UK Civil Aviation Authority

The CAA's Response to Emerging AI-Enabled Automation

Part A: Strategy for Regulating AI and Advanced Automation in Aerospace

CAP3064A

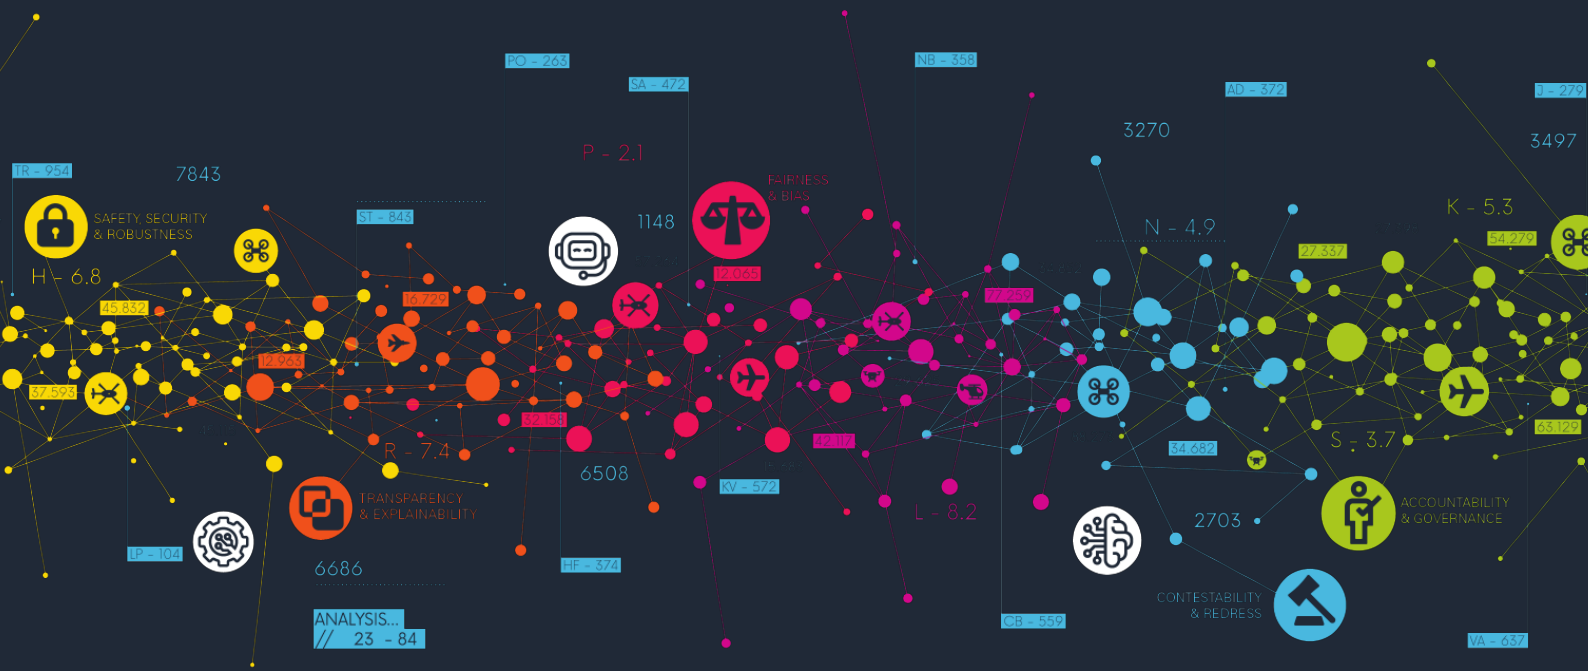

PROTECTING PEOPLE, ENABLING AEROSPACE

CAA Ai Innovation Strategy

v1.3.2 Nov 2024

Published by the Civil Aviation Authority 2024

Civil Aviation Authority
Aviation House
Beehive Ring Road
Crawley
West Sussex
RH6 0YR

First published November 2024

Enquiries regarding the content of this publication should be addressed to: strategyforai@caa.co.uk

The latest version of this document is available in electronic format at: www.caa.co.uk/AI

Executive Summary

The Civil Aviation Authority (CAA) has developed a comprehensive strategy to address the challenges and opportunities presented by Artificial Intelligence (AI) in the aerospace sector. This strategy aims to enable safe innovation while maintaining public confidence and aligning with global standards.

We recognise AI as a powerful new technology that can automate complex tasks in ways never before possible. This brings two main challenges: first, we need to ensure the AI software itself is safe and reliable - for example, if it's helping to control an aircraft, we must be absolutely certain it will make the right decisions every time. Second, we need to carefully consider how roles change when AI takes over tasks - for instance, if pilots move from actively flying to supervising AI systems, we need to ensure this new way of working remains safe. To address these challenges, we will oversee AI across all areas of the CAA's remit by building on our established regulatory frameworks - using proven rules where they work well and carefully updating our requirements where needed. We'll take a measured approach, learning from focused research and real-world projects while developing our expertise in regulating both AI systems and increasing levels of automation.

Given the extensive nature of the challenge that is expected to span the next decade, we need an approach that will create the foundations for long-term change, while also being reactive to short term demands. The initial activity over the next 3 years is therefore enabled by 3 pillars. Pillar 1 provides valuable insights and evidence through horizon scanning and market engagement. This will inform pillar 2 which defines our strategic directions – the hypotheses and regulatory questions that we must explore and act upon. Pillar 3 will then identify and provide the CAA with the skills and capabilities to deliver this extensive change programme.

Central to this strategy is the establishment of an AI Strategy & Portfolio Hub, which will provide centralised oversight and expertise, manage the AI Portfolio, implement the UK's AI regulatory principles, and collaborate across CAA teams to integrate AI into internal operations. We also emphasise international collaboration, aiming to influence ICAO's position on AI in aviation, collaborate with other National Aviation Authorities, and contribute to industry standards development.

