

All-Party Parliamentary Group on Artificial Intelligence

AI & Government

Al in the Public Sector -Redefining Government & Welfare With Al



BIG INNOVATION CENTRE Secretariat

20 January 2025 Policy Forum

Table of Contents

Introduction: Page 3 Aim of Session: Page 4 Findings: Page 6

Evidence: Page 11

- Jon Atkin: Co-founder and Chief Executive Officer at Favom: Page 12
- Rachel Astall: Chief Customer Officer (CCO) at Beam: Page 19
- Morgan Rees: Vice President of Data, Analytics & Al at Capgemini: Page 26
- Dr Cosmina Dorobantu: Co-Director, Public Policy Programme, The Alan Turing Institute, and Turing OII Fellow, Oxford Internet Institute, University of Oxford: Page 30
- David Fearne: Global Director of Generative Al at Cognizant: Page 37
- Hugh Eaton, Colonel HRA Eaton OBE, MA, VR, UK Ministry of Defence: Page 42

BIOs of Evidence Givers: Page 48 About APPG AI: Page 55 Contact: Page 59



Title: Al & Government: Al in the Public Sector - Redefining Government & Welfare With Al Publication Details: Authoring Organisation: All-Party Parliamentary Group on Artificial Intelligence (APPG AI) Date of Publication: January 2025 Publication Type: Parliamentary Brief | Policy Brief Publisher: Big Innovation Centre



INTRODUCTION

This document is a transcript and summary of an APPG AI evidence meeting that took place on 20 January 2025 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 & Government: AI in the Public Sector - Redefining Government & Welfare With AI
- Time 5:30 pm 7:00 pm (GMT)
- Date: Monday, 20 January 2025
- Venue: Committee Room 1 in the House of Lords.

CONTACT THE SECRETARIAT

appg@biginnovationcentre.com APPG AI Secretariat Big Innovation Centre

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

EVIDENCE GIVERS

- Jon Atkin: Co-founder and Chief Executive Officer at Favom
- Rachel Astall: Chief Customer Officer (CCO) at Beam
- Morgan Rees: Vice President of Data, Analytics & Al at Capgemini
- Dr Cosmina Dorobantu: Co-Director, Public Policy Programme, The Alan Turing Institute, and Turing OII Fellow, Oxford Internet Institute, University of Oxford
- **David Fearne:** Global Director of Generative AI at Cognizant
- Hugh Eaton, Colonel HRA Eaton OBE, MA, VR, UK Ministry of Defence

MEETING CHAIRS AND RAPPORTEUR

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



Aim of Session

Al & Government: Al in the Public Sector - Redefining Government & Welfare With Al

In this Evidence Meeting, we addressed the potential and challenges of integrating Al into the public sector to enhance government services and welfare. We focused on identifying areas where AI can improve efficiency, streamline processes, and support decision-making while addressing risks such as bias, privacy concerns, and security issues. The meeting also examined the development of trustworthy data infrastructures and AI tools tailored to civic needs, emphasising transparency, inclusivity, and adaptability. These discussions aimed to provide actionable insights for leveraging AI responsibly within the government and the public sector while maintaining public trust and ethical integrity.

Questions were raised to inspire the discussion:

- Key theme: In what areas can AI meaningfully contribute to efficiency and effective government services and welfare programs? What risks require mitigation?
- How can the public sector build trustworthy data infrastructure and AI tools suited for civic needs?
- Where should the boundary between human decision-making and AI intervention be in public sector operations?
- What guardrails and ethical considerations should guide AI adoption in government decision-making for transparency and accountability, particularly in areas impacting basic welfare?



From left to right: David Fearne (Global Director of Generative AI at Cognizant), The Rt. Rev the Lord Bishop of Oxford (APPG AI Parliamentary Member), Baroness Uddin (APPG AI Parliamentary Member), Professor Birgitte Andersen (CEO Big Innovation Centre and APPG AI Secretariat), Lord Clement-Jones CBE (APPG AI Co-Chair), Allison Gardner MP (APPG AI Co-Chair), Jon Atkin (Co-founder and Chief Executive Officer at Favom), Rachel Astall (Chief Customer Officer (CCO) at Beam), Dr Cosmina Dorobantu (Co-Director, Public Policy Programme, The Alan Turing Institute, and Turing Oll Fellow, Oxford Internet Institute, University of Oxford), Morgan Rees (Vice President of Data, Analytics & AI at Capgemini) and Hugh Eaton (Colonel HRA Eaton OBE, MA, VR, UK Ministry of Defence)

Full Committee Room in excess of 70 participants. Images from APPG AI events are available on the APPG AI Pavilion - please refer to page 58.

FINDINGS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

ACTION FIELDS FOR POLICY AND STAKEHOLDER GROUPS

Summary of Action Points for Government and Policymakers:

1. Leverage AI for Public Service Efficiency & Improvement

- Expand Al-driven digitisation in welfare and government services (such as health, defence, security and so forth) to improve efficiency.
- Use Al-powered data analytics to enhance decision-making in welfare programs and public sector operations.
- Invest in Al solutions that streamline bureaucracy while ensuring human oversight in critical decision-making.

2. Ensure AI Transparency & Managing Risks

- Mandate explainable AI in public sector applications to prevent "black box" decision-making.
- Implement regular audits for AI bias and fairness, particularly in welfare and government services.
- Encourage academic and scientific engagement to increase public understanding of Al technologies.

3. Reform AI Procurement for Innovation & Competition: Encourage a level playing field for competition and innovation beyond large incumbents.

- Remove barriers that prevent AI startups and SMEs from competing for public sector contracts.
- Develop innovation-driven procurement frameworks to promote competition and reduce reliance on large tech firms.
- Ensure AI procurement prioritises transparency, accountability, and long-term national AI capabilities.
- 4. Strengthen AI Governance, Ethics and Security Frameworks
- Establish clear boundaries between Al assistance and human decision-making in public services.
- Implement guardrails to ensure Al does not make final decisions on welfare and human rights issues.
- Develop ethical AI guidelines enforcing transparency, accountability, and human oversight.
- Establish clear guidelines ensuring Al is ethical, legal, robust, and sustainable:
 - Ethical Al: Maintain human agency and alignment with societal values.
 - Legal AI: Ensure compliance with data protection and copyright laws.
 - Robust AI: Monitor AI performance and require explainable decisions.
 - Sustainable AI: Reduce environmental impact of AI models.
- Promote responsible AI use while educating stakeholders about risks and uncertainties.

5. Build a Trustworthy AI & Data Infrastructure

- Invest in secure, transparent, and ethical AI solutions for public services.
- Promote international AI collaboration while maintaining control over critical national AI infrastructure (including defence and security).
- Support open-source AI development to ensure public sector AI remains accountable and explainable.

6. Facilitate Safe and Responsible Al Adoption

- Adopt a 'test and learn' model, balancing domain expertise with technical knowledge.
- Provide hands-on AI training for frontline workers to ensure effective adoption.
- Create testing environments for piloting AI technologies within the public sector and publish AI pilot results and lessons learned to inform government-wide implementation.
- Prioritise AI security to protect sensitive data and ensure responsible usage.

7. Automate Government Processes To Enhance Public Services

- Identify key areas where AI can improve communication and streamline bureaucratic processes.
- Automate administrative tasks such as passport applications and benefit eligibility checks.
- Use AI for economic modelling and policy decisions, such as assessing workforce skill gaps.

8. Scale Al Across Government Services

- Invest in AI tooling that enhances both productivity and social service effectiveness.
- Collaborate with industry to establish strategies for scaling AI solutions beyond pilot projects.

Build the necessary infrastructure, processes, and data foundations for large-scale Al adoption.

9. Positioning the UK as a Global AI Leader

- Leverage UK strengths in mathematics, research, and AI to establish leadership in AI development.
- Promote UK AI capabilities globally to attract investment and strategic partnerships.
- Treat AI as a national strategic asset, focusing on homegrown innovation rather than foreign dependency.

10. Develop a whole-of-government AI strategy

- Implement key recommendations from the Matt Clifford Report on Al governance.
- Strengthen central government capabilities for managing large-scale Al initiatives.
- Develop and enforce ethical frameworks for Al in defence, ensuring human oversight in decisionmaking.
- Pursue national and international AI partnerships to strengthen defence and security capabilities.
- Provide Al-focused training for government executives to enhance strategic decision-making.

