Aeroworkshop: artificial intelligence in the aviation industry

Current status, trends and challenges, opportunities

HumaneAI event

16 MAY 2024

Foreword

In collaboration with the Humane AI (https://www.humane-ai.eu/) project, Airbus & DFKI have organised a closed-door industry online workshop on Artificial Intelligence research areas for aeronautics. Under the patronage of Paul Lukowicz (DFKI) and Marc Contat (Airbus), this workshop has been organised by Jakob Karolus (DFKI), Guillaume Gadek (Airbus) and Louis Lefebvre (Airbus).

Humane AI explores ethics issues, the role of humans in the operation of AIs, as well as the impact of AI on society and the environment. Indeed, AI is already modifying all lines of businesses and disrupting areas of the economy. This report, resulting from the workshop, aims to provide an informed insight about the role of AI in the aeronautics industry.

The workshop is structured around three pillars. Each pillar includes presentations from selected speakers to provide ground for discussion, followed by an interactive session using a digital whiteboard (Miro), during which participants were asked to contribute and discuss their experiences, insights and opinions. To conclude the pillar, a question & answer session enabled the participants to clarify doubts and open new perspectives.

A first pillar of discussion discusses the current, and future, uses of AI in any area of the industry, such as conception, manufacturing, automation of the engineering and assembly phases, aircraft in operations (incl. Predictive maintenance), flight and navigation tools (incl. towards unmanned aircraft), use of AI for logistics (cargo) and services (to passengers and/or airliners). The participants were welcomed to expose both AI “success stories” or lessons learned, and to propose potential emerging technologies and applications.

A second pillar covers the certification and regulation aspects. The aerospace industry is already highly regulated; new AI regulations such as the EU AI Act are likely to influence or impact this domain. An overview of the challenges, regulatory constraints and potential solutions will enable the participants to better understand and recommend efforts towards certifiable and trustable AI for the aeronautics industry. The objective is to highlight potential gaps between airworthiness certification & regulation requirements, and desirable features for aerospace humane AI.

Last but not least, the third pillar is named “Better through AI”, to take into account new concerns and considerations about AI in the aeronautical domain. The participants are welcome to expose AI-related ideas, prototypes or efforts a) towards a greener aeronautics industry, and b) about ethical considerations for AI in the aeronautics industry. Notably, the workshop aims to discover and recommend future activities to address the environmental and societal impacts of AI.

Table of Contents

Speakers

Name	Organisation
Pooja Narayan	Airbus Defence and SPace
Romaric Redon	ANITI
Alexis De Cacqueray	Airbus Defence and Space
Benjamin Laroche	Airbus
Antonio Monzon Diaz	Airbus Defence and Space
Thirizi Belkacem	Airbus Protect
Guillaume Soudain	EASA
Nicolas Schneider	Airbus
James Crowley	Grenoble INP
Arnault Ioualalen	Numalis
Vanessa Vohs	Bundeswehr University Munich
Mohamed Chetouani	Sorbonne University

Attendees

In addition to the speakers, more than 20 participants attended the workshop and provided valuable insights and inputs during the discussion session. They were affiliated with the following organisations.

Organisation
Grenoble Institut Polytechnique
DFKI GmbH
Fortiss GmbH
DLR
Universidad Politécnica de Madrid
Deep Blue
Diehl Aviation
IABG

Detailed Agenda of the Workshop

Topic	Participant	Subject	Time slot
Introduction	**Organisers**	**Welcome word**	**9:00**
Current, and future, uses of AI - 9:15	Pooja Narayan	Airbus AI roadmap	9:15
	Romaric Redon	ANITI Trustworthy AI methods for introduction of AI in aviation.	9:35
	Alexis de Cacqueray	Project EICACS	9:55
	James Crowley	A Research Roadmap for the AI Copilot	10:15
10:45 - Break - 15 min			
Interactive Breakout Session	MIRO	Brainstorming & gathering insights	11:00
Panel Discussion	Moderator: Guillaume Gadek Panel: Presenters of the Pillar	Q&A and Discussion	11:30
12:00 - Lunch Break - 1 hour			
Certification and regulation aspects - 13:00	Benjamin Laroche	Regulating AI @Airbus	13:00
	Antonio Monzon Diaz	EUROCAE WG114	13:20
	Arnault Ioualalen & Vanessa Vohs	AI & Robustness	13:40
	Thiziri Belkacem, Guillaume Soudain	MLEAP: insights from the project	14:00
Interactive Breakout Session	MIRO	Brainstorming & gathering insights	14:20
Panel Discussion	Moderator Marc Contat Panel: Presenters of the Pillar	Q&A and Discussion	14:50
15:20 - Break - 10 min			

Ethical and societal concerns - 15:30	Mohamed Chetouani	Human/Robot interaction	15:30
	Nicolas Schneider	Sociological acceptance of new technologies	15:50
Interactive Breakout Session	MIRO	Brainstorming & gathering insights	16:10
Panel Discussion	Moderator: Jakob Karolus Panel: Presenters of the Pillar	Q&A and Discussion	16:40
Conclusion	**Organisers**	**Wrap-up word**	17:10

Pillar A: Current, and future, uses of AI

This pillar aims to provide an up-to-date shared picture of the current applications of Artificial Intelligence within the aviation industry, to discuss the still relevant challenges of conception, development and deployment of AI in the domain, and to identify the newest trends to come in the next few years.

AI is already pervasive in the domain

In 2024, all steps of the life cycle are already playgrounds for AIs. The Airbus AI roadmap structures the domain along Design, Test and certification (of the aircraft), Production, Delivery and Operations. The aim is to deliver enhanced performances within the products by improving the way to *make* the products, using AI.

