AI Horizon Scanning

The Next Wave of AI Technologies and Their Impact on Our World

All-Party Parliamentary Group on Artificial Intelligence

BIG INNOVATION CENTRE
Secretariat

February 2026
Policy Forum

Table of Contents

Rapporteur Preface: Page 3
Introduction: Page 5
Aim of Session: Page 6
Findings: Page 8
UK AI Strategy: Immediate Priorities: Page 14

Evidence: Page 15

- Yann LeCun, AMI Labs and New York University. Page 16
- Prof. Viv Kendon, University of Strathclyde. Page 20
- Prof. Mihaela van der Schaar, University of Cambridge. Page 25
- Bob De Caux, Chief AI Officer, IFS. Page 30
- Mark Taylor, Osborne Clarke. Page 35

BIOs of Evidence Givers: Page 40
About APPG AI: Page 45
Contact: Page 49

Title: AI Horizon Scanning: The Next Wave of AI Technologies and Their Impact on Our World
All-Party Parliamentary Group on Artificial Intelligence (APPG AI)
Date of Publication: February 2026
Publication Type: Parliamentary Brief | Policy Brief
Publisher: Big Innovation Centre

Rapporteur's Preface
AI Horizon Scanning:
The Next Wave of AI Technologies and Their Impact on Our World
Rapporteur: Professor Birgitte Andersen

Purpose of the Session

Artificial intelligence is evolving at a pace that challenges not only technological understanding but also policy, governance, and institutional readiness. This evidence session on AI Horizon Scanning was convened to examine how AI is changing, what lies ahead, and how the UK can position itself in response to these developments. Drawing on perspectives from leading experts across academia and industry, the discussion sought to move beyond immediate developments and consider the broader trajectory of AI.

A Shift in the Nature of AI

A central insight emerging from the discussion is that AI is undergoing a fundamental transition. The recent wave of progress has been driven largely by large language models, which have demonstrated impressive capabilities in generating and manipulating text. However, the next phase of AI will not simply be an extension of these systems. It will involve a shift towards systems that can engage with and act in the real world.

These more advanced, "agentic" systems are expected to move from responding to instructions to making decisions, coordinating actions, and interacting with complex environments. This shift has far-reaching implications. It expands the potential of AI beyond communication and information processing into domains such as healthcare delivery, industrial operations, and public services, where decisions have tangible and sometimes irreversible consequences.

Key Themes: Opportunity, Risk, and Constraints

The discussion made clear that this transition presents both significant opportunities and material risks. AI has the potential to accelerate scientific discovery, improve the effectiveness of healthcare systems, and enhance productivity across the economy. At the same time, the increasing autonomy and interconnectedness of AI systems introduce new forms of risk that existing governance frameworks are not yet fully equipped to address.

Data emerged as a critical constraint. Access to high-quality training data remains uneven, particularly for smaller organisations and for languages and regions that are underrepresented in global datasets. Questions of ownership, intellectual property, and control over data are therefore likely to become more, rather than less, significant. At the same time, approaches such as 'federated learning' *(defined below)* and synthetic data point towards new ways of reconciling innovation with privacy and sovereignty.

The discussion also highlighted the growing importance of infrastructure and energy, including the potential role of emerging technologies such as photonics, as well as the need to invest in skills and education to ensure that individuals and institutions can adapt to these changes.

Definitions:

Federated learning is a way of training AI models without moving or sharing the underlying data. Instead of collecting all data in one central place, the model is sent to where the data already exists (for example, in different organisations or countries). Each participant trains the model locally on its own data and then shares only the model updates (such as parameters or weights)—not the raw data. These updates are then combined to improve a single, global model.

Why it matters:

- *Privacy: Sensitive data (e.g. medical records) does not leave its source*
- *Data sovereignty: Countries or organisations retain control over their data*
- *Access: Enables training on diverse datasets that cannot be pooled centrally*
- *Collaboration: Allows multiple parties to contribute to a shared AI model without exposing their data*

Scope of This Report

This report reflects the perspectives expressed during the session and sets out the key themes and implications for policymakers, regulators, industry, the research community, and civil society. It is intended to support a more informed and forward-looking approach to AI, recognising that the challenges ahead are not solely technical, but institutional, economic, and societal.

.

INTRODUCTION

This document is a transcript and summary of an APPG AI evidence meeting held on 26 January in the House of Lords, Committee Room 1, UK Parliament. It exclusively contains crucial discussion elements; not all points are addressed.

DETAILS

- Evidence Session: AI Horizon Scanning: The Next Wave of AI Technologies and Their Impact on Our World, Work, and Play
- Time 5:00 pm – 7:00 pm (GMT)
- Date: Monday, 26 January 2026
- Venue: Committee Room 1, House of Lords.

CONTACT THE SECRETARIAT

appg@biginnovationcentre.com
APPG AI Secretariat

BIG INNOVATION CENTRE

EVIDENCE GIVERS

1. **Yann LeCun.** Executive Chairman, AMI Labs; Professor, New York University; Founding Director, Meta FAIR; former Chief AI Scientist, Meta.
2. **Prof. Viv Kendon,** Professor of Quantum Technology, Department of Physics, University of Strathclyde
3. **Prof. Mihaela van der Schaar,** John Humphrey Plummer Professor of Machine Learning, Department of Engineering, University of Cambridge
4. **Bob De Caux,** Chief AI Officer, IFS
5. **Mark Taylor,** Partner, Head of Digitalisation (UK), Osborne Clarke

APPG AI CO-CHAIRs AND RAPPORTEUR

The All-Party Parliamentary Group on Artificial Intelligence (APPG AI) is co-chaired by **Allison Gardner MP** and **Lord Clement-Jones CBE**.

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

Aim of Session: AI Horizon Scanning

The session provided a forward-looking, evidence-based view of where artificial intelligence is heading over the coming decades — technologically, scientifically, economically and legally. It moved deliberately from frontier science through disruptive futures to real-world systems, and finally to industry and governance implications. The session focused on emerging AI frontiers, including agentic AI systems, AGI-orientated models, post-LLM architectures, quantum–AI interfaces, synthetic data and autonomous systems.

Discussion covered

Capability:

- What will actually drive the next major leap in AI capability - scaling today's models, or fundamentally new architectures such as world-models, post-LLM systems and self-improving AI?
- How important will new forms of computation, including specialised chips, advanced accelerators and quantum technologies, be in unlocking the next phase of AI capability?
- How will synthetic data, simulation and AI-generated environments change what AI can learn and how fast it can improve?
- What happens when AI systems move from responding to prompts to independently setting goals, taking actions and coordinating with other systems?

Power

- What becomes scarce and strategically decisive in an AI-driven economy — data, compute, energy, talent, or control of platforms?

Governance

- When AI systems act autonomously in the real world, who is responsible and how should they be governed?

National Strategy

- What must the UK decide now to remain a serious global AI power?

Above (from left to right): Prof. Viv Kendon (University of Strathclyde), Lord Ranger (APPG AI Vice Chair), Lord Clement-Jones (APPG AI Co-Chair), Allison Gardner MP (APPG AI Co-Chair), Lord Taylor of Warwick (APPG AI Honorary Vice Chair), Prof. Birgitte Andersen (APPG AI Secretariat, Big Innovation Centre), Mark Taylor (Osborne Clarke), Prof. Mihaela van der Schaar (University of Cambridge), Bob De Caux (IFS), Iqbal Mohamed MP (APPG AI Member).

FINDINGS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

EXECUTIVE SUMMARY - OVERALL FINDINGS

This session provided a forward-looking, evidence-based discussion on the development, deployment, and governance of artificial intelligence, with particular attention to the UK's strategic position.

1. The Next Phase of AI: From Large Language Models to Agentic Systems

Speakers emphasised that AI is entering a new phase. Current systems, largely based on language models, will be superseded by more advanced systems capable of understanding and interacting with the real world. This shift is expected to drive significant advances in science, healthcare, and industry, while also introducing new risks and governance challenges.

