All-Party Parliamentary Group on Artificial Intelligence

AI in Military and Defence

Shaping the Future of National Security and AI-enabled Weapons in the UK

BIG INNOVATION CENTRE
Secretariat

8 September 2025
Policy Forum

Table of Contents

Rapporteur Preface: Page 3
Introduction: Page 4
Aim of Session: Page 5
Findings: Page 7

Evidence: Page 11
- Dr Dariusz Standerski, Ministry of Digital Affairs, Poland: Page 12
- Colonel Hugh Eaton OBE, Boston Consulting Group & UK Ministry of Defence: Page 18
- Group Captain Andrew M. Tidmarsh, Royal Air Force: Page 23
- Liberty Hunter, Ax:son Johnson Institute for Statecraft and Diplomacy: Page 29
- Rob Bassett Cross MC, Adarga: Page 35
- Professor David Whetham, King's College London: Page 41

BIOs of Evidence Givers: Page 47
About APPG AI: Page 51
Contact: Page 55

Title: AI in Military and Defence: Shaping the Future of National Security and AI-enabled Weapons in the UK
All-Party Parliamentary Group on Artificial Intelligence (APPG AI)
Date of Publication: September 2025
Publication Type: Parliamentary Brief | Policy Brief
Publisher: Big Innovation Centre

Rapporteur Preface from the APPG AI Secretariat
Professor Birgitte Andersen

This parliamentary policy brief synthesises and reports the insights presented to Parliament on 8 September 2025 on the subject: “AI in Military and Defence: Shaping the Future of National Security and AI-enabled Weapons in the UK”. The evidence session was co-chaired by Allison Gardner MP and The Lord Clement-Jones CBE.

This session was not a speculative discussion about the future.

It was a frank assessment of the strategic present — and the reality that AI has already moved into the centre of geopolitical power, military deterrence, and national resilience.

We brought together:

- Government leadership responsible for national cyber security, dual-use AI and sovereign capability (Minister Stański – Government of Poland)
- Strategic defence policy experts assessing national posture, industrial dependency and alliance interoperability (Colonel Hugh Eaton – BCG and British Army)
- Serving military commanders implementing AI at operational level (Gp Capt Tidmarsh – RAF)
- Infrastructure security specialists on subsea infrastructure protection (Liberty Hunter – Ax:son Johnson Institute)
- Defence practitioners, innovators and strategic operators (Rob Bassett Cross – Adarga)
- Military ethicists and policy advisors working inside Defence doctrine and education (Professor David Whetham – Kings College London)

Each witness addressed a different vector of the same problem:

- how the UK can retain sovereign agency when the “weapon” is no longer the platform, but the AI-enabled information infrastructure beneath it.

This report therefore covers:

1. How adversaries treat peace/war, influence/attack, civilian/military as one continuum.
2. The structural barriers blocking the UK’s ability to absorb AI at speed.
3. The procurement reforms required to deploy capability, not only discover it.
4. The infrastructure vulnerability created when compute and cloud are privately or foreign-owned.
5. The emergence of AI as a determinant of deterrence — and the consequences if the UK lags.
6. The ethical and accountability frameworks needed to retain meaningful human control.

This is not an APPG AI policy brief arguing for more innovation.

- The evidence shows we already have innovation.
- The question before Parliament is whether the UK can industrialise that innovation fast enough to protect the nation’s sovereignty, our alliances, and the credibility of our deterrence posture — and the digital infrastructure of our national security, which is now the principal attack surface in hybrid conflict.
- This is the central thesis running across all contributions which follow.

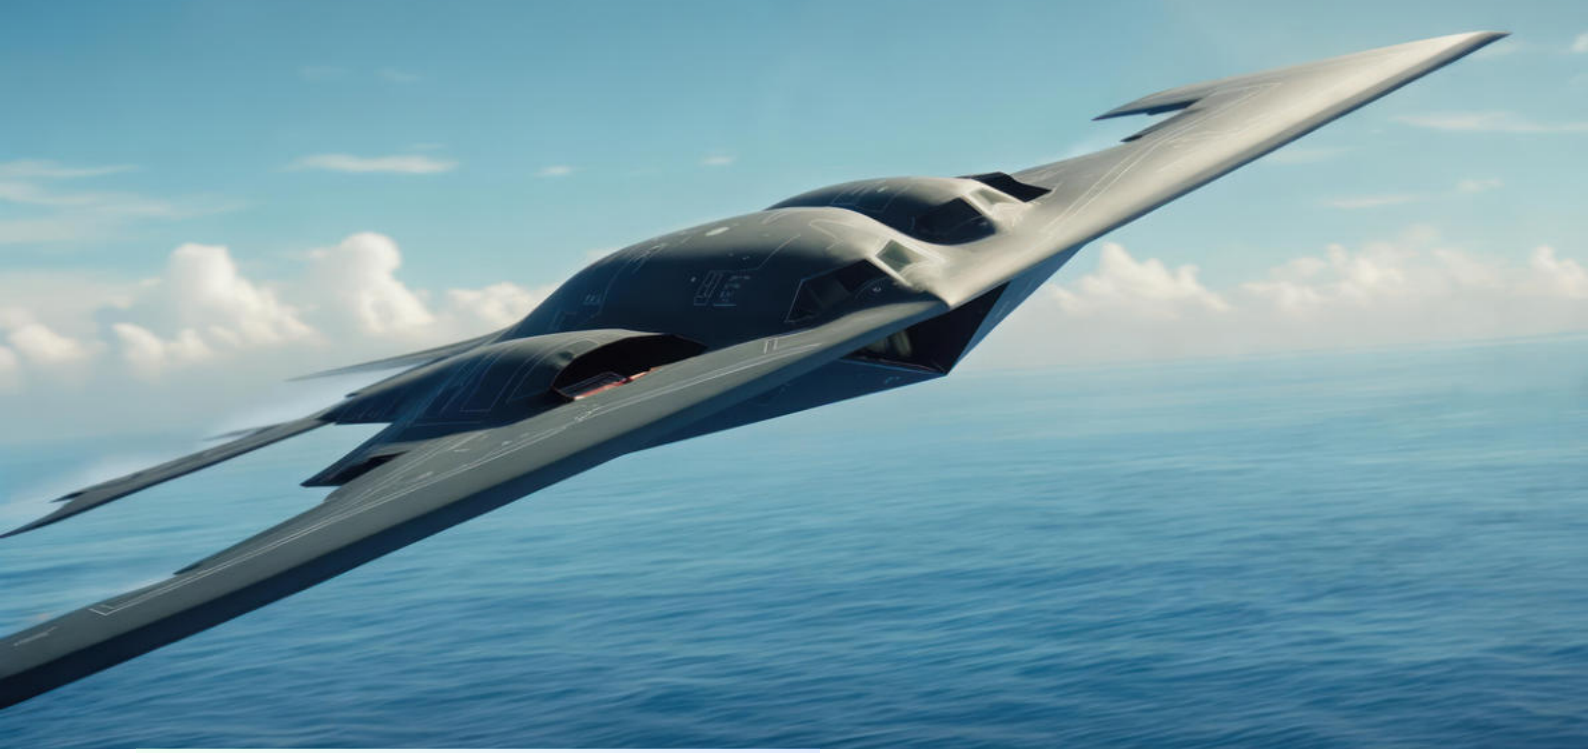

INTRODUCTION

This document is a transcript and summary of an APPG AI evidence meeting held on 8 September in the House of Lords, Committee Room 2a, UK Parliament. It exclusively contains crucial discussion elements; not all points are addressed.

DETAILS

- Evidence Session: AI in Military and Defence:
- Shaping the Future of National Security and AI-enabled Weapons in the UK
- Time 5:00 pm – 7:00 pm (GMT)
- Date: Monday, 8 September 2025
- Venue: Committee Room 1 House of Lords.

CONTACT THE SECRETARIAT

appg@biginnovationcentre.com
APPG AI Secretariat
Big Innovation Centre

BIG INNOVATION CENTRE

EVIDENCE GIVERS

1. **Dr Dariusz Standerski** - Secretary of State, Ministry of Digital Affairs, Poland
2. **Colonel Hugh Eaton OBE** - Special Operations Task Group, UK Ministry of Defence
3. **Group Captain Andrew M. Tidmarsh** - Deputy Head, Strategy & Plans, UK Strategic Command Headquarters | Royal Air Force
4. **Liberty Hunter**, Project Manager & Researcher, Ax:son Johnson Institute for Statecraft and Diplomacy
5. **Rob Bassett Cross MC** - CEO & Founder, Adarga
6. **Professor David Whetham** - Professor of Ethics and the Military Profession, King's College London | Director, King's Centre for Military Ethics

MEETING CHAIRS AND RAPPORTEUR

The Meeting was co-chaired by **Allison Gardner MP** and **Lord Clement-Jones CBE;**
Co-Chairs of the All-Party Parliamentary Group on Artificial Intelligence.

Rapporteur for this meeting: **Professor Birgitte Andersen**, CEO Big Innovation Centre

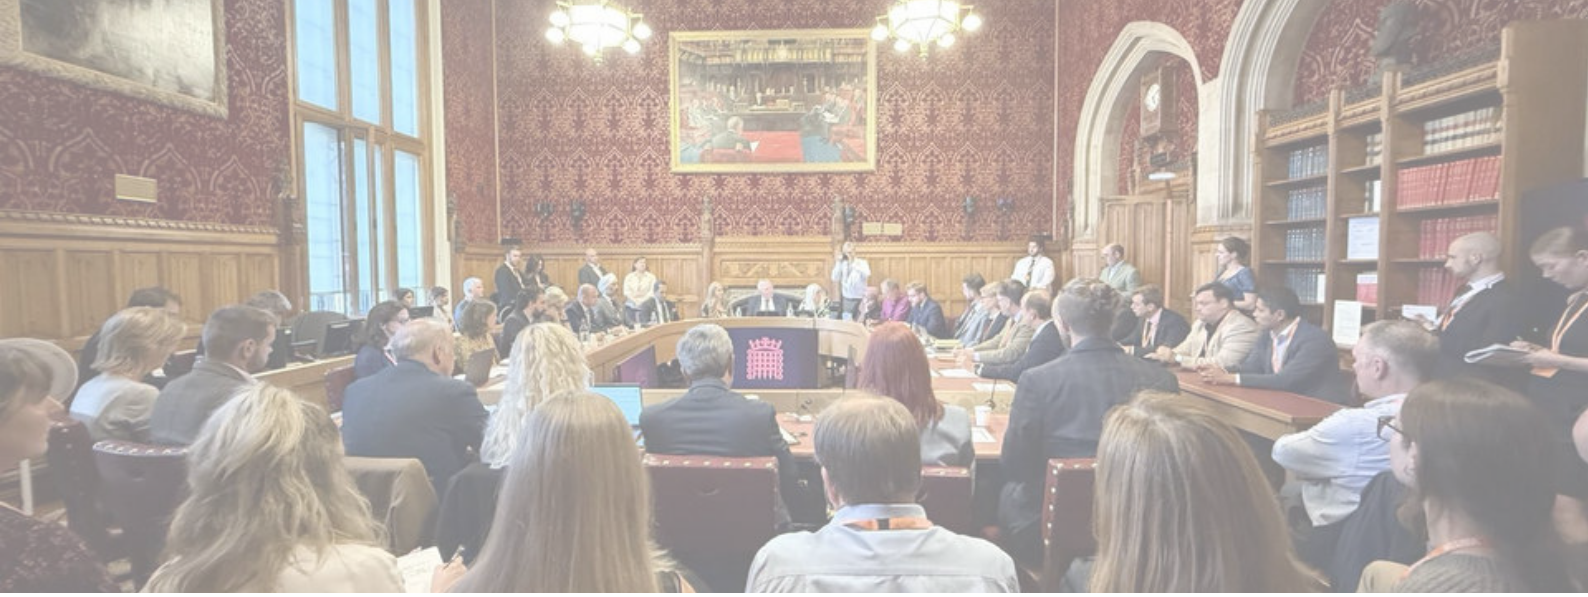

Aim of Session

AI in Military and Defence: Shaping the Future of National Security and AI-enabled Weapons in the UK

Artificial Intelligence is redefining the landscape of modern warfare, from real-time intelligence to autonomous weapons systems. This high-level meeting was convened to explore how AI can enhance the UK's national security, strengthen strategic decision-making, and drive defence innovation—while addressing the profound ethical, legal, and geopolitical questions it raises.

Discussion will cover:

- Strategic integration of AI in surveillance, intelligence, and combat operations.
- Autonomous systems and robotics—balancing military advantage with ethical safeguards.
- AI's role in strengthening cybersecurity and protecting critical infrastructure.
- International collaboration, interoperability, and governance frameworks.
- Ethical and legal implications of AI in lethal and non-lethal applications.
- Predictive analytics and threat detection for proactive defence strategy.
- Fostering innovation through public-private partnerships and investment.

With global military AI adoption accelerating, the UK faces a critical moment to define its role as a leader in responsible, practical, and strategically advantageous AI in defence. This evidence session brought together defence leadership, policymakers, industry, military ethicists, and security practitioners to assess how AI will reshape the UK's defence posture, national sovereignty, and the ethics of autonomous capabilities.

