

Automatic Control Engineering

A Strategic National Capability for the UK's
Prosperity, Security and Sustainable Growth

A Position Paper of the UK ACE Network

April 2026

This position paper draws on activities of the ACE Network funded by the Engineering and Physical Sciences Research Council (EPSRC). The views expressed are those of the authors and do not necessarily reflect those of EPSRC.

Executive Summary

The United Kingdom's prosperity, security and long-term competitiveness increasingly depend on its ability to build, govern and operate complex automated systems. Automatic Control Engineering (ACE) is the discipline that makes this possible.

ACE provides the mathematical, computational and systems-engineering foundations that turn data and computation into reliable physical action. It is the hidden infrastructure and technology on which modern society depends.

The UK has national strategies for artificial intelligence, quantum technologies and engineering biology. Yet the engineering discipline that enables AI and automation to operate safely in the physical world, ACE, currently has no coordinated national programme.

As the UK becomes increasingly reliant on AI-enabled, interconnected and automated systems, the role of ACE is becoming more critical. **AI models alone cannot guarantee safe behaviour under novel or uncertain conditions.** Learning-based components must operate within stable, verifiable and constrained architectures. ACE provides the tools that make this possible, ensuring stability, fault detection, safety enforcement and system-level certification. Without strong control capability, AI cannot be deployed at scale in regulated sectors.

The UK's Automation and Productivity Challenge

Control engineering is embedded across sectors that together contribute over £500 billion in gross value added to the UK economy - including energy, manufacturing, transport, healthcare, telecommunications and defence - yet it receives no dedicated strategic investment at national level.

The UK has approximately 111 robots per 10,000 manufacturing workers, (IFR, 2024) compared with 429 in Germany, 419 in Japan, 295 in US and 1,012 in South Korea.

UK labour productivity has grown at just 0.4% per year over the past decade (The Productivity Institute, 2024), less than a third of the G7 average. Multiple analyses attribute a significant proportion of this shortfall to comparatively low investment in automation, digitalisation and advanced process control (OECD, 2024; McKinsey, 2019).

The Made Smarter Review identified £183.6 billion of value at stake from the application of automation and robotics within UK industry. ACE provides the engineering foundation that unlocks this value.

Together, these pressures shape the priority domains where ACE capability is now strategically essential for the UK, leading directly to the five missions set out below.

Five Mission Domains Where ACE Has National Impact

This Position Paper identifies five mission domains where strengthened UK capability in ACE would deliver the greatest national impact:

- **Trusted Intelligent Autonomous Systems**, ensuring AI-enabled systems operate safely, predictably and are certifiable in real-world environments.
- **Clean and Secure Growth**, deploying advanced automation and control to deliver net-zero energy systems, boost manufacturing productivity and improve supply-chain resilience.

- **Health, Life Sciences and Ageing**, supporting adaptive medical technologies, efficient bioprocessing, robotic surgery and assistive systems operating safely under uncertainty.
- **Resilient Mobility and Infrastructure**, providing the real-time optimisation, coordination and security needed for future transport and interconnected national infrastructure.
- **Defence, Space and Sovereign Capability**, embedding robust, intelligent control within defence platforms, critical infrastructure and strategically important technologies.

What the UK Risks Without Investment

Without strengthened capability in Automatic Control Engineering (ACE), the UK faces three strategic risks. First, continued productivity stagnation driven by slow adoption of advanced automation across industry and infrastructure. Second, limited ability to deploy AI safely in regulated and safety-critical sectors. Third, growing dependence on external technologies, standards and suppliers, reducing resilience across energy, transport and digital infrastructure and weakening sovereign capability in strategically important technologies.

Three National Actions Required

This paper calls for three coordinated national actions, supported by enabling priorities:

1. **ACE Network+: Community coordination and roadmap delivery**

Transition the existing community activity to an ACE Network+ acting as the national coordination body, providing community coordination, seed funding, doctoral training support and organisational capacity for the 2025–2035 Roadmap. Indicative investment: **~£3M** over three years.

2. **UK Centre of Excellence for Future Automation and Control Engineering: National infrastructure and technical leadership**

Create a federated Centre providing shared infrastructure and technical leadership for advanced control engineering research, system-level verification and certification, and the safe deployment of autonomous and AI-enabled systems in regulated sectors. The Centre would operate national-scale cyber-physical testbeds, develop next-generation control and optimisation methods, support regulators and standards bodies, and coordinate national workforce development in automation and control engineering. Indicative investment: **~£20M** over 7 years.

3. **UK National Programme in Future Automation and Control Engineering: Long-term, mission-driven investment**

Position ACE as a recognised strategic capability within UKRI's mission-aligned programmes. Fund competitive research, industrial demonstrators and cross-sector pilots across Digital and Technologies, Clean Energy, Life Sciences and Defence, the sectors where control engineering is critical to delivery. Indicative investment: **~£20M** per annum over ten years.

Together, these actions would establish the institutional infrastructure, community coordination and strategic investment required for the UK to develop world-leading capability in automation and control engineering.

Enabling Priorities

- **Invest in national ACE-aligned testbeds and regulatory sandboxes** for energy systems, healthcare technologies, autonomous mobility and cyber-physical infrastructure security, enabling the safe development and validation of advanced automated systems.
- **Build a sustained national skills pipeline in automation and control engineering**, including modernised university curricula, doctoral training, technician education, continuing professional development (CPD) for industry, and broader understanding of control engineering principles across all engineering disciplines.
- **Increase UK participation and leadership in international standards bodies** (ISO, IEC, IEEE, UNECE) for autonomy and cyber-physical systems.
- **Accelerate industrial adoption of automation and control technologies** through Catapult partnerships, translational engineering teams, Robotics Adoption Hubs and early-stage demonstrators, drawing on proven international models.
- **Improve national supply-chain resilience** for sensors, actuators, control hardware and power electronics, aligned with the UK Semiconductor Strategy.

A coordinated national effort in automation and control will accelerate progress towards net zero and enable safe AI-enabled autonomy. It will also strengthen UK life sciences, improve infrastructure resilience and boost manufacturing productivity. Together, these advances will enhance sovereign capability in defence and critical technologies. Investment is needed now.

1. Introduction: Why Automatic Control Engineering Matters

When a SpaceX booster lands itself on a barge in the Atlantic, it is not a pilot bringing it down. Control algorithms are computing thrust, orientation and position hundreds of times a second, adjusting for wind, fuel slosh and engine variance, all the way to touchdown. When you fly through a storm and barely feel it, a fly-by-wire system is correcting the aircraft's attitude thousands of times a second. When a surgeon operates through a da Vinci robot, control engineering filters the tremor from their hands and scales their movements down to fractions of a millimetre. In Japan, FANUC runs factories with the lights switched off, in effect robots building robots for weeks at a time, governed entirely by control systems.

All of these depend on the same engineering discipline: automatic control.

Control engineering is concerned with making physical systems do what they are supposed to do, reliably, safely, in real time, even when conditions are uncertain. It combines sensing, mathematical modelling, computation and actuation: measure what is happening, work out what to do about it, act, and repeat.

The discipline is not new, but the demands on it are growing fast. Energy grids with heavy renewable penetration are harder to stabilise than traditional thermal systems. Vehicles, surgical robots and defence platforms are all being asked to act autonomously in situations where the consequences of failure are serious. Manufacturing is mixing physical machinery with AI in ways that create new control problems. These systems are more complex, more interconnected and more safety-critical than anything control engineers dealt with a generation ago. The UK needs the capability to build and govern them.

That is what automatic control engineering does, and why it must be treated as a national capability.