Elaboration

Where we are now: Artificial intelligence is entering a new phase. Most current AI systems are based on language models, meaning they are designed to process and generate text. They can answer questions, write documents, and summarise information, but their understanding of the world is indirect—they rely on patterns in data rather than real-world experience.

Where we are heading: The next generation of AI is expected to move beyond this. These systems will be able to work with real-world data such as images, video, and signals from physical systems, allowing them not only to interpret information but also to act on it. In practice, this means AI could monitor environments, make decisions over time, and support or automate real-world tasks in areas such as healthcare, manufacturing, and infrastructure.

Potential Benefits to Society: This shift has the potential to deliver significant benefits. In healthcare, AI could support more personalised and timely treatment decisions. In science, it may accelerate discovery by identifying patterns in complex data. In industry, it could improve efficiency, safety, and productivity by optimising operations and predicting failures before they occur.

New risks introduced by greater autonomy: However, this transition also introduces new risks. As AI systems begin to act in the real world, errors can have direct consequences, raising concerns about safety and reliability. Questions of accountability become more complex when systems operate with greater autonomy. There are also ongoing challenges around bias, data quality, and the concentration of technological power in a small number of organisations.

The case for effective governance: As a result, effective governance will be essential:

- Clearer rules: Clearer rules for how AI is applied in high-impact areas
- Improved oversight: Improved systems for monitoring and oversight
- Greater experience: Build technical understanding among regulators.
- Focus on Regulating Applications, Not the Technology: Many experts emphasise that the focus should be on regulating the use of AI, rather than restricting the underlying technology itself.

Overall, this shift represents a major step forward in AI capability, with the potential for substantial societal benefit—but it will require careful management to ensure those benefits are realised safely and responsibly.

2. Data as the Central Strategic Asset

Data was identified as the foundation of AI capability.

- There is a growing trend towards treating data as a national asset, raising questions of sovereignty and control.
- Approaches such as federated learning and synthetic data may allow innovation while protecting privacy and national interests.
- However, legal uncertainty—particularly around intellectual property—remains a major constraint.

3. Emergence of Agentic AI

AI systems are increasingly moving from passive tools to "agentic" systems capable of taking actions and making decisions.

- These systems are already being deployed in industry.
- They raise new issues of accountability, safety, and coordination—particularly when multiple agents interact.
- Governance frameworks will need to evolve to address autonomy, auditability, and responsibility.

4. Infrastructure, Energy, and Technology Pathways

The future of AI will depend on underlying technologies and infrastructure.

- Photonics and quantum-related technologies were highlighted as potential solutions to the growing energy demands of AI.
- The UK has existing strengths in these areas but will require sustained investment to remain competitive.

5. Regulation: Focus on Applications, Not Technology

There was broad agreement that:

- Regulation should focus on how AI is used, particularly in high-impact applications, rather than the underlying technology itself.
- Premature or overly restrictive regulation could hinder innovation and reinforce dominance by a small number of global players.
- At the same time, regulators must develop deeper technical understanding and more adaptive approaches.

6. Long-Term Risks vs Immediate Priorities

While issues such as artificial general intelligence (AGI) and superintelligence were discussed, most speakers considered these to be longer-term concerns.

- The immediate priority should be addressing current challenges: deployment, governance, reliability, and economic impact.
- There was caution against allowing speculative future risks to distract from present policy needs

.

7. Human Impact: Skills, Work, and Society

A key theme was the need to refocus on people:

- AI will reshape labour markets, potentially reducing entry-level roles and changing skill requirements.
- There is a pressing need for education, reskilling, and workforce adaptation.
- AI also presents opportunities, particularly in personalised education and healthcare, if directed towards public benefit.

Overall, the meeting highlighted both significant opportunities and material risks:

- The UK has strengths in research, data, and emerging technologies that could support global leadership.
- However, there are challenges around governance, infrastructure, legal frameworks, and skills.
- A consistent message was the need for a balanced approach—supporting innovation while ensuring accountability, diversity, and public benefit

EXECUTIVE SUMMARY - IMPLICATIONS FOR STAKEHOLDERS

The discussion highlighted a number of implications for different stakeholder groups as AI systems evolve from language-based models towards more advanced, agentic forms that can act in real-world environments. The points below reflect the themes and perspectives expressed during the evidence session.

1. Policymakers (Government and Parliament)

Implication:
Policymakers will need to respond to the transition towards agentic AI by prioritising high-impact applications, addressing constraints in access to training data, and balancing national data sovereignty with international collaboration.

- Focus on regulating **applications**, particularly where AI has significant impact on people's lives
- Avoid regulating **research and core technologies**, which may hinder innovation and open collaboration
- Address **data-related challenges**, including access, ownership, and intellectual property
- Recognise that access to **globally representative training data is uneven and often limited**, particularly for SMEs and less widely represented languages and regions
- Consider the role of **data sovereignty**, including whether training data should be treated as a national asset while enabling collaboration through approaches such as federated model development
- Support **investment in emerging technologies** (e.g. photonics, quantum-related computing)
- Recognise the importance of **education, skills, and workforce adaptation**

2. Regulators

Implication:
Regulators will need to adapt to the emergence of **agentic AI systems**, particularly those that act autonomously or interact with other systems, where failures may have direct real-world consequences.

- Focus on **outcomes and real-world impact**, particularly in high-risk applications
- Develop stronger **technical understanding** of agentic AI systems
- Support the development of **standards, benchmarking, and verification methods**
- Address **accountability and responsibility** for AI-driven decisions
- Consider how to oversee **multi-agent interactions** and system-level risks
- Ensure that regulatory approaches remain **adaptive** as technologies evolve

3. Industry (Technology Firms, Corporations, SMEs)

Implication:
Industry will be responsible for deploying agentic AI systems in real-world contexts, requiring greater attention to accountability, safety, and system integration.

- Ensure clear responsibility and accountability for AI system outcomes
- Manage risks associated with autonomous and decision-making systems
- Address challenges in human–AI interaction, particularly handoffs between systems and people
- Invest in data management, including addressing limitations in training data availability
- Prepare for workforce transformation, including changes in roles and required skills
- Recognise emerging dynamics such as machine-to-machine interactions

4. Research Community (Universities and Public Research Organisations)

Implication:
The research community will play a critical role in advancing AI beyond language-based models while addressing challenges in reliability, coordination, and data access.

- Continue developing AI systems that interact with the real world, beyond language-based models
- Address challenges in reliability, reasoning, and system behaviour
- Advance research on multi-agent systems, including coordination and failure risks
- Explore approaches such as federated learning and synthetic data to address limitations in access to training data, particularly for SMEs and underrepresented languages and regions, while supporting privacy and data sovereignty
- Contribute to standards, evaluation, and technical understanding

5. Civil Society (Public Interest Groups and Communities)

Implication:
Civil society will play an important role in shaping how agentic AI is understood, governed, and accepted, particularly in relation to trust, fairness, and societal impact.

- Raise awareness of AI's impacts on individuals and society
- Highlight issues of fairness, bias, and inclusion
- Engage in discussions around data use, privacy, and trust
- Monitor the broader societal effects of AI, including impacts on employment, education, and public services

Cross-Cutting Implication

Across all stakeholders, a central challenge is managing the transition from language-based AI to agentic systems that act in the real world, where failures have direct consequences and require clearer accountability, stronger technical oversight, and new approaches to data governance.

Overall Assessment

The discussion underscored that while AI presents significant opportunities across sectors, the shift towards more autonomous and interconnected systems introduces new and more complex risks.

Addressing these will require coordinated action across policy, regulation, industry, research, and civil society, with a focus on maintaining innovation while ensuring safety, accountability, and public trust.

