



All Party Parliamentary Group on
Artificial Intelligence

AI AND TECHNOLOGY SOVEREIGNTY

**STRATEGIC DEPENDENCE &
NATIONAL CAPABILITY**

MAY 2026
POLICY FORUM



**BIG
INNOVATION
CENTRE**

APPG AI Secretariat

Table of Contents

Rapporteur Preface	3
Introduction	4
Aim of Session	5
Findings	8
Evidence	16
• Professor. Ashley Braganza, Brunel University of London	17
• Dr Neel Savani, Cognizant	20
• Dan Patten, Scale AI	23
• Dr Susan Oman, University of Sheffield	27
• Leo Rogers, Curvo AI	31
• Joséphine Kant, Sovereign AI Unit, UK Government	35
Discussion	40
Biographies of Evidence Givers	44
About APPG AI	49
Contact	52



appg@biginnovationcentre.com
APPG AI Secretariat



Rapporteur's Preface

Artificial intelligence is becoming a foundation of economic competitiveness, public services and national capability. Yet much of the infrastructure on which AI depends, from compute and cloud to foundation models, is concentrated in a small number of companies and jurisdictions outside the UK. This creates a fundamental strategic question for Britain: how can we remain open to global innovation while retaining sufficient capability and control to shape our own AI future?

The answer is not technological self-sufficiency. In an interconnected global AI ecosystem, that would neither be realistic nor necessarily desirable. The more important question is where the UK must retain capability, resilience and strategic choice, and where international interdependence risks becoming strategic vulnerability.

This APPG AI evidence session examined these issues across the AI stack, from compute, cloud and models to data, investment, skills and public trust. The evidence demonstrated that sovereignty is not simply about where technology is built or hosted. It is about who controls critical capabilities, who governs the data, where economic value is created and retained, whether British companies can scale, and whether the UK can continue to act in its own national interest when circumstances change.

The discussion also highlighted significant UK strengths. Britain has world-leading universities and research, exceptional public data assets, strong institutions, innovative AI companies and internationally recognised capabilities in AI safety, assurance and governance.

The strategic task is to connect these strengths more effectively and translate them into enduring national capability and global competitive advantage.

A particularly important message from the meeting was that **openness and sovereignty should not be treated as opposites**. The UK can remain an open, internationally connected technology economy while being much more deliberate about the capabilities it develops, the dependencies it accepts and the strategic assets it protects.

The evidence presented in this report therefore points towards a practical conception of AI sovereignty: **not isolation, but the ability to exercise strategic choice - to understand our dependencies, build strength where the UK can create strategic advantage, maintain meaningful alternatives, and ensure that AI ultimately serves the UK's prosperity, resilience, public institutions and citizens.**

Professor Birgitte Andersen

Rapporteur

CEO, Big Innovation Centre

Secretariat, All-Party Parliamentary Group on Artificial Intelligence

Introduction

About this Evidence Session

INTRODUCTION

This report includes a structured transcript and summary of the principal evidence and discussion at the APPG AI evidence meeting held on 18 May 2026 in the House of Lords, Committee Room 1, UK Parliament.

It captures the key elements of the evidence and discussion and is not intended to provide a complete verbatim record of all contributions.

DETAILS

- Evidence Session: AI and Technology Sovereignty: Strategic Dependence & National Capability
- Time: 5:30pm – 7:00pm
- Date: Monday 18 May 2026
- Venue: UK Parliament: House of Lords – Committee Room 1

CONTACT THE SECRETARIAT

appg@biginnovationcentre.com
APPG AI Secretariat
BIG INNOVATION CENTRE

EVIDENCE GIVERS

- **Prof. Ashley Braganza**, Director, Centre for Artificial Intelligence, Brunel University of London
- **Dr Neel Savani**, Responsible AI, AI Labs, Cognizant
- **Dan Patten**, Government Affairs Lead (UK and Europe), Scale AI
- **Dr Susan Oman**, Senior Lecturer in Data, AI and Society, University of Sheffield (Centre for Machine Intelligence)
- **Leo Rogers**, CEO & Co-Founder, Curvo AI
- **Joséphine Kant**, Head of Ventures, Sovereign AI Unit, UK Government

APPG AI CO-CHAIRS

The All-Party Parliamentary Group on Artificial Intelligence (APPG AI) is co-chaired by **Allison Gardner MP** and **Lord Clement-Jones CBE**. Lord Clement-Jones CBE chaired this evidence meeting held on 18 May 2026.

RAPPORTEUR

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

AI and Technology Sovereignty

Strategic Dependence & National Capability



From left to right: David Krucik (OC&C Strategy), Dan Patten (Scale AI), Ben Johnson (Uptitude), Lord Tarassenko (House of Lords), Prof. Birgitte Andersen (APPG AI Secretariat, Big Innovation Centre), Dr Susan Oman (Univ. of Sheffield), Lord Ranger (House of Lords), Lord Taylor (House of Lords), Lord Tim Clement-CBE (House of Lords), Joséphine Kant, (UK Government), Prof. Ashley Braganza (Brunel Univ. of London), Dr Neel Savani (Cognizant), Leo Rogers (Curvo AI), Moushami Verma (Santander UK), Laura Bishop (British Standards Institution - BSI)

Aim of Session

This session examined how the UK's dependence on global AI infrastructure, particularly US hyperscalers, compute, cloud, foundation models and capital, affects national resilience, strategic autonomy and the UK's ability to capture economic value. It explored where critical dependencies and vulnerabilities lie across the AI stack, where the UK should build or strengthen sovereign capability, and how this can be achieved while remaining open to international collaboration, investment and innovation.

Discussion Covered

- Strategic Dependence, Resilience & AI Sovereignty
- US Hyperscalers & Concentration across the AI Stack
- Building UK Sovereign Capability & Strategic Advantage
- Data Sovereignty, Public AI & Trusted Infrastructure
- UK Startups, Investment, Compute & Public Procurement
- AI Governance, Regulation, Standards & Assurance
- Public Trust, AI Literacy & Citizen Participation
- Openness, International Collaboration & Strategic Choice

Setting the scene:

AI and Technology Sovereignty

I'm absolutely delighted that we've got a session today on AI and technology sovereignty, which I think has become really quite a big issue over the past few months.



Lord Tim Clement-Jones CBE

*Co-Chair, All-Party
Parliamentary Group on
Artificial Intelligence*



UK Ambition and the Risk of Dependency

The UK is rightly ambitious about being a leading AI nation, but ambition without capability risks becoming what we might describe as dependency across the AI stack, whether it's compute and cloud infrastructure, foundation models, data and talent. We're deeply embedded in a global ecosystem that increasingly is concentrated, particularly in the U.S. And I think that's one of the areas where people are nervous.

That brings enormous benefits, access to world-class technology, investment, and innovation, but it also raises serious questions about resilience, autonomy, and where the value ultimately accrues.

And if the UK is primarily a consumer, rather than a producer of advanced AI systems, we risk not only economic leakage, but diminished strategic control over technologies that will shape our public services, our security, and our democratic processes.



Openness Must Be Matched by Capability

But this is not about digital isolationism or retreating from international collaboration. Our strength in the UK has always been as an open, connected economy. But openness must be matched by capability, in my view. Sovereignty in this context means having meaningful choices, ensuring we're not locked in, not overly dependent, and not unable to act in our own national interest.



Identifying the UK's Critical Dependencies

So today, we wanted the APPG to explore some key questions. Where are our critical dependencies in compute, in cloud, in data infrastructure, in intellectual property, and how exposed are we?

What would happen if access were constrained? through commercial, geopolitical or regulatory pressures?



What Would a Credible UK Sovereign Capability Look Like?

What does a credible UK sovereign capability actually look like, not necessarily building everything ourselves, but being strategic about the layers that matter most, whether that's advanced compute, trusted data environments or key parts of the model ecosystem?



Policy Levers to Support Domestic Capability

And crucially, what policy levers do we have, from procurement and standards to investment and regulation? shape markets to support domestic capability while remaining globally competitive.



Open Source, Regulation and Sovereignty

And we should also consider the role of open source as a potential counterbalance to concentration and whether the UK's current regulatory approach is sufficiently clear and purposeful to support both innovation and sovereignty.



From Diagnosis to Action

Now, we're extremely fortunate to have a very distinguished panel bringing perspectives from academia, industry, and policy. And I look forward to a very rigorous and practical discussion focused not just on diagnosing the challenge, but on identifying actionable steps for Parliament and government.



FINDINGS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

Executive Summary

The meeting examined what **AI sovereignty should mean for the UK in practice**, at a time when critical elements of the AI ecosystem - including compute, cloud infrastructure, foundation models, data and investment - are increasingly concentrated in a small number of global technology companies and jurisdictions.

