

AI in Aviation

Andrew Lea MA(Cantab) FBCS FRSA



Contents

1. Introduction
2. A crash course in Artificial Intelligence
3. Safety Neutral AI
4. Safety Enhancing AI
5. Safety Risk AI
6. Summary and Conclusion

1. Introduction

This paper considers Artificial Intelligence in Aviation: how it can help, and how it can do so safely. We assume our reader is familiar with commercial or general aviation, but not particularly familiar with writing software or how artificial intelligence actually works.

Whilst it is obvious that aviation is an expert's domain, it is not so obvious that artificial intelligence is too. Many people, familiar with using software, believe AI is simply "clever software" and that they are therefore well-placed to make decisions about it. Misunderstanding the implications of the particular characteristics of AI, or even not realising that AI *has* particular characteristics, may lead to incorrect - even dangerous - regulatory or design decisions.

For our discussion, the crucial attribute of aviation is that it is a safety-critical activity. The safe nature of airline travel is not a reflection of aviation being intrinsically safe, but on the extraordinary efforts of accident investigators and manufacturers that have made it so. Much of this discussion could apply to self-driving cars, driverless trains, and crewless ships.

We will take a brief crash course in the characteristics of AI. We will then look at how AI can be used in aviation, and the particular considerations needed to make sure it is safe, by considering three categories of AI in aviation: safety neutral, safety enhancing, and safety risk. Finally, we will summarise with some conclusions as to how AI should be used.

Disclaimer

This article is, essentially, an opinion piece. If you use any of the ideas or techniques expressed or explained, you must perform your own evaluation and testing as to their validity and safety. No representation or warranty is given that any of the ideas or techniques are fit or safe for any particular purpose.

About the Author

Andrew Lea has worked in the exploitation of AI for several decades in many sectors, including aerospace, law enforcement, industrial, marketing, financial and the the public sector. Although not a commercial aviation professional, he has a private pilot's licence and is a paragliding coach. Andrew is extremely familiar with all these AI approaches, in both their use and implementation.

2. A crash course in Artificial Intelligence

This brief primer should give you enough understanding of “AI” to appreciate the discussion on its use in aviation.

We will introduce *specific* example AI techniques, in a box like this, either as we encounter their application, or as needed for understanding.

What is Artificial Intelligence?

For our purposes, let us agree to regard AI as “a computer system which undertakes a task which, were it be undertaken by a human, *would* require reasoning”. This is similar to the definition adopted by a well-known regulator, but which disastrously omits the requirement for intelligence, and thus misses the point entirely. We will, reluctantly, put discussions of artificial sentience and similar fascinating ideas to one side.

It's helpful to know that an **algorithm** is a method of calculating an exact and correct answer. A **heuristic** is a rule of thumb, a method of calculating a good, but not necessarily perfect, answer.

One difficulty the layman faces is that, sadly, the software industry labels nearly everything as “AI”, because it has become the buzz-word de jour. An autopilot that “intelligently” follows a route is neither AI nor intelligent - it's just very good software, written by intelligent capable people.

Another difficulty is that the phrase ‘Artificial Intelligence’ is an oxymoron. Although artificial, an aeroplane flies; it does not ‘artificially’ fly — the flying is real. Just so with intelligence: if it is exhibited, it's real, and if it's a machine exhibiting that intelligence, it is no less real. Calling it *artificial* invokes a cognitive dissonance. A better term would have been “machine intelligence”.

Characteristics of Artificial Intelligence

AI is not just “clever” software: the very characteristics that make it successful are also those which differentiate it from ordinary software, and which demand unusual care in safety-related areas. These characteristics are that **AI may...**:

- **give different answers to the same question *even* under the same circumstances.** It is non-deterministic, and with AI this is not an error. Many important techniques need randomness to work, so output varies even with the same input. Formal logic methods can prove the predictability of deterministic programs, but not of AI.
- **be intended to work in new circumstances:** both a boon, because it can work in unforeseen circumstances, and a curse because full testing is impossible.
- **learn from experience:** this is incredibly useful, but this ability to autonomously change its responses based on what it learns makes full testing difficult if not impossible.
- **teach itself how to think** (technically, it may infer its own knowledge or reasoning representation) so we may not know how it works.

We will call software with *any* of these characteristics as AI, and other software as “conventional”.

Knowledge in AI

We can also think of AI as “the acquisition, manipulation, and exploitation of knowledge”. All AI contains or acquires knowledge, which can come from three possible sources:

1. written in the code itself by the designers
2. explicit knowledge, carefully gleaned from experts.
3. learnt through training or trial and error. Not all AI is machine learning.

