July 2023
APPG AI Evidence Meeting

BIG INNOVATION CENTRE

Feedback on the AI White Paper

PARLIAMENTARY BRIEF

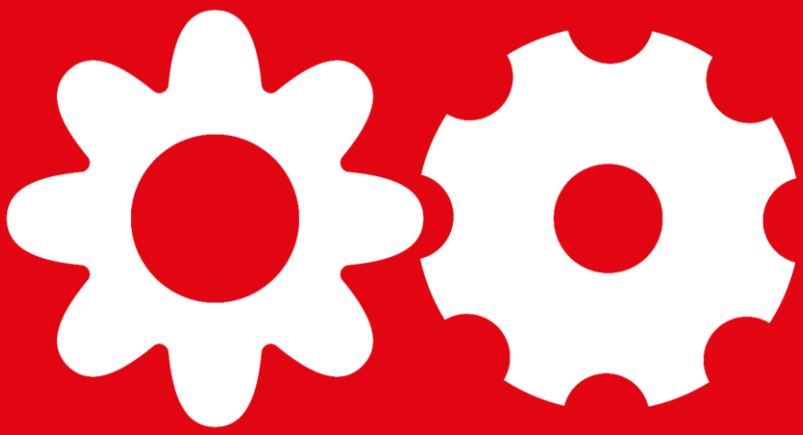

Feedback on the AI White Paper *is a* Parliamentary Brief based upon the All-Party Parliamentary Group on Artificial Intelligence (APPG AI) Evidence Meeting held in House of Commons: Committee Room 10 on the 13th of June 2023.

This APPG AI is co-Chaired by **Stephen Metcalfe MP** and **Lord Clement-Jones CBE**.

We would like to express our appreciation to the following people for their oral evidence:

- **Anthony Walker,** Deputy CEO, **TechUK**
- **Prof. David Leslie,** Director of Ethics & Responsible Innovation Research, **Alan Turing Institute**
- **Sulabh Soral,** Chief AI Officer, **Deloitte**
- **Roger Taylor,** Advisor to Responsible AI Programme, **Accenture**
- **Fran Bennett,** Interim Director, **Ada Lovelace Institute**
- **Daniel Hulme,** CEO, **Satalia**

Big Innovation Centre is the appointed Secretariat for APPG AI

- CEO, **Professor Birgitte Andersen**
- Rapporteur, **George Farrer**

www.biginnovationcentre.com | Email: info@biginnovationcentre.com | @BigInnovCentre

https://bicpavilion.com/about/appg-artificial-intelligence |
Email: appg@biginnovationcentre.com | @APPG_AI

PARLIAMENTARY BRIEF

Feedback on the AI White Paper

All Party Parliamentary Group on
Artificial Intelligence

APPG AI Sponsors

The Group supporters – Ami Limited, British Standards Institution, Brunel University London, CMS Cameron McKenna Nabarro Olswang, Deloitte, Ernst & Young, Innovate UK – UKRI, Osborne Clarke, and Rialto – enable us to raise the ambition of what we can achieve.

Contents

1. Introduction **6**

2. APPG AI Pavilion Survey **8**

3. Recommendations for policymakers **12**

4. Evidence statements **18**

Anthony Walker, TechUK 18

Prof. David Leslie, Alan Turing Institute 21

Sulabh Soral, Deloitte 25

Roger Taylor, Accenture 28

Fran Bennett, Ada Lovelace Institute 31

Daniel Hulme, Satalia 34

5. Institute of Directors Submission **37**

6. Speaker Bios **41**

7. Contact **45**

1. Introduction

“Feedback on the AI White Paper”: Evidence Meeting of the All-Party Parliamentary Group on Artificial Intelligence (APPG AI) held in House of Commons: Committee Room 10 on the 13th of June 2023.

During the APPG AI meeting, the focal point of discussion centred on the UK Government's Artificial Intelligence (AI) White Paper titled **'A Pro-Innovation Approach to AI Innovation'**[1], which was published in March 2023. Expert speakers offered their diverse perspectives on the Government's proposed framework, highlighting its four key components: defining AI based on its unique characteristics, adopting a context-specific approach, providing cross-sectorial principles for regulator responses, and delivering central functions to support coherence. The speakers delved into their insights regarding the practical implementation of the framework, suggested additions, and discussed the anticipated life cycle for its implementation.

AI regulation is an immensely important and pressing issue that demands thorough discussion. As the implementation and deployment of AI technologies proliferate across various sectors, it is crucial to establish comprehensive regulations that ensure responsible and ethical use. Discussions around AI regulation enable the development of frameworks that promote transparency, accountability, and fairness, allowing us to harness the benefits of AI while safeguarding the rights and well-being of individuals and communities.

Main questions:

- *What is included, and what is missing in the UK Government’s AI White Paper?*
- *What is required for more effective and efficient adoption of AI in the UK?*
- *How can the UK become a world leader in AI Governance?*

List of panellists:

- **Anthony Walker,** Deputy CEO, **TechUK**
- **Prof. David Leslie,** Director of Ethics & Responsible Innovation Research, **Alan Turing Institute**
- **Sulabh Soral,** Chief AI Officer, **Deloitte**
- **Roger Taylor,** Advisor to Responsible AI Programme, **Accenture**
- **Fran Bennett,** Interim Director, **Ada Lovelace Institute**
- **Daniel Hulme,** CEO, **Satalia**

[1] UK Government, **‘A Pro-Innovation Approach to AI Regulation’** (March, 2023). https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper

(From L-R: Fran Bennett, Daniel Hulme, Roger Taylor, Prof. Birgitte Andersen, Stephen Metcalfe MP, Antony Walker, Prof. David Leslie, Sulabh Soral)

This meeting was chaired by Co-Chairs **Lord Clement-Jones CBE** and **Stephen Metcalfe MP**.