UK AI STRATEGY: IMMEDIATE PRIORITIES

The UK faces a narrow window to shape the next phase of artificial intelligence.
Based on the evidence presented, five strategic priorities emerge

1. Position the UK as a Global Leader in Agentic AI Deployment

The next wave of AI will be defined not by larger models, but by systems that act in the real world.
The UK should prioritise:

- Deployment of agentic AI in high-impact sectors (healthcare, infrastructure, public services)
- Regulatory sandboxes for real-world AI systems
- Public-sector adoption as a catalyst for innovation

2. Treat Data as Strategic Infrastructure

Access to high-quality, representative data is now the primary constraint on AI capability.
The UK should:

- Develop national data infrastructures in key sectors (e.g. NHS, energy, mobility)
- Enable federated data ecosystems to balance sovereignty and collaboration
- Clarify intellectual property frameworks for training data and synthetic data

3. Build Sovereign Capability in AI Infrastructure

AI competitiveness will increasingly depend on compute, energy, and specialised hardware.
The UK should:

- Invest in next-generation compute (including photonics and quantum-adjacent technologies)
- Ensure access to scalable, energy-efficient infrastructure
- Support domestic capabilities in critical AI supply chains

4. Lead in Governance of Autonomous Systems

As AI systems become more autonomous, governance must evolve from static approval to continuous oversight.
The UK should:

- Focus regulation on applications and outcomes, not core technologies
- Develop frameworks for auditability, accountability, and system-level risk
- Establish standards for multi-agent and cross-organisational systems

5. Invest in Human Capital and Institutional Readiness

AI will reshape labour markets and decision-making systems.
The UK should:

- Expand investment in skills, reskilling, and interdisciplinary education
- Build technical capability within regulatory institutions
- Support new organisational models for human–AI collaboration

Strategic Imperative:

The UK's opportunity lies not in competing on scale alone, but in leading the safe, effective deployment of AI in complex real-world systems.

APPG AI Co-Chair:
Allison Gardner MP

APPG AI Co-Chair:
Lord Clement-Jones CBE

Secretariat & Rapporteur:
Professor Birgitte Andersen

EVIDENCE

Evidence Giver:
Yann LeCun

Evidence Giver:
Prof. Viv Kendon

Evidence Giver:
Prof. Mihaela van der Schaar

Evidence Giver:
Bob De Caux

Evidence Giver:
Mark Taylor

STRATEGIC TAKEAWAY

AI is heading towards a completely new paradigm of general, real-world intelligence, and must remain 'open' and decentralised to avoid systemic risks.

Yann LeCun

Executive Chairman, AMI Labs; Professor, New York University; Founding Director, Meta FAIR & former Chief AI Scientist, Meta.

Opening Introduction

My name is Yann Le Cun. I am the I'm a professor at New York University. I'm also executive chairman of a new startup company called Advanced Machine Intelligence, AMI Labs. I'm the former chief AI scientist at Meta, which I left about a month ago. I've been interested in a number of different policy issues related to AI, as well as the design of future AI systems.

The Future of AI Will Be Fundamentally Different

So the first thing I want to say is that AI technology is going to change in major ways over the next few years. And so it's very dangerous to extrapolate the capabilities of current AI systems and imagine that future AI systems would just be more powerful versions of current systems. They will be very different. They promise to accelerate progress in science and medicine. and also automate a lot of tasks in industry in the future.

Currently, the capabilities of AI systems is largely focused, at least if we're talking about large language models, is mostly focused on language. So manipulating and generating language. But they're not appropriate for the real world. So far, approaches, AI approaches to deal with the real world, we're talking about video, sensor data, Data from industrial processes, from physical systems, robots, whatever they are, have been piecemeal. It's not been sort of an approach that treats all of those systems as kind of with a single methodology.

We have specialised systems for, example, predicting the structure of proteins or predicting the properties of materials or things like that. but we don't have sort of a general approach. I think in the future, one revolution that will occur over the next few years is systems that understand the real world, but have the same level of generality as LLMs have in the context of language. So we're going to have a new revolution in AI, and in fact, that's kind of the purpose of the company I'm building.

AI as the Mediator of Information

The second point is that most of our digital diet for everyday citizen is going to be mediated by AI systems. To some extent it is really the case that much of the information that we consume is being picked and chosen and filtered by AI systems. But in the future, I mean even more so, we'll be talking to our AI assistant and it will be telling us, you know, answer questions that we have, help us in our daily lives at all times. Perhaps it will live in our wearable devices, mobile devices, and things like this.

Risks of Centralised Control

So in a future where all of our digital diet is mediated by AI systems, we cannot afford to have those AI systems be under the proprietary control of a handful of companies on the west coast of the US, or coming from China. So there is a need for sovereignty, but also for the preservation of cultural diversity, linguistic diversity, value systems, et cetera, centers of interest, political biases, et cetera, for AI platforms to be diverse.

We're going to need a high diversity of AI systems for the same reason we need a high diversity of news media, because they're all going to be biased. This is the biggest venture I see in the future of AI, which is capture of information by a small number of companies through proprietary systems.

The Case for Open AI Platforms

So this would require open platforms, platforms that are open source, open weight, in the context of AI, so that anybody can fine-tune them for their own interest or vertical applications or biases regarding whatever they are. And currently, unfortunately, Beta used to be a leader in providing open source systems and over the last year has lost ground and now it's not clear who's going to pursue this.

So there is a need, I think, for governments around the world that are neither the US nor China to perhaps federate under a program to promote open source AI platforms. that may eventually constitute a repository of all human knowledge, but because of their openness, will be customisable by anybody and provide the diversity that we seek.

Closing

That was my main points, so I'll stop here. Thank you very much.

Views of Yann LeCun

Given that this was an online presentation and mindful of potential technical instability, the presentation proceeded directly to questions.

The exchange explores data access, privacy, and regulation in AI, with Yann LeCun responding to questions from policymakers and panelists.

1. Open Source and Access to Training Data

A concern was raised about whether open-source AI is enough, given limited access to high-quality and diverse training data (e.g. news and copyrighted material).

Yann LeCun's view:

- The difficulty of accessing global data reinforces the need for open, collaborative AI.
- No single company can gather all necessary data, especially for:
  - Underrepresented languages (e.g. in India, Africa, Indonesia, parts of Europe).
- He proposes a federated approach:
  - Countries/regions collect and keep their own data locally.
  - They contribute to training a shared global model.
  - The final system benefits from global knowledge without centralising raw data.

2. Privacy and Personal AI Assistants

The discussion shifts to the individual level, focusing on privacy in a future with AI assistants.

Yann LeCun's view:

- AI assistants will become deeply integrated into daily life—effectively "your best friend", knowing what you see and hear.
- This creates serious privacy concerns.
- Technically:
  - Some AI can run locally on devices.
  - But most powerful capabilities will still require cloud-based processing.
- Therefore, the key issue is:
  - Where computation happens (local vs cloud)
  - Who controls the data, shaped by national rules and sovereignty

3. Regulation and Legislation of AI

A policymaker asks how governments should regulate AI.

Yann LeCun's view:

a) Regulate applications, not research

- AI systems that impact people's lives should be regulated.
- But research and development must remain open and unrestricted.

b) Do not restrict scientific exchange

- Regulating R&D would:
  - Slow innovation
  - Prevent collaboration
  - Damage open-source ecosystems

c) Avoid regulatory capture

- Overregulation risks favouring large tech firms, who can comply more easily.
- This would increase concentration of power - a problem he warns against.

d) Encourage open-source AI

- Governments should create incentives for open platforms.
- This supports diversity, competition, and wider participation.

In the feedback from parliamentarians Yann LeCun consistently argues that:

- Open, federated AI systems are essential for global inclusion
- Data sovereignty + collaboration can coexist
- Privacy will be a major challenge in AI adoption
- Regulation should target usage, not innovation
- Overregulation risks concentrating power in a few companies

In short: He advocates for a decentralised, open, and collaborative AI ecosystem, supported—but not stifled—by policy.

Key Takeaways from Yann LeCun's Evidence

Core view: AI is heading towards a completely new paradigm—general, real-world intelligence—and must remain open and decentralised to avoid systemic risks.

LeCun's argument is fundamentally about trajectory and power.

- He stresses that it is misleading to extrapolate from today's LLMs —future AI systems will not just be "bigger chatbots" but systems that understand, reason about, and act in the physical world.
- This implies a shift toward unified models that handle perception, action, and reasoning, rather than fragmented, task-specific systems.

