

Artificial Intelligence



Table of contents

Acknowledgements	2
Foreword.....	3
1. Introduction.....	4
1.1 General Description	4
1.1.1 AI Definition.....	4
1.1.2 AI Approaches.....	4
1.1.3 AI Classification	5
1.2 Ongoing Work on AI in Aviation.....	7
2. Enablers	9
2.1 ML Lifecycle and Explainability.....	9
2.2 Safety Aspects.....	11
2.3 Certification.....	13
2.4 Common Data Infrastructure Initiatives	14
2.5 Human-Machine Collaboration	16
3. Fields of Application.....	19
3.1 General Applications of AI	19
3.2 Application Examples for ATM/ANSPs	20
4. Conclusions.....	25
References.....	26

Acknowledgements

This whitepaper was produced by the CANSO Strategic Technology Workgroup (STWG) and industry partners. CANSO would like to thank the following individuals and organisations for their dedication and commitment in developing this document:

- Beatrice Pesquet, Thales (lead)
- Firdu Bati, FAA
- Alfonso Nicotera, IDS AirNav
- Adam Gyarmati, HungaroControl
- Eduard Gringinger, Frequentis
- Zoltan Molnar, HungaroControl
- Roland Tichy, HungaroControl
- Maria Tamm, EANS
- Julian Titmuss, NATS

The authors would like to thank to all the reviewers of this paper, in particular:

- Billy Josefsson, LfV
- Valentin Polishchuk, LIU
- Jean-Marc Loscos, DGAC
- Tomasz Klosowicz, PANSA
- Matteo Pace, ENAV
- Eduardo Garcia, CANSO

Foreword

The concept of Artificial Intelligence (AI) dates back to the 1950s, but the expansion of its use in solving previously unsolvable problems is driven by three main factors:

- The availability of huge amounts of data thanks to the digital revolution, which is the main ingredient for using neural networks;
- The processing capacity of computers is increasing constantly thus making the technology available for everyday use;
- The previous two effects have motivated and accelerated the emergence of powerful algorithms to solve problems of increasing complexity.

AI use is spreading rapidly in every industry and aviation is no exception. However, new technologies come with some associated risk, and as aviation is a safety critical domain, there is a need to examine the use cases and capabilities of AI carefully. In this paper, we present an overview of the potential use cases, challenges and possible solutions of implementing AI, focused on the air traffic management (ATM) industry.

1. Introduction

1.1 General Description

The significant and continued growth in air traffic in the years prior to the pandemic has prompted considerable exploration of the use of AI in aviation, especially in ATM. It is expected AI will have the capacity to meet the challenges of increasing air traffic complexity through growth and new airspace users and support more efficient operations in order to reduce the environmental impact of aviation.

The essential question is not whether, or if these AI-based applications will be used, but when they will be in operation. While the decline in air traffic caused by the COVID-19 pandemic has reduced the immediate capacity pressure on the ATM system, it is worth using the years ahead to reduce potential barriers (including legal, safety and technological) to the effective use of AI.

1.1.1 AI Definition

AI is a generic term, covering many disciplines that have expanded with many technologies over the past 60 years and today covers a wide range from machine learning to vision, natural language processing, machine reasoning or robotics. In essence, AI systems solve problems which, in humans requires high levels of intelligence, or “heavy thinking” to solve. The concept was first discussed at the Dartmouth Summer Research Project on Artificial Intelligence, in 1956, by a range of mathematicians and scientists. It concluded that any feature of intelligence can be described so precisely that a machine can be constructed to simulate it. According to the European Commission’s definition, “AI systems are software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions”. The most widely accepted definition is based on this: AI is any technology that appears to emulate the performance of a human.

1.1.2 AI Approaches

The current level of AI is described as Artificial Narrow Intelligence (ANI), which means that we use the technology for some specific tasks such as voice recognition or weather forecast. The broader term, Artificial General Intelligence (AGI), an AI that would be able to plan and execute any tasks, remains under development.

The common element in today’s AI systems is they use data as input and, without the need of a priori models, they provide an appropriate output. They can do this thanks to their internal intelligence.

There are two main schools of AI, which have co-existed through its development to date. The first one is model-driven AI where the system is traditionally programmed and is based on the

knowledge of experts in the field. The model can provide answers based on this domain knowledge. The second and dominant model today is the data-driven AI where large amounts of data is the basis of learning. The model learns implicit rules and connections in the data and provides output based on that data, without explicitly providing any human understandable rules.

1.1.3 AI Classification

Within AI a large number of sub-areas can be identified. These can be grouped in several dimensions as indicated in Figure 1. The image was drawn by Francesco Corea to give an overview about the complex field of AI technologies. On the y-axis the AI problems are described, whereas on the x-axis the AI paradigms, technical solution for these problems, are defined. The six paradigms are:

1. Logic-based tools
2. Knowledge-based tools
3. Probabilistic methods
4. Machine learning
5. Embodied intelligence
6. Search and optimisation

These paradigms can be group into three different approaches:

1. Symbolic
2. Sub-symbolic
3. Statistical

In this whitepaper we highlight some that could play important roles in ATM industry.

