that computers are everywhere, so computers affect everything.

The slightly longer answer is that, unlike computer security, aviation is a field that has a robust structure for studying how systems in production fail. Organizations like the BEA

robust structure for studying how systems in production fail. Organizations like the BEA and NTSB have decades of experience conducting post-mortem investigations to identify the variety of causes of failure and making recommendations to prevent this failure from happening in the future. This means that aviation disasters, as well as near-misses, are often thoroughly studied and documented. That is, there is a large body of work that people from other fields can reference when learning about how systems fail.

That said, change is happening. In 2021 President Biden signed the Executive Order on Improving the Nation's Cybersecurity, which established the Cyber Safety Review Board [2]. The board's goal is to "focus on learning lessons and sharing them with those that need them to enable advances in national cybersecurity" [3]. They were envisioned to be an organization that looks and operates very similar to the NTSB, and their initial report, on Log4j, reflects this.

What is the Benefit of Autopilot?

When looking at the example of AirFrance Flight 447, you may be wondering why airplanes have features that can disable themselves. That is, why have autopilot in the first place if it does not always work?

In a system that involves people, people are often one of the less deterministic parts of the system; we do not always behave consistently. Our judgement becomes poorer when we are hungry, tired, overworked, or upset. We can be distracted when there is an alarm going off, or are unable to identify all of the multiple alarms sounding at once. There are also limitations to our physical senses that mean we can't see through dense fog or directly read the beacons that aircraft use to broadcast their locations.

Computers in general, and assistive technologies like autopilot specifically, help people to overcome those limitations. There are some types of functions, like mathematical calculations, that computers will usually perform more accurately than most people. Other functions, like reading data from sensors that can detect things that human senses cannot, can only be done by computers. And, in general, computers are much more deterministic than humans. However, their behavior is not entirely predictable. They are subject to power surges, bit flips, resource exhaustion via memory leak, but will suddenly

behave correctly when there is a technician nearby.

As an example, in January, two Alaska Airlines flights had tail strikes (their tails scraped the runway during takeoff) within six minutes. The problem was quickly traced to the software that calculated the takeoff weights of the aircraft [4]. This is the sort of function that is very good to have a computer do: it is relatively easy to program as a "mechanical" process, but people may be inconsistent when performing that sort of math. Unfortunately, as this incident showed, even the simplest software can fail.

Things get more complicated when adding more advanced computation, such as those pieces of software being sold as artificial intelligence. As they are designed to more closely replicate human methods of learning, they also have some of the human tendency towards a more nondeterministic nature. Two people asking the same question of an AI-powered chatbot may receive wildly different answers.

While software is not infallible, it can still be better than having people performing the same functions. A person who is able to focus on only the tasks that software cannot do well will have a lighter cognitive load than a person who is trying to do everything. This means that, if and when the software fails, that person is not already fatigued; they are generally better at making good decisions. We benefit greatly by offloading as much as possible to computers, and having people focus on the tasks that people do better than computers.

This is the approach that has guided aircraft manufacturers when designing their products: build technologies that focus on doing the core tasks well and leave the edge cases to the pilots. As has been hinted at with AirFrance Flight 447, there is more than just autopilot. There can be entire software suites of assistive technologies that integrate data from sensors all over the aircraft to make it easier to operate. But these software suites can sometimes have too much control.

What Happens When the Software Takes Command?

The two 737-MAX crashes in 2019 were a chilling reminder of the fact that it is possible for software to have too much control. To compensate for the significantly different

aerodynamics caused by needing to mount larger engines more forward on the wings, Boeing created some software (called “Maneuvering Characteristics Augmentation System”, or MCAS) that would make it seem to the pilot like nothing had changed from previous 737 versions.

MCAS worked by reading the state of the angle-of-attack sensor (which was one of the changes that came about from AirFrance Flight 447, so the pilots had a better sense of whether the aircraft was flying level, nose up, or nose down) and pulling the nose of the aircraft down if it thought the plane was likely to stall. Unfortunately, the flight management computer was not programmed to check the sensors on both sides of the aircraft; it would only take readings from one side. This is a problem when dealing with a sensor that is exposed to extreme temperature shifts, as sensor failures happen.

Without assistive software, the pilots would be able to check readings across the two sensors and determine if a sensor was likely malfunctioning. There are also other instruments the pilots can check to determine pitch. They can also, when visibility permits, simply look out their windows. With MCAS, the plane would read a single faulty sensor, take it on blind faith that it was correct, and push the aircraft nose down, to its doom.

Crucially, MCAS was still engaged even when the pilots disengaged autopilot. This may not have been terrible on its own, but Boeing worked to keep the 737-MAX certified as any other 737. As such, pilots did not have to undergo any additional training to operate the aircraft. They not only did not know MCAS remained engaged when autopilot disengaged, but they did not know MCAS existed at all. It is very difficult, if not impossible, for an operator to compensate for a technology that they don’t even know is part of the equipment they are operating.

The Autopilot Problem

This is the autopilot problem: How do you know when the computer is giving you a wrong answer? Ideally, you know because you have enough experience to know what the right answer looks like, but this assumes you are paying attention to the answer you’ve been given. In the case of the Alaska Airlines flights, the crew should have noticed that the weight estimates were about 20-30,000 pounds lighter than normal. Why didn’t they?

weight estimates were about 20-30,000 pounds lighter than normal. Why didn't they?

To speculate, one reason why the crew did not recognize that the calculations were off is because they've gotten used to that software always being correct. They have grown complacent. Another reason is that they've never had to perform those calculations manually. Even if there was no expectation that performing those calculations would be a daily part of their job, learning what goes in to them would give a much more visceral sense of what right looks like.