Knowledge and explainability

When validating AI we must validate both the knowledge and the underlying software, so it is highly desirable that safety-critical AI can *explain* why it made a particular decision - that it be explainable. Some techniques, even some learning techniques such as **decision trees**, lend themselves well to explainability.

Knowledge without explainability

Other machine-learning techniques, such the popular “**neural nets**” are intrinsically obscure, and cannot explain “why”.

We could simply insist on explainable-only techniques, but unfortunately sometimes non-explainable techniques may be more accurate or useful.

“**Large Language Models**”, such as the ubiquitous chat engines, are based on neural nets. Consequently they too, whilst often generating useful text, cannot explain why an answer was generated. It is easy to get a large language model to list the critical factors in a situation, but they are frequently wrong.

That is the end of our AI primer. The next three sections look at the application of AI to aviation in safety-neutral, safety-enhancing, and safety-critical contexts.

There is more on “Why is AI so hard to define” at <https://www.bcs.org/articles-opinion-and-research/why-is-ai-hard-to-define/>

Decision Trees learn a branching set of rules from pairs of situation-and-answer examples. Each rule is a question (eg “is the airspace class A?”) and either an answer (eg “VFR flight not allowed”) or links to deeper rules.

Later on they can predict the answer to a situation (eg “I’m in class A airspace can I fly VFR?”) by following those rules from the top until an answer is reached.

By reciting the rules they followed, decision trees can explain how they got to an answer, and can cope with uncertainty.

A decision tree is similar to the process humans use when diagnosing a problem. Why has the engine failed? Is the tank empty? No - in that case, is the fuel pump off? No - is the engine oil temperature in the red? Yes - oil pressure down? Yes it is - we have an oil leak.

Artificial Neural Nets (ANNs or sometimes just Neural Nets) are inspired by the connections in a human brain.

They learn from examples how to recognise a pattern, by adjusting the weight or importance of each connection. This training can be very expensive, starts with random guesses, and uses techniques such as “back propagation” to gradually improve those weights.

Neural nets recognise things in images in a similar way to people. The pattern in the sky fits clouds I have seen before, and was told were cumulous.

Although they cannot explain their reasoning (I don’t know how my brain recognised cumulous, either), they are still the foundation of many machine learning applications,

Large Language Models, also known as “**LLMs**” or **Generative Models**, are trained by reading, basically, the internet. They are huge, and learn, for a list of words, what words tends to follow next.

They work by predicting, given the preceding words, what is the most likely word to follow. In other words, what would the man on the internet probably say next.

We have similar patterns in our heads, too. If I asked you what word follows “the delay was caused by bad” you would probably guess “weather” - its the most likely.

Like many people, when they don’t know, LLMs they make it up. In fact, this is what they do all the time, it’s simply that what they make up is so plausible it seems to be useful.

3. Safety Neutral AI

In this category, errors have no safety implications, and are most likely to manifest as inefficiencies. Possible non safety-critical activities include airline commercial and operations, as well as aircraft design by manufacturers.

Fraud Detection

Machine learning **classification systems** can detect fraud or other financial loss in the supply chain or customer purchases (e.g. with stolen or forged credit cards, or loyalty-card theft). For some techniques, we can inspect that learning to find out how to adjust the business to make it less vulnerable to fraud in the first place.

Customer Feedback

Open question surveys can uncover new, unexpected issues, so are more likely to yield helpful insight than “rate your satisfaction with...” surveys. This can be achieved with AI which understands natural language or large language models.

Customer Interaction

It is tempting to hook up a large language model to your web pages, and let your passengers chat away. However, if that system makes assertions to an airline’s commercial detriment, courts have (very reasonably) ruled that the airline is contractually bound by them. LLMs do just make things up, especially if they are stuck, because they are *chat* systems. Whilst safety-neutral, such an AI deployment would be a commercial risk.

Marketing

Recommendation systems can generate propositions which, by being focused on the individual passenger’s likes and dislikes, are very much more likely to be taken up.

Classification Systems

Many techniques exist to learn from example classifications, and then classify new instances. For example, you probably use a “spam filter” in your email: it gradually learns what is or isn’t spam.

Humans classify text by looking at the contents. Oh, this e-mail contains “double-glazing” - it must be spam. Of course, it’s really more complicated than that, depending on combinations of phrases.

Similarly, classification systems regard the incoming instance as a collection of “features”, such as a word or phrase, and learn what combinations of features are associated with which classification.

The techniques vary from simple (**nearest neighbour**: which example is most similar?), through probability statistics (**Bayesian classifiers**) and **decision trees** to the complicated **support vector machines**.

