Capgemini x AEROSPACE TECHNOLOGY INSTITUTE

AI FOR AEROSPACE

Make it real.

Foreword

Across civil aerospace, organisations are operating under increased delivery pressure while also navigating structural change. Production rate increases, workforce transition, and readiness for future aircraft and propulsion programmes are all unfolding in parallel. In this environment, competitiveness is shaped less by individual technical breakthroughs and more by the ability to consistently convert engineering intent into assured, repeatable outcomes.

The recent UK national aerospace technology strategy, Engineering Growth, produced by the ATI, highlighted the importance of adopting Artificial Intelligence (AI) to help solve some of the challenges the sector faces. In a sense, AI is both overpromised and underrated: yet to be proven in full-scale implementation, but underestimated in terms of its transformative potential, even in high-barrier industries. Across aerospace its adoption to date is uneven, and many manufacturers have yet to explore how AI can support delivery, consistency, quality and scale. Much of this caution comes back to the question of trust. In a safety-first sector like aerospace, how do we responsibly pursue the opportunity AI presents while being clear about and mitigating the risks? And how do we build adoption pathways that work alongside existing engineering and assurance practices?

AI is no longer confined to research roadmaps or isolated demonstrations. It is becoming a practical capability that will increasingly shape how aerospace organisations deliver complex programmes with confidence.

AI is not an end in itself, but a toolbox that creates value. Its adoption in aerospace will be an iterative journey that builds capability over time – technically, organisationally, and culturally. This paper, a collaboration between the Aerospace Technology Institute and Capgemini, seeks to support aerospace organisations as they navigate the practical adoption of AI in environments where safety, assurance, and long programme lifecycles are defining constraints. It illustrates credible applications of AI across design, manufacturing and through-life support today and how these capabilities will evolve, as assurance, integration, and confidence increase, to support better use of time, expertise and resources in the future. Most importantly, it cuts through the noise, offering an aerospace-specific, technology-literate perspective that supports confident, informed decision-making when considering your own adoption of AI.

Gary Elliott,
CEO, Aerospace Technology Institute

Keith Williams,
Global Head of Technology & Innovation, Capgemini

Content

Introduction

Why now?

Aerospace is navigating a period of significant change, shaped by emerging digital capabilities and evolving industrial priorities. Delivering unprecedented production rates while pursuing ambitious innovation and sustainability goals poses a significant challenge. Against this backdrop, the opportunity AI offers is significant. How can this diverse cluster of technologies – ranging from well-established methods to emerging capabilities – help the sector meet its immediate and long-term ambitions?

Advancements in the last decade have propelled AI to the forefront of technological innovation and real-world applications. Enhanced computational power, vast amounts of data from an increased number of sources, improved algorithms, accessible frameworks, commercial interest, regulatory support and interdisciplinary integration are all coming together to create a fertile ground for rapid AI innovation. Generative AI (GenAI), in particular, has rapidly entered the business world, bringing an unprecedented level of AI visibility into offices worldwide. Agentic AI, where AI moves from generating outputs to taking on multi step work on a user's behalf, is the next step, providing planning, using tools, and coordinating actions within defined boundaries. When considering the potential of AI, it is essential to recognise the distinctive starting point of aerospace: a sector where engineering, including digital engineering, is fundamentally grounded in physical reality. For decades, physics-based modelling, simulation and analysis have been used to understand, predict, validate and improve real-world performance. AI adoption should build on this foundation: augmenting engineering expertise while remaining anchored in physics, evidence, validation and human accountability.

Making sense of the transformation

A useful way to think about AI is within the wider framework of digitalisation: as an emerging capability within a broader spectrum of digital technologies. While AI brings distinctive opportunities and risks, its adoption shares many of the same underlying requirements as digital transformation more broadly: robust data management, appropriate IT infrastructure, effective integration across systems, and clear strategic intent. In engineering and manufacturing environments, this also means developing shared definitions and process understanding so that data from different tools, teams and stages of the lifecycle can be interpreted consistently. Even the most advanced AI tools require strong digital foundations.

Organisations need to be explicit about where they want AI to make a difference, which problems are worth prioritising, and how success will be measured in terms of value, risk and cost. From there, adoption is likely to be iterative: building practical experience while strengthening the digital foundations, governance and culture needed to scale. The journey must encompass workforce development, change management, and strategies for ethical and secure deployment as much as data management and IT strategy. Without this alignment between people, culture, digital capability, and strategy, AI adoption is unlikely to deliver sustained value.

A competitiveness shift?

The nature of aerospace competition is changing. Success will depend not only on what organisations can design, but also on how reliably they can validate, produce and support it. AI has the potential to support organisations across the aerospace value chain as they respond to this shift. In design and early engineering phases, AI can enable engineers to explore a wider range of viable design options earlier, before narrowing them down through detailed analysis and physical validation. In manufacturing, AI-enabled optimisation can reduce lead times and improve consistency. In assembly and inspection, it can support defect detection and reduce variability. In through-life support and maintenance, intelligent systems can support failure prediction, inform maintenance scheduling, and improve the capture and reuse of inspection and operational knowledge as workforces evolve.

This shift is already visible in both aerospace and adjacent sectors, where AI is being used to accelerate analysis, improve consistency and support complex decision-making. Examples across aerospace and adjacent sectors, including automotive, consumer electronics, and Formula 1, illustrate this. Physics-based AI tools simulate complex behaviours in minutes, complementing traditional methods and accelerating innovation. GenAI is also being piloted in areas such as predictive maintenance, digital twin modelling, and intelligent documentation. In operational and knowledge-intensive workflows, emerging applications include AI assistants that help users navigate technical documentation and tools that support service, engineering and production teams.

The opportunity is not only efficiency, but a stronger ability to learn, validate and scale with confidence. This distinction is particularly important for the UK aerospace sector, where competitiveness depends on high value engineering capability, ability to deliver reliably while increasing rate, and resilience to workforce and supply chain pressures. The UK aerospace sector's ability to double in value over the next decade, and reach $41bn by 2050, depends on how successfully we embrace new technologies to accelerate our productivity and increase efficiency. AI clearly offers an opportunity to do that, acting as a multiplier for our workforce expertise and enabling us to do more, more quickly, and with greater precision. Organisations that harness AI-enabled capability safely and at scale stand to gain a significant advantage when competing for the next-generation aircraft programmes.

This paper aims to address both the practical examples and the broader enablers of AI in aerospace by examining:

- The cluster of technologies commonly described as AI,
- Case studies of current uses,
- Applications that support rate ramp-up and industrial readiness ahead of the next major aircraft programmes, and
- Potential game-changing capabilities for the future

It supports aerospace organisations to consider where AI fits within their strategy, which problems are genuinely worth addressing, and what needs to be in place before acting.

Demystifying AI

Defining AI

A broad definition of AI is the pursuit of building systems that perform tasks which typically require human intelligence. Within this field lies a wide spectrum of techniques: at one end, purely statistical methods for modelling data with no pretence of true intelligence; at the other, systems with broad capabilities to act in the world and an ambition to achieve genuine intelligence.

Two commonly cited definitions help illustrate this distinction:

"**Machine learning** is the field of study that gives computers the ability to model and solve narrowly specified problems without being explicitly programmed, and to improve their performance through experience" (Arthur Samuel, 1959)

"**Artificial intelligence** is the field concerned with designing and building of intelligent agents that receive percepts from the environment and take actions that affect that environment" (Russell & Norvig, 1995).

What all types of AI have in common is an iterative process, typically referred to as learning, used to build and refine a model of the data the AI system works with. Building models through learning by example is what differentiates AI and machine learning (ML) from more traditional, bottom-up mathematical or deterministic approaches.

A defining characteristic of many modern AI systems is that they are probabilistic rather than deterministic.

The goal is not to uncover a perfect underlying logic, but to build a model that is mostly right, most of the time. The outputs of these models may vary from run to run, with no guarantees of correctness, only a statistical likelihood that the result is right. This probabilistic nature is both a strength, enabling AI to tackle complex, generic problems, and a weakness that must be understood and managed when designing and using AI systems.