We will integrate this strategy into our established governance structures, ensuring it aligns with other CAA strategies while following the government's approach to regulating AI across UK industries. The CAA has outlined an "AI Flight Plan" with specific goals for each phase: "Pre-flight Checks" (2024-25) to deliver the initial strategy and develop business plans, "Taxi" (2025-26) to establish the AI Strategy & Portfolio Hub and initiate collaborative projects, and "Take-off" (2026-28) to conduct systemic reviews and develop rulemaking mandates.

Through this comprehensive approach, the CAA aims to lead in harnessing AI to enhance aerospace safety, security, consumer protection, airspace, and environmental sustainability while fostering innovation and public confidence in the rapidly evolving field of AI in aerospace.

Chapter 1

The Strategic Challenge

1.1 Modern Artificial intelligence (AI) is being explored across the aerospace industry with the potential to enable transformative opportunities for improved safety, security, efficiency, the consumer experience, and environmental sustainability. Applications span across aircraft, airports, ground infrastructure, airspace, space, consumer services and many more. For aircraft, AI could enhance flight control systems, provide training simulations, assist pilots, and reduce workload. For airports and infrastructure, AI shows potential for streamlining passenger processing, baggage handling, security screening, and ground operations. And in managing airspace, it could optimise air traffic management, weather prediction, flight routing, and overall airspace utilisation through advanced analysis of real-time data. Our publication CAP3019 “AI Technology Outlook” explores the potential future of AI in aviation in more detail and will continue to be reviewed and updated as we learn more from the market and research.

1.2 High levels of aircraft autonomy currently remain a concept, but an exciting path lays ahead with AI and automation expected to evolve rapidly across all parts of the aerospace ecosystem. However, we also don’t know precisely how this technology will evolve. Regulatory priorities will need to evolve in lockstep with the maturation of these AI applications to ensure their safe, responsible, and scalable integration into UK aerospace over the coming decade.

1.3 The CAA has a broad scope of responsibilities for safety, security, consumer protection, airspace, and the environmental sustainability of aviation. Looking at this AI-enabled future through this lens, we see that AI is the technology that will unlock complex automation which has not been possible before. It will deliver the potential for high levels of autonomy, and in turn present the aerospace sector and the CAA with new hazards and risks to manage while enabling innovation.

1.4 This strategy aims to outline the opportunities and challenges that we understand today, and provide foundations to establish a proactive, engaging, and agile approach that will unlock the benefits of artificial intelligence and advanced automation in aerospace, while maintaining trust.

The Regulatory Context

1.5 Recent years have seen an increasing awareness of this next phase in AI technology and the resultant challenges. In 2018 the Organisation for Economic Co-operation & Development (OECD) published its principles for AI, providing an approach that is focused on the outcomes of AI systems. Shortly after, the European Union Aviation Safety Agency (EASA) shared its first interpretation of

an AI roadmap. Subsequently, we have seen a proliferation of technological, governmental, societal, and regulatory strategies globally.

1.6 The UK Department for Science, Innovation and Technology (DSIT) confirmed on 6th February 2024[1] its pro-innovation approach to regulating AI. This centres on the adoption of five principles that will form the cornerstone of AI regulation in all sectors including aerospace. Due to the continued growth in investment in AI research and development globally, there is a widely recognised belief that "frontier AI" – the kind that could enable a much broader and deeper scope of autonomy than we see today – may be possible soon. With the pro-innovation principles establishing a performance-based approach, AI risks will continue to be assessed and monitored closely by UK Government and additional responsibilities on developers and suppliers are likely to be introduced in future where necessary to mitigate specific risks.

1.7 The international nature of aerospace means that products and systems are developed to international standards and framed by laws across the globe. In 2024 the European Union (EU) adopted its AI Act. This has a scope much wider than aerospace, but in the same way that the General Data Protection Regulation (GDPR) influences legislation in other jurisdictions, there is some expectation that the EU AI Act[2], which is expected to apply in mid-2026, will have a knock-on impact on the development of AI-based systems here in the United Kingdom and around the world. The effect on international technical and performance standards, such as those developed by Eurocae[3], ASTM International[4], and the British Standards Institute[5], may indirectly bring elements of the EU's legislation into non-EU products. It is currently unclear, but we will monitor this carefully.

1.8 While there are currently limited standards and regulations that are appropriate and effective for the next era in AI, automation, and autonomy, standards development organisations both in the UK and globally are working on new and amended standards today. This activity is only likely to increase over the coming years, and there is therefore a growing need for the UK CAA and UK

[1] "A pro-innovation approach to AI regulation: government response", Department for Science Innovation & Technology; available at https://www.gov.uk/government/consultations/ai-regulation-a-pro-innovation-approach-policy-proposals/outcome/a-pro-innovation-approach-to-ai-regulation-government-response

[2] Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)

[3] Eurocae, eurocae.net

[4] ASTM International, astm.org

[5] British Standards Institute (BSI), bsigroup.com

Government to establish a position to be able to influence this international debate.

1.9 Across the portfolio and responsibilities of the International Civil Aviation Organisation (ICAO), AI is considered to greater or lesser degrees. The ICAO's strategy for AI is largely focused on enabling collaboration both within aerospace, and with other global sectors through other UN bodies, academia, and standards development organisations. The UK CAA continues to have a role in defining and influencing the ICAO approaches to AI. More generally, we have applied ICAO's "Critical Elements" (CEs)[6] of both state safety and security oversight systems within the development of this strategy to start to identify areas where the CAA's AI Capability needs to be developed. This is detailed further in section 0.

1.10 It is quite clear that we need a suitably agile and robust regulatory strategy, supported by a competent AI capability across the CAA, to be able to respond to the continually and rapidly evolving context of AI.

Our Vision for Regulating AI and Advanced Automation

1.11 We will enable AI to enhance aerospace efficiency, sustainability and scalability, while ensuring safety, security, consumer protection, and environmental sustainability through proportionate governance.