Parliament has appointed Big Innovation Centre as the **Secretariat of the APPG AI**, led by **Professor Birgitte Andersen (CEO)**. The Project Manager and Rapporteur for this meeting is **George Farrer**.

2. APPG AI Pavilion Survey

Prior to the APPG AI meeting, a survey was issued on the **APPG AI's Pavilion Platform**.

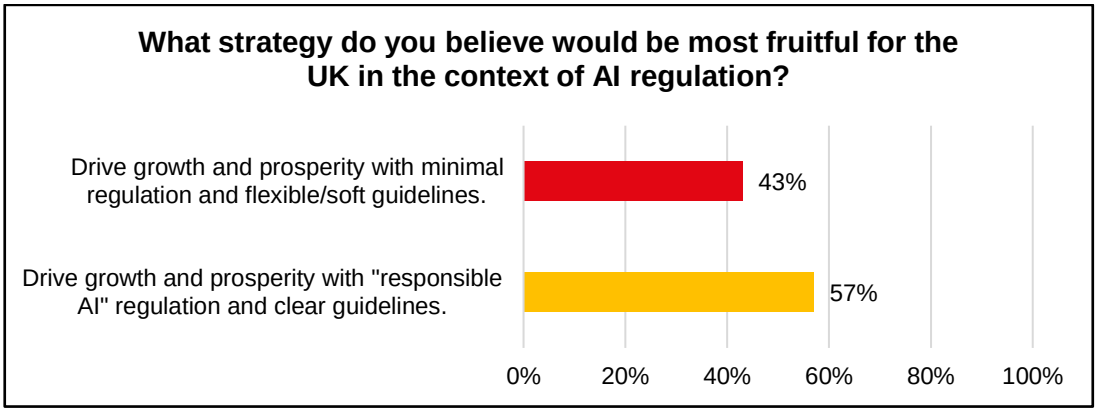

Question 1 asked members of the APPG AI *what regulatory strategy they believe is the most fruitful for the UK*. The majority of respondents, comprising 57%, advocate for **'driving growth and prosperity through "responsible AI" regulation and clear guidelines'**. This indicates a preference for a regulatory approach that encourages the development and adoption of AI technologies while emphasising ethical considerations and responsible practices.

On the other hand, 43% of respondents support a **'strategy that prioritises driving growth and prosperity with minimal regulation and flexible/soft guidelines'**. This viewpoint suggests a desire for a more hands-off regulatory environment that allows for greater flexibility and autonomy in AI development and implementation.

The survey findings highlights that it is essential to strike a balance between fostering innovation and ensuring ethical AI practices when crafting a regulatory framework. This balance is crucial to maximise AI's benefits while addressing potential risks and concerns.

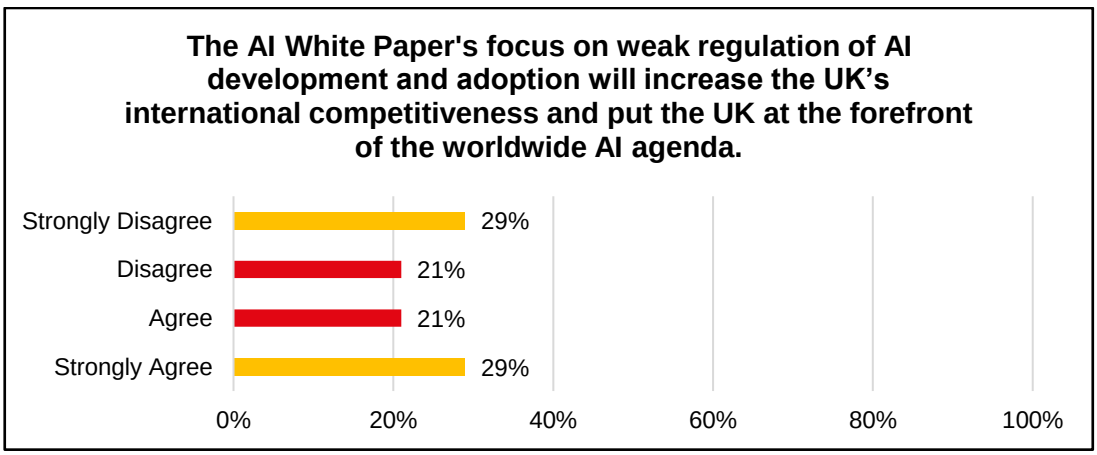

Question 2 proposed the *statement that the AI White Paper's focus on weak regulation of AI development and adoption will increase the UK's international competitiveness and position the country at the forefront of the worldwide AI agenda.*

An equal percentage of respondents, 29%, both **'strongly disagree'** and **'strongly agree'** with the statement, indicating a significant level of conviction in their positions. This suggests that there are individuals who strongly believe that weak regulation would have a positive impact on the UK's international competitiveness and AI prominence, as well as those who strongly disagree and hold the opposite view. Additionally, 21% of respondents both **'disagree'** and **'agree'** with the statement, indicating a moderate level of agreement or disagreement.

These results demonstrate a diversity of opinions among respondents regarding the potential impact of weak regulation on the UK's international competitiveness and its positioning within the global AI landscape.