At the same time, he introduces a political economy concern:

- AI will increasingly mediate all human information consumption (via assistants, devices, etc.), effectively shaping reality for individuals.
- If controlled by a few firms or countries, this creates risks of bias, cultural homogenisation, and concentration of power.

His solution is strongly normative:

- Build open, federated AI ecosystems, where:
  - Data remains locally controlled (sovereignty)
  - Models are collaboratively improved
- Governments (especially outside US/China) should actively support open AI platforms.

On regulation:

- Regulate applications (impact), not research (innovation)
- Overregulation risks entrenching big tech dominance

In essence:
LeCun sees AI as both a technical revolution and a geopolitical infrastructure and argues for openness as the key safeguard.

STRATEGIC TAKEAWAY

Quantum computing will play a supporting, specialised role in AI, but expectations should be realistic and grounded in current scientific limits.

Prof. Viv Kendon

Professor of Quantum Technology, Department of Physics, University of Strathclyde

Introduction and Background

I'm a professor in quantum technology at the University of Strathclyde in Glasgow. I've been developing underpinning science for quantum computing for the past 25 years. And I've been also at Imperial, Leeds, Durham in that time, so I've moved around a bit. And before this, that So way back, I ran a really early internet service provider. So I helped to kickstart global networking back in those days. So I have a very long perspective, if you like, on the evolution of information technology.

But anyway, what I'm here to do today is to explain how quantum computing specifically intersects with future AI, insofar as we know this. And there are many open questions.

The UK Quantum Strategy Vision

So As you probably know, there's a national quantum strategy. And Mission 1, this is government strategy, Mission 1 is a 10-year plan for quantum computing. And it is summarised, and I'm quoting the wording, by 2035, there will be accessible UK-based quantum computers capable of running 1 trillion operations and supporting applications that will provide benefits well in excess of classical supercomputers across key sectors of the economy.
So that's the vision. And it's more than a prediction. It's a plan. People are putting money behind this. So given how much AI is running on supercomputers now, this implies we expect quantum should be key players in the future. taking it beyond supercomputing capabilities.

Current State of Quantum Computing

Right, so that's my starter, but now let's see how many of my points I can get in as to whether I think this is likely. The current situation with quantum computing in general is that despite the hype, we're still in discovery science and early TRL engineering. We maybe think 1980s classical. We're kind of at that stage with quantum. as you're thinking about where we are today, what our phones and computers now do.
But it's worth remembering that the science and engineering for both quantum and AI have been developing over many decades. They are at a steady pace and there is a solid foundation for both of them. It's not, the hype cycle makes it seem like it's come suddenly, but that's not the case. We do have good foundations for that.

Quantum Machine Learning Today

The current situation for quantum machine learning, we tend to talk about machine learning rather than AI at the moment. There's a lot of good science being done, but it's very much at the science stage. They're looking at many of the different possible applications like classification, generative data, time series analysis, that sort of thing.
There is so far scant evidence that quantum provides a step advantage over AI in applications. There are some grounds for thinking we may be able to speed up training, and that's often the costliest part. So if that's the case, and then we can even deploy it on classical computers, that will be a win.

The training process means it's less sensitive to the errors that quantum computers are plagued by, so it may be an early win if it works. It also means that room temperature photonic devices are currently the front runner for looking at AI applications. There are UK companies exploring this. It would be really good for the UK because we've got a huge photonics industry that's very high tech and very much into this. So this is to be encouraged.

Data Challenges

Data is everything for AI, right? So, and current and near future quantum computers are going to be too small to process data, large volumes of data. So the compression algorithms are an important area of research for both quantum and classical.
So you might say, okay, well, what about quantum data if we can't do large volumes of classical? But we don't have quantum data. We don't have long-term stable quantum data storage. We have small amounts of data storage that are related to doing quantum communications, but it's really hard, and I don't think the drivers are there to try to get further with that.

Possible Future Trajectories

Possible future trajectories, very speculative, but if we assume that the mission is met by 2035, quantum computers then become cost-effective for large-scale supercomputing. It doesn't imply AI is going to be one of the applications they're good for, but that doesn't follow.

It's more likely they will be accelerators for more specific tasks. And as I said, the photonic angle is more likely to be the hardware that's of interest.

Energy and Hardware Considerations

So I'll skip to some You've got about a minute and a half left. Yeah, I've got a little more time here, but I want to make some actual recommendations for what needs to be done. So the Just maybe an aside, that optical computing didn't make it during the 80s, but that's because the hardware wasn't good enough. It is now, and off the back of the quantum computing photonic developments, we may see a second coming for what we can do.

And the point is, what we think it will be lower energy, and that's going to be absolutely critical, is getting more compute for less energy. The idea of putting a small nuclear reactor in a data center is not one direction we want to be going in really.

Recommendations and Conclusion

Okay, so wrap up. Fund the science and engineering, well I would say that, but why? So that we are leading the discoveries, we're out there world-leading, collaborating with the other top scientists, so we're agile, so we can pivot because we know what's coming and we understand what's coming.

And we've already had China mentioned and India. We've been training their top scientists for decades. They are putting huge amounts of investment into the science and engineering and we need to keep our place in that. It's not just about being able to buy turnkey solutions; you have to have the people who really understand how to use them.

So that training and knowledge is crucial, especially when data security and data integrity is involved. So data is crucial, as I said. So we also need to fund data curation and storage in the UK. And the data ownership needs strong protections, both in law and in practice, which means we need the technical capability to look after it, not just the law saying this is forbidden, but you can actually, with the technology, make sure that the law is carried out.

The role of quantum computing alongside classical supercomputers

Data Processing Limits of Quantum Computing

Prof. Viv Kendon's view on whether quantum computers will be able to handle the large volumes of data that current supercomputers process:

- Even looking 10 years ahead, quantum computers will not handle large data volumes.
- Their advantage comes from processing power, not data capacity.
- This makes them fundamentally different from classical supercomputers.

Will Quantum Replace Supercomputers?

Prof. Viv Kendon's view on whether quantum computing could replace classical systems.

- Quantum computers will not replace supercomputers.
- The UK strategy itself does not suggest replacement, but rather integration.
- The future lies in:
  - Hybrid systems, combining quantum and classical computing
  - Using each where it is most efficient
- Key principle:
  - Do not use quantum computing for tasks that classical systems already handle efficiently
  - Unless quantum hardware becomes significantly more energy-efficient

Efficiency and Complementarity

- Quantum computing may offer advantages in:
  - Specialised, high-complexity processing tasks
- Classical computing remains superior for:
  - Large-scale data handling

Key Takeaway: The role of quantum computing alongside classical supercomputers

Prof. Viv Kendon emphasises that:

- Quantum computing is not a replacement technology
- Its strength lies in processing, not data volume
- The future is integration with classical systems
- Value will come from targeted applications and efficiency gains, not wholesale disruption

In short: Quantum computing will complement—not replace—existing supercomputing infrastructure.

Prof. Viv Kendon – Key Takeaways

Core view: Quantum computing will play a supporting, specialised role in AI, but expectations should be realistic and grounded in current scientific limits.

Kendon's perspective is deliberately measured and corrective.

- She emphasises that quantum computing is still in a "1980s stage" of development, despite hype.
- There is no clear evidence yet that quantum machine learning outperforms classical AI in real applications.

Her key distinction:

- Quantum strength = processing power (specific tasks)
- Classical strength = data handling and scalability

This leads to a hybrid future:

- Quantum systems may:
  - Accelerate certain computations (e.g. training optimisation)
  - Improve energy efficiency (e.g. photonics)
- But they will not replace classical AI systems, especially for data-heavy tasks.

She also highlights strategic implications:

- The UK must invest in:
  - Science and engineering capability
  - Data infrastructure and ownership
  - Skills and talent pipelines
- Without this, countries will depend on external technologies they don't understand or control

In essence:
Kendon reframes quantum AI from "disruption" to "complementary infrastructure" and stresses long-term capability building over hype.

STRATEGIC TAKEAWAY

The real breakthrough in AI will come from agentic systems that make decisions in complex real-world environments, with healthcare as the prime opportunity.