The speakers approached AI sovereignty from complementary but distinct perspectives: Prof. Ashley Braganza emphasised governance and public AI; Dr Neel Savani focused on infrastructure resilience and technological dependencies; Dan Patten on sovereign data and trusted AI; Dr Susan Oman citizens, trust and AI literacy; Leo Rogers the ability of British startups to commercialise and scale; and Joséphine Kant emphasised strategic state investment and the creation of a UK “sovereign edge”.

Together, their evidence suggested that **AI sovereignty is not a single technological objective, nor does it require technological isolation or complete self-sufficiency. It is a multidimensional challenge spanning governance, infrastructure, data, society, markets and state capability.** The central question is where the UK needs sufficient capability, control and strategic leverage to retain meaningful choices and avoid critical dependencies.

Expert Evidence Main Contributions



Governance and Public AI **Prof Ashley Braganza**

Britain cannot realistically outspend or outcompute the US and China at the AI frontier, but it has an opportunity to “outgovern” them. Braganza argued that sovereignty requires clearer definitions, stronger governance of technological dependencies, and a public-interest dimension alongside commercial AI.

- Establish a clear UK sovereignty baseline, including infrastructure control, data residency, IP ownership and exit provisions.
- Develop public AI alongside commercial AI, with resources available to public institutions and civil society.
- Treat citizen-generated public-sector data as a strategic national asset, governed for public benefit.



Infrastructure and Resilience

Dr Neel Savani

Dr Neel Savani focused on resilience across the whole AI technology stack. Even a public-interest model may fail the sovereignty test if the UK cannot run, audit or evolve it without dependence on a particular cloud provider, chip architecture or technology vendor.

- Identify choke points and dependencies across the AI stack, not simply at the model layer.
- Reduce single points of failure and understand exposure to changes in pricing, export controls and service terms.
- Use procurement, assurance and strategic investment to strengthen interconnected UK capabilities.



Data Sovereignty and Trusted AI

Dan Patten

Dan Patten identified the data layer as one of Britain's strongest potential sovereign advantages. The UK's exceptional public data assets could underpin AI systems that better reflect British laws, values and society.

- Treat assets such as NHS and other public data as strategic sovereign infrastructure.
- Ensure consequential AI is based on data that is appropriately governed, auditable and traceable.
- Compete not necessarily by building the world's biggest model, but by making Britain a leader in the world's most trusted AI.



People, Trust and AI Literacy

Dr Susan Oman

Dr Susan Oman argued that AI sovereignty is also a **societal question**. It cannot be understood solely through infrastructure, models and investment without considering how people understand, trust and use AI, who benefits from it and who may be placed at risk.

- Make people and communities part of sovereign AI strategy.
- Recognise **AI literacy as national infrastructure**, reflecting UK law, society and needs.
- Build public trust through inclusive engagement, transparency and meaningful participation.



Startups, Commercialisation and Procurement

Leo Rogers

Leo Rogers approached sovereignty from the perspective of a British AI entrepreneur. For UK companies, sovereignty also depends on whether they can start, commercialise and scale successfully from Britain.

- Open public procurement more effectively to innovative UK startups.
- Improve access to compute, R&D support and specialist early-stage funding.
- Create conditions that enable British companies and talent to remain and scale in the UK rather than moving towards Silicon Valley for capital and commercial opportunities.



Strategic Leverage and the UK's "Sovereign Edge"

Joséphine Kant

Joséphine Kant presented the government investment perspective and rejected the idea that sovereignty requires complete technological self-sufficiency. Instead, Britain should seek a "sovereign edge" by concentrating resources where the UK already possesses capabilities, depth and international relevance.

- Focus on strategic leverage rather than complete self-sufficiency.
- Combine venture investment, sovereign compute, strategic data, procurement and talent to strengthen UK capability.
- Help British AI companies start here, scale here and compete globally, while maintaining Britain's openness to international technology and investment.

A Shared Direction

Across the evidence, a common position emerged: UK AI sovereignty should be about strategic choice rather than isolation. Britain does not need to own or reproduce every component of the global AI stack, but it does need to understand where its critical dependencies lie and ensure that it cannot easily be locked out, priced out or left unable to govern technologies essential to public services, economic prosperity and national resilience.

The evidence therefore pointed towards a sovereign AI strategy built around **strategic capability, resilient infrastructure, trusted data, effective governance, public trust, stronger commercialisation and purposeful state investment and procurement.**

Implications for Stakeholders

The evidence has implications across government, the public sector, industry, research and society. Taken together, the discussion suggests that AI sovereignty will require coordinated action across these groups, rather than a standalone government technology policy. Implications include:



Government and Policymakers

- Define more clearly what constitutes **UK sovereign or British-anchored AI**.
- Map critical dependencies across compute, cloud, chips, models and data.
- Concentrate investment where Britain can develop a genuine “**sovereign edge**”, rather than seeking complete technological self-sufficiency.
- Use **procurement, compute access and strategic data assets** more deliberately to strengthen domestic capability.
- Balance economic growth with resilience, public benefit and the ability of the UK to act in its own national interest.



Parliament and the Public Sector

- Scrutinise critical AI dependencies and the long-term implications of major technology procurement decisions.
- Seek appropriate transparency over **data provenance, infrastructure dependencies and AI supply chains**.
- Ensure public bodies retain sufficient ability to **audit, govern and change** consequential AI systems.
- Consider whether AI used in public services appropriately reflects UK law, rights and societal context.
- Create credible procurement opportunities for capable UK startups and SMEs.



Industry, Startups and Investors

- Strengthen access to **compute, capital, customers and public-sector markets** so promising UK AI companies can scale from Britain.
- Design AI systems around transparency, interoperability and assurance, reducing unnecessary technological lock-in.
- Recognise **trusted and sovereign AI capabilities as potential competitive advantages**.
- Strengthen connections between startups, investors, government customers and the UK's research base.
- Create conditions that make it possible for strategically important companies and talent to remain anchored in the UK as they grow.



Universities and the Research Community

- Strengthen pathways from **research to commercialisation and scale**.
- Develop approaches that reduce unnecessary dependence on individual vendors and technology stacks.
- Build further UK leadership in **AI safety, evaluation, assurance and governance**.
- Support responsible use of strategic UK data for research and innovation.
- Ensure AI education and literacy initiatives reflect UK law, institutions and societal needs.



Regulators and Standards Bodies

- Strengthen proportionate transparency and assurance around consequential AI systems and their supply chains.
- Promote **interoperability, auditability and resilience** while avoiding unnecessary barriers to innovation.
- Build on UK strengths in AI evaluation and assurance as a potential area of international leadership.
- Maintain a balance between accountability, public protection and an internationally competitive innovation environment.



Citizens and Civil Society

- Treat **AI literacy and public understanding as part of national AI infrastructure**.
- Involve citizens and communities in discussions about consequential uses of AI.
- Ensure fairness, accountability and public benefit remain central alongside economic growth.
- Recognise that **public trust may itself become part of Britain's competitive advantage** in trusted AI.

Overall Implication

The meeting pointed towards a conception of **AI sovereignty based on strategic choice, resilience and agency (i.e. the ability to make and act on decisions in the UK's own interests)** rather than technological isolation.

The UK does not need to build or own every component of the AI stack. It does, however, need to understand its critical dependencies, maintain credible capabilities and alternatives where they matter most, and retain sufficient control to govern technologies affecting its economy, public services and citizens.

The challenge is therefore to combine **openness to global technology, investment and collaboration with sufficient domestic capability to protect UK resilience, prosperity and democratic agency**.

EVIDENCE

EVIDENCE GIVERS

- Prof. Ashley Braganza, Brunel Univ. of London
- Dr Neel Savani, Cognizant
- Dan Patten, Scale AI
- Dr Susan Oman, Univ. of Sheffield
- Leo Rogers, Curvo AI
- Joséphine Kant, UK Government





Ashley Braganza

Director, Centre for Artificial Intelligence,
Brunel University of London

Britain cannot outspend or outcompute the US and China at the AI frontier, but it can “outgovern” them by defining sovereign AI clearly, treating AI and citizen-generated data as public infrastructure, and making technological dependencies transparent and governable.

A Consequential Moment for Britain and AI

We are at a consequential moment. Artificial intelligence is becoming the infrastructure through which economies function, public services are delivered, and national power is exercised. The decisions this country makes over the next few years will, I believe, shape Britain's AI position for a generation.

Britain as an AI Middle Power: The Opportunity to Outgovern

My argument is fairly straightforward. Britain is a middle power in a world being shaped by two AI superpowers, the U.S. and China. I don't believe we can outspend them, and I don't believe we can outcompute them at the frontier. But what I do think we can do is outgovern them. And right now I'm not sure we are doing so fully.

Defining Digital Sovereignty

So let me start with definitions. I feel as though this word ‘sovereignty’ is doing an awful lot of work in policy debates and I'm not sure we all mean the same thing when we use the word. A systematic review of 271 peer-reviewed studies by researchers at Oxford, Bologna, and Yale identify 4 distinct models of digital sovereignty.