1.12 This vision will guide the CAA by describing a future state where we have intentions to:

- Build an agile and robust regulatory approach centred on the 5 AI Principles that proactively enables innovation while fortifying public confidence in AI-powered aerospace systems and ensuring harmonisation with global standards.
- Provide a regulatory path for AI to automate complex tasks and unlock unprecedented levels of autonomy in a safe and secure manner, through robust governance mechanisms anchored in ethics and human oversight.
- Establish world-class AI regulatory capabilities to scrutinise these systems through their entire lifecycle, setting high standards for reliability, cyber-resilience, and intelligent human-machine teaming.

1.13 "The CAA's Response to Emerging AI-enabled Automation" (CAP3064) describes the foundational AI Framework that underpins this strategy for

[6] Universal Safety Oversight Audit Programme (USOAP) focuses on a State's capability in providing safety oversight by assessing whether the State has implemented the critical elements (CEs) of a safety oversight system effectively and consistently, https://www.icao.int/safety/CMAForum/Pages/default.aspx

regulating AI and advanced automation. This includes descriptions of the common elements listed below which are referred to within this document.

- The purpose and structure of our approach.
- The AI framework: Common language, building trust and technology outlook.
- The CAA's AI Strategy & Portfolio Hub.
- The AI Portfolio.
- The UK and international context.

1.14 It is therefore recommended to read this strategy in conjunction with CAP3064. This is available at www.caa.co.uk/CAP3064.

Chapter 2

The Regulatory Challenges

2.1 Software is introduced into a system to carry out specific tasks or functions. Artificial Intelligence enables the development of a new generation of software-based systems which are fundamentally different and significantly more powerful compared to 'conventional' software systems. Conventional software can address relatively simple and well-defined computational tasks, which can be accomplished by following a well-defined sequence of steps. This sequence of simple straightforward steps is known as an *algorithm*. Therefore, conventional software is well-suited for tasks with clear rules and procedures. In contrast, AI, and particularly Machine Learning (ML), excels in tasks where the rules are difficult or impossible to define explicitly, such as recognising patterns in images or understanding natural language. On the other hand, AI offers a powerful class of methods for more complex tasks that typically require elements of human-like intelligence, like learning and extrapolating from data to adapt to new situations.

2.2 As a result, AI presents 2 main challenges for regulation:

Assuring the robustness of AI-based software

AI can perform more complex tasks than traditional software through advanced algorithms and machine learning. However, this sophistication makes it harder to verify its decision-making processes and prove it will operate safely and securely across all scenarios. Additionally, AI systems can be designed to learn from operational data and adapt their behaviour and internal design to optimise performance. While this adaptability offers benefits, we must ensure robust safety controls remain effective even as the system evolves through learning.

People & AI working together

For systems that rely on human skills and competences to ensure safety and security, introducing AI may affect the human's role. How can we ensure that people and AI work effectively together, and ensure that there is adequate trust in the AI's capabilities?

2.3 AI systems can also be more difficult and costly to develop, verify, maintain, and operate. Deciding to use an AI technology like machine learning to implement a function of a system is a conscious choice made by the developers, weighing up the benefits and costs of doing so. Most likely this will be because the function is particularly complex and so demands a complex technological solution to achieve automation.

Machine Learning: the driving force of the AI boom

2.4 Machine Learning (ML) has become the prevalent approach to developing AI systems. A ML system learns to perform a task through data and without being explicitly programmed *how* to do so.

2.5 Advances in ML have been significantly accelerated in recent years due to three main driving factors: big data (unprecedented volumes and variety of data), cloud computing (unprecedented access to high-power computation) and algorithms (novel methods to use the data and computing to address computational tasks of unprecedented complexity). This has led to increasing adoption of the technology across all sectors of the economy, which comes with excitement about the potential benefits, coupled with concerns about its novel challenges.

The characteristics of Machine Learning

2.6 The fundamental component of these systems is the *ML model*: a mathematical function that aims to create connections between system inputs and corresponding outputs to meet a prescribed objective. This function has a fixed structure with adjustable (or *learnable)* parameters. Training data is provided to the model and patterns in the data are identified. Through the process of *training*, the model adjusts its parameters to match the patterns it learned from its training data.

2.7 At a high level, the process of ML can be divided into two stages, *Training* and *Deployment*. In training, the untrained model learns from training data. In deployment, the trained model is used to perform its trained task, for example to control a robot in operation. There are many different training schemes, but it is important to highlight the two main modes of ML system deployment:

- Frozen Mode: The AI only uses what it learned during its *training* and remains constant in its design during *deployment*. The AI can be re-trained with new data which has been captured from its environment, but this re-training is completed offline while it is not in operation. This can be considered a *deterministic* system.
- Online Learning Mode: The model continues learning from the new data during the deployment stage, and re-training is completed while is operation. This can be considered a *non-deterministic* system.

2.8 This is a fundamental difference, and of vital importance in the regulation of AI. Frozen ML systems are powerful, complex, flexible, and yet, under certain conditions and assumptions, predictable. This is a particularly important characteristic for any system intended for safety critical applications, as it is much more likely to be possible to establish the safety of a predictable system.

2.9 Through our engagements with industry, international regulators, and standard bodies, the emphasis for AI in aerospace currently lies in deterministic ML systems. It is our view that this is reasonable: the emphasis in ML is justified by the technology's success in other industries and across the economy, while the preference towards deterministic ML systems is in alignment with the industry's imperative prioritisation of safety and security.

Understanding the language of ML

2.10 There are several key concepts in ML that are important to clarify in this strategy:

- Supervised Learning: In this type of ML, the model is trained on a labelled dataset. A labelled dataset is one where the target outcome or the correct answer is provided. This method is used for tasks such as:
  - Classification: Predicting discrete responses, for example, whether an email is spam or not spam.
  - Regression: Predicting continuous responses, for example, predicting house prices based on various features.
- Unsupervised Learning: This type of ML involves training the model using an unlabelled dataset. The model tries to find patterns and relationships in the data. It's used for tasks such as:
  - Clustering: Grouping the data into clusters based on similarity. For example, customer segmentation in marketing.
  - Dimensionality Reduction: Reducing the number of variables under consideration, by obtaining a set of principal variables that sufficiently capture the variance of a dataset.
- Reinforcement Learning: In this type of ML, the AI learns to behave in an environment by performing certain actions and observing the results it gets from those actions. This is used in applications where a system needs to perform decision-making functions within an operating environment, such as robotics, gaming, and navigation.
- Generative AI: In this type of ML the aim is not just to learn from data but also to generate new data that has the same statistical properties as the training data. These models have been used in a variety of applications, including to produce images, human-like text responses, music, molecules for drug discovery, and even art.