AI foundations

Subfields of AI

To make sense of the AI toolbox, we look first at the underlying approaches used to build AI systems, then at the capabilities these systems enable in practice, and finally at how models learn and specialise. GenAI is a useful starting point because it has driven recent visibility, but it is only one part of a wider AI toolbox. It is a prominent example of statistical AI, one of four broad subfields, alongside symbolic, stochastic and probabilistic AI. GenAI models learn patterns in data so well that they can generate new outputs with similar structure and characteristics. Large language models (LLMs) are an example of this, learning patterns in written language well enough to generate text that, in many situations, reads as fluently as human writing. More broadly, GenAI can be applied to any data with structure – for example, learning how pixels combine to form images, or how air traffic patterns evolve over time to generate realistic synthetic traffic data.

Focusing only on GenAI risks narrowing the view of what AI can do for aerospace. Different AI approaches are suited to different problems. Statistical approaches work best when you have lots of data, but the underlying behaviour is hard to write down as rules. Symbolic approaches are useful when expert knowledge can be expressed explicitly, and decisions need to be efficient and explainable. Stochastic methods can be powerful for search and optimisation problems where "good enough" solutions emerge through iteration, while probabilistic methods help when uncertainty is part of the problem, not just noise in the data.

These approaches can be combined into hybrid systems that can draw on their complementary strengths. Over-emphasising the capabilities of one subfield at the expense of others is a mistake, as is assuming a single approach is the most effective solution: complex problems will often require hybrid approaches that combine pattern recognition, explicit knowledge, optimisation and uncertainty management.

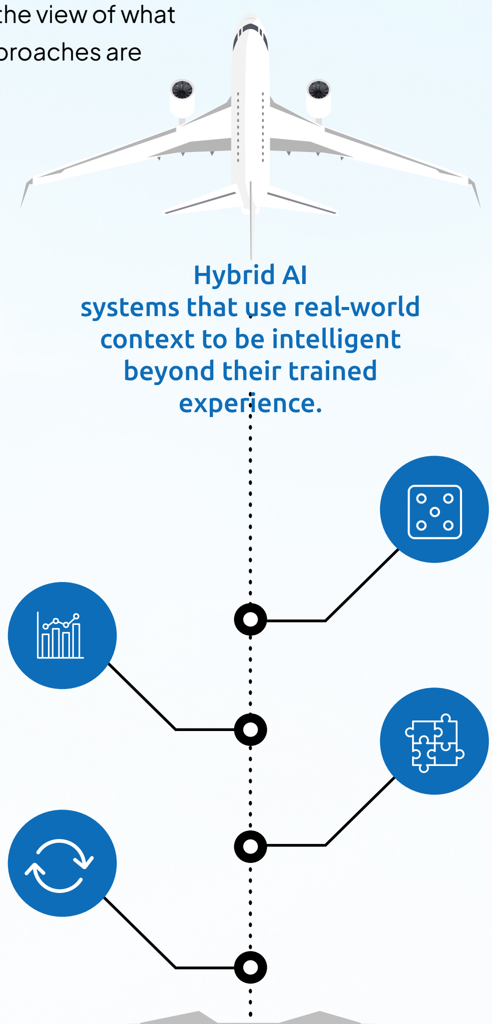

Statistical AI
where learning involves iteratively refining the parameters of mathematical functions so that they better fit their training data. This branch of AI includes techniques such as linear and logistic regression, neural networks (including deep learning and convolutional neural networks, CNNs), Support Vector Machines (SVM) and all of Generative AI, including Large Language Models.

Stochastic AI
where we harness the complex emergent intelligence generated by populations of simple entities. Techniques include evolutionary computing such as genetic algorithms, swarm intelligence, particle swarm optimisation.

Probabilistic AI
is very similar to symbolic techniques but explicitly model uncertainty and the probabilistic nature of many real-world decisions. Techniques include Naïve Bayes classifiers, Bayesian Belief Networks, Markov Modelling, Causal and Counterfactual reasoning.

Symbolic AI
where knowledge is represented in explicit symbolic form, and then logical operations and inference are used to make decisions. This branch of AI includes techniques such as decision trees, random forests, planning & reasoning, and logical inference.

Figure 2: Four main subfields of AI

AI researchers aim to combine the subfields, creating hybrid AI systems that blend complementary strengths. Neuro-symbolic systems are one example of this broader hybrid approach. To address limitations of purely statistical AI (such as hallucination, weak logical reasoning, and high data or compute requirements), developers are building neuro symbolic systems that blend statistical and symbolic approaches to deliver the fast, intuitive, pattern-based intelligence of GenAI, with the rational, explainable, knowledge-based logic of symbolic models.

AI capabilities and use cases

The subfields above describe the underlying approaches used to build AI systems; Figure 3 takes a more practical view, grouping AI by capability or method, and use case. The purpose of the table is to show the range of AI tools and examples of where they may be applied.

AI capabilities, methods and use cases	Example use cases
Understanding and decision making	
Supervised learning	Email classification, credit scoring, medical diagnosis
Unsupervised learning	Customer segmentation, anomaly detection, topic modelling
Reinforcement learning	Robotics, game AI, energy optimisation
Generative AI	Content creation, design, simulation, storytelling
Causal inference	Policy evaluation, healthcare outcomes, decision science
Evolutionary algorithms	Scheduling problems, structural & physical optimisation, symbolic search
Sensing the physical world	
Computer vision	Object detection, facial recognition, medical imaging, autonomous vehicles
Speech recognition	Voice assistants, transcription, call centre automation
Audio processing	Predictive maintenance, acoustic monitoring, wildlife tracking
Gesture recognition	AR/VR interfaces, gaming, touchless control
Tactile AI (haptics)	Robotic surgery, prosthetics, haptic feedback systems
Sensor fusion	Robotics, drones, autonomous vehicles
Physical AI / acting in the physical world	
Robotics and control systems	Manufacturing, logistics, healthcare robotics
Smart automation	Workflow automation, real time monitoring, industrial control
Agentic AI	AI copilots, autonomous research agents, multi step task automation, digital workers

Figure 3: AI capabilities/methods and use cases

The breadth of approaches means there is no single platform or model that will be right for every context or application. AI is better understood as a toolbox, from which different capabilities are selected depending on the problem being addressed, the data available, and the level of risk that can be tolerated. This also means AI adoption should progress iteratively through experimentation, validation and trust-building, rather than being treated as a one-off implementation.

Learning and specialisation

Having looked at the underlying approaches used to build AI systems and the capabilities they can enable, it is important to consider how those systems learn, specialise and are configured for use. The way an AI system learns or is configured is central to what it can be trusted to do. Not all AI systems are trained in the same way: some learn from labelled examples, where the desired answer is known in advance; others identify structure in unlabelled data, learn through feedback from an environment, search iteratively for better solutions, or combine learned patterns with explicit domain knowledge.

In some cases, capability is embedded during training; in others, it is added later through fine-tuning, retrieval, reasoning, or integration with external tools and rules. These distinctions affect explainability, reliability, data requirements, validation burden, and the degree of human oversight required. A classification model used to support inspection, a generative model used to explore designs, and an agentic system coordinating maintenance tasks may all be described as AI, but they learn, operate, and need to be governed in different ways.

GenAI has become the most visible family of approaches in recent business adoption. Because generative models are general by design, a key challenge in enterprise and engineering settings is grounding them in trusted, domain-specific information and reducing the risk of confident but incorrect outputs. For this reason, various techniques exist to instil specialist capabilities in a generative model.

Two common approaches are fine tuning (additional training on domain data) and retrieval-augmented generation (RAG), where the model is given relevant information from your own sources at the time you ask the question. Some models also offer "reasoning" modes where the model iterates and refines an answer before responding. The table below compares the pros and cons of adding specialist capability to generative models at different stages: during pre-training, after training, or at runtime.