The UK sits in the market-oriented model, light regulation, commercial competition, and innovation through market dynamics. This model has generated world-class companies, for example, Arm and DeepMind. What I don't believe it generated is digital sovereignty, because both those companies are now foreign-owned. The market-oriented model left to itself builds things and then sells them.

The Emergence of “Sovereignty as a Service”

Meanwhile, I think there is a concept which is starting to emerge, which I think we should all be a bit concerned about, which is sovereignty as a service. Companies like Microsoft, Amazon, Anthropic, and Google have responded to the pressure for digital sovereignty by rebranding some of their products as sovereignty solutions, offering data residency, audit layers and compliance certificates. The result is that sovereignty is provisioned by the very platforms on their terms. You can buy what I believe is an appearance of sovereignty from the very companies that are the source, Tim, as you said, of your dependency. Britain must make sure that they don't fall into that trap.

Four Propositions for Sovereign AI

So what does sovereignty in AI look like? And I have 4 propositions I'd like to put forward.

1. Define What British AI Means

The first is to define what we mean. The AI Sovereign Fund is genuinely a welcome development. half a billion of investments, supercomputer access, fast track visas, and government as the first customer. These are really important policy developments. But the fund has no apparent working definition of what British AI actually means. Not one that I could find. A company registered in the UK, running on AWS, training on American APIs with a Delaware parent. Is that sovereign AI? We need, in my view, a published sovereignty baseline. What are the minimum conditions around infrastructure control, data residency, IP ownership, and exit provisions as examples? Without it, we risk funding a generation of British startups that will be acquired before they deliver national benefit. We've done this before, and I don't think we should do it again.

2. Introduce Public AI Alongside Commercial AI

The second is to introduce public AI alongside commercial AI. The Sovereign AI Fund, if you ask me, is building one engine, which is around commercial competition. Other countries are building a second engine. Switzerland, for instance, released Appertis, a fully open, publicly funded language model free to any hospital, school, or library, i.e. the Swiss public, to deploy. AI, in my view, should be like public infrastructure, a bit like the highways, the railways, and energy.

Sovereign AI has a strategic dimension which addresses who controls the infrastructure. But I think it also has a civic dimension that asks who benefits from it. Current policy addresses the first. It does not, in my view, address the second.

Citizen-Generated Data as a National Strategic Asset

In my work with my colleague at Brunel, Dr. Tabaghetti, we argued that citizen-generated data from the NHS, local authorities, public services must be treated as a national strategic asset and not given away to foreign-owned entities. A sovereign data fund that governs and monetizes citizen data ethically with profits reinvested for public benefit is one concrete way to fire up that second engine. Public AI resources should be available to NHS trusts, universities, libraries, the third sector, and not just to commercial startups.

3. Regulate to Reveal Dependency

Third, regulate to reveal dependency. And what I mean by that is the EU Act requires AI supply chain transparency, audit rights, documentation of training data, and infrastructure provenance. Britain's current regulatory approach doesn't seem to do that. And without equivalent requirements, we cannot know what we depend on. We cannot audit systems deployed in our public services, and I don't believe we can claim to govern that what we cannot see. So high-value, high-risk AI procurement should require disclosure in terms of infrastructure dependencies, training data provenance, and foreign ownership. As Tim said earlier, we cannot stop dependency, but we can make it visible so it can be governed.

4. Sovereignty Is a Direction of Travel

And lastly, sovereignty is a direction of travel rather than a destination. We need sovereign capability at the layers that matter for public life, around deployment, fine-tuning, inference and civic application. This doesn't require winning the frontier race, but it does require policy choices.

From a Sovereign AI Venture Fund to a Sovereign AI Institution

The Sovereign AI Fund is the beginning of something really important, but I think the question for parliamentarians is whether it becomes Britain's sovereign AI venture fund or Britain's sovereign AI institution.

AI as Public Infrastructure

If I go back a little bit, the NHS was built because we decided health care was too important to be rationed by ability to pay. Public libraries were built because we decided access to knowledge should not depend on wealth. Public institutions were built because the market alone could not serve every citizen equally. And I think AI now falls into that category.

Ensuring Sovereign AI Serves British Citizens

The choices are not around sovereignty or openness, innovation or regulation, public good or commercial success. I think we must free ourselves from the tyranny of awe. The question, I think, is around whether the UK's sovereign AI strategy that we build serves British citizens and the interests of companies that we the public are funding to design it.

And I'd just like to commend both the APPG AI for facilitating this discussion. And I'd also like to acknowledge that Dr Neel Savani (next speaker) and I are doing some very exciting work in this area.

Thank you.



Dr Neel Savani
Responsible AI, AI Labs, Cognizant

UK AI sovereignty depends on resilience across the entire technology stack - not simply the models we use - and requires Britain to identify critical choke points, avoid single-vendor dependency, and use procurement and investment to strengthen strategic domestic capabilities.

Introduction and Background

My name is Neel Savani. I'm proudly British and I've graduated from Imperial College three times. I currently lead our Responsible AI Center of Excellence at Cognizant and work within our Responsible AI Office to turn governance, safety research and standards into deployable practice. Currently, I serve on the UK Government's Digital Sustainability Alliance AI Working Group and remain an Affiliate Professor at the University of Maryland, Baltimore County. My earlier career was as a research scientist at NASA Goddard and included areas around future mission design and space situational awareness for government policy.

Strategic Dependence on Vertically Integrated AI Infrastructure

Today, I've chosen to speak about strategic dependence on vertically integrated AI infrastructure and what that implies for UK autonomy. It sits directly under the terms of reference: where dependence concentrates, what breaks when access or pricing shifts, and what is our credible UK response.

AI Sovereignty Is About the Entire Technology Stack

Today, the central point that I would like to make is that sovereignty for AI is not only about the model we buy; it's about whether we can choose, govern or sustain the compute, the data, the models and the assurance over time without a single vendor's silicon, compute, middleware or cloud defining our economics and our risk.

The Shift from the Model Race to Infrastructure

Recent industry commentary, much of it from the US, describes a pivot from the pure model race towards an infrastructure-led strategy, where energy, connectivity and hyperscale compute are bundled so tightly that the long-term advantage sits below the application layer.

And it is this pattern that matters to this committee. Whoever controls the physical stack can tax, prioritise and constrain everyone else, including the frontier labs, which still depend on someone else's machines.

Vertical Integration Across Compute and Energy

For example, if you look at this year's announcement with SpaceX and the Colossus 1 data centre, in the order of 200,000 GPUs and roughly 300 megawatts, they are reportedly providing dedicated capacity to Anthropic. And it illustrates the importance of vertical integration across the compute and energy layers, and not just the application and model layers.

In this scenario, a potential future could emerge where a lockout at the governance layer is realised, where parliaments do not set the prices or the priority rules, but commercial contracts rule.

GraphCast: A Practical Example of Infrastructure Dependency

And this pattern is not abstract to me. During this last year, my research team and I have carried out our work and research using DeepMind's GraphCast weather model. We set out to validate the benefits to the UK economy in a way that would be useful all the way through the spectrum, not just to sole traders, but all the way to the farming industry, SMEs and large renewable operators.

But as we pushed forward to create that infrastructure-agnostic deployment, the story dramatically changed. Outside of the Google stack, we saw that material cost and latency penalties existed. And so that design pattern created a deep coupling to the particular cloud, its particular custom TPU chips and its compute infrastructure.

The Sovereignty Test: Can the UK Run, Audit and Evolve AI?

The policy lesson was relatively simple. A brilliant public-interest model can still fail that sovereignty test if the UK cannot run, audit or evolve it in a contestable manner. That's why I align with the public AI framing rather than the buying-more-licensing framing.

Three Approaches to AI Infrastructure

The white paper we've developed at Cognizant distinguishes three postures: **commercial concentration, fast but structurally dependent; state monopoly, autonomous but brittle to transparency; and public services AI, autonomy with democratic accountability and inclusive access.**

For a middle power such as the United Kingdom, the third path is the honest centre of gravity, in my opinion.

A Seven-Layer Public AI Stack

Operationally, that path is organised by Cognizant as a seven-layer public AI stack, with the layers tightly intertwined and choke points emerging as we move between adjacent layers.

Designing for Resilience: Lessons from NASA

And it is my NASA experience in designing missions that reinforces one discipline I think Parliament should consider. I've learned the importance of minimising single points of failure to maximise resilience.

Editorial explanation

Dr Savani is setting out three alternative models or “postures” for organising AI infrastructure / governance:

1. **Commercial concentration** — fast, but structurally dependent. This is essentially the current hyperscaler-led model: private companies provide the infrastructure and capability, allowing rapid innovation and deployment, but creating dependence on a small number of commercial providers.
2. **State monopoly** — autonomous, but brittle to transparency. The state owns or controls the AI infrastructure, providing greater national autonomy, but potentially creating problems around openness, scrutiny, flexibility and accountability.
3. **Public services AI** — autonomy with democratic accountability and inclusive access. This is the model Dr Savani favours for the UK: sufficient public/national capability to avoid excessive commercial dependence, but with democratic governance and broad access rather than a state-controlled monopoly.