2.11 Our guidance on AI language, CAP2966, will be updated to reflect the terms used in this strategy.

AI's Impact in Aerospace

2.12 The application of AI in the aerospace sector has an opportunity to bring substantial benefits across safety, security, consumer experience and protection, and environmental sustainability. Our horizon scanning report, CAP3019, discusses some of these, but we need to understand what applications the CAA should be concerned with, as not all will require regulatory consideration.

2.13 To aid in this, we have assessed the CAA's wide range of regulatory roles and responsibilities. We also use ICAO's Critical Elements as a framework to identify areas of the CAA where the regulation of AI will be affected.

The CAA's Readiness for regulating AI

2.14 In regulating the aerospace sector, the CAA continually strives to ensure that the eight critical elements of effective safety and security oversight systems are in place, as described by ICAO. The impact of AI in aerospace is likely to affect many of these critical elements. We can use them to assess our readiness for regulating AI.

2.15 An initial assessment of the potential impact is outlined below, and the AI Portfolio will be informed by a more thorough assessment and ongoing review against these 8 critical elements.

ICAO Critical Element 1: Primary aviation legislation / Aviation security legislation

2.16 Description: Establishing comprehensive laws that cover both safety and security aspects of civil aviation.

2.17 Discussion: While AI itself may not require amendments to primary legislation, it is very likely that the resultant high degrees of autonomy might. In 2023, the CAA and DfT appointed the Law Commission of England and Wales to conduct a thorough review of aerospace legislation with regards to autonomy in future of flight aviation operations.

ICAO Critical Element 2: Specific operating regulations / Aviation Security programmes and regulations

2.18 Description: Enacting regulations that address safety and security requirements for aviation activities.

2.19 Discussion: It is very likely that AI will have impacts across all regulatory remits of the CAA. An extensive review of policy, regulations, and guidance is required. This is proposed to be part of the AI Portfolio.

ICAO Critical Element 3: State safety functions / State appropriate Authority for aviation security and its responsibilities

2.20 Description: Establishing competent authorities responsible for both safety and security oversight.

2.21 Discussion: To deliver the CAA's responsibilities it has structured its organisation to meet these obligations. The need for regulation and oversight of AI is likely to require some organisational design changes to accommodate the emerging responsibilities. Section 3.3 of the Strategy for AI: Enabling Responsible AI Adoption (CAP3064) outlines our proposed operating model to support these emerging responsibilities. Over and above the usual state safety and security functions, the CAA's innovation system also plays a critical role in enabling AI and autonomy. Our functions for horizon scanning, innovation strategy, sandboxing, test & evaluation, and future policy development place us in a good position for responding to emerging innovations such as AI.

ICAO Critical Element 4: Qualified technical personnel / Personnel qualifications and training

2.22 Description: Ensuring personnel involved in oversight are adequately qualified and trained covering both safety and security aspects.

2.23 Discussion: To be a competent regulator the CAA needs to ensure that we can adequately challenge and support the sector to achieve the upmost performance in safety, security, airspace, consumer protection and environmental sustainability. AI skills and expertise will undoubtedly be required within the CAA as its usage grows across the sector. Within the AI Portfolio, a future workforce skills plan will outline the needs for education, training, recruitment, and retention of skills.

ICAO Critical Element 5: Technical guidance, tools and provision of safety-critical information / Provision of technical guidance, tools and security critical information

2.24 Description: Providing guidance materials to both the industry and inspectorate while providing the tools that cover safety and security oversight activities.

2.25 Discussion: The CAA holds an extensive library of policy, regulations, guidance, tools, and safety-critical information. As AI introduces new challenges both from AI-enabled software and in the resultant system autonomy, we may need new tools and guidance to support the sector and CAA colleagues to effectively regulate and oversee activity. The CAA's Strategy & Portfolio Hub is designed to support the whole organisation to review its current guidance, tools, and safety information, and lead the AI Portfolio to deliver the necessary updates.

ICAO Critical Element 6: Licencing, certification, authorisation, and approval obligations / Certification and approval obligations

2.26 Description: Implementing processes for licensing, certifying, and approving entities, personnel, and equipment regarding both safety and security.

2.27 Discussion: The extent of the impact from AI on these activities is yet to be fully assessed, but it is expected to involve a high degree of cooperation with industry standards and other authorities. The 5 AI Principles for building trust (CAP2970) will be important here in providing consistency to our regulatory approaches. The AI Portfolio includes activity to map existing regulations and standards, allowing us to prioritise and direct activities.

ICAO Critical Element 7: Surveillance obligations / Quality control obligations

2.28 Description: Conducting regular monitoring and inspection activities to ensure compliance with both safety and security regulations.

2.29 Discussion: The challenge presented by learning AI systems, as described in a previous section, may have some interesting effects in the way we define and carry out our monitoring and oversight obligations. The ongoing assurance of learning software is yet to be determined, but research across all sectors is working to tackle this. Industry standards will continue to play a strong part here, but research and collaboration is also fundamental. Note that the CAA's capability in this regard may in fact be improved in future by the adoption of AI technologies within the CAA itself – this is discussed in Part B of this strategy.

ICAO Critical Element 8: Resolution of safety issues / Resolution of security concerns

2.30 Description: Establishing mechanisms for addressing safety and security issues, incidents, and accidents promptly and effectively.

2.31 Discussion: The 5 AI Principles for building trust (CAP2970) include a need to consider contestability and redress. In the aerospace context, there is already an extensive system in place for reporting safety occurrences, investigating incidents, and resolving safety issues. The extent of impact from AI and autonomy is yet to be assessed but is likely to be high. Note that the CAA's capability in this regard may in fact be improved in future by the adoption of AI technologies within the CAA itself – this is discussed in Part B of this strategy.