The UK has genuine research strength in this field, but that capability is fragmented across universities and sectors and largely absent from national technology policy.

1.1 Why Control Matters Even More in the Age of AI

AI can tell you what is likely to happen. It cannot by itself guarantee what will happen next. A neural network might recognise a pedestrian, but that is not the same as ensuring the car stops in time. It might predict a turbine fault, but keeping the grid stable while the fault is handled is a different kind of problem entirely, one that requires control engineering.

That gap, **between prediction and safe physical action**, is what control engineering addresses. Control algorithms enforce safety limits, maintain stability when conditions change and ensure that systems behave predictably even when the AI component encounters something it has not seen before.

This distinction matters for regulation. In transport, energy, healthcare and defence, regulators require evidence that a system will behave safely under defined conditions. That evidence comes from control engineering: stability analysis, robustness guarantees, formal verification. Without it, AI-enabled systems cannot be certified for use in any domain where failure has serious consequences.

The UK has positioned itself as a leader in AI safety through the *Bletchley Declaration* (HM Government, 2023b) and the **AI Safety Institute**. But AI safety in physical systems is not purely an AI problem. It requires combining advances in artificial intelligence with the methods of automatic control engineering and systems theory.

There is growing investment in AI systems designed to learn models of the physical world from sensor data. Building mathematical models of physical systems from experimental data has been central to engineering and science for centuries. Within control engineering, system identification has developed a rigorous theoretical and practical framework for this. It ensures that models respect physical structure and are suitable for use within closed-loop control systems.

Learned models that predict how a system will behave are not the same as control systems that ensure it behaves safely. As AI moves from generating text to governing physical systems, the integration of learning with principled control architectures becomes essential.

Recent Global Investment in AI for Physical Systems

Company	Focus	Funding	Year	HQ
AMI Labs (Y. LeCun)	World models for physical systems	\$1.03bn	2026	France
World Labs (F-F. Li)	Spatial intelligence, 3D world generation	\$1.23bn*	2024-26	US
Physical Intelligence	Foundation models for robot control	\$1.1bn	2024–25	US
Figure AI	Humanoid robotics	\$675m	2024	US
Skild AI	General-purpose robot intelligence	\$300m	2024	US
Waymo (Alphabet)	Autonomous driving	\$27.1bn†	2020-26	US

*Table 1. Selected private investment in autonomous and AI-enabled physical systems, 2020–2026. Sources: WIRED, Bloomberg, TechCrunch, company announcements. Figures as reported at the time of funding. *Cumulative across two rounds (\$230m, 2024; \$1bn, 2026). † Cumulative total funding 2020–2026.*

Current approaches to AI safety and assurance often focus on the behaviour of models, whether a system produces reliable outputs under testing, verification or adversarial evaluation. However, when AI systems operate in the physical world, safety depends on the behaviour of the entire closed-loop system.

Once decisions are executed through sensors, actuators and physical processes, the outcome is determined by system dynamics, actuator limits, sensing uncertainty and environmental disturbances. A model that performs well in isolation may still lead to unsafe behaviour if these physical factors are not properly accounted for.

Ensuring safe autonomy therefore requires guarantees at the system level: stability, robustness, constraint satisfaction and predictable behaviour under disturbance. These questions sit at the core of automatic control engineering.

For autonomous systems operating in energy networks, mobility platforms, medical devices and defence systems, assurance must therefore integrate advances in AI with the control architectures that govern physical behaviour.

There is also a sovereign dimension. Trust in critical systems cannot be outsourced. For systems operating in defence, critical infrastructure and regulated sectors, the ability to independently verify, certify and assure the complete control architecture is not a service that can be safely procured from others. It is a strategic capability that must be maintained domestically.

Case Study: When AI Meets the Physical World — Lessons from Autonomous Vehicle Safety

In March 2018, a woman crossing a road in Tempe, Arizona was struck and killed by an Uber test vehicle operating in autonomous mode. The car's software detected her six seconds before impact but kept changing its mind about what she was - a person, a vehicle, a bicycle - and each time it reclassified her, it lost track of where she was heading. The factory emergency braking had been switched off. The safety driver was not watching. By the time the system committed to a response, it was too late.

What was missing was a system-level control architecture that could act on uncertain but safety-critical information regardless of classification. A control engineer would recognise this as a problem of robust decision-making under uncertainty, triggering a safe response when any plausible interpretation of the sensor data indicates danger, rather than waiting for a confident classification that never came.

The investigation found no independent safety layer capable of overriding the autonomy software. The system lacked the kind of supervisory control architecture, including stability constraints, safety envelopes, fall back behaviours, which control engineering provides for safety-critical systems in aviation, energy and defence. AI safety in physical systems requires more than model accuracy. It requires control architectures that guarantee safe behaviour even when the AI is uncertain or wrong.

In response to such incidents, ISO 21448 (SOTIF), released in 2022, was developed to address “unknown-unsafe” scenarios where systems behave as designed but still lead to unsafe outcomes. SOTIF marks a shift for control engineering. The problem is no longer just control of dynamics, but control of decisions under uncertain interpretation of the environment.

1.2 The Economic Value of Control Engineering

Automatic control engineering is deeply embedded in the UK economy, yet its contribution has never been explicitly recognised or quantified at a national level. A conservative estimate, based on the sectors in which control systems are critical operational infrastructure — energy (over £100bn+ GVA), manufacturing (over £190bn+ GVA), transport (over £70bn+ GVA), healthcare (over £170bn+ GVA), telecoms (over £40bn+ GVA) and defence (over £25bn+ GVA) — indicates that ACE-dependent systems underpin economic activity exceeding £500 billion in gross value added.

Even attributing a modest proportion, for example 10–20%, of sectoral productivity, operational efficiency and safety performance to the control systems that regulate them implies an ACE contribution of the order of £50–100 billion annually. This figure is consistent with international estimates: the US National Science Foundation's analysis of cyber-physical systems and the German government's assessments of Industrie 4.0 both attribute comparable proportions of industrial GVA to automation and control.

The Made Smarter Review (2017) identified **£183.6 billion of value at stake** from the application of automation and robotics within UK industry. ACE provides the engineering foundations that unlock this value. Without strong ACE capability, the UK cannot realise the productivity and competitiveness gains that automation promises.

1.3 Structural Challenges

The countries the UK competes with are not standing still (European Commission, 2025). The US, Germany, Japan, South Korea and China are all running sustained national programmes in industrial automation and autonomous systems. These are not short-term programmes. They are long-term national programmes designed to build industrial strength and reduce reliance on others.

The UK has strong research in control engineering, but its factories, infrastructure and industrial systems are far less automated than those of the countries it competes with.