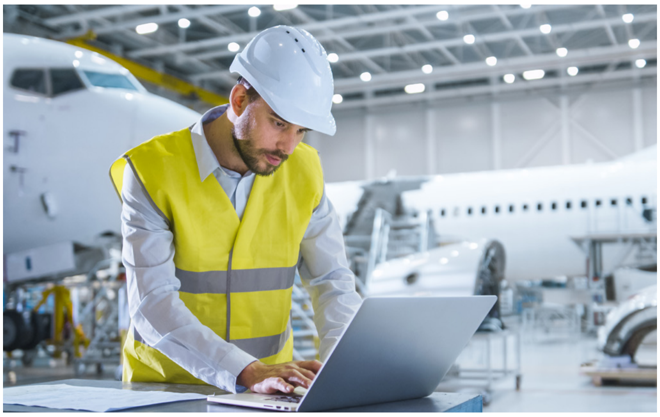

Stage	Pre-training	Post-training (e.g., fine-tuning, distillation, transfer-learning)	Test-time / Runtime (e.g., RAG, reasoning)
Advantages	• Requires no furthe modification to use, so simple to deploy.	• Model will be faster at runtime • A one-off exercise to embed non-changing static knowledge.	• Adaptable beyond pretrained experience • Can incorporate dynamically changing domain-specific knowledge
Disadvantages	• High cost • Training time	• Maintenance of derivative models • Catastrophic forgetting • Risk of over-fitting	• Higher runtime cost • Slower run-time speed

Table 1: Specialisation of generative AI models

Deploying AI

Physical AI

Having looked at AI foundations and practical capabilities, the next question is how these systems are deployed. While many AI systems operate in purely digital environments, supporting analysis and decision-making, physical AI refers to AI deployed into physical systems, such as robotics or embodied systems, where sensing, movement, control and interaction with the real world are key. This introduces additional challenges, particularly when systems must interact with people or other AI systems, requiring sensing technologies, such as computer vision and speech recognition, alongside the engineering capability to perceive, move, and act in the real world.

Agentic AI

For much of our history with technology, making a computer solve a particular task has required highly specialist skills: first to determine the solution, and then to implement it through programming. Agentic AI systems can instead allow users to express higher-level objectives and achieve a goal, even where they do not know the detailed steps involved. This introduces a more proactive form of computing where autonomous agents operate within defined boundaries to support complex decision-making. These agents synthesise data from Internet of Things (IoT) sensors, workforce systems, and planning tools to dynamically coordinate tasks such as maintenance scheduling. By interpreting structured domain knowledge through ontologies, they ensure decisions align with safety, regulatory, and operational requirements.

This marks a shift toward more adaptive, intelligent operations – where agents act independently but responsibly, improving agility and precision in environments where compliance and responsiveness are critical. In manufacturing terms, this makes the leap from following predefined steps to a system that decides its own path to an outcome within set limits.

There are three considerations when deciding how much responsibility to give an AI system:

Autonomy – the degree to which an agent can make independent decisions without outside direction

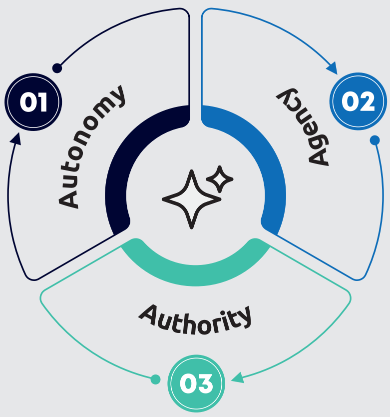

Agency – an agent's capacity to act on those decisions and take actions that affect its environment.

Authority – the limits of the actions an agent is allowed or not allowed to take.

The term "agent" refers to any entity that works on behalf of another entity, working to accomplish high-level objectives often using specialist capabilities. Agents have a degree of autonomy and authority to take actions that modify their world. https://www.capgemini.com/insights/research-library/business-meet-agentic-ai/

The question is not simply whether an agent can act, but what it is allowed to do, how its decisions are constrained, and where human judgement remains essential. The decisions made about the levels of autonomy, agency and authority given to an agent will determine the risks and benefits of delegating the task. As autonomy grows, the importance of good problem framing, ethical oversight and collaborative, informed decision-making increases. Applied responsibly, these technologies are likely to shift engineering sectors from tool-centric to outcome-driven processes and roles, fundamentally reshaping how we work. In practice, most systems will also involve multiple agents working together to solve tasks, although this may not be visible to the end user.

AI-augmented engineering

The aerospace industry is already concerned about the future workforce and, as engineering resources become stretched, AI will enable engineers to focus on core problems, rather than repetitive tasks, and enhance organisational knowledge retention in a fluid, and often geographically distributed, workforce and supply chain. As AI tools evolve, organisations will have to consider both how these tools are integrated into engineering workflows and how engineers develop the literacy to use them effectively.

AI-augmented engineering could become a core capability for supporting traditional engineering tasks and decision-making. Engineers who master these tools will be in high demand across all sectors. Using AI to augment traditional engineering tasks will create a force multiplier that enables faster, more informed, and more context-aware choices by reducing cognitive load and unlocking new opportunities for innovation. To ensure safe adoption, engineers trained in the underlying physics and fluent in the key industry challenges will be key to implementation, and will also need to oversee and validate AI outputs and monitor deployment. This approach can accelerate the onboarding of new talent and preserve decades of engineering experience, driving collaboration, quality and productivity.

Adjacent industries are adopting similar approaches. In the automotive sector, GenAI supports diagnostics and repair; in defence, autonomous agents assist with mission planning and logistics; in energy, GenAI enhances field service and safety compliance; in the life sciences, AI accelerates biological sequence generation and targeted molecule design. While many of these implementations are still anonymised, hypothetical, or in pilot stages, they demonstrate the versatility and scalability of AI in engineering. In all these cases AI enhances, rather than replaces, human expertise. Applied in this way, it supports smarter, faster, and more resilient engineering in domains characterised by complexity, risk, and long development cycles.

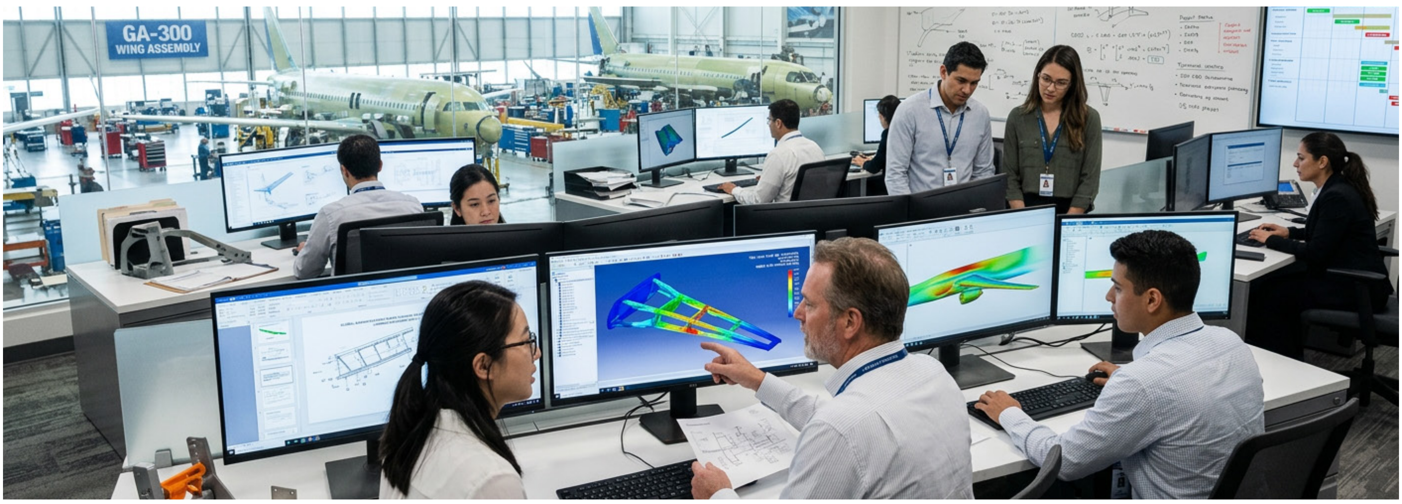

AI applications across the value chain

AI is being applied across different stages of the aerospace value chain, from design to through-life support. For each stage, AI has the potential to create value, illustrated by case studies of current activity in the sector in this section. These examples show how organisations are exploring where AI can create value, what is required for responsible deployment, and how confidence is built over time.

The case studies below span funded R&D projects, validation of commercial products for aerospace use, and operational applications where AI capability continues to evolve as datasets grow. Taken together, they reflect the range of maturity that exists across the sector today.

Pilot projects play an important role in this process as a mechanism for learning and capability building, particularly when they are undertaken with a clear view of how insights can scale, integrate, or inform longer-term organisational capability. In all cases, the ability of AI applications to drive value reflects the strength of the underlying foundations in data, process understanding, and digital capability.

Design, verification and validation

Context

Design, verification and validation are essential stages in the aerospace value chain and closely linked to certification. In a historically linear engineering process, high-level product requirements are broken down into specifications for sub-systems and components. Engineered solutions are then designed and tested to ensure they meet the strict safety, performance, and regulatory standards of the aerospace sector. Validation is critical to avoid lengthy and costly redesigns. Striking the right balance between quality and speed is therefore essential.