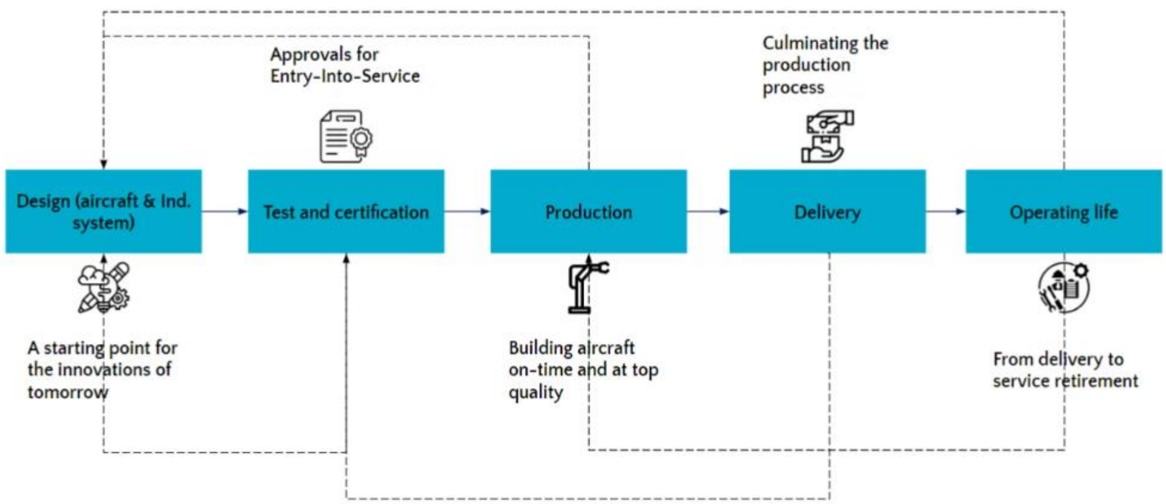

The life cycle of an aircraft: opportunities for AI at each stage. © Airbus

Airbus relies on the development of “demonstrators”, notably through the UpNext entity. Current initiatives include an AI help to land the aircraft in case of emergency, and assistance for air-to-air refuelling, a classic example of high-risk manoeuvre in flight.

Another breakdown of the exploration of using AI in the domain focuses on the type of task and the specific technology that provides the solution. Computer Vision receives much attention to provide cognitive help to pilots, during navigation and during taxi; and also for manufacturing, notably with the provided example of default detection in solar panel cells.

All the technologies of AI seem to be relevant, including Digital Twins, Hybrid Modelling, Knowledge Engineering. The participants noted that AI has actually been present in the domain since decades, with Case Based Reasoning and Expert Systems.

The AI tasks in aeronautics and aviation industry. © Aniti

The effort bears on providing satisfactory instantiations of AI-based solutions to safely fulfil a task. As such, a special care is taken in terms of predictability and reliability. Proofs of Concept are ambitious, but deployment in production requires to eliminate all doubts.

Success stories and applications

It is not possible anymore to exhaustively relay all the initiatives of integrating AI in the aviation domain. To give but an overview, we briefly introduce the following projects. Focusing on the manufacturing process, a Computer Vision algorithm analyses the photographies (either in the visible spectrum or others, such as X-ray) of components such as Solar Panel Cells. The algorithm looks for defaults, here named "cracks". The human analysis of such pictures is tedious and time-consuming, as small cracks require careful attention to be detected. The algorithm enables to increase the exhaustivity of the detection, and to guide the human expert directly on the suspicions.

AI in manufacturing: detecting cracks in solar panel cells. © Airbus

In a totally different scale and application, the project EICACS[1] is grounded in the military aspect of the aviation industry. The main intuition is to conceive the whole system for collaborative air combat, including the relationship between pilots in human-guided aircraft, and unmanned aerial systems. This life-and-death stakes highlight the need for trust during the operations of such a system: the project includes the definition of the required properties of the system in terms of Trustworthiness, as illustrated below.

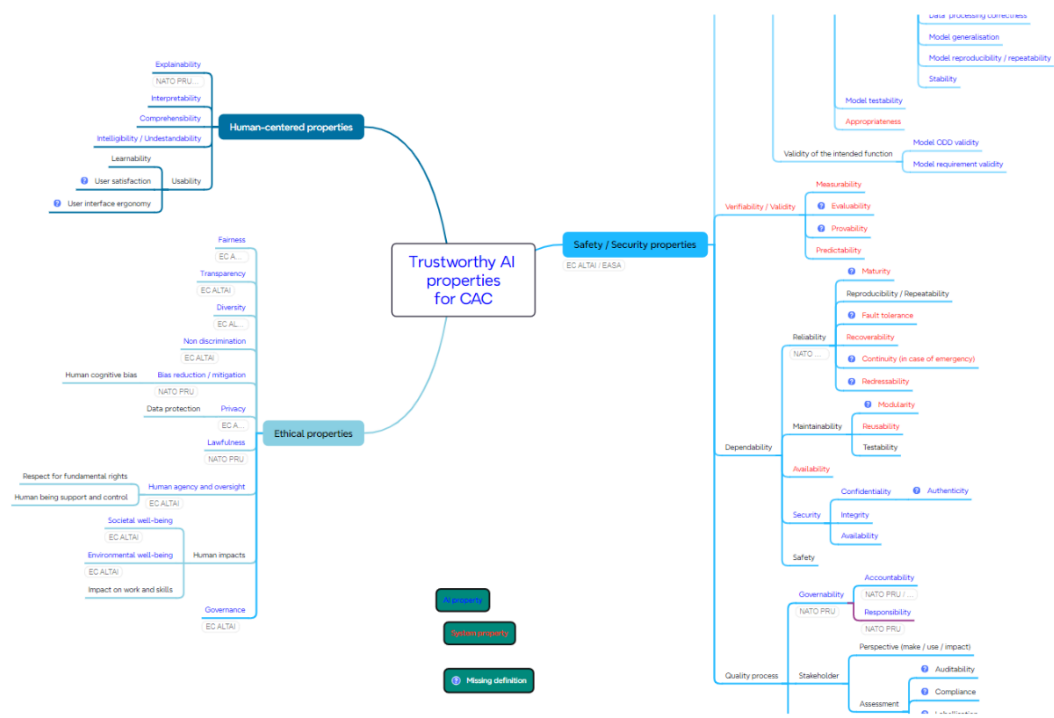

Desired specific properties of AI for a trustworthy usage in collaborative air combat applications. © Airbus