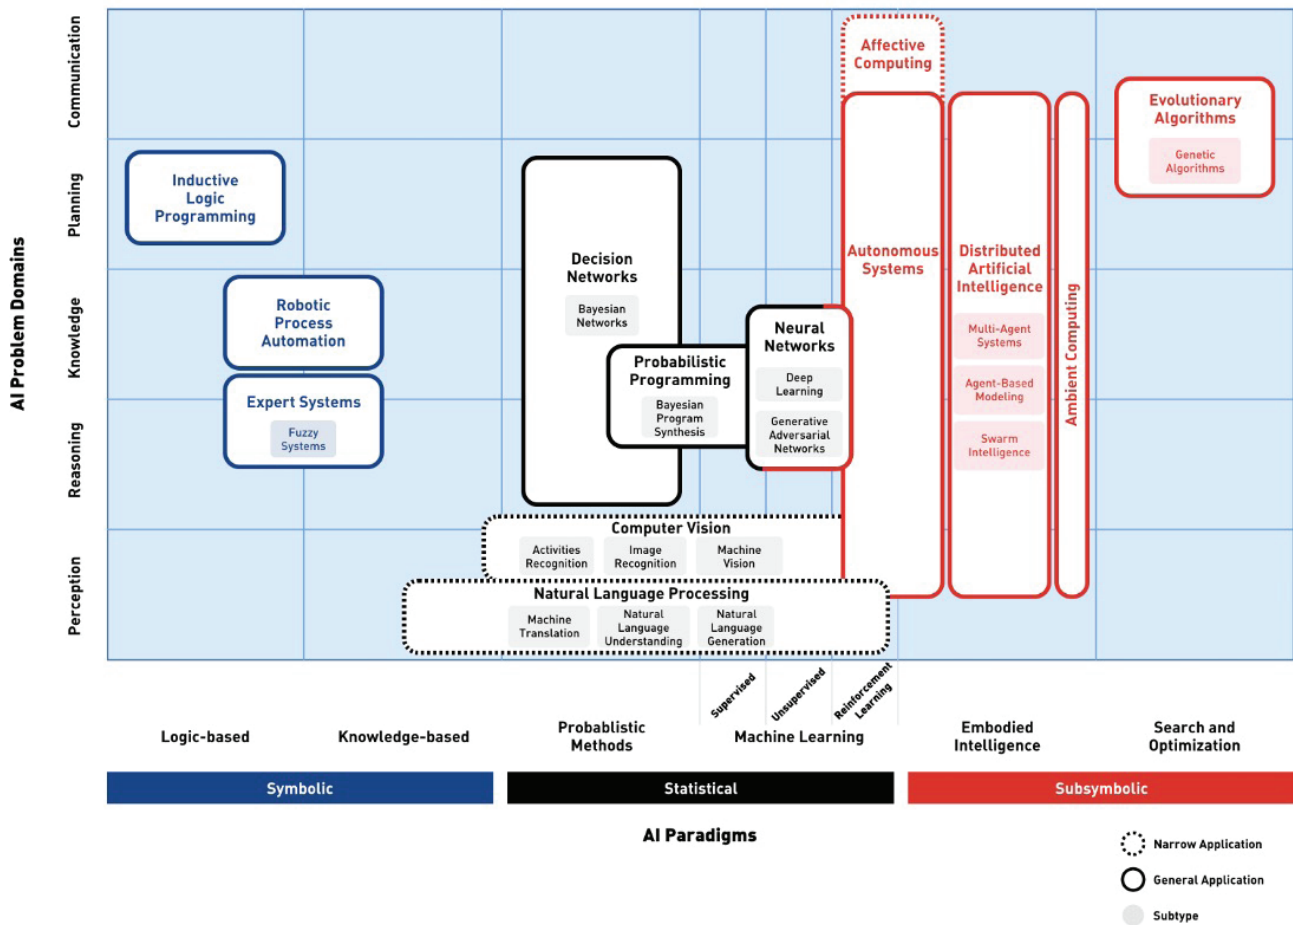


Figure 1: AI Knowledge Map: How to Classify AI Technologies ([forbes.com](https://www.forbes.com))

Expert Systems

The most common form of expert system is a programme that consists of a set of rules that analyse information (provided by the user) in a specific area and suggests a sequence of actions to be performed. Such systems have the capability to draw conclusions by seemingly logical reasoning.

Machine Learning

Machine learning (ML) algorithms are capable of learning, generating knowledge from experience. The system is able to recognise regularities/rules independently or with human help based on example data. The system can make good decisions on unseen data after learning. Machine learning has four main areas:

- **Supervised learning:** In this case the expected output labels are available for the training data. The system tries to learn a function that gives a proper mapping from input to output variables and has a proper generalisation ability. Examples of such methods are regression (for numeric output variables) and classification (for categorical output variables).
- **Unsupervised learning:** In this case output variable labels are not available. The system needs to develop some behaviour without the desired responses. Examples for this are clustering and dimension reduction techniques.

- Semisupervised learning: Deals with data that is partially labelled via unsupervised learning, requiring knowledge to tag/label/confirm labels (supervised) and that then automatically propagates the labels to the full set of data. They are trained sequentially in an unsupervised manner and then fine-tuned using supervised learning machines. Examples: DBF (Deep Belief Networks based on unsupervised components RBM (Restricted Boltzman Machines)
- Reinforcement learning: The AI has the controls in the system and tries to reach the best score (defined by a function integrating measurable objectives) depending on the rewards and penalties, which are pre-set to certain actions/ happenings. A good example is robotics, where it is commonly used to learn complex behaviour.

Deep learning

Deep learning (DL) is a subset of ML. The main difference is that, in ML the representation of data is part of the design process (the “model” is learnt based on the training dataset and validated and tested on an independent dataset), while in DL the representation part is done by the algorithm (with hidden layers and non-linear transformations like Rectified Linear Unit (ReLU)). This relieves the user of the task of writing explicit rules and can generate state-of-the-art results, but on the other hand acts like a black box. In other words, it can cause safety challenges due to explainability issues. Some very popular fields in DL are:

- Computer vision which can be used for image recognition or activities recognition
- Natural Language Processing (NLP) which can be used for automatic voice recognition or automatic translation
- Optimisation, planning and resource allocation tasks
- Automation, which can include robotics, but also any kind of automation tools

1.2 Ongoing Work on AI in Aviation

AI is evolving rapidly and changing our lives through its use in a variety of areas. Despite its huge potential, aviation is a safety critical industry and implementing AI will come with challenges. There is a large effort by major agencies and working groups to standardise the use of AI, develop standards for AI applications and demystify AI for a broader applicability.

- *FLY AI Report* by EUROCONTROL. Provides an overview of how AI can be applied in the industry, its potential and how it can transform the sector.
- *EASA Artificial Intelligence Roadmap 1.0* by the European Union Aviation Safety Agency. Visions on the safety and ethical dimensions are written in the roadmap. The Agency will amend it each year with thoughtful views and lessons from practical development.
- *AI in Aviation* whitepaper by International Air Transport Association (IATA). This document highlights some of the threats and opportunities in AI.
- ICAO has engaged in various activities supporting multiple initiatives in the domain of artificial intelligence, e.g. partnering in the AI for Good Annual Global Summit. We have less than 10 years to solve the United Nations’ Sustainable Development Goals (SDG) and the AI for

Good's goal is to use AI for help.

- SAE G-34 and the EUROCAE WG 114 working groups are currently jointly working on AI standards of the industry. All documents are for the moment work in progress.
- *Artificial Intelligence Cybersecurity Challenges* by ENISA. Even though it is not ATM specific but general, it is worth mentioning due to its applicability.

A very broad range of AI applications are targeted under the umbrella of the ISO/IEC JTC1/SC42, which is a reference for all the industries developing AI-related standards.

2. Enablers

In order to see successful applications of AI in ATM, several enablers are necessary. Firstly, trust needs to be built for use of this new technology. There are many components of this trust. Cybersecurity and ethics will be important but will not be discussed here since they have been thoroughly examined in the previous references. Our focus will be on the explainability of the algorithms (Section 2.1) and the safety issues (Section 2.2). These are all strongly related to ongoing activities aiming at their certification (Section 2.3).

A second enabler is the availability and the interoperability of the operational data on which the ML models are built. We describe a few initiatives for common data infrastructures (Section 2.4).

Finally, the human machine collaboration (Section 2.5) is of paramount importance and is related to the previously mentioned issues of trust and safety. Indeed, there is a good trade-off to be found between machine autonomy and human supervision, in other terms a “shared autonomy”, in the form of a fluid cooperation framework between human and AI to minimise operational risks. This passes through the definition of who is the last resort in each situation. Currently, it is always the ATCO who has the final decision; the system proposes, the ATCO accepts or rejects the suggested solution. Moreover, there is a positive loop of human feedback to machine improvement, since the machine can learn from human interaction. Trust can also be improved by upskilling human operators, to better understand the new data-based algorithms paradigm and adjust their expectations to the existing limits.

2.1 ML Lifecycle and Explainability

A main difference *between data-based* ML development and *rules-based* (or model-based) software is related to the black-box aspect and the explainability of the results. The black box aspect involves three issues:

- It is difficult or almost impossible to link functional requirements to written code in data-based methods development, and this is intrinsic to the method;
- There is limited theoretical capacity right now to prove all the properties that we would like about ML models (but we can expect this aspect will be solved by the academic community);
- The testability of the data-based ML algorithms, which is related to the previous two issues, and amounts to design adequate test datasets to prove robustness of the neural networks and ensure all the corner cases are covered.