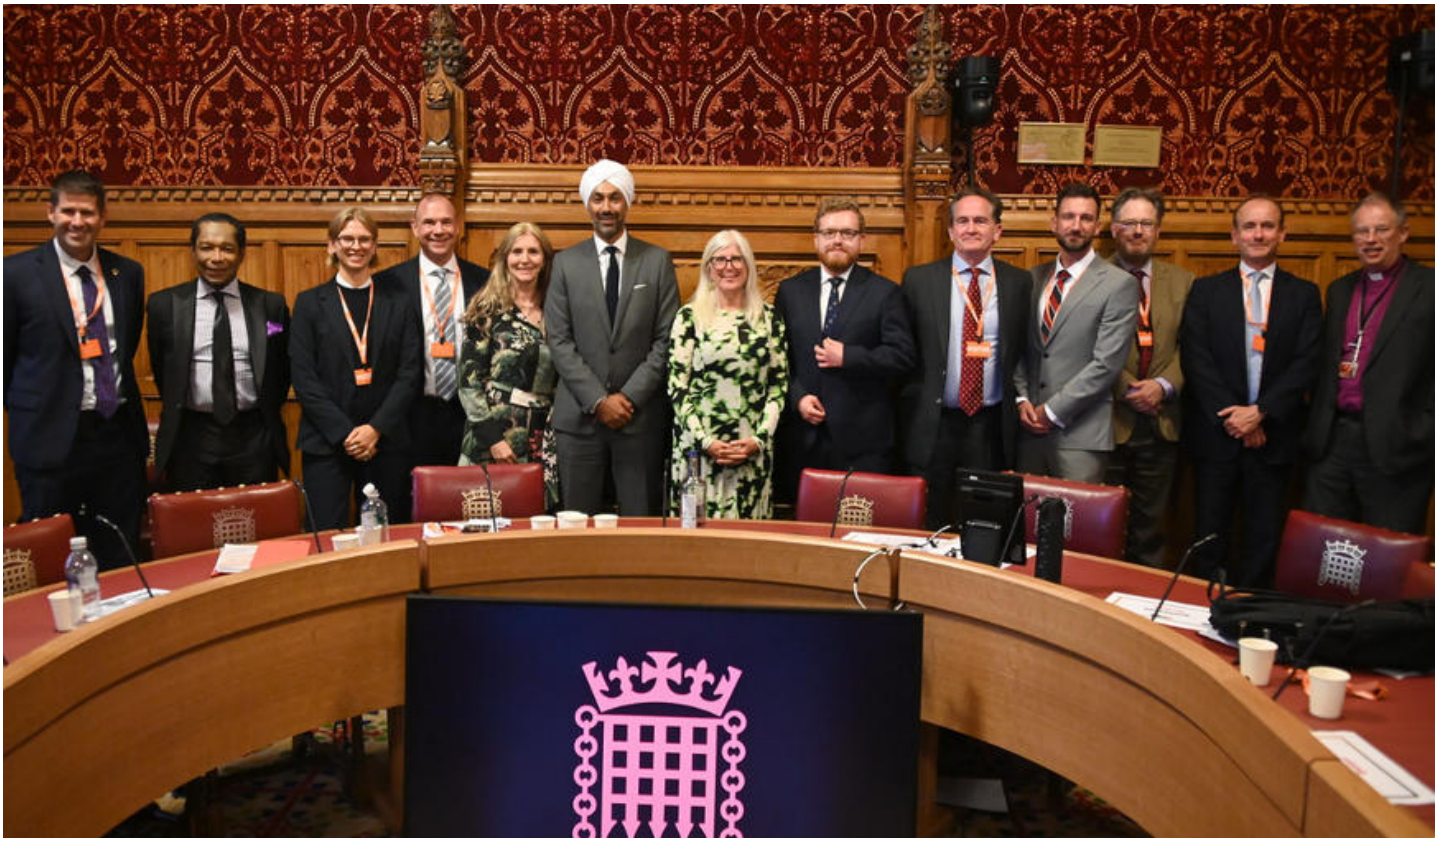

Above (from left to right): Chris Morton, Lord Taylor of Warwick, Liberty Hunter, Mark Buongiorno, Prof. Birgitte Andersen, Lord Kulveer Ranger, Allison Gardner MP, Secretary of State Poland Dariusz Standerski, Colonel Hugh Eaton OBE, Group Captain Andrew Tidmarsh, Prof. David Whetham, Robert Bassett Cross MC, Bishop of Oxford Steven Croft.

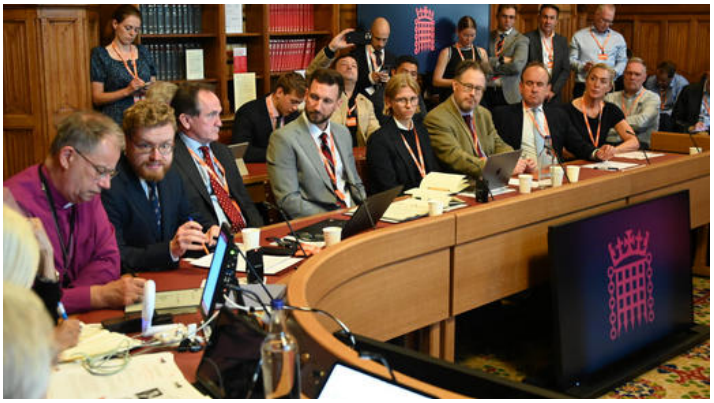

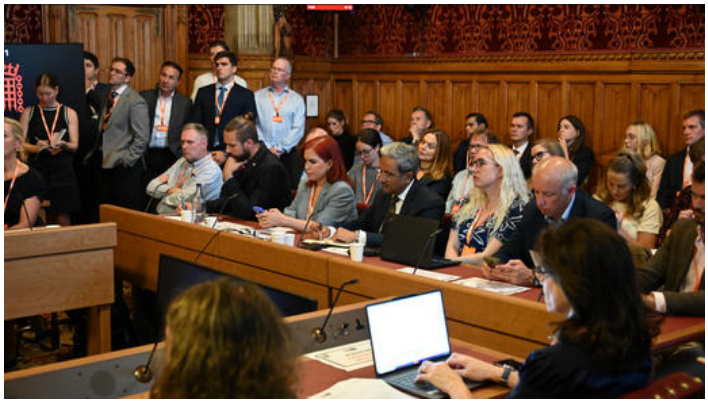

FINDINGS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

EXECUTIVE SUMMARY

This evidence session brought together defence leadership, policymakers, industry, military ethicists, and security practitioners to assess how AI will reshape the UK's defence posture, national sovereignty, and the ethics of autonomous capabilities.

Across all evidence provided, four themes emerged:

1. AI is now a strategic determinant of national power

AI is not a defence "add-on" — it is becoming the infrastructure of sovereignty:

- Nations that integrate AI fastest will set the strategic balance for the next decade.
- If Big Tech owns the digital infrastructure and with closed standards (which is the current trend), then Big Tech exercises strategic power — not the government.

2. The battlespace is already hybrid and continuous and based upon a 'grey-zone strategy'

- Adversaries (Russia/China) treat peace/war, civilian/military, and information/kinetic attack as a single continuum.
- Defence must now operate at machine-speed to deter, not just to fight.

'Grey-zone strategy' explanation:

Russia and China do not draw the same boundaries between peace and war that the West historically does; they see peace and war as a continuous spectrum, where civilian and military tools, and information and physical force, are all blended.

Meaning:

- They do not believe war only starts when bombs fall.
- They see influence campaigns, election interference, disinformation, cyber hacking, economic leverage, energy leverage, AI manipulation, etc. as part of conflict — even if no bullets are fired.
- Civilians, commercial platforms, and private data are all in the battlespace.
- They "prepare and shape" the battlefield continuously, not only during declared war.

This is called 'grey-zone strategy' (by NATO, U.S. Department of Defense, and UK MoD)

Put differently:

- In their doctrine, war has no sharp beginning — it is always happening, at different intensities, across multiple domains.

3. The UK has world-class talent — but an adoption bottleneck

The UK excels in research and innovation but lacks the machinery of absorption:

- procurement rules and processes are built for hardware,
- insufficient sovereign cloud (where data and models are stored and accessed) and compute (the physical processing power that runs the AI)
- siloed policymaking, and
- limited AI literacy among decision-makers.

4. Ethics is about responsibility and sovereignty, not prohibition

Ethics cannot be reduced to banning autonomy — it is about ensuring:

- human accountability over command intent,
- meaningful human control,
- and sovereign control over infrastructure.

This meeting established that the UK's strategic risk is not being "out-innovated" — it is being out-deployed; sovereignty now depends on the ability to industrialise AI faster than adversaries can exploit our delay.

STAKEHOLDER ACTIONS – MEETING OUTCOMES

GOVERNMENT / NO.10:

- Treat AI not as "digital transformation" but as core national security or state infrastructure
- Lead a sovereignty reset (repeating 2010's tech rebalancing (Efficiency and Reform Group (ERG), but for AI)
- Reclaim strategic control of infrastructure — because the entity that owns the platform owns the terms of use (if critical AI infrastructure is privately or foreign-owned, power shifts away from the state and toward Big Tech — government cannot exercise sovereignty it does not own). Alternatively, request open architectures for plug-ins.
- Prevent permanent dependency on privately-owned or foreign cloud (where data and models are stored and accessed) / compute (the physical processing power that runs the AI)
- Build UK sovereign capacity before quantum-era lock-in

MINISTRY OF DEFENCE

- Move from prototype → scale.
- Procure software as a living capability, not a static programme.
- Embed AI in C2 (command and control); train officers in its oversight.

INDUSTRY

- Provide open architectures to avoid dependency-lock.
- Co-invest when Defence provides a credible adoption pathway.
- Engineer for traceability and accountability from the outset.

RESEARCH & ACADEMIA

- Expand AI literacy as part of defence and C2 (command and control) professionalisation.
- Support doctrinal integration of AI (governance rules, standards, norms, KPIs, processes), not just technical integration of AI.
- Continue ethics-performance validation of AI vs human judgement.

NATO / ALLIED COOPERATION

- Subsea infrastructure & cyber defence must be internationally defended.
- AI-enabled surveillance (air, land, sub-sea) becomes an alliance responsibility, not a national one.

CIVIL SOCIETY / DEMOCRATIC OVERSIGHT

- The real frontier threat is information coercion (using power or leverage to force compliance), not about hardware (such as robotics, drones, or tankers).
- Legitimacy flows from how AI is used, not whether it exists → Public AI assurance requires transparency of doctrine (such as AI governance rules), not disclosure of code.

APPG AI Co-Chair:
Allison Gardner MP

APPG AI Co-Chair:
Lord Clement-Jones CBE

Secretariat & Rapporteur:
Professor Birgitte Andersen

EVIDENCE

Evidence Giver:
Dr Dariusz
Standerski

Evidence Giver:
Colonel Hugh Eaton
OBE

Evidence Giver:
Group Captain
Andrew M. Tidmarsh

Evidence Giver:
Liberty Hunter

Evidence Giver:
Rob Bassett Cross
MC

Evidence Giver:
Professor David
Whetham

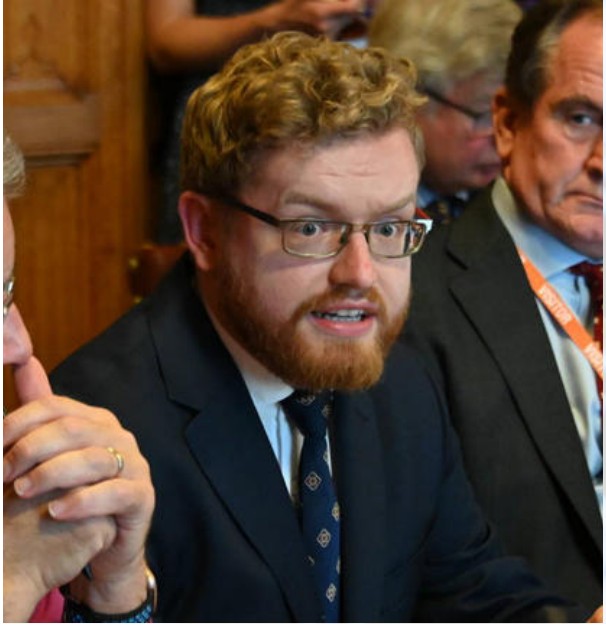

STRATEGIC TAKEAWAY

Poland's model shows that in the AI era, sovereignty does not begin at the border — it begins in the information layer of society, and whoever controls that layer controls the future security environment.

Minister Dariusz Stański
Poland's Secretary of State for Digital Affairs

Opening Remarks

Thank you very much, Chair, distinguished Members of Parliament and Peers, ladies and gentlemen. It is a great honour to have the opportunity to address you today. And if I may add, although my name seems difficult in English, it is actually one of the easiest in Polish — so please be prepared for the truly difficult ones in future.

Role and Strategic Scope

As Poland's Secretary of State for Digital Affairs, my portfolio covers not only artificial intelligence, but quantum technologies, digital infrastructure, and cybersecurity. Our government has deliberately structured these under one strategic umbrella so that we can treat civilian and military digital security as a single national ecosystem.

Cybersecurity Threat Landscape

Poland, like the United Kingdom, is under continuous attack in cyberspace. Only yesterday we registered 1,800 recorded cyberattacks inside Poland — including a dozen attempts against critical infrastructure. None of them were successful. The reason is not luck: it is a combination of highly trained personnel, heavy investment, and sovereign technological capability.

Civil–Military Integration

We work hand in hand with our Ministry of Defence to accelerate AI deployment across the Polish armed forces. This year alone, Poland will invest approximately £700 million in cybersecurity and AI integration. Our armed forces currently comprise 200,000 soldiers, including 7,000 serving in cyber and electronic defence units. And at the governmental level, there is daily cooperation between civilian digital authorities and the military cyber command.

IDEAS Institute: A Dual-Use First

In March of this year, we established the IDEAS Institute — a world-first dual-use AI research institution jointly overseen by both ministries. It serves both civilian administration and the armed forces, with research mandates coming from both. It is, to our knowledge, the only example in the European Union of an AI institute designed for dual-use from the very outset.

Scientific Talent and Sovereignty

The Institute is staffed by world-class Polish AI scientists, including four-time European Research Council grant-winners — people who worked for major global AI firms abroad and have now returned home to help build Poland's sovereign capabilities.

National AI Strategy Alignment

Policy and strategy are aligned accordingly. The Ministry of Defence has already adopted Poland's AI Strategy for the Armed Forces, with planning out to 2039. In parallel, my Ministry has delivered a national AI policy for the civilian sphere. Both strategies form a single blueprint for AI development and integration across the state.

Defence Vision for 2039

By 2039, our vision for defence includes:

- Autonomous reconnaissance and precision missions with minimal risk to soldiers;
- AI-driven intelligence able to predict adversary actions and recommend decisions in real time;
- Predictive logistics and maintenance systems;
- Cyberdefence capable of detecting and neutralising attacks instantly — particularly vital along our eastern and northern borders with Belarus and Russia.

We also expect AI to transform military training through hyper-realistic simulation and personalised adaptation, ensuring forces are prepared for dynamic threats.

Responsible and Ethical AI

However, we are fully aware that with opportunity comes risk. Autonomous weapons raise ethical questions. Algorithms can be manipulated. This is why Poland focuses on strict governance, rigorous testing, and robust assurance frameworks to ensure human accountability at every stage.

Sovereign Large Language Model

We have also developed our own national large language model, trained entirely on Polish infrastructure, Polish language corpora, and by Polish scientists. Sovereign AI is not a slogan for us — it is a national-security necessity. Technology that determines how a society thinks and communicates must be domestically rooted, not borrowed.

The Three Pillars of Sovereignty

Our sovereignty roadmap has three pillars:

1. Preparation — infrastructure, data integrity, skills
2. Implementation — scaling AI responsibly and securely
3. Protection — defending strategic assets and setting ethical guardrails

At the centre is the AI Implementation Centre embedded inside the Polish Cyber Defence Forces, ensuring research turns rapidly into operational capability.

Investment and Infrastructure

Poland now invests 4.7% of GDP on defence. A major portion of that supports dual-use digital infrastructure, including:

- a National Centre for Data Processing (£350m),
- a new sovereign quantum computer for our Military Technical University,
- and our Baltic AI Gigafactory project (£2.5bn) to house sovereign compute capacity in Europe.

We are proud that two British companies are already partners in this project.

European and Allied Cooperation

However, we fully recognise that no European country can do this alone. Sovereign capability must be networked across allies. We are here today because we want to deepen cooperation with the United Kingdom — in investment, in research, in infrastructure, and in interoperability of sovereign AI projects.