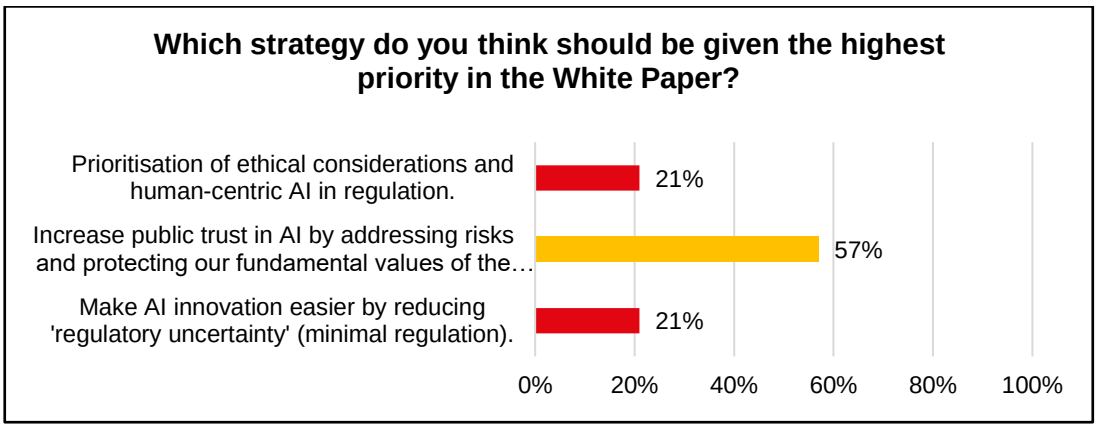

Question 3 asked respondents to *select the strategy they believe should be given the highest priority in the White Paper*. The findings indicate that a majority of respondents, 57%, prioritise **'increasing public trust in AI by addressing risks and safeguarding fundamental values for the safe adoption of AI'**. This suggests a strong emphasis on building trust and ensuring responsible AI practices as a key focus area.

On the other hand, 21% of respondents support two alternative strategies: **'making AI innovation easier by reducing regulatory uncertainty (minimal regulation)'** and **'prioritising ethical considerations and human-centric AI in regulation'**. However, these options garnered a smaller percentage of support compared to the strategy aimed at increasing public trust.

In conclusion, the survey findings suggest that prioritising the establishment of public trust and addressing risks for the safe adoption of AI should be given the highest priority in the White Paper. This underscores the significance of ethical considerations and responsible AI

practices in building a robust and trustworthy AI ecosystem.

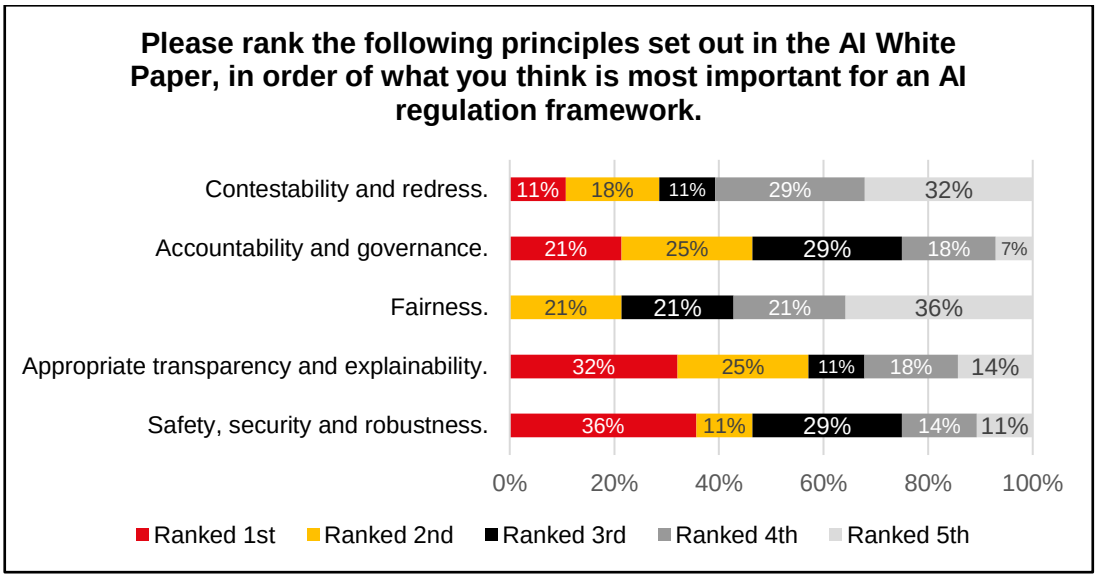

Question 4 asked members of the APPG AI Community to *rank the principles set out in the AI White Paper in order of importance for an AI regulation framework,* which were as follows: safety, security and robustness; appropriate transparency and explainibility; fairness; accountability and governance; contestability and redress.

'Safety, security, and robustness' received the highest number of first-ranked votes at 36%, closely followed by **'appropriate transparency and explainibility'** with 32% of first-place votes. These two principles can be considered the most important for an AI regulation framework. Additionally, **'Safety, security, and robustness'** received 25% of second-placed votes. On the other hand, **'Accountability and governance'** received 21% of first-ranked votes and the fewest fifth-ranked votes, indicating a moderate level of importance among members of the APPG AI Community.

'Contestability and redress' and **'fairness'** received very few first and second-ranked votes but had a higher number of fourth and fifth-ranked votes. This suggests that these principles are considered less important compared to the others.