[1] https://www.eicacs.eu/

Among the numerous applications to help the pilots handle their cognitive load, significant effort and progress have been observed in the transcription of Air Traffic Control (ATC) communications[2]. Speech-to-text technologies work very well in most cases; however the very specific vocabulary, extremely noisy environment and highly diverse accents[3] constituted obstacles to their deployment in aircrafts.

Indeed the experience and the practitioners of AI are already present in the aviation industry. However there are still open challenges and usual difficulties to be expected, and the future work remains significant before having a full and wide adoption.

Recurring challenges

The recent years have seen the formalisation and standardisation of the notions of robustness, explainability and acceptability of an AI system, resulting in a shared framework of **Trustworthiness of AI systems**, addressing all the desired “quality” dimensions.

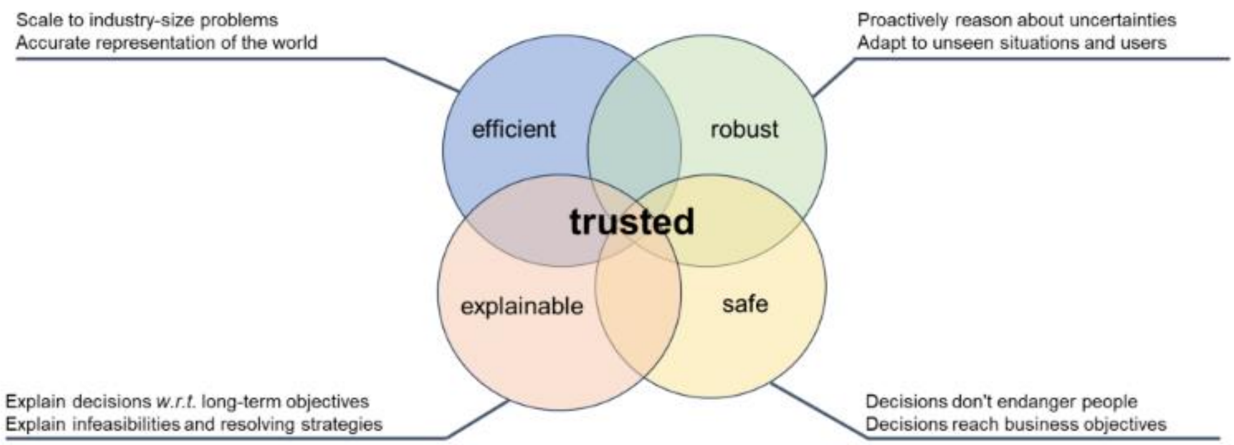

The four characteristics of Trusted AI. © Aniti

An AI system provides the promised value if it is *trusted*. To gain this trust, a number of qualities must be met: the system must be efficient, robust, safe and provide a satisfactory level of explanation: a frequent pitfall appears when it is hard to understand the AI mistakes.

The European Union promotes the notion of “Trustworthiness”, that is the ability to rightfully gain this trust in the AI system. The following diagram includes the seven pillars, also referred to as ALTAI[4] (Assessment List for Trustworthy Artificial Intelligence). The figure also relays other areas not to be forgotten: the specificities of the technical family of artificial intelligence subdomains,

[2] https://arxiv.org/pdf/1810.12614
[3] https://www.isca-archive.org/interspeech_2021/jahchan21_interspeech.pdf
[4] https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment

and the different expectations in the integration tasks; as an example a cultural gap remains between cloud-based and embedded AI components.

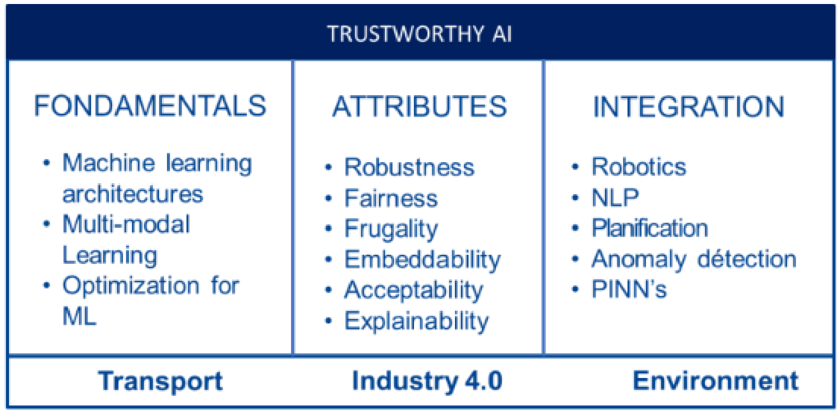

Encompassing view of Trustworthy AI. © Aniti

The participants highlighted the difficulty in finding the correct balance in the human and AI collaboration. AIs generally have better performance than humans but they are becoming more complex. Also, their failures become unforgivable. As an example, the design of an AI assistant like an aircraft co-pilot requires taking into account the available pilot attention in *all* possible situations. A tactile screen only provides visual feedback while the feedback of a lever can be felt through the hands. The tradeoff positions between humans and AIs have been formalised during the HumaneAI project; the following figure ranges from reactive to creative, spanning the whole spectrum of complexity, modality and intensity of the humane-AI relationship.

