aippg

appg

All-Party Parliamentary Group on Artificial Intelligence

AI in Science

AI's Transformative Impact on Scientific Research

BIG INNOVATION CENTRE
Secretariat

11 November 2024
Policy Forum

Table of Contents

Title: AI in Science: AI's Transformative Impact on Scientific Research
Publication Details:
Authoring Organisation: All-Party Parliamentary Group on Artificial Intelligence (APPG AI)
Date of Publication: November 2024
Publication Type: Parliamentary Brief | Policy Brief
Publisher: Big Innovation Centre

INTRODUCTION

This document is a transcript with summary of an APPG AI evidence meeting that took place on 11 November 2024 in the House of Lords Committee Room 1, UK Parliament. The transcript exclusively contains crucial discussion elements; not all points are addressed.

DETAILS

- Evidence Session: AI in Science: AI's Transformative Impact on Scientific Research
- Time 5:30 pm – 7:00 pm (GMT)
- Date: Monday 11 November 2024
- Venue: Committee Room 1 in the House of Lords.

CONTACT SECRTARIAT

appg@biginnovationcentre.com
APPG AI Secretariat
Big Innovation Centre

EVIDENCE GIVERS

- **Professor Charlotte Deane**, Executive Chair at the Engineering and Physical Sciences Research Council (EPSRC), UK Research and Innovation (UKRI)
- **Professor Lord Tarassenko**, Professor of Biomedical Engineering at the University of Oxford and Director of TTO, Oxford University Innovation.
- **Conor Griffin**, AI Policy Research Lead, Google DeepMind
- **Richard Traherne**, Next Frontiers, Strategy and Portfolio, Capgemini

MEETING CHAIR AND RAPPORTEUR

The Meeting was chaired by **Lord Clement-Jones;** Co-Chairs of the All-Party Parliamentary Group on Artificial Intelligence.

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

Aim of Session:
AI in Science: AI's Transformative Impact on Scientific Research

This Evidence Meeting explored the transformative role of AI in accelerating scientific discoveries and breakthroughs, along with the limitations and risks of over-relying on AI for scientific research. Key topics also included creating policies to guide the ethical use of AI in experimentation, as well as ensuring transparency and accountability in AI-driven scientific processes.

Questions were raised to inspire the discussion:

- How can AI transform and accelerate scientific research, discoveries and breakthroughs in fields like biology, chemistry, and medicine? What are the examples demonstrate this potential?
- What are the limitations or potential pitfalls of relying on AI for scientific research? How can we mitigate these pitfalls?
- What policies or guidelines might be beneficial for the ethical application of AI in the sciences, particularly with regard to experimentation?
- How do we ensure transparency and accountability in AI-driven scientific processes?

FINDINGS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

The APPG AI Evidence Session found that a multi-dimensional and cooperative approach is essential to ensure that AI serves as a powerful tool for scientific advancement while also supporting societal well-being.

Recommendations for Policymakers, Regulators, Businesses, and Society:

1. Policy and Regulation

- **Develop Adaptive Policy Frameworks**: Create and regularly update policy frameworks that keep pace with rapid advancements in AI technology, including clear guidelines on data reuse, privacy protection, and ethical standards for AI applications, particularly in scientific research.
- **Encourage Rapid Policy Development**: Facilitate close collaboration between regulators and scientists to quickly integrate scientific insights into effective and flexible policymaking for emerging AI technologies.
- **Implement Proactive Regulation or Protocols:** Develop regulations or protocols that address current risks associated with AI and anticipate future challenges, including mechanisms for monitoring AI systems and addressing dual-use potential.
- **Establish Evaluation Frameworks**: Create frameworks to assess the effectiveness, safety, and societal impact of AI systems, collaborating with third parties to conduct independent assessments for objectivity.
- **Develop Whistleblower Protections:** Establish safe channels for reporting concerns regarding AI practices, ensuring protection for individuals who raise red flags about potential risks.

2. Data Governance and Transparency

- **Prioritize Data Governance**: Establish robust data governance structures that balance the need for scientific advancement with individual privacy rights, encouraging transparency in data utilization for AI training to build public trust.
- **Embrace Transparency and Accountability**: Adopt transparent practices in AI development and implementation to ensure accountability for the outcomes resulting from AI technologies.
- **Establish a National Data Service:** Create a centralized national data service to unite and provide access to various health and scientific data sources. This service should be designed to enhance research capabilities and improve public health outcomes while ensuring appropriate governance.

3. Interdisciplinary Collaboration and Innovation

- **Support Interdisciplinary Collaboration:** Foster collaboration across various fields, particularly between the social sciences and natural sciences, to leverage diverse expertise and promote comprehensive solutions to the complex challenges posed by AI technologies.
- **Invest in Research Environments and Digital Resources:** Advocate for substantial infrastructure investments that support AI and scientific initiatives, including funding for research environments and digital resources that facilitate effective data sharing and collaboration.
- **Explore Funding Programmes and Incentives:** Actively explore models of grants, subsidies, and tax incentives to support innovative AI projects, especially in high-risk, high-reward areas of research.
- **Invest in T-shaped (Problem-solvers) Talent Development:** Build multidisciplinary teams that include experts and T-shaped talents integrating various fields such as biology, engineering, AI, and social sciences to foster innovation and holistic approaches to problem-solving.

4. Education and Community Engagement

- **Enhance AI Literacy for AI Engagement:** Promote AI education and training programs to improve AI literacy among policymakers, researchers, and the public, equipping individuals with the necessary skills to engage with AI tools and AI systems and make informed decisions about when and how to use AI technologies.
- **Promote Public Engagement and Open Dialogue:** Actively engage with the public to raise awareness about AI's potential and implications, countering misinformation by providing clear information on AI technologies and their applications.