The aim of Explainable AI (XAI) is to “describe the purpose, rationale and decision-making process of the AI tool in a way that can be understood by the average person”. While the AI tool focuses on the *what* aspect the problem, explainability deals with the *why*. However, the “average person” to whom this explanation is intended depends on the context in which the explanation is used. It can be the data scientist, the expert in the operational field, the customer (end user), or the regulatory body. The explanation has to be useful for those who use it. The need of explainability comes from several reasons:

- Need for a human readable justification of a decision: explain to build trust
- Need for identification of system flaws: explain to improve the system and prevent operational risks
- Need to ensure compliance to regulations: explain to control the behaviour and prevent conformity risks
- Need to extract information: in this case, the explanation aims at providing new insights in the data for the data scientist to build better models.

The lack of direct explainability of data-driven models is primarily related to the nature of machine learning algorithms, which, in general, take training datasets as inputs and produce predictive models, sometimes overly complex ones. Some AI models are easier to explain than others: for example, linear models (linear and logistic regressions) decision trees, Naive Bayes, probabilistic graphical models. These will be easier to interpret than more complex models like deep neural networks (DNNs) that exhibit a very high level of complexity in structure and parameters. DNNs often involve millions of parameters and functional behaviours cannot be strictly related to some of these neurons. Interpretable models, in general, have desired properties such as unbiasedness, protecting sensitive information, reliability, and trust by humans.

However, complex algorithms typically produce models with high prediction performance as they can capture non-linear relationships and are able to adapt to massive training datasets. Contrary to interpretable model-based methods, where the performance limitations come from the generality and assumptions of the model, in data-based DNN the performance can theoretically increase with the size of training dataset and the complexity of the model. In addition, these models do not require manual feature engineering. For these data-driven models, the focus is on providing a human understanding of how the underlying algorithm makes prediction decisions, either in terms of the main factors that led to the decision, in terms of features (or high-level concepts that are activated when the prediction was given), but the explanation can also be obtained by counterfactual reasoning: if one feature value is modified, how does the prediction change?

The content and format of the explanation is also a function of the type of data involved; for visual or audio data the interpretation will be related to some human-interpretable features in these data, for natural language processing the interpretation will be both related to the syntax and semantic of the data, while for tabular data, the interpretation will be highly dependent on the various variables and their correlations.

Explainability can be included in each of the steps of the ML lifecycle: from the conception to model design and training (analytical evaluation through code inspection), validation of the model (performance assessment, conformity tests, evaluation of explanations), deployment into operations (performance monitoring and drift monitoring, monitoring of the bias and fairness of the model, but also of the explanations), and finally the corrective maintenance (where the explanations should take the form of a diagnosis). This control cycle should be permanent for the conception, validation, and runtime, but periodical audits should equally take place for the entire system.

An explainable model can be obtained either through a normative approach (train an interpretable simple model, conceived as such: for example, identify a specific neural network

(NN) to a set of rules, or modify the architecture of a NN to make it interpretable) or by a descriptive approach (produce explanations post-hoc, especially for complex models). Some post-modelling explanatory approaches are the Partial Dependency Plots¹, the SHAP method with variants², the local surrogate or LIME method³, Global Surrogate models (Surrogate), adversarial and counter-factual, or Anchors⁴.

The different explanation methods are classified based on various elements including criteria, result, type, and scope of the method. According to criteria employed, explanation methods are categorised into intrinsic (using constraints) and post hoc (after models are trained). In terms of results, explanations are categorised as features summary statistics, feature summary visualisation, model internals (structure and parameters), and intrinsically interpretable models. Based on types or classes of methods, explanations are categorised as model specific (limited to specific type of models like regression) and model-agnostic (tools that are generic and work with any model). Finally, explanation methods can be categorised by scope as local (explaining individual predictions) and global (explaining entire model behaviour).

Depending on the different stakeholders involved in the lifecycle of the ML models, the goals of the XAI may be different. For the data scientist, the intelligibility is related to the fidelity of the model to the data, for the customer, the satisfaction may be related to the capacity to provide a simple aid in the decision, while for the regulator conformity is the main concern. Finally yet importantly, the explainability is also related to the acceptability of a given technique, so to a larger societal challenge and an ethical dimension.

However, depending on the application problem area, explainability may not always be a critical requirement. For instance, if the application is a simple classification of sparse unstructured data, such as event descriptions in safety reports, then what may be more important is the accurate detection of the event type as opposed to what combination of terms used in the detection algorithm led the machine to identify the correct event type. Similarly, if the application is to transcribe voice communications between pilots and controllers, getting the most accurate transcription and deciphering the meaning are more critical than understanding the speech patterns and communication style of the underlying voice interactions that enabled the model to render highly accurate transcription. In general, there are few scenarios in which model explainability is less important, which include application with no significant impact (trivial application), problems that are well studied, and in areas where there exists a potential for system manipulation in which rogue users can exploit the system.

2.2 Safety Aspects

The safety assurance of ML-based algorithms is an emerging question. Safety assessment processes (based on ICAO SMM as a common reference and defined by EU regulation 2017/373 in Europe) clearly set out the main steps to be followed, although seldom provide detailed methodology for software safety aspects. Standards⁵, used to allocate “assurance levels” to specific functions based on the potential severity and probability of safety effects, but these are not fully applicable in case of ML. There is a need to fill the gap of safety aspects not covered by existing standards.

Safety approaches that rely on existing software are based on the assumption that the software is fully explainable. Traceability of the requirements is provided throughout the system development lifecycle, while fulfilment of these requirements can be justified with clear evidence. However,

it is not always true for ML, as almost every popular ML method contains a “black box” part, hardly constraining the explainability (see section 2.1). Among the challenges related to ML, which make it different from previous software development techniques, is that the conception of the ML algorithms is fully embedded with the data, so the data becomes part of the algorithm specifications and requirements and these new requirements need to be carefully designed. Hence, there is a need to ensure clear specifications about data accuracy, accessibility, relevance, consistency, completeness, and all these can be gathered under the umbrella of data representativity. This represents a new paradigm for software assurance, where we go from certifiable algorithms to certifiable data sets and training programs. So even the most robust ML algorithm’s performance is highly dependent on the training dataset. It might not be a problem in the case of applications with low safety effect but can be a strong constraint in the development of safety critical applications. The situation is even more complex in case of online learning methods, so these have not been in the focus of current standardisation works yet.

Concerning AI assurance, currently the best example from industry is from Microsoft who have come up with the following key principles:

- Fairness – treats all users equally
- Reliability and safety – recognises any safety critical concerns within the application;
- Privacy and security – assurance that all legal, regulatory and cyber security secure development principles have been adhered to
- Inclusiveness – systems should be introduced in such a way as to ensure that all affected users are engaged throughout the process
- Transparency – all systems are understandable
- Accountability – the scope of the AI responsibility should be assessed and the ultimate decision making responsibility that results must be clearly identified

These principles could be used as the basis of ATM industry research questions and may possibly feature as a minimum standard to be adhered to by suppliers. Evidence that suppliers have looked at these or similar principles will aid in a level of assurance that the AI solution at its most basic level is fit for inclusion within the ATM system. This initial building block would then, in turn drive a more detailed piece of work to develop a framework around what the ATM industry needs from AI solutions, where necessary bespoke to the ways in which AI solutions are envisaged to be applied. The outputs of this would then enable the ATM industry and suppliers to identify the appropriate solution space for AI.