Opportunity

Advanced simulation, digital engineering, and model-based systems are now leveraging data and advanced modelling capability, using ML as an extension of advanced statistical methods, to accelerate development and response times to changing customer and regulatory demands. Incorporating such digital innovation – for example, through optimised aerodynamics, materials, and systems integration – enhances product performance and competitive advantage.

Looking ahead, more advanced AI technologies have the potential to evaluate a wider range of design options while significantly accelerating time-to-market. AI-driven generative tools can examine thousands of design options in minutes, optimising for weight, strength, manufacturability and other key design parameters. In addition to speeding up the traditional design-simulation loop, AI tools will also be able to generate fundamentally new designs and optimise test plans.

Despite the aerospace sector's complex and safety-critical nature, successful application of AI tools in other sectors shows how they could further transform our industry. More cutting-edge applications in aerospace could include adversarial testing and failure

Mode simulation, already transforming software testing in the financial and professional services sectors. AI has also fundamentally transformed design workflows, predictive analytics, and manufacturing efficiency in sectors such as automotive and consumer electronics. In life sciences, generative models and specialist domain models combine to use specific scientific constraints to generate valid and probable solutions.

Adoption considerations

In design, verification and validation, AI adoption depends on ensuring that faster design exploration does not come at the expense of technical validity, physical realism or engineering accountability. For AI tools to offer value in aerospace design, they must understand the competing trades that need to be made, and the long-reaching impact of decisions taken at this stage on the overall success of the solution. Organisations need to recognise that while AI offers opportunities to increase speed and optimisation, it remains essential that outputs are grounded in scientific and engineering reality.

For AI to support aerospace design and optimisation, its outputs must be technically valid, not merely plausible. Large language models and other generative systems can produce convincing outputs without reliably representing the underlying physics, constraints, or engineering trade-offs, creating risk if used to inform decisions without appropriate validation. Because many data driven AI approaches are probabilistic rather than deterministic, hybrid approaches that combine learned patters with domain knowledge, constraints and models, can help improve validity and build confidence in outputs. A well-known example of such a hybrid approach from biology is AlphaFold2, developed at Google DeepMind and recognised with the 2024 Nobel Prize in Chemistry for its contribution to protein structure prediction. The case study below examines how a similar principle can be applied in aerospace practice.

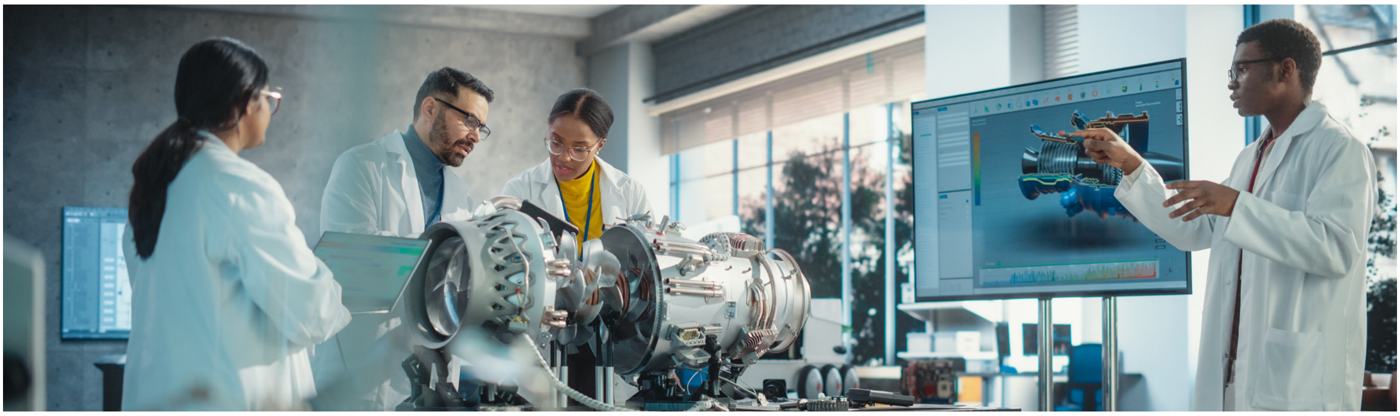

BeyondMath: Optimising design generation and evaluation

Company context

BeyondMath® is a UK-based SME developing hybrid AI tools that blend generative models with physics-based techniques. To date, the company's approach has been applied to external aerodynamics in automotive and motorsports such as Formula 1, where the tool allows engineers to complete tens of thousands of design variations in short timeframes. BeyondMath is now exploring the implementation of this capability in the aerospace sector. An example of the company's ongoing development in aerospace includes its involvement in the recently funded ATI STRATA project in partnership with Honeywell. In this project, BeyondMath will implement its generative physics platform to support the optimisation of additively manufactured components for future Environmental Control and Cabin Pressure Control Systems.

Challenge

The early design stages often present the most significant opportunities to shape concepts and maximise benefits by optimally positioning designs to balance performance across a wide range of competing constraints. However, engineers often operate with limited data and time to explore the full breadth of the design space. Here, AI methods that support rapid design space exploration can provide significant benefits by reducing the time to iterate and generate a multitude of ideas. This is where BeyondMath's approach can offer significant value, applying a generalised physics approach to steer development towards solutions grounded in real-world engineering requirements.
The challenge is maintaining accuracy and validity while maximising options. This is where many ML methods fall short. Traditional ML and surrogate models require case-specific datasets. When designs deviate from this space, accuracy drops sharply because ML tools rely solely on learned patterns and lack the ability to extrapolate using the underlying principles of physics. Teams then spend valuable time retraining models, diverting engineers away from innovation to training datasets. BeyondMath's physics-based approach addresses this by training the AI on fundamental physics, aiming to provide better model flexibility while maintaining accuracy across a wide range of use cases. This approach aims to reduce the need for more training data, freeing up engineers to focus on design development.

Solution

BeyondMath addresses the challenge by grounding the tool's capabilities in fundamental physics. In aerodynamics applications within the Formula 1 industry, this involves teaching the tool the governing physical principles that determine aerodynamic performance, such as laminar-to-turbulent flow transition behaviour, Reynolds number effects, curvature impacts, and separated flow structures. By understanding these core rules, the tool can adapt to new, previously unseen configurations while remaining directionally accurate.

Expanding this approach into aerospace creates new development opportunities. As aerospace customers

Show interest in multi-physics approaches, the goal is to continue developing the tool to encompass a broader understanding of the underlying principles of physics, enabling it to adapt to aerodynamic, structural, thermal, and even acoustic problems. The key is achieving the accuracy required to give engineers directional confidence, while continually broadening the understanding of the physical world.

Implementation

The philosophy implemented by BeyondMath is that a foundational model should be pretrained on the underlying principles of physics. These underlying principles are governed by partial differential equations, and therefore, the AI model must be capable of predicting the outcomes of these equations. To achieve this, the training cases are designed to build a library of physical understanding that is not tied to specific use cases, allowing the tool to predict underlying behaviour. As the model expands its physical knowledge, the ultimate goal is to develop a system that can maintain accuracy when applied to any new problem without requiring further training.

When implemented, the tool can interpret any design geometry and intelligently predict physical interactions by recognising the problem's boundary conditions and constraints, generating solutions that closely match those calculated from the differential equations. The solutions produced are also time-dependent, reflecting the true nature of physics; in aerodynamics, this means that fully unsteady flow velocity and pressure fields are generated for further analysis.

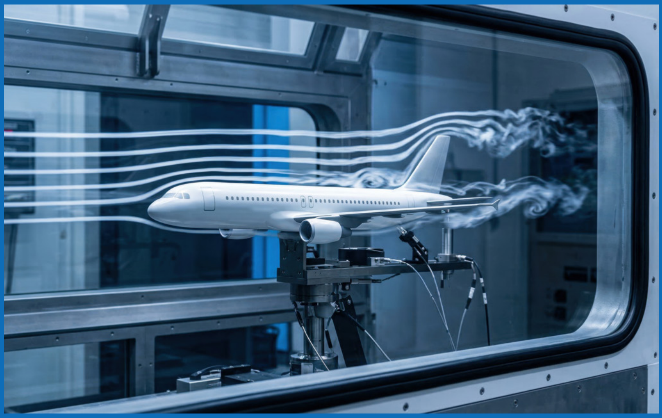

Results and impact

In an automotive aerodynamics case study, BeyondMath's tool simulated the full flow in 1.5 minutes - compared to 24 hours on a 720-core high-performance computing system - while maintaining accuracy within 2% of a high-fidelity Detached Eddy Simulation.