Thank you.

Q&A #1 — Human control in autonomous systems

"Where should the human be inserted in the loop? At which point in the chain of command do you believe human oversight must occur?"

Minister Stański:

This is a point of very active discussion in both our government and our military. Poland's approach is strongly anchored in the European Artificial Intelligence Act, which we are now implementing domestically. That gives us a legal baseline that guarantees meaningful human supervision over algorithmic processes.

For defence specifically, we have made a clear constitutional decision: AI may transform administrative procedures, logistics, sensing, targeting support, and simulation — but not the chain of command itself.

The chain of command remains human, because it is rooted in the civil and penal codes, which we consider technology-invariant. These are universal foundations — and they do not change regardless of innovation.

So, in Poland:

- AI can inform command,
- but it cannot become command.

That is our position.

Follow-up Q&A #2 — Sovereign AI / Polish LLM

Minister Stański continues:
We are also very intentional about sovereign AI. We are not building a large language model (LLM) because it is fashionable — we are building it because it provides strategic cognitive independence.

Our large language model uses:

- Polish infrastructure,
- Polish data,
- Polish scientists.

This is a matter of sovereignty. If the "thinking layer" of society is outsourced to a foreign system, then the nation does not control its own informational space. This is why sovereign AI is not just technical — it is statecraft.

Follow-up Q&A #3 — Procurement exemption: Later, when questioned about whether Poland has the same procurement bottlenecks that slow the UK:

Minister Stański:
In Poland we do not criticise procurement rules — because for defence and security matters, we do not use them.

We have an exemption for national-security investment. Instead of civilian procurement law, we use direct security oversight mechanisms from our anti-corruption and internal security agencies.

This gives us:

- speed,
- sovereignty of supply chain,
- infrastructure security,
- and full lifecycle visibility.

So, in our system:

- It is not procurement rules that safeguard national security —
- It is security oversight that safeguards national security.

Follow-up Q&A #4 — Populism, algorithmic warfare & information threat

Toward the end of the session, he made his strongest intervention:

The greatest threat is not cyberattacks or drones — it is algorithmic influence over populations.

Europe is now experiencing the rise of populism driven not by tanks or missiles, but by information operations. These are engineered, not organic. They are targeted psychological shaping of democratic opinion.

Poland recently disrupted a Russian FSB-directed operation recruiting Polish-language speakers for €5,000 each to spread anti-Ukrainian and anti-European narratives. If this is happening in Poland, it is certainly happening everywhere in Europe.

And Europe has no sovereign platforms:

- no sovereign social media,
- no sovereign search engines,
- and therefore no control of the public discourse infrastructure.

Our adversaries control their infosphere. We do not control ours. This is the battlefield that will define whether democracies remain cohesive.

Summary of Key Points from Minister Stański

1. AI and defence are inseparable

Poland treats digital security (AI, cyber, quantum, data infrastructure) as a single national-security ecosystem, not separate silos.

2. Cyberwar is already here

Poland faces continuous cyberattacks (1,800 in a single day). Defence works jointly with civilian AI infrastructure — daily coordination.

3. Dual-use AI is the future of defence

The IDEAS Institute is Europe's first deliberately dual-use AI research centre – civilian + military, from inception.

4. Sovereignty requires sovereign compute + sovereign models

Poland built its own LLM because control over language = control over cognition. AI is not only technical; it is statecraft.

5. Human remains in command

AI can inform, accelerate, target, and simulate —
but cannot ever become command or authority.
This is protected in constitutional/legal doctrine.

6. Procurement exemptions enable speed

National-security projects bypass normal procurement law entirely — replaced by security oversight instead of commercial procedure.

7. The real battlefield is informational

The gravest threat is not kinetic weapons — it is algorithmic population-shaping (disinformation pipelines, foreign influence).

8. Europe lacks sovereign platforms

Because Europe has:

- no sovereign social media,
- no sovereign search engine,
- no sovereign data platforms,
- it currently cannot protect its infosphere.

This is exactly where hostile actors exploit weakness.

Action points and implications:

For Government

- Treat AI, cyber, defence, and data as one strategic security domain to break siloed governance.
- Introduce national-security procurement exemptions to speed deployment, ensure sovereignty, and avoid MoD bottlenecks.
- Build a sovereign LLM (or co-develop with allies) to ensure cognitive independence.
- Create dual-use R&D infrastructure serving both civilian and military needs for faster innovation diffusion.
- Establish state oversight of information pipelines to counter algorithmic interference and protect the infosphere.

For Business / Industry

- Join sovereign infrastructure projects (compute, cloud, datacentres) to align economic activity with national security.
- Ensure supply-chain transparency and localisation to secure capability sovereignty.
- Co-develop dual-use applications, as the defence market now underpins national resilience.
- Participate in sovereign AI ecosystems instead of over-relying on foreign digital platforms, ensuring long-term strategic relevance.

For Research / Academia

- Build dual-use research models to accelerate translation from labs to state and defence applications.
- Help develop the sovereign cognitive stack (LLMs, datasets, ethics), recognising that language is sovereignty.
- Provide ethics-by-design frameworks to support a human-in-command doctrine.
- Support national talent retention so researchers stay and contribute domestically, reversing the brain drain.

For Civil Society / Democracy Protection

- Recognise algorithmic influence as a national security threat capable of destabilising democracies more than physical weapons.
- Advocate for public transparency in AI deployment to build trust and legitimacy.
- Support sovereign or federated platforms to reduce dependency on U.S./Chinese infospheres.
- Strengthen digital literacy on manipulation to reinforce democratic resilience.

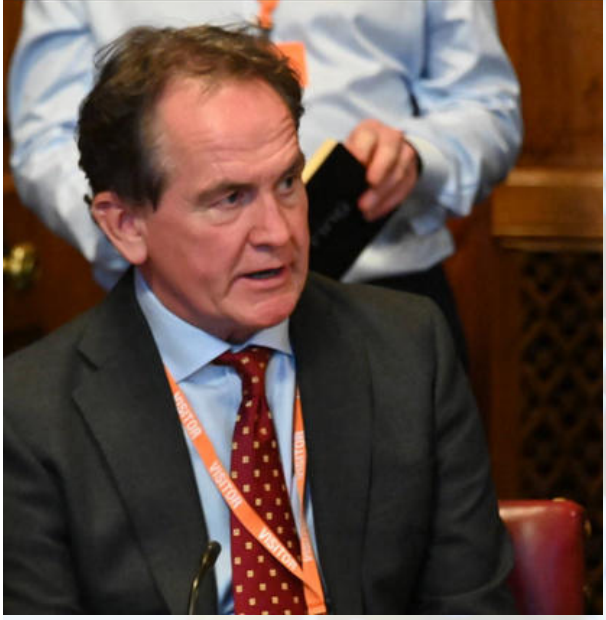

STRATEGIC TAKEAWAY

The UK is not losing because it lacks technology — it is losing because it has not reorganised the state to govern technology.

Colonel Hugh Eaton OBE

Boston Consulting Group &
Special Operations Task Group, UK Ministry of Defence

When I last spoke here six months ago, I warned that our ability to govern technology was falling behind our ability to develop it. That was true then — it is even more true now.

I want to cover three things:

1. What we actually mean by AI and defence
2. Our dangerously unbalanced relationship with Big Tech
3. NATO spending and the illusion of percentage benchmarks

And I will finish with a practical recommendation.

1. AI in defence and defence in AI

AI is not primarily a defence issue — it is a whole-of-state issue which also touches defence. Other regions — particularly in Asia and the Middle East — already understand this and are developing whole-of-government architectures around AI, critical resilience and digital sovereignty.

By contrast, the UK still treats AI as something fragmented: MoD (Ministry of Defence) here, DSIT (Department for Science, Innovation and Technology) there, industry somewhere else, and security somewhere separate again. We build bilaterals, even multilaterals, but we do not build coherence inside our own system.

The contest we are in is already under way — and yet we still behave as though conflict begins only once the first shot is fired. Our adversaries — Russia and China — do not share this threshold mentality.

They do not distinguish between:

- peace vs. war
- civilian vs. military
- influence vs. conflict

To them, it is all one continuum: the spectrum of pressure, coercion and shaping. We are late in recognising this — and it is hurting us.

2. Our relationship with Big Tech is badly out of balance

Right now, the British Government is not directing Big Tech — Big Tech is directing government. Our entire digital operating environment is critically dependent on a handful of US tech firms. And while these are friendly to us today, they are not sovereign, nor permanent, nor under national obligation. The next strategic technology after AI will be quantum — and if we repeat the same dependency mistake, we will permanently lose strategic autonomy.

If the state does not set terms, then the platform companies will. We have done this properly before. In 2010, during the austerity transformation, the government created the Efficiency and Reform Group (ERG), which forced Big Tech to the table. It worked. It saved billions. It rebalanced the relationship. We must do this again — not as a procurement exercise, but as a sovereignty-reset exercise.

3. The NATO 2.5% spending target is a distraction

Whether a country hits 2.5%, 2.7% or 3.1% defence spending is almost irrelevant if:

- the spending is not allocated strategically,
- it does not include the AI backbone, or if,
- it cannot be repurposed toward new forms of conflict.

Percentages will not win the next war. Capability, agility and sovereignty will.

We have a continuous-at-sea nuclear deterrent that few others fund — yet there is no mechanism in NATO accounting that offsets that contribution or compensates for it in innovation capability. We are measuring the wrong thing.

Practical recommendation

In 2010, the UK 2010 created the Efficiency and Reform Group (ERG), and it forced Big Tech to adjust to the UK government's terms. We need the equivalent now for AI and sovereignty: A national AI capability reset, led from No. 10, not delegated downwards.

We are not helpless — we simply are not yet organising ourselves to use the leverage we already have.

Q&A #1 — On interoperability versus sovereignty
"Doesn't interoperability with allies require using big US platforms?"

Hugh Eaton:

- No. Interoperability is not the same thing as dependency.
- You do not need to buy someone else's system to connect with it. You need open standards and open architectures. Sovereign capability plus interoperability is the correct model. Dependency is the lazy substitute.

Q&A #2 — Who holds the levers of change?
"Which person, group, or entity has the power or influence to initiate significant national security transformation and capability?"

Hugh Eaton:
This is not a Ministry of Defence problem. It is a national posture problem.

The levers lie with:

- No. 10
- The Cabinet Office
- Treasury (because tempo (flow and speed of activity) follows money)
- Parliament

Defence cannot fix this alone — the centre of government must redefine the national information posture.

Q&A #3 — Algorithmic disinformation as warfare

Hugh Eaton:
Poland's view is correct: the frontier is informational. We are under attack now, every day — in ways that fall below our own political threshold for "war." We are trying to fight 21st century state power with a 20th century peacetime mindset.

Q&A #4 — On preventing misuse of AI by government internally

Hugh Eaton:
There are currently not enough guardrails — but the bigger immediate danger is foreign misuse against UK citizens, not domestic misuse by the British state.

The UK currently has no clear owner of disinformation defence — that is the structural danger.

Summary of Colonel Hugh Eaton's Key Points

1. The UK misunderstands the nature of AI in defence

AI is not just a military technology — it is a whole-of-state capability that underpins national power and resilience. The UK is still treating AI in fragmented silos (Defence vs DSIT vs industry), unlike adversaries who use integrated strategy.

2. The UK is strategically dependent on Big Tech

Currently Big Tech sets the rules, not government. The UK relies on a small group of US platform companies for critical infrastructure. This dependency gives the UK no sovereignty over the digital environment. With quantum coming next, repeating this mistake means permanent dependence.

3. NATO spending metrics are misleading

Defence "percentages" (2%, 2.5%, 3%) don't measure capability. What matters is where money is directed and whether it builds strategic autonomy. Current NATO metrics disguise structural weaknesses.

4. The real frontier is informational conflict

Russia and China view peace/war, civilian/military, influence/conflict as one continuum. The UK still believes war begins when shots are fired — adversaries believe it begins when information is shaped.

5. The UK has leverage but is not using it

The UK successfully rebalanced Big Tech influence in 2010. The same approach is now needed for AI and digital sovereignty — centrally led from No. 10.

ACTION FIELDS BY STAKEHOLDER

Government

- Establish a National AI Sovereignty Directorate under No. 10 — because AI is now a whole-of-state capability.
- Replace fragmented oversight with a central strategic command — to avoid silo competition and policy duplication.
- Rebalance the relationship with Big Tech through a new Efficiency & Reform 2.0 approach — to restore sovereignty leverage.
- Define a formal institutional owner for disinformation defence — because currently no body is responsible for it.
- Shift from spend-percentage metrics to capability auditing — so the UK builds real defence resilience rather than budget optics.

Business / Technology Providers

- Build sovereign alternatives to foreign hyperscalers (compute + infrastructure) — to reduce strategic dependency.
- Adopt open standards and architectures — to prevent long-term platform lock-in.
- Participate in UK-led AI sovereignty frameworks — to align industry incentives with national strategy.

Research & Academia

- Support whole-of-state AI models, not siloed R&D — this matches adversarial doctrine and real-world threat posture.
- Study information warfare doctrine in applied form — preparing the UK for a constant-shaping environment.
- Develop governance frameworks that can actually be operationalised — moving ethics from theory into command posture.

Civil Society / Democracy Protection

- Recognise disinformation as a national security threat — not merely “online harm”, but hostile influence operations.
- Push for transparency in digital–state relationships — as a guardrail against captured sovereignty.
- Demand strategic accountability from Big Tech platforms — which currently shape the civic space without democratic consent.

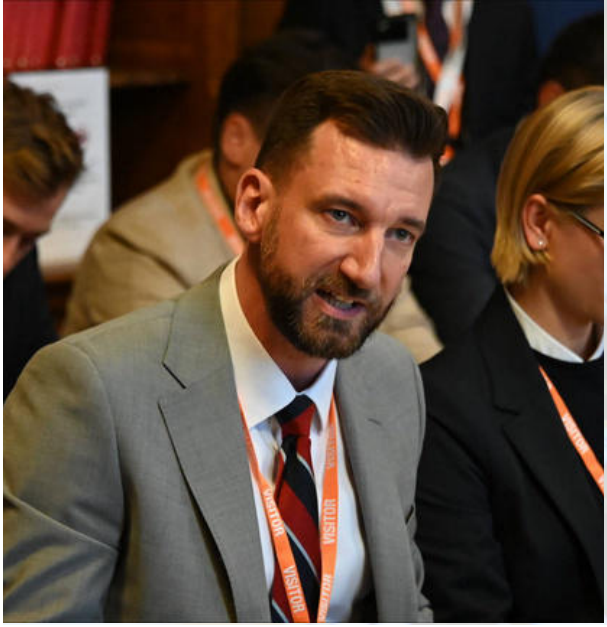

STRATEGIC TAKEAWAY

The UK does not need to fear AI in defence — it needs to fear being slower than those who already use AI decisively.