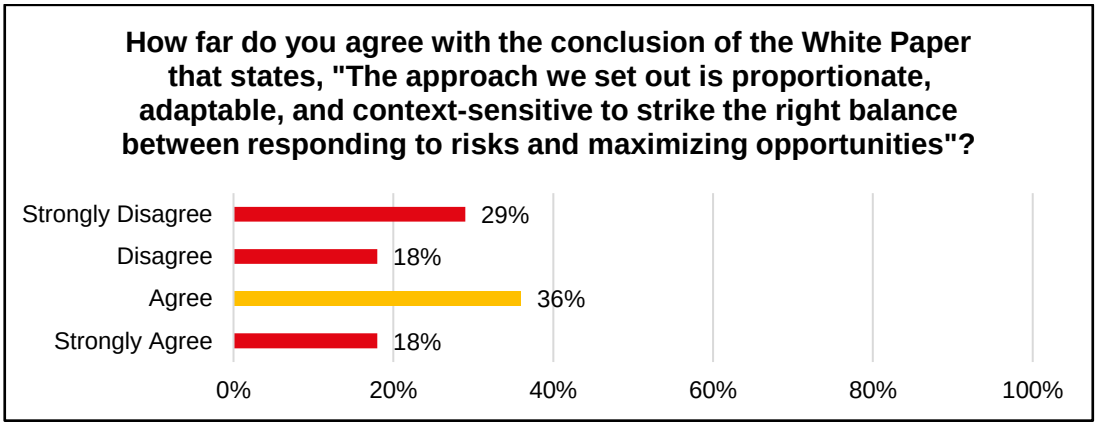

Question 5 questioned members of the APPG AI Community *how far they agree with the conclusion of the AI White Paper.* 29% of respondents **'strongly disagree'** and 18% **'disagree'** with the statement, indicating a significant portion who do not believe that the proposed approach strikes the right balance.

On the other hand, 36% of respondents **'agree'** and 18% **'strongly agree'**, suggesting a notable proportion who support the notion that the approach outlined in the White Paper is proportionate, adaptable, and context sensitive. These findings highlight the presence of divergent perspectives among the survey participants regarding the effectiveness of the proposed approach in responding to risks and maximising opportunities.

Overall, the survey results demonstrate a range of opinions on the conclusion of the White Paper. It is evident that there are varying interpretations of whether the proposed approach successfully achieves the desired balance. These findings underscore the need for further exploration and deliberation on the topic to address the diverse perspectives and ensure the development of a comprehensive and well-rounded AI framework.

3. Recommendations for policymakers

1. Policymakers should implement regulations that **account for the diverse contexts and risks of AI applications**. By differentiating regulatory approaches based on specific sectors and use cases, policymakers can **effectively address challenges**. Additionally, anticipating and mitigating potential societal impacts, such as post-truth environments and surveillance concerns, will promote a positive and inclusive AI ecosystem.
2. Engage in discussions with regulators to **transform the 5 principles outlined in the White Paper into practical guidelines and enforceable laws**. This collaboration will help bridge the gap between the abstract principles listed in the White Paper and their implementation.
3. Recognise that the **increasing industrialisation and commercialisation of AI**, exemplified by technologies like **ChatGPT**[2], **necessitates the implementation of robust regulatory measures**. It is crucial to prioritise the development and enforcement of effective regulations to address potential risks, ensure responsible use, and maintain public trust in AI systems.
4. **Proactively lead efforts to align principles and standards for AI regulation at the international level**. By promoting collaboration and cooperation among countries, the UK can **ensure consistency, avoid fragmentation, and establish a global framework** for responsible AI governance.
5. Prioritise the allocation of adequate funding and resources to **empower AI regulators**. This will enable them to effectively keep pace with the rapidly evolving AI landscape, ensuring they have the necessary tools and capabilities to regulate AI technologies responsibly.
6. Policymakers should **engage academia, civil society, and industry stakeholders in the AI policy debate to develop inclusive and impactful regulations**. By promoting collaboration between these diverse stakeholders, policymakers can leverage their expertise and perspectives to create a comprehensive regulatory framework that addresses practical implementation challenges and aligns with societal needs.

During the discussions, experts expressed a consensus that the principles outlined in the AI White Paper are indeed the right principles to guide AI regulation. However, there was a shared concern that these principles have been treated in a homogenous manner, which overlooks the critical distinctions necessary for effective and certain regulation. It was

[2] **ChatGPT**. https://chat.openai.com/

highlighted that while some of the principles are well defined in law, others lack analogous statutory redress or an existing juridical basis. This flattening of the principles into a simplified framework has inadvertently limited the understanding and complexity of AI. Participants agreed that a more nuanced approach is needed, one that engages regulators in a dialogue to translate these principles into practical guidelines and actionable laws. This agile and context-specific approach would allow sectoral regulators to take the lead, fostering an environment that encourages innovation while ensuring regulatory efficacy. By recognising the multifaceted nature of AI and embracing a more comprehensive framework, regulators can better address the complexities and challenges posed by AI technologies within their respective domains.

Figure 1: Principles in the AI White Paper, to guide and inform the responsible development and use of AI in all sectors of the economy.

Furthermore, the discussions highlighted the need for a more global perspective in aligning principles and standards for AI regulation. With the rapid commercialisation of foundational AI models, such as **ChatGPT**, in recent months, the focus has shifted from whether to regulate AI to how to regulate it effectively, both within the UK and internationally. To maintain the UK's position as a thought leader in AI policy and governance, it is vital to strengthen the White Paper and address the growing demand for statutory assurance of safe and responsible AI at the national and international levels. This entails achieving convergence with the international community, particularly in customer-centric aspects, as businesses operate across various jurisdictions. Regulation must be interoperable and coherent throughout the country, ensuring alignment with international frameworks for consistent and effective AI governance.