Prof. Mihaela van der Schaar

John Humphrey Plummer Professor of Machine Learning, Department of Engineering, University of Cambridge

Introduction and Background

I'm Mihaela van der Schar. I'm the John Humphrey Plummer Professor of Machine Learning, AI and Medicine at the University of Cambridge, and I direct the Cambridge Centre for AI and Medicine. For over 2 decades, my research has focused on developing cutting-edge AI and machine learning that can make an impact in the real world. And I call this reality-centric AI.

A Provocative Starting Point

So let me start with a slightly provocative statement. Most generally, good ideas in AI and machine learning don't come from inside machine learning itself. They come from engaging with ambitious, real-world problems and the challenges they pose.
So it is from this perspective that I want to focus today on what I believe will drive the next major leap in AI capabilities and why the UK has a unique opportunity to lead this shift.

AI, the NHS, and a National Opportunity

My core argument is that by tackling the real-world challenge of improving how the NHS functions, and by making clinical trials smarter and more adaptive, we are not just applying AI to healthcare, but we are creating one of the world's best testbeds for developing the next generation of agentic AI systems, which can improve both NHS, but also AI itself.

So crucially, this is not about experimenting on patients, but about improving the healthcare system itself, its processes, pathways, and decisions, using AI for simulation, digital twins, as well as careful governance to learn safely and responsibly.
So I choose this focus because healthcare is where AI capability, governance, and national strategy collide, in my opinion, and where the UK can realistically become world-leading.

Rethinking the Focus of AI in Healthcare

The central point I want to make is that much of the current debate about AI is overly focused on two things. Building ever larger models and using AI primarily for drug discovery. Both matters, but neither is where the biggest gains for society now lie.

We already have a vast and growing arsenal of drugs. The binding constraint is no longer discovering more molecules but delivering better healthcare, getting the right treatment to the right patient at the right time using limited resources wisely.

So this is why the focus must shift, in my opinion, towards the NHS and healthcare systems themselves and toward rethinking how we can generate evidence about which treatments are useful and for which patients.

From Passive AI to Agentic Systems

So the next major advance in AI for healthcare will therefore not come from models that are simply larger or better at predicting risk or answering medical questions. It will come from smarter, authentic AI systems that can decide when to act, how to act, and when not to act in complex resource-constrained environments like the NHS.

So this is a shift away from scaling models for their own sake and toward AI system for decision-making in the messy real world.
So most AI used in healthcare today is passive. It flags risks, scores patients, detects patterns, but the burden of deciding what to do, when to do it, and whether it is worth doing still sits entirely with clinicians and managers. This is where delay, inefficiency, and unnecessary costs arise.

So high risk for the patient does not automatically mean intervention is beneficial, and low risk does not always benefit, mean we should wait. Healthcare is fundamentally a problem of timing, trade-offs and uncertainty over time, and challenges that larger models alone in AI cannot solve.

Introducing “Genius” and Digital Twins

So next generation agentic AI, which I call Genius, are designed specifically to address this gap. A genie is not a chatbot and not simply a powerful AI agent. It is an AI system designed to reason over time, know what it knows and what it doesn't know, when it's certain and it's not, evaluate alternative actions, and learn from outcomes after deployment.

Instead of asking what is the risk, a genie asks what happens if we intervene now versus later, which intervention will actually change the outcome, and when does intervention stop being worth in the cost for the burden of the patient.

So to do this safely, Genius will rely potentially on simulation models. Let's call this AI-enabled digital twins. And this combination of authentic AI and digital twins, I believe, is particularly powerful for the NHS.

Consider, for instance, a patient with multiple long-term conditions. A traditional AI model might tell us that this patient is at risk of hospitalisation. A GENI, though, using a digital twin of the patient and the surrounding care pathway, can go further. It can evaluate whether an early intervention will actually prevent admission, whether community care is sufficient, or whether intervention at this point will make little difference and consume scarce resources.

Crucially, it can adapt its recommendation as the patient's condition evolves, as the hospitalization pressure change, and new evidence emerges.

Transforming Clinical Trials

The same approach has profound implications, in my opinion, for how we think clinical trials, and this is where regulation also matters.
As medicine becomes more personalised, the question is no longer how to test more drugs faster, but how to determine which drugs are truly useful for which patients and under which conditions.
And many trials fail today, not because the underlying therapy is ineffective, but because the trial design was fragile, the wrong patients were recruited, assumptions about progression or utterance were wrong, or operational constraints were underestimated.

So Agenda AI allow us to rethink clinical trials as adaptive simulated systems rather than fixed protocols. And using digital twins of patient population and trial populations, Genius can generate candidate trial designs, stress test them under realistic conditions, explore failures modes, and recommend designs that are robust rather than just optimistic.

Governance and Accountability

So A crucial point for this committee is governance. There is a concern that AI can become more autonomous. As it becomes more autonomous, responsibility becomes blurred.

In fact, well-designed agentic AI systems make governance, in my opinion, easier, not harder, because such next-generation AI agents, genius, can simulate decisions before acting. They can track uncertainty, unlock the reasoning. The recommendations can be audited, stress-tested, and monitored continuously after deployment.
This is far safer than relying on everlasting static models that are approved once and then quietly drift out of validity as conditions change.

Conclusion and Recommendations

So to conclude, the UK has a unique opportunity here in my opinion. The NHS, which is scale, longitudinal data and integrated structure, can become a global exemplar for how authentic AI and digital twins are deployed responsibly to improve healthcare outcomes as well as system efficiency.

The same is true for clinical trial, trials where the UK could lead globally in defining how adaptive AI-enabled clinical trials are designed, overseen, and regulated.

And if the UK chooses to invest strategically in this next generation agentic systems, it can shape not just how powerful AI becomes, but how trustworthy and beneficial it becomes in practice.

So my 3 recommendations are, first, UK should make the NHS the world's leading testbed for agentic healthcare AI, focused on decision-making, not just prediction, focusing on improving healthcare systems.
Second, the UK should invest in agentic AI for clinical trials, positioning itself as a global leader in safe, adaptive, evidence-generating AI for drug approval.

And finally, UK regulations should move away from treating AI as a medical device that is approved once and left unchanged, but instead it should think more creatively about regulation to support safe monitored learning after AIs are introduced with clear accountability for how systems adapt and are audited and are corrected over time.

Thank you.

Prof. Mihaela van der Schaar – Key Take-aways

Core view: The real breakthrough in AI will come from agentic systems that make decisions in complex real-world environments, with healthcare as the prime opportunity.

Prof. van der Schaar's argument is both conceptual and strategic.

1. Critique of current AI focus

- Too much emphasis on:
  - Scaling models
  - Drug discovery
- Not enough on improving real-world systems (e.g. healthcare delivery)

2. Shift to "decision intelligence"

She argues the next leap is:

- From prediction → action
- From static models → adaptive systems over time

Her concept of "GENIUS" systems (agentic AI):

- Evaluate:
  - What happens if we act now vs later?
  - Which intervention actually changes outcomes?
- Incorporate:
  - Uncertainty
  - Resource constraints
  - Dynamic environments

3. Role of digital twins

- AI systems should simulate outcomes using digital twins before acting
- This allows safe experimentation and better decisions

4. UK opportunity (NHS)

- The NHS offers:
  - Scale
  - Longitudinal data
  - Integrated system
- This makes it a global testbed for next-gen AI

5. Governance rethink

- Agentic AI can actually improve governance:
  - Decisions can be simulated, audited, and monitored continuously
- Calls for regulation that supports adaptive, evolving AI systems

In essence:

She shifts the debate from "how powerful AI is" to "how useful and accountable it is in real systems", positioning the UK to lead via healthcare.

STRATEGIC TAKEAWAY

Agentic AI is already becoming operational in industry, and the central challenge is engineering safe, accountable autonomy at scale.

Bob De Caux, Chief AI Officer, IFS

Introduction and Background

I'm Bob De Caux, Chief AI Officer at IFS. And today I'm going to build, I think, a lot on what Mahela talked about with Gentic AI and its use within industry at scale.