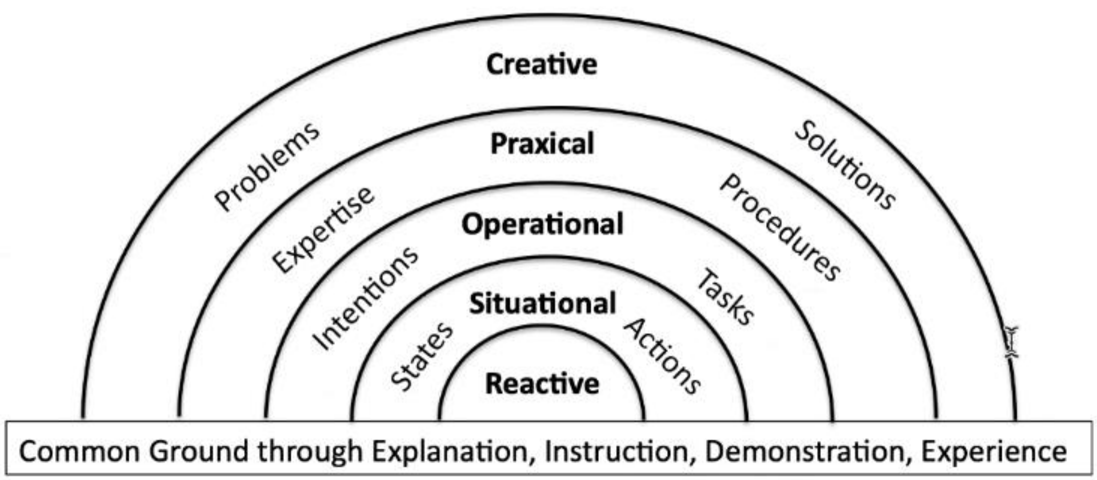

Hierarchical Framework for Collaborative AI, © HumaneAI

Most participants with PoC experience mentioned that each project begins with a phase of data collection or even dataset creation. This means there are still difficulties in accessing or obtaining the required data to perform data science or AI development: hopefully such projects would start from a continuous stream of data as the GAFAM claim to do. Secrecy, security and business constraints (silos) seem to be the key obstacles.

Way ahead and readings recommendations

The Airbus AI roadmap provided the following illustration to structure the way ahead in terms of technological blocks. In the centre, trustworthiness and AI regulation are likely to strongly impact all activities in the field.

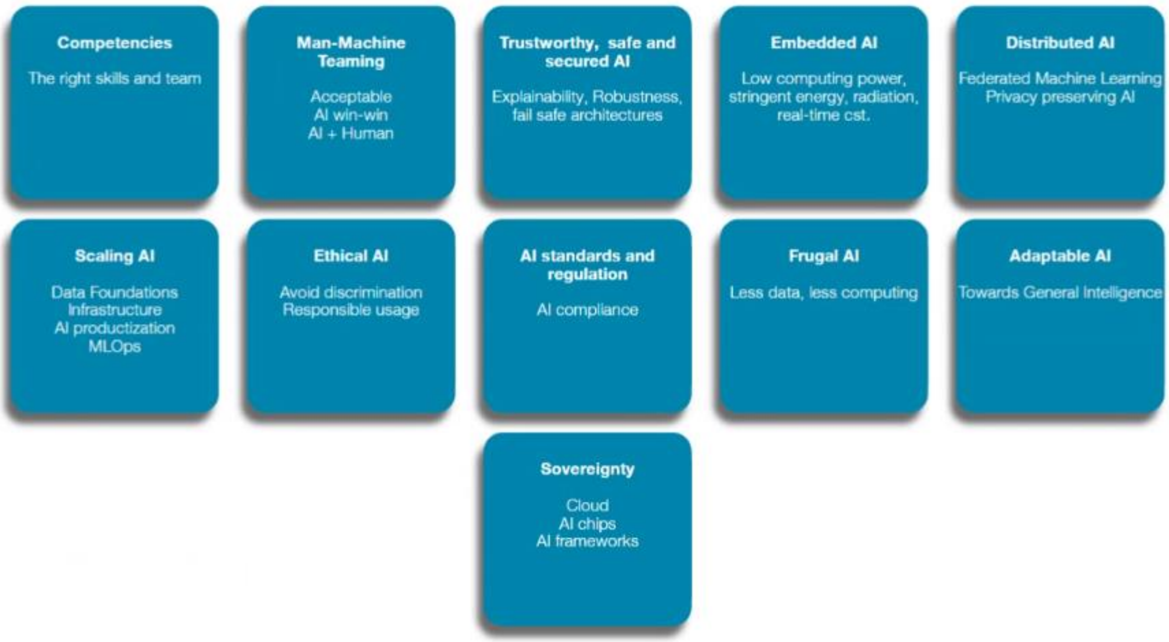

AI roadmap: Technological challenges of AI. © Airbus

An interesting posture consists in exploring both the scalability, including LLMs and Foundation models, and the reduction in size (Frugal AI & Embedded AI). This is a reflection of the constraints bearing either on the ground-based or not-time-constrained applications, or on the on-flight, time-pressured, failsafe setups (the qualifiers do not add up; more than two situations).

Through the interactive session, some participants shared interesting documents and references. This includes the Process Standard for Development and Certification/Approval of Aeronautical Safety-Related Products Implementing AI, ARP6983[5]. A successful Proof of Concept prompting an AI-based collision avoidance system[6] was further relayed. The community also supports a Guidance Concept Paper[7] from EASA.