5. Social and Environmental Responsibility

- **Promote Environmental and Social Responsibility:** Advocate for AI applications that contribute positively to society, addressing critical issues such as climate change, healthcare, and sustainability.
- **Recognize data as a strategic national asset that can drive progress:** Engage in discussions about the importance of data sharing and collaboration in scientific research, recognizing data as a strategic national asset that can drive progress (relates to the recommendation "Establish a National Data Service")

Evidence Giver:
Professor Lord Tarassenko

Evidence Giver:
Conor Griffin

Evidence Giver:
Professor Charlotte Deane

Evidence Giver:
Richard Traherne

APPG AI Chair:
Allison Gardner MP

APPG AI Chair: Lord
Clement-Jones CBE

Secretariat & Rapporteur:
Professor Birgitte Andersen

EVIDENCE

Professor Charlotte Deane

Executive Chair of the Engineering and Physical Sciences Research Council (EPSRC), UK Research and Innovation (UKRI).

Introduction to My Roles and Responsibilities

I am the Executive Chair of the Engineering and Physical Sciences Research Council, which is part of UK Research and Innovation(UKRI). This basically is an arms-length government body that puts public funding into universities, into research and innovation, but also into small and medium enterprises, or SMEs.

There are two reasons that I'm here related to that role. One is that I Head the Engineering and Physical Sciences Research Council, so obviously, AI sits underneath that. But I am also UKRI's AI champion, overseeing the portfolio across all of UKRI. The second part of my role is that I am also an academic at the University of Oxford, where I lead a group of about 40 people developing AI methods for drug discovery for large and small molecules. Throughout my time in this position, I have worked a lot with industry, both in and with industry.

AI is Already Transforming: Nobel Prizes Recognition

The first point I want to make is that AI is already transforming how we do science and innovation. I hope most people have noticed this. For instance, a couple of Nobel Prizes were awarded recently for foundational discoveries and innovations, with one being given for an entirely computational technique for protein design and structure prediction.

Changing Research Approaches

One analogy I like to use to describe this transformation is to think about the world before we had computers and to ask whether computers would transform how we conduct research. The answer to that question is clearly yes. It's not just that computers speed up our processes; they change the types of questions we can ask and how we approach them. AI doesn't merely enable better predictions; it fundamentally changes how we conduct research.

AI's Data Processing Capabilities

The most obvious way AI is transforming research is through its remarkable ability to absorb and synthesize vast amounts of data. These algorithms can digest data in ways humans cannot. I can't read every research paper in my field, but I can get algorithms to summarize the relevant research effectively. This capability represents a significant advancement for researchers, yet many are still not utilizing it to its full potential.

Making Connections and Recognizing Patterns

AI can also make connections within that data, linking potential drugs to all existing targets. This is more than just making predictions—it's about how we structure and hold data. Next, AI's ability to process that data allows it to recognize patterns. For instance, in medical imaging for disease diagnosis, AI can identify important and repeatable patterns without even moving into predictive territory yet.

Multi-modal Data Integration

Another exciting aspect is that these models can process multi-modal, unstructured data, including text, images, and numbers. This capability enables us to connect datasets that were previously considered completely separate, such as health and social data, without needing to format each dataset meticulously for traditional statistical models. AI can handle the unstructured nuances between these data types and streamline the process.

Impact on Programming and Research

Importantly, one significant impact of AI in the engineering and physical sciences is how well these models can run code. This is noteworthy because they can communicate with other computer programs. For instance, in a simple scenario, I can ask an AI to transform numbers on a page into a graph. On a more complex level, I can instruct it to integrate methodologies for drug discovery and execute that code without writing extensive lines of code myself; I can simply use natural language. This is a game changer for scientific research, especially as many scientists are not code literate.

Upskilling for Scientific Progress

However, I want to emphasize that AI is not magic. This point is vital. The major pitfall lies in user mistakes. We must recognize that if we want to reap the benefits of AI and see scientific progress, we need to upskill many individuals. This situation is reminiscent of the time when statistics became prevalent; while they offered tools for improvement, they were often misused.

Building an AI-Literate Community

Today, we have two communities—naive AI users who are excited but may not understand what they're utilizing, and naive AI experts who might run their code incorrectly. The final takeaway is that we need to cultivate an AI-literate community and think continuously about co-creation in science. Right now, there aren’t many experts who excel in both AI and their scientific domains, so we must enhance our training capabilities. Continuous professional development is necessary for skilling everyone, alongside training more AI experts to integrate into multidisciplinary teams.

Final remarks

AI is a remarkable tool and an exciting resource poised to answer questions we haven't even begun to ask. However, success will only come if we have the right people equipped to leverage it effectively across all scientific fields. Thank you.