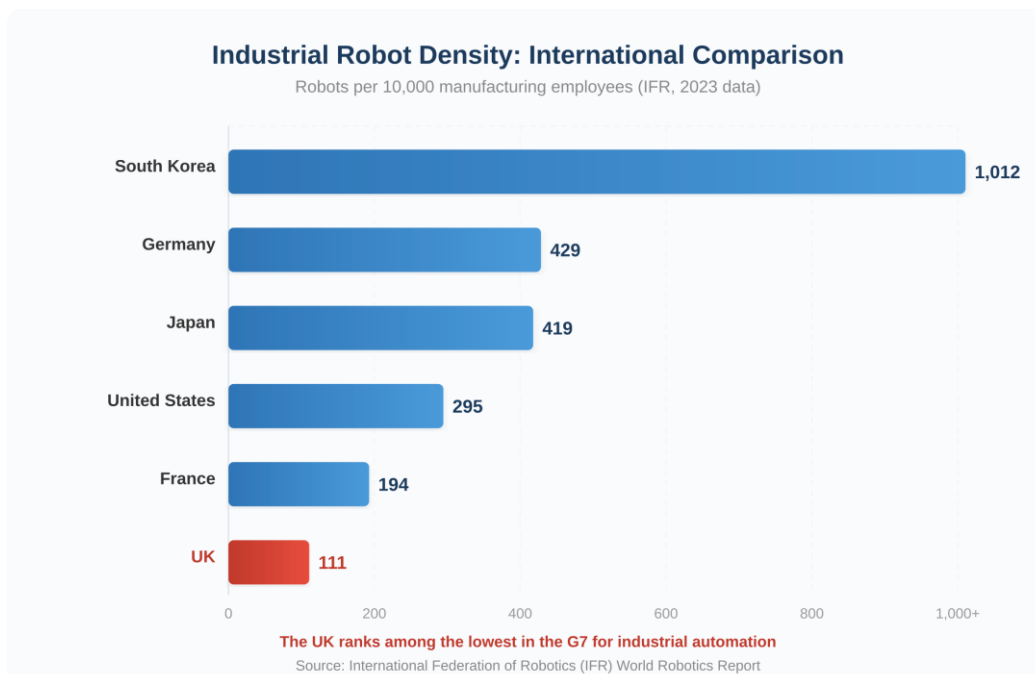


Figure 1. Industrial robot density (robots per 10,000 manufacturing employees), selected advanced economies. Source: IFR, World Robotics Report.

- **Slow adoption of automation.** The UK has just 111 industrial robots per 10,000 manufacturing workers. Germany has 429, Japan 419, South Korea over 1,000 (IFR, 2024). But the gap is not just about buying robots. Making automation work reliably, productively and safely, requires the control engineering, systems integration and process optimisation that sits behind every automated line. That is where the UK's deeper deficit lies.
- **Fragmented expertise.** There are world-class control engineering groups at universities across the country, but they operate largely independently. Unlike AI, quantum or even robotics, there is no national programme, no shared infrastructure and no coordinated investment bringing them together.
- **Control engineers are in short supply.** Industry struggles to recruit enough specialists. The wider engineering workforce, mechanical, electrical, civil, biomedical, often lacks

basic understanding of control principles. Systems get commissioned that underperform, cost more to maintain, or never deliver the productivity gains they promised.

- **Key supply chains are foreign.** Most sensors, actuators, control hardware and semiconductors are imported. In energy, defence and telecoms, that is a strategic exposure.
- **National-scale validation infrastructure remains limited.** The UK has valuable test and demonstration facilities across transport, energy and robotics, but they tend to be sector-specific and project-funded. Competitor economies invest differently. Germany's Fraunhofer institutes and Mittelstand-Digital centres provide sustained federal support for industrial automation. The US runs Manufacturing USA institutes and dedicated DoD autonomy test ranges. China designates national intelligent manufacturing pilot zones with long-term funding. The UK has no equivalent that is nationally coordinated, cross-sector and built to last.

2. Where Control Engineering Meets National Policy

Every critical technology identified in the DSIT Science and Technology Framework (HM Government, 2021), including AI, engineering biology, quantum technologies, telecommunications and semiconductors, depends on automation and control to operate safely and at scale.

Every UKRI mission priority, Engineering Net Zero, AI and Digitalisation, Transforming Health and Quantum Technologies (UKRI, 2022), also relies on control engineering to deliver. Yet control engineering is not explicitly recognised as a strategic capability in these frameworks. This paper argues that it should be.

2.1 Net Zero and Energy Transformation

The UK's commitment to net zero by 2050 (DESNZ, 2024) requires unprecedented system transformation. As fossil-based infrastructure gives way to distributed, variable and low-inertia systems, control becomes the determining factor in stability, affordability and resilience (CBI Economics and ECIU, 2025). With renewable generation exceeding 40% of UK electricity supply (National Energy System Operator (NESO), 2024), system inertia has fallen substantially, increasing volatility and tightening stability margins.

Multi-vector energy systems, linking electricity, heat, hydrogen, mobility and storage, require real-time optimisation and advanced control to operate safely and efficiently. Industrial decarbonisation depends on energy and emissions optimisation, integration of carbon-capture systems and process intensification. Evidence from sectors such as wastewater treatment and process industries shows that improved control strategies can reduce energy use by 5–20% (Lamnabhi-Lagarrigue et al., 2017).

Case Study: Keeping the Lights On — Control Engineering and the National Grid

The national electricity grid operates at a frequency of 50Hz, held there by continuously matching generation to demand. If frequency drops too far, protection systems start disconnecting load. Factories lose power. If it cascades, hospitals do too.

Traditional coal and gas plants provided natural inertia; renewables do not. When generation drops or demand spikes, there is less to cushion the shock.

As the UK moves toward a renewables-dominated grid, control engineering fills the gap. Advanced algorithms now manage grid-forming inverters, battery storage and millions of distributed energy assets in real time, making the split-second decisions that keep 67 million people's electricity stable. Without these control systems, a renewables-dominated grid cannot function.

2.2 AI Safety

The UK has positioned itself as a prominent international actor in responsible AI through the *Bletchley Declaration* and the AI Safety Institute (HM Government, 2023b). However, AI safety cannot be achieved without system-level safety mechanisms rooted in control engineering. When an AI model operates a vehicle, a surgical robot, a power grid or defence platforms, it must do so under uncertainty, within safety constraints and subject to external disturbances. Control engineering provides the stability guarantees, safety constraints and certifiable architectures that make this possible. Regulators increasingly require evidence of bounded system-level risk and

demonstrably safe fallback behaviour. That evidence comes from control engineering, not from the AI model itself.

2.3 Health and Life Sciences

Many modern medical technologies, including insulin pumps, ventilators, neurostimulators and robotic surgical systems are fundamentally control systems. Their safety and regulatory approval depend on rigorous control design and verification. Robotic surgery is growing rapidly across the NHS, and every platform requires precise, stable, fail-safe control of multi-axis manipulators operating millimetres from living tissue.

Beyond the bedside, biomanufacturing for vaccines, biologics and advanced therapies depends on feedback control and model-based optimisation. Uncontrolled variability can account for up to 60% of failure-related costs in cell and gene therapy manufacturing; closed-loop control can reduce that variability by 30–50% (Schaefer, G. et al., 2023).

Healthcare delivery itself including patient flow, workforce allocation, diagnostics bottlenecks and elective-care scheduling, benefits from control and optimisation approaches aligned with NHS Long Term Plan goals (Office for Life Sciences, 2021).

Case Study: The Artificial Pancreas — A Control System That Saves Lives

For millions of people with Type 1 diabetes, managing blood sugar is relentless and dangerous. Every meal, every exercise session, every night's sleep carries risk. The artificial pancreas changes this. A sensor measures glucose continuously. A control algorithm predicts what will happen next and computes the precise insulin dose. A pump delivers it automatically. The patient sleeps safely.

UK researchers and NHS clinicians have been at the forefront of developing and trialling these closed-loop systems, with clinical studies showing dramatic reductions in dangerous blood sugar episodes. This is control engineering at its most human, sensing, predicting, deciding and acting to keep someone alive.

The regulatory pathway to wider adoption depends directly on the ability to verify and certify the control algorithms that make treatment decisions. Similar approaches are now being developed for automated ventilator management, where digital twins of patient lung mechanics allow control algorithms to be trained and safety-tested in simulation before clinical deployment.