The tool enables engineers to screen and iterate across a vast design space rapidly, then validate selected designs using CFD. As this capability continues to develop, the key question is: how will engineers utilise and integrate such tools to augment their skills and gain maximum benefit? Tools like BeyondMath's could significantly increase iteration speed, improve design exploration, and enhance the overall performance of future aerospace systems. Balancing the trade-offs between the different capabilities of hybrid AI systems to drive the right engineering outcomes will be key to implementing AI effectively.

Manufacturing and assembly

Context

The high-value jobs associated with aerospace manufacturing and assembly are crucial both to the UK aerospace sector and the broader UK economy. Aerospace manufacturing transforms a pre-processed raw material into a finished part, such as milling aluminium billets or laying up and curing carbon fibre pre-pregs, while assembly processes handle individual components before their subsequent fitting to produce a product, such as landing gear.

Aerospace is a sector with high barriers to entry and exacting quality, traceability, and manufacturing tolerance requirements. The rigour required can foster a level of conservatism, especially for high-rate aircraft programmes, where disrupting proven processes could jeopardise delivery schedules. Aerospace manufacturing supply chain challenges are well-documented and continue to pose a risk for current and future aircraft programmes. Aerospace manufacturing also faces workforce challenges, including a demographic wave of retirements on the horizon and limited interest from the next generation in repetitive, manual work. By embedding intelligence into manufacturing operations, aerospace organisations can build greater resilience and responsiveness, even in highly complex production environments.

Opportunity

Across the diverse range of processes that comprise aerospace manufacturing and assembly, AI presents significant opportunities to boost productivity while upskilling the workforce for advanced, digitally enabled roles. AI is used today to augment existing workflows, reduce lead times, improve quality, and enable solutions that were previously out of reach in areas such as augmenting computer-aided manufacturing (CAM) to optimise tool paths.

AI predictive models are increasingly being used to identify production bottlenecks weeks in advance. These insights are directly integrated into operational systems, enabling dynamic scheduling, resource allocation, and proactive corrective actions, delivering productivity gains of up to 50%. Defect detection and identification is another use case gaining traction. Computer vision methods, trained for medical applications, have been successfully applied in aerospace to identify defects in processes such as additive manufacturing. Given inspection is typically an offline, non-value-added activity, reducing the time it takes while increasing accuracy offers significant benefits.

In the longer term, the greatest opportunity lies in optimising human involvement within manufacturing process control to achieve more, faster, with less manual intervention. Skilled human operators will focus on the most critical aspects of production while AI can support delivery of repetitive, low-value tasks by identifying and rectifying deviations from strict manufacturing tolerances.

Adoption considerations

In manufacturing and assembly, successful AI adoption will depend on proving value early while building the data, digital infrastructure and workforce confidence needed to scale into more critical production processes. In near-term applications, active R&D is accelerating adoption with several technologies already approaching production readiness. Capturing value from these early implementations – whether that be increased productivity, improved yields, or reduced scrap – will be critical to justify scaling up and further investments.

Long-term adoption of AI in aerospace hinges on building trust over time. This will require comprehensive, well-governed data models, continuous workforce upskilling, and iterative collaboration between AI systems and human experts. Confidence can be developed through a staged approach, beginning with lower value, non-safety critical parts, such as bracketry, and progressively advancing to higher-criticality components and complex assemblies.

A key enabler for AI in manufacturing and assembly is continued digitalisation across aerospace operations. Today, many processes remain paper-based and reliant on legacy, non-connected equipment. Building AI capability means improving not only data capture and system connectivity, but also the shared understanding of processes, tolerances, defects and quality outcomes that allows information to be used consistently. Without investments to connect these disparate systems and capture high-quality data, the full benefits of AI cannot be realised. AI adoption does not depend on perfect data from the outset, but early, well-scoped use cases can help reveal which foundations need to mature as capability develops.

CAM Assist: Transforming computer-aided manufacturing

Company context

CloudNC, a UK-based technology company, aims to transform global manufacturing by applying AI to computer-aided manufacturing (CAM) programming. Its software, CAM Assist, automates CAM programming by generating and supporting tool selections, approach directions and cutting parameters for machined parts, seamlessly interfacing with industry-standard CAM packages. CloudNC supplies its software as well as machined parts from its factory in Chelmsford. CloudNC has been working with GKN Aerospace and the Advanced Manufacturing Research Centre (AMRC) in the Autonomous Robust & Rapid Processes for the Machining of Aerospace-Specific Parts & Components (ARRP) project to develop the CAM Assist program to better serve the needs of aerospace manufacturers.

Challenge

CAM programming is time-consuming; aerospace organisations can spend over 60 hours of programming time preparing a part for manufacturing. The process relies heavily on the knowledge of highly experienced individuals, requires several iterations to ensure that the program both works and is optimised, and is susceptible to mistakes. This wastes both material and time and can damage expensive tools and equipment. Moreover, an ageing workforce and recruitment challenges have made CAM programming a bottleneck in manufacturing.

CloudNC initially focussed on the mass market, where its CAM Assist software has delivered strong results on small to medium-complexity parts made from “soft” metals such as aluminium. Aerospace parts, however, present far greater challenges: complex geometries, challenging features, tight tolerances, and materials like titanium that are significantly harder to machine. Through the ATI-funded ARRP project, CloudNC has been able to focus on these aerospace-specific requirements, developing advanced capabilities to meet the sector’s demanding standards.

Solution

CAM Assist generates CAM programs using a feature-based approach that mirrors the decision-making process of an experienced CAM programmer. It can assess the part geometry alongside the available machines, tools and fixtures to create an optimised machining strategy. While completeness of the program varies by part complexity, CAM Assist can produce up to 80% of the program even for highly complex parts, significantly reducing manual programming time. CAM Assist works as a plugin to major CAM packages, seamlessly gathering the necessary data, creating the program, and returning for final optimisation. Features of the program, including cutting models that determine machining parameters such as "feeds and speeds", have been developed by CloudNC using a combination of mathematical models and empirical data generated in its own factory.

The ARRP project has enabled the CloudNC team to develop more sophisticated strategies for machining – both roughing and finishing – required to tackle more complex aerospace parts. The project also generated data that supports the machining of titanium, including the transition from 3-axis machines to 3+2 axis machines, adding in rotation. The CloudNC team were also able to develop the ability to read model-based definition (MBD) files and to probe the part during build, measuring and compensating.

Implementation

The CAM Assist program is commercially available today, though its current version is somewhat limited for aerospace applications. Through the ARRP project, CloudNC developed a customised version for project partners GKN Aerospace and the AMRC. Both partners have been comprehensively validating the software on real components, which has involved significant physical experimentation in CloudNC, GKN and the AMRC. At the conclusion of the project in 2026, advanced new aerospace-appropriate features will be rolled into CAM Assist.

One of the biggest challenges on the implementation journey has been producing an output that a CAM programmer would expect to see – this is critical for building trust in what the AI has generated. For adoption to succeed, users have to be able to accept and trust the AI output and resist the urge to override it, even if the approach differs from their own.
To achieve this, significant effort within the system's design has gone into capturing the thought process of an experienced CAM programmer.

Results and impact

Against the ARRP project objectives, the team has delivered strong results. The collaboration between CloudNC and GKN Aerospace has been impactful, and initial experimental results on an actual GKN part – machined at CloudNC – have validated that the AI solution can generate a CAM program in half the usual time while meeting the required quality.

Although the journey has not been without challenges, particularly as the AI approaches the machining strategy differently to humans, the project has demonstrated that the CAM programmer's role can evolve. Routine tasks can be automated, shifting the focus to optimisation and refinement, making the role more rewarding.
Looking ahead, there is scope for future innovation, including new features and introducing AI-driven decision-making earlier in the process to influence setup, fixtures, nesting, and orientations. These are currently defined manually and vary by individual experience. Further material advancements are also on the horizon, with titanium still in progress and nickel alloys emerging as critical for aerospace applications

Through-life support and maintenance, repair and overhaul (MRO)

Context

Further along the value chain, through-life support has been an early adopter of digital technologies and AI in aerospace. This process-centric, labour-intensive, and documentation-heavy subsector ensures continued airworthiness requirements and lifecycle servicing of airframes, engines, and components.