Summary: Professor Charlotte Deane ‘s Evidence Statement

- **Charlotte Deane:** Serves as the Executive Chair of the Engineering and Physical Sciences Research Council and is also the AI Champion for UK Research and Innovation (UKRI). She is an academic at the University of Oxford, leading a team focused on AI methods for drug discovery.
- **AI's Transformative Impact:** AI is significantly changing how science and innovation are conducted, influencing the types of questions researchers can ask and the methodologies they use.
- **Nobel Prizes Recognition:** Recent Nobel Prizes in physics and chemistry highlight foundational discoveries and computational techniques in research, indicating AI's growing importance.
- **Data Synthesis:** AI excels at absorbing and synthesizing large datasets, providing summaries and insights that researchers may struggle to gather manually.
- **Making Connections:** AI can identify connections between data points, such as linking potential drugs to multiple targets in drug discovery.
- **Pattern Recognition:** Advanced AI methods can recognize patterns in data, which is particularly useful in areas like medical imaging for diagnosis.
- **Multi-Modal Data Processing:** AI can effectively handle diverse data types (text, images, numbers), enabling connections between previously disjointed datasets like health and social data.
- **Code Execution:** AI's capability to run code and execute complex tasks using natural language is revolutionizing research methodologies, allowing non-coders to execute programming tasks.
- **Need for Upskilling:** There's a critical need for the scientific community to become more AI literate to avoid misusing AI technologies and to facilitate effective collaboration.
- **Co-Creation in Science:** The expansion of expertise in AI and its integration into scientific teams is essential to fully leverage its potential across different disciplines.
- **Exciting Future:** AI offers vast possibilities for enhancing research capabilities, but success will depend on the cultivation of a skilled, knowledgeable workforce suitable for its integration into scientific practice.

Professor Lord Tarassenko

Professor of Biomedical Engineering at Oxford University and Director of the University's wholly-owned technology transfer company, Oxford University Innovation.
Also, Member of the House of Lords

Historical Context of Machine Learning

Machine learning (ML) runs AI, but let's carry on talking about AI. I was doing machine learning before it was called machine learning because, of course, it wasn't always called machine learning. I started doing it around when Geoffrey Hinton was one of the three authors of this key paper in 1986:

With David Rumelhart and Ronald J. Williams, Hinton co-authored a highly cited paper published in 1986 titled "Learning Representations by Back-Propagating Errors" (Nature, 323(6088), 533–536) that popularized the backpropagation algorithm for training multi-layer neural networks.

Advancements in AI Generations

My career has followed the evolution of different generations of AI. The first generation was ML followed by the second generation, which is deep learning. We have now conducted sufficient research into deep learning and its application to real-world problems to confirm that a machine learning algorithm can perform at a level comparable to that of a radiologist. The key consideration is how we leverage this capability to replicate and enhance the abilities of a radiologist.

The Promise of Generative AI

Finally, for the last three or four years, generative AI, and generative AI is brilliant for scientific research. In my own field, as you can see from my biography, I work very closely with Royal Society for 15 years on building ML applications for monitoring engines. But I don't have time to go into that. I'm going to talk just about healthcare, which is my main field. For the first time, we can deal with multimodal data (i.e. data that comes from different sources and formats that is then combined together), which is really a huge benefit to healthcare because you've got doctors' notes, discharge summaries, images, vital signs, blood counts. What have I missed? Genomic data, eventually wearable data that people collect themselves, etc. And we have a tool for combining all of that in a relatively straightforward way.

Advancements in Healthcare Applications

In addition to that, which is also very important for healthcare applications, we can model the time-varying aspects through the attention mechanism. Everything is possible with generative AI, and when I first looked at it, I was sceptic. I'm now a believer because we've actually proved it on a number of sample problems within my group, and that is transformative. So, the big question (especially with my friend from Google DeepMind on the panel), is when you come to use machine learning to solve a new healthcare problem – do you really have a choice? Do you take one of the existing tools, brilliant tools like Gemini, for example, picking one at random, and do you fine-tune it for your healthcare application? That's what we do. Who indeed might have done it for us? They've got Med-Gemini. I have a research group working on using these techniques.

The Debate on Custom vs. Existing Tools

Now somebody can take the available code on GitHub or other open sources that are available and program a brand-new model from scratch, and it's good in many ways because then you understand it better. You can minimize the chance of hallucination on your very specific dataset.

So, we have this debate: do we take the big tools available from big tech and fine-tune them, or do we redesign on our own dataset with 200 lines of code using Python, for example, and control the data, control the code, and so on? It's an open question. I'm not going to say which way you should go; it's an open question, but the beauty of it in this country is we have students and researchers in academia and industry able to do either.

Real-World Application Example

Just to give you an example of what we're working on—and it is kind of in the public domain—Drug GPT. So, a 30-year-old woman presents with swelling in her hands. She's got a positive rheumatoid factor and a negative anti-neutrophil antibody test. That's going to be quite tough for the GP. It's probably the edge of their knowledge. With the right Drug GPT model, plus the patient history trained on all the papers on rheumatoid arthritis and autoimmune diseases and so on, you can come up not just a diagnosis but also drug recommendations, dosage recommendations, any likelihood of adverse reactions, and any drug-drug interaction with any drugs that patients are thinking about. Can you imagine how this would change the way a doctor does consultation? Not for the 90% of the cases, which are straightforward, but for the 10% that are really difficult. That is available today, just about; it's not just for the future. It's today, and it will lead to personalized treatment plans.

The Paradox of Personalization

Now, this is true of all scientific research. There's a paradox here: to get a personalized recommendation, you need big data. You need to train over the whole input space, the variations of all these multimodal parameters. By the time you apply it to you or me, or to my patient, then it can be an interpolation within this human space of input data, not extrapolation.

So, it is a paradox. I've tried to explain it in 20 seconds; I may have failed. But to do personalized treatment, you need big data to train your models, and once you go to personalized treatment, you start to think about prevention and so on. So, all of that is beginning to be possible.

Importance of Accessible Data

The main issue I was asked for Parliamentarians to consider is actually not the algorithms but the availability of the data. There is a really important report that came out on Friday—the "Sudlow Review: uniting the UK's health data". It discusses uniting the UK's health data, which presents a huge opportunity for society as it talks about linking multiple data sources and accessing them to give a step change in prevention, diagnosis, and treatment of multiple diseases, including cancer. Professor Cathie Sudlow is proposing a National Health Data Service built around data access through the secure data environments that we're currently designing within the NHS, and a UK-wide approach for appropriate data governance.