I did my PhD in multi-agent systems back before agents were quite so fashionable. I was CEO of a machine learning company in London, focusing on explainable AI. And for the last seven years, I've headed AI strategy at IFS, where we make enterprise software that's used in the real economy, energy and utilities, construction, manufacturing.

Agentic AI in Industry

Agentic AI is critical to many of the businesses we deal with because our software is where agents become real through work orders, schedules, parts, systems that connect decisions to actual operational outcomes.

In these environments, AI is not just generating text, it's often initiating actions that affect safety, uptime, cost, and public trust.

What Are AI Agents?

So, what are AI agents?

They're software systems typically based around an LLM that can take a goal, break it into steps, and act to achieve it, often autonomously, observing and adjusting along the way.

And the reason they've exploded in popularity is because the potential is enormous. The economic value of generative AI could be in the trillions globally, with an increasing share coming not from chatbots, but from agentic workflows.

Current Use: Controlled Autonomy

In industry today, agents setting goals does not mean unconstrained autonomy. Many production deployments are bound to workflow agents, drafting or summarising content, orchestrating steps under supervision, or co-piloting specialist roles such as maintenance engineers.

There remains huge potential for growth in these areas for businesses and for personal productivity.

But organisations are already moving to multi-agent patterns. A planner proposes steps, a policy agent checks constraints, an executor takes actions. The leap to real system-level autonomy is driven as much by architecture, tools, and governance as it is by bigger and better AI models.

Tool Selection and Standards

To help make the leap at scale, a key development area for agents is tool selection. And that's because, as Jan mentioned, large language models are language engines. They're not inherently reliable at precise calculation.

Industrial planning problems, such as scheduling or inventory trade-offs, need deterministic correctness. So while LLMs can reason and propose, they need other tools to be able to fulfil their goals.

Industrial environments can have thousands of these tools, creating a work order or scheduling a crew. So an agent must be able to reliably choose the right tool for the right job at the right time.

And this is where standards such as Model Context Protocol or MCP matter. They make tool access more consistent and governable. But we are moving towards a world of tool discovery through standardised interfaces.

That will unlock scale, but it introduces new risks. Power tools can be found. Tool metadata can mislead, and tool ecosystems can become the new supply chain. So capability has to be paired with tool governance.

Reliability, Safety, and Determinism

Now this tool using autonomy can only scale if it is consistent and auditable. In critical systems, even a very low error rate can be unacceptable.

So another key area of development for agentic AI is determinism by design, grounding actions and systems of record and putting in deterministic guards such as approvals and rollback patterns so the agent can't guess its way into causing harm.

Now, incomplete information is very normal in these systems, so a production agent must understand what's missing or be able to ask a human. And if it can't do that, it must degrade safely to recommending rather than executing actions.

We need observability and auditability of agents, especially across the handoff between agents and humans.

Scaling, Cost, and Architecture

Now, looking forward, a key consideration for Agentic AI will be the use of runtime compute, because agents don't just answer once, they plan, they call tools, and they retry.
So scaling this across industry has real cost and energy implications. And that's why we're already seeing agent stacks becoming more tiered. Industry-specialised small language models, or SLMs, handle routine steps cheaply, while larger models are used selectively for ambiguity and complex trade-offs.

Agents will improve fastest, as Mahela said, when we train and test them in simulated environments, so recovery behavior is safe before they enter production.

And to go further, especially in critical infrastructure, we'll likely need a version of the world models that I know are beloved of Jan. The impractical terms are probably more like a digital twin of the enterprise, so agents can plan and anticipate consequences before they act.

Governance and Security

So how do we govern this new level of autonomy, especially when agents transact across organisations? The model should be workforce governance. Responsibility must be explicit. The organisation that deploys and benefits from agents remains responsible for outcomes, as it is for employees and operational software.

In practice, agents need identity and sponsorship, role-based access, separation of duties, and complete audit trails. And at scale, governance also means operational discipline for agents, assurance metrics for behaviour over time, change control tools, and data governance, including data provenance.

Now, the standout new security risk with agents is prompt injection, agents ingesting untrusted content that can carry hidden instructions. With tool access, this becomes not just a bad answer to a question but a bad action, especially at agents' hands-offs and cross-organisational boundaries.

Human Workforce Implications

Finally, I'd like to mention the effect on the human workforce. Agentic AI will change the shape of many roles. Less routine administration and coordination, more supervision and exception handling, which can raise productivity and free experienced staff for safety-critical judgment.

But we must manage the human factors. The risk often isn't the model itself, it's the handoff to humans through over-trusting agents or automation complacency.
We must also manage the risk to entry-level pathways for the workforce, as junior tasks become automated, while still building domain expertise for the future, which is already a challenge for many companies.

Conclusion and Recommendations

To close, - the opportunity is huge. Agents are the practical route from AI proofs of concept to real productivity.

Here are three recommendations.

1. Firstly, make accountability for autonomous action explicit. High-impact agents need an operator of record, auditable tool actions, and incident reporting.
2. Secondly, drive a UK baseline for secure, interoperable, tool-connected agents, with separation of duties and portability across systems and supply chains.
3. And thirdly, treat agent adoption as change management, not as an IT installation, with practical programs to help organisations redesign controls, train supervisors, and prevent uncontrolled agent sprawl across the enterprise. whilst also focusing on entry-level pathways to build domain knowledge.

Done well, the UK can lead in making agents work safely in the real economy.

Thank you.

Key Take-aways – Bob De Caux

Core view: Agentic AI is already becoming operational in industry, and the central challenge is engineering safe, accountable autonomy at scale.

De Caux provides a pragmatic, deployment-focused perspective.

1. What agentic AI really means in practice
- Not just chatbots, but systems that:
  - Plan tasks
  - Call tools
  - Execute actions in real workflows
- Already used in sectors like:
  - Energy
  - Manufacturing
  - Construction

2. Controlled autonomy (current reality)
- AI agents are:
  - Supervised
  - Embedded in workflows
- Moving toward multi-agent systems:
  - Planner → policy checker → executor

3. Key technical challenges
- Tool selection: choosing the right system/action reliably
- Determinism: avoiding probabilistic errors in critical tasks
- Observability: tracking what agents do
- Safety fallback: agents must defer or degrade when uncertain

4. New risks
- Prompt injection → bad actions (not just bad outputs)
- Tool ecosystems = new supply chain risks
- Cross-organisational autonomy increases complexity

5. Governance model
- Treat AI like a workforce:
  - Assign responsibility
  - Track actions
  - Audit decisions
- Organisation deploying AI remains accountable

6. Workforce impact
- Jobs shift toward:
  - Supervision
  - Exception handling
- Risk: erosion of entry-level pathways

In essence:
De Caux translates "agentic AI" into real-world systems engineering, emphasising control, safety, and organisational accountability.

STRATEGIC TAKEAWAY

AI progress is increasingly constrained by data access and legal frameworks, making synthetic data a critical, but complex, solution.

Mark Taylor
Partner - Head of Digitalisation (UK)
Osborne Clarke

Introduction and Background

Thank you and good evening. I'm head of digitalisation, but the other role I fulfil at the firm is leading our data practice. And as such, I work regularly with clients in a range of sectors on data issues, including the use of data with AI systems. And that's why I'm going to talk a little bit about synthetic data.

The Growing Importance of Data in AI

And I think that fits with some of what we've had from today. Jan has mentioned some of the questions around data. Other speakers, I think, have also touched on this.

And we're currently in a period of quite rapid development and deployment of AI tools and solutions. And as everybody's made clear, and I think is well known, a key element in their creation and training of those solutions, whatever their form, is data, but there are reality practical issues now emerging, certainly I find, with obtaining sufficient high quality and human-generated data to train models and solutions.

And rightly, privacy law imposes limits on how personal data can be used.

What Is Synthetic Data?

So one solution that is not necessarily novel, but I think is going to come more to the fore is the use of synthetic data. And so that's what I wanted to cover.
What is synthetic data? Effectively, it's data that's been artificially generated by a computer algorithm rather than being recorded from real-world events. So it's almost like a data twin of real-world information.