Group Captain Andrew M. Tidmarsh

Deputy Head, Strategy & Plans, UK Strategic Command Headquarters, Royal Air Force

Opening and Introduction

I am Group Captain Andrew Tidmarsh, a serving Royal Air Force officer. My previous roles include UK National Representative to NATO's Integrated Air and Missile Defence System; Commanding Officer of the UK's Electronic Warfare Testing and Training Facility; and Head of the Air Information Experimentation Group, charged with pioneering cutting-edge innovation for the front-line.

I am now Deputy Head Strategy and Plans in Cyber and Specialist Operations Command. My PhD research explores potential futures for Lethal Autonomous Weapon Systems.

The Reality of the Modern Battlefield

On today's battlefields, AI and autonomy are driving rapid changes in speed, mass, performance and insight. In Ukraine, RUSI has reported that drones now account for an estimated 60-70% of damaged or destroyed Russian systems – a staggering shift in the character of conflict. AI-based targeting systems are multiplying lethality on an unprecedented scale. This is the reality of the world today.

The UK Strategic Posture

The UK's National Security Strategy and the Strategic Defence Review mobilise science and technology: through the Integrated Force, a Digital Targeting Web, and rapid adoption of autonomy and AI. On 1 September 2025, Cyber and Specialist Operations Command stood up to unite cyber and specialist capabilities; counter cyber and electromagnetic threats; and deliver integrated operations and precision targeting.

Harnessing the electromagnetic spectrum, delivering machine-speed Command and Control (C2), and integrating autonomy are vital so that we may out-think, out-pace and out-fight adversaries in a contested, data-saturated environment.

Project DEEPTHOUGHT

Project DEEPTHOUGHT is a capability prototype that fuses electronic warfare data into a single analytical space. Machine Learning assists with parametric and mission-data refinement; visualisation of the electromagnetic domain at scale; and highlights anomalies faster for mission-data teams. It can be rapidly iterated to the edge – vital to win the cyber and electromagnetic contest that underwrites the Digital Targeting Web.

Project BOYD and Machine-Speed C2

The tempo of modern C2 (Command and Control) simply overwhelms human capacity. Project BOYD introduces the prototype Air Power 4.0 system, which has demonstrated explainable decision support to help commanders rapidly generate courses of action; compare risk, cost, effect and resilience; build resourced plans; monitor execution and optimise in real-time.

This is how advantage emerges from networks, not isolated platforms. Through clouds and sensor fusion, it will be the mesh that absorbs and processes information from all sensors; to enable rapid decision-making for employment of the best effector at the right time. That's the basis of the Digital Targeting Web. AI will be laced throughout. Indeed, it must be if we are to not only maintain pace with peer adversaries, but also to ensure safe and effective performance.

The Ethics and Autonomy Debate

AI in war is polarising, perhaps most apparent in the debate on autonomous weapon systems (AWS): often reduced to the idea of dystopian "killer robots". Strawman examples regularly tend to the extreme or mischaracterise conflict.

UN Secretary-General Antonio Guterres has called for agreement on a legally binding instrument that would prohibit lethal AWS that function without human control by 2026.

Here's the difficulty: what defines an AWS, and what defines human control?

Divergent State Positions

States' positions on a binding instrument, and their formal definitions or working descriptions vary. The House of Lords AI in Weapon Systems Committee report in 2023 recommended that the UK government adopt an operational definition of AWS. The previous government declined. Canada has not endorsed a new instrument, arguing that systems that lacked appropriate human control would already be outside International Humanitarian Law (IHL), though without a specific definition per se. France has defined partial and fully autonomous: with "fully" appearing to possess far greater autonomy than any human in the battlespace, and "partial" aligning quite closely with NATO C2 (Command and Control) states.

Signs of Convergence

There is growing convergence. At last week's gathering of the Group of Governmental Experts on technologies relating to LAWS, papers were submitted that sought to clarify the characterisation of LAWS as functional combinations of one or more weapons that may select and engage targets without human intervention; and to work within existing IHL, recognising life-cycle assurance, Commanders' intent and contextual restraints. This is not dissimilar to the position of Automated Decision Research, an initiative of the Stop Killer Robots campaign. They support context-dependent requirements for predictability, understandability, and temporal and geographic limitations.

The fundamental tension is that for definitions too vague or too extreme, any instrument would be meaningless; too restrictive, and they limit the potential for good.

The debate on laws hangs on IHL, human dignity and stability; underpinned by the assumption that only humans, not machines, can morally reason.

Research on Machine Moral Reasoning

Following early evidence of the emergence of Theory of Mind and fluid intelligence in Large Language Models, my own research has explored emerging psychological and attitudinal traits in AI that have been shown to directly predict human ethical performance on operations. I have explored the latent light and darkness of these traits and how they impact moral reasoning in military ethical scenarios, including the challenging British Army's Intermediate Concepts Measure, which has been used to test British Army Officers' alignment with Army values. In dialogue with Harvard Professors Julian De Freitas and Alon Hafri, I have tested AI ability to rapidly extract morally salient features defined by humans in visual scenes.

Findings

In all these cases, AI have out-performed the human comparisons. They have demonstrated higher levels of empathy, humility and honesty; lower levels of psychopathy and narcissism. They have been more sensitive to rapidly displayed moral features in visual scenes; adherent to IHL; and, within the experimental scope of the AICM, more closely aligned to British Army values than British Army Officers. In some instances, nearly perfectly aligned.

Concluding Reflection

This is absolutely not me suggesting we abandon human control, but that we should think carefully about where to apply it. The distinction between what it means to understand and whether it matters is something we can, perhaps, discuss.

But pandora's box is already open, the world is moving quickly and we have to be in to make a difference. We must remember who we are trying to protect, and be open and honest about how best to do that.

Q&A #1 — Does the UK still believe existing IHL (International Humanitarian Law) is enough?

Is the UK Government still taking the view that existing international humanitarian law is sufficient to govern autonomous systems?

Andrew M. Tidmarsh:

Yes — that remains the position. The IHL (International Humanitarian Law) framework is considered technologically agnostic, and the UK's position has not diverged from that. What is changing, however, is the implementation detail — stages are starting to move from abstract legal principles toward much more operationalised definitions of what constitutes sufficient human control, accountability, and constraints. So, the law itself remains unchanged, but its characterisation is becoming more functional.

Q&A #2 — Does the UK have a working definition of lethal autonomy?

Has the UK now adopted a working definition of an autonomous weapon system?

Andrew M. Tidmarsh:

To the best of my knowledge — no, not formally. There is movement toward definitional convergence internationally, but I am not aware of any confirmed UK wording being adopted domestically yet. The UK is aligning conceptually with that international direction of travel, but it has not crystallised into a national legal definition.

Q&A #3 — Can AI sometimes outperform humans ethically? - Is there a human in the loop?

If AI outperforms humans on empathy and moral alignment, does that not undermine the case for human control?

Andrew M. Tidmarsh:

Not undermine — but reframe. The question should not be "Is there a human in the loop?" but "Is the human in the right part of the loop?" In some time-critical decision environments, inserting a human at the wrong point may actually increase risk, not reduce it. The issue is quality of control, not symbolism of control. Human judgment has value, but it must be well-positioned in the decision-making chain.

Q&A #4 — Are lethal AWS (autonomous weapon systems) inevitable?

Andrew M. Tidmarsh:

We are rapidly moving toward inevitability in the sense that machine-speed operations leave no alternative. Without AI assistance, we will fall behind tempo (the relative speed and rhythm of actions with respect to the enemy, aimed at keeping them off balance and preventing effective countermeasures). However, inevitability does not remove governance: the future is not autonomy without constraint; it is autonomy within a disciplined context. The barrier is not technology — it is calibrated trust and guardrails.

Q&A #5 — Where do humans still matter most?

If AI can sometimes make better moral judgements, where does that leave the human role?

Andrew M. Tidmarsh:

Humans remain responsible for command intent, boundaries, escalation logic and normative accountability. Machines can excel at compliance — they cannot yet own responsibility. The human role shifts upward: from technician to custodian of legitimacy.

Summary of Key Points — Group Captain Andrew Tidmarsh

Definition: "Tempo" means the relative speed and rhythm of actions with respect to the enemy, aimed at keeping them off balance and preventing effective countermeasures.

1. AI is now central to modern warfare

- The tempo of cyber and electromagnetic conflict vastly exceeds unaided human decision-making.
- Projects like DEEPTHOUGHT (data fusion) and BOYD (machine-speed decision support) are already demonstrating battlefield-ready military AI.

2. The nature of conflict has changed

- The battlespace is now networked and data-driven.
- Success no longer comes from superior single platforms, but from fusion: clouds, sensors, autonomy, C2 (command and control).
- The UK must adapt to "machine-speed command and control" or become tactically outpaced.

3. The lethal autonomy debate is being reframed

- The old question, "Should there be a human in the loop?" is too crude.
- The real question becomes: "Where in the loop should the human sit?"
- Human oversight must be meaningful, not ceremonial.

4. AI may sometimes be more ethically reliable than humans

- Experimental evidence shows AI can:
  - Extract morally relevant features faster,
  - Show greater consistency with IHL (International Humanitarian Law) and values,
  - Display lower "dark triad" psychological traits.
- This does not remove responsibility — but it changes who performs which part of the decision process.

5. Governance model must match tempo

- Slowed-down human control is not ethical if it causes harm to civilians or harms friendly forces.
- Opt for machine-speed partners with human judgement rather than replacing it (i.e., we shall let machines accelerate parts of the decision cycle, but humans still make the authoritative moral/legal judgement).

ACTION FIELDS BY STAKEHOLDER

Government / Centre of State

- Adopt operational guidance on “meaningful human control” (not only a legal baseline).
- Accelerate alignment with emerging GGE international frameworks.
- Update the defence doctrine to reflect machine-speed operations.

Ministry of Defence / Armed Forces

- Embed AI throughout C2 (command and control) and battlespace intelligence, not as bolt-ons.
- Expand programmes like DEEPTHOUGHT and BOYD into full-scale capability.
- Train commanders in where human control is best applied.

Industry

- Provide open-architecture, modular battlefield AI infrastructure.
- Support explainability and mission traceability in operational systems.

Shift from selling platforms to selling plug-in decision capabilities. Instead of defence companies saying, “Buy my drone,” they should say, “Buy my AI module that turns any drone into a smart ISR (Intelligence, Surveillance and Reconnaissance) asset.” Why? Because war advantage is now information + decision superiority, not the metal.

Research & Academia

- Continue ethics-performance testing of AI vs. human operators.
- Develop models for calibrated trust and role allocation.
- Provide doctrinal research on life-cycle assurance and accountability.

Civil Society (where relevant)

- Demand transparency of doctrine, not just transparency of code.
- Understand that restraint can be designed into systems, not only imposed by treaties.
- Shift the debate away from Hollywood “killer robots” tropes toward real governance questions.

STRATEGIC TAKEAWAY

AI does not just protect subsea infrastructure — subsea infrastructure is what makes AI possible. Defence must begin below the waterline.

Liberty Hunter,
Project Manager & Researcher,
Ax:son Johnson Institute for Statecraft and Diplomacy

Opening Remarks

My name is Liberty Hunter I work for the Ax:son Johnson Institute. I have two degrees from Oxford University where I specialised in British Maritime Strategy in the 20th century and specifically the defence of submarine infrastructure which I will be talking about today as a case study in a modern context, asking how AI can help better defend these cables.

Why Submarine Cables Matter

Just to set the scene slightly. Although this presentation focuses on how AI can help these cables, it is worth reflecting on the fact that. AI in Britain relies on these cables to work. Generative AI system might use around 33 times more energy than usual search engines. Britain imports 44% of its energy. Imports along critical maritime infrastructure.

So, why do we care about submarine cables?

Since the late 19th century, submarine critical infrastructure has been both a target and a weapon in war. Today, as they were during the First and Second World War, submarine cables (and other maritime infrastructure) are vital to everyday communications and transactions.

They are however, now more valuable and more vulnerable. They are more valuable - upwards of 95% of international data traffic is dependent on submarine cables rather than satellite transmissions. And Secondly, because the nature of the cargo carried by the cables is much more extensive and critical to everyday life – Damage to a significant number of UK data cables, or to significant cable landing stations, would not only test the threshold for a declaration of war. It would also be catastrophic for personal communications, business transactions that keep our country afloat, and national vulnerability as satellites do not have the capacity to carry the weight of traffic and are not a viable alternative, especially for financial transactions.

They are more vulnerable, firstly, because International Maritime Law has not been updated to meet the security needs of todays, and secondly the capabilities of adversaries to sabotage cables has increased significantly.

The Scale of the Threat

There are 150 to 200 instances of damage to the global network each year. most damage to subsea infrastructure is accidental, but purposeful damage is increasingly common and is extremely useful as a form of hybrid warfare. Attributing responsibility for cable damage is difficult but international tensions are increasing as owners of damaged cables – 98% of which are privately owned- are taking action. In response to the severing of the Finnish -Estonia cable in late 2024, Finnish authorities have since filed charges against three crew members of the oil tanker suspected of damaging the cable.

The UK's Current Posture

In light of recent cable sabotage we are starting to realise the vulnerabilities of these cables. In the most recent Strategic Defence Review, the Royal Navy has been tasked with surveillance and protection of subsea infrastructure. But the Royal Navy does not have the capacity to defend or even maintain our CMI – relying almost solely on one repair ship CS Sovereign for repairs of our 60 cables.

The Role of AI

The additional use of AI has so far and increasingly will improve productivity and efficiency of defending CMI. So far, the Royal Navy has developed and deployed some forms of AI to enhance its maritime surveillance capacity using, for example, the SeaCat, Defender ROV, and Gavia on RFA Proteus, which are equipped with AI. Although AI is not necessarily the solution to all Royal Navy capability gaps - There are several clear ways in which the nature of AI lends itself to submarine infrastructure surveillance and threat detection.

How AI Assists Maritime Surveillance

What can AI do to help the Royal Navy mandate? AI can help to improve maritime surveillance which has the capacity to prevent cable damage. AI can process data much faster – and a much higher volume - than humans can and would therefore increase the amount of data that the Royal Navy can use to assess maritime traffic and threats to infrastructure. The UK Maritime Domain Awareness team receive in excess of 160 million data points a day so understanding that data through traditional means is impossible and therefore makes us reactive to threats. Additionally, to understand the Pattern of Life, which in our case is 14 years' worth of data (14 x 365 x 160,000,000), makes accurate human understanding impossible – hence use of AI.

AI has the capacity to speed up data processing in order to create a more comprehensive 'Recognised Maritime picture' that involves processing and fusing multiple data sources. By learning how ships act, what their patterns of travel are, and surveying current maritime traffic, AI can be harnessed to detect 'Maritime Anomalies' and subsequent potential threats to subsea infrastructure, these aren't just abstract technologies – they have been successfully used by Windward, Lockheed Martin, and Lloyd's List Intelligence. To give an example, just yesterday Microsoft confirmed that two key submarine cables were severed near Jeddah, Saudi Arabia. Windward's Maritime AI intelligence platform, which layers undersea cable locations with real-world vessel behaviour, was able to identify two candidate vessels whose activities at the time and place match potential involvement which will help governments and private companies to attitude responsibility better.