Recommendation Systems learn what customers like and dislike, combine this with (anonymised) data of what similar and dissimilar customers also like and dislike, and come up with very powerful recommendations, typically of what to buy next.

This is not dissimilar to how people recommend films, for example. “You also like adventure films and dislike romances? In that case you would probably like ...”

Operations Management

An operations graph can answer general “what do we need to know” questions. These might be quite complex, such as “which parts can repair this fuel-line, and are there any suitable parts where we already have an aircraft flying from, preferably from stores where we have several, or if not in a location where we will need to ship other parts from to here for our other aircraft needing maintenance?”.

A **graph** is a network of noun-link-noun, which can represent most complex knowledge domains. This tiny example tells us that G-OWAK is at JFK, but unfortunately has a fuel-line fault, but the repair part is at Heathrow:

```
“G-OWAK” -is-at-> “JFK”      “LHR”
      |                   |
      has-fault            is-at
      |                   |
“fuel-line” <-repairs- “fuel line part #123”
```

Powerful queries rapidly explore huge graphs.

Maintenance Optimisation

A probabilistic model of aircraft-component failure could optimise maintenance schedules, both in reducing aircraft outage and in minimising the probability of needing to ship parts to remote airports. These techniques are well understood, having been applied in the electricity industry four decades ago. Some techniques could also explain not just when a piece of equipment will need overhaul, but what the factors are that shorten its life-time or time between maintenance.

Optimisation is an AI strength, and it has the same meaning in AI as it does commercially - minimising cost (time, distance, fuel) or maximising benefit (profit, aircraft utilisation).

An AI technique, like a person, might start with an imperfect schedule, and then try lots of changes gradually make it better. This is called “hill climbing” - trying to find the peak in fog by always taking uphill steps.

Unfortunately there is no guarantee that the hill we start on is the highest in the range - we may need to travel downhill (make our solution worse) to get to a better hill to climb. In other words there may be no incremental set of changes from *this* starting schedule to a perfect one that satisfies all my constraints on instructor, student, and equipment availability.

AI has many cunning techniques to find one of the best hills, and then climb that.

Provisioning Optimisation

A model trained on actual flights’, provision manifest and residuals, could rapidly learn to optimise future flights’ provisioning, so pax are most likely to get what they like (increasing satisfaction), whilst reducing the costs of over-provision needed to ordinarily achieve this.

Aircraft Design

The design of aircraft already uses very complex conventional software, in which the airflow over a wing or entire plane is modelled. When designs have passed this stage, they are then tested in wind tunnels, and finally with flying prototypes, and performance measurements made.

Often the underlying mechanics giving rise to those measurements are not understood, and this is where AI techniques could help. Symbolic regression could be used to take observational data from wind tunnel experiments or data captured in real flight, and find the closed-form equations that describe the situation. These can be validated, used as insight, and used as models to create better (for example, more fuel efficient) aircraft.

Symbolic Regression is likely to prove the mathematical equivalent of large language models.

Using a multitude of techniques, symbolic regression takes observational data, and uncovers the underlying equations, often yielding deep insight. This has been tested on the highly complex laws of advanced physics.

4. Safety Enhancing AI

Training

Whilst not safety-critical of itself, training systems are indirectly related to safety, since well-trained crew are, one assumes, safer than poorly trained crew.

AI based training could be both effective and cost efficient. For example, suppose the student needs to learn and internalise a particular skill, such as keeping the aircraft's aspect and speed (following the flight director), or how to automatically adjust the rudder when changing the throttle in an engine-out situation, or following the procedures at a hold. The AI might act as an instructor, initially presenting the student with easy (yet realistic) tasks. As the student progresses, the AI gradually makes the task harder, developing their skills. If the student finds one aspect easier than another (say keeping altitude better than heading), the tasks are adjusted (in flight) to build the weaker skills. Hardware to accomplish this might be much cheaper than a full simulator, requiring only the controls in question and a conventional computer, removing anxiety of the training cost and the temptation to skip from the student.

Similarly crew, being human (for now, anyway), will make mistakes. Patterns could be identified and used to identify mitigation measures, such as making a particular training element clearer, or (by adjusting other duties) reducing crew overload.

Simulators

Professional flight simulators are very expensive. Using a technique we do not wish to disclose here¹, AI could drastically reduce this cost. This would enhance safety not by producing a better simulator, although it might, but by reducing the cost of running simulators and thereby increasing their accessibility.

In-Flight Advisory

In this half-way house the AI is used in a safety-critical environment, in an advisory capacity². The aim is to reduce the pilots' cognitive load. Simple examples might be informing the pilot the aircraft is approaching airspace, or that the local windspeed is increasing such that landing might be challenging for a light aircraft. It may not be a co-incidence that both these examples could be implemented using explicit knowledge AI.