Al can revolutionise government services by improving efficiency, transparency, and responsiveness. However, responsible adoption requires strong governance, transparency, and ethical oversight. Policymakers must prioritise Al investment, security, and public trust, ensuring Al benefits society. By leading in Al governance and infrastructure, the UK can become a global Al leader while safeguarding national interests. There are also action points for stakeholder groups beyond just the government in the context of the AI in the Public Sector discussion. These groups would include:

1. SMEs & Startups (AI Providers)

- Al Procurement: Remove barriers preventing Al startups and SMEs from participating in public sector procurement.
- Innovation-driven Frameworks: Collaborate with government agencies to develop procurement frameworks that encourage competition, ensuring that smaller tech providers have the opportunity to innovate and contribute.
- Ethical Al Development: Ensure that Al solutions developed by SMEs align with ethical, transparent, and accountable guidelines to maintain public trust.

2. Industry Partners & Technology Providers

- Al Tooling and Infrastructure: Invest in Al tools that both enhance productivity and serve public good, with a focus on transparency and effectiveness.
- Collaboration with Government: Partner with the government to scale Al solutions from pilot projects to full implementation, providing the necessary infrastructure and data foundations for Al adoption at a national level.
- Responsible AI: Work closely with the government and other stakeholders to ensure the responsible development and deployment of AI systems, addressing AI risks, fairness, and accountability.

3. Frontline Workers & Public Sector Employees

- Training & Support: Frontline workers should receive practical, hands-on training for successful adoption of AI tools in their daily tasks, such as in healthcare and welfare programs.
- Adaptation to Al-driven Systems: Invest in upskilling programs to help workers adapt to Al systems, ensuring that they can leverage technology to improve efficiency while maintaining their role in decision-making processes.

4. Academia & Research Institutions

- Public Understanding of Al: Engage in academic research to promote transparency in Al systems and develop explainable models for public sector applications.
- Bias Audits & Ethical Research: Contribute to studies that evaluate the ethical and bias aspects of AI applications, especially in government services and welfare programs.
- Training Programs for Policymakers & Government Officials: Collaborate with government bodies to design educational initiatives that help policymakers better understand AI technologies and their societal implications.

5. Citizens & Public (End Users of Al in Public Services)

- Transparency & Accountability: Advocate for AI systems that prioritize clear, understandable outputs and give citizens confidence in how decisions are made (e.g., explainable AI).
- Feedback Mechanisms: Encourage feedback loops that allow citizens to share their experiences and concerns about AI applications in public services. This ensures AI solutions are continually refined to meet public needs.
- Public Awareness Campaigns: Raise awareness about the benefits and risks of Al in government services to foster public trust and promote responsible usage.

6. International Partners & Collaborators

- Global Al Collaboration: Engage with international partners to share Al knowledge, technology, and strategies. Collaborate on setting global standards for Al governance and ethics.
- Technology & Defence Partnerships: Pursue strategic partnerships with other nations to enhance Al capabilities, especially in national security and defence.

These action points reflect how various stakeholders—beyond the government—can play a critical role in the successful implementation of AI in the public sector, contributing to a more responsible, inclusive, and effective AI-driven future.





Evidence Giver: Jon Atkin



Evidence Giver:



Morgan Rees Dr Cosmina Dorobantu

Evidence Giver: Rachel Astall



Evidence Giver: David Fearne



Evidence Giver: Hugh Eaton



APPG AI Chair: Allison Gardner MP



Evidence Giver:

APPG AI Chair: Lord Clement-Jones CBE



Secretariat & Rapporteur: Professor Birgitte Andersen

EVIDENCE

Jon Atkin

Co-founder and Chief Executive Officer at Favom



Introduction

I am the CEO of a UK-based SME specialising in business technology with a focus on artificial intelligence. I'm joined today by my colleague and co-founder, Professor Stefan Zohren of the University of Oxford. I'd like to begin by addressing a topic raised recently by the Prime Minister in his speech on AI, an address that greatly inspired both myself and our team. Specifically, he posed the question: how can the UK position itself as a global leader in artificial intelligence? I would like to share some insights and evidence in support of this ambition, drawing on our experience over the years working with the public sector, both domestically and internationally, in the field of AI.

Personal Experience

In preparing for today, one of the questions raised was about the kind of evidence we can bring, and I believe personal experience plays an important role in that. My passion for applying technology to real-world challenges began around 25 years ago, when I was invited by the NHS to visit a medical records library, an experience that proved foundational to my career.

This particular hospital, which has since become a long-standing client of ours, managed approximately 1.2 million patient records at the time, not pages, but full medical records, supported by over 130 staff. Stepping into the facility felt more like entering an aircraft hangar due to its vast size. At the end of the tour, I noticed two large boxes stacked with files, ready to be sent to clinics the following day. Assuming they contained all the records needed for the hospital's operations, I was astonished to learn they were for a single sixmonth-old baby.

The medical records manager explained that because infants can't verbalise their symptoms, they often undergo multiple consultations and tests, generating substantial documentation. They explained to me how babies, obviously because they can't communicate and say, "I've got a bit of a pain here," or "I didn't sleep well last night," it's very difficult to diagnose a baby, so they tend to go round the hospital, and lots of paper is created.

I later spoke with the paediatrician who was scheduled to see the baby and their parents. I asked how he managed to process that volume of information. His response was candid: "I can't. It's impossible to read and digest that much paper-based data." That moment crystallised for me the need for more intelligent, accessible, and efficient systems, an insight that continues to guide our work in AI to this day.

Digitisation of Records

We began by digitising those paper records and have since evolved our approach to incorporate the latest advancements in artificial intelligence. Today, clinicians can analyse vast amounts of patient data in real time, identifying patterns and trends across thousands of data points. Our technology provides a summarised, prioritised view of the patient's record, supporting more informed and efficient clinical decision-making.

To be clear, we do not attempt to diagnose. All is not yet at a stage where it can fully replace clinical judgement, but we aim to assist, to guide, and to present relevant information at the right time, enhancing the clinician's ability to deliver care.

Global Position

That's the personal foundation of my passion. Turning briefly to the broader picture, I'd like to touch on the UK's global position in AI. While the United States currently appears to hold a dominant role, I find this somewhat frustrating because I believe the UK has all the essential ingredients to lead in this space.

We are home to some of the world's most brilliant mathematicians, a rich legacy of scientific and technological innovation, from pioneers like Tim Berners-Lee, and a network of world-class research institutions. The foundation is here. With the right focus and support, I believe the UK is well-positioned to become a global leader in AI.

Importance of International Collaboration

International collaboration is absolutely vital to the advancement of AI and technology more broadly. Although we are an SME rather than a large corporation, we have established a successful partnership agreement with Singapore, which has been instrumental in embedding advanced algorithms into eye care and facilitating the mutual exchange of innovation and expertise.

What we've seen through this collaboration, and in discussions with partners in the Middle East, is a shared enthusiasm and commitment to working together. It's clear that the UK is well-positioned to strengthen its role on the global stage by fostering such international relationships and leveraging them to drive innovation.

Investment Frustrations

To be candid, investment in AI remains one of my key frustrations, particularly within public sector procurement. Too often, purchasing decisions are driven by familiarity with established platforms and well-known brands, rather than by the underlying innovation or long-term value a solution may offer.

One of the reasons I'm delighted to be a part of this APPG is the belief that AI fundamentally stems from mathematics. If we're serious about building long-term capability and leadership in this field, we must look beyond software alone. A critical part of that strategy should involve investing in the next generation of mathematicians and ensuring a strong pipeline of talent that underpins the future of AI development in the UK.

The Future of AI

The future of AI is both dynamic and full of promise. When new capabilities, such as those introduced by organisations like OpenAI, are released, they often generate surprise and excitement among the public. However, for those of us working at the forefront of this field, these developments are rarely unexpected. We typically have a clear view of the direction AI is taking and the innovations that are on the horizon.

That said, our clients and partners may not share the same level of insight. This creates a responsibility for us, not only to innovate, but also to communicate. As subject matter experts, we must bridge the gap between what we know is coming and what others may be anticipating. This means being transparent about our roadmaps, sharing insights into emerging trends, and engaging proactively before new products and capabilities materialise.

By doing so, we can reduce uncertainty, foster trust, and help organisations prepare for change, ultimately ensuring that AI technologies are adopted in a way that is thoughtful, strategic, and beneficial to all stakeholders.