The introduction of the 'Power by the Hour' model in the 1990s marked a pivotal shift, moving the industry from reactive maintenance to a service-orientated model prioritising active management of reliability, efficiency, and customer value. MRO became a strategic enabler underpinning safety, profitability, and future innovation, with aerospace OEMs generating more than 50% of their revenue from MRO and aftermarket activities. Under this business model, consistent, reliable and transferable digitalised information has become paramount.

Opportunity

Leveraging AI capabilities in MRO represents a significant opportunity for the aerospace industry today. Durability issues with new narrowbody aircraft are delaying the retirement of existing fleets and driving up demand for MRO. In this context, AI can increase profitability through its potential to enhance efficiency, reduce turnaround times, and improve reliability. For businesses of all sizes in this area of the aerospace value chain, planning how to integrate AI into operations will be essential to maintain global competitiveness and secure a strong position in an increasingly contested MRO market.

AI can unlock value from large amounts of historical data to drive productivity and efficiency while continuing to meet regulatory requirements.
By enabling reliable and effective data interrogation,

it will overcome the limitations of previously fragmented systems and equip licensed aircraft mechanics with advanced capabilities including high-precision sensing and diagnostics, intelligent robotics, digital twins, and software-based diagnostic tools. Effective use of AI will also draw on methods and tools for accessing previously unreachable areas, rapid scanning and automated damage detection, as well as seamless.

digital threading of validated MRO data back to the original equipment design process, ensuring durability is prioritised alongside performance at the outset. The convergence of digital technologies and new AI capabilities will enable MRO operators to mitigate future workforce shortages as they struggle to compete against other modern technology and innovation sectors.

Adoption considerations

In through-life support and MRO, AI adoption depends on ensuring that efficiency gains are achieved without compromising continued airworthiness, regulatory confidence or trust in maintenance decisions. MRO organisations are operating under intense workload pressure, making it challenging to balance near-term turnaround requirements with long-term technology acquisition. A unified, strategic approach to AI, supported by agreed standards and active regulatory engagement, would help reduce uncertainty and provide greater confidence for investment in these capabilities.

Successful adoption of AI will depend on managing both technical risks, including the reliability and traceability of AI-enabled outputs, and commercial risks around investment and integration. It will also require workforce upskilling, so that maintenance teams can use AI responsibly and understand when to trust, challenge or override its outputs. Companies that build this capability early will be better placed to capture the opportunity and establish market leadership.

Intelligent Borescope: Improving performance through augmented inspection

Company context

Rolls-Royce has been at the forefront of aero engine innovation for over 100 years. Digital and AI transformation is underway within the company to build a high-performing and competitive entity that will continue to deliver mission-critical solutions to its customers. Engine maintenance and repair (MRO) has been among the early adopters of digital transformation and remains a critical sector of the aerospace industry, where further AI implementation can unlock greater efficiency, cost reduction, and safety assurance.

Solution

Rolls-Royce partnered with Waygate Technologies and RH Aero to deliver the Intelligent Borescope Method, introducing a high-resolution, templated imaging system for capturing detailed HPT blade visuals during inspection. This initiative establishes a structured approach to data collection, creating a foundation for future AI-driven tools that will support decision-making in inspection and engine maintenance planning. For airlines, the solution includes an app that converts high-quality video into still images for automatic inclusion in standard reports, saving hours of manual report writing and enabling instant sharing of results with Rolls-Royce.

Implementation

The Intelligent Borescope Method integrates advanced borescope and engine positioning hardware with digital and AI intelligence to tackle one of MRO's most critical metrics: engine downtime. The initiative began as a proof of concept using actual engine borescope data to develop and train AI algorithms capable of analysing blade condition. This initial phase validated the approach before progressing to a pilot and ultimately full production deployment.

The core objective was to extend engine on-wing time by correlating in-service usage with borescope video images of each blade. By linking the condition of every blade, the system could accurately predict how much longer the engine could remain operational before requiring removal. AI served as the foundation of this solution, enabling automated damage detection and predictive modelling through continuous algorithm training.

Three key pillars were developed:

- **Consistency –** Achieved by digitally positioning high-pressure turbine (HPT) blades and identifying damage features, creating a standardised approach to problem definition and location mapping.
- **Data management –** Implemented through cloud-based storage, ensuring immediate global accessibility for Rolls-Royce facilities and affiliates. This supports accurate, data-driven maintenance decisions to optimise time-on-wing.
- **AI capability –** As the cloud database grows, AI models evolve to deliver greater precision in damage identification, predictive analytics, and decision-making. This progression enhances certainty for airlines managing fleets, inspectors assessing MRO components, and Rolls-Royce in planning parts and services to minimise downtime and operational disruption.

This end-to-end approach from proof of concept to production has transformed borescope inspections into an intelligent, predictive process, delivering measurable improvements in engine lifecycle management and operational efficiency.

Results and impact

The Intelligent Borescope Method is making the HPT blade inspection process more efficient, minimising rework, reducing human factor risks, and limiting variability by providing digital augmented assistance to MRO inspectors. It has helped identify HPT blade wear patterns early, reducing re-inspection and increasing the certainty of maintenance planning and fleet management. The digitalisation of maintenance data in a structured form further opens its usage in future digital twin-based engineering design environments, interweaving complete lifecycle data for every engine in service.

The new system also enhances training and knowledge retention, particularly in instances of staff turnover, by simplifying the gathering of audit evidence and historical data searches. Rolls-Royce has built a 'maintenance brain' – a repository of validated inspection records, comparative analyses, and diagnostic benchmarks. Images captured through the Intelligent Borescope Method continuously enrich this system with new data, further training the system and enhancing its ability to provide AI-augmented decision-making support during inspection. Emerging AI capabilities of the system include damage recognition, autonomous measurements, the capability to recommend actions, predictive analytics for damage progression, and optimised scheduling of engine inspection and removal to improve planning stability.

Beyond this initiative, Rolls-Royce recognises the transformational opportunity AI technologies present for the aerospace sector and is supporting the broader adoption of AI in aerospace and beyond. To promote responsible implementation, it has made the Rolls-Royce Aletheia framework on AI ethics freely accessible in the public domain.

"Rolls-Royce is committed to pioneering the use of new technologies for the benefit of Service Operation for our Customers. Standardisation and quality of inspection images, starting with High Pressure Turbine Blades, deliver immediate operational benefits in sentencing efficiency and reducing human factor risk. Processing data shared in a Cloud system will then allow app development and set the foundation for future improvements."

Tony Bell, Director of Commercial Aviation Aftermarket Engineering

Getting ready for AI

As the case studies indicate, implementing AI and integrating it at full scale can take different paths. There is no single starting point or finish line for AI adoption in practice. Different organisations will move at different speeds, depending on what problems they are trying to solve and what foundations they already have in place. While iterative experimentation is essential for building capability, confidence and trust, it must still be aligned with clear strategic intent. Successful AI integration needs a clear strategy to ensure it delivers value, aligns with regulatory requirements, and adheres to ethical standards.

Talent, data, infrastructure

Successful AI deployment builds on three pillars: talent, data, and infrastructure.
As AI increasingly affects how work is done, organisations need to invest in the skills and understanding required to use these technologies effectively and responsibly. This includes not only deep technical expertise, but also broader AI literacy across the workforce, as AI is likely to influence a wide range of roles rather than remaining a niche capability held by a small group of specialists.

In aerospace, it is particularly important that teams understand both the benefits and limitations of AI, given the safety critical nature of many applications and the need to integrate new capabilities reliably into existing systems. Organisations will need to make deliberate choices about how they build AI capability: whether through upskilling existing staff, recruiting new talent, partnering, or a combination of all three. A key strategic decision is whether AI capability is treated as a peripheral skillset or embedded as a core organisational capability over time.

Organisations will also need a well-articulated data strategy. AI relies on data, but the limiting factor is often data quality and governance rather than volume. AI models reflect the data they are trained on, including any inherent bias within it, which makes the choice and quality of input data particularly important. For many smaller manufacturers, the main challenge might be inconsistent, fragmented, or hard to access data from existing processes. Establishing data governance frameworks that address issues such as bias is crucial.