Privacy and Regulatory Considerations

There are many things we could think about, including how people opt in or opt out and the privacy issues. She addresses that in her report. It's about two decades since I've used these types of data for research; thousands of researchers up and down the country have been using this data on millions of records. There's not been one single reidentification after anonymizing the data. So let's not just focus on the negatives of privacy, which always comes out first. It has never happened in this country, and we could build ways of opting out, which I'm happy to talk about in the discussion.

Regulatory Frameworks for AI

It also opens the door to regulation. How do we regulate? Actually, there's a lot of regulation already in AI for healthcare. Both the FDA (Food and Drug Administration) in the US and MHRA (Medicines and Healthcare products Regulatory Agency) in this country have taken three products through, two through the FDA and 13 through the MHRA. They are all three machine learning products. So for parliamentarians, I think it's important to consider whether we need provisions in the current data bill that's going through Parliament to ensure that we can reach the benefits of uniting the UK's health data.

Summary: Professor Lord Tarasenko’s Evidence Statement

- **Machine Learning Roots:** Lord Tarasenko mentions his experience with machine learning (ML) before it was formally termed as such, tracing back to influential work from Geoffrey Hinton and team in 1986.
- **Generations of AI:** His career has followed the evolution of AI, from the first generation of basic ML to the current advancements in deep learning and generative AI, particularly in healthcare.
- **Transformative Healthcare Potential:** Generative AI enables the handling of multimodal data in healthcare, allowing integration of various data types (doctor's notes, images, vital signs, genomic data, etc.) into a cohesive analysis tool.
- **Capability of ML Tools:** He discusses the choice in healthcare applications—whether to use existing advanced tools (like Gemini) for specific applications or to develop new models tailored to specific datasets, emphasizing both options are valid.
- **Example of AI in Practice:** He shares a scenario involving a 30-year-old woman with rheumatoid arthritis where utilizing a drug AI-recommendation model could significantly enhance diagnostic and treatment processes for difficult cases.
- **Personalized Treatment Paradigm:** The paradox of personalized treatment requiring extensive data training is highlighted. Adequate data can lead to better prevention and treatment strategies.
- **Data Availability Challenge:** The key challenge for AI and ML in healthcare lies in the accessibility of data rather than the algorithms themselves.
- **Sudlow Report Emphasis:** Lord Tarasenko refers to a recent report proposing a National Health Data Service to unite and access various health data sources, which can improve prevention, diagnosis, and treatment.
- **Data Privacy Assurance:** Lord Tarasenko stresses the importance of privacy and the long history of using anonymized health data in the UK without any reidentification incidents.
- **Regulatory Considerations:** He points out existing regulations for AI in healthcare and raises questions about provisions in the current data bill that can facilitate the benefits of united UK health data while ensuring proper governance and privacy measures.

Richard Traherne
Next Frontiers, Strategy and Portfolio, Capgemini.

Introduction to Capgemini and Future Focus

I'm Richard Traherne from Capgemini, where I sit within the global leadership team driving the next frontiers. I'll unpack that. Essentially, we're looking at the science and technology areas that are going to drive our clients' businesses in the next 5 to 10 years. So, looking ahead and making that a commercial reality in a sense.

Commitment to Economic Growth Through Innovation

I'm a passionate believer in economic growth through innovation, not just incremental productivity, but fundamental new additive breakthroughs. I am based in Cambridge, which is where we have our labs—our science labs, engineering labs, and others.

The Potential of Engineering Biology

I also sit within the World Economic Forum Steering Group for the Bioeconomy Initiative and lead it within Capgemini. It's a major new industrial revolution. I describe it that way because it's not incremental; it's something that's going to fundamentally change the way we operate in industries and in manufacturing.

Recognizing Industry Disruption

Thought leaders like Professor Dame Angela McLean and Craig Venter, and many others, have described this century as the biological century, where we're going to make advances in a way that we haven't before in this area. Earlier this year, we interviewed 1,100 corporates and 500 startups, and half of those respondents saw engineering biology disrupting their industry within five years. The other half sees it happening within the next ten years.

AI as a Tool for Addressing Sustainability

So, it's a major and imminent opportunity for the UK. For those that are less familiar, it uses the power of nature to grow valuable new products. That can be chemicals, ingredients, pharmaceuticals, fuels, fertilizers, and also provide entirely new methods to accomplish tasks such as capturing carbon by remediating and degrading. It uses both physical and digital AI techniques to advance the biological world.

The Role of AI in Innovation

A simple analogy is that software developers string ones and zeros together to create software code; engineering biologists string together DNA sequences to create new forms of organisms. It fundamentally provides new ways to address many of our sustainability imperatives around the world—not just incremental changes, but entirely new methodologies.

Transforming Research Processes

Currently, we're losing the race, so this is good news. It presents a great opportunity, but the design space is massive. There are more combinations of DNA in the world than there are atoms. This is why AI is an incredibly valuable technique for breaking this down. Historically, we've used a trial and error approach: lots of people in white coats with 15 labs would make something, test it, refine it, and repeat. It was slow.

Accelerating Scientific Development

Now we're using engineering and digital methods, especially AI, to cut costs and time in our labs. We've been taking year-long processes and reducing them down to months or weeks. Most recently, we developed an enzyme called Cutinase, which breaks down PET (Polyethylene Terephthalate) plastic. For years, that has been developed traditionally, but we created the highest-performing Cutinase in the world using AI-based techniques in a matter of weeks. AI is fundamentally making this world commercially viable when it previously hasn't been.