As an example from the military: in 2015 the United States Naval Academy reintroduced celestial navigation as part of its curriculum. The technique had been dropped due to the ubiquitous nature of GPS and other navigation systems. However, the possibility of electromagnetic attacks against GPS means that it is vitally important to know how to navigate without such technology [6]. And, as BEA reported on AirFrance Flight 447, there is value in understanding the theoretical knowledge. Even if the students do not use celestial navigation on a regular basis, having the knowledge of how it works means that they are more likely to be able to tell when their GPS is giving an implausible result.

As an example from the military: in 2015 the United States Naval Academy reintroduced celestial navigation as part of its curriculum. The technique had been dropped due to the ubiquitous nature of GPS and other navigation systems. However, the possibility of electromagnetic attacks against GPS means that it is vitally important to know how to navigate without such technology. And, as BEA reported on AirFrance Flight 447, there is value in understanding the theoretical knowledge. Even if the students do not use celestial navigation on a regular basis, having the knowledge of how it works means that they are more likely to be able to tell when their GPS is giving an implausible result.

Psychological Safety

In a healthy workplace, people feel comfortable enough to speak up when they see that something bad is likely to happen. There is a well-documented history of what happens when air crews and ground crews do not experience this level of psychological safety [7]. But we also have examples of what happens when these crews have it.

During the Alaska Airlines incident, Bret Peyton, the on-duty director of operations, called

During the Alaska Airlines incident, Bret Peyton, the on-duty director of operations, called a nationwide halt based only on the two tail strikes. This was not something that he did lightly, but he was empowered to do what he felt was right based on his assessment of the risk. With he 737-MAX, the pilots were not able to take any action because they were not empowered; they did not even know what software was causing the aircraft's behavior, much less how to override it.

In a world where technology is making decisions, workers need to have not only the psychological safety to speak up against their coworkers, but also to speak up against their technology. They need the ability to choose whether or not they will accept the decision a computer makes. In the famous words of a 1979 IBM slide, "A computer can never be held accountable. Therefore, a computer must never make a management decision" [8].

One of the primary benefits of computers is that they are often able to react more quickly than a person. However, outside of combat, rarely are things quite as time-sensitive as we tend to think they are. In aviation-speak: you can always make another pass.

Key Takeaways

Automations make our lives easier in thousands of ways. However, they sometimes fail. This does not mean that we should never use these automations. It means that we need to properly plan for them failing. In practice, this means we need to not only train personnel on the operation of the technology, but educate them on how the technology works. This education will better give them the ability to know when the technology is not behaving correctly, and how to perform the task manually.

In addition to initial education, we all need to invest in our own development. We need to maintain currency, to continue practicing our crafts to ensure that we can perform the full range of tasks that we may be expected to do. In the case of pilots, this means spending time in simulators where they can focus on the edge cases that require a person's judgement over the computer's programming. For those in computer security, it means things like practicing performing packet analysis manually, or looking at the Windows registry without tools that automatically pull out the keys that are usually interesting. This

takes time, and will be slow, but ensures that you can continue operating in the middle of an incident if your tools no longer work.

Speaking of tools: understand how they work. Verify that they perform the tasks that you think they do, preferably by comparing to manual execution of that task. Check to see if they have any features that change how the tools behave under certain conditions. Validate your assumptions.

Finally, cultivate psychological safety in your teams. Make sure that your people feel like they have the ability to speak up when they see something that is not quite right, whether it is a person's behavior or a computer's behavior.

If you keep these tips in mind, and implement them on your teams, you can hopefully learn from aviation that while automation is good, it is not perfect. We need to ensure that we are never so dependent on our automations that we completely give up control of the aircraft.

References

[1] Bureau d'Enquetes et d'Analyses. Final report on the accident on 1st June 2009 to the Airbus A330-203. 2012. https://bea.aero/docspa/2009/f-cp090601.en/pdf/f-cp090601.en.pdf

[2] The White House. Executive Order on Improving the Nation's Cybersecurity. 2021. https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/

[3] Cybersecurity & Infrastructure Security Agency. Cyber Safety Review Board (CSRB). n.d. https://www.cisa.gov/resources-tools/groups/cyber-safety-review-board-csrb

[4] Gates, Dominic. After Alaska Airlines planes bump runway while taking off from Seattle, a scramble to 'pull the plug'. Anchorage Daily News. 2023. https://www.adn.com/alaska-news/aviation/2023/02/20/after-alaska-airlines-planes-bump-runway-a-scramble-to-pull-the-plug/

[5] Travis, Gregory. How the Boeing 737 MAX disaster looks to a software developer. IEEE Spectrum. 2019. https://spectrum.ieee.org/how-the-boeing-737-max-disaster-looks-to-a-software-developer

[6] Brumfiel, Geoff. U.S. Navy Brings Back Navigation By The Stars For Officers. NPR Morning Edition. 2016. https://www.npr.org/2016/02/22/467210492/u-s-navy-brings-back-navigation-by-the-stars-for-officers

[7] Psychological Safety. Psychological Safety in Aviation. 2023. https://psychsafety.co.uk/psychological-safety-aviation/

[8] Goetze, Trystan. Responsibility and Automated Decision-Making. Blog of the APA. 2023. https://blog.apaonline.org/2023/04/13/responsibility-and-automated-decision-making-draft/

About the Author

Amanda Draeger is currently the Sergeant Major of the Army Cyber Institute. She holds a Masters of Science in Information Security Engineering from SANS Technology Institute, is a GIAC Security Expert (GSE), and a member of the Sergeant Audie Murphy Club.

The MCPA thanks its national sponsors: Capitol Technology University, Google Cloud, Lockheed Martin, Virginia Tech's National Security Institute, NYU's Center for Global Affairs, and Cisco.