The third pillar is infrastructure. In practice, deployments risk failure if the underlying systems cannot support them reliably. That means being realistic about what existing systems are capable of. As AI deployments scale, technology stacks will need to support data processing, analytics, and AI workloads appropriate to the use case. This may require

investment in high-performance computing systems, scalable cloud solutions, and secure, reliable storage infrastructures to ensure resilience and flexibility.

AI is therefore not a standalone solution, but part of a broader digital transformation journey. Advanced AI applications build on broader digital capabilities, such as robust data capture, storage, cleaning, integration, analytics, modelling, and ML, enabling them to operate reliably. This means that value-driven AI integration is rarely a simple, turnkey process as it often requires significant effort to build the core organisational and digital capabilities required. These foundations can have wider benefits, accelerating digitalisation and unlocking wider benefits across the organisation.

Change management and digital culture

Successfully building AI capability fundamentally depends on cultivating an organisational and sectoral culture to enable this transformation. Discussions around AI adoption often centre on IT strategy and business value, yet it is equally important to develop an environment that can facilitate and absorb innovation and change.

For many companies, overcoming resistance to change can be challenging. Leaders must actively engage employees across all levels, demystifying AI and demonstrating where it can add value. In safety-critical industries like aerospace, AI needs to be deployed as a tool that empowers the workforce, enhancing their expertise and enabling focus on higher-value activities. This principle is fundamental to building trust, fostering innovation, and safe and responsible adoption.

Change management plays a key role in whether AI adoption succeeds. It involves aligning on strategic business objectives, identifying opportunities for innovation, setting clear milestones, maintaining open communication channels, and providing robust support systems to help teams feed back and adapt. As with any disruptive technology, there will be tension between the urgency to adopt and the need to protect ongoing, often safety-critical operations. A gradual, well-structured rollout, combined with ongoing training and feedback loops, can help ensure that both people and processes, transition smoothly into an AI-augmented future.

Regulatory considerations and ethical AI deployment

The regulatory landscape governing AI is continually evolving and remains uneven, creating challenges for a highly regulated, global industry like aerospace. Aerospace companies need to stay informed about current and emerging regulations and engage proactively with regulatory bodies as AI standards and certification protocols develop. This differs from more established technology areas, where expectations and requirements are clearer from the outset.

In many instances, standards and adoption will need to develop together. Practical, well-governed experimentation can generate the evidence needed to shape future assurance approaches, while emerging standards provide the confidence required to scale. The absence of fully mature AI standards should not be a reason to pause innovation, but it does reinforce the need to start with bounded, transparent and well-supervised use cases.

Integrating AI into critical decision-making systems raises concerns around autonomy, accountability, and transparency. Ethical guidelines should ensure that AI technologies are used responsibly while also addressing issues such as potential degradation of human skills, workforce impacts and the societal consequences of over-reliance on technology. Larger organisations may choose to formalise this through ethics boards or advisory groups, while smaller organisations are more likely to rely on existing governance structures.

Collaboration across the ecosystem

AI adoption is not a journey that companies can take in isolation. Effective collaboration between industry players, technology providers and regulatory bodies is essential to accelerate AI readiness and innovation. Across aerospace, strategic partnerships or consortia may enable shared research and development, pooled infrastructure investment and more holistic approaches to cybersecurity.

Collaboration will be particularly important while AI standards continue to mature. Shared pilots and applied research can help generate the evidence needed to shape sector-appropriate assurance, standards and certification approaches, while emerging frameworks can give organisations greater confidence to scale. In this way, innovation and assurance can develop together, without waiting for standards to be fully settled before progress begins.

Cybersecurity and reliability

As AI becomes more integrated into engineering and manufacturing operations, exposure to cyber risks also increases. Aerospace organisations often operate in complex, highly connected environments, where AI systems interact with design tools, production equipment, quality systems, and supply-chain data. Robust cybersecurity needs to maintain production continuity, quality, and operational resilience. This means extending beyond securing IT systems and perimeters to include protecting data pipelines, models, and decision logic.

Deliberate external threats are not the only issue to consider. AI models used in engineering and manufacturing can perform well on average while still behaving unpredictably in specific cases. In many production settings, being "mostly right" is not sufficient; reliability and robustness set the baseline for acceptance. This places emphasis on testing AI systems across realistic production scenarios, including edge cases, data drift, and changes in process conditions. Manufacturers must consider whether an AI model is accurate, and also, what it means for that system to be "good" in its specific operational context. This should include how it handles uncertainty, degrades gracefully, and signals when human intervention is required.

Balancing investment and innovation

Scaling AI technologies requires significant upfront investments in infrastructure, talent, and ongoing maintenance; however, these investments must be made with the understanding that this technology is still in a period of rapid change and will require maintenance over time. Organisations have to balance short term value with strategic business objectives, ensuring that the long-term strategic advantage that AI integration offers is delivered. The focus should be on targeting use cases driven by clear business needs to move the organisation towards its longer-term goal.

While AI can improve effectiveness and decision quality, it should not be assumed to deliver immediate cost reduction. For aerospace engineering organisations, AI's value will likely be realised through improved quality, better utilisation of assets and engineering effort, reduced rework and downtime, and more effective planning – rather than through direct short term cost savings. Over time, progressively expanding a portfolio of AI capabilities aligned to evolving business needs, data maturity, and acceptable risk will require investment in ongoing oversight, potentially creating a new operational function.

Sustainability

The aerospace industry has a stretching decarbonisation target, and sustainability implications need to be considered in the context of how AI systems are selected, deployed, and used. Some AI techniques, particularly large, general purpose models, can be energy intensive during both training and operation. Organisations should therefore assess AI use cases not only by productivity or performance gains, but also by the compute, energy and lifecycle impacts associated with deployment.

Many domain-specific applications rely on smaller, more specialised models that operate within constrained environments and require significantly less compute than general-purpose systems. Commercial pressures, alongside regulatory and sustainability requirements, are already shaping how such models are designed and deployed. AI also has the potential to contribute directly to aerospace sustainability objectives by driving efficiency through improved design exploration, optimising flight paths and smarter airport and airspace operations. The potential to reduce emissions at a systems level could deliver a net benefit across the full operational lifecycle.

Managing the risk of over-reliance

While intelligent systems can significantly boost operational efficiency, unchecked over-reliance on AI systems risks eroding human expertise and weakening our ability to sustain a skilled workforce. Despite pressures for short-term gains, human judgement and intervention must remain central to ensure safety, reliability, and ethical integrity as AI is adopted. In other sectors, over-reliance has caused incidents, such as biased hiring, self-driving car crashes, and incorrect medical diagnoses, highlighting the importance of rigorous oversight throughout the lifecycle of the tools. Success depends on iterative capability growth, executing design processes with precision, selecting the right AI tool for the right problem and providing sufficient oversight and transparency to avoid costly and unintended consequences.

What does the future hold?

Across the near and medium term, there are numerous applications of AI in aerospace that hold the potential to drive value and help address some of the most critical challenges and opportunities facing the sector. Beyond these, understanding future capability helps organisations anticipate their impact and incorporate a long-term perspective into near- and mid-term strategies. Among the most transformative developments will be the integration of robotics and AI, and a fundamental shift in engineering approaches, from traditional, linear, single-company processes to more iterative, collaborative, and systems-based approaches.

Industrial robotics

Industrial robotics is scaling rapidly worldwide, with China alone installing around 300,000 industrial robots per year, almost twice as many as the following four countries combined. In aerospace manufacturing, robotics and AI are already present in areas such as drilling and fastening, composite layup, and inspection, but deployment remains fragmented and often constrained to automating pilot lines or high volume, repeatable work. Over time, greater adoption of AI enabled robotics, especially systems that can perceive, adapt, coordinate and dynamically adjust, is likely to become critical to competitiveness.

This provides the building blocks for more autonomous production environments, from highly automated cells all the way to “lights out” factories in which lines can run with minimal human presence for extended periods. In aerospace, full lights out operation across an entire final assembly or central structure line may be a longer term prospect, given safety, quality, and configuration complexity requirements. Even so, more AI-driven scheduling, anomaly detection, in-line inspection, and self-optimising processes, with humans increasingly focused on oversight, certification, and high-value problem-solving rather than repetitive shop-floor tasks, could offer significant opportunities for the sector.