Some ML based solutions are expected to directly assist ATCOs in their tasks (see section 3 – e.g. providing decision-making support or resolution advisory). This means that safety critical ML applications are not far away from deployment with a potential to also increase the level of automation. Combining new solutions and increased automation significantly affects the human-machine interface of the functional system. Human performance aspects, such as situational awareness, trust in the system, roles and responsibilities and team composition need careful consideration.

The following list provides some considerable potential challenges for ML:

- Learning dataset is inadequate – as the data itself can drive the ML function, quality, size, and composition of the datasets highly influence the behaviour of the system. Incorrect, incomplete or non-representative data, poor quality data or irrelevant features can significantly decrease the performance of the ML item even before training.
- Learning process is the subject of bias or variance – high bias (underfitting), or high variance (overfitting) can significantly decrease the performance of the ML item.
- Test and validation datasets are inadequate – if independency from the training dataset is not assured, or the datasets are either incorrect or incomplete, reliability of system performance metrics is questionable, and the desired level of assurance cannot be guaranteed.
- Decisions are based on unexplainable patterns:
 - Decision making mechanism is not straightforward for the user which might decrease trust
 - The performance of the system might be different in specific circumstances (depending on the composition of the training dataset), negatively affecting the integrity of the system
 - Unpredictable behaviour in unprecedented situations, sensitivity to disturbance

Potential solutions to maintain or enhance safety level:

- Enhancing explainability – understand the behaviour and the decision mechanisms of the ML tool and prove that it behaves as it is expected. Depending on the safety criticality (safety assurance level) of the application, the use of more explainable methodologies should be considered.
- Assuring that the quality, size, and composition of the datasets are adequate, all datasets are complete and correct, and that appropriate safeguards are in place to protect the datasets from external, malicious manipulation, taking into consideration the safety impact of the function. Special care should be taken when the explainability of the model is limited.
- Adequate verification and validation in simulated or real (e.g. passive “shadow” mode) environment for a given time period.
- Verification and validation on functional system level, taking into consideration the integration of the ML component to the functional system, and the effects to human performance.

2.3 Certification

Competent Authority certification and approval of AI (ML) applications is a key issue, as specific regulations and standards to support the process are not yet widespread. A well-founded safety assessment (or standard) for ML is definitely a prerequisite of approval, but there are other important questions to answer. Among the regulators, EASA has already published a roadmap on AI-related issues⁶ which highlights the importance of ethical AI and trustworthiness as key factors in AI development (referring to “Ethics Guidelines on Artificial Intelligence” report of European Commission High Level Expert Group). Challenges and building blocks of trustworthiness, especially learning assurance, explainability and AI safety risk mitigation are discussed in specific chapters. It also gives examples for potential applications in the ATM field such as

supporting ATCOs with enhanced automation in case of repetitive tasks. Safety risk management, cybersecurity and environment are also potential fields of application.

The document by EASA provides a classification of AI/ML applications in three levels:

- Level 1 AI/ML: assistance to human
- Level 2 AI/ML: human/machine collaboration
- Level 3 AI/ML: more autonomous machine

According to the EASA AI roadmap, publication of the first usable guidance for human assistance/augmentation can be expected in 2021, and for human/machine collaboration in 2022, while the finalisation of EASA AI/ML policy is planned for 2028.

Current European level regulations provide an open framework for the use of ML in ATM, although further adaptation of these regulations to specific applications is necessary. Application of performance-based regulations (focusing on objectives to be achieved) providing a more flexible regulatory framework, and supporting innovation, is a key in ML certification. The latest amendment of the EASA Basic Regulation (EU 2018/1139) facilitates this process.

The role of standards is very important in safety assurance of conventional software, and will be equally important in case of ML applications. Standardisation work is already ongoing in a joint working group of EUROCAE Working Group 114 “Artificial Intelligence in Aeronautical Systems” and SAE G-34 ‘Artificial Intelligence in Aviation’. The main deliverable will be the standard document, providing guidance on assuring safe and secure AI-based systems, and regulatory approval. The scope of the document covers ground-based equipment and airborne vehicles, including unmanned aircraft systems (UAS) products, and the whole lifecycle of the ML models. Publication can be expected end of 2022. A *Statement of Concerns and a Taxonomy* document is expected to be released in spring 2021, which provides relevant definitions, describe the scope of the work, and the main concerns and considerations for AI (ML) applications in civil aviation.

From a cyber-security point of view, the main focus is going to be looking at standards for how AI is designed and used; we believe that these standards have yet to be agreed. Should the solutions be software based, then there are principles already agreed for secure development cycles. The UK National Cyber Security Centre discuss the who, what⁷ and the how⁸. The context here is similar to the approach used in other areas; a multi-layered approach encompassing socio, technological, physical and all the other techniques available to us. In essence, rather than a one size fits all solution there is a “toolkit” that is adaptive and lends itself to starting to look at AI based security tooling..

2.4 Common Data Infrastructure Initiatives

The strength of AI relies on the availability of a huge amount of structured and unstructured data that, typically, converge into a well-known infrastructure named *Data Lake*.

Availability, in this context, must be defined as per two main axes:

- Ability to access data
- Trust data, as defined by the CIA (Confidentiality, Integrity and Availability) model

Therefore, the need for a common global, reliable and fast Data Lake infrastructure reflecting those two main axes is crucial.

It must be emphasised that even if there are different initiatives around the world providing platforms for ATM driven data, the approach seems to be fragmentary and localised to most digital advanced areas and, in some cases, limited to few countries inside them.

This is in contrast with the essence of a Data Lake where even raw data are pushed and from which ML/AI algorithms can take advantage. For further details about the main features of a Data Lake infrastructure, there is a set of rules collected into the “Data Lake manifesto”⁹.

The following describes some widespread initiatives that represent a starting point for boosting productive usage of AI in the aviation industry. All of them share the principle of building a common data infrastructure starting from standardisation aiming to implement platforms with this principle in mind.

- GAIA-X¹⁰ is a project, initiated by Europe for Europe, aimed to develop common requirements for a European infrastructure providing data and services in an open, transparent, federated and trusting environment. Currently, representatives from seven European countries are involved.
- NM B2B¹¹ web services is an interface provided by the EUROCONTROL Network Manager (NM) for system-to-system access to its services and data, allowing users to retrieve and use the information in their own systems. These services are at the core of the NM Interoperability Strategy, and follow the Single European Sky Aviation Research (SESAR) and the International Civil Aviation Organization (ICAO) system-wide information management (SWIM) principles, which aim at achieving real-time information exchange at a global level for air traffic flow and capacity management (ATFCM). The benefits of such an open ATM digital collaborative environment include data quality and accuracy, timeliness of information, simplification of processes, cost reduction and safety. These services are made available mainly to ANSPs, aircraft operators, airports, ground handling agents, computerised flight plan service providers (CFSPs) and airspace management cells (AMC).
- Data4Safety¹² Programme was launched in 2017 by the European Aviation Safety Agency (EASA) which aims at changing from the current reactive approach to safety – based on accident and occurrence reports – to a proactive approach, analysing trends and situations which can lead to a safety hazard. Stakeholders within the aviation community will share data, information and knowledge in order to deliver safety benefits. Data4Safety will merge the very large and growing number of data sources available in the European aviation system and provide a critical mass of data. Analysis capabilities will be significantly increased thanks to the sharing of large amount of data and the use of the most advanced information and communication technologies in the fields of Big Data and Data Mining. The Data4Safety members signed a programme charter which lays down the principles of collaboration, voluntariness, confidentiality and Just Culture.
- Data Beacon¹³ is an open platform focusing on data ownership, confidentiality and data protection. This research platform was developed in the framework of the SAFECLOUD H2020 project. It facilitates collaboration with third parties without compromising control of private data, and allows airlines, airports, ANSPs and regulators to explore together data solutions without sharing their different data assets.