Key Areas for Support and Development

I can't underline our commitment to this enough. I would argue that of all the application areas of AI, engineering biology is one that the UK should be focused on, not only because of its power and range but also because it's nascent. We have an opportunity to build a position with the skills that we have to do something very special.

Areas of Action for Advancement

Now, if I turn to what's the ask for this group. I've been greedy and have five points. I'll start with some softer topics. The first is bio-literacy. As Carl Sagan, the US scientist, said, we live in a society exquisitely dependent on science and technology, in which hardly anyone knows anything about science and technology. For any new revolution, it's of course a concern that we have to overcome.

(1) Importance of Multidisciplinary Skills

Over 60% of the respondents to the survey I mentioned earlier said that bio-literacy was a key obstacle in creating demand for what they want and overcoming the concerns that otherwise exist. So, I would say that the government needs to be proactive in taking the nation on that journey. It's about opening people's eyes to the potential and opportunity and counteracting misinformation, which otherwise artificially stifles progress.

(2) Investment Needs in Engineering Biology

The second area would be multidisciplinary talents. To make this work, you need biology, but you also need physical engineering capabilities, digital and AI skills, and physical science expertise, among others. This is crucial because historically, technologies have been siloed, and gaining expertise in those areas was possible. Today, the challenge is to find individuals who are T-shaped, those who can sit across these domains and innovate collaboratively. We need to reinvent our university systems and for this multi-domain technology world.

(3) Economic Support for Innovation and the Strategic Selection of Industries

The third point would be subsidies, as engineering biology is really expensive. This new area requires investment in talent and infrastructure, so we need to bootstrap this as an industry and create a level playing field. We all know the 10X rule: that a new disruption typically needs to be at least ten times better and ten times cheaper for people to switch to that new approach. The same principle applies here, too. We need grants, subsidies, piloting schemes, tax incentives, and disincentives that can help incentivize growth in this sector.

We can look to the US for examples and draw our own conclusions about the current landscape. Controversially, do we need to be bolder in picking our winners? This is a theme seen elsewhere in the world, where governments are more willing to identify winners and put resources toward stimulating progress as a nation.

(4) Regulatory Frameworks for New Innovations

The fourth point is about creating entrepreneurially friendly regulation and assurance. New innovations carry risk; they must be understood and regulated effectively. There is a lot of dual-use potential in areas like this. In our survey, 96% of respondents indicated they were already working on biosolutions in some form, making it clear that this is a current issue, not something that is on the horizon. We need to find ways to support that. It involves balancing regulation with maturity, creating sandboxes and controlled environments to stimulate early-phase research. Establishing buyer assurance frameworks and standards can help mitigate risks. At this point, I would highlight that there are several organizations, including UKRI, the Regulatory Horizons Council, and BSI (The British Standards Institution), among others, already tackling these issues from various perspectives, and they are doing a great job.

(5) Need for Rapid Policy Development

The final point is the need for rapid and effective policy developments. This requires close collaboration between regulators and scientists to integrate deep scientific insights into understanding and managing risks quickly. We work with our clients in policy labs, using a typical structure that is very expert-driven with user-centred design approaches in policy-making.

Conclusion and Invitation for Collaboration

So, that's my final point. In conclusion, I mentioned a report several times. It contains a wealth of information and practical examples:

Report from the Capgemini Research Institute referred to is this presentation: *"Unlocking the potential of engineering biology: The time is now"*. This report looks at how engineering biology can be applied in different sectors and how it has the power to transform industries. It also examines biosolutions from a sustainability perspective.)

Summary: Richard Traherne's Evidence Statement

- **Introduction and Role:** Richard Traherne from Capgemini discusses his position within the global leadership team, focusing on science and technology areas set to drive client businesses in the next 5 to 10 years.
- **Belief in Innovation:** He emphasizes the importance of economic growth through innovation, advocating for substantial breakthroughs rather than just incremental improvements in productivity.
- **Engineering Biology:** Richard Traherne highlights engineering biology as a transformative field within what is described as the "biological century," identifying it as a major industrial revolution that will change operations in various industries.
- **Industry Disruption Potential:** His organization's survey indicates that half of corporate and startup respondents expect engineering biology to disrupt their industries within five years, showing significant imminent opportunities for the UK.
- **Harnessing Nature:** Engineering biology harnesses natural processes to create valuable products like chemicals, pharmaceuticals, and fertilizers, as well as new methodologies for sustainability challenges such as carbon capture.
- **Role of AI:** AI is highlighted as a crucial tool in expediting the engineering biology process, moving away from traditional trial-and-error methods to faster, more efficient configurations using digital techniques.
- **Commitment to the Field:** Trahan stresses Capgemini's dedication to making engineering biology commercially viable and urges the UK to focus on this area for its potential to foster innovation.

Five Key Requests:

1. **Bio-literacy:** Increase public understanding of science and technology to overcome obstacles and stimulate progress.
2. **Multidisciplinary Talent:** Develop professionals who possess expertise across biology, engineering, AI, and physical sciences to encourage innovation.
3. **Subsidies and Funding:** Advocate for grants, subsidies, and tax incentives to support the expensive nature of engineering biology and encourage industry growth.
4. **Entrepreneurially Friendly Regulation:** Create a regulatory framework that balances risk management with the need for innovation, including sandboxes for early-phase research.
5. **Rapid Policy Development:** Foster collaboration between regulators and scientists to quickly translate scientific insights into effective policies.

- **Encouragement of Engagement:** Trahan invites attendees to discuss the report he referenced, which contains valuable insights and examples of active initiatives in engineering biology.