Procurement Challenges for SMEs

From a procurement perspective, SMEs and startups, both in the UK and internationally, often face significant barriers to entry when it comes to engaging with the public sector on technology initiatives. Many of you will be familiar with these challenges, which can result in innovative, high-potential companies being excluded simply because they cannot meet initial thresholds or navigate complex procurement processes.

This is something that must change. Embracing SMEs in public sector procurement is essential for driving innovation and achieving better outcomes. One of our sister companies operates a public sector procurement framework specifically designed to address this issue. It enables smaller, agile technology firms to compete on a level playing field, and I can say with confidence that we've seen considerable global interest from organisations eager to work with the public sector, but previously unable to gain access through traditional channels.

Black Box Marketing

A key challenge associated with many of today's widely recognised AI platforms is what I refer to as "black box marketing". These technologies are often promoted based on their powerful capabilities, yet offer limited transparency into how they actually function. Beyond data scientists and mathematicians, very few individuals fully understand the underlying mechanisms. This opacity fosters uncertainty, anxiety, and mistrust, barriers that must be overcome if we are to build lasting public confidence in AI.

As we engage in vital discussions around ethics, regulation, and the responsible deployment of AI, transparency must remain a core principle. The UK is well-positioned in this regard, with a strong tradition of academic excellence and regulatory thought leadership. What is needed now is greater collaboration across government, academia, and industry to empower experts to communicate how these systems work in an accessible, meaningful way. By doing so, we can demystify AI and encourage a more informed, constructive public discourse.

Concluding Remarks

In closing, I believe there is a critical need to strengthen collaboration between the public and private sectors to fully realise the UK's potential in AI. The work of the AI All-Party Parliamentary Group (APPG) is a significant and welcome step in that direction, and we were honoured to be invited to join last year.

Looking ahead, we must do more to promote the UK's capabilities on the global stage, clearly and confidently showcasing who we are and what we can achieve. Finally, I believe AI is fast becoming an essential utility, much like the internet and Wi-Fi. As such, it is vital that we remain focused on innovation. If we do not lead in the development of these technologies, we risk becoming passive consumers rather than active creators.

Summary of Jon Atkin's Key Points

1. Al's Contribution to Government Efficiency & Welfare Services

- Healthcare Innovation: Al has the potential to transform healthcare delivery. The digitisation of medical records, combined with Al-driven data analysis, enables clinicians to identify patterns, prioritise care, and make more informed decisions, ultimately improving patient outcomes and operational efficiency.
- Public Sector Automation: Al can streamline administrative functions across government departments, reducing bureaucratic overhead and enabling faster, more accurate service delivery to citizens.
- Data-Driven Welfare Programs: Advanced analytics powered by AI can support better decisionmaking in welfare and social services. By uncovering trends and optimising resource allocation, AI helps ensure support is targeted where it is most needed.

2. Risks & Mitigation Strategies

- Transparency and Trust: The opaque nature of many Al systems, often referred to as "black box" models, can undermine public trust and accountability. To address this, it is essential to promote Al literacy and ensure that experts are actively involved in explaining how these technologies function in accessible and transparent ways.
- Bias and Fairness: Al models are susceptible to bias if not properly designed and monitored. Mitigating this risk requires the use of diverse and representative data sets, regular auditing of algorithms, and ongoing human oversight to ensure fairness, equity, and ethical integrity in decisionmaking.
- Procurement Barriers for SMEs: The exclusion of SMEs and startups from public sector Al
 procurement risks creating an over-reliance on a small number of dominant technology providers,
 stifling innovation and competition. Reforming procurement frameworks to be more inclusive and
 accessible can help foster a more dynamic and innovative Al ecosystem.

3. Building Trustworthy AI & Data Infrastructure for Civic Needs

- Leveraging National Strengths: The UK is uniquely positioned to lead in the development of transparent, reliable AI tools for the public sector, drawing on its world-class capabilities in mathematics, academic research, and technological innovation.
- Promoting International Collaboration: Strategic partnerships, such as ongoing collaborations with Singapore, can facilitate the exchange of knowledge, accelerate innovation, and help establish globally recognised best practices for deploying Al in civic contexts.
- Prioritising Transparency and Accessibility: Public sector Al investment should focus on open-source, explainable, and ethically designed solutions. Doing so will enhance trust, support accountability, and ensure that Al systems remain inclusive and accessible to all stakeholders.

- 4. Boundary Between AI and Human Decision-Making
 - Al Supporting, Not Replacing, Human Judgment: In critical domains such as healthcare and welfare, Al should be used to support, not substitute, human decision-making. These areas require empathy, contextual understanding, and ethical sensitivity that Al alone cannot replicate.
 - Enhancing Insight, Not Authority: AI can play a valuable role in surfacing trends, identifying risks, and providing prioritised insights. However, final decisions, particularly in sensitive or high-stakes scenarios, must remain the responsibility of qualified professionals.
 - Policy Focus on Human-Centric Al: Government Al strategies should prioritise the augmentation of human expertise. This includes ensuring that Al tools are designed to empower professionals, not to automate decisions where human oversight is essential for fairness, accountability, and ethical integrity.

5. Guardrails & Ethical Considerations for AI in Government

- Transparency and Explainability: To build and maintain public trust, transparency and explainability must be foundational to any Al adopted in government. Citizens have a right to understand how decisions are being made, particularly when those decisions impact their lives directly.
- Robust Ethical Frameworks: Ethical AI deployment requires clear frameworks that incorporate bias detection, regular auditing, and consistent human oversight. These measures are essential to ensure fairness, accountability, and compliance with democratic values.
- Strategic National Asset: Al should be recognised and treated as a strategic national asset. For the UK to remain a leader, not merely a consumer, of Al technologies, it must invest in domestic capabilities and innovation.
- Stronger Public-Private Collaboration: Meaningful engagement between the public and private sectors is essential to align Al development with civic priorities and ethical standards. Collaboration can ensure that innovation serves the public good while adhering to principles of transparency, equity, and trust.



Summary of Policy Application of Jon Atkin's Evidence Statement

1. Leveraging AI to Enhance Efficiency and Public Service Delivery

- Expand Al-Enabled Digitisation: Accelerate the adoption of Al-driven digitisation across healthcare and government services to improve operational efficiency, reduce administrative burden, and enhance service accessibility.
- Deploy Advanced Data Analytics: Utilise Al-powered analytics to support evidence-based decision-making in welfare programmes and broader public sector operations, enabling more targeted and effective resource allocation.
- Invest in Human-Centric Automation: Prioritise investment in Al solutions that streamline bureaucratic processes while maintaining human oversight, particularly in high-stakes areas such as healthcare and welfare assessments.

2. Addressing Risks and Ensuring AI Transparency

- Mandate Explainability in Public Sector AI: Require the use of explainable AI in all government applications to mitigate concerns associated with opaque, "black box" systems and to promote accountability in public decision-making.
- Implement Rigorous Auditing Protocols: Establish regular and independent audits to assess Al systems for bias, fairness, and ethical compliance, particularly in sensitive areas such as welfare and public service delivery.
- Promote Scientific and Academic Engagement: Foster closer collaboration with academic and research communities to improve public understanding of AI technologies, demystify their functionality, and support informed policy development.

3. Reforming AI Procurement to Include SMEs and Startups

- Eliminate Barriers to Entry: Address structural and procedural barriers that prevent Al startups and SMEs from accessing public sector procurement opportunities, ensuring a more inclusive and competitive supplier ecosystem.
- Adopt Innovation-Led Procurement Frameworks: Design procurement processes that actively encourage innovation, support emerging technologies, and reduce reliance on a small number of dominant technology providers.
- Align Procurement with Strategic National Interests: Ensure that public sector AI procurement prioritises transparency, accountability, and the development of long-term national AI capabilities, reinforcing the UK's position as a leader in ethical and sovereign AI innovation.



4. Strengthening Al Governance and Ethical Frameworks

- Define Clear Human-Al Boundaries: Establish clear guidelines delineating the role of Al as a decision-support tool, ensuring that final authority in public services, particularly in areas affecting welfare and fundamental rights, remains with human professionals.
- Implement Robust Safeguards: Introduce enforceable guardrails to prevent Al systems from making autonomous decisions in high-stakes or ethically sensitive contexts, reinforcing public trust and accountability.
- Develop Comprehensive Ethical Standards: Create and uphold ethical Al frameworks that mandate human oversight, system transparency, and accountability throughout the lifecycle of Al deployment in the public sector.