[5] https://www.sae.org/standards/content/arp6983/

[6] https://www.intel.com/content/dam/www/central-libraries/us/en/documents/2023-02/the-future-of-avionics.pdf

[7] https://www.easa.europa.eu/en/newsroom-and-events/news/easa-publishes-artificial-intelligence-concept-paper-issue-2-guidance

Pillar B: Certification and regulation aspects

This pillar focuses on the legal and regulatory requirements concerning AI and AI in aviation. Some texts are already applicable; a significant number of projects are discussing, developing, implementing and promoting other texts, standards and norms.

A mature regulatory framework

Notably through the EU AI Act, the participants highlighted the strong will from the legislator to regulate artificial intelligence deployments in an encompassing manner. The AI act provides a complete framework around the use of AI, beginning with the notion of Risk Level.

EU Artificial Intelligence Act: Risk Levels. © Airbus

In aviation, and even more for on-flight AI components, the assessed risk level is generally set at "high risk". Unacceptable risk applications are not developed. To comply with the law, the development team and the stakeholders must be able to understand AI, both in terms of technology and regulations.

The following figure provides a summary of the main obligations, spanning all stages of any AI project. Such projects now require strong support in terms of risk management and compliance with the enterprise data governance policy. They must also pay special attention to the role of humans, their oversight capabilities, the auditability mechanisms and the monitoring functions. This comes on top of "traditional" tasks, consisting in technical documentation, focus on accuracy and robustness, and explainability, all targeting the good technical functioning of the AI capability on its task.

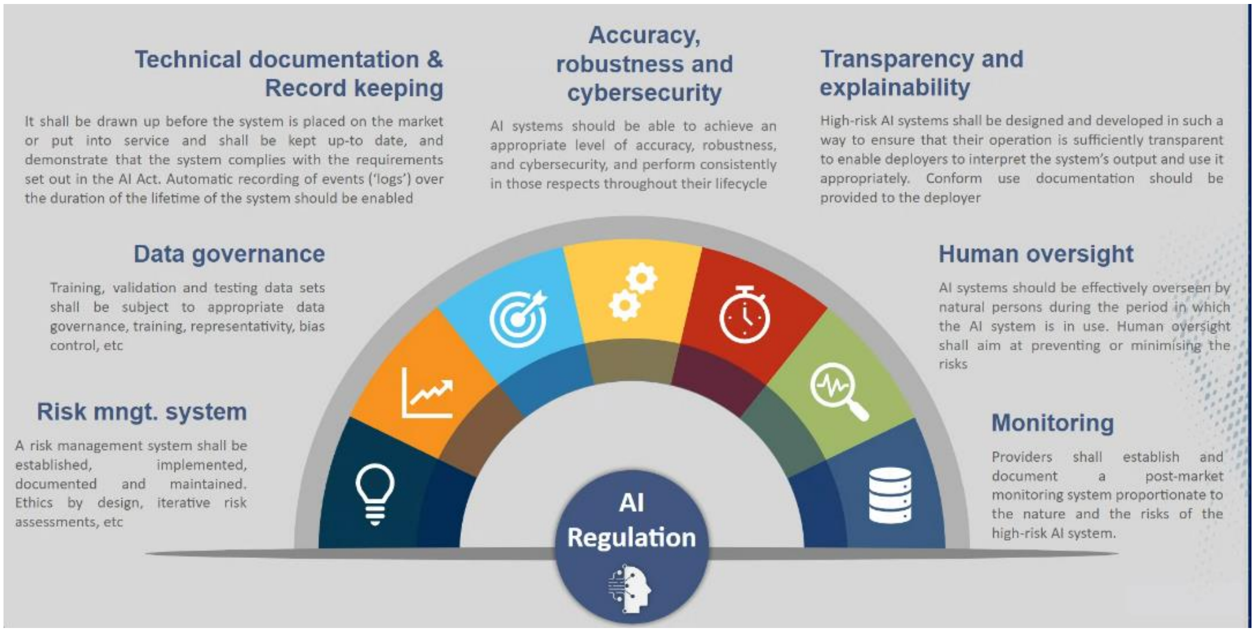

High risk AI systems - ex ante *and* ex post *obligations. © Airbus*

While the AI act provides a framework valid for all AIs, there are more aviation-specific texts to regulate aeronautical applications. The next figure provides the currently released, and work-in-progress, references that already enable AI compliance in this industry. These texts include European Union laws; regulations issued by the EASA; and recommended practices and standards for hardware, software and systems. All these texts include safety as the top quality to be focussed on.