Three Calls to Action for Parliament

And so, if I may, I would like to leave the audience with three calls to action that I think are important.

One, identify those choke points across the layers, not just the model layer. Map where the UK's strengths are and how it could be locked in or locked out if pricing, export controls and service terms change and shift.

Second, leverage the procurement process to strengthen our stack. Create and apply minimum assurance cases and risk scoring as part of that contract process.

And third, build upon our unfair advantage, much like the Sovereign AI Fund is trying to do. Define our current layers of strength, financially support those innovations and strengthen those interconnected layers. Create strategic resilience by doing so.

And maintain participation across industry, government and research, much as I can continue doing research with Ashley.

Strategic Dependence as a Portfolio Risk

And it is this strategic dependence on foreign hyperscalers. It's not an indictment of global trade. It's a portfolio risk. And in that regard, the UK can narrow that risk without closing the world off.

Thank you.



Dan Patten

Government Affairs Lead (UK and Europe), Scale AI

“**Britain’s greatest opportunity for genuine AI sovereignty lies in its data: governing sovereign assets such as NHS and other public data domestically could ensure AI reflects British laws and values while making trusted, auditable AI a UK competitive advantage.**”

Who Controls AI Decisions About British Citizens?

AI is increasingly making or directly informing decisions about British people, about their health care, their benefits and their access to public services. The question I want to put to this APPG is a simple one. Who controls those decisions and on whose behalf are they being made?

Scale AI: Building and Evaluating AI Systems

My name is Dan Patten and I lead policy for Scale AI here in London. Scale AI's mission is to develop reliable AI systems for the world's most advanced AI models. We provide the data infrastructure that underpins the world's most advanced, sorry, and we build and deploy AI applications for governments and enterprises, for national public sector programs and Fortune 100 companies.

Through Scale AI Labs, we are the world's leading independent AI evaluator, measuring whether AI models actually do what they claim to do and for whom. That combination of building systems and testing them gives us a distinctive view of where sovereign AI works and where it doesn't.

What Does Sovereign AI Actually Mean?

Sovereign AI has become a term that covers a great deal of ground, and it's worth being clear about what the term is actually being used to mean, because different things are being conflated in ways that matter for policy.

Sometimes it means security, national resilience, ensuring foreign actors cannot degrade critical systems at a moment of crisis.

Sometimes it means economic participation, ensuring that productivity gains of AI accrue to British industries and workers, rather than flowing to a handful of overseas companies.

And sometimes it means a political signal, the visible act of building something at home, satisfying the legitimate public expectation that the country is shaping its own future.

All 3 matter, but behind them is a question that rarely gets addressed directly.

Do AI Systems Reflect British Laws, Society and Values?

When AI systems are making consequential decisions about British people, do those systems actually reflect British society laws and values? The standards and premises governing a decision made in this jurisdiction are distinct from those guiding the same decision elsewhere.

That is the dimension I want to focus on, because it is simultaneously one of the most important and the most underinvested.

Sovereignty as Choices Across the AI Stack

To understand where the UK should focus, It helps to think about sovereignty not as a binary, but as a set of choices made at every layer of the AI stack, from energy, compute, data, models, application and talent.

At each layer, a government can choose to control directly, to steer through regulation, investment and partnership, or to consciously depend on external actors. None of these are universally right or wrong. What matters is that the choices are deliberate, the risks are understood and there is a plan for disruption.

The UK's Structural Dependency on Data

If we apply this measure to the UK, on the computer model layers the government is rightly investing to steer through partnership with companies whilst to a certain degree setting the terms. But on data, the layer that determines what values, norms and assumptions get encoded into every AI system built on top of it, the UK has decided to depend. And unlike compute, which can be procured or models which can be switched, data dependency is structural. You cannot retrofit British values into a system trained on someone else's assumptions of the world.

This is why the data layer is not just the most important layer for the UK. It is the one where the UK has the greatest capacity to exercise genuine control and the most to gain from doing so.

AI in British Public Services and the Risk of Foreign Assumptions

The AI systems currently deployed in British public services were predominantly trained on internet data fine-tuned, whether by design or by distribution, to American legal norms, cultural assumptions and patterns of life. For many AI systems, that doesn't necessarily matter, but for systems making decisions about NHS patients, benefit entitlements or planning applications, systems that are expressing something about what British society values and how it treats its citizens, it matters enormously.

NHS Values, Accountability and Data Provenance

The NHS operates under a statutory constitution with patient rights and values. Do the AI tools being used by the NHS operate under the same values and rights?

When those systems produce a wrong outcome, the accountability trail goes cold the moment the training data sits outside British jurisdiction. No minister can answer a parliamentary question about a decision they cannot trace. No ombudsman can investigate harm they cannot source.

Democratic accountability for AI assisted decisions requires data provenance, and data provenance requires the data to be here, governed and auditable.

Underrepresented Communities and Systematic Bias

The community that depends most heavily on NHS care, on welfare and public housing decisions, older people, people in rural areas, ethnic minorities, are also the communities most systematically underrepresented in the data those systems were trained on.

The bias is not random. It runs in a predictable direction, away from the people the system is supposed to serve.

Sovereign Data as an Exportable Asset

But this is not a defensive argument. If the UK governs its data well, if it builds its AI on NHS records, on its legal corpus and its cultural and linguistic heritage, it is not merely protecting its citizens from foreign assumptions.

It is encoding something the other countries, particularly those that share British legal traditions and democratic norms, may want to build on. Sovereign data governance done properly is an exportable asset.

Britain's Competitive Advantage: Building the Most Trusted AI

The UK's AI competitive advantage does not have to rest on building the biggest model, and nor should it. can rest on building the most trusted one.

Trusted because it's been built on data that is governed, auditable and reflective of the society it serves.

The UK Already Owns World-Class Data Assets

The UK already owns what it needs. The NHS data set 7 decades of universal health records covering an entire national population. There is no equivalent anywhere in the world.

HMRC, the BBC Archive, Ordnance Survey, these are sovereign, world-class, already paid for, and copyright clearable.

The governance frameworks to deploy them safely are incomplete. That is the gap. (Editorial note: Dan Patten is saying that UK's weakness is **not necessarily a lack of valuable data**. The gap is having the governance, legal arrangements, access mechanisms and safeguards that would allow these national assets to be used safely and legitimately for AI.)

Data Governance Is the Sovereignty Question

Governing those assets, making them interoperable and under the domestic jurisdiction is not a technical question, it is the sovereignty question.

It is the difference between a country that has made deliberate choices about its AI future and one that has accepted the defaults. between steering and depending, between a system that reflects British society and one that was built for someone else's.

Sovereign AI Is the Cumulative Result of Policy Choices

Sovereign AI is not a destination you announce. It is the cumulative result of choices made at every layer of the stack.

The UK has made those choices on compute, it has made them on models. The data layer is where those choices remain unmade, and it's the layer that determines whether everything built on top of it is genuinely British or just hosted here.

Thank you.



Dr Susan Oman

Senior Lecturer in Data, AI and Society,
University of Sheffield (Centre for
Machine Intelligence)



AI sovereignty is fundamentally a societal as well as a technological and economic question, requiring a people-centred approach that builds public trust, strengthens UK-relevant AI literacy, and asks who develops AI, who benefits from it and who may be at risk.

Introduction: A Human-Centric Perspective on AI

I'm Susan Oman, and I'm from the University of Sheffield, where I lead on human-centric AI through its Centre for Machine Intelligence. My expertise specifically lies in understanding how different parts of society respond to technological progress and the policy decisions surrounding it. It speaks to questions of what public AI is, or could be, and how this depends on what we mean by "the public".

My research spans different populations, from those we might consider non-experts among the citizenry to professionals including those working in national statistics offices, senior policymakers, local government officials and small business owners.

AI sovereignty emerges as a concern in the research I do and in the conversations that I have daily across all of these stakeholders.

The Sovereign AI Fund: Confusion Over Who and What Counts

The most recent example of these conversations relates to the Sovereign AI Fund, of course. In conversations I've had with local businesses in the last week or so since the announcement, they have expressed deep concern about what they see as the limited potential of the funding, rather than the excitement that we were hoping for, and, crucially, confusion over eligibility. There are rumours of one recipient of compute credits being incorporated in the US, with senior team members based in Silicon Valley rather than here in the UK.

For the small businesses and startup community in Sheffield that were really hopeful about the Sovereign AI Fund - and have been waiting for this kind of investment and initiative to be announced - there is a feeling of alienation and confusion over who counts and what counts in this sovereign AI initiative.

They also raise concerns about how it actually contributes to these questions of national resilience or competition.

AI Sovereignty as a Deeply Human and Societal Question

I want to pause and think again about definitions: what do we mean when we talk about AI sovereignty?