In-Flight Safety Monitor

An AI could also be a safety monitor. With full access to the 'plane's control settings and flight instrumentation, it could watch out for and warn of anomalous flight conditions, such as (by way of a simple example) climbing away with the undercarriage down. With a suitable knowledge base, it could also infer complex anomalies, in addition to simple direct ones. Further sensors could give further capability: speech recognition could monitor ATC

An **Inferencing System**, also known as an **Expert System**, is provided with rules, often taken from an expert.

In use, like a rational human, it can make deductions and reason on the situation it faces. It can probably explain why it reaches the conclusions it did.

There are two forms, and the easier to understand is forward chaining. It asks, like a detective, given all that I know, and this set of rules, what can I deduce? And what can I deduce from that...?

¹ Because it is new, potentially commercially viable and valuable, and patentable.

² Disclosure: your author is developing a system to do exactly this, and a prototype exists.

instructions, and visual sensors could look out for other aircraft or drones.

An AI could also, in the same way, monitor pilot performance, much as the pilot-in-command is monitored by their colleague. This would be valuable in single-crew operations: there are still two crew, but only one is human.

Cognitive Load Reduction

Augmented Intelligence uses techniques similar to AI combined with visualisation techniques, with the aim of enabling the human to rapidly assimilate lots of information. In other words, rather than have the pilot integrate disparate observations to obtain the overall situation, present that situation already integrated.

In-Flight Fallback

Even though we cannot prove our AI will be correct, we may well suspect that its error rate is lower than a human's. (In other domains, experiments show that the *type* of errors AI and people make are different.)

There may be circumstances where the pilot is confused, or stuck, and an AI might well "get it right". An inferencing system might be much better than a person under stress, when faced with multiple instrumentation errors, in determining which readings are valid. For these situations we might consider allowing the pilot to hand over control to the AI, presumably temporarily.

Post-Flight Performance Analysis

By recording the flight, and using an expert system, after landing the AI could draw the pilot's attention to aspects that might need improving. It would not be critical that the AI was correct - all that is needed is that the pilot consider what happened.

Software Testing

AI can help test aviation code. The aim is to find more ways to stress the software being written, and to find errors either in the code or in its specification. If the AI wrongly mis-identifies code as faulty, that code will simply endure unnecessary inspection.

5. AI with Safety Risk

Software Writing

AI can write, or help write, software.

The implications of this are considerable. It means that non-AI software may have been written by AI. It may be that either no-one understands how it works, or that they only partially understand. AI assisted code will become more prevalent, as it is a cheaper to write. Manufacturers and operators might consider it prudent to ask their software suppliers to disclose what AI tools were used in writing the software.

In the command loop

Regulators manufacturers and even operators have to balance two conflicting drivers when it comes to using AI in the command loop, that is making decisions.

In its favour, AI may handle unexpected situations on which neither it nor the pilot has been trained and practiced, ie outside the training envelope, and it may do so well.

The jumbo in the jet here is that we cannot prove, with non-deterministic software - AI - that it will do so correctly. We cannot, in general, even prove that it will handle all situations within the training envelope correctly, and this is especially true with AI which has used machine learning.

Conventional, deterministic, software can however be tested or even be written with formal, design-proving techniques. It will operate “within the specification envelope” correctly. Moreover, conventional - but still very clever - software can also be made, to an extent, to handle unexpected circumstances. The more that scope is increased, however, the more like AI it becomes.

Software-writing AI

There are two approaches to this. One is to regard programs as text, and ask a suitably trained **large language model** to help. They do not always get it right.

The other approach is **genetic programming**, which mimics Darwinian natural selection to evolve a program to solve a problem.

A population of programs compete at solving the problem. The mutated or mated offspring of the best survive into the next generation. “Survival of the fittest” means that all being well, a program which can solve the problem evolves.

There is no guarantee that *anyone* understands how these evolved programs work, but they may solve problems which people cannot.

6. Summary and Conclusion

AI is not like ordinary software - the very characteristics that make it powerful also make it impossible to verify. Because of this safety aspects must be carefully considered, and we may consider there to be three categories:

1. **Safety Neutral.** Many commercial and operational aspects of an airline have little direct safety implications, and the use of AI is subject to normal considerations of privacy ethics and commercial cost-effectiveness.
2. **Safety Enhancing.** In many other assets of aviation, AI could well enhance safety.
3. **Safety Risk.** Finally, the difficulty in formally verifying AI may well make it unsuitable for use in the command loop whilst flying. AI used in the development of safety-critical software should also be carefully considered.

Regulators, manufacturers, and operators need to be mindful of which category an AI application fits.