5. Building Trustworthy AI and Data Infrastructure

- Invest in Ethical, Transparent Al Solutions: Prioritise the development and deployment of secure, transparent, and ethically designed Al tools that are purpose-built to address the unique needs of the public sector.
- Foster International Collaboration with Strategic Oversight: Promote global cooperation in AI research and application while ensuring that the UK retains strategic control over critical AI infrastructure and capabilities.
- Support Open-Source and Explainable AI: Encourage the use of open-source AI solutions to enhance transparency, foster accountability, and ensure that public sector AI systems remain understandable, auditable, and in the public interest.

6. Positioning the UK as a Global Leader in Al

- Leverage National Strengths: Capitalise on the UK's world-class expertise in mathematics, scientific research, and AI development to establish a leadership position in the global AI landscape.
- Promote Global Visibility and Collaboration: Proactively showcase the UK's Al capabilities on the international stage to attract investment, foster strategic partnerships, and enhance global influence in Al governance and innovation.
- Treat AI as a Strategic National Asset: Recognise AI as a critical component of national infrastructure and competitiveness, ensuring that the UK drives innovation rather than becomes dependent on foreign-developed technologies.



Rachel Astall:

Chief Customer Officer (CCO) at Beam

Introduction

My name is Rachel Astall. I'm the Chief Customer Officer at Beam, where I oversee our work with Local Authorities. My evidence today will be focused on how AI can support the efficiency and effectiveness of welfare services.

About Beam

Beam is a British social enterprise bringing AI to welfare services for more than 100 government partners, across both local authorities and central government departments. We have a team of almost 100 case workers helping people up and down the country get into stable housing, secure long-term work, learn English, and integrate into communities. Through this work, we have a direct impact on 5 of the current government's 6 missions.

The Impact of AI on Welfare Services

Al has already enabled our team of expert caseworkers to cut paperwork, surface information rapidly, and improve response times.

We're now using our direct frontline experience to roll out transformative AI software to thousands of frontline workers across the UK.

What is Magic Notes?

Our first tool, Magic Notes, uses AI to draft paperwork for frontline workers. Magic Notes is used to record conversations and produce high-quality draft reports in minutes, which they check and refine.

It sounds simple. The impact of a tool like this may not be immediately apparent to those who do not work in the field.

The Administrative Burden

Social workers spend more than three days of every working week on admin and paperwork.

Whilst a portion of this is high value—for example, using professional expertise to draw insights and make recommendations for an individual's care—the sad truth is that most of this time is spent simply transcribing notes from meetings and formatting them into reports.

Clearly, this is incredibly time-consuming and financially costly.

But, more than that, I hear time and time again from social workers that it also means they struggle to deliver the best possible care. That the relentless need to manually keep records keeps their heads in a notebook, rather than focused on the person in front of them. That they spend long hours catching up on admin.

This is not what any social worker came into the profession to do. And it is adding to the widespread burnout and poor retention of frontline workers, at a time when demand has never been higher.

A Huge Opportunity for Change

Al gives us a huge opportunity to change this.

Through multiple independent evaluations, conducted by both Local Authorities and an academic team headed by Professor Rob Procter of the University of Warwick, we have seen third-party evidence that Magic Notes can save each frontline worker more than 8 hours per week in admin time. A full day. Every week.

To put that in perspective, if just 100,000 frontline workers used Magic Notes, we would save the public sector more than 7,500 years of time, every single year.

That's worth more than £40 million pounds per week to the taxpayer.

Beyond Time Savings

As staggering as that might be, the impact goes well beyond direct time savings.

Those using Magic Notes consistently report improved connection with those they support.

Managers report that records are high quality and submitted faster. One manager in Wales told us that "it saves time, but most importantly improves the quality of practice and engagement with service users. I've worked in social work for 20 years and this is the biggest advancement for us as a profession I've ever known". Residents report that they feel more heard. One person told us that they feel like their social worker "is actually interested in what I'm saying ... not just taking notes like a robot", and not "just ticking boxes like other caseworkers".

This all means higher quality, more effective care.

Reconnecting with Their Passion

And Social Workers are getting back to what they love.

That includes the occupational therapist from Thurrock who told us that they "thought the admin side of the job would make me want to quit" but not anymore.

Or the social worker from Sheffield, who said Magic Notes "provided me with the confidence and support I desperately needed.... I no longer spend days worrying about my write-ups, and ... work-life balance has improved significantly".

Or CJ, who shared that "this has been one of the most productive weeks I have had ... I can focus on being a social worker, and not spend as much time writing up the discussions I have had."

The Transformational Potential

A tool this impactful has the potential to fundamentally change welfare for the better.

The Responsibilities of AI Development

But of course, this is not without risk.

It is an enormous privilege to be living in an age of such profound technological advancement. It is an even greater privilege to be working on building the tools of the future. This comes with responsibilities we take extremely seriously.

Our strong belief is that we need to leverage technology, including AI, to enable humans to do the things which only humans can do, and things we all only want humans to do.

This means that technology, including AI, must be deployed both safely and responsibly.



Risks to Address

Big tech could be in such a race to win with AI that the deep and important questions around how data is actually used and how users themselves interact with tools become secondary.

Security is clearly of paramount importance when handling sensitive data. It's good that security accreditations exist, and even better that the government is creating fertile ground for British tech companies, so data hopefully won't ever have to cross the Atlantic for British citizens to receive world-class services.

And we hope security continues to be very high up the government agenda.

Beyond technical security, we must also consider how we use Al responsibly in this setting.

Our Commitment to Responsibility

At Beam, we invest heavily in ensuring our AI tools produce clear, accurate outputs. We design our products so that a qualified, human worker needs to check every output. We seek constant feedback from frontline workers. And our organisation has people with direct frontline experience at every level—from our co-founder, to our engineers writing code and prompts, to our sales team working with customers.

Deep contextual understanding runs through our DNA.

The Risk of Poorly Designed Technology

This highlights another risk we must draw the government's attention to.

If we fail to design this transformative technology with the complex, real-world scenarios of welfare settings in mind, we risk seeing it gather dust on the bookshelf, unused or ineffective.

The needs of an occupational therapist in the Orkney Islands are different to an educational psychologist in Hackney—or an MP in Westminster. We must design and build tools specifically for the welfare settings they will be used in.

Equally, we could build the best tools in the world, but if we do not invest in on-the-ground support, we will fail to see them adopted.



We've worked with councils where social workers still use Nokia 3310s. Where social workers have used pen and paper for decades. Successful adoption of AI in welfare settings will not be achieved through a centralised software update. We need to connect tools to people with human support, practical and handson training and deep understanding of local context. We need to build tools that frontline workers truly love to use.

The Role of Government in Supporting AI Adoption

The only way we will achieve this is through deep partnership between technology and service delivery teams.

Government Recommendations

So how can the government help?

We recommend that the government fully endorse the 'test and learn' model for responsible AI in welfare services through four areas:

- Firstly, acknowledging the importance of domain expertise ensures that frontline experience is given equal weight alongside technical expertise.
- Secondly, promoting a level playing field. Whilst it is tempting to favour large incumbents who are well known, true "test-and-learn" will require thinking beyond the familiar. Open and honest competition will drive the innovation we need.
- Third, the creation of a procurement landscape which supports government bodies to trial innovative solutions safely, responsibly, and quickly.
- And finally, a commitment to publishing pilot results, so that all government departments can benefit from shared lessons and collective experience.





Summary of Rachel Astall's Key Points

- **1.Al in Welfare Services:**
 - Beam is a British social enterprise using AI to enhance welfare services for over 100 government partners, including Local Authorities and central government departments.
 - The organisation supports individuals in securing housing, employment, learning English, and community integration.
- 2. The Role of Al Tools (Magic Notes):
 - Magic Notes is an Al-powered tool that drafts paperwork for frontline workers, saving them significant time and improving the quality of their work.
 - Independent evaluations show that Magic Notes saves over 8 hours of admin time per week per frontline worker, equating to significant cost and time savings for the public sector.

3. Impact Beyond Time Savings:

- Al tools like Magic Notes enhance the connection between social workers and service users, improve record quality, and reduce burnout among workers.
- Social workers report improved work-life balance and greater job satisfaction.

4. Risks and Responsibilities of Al Adoption:

- Al tools must be designed specifically for the complex, real-world scenarios of welfare settings.
- Successful adoption requires investment in training, human support, and understanding local contexts.
- The government must prioritise security and ensure Al tools are deployed responsibly.