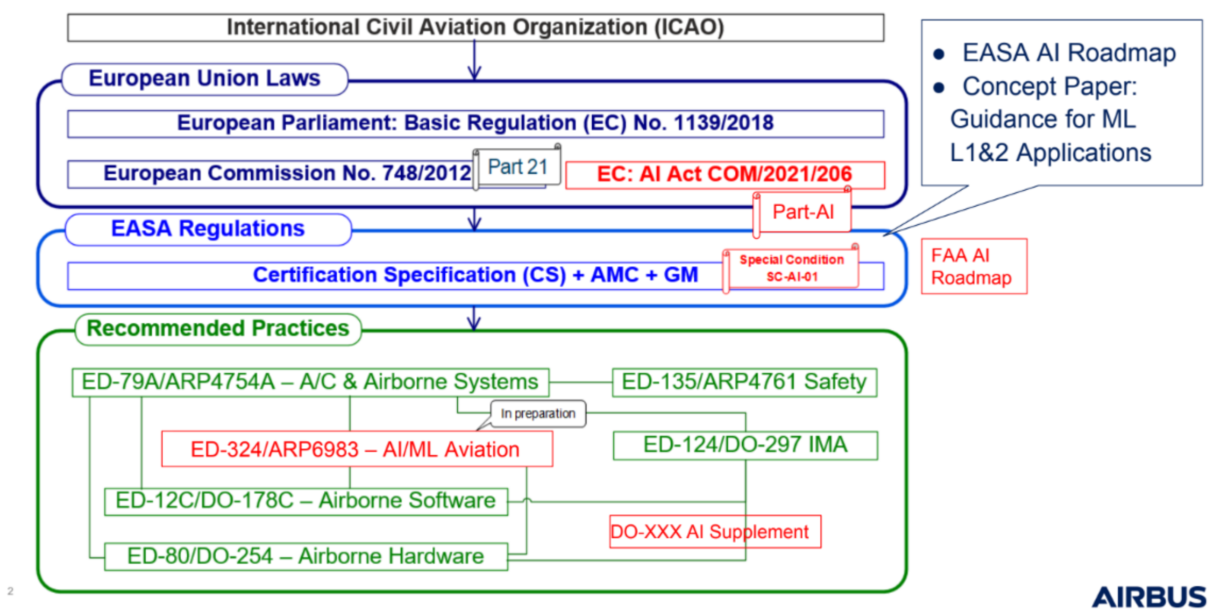

Aviation regulatory framework for AI applications. © Airbus

Initiatives to adopt, enforce, and develop the regulatory framework

The defence-only applications are not regulated by the AI Act; however the customers already require AI providers to deliver their AI with the same documentation package and level of exigence at documenting and characterising all the aspects of the AI component. The participants mentioned a regulatory effort which is “foreseen” in the military domain, which can be put in relation with the diplomatic effort on responsible use of AI[8].

As a mirror to the on-going civil aviation regulatory framework, encompassing efforts are being undertaken to build the military aviation regulatory landscape.

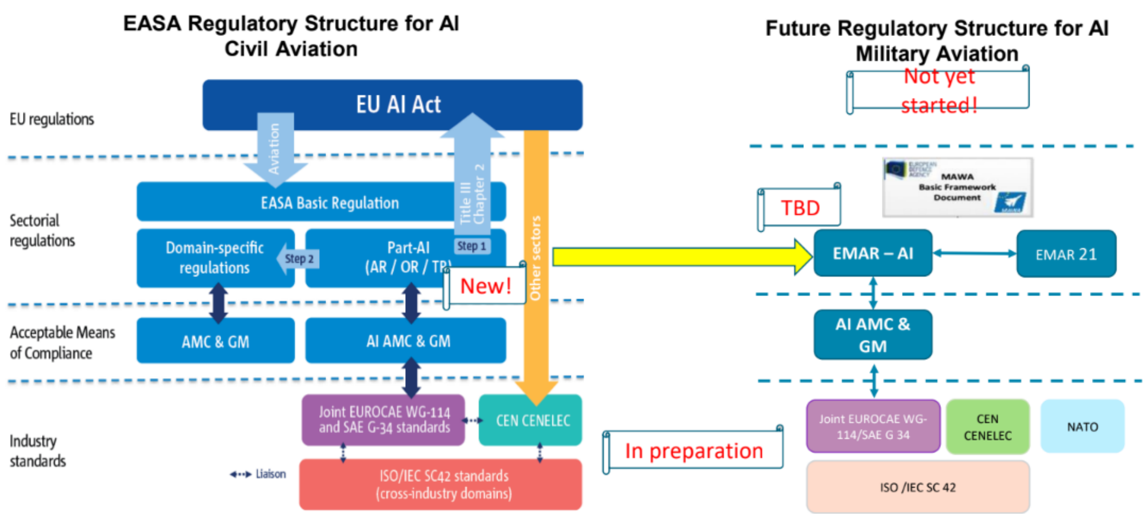

Future regulatory framework for AI-based systems in aviation. © Airbus

In order to promote and adopt the new best practices of ethics within AI, the AI4DEF project explores four use cases of military applications of AI. The AI4DEF project[9] is an EIDP set to pave the way for accelerated development and application of AI in defence. In these applications, the speed and data volumes overwhelm the humans: they definitely need automation in support. However, the human tendency to automation bias is clearly identified as a problem, in terms of skillset loss, misplaced trust and broken chain of responsibility in a high-risk environment.

The military domain comes with its own stages of operations, from the political, strategic level through operations, planning and then execution at the tactical level. Within this last stage, the “killchain” is composed of six steps: find, fix, track, target, engage and assess the impact of the action. This breakdown enables the designers to clarify the role of an AI and its associated risk level. The relevant control level differs with respect to preparation or attack, and also with respect to the step of the killchain.

[8] https://reaim2023.org/
[9] https://ai4def.com/

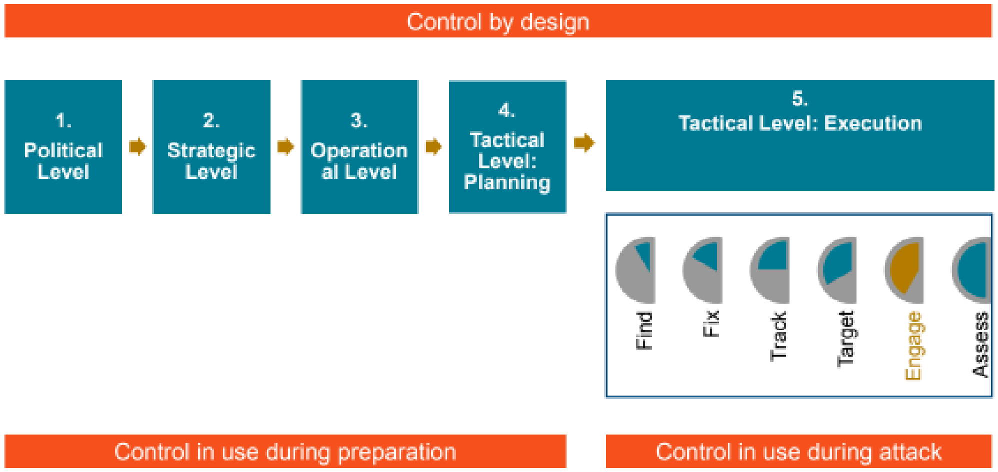

Exploring the use of AI in the killchain: towards meaningful human control. © AI4DEF