2.4 Defence and National Security

The 2025 Defence Industrial Strategy (MoD, 2025) puts autonomous platforms and secure infrastructure at the centre of UK defence. Every uncrewed aircraft, vessel and ground vehicle depends on control engineering for guidance, navigation and safe operation. As these systems become more autonomous, the ability to certify their behaviour becomes critical, and certification depends on control theory. In contested environments, the difference between platforms is increasingly the quality of the control algorithms inside them. Two vehicles with identical hardware but different control software will not perform the same. Better estimation, faster adaptation and stronger stability guarantees provide decisive operational advantages. **The countries that invest in control engineering research do not just build autonomous systems; they build better ones.**

The UK cannot maintain sovereignty in defence autonomy if it relies on imported expertise and foreign supply chains to build the control systems inside its own platforms.

2.5 Industrial Strategy and Productivity

Economic growth is the government's central mission (DBT, 2024), and productivity is its main constraint. UK labour productivity has grown at just 0.4% per year over the past decade, compared with 2.2% before 2008 and a G7 average of 1.2% (The Productivity Institute, 2024; OECD, 2024; Resolution Foundation, 2023). Multiple analyses attribute 25–40% of this shortfall to low investment in automation and advanced process control (OECD, 2024; McKinsey, 2019; The Productivity Institute, 2024). International competitors investing in these technologies report productivity gains of 15–30% (McKinsey, 2019).

The Industrial Strategy's Robotics Adoption Hubs are a welcome step. For them to deliver lasting impact, they must be underpinned by the control engineering, systems integration and optimisation expertise needed to make automation work on the factory floor, not just arrive there.

Case Study: Lights-Out Manufacturing — The Productivity Opportunity

In Oshino, Japan, FANUC runs factories around the clock with the lights off. Robots building robots for extended periods without human intervention. In Beijing, Xiaomi opened a highly automated “dark factory” in 2024 where AI-driven production systems assemble smartphones with very limited human involvement, producing roughly one device per second.

In Drachten, the Netherlands, Philips produces millions of electric shavers each year in one of the world's most highly automated consumer electronics factories, where robotics, machine vision and precision control coordinate much of the production process.

In Amberg, Germany, Siemens operates one of the world's best-known smart factories, producing industrial control components through tightly integrated automation, digital manufacturing systems and real-time process control.

3. Five Missions for UK Control Engineering

These challenges point to the need for a mission-driven framework that positions ACE as a strategic national capability. The five missions below define where the UK must build capability and form the backbone of the forthcoming 2025–2035 ACE Network Research and Innovation Roadmap.

3.1 Mission 1: Trusted Intelligent Autonomous Systems

Autonomous systems are no longer confined to laboratories. They are being deployed in logistics centres, transport networks, clinical settings and defence environments. Safe operation in these domains does not follow automatically from advances in AI. Perception and learning components must operate within closed-loop control architectures that enforce stability, respect safety constraints and manage uncertainty in real time. What ultimately matters is not model accuracy in isolation, but whether the overall system behaves predictably under disturbance and edge conditions.

Regulators therefore assess system-level behaviour, not just algorithmic performance, and require demonstrable assurance of safety. As reflected in the U.S. Department of Defense guidance on the developmental test and evaluation of autonomous systems, assurance must be demonstrated at the level of system behaviour under real-world conditions. (U.S. Department of Defense, 2023; US . Department of Defense, 2025).

Trusted autonomy requires advances in stability and robustness of learning-enabled control systems, verifiable autonomy architectures, human–machine interaction, simulation-to-reality transfer and real-time fault detection. The UK has international strengths in AI safety, formal verification, software engineering and robotics. Embedding ACE creates globally recognised leadership in safe, certifiable autonomy for manufacturing, mobility, nuclear robotics, defence, medical and assistive robotics, and infrastructure inspection.

Case Study: Landing a Rocket on a Barge — SpaceX and the Power of Control

When a SpaceX Falcon 9 booster lands itself on a floating platform in the ocean, what the world sees is a spectacle. What makes it possible is automatic control engineering. During descent, the vehicle's guidance, navigation and control system continuously estimates position, velocity and attitude using onboard sensors and state observers.

Hundreds of times per second, the flight controller computes the thrust and control inputs needed to stabilise the vehicle. These commands are then issued to the actuators in a closed loop, stabilising a 40-metre structure descending on a single engine through turbulent air. No pilot intervenes; the landing sequence is executed entirely by the onboard control system.

This is the same discipline, the same mathematics and engineering principles that the UK needs to deploy autonomous systems safely across manufacturing, transport, energy, defence and healthcare. The difference is not the science; it is the scale of national investment in making it happen.

3.2 Mission 2: Clean and Secure Growth

This mission addresses both energy system transformation and industrial productivity. Control engineering is central to both.

On the energy side, the transition to a renewables-dominated grid requires real-time balancing of supply and demand across increasingly complex and distributed assets, including battery

storage, flexible demand, hydrogen electrolysis and heat networks. As fossil generation is retired, the control systems that maintain grid stability become more important, not less. Similar challenges arise in coordinating multi-vector energy systems and in optimising industrial energy use alongside emissions reduction.

In manufacturing, control engineering enables production systems to maintain tighter tolerances with less waste, to detect and correct quality problems during production rather than after it, and to reconfigure between product runs with less manual intervention.

Case Study: Nuclear Fusion — Controlling a Star on Earth

Nuclear fusion requires the confinement of plasma at temperatures exceeding 150 million degrees Celsius using powerful magnetic fields. The plasma is inherently unstable. Small perturbations can grow rapidly, leading to disruptions that terminate the reaction or damage reactor components. Maintaining a stable plasma therefore depends on continuous measurement, prediction and control of its behaviour in real time.

The UK's STEP programme, led by the UK Atomic Energy Authority (UKAEA), aims to deliver a prototype fusion power plant by the early 2040s (Lennholm et al., 2024). Experiments on the MAST Upgrade spherical tokamak at Culham are exploring advanced plasma control techniques, including active magnetic control to detect and suppress instabilities as they develop.

Fusion reactors will rely on fast diagnostics, predictive modelling and high-bandwidth control of magnetic fields and plasma current. As emphasised in international programmes such as ITER (ITER, 2019), sustained fusion power depends not only on materials and physics but on the ability to stabilise plasma through sophisticated real-time control systems.

3.3 Mission 3: Health, Life Sciences and Ageing

This mission embeds advanced control in clinical technologies, biomanufacturing and healthcare operations.

It includes the development of closed-loop medical devices that adapt therapy to individual patients, the progression of robotic surgery towards greater autonomy under verified safety constraints, model-based control of biomanufacturing processes, and assistive robotics capable of safe interaction with patients. It also applies optimisation and control methods to NHS operations, including patient flow, theatre scheduling and workforce allocation.

The unifying engineering challenge is to deliver adaptive, safety-critical control in biological, human-facing and highly regulated environments.

Case Study: Deep Brain Stimulation — Closing the Loop in the Brain

Deep brain stimulation (DBS) delivers electrical pulses to specific regions of the brain to treat Parkinson's disease. Current devices are calibrated by a clinician in a clinic and run at fixed settings until the next visit, regardless of how the patient is doing.

Adaptive DBS changes this by reading neural signals that reflect symptoms that reflect symptoms like tremor or stiffness, to adjust the stimulation in real time. This “closed-loop” approach can improve symptom control, reduce side effects, and extend battery life.

But making this work is far from straightforward. The system has to respond safely and reliably to a brain that behaves differently from one patient to another, and even from moment to moment. The underlying dynamics are complex and not fully understood, and everything has to run on a tiny implanted device where failure is simply not an option.

3.4 Mission 4: Resilient Mobility and Infrastructure

This mission addresses the control engineering required for future transport, logistics, critical infrastructure and space systems.