Conor Griffin

AI Policy Research Lead, Google DeepMind

Acknowledgment of Nobel Prize Contributions

Yes, my boss (Demis Hassabis) is a Nobel Prize winner. We all claim to be part of it, but we're very chuffed for our colleagues who won the Nobel recently and for David Baker, who's an excellent scientist who also won for using AI to design proteins.

Role in AI Policy Research

I work on AI Policy Research at Google DeepMind. In practical terms, this means I work with two main teams. One is our internal responsible innovation team, which primarily works with the teams developing AI models to think through the risks and benefits to society from those models and how we should respond internally. The other team is the public policy team that engages with governments about the opportunities and risks from our work, and what they and we should do in response.

Focus on Foundational AI Technologies

For those who don't know, Google DeepMind is an AI research lab within Google that focuses on foundational AI technologies, but we also have a dedicated science team that researches various scientific problems. This ranges from using AI to predict protein structures—what we won the Nobel Prize for—to designing proteins, as well as other areas like weather forecasting, mathematics, and computer science where we are seeing interesting breakthroughs.

Complex Scientific Endeavours

One challenge I personally face is summarizing what our science team is doing into an overarching narrative on scientific advancement that we can communicate to policymakers. It's a complex problem because we've been present at various 'AI for science' events. It tends to be challenging to pinpoint specific AI for science opportunities; often, discussions lead to extensive lists of problems that AI could address.

Identifying the Top Opportunities

Another challenge arises when scientists rarely agree unanimously on anything, so attempting to get them to prioritize the top five opportunities is a daunting task. Because we are a company that receives pushback rightly so, we ask not just which areas could use AI, but in which areas must AI be utilized given the broader challenges in science. Therefore, we've compiled a few key opportunities where we see real potential for using AI across different disciplines.

Common Challenges in Scientific Research

These opportunities differ, but they all relate to the common theme of scale and complexity. Scientists today face challenges tied to the growing size of the literature base, which becomes increasingly specialized. This growth creates difficulties for researchers trying to make breakthroughs or deploy the breakthroughs as useful products, as they struggle to keep up with the ever-expanding body of knowledge and increasingly complex experiments.

AI as a Solution to Scale and Complexity

We see AI as an opportunity to tackle these challenges. The simplest way to think about it is that AI can aid in overcoming issues of scale and complexity, particularly through deep learning. So, concretely, the first opportunity concerns scientific knowledge. The scientific literature continues to grow larger and more specialized, making it increasingly challenging for scientists—and indeed anyone outside of science such as policymakers and business professionals—to stay informed and digest relevant findings.

Improving Literature Reviews

There are already examples of scientists using AI to improve literature reviews and policymakers employing these technologies to help with evidence synthesis in their respective fields. Moreover, there is potential for scientists to explore new methods of sharing their findings. While traditional academic papers have their merits, adapting content to different audiences and levels could enhance comprehension. For example, using AI to tailor explanations based on age or expertise can help make science more accessible.

Transforming Experimental Design

The second significant opportunity pertains to experimental design. Scientists today across disciplines often struggle to conduct the experiments they wish to due to cost, time constraints, or access to necessary facilities or resources. One area I want to touch on is fusion energy. If we could achieve it, fusion could provide an essentially limitless source of emission-free energy, enabling many downstream innovations, including desalination for fresh water.

AI in Fusion Research

However, practical challenges exist in accessing reactors for experiments, which are expensive and time-consuming to construct. A few years ago, we collaborated with colleagues at the EPFL Swiss Plasma Centre to utilize AI in managing plasma shape within a simulated reactor, demonstrating how AI can inform experimental design rather than simply replacing the experimentation itself.

Exploring the Search Space in Scientific Problem-Solving

One final opportunity I want to mention relates to the search space in scientific problem-solving. We often see multiple potential solutions for any given problem. For instance, in our work designing new proteins capable of degrading plastics in the environment, the vast number of potential amino acid arrangements presents a significant challenge. Historically, scientists relied on intuition, trial and error, iteration, or brute force computations to navigate this space.

Novel Ideas Through AI

AI can help explore new parts of this search space through large language models, generating novel ideas that hadn't previously been considered. However, newness alone isn't sufficient; it is crucial to ascertain the viability and utility of these ideas. This is where collaboration with real-world experimentalists becomes essential. We've established a wet lab with our partners at the Francis Crick Institute, our neighbours in Kings Cross London, enabling us to combine novelty with real-world applicability.
Success at the Mathematical Olympiad

Additionally, I'd like to highlight an exciting aspect of our work in mathematics. At this year's International Mathematical Olympiad—an incredible competition showcasing young talent—our AI system successfully answered four of the six questions by utilizing our Gemini large language model alongside another AI grounded in formal mathematical logic. This collaboration allowed us to validate suggestions from Gemini systematically and arrive at the most likely correct answers. This illustrates how stitching together different AI systems can help address the limitations of individual models.

Addressing Risks in AI Deployment

Now, I want to quickly touch on the topic of risks. While my focus has been mainly on opportunities, I am happy to discuss risks if questions arise. We think about risks in two primary ways: risks to the practice of science and broader societal concerns. For example, issues such as the reliability of scientific outputs, including the potential for hallucinations from large language models, are significant considerations. There are also environmental risks associated with the implementation of AI in scientific fields.

Mitigation Strategies for AI Risks

What I want to emphasize about risks is that there are many exciting mitigation strategies being developed. For instance, concerning hallucinations, we have been working on what we term "AI factuality," which involves creating new methods to ground the outputs of large language models. However, we confront challenges in determining which sources to rely on, which is not a straightforward issue.