I want to suggest that we are ultimately asking a set of deeply human and societal questions. They are questions about who is making what, for what purposes and why. They are questions surrounding how the technology will be used, but also, crucially, not just how it will be used in the short term, but what will happen next and what that will enable.

So part of the question for policymakers, of course, has always been who benefits and where the benefits accrue. AI sovereignty is key to these considerations.

The AI Skills Hub and Sovereign AI Infrastructure

Another recent policy initiative, the AI Skills Hub, was an attempt to see beyond national AI infrastructure as simply technical or economic and consider the importance of future AI infrastructure as people-centred and focused on skills.

At its launch in February this year, colleagues and I wrote an op-ed in *Computer Weekly* outlining our concerns, from the perspective of sovereignty, about a number of the decisions surrounding the £4.1 million investment.

The launch included 595 courses, with only 14 benchmarked by Skills England. All of these foundational courses were from US providers, essentially training our future AI skills using US products.

The press releases foregrounded free training, but 60% of the courses required payment, including some of those that were marked as free. There has been a lot of public criticism from educators and others, so I'm not going to linger too long.

But, crucially, some of the courses promoted aren't relevant to the UK or, more importantly, misrepresent key UK law on issues such as intellectual property.

AI Literacy Infrastructure and Dependence on US Products

Therefore, I want to make the point that we didn't consider sovereignty in the delivery of AI literacy infrastructure. We've invested £4.1 million in infrastructure that claims to offer access to AI for all of the UK, but it isn't accessible to everybody, and what can be accessed trains the UK on US products.

This is a fundamental issue for AI sovereignty in and of itself, of course.

“Let's Talk AI”: A People and Community-Centred Alternative

Now, there are alternative models, and I've spent the last year working on a national AI awareness campaign that was devised by, with and for parts of the citizenry with very low digital literacy. This launched last month, and our efforts towards people- and community-centred AI literacy as infrastructure can be found on our website, *Let's Talk AI*.

The process has enabled an even deeper understanding of what people want to gain from AI literacy as infrastructure. What people want from AI literacy is deeply rooted in what they're concerned about, specifically what they trust and their lack of confidence in what they should trust.

These concerns are related to issues of AI sovereignty: who is using what, who benefits and who could be at risk.

Public Concern About AI Is Rising

As we know from work that I've done on another research project called *Public Voices in AI*, an £850,000 UKRI investment, concern is rising. It's rising fast. It's gone up 10% between two waves of our nationally representative survey, conducted between 2023 and 2025.

We also know from follow-on research delivered by the Ada Lovelace Institute at Christmas that the UK public don't trust private companies to self-regulate, with 89% wanting an independent regulator for AI equipped with enforcement powers.

We also know that 91% of the public feel it's important that AI systems are developed and used in ways that treat people fairly, and that this should be prioritised over economic gains, speed of innovation and international competition.

Public Support for Sovereign Technology

When you ask the public specific questions related to sovereignty, the most recent example being commercial research from FSS Global, it seems that, unsurprisingly, the public explicitly supports sovereignty, especially in light of recent geopolitical tensions.

What is interesting is that it seems the public is also willing to pay extra for technology developed by domestic companies, with two in five stating that they would be prepared to pay up to 50% more.

The Limitations of Data and AI Sovereignty

It is important to note that data and AI sovereignty have limitations. Research I conducted during COVID and shortly after the launch of UK GDPR involved speaking to BBC R&D departments, where they were trialling different models of reimagining data sovereignty. These included models where local organisations or individuals had sovereignty over their own data - data and AI sovereignty at the personal level rather than the national level.

What was found was that the public didn't really engage with these ideas of data sovereignty because they didn't look enough like the US platforms that we've all become accustomed to using every day. Crucially, our data literacy, AI literacy and digital literacy are all embedded in another model.

Building Meaningful Data, AI and Digital Literacy

Therefore, addressing what data, AI and digital literacy mean requires much care and consideration if we're to make a meaningful intervention, as well as investment and backing from central government.

A People-Centred Approach to AI Sovereignty

So then, to conclude, sovereignty is crucial to all considerations around UK AI and its future. However, we won't make the decisions that serve the UK without thinking about AI sovereignty as a profoundly societal question and one that demands multiple perspectives.

If we want a resilient AI ecosystem, we must stop asking what the technology can do for our economy at the macro scale alone, without considering the economics of everyday human choice and behaviour.

There are indications that AI sovereignty increases public trust and therefore support, including purchasing power. But there are also indications that AI sovereignty is limited by how we already live and work within systems that are not sovereign.

For that reason, I propose that we need to take a people-centred approach to AI sovereignty.

People-Centred Infrastructure Before AI Becomes Entrenched

Let's Talk AI is one model of a people-centred approach to AI infrastructure that encourages healthy conversations around AI.

Our evaluations indicate that this people-centred approach to infrastructure needs to happen before AI becomes too deeply entrenched. It is one of the best ways of intervening in response to people's increasing concerns surrounding AI.

These alternative models should inform how we communicate models of AI sovereignty across different forms of infrastructure, and address those key questions of the who, what and where of AI that drive policy decisions around AI and its sovereignty.

Thank you.



Leo Rogers
CEO & Co-Founder, Curvo AI

Britain should create the conditions for UK AI startups to compete globally by opening public procurement to smaller British companies, improving access to compute and growth capital, and using the public sector as a route to commercialisation rather than allowing talent and companies to migrate to the US.

Introduction: A British AI Company and the Sovereignty Debate

Curvo AI is an agentic AI platform that supports enterprise sales leaders, giving real-time suggestions during virtual meetings on handling objections, asking questions and navigating the most high-stakes conversations.

As a British company, the concept of sovereignty is very close to home and very meaningful in terms of data residency and ensuring that the UK is at the frontier of what's possible.

A Realistic Approach to AI Sovereignty

To open my piece, I'll start by saying that the sovereignty debate needs a realistic approach. We're not here to outscale the US or to outscale China, but we are here to make Britain's proposition competitive and enable British companies to excel alongside their US counterparts, which are equally gaining market share.

Enabling British Startups to Compete

British startups should be able to compete fairly in their own country. They should have the resources to do so, and they should have access to the capital needed to generate truly global and competitive propositions.

And I think the state itself should be confident in supporting that innovation, confident in taking risks and confident in backing the companies that are building the future of the country and accelerating GDP in one of the fastest-growing sectors in the world.

Britain historically has excelled when opening markets with strong institutional backing on its own shores. And now we need to reach a point where talent isn't leaving for Silicon Valley; talent is staying here in London, enabled and incentivised to do its best work and to give back to the British economy.

A Sovereign AI Product and Services Strategy

Now I want to discuss what the Government could practically do to generate a sovereign AI product and services strategy. This is largely about the public sector, but also about helping companies while ensuring national security and that UK companies domiciled here are indeed giving back to the economy.

The Sovereign AI Unit and Britain's AI Foundations

Firstly, a big testament to Joséphine Kant, our next speaker, and the work at the Sovereign AI Unit, which has raised half a billion to back and promote frontier AI labs and infrastructure-layer companies.

I think, first and foremost, this is a signal that Britain is moving in the right direction and recognising the strategic importance of domiciling AI capability here on our own shores.

Today, 200 tech unicorns have been built here. We have some of the best talent in the world and people who are at the frontier. And the first modern computer (i.e. theoretical mathematical model for the universal computing machine on paper) was built here in by Alan Turing (in 1936). We have the foundation now to excel and to build further.

Using British AI to Transform the Public Sector

Now, we feel at Curvo AI that there is an argument to take this much further. What are the use cases where AI startups can help with transformation within the public sector itself?

A lot of that is around our belief that civil servants face growing case complexity, particularly in relation to immigration policy and the asylum process.

As we know it today, the Home Office spent £4.76 billion between 2024 and 2025, and decision timelines were at 12 months from asylum cases being processed to decisions being executed. (Editorial: Of that total £ 2.1. billion was specifically used for temporary hotel accommodations.)

This is a very challenging issue for everyone, and everyone indeed wants to get it right in a fair and transparent way.

AI and the Immigration and Asylum Process

Curvo AI is poised to solve this problem by streamlining the immigration interview process so that it can be fairer and more transparent, and enable immigration officers to make decisions in real time - in minutes and days as opposed to weeks and months - to get decisions over the line.

And this is one example of where British companies should be poised to support the public sector, where we have patriotic interests at heart that no US provider can reciprocate.

We understand the sensitivity and nuance behind deploying into public services, and I believe that British companies, particularly at their early stages, should be brought into the public sector in order to win tenders and transform industries in the right way.

Incentivising Innovation and Investment

There are many incentives that could be put in place to help accelerate growth. The UK has some fantastic schemes, including the Seed Enterprise Investment Scheme and the Enterprise Investment Scheme, which stimulate growth among startups and provide fantastic tax reliefs to investors.