British allies are already harnessing AI on offer. For example, Denmark has deployed 'saildrones', fitted with onboard AI, to compile data using multiple sensors, cameras and radar in the Baltic. This produces a more detailed picture of maritime activity than satellites can provide. There has been significant recent co-operation between NATO allies for data sharing and joint CMI surveillance missions that harness AI, including the Operation Mainsail.

AI on the Seabed: Threat Detection Below Surface

A second application of AI for enhancing maritime security is in its ability to aid detection of threats on the seafloor itself. Recently, Thales and US companies have produced an autonomous mine hunting system that is equipped with AI. The equipment, the AI-led 'Mi-Map' sonar analysis is four times faster than previous tools and is more precise.

AI is, similarly, analysing sonar data in a NATO Operation Mainsail to detect the acoustic signature of a ship's anchor hitting the seabed. Static seabed sensors detected the anchor drop and transmitted the acoustic data and generated an automatic alert suggesting possible sabotage. This AI could be used to better detect both accidental and purposeful damage to submarine cables, and alert a deterrent response, saving millions of pounds in repairs.

What Next for Parliamentarians? (Concluding Thoughts)

It is amazing to be speaking to people who are able to influence policy. If I were able to offer three key take ways they would be as follows

1. Firstly, AI-enabled fusion is essential for maritime infrastructure security. Royal Navy physical assets should be seen as complementary to AI. AI is not necessarily a capability gap filler on all levels in British defence. However, the RN does not currently have the capacity to defence our infrastructure physically. While waiting for the RN to expand fleet of vessels like RFA Proteus, fitted sensors are an excellent stopgap.
2. Second, this is a global issue. There are almost 1.5 million kilometres of cables worldwide. Subsea infrastructure is a British national security issue, but these cables are, in practice, international entities, and should be defended as such. Use of NATO missions and shared AI capabilities is crucial, as seen in the Baltic Sentry, Operation Mainsail, and Nordic Warden already. The sharing of resources and cross-national data sets should be encouraged.
3. Third, it is clear from the recent Strategic Defence Review (SDR) and National AI Action Plan that the government is endorsing an increasing use of AI. However, policy makers are inherently risk-averse, but the Royal Navy does not have sufficient networks necessary to support the scale of integration of AI into their defence operations. Recognising this shortfall is critical. Conversations with policymakers need to devise a strategy for the Royal Navy to develop a network capable of managing AI integration and holding the extensive amount of data necessary for AI to operate effectively. Whilst Britain builds up its physical defence assets, AI in the short term can and should act as an effective and bold step towards the modernisation of the defence of British CNI (Critical National Infrastructure). Looking towards the long term, there are huge potential wins, not only for efficient defence capabilities but also for reallocation of Royal Navy resources serving to improve British Homeland defence.

Thank you so much for listening.

Q&A #1 — Big Tech dependency in maritime security

Earlier we heard from the Colonel Hugh Eaton OBE that Big Tech dominance is a strategic risk. In the subsea infrastructure space, is the UK also dependent on US tech in the same way?

Liberty Hunter:

- There has been significant cooperation recently between European governments and private maritime data companies, so I would not say the UK only relies on American systems — the dominant providers are still overseas. There is a European capability emerging, but it is fragmented and not at the same scale.
- The UK currently depends on private-sector platforms for subsea situational awareness, and most of those platforms are not British. So yes — the sovereignty issue does apply in this area as well, even though we do have some allied capacity.

Q&A #2 — Institutional risk aversion

You mentioned policymakers are risk-averse. Do you see this as a cultural or structural barrier?

Liberty Hunter:

- Both. The Navy does not currently have the technical support networks that would allow it to fully operationalise AI at scale, and until that is in place, risk aversion is almost inevitable. Policymakers are also hesitant because the infrastructure we are protecting is privately owned and internationally distributed, which complicates responsibility.
- So this is partly cultural — but mostly structural. The people responsible for defending subsea cables do not yet have the architecture needed to deploy AI confidently.

Q&A #3 — On NATO burden sharing

Should subsea infrastructure really be considered "national" security when 98% of it is privately owned and most is internationally routed?

Liberty Hunter:

- This is exactly why I stressed that this is not a purely British problem. For example, (i) legally, subsea cables are private; (ii) operationally, they are critical national infrastructure; (iii) strategically, they are international.
- That means the defence burden cannot just fall on one country. NATO-level AI-enabled surveillance is the only realistic way to secure them in practice.

Q&A #4 — Maritime domain awareness as deterrence

Does AI deter sabotage or merely help detect it?

Liberty Hunter:

- It does both. Detection is the first layer, but attribution is the deterrent; - just spotting that something bad is happening is not enough — you need to be able to prove who did it, and show you can respond.
- If nations believe sabotage will be identified, they are less likely to attempt it. That is the strategic value of AI — it changes the risk calculus for the aggressor.

Summary of Key Points — Liberty Hunter

The UK's most vulnerable infrastructure is under the sea

- Subsea cables carry ~95% of global data, and the UK imports ~44% of its energy through undersea infrastructure. If cable systems fail, the AI stack collapses with them — AI depends on what lies beneath.

Subsea sabotage risk is rising

- Damage was once accidental — now it is increasingly deliberate. Attribution is difficult, making it ideal for hybrid warfare and grey-zone coercion.

The Royal Navy cannot physically protect the cables

- The UK has ~60 subsea infrastructure cables, but effectively relies on one repair ship. Physical defence capacity is insufficient.

AI is already proving effective

- AI provides maritime domain awareness and anomaly detection on a scale that humans cannot process — 160 million data points per day.
- Systems like Windward and NATO's Operation Mainsail have already identified likely saboteurs. Denmark's saildrones demonstrate that AI outperforms satellites for situational awareness.

This is a global problem, not a national one

- Cables are privately owned and internationally routed. No single nation can defend them alone — NATO Intelligence Fusion Centre (NIFC) is an operational necessity.

The policy bottleneck is not technology — it is integration capacity

- The Royal Navy lacks the digital backbone and support networks to scale AI. Policymakers are risk-averse because they lack the infrastructure needed to act confidently.

Stakeholder Action Fields

1. Government / Parliament

- Establish a national maritime AI infrastructure strategy → Current approach is fragmented and reactive
- Fund digital maritime surveillance as a complement to ships → The Royal Navy fleet cannot meet the defence burden alone
- Clarify legal responsibility across Nation state vs. private ownership → 98% of cables are privately owned
- Deepen NATO integration for infrastructure defence → Single-state defence is structurally impossible

2. Defence & Industry

- Co-develop sovereign or allied maritime AI tools → Reduces dependency on non-UK platforms
- Expand AI-enabled domain awareness networks → Early attribution is the real deterrent
- Incentivise dual-use technologies → Defence benefits from commercial innovation cycles
- Build resilience partnerships with cable operators → Physical risk sits on private infrastructure

3. Research & Technical Community

- Develop AI for acoustic/sonar pattern recognition → Enables faster threat detection on the seabed
- Expand modelling for attribution frameworks → Legal deterrence depends on evidence
- Improve fusion of multi-source maritime data → 160 million datapoints a day is beyond human analytical capacity
- Conduct research on hybrid warfare thresholds → Clarifies escalation boundaries and legal response doctrine

4. Civil Society / Strategic Communications

- Raise public understanding of seabed vulnerability → Most people think AI is “in the cloud”, not “under the sea”
- Build support for NATO shared protection → Infrastructure is international by nature
- Normalise dual-use defence framing → Civilian + military = national continuity
- Push for transparency in infrastructure defence → Builds legitimacy around AI-enabled monitoring

STRATEGIC TAKEAWAY

The UK does not lack capability — it lacks the infrastructure and governance to absorb capability at the speed required to deter modern threats.

Rob Bassett Cross MC
CEO & Founder,
Adarga

I am Rob Bassett Cross, CEO and founder of Adarga — a UK defence artificial intelligence company operating in both the UK and the US.

My career has focused on applying emerging technologies to real-world defence missions. My first exposure to this was in Iraq, where I helped introduce computational linguistics and network-analysis tools into intelligence and targeting operations. Later, in the private sector, I watched JP Morgan transform itself using AI at enterprise scale. What mattered there was not slides, not pilots, not declarations of "innovation" — but infrastructure, culture, and velocity of deployment.

And my central message to you today is this:

- The UK has the AI talent. What it lacks is an operating system capable of absorbing that talent at speed.
- We are one of the strongest AI ecosystems in the world — but Defence is still not AI-ready. The gap between potential power and realised power is widening.
- AI will determine future military strength, economic leverage and geopolitical influence. The race we are in is not "who has the best model" — it is who can deploy capability fastest. Right now, Defence cannot deploy fast enough because the adoption machinery is still structured for the industrial age, not the software age.

Procurement is the central bottleneck

It is still built for hardware-era programmes: multi-year cycles, rigid specifications, technology "freeze points", and the illusion that certainty is possible up front. AI is a living capability — it must evolve continuously. The commercial sector has already restructured around this reality. Defence has not.

Capabilities routinely die in experimentation because there is no route to scaling them. The effect is that we keep "discovering" innovation we never industrialised. Meanwhile, our adversaries — and even some of our allies — are industrialising theirs.

Infrastructure is the second bottleneck

We do not yet have sovereign, defence-grade compute capacity at the scale required. The French Government recently purchased a thousand high-end GPUs to seed their sovereign stack. UK Defence has access to only a tiny fraction of that capacity.

Banks today are running thousands of GPUs every day for commercial inference. Defence — which is supposed to deter war — cannot currently match that. This is strategically backwards.

The window of risk

We are approaching a convergence of two curves:

1. The emergence of far more capable machine intelligence, and
2. A reconstituting threat posture from hostile states.

If that convergence arrives before the UK can industrialise AI, then we will not deter — we will invite coercion. Failure to adopt is not neutral; it is actively dangerous. It signals incapacity.
Not adopting AI has become an anti-deterrent.

Defence does not have to build everything — but it does have to build the conditions for others to build with it. That is the missing layer.

The UK has done this before. In 2010, the Government forced Big Tech to the table through the Efficiency and Reform Group (ERG). It rebalanced power. It saved billions. It reset the relationship. We now need a sovereignty-reset for AI — led from the centre of government, not delegated downwards. If we do this, the UK can lead. If we do not, we will fall behind — not because we lacked ingenuity, but because we lacked the ability to absorb it.

Thank you.

Definitions:

- Velocity issue in defence refers to is that the speed of events / attacks / decisions is now so fast that human-only processes cannot keep up
- Deterrence is the strategy of preventing someone from doing something, not by stopping them physically, but by making them believe that doing it will cost them more than it's worth.

Q&A #1 — Procurement failure and "why hasn't Defence fixed this yet?"

"Your criticisms of procurement are very familiar. Is Defence repeating with AI the same mistakes it made with legacy technology?"

Rob Bassett Cross:

Yes — structurally, yes. The procurement system is still treating software like hardware. It assumes you are buying a finished product rather than a continuously evolving capability. That mindset kills speed. It also kills adoption. Defence doesn't lack innovation. It lacks a pathway to industrialise innovation.

Q&A #2 — Interoperability vs Sovereignty

"Big Tech argues interoperability requires us to use their platforms. Is that correct?"

Rob Bassett Cross:

No — interoperability is not the same as dependency. You don't need to buy someone else's stack to connect to it. If you build on an open architecture with sovereign control, you can still integrate with allies. The dependency narrative is convenient for vendors — but strategically very dangerous if we accept it uncritically.

- To work with allies, you only need standard interfaces (open standards, open APIs, published data schemas, interface contracts, modular plug-in architecture) — you do not need to buy their full tech stack. Vendors tell you otherwise because it makes them money, but strategically it is very risky to believe them:
- Interoperability = ability to connect.
- Dependency = losing sovereignty and choice.

Q&A #3 — Why industry is holding back investment

"Why are companies not stepping in more aggressively?"

Rob Bassett Cross:

Because the demand signal is still weak. Industry will invest when Defence shows it can absorb capability and scale it. Right now, AI companies look at Defence and think: "I will spend three years in pre-procurement, deliver a prototype, and never see it fielded." That's not a market — that's a dead end. Fix absorption, and capital comes in.

Q&A #4 — Is this just a resourcing issue?

"Is this about not spending enough?"

Rob Bassett Cross:

- It's not primarily a money issue — it's a velocity issue. We could double the spend tomorrow and still fail if the operating model doesn't change.
- Capability now depends on iteration speed, not budget size. AI punishes slow systems.

Q&A #5 — Who actually holds the levers?

"If the system is stuck, who can unstick it?"

Rob Bassett Cross:

Not the MoD (Ministry of Defence) alone. This is a national posture question, not an internal procurement tweak. The levers are at the centre of government:

- No. 10 (Prime Minister)
- The Cabinet Office
- Treasury
- MoD cannot reform the national AI backbone by itself. It needs No. 10 to reset the relationship with Big Tech and signal sovereign intent.

Q&A #6 — Dependency on US technology

"Does this trajectory make us permanently dependent on the United States?"

Rob Bassett Cross:

If we fail to build sovereign capacity — yes. And that dependency will be locked in permanently at the quantum stage. We still have a sovereign window in AI. But it is closing. Once you're reliant at the compute layer, you don't get sovereignty back.

Q&A #7 — Deterrence consequences

"What happens if we fail to adopt in time?"

Rob Bassett Cross:

Then we don't just fall behind — we weaken deterrence. Adversaries don't look at intentions, they look at capability velocity. If a state can't integrate AI fast enough, it signals vulnerability. That invites pressure. That is why I say: "Not adopting AI is now an anti-deterrent".

Summary of Key Points – Rob Bassett Cross (CEO, Adarga)

The UK has world-class AI talent, but Defence cannot currently absorb it.

- The limiting factor is not innovation supply — it is innovation absorption. The operating system of Defence (procurement, infrastructure, and data access) remains industrial-era, not software-era.

AI is becoming a determinant of national power and deterrence.

- Geostrategic influence will divide along who can scale AI first. Lagging adoption is not neutral — it actively weakens sovereignty and military credibility.

Procurement is the structural failure mode.

- The system still assumes "capability as a finished product," not "capability as a continuously evolving software service." This kills deployment speed and prevents industrialisation of prototypes.

Infrastructure is the second systemic bottleneck.

- UK Defence lacks sovereign high-grade compute at scale. Allies and adversaries are investing heavily; the UK risks structural dependency if it does not accelerate.

Failure to adopt creates an 'anti-deterrent' effect.

- Capability velocity is now part of deterrence signalling. If the UK cannot industrialise AI, it signals vulnerability and invites coercion.

This cannot be fixed by MoD alone.