- Airbus Skywise¹⁴ is fast becoming the reference platform for all major aviation players (more than 100 airlines and major suppliers sharing data) to improve their operational performance and business results, as well as to support their own digital transformation. Skywise has been designed to allow data integration, access, and analysis seamless for users while ensuring highest levels of security.
- IATA Open Air¹⁵ (former IATA Open API) aims to provide standards that enable interoperability between and among airlines and other participants in the aviation supply chain in order to make possible a global interconnected aviation system. By facilitating the development of industry standards and best practices on data exchanges, Open Air initiative aims to provide capability to consistently share data by promoting the adoption of an open API technology.
- FAA NextGen¹⁶ is the FAA-led modernisation of America's air transportation system aiming to make flying even faster, more efficient and more predictable. It encompasses innovative and transformative technologies that are being developed and deployed after thorough safety testing. Under NextGen umbrella, in terms of Common Data Infrastructure, several assets can be enumerated as SWIM enabled data and services: NSRR (NAS Service Registry and Repository), TFMData (Traffic Flow Management Data), SFDPS (SWIM Flight Data Publication Service), STDDS (SWIM Terminal Data Distribution System), ESM (Enterprise Service Monitoring).
- ICAO iCADS¹⁷ (CIVIL AVIATION DATA SOLUTIONS) is a portal built with the needs of aviation industry in mind, integrating the entirety of data, business intelligence products and curated ICAO reports and documents. Users are allowed to access to multitude of data sets and business intelligence solutions.

It is worth noting that where a variety of actors use common data in order to support and drive their activities, the cyber security and assurance of these data sets must be ensured through robust safeguards and a series of checks and balances. The data must be protected from interference, (both internal and external sourced) as much as possible and tracking mechanisms must be in place to identify and flag deltas in the data should they occur.

2.5 Human-Machine Collaboration

In a human-machine collaboration paradigm, the AI model is not used as a mere tool. Instead, a human expert can focus in areas where additional operational perspective is needed while the machine can automate a large and repetitive task that it can handle more accurately in a scalable fashion. This is arguably the most impactful application of AI as it can complement and augment human capabilities. This paradigm can also be leveraged with AI models that can provide semi-automation solutions.

For instance, if an AI model can provide the level of confidence in its classification task, human experts can focus in areas where the machine displays lower confidence in its categorisation. A human-machine collaboration can also be deployed as a feedback loop to continuously improve the AI system with more incoming data, which would be particularly useful for domains in which the operational dynamics change through time.

Sometimes, an AI model needs to be developed using data that may not be sufficient for all scenarios, and it can continuously be refined with the help of humans providing feedback

(corrections) until it reaches a point where the model has an acceptable level of performance.

A different application scenario is that some aspects of the operational context change as rules and procedures change and/or new ones are deployed. To account for these changes, humans may need to continue to provide feedback to enable the system to adapt.

An example of machine-human collaboration in aviation safety is for an AI model to identify operations with inherent risks, such as violations of rules, procedures etc. while a domain expert can validate the risky behaviour from an operational perspective. Air traffic controllers can override some procedures in a given context while still maintaining an acceptable level of safety. For example, in some terminal environments, controllers sequence aircraft to optimise efficiency, where a machine can detect potential risk due to an increase in frequency of airspace excursions, but the risk may be deemed low based on domain expertise.

Another example is the anomaly detection application referenced in section 3.2. In any reasonably complex system the volume of operations makes it resource intensive to evaluate all flights. An ML model can be deployed to analyse all flights and pre-classify the anomalous flights, typically involving a very small fraction of the total operations. However, this pre-classification would not be able to provide additional contextual information that will enable decision makers to understand the issue further and mitigate as necessary. Given the anomalous cases are small in size, human experts can study those cases and provide additional contextual information, such as causal and contributory factors that will facilitate decision making.

Roles and responsibilities of human actors in the functional system will change with the introduction of AI-based models. Applications would take over certain tasks from ATCOs (and from other human operators), changing the way of working. An example is the ACAS function, which can be fully delegated to the machine¹⁸. Some current tasks might disappear, but in the meantime, new tasks will emerge. Decision-making support (e.g. resolution advisory), or direct implementation of actions modify the sharing of responsibility between human and machine, requiring the re-evaluation and update of regulatory background. Teamwork and team composition will also be affected.

It is important to emphasise that automation without AI (ML) components can have very similar effects. On the other hand, effectiveness of ML models in certain well restricted but complex tasks can significantly accelerate the automation process in ATM. Keeping the human in the loop is essential in a synergic human-machine collaboration and contributes to the resilience of the functional system. It is even more important because applications increasing the level of automation and support to ATCOs to handle traffic with higher complexity, might decrease situational awareness. Proper assessment of this issue is important to make the ATCO self-confident in his/her new role. Explainability (and indirectly, proper safety assurance) of AI is also a key factor to build trust and support human-machine collaboration.

The role of automation (e.g. human in the loop) and the trust in technological agents, will necessitate further investigation based on the corpus of knowledge already available from other domains, but rarely addressed as such in ATM. Metrics and event classification to define potential risks are issues that must be addressed with all stakeholders.

The issue of standards development is already there, including in cyber-security, where the shortage of experts from ATM is a real problem. The risk is that ATM specific issues may not be addressed adequately.

Finally, AI will have to be integrated in a socio-technological ecosystem in ATM that is changing rapidly (legacy modernisation, more information – SWIM, DAT, drones, etc). With the development of a Statement of Concerns (security operations centres), we need to account for another layer of protection, but one that needs to be articulated with safety barriers, and account for human actions as well as technological / AI activities. An overall perspective that brings experts with various backgrounds is needed to progress our safety / security culture.

3. Fields of Application

3.1 General Applications of AI

The use of AI and ML have permeated many industries. However, aviation is a safety-critical domain in which there is very little tolerance for failures. The deployment of any automation, particularly those driven by AI systems, will need to go through a rigorous verification, or certification process, before deployment.

In general, the stringent requirements of aviation have contributed to the slow adoption of AI and ML in aviation. However, many aviation sectors are increasingly adopting AI systems, ranging from automating simple, yet tedious and repetitive tasks to a more sophisticated application of a complex autonomous air traffic control system to de-conflict traffic¹⁹.