Antony Walker, Deputy CEO of **TechUK**[3], asserts that the White Paper aims to outline an agile

[3] **TechUK**. https://www.techuk.org/

and responsive approach to AI, a fundamentally dynamic technology. In his viewpoint, the necessity of being adaptable and swift holds great significance. Walker expresses his belief that the White Paper serves as a positive starting point, emphasising the need for various actors and regulators to address the associated risks. He further supports the notion of a federated approach to AI regulation, considering it to be the correct path forward.

In addition, Walker elaborates on the metaphor of AI regulation, contemplating whether it resembles a marathon or a sprint, and concludes that it is indeed a marathon. Nonetheless, he highlights that the world's top marathon runners exhibit exceptional speed, necessitating us to sprint in order to keep pace with them. Reflecting on this notion, Walker asserts that for the Government to showcase the viability of their approach, they must forge ahead with great urgency and rapid progress.

Walker, in addition, advises the Government to articulate the array of challenges posed by AI. These challenges encompass short-term and immediate concerns, as well as longer-term implications and even hypothetical yet significant scenarios. Walker urges the Government to delineate these challenges more explicitly, enabling a clearer understanding of their nature and scope. Lastly, he emphasises the need for the Government to allocate funding and resources to support the regulators engaged in this critical work. Furthermore, he requests that the principles outlined in the White Paper be translated into practical implementation, ensuring effective governance of AI.

Professor David Leslie, Director of Ethics and Responsible Innovation Research at the **Alan Turing Institute**[4], begins by discussing the international dimension, highlighting the significant discussions held at the recent G7 summit[5] regarding the importance of achieving interoperability among international standards and regulatory approaches for AI. He emphasises the need to avoid the siloing of nation states, as such isolation would worsen geopolitical challenges and restrict the vast array of public interest and commercial opportunities available. Professor Leslie concurs with the view that globally aligned principles and standards are essential to guide statutory interventions in the AI domain.

Furthermore, in reference to the White Paper, Prof. Leslie expresses his desire for a more nuanced approach that encompasses the pressing need for targeted and well-informed statutory interventions. He emphasises the importance of directly addressing the imbalanced power dynamics within the AI ecosystem of the UK. Regarding the principles outlined in the White Paper, Professor Leslie argues that treating them uniformly diminishes the crucial distinctions necessary for regulatory effectiveness and clarity.

Prof. Leslie contends that the White Paper exhibits a lack of nuance and differentiation in its treatment of the private sector, particularly in relation to AI. He highlights the absence of clear distinctions between large tech corporations, which wield substantial infrastructural power within the AI innovation ecosystem, and SMEs, who often find themselves subject to the

[4] **Alan Turing Institute**. https://www.turing.ac.uk/

[5] The 2023 G7 Summit was held in Hiroshima, Japan, from the 19th to the 21st of May 2023.

influence and service provisions of these major players. Instead, the White Paper tends to portray the private sector as a homogeneous entity. According to Prof. Leslie, this approach, characterised by a power imbalance, undermines the nuanced considerations required for fostering innovation.