Positive Contributions of AI

Moreover, I want to highlight that while certain areas pose risks, they also present genuine opportunities for AI to contribute positively. Climate change is a prime example, where AI can play a critical role in developing solutions. Additionally, AI can enhance scientific creativity, which may seem inherently human, but we can leverage AI to stimulate and support creative processes in research, potentially leading to more reliable scientific outputs.

Policy Responses to AI Challenges

Lastly, on the topic of policy responses: One approach we advocate for policymakers is to work backward and identify what conditions must be in place for successful AI-for-science initiatives to occur. Many of the enabling factors stem from public investments, whether in data sets, academic research, or developing domain expertise necessary for addressing the right scientific problems.

Evaluation of AI Models

Another important role for scientists involves evaluating the scientific capabilities of AI models. For instance, while large models could assist with experimental design, it's crucial to demonstrate and prove their effectiveness. This poses a challenge given that scientists can be sceptical, and convincing them requires strong evidence.

Collaboration for Advancement

In conclusion, there is a growing field of work focusing on how to evaluate and compare different AI approaches, particularly in protein design and related areas. Ideally, this would involve collaboration among domain specialists from academia and AI labs. We see this as a significant opportunity for advancement, and we hope policymakers can facilitate and support these collaborative efforts.

Thank you for your attention.

Summary: Connor Gryphon's Evidence Statement

- **Introduction and Role:** Connor Gryphon works on AI Policy Research at Google DeepMind, focusing on responsible innovation and public policy regarding the opportunities and risks associated with AI in scientific research.
- **AI in Science:** DeepMind conducts foundational AI research and has a dedicated science team that utilizes AI to address a variety of scientific challenges, including protein structure prediction and fusion energy.
- **Challenges of Communication:** Connor Gryphon faces the challenge of translating complex scientific activities into understandable narratives for policymakers, which involves identifying key AI for science opportunities.

Opportunities for AI in Science:

- **Enhancing Scientific Knowledge:** AI can assist scientists and policymakers by improving literature reviews and synthesizing evidence, making it easier to understand an ever-expanding body of specialized knowledge.
- **Streamlining Experimental Design:** AI can help scientists design and manage experiments more effectively, as seen in fusion energy research where AI aids in simulating plasma conditions without needing physical reactors.
- **Navigating Complex Search Spaces:** AI can assist in exploring numerous potential solutions to scientific problems, particularly in protein design, by using large language models to develop novel ideas while ensuring their viability through experimental collaboration.
- **Collaboration with Experimentalists:** DeepMind has partnered with the Francis Crick Institute to foster real-world applicability of AI-generated ideas, demonstrating the importance of collaboration between AI researchers and experimental scientists.
- **AI in Mathematics**: Connor Gryphon shares a success story from the International Mathematical Olympiad, where their AI system effectively answered complex problems, illustrating the potential of combining various AI models to enhance creativity and problem-solving.
- **Addressing Risks:** He acknowledges the importance of discussing risks associated with AI in science, including issues of reliability and environmental impact, but emphasizes the existence of promising mitigation strategies, such as AI factuality and responsible sourcing.
- **Policy Recommendations:** Connor Gryphon encourages policymakers to facilitate public investments and conditions necessary for effective AI in science initiatives, while also advocating for enhanced evaluation methods to assess the capabilities of AI models in scientific contexts.
- **Call for Collaboration:** He highlights the need for ongoing collaboration between scientists, AI researchers, and policymakers to ensure effective integration of AI in scientific research and to address risks and opportunities in a balanced manner.

BIOs of Evidence Givers

Professor Charlotte Deane MBE

Executive Chair at the
Engineering and Physical Sciences Research Council (EPSRC),
UK Research and Innovation (UKRI).

Charlotte is a Professor in the Department of Statistics at the University of Oxford, Charlotte is also Executive Chair of the Engineering and Physical Sciences Research Council (EPSRC).

From 2022 to 2023, Charlotte was Chief AI Officer at Exscientia, a biotech with ~450 employees, where she led its computational scientific development. She served on SAGE, the UK Government's Scientific Advisory Group for Emergencies, during the COVID-19 pandemic, and acted as UK Research and Innovation's COVID-19 Response Director.

At Oxford, Charlotte leads the Oxford Protein Informatics Group (OPIG), who work on diverse problems across immunoinformatics, protein structure and small molecule drug discovery; using statistics, AI and computation to generate biological and medical insight.

Her work focuses on the development of novel algorithms, tools and databases that are openly available to the community. These tools are widely used web resources and are also part of several Pharma drug discovery pipelines. Charlotte is on several advisory boards and has consulted extensively with industry. She has set up a consulting arm within her own research group as a way of promoting industrial interaction and use of the group's software tools.

Professor Lord Tarassenko CBE FREng FMedSci

Professor of Biomedical Engineering at Oxford University and Director of the University's wholly-owned technology transfer company, Oxford University Innovation.
Also, Member of the House of Lords.

Professor of Biomedical Engineering at the University of Oxford and Director of Oxford University Innovation, Professor Lord Lionel Tarassenko is a prominent figure in applying signal processing and machine learning to healthcare. He is the founding president of Reuben College and leads the AI & Machine Learning research cluster.

He pioneered the first FDA-approved machine learning system for patient monitoring in critical care, significantly enhancing patient safety. Additionally, his work on jet engine monitoring software for Rolls-Royce earned him the Chairman's Award for Technical Innovation in 2001 and the Sir Henry Royce High Value Patent Award in 2008.

Born in Paris in 1957, Professor Tarassenko completed a BA in Engineering Science (1978) and a DPhil in Medical Electronics (1985) at Oxford. He joined the university in 1988 after a stint at Racal Electronics and was appointed Professor of Electrical Engineering in 1997. He played a key role in establishing the Institute of Biomedical Engineering (IBME), which grew significantly under his leadership and received a Queen’s Anniversary Prize in 2015.