Keeping a lookout for new risks

2.32 The unique nature of AI and its ability to enable automation of complex processes in aerospace may introduce new risks that we have not seen previously. It may bring about risks in areas that are not currently regulated. Therefore, it is important that the CAA monitors the use of this technology

closely, working collaboratively with the sector and other authorities ideally to forecast these risks before they cause harm.

2.33 We must also be willing and able to adapt our approach based on this intelligence if a change is required. The CAA's Strategy for AI will continue to evolve to reflect the latest insights from aerospace and other sectors.

Autonomy: The Regulatory Challenge of Automation

Increasing Autonomy

2.34 An activity is typically formed of several discrete or related tasks, each in some logical order, necessary to achieve an outcome. This can be applied to any activity or system, small or large. For example, making a hot drink for a passenger might involve filling a kettle, turning it on, adding tea or coffee to a cup, pouring on the hot water, and adding milk. Alternatively, the activity might be flying passengers from an origin airport to their destination, with tasks including the aircraft walk-around, starting engines, taxi, take-off, cruise, landing, taxi again, and shut down.

2.35 The first is focused on a more granular example and the second has a much wider scope, but in both cases, it is possible to see how an outcome is generally achieved by completing several tasks. If all these tasks required human operation there would be no automation in the system. As each of the tasks are automated by introducing some form of technology, the need for human control, intervention, and oversight changes. As a result, the system's *capability for autonomy* increases.

2.36 In both examples, while the system could be fully automated it can only operate with autonomy (without the need for human intervention or oversight) if the human can trust that the system will operate effectively. Only once there is trust, will a human be willing to disengage from the system and trust that it will complete the tasks effectively. This emphasises the critical importance of *creating and maintaining trust* which we believe is achieved through rigorous development, testing, verification and validation, transparency, and communication.

The Law Commission of England & Wales Review: Aviation Autonomy

2.37 The Law Commission has been asked by the CAA and Department for Transport to review the law around autonomous flight, to support the safe development of rapidly advancing technology. Information is available on the Law Commission's website at lawcom.gov.uk/project/aviation-autonomy/

2.38 The 3-year review is partly funded by UK Research & Innovation (UKRI) through the Future Flight Challenge[7], delivered by Innovate UK and the Economic and Social Research Council. It examines the existing legal framework to identify the challenges and opportunities linked to the introduction of highly automated systems into the aviation sector.

2.39 The focus relates to electric vertical take-off and landing (eVTOL) aircraft and remotely piloted aircraft systems (RPAS), alongside the novel air traffic management (ATM) technologies that are anticipated to support these future aviation operations.

2.40 The Law Commission's review includes 2 consultations, culminating in a series of recommendations to the DfT Secretary of State for law reforms that will ensure the UK is ready to take advantage of oncoming advances in automation.

The importance of Trust

2.41 In aerospace we have seen tasks become automated using software and hardware for decades. In an airliner of the 1960s era the fuel on the aircraft would have been closely monitored and controlled throughout the flight by a qualified and experienced engineer. In a modern aircraft the fuel management system takes care of that role almost entirely. The pilot trusts that the fuel management system will operate effectively, perhaps only monitoring the fuel burn. So, in aerospace, there is already experience of creating and maintaining this trust. But there is a risk of automation bias, and we have also seen unfortunate examples where trust in automated functions has been eroded following an incident that has resulted in the loss of life.

2.42 Communicating trust merits some further emphasis, as it does not only depend on trust from the human operator of a system but from everyone directly or indirectly involved in and affected by it. Designers and developers, users, broader society, operators of co-existing systems such as pilots of aircraft sharing the airspace with an autonomous aircraft. All will require different levels of assurances with regards to the system's trustworthiness.

2.43 Moreover, we should draw clear lines between the desired dimensions of trust (safety, security, robustness, resilience) and methods to establish trust (verification, validation, testing, transparency, explainability). These methods are not yet mature for most applications of AI in aerospace, and so there is currently a gap that should be filled before trust can be built effectively.

[7] https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/future-flight/

Human Factors

2.44 As functions and tasks within an aerospace system move into software, enabled by AI technology, human factors will become a cornerstone for the regulation of AI and the resultant autonomy of the systems it is used in.

2.45 An aerospace system typically combines three elements in partnership to achieve a task: the software, the hardware, and the human. The role of the human involves interpreting information, making decisions, and providing oversight and control to ensure the safe and efficient operation of the system, while considering factors such as situational awareness, workload management, and cognitive limitations. The environment also sets an important context for the task, for example congested airspace around a major airport presents greater challenge than quiet airspace away from regular traffic.

2.46 The human factors of a system are therefore critical to consider when choosing to use AI, as it will likely have a substantial impact on how the human operator carries out its role in relation to the rest of the system. Our approach to regulating autonomy will shift the current focus more towards the role of the human and the way they interact with a system.

2.47 The CAA's Human Factors strategy[8] provides an important framework against which the impacts of AI should be assessed and managed. Research and intelligence gathering in this area will form a substantial part of our AI Portfolio specifically with regards to the high degrees of autonomy we expect to experience from the introduction of AI.

Liability Implications

2.48 As AI enables higher levels of autonomy in aerospace systems, the traditional notions of liability and accountability will inevitably be challenged. Historically, liability has primarily rested on human operators, such as pilots or air traffic controllers, whose actions or inactions directly impacted the safe operation of aircraft. However, with the increasing autonomy facilitated by AI, the allocation of responsibility and liability in these areas becomes more complex.

2.49 The CAA recognises the need to proactively address these evolving liability considerations. As autonomous systems take on a greater role in decision-making and execution of critical tasks, it becomes necessary to re-evaluate the boundaries between human and system responsibilities. This exercise will inform the development of appropriate legal and regulatory frameworks to delineate liability in circumstances where an autonomous system's actions or decisions contribute to an adverse event.