They are among the best schemes in the world, and the US doesn't have an equivalent that stimulates economic growth for high-risk startups in the same way.

Compute Scarcity and R&D Tax Credit Support for compute costs

But what we don't have is R&D tax credits for compute costs. We're in an economy where there is compute scarcity at the moment. Every country is racing to build data centres and expand its national grids.

And startups should be provided with the compute needed to optimise for quality as opposed to cost.

Quality creates innovation and enables us to compete at the frontier, whereas if our AI companies are optimising for cost and unit economics, it means that we are focusing on the wrong metric and the wrong measure of success.

Creating a Competitive Environment for British AI

For us to build this competitive edge against our US counterparts, British talent needs to be motivated, the right funding needs to be in place, and we need to support those who are striving for change and building the economy.

Three Practical Asks

As for my closing remarks, there are three practical asks from Curvo AI.

(1) Open Public Procurement to UK Startups

The first is procurement policy within the public sector. How can we genuinely open the doors for UK startups to bid for and win tenders in a way that is transparent, with clear eligibility criteria, and in a way that allows startups to progress, rather than tenders being won by the Big Four, Accenture, Deloitte and the rest? And how can we enable British firms to compete on a level playing field, rather than tenders being won by US big tech?

(2) Create a High-Trust Environment Between Government and Startups

The second point is about government enabling more opportunities and creating a high-trust environment so that startups can understand the real problems within the public sector from the inside out and solve them with the most appropriate solutions.

(3) Build a Competitive Innovation Environment

And thirdly, we need a competitive innovation environment. That means R&D tax credits for compute costs, so that inference is subsidised by the state. (Leo Rogers makes clear that he is specifically talking about **R&D tax support for compute**, not claiming that the UK has no R&D tax relief system at all.) And it also means more specialised funding for early-stage companies: funding that focuses on commercial outcomes and enterprise products, and funding that stimulates growth to a level that is globally competitive.

Building a Template for National Excellence

We're revolutionising what it means to be a British startup, and now we need the conditions and the right mechanisms put in place to build a template for national excellence.

Thank you very much.



Joséphine Kant

Head of Ventures, Sovereign AI Unit, UK Government

“The UK should pursue a “sovereign edge” rather than technological self-sufficiency, concentrating investment, compute, procurement, data and talent on strategic areas where Britain can gain leverage and become an indispensable global AI partner while retaining real agency over critical capabilities.

Introduction: Why the UK Can Become a Global AI Leader

Thank you very much, and thank you for the invitation to speak today. I did not expect quite as many people. I'm sorry I can't make eye contact with all of you directly, but I hope this is just the beginning of the conversation.

My name is Joséphine Kant. I'm Head of Ventures for the UK Sovereign AI Fund. Before joining Sovereign AI in January to help lead the fund, I spent my career working across technology and venture investing. I was at Google and Y Combinator, later helped launch an AI-focused fund in the US, and came back from Silicon Valley to the UK because I think we all realise we have a very, very important question to answer. And before I talk about what that question is, what drew me to the UK Government and to this tour of duty was a belief that the UK has all of the ingredients required to actually become a global leader in AI.

We have exceptional universities, we have frontier research capabilities, we have strong public institutions, and we have incredibly ambitious founders, whom I meet every single day.

But historically, we have not always translated those strengths into globally scaled companies that are strategically anchored in Britain.

That is precisely why Sovereign AI was created.

The Sovereign AI Fund: Supporting British Startups to Start, Scale and Win

I chose to do this tour of duty, which I can assure you is quite painful. It has me now 33 days into government, learning what the acronyms are and learning how to work in the Civil Service. And this will probably be the largest group project that I think we, as a country, will work on together to figure out how to do this right. I chose to do it because I really think we can get it right.

And I have a part here in my speech about what Sovereign AI is, but based on the conversation so far, by a show of hands, how many of you are familiar with the work that we do at Sovereign AI? Quite a few. I will just reintroduce it for those who don't know what we do.

We are the Government's new £500 million venture fund focused on supporting British startups to start here, to scale here and to win everywhere.

That means we invest directly in British-anchored AI startups and give them access to the things that increasingly determine their success in AI, whether it's compute, whether it's data, as Dan just mentioned, whether it's talent, or whether it's procurement, as Leo just mentioned. It is about making sure that we're able to operate at the speed and ambition of venture capital while using the unique levers of the state to strengthen Britain's long-term strategic position in AI.

Early Investments and Access to Sovereign Compute

We launched on 16 April. In the first couple of weeks, we've made three investments so far: Callosum, which is a compute orchestration company founded by two Cambridge scientists; Ineffable, which is the new startup founded by David Silver, who is a British founder and one of the forefathers of reinforcement learning; and Isomorphic Labs, founded by Sir Demis Hassabis.

And alongside that, we've enabled six startups to gain access to the UK's largest supercomputer through a compute programme.

Building Sovereign AI in Public and Across the UK

There are three things that I want to say today. Before I say them, I want to acknowledge all of the really good feedback from the panel so far, on which I took many, many notes.

We are really committed to building in public and doing so with your input. We have committed to providing quarterly updates through our Chair, James Wise. We have a lot more to do.

We are just embarking on our regional tour to make sure that we are looking at other cities and we're not just looking at London. We were in Newcastle last week, I'll be in Glasgow on Thursday, and I'll be in Cardiff next Thursday. We have many, many more cities to come.

I welcome your ideas on where else we should go across this country to find the most promising UK AI startups, because I know they're not just in London.

And we really need to talk more about how we actually define eligibility, how we talk about what a UK-anchored company is, what sovereignty is and what AI is. These are not simple questions.

So we are really committed to building in public and doing so with your input. Please find me if you have thoughts on that.

Three Questions for the UK's Sovereign AI Strategy

The three things that I want to address today are: which elements of the AI stack do we think the UK should focus on and gain leverage in? How do we actually build a strategy to fulfil that? And then the third thing I want to end on is: why now? Why are we all choosing to be here on a Monday evening? Why is this such an important conversation to have now?

Building a "Sovereign Edge" Rather Than Complete Self-Sufficiency

On the first question, when we speak of sovereign capability - and a few of you mentioned this earlier - we are not talking about turning inward.

We are talking about ensuring that Britain has strategic depth and access options at the most important points of the value chain. And we should focus on gaining leverage in the most strategically valuable layers of the stack rather than attempting to achieve complete self-sufficiency across every component.

So this is not about isolationism. This is about building what we describe as a “sovereign edge”: concentrating public resources in areas where we already have depth, where we have capabilities and where we have relevance.

Compute as Strategic National Infrastructure

One of those areas for us is compute infrastructure. Compute is increasingly becoming strategic infrastructure in the same way that energy or telecommunications infrastructure has been in previous eras.

So the UK needs domestic access to sovereign compute capacity, including national supercomputers, and we need to make sure that we have the right orchestration and optimisation software.

That is why we invested in Callosum, a company out of Cambridge that is building the orchestration layer across diverse chip architectures.

That is why compute access is a really important part of our sovereign AI offer. Many promising British AI startups cannot compete globally if they can't actually access the right compute.

Using Compute Access to Increase UK Influence

One of the companies that we invested in is a company called Odyssey, which prompted some of the panel members to ask: why are we investing in a company that is a US startup?

Odyssey is a company - if I can just speak kind of off script - that employs half of its workforce in the UK. The majority of its core research team is in London, and by granting them access to build their models on local infrastructure, it helps us increase the UK's influence over where their capabilities are developed.

Frontier AI Research: Remaining a Maker, Not Just an Adopter

We have made these six commitments so far, and there is another question about how we actually figure out, beyond orchestration, what we do. This brings us to our research challenge and to frontier AI research.

If Britain wants to remain an AI maker instead of just an adopter, we need enduring capabilities in these next-generation models. That is why we invested in Ineffable Intelligence, David Silver's company, which I can tell you was a very competitive round for us to get an allocation into.

And we need to build these enduring capabilities.

AI Safety, Assurance and Governance as a UK Strength

And it is really important here for me to call out the excellent work that the AI Security Institute has done. If anything, we need to do this safely.

So we place a lot of emphasis on AI safety, governance and assurance. This is an area where the UK has really strong international credibility, and we could become the globally indispensable leader in AI assurance, evaluation and governance tooling. So our overarching philosophy is: can we build a sovereign edge? Can we become a strategically indispensable ally rather than trying to replicate every single layer?

Combining Venture Investment with the Levers of the State

Now the next question is how we actually do this, which is the hardest thing to get right.

The way that we think about it (and I'm very open to ideas on what we did not get right, what we did get right and what else we should think about) is a combination of venture investment, public infrastructure, procurement and regulatory support designed to help UK firms build globally competitive AI companies from Britain.

Venture Funding at the Speed of AI

The first thing is venture funding. And I love the comments about whether we could be a sovereign AI institution instead of just a sovereign AI venture fund.

One of the reasons why we are starting as a venture fund is the recognition that these frontier AI markets move at extraordinary speed. Strategic companies need rapid access to capital, whether they're pre-seed or growth-stage.