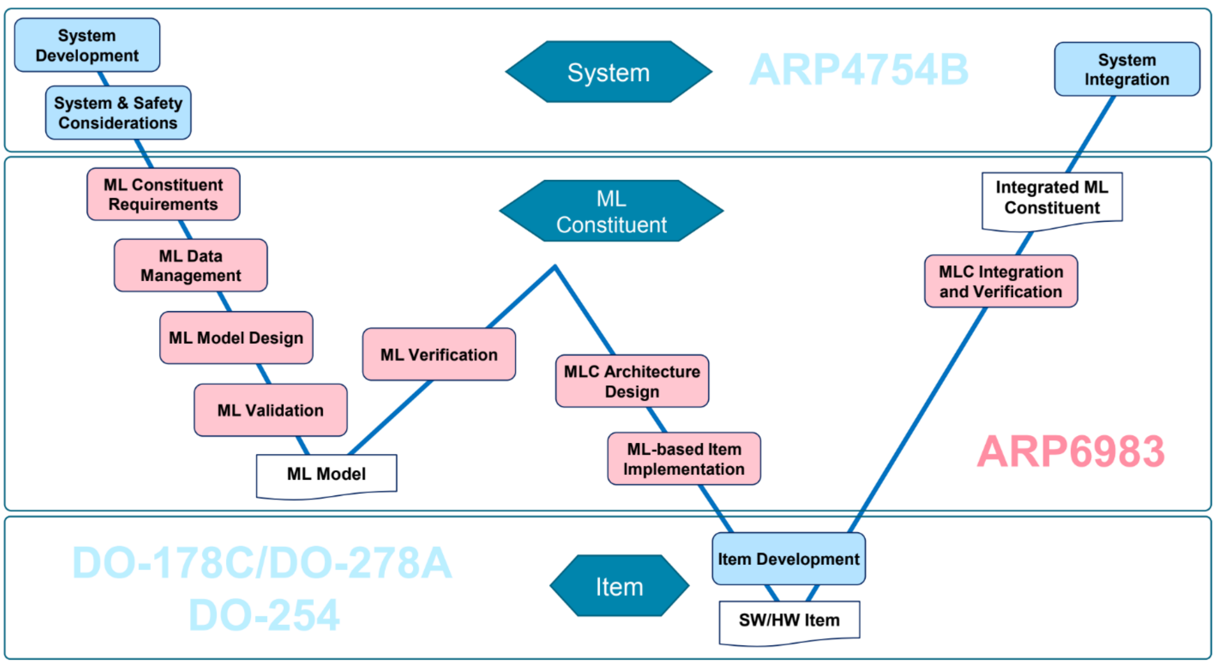

ARP 6983: ML Development Lifecycle (W-Shape). © Airbus

The document ARP6983 describes a Machine Learning development lifecycle, with a "W" shape, enabling to define when and who must take into account all the desired qualities of trustworthiness. While not a complete novelty in itself, such a description consolidates the feasibility and the concrete aspects of Trustworthiness in a manner that is understandable for organisational management, improving the adoption of the new best practices.

In aviation, usually each component within a system has bounded inputs to bound the mathematical space of output. This behaviour provides solid elements to guarantee the predictability of the component and its certifiability. AI often comes with an internal stochastic behaviour, and is furthermore ingesting wide ranges of inputs (e.g. recent research on context-awareness, or even just language), reducing this guarantee as the input space cannot be exhaustively explored during verification. The participants of the workshop highlighted a lack of machine learning specific verification & validation tools (V&V). Non-deterministic AI components are likely to be embedded into safety layers in order to reduce the out-of-distribution effects down to an acceptable level.

Towards more certifications

The participants underlined the tendency of human control to decrease over time, i.e. there is no permanent guarantee to maintain a full concrete understanding of the situation and possibilities for human intervention (design and use). This is not a technical aspect and must be addressed through organisational measures, as part of an AI governance requirement.

The multiplication of AI applications and their general openness results in the emergence of new types of cyber-attacks, dedicated to take benefit from the weaknesses of artificial intelligence software. Common attack goals include data leakage, mispredictions and model poisoning, resulting in unexpected (and thus, unsafe) model behaviour. This field is likely to receive much attention and effort, likely with a significant impact in the qualification and release steps of future AI-based components. The question remains open as to whether a military application should rely on a foundational model, without the ability to access its training process and data.

Participants already familiar with ALTAI mentioned their expectations to enhance this assessment, in order to include more metrics. Doing so, AI quality would not only be qualitative as a declaration, but also measurable through KPIs. In a similar fashion, practitioners would appreciate more tools to translate the rules and regulations into development best practices.

Pillar C: Ethical and societal concerns

The third pillar opens to ethics and general impacts of artificial intelligence in aviation on the whole of society. This approach includes the current status and effort on this topic, but also triggers insightful reflections, unsolved questions and future challenges and opportunities to use AI for good.

Ethics everywhere

As the first overarching ethical framework, ALTAI is already identified as a key reference by the practitioners. Participants underlined the importance of expliciting the ethical aspects of their work. A considerable amount of discussion is rooted in the general data protection regulation (GDPR, adopted by the EU), which has generally increased the level of awareness about data ownership and usage limits.