- The lever is national posture — requiring action from No.10, Cabinet Office and Treasury. Defence cannot reform sovereign infrastructure by itself.

STAKEHOLDER ACTION FIELDS

Centre of Government (No.10, Cabinet Office, HMT)

- Issue a national AI capability reset, as was done in 2010 with the Efficiency & Reform Group (ERG).
- Treat AI as sovereignty infrastructure, not merely "digital transformation."
- Provide a long-term demand signal to unlock private co-investment.

Ministry of Defence

- Shift from prototype-to-nowhere → prototype-to-scale.
- Procure software as a living capability (continuous upgrade model).
- Reform classification-era IT bottlenecks to enable rapid deployment.

Industry

- Deliver open-architecture AI that can plug into sovereign stacks.
- Co-invest once there is a credible demand signal from the government.
- Engineer for trust, auditability and rapid integration — not bespoke lock-in.

Research & Academia

- Develop methods to translate AI capability from lab to deployment.
- Produce frameworks for measured trust and operational assurance.

Civil Society / Legitimacy Layer

- Shift public debate from "stop AI" → "govern AI deployment."
- Recognise that deterrence now includes technology velocity.

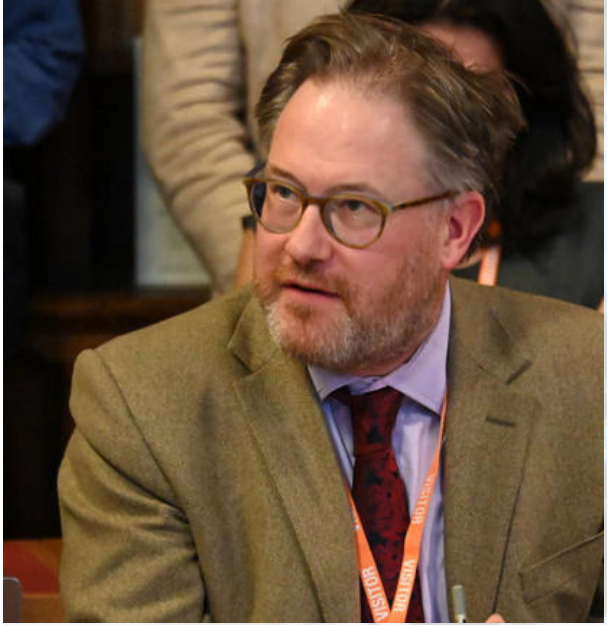

STRATEGIC TAKEAWAY

Ethical AI in Defence is not about constraining capability — it is about ensuring sovereignty, responsibility, and decision-making remain human.

Professor David Whetham
Professor of Ethics and the Military Profession
King's College London | Director, King's Centre for Military Ethics

I'm David Whetham, Professor of Ethics and the Military Profession at King's College London. I have just taken over as academic director at the Royal College of Defence Studies and have been more broadly based at the UK Defence Academy for the last 23 years where I lead on the ethics components of courses for up to 3,000 British and international officers a year. I also support wider UK Defence in a number of other capacities, including providing advice to 77 Brigade, Defence Medical Services and I sit on the MoD AI Ethics Advisory Panel that helped introduce the Responsible AI policy for Defence. Today I speak in a personal capacity.

Not the Expected Warning

You might expect me as an academic and a military ethicist to spend the next 5 mins giving you warnings about how terrible AI and autonomous weapons are and how we need to massively restrict what we are doing now and in the future.
I'm not going to do that.

A Cautious Optimism

I'm cautiously optimistic about the opportunities provided by embracing AI for UK Defence, (as we have heard) including some genuine opportunities to act in a more ethical way while enhancing rather than compromising military effectiveness. For example, I'm part of a project researching the very real capability enhancements afforded by RPAS/drone Casualty Evacuation platforms, incorporating AI triage and ai-supported medical techniques to keep our people alive when they are at their most vulnerable.

However, from my own experience, I have seen some very real challenges to maximising the potential of AI.

Challenge 1: Defence Education

I'll start with some thoughts on military education. AI poses a challenge for education more broadly – how to integrate AI into what we do, how to use it to enhance rather than dull critical thinking and reflective skills – but it's a particular challenge for Defence Education, where people need to be equipped for a highly competitive environment, as noted in SDR 2025. Clearly, this needs to be reflected in what is taught, and also how it is taught at the Defence Academy, and you won't be surprised to know that curricula are evolving very fast in response to the changes.

However, worryingly, evidence right now suggests that there is a real gap between what staff officers think they know about AI, and their actual knowledge of it as a capability. Given the stubbornly consistent demographics, maybe it's just a case of "lord, give me the confidence of a middle-aged white man", but I think there's more to it than that.

A study completed this year showed that whilst students on the Advanced Command and Staff Course value and recognise the relevance of AI for their military functions, they lack the fundamental AI literacy that could lead to genuine understanding and the capacity to take advantage of the opportunities of AI; or to deploy it in an appropriate and ethical manner. So, one of the challenges is a general confusion, or at least fuzziness, about the whole subject area and what it means. This is a problem for defence as a whole as we grapple with what new technology can offer. But, given the decision-making roles where many of these staff officers end up, it also a challenge for procurement.

I believe our decision-makers are particularly vulnerable to "snake oil salesmen" who promise that only 'their' bespoke system can deliver the magic bullet that defence desperately needs.

Challenge 2: Legal and Policy "Fuzziness"

I think that this fuzziness in understanding extends into, and in someways is made worse by, the policy and legal areas. I note the UK policy position that the Law of Armed Conflict is "technologically agnostic". The law is the law. As such, our official position is that autonomous weapon systems do not require any additional legal frameworks.

On the one hand, I do think that our current legal and policy frameworks are a lot more adaptable than many give them credit for. For example, let's look at the issue of accountability. There are suggestions in some quarters that as you can't court martial an AI, its going to be down to the developers to "carry the can". However, this isn't the way it works in any other area, so why should AI enabled capabilities be magically different? For example, nobody argues that a military dog is not a sentient creature capable of acting under its own initiative. At the same time, nobody has a problem with ascribing responsibility or accountability if something goes wrong – you don't court martial the dog (although it may face other more permanent sanctions instead).

Most people would also find it peculiar at the very least to blame whatever the dog does entirely on the dog breeder, the trainer or even the environment. There may be situations in which any of those factors are highly relevant, but that doesn't mean the command decision to deploy the capability is immune from criticism or that the team or dog handler's actions are somehow devoid of accountability assessments.

That doesn't mean we don't need some robust thinking about how to operationalise this understanding with the new technologies.

Challenge 3: Narrow Definitions of Autonomy

There are other areas where policy appears so black and white that it should banish all notions of fuzziness, but the effect is actually the opposite.

I have noted in some settings a worrying tendency for some MOD Lawyers to use an obscure definition of autonomy, which I'd describe as relating only to truly 'self-adaptive' systems. The only thing they acknowledge as really autonomous is a system that can completely self-task in an unbounded environment, choosing how, when and where it acts without constraints and without any human input at all.

If we are relying upon such a narrow definition simply so the UK can say "we don't do autonomous weapons", "we don't do killer robots", we are adding to the fuzziness of the subject and stifling engagement, understanding and policy development. It is far more useful to think of autonomous systems within a framework of meaningful or appropriate human control, while acknowledging that there may be a degree of autonomy within certain bounded areas.

For an example of a tool designed to try and remove some of this fuzziness, see www.ai.militaryethics.uk – an education tool we developed in partnership with Dstl to socialise Responsible AI into the developer community, but now also being used by new RAISOs (Responsible AI Senior Officers) etc to help understand the context of their new jobs.

Challenge 4: New Vulnerabilities

My final concern is the potential creation of brand-new vulnerabilities and critical points of failure. Who controls access to the new capabilities we are developing? Military technology is rarely sovereign in the sense that HMG completely owns the tool or just as importantly, the enabling infrastructure.

I know that bespoke military capabilities are hardened and designed to be robust against external influence, but try doing anything with the British military without Microsoft. Imagine removing WhatsApp and Signal and your chances of coordinating anything at the organisational level has just become vanishingly small.

In a world where foundational assumptions are being challenged, and a US administration can compel a private company to lock out judges at the International Criminal Court preventing them from accessing essential court materials because they are doing their job, what less obvious or exquisite capabilities rely on foundational assumptions that were formed when the geopolitical situation looked rather different?

Q&A #1 — Are existing legal frameworks sufficient?

"The MOD (Ministry of Defence)'s longstanding position is that existing IHL (International Humanitarian Law) is sufficient. In your view, is that still credible as AI develops?"

David Whetham:

I believe IHL can be sufficient — but only if we operationalise it properly. The danger is not the absence of law, but the absence of clarity about how human responsibility is embedded in practice. If we hide behind a slogan of "technological agnosticism" without doing that work, then the law becomes a comfort blanket rather than a protection.

Q&A #2 — Should the UK adopt a clear definition of autonomy?

"Earlier speakers noted growing convergence on definitions. Should the UK now adopt one?"

David Whetham:

Yes — because right now we are using an artificially narrow definition in order to say "we don't do autonomous weapons." That approach is disingenuous and unhelpful. It freezes policy discussion. We should acknowledge autonomy within bounded contexts and focus on appropriate human control, not deny autonomy exists at all.

Q&A #3 — Liability versus accountability

"You gave the example of the "military dog". Does that mean developers are not liable?"

David Whetham:

Developers can be relevant — but they are not the sole or default bearer of blame. We don't court-martial the "dog breeder". The chain of accountability remains with commanders and users. The fact that a system contains AI does not magically relocate responsibility outside the military decision-making structure.

Q&A #4 — The deeper risk: infrastructure dependency

"You warned about new vulnerabilities. How serious is this?"

David Whetham:

Extremely serious. We are focusing on "weapons," but the real dependency risk is infrastructure. If access to the underlying platforms is not sovereign, then in extremis someone else controls the "off switch." Lose infrastructure, and capability collapses — AI, communications, even command unity.

Q&A #5 — Guardrails against domestic misuse vs hostile misuse

"Are we more at risk from our own government misusing AI, or external actors?"

David Whetham:

In the immediate term, hostile external misuse is the greater danger. But the absence of clear lines of accountability creates space for internal misuse later. That is why ethics has to be built in now, not after crisis pressure.

Summary of Key Points — Professor David Whetham

1.Cautiously optimistic, not alarmist.

AI can enable more ethical outcomes when correctly integrated, including life-saving use cases (e.g., casualty evacuation).

2.The biggest barrier is conceptual "fuzziness."

Defence actors think they understand AI but lack genuine literacy, which creates:

- opoor doctrine,
- ovulnerability to vendor hype,
- oand bad procurement choices.

3.Ethics cannot be bolted on.

Human control must be meaningful, not symbolic.
Accountability must remain within the military chain of command.

4.The UK is hiding behind an artificially narrow definition of "autonomy."

This protects reputation, not reality — and stalls policy development.

5. National security implications: The greatest long-term ethical risk is dependency on non-sovereign infrastructure.

If AI systems — or communication systems — can be shut down by a foreign entity, ethics are irrelevant: Before you can talk about "ethics" of how you use AI, you must first own and control the AI.

- Sovereignty > Security > Ethics.

STAKEHOLDER ACTION FIELDS

Government (Cabinet Office / MoD Centre)

- Stop using artificially narrow definitions to avoid policy debate.
- Officially adopt a meaningful human control doctrine.
- Treat infrastructure sovereignty as an ethical requirement, not only a technical one.

Example of the problem of too narrow definitions:

If you define an autonomous weapon as: “something that can set its own mission, in an unbounded environment, with zero human input at all” → then almost nothing currently exists that fits that definition.

So if you choose that definition, you can claim:

- the UK has no autonomous weapons
- therefore the UK never needs to have a policy debate about them
- therefore the UK doesn’t need to update doctrine

Professor David Whetham is saying that too narrow definitions stop honest discussion. Because in real life, many systems have some autonomy inside boundaries, and that’s where the real risks and real decisions are. Plain language version of Whetham’s point: “Don’t pretend we have no autonomous systems by defining autonomy so narrowly that nothing counts. That stops us having the real policy debate we need.”

Defence Academy / Military Training System

- Scale AI literacy training across the officer pipeline.
- Embed ethics alongside capability, not as an afterthought.
- Train commanders in accountability pathways for AI-enabled decisions.

Procurement & Capability Teams

- Guard against “snake oil” sales through literacy, not bureaucracy.
- Require clarity on responsibility, not only performance metrics.
- Ensure lifecycle accountability is traceable.

Industry

- Design transparency-by-default to support military accountability.
- Stop hiding autonomy behind euphemistic labels; be explicit about functionality.

Research & Academia

- Provide tools to demystify AI (as per ai.militaryethics.uk).
- Help Defence operationalise ethics rather than theorise it.
- Support doctrinal, not purely technical, integration.