Access to Sovereign Compute

The second for us is access to sovereign compute. Access to compute is now one of the largest barriers facing frontier AI startups. Publicly backed compute access through national infrastructure can lower barriers to innovation.

One thing about that is that I think we have to be really honest, when we talk about sovereignty, about what that actually means.

Sovereignty does not mean complete control. I could give a company a million GPU hours to train the next generation of its model on supercomputers hosted here in the UK. But those supercomputers run on NVIDIA chips, while the semiconductor supply chain itself extends internationally, including through ASML in the Netherlands and manufacturing in Taiwan.

If you really start to go down the stack, this idea of true sovereign control and a truly protectionist approach would be a silly thing for us to do.

So we have to think about leverage.

Public Procurement as an Anchor Customer

There's a third thing around procurement. The public sector, as Leo said, can play a really important role as an anchor customer.

That's why we're very excited that, through our fund, we are able to do procurement in a way that is fast, in a way that is fair and in a way that actually serves companies that are building here in the UK.

Creating Strategic Data Assets

And then there's another point about strategic asset creation. To the point that Dan made, if data is our most valuable asset — and particularly in biotech, where we have access to arguably the best data in the world — can we actually fund these shared datasets, autonomous labs and scientific infrastructure?

Our ambition is not protectionism.

Attracting and Retaining Global AI Talent

And the last thing is talent. The UK needs to continue to attract the very best talent in the world and make sure that the highest-value R&D activity is here.

Focusing on Sectors Where Britain Has Genuine Advantage

In practical terms, we believe that the UK should focus on sectors where it has a genuine advantage.

So part of how we determine eligibility and pick which companies we want to invest in is to look at areas including health and life sciences, scientific discovery, compute efficiency, AI safety and assurance, advanced infrastructure software and national security applications.

Why Sovereign AI Matters Now

To end, I want to step back from what I just said and address a more fundamental question, which is: why now? Why sovereign AI? And why did I choose to do this?

We are at a turning point. Technology, and especially AI, is now a form of hard power. In the past, industrial capacity underpinned and helped determine whether a country was secure. Now, whoever controls the digital systems increasingly influences our day-to-day lives.

And if another country controls your technology and controls your public services, you're exposed.

You can be pressured, you can be priced out, you can be cut off. That means that your way of life is being shaped by decisions that are made elsewhere. No responsible government can accept that without careful judgement.

Britain Needs Real Agency Over AI

That's why Britain needs real agency over how AI is developed and how it is used. Not by shutting ourselves off from the world, and not by trying to do everything alone, but by controlling the critical capabilities that we depend on.

A Narrowing Window for Britain

And I should end by saying that, if we believe that backing British AI is about security and is about prosperity, I think we have a very narrow window in which to make the most important decisions.

They say that it takes a village to raise a child. I think it takes a country to get sovereign AI right.

And I hope to collaborate with all of you in this room to secure the most favourable outcome for this nation.

Thank you.

DISCUSSION

DISCUSSION AND QUESTIONS FROM PARLIAMENTARIANS AND THE AUDIENCE



1. Can the Sovereign AI Fund Drive Culture Change Across Government?

Lord Tim Clement-Jones asked whether the Sovereign AI Fund could influence wider government culture, particularly by encouraging a more strategic and sovereign approach to public procurement.

Joséphine Kant agreed that this was part of the ambition. She argued that many previous procurement decisions had been made without recognising technology as a form of “hard power”. Government should therefore consider the strategic consequences of technology procurement more deliberately. She also emphasised that taxpayers should share in the financial upside from successful AI investments: the Fund should seek both strategically important investments and strong commercial returns, potentially developing into a sovereign wealth fund-style institution.

2. Should We Separate Public-Sector Sovereignty from Commercial AI Sovereignty?

A parliamentarian questioned whether the debate should distinguish more clearly between two dimensions of sovereignty: government and public-sector transformation, including procurement and buying British technology; and the commercial challenge, where UK companies need capital and speed to compete in a rapidly developing global AI market.

Ashley Braganza argued against drawing too sharp a distinction. Public and commercial AI are increasingly intertwined, and the important issue is the governance governing their interaction. Proper governance could allow government and private-sector capabilities to reinforce each other while ensuring transparency over how public contracts and opportunities are awarded.

Susan Oman similarly argued that international technology companies and public-sector procurement are already deeply interconnected. From her research, she questioned whether the two could realistically be separated.

3. Should Britain Build Its Own Sovereign AI Model; and Who Should Build It?

Professor Lionel Tarassenko argued that Britain's exceptional sovereign data assets (including the NHS, BBC, Met Office and others) could be used to develop a sovereign UK AI model. He suggested that application-specific frontier-level models need not require the enormous expenditure associated with general-purpose frontier models and asked the panel **who should build such a British model**.

The panel offered different approaches.

Ashley Braganza favoured a private-sector-led model, particularly involving SMEs with relevant sector expertise, rather than a top-down programme led by large companies.

Dr Neel Savani saw value in maintaining UK models partly to preserve Britain's distinctive cultures and languages, pointing to lessons from Switzerland and the importance of Welsh, Scottish and Cornish cultural and linguistic characteristics.

Dan Patten was more cautious about creating a single sovereign model. He argued for maintaining **interchangeability at the model layer**, allowing different commercial and open-source models to be selected according to the use case, sensitivity of the data and deployment environment.

Joséphine Kant said government should not necessarily build the model itself. She described sovereign AI as a “group project”, noting that several companies were already developing UK sovereign AI models. She suggested that specialist startups could provide the expertise, while government might play an important **convening and safety role**.

Professor Tarassenko, returning to the question at the end, proposed that the **Alan Turing Institute** should lead such an initiative as the UK's National AI Institute, working with industry. His central argument was that a sovereign model requires full control and transparency over its training data and model weights. This could improve security against data poisoning and backdoors while also creating a transparent framework through which creative industries could understand - and potentially be remunerated for - the use of their material in training.

4. What Should Britain Learn from the EU AI Act?

Lord Taylor of Warwick asked whether there were useful lessons, positive or negative, from the EU AI Act, while expressing concern that the legislation was overly centralised and bureaucratic.

Dr Neel Savani agreed that the EU approach could be highly prescriptive but argued that its clear categorisation of high-, medium- and low-risk AI provides useful lessons. Britain could use similar risk profiles more flexibly to identify where behavioural change and governance intervention are actually required.

Susan Oman highlighted the Act's approach to **accountability and AI literacy**. She particularly welcomed the way responsibility is embedded across different levels of organisations deploying AI and the expectation that workers understand the potential positive and negative social consequences of AI systems.

Leo Rogers warned against allowing regulation to undermine innovation. He contrasted Europe's more regulatory approach with the US emphasis on speed and innovation, arguing that Britain needs to find a balance that allows companies to commercialise world-class UK research without feeling compelled to relocate to the United States.

Joséphine Kant argued that Britain's opportunity may lie precisely in combining innovation with its institutional strengths. Rather than simply choosing between US-style innovation and European-style regulation, the UK could use its expertise in institutions, law, AI safety and security to help establish an international “operating system” for the safe deployment of AI, without falling into overregulation.

5. How Can Britain Stop AI Startups Moving to the United States?

The discussion also addressed the recurring problem of promising British technology companies moving to the US or being acquired by American companies.

Leo Rogers described the familiar startup trajectory: reach Series A, undertake a Delaware flip, move towards Silicon Valley and raise capital there. He attributed this largely to the greater availability of funding and the faster pace of the US ecosystem.

Britain therefore needs not only excellent research but much stronger **commercialisation capability**, connecting its universities and deep-tech expertise with capital, commercial talent and incentives that allow companies to scale from the UK.

Overall Message from the Discussion

The discussion revealed broad agreement that **UK AI sovereignty should not mean technological isolation or attempting to reproduce every layer of the global AI stack domestically**. Instead, sovereignty should provide Britain with meaningful strategic choices and reduce dangerous dependencies.

There was also a strong emerging consensus around several priorities: using procurement strategically; protecting and exploiting sovereign data assets; strengthening the commercialisation of British research; retaining AI companies and talent in the UK; maintaining transparency and accountability; and developing domestic capability where it creates genuine strategic advantage.

The principal area of debate was **how far Britain should go in developing its own sovereign AI model**. Views ranged from maintaining an interchangeable ecosystem of commercial and open-source models to creating a dedicated British model led by startups or, as Professor Tarassenko proposed, by the Alan Turing Institute in partnership with industry



BIOGRAPHIES

OF EVIDENCE GIVERS



From left to right: Dan Patten (Scale AI), Lord Tarassenko (House of Lords), Prof. Birgitte Andersen (APPG AI Secretariat, Big Innovation Centre), Dr Susan Oman (Univ. of Sheffield), Lord Ranger (House of Lords), Lord Taylor (House of Lords), Lord Tim Clement-CBE (House of Lords), Joséphine Kant, (UK Government), Prof. Ashley Braganza (Brunel Univ. of London), Dr Neel Savani (Cognizant), Leo Rogers (Curvo AI).