It will develop distributed control architectures for coordinating complex infrastructure networks, predictive optimisation for transport systems, real-time detection of cascading failures, digital twins for infrastructure management and secure control systems resilient to cyber-physical attack.

Case Study: Railway Energy Hubs —Turning the Railway into a Virtual Power Plant

Network Rail consumes over 4 TWh of electricity annually, yet only 39% of the UK rail network is electrified. At the same time, wind curtailment reached 8.3 TWh in 2024, costing nearly £1 billion in wasted renewable generation.

Railway Energy Hubs address both problems. These are modular microgrids installed along rail corridors, combining battery storage, solar generation and advanced control systems. They supply traction power to the railway while providing services to the national grid — frequency response, peak shaving and inertia support. Electrified rolling stock alone can deliver up to 550 MW of inertia support during grid disturbances.

Demonstrator programmes funded through Ofgem's Strategic Innovation Fund, in partnership with Network Rail, SP Energy Networks, GE and Ricardo, are already validating the concept. By networking multiple hubs, the railway shifts from passive energy consumer to active contributor to national energy flexibility.

The mission would also develop the control foundations for next-generation mobility. It would enable autonomous vehicles to operate safely in mixed traffic with the certified guarantees regulators require. It would provide real-time coordination for drone corridors and urban air mobility. This is a distributed control problem that does not yet have a proven solution at scale. It would support autonomous maritime navigation in congested shipping lanes, and enable higher levels of railway automation, increasing both capacity and safety.

In space, the mission would strengthen UK capability in autonomous satellite operations, constellation management and debris avoidance. These are areas where the UK's commercial ambitions are growing faster than its control engineering capacity to support them.

Case Study: Fly-by-Wire and the Safety of Modern Aviation

Modern commercial aircraft are not flown directly by cables and mechanical linkages. A pilot's inputs are interpreted by onboard computers which translate them into control surface movements while simultaneously correcting for wind, turbulence and pressure changes, thousands of times per second. Much of the smoothness passengers experience, even when flying through rough weather, is the result of these control systems working quietly in the background.

Some aircraft are deliberately designed with reduced natural stability to improve fuel efficiency and agility. In those cases, it is the control system alone that makes them safe to fly. The UK has longstanding expertise in flight control. The next phase of aviation, including autonomous air mobility, uncrewed platforms and managed urban drone corridors, will require extending this capability to increasingly distributed, software-defined and certifiable control architectures.

3.5 Mission 5: Defence, Space and Sovereign Capability

This mission will strengthen the UK's sovereign capability in the control architectures that underpin defence platforms across air, maritime and land domains, as well as military space operations (IISS, 2025).

It will advance the control methods, algorithms and system architectures required for uncrewed defence platforms to execute missions independently in contested and communications-denied environments (Black et al., 2024). These systems must adapt to deception, interference and uncertainty while maintaining stability, safety and mission performance without continuous human oversight. The mission will reinforce UK capability in high-precision guidance, interception and coordinated multi-platform autonomy. In the space domain, it will establish domestic capability in satellite survivability, manoeuvre, attitude regulation and resilient constellation coordination.

The quality of the control engineering inside a platform, including estimation, adaptation and stability guarantees, is an increasingly decisive factor in defence capability.

A central component of this mission is the creation of sovereign UK capacity to verify, validate and certify autonomous defence and space systems at the level of physical system behaviour. This extends beyond software correctness to encompass stability guarantees, safety bounds and fault tolerance of the complete closed-loop system operating under adversarial conditions. Such assurance capability must be developed and sustained domestically (NATO STO, 2024). Dependence on externally developed certification tools or opaque assurance frameworks introduces strategic risk in systems where resilience and trust are non-negotiable.

Case Study: Defensive Interception

When a missile defence system detects an incoming threat, the interceptor system must estimate the target's trajectory from noisy sensor data, compute a guidance solution and continuously adjust its course as it closes in a matter of second. At these speeds, there is no role for human intervention. The system operates autonomously from launch to intercept.

Whether interception succeeds depends on the control system, the accuracy of its tracking, the speed of its response, its ability to maintain performance when sensor data is degraded or the target behaves unexpectedly. This is a control engineering problem operating at the limits of physical performance. The credibility of any defensive interception capability depends on the robustness and certifiability of the control architecture inside it.

Sovereign expertise in these systems is essential if the UK is to independently evaluate, certify and sustain its defensive platforms.

3.6 Interdependencies

The five missions are mutually reinforcing. Trusted autonomy depends on resilient infrastructure and sector-specific regulation. Clean growth requires autonomous inspection and adaptive robotics from Mission 1. Health innovation strengthens embodied intelligence and safety-critical control. Mobility relies on distributed optimisation and cyber-physical security. National security depends on advances in manufacturing, semiconductors, quantum and safe autonomy. Investment in ACE yields system-wide benefits, compounding progress across all domains. A mission-led approach requires coordinated capability building across these domains.

Figure 2 maps the interdependencies across the five missions. It shows that advances in autonomy, energy systems, health technologies, mobility and defence rely on shared control engineering foundations, reinforcing the need for an integrated national programme rather than sector specific investments.

Five ACE Missions: Interdependencies

Investment in ACE yields system-wide benefits across all national priorities

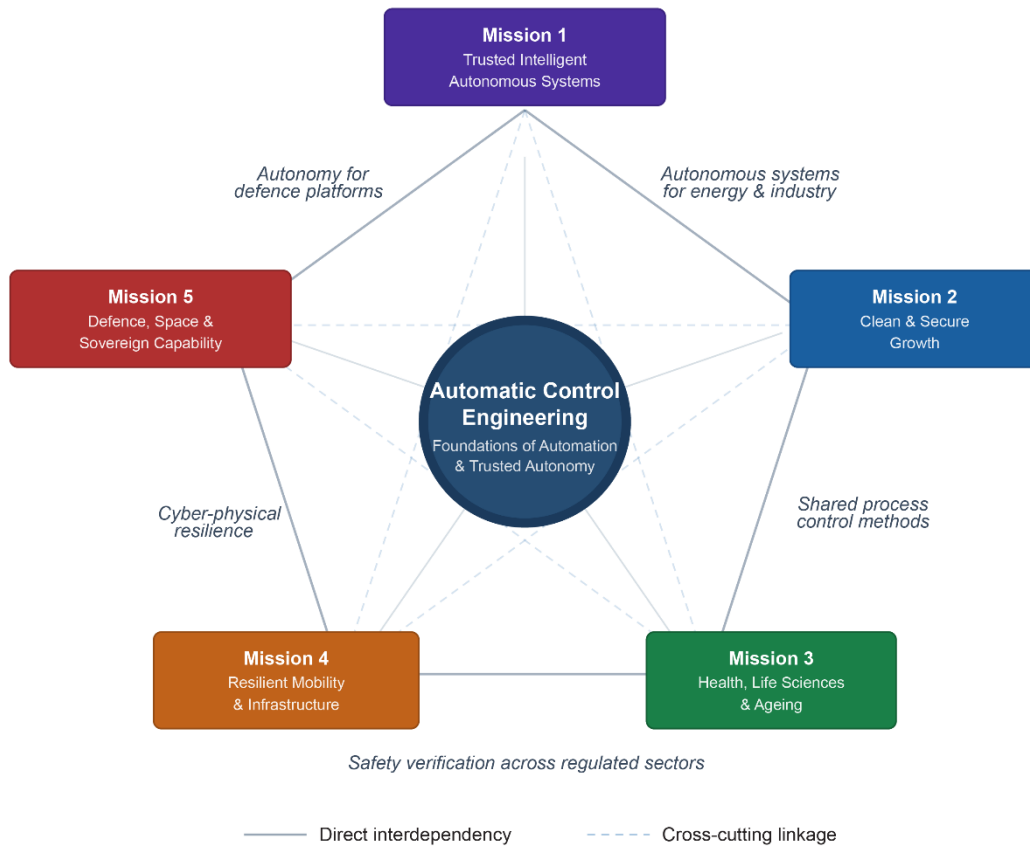


Figure 2: The five ACE missions and their interdependencies.