Action Points for Government and Policymakers

- **1. Support the 'Test and Learn' Model:**
 - Acknowledge the importance of domain expertise, giving equal weight to frontline and technical expertise.
 - Promote a level playing field by encouraging competition and innovation beyond large incumbents.
- 2. Facilitate Safe and Responsible Adoption:
 - Create a procurement framework that enables government bodies to trial innovative solutions quickly and responsibly.
 - Provide practical, hands-on training for frontline workers to ensure successful adoption of AI tools.
- **3. Encourage Evidence Sharing:**
 - Commit to publishing pilot results and lessons learned to benefit all government departments.
- 4. Ensure Security and Accountability:
 - Prioritise technical security to protect sensitive data.
 - Design tools with clear, accurate outputs requiring human review to ensure responsible Al usage.

According to Rachel Astall's evidence:

- Savings Per Worker: Magic Notes saves each frontline worker over 8 hours of admin time per week—equivalent to a full working day.
- National Impact: If 100,000 frontline workers use Magic Notes, this equates to saving over 7,500 years of time annually.
- Financial Savings: These time savings translate to more than £40 million per week for the public sector.
- Annual Savings: Over a year, this would amount to approximately £2 billion in potential savings.
- These figures underscore the significant financial benefits of adopting Alpowered solutions like Magic Notes across welfare services, alongside the operational and societal benefits.



Morgan Rees

Vice President of Data, Analytics & Al at Capgemini

AI Adoption Is Inevitable

Al adoption is inevitable, but it is not a runaway train. Humans will decide how, when, and why Al is used. However, we need to be proactive and conscious in the choices we make to design the future we want.

There is too much to lose by not making the most of Al.

The Spectrum of AI

For me, AI is a spectrum. On one end, it is a tool to automate the boring tasks, such as reading and summarising documents. On the other end, it is a tool that can rethink the role of government in society.

The value of AI can be broken down into two main categories:

- 1. Al as a Productivity Engine: Improving speed, accuracy, and throughput.
- 2. Al as a Service Improvement Engine: Making services more effective, not just more efficient.

Examples of Productivity Gains

- Defence Sector:
- Large physical assets require complex and costly maintenance. Drones capture hundreds of hours of footage, which traditionally takes a long time to review. Using generative AI, inspection times have been reduced by 75%, allowing for near-instant turnaround between inspections and repairs.
 - This reduces costs.
 - It ensures repair work happens earlier, improving asset availability.

Examples of Service Improvement

Tax Recovery:

 In one country (not the UK), swimming pools are often undeclared because they lead to higher property taxes. Using AI-powered computer vision on geospatial images, over 20,000 undeclared swimming pools were identified in pilot regions. This uncovered nearly €100 million in recoverable tax revenue.

Transparency:

- A public sector institution has 60,000 hours of historical interviews of public interest. Al is being used to make these searchable by the public—something that would have been impossible without Al.
- This is good for society, democracy, and the economy.

The Future of AI – Large Action Models (LAMs)

Generative AI so far has largely focused on Large Language Models (LLMs), which are good at reading and summarising information. The next wave of models, Large Action Models (LAMs), goes further—they can "do" things.

We are entering a phase where AI will not just automate boring tasks but could also automate hard tasks.

For example, if AI can identify undeclared swimming pools, a LAM could go a step further and initiate legal proceedings automatically.



Why AI Tools Succeed or Fail

Delivering AI in government does not just happen. There are reasons why some AI tools succeed and others fail. To build effective AI tools, you need four key components:

1. Strategy:

- Set a clear direction for how AI can create high-value change in public services.
- Use AI to do things differently, not just faster.
- 2. Delivery Confidence:
 - Work with industry to navigate the evolving technology landscape.
 - Understand what can be procured as a service and what should be built in-house.
 - Proofs of Concept (PoCs) are easy; scaling AI is hard without the right infrastructure, processes, and data foundations.
- 3. Adoption Culture:
 - Focus on people first. Integrate AI safely and compassionately.
 - Recognise that there will be mistrust, misunderstanding, and resistance to AI.
 - Identify the skills needed today and how they will need to evolve over time.

4. Governance:

- Establish clear guardrails across four dimensions:
 - Ethical AI: Ensure human agency and alignment with values.
 - Legal AI: Ensure compliance with laws, from data protection to copyright.
 - Robust AI: Monitor AI to ensure it behaves as intended and explain decisions.
 - Sustainable AI: Consider the environmental impact of AI design.

The Challenges of Scaling AI in Government

Proof of concept is easy. Scaling AI in government is hard. Too often, the barriers are a lack of strategy, delivery confidence, adoption culture, or governance.

Call to Action for Government and Policymakers

- Invest in AI Tooling:
- The government must continue to invest in AI tools and set a clear ambition for service improvement, social benefits, and productivity gains.
- Build AI Capability at Scale:
- Work with industry to strengthen strategy, delivery confidence, governance, and adoption culture.
- Lead on Al Safety:
- The government must address safety concerns head-on, educate people about Al's uncertainties, and realise its full potential benefits.

Al offers extraordinary opportunities, but its risks and challenges must be met with action.

Summary of Morgan Rees's Key Points

The Role of Al in the Public Sector

- Al adoption is inevitable, but humans must guide its implementation consciously and proactively to shape the desired future.
- Al serves as both a productivity engine (e.g., speeding up tasks and improving accuracy) and a service improvement engine (e.g., enhancing transparency and effectiveness).

Examples of AI Applications

- Defence: Al reduces inspection times for assets by 75%, enabling quicker repairs and cost savings.
- Tax Services: Al identified 20,000 undeclared swimming pools, leading to an estimated tax recovery of €100 million.
- Transparency: Al made 60,000 hours of public-interest interviews searchable, fostering transparency and democracy.
- Emerging technologies like Large Action Models (LAMs) could automate more complex tasks, such as initiating legal proceedings.

Key Challenges in Scaling Al

- Moving from Proofs of Concept (PoCs) to widespread implementation is difficult without the right strategy, infrastructure, and data foundations.
- Public adoption of Al requires addressing mistrust and building skills for the evolving landscape.
- Governance is critical to ensure ethical, legal, robust, and sustainable AI deployment.



Summary of Morgan Rees's Key Points

Action Points for Governments and Policymakers

1. Invest in AI Tooling and Benefits:

• Focus on both productivity gains and social service improvements, emphasising transparency, effectiveness, and citizen benefits.

2. Develop Capability to Scale Al:

- Collaborate with industry to establish strategies for scaling AI solutions beyond pilots.
- Build infrastructure, processes, and data foundations to support Al adoption at scale.

3. Address Al Governance:

- Implement clear guidelines for ethical, legal, robust, and sustainable AI practices:
 - Ethical AI: Maintain human agency and alignment with values.
 - Legal AI: Ensure compliance with data protection and copyright laws.
 - Robust AI: Monitor performance and explain decisions.
 - Sustainable Al: Reduce environmental impact.

4. Foster an Adoption Culture:

- Put people first, addressing resistance and mistrust of Al.
- Invest in upskilling workers to align with the demands of Al-driven systems.

5. Lead on Al Safety:

- Continue educating stakeholders about AI risks and uncertainties.
- Promote responsible use of AI to maximise benefits while mitigating risks.

By addressing these points, governments can harness AI's transformative potential while ensuring its responsible and equitable implementation.



Dr Cosmina Dorobantu

Co-Director, Public Policy Programme, The Alan Turing Institute, and Turing OII Fellow, Oxford Internet Institute, University of Oxford

Introduction

Just a few words about me – I started my career at Google, where I had access to their amazing databases. I went on to do a PhD using data from online platforms to uncover digital trade patterns and saw how powerful it can be when you use these vast datasets to answer questions that matter to the public sector. I spent the past seven years setting up and then directing, alongside Professor Helen Margetts, a programme of research focused on empowering the public sector to use the latest generation of data-intensive technologies and to understand how to govern them responsibly. We helped over 100 public sector organisations with their AI questions and problems, from local authorities to regulators, government departments, and international entities like the United Nations and the Council of Europe.

The Potential of AI in the Public Sector

I will focus on the first of today's questions – in what areas can AI meaningfully contribute to efficiency? – because it is a hugely exciting area of work.