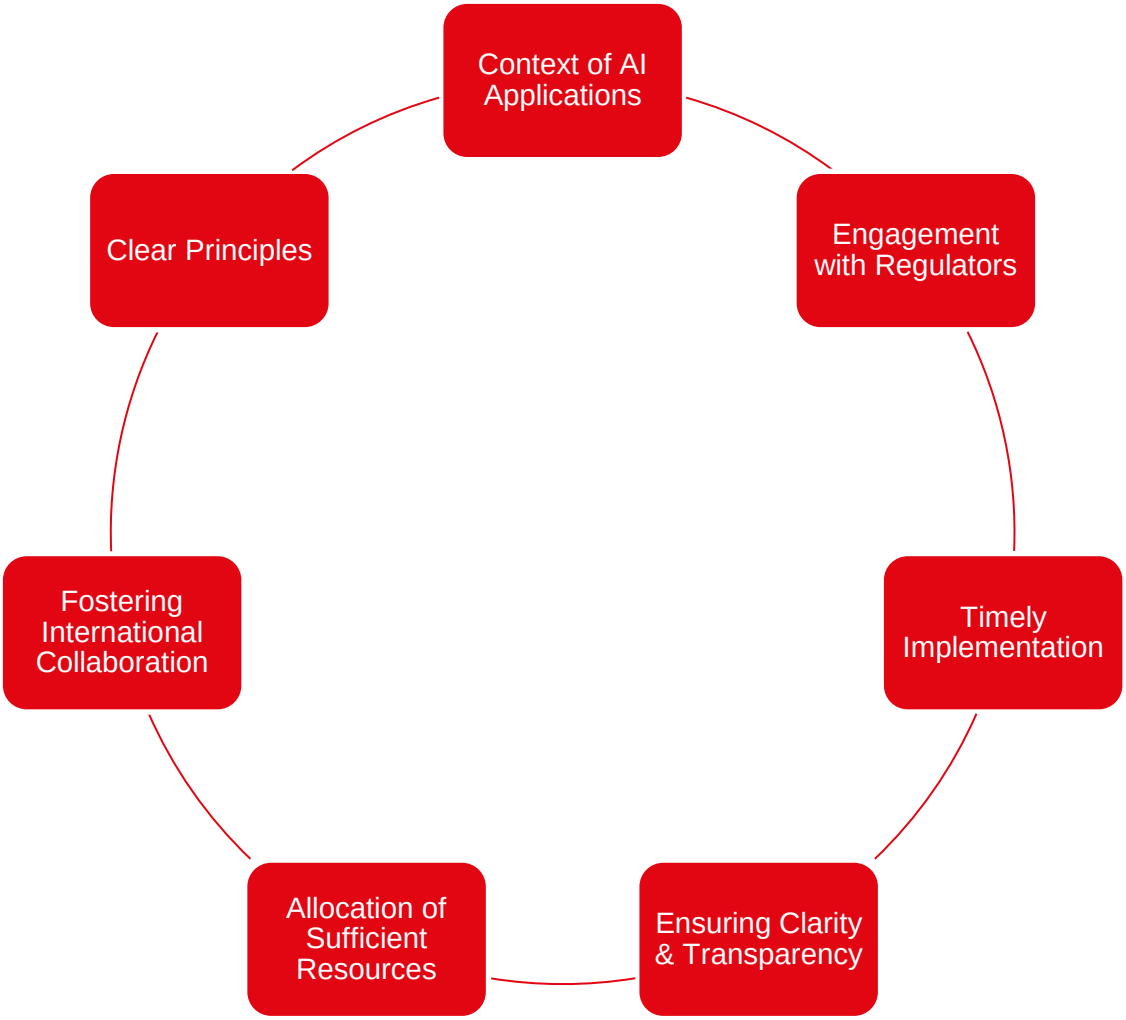

Figure 2: AI Regulation: What to Consider

Sulabh Soral, Chief AI Officer at **Deloitte**[6], commences by expressing his conviction that the pro-innovation approach outlined in the White Paper is indeed the right path to follow. He clarifies that the White Paper's focus lies in mitigating the potential negative outcomes that AI may yield, rather than seeking to regulate the technology itself. Soral appreciates this approach, as it does not aim to exert control over the technology per se, but rather considers the specific use-cases and the contextual framework within which AI, particularly generative AI, is predominantly employed.

[6] **Deloitte**. https://www.deloitte.com/global/en.html

Continuing the discussion, Soral elaborates on the suitability of the UK for adopting a pro-innovation approach. He asserts that a pro-innovation approach requires a robust regulatory framework, supported by strong regulatory bodies and a foundation of rule of law, all of which the UK possesses. Consequently, Soral firmly believes that the UK possesses the necessary infrastructure to facilitate a context-driven approach to AI regulation.

In conclusion, Soral emphasises the crucial need for convergence within the international community on AI matters. Given the customer-centric nature of AI, any actions or developments with potential impacts on customers necessitate international convergence. With numerous companies conducting business in Europe, China, and beyond, the effects of AI extend globally. Therefore, international alignment is vital to ensure harmonisation and consistency in addressing these customer-centric aspects of AI.

Roger Taylor, Advisor to the Responsible AI Programme at **Accenture**[7], begins by expressing his concurrence with the White Paper's emphasis on growth and its outlined objectives. However, he puts forth the assertion that the proposed approach will ultimately prove ineffective. While Taylor agrees with the principles and acknowledges the necessity for a pro-innovation and agile approach, as opposed to hasty legislation, he contends that the White Paper, as it stands, will not achieve the desired outcomes.

Furthermore, Taylor highlights the fundamental reason why regulations are effective: they provide clarity regarding where individuals and entities stand. If the delineation is unclear, it becomes a futile endeavour and only serves to increase friction within the economy. Taylor poses a critical question regarding the practical implications of the principles outlined in the White Paper. He asserts that their ambiguity and lack of clarity will inevitably lead to increased friction and hinder their practical implementation.

Concluding his evidence, Taylor emphasises the crucial importance of providing regulators with the necessary resources to carry out their responsibilities effectively. Failure to regulate properly would burden the industry with the onus of navigating unclear regulations, resulting in an immense workload that ultimately yields limited progress. Such a scenario would introduce a degree of uncertainty, thereby allowing less responsible organisations to thrive and leading to detrimental outcomes. It is therefore essential to empower regulators adequately to ensure clarity, foster responsible practices, and achieve desirable outcomes.

Fran Bennett, Interim Director of the **Ada Lovelace Institute**[8], begins her evidence by outlining the three tests that the institute believes a regulatory framework should successfully meet. The first test pertains to the coverage of the framework. Bennett explains that the principles outlined in the White Paper describe the desired outcomes that regulators should strive for when AI is employed within their respective domains. However, she notes that there are instances, such as AI usage in job interviews, where there is no specific regulator assigned to ensure the safety and fairness of such systems. Consequently, individuals may lack

[7] **Accenture**. https://www.accenture.com/us-en

[8] **Ada Lovelace Institute**. https://www.adalovelaceinstitute.org/

avenues for redress in case of system failure. While equalities law may provide some recourse, it is applicable across the entire economy rather than being domain specific.

The second test concerns the capability of the regulators themselves. Bennett emphasises the necessity for regulators to possess the necessary authority to mandate companies to disclose the inner workings of their AI systems. As foundational models gain wider adoption, there is a growing need for enhanced governance at earlier stages of the AI development value chain. Additionally, mechanisms must be in place to effectively manage the distribution of liability among the various actors involved. Bennett underscores the importance of ensuring that regulators are adequately equipped with the required resources to fulfil their responsibilities effectively.

Bennett's final test examines the criticality of implementing regulation in a timely manner. Given the pressing risks that arise with the advancement of AI, particularly concerning the increasing capabilities of general-purpose AI, it is imperative that the initial iteration of the AI framework be expedited. Bennett stresses that without timely intervention from the government, the unregulated proliferation of AI systems will only compound the challenges of addressing them in the future. Acting promptly is crucial to mitigate potential risks and prevent the regulatory challenge from becoming more complex and difficult to resolve.

Daniel Hulme, CEO of **Satalia**[9], initiates his remarks by expressing his encouragement that the AI White Paper takes a different approach by focusing on the applications of AI rather than solely on technologies and definitions. However, Hulme highlights a key deficiency in the White Paper, namely the absence of a comprehensive categorisation and framework for AI applications. Hulme proceeds to discuss several notable applications, including task automation, content generation, complex decision-making, and insight extraction. These applications represent crucial domains where AI can have a significant impact.