So for example, instead of using real people's records, say medical records, to train a solution, as Mihedra was talking about, a mathematical model could be used to create a statistical ghost of those records. That doesn't, the ghost record doesn't directly represent individual people, but it could be synthesised to behave like a group of real people. It might have height, weight, blood pressure that are statistically plausible within a specific population.

And synthetic data can take different forms. It can be fully synthetic, entirely new data points. It can be partially synthetic, where you're mixing real world and other identifiers as well.

Why Synthetic Data Matters

Why is this important? As I've sort of touched on already, in some areas there is a limited supply of high quality public data, or obtaining that data is hard. Synthetic data effectively allows organizations to create a greater quantity of representative data which can be used to train AI.

And in some scenarios, it's essential. For example, if you're driving, in a simplistic example, training a self-driving car, you cannot... and wait for real-life accidents to occur, you need to teach the car how to react. You need to synthesise that data to do that.

That's a very simple example, but if you look into some of the solutions that we've been talking about, hinting at this evening, that will become a much more, I think, relevant scenario than ability to synthesise data.

Legal and Privacy Advantages

It also offers a path to sidestep some of the issues around personal data. If data is truly synthetic, then it steps outside the scope of UK GDPR, doesn't relate to an identifiable individual, and that also facilitates the use and sharing of data sets with reduced legal friction.

And legal friction around sharing of data sets is one of the other issues that businesses are finding is an issue in practice.

Effectively, with a synthetic data approach, you can use a smaller, anonymised data set to seed and generate a broader population of synthetic records. Medical records, again, might be another example. And those records, if done correctly, should contain subtle but statistical tells of the items that might prompt the AI to find what you're looking for.

Accelerating Innovation

It's not really just therefore about legal convenience, it's also about rapidly speeding up research and innovation. If you can synthesise data sets more quickly, you can sidestep some of the problems with identifying, labelling human world data.
You can avoid some of the issues around, for example, in the medical space, consents and patient information and the like and ethical considerations. It's potentially a way of speeding life up.

The other reality is that our real world data is often messy and prejudiced. And so synthetic data does allow a means for organisations to rebalance the scales by generating data points for underrepresented groups. So to build a data set that is perhaps more rounded than might otherwise be the case.

Risks and Challenges

However, it's fair to say there are problems and there are problems that need thinking about.
The quality of synthetic data is clearly important and it's closely tied to the method used to generate it. And each technique I haven't got time to go into tonight offers strengths and limitations. And selecting the correct technique takes care. You need to both balance data diversity and realism in the data set. And it's important to maintain that balance.

There is also a risk that in practice a lot of synthetic data is produced by other AIs or other statistical models and so care is going to be needed to guard against creating a digital echo chamber where AI is becoming detached from reality and effectively is almost talking to itself.

There are issues, despite the benefits from privacy of anonymity. Is synthetic data truly anonymous? So there is plenty of guidance on this topic indeed from the UK Information Commissioner around some of the risks around it and it's something that needs I think to be better understood by businesses and other organisations.

And then lastly there's the question of copyright and other intellectual property rights. Now that's a hot case in the courts at the present and of course the UK government's own approach is under review and discussion.

But I think if we're interested in the benefits of synthetic data, we really lack legal sensitivity on the IP position in that respect. Whether synthetic data is seeded by copyrighted data, does that constitute IP infringement?

And I think there is a possible view that you might reasonably distinguish between derivation from data sets, what you might call copying, and synthesis, learning. And there may be a rationale for considering a simpler regime to permit generation of synthetic data regardless of the position reached on using data sets to train AI more generally.

Conclusion

So to come to conclusion, synthetic data is not a fake version of the truth. It's a powerful tool to protect privacy and accelerate innovation.
It's not a panacea. It requires care in its use and deployment. and it should be nurtured by government and regulators accordingly.
We need to encourage organisations to create it with a care and ethical approach. And we also need to ensure that data protection and intellectual property law do not unduly block the use and deployment of synthetic as a tool.
It's a tool which can limit the need to make broader and more intrusive use of real world data centres. So it's something that can be potentially enhancing some of the legal problems we face.

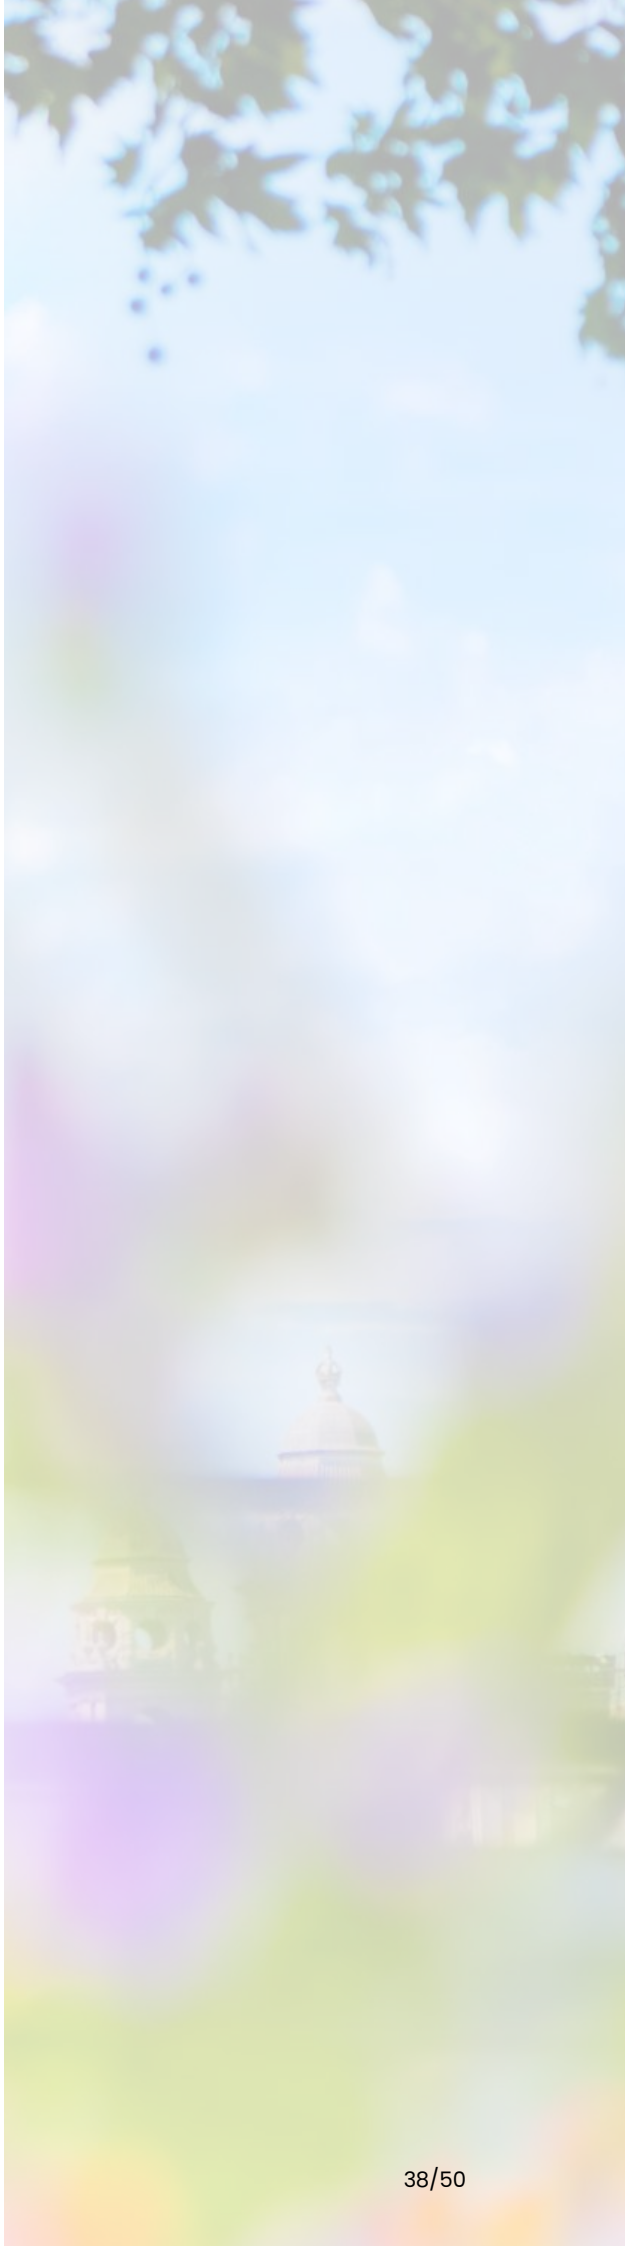

Key Take-aways – Mark Taylor (Osborne Clarke)

Core view: AI progress is increasingly constrained by data access and legal frameworks, making synthetic data a critical—but complex—solution.