When we look at the public sector in the UK, it employs almost 6 million people, and it is responsible for 20% of the country's economic output. The sheer size means that, in theory, the scope for benefits from Al-driven innovation is huge.

But what does the practice look like? This is a hard question, and one way to simplify it is to think about what government does – what its day-to-day operations entail – and where AI can really help.

Three Areas Where AI Can Help

I want to propose three areas to you today corresponding to three massive government tasks:

- Communicating with the public.
- Conferring rights to individual citizens.
- Making policy decisions.

1. Communicating with the Public

Let's start with the first one. Since generative AI hit the big screens with the launch of ChatGPT, a lot of the questions that we get from various bits of the public sector are about using large language models to communicate with the public. The dream of many officials is to use these technologies to perform today's tasks – respond to citizens' queries, summarise consultation responses, generate press releases, and so on.

When carefully designed, developed, and implemented, Al technologies can lead to important efficiency gains in this area. But some of the key questions to answer here are:

- How can you build on strengths? For example, where in the public sector do you have repositories of high-quality content?
- How can you avoid duplication? So how can you avoid a situation where each public agency builds its own tool to summarise consultation responses?
- How can we encourage public agencies to think beyond what they currently do? For example, for the first time ever, we can use generative AI to create feedback loops between the public and government. This has the potential to transform governmentcitizen interactions. How can we get these ideas going?



2. Conferring Rights to Individual Citizens

The second area that I want to put on the table today corresponds to another massive chunk of government work, which involves conferring rights to individual citizens, such as rights to residency or citizenship, a license or passport, or entitlement to a benefit or service.

The dream for large parts of the public sector is to automate entire processes. In general, we see a lot of public backlash and poor design, development, and deployment practices in precisely these types of projects, where the aim is to automate everything.

What we advise government to do is to think of any administrative process as a string of microtransactions. If you apply for a passport, for example, there is a string of micro-transactions that need to be made to get one. These are small things like validating a photo or cross-checking text data from the application.

We did a study and found that central government alone completes about one billion of these microtransactions per year. And out of these one billion micro-transactions, 120 million have a very high automation potential with AI. These are things like checking photos, extracting relevant bits of text, and so on. This is where a lot of the potential for AI technologies in government currently is.

The key question here is how do we get officials excited about automating the minutiae of government's bureaucratic practices? The problem in this area of work is that this is not the stuff that dreams are made of. It won't make front-page news – but this is where AI – at the moment – can lead to substantial efficiency gains.



3. Making Policy Decisions

This brings me to the final area of government work where AI can be a game-changer – and that is improving policy decisions.

In the UK, more than 52,000 policy decisions underpin how over a trillion pounds of public money gets allocated each year. Improving the way in which policy decisions are made can lead to substantial benefits, and data science and AI can do that in ways in which traditional statistical techniques cannot do.

We have a tremendous opportunity to adapt our economic models for an age when every individual and company generate massive amounts of data on a daily basis and for a time when the technologies that we have are starting to allow us to develop better answers to the questions that policy-makers have been asking for years. What happens if a big company closes? Can the excess capacity be absorbed by the local economy? Where are the skill shortages, and who are the workers that can most easily be retrained to fill them?

Conclusion

I have given you today three areas of opportunity corresponding to three massive government tasks:

- 1. Communicating with the public.
- 2. Conferring rights to individual citizens.
- 3. Making policy decisions.

If efficiency gains are the ultimate goal, these three areas are ripe with potential.



Summary of Dr Cosmina Dorobantu's Key Points

Dr Cosmina Dorobantu shared insights on the potential of AI to improve efficiency in the public sector, drawing from her extensive experience in data-driven technologies. Her evidence focused on three key areas where AI can significantly contribute to the government's work.

Communicating with the Public:

- Al, particularly generative models like ChatGPT, has the potential to transform how the government interacts with citizens. From responding to queries to summarising consultation responses, Al can enhance efficiency.
- However, the key challenges are identifying high-quality content repositories, avoiding duplication, and innovating beyond current practices by creating feedback loops between the public and the government.

Conferring Rights to Individual Citizens:

- Al can improve bureaucratic processes by automating microtransactions involved in tasks such as passport applications or benefit entitlements. A study revealed that a significant portion of these transactions is highly automatable with Al.
- The challenge for government officials is to embrace automating mundane tasks, which might not be newsworthy but can deliver substantial efficiency gains.

Making Policy Decisions:

- Al and data science can enhance policy-making by providing better insights into economic models and societal trends.
- By leveraging vast amounts of data generated daily by individuals and companies, AI can help answer long-standing questions, such as how local economies absorb the closure of large companies or where skill shortages and retraining opportunities lie.

Action Points for Government and Policymakers

1. Enhance Communication with Citizens:

- Identify areas within the public sector where AI can improve communication and streamline processes.
- Encourage public agencies to think beyond current methods and create innovative ways to interact with citizens, such as feedback loops enabled by Al.
- 2. Automate Administrative Processes:
 - Focus on automating the micro-transactions in bureaucratic processes (e.g., passport applications, benefit entitlement checks).
 - Explore the automation potential within existing government functions, aiming for efficiency gains, even in areas that may seem mundane.
- 3. Leverage Al for Policy Decisions:
 - Invest in AI tools that help in improving economic models and policy decisions.
 - Use AI to answer key policy questions, such as the effects of company closures on local economies and how to address skill shortages and retraining needs.

In conclusion, AI presents significant opportunities for enhancing efficiency in government tasks related to communication, conferring rights, and policy-making. For policymakers, embracing AI can lead to better service delivery, improved decision-making, and more streamlined bureaucratic operations.



David Fearne

Global Director of Generative AI, Cognizant

Introduction:

Good evening. My name is David Fearne, and I lead Generative AI at Cognizant, focusing on both client engagements and cutting-edge applied research.

In regulated industries, transparency, fairness, and ethics are non-negotiable. The UK's "Ethics, Transparency and Accountability Framework for Automated Decision-Making," published in November 2023, this provides foundational guidance for the ethical deployment of AI systems in the public sector. However, this framework, lacks detailed and practical advice on effectively embracing AI. Similarly, the Centre for Data Ethics and Innovation's "Review into Bias in Algorithmic Decision-Making," conducted in 2020, highlights the need for updated strategies to address bias in AI systems.

Tonight, I want to share some evidence of our applied research from our work across multiple sectors, demonstrating how we've delivered high-impact outcomes under intense scrutiny. My goal is to provide actionable recommendations to enhance existing frameworks, to better helping welfarecentric government departments embrace AI more confidently.

I will explore how AI can be responsibly integrated into public services by focusing on our work in three critical areas: transparency, bias, and ethics.

Transparency:

Transparency is the foundation of trust in AI. Without it, citizens and policymakers cannot trust AI-driven decisions. But true transparency isn't just about outcomes—it's about the process behind them.

Our research shows that AI systems can appear accurate in testing but fail in real-world challenges if their reasoning is flawed. By focusing on the steps in the decision-making process, we can surpass today's transparency standards with auditable workflows and robust decision documentation.

For example, we worked with a leading UK mutual pension provider to address their challenge of auditing client-facing materials to comply with the FCA's Consumer Duty regulation.

This was a complex, multi-week process involving 100's of emails and 10's people. Using Generative AI, we distilled the expertise of SMEs into clear, actionable instructions for a large language model, automating the process and reducing it to just three minutes.

Crucially, the AI workflow explained and documented every decision point, providing unparalleled traceability. The result? Faster processes, but more importantly surpass previous human standards for accountability and trust leading to superior transparency.

Bias:

Al models encode patterns and beliefs from their underlying training data. As we demand more intelligent Al models, a consequence is the consumption of larger volumes of inevitably biased data.

The risk is that these biases can unintentionally exclude or disadvantage certain groups.

But bias is not always a bad thing. Few people realize there are over 180 cognitive biases, some of which can be harnessed to our advantage. In our research, we've developed and deployed two practical techniques to manage bias through AI instructions:

- Bias Over-Indexing: This involves instructing an AI to maximize for a certain heuristic, intentionally biasing the LLM toward a specific belief. This is effective when seeking novel solutions and avoiding biases like affinity or status quo bias.
- Contradictory Bias Meta-Prompting: This technique pairs opposing biases to balance decision-making, leveraging Al's statistical imperative to find the most balanced approach while avoiding emotional or unintended human biases.



Examples include:

- Confirmation Bias vs. Devil's Advocate: Encourages the AI to challenge its own assumptions.
- In-Group Bias vs. Out-Group Bias: Promotes inclusivity and diverse perspectives.