Similarly to the network of data protection points of contact, activated for the GDPR, it is expected that a comparable network dedicated to artificial intelligence compliance will be set up in the next months. In parallel, the understanding of what is to be included in the ethics concept, and what cannot be included, will be clarified.

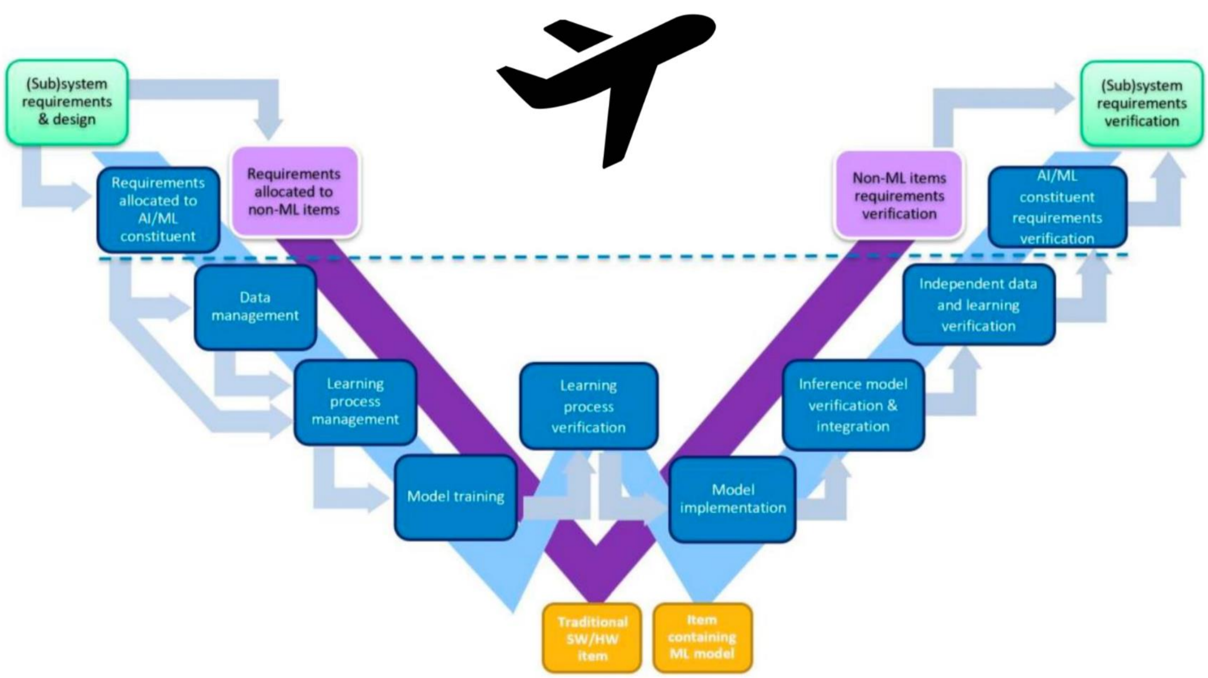

EASA Artificial Intelligence Concept Paper Issue 2: Guidance for Level 1 & 2 machine-learning applications). © EASA

Unsolved questions

Practitioners insisted on the data access problem, here branded as the legitimacy of data ownership. Dilemmas were proposed: should an AI copilot be loyal to its human pilot, or to its maker, or to the police? Is it correct to use air traffic control recordings which include the voices of individuals (namely, the pilots and the controllers)? Who owns the data for predictive maintenance? Who should own the training results of machine learning algorithms that run locally in an embedded system? Instead of one general answer, the solution seems to stem from the practical description of each precise application.

The cultural differences between the countries have an impact on the vision of the people on AI, mainly guided by the US-based companies. The release of new AIs in the European market must be fully compliant with the trustworthiness approach in order to avoid market distortions. Also, the ethics mimicked by an AI should be adapted to the country where it will be used to respect their culture: ethics actually are not universal, and the behaviour of an AI agent may be shocking across different contexts.

The environments where artificial intelligence is used differ between civilian and military applications. By definition, AI systems in defence are used in engagement versus other entities. As a result, each party will try to deceive the systems used by their opponents and to act unexpectedly. An open question consists in qualifying the operational robustness in face of these unexpected threats.

Examples of unexpected threats in a military application of computer vision. © AI4DEF

This problem also covers the regularity and boundedness of the domain, observed in the civilian applications: practitioners are accustomed to measure the generalisation abilities of their AI. The military domain already plans to surprise such components, and also to repurpose civilian-designed components for other applications.

Insights for the future

Artificial intelligence has, in the past, already suffered two “winters of AI” due to the hardware limitations of its time. A participant suggested that the next winter of AI would not be technological but societal: society may change to be able to accept AI, but the ability for people to get used to living with automatons is still unpredictable. In the meantime, people tend to more acutely refuse to have their data used to train new artificial intelligence software.

Social acceptance is already well integrated into ethical AI frameworks such as ALTAI, which is however not fully adopted by the community of practice as of today. Among the hints to cover this challenge was a suggestion to create an AI to encode ethics and morals as rules, enabling a clearer diffusion and uniformisation on this topic (with its own challenges). Last but not least, suggestions about the criticality of AI decisions proposed to modify the weights applied to short-term versus long-term AI decisions: should long-term influence become high criticality?

Acknowledgements

This workshop and report have been supported by the HumaneAI project under grant number 952026.

Special thanks to Paul Lukowicz, coordinator of the HumaneAI project, for its support and welcoming words to the participants.

Congratulations to the participants, speakers and contributors, for having exposed, shared and discussed their views and expertise about the future of artificial intelligence in the aviation industry.