4. Technological Trends and Methodological Advances

Automatic Control Engineering is entering a new phase. Systems are becoming more interconnected, data-rich and software-defined. They combine physical dynamics with learning components, operate across cloud–edge hierarchies, and interact with uncertain, sometimes adversarial environments.

The foundational principles of the discipline remain essential (EPSRC, 2022, Lamnabhi-Lagarrigue et al., 2017). What is changing is the context in which they must operate. Control must now be reconciled with machine learning, distributed architectures, scalable optimisation, cyber-physical security and real-time computation at unprecedented scale.

4.1 Learning-Enabled Control

Learning is increasingly embedded within control loops, with reinforcement learning and adaptive policy optimisation deployed in real-time systems. However, learning components alone do not guarantee safe behaviour. This distinction is well recognised in the emerging literature on data-driven and learning-based control, where models learned from data must be embedded within control frameworks that ensure stability and constraint satisfaction (Brunton and Kutz, 2019; Dean et al., 2020).

For safety-critical autonomy, systems must satisfy stability under uncertainty, constraint enforcement, fault tolerance and certifiability. Hybrid architectures, where learning operates within provably safe supervisory structures, are most promising. Key challenges include safe closed-loop exploration, real-time uncertainty quantification, combined formal–statistical verification and scalable simulation for rare-event analysis. Reinforcement learning approaches, while powerful, face well-documented challenges in ensuring stability, safety and robustness in continuous control settings (Recht, 2019).

4.2 Distributed and Multi-Agent Control

Future infrastructure is inherently distributed: energy networks, logistics systems, mobility platforms, digital twins and robotic swarms. The core methodological challenge is to guarantee coherent global behaviour from locally informed decisions. These challenges have been studied extensively in the context of consensus and distributed optimisation in networked systems (Olfati-Saber et al., 2007; Nedić and Ozdaglar, 2009). Coordinating heterogeneous agents requires stable, scalable and robust systems despite partial information and dynamic network conditions

Addressing this requires advances in distributed model predictive control, consensus and agreement protocols, event-triggered coordination and optimisation over dynamic network graphs. Stability and constraint satisfaction must be maintained under topology changes, communication delays and packet loss. At the same time, performance must scale to thousands or millions of interacting nodes without reliance on fragile centralised coordination.

4.3 Cyber-Physical Security

Control systems that are networked and remotely accessible can be attacked through their data, their communications or their software. A compromised sensor reading or a manipulated control command can cause physical damage that no amount of IT security can reverse. Such vulnerabilities are well established in the study of cyber-physical systems, where attacks on

sensing and actuation can destabilise closed-loop behaviour (Amin et al., 2009; Pasqualetti et al., 2013). Resilient control requires detection of anomalies, diagnosis of attack type, isolation of compromised subsystems, reconfiguration of control laws and fail-operational modes. This depends on understanding closed-loop dynamics and how perturbations propagate through physical systems.

4.4 Digital Twins and Simulation

Digital twins are becoming indispensable for design, certification and operation. Their role in integrating data-driven models with physics-based simulation for control and optimisation is increasingly recognised in advanced manufacturing and cyber-physical systems (Tao et al., 2019). Their value arises in closed loop: real-time state estimation, predictive simulation, virtual commissioning and support for AI-in-the-loop certification. Major advances are needed in hybrid simulations combining physics and data-driven models, high dimensional real-time optimisation, rare event modelling and uncertainty quantification. A national closed-loop simulation and certification environment, similar to US DoD and German Industrie 4.0 platforms, would accelerate progress across all missions.

4.5 Emerging Hardware and Computation

Quantum systems, semiconductor manufacturing and soft/biological robotics all demand advanced control to manage fragile, nonlinear and hybrid behaviours. These domains are central to sovereign capability, aligning directly with DSIT priorities.

Embedded GPUs, edge/cloud hierarchies and neuromorphic processors enable fast optimisation and fleet-level coordination, but introduce certification challenges and reliance on international hardware supply chains.

4.6 Embodied Intelligence

Embodied intelligence concerns systems in which perception, learning and control are integrated directly with physical form. Performance no longer depends only on algorithms, but on the joint design of mechanical structure, sensing configuration and control architecture. This perspective aligns with the established concept of embodied intelligence, where physical morphology, sensing and control are co-designed to achieve robust behaviour (Pfeifer et al., 2012).

Control engineering enables this integration. It supports co-design of morphology and actuation, ensures that learned components preserve stability and safety, governs physical interaction with people and the environment, and allows systems to adapt in real time to changing conditions.

This integration is essential for advanced manufacturing robots, surgical and assistive devices, inspection and maintenance drones, space robotics and next-generation defence platforms. In each case, intelligence is not separate from the body of the system, it is embedded within it.

5. The UK ACE Ecosystem: Strengths and Gaps

Understanding the existing landscape is essential for designing interventions that complement rather than duplicate. The UK's ACE ecosystem has significant strengths but also critical gaps that the proposed national actions are designed to address.

5.1 Existing Strengths

- **World-class research:** Over 1,000 active researchers across 30+ research groups and centres at 20+ universities, with internationally recognised contributions to adaptive control, multivariable control, robust control, model predictive control, system identification, nonlinear systems and optimisation.
- **Professional community:** UKACC (IFAC National Member Organisation), IET Control and Automation Network, InstMC, IChemE and IMechE specialist sections provide professional coordination, events and standards engagement.
- **ACE Network:** Building on the wider UK community, over 500 researchers and 23 industry partners connected through the EPSRC-funded network, with demonstrated capacity for community coordination, roadmap development and feasibility studies.
- **Adjacent national programmes:** Alan Turing Institute (AI), UK-RAS Network (robotics), Quantum Technology Hubs, Engineering Biology Mission Hubs, Henry Royce Institute (materials), Catapult network.
- **Industrial base:** Major industrial players in sectors where automation and control are critical, including aerospace (Rolls-Royce, BAE Systems, Airbus), automotive (JLR), energy (National Grid, SSE, Siemens Energy), pharmaceuticals (GSK, AstraZeneca), defence (QinetiQ, MBDA, Leonardo) and technology (ARM, Dyson).

5.2 Critical Gaps

- **No national coordinating centre:** Unlike AI (Turing), quantum (NQCC) and materials (Royce), control engineering lacks a nationally coordinated capability providing leadership, integration across the community, and sustained support for translation, standards and regulatory engagement.
- **No dedicated national programme:** ACE has no equivalent of the National Quantum Technologies Programme or the AI Strategy's investment commitments.
- **Limited testbed infrastructure:** The UK has limited nationally coordinated test environments for validating autonomy, cyber-physical systems and large-scale infrastructure control.
- **Weak translational pathways:** The UK has strong research capability but weaker mechanisms for translating automation and control technologies into widespread industrial adoption, particularly among SMEs. International competitors operate structured programmes that embed automation expertise directly within industrial clusters.
- **Supply-chain vulnerability:** The UK remains highly dependent on imported sensors, actuators, control hardware and semiconductor components.
- **Skills pipeline constraints:** Insufficient specialist supply and inadequate ACE literacy across the broader engineering workforce.