By balancing competing heuristics, we create fairer, more balanced outcomes—essential when deploying AI in public services that impact millions.

Ethics:

Finally, ethics. AI must reflect our values—it cannot simply automate processes. Fairness, accountability, and privacy must be embedded into every system.

Ethics in AI requires actively designing systems to promote societal good. This isn't a set of rules programmed into the AI but a design approach. It means engaging diverse stakeholders, including policymakers, domain experts, and affected communities, to define goals, set standards, and evaluate solutions early.

We applied this approach when developing an AI medical triaging service for the UAE's primary healthcare system.

We ensured the system—from its methods of interaction to the advice it provided—didn't exclude underrepresented groups. For instance, users requested that the system had an option to operate anonymously, avoiding personal details like name, age, or medical history, which were seen as indicators of social status and potential inequality. Incorporating these insights resulted in a system that was more inclusive and aligned with the community's needs.

Recommendations to the Committee:

To sum up:

- Transparency: Set a standard of surpassing human expectations of transparency in AI reasoning while enhancing the standard of broad accessibility.
- Bias: Develop and apply updated frameworks to give confidence in how to effectively manage bias in AI systems at scale.
- Ethics: Establish best practices for designing AI systems that actively promote societal good through collaboration, innovation, and inclusivity.

The public sector has a unique opportunity to lead in responsible AI adoption, but the devil to adoption at scale is in the detail.

Thank you.



Summary of David Fearne's Key Points

David Fearne leads Generative AI at Cognizant and emphasises the need for transparency, fairness, and ethics in AI systems within regulated industries. He highlights the UK's "Ethics, Transparency and Accountability Framework for Automated Decision-Making" (2023), noting its lack of practical advice for effective AI adoption. He aims to provide recommendations to help government departments adopt AI confidently, focusing on transparency, bias, and ethics.

Transparency:

- Key Point: Transparency is essential to building trust in Al. It's not just about outcomes but the decision-making process behind Al systems.
- Example: Fearne's team used Generative AI to automate a complex auditing process for a UK pension provider, reducing a multi-week process to three minutes. The AI workflow explained and documented every decision point, ensuring transparency and accountability.

Bias:

- Key Point: Al models can inherit biases from their training data, which can unintentionally disadvantage certain groups. However, some biases can be strategically used for better outcomes.
- Techniques:
 - a. Bias Over-Indexing: Directing AI to focus on a specific belief or heuristic to encourage novel solutions and avoid biases like affinity or status quo bias.
 - b.Contradictory Bias Meta-Prompting: Pairing opposing biases to achieve balanced decision-making.
- Examples: Using AI to counter confirmation bias or promote inclusivity by balancing in-group and out-group biases.

Ethics:

- Key Point: Al systems must align with societal values, incorporating fairness, accountability, and privacy. Ethics in Al is not just about rules but the design process, which should involve diverse stakeholders.
- Example: In the UAE, Fearne's team developed an AI medical triaging system that prioritised inclusivity by allowing anonymous interactions to avoid biases associated with personal information, ensuring the system met community needs.

Recommendations for Government and Policymakers:

- 1. Transparency: Set standards to exceed human expectations for transparency in AI reasoning and enhance accessibility.
- 2. Bias: Develop frameworks that provide confidence in managing AI bias at scale.
- 3. Ethics: Establish best practices for Al design that prioritise societal good through collaboration, innovation, and inclusivity.

Conclusion:

The public sector has a unique opportunity to lead in responsible Al adoption, but scaling Al responsibly requires attention to detail in these areas.



Hugh Eaton

Colonel HRA Eaton OBE, MA, VR, UK Ministry of Defence

Introduction:

I am a special advisor to Boston Consulting Group and serve as a Colonel and advisor to the UK Ministry of Defence. I used to run the global government business at Microsoft, so I'm hoping we are going to have a robust conversation about big tech because the government needs to have that conversation. Before that, I ran the public sector at Cisco for Europe and the Middle East.

Practical Approach

I'm going to try and make this as practical as I can. I think that Prime Minister Keir Starmer's speech on AI really demonstrates a level of ambition that I admire, and I think there are some pieces that we need to address in order to achieve faster, more effective implementation of AI. I think the Matt Clifford report is a great start, and if people in the room who are interested would read one report, it's his, about AI opportunities in the UK Government.

- "AI Opportunities Action Plan: Recommendations for the government to capture the opportunities of AI to enhance growth and productivity and create tangible benefits for UK citizens".
- <u>https://www.gov.uk/government/publications/ai-opportunities-action-plan</u>

(i) Recommendations from the Matt Clifford Report

I don't know Matt Clifford, but it's well written, it's intelligent, and it's practical. However, it has 50 recommendations, which is a lot, and the Government has accepted them all, and then awarded 44 of them to the Department for Science, Innovation and Technology, one of the smallest Ministries in the UK Government with one of ethe smallest budget. I find that difficult to understand how that's going to work?

(ii) Distribution of Effort

Eighty-eight percent of the recommendations going to one of the smallest departments in the Government, I think that's going to need somebody to have a look at how that's actually going to happen.

I'd like to recommend a book to any of you here, written by Tony Blair. He's written a book on leadership where he states that the essence of strong, capable and efficient government is about making the centre strong.

The idea of assigning someone the task of conducting thorough analysis around the application of AI and then parcelling that out to a department that seems not entirely equipped to deliver it raises concerns about how effectively that delivers on that key proposition of keeping the centre strong?

So that's one of my recommendations; We must have a look at the distribution of effort.



Ethical Considerations

Einstein talked in 1932 about how he entreated humanity to try and ensure that we 'had the skills and governance to keep pace with the technology that we created', but I'm not sure how well we're doing that.

I admire Keir Starmer's ambition; however, I think we've got a way to go in terms of implementation, particularly on the ethics side. We have a clear responsibility to be very careful about who we apply kinetic energy to, and we need to be clear-eyed about non-kinetic effects.

(i) Human Oversight

I was speaking earlier with Bishop Steven Croft about the ethical conundrum of wanting to keep a human in the loop so that the process between the sensor, the decider and the effector—so for instance, a drone, the headquarters, and a missile system—was informed by a human being. It's what we would all wish to have, so that there isn't hallucination in the AI system or some technical glitch.

(ii) Adversarial Concerns

And that is noble and good - until we come up against an adversary who is less queasy about ethical things, and he or she is operating at machine speed, and we start to have lots of dead soldiers coming back to RAF Lyneham – and so what do we do then? Because our ethical case, it seems to me—and Bishop Steven Croft and others are here who know much more about this than I do—that some principles need to be revisited. And so in Defence we have some complications there that may not exist across the rest of the public sector.

This argument could imply that amidst the complexities and rapid advancements in AI and technology, the foundational ethical principles may not be receiving the attention and clarity they require, leading to potential misalignments in practice and understanding. We may need to introduce the notion of a range of ethical and legal positions.



National and International Partnerships

The second thing I would invite people to do is to look at partnerships, and I'm not just talking about industry partnerships; I'm also talking about international partnerships. We've talked a lot today about prioritisation and focus and partnering with industry. But for our soft power to be backed up by capability, capacity, and the edge of hard power, we need money—lots of it— to invest in this country. We are working on the Strategic Defence Review at the moment; we're hopeful for 2.5%, gusting 3% of GDP to the defence budget. We're probably planning on not getting any more, but where is that money coming from?

NHS is quite rightly making their demands and the Department of Education is making their demands, and so it goes, but we needs to think creatively, and we don't need AI for this! Our race has not yet run as humans, so we can be smart too, and we need to think creatively about what those partnerships might look like and who they might be with.

(i) Examples of International Cooperation

Two data points: Firstly, the relationship that the Ministry of Defence has with the British government, Australia, and the United States in AUKUS.

- Pillar one is to do with nuclear-powered (not nuclear armed) submarines; you all know about that.
- Pillar two is about almost everything else that has any sort of defence connotation. Artificial intelligence is part of that.

Secondly, there's the global combat air programme, where we're partnered for the first time with a non-NATO member, Japan. That is not standard practice, but the belief is that's where we're going to get the best bang for our buck. Well, why don't we partner with a country to help fund our organic capability and capacity in AI?