ANSPs can benefit from employing AI systems to improve both the efficiency and safety aspects of air traffic control. It is unlikely machines will replace humans in such safety critical domain for the foreseeable future. But they can be deployed to enhance situational awareness and provide decision support to increase safety and efficiency by enabling human experts to focus on more complex decisions. The role of Digital Assistants and decision-aid systems, will be discussed further below. A later paper in this series by the CANSO Strategic Technology Workgroup will examine the role of AI in aircraft systems, and their impact on ATM operations²⁰.

As noted throughout this paper, safety concerns are paramount in our industry and most applications of AI will have to account for this aspect and contribute to it. Several specific actions should be taken to understand potential risk and increase safety:

Risk Metrics: Metrics is one of the key tools to monitor the safety performance of complex systems like air traffic control. Developing risk metrics is a complex process which involves prediction of risk and identification of different outcomes of undesired safety events. For instance, a metric that helps monitor the risk on the surface of an airport environment needs to detect and identify events such as runway excursions, runway incursions, and taxiway incidents and assign appropriate numerical indices proportional to the outcome of each event. An accident with a fatality, injury and aircraft damage will have a corresponding weight for each outcome and an incident that involves no outcome can be assigned a severity weight based on its probability of becoming an accident. Models that support such metrics need to be able to process data from diverse sources. An example of such application is surface safety and airborne safety metrics deployed by the FAA²¹.

Event Classification: Air traffic systems have a safety management system in which one of the critical activities is collecting and analysing safety events through different means. The two primary data collection processes include a reporting system and automated capturing of rule violations and separation losses (close encounters). The main goal of collecting safety events is to identify potential risk in the system by proactively detecting unsafe practices and operational violations. The process requires the categorisation of different safety events by operation types and environments such as surface incidents and airborne encounters. The manual classification of these events is not only laborious, but it is also error prone. In large reporting systems, like NASA ASRS²², which process a large number of aviation safety reports from many airlines and operators,

the manual categorisation is so intensive that they only process a very small percentage of the reported safety events. Machine learning can be a critical tool in this area both from efficiency and safety perspectives. Supervised learning models can be trained and deployed using large amount of data collected over several decades by many safety action programs. These models can categorise not only the type of events but also severity levels. One such example application is a system recently deployed by the FAA Air Traffic Organization to classify runway incursions into four severity categories (A, B, C, D) according to ICAO rules. The identification process largely relies on event descriptions as reported by pilots and controllers.

Anomaly Detection: The majority of aviation data is not annotated (labelled). In addition, an overwhelming amount of the data is for normal operations due to the simple fact that aviation is a very safe domain in which accidents and serious incidents are very rare. Therefore, in some segments of the system, supervised learning is impractical due to the difficulty of rare event modelling, which is caused by the data imbalance issues. Hence, to identify a small number of safety events from a large amount of normal operations, anomaly detection is the most viable alternative. One application of anomaly detection in safety is to analyse all flights and identify those that manifest different characteristics, which would be considered anomalous compared to a large portion of the flights. Domain experts can validate once a smaller subset of events is identified. NASA recently demonstrated this application on several airports in the U.S National Airspace System using a kernel method²³. Another application scenario of this algorithm is called Sequence Miner, which detects and characterises anomalies in large sets of high-dimensional symbol sequences that arise from recordings of switch sensors in the cockpits of commercial airliners²⁴.

Safety Key Performance Indicators (KPI): Many of the known safety KPIs are either collected through manual reporting or detected through rule-based deterministic algorithms that are rigid; these algorithms also tend to generate a higher rate of false positives. Some of commonly monitored safety KPIs include, airspace excursions, missed approaches, unstable (low/high energy) approaches, opposite direction operations, runway overruns, rejected take-offs, and approach overshoots/undershoots. There is a huge potential to leverage machine learning to automate the detection of these KPIs given that there is large amount of data collected over time. In addition to providing a scalable solution that doesn't rely on manual detection, ML-based solutions will provide a stable detection mechanism that is consistent with existing knowledge about these KPIs encoded in data. Recently, the FAA Air Traffic Safety Organization demonstrated a voice-based ML model to detect UAS encounters with manned aircraft which provided a highly accurate result. This suggests that some of the KPIs can benefit from leveraging voice data in addition to events/accidents collected over the years.

3.2 Application Examples for ATM/ANSPs

This section describes some recent applications of AI/ML for ATM/ANSPs, in particular related to ATC, Tower, ATFM, and ground operations.

Conflict Resolution Advisory: AI can provide tools for helping ATCOs resolve (especially) mid-term conflicts, once these are detected. Indeed, ML algorithms can provide optimal solutions as a sequence of actions solving the conflict and satisfying a certain number of intended qualities: for example, ensuring preferred safety distance, having a minimal number of commands, reducing fuel consumption or delay and others. These methods can learn from ATCO experience and also explore billions of possible solutions and choose among them the best according to the above-

mentioned criteria. An interesting observation is that mid-term conflict resolution puts human in the decision loop (ATCO can change the suggestion provided by the algorithm), while in Short-Term Conflict Alerts (STCA) the trend is rather to empower automatic resolutions. These are interesting new developments for airborne collision avoidance systems, also achieved through ML in recent standards delivered by RTCA SC-147 under the ACAS suite (ACAS, ACAS-Xa, -Xu, -sXu) for airborne collision avoidance with or without an on-board pilot²⁵, which enable the ACAS resolution advisories to be automatically followed by the autopilot.

The first step in this automation approach is to use the ATCO experience to train an algorithm that after a conflict identification can suggest the solution ranked on the basis of the ATCO behavior.

Voice Recognition: AI algorithms and methods can be used in the areas of Speech-to-Text, context-sensitive AI dialogues (Natural Language Processing [NLP] and Understanding [NLU]) and Text-to-Speech to automate communication between air traffic controllers and pilots, as well as a decision automation system ATM process automation. It should be noted that NLP is quite difficult on narrow-band audio signals received over analogue radio. Most work currently uses local recordings of controller audio that is available in much better quality.

Machine learning and deep learning algorithms and language models can be used for optimising the audio quality of analog audio data from ATC radio. The identification and recognition of speakers can be used for assigning flight identification and pilot(s) phrasing and language. Transcription of voice into text with a low Word Error Rate (WER) can be used for documentation purposes and for downstream NLU processing and content recognition²⁶.

Machine learning, deep learning algorithms, and NLU models can be used for air traffic control for specially formulations, dialogues and phraseology of the standardised and “common (deviating from the standard)” air traffic controller and pilot language²⁷.

Content recognition and semantic can be used to understand radio messages, i.e. detection of entities (e.g. flight number, altitude, etc.) and intent (target of the radio message, e.g. clearance of runway, approach, entry of airspace, clearance of taxiway, crossing of runway, etc.).

Automation of communication between air traffic controllers and pilots, monitoring and corrective action identification will enable air traffic controllers to focus more on safety-critical, complex situations.

A step further could be the interface between the voice recognition system and the working positions software for both ATCOs and Pseudo-pilots. This approach could optimise the workload of the operators in terms of HMI interaction.

Optimal Sectorisation: In the Dynamic Airspace Management domain, ML/AI techniques can introduce a significant improvement. The basic idea is to train a machine with traffic data and sector configuration and let it learn how to re-shape sectors boundaries in order to better accommodate the traffic.