Selected Images from the discussion



Lord Tarassenko
(APPG AI Parliamentary Member)



Lord Clement-Jones CBE
(APPG AI Co-Chair)

Lord XXXX (APPG AI)
Parliamentary Member



Lord Taylor of Warwick
(APPG AI Honorary Vice Chair)



Lord Ranger of Northwood
(APPG AI Vice Chair)



Professor Ashley Braganza

*Director, Centre for Artificial Intelligence,
Brunel University of London*

Professor Ashley Braganza is Director of the Centre for Artificial Intelligence and Chair of Business Transformation at Brunel University of London. He is a recognised expert in AI and digital transformation, founder of Brunel's interdisciplinary Research Centre for Artificial Intelligence and leader of the university's wider AI ecosystem. His work spans AI strategy, policy, research and industry collaboration, including contributions to the All-Party Parliamentary Group on Artificial Intelligence.

Professor Braganza has secured more than £3.5 million in research funding and is part of the consortium behind the €5.8 million ELOQUENCE EU project (2024–2027). He launched Brunel's Master's in AI Strategy and has published more than 100 papers alongside completing over 50 international consultancy assignments. He is a Fellow of the British Academy of Management, a Trustee of the Industry and Parliament Trust, and a recipient of the 2024 Sustainability Leadership in Education Award.



Dr Neel Savani

*Responsible AI,
AI Labs, Cognizant*

Dr Neel Savani leads Responsible AI within Cognizant's AI Labs, working at the intersection of AI governance, safety research, standards and enterprise deployment. His work focuses on translating cutting-edge AI safety research and governance approaches into practical applications that can be deployed by organisations.

Within Cognizant, his previous roles have included leading EMEA pre-sales for Generative AI and regional data science delivery, giving him experience across both the technical and commercial application of AI.

Earlier in his career, Neel was a research scientist at NASA Goddard, where his work included future mission design and space situational awareness for government policy. This gave him experience in operationalising advanced research in mission-critical environments, which he now applies to enterprise-scale AI challenges.

Neel is a member of the UK Government's Digital Sustainability Alliance AI Working Group and an Affiliate Research Professor at the University of Maryland, Baltimore County. He is also a three-time graduate of Imperial College London.



Dan Patten

*Government Affairs Lead
(UK and Europe), Scale AI*

Dan Patten leads Government Affairs for the UK and Europe at Scale AI, where he is responsible for the company's engagement with the UK Government and its policy strategy across the region. Before joining Scale AI, Dan worked in the UK Parliament and in consultancy, focusing on the intersection of public policy and technology.

Based in London, Dan represents Scale AI as the company expands its work in model evaluation and AI safety testing. Scale AI develops AI applications, provides data infrastructure and operates evaluation frameworks designed to assess the capabilities and reliability of advanced AI systems.

Its work spans leading AI companies, governments and major corporations, including work with organisations such as OpenAI, Meta and Microsoft. This combination of building AI systems and independently evaluating them informs Scale AI's perspective on AI reliability, safety and sovereign capability.



Dr Susan Oman

*Senior Lecturer in Data, AI and Society,
University of Sheffield, Centre for Machine
Intelligence*

Dr Susan Oman researches digital, AI and data issues, with a particular focus on well-being, inequalities and improving organisational and policy practices. At the University of Sheffield, she leads work on human-centric AI through the Centre for Machine Intelligence. She is recognised for communicating complex ideas to diverse audiences and has received awards for open research, impact and public engagement. Her expert commentary regularly appears across BBC news outlets.

Susan completed a two-year policy research fellowship within the UK Government's Department for Digital, Culture, Media and Sport (DCMS), which led to her membership of its College of Experts.

Her 2021 monograph, *Understanding Well-being Data*, has been downloaded more than 70,000 times and developed into an interactive resource for policymakers, practitioners and students.

Susan has delivered presentations, workshops and advice to international policymakers and practitioners, as well as across Whitehall, devolved and local governments and public bodies. Her research includes public attitudes towards AI, AI literacy and the societal dimensions of technological change.



Leo Rogers

*CEO & Co-Founder,
Curvo AI*

Leo Rogers is a London-based technology entrepreneur and Co-Founder and CEO of Curvo AI, a real-time agentic AI sales execution platform that supports enterprise sales teams during high-stakes virtual meetings. Curvo AI has raised \$1.2 million in funding backed by two venture capital investors.

Leo previously founded a seven-figure UX design agency, working with clients including HSBC, NS&I and Spacelabs. The agency delivered more than 100 products collectively used by over 30 million people. He is also Co-Founder of Linkable, a venture-backed platform using AI to connect brands with relevant influencers.

Across three ventures, Leo has developed experience in agentic AI, product development, enterprise sales and company building. He has closed millions in sales, won Fortune 500 customers and built go-to-market operations from the ground up. His experience brings the perspective of a British technology entrepreneur seeking to build and Scale AI businesses in an increasingly competitive international market.



Joséphine Kant


*Head of Ventures, Sovereign AI Unit, UK
Government*

Joséphine Kant is Head of Ventures at the UK Government's Sovereign AI Unit, where she is helping build the Government's venture investment capability to back and scale strategically important UK AI companies.

Before joining the UK Government, Joséphine spent four years in the United States investing in technology companies, first with Y Combinator and subsequently building Dogwood Ventures. Earlier in her career, she spent eight years at Google, including work with Google for Startups supporting entrepreneurs and startup ecosystems internationally.

At the Sovereign AI Unit, her work focuses on combining venture investment with the wider strategic capabilities of government, including access to compute, data, talent and procurement, to support high-potential AI companies and strengthen the UK's position in strategically important areas of artificial intelligence.

Her perspective combines experience in global technology, venture investment and government, with a focus on enabling leading AI companies to start, scale and remain anchored in the UK while competing internationally.



About APPG AI

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

- Dr 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
- The Rt Hon Damian Hinds MP Conservative
- Dr Danny Chambers MP Liberal Democrat
- Dawn Butler MP Labour (APPG AI Vice-Chair)
- David Reed MP Conservative
- Dave Robertson MP Labour
- The Rt Hon Esther McVey MP Conservative
- Emily Darlington MP Labour
- George Freeman MP Conservative
- Gordon McKee MP Labour
- Graham Leadbitter MP SNP
- Iqbal Mohamed MP Independent
- Juliet Campbell MP Labour
- Leigh Ingham MP Labour
- The Rt Hon Liam Byrne MP Labour
- Sir Mark Hendrick MP Labour
- Martin Wrigley MP Liberal Democrat
- Mike Martin MP Liberal Democrat
- Maureen Burke MP Labour
- Neil Duncan-Jordan MP Labour
- Peter Fortune MP Conservative
- Samantha Niblett MP Labour
- Sarah Edwards MP Labour
- Tim Roca MP Labour
- Tom Collins MP Labour
- Tom Gorden MP Liberal Democrat
- Tony Vaughan MP Labour
- Zöe Franklin MP Liberal Democrat
- Dr Zubir Ahmed MP Labour

Parliamentary APPG AI Members: House of Lords

- Lord Clement-Jones (Tim Clement-Jones) Liberal Democrat (APPG AI Co-Chair)
- Lord Bilimoria (Karan Faridoon Bilimoria) - Crossbench
- Viscount Colville Of Culross (Charles Mark Townshend Colville) Crossbench
- Lord Craig of Radley (David Brownrigg Craig) Crossbench
- Lord Cromwell (Godfrey Cromwell) Crossbench
- Viscount Camrose (Jonathan William Berry) Conservative
- Lord Hunt of Kings Heath (Philip Hunt) Labour
- Lord Janvrin (Robin Berry Janvrin) Crossbench
- Baroness Kramer (Susan Veronica Kramer) Liberal Democrat
- Baroness McGregor-Smith (Ruby McGregor-Smith) Non-affiliated
- Lord McNicol (Iain Mackenzie McNicol) Labour
- Lord Pitkeathley of Camden Town (Simon Pitkeathley) Labour
- Lord Ranger of Northwood (Kulveer Ranger) Conservative (APPG AI Vice-Chair)
- The Lord Bishop of Oxford (Stephen Croft) Bishops
- Viscount Stansgate (Stephen Stansgate) Labour
- Baroness Thangam Debbonaire (Thangam Elizabeth Rachel Debbonaire) 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 (APPG AI honorary Vice-Chair)



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.

[Download your Pavilion App Now!](#)



Annual APPG AI 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.



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





SECRETARIAT

appg@biginnovationcentre.com

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



**BIG
INNOVATION
CENTRE**

All rights reserved © Big Innovation Centre. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form without prior written permission of the publishers. Big Innovation Centre is the trading name of BIC London Limited, a company registered in England and Wales.