Hulme proceeds to shift his focus to the five principles outlined in the White Paper. He asserts that by condensing the complexity of deploying AI into a mere five principles, the intricacies of the issue are oversimplified, rendering the matter nebulous. Hulme expresses concern that this flattening of principles will hinder the possibility of engaging in substantive discussions about AI. Without a nuanced understanding and a comprehensive framework, meaningful debates about AI will become increasingly challenging to facilitate.

[9] **Satalia**. https://www.satalia.com/

4. Evidence statements

Anthony Walker, Deputy CEO, TechUK

Introductory Remarks

I'm Anthony Walker, I'm Deputy CEO of **TechUK**. There has been a huge amount of activity and debate and discussion, and some of it rather animated around AI over the last couple of months. At the same time of course the UK Government's AI White Paper has been circulating for deliberation, and the consultation closes next week (Wednesday 21st June)[10] on the White Paper itself.

AI White Paper

First of all, timing is everything in politics. If anything, I think from a Government perspective the White Paper perhaps arrived a little bit either too late or too early. It has certainly arrived at a slightly difficult time as the impact and potential of generative AI was becoming really apparent around the world with **ChatGPT**.

I think that has led to a certain shakiness in terms of confidence around the White Paper. However, as we've looked at the White Paper from a perspective of a broad swathe of the technology companies, in terms of what the White Paper is fundamentally trying to do which I

[10] **AI regulation: a pro-innovation approach – policy proposals.**
https://www.gov.uk/government/consultations/ai-regulation-a-pro-innovation-approach-policy-proposals

think is to try to set out a quite agile and responsive approach to what is likely to be a highly dynamic technology. Also, a general-purpose technology that's going to be very widely dispersed in terms of usage. I think the requirement to be to be agile and fleet of foot I think remains quite important. Our sense is in terms of the starting point of the White Paper we think that's quite positive. We think that it is going to be vital that many aspects of concern and many of the risks are going to have to be addressed by a wide variety of actors and a wide variety regulators, and this idea of this more kind of federated approach to the regulation of AI is certainly right.

I think the question that everybody's coming back to is, isn't there a role for more fundamental harder regulation at the heart of this new emerging AI ecosystem, in particular that can address some of the more substantive challenges that come out of these foundation models. I think that's where the White Paper is behind the debate already. In that respect I think in some ways this White Paper looks a little greener than I think maybe government might have hoped at the outset, and it seems to me it's something that is going to require quite a lot of iteration. This comes back to my original point which is I think to get policy right in this area we are going to have to iterate, and we are going to have to develop as we go.

One thing I would stress is the fact that it's not like we've just started this work in the last six months. There's been an enormous amount of work done around AI ethics, AI principles, work around thinking about how you turn those principles into action within organisations, and so that's work that **TechUK** have been involved in for a long time. We've been running our AI ethics summit for six years now, where we've really brought together a diverse group of people to talk about those issues.

AI Regulation

There's a huge amount that we can build on and I think just we were discussing the metaphor of you know is AI regulation a marathon or a sprint? I think for those of us who've been working on these issues for quite a long time I think we say well we think it's more like a marathon even if at the moment the media is making it sound like it needs to be a sprint. One thing to think about that metaphor is the world's best marathon runners actually run really quickly and for most of us we've got to sprint to keep up with them anyway. I think there's a real need for us, generally, to accelerate the discussion. If government really wants to and needs to demonstrate that this approach can work, therefore I think we need to kind of push forward quite quickly.

How can the Government build confidence in their approach?

I think there are there are five things that we think government could do very quickly just to help build confidence that this approach can work overall.

First thing is just segmenting the debate. We've got some incredibly confused debates today that mix up very immediate concerns around about the use of deep fakes or the challenges that we have in terms of bias and data sets, and so on. These long-term, hypothetical

existential risks, all gets mixed up. The first thing government could do is actually help to set out the range of challenges: some short-term immediate challenges, some longer-term medium-term challenges and then longer-term, potentially hypothetical, but incredibly serious challenges. Maybe we should segment those out a little bit more clearly. Therefore, let's be more explicit about the risks, but more explicit that they are these risks are over different timescales and they need to be thought about differently.

Secondly, Government could get on and start funding and resourcing the entities that need to do this work. They need to release funding to make sure that our regulators have got the resources they need. Release some funding to start building this central function within Government, who knows what that might turn into, but it needs some resource it needs to go on and do its work now.

Third, do more to define some of the roles of the key bodies: **AI Council**[11], **Centre for Data Ethics and Innovation**[12], the **Foundation Model Taskforce**[13]. Let's get on and work out what those are.

Fourth, there's a real job when it comes to talking to regulators about how you turn some of these principles in the White Paper into practice. Some of these principles are quite well defined in law, others are not, so I think there's a range of questions there.

Finally, the international engagement piece is really important. I absolutely welcome the initiative by the Prime Minister to hold a summit in the UK[14] later this year, I think we need to lead the debate. However, I would stress that one of the reasons why the UK should be leading the debate is because of the strength of knowledge and expertise that we have in academia, civil society, and in many of the other kind of institutes and organisations that we have here. It's vital that those bodies are also involved in that international summit. If it's just governments talking to industry, then you won't have the expertise in the room that you need to make real progress.