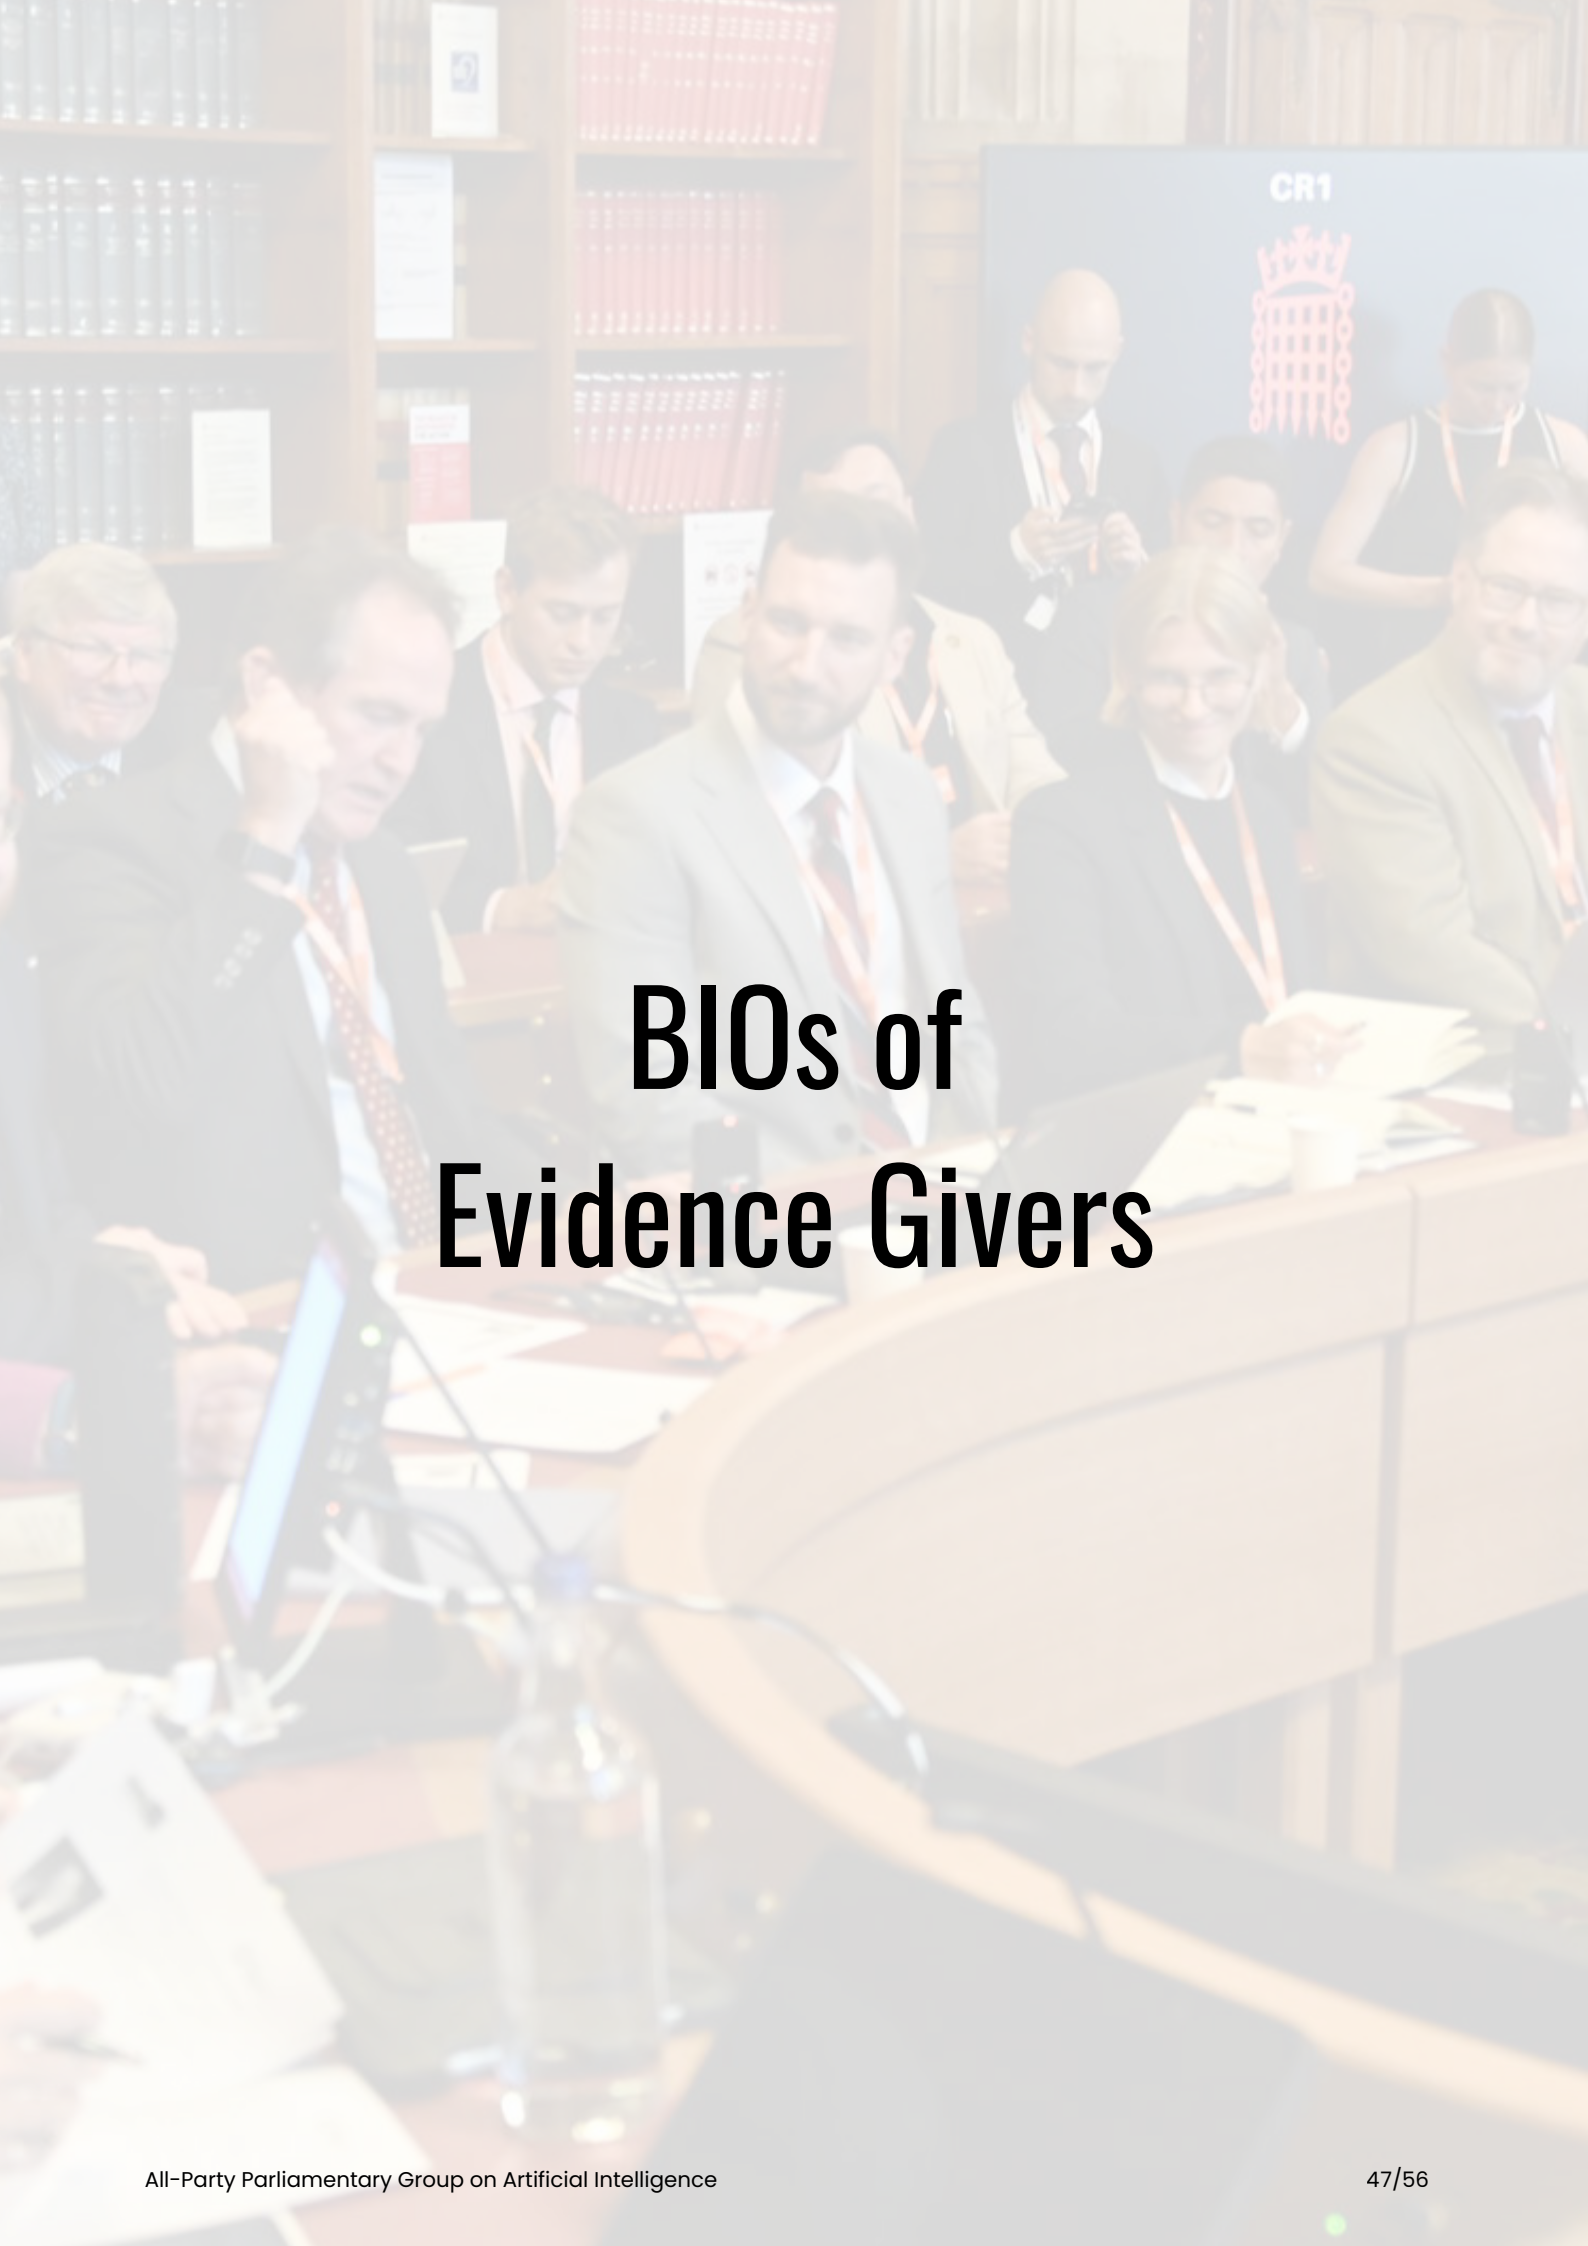

BIOs of Evidence Givers

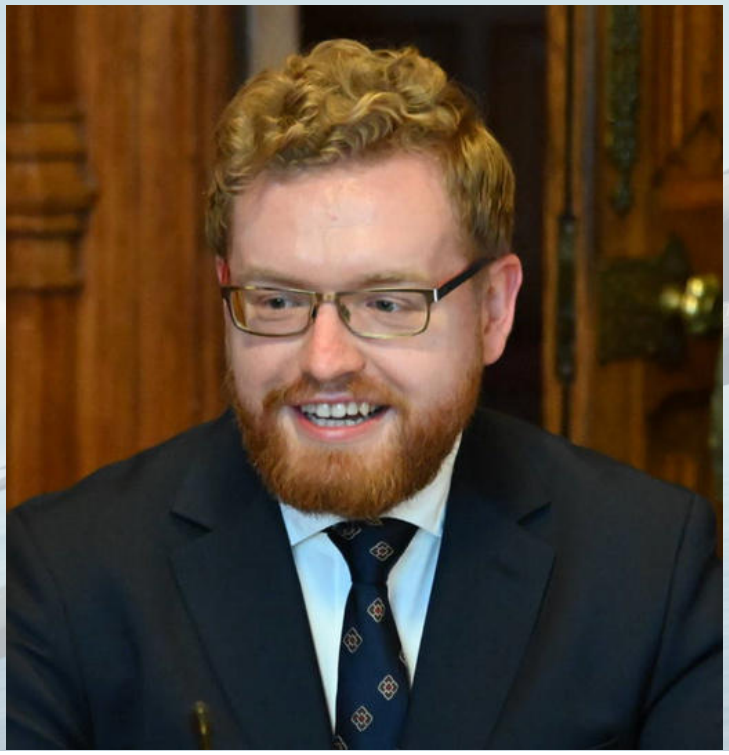

Dr Dariusz Standerski
Secretary of State,
Ministry of Digital Affairs, Poland
Economist & Lawyer | Digital Transformation Leader

Dr Dariusz Standerski is the Secretary of State at Poland's Ministry of Digital Affairs, where he leads the country's digital transformation strategy. His portfolio includes artificial intelligence, quantum technologies, EU digital policy, internet governance, cybersecurity, and the legislative frameworks that underpin innovation and national security.

A Doctor of Economics and qualified lawyer, Dr Standerski lectures at the University of Warsaw's Department of Economic Sciences. He previously served as Lead Economist and Management Board member at the Kalecki Foundation and as Director of Legislation for the Left Parliamentary Club, where he authored over 200 legislative proposals in economic, social, and digital policy.

Co-author of the "Digital State. Strategy for Poland", he plays a central role in shaping Poland's position as a competitive, innovative, and ethically responsible leader in the digital era.

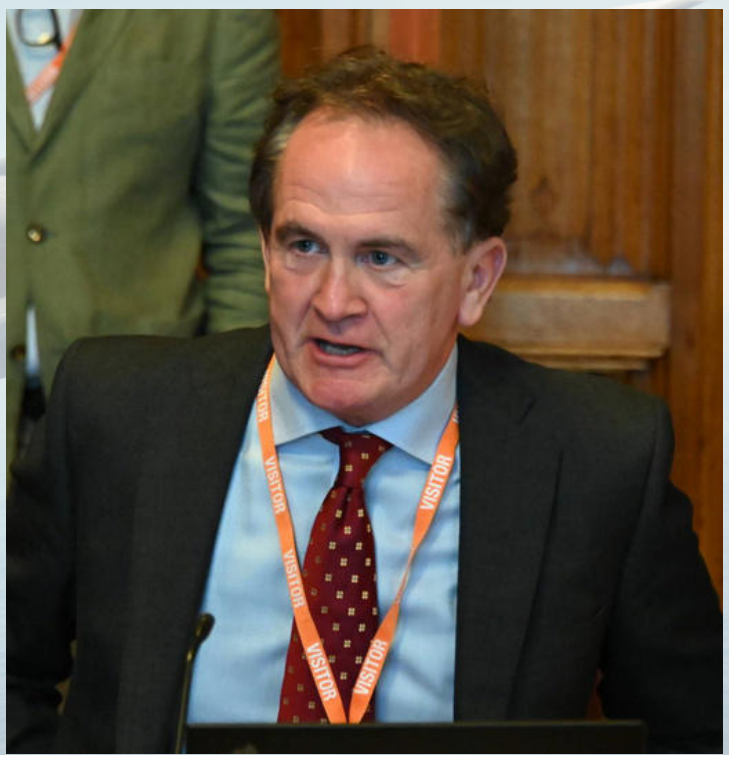

Colonel Hugh Eaton OBE
Boston Consulting Group &
Special Operations Task Group,
UK Ministry of Defence
International Defence & Technology Strategist

Colonel Hugh Eaton OBE is a senior military officer and internationally recognised strategist at the intersection of defence and emerging technology. As a Non-Executive Adviser to the UK Ministry of Defence, he supports the Strategic Defence Review, guiding the integration of advanced capabilities such as artificial intelligence and cyber technologies into UK national security strategy.

With over 18 years' operational command experience in the Army Reserve, including high-intensity combat deployments, Colonel Eaton has also held senior leadership positions as Global Head of Government at Microsoft and Regional Director for Public Sector at Cisco, shaping government digital strategy across Europe, the Middle East, and Africa. He has served as Managing Director in the Defence and Intelligence Practice at PA Consulting Group and is currently an Expert Adviser to the Boston Consulting Group.

He has advised the UN, EU, NATO, and national governments worldwide on AI, cybersecurity, smart cities, and digital transformation. A recipient of both the MBE and OBE for outstanding service to the United Kingdom, he brings deep geopolitical insight to the future of AI in modern warfare.