Taylor's perspective is legal–practical.

1. The data bottleneck
   - AI depends on:
     - High-quality, human-generated data
   - But access is limited by:
     - Privacy law
     - Availability
     - Cost

2. Synthetic data as a solution
   - Artificially generated data that:
     - Mimics real-world patterns
   - Benefits:
     - Scalability
     - Privacy protection
     - Faster experimentation
   - Essential in cases where real data is unavailable (e.g. rare events)

3. Strategic advantages
   - Enables:
     - Broader data sharing
     - Reduced legal friction
   - Can correct biases by filling gaps in datasets

4. Risks and limits
   - Quality depends on generation method
   - Risk of:
     - "AI training on AI" (echo chambers)
     - Loss of real-world grounding
   - Legal uncertainty:
     - Is synthetic data derived from copyrighted data lawful?

5. Policy implication
   - Need:
     - Clear IP frameworks
     - Balanced regulation
   - Synthetic data should be enabled, but carefully governed

In essence:
Taylor highlights that data (not models) is becoming the real constraint, and resolving this requires legal as much as technical innovation.

BIOs of Evidence Givers

Selected Images from the discussion

Allison Gardner MP
(APPG AI Co-Chair)

Lord Ranger of Northwood
(APPG AI Vice Chair)

Lord Tarassenko
(APPG AI Parliamenarthy Member)

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

Iqbal Mohamed MP
(APPG AI Member)

Lord Taylor of Warwick
(APPG AI Honorary Vice Chair)

Yann LeCun,
Executive Chairman, AMI Labs; Professor, New York University; Founding Director, Meta FAIR; former Chief AI Scientist, Meta.

Yann LeCun is the Executive Chairman of AMI Labs and the Jacob T. Schwartz Professor at NYU affiliated with the Courant Institute of Mathematical Sciences & the Center for Data Science. He was the Chief AI Scientist at Meta, the founding Director of Meta-FAIR and of the NYU Center for Data Science. He received an Engineering Diploma from ESIEE (Paris) and a PhD from Sorbonne Université. After a postdoc in Toronto he joined AT&T Bell Labs in 1988, and AT&T Labs in 1996 as Head of Image Processing Research. He joined NYU as a professor in 2003 and Meta/Facebook in 2013. His interests include AI, machine learning, computer perception, robotics, and computational neuroscience. He is the recipient of the 2018 ACM Turing Award (with Geoffrey Hinton and Yoshua Bengio) for "conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing", a member of the National Academy of Sciences, the National Academy of Engineering, the French Académie des Sciences.

Prof. Viv Kendon,
Professor of Quantum Technology, Department of Physics, University of Strathclyde

Viv Kendon is Professor of Quantum Technology in the Department of Physics at the University of Strathclyde, where she rejoined the Computational Nonlinear & Quantum Optics group in November 2021. Prior to returning to Strathclyde, Professor Kendon held academic appointments at Imperial College London, The University of Leeds, and Durham University. She also serves on the UKRI STFC Computing Advisory Panel, and the NQCC Technical Advisory Group.

Profesor Kendon is an internationally recognised physicist specialising in quantum computing, quantum algorithms, and complex computational systems, with a strong record of interdisciplinary research spanning physics, computer science, and emerging technologies. She brings a rigorous, evidence-based approach to horizon scanning, which positions her well to contribute to discussions on AI foresight, research prioritisation, and the long-term governance of advanced technologies.

Prof. Mihaela van der Schaar
John Humphrey Plummer Professor of Machine Learning, Department of Engineering, University of Cambridge

Mihaela van der Schaar is the John Humphrey Plummer Professor of Machine Learning, Artificial Intelligence and Medicine at the University of Cambridge. In addition to leading the van der Schaar Lab, Mihaela is founder and director of the Cambridge Centre for AI in Medicine (CCAIM).

Mihaela was elected IEEE Fellow in 2009 and Fellow of the Royal Society in 2024. She has received numerous awards, including the Johann Anton Merck Award (2024), the Oon Prize on Preventative Medicine from the University of Cambridge (2018), a National Science Foundation CAREER Award (2004), 3 IBM Faculty Awards, the IBM Exploratory Stream Analytics Innovation Award, the Philips Make a Difference Award and several best paper awards, including the IEEE Darlington Award. She was a Turing Fellow at The Alan Turing Institute in London between 2016 and 2024. In 2025, she was appointed as Spinoza Guest Professor at Amsterdam University Medical Center.

Mihaela is personally credited as inventor on 35 USA patents (the majority of which are listed here), many of which are still frequently cited and adopted in standards. She has made over 45 contributions to international standards for which she received 3 ISO Awards. In 2019, a Nesta report determined that Mihaela was the most-cited female AI researcher in the U.K.

Bob De Caux,
Chief AI Officer, IFS

Bob De Caux, PhD is an experienced technology executive combining deep technical expertise in artificial intelligence with extensive leadership experience in deploying AI at scale. He holds a PhD in Artificial Intelligence and complex systems simulation and has built a career translating advanced research into robust, commercially viable products and customer solutions across highly regulated sectors, including financial services, insurance, and healthcare.

Bob is Chief AI Officer at IFS, a global provider of cloud enterprise software for asset- and service-intensive organisations. In this role, he leads the strategic development and responsible deployment of AI across products and platforms, with a particular focus on governance, assurance, and delivering trustworthy outcomes.

He is a Leadership Fellow at St George's House, contributing to cross-sector dialogue on leadership, ethics, and public trust, and serves on the Data Board of MIT CISR, advising on research into data, AI, and digital transformation of relevance to policymakers.

Mark Taylor,
Partner, Head of Digitalisation (UK),
Osborne Clarke

Mark is a lawyer with extensive experience in technology and data matters. He has expertise advising clients in the financial services and tech, media and comms sectors in these areas, and is Osborne Clarke's International Head of Digitalisation. Mark regularly advises clients on technology and outsourcing contracts of all shapes and sizes, as well as a broad range of data matters. He has long standing experience of advising clients in both the financial services and technology sectors, and at the intersection of the two. As such, he frequently advises on blockchain, digital payments and cybersecurity issues, with FinTech being a particular area of interest.

From a data perspective, Mark advises clients on a wide range of data protection issues, from GDPR matters through to the use and exploitation of large data sets. His work includes advising clients on the creation and use of synthetic data, and the use of data in AI systems. As Head of Digitalisation, Mark leads Osborne Clarke's engagement with clients on Digitalisation themes, including initiatives covering AI and digital regulation.

Mark is a Fellow of the Society for Computers and Law, and is recommended for both Technology and Data matters in the leading legal directories.

ABOUT APPG AI

ABOUT:

All Party Parliamentary Group on
Artificial Intelligence

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
- Iqbal Mohamed MP Independent
- Leigh Ingham MP Labour
- 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
- Baroness Debbonaire (Thangam Elizabeth Rachel Debbonaire) Labour
- The Earl of Erroll (Merlin Hay) Crossbench
- Lord Fairfax of Cameron (Nicholas Fairfax) Conservative
- Lord Freyberg (Valerian Bernard Freyberg) Crossbench
- Lord Hunt of Kings Heath (Labour)
- 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 (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.

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.

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