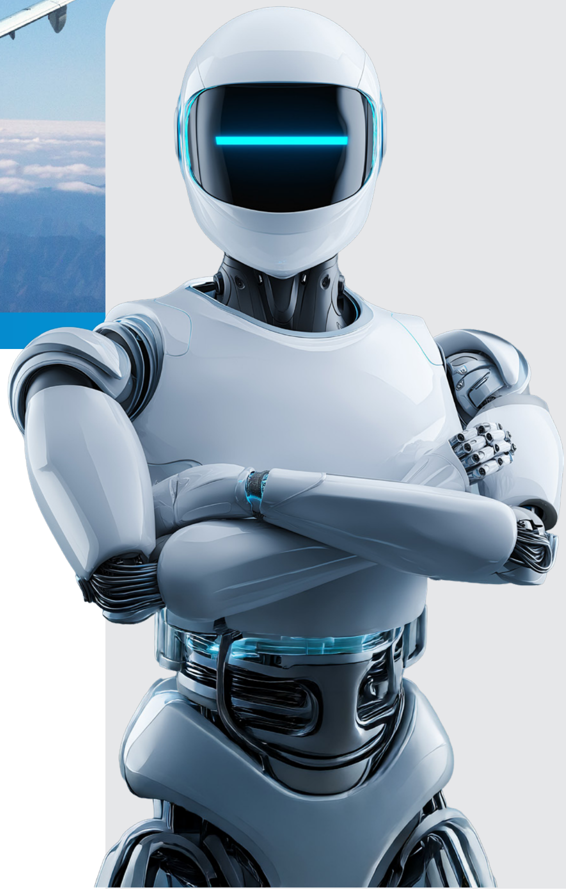

Humanoids

Human-like robots are already in development, with many approaching commercial viability as a response to demographic workforce challenges across sectors. Currently, most humanoids can reliably perform only a few specific tasks, with a limited ability to handle significant task variation or unexpected scenarios. As they evolve, these platforms exemplify how sophisticated hybrid AI technologies can be layered, with a physical body acting as a container for multiple AI systems that coordinate with real time control and hardware to execute dexterous, human like tasks.

As these sophisticated systems start fusing the whole range of AI capabilities discussed earlier, they will develop into reliable systems with the potential to work in human environments, create real-world impact and seamlessly collaborate, with humans and with each other. While robotics is still at a stage where it alternates between impressive demonstration and visible failure, the technology is rapidly evolving and several manufacturers have intentions to move towards larger scale deployment.

Autonomous optimisation of supply chains

Agentic AI solutions have the potential to radically reinvent how supply chains operate. Traditionally linear, today's supply chain is slow to converge and prone to fragility. Emerging AI supply platforms, such as Capgemini's ReFlow, use multi-agent ideas to monitor supplier performance, geopolitical risks, and logistics data to classify and respond to disruptions. While this is not yet a fully autonomous supply chain, this augmented AI platform creates a more proactive, adaptive supply chain that balances automation with expert oversight.

In future, the whole supply network can be considered as a single system, in which agents optimise their local goals as well as the global objectives for the health of the whole network, and where decision cycles could be accelerated. This transforms supply from a slow, manual process to a more complex living entity built as much upon mathematics, game theory and economics as on AI and technology.

Closing perspective

AI offers significant potential to support transformation across the aerospace sector, driving efficiency and safeguarding competitive advantage. Yet it may still seem too risky, with perception swinging between hype and scepticism. Successful AI adoption will require discipline: choosing the right problems, building the right foundations and maintaining human accountability. The fundamental reality remains that advanced digital capabilities are coming, and fast. The key is not simply to invest in the right technologies, but to build the organisational and digital foundations needed to deploy them responsibly. In safety-critical environments, data strategy, cybersecurity and reliability must remain central to AI deployment.

Preparing for AI involves more than technology choices. For organisations at an early stage of adoption, the starting point is not choosing an AI tool but being clear about the problem you are trying to solve and the value you expect to create. Progress is likely to depend on practical, staged action: developing skills, improving data integrity, strengthening infrastructure, and building confidence through well-governed use cases. Organisations that take a deliberate approach will be better placed to capture AI's benefits and manage its risks. Thoughtful change management and a commitment to ongoing education will help protect against dependency risks, ensuring that human expertise and AI capabilities can scale together responsibly.

Next steps

AI use case canvas

To conduct more informed exploration of how your organisation can begin building towards AI adoption, the AI Use Case Canvas at Appendix A is designed to help you shape initial ideas for your own AI use case, identify what is needed, and outline the elements required to make a compelling case for internal investment.

More information

A number of commercial and non-commercial organisations across the UK offer education and training, application support, adoption frameworks, policy guidance, and testing and demonstration environments. Contact the Catapult Network, including the High Value Manufacturing, Digital, and Connected Places Catapults, the Alan Turing Institute, and Innovate UK, including the AI Skills Hub for more information. The Civil Aviation Authority can provide information about emerging regulation and certification standards and can support engagement on responsible AI adoption.

Appendix: AI Use Case Canvas

The AI Use Case Canvas is a practical tool to support early exploration of AI use cases within an organisation. It is designed to help teams move from an initial challenge or opportunity to a clearer, more structured view of the problem, the value case, the delivery requirements and the adoption considerations.
The canvas is intended for both technical and non-technical users, including engineers, project managers, operational teams and decision-makers. Used well, it can help frame AI as a practical capability within a wider engineering or business context, rather than as a technology-led intervention.

How to use the canvas

The canvas is designed to help teams answer five core questions:

- What problem are we trying to solve, and why does it matter?
- Why might AI be an appropriate part of the solution?
- What value could the use case create, and how would success be measured?
- What data, digital infrastructure, skills and governance would be required?
- What risks, constraints, validation requirements and human oversight need to be considered?

It is recommended to use the canvas in the order suggested through the numbering. Importantly, the canvas encourages a challenge-led rather than a technology-led approach.

It can be used during:

- Early discovery and ideation;
- Project framing and internal alignment;
- Project gate or investment review discussions;
- Implementation planning;
- Project assessment, lessons learned and closure.

Who should be involved

Aim to involve a small cross-functional group, typically around 4–10 people. This may include domain experts, engineers, operational users, product or programme representatives, data or digital specialists, assurance or safety representatives where relevant, and the person accountable for the outcome. A first draft can also be completed individually to clarify the initial thinking before wider discussion.

Working through the prompts helps teams build organisational buy-in by creating a shared view of how an AI use case could develop in practice. By linking the problem and value case to delivery, risk, governance and adoption considerations, the canvas helps teams align around an initial idea and define practical pathways for implementation.

AI use case canvas

A structured tool for framing AI opportunities in aerospace

Note: It is recommended to follow the numbering and start with the problem and value case, not the technology.

4. Users and stakeholders

Who will use or be affected by the system?

Who owns the outcome, and who needs confidence in the outputs ?

2. Current process (As-Is)

How is the task or decision managed today?

Where are there challenges or limitations?

Consider bottlenecks, sources of variation, rework, data gaps or knowledge dependencies.

3. Future state (To-Be)

What would a better process look like?

What should remain under human control, and where could AI provide support?

1. Value proposition

What problem are we trying to solve and why does it matter?

What value could AI help create?

Consider, for example, improved quality, consistency, productivity, resilience, decision support or knowledge retention. Outline the measurable benefits expected.

8. Delivery team and capabilities

Define the required delivery capabilities, for example, domain, engineering, data, digital, assurance, safety, cybersecurity and change-management capabilities. Highlight areas that may require capacity building or upskilling.

9. Implementation pathway

How could the use case be tested, validated and scaled?

What would be the first bounded pilot or proof point?

5. Constraints , risks and assurance considerations

What safety, regulatory, technical, operational, data, cybersecurity or ethical constraints apply? How will risks be managed?

6. Relevant precedents and lessons

Have similar use cases been tried internally, in aerospace or in adjacent sectors?

What evidence, standards or lessons can be drawn from them?

7. AI role and candidate approaches

What role would AI play: classification, optimisation, prediction, generation, decision support, automation or agentic coordination?

Why is AI appropriate compared with other digital or process improvements?

About ATI

The Aerospace Technology Institute (ATI) is an independent not-for-profit organisation working in partnership with the UK government and industry to enhance the competitiveness of the UK aerospace sector through technology leadership and funding. Since 2014, the ATI Programme has invested over £4.2 billion of joint government-industry funding in a balanced portfolio of research and technology projects. The ATI Programme is delivered in partnership with the Department for Business and Trade and Innovate UK.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real.
www.capgemini.com

Capgemini x AEROSPACE TECHNOLOGY INSTITUTE