A fellow of the Royal Academy of Engineering (2000) and the Academy of Medical Sciences (2013), he has received awards such as the British Computer Society Medal (1996) for his work on sleep disorders, the E-health Innovation Award (2005), and the Silver Medal of the Royal Academy of Engineering (2006). He has published 320 journal articles, 280 conference papers, and holds 32 patents.

In May 2024, he was appointed as a non-party-political peer in the House of Lords, continuing his influence in the fields of engineering, technology, and healthcare.

Richard Traherne

Next Frontiers, Strategy and Portfolio, Capgemini.

In his role at Capgemini Invent, Richard Traherne leads the Next Frontiers initiative, which spans emerging technologies, including quantum computing, AI, and Engineering Biology.

Under his guidance, Capgemini has launched an AI-driven biotechnology lab in Cambridge, where advanced AI applications are delivering world-leading innovations and products to accelerate the bioeconomy.

Conor Griffin

AI Policy Research Lead, Google DeepMind

Conor is an AI policy research lead at Google DeepMind. He carries out research on topics related to the safe and responsible development of AI, both with respect to the actions that AI labs should take, and potential public policy interventions. In recent years, he has carried out research on various topics relating to the intersection of AI and science, for example on the

potential biosecurity benefits and risks posed by emerging AI models. He sits on the UK Biosecurity Leadership Panel. Prior to joining Google DeepMind, Conor spent a decade at The Economist Intelligence Unit (EIU) where he worked on public policy research, with a focus on science and technology policy. He is currently based in London, and previously worked for many years in China and the Middle East.

ABOUT APPG AI

All Party Parliamentary Group on
Artificial Intelligence

ABOUT:

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

Parliamentary APPG AI Members: House of Commons

- Allison Gardner MP Labour (APPG AI Co-Chair)
- Alison GRIFFITHS MP Conservative
- Andrew Pakes MP Labour
- Daniel Aldridge MP Labour
- David Reed MP Conservative
- Dawn Butler MP Labour (APPG AI Vice-Chair)
- Esther McVey MP Conservative
- George Freeman MP Conservative
- Gordon McKee MP Labour
- Graham Leadbitter MP SNP
- Liam Byrne MP Labour
- Mike Martin MP Liberal Democrat
- Martin Wrigley MP Liberal Democrat
- Peter Fortune MP Conservative
- Samantha Niblett MP Labour
- Tom Collins MP Labour
- Tony Vaughan MP Labour
- Sir Mark Hendrick MP Labour
- Zöe Franklin MP Liberal Democrat
- Dr Zubir Ahmed Labour

Parliamentary APPG AI Members: House of Lords

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

aippg appg

All Party Parliamentary Group on
Artificial Intelligence

THANK YOU TO OUR SUPPORTORS

Helping Us Raise Our Ambition for What Can Be Achieved

Meet Ami Love Ami

bsi.

Brunel University London

BT Group

Capgemini

Centre for the Governance of AI

CMS law·tax·future

cognizant

Deloitte.

duolingo

EY

FAVOM

Hewlett Packard Enterprise

UKRI Innovate UK

onfido

Osborne Clarke

Queen Mary University of London

Santander

UPTITUDE

WINDWORKX

ACCESS APPG AI RESOURCES, EVENTS AND FULL PROGRAMME

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

Annual Programme

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

Networking

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

Resources

Reports, transcripts, videos, and photo albums.

Download your Pavilion App Now!

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

From your computer:

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

From your mobile:

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

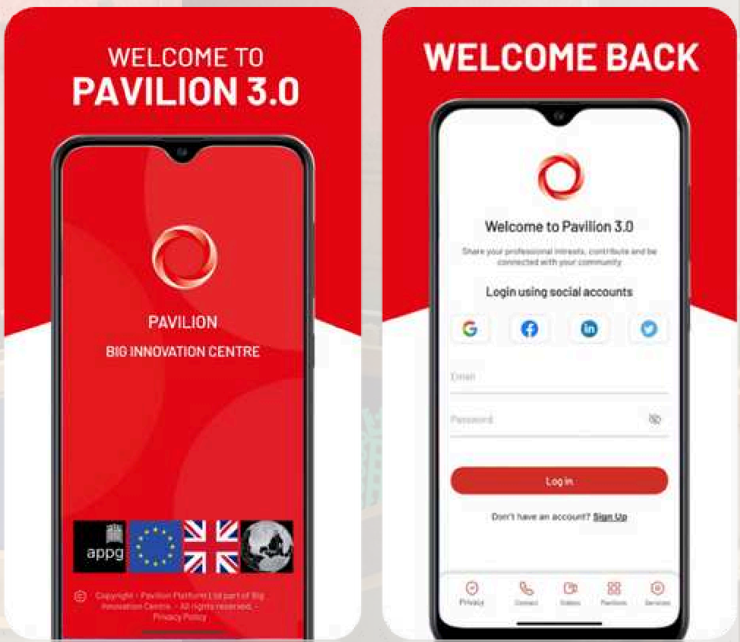

Please use the same username and password across all web and mobile app devices, avoiding the hassle of multiple accounts.
Click below:

visit our website bicpavilion.com

Download on the App Store

GET IT ON Google Play

All Party Parliamentary Group on
Artificial Intelligence

CONTACT

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

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

Contact:

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

All-Party Parliamentary Group on
Artificial Intelligence

appg@biginnovationcentre.com

SECRETARIAT

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

BIG INNOVATION CENTRE