6. National Actions

The five missions and the gaps identified in the previous section point to three coordinated national actions, supported by enabling priorities. The investment figures are indicative and reflect the scale of commitment required.

6.1 Action 1: ACE Network+

The ACE Network has connected more than 500 researchers and industry partners across 30+ research groups at over 20 universities.

To support the missions and Centre of Excellence, it must transition to a sustained ACE Network+ providing community coordination, horizon scanning, seed funding, doctoral training support, organisational capacity required to support the **2025–2035 ACE Network Research and Innovation Roadmap**, and the precursor governance platform for the Centre.

Indicative investment: ~£3M over three years.

6.2 Action 2: A UK Centre of Excellence for Future Automation and Control Engineering

The technological trends outlined above require capabilities that cannot be developed through fragmented research alone. They demand coordinated national infrastructure, system-level validation environments and sustained technical leadership.

A federated Centre of Excellence would provide the institutional anchor for advanced control engineering research, system-level verification and certification, and the safe deployment of autonomous systems in regulated sectors. It would operate national-scale cyber-physical testbeds, support regulators and standards bodies, and coordinate national workforce development.

The Centre would be hosted by a consortium of universities with satellite hubs aligned to mission domains, joint governance with UKRI, co-location with selected Catapults and an industry advisory board.

Indicative investment: ~£20M over seven years.

6.3 UK National Programme in Future Automation and Control Engineering

The UK invests in AI, quantum, engineering biology and semiconductors. There is no dedicated programme for the engineering that connects all of them to safe physical action.

Sustained national investment in ACE would fund mission-aligned competitive research across all five domains, large-scale demonstrator programmes, cross-sector pilots in energy, defence, mobility and health, and industrial co-investment mechanisms. The Centre of Excellence provides the institutional anchor. The National Programme provides the funding that flows through it and across the wider community. The US, Germany and China already operate comparable programmes.

Investment would be delivered through UKRI's strategic programmes in Digital and Technologies, Clean Energy, Life Sciences and Defence, coordinated with DSIT, DESNZ, DfT, MoD and DHSC.

Indicative investment: ~£20M per annum over ten years, delivered through UKRI's strategic programmes across the five mission domains.

Together, the ACE Network+, Centre of Excellence and UK National Programme in Future Automation and Control Engineering form a single integrated structure: the Network+ coordinates the community, the Centre provides national scale infrastructure and technical leadership, and the National Programme funds mission aligned research and innovation.

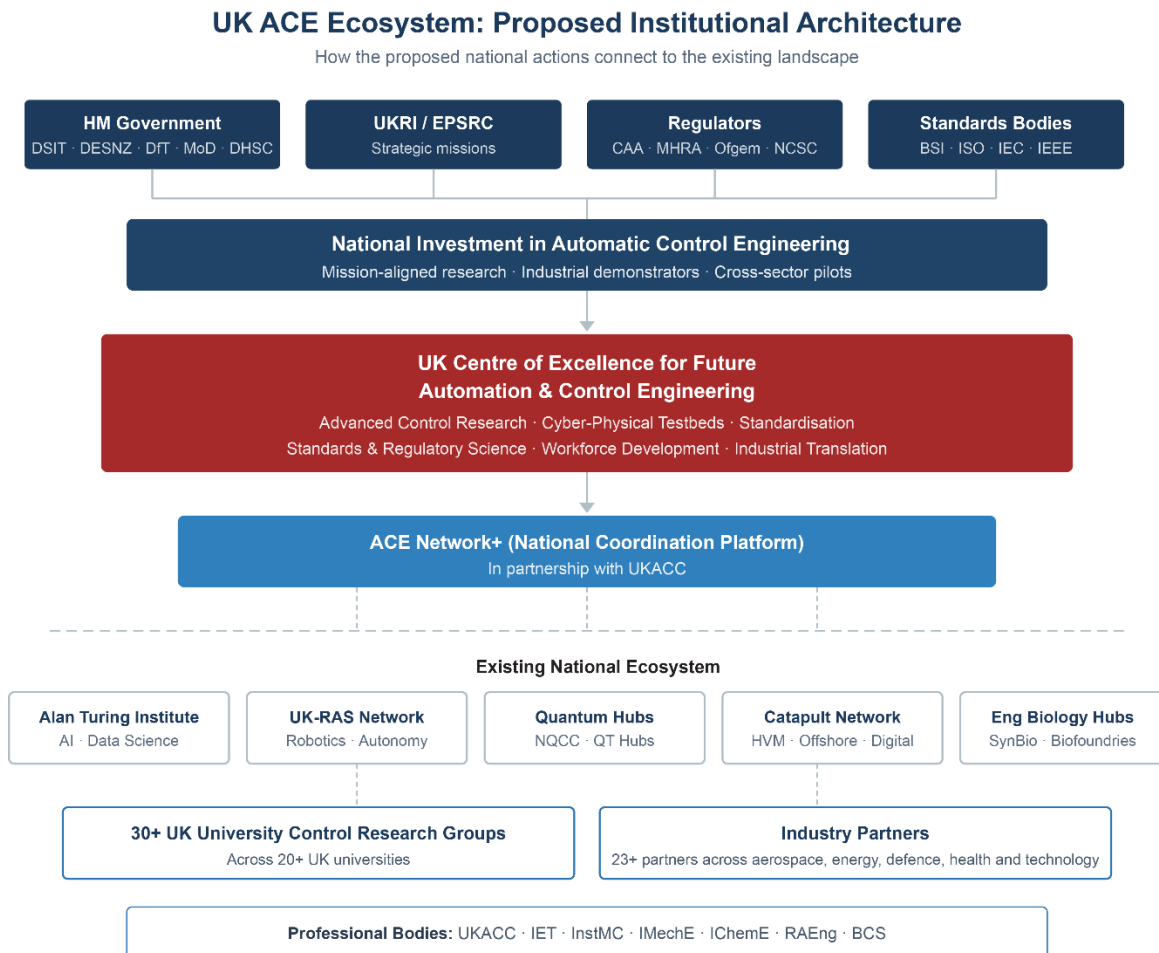


Figure 3: Proposed UK ACE institutional architecture.

6.4 Enabling Priorities

National Testbeds and Regulatory Sandboxes

The UK urgently needs integrated, **closed-loop test environments**. Priorities include autonomous mobility corridors (road, rail, maritime and air) with safety-critical Verification & Validation environments; whole-system energy simulators integrating electricity, heat, hydrogen and storage; healthcare robotics and digital therapeutics labs; bioprocess and engineering-biology pilot plants with real-time optimisation; cyber-physical security demonstrators for infrastructure resilience; and federated digital-twin platforms with real-time control interfaces.

Regulatory sandboxes co-designed with the Civil Aviation Authority, Driver and Vehicle Standards Agency, Medicines and Healthcare products Regulatory Agency, Office of Gas and Electricity Markets, Water Services Regulation Authority, National Cyber Security Centre and the AI Safety Institute would enable testing of AI-enabled control under monitored risk, new certification pathways and adaptive regulation aligned with innovation.

Skills, Workforce and National Capability

A national skills strategy for automation and control must address both **specialist depth and broader engineering literacy**. Beyond training more control engineers, the UK needs deeper understanding of control engineering principles across all engineering disciplines, mechanical, electrical, civil, chemical, biomedical, so that engineers across the economy can specify, commission and work effectively with advanced automation and control systems.