The training phase can be based on Reinforcement Learning approach according to which the Agent learns how to change the sector boundaries in order to solve the hotspots where more Capacity is required. The whole process will lead to a mix of already defined volumes plus new sector shapes which determine the entire sectors configuration for use during next operational timeslot. The dynamic configuration of the airspace through different timeslots could lead to an optimisation of the workload for the ATCOs.

In this context of improved performance for flow management at ACC level, the SINAPS project (SWIM Integrated Network management and extended ATC Planning Services) has delivered a solution for DSNAs which was deployed in 2019 in an operational environment in France. The SINAPS project supports both tactical and strategic operations²⁸. The next step will consist of addressing cross ACC optimisation.

Demand Prediction: ML/AI algorithms can be applied to improve the accuracy and predictability of traffic demand by adding to the current calculation methods, the complementary stochastic information deriving from historical data intelligent processing aimed to provide hidden patterns between flight plans and actual flown trajectories.

This kind of AI-augmented demand enables the improvement and the efficiency of the entire Demand and Capacity Balance (DCB) process.

Remote Towers: AI based pattern recognition can be used to combine multiple video tracks of one detected object (separate tracks on wings, body, etc.). The entire shape of an aircraft or vehicle can be recognised from different viewing angles (e.g. back profile, side profile, front profile). It can also be used to detect standing aircraft or vehicles on the RWY or holding point, stands or parking position and continue tracking if such objects start moving. Or in a close-up view (pan-tilt-zoom or static hot spot cameras) to recognise no landing gear is down or special situations (e.g. smoke or engine fire) to alert the controller. It can also be used to filter output data of conventional video tracker and remove false bounding triggered by cloud movements.

Unmanned Traffic Management: UAVs with high levels of automation and advanced features will require AI to drive the development of new UTM/U-space services and address the ATM integration challenge. AI will support ATM actors and increase safety from flight planning to operations by providing new solutions for conflict detection, traffic advisory and resolution, and cybersecurity. The change in both UTM and ATM will entail increased levels of automation and applying AI/machine learning at the strategic, pre-tactical and tactical aspects to enable the transition from airspace-based to trajectory-based operations.

Scalable and reliable performance of the UTM framework requires real-time processing of large amounts of shared data. ML algorithms are effective in assisting UTM functions like Operation Planning (prediction, classification, and optimisation of trajectories), Situational Awareness (weather, terrain and obstacle awareness), Failure Detection and Recovery (detect, assess and classify the failure, apply a failure recovery plan), and Remote Identification (data security and analysis)^{29 30}. Similar to ATM, one of the main challenges in UTM is to design efficient Conflict Detection and Resolution (CDR) approaches, ensuring conflict-free paths for all UAVs in both pre-flight and in-flight phase, that can be supported with multi-agent pathfinding (MAPF) algorithms. Highly autonomous decision making in UTM is managed by AI and dependent on the sophistication of the solutions.

This will be further explored in a later whitepaper in this series by the CANSO Strategic Technology Workgroup examining UTM and its implications for ATM³¹.

Collaborative Decision Making: Decision-making is a central aspect of the ATM industry and, there is an increasing ambition for building AI that will support people in their everyday life and work. So it is more important than ever that AI algorithms are designed and optimised in a way that simplifies and accounts for human-AI collaboration. Addressing these challenges requires methods and techniques at the intersection of cognition, sociology, and artificial intelligence. Several

research initiatives demonstrate the benefits of the combination of humans and artificial intelligence to solve very complex problems. This new kind of heterogeneous collaboration can improve even more the collaborative decision making (CDM) process inside air traffic flow management (ATFM) application field. Currently, the CDM process is restricted to humans only who, supported by collaboration tools more or less integrated into ATFM platforms, try to make the best decision with respect to the traffic flow hotspots happening during daily operations. The information they share starts with a set of potential Traffic Management Initiatives (TMIs provided by AI) to apply and a set of preconfigured variables extracted from the traffic issues they want to solve. Here is where AI can assist by providing 'hidden layers' and related synthetic descriptors thus enhancing the scenario with more relevant indicators that can help flow managers and their stakeholders to easily agree on a common shared and, even more realistic picture. Last, but not the least, it is important to stress that fairness in human-AI collaboration processes is crucial because it enables trust between humans and machines and between the stakeholders around the CDM process.

Digital ATCO & Digital Twin: Powerful ML techniques can provide digital assistance to ATCOs in many ways:

- By a natural human machine dialogue, enabled by automatic speech, natural language processing (NLP) and intent recognition. Thus, ATCO orders as well as pilot voice can be directly transcribed into the system, to increase ATCO efficiency and globally the safety of the dialogue (SESAR projects MALORCA³²).
- By increasing situation awareness, through an increased precision of data surveillance.
- By learning from previous ATCO actions and providing decision-aid for increased efficiency AI will complement ATCOs enabling them to enhance their capabilities, reducing their cognitive involvement in repetitive tasks and allowing them to focus on highly value-added tasks. This requires a high degree of system stability and backups to guard against AI failure. It must be recognised that the more you automate the greater the potential for skill fade (or even skill absence if the ATCO has never been trained to work without AI). In the event of a system failure an ATCO may find themselves in a situation they are not equipped to get themselves out of. There needs to be a focus on the scope of AI functionality and expected cross-over of skills and capabilities between the human and AI.
- By providing simulation power and modelling through smart Digital Twins, to experiment What-ifs in tactical and strategical situations.

Intelligent Label Placement: During aeronautical charts production and particularly in the cleaning phase, there are a lot of manual activities aimed at producing clean and readable charts by moving, resizing or placing labels without overlap or related features. Those activities are very time consuming and repetitive at each production cycle.

Moreover, in ATM situational awareness applications, where a visual component is present, the same problem described for aeronautical charts takes place and is even more challenging to solve because of the (near) real-time processing needed to finally display the situation without labels conflicting.

By using AI algorithms, it is possible to support both use cases thus:

- Giving a boost to chart production in terms of both efficiency and reliability.

- Providing a better user experience while using air situation display consoles because the optimal de-conflicted placement of the labels is done in real-time while the user performs pan/zoom-in/out actions.

In the end, the use cases referenced can be integrated by using an ML approach, thus obtaining for situation awareness case an intelligent label placement engine able to reproduce in real-time fashion the cartographic expertise used in charts aeronautical charts production.

Predictive Maintenance: ATM operations are supported by disparate systems, built on electronic/mechanical integration. Each of these systems often provide data by means of sensors on some of their components Typically this information is collected and, where applicable, used to provide alerting mechanisms when something is already going wrong.

Considering that even the smallest system has the possibility to provide its status constantly via Internet Protocol connectivity, the idea is to collect data from so-called “internet-of-things” devices, inside ATM domain, and apply ML techniques in order to extract incrementally:

- Behaviour model
- Strange/unwanted behaviour
- Hidden failure patterns

With the availability of this new abstract layer of information, it is possible to anticipate issues deriving from failures and to plan maintenance in a more reliable and preventive way.

Weather Impact Prediction: In the areas where weather conditions suddenly change and/or affect significantly the capacity of air traffic flow network, ML/AI can play an important role by predicting capacity reduction of specific network resources like runways, sectors, airways, waypoints, airports.

By correlating historical weather data and capacity numbers, ML can provide improved solutions around predicted severe weather events by assigning reduced capacity numbers affected by weather phenomena.