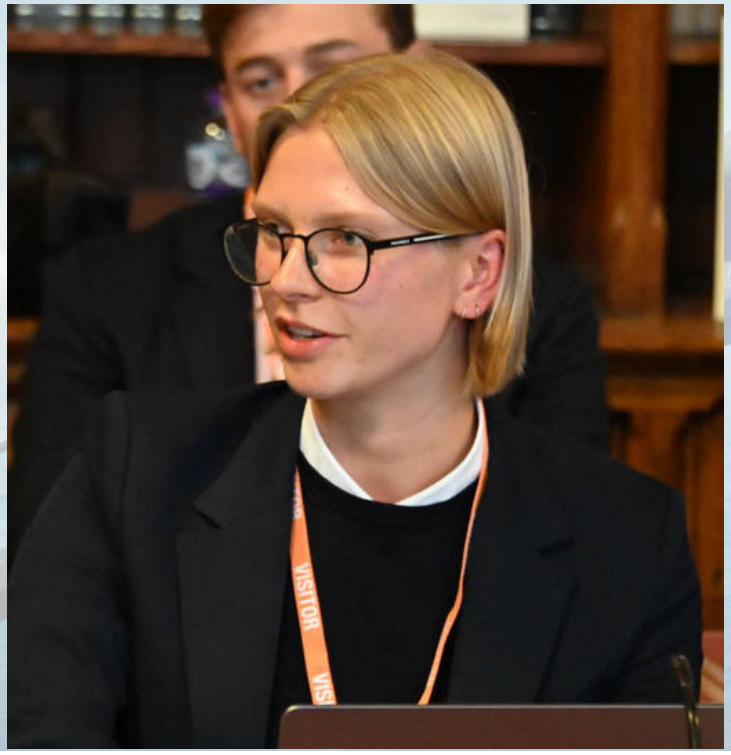

Liberty Hunter
Project Manager & Researcher
Ax:son Johnson Institute for Statecraft and Diplomacy
Critical National Infrastructure (CNI) damage, cognitive warfare, and AI-led solutions

Liberty Hunter is a Project Manager and Researcher at the Ax:son Johnson Institute for Statecraft and Diplomacy and a Research Assistant at the Hertford College Diplomacy Centre. She has completed a BA in History and Politics at Hertford College, University of Oxford and an Mst in History of War at Lincoln College, University of Oxford.

Throughout her time at university, Liberty focused on the history of British maritime strategy – specifically looking at British submarine cable strategy in the 20th century – and its application to contemporary security issues.

She is currently involved in two research projects. The first focuses on the relationship between Critical National Infrastructure (CNI) damage and cognitive warfare, the second looks at British CNI defence including the integration of AI-led solutions into Royal Navy capabilities.

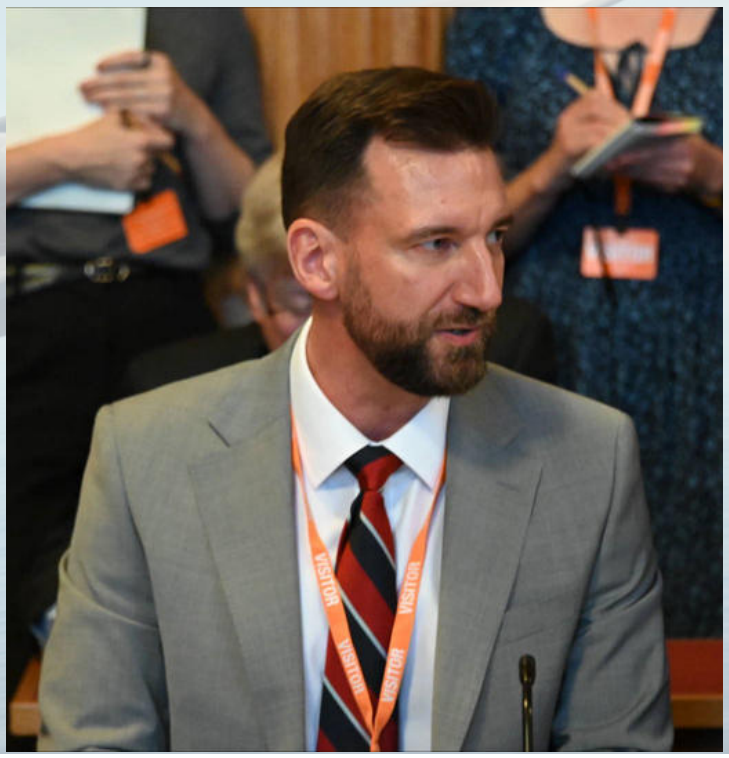

Group Captain Andrew M. Tidmarsh
Deputy Head, Strategy & Plans, UK Strategic Command Headquarters
Royal Air Force
Defence Futures Leader

Group Captain Andrew M. Tidmarsh is Deputy Head of Strategy & Plans at the UK Strategic Command Headquarters, where he helps shape the future force design and operational integration of advanced capabilities across the UK Armed Forces.

With extensive experience in operational command, policy development, and joint service coordination, he is instrumental in ensuring that emerging technologies —including AI—are effectively embedded into military planning, readiness, and multinational interoperability. His leadership ensures that the UK's strategic command structure remains agile, innovative, and capable of addressing the evolving challenges of modern conflict.

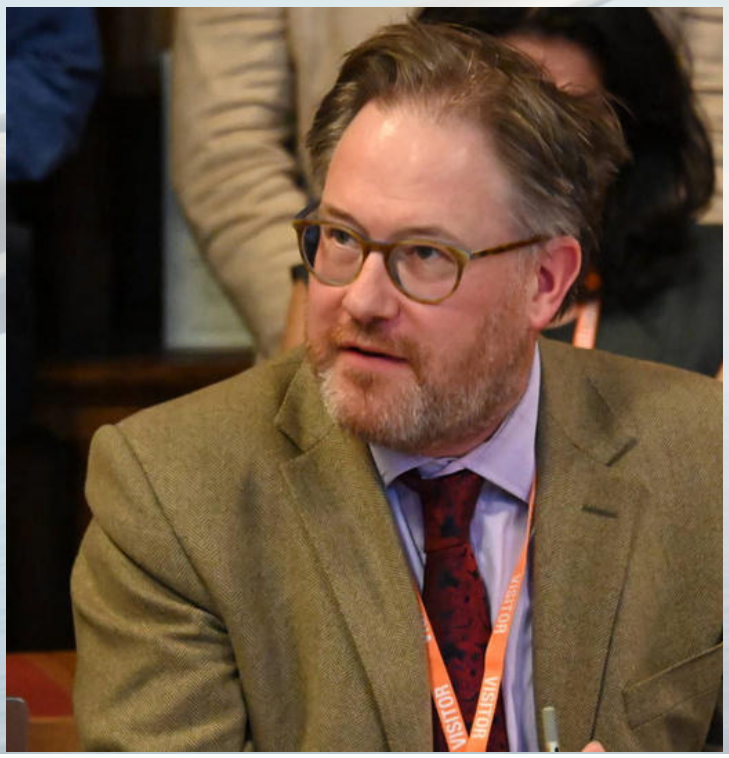

Rob Bassett Cross MC
CEO & Founder
Adarga
Military Cross Recipient | Defence AI Innovator

Rob Bassett Cross MC is the CEO and Founder of Adarga, a leading UK developer of advanced AI software for the defence and national security sectors. A decorated former British Army officer awarded the Military Cross for exemplary gallantry in combat, he combines operational leadership experience with technological vision to transform how defence organisations harness data for strategic advantage.

Under his leadership, Adarga has become a trusted partner to the Ministry of Defence and allied forces, delivering AI systems that enhance intelligence analysis, decision-making speed, and situational awareness. Rob's unique perspective—shaped by service on the front line and in the boardroom—positions him as a leading authority on the operational application of AI in modern defence.

Professor David Whetham
Professor of Ethics and the Military Profession,
King's College London
Director, King's Centre for Military Ethics

Professor David Whetham is Professor of Ethics and the Military Profession in the Defence Studies Department at King's College London and Director of the King's Centre for Military Ethics. He is a leading authority on the ethical and legal frameworks that govern armed conflict, specialising in the Just War tradition, the law of armed conflict, and military ethics education.

He advises and trains armed forces worldwide, ensuring that technological and tactical innovation—such as the integration of AI and autonomous systems—aligns with robust ethical principles and legal safeguards. His work supports defence organisations in navigating the moral complexities of modern warfare.

ABOUT APPG AI

All Party Parliamentary Group on
Artificial Intelligence

ABOUT:

APPGs are informal cross-party groups in the UK Parliament. They are run by and for Members of the Commons and Lords. The All-Party Parliamentary Group on Artificial Intelligence (APPG AI) functions as the permanent, authoritative voice within the UK Parliament (House of Commons and House of Lords) on all AI-related matters, and it has also become a recognisable forum in the AI policy ecosystem both in the UK and internationally.

Parliamentary APPG AI Members: House of Commons

- Allison Gardner MP Labour (APPG AI Co-Chair)
- Alison GRIFFITHS MP Conservative
- Andrew Pakes MP Labour
- Bell Ribeiro-Addy MP Labour
- Chris Kane MP Labour
- Daniel Aldridge MP Labour
- Danny Chambers MP Liberal Democrat
- Dave Robertson MP Labour
- David Reed MP Conservative
- Dawn Butler MP Labour (APPG AI Vice-Chair)
- Esther McVey MP Conservative
- George Freeman MP Conservative
- Gordon McKee MP Labour
- Graham Leadbitter MP SNP
- Liam Byrne MP Labour
- Mike Martin MP Liberal Democrat
- Martin Wrigley MP Liberal Democrat
- Maureen Burke MP Labour
- Peter Fortune MP Conservative
- Samantha Niblett MP Labour
- Sarah Edwards MP Labour
- Tom Collins MP Labour
- Tom Gorden MP Liberal Democrat
- Tony Vaughan MP Labour
- Sir Mark Hendrick MP Labour
- Zöe Franklin MP Liberal Democrat
- Dr Zubir Ahmed Labour

Parliamentary APPG AI Members: House of Lords

- Lord Clement-Jones (Tim Clement-Jones) Liberal Democrat (APPG AI Co-Chair)
- Viscount Camrose (Jonathan Camrose) Conservative
- Viscount Colville Of Culross (Charles Mark Townshend Colville) Crossbench
- Lord Craig of Radley (David Brownrigg Craig) Crossbench
- Lord Cromwell (Godfrey Cromwell) Crossbench
- The Earl of Erroll (Merlin Hay) Crossbench
- Lord Fairfax of Cameron (Nicholas Fairfax) Conservative
- Lord Freyberg (Valerian Bernard Freyberg) Crossbench
- Lord Strathcarron (Ian David Patrick Macpherson) Conservative
- Lord Janvrin (Robin Berry Janvrin) Crossbench
- Baroness Kramer (Susan Veronica Kramer) Liberal Democrat
- Baroness McGregor-Smith (Ruby McGregor-Smith) Non-affiliated
- Lord Ranger of Northwood (Kulveer Ranger) Conservative (APPG AI Vice-Chair)
- The Lord Bishop of Oxford Stephen Croft Bishops
- Lord Pitkeathley (Simon Pitkeathley) Labour
- Viscount Stansgate (Stephen Stansgate) Labour
- Professor Lord Tarassenko (Lionel Tarassenko) Crossbench
- Lord Taylor of Warwick (John David Beckett Taylor) Non-affiliated (APPG AI honorary Vice-Chair)
- Baroness Uddin (Manzila Pola Uddin) Non-affiliated

All Party Parliamentary Group on
Artificial Intelligence

THANK YOU TO OUR SUPPORTORS

Helping Us Raise Our Ambition for What Can Be Achieved

ACCESS APPG AI RESOURCES, EVENTS AND FULL PROGRAMME

Pavilion proudly hosts the All-Party Parliamentary Group on Artificial Intelligence (APPG AI), providing a centralised hub for all its resources, including publications, event registrations, and more.

Annual Programme

At least 6 Round Table Evidence Sessions.
4 Advisory Board Meetings.
Special Policy Briefings.

Networking

All events are held in the UK Parliament and chaired by the APPG AI Co-Chairs and the Parliamentarians.

Resources

Reports, transcripts, videos, and photo albums.

Download your Pavilion App Now!

Go to APPG AI Pavilion and click on what you are looking for.

From your computer:

Pavilion on PC website: https://bicpavilion.com/

From your mobile:

Pavilion on App Store https://apple.co/4dCawaW
Pavilion on Google Play https://bit.ly/44Da6N3

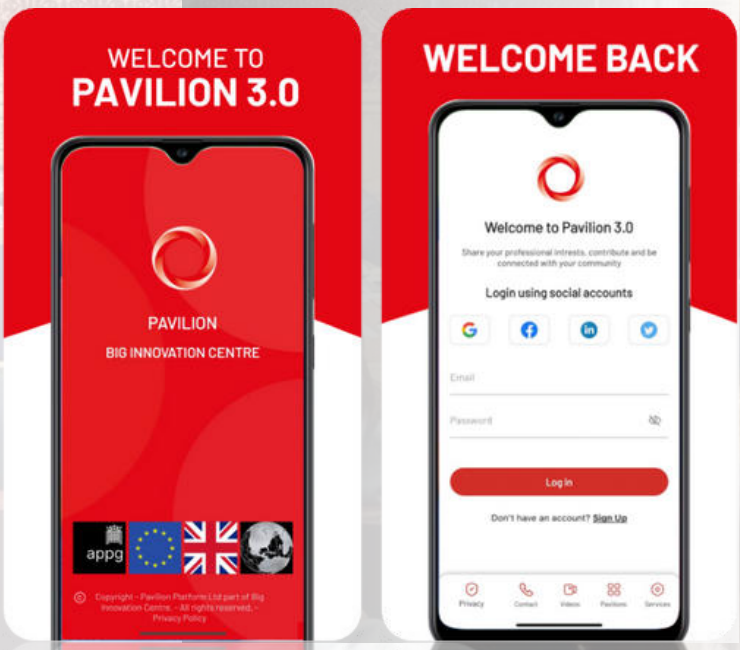

Please use the same username and password across all web and mobile app devices, avoiding the hassle of multiple accounts.
Click below:

visit our website bicpavilion.com

Download on the App Store

GET IT ON Google Play

All Party Parliamentary Group on
Artificial Intelligence

CONTACT

Secretariat:
Big Innovation Centre is appointed as the Group's Secretariat.

The Secretariat is responsible for delivering the programme for the APPG AI, organising the outputs, advocacy and outreach, and managing stakeholder relationships and partnerships.

Contact:
Professor Birgitte Andersen, CEO, Big Innovation Centre
appg@biginnovationcentre.com

All-Party Parliamentary Group on
Artificial Intelligence
appg@biginnovationcentre.com

SECRETARIAT

Big Innovation Centre is appointed by the UK Parliament as the Group's Secretariat.

BIG INNOVATION CENTRE