[8] https://www.caa.co.uk/safety-initiatives-and-resources/working-with-industry/human-factors/strategy-and-action-plan/

2.50 In collaboration with the CAA, the legal experts at the Law Commission of England and Wales are reviewing existing liability models relating to future of flight modes including drones, eVTOLs, and novel air traffic management. This work is exploring potential regulatory and legislative revisions to accommodate the unique challenges posed by AI-enabled autonomy. Alongside similar reviews across other aspects of the aerospace sector, this effort will aim to strike a balance between fostering innovation and ensuring fair and equitable allocation of accountability.

2.51 Factors to be considered include:

i) Explainability and transparency of AI systems, enabling comprehensive investigations and root cause analyses in the event of incidents or accidents.

ii) The degree of human oversight and intervention capabilities within autonomous systems.

iii) The role of developers, manufacturers, and operators in ensuring the safe and responsible deployment of autonomous systems.

iv) The impact of environmental factors, such as degraded sensor data or unforeseen operating conditions, on autonomous system performance.

v) Protecting against prohibited commercial practices and anticompetitive behaviour.

2.52 By proactively addressing liability implications, the CAA aims to create a regulatory environment that provides clarity and certainty for all stakeholders, including manufacturers, operators, and the broader aviation ecosystem. This approach will foster trust and confidence in the adoption of AI-enabled autonomous systems while upholding the highest standards of safety and accountability.

Chapter 3

Our Initial Regulatory Approach

The Strategic Imperatives

3.1 Our strategy for regulating AI is bound by several strategic imperatives that will permeate through the entire AI Portfolio:

- We will regulate AI and autonomy across each area of our regulatory responsibilities[9] and in accordance with our regulatory principles[10].
- We will use existing regulation wherever possible, reducing the need for lengthy legislative processes, instead amending, adjusting or otherwise reusing policy, guidance, and acceptable means of compliance to achieve the desired outcomes.
- We will be steered by the priorities of central government, CAA strategies, and the market, while ensuring consumer and public interests remain central to our decisions.
- We will take an incremental approach, initially learning from specific research and projects targeted at areas of interest, prioritised in part by market intelligence.
- In turn this will help to inform more systemic regulatory reviews across all our responsibilities, enabled by dedicated resources, to identify areas of the regulatory frameworks that need adaptation for AI.
- We will build and grow the CAA's capability and maturity for regulating AI and autonomy, from establishing a foundational AI awareness for all colleagues, through to world-leading technical regulatory skills and expertise.

Our Initial Priorities

3.2 We plan to deliver this strategy for regulating AI in aerospace through the AI Portfolio across three pillars between 2024 and 2028 that aim to increase the CAA's maturity for AI regulation:

[9] The CAA's role, https://www.caa.co.uk/Our-work/About-us/Our-role/

[10] The CAA's regulatory approach, https://www.caa.co.uk/our-work/about-us/our-regulatory-approach/our-regulatory-approach/

Pillar 1: Horizon Scanning & Market Insights.

3.3 We will work with the sector to identify emerging technologies and models, making initial assessments of the specific regulatory challenges and informing the regulatory change programmes of the AI Portfolio.

Pillar 2: Defining the Strategic Directions.

3.4 We will establish strategic directions within the strategy that inform and prioritise our activities based on the intelligence and insights gathered in Pillar 1. The AI Portfolio will be continuously updated with these strategic directions, as well as influencing the development of our skills and capabilities for regulating AI and autonomy.

Pillar 3: Building CAA Capability.

3.5 Over time, our established AI regulation capability will enable us to conduct comprehensive systemic reviews of regulations across all CAA's regulatory responsibilities. These reviews will align with the technology, market needs, and consumer trends of the day. This proactive approach will inform necessary legislative and regulatory reforms to ensure that the UK's aerospace regulatory system remains fit for purpose as AI and autonomous systems continue to evolve.

3.6 Throughout the phases outlined above, we anticipate fostering and supporting collaborative projects with innovators across the sector. These collaborative endeavours will enable us to identify potential regulatory challenges proactively and test practical regulatory approaches in a timely manner.

Our Long-Term Plan

3.7 Regarding AI in software, the CAA is unlikely to certify specific AI algorithms or models directly. Instead, our regulatory role will extend to the software systems that contain AI and their intended functions – aligning with established practices for software certification and assurance.

3.8 With the increased autonomy in complex aviation tasks enabled by AI, the CAA will need to undertake a comprehensive review of existing regulations and policies to ensure their continued relevance and effectiveness.

3.9 This review will encompass several crucial areas:

Research & Development

3.10 The CAA will support and conduct research initiatives advancing our understanding of AI's implications across aerospace. This research will inform our regulatory approach, ensuring that it is grounded in the latest scientific and

technological advancements while upholding our regulatory principles and responsibilities. Areas of research are likely to include:

Human Factors

3.11 The introduction of autonomous systems necessitates a thorough evaluation of human-machine interactions and the evolving roles and responsibilities of human operators. The CAA will assess the human factors implications of increased autonomy, addressing issues such as skill degradation, situational awareness, and trust in automated systems.

Legal Implications

3.12 In collaboration with the Law Commission, the CAA will review existing legal frameworks to identify areas requiring clarification or amendment to accommodate the challenges and opportunities presented by autonomous systems. This effort will ensure that the regulatory and legal landscapes remain aligned and supportive of responsible innovation.

Consumer Trust

3.13 The adoption of AI and autonomous systems in aviation has the potential to significantly impact consumer confidence and trust in air travel and in commercial practices within the sector. The CAA will conduct comprehensive reviews to understand how increased automation and AI-enabled decision-making affect consumer perceptions and behaviours. We will develop policies and guidelines that ensure transparency in AI applications, particularly in consumer-facing areas such as booking systems, customer service, and in-flight experiences. This will draw on insights from our Consumer Strategy and the expertise of our Consumer Panel. We will also draw on expertise from and collaborate with the Consumer and Markets Authority (CMA) through their own AI strategy. These efforts will aim to strike a balance between promoting innovation and safeguarding consumer interests, ensuring that the benefits of AI in aviation are realised without compromising public confidence in the safety and reliability of air travel.