(iii) Convening Power and Organic Capability

For us to be an effective convening power—perhaps a great convening power—requires organic capability and capacity for credibility and coherence. It's the same way that soft power needs credible hard power for the soft power to be credible; otherwise, the whole thing unravels. There are less friendly nations out there who are looking at us and examining what our organic capability and capacity are.

(iv) Britain's Image Abroad

Here's the thing about being a Brit abroad. It may not look or feel like this if you've got a nice title or fancy uniform, or protected and flattered by governmental protocosl, but if you're a Brit business person abroad, I would argue nobody is interested in what British people have to say in terms of, "Go off and do this," or "Go off and do that." At least that hasn't been my experience.

What our friends overseas are genuinely interested in, is where the UK has done something that is successful and manifestly effective, scalable, transferable, portable, and relevant to the rest of the world - so we need to ensure that we're generating that organic capability.

Final Thoughts:

(i) Piloting - permission to fail Sir!

The final thing I'd push on is around piloting. With the SDR complete we will not need any more reports or reviews for a while. For the members of the new administration and members of both houses to understand how we operate in Defence I would propose the Armed Forces Parliamentary Scheme. In its Defence Academy, we would like to take the opportunity to educate some of the special advisors and some of the Chiefs of Staff in the new administration, so they understand a little bit more—and not just the what, but also the how, because what defence needs are many more pilot projects focussed on adaptation and pace.

(ii) Implementation and Permission to Fail

In order to have a successful pilot project, defence is the same as anywhere else in the world. You need the **authority to operate**. You need **permission to fail**—it doesn't matter if you screw this up—and you need the **mechanism to scale** if it does work.

Conclusion

So those are the areas that I think defence needs to look at. As I say, the offer is there from the Defence Academy to talk to some special advisors and some Chiefs of Staff about what we're doing and what we need.

But we are a whole-of-government enterprise; it's about our national security strategy activity. It's not just for defence alone. Thank you.



Summary of Hugh Eaton's Evidence Statement and Action Points for Government and Policymakers

- Implement Recommendations from the Matt Clifford Report: Prioritise the effective implementation of the recommendations in the Matt Clifford report, ensuring adequate resources and support are allocated.
- Strengthen Central Government Capability: Review the distribution of responsibilities and resources among government departments to enhance the central government's ability to manage major technological initiatives effectively.
- Establish Ethical Frameworks: Develop and enforce ethical guidelines for AI deployment, particularly within defence, ensuring that human oversight is maintained in critical decision-making processes.
- Encourage National and International Partnerships: Actively pursue and fund strategic partnerships with other nations to enhance defence capabilities and share technology, leveraging existing collaborations effectively.
- Promote Organic Capability Growth: Invest in developing a robust domestic capability in AI and defence that is both innovative and effective, showcasing successful UK initiatives to build a positive international reputation.
- Pilot and Test New Technologies: Create environments for piloting Al initiatives within the defence sector, allowing for experimentation with new technologies while fostering a culture of learning from failures.
- Enhance Training for Government Executives: Provide educational opportunities for members of the new administration, focusing on the practical applications of AI in various sectors, including defence.
- Develop a Whole-of-Government Strategy: Formulate a comprehensive national security strategy that integrates Al deployment across government sectors, ensuring coherence and collaboration in efforts.

BIOs of Evidence Givers



Jon Atkin Co-founder and Chief Executive Officer Favom

Jon has an extensive background in IT strategy, AI, and business development across multiple vertical markets. He has led technical project teams comprising several hundred staff and managed accounts across both the Public and Private Sectors internationally. He has been responsible for delivery team annual budgets exceeding £18 million, held individual account P&L responsibility of £320 million, and achieved growth targets of more than 80%.

Jon has led teams in delivering innovative AI solutions in healthcare, predictive techniques, and intuitive data entry and capture technologies using AI and machine learning, with a focus on healthcare and policing.

For the past 25 years, Jon has worked on IT programs within the Public Sector, starting with the NHS National Programme for IT, deploying Electronic Patient Record solutions, and learning the lessons that contributed to its downfall. He then applied this knowledge to new, innovative UX-led Public Sector application projects over the past 15 years, including large-scale Global Digital Exemplar programs, bespoke AI/Machine Learning clinical research applications, and innovative procurement and commercial modelling.

Jon is also a member of the leadership team for various start-up businesses.



Rachel Astall Chief Customer Officer (CCO) Beam

Rachel Astall is an expert in Al-driven welfare services, as the Chief Customer Officer (CCO) at Beam, a social enterprise that works with 100+ government partners across the UK.

Rachel's dedication to social impact began at respected charities like Macmillan, British Heart Foundation, and The Prince's Trust, where she championed meaningful change for those in need. Transitioning to the technology sector, she rose to be Director of Global Partnerships at GoCardless, a leading British fintech, over a period of five years.

At Beam, Rachel has always led with purpose, and she began by growing and developing Beam's frontline services to have maximum impact through the use of technology.

In her current role, Rachel oversees all of Beam's government partnerships, which includes the 70+ Local Authorities using our Al-powered tool, Magic Notes, to drastically save time for social care teams, across England, Scotland and Wales.



Morgan Rees Vice President of Data, Analytics & Al Capgemini

Morgan Rees, Vice President of Data, Analytics & AI at Capgemini, holds two key roles within Capgemini Invent.

First, he leads a team of approximately 280 professionals across Capgemini's full spectrum of Data and AI consulting capabilities, focusing on strategy, innovation, and scaling across industries. Second, he oversees Invent's Public Sector Data and AI portfolio, driving impactful solutions for government and public sector clients. He has 15 years' experience working across central government helping his clients use data and AI to deliver policy outcomes and improve services.





Dr Cosmina Dorobantu

Co-Director, Public Policy Programme, The Alan Turing Institute

Turing OII Fellow Oxford Internet Institute University of Oxford

Dr Cosmina Dorobantu is the Co-Director and Co-Founder of the Public Policy Programme at The Alan Turing Institute, the UK's national institute for data science and artificial intelligence. The programme has advised over 100 public sector organisations, from government departments and international organisations to regulators and local authorities, on how they can harness the power of AI for the public good.

Dr Dorobantu is also a Turing Oll Fellow at the University of Oxford, and a member of the Mayor of London's Data for London Board. Prior to joining the Turing, Dr Dorobantu was an Executive Director and Co-Founder of Aurora Energy Research, which is now the largest power analytics provider in Europe, with over 600 employees and offices in 13 countries. She also spent five years at Google, defining the country's investment and business strategy in Eastern Europe, Middle East, and Africa.

She has a MPhil, obtained with Distinction, and a DPhil in Economics from the University of Oxford.



David Fearne Global Director of Generative Al Cognizant

David heads up Generative AI for Cognizant's global Artificial Intelligence and Analytics (AIA) practice, leading our efforts to deliver value from Generative AI. Pioneering the application of enterprise cognitive architecture with generative AI and currently working on a number of papers to share Cognizant's research and innovation in this space.

He joined Cognizant from PA Consulting in June 2021 as European Head of Applied Innovation within the Microsoft Business Group, building on a distinguished career in data. David has operated in other consultancies, lead data intelligence globally for a fortune 500 company and worked at leading technology vendors.

In 2017 he was recognised as one of the 50 most influential people in data, wining his category with a project to raise awareness of large-scale data analytics and AI through the innovative How Happy is London? project. A technology platform that, at the time, was one of the world's largest real time open source implementations.



Hugh Eaton Colonel HRA Eaton OBE, MA, VR, Special Operations Task Group, UK Ministry of Defence

Hugh Eaton is currently supporting the UK Government's Strategic Defence Review in his role as Non-Exec Adviser to the UK MOD.

As a Reservist in the British Army, Hugh has had long exposure to the rest of the UK Public Sector and previously led the Global Government business for Microsoft. He is also currently a Senior Adviser at Boston Consulting Group, consulting into the UK and Saudi Arabian Governments on Digital Transformation and emerging technology.

Prior to Microsoft he led the Public Sector business for Cisco in Europe and the Middle East, before which he commanded soldiers in the Regular Army on combat missions for 16 years. Hugh Eaton is a leading global expert on digital transformation and cognitive government.

ABOUT APPG AI



ABOUT:

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

Parliamentary APPG AI Members: House of Commons

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

Parliamentary APPG AI Members: House of Lords

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



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 Al), 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.



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**







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



HEADER OF CONTRACT OF CONTRAC

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

All rights reserved © Big Innovation Centre. No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form without prior written permission of the publishers.