This requires modernising undergraduate control curricula; embedding AI-in-the-loop safety, optimisation and cyber-physical security as standard elements; creating new Centres for Doctoral Training aligned with missions; establishing CPD programmes for mid-career engineers; supporting technician training and apprenticeships; and strengthening diversity and inclusion. Work with professional bodies, including IET, RAEng, BCS, IMechE, InstMC, IChemE, will be essential for accreditation and reach.

Accelerating Industrial Adoption

Catapult partnerships, translational engineering teams and the Industrial Strategy's Robotics Adoption Hubs provide important foundations for industrial adoption. Embedding stronger control engineering capability within these mechanisms would increase their impact. International models worth considering:

- **Germany's Mittelstand-Digital centres** embed automation expertise directly within SME clusters through regional competence centres and model factories.
- **Japan's Robot Revolution Initiative** provides structured national support including shared testbeds and dedicated integration engineers who work alongside companies adopting automation for the first time.
- **Singapore's Advanced Manufacturing Centre of Excellence** offers open-access pilot lines where companies can test and validate advanced control technologies before committing to full-scale investment.

The UK should draw on these models to develop low-barrier demonstrator funding, automation-readiness diagnostics, shared simulation tools and open-access pilot environments.

International Standards

UK participation in ISO/IEC, IEEE, UNECE WP.29, IEC TC65 and ISO TC299 determines whether the UK shapes the rules for autonomy, automation and cyber-physical systems or follows rules set by others.

Supply Chain Resilience

Resilient domestic supply of sensors, actuators, control hardware and power electronics, aligned with the UK Semiconductor Strategy.

7. Conclusion

This paper has set out the case for recognising Automatic Control Engineering as a strategic national capability. It proposes three coordinated actions:

1. ACE Network+ as the immediate enabling step. Indicative investment: **~£3M** over three years.
2. UK Centre of Excellence for Future Automation and Control Engineering, with indicative investment **~£20M** over seven years.
3. UK National Programme in Future Automation Control Engineering, across UKRI strategic priorities with indicative investment **~£200M** over ten years.

The forthcoming 2025–2035 ACE Network Research and Innovation Roadmap will provide the detailed evidence base, programme milestones and investment case for these actions.

The UK ACE Network will work with UKRI, relevant government departments and industry partners to develop and deliver these priorities.

References

UK Policy, Strategy and National Capability

Department for Business and Trade (2023) *Advanced Manufacturing Plan*. London: HM Government.

Department for Business and Trade (2024) *Invest 2035: The UK's modern industrial strategy*. London: HM Government.

Department for Business and Trade (2025) *The UK's Modern Industrial Strategy*. London: HM Government.

Department for Science, Innovation and Technology (2021) *UK Innovation Strategy*. London: HM Government.

Government Office for Science and DSIT (2025) *The Wider Economic Impacts of Emerging Technologies in the UK*. London: HM Government.

HM Government (2021) *UK Innovation Strategy: Leading the Future by Creating It*. London: Department for Science, Innovation and Technology and Department for Business, Energy & Industrial Strategy.

HM Government (2023a) *The Integrated Review Refresh 2023: Responding to a More Contested and Volatile World*. London: HM Government.

HM Government (2023b) *The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023*. London: HM Government.

HM Government (2023c) *National Quantum Strategy*. London: Department for Science, Innovation and Technology.

Ministry of Defence (2025) *Defence Industrial Strategy 2025: Making Defence an Engine for Growth*. London: MoD.

National Energy System Operator (2024) *Clean Power 2030 Implementation Plan*. London: NESO.

Research Landscape

UK Research and Innovation (2022) *UKRI Strategy 2022–2027: Transforming Tomorrow Together*. Swindon: UKRI.

EPSRC (2022) *Tomorrow's Engineering Research Challenges (TERC) Report*. Swindon: EPSRC.

European Commission (2025) *European Innovation Scoreboard 2025*. Brussels: European Commission.

International Benchmarking, Defence and Strategic Competition

International Institute for Strategic Studies (2025) *Military Balance 2025*. London: IISS.

Black, J. et al. (2024) *Strategic Competition in the Age of AI*. Santa Monica, CA: RAND Corporation.

NATO Science & Technology Organization (2024) *Science & Technology Trends 2025-2045*. Brussels: NATO.

U.S. Department of Defense (2023) *DoD Directive 3000.09: Autonomy in Weapon Systems*. Washington, DC: DoD.

U.S. Department of Defense (2025) *Developmental Test and Evaluation of Autonomous Systems Guidebook*. Office of the Under Secretary of Defense for Research and Engineering.

International Federation of Robotics (2024) *World Robotics Report 2024*. Frankfurt: IFR.

Productivity, Competitiveness and Automation

The Productivity Institute (2024) *What Explains the UK's Productivity Problem?* Manchester: The Productivity Institute.

Resolution Foundation (2023) *The Economy 2030 Inquiry: Ending Stagnation*. London: Resolution Foundation.

Made Smarter Review (2017) *Review Report*. London: HM Government.

OECD (2024) *Economic Surveys: United Kingdom 2024*. Paris: OECD.

McKinsey Global Institute (2019) *The Future of Work in the United Kingdom*. London: McKinsey.

CBI Economics and ECIU (2025) *The Future is Green: The Economic Opportunities Brought by the UK's Net Zero Economy*. London: CBI.

Health, Bioprocessing and Life Sciences

Office for Life Sciences (2021) *Life Sciences Vision*. London: HM Government.

Schaefer, G. et al. (2023) Managing integrated continuous bioprocesses in real time: Highlights from N-mAb. *Biotechnology Progress*, 39(6), e3414.

Energy Systems and Net Zero

Department for Energy Security and Net Zero (2024) *Clean Power 2030 Action Plan*. London: HM Government.

CBI Economics and ECIU (2025) *The Future is Green: The Economic Opportunities Brought by the UK's Net Zero Economy*. London: CBI.

Lamnabhi-Lagarrigue, F. et al. (2017) Systems & control for the future of humanity. *Annual Reviews in Control*, 43, pp. 1–64.

Control Systems, Automation and Intelligent Systems

Brunton, S.L. and Kutz, J.N. (2019) *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge: Cambridge University Press.

Dean, S., Mania, H., Matni, N., Recht, B. and Tu, S. (2020) On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*.

Recht, B. (2019) A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2, pp. 253–279.

Olfati-Saber, R., Fax, J.A. and Murray, R.M. (2007) Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1), pp. 215–233.

Nedić, A. and Ozdaglar, A. (2009) Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1), pp. 48–61.

Amin, S., Cárdenas, A.A. and Sastry, S.S. (2009) Safe and secure networked control systems under denial-of-service attacks. In: Majumdar, R. and Tabuada, P. (eds.) *Hybrid Systems: Computation and Control (HSCC 2009)*. Lecture Notes in Computer Science, vol. 5469. Berlin: Springer, pp. 31–45.

Pasqualetti, F., Dörfler, F. and Bullo, F. (2013) Attack detection and identification in cyber-physical systems. *IEEE Transactions on Automatic Control*, 58(11), pp. 2715–2729.

Pfeifer, R., Lungarella, M. and Iida, F. (2012) The challenges ahead for bio-inspired 'soft' robotics. *Communications of the ACM*, 55(11), pp. 76–87.

Tao, F., Zhang, H., Liu, A. and Nee, A.Y.C. (2019) Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), pp. 2405–2415.

Nuclear Fusion

Lenholm, M. et al. (2024) Plasma control for the STEP prototype power plant. *Nuclear Fusion*, 64, 096036.

ITER (2019) Sustained nuclear fusion? Not without control engineering. Written by Pat Brans, 1 July 2019. Available at: <https://www.iter.org/node/20687/sustained-nuclear-fusion-not-without-control-engineering> [Accessed March 2026].