3.14 Other areas of focus may include System Theoretic Process Analysis (STPA) for autonomous systems, explainability in ML systems, and the exploration of digital flight rules tailored to autonomous operations. This research will inform evidence-based policy development and regulatory decision-making.

Standards & Regulatory Mapping

3.15 The CAA recognises the array of applicable AI development and assurance standards and guidelines being developed. To navigate this landscape, we will develop a comprehensive map of standards and regulations against the

regulatory areas across the CAA, enabling us to monitor, influence, and selectively implement or adopt those most pertinent to our responsibilities.

Collaboration with Partners

3.16 At the forefront of defining AI development and assurance standards and best practices, the work of standards development organisations is crucial. To ensure the safe and responsible development of AI-enabled aviation software, the CAA will closely monitor and actively contribute to relevant industry standards. While not developing standards directly, we will actively participate, providing aviation regulatory expertise and insights. We will continue to work closely with our international NAA partners to align our approach with the aim to enhance harmonisation. We will also collaborate with industry and academic partners, where there is an opportunity to drive progress against the strategic directives.

3.17 Through this multifaceted approach – aligning with national AI principles, collaborating with international partners, mapping relevant standards, and targeted research – the CAA will establish a robust regulatory framework for AI-enabled aviation software and systems. This framework will foster innovation while maintaining the highest levels of safety and operational integrity.

The AI Strategy & Portfolio Hub

3.18 To ensure a consistent and comprehensive approach to regulating and utilising AI throughout the CAA, a dedicated AI Strategy & Portfolio Hub will be established as outlined in CAP3064, “The CAA’s Response to Emerging AI-Enabled Automation”.

3.19 Regarding the regulation of AI and autonomy, the AI Strategy & Portfolio Hub will play a crucial role in several areas:

- Providing leadership and strategic guidance on AI and autonomy initiatives, both within the CAA and for the broader aviation sector under the CAA's regulatory purview.
- Promotion of AI in aerospace through initiatives such as sandboxing, working groups, and collaborative projects.
- Overseeing and managing the AI Portfolio, encompassing the various programmes, projects, initiatives, and regulatory activities related to the regulation of AI and autonomous systems.
- Ensuring the effective implementation and adherence to the UK's AI regulatory principles, as outlined by UK government[11].

[11] "Implementing the UK's AI Regulatory Principles, Initial Guidance for Regulators" (February 2024).

- Collaborating closely with teams across the CAA, such as Human Factors, Flight Operations, Cyber Security, Airworthiness, Design and Certification, and the Consumer & Markets Group. This cross-functional integration will be essential to address the multidisciplinary challenges and opportunities presented by AI and autonomy in aviation.

3.20 The AI Strategy & Portfolio Hub's oversight and expertise will also be instrumental in integrating AI into the CAA's internal operations, for both regulatory and business functions. This integration will enable our colleagues to leverage the benefits of these technologies when delivering their regulatory duties, such as application assessments, risk analyses, inspections, and reporting. Our strategy for using AI in the CAA is outlined in CAP3064B.

The AI Portfolio

3.21 The activities and initiatives necessary to deliver against this strategy are far reaching across the organisation and the sector. This needs to be effectively managed and reported on according to our governance proposals.

3.22 To enable this, the AI Portfolio encompasses strategically significant activity from across the CAA that is associated with AI and autonomy. Managing this within the AI Strategy & Portfolio Hub enables the CAA to define specific outcomes, promote AI in aerospace and in the CAA, ensure a consistent approach that aligns to the strategy, efficiently monitor progress against the strategic plans, and provide transparent and coherent reporting to all stakeholders.

International Collaboration

3.23 The CAA recognises the global nature of the aviation industry and the importance of international collaboration in effectively regulating AI and autonomy. As such, we will actively engage with international organisations and partner agencies to ensure a harmonised and coordinated approach.

Influencing ICAO.

3.24 The CAA will contribute to the development of ICAO's position on AI within aviation. By providing our expertise and insights, we aim to shape a comprehensive global framework that promotes safe and responsible AI integration while facilitating innovation.

Collaboration with other NAAs.

3.25 The CAA will collaborate closely with other National Aviation Authorities (NAAs) through various channels, including the NAA Network, international working groups, and bilateral relationships. This collaborative effort will foster the exchange of best practices, shared learnings, and the alignment of regulatory

approaches, ultimately enhancing aviation safety and efficiency on a global scale.

Influencing Industry Standards.

3.26 The CAA will actively monitor, contribute to, and influence the development of industry standards related to AI and autonomy in aviation. By engaging with relevant standards bodies and organisations, we will ensure that our regulatory framework remains consistent with globally recognised best practices and standards.

Thought Leadership.

3.27 The CAA aspires to be a thought leader in the international arena, sharing our knowledge, experiences, and insights on the regulation of AI and autonomy. Through participation in conferences, workshops, and forums, we will contribute to shaping the global discourse and advancing the collective understanding of this rapidly evolving field.

3.28 By fostering international collaboration, the CAA will play an active role in establishing a harmonised and coordinated global approach to regulating AI and autonomy in aviation. This collaborative effort will not only promote aviation safety and security but also facilitate the responsible adoption and integration of these transformative technologies across the international aerospace ecosystem.

Chapter 4

Governance Integration

4.1 The strategy for AI and the frameworks introduced by it need to permeate throughout the CAA to ensure that, wherever AI technology affects the safety, security, consumer protection, airspace, or environmental sustainability of aerospace, we take a consistent and responsible approach to regulating it.

4.2 The CAA's numerous strategies and programmes will be affected by AI to varying degrees. It is not the intention to create a whole new pillar of governance centred on AI that draws these strategies and programmes into it, but instead to integrate AI governance into them. A key role of the AI Strategy & Portfolio Hub will therefore be to manage this governance integration and ensure that the approach and principles of the Strategy for AI are effectively represented. This will undoubtedly be enabled further by the AI Community of Practice.

4.3 Senior responsibility for AI regulation will need to sit with the appropriate members of the Executive Committee, according to the nature of the impact. The Head of the AI Strategy & Portfolio Hub will directly support these members by providing expertise and strategic guidance.