Some experiments have already been conducted with successful results demonstrating the ability to predict reduced airport capacities because, among the others, of severe weather phenomena (icing).

Performance Data to Improve Trajectory Prediction: The table-based performance data (i.e. BADA) suffers from insufficient climb/descent rate precision. These data are very important for the trajectory calculation algorithm. Constructing a Data Lake with the flight data of different companies in different geographic areas taking into account the different elements that contribute to define the flight profile could enhance the methodology for calculating aircraft performance. This could provide more precise trajectory calculation for flight data processing and could help to generate a more precise simulated traffic.

System Engineering: To have a continuous verification of the ATM system developments it is important to define a common toolchain with the industry in order to share software engineering artifacts. Currently, the verification of the artifacts requires a lot of effort. The usage of AI based automation tools can simplify the process. The objective is to have processes of verification that are less error prone and with high level of automation for the formal artifacts, and the integration of the formal artifacts for the (semi) automatic generation of non-formal software engineering documents.

4. Conclusions

The development of AI has come a long way in recent years. The availability of enormous quantities of data allied to the growth of highly effective and efficient computing power has enabled the use of AI in ATM and related fields, as has been explored in this paper.

The use of machine learning and deep learning systems is exciting and offers opportunities to increase the safety and efficiency of services provided by ANSPs through a variety of different tools and systems. The expanded use of AI in ATM has the potential to bring significant benefits, including: improving safety (if properly assured) by understanding anomalous traffic patterns; improving capacity by supporting decision making and understanding changes in flows; and finally, bringing environmental benefits by optimising flight plans.

The modern development of AI must give sufficient attention to the trust considerations if it is to achieve broad use in ATM. Trustworthy AI has to begin by using sound engineering practices, mandated by agreed and accepted standards. Too much of AI in the commercial world thus far has consisted of short-term solutions, algorithms that get a system to work immediately, without the critical layer of safety assurance and engineering guarantees that are taken for granted within the aviation world. The focus on the development of accepted standards for both systems and datasets is a strong enabler for AI adoption by manufacturers and end users. It is clear there has to be a focus on developing appropriate safety metrics (based on standardised approaches), complemented by cyber security standards, and different approaches to safety and security assurance. We describe the challenges of the 'black box' output paradigm which limits 'explainability'; this has to be overcome in order to be able to make an effective safety case.

Notwithstanding the above, this paper describes only a small set of areas where AI is already making a difference. We are convinced that AI will continue to be adopted in a wide range of applications in ATM.

References

- 1 Friedman, Jerome H. "Greedy function approximation: A gradient boosting machine." *Annals of statistics* (2001): 1189-1232
- 2 Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." *Advances in Neural Information Processing Systems*, 2017
- 3 Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "Why should I trust you?: Explaining the predictions of any classifier." *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM 2016
- 4 Marco Tulio Ribeiro, Sameer Singh, Carlos Guestrin. "Anchors: High-Precision Model-Agnostic Explanations." *AAAI Conference on Artificial Intelligence (AAAI)*, 2018
- 5 ED-109A/DO-278A & ED-153 with ED-218/DO-331 for Ground or ED-12C/DO-178C with ED-218/DO-331 for Airborne
- 6 EASA 2020, Artificial Intelligence Roadmap 1.0, <https://www.easa.europa.eu/newsroom-and-events/news/easa-artificial-intelligence-roadmap-10-published>
- 7 Intelligent security tools – ncsc.gov.uk
- 8 A sociotechnical approach to cyber security – ncsc.gov.uk
- 9 DataLake Manifesto, <https://tdwi.org/Articles/2017/10/16/ARCH-ALL-Data-Lake-Manifesto-10-Best-Practices.aspx?Page=1>
- 10 GAIA-X initiative, <https://data-infrastructure.eu/>
- 11 NM B2B, <https://www.eurocontrol.int/service/network-manager-business-business-b2b-web-services/>
- 12 Data4Safety, <https://www.easa.europa.eu/newsroom-and-events/news/data4safety-partnership-data-driven-aviation-safety-analysis-europe/>
- 13 Data Beacon project, <https://databeacon.aero/>
- 14 SkyWise, <https://skywise.airbus.com/>
- 15 IATA Open Air, <https://www.iata.org/en/programs/passenger/open-air/>
- 16 FAA NextGen, <https://www.faa.gov/nextgen/>
- 17 ICAO product, <https://data.icao.int/icads/Product/View/27>
- 18 RTCA SC-147, <https://www.rtca.org/sc-147/>
- 19 Brittain, M., & Wei, P. (2019). Autonomous Air Traffic Controller: A Deep Multi-Agent Reinforcement Learning Approach. *ArXiv*, *abs/1905.01303*.
- 20 CANSO Strategic Working Group (S&I STWG), "Airborne tools impacting ATM operations", White Paper, to appear, 2021
- 21 F. Bati and L. Withington, "Application of Machine Learning for Aviation Safety Risk Metric," 2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC), San Diego, CA, USA, 2019, pp. 1-9, doi: 10.1109/DASC43569.2019.9081657.
- 22 NASA ASRS, <https://asrs.arc.nasa.gov/>
- 23 S. Das, B. L. Matthews and R. Lawrence, "Fleet level anomaly detection of aviation safety data," 2011 IEEE Conference on Prognostics and Health Management, Montreal, QC, 2011, pp. 1-10, doi: 10.1109/ICPHM.2011.6024356.
- 24 NASA open-source projects, <https://ti.arc.nasa.gov/opensource/projects/sequenceminer/>
- 25 RTCA SC-147, <https://www.rtca.org/sc-147/>
- 26 Shuo Chen, Hunter Kopald, Rob Tarakan, Gaurish Anand, Karl Meyer (2019) Characterizing National Airspace System Operations Using Automated Voice Data Processing 13th USA/EUROPE Air Traffic Management R&D Seminar, 17.-21. June 2019, Vienna, Austria.
- 27 Helmke, Hartmut and Kleinert, Matthias and Rataj, Jürgen and Motlicek, Petr and Klakow, Dietrich and Kern, Christian and Hlousek, Petr (2019) Cost Reductions Enabled by Machine Learning in ATM How can Automatic Speech Recognition enrich human operators performance? 13th USA/EUROPE Air Traffic Management R&D Seminar, 17.-21. June 2019, Vienna, Austria.
- 28 Presentation video of the project: "SINAPS DYNAMIC CONFIGURATION at DSN, France" <https://www.youtube.com/watch?v=W5s4YbkGPLg>
- 29 Ho, Florence, Goncalves, Artur, Salta, Ana, Cavazza, Marc, Gerald, Ruben and Prendinger, Helmut (2019) Multi-agent path finding for UAV traffic management: Robotics track. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2019)*. Association.
- 30 Abdalla, A. S., & Marojevic, V. (2020). Machine Learning-Assisted UAV Operations with UTM: Requirements, Challenges, and Solutions. *arXiv preprint arXiv:2006.14544*.
- 31 CANSO Strategic Working Group (S&I STWG), "Impact of new concepts stemming from UTM to ATM", White Paper, to appear, 2021
- 32 SESAR project MALORCA, <https://www.sesarju.eu/projects/malorca/>



Visit us:
canso.org