4.4 Additionally, there must also be consideration for the governance of using AI, as described in Part B of this strategy. Where AI is being used by the CAA in regulatory functions, there will likely be a resultant impact on how we regulate AI. The controls, guidelines, and tools that will be developed and promoted through the governance for using AI will therefore impact on regulatory activities as well.

4.5 It is anticipated that central government will continue to develop the governance that enables central guidance for UK sector regulators aligned to their pro-innovation approach to AI regulation, while also creating responsibilities on regulators to inform government on AI risks, implementation of the 5 AI Principles, and overall effectiveness of the approach. Therefore, the CAA will reflect these new responsibilities and requirements within its own structure and governance for AI regulation.

Chapter 5

Our Initial Flight Plan for Regulating AI

5.1 The introduction of the CAA's first Strategy for AI represents the first step on a long but exciting journey towards enabling innovation with AI and the future benefits from high degrees of autonomy across the aerospace sector. The current plan outlines the actions we expect to take to begin this journey.

Pre-flight checks (2025-26)

Developing our capability

5.2 Identify, develop, and deliver the requirements for an AI regulatory capability, and any necessary organisational design and governance, integrating this AI strategy into relevant CAA programmes and strategies.

5.3 Identify and deliver initial and future training requirements and scope the options for delivery.

Regulatory Research

5.4 Ongoing development of the initial strategy for regulating AI and the relevant aspects of the AI Portfolio.

5.5 Research and determine the structure and prioritisation of the strategic directions for regulating AI.

5.6 Iterate on the AI framework to update based on intelligence and insights.

Standards & Regulatory Mapping

5.7 Identify our strategy for standards, maintain a register of relevant standards development groups, conduct an initial analysis, and map of international standards to enable us to coordinate our representation.

5.8 Map key use cases to the ICAO Critical Elements to assess regulatory readiness.

Collaboration & Partnering

5.9 Establish open engagement opportunities for collaborating routinely with industry and academia to steer the AI Portfolio and prioritisation of CAA activity by continuing to identify the market and research priorities in AI and autonomy.

5.10 Continue to support and enable the Law Commission's Aviation Autonomy review, and of the Department for Transport's response.

5.11 Develop our engagement and communications strategy for the multitude of stakeholders not already captured above, including other National Aviation Authorities (NAA), Non-Governmental Organisations (NGO), Standards Development Organisations (SDO), and other UK sector regulators.

Taxi & Take-off (2026-28)

5.12 Following our initial year of activity in the pre-flight phase and depending on the pace of developments and funding available, we hope to move into the taxi and take-off phases. Some examples of the types of activity we aspire to deliver include:

i. Launching regulatory projects in collaboration with the sector to explore regulatory challenges from AI and autonomy based on the government, market, and research priorities.

ii. Sponsoring regulatory and policy research to support emerging areas of risk, particularly in human factors and software safety assurance.

iii. Continued growth of the CAA's capability for regulating AI, working closely with the CAA's future skills strategy and information strategy.

iv. Systemic reviews across the CAA's regulatory responsibilities, guided by our regulatory principles, to identify gaps in the frameworks and to define future regulatory development activities, including research, testing, and assurance.

v. Ongoing assessment of the Law Commission's final recommendations on aviation autonomy and support to the development of any actions as a result.

vi. Identify and develop appropriate mandates for rulemaking in areas identified for change, guided by evidence from across the AI Portfolio through research, sandboxing, and engagement.

Chapter 6

Getting involved, providing feedback and more information

6.1 We are committed to maintaining open channels of communication and collaboration as we implement and refine our AI strategy. We encourage engagement from all stakeholders and offer several ways to stay informed, provide input, and get involved:

CAA Website

6.2 Our website now includes a dedicated AI page at www.caa.co.uk/AI, serving as a central portal for all AI-related information. Here, you'll find the latest guidance, updates on regulatory developments, and links to relevant resources from across the CAA and beyond. We encourage you to visit this page regularly for the most up-to-date information on AI in aviation.

Email Contact

6.3 For enquiries, requests, or to share insights, please contact us at StrategyforAI@caa.co.uk. This central mailbox is monitored regularly, ensuring timely responses to your queries and contributions.

Feedback

6.4 We welcome feedback on all aspects of this strategy. Your input is invaluable in helping us refine our approach and ensure it remains relevant and effective. We aim to be collaborative, accessible, and supportive in enabling innovation in AI in aerospace and in the CAA, and your feedback plays a crucial role in achieving this goal.

A Note on AI Assistance

6.5 In the spirit of embracing the technologies we are discussing we would like to acknowledge that generative AI tools were used in the drafting of this strategy document. This approach has brought several benefits, including improved use of time, effective summarisation of complex information, and improved accessibility of the content. All content has been carefully reviewed and validated by our team to ensure accuracy and alignment with CAA policies and standards.

6.6 We look forward to working with you as we navigate the exciting and challenging landscape of AI in aviation. Together, we can ensure that the UK remains at the forefront of safe, innovative, and responsible AI adoption in the aerospace sector.

APPENDIX A

Abbreviations

AI	Artificial Intelligence
ASTM	American Society for Testing and Materials
ATM	Air Traffic Management
CAA	Civil Aviation Authority
CAP	Civil Aviation Publication
CE	Conformité Européene, or European Conformity marking
CMA	Competition and Markets Authority
DSIT	Department for Science, Innovation & Technology
EASA	European Union Aviation Safety Agency
EU	European Union Aviation Safety Agency
eVTOL	Electric Vertical Take-off and Landing
GDPR	General Data Protection Act
ICAO	International Civil Aviation Organisation
ML	Machine Learning
NAA	National Aviation Authority
NGO	Non-Governmental Organisation
OECD	Organisation for Economic and Cooperative Development
RPAS	Remotely Piloted Air System
SDO	Standards Development Organisation
STPA	System Theoretic Process Analysis
UK	United Kingdom
UKRI	UK Research & Innovation

UN	United Nations