[11] **AI Council.** https://www.gov.uk/government/groups/ai-council

[12] **Centre for Data Ethics & Innovation.** https://www.gov.uk/government/organisations/centre-for-data-ethics-and-innovation

[13] **Foundation Model Taskforce.** https://www.gov.uk/government/news/initial-100-million-for-expert-taskforce-to-help-uk-build-and-adopt-next-generation-of-safe-ai

[14] **UK to host first global summit on Artificial Intelligence.**
https://www.gov.uk/government/news/uk-to-host-first-global-summit-on-artificial-intelligence

Prof. David Leslie, Director of Ethics and Responsible Innovation Research, Alan Turing Institute

International Perspective

In the G7 that's just happened there was a lot of talk about the need for interoperability among international standards and regulatory approaches on AI. We are at an inflection point when we think of the ways in which the inward-looking approaches to standards and regulation of these kinds of technologies are just insufficient to the scale and scope of the impacts of the technology. What we don't want is siloing among nation states which would create not only further troubles geopolitically, but it would in a sense curtail the tremendous set of public interest and commercial opportunities, that would benefit humanity. We do at this point I think need to think in a more global way about how we can have at least aligned principles, and at least aligned standards behind the way statutory interventions are inevitably going to occur in this space.

AI Regulation

To this metaphor about marathon running and sprinting (referencing Antony Walker), I think it's a nice point in so far as what we really need to do is to close the distance between what has been a sprint on the side of a breakneck industrialisation of very impactful AI technologies. We see this now with the rush to market of the foundation models that has occurred in the last six months amidst many issues that that should have been anticipated amongst those who have built those systems. There is a move fast and break things sprint on the innovation side, and a more marathon runner attitude that you see in standards development organisations, regulators, and in policy making where there's a tendency towards airing on the side of inertia

because social norms just change at a different pace, than innovation itself changes.

I think the challenge that we face is closing the distance right between the sprint and the marathon, if we don't close that distance we'll be subject to what the sociologist Alvin Tofler called 'future shock'[15] where you know we're in a sense victimised by the premature arrival of the future, because we simply don't have the normative, regulatory and legal categories to put controls around technology, that has in a sense gotten out of our hands, or at least is being developed in ways that we would otherwise want to control.

The UK Public's View on AI

The explosion of attention to issues of responsible and safe AI that have been triggered by this breakneck commercialisation of foundation models like **ChatGPT** over the last several months has meant that we now are moving from a period of asking the question, 'should we regulate AI in the UK', to the period of 'how do we effectively regulate AI both within the UK and globally'?

That shift that I just talked about has already been signaled by the public. Last week, the Alan Turing Institute and the Ada Lovelace Institute published our national survey[16] on public attitudes towards AI and we found that members of the public were strongly supportive of a new AI-centered regulatory regime. In fact, when asked what would make them more comfortable with the use of AI, 62% of our respondents chose laws and regulations that prohibit certain uses of AI, and that guide the use of all AI technologies. Two thirds of the public asking for regulatory intervention is significant. It's with this in mind that I would say we must consider how to appropriately strengthen the regulatory White Paper so that the UK thought leadership in AI policy and governance can sufficiently keep pace with rising national and international demands for statutory assurance of safe and responsible AI.

The AI White Paper

I'm going to explore a couple of instances in which the White Paper lacks the nuance that could lead to regulatory narrowness and ineffectiveness that would actually be at cross purposes with a sufficient regulatory robustness. I'll stress the need for a more differentiated path forward, one that addresses the urgency for targeted and well-informed statutory intervention, and that confronts head-on imbalanced power dynamics within the UK's AI ecosystem.

The White Paper framework presents five cross-cutting principles intended to inform the UK's regulatory regime, and to be implemented by existing regulators. These are:

[15] Tofler's future shock is the feeling of being overwhelmed by rapid technological advancements and disruptive changes, necessitating adaptive strategies for success in an uncertain future.

[16] Alan Turing Institute & Ada Lovelace Institute, '**How do People Feel About AI'** (June 2023)**.** https://www.adalovelaceinstitute.org/report/public-attitudes-ai/

- Safety, security, and robustness.
- Appropriate transparency and explainibility.
- Fairness.
- Accountability and governance.
- Contestability and redress.

These are all very universally accepted norms that are important in this space. However, the White Paper treats all of these principles in a kind of homogenous way, flattening crucial distinctions that are of high importance for regulatory efficacy and certainty. For instance the principle of fairness has a strong pre-existence existing statutory footing, in say the **Equality Act**[17] and the **Human Rights Act**[18], allowing regulators to bootstrap their AI compliance activities on the basis of this legislation. Likewise in operationalising the principle of safety, regulators can draw on a wide range of statutes that apply to the protection of people in high impact or safety critical AI domains in context.

However, principles like transparency, accountability and governance, and redress have no such analogous statutory redress or existing juridical basis. This is the difference that makes an enormous difference in the effective governance of AI as has been recognised and by multiple initiatives in the UK. Thinking here for instance the CDDO's transparency standard[19], and around the world to take steps to legally codify algorithmic transparency and accountability requirements. We should note here that legislators in the US have already introduced a new updated **Algorithmic Accountability Act**[20] that would create a public repository at the FTC of automated systems, that make critical decisions to ensure a degree of public transparency. In our context Lord Clement-Jones, a while ago, introduced the **Public Authority Algorithm Bill**[21] which in a public sector context contains transparency and system logging mechanisms and ex ante accountability requirements like impact assessment.

Lack of Nuance & Differentiation

My second point has to do with a sufficient lack of nuance and differentiation in the White Paper in how the White Paper itself addresses the private sector. Rather than making a distinction between large tech corporations which possess immense amounts of infrastructural and economic power in the AI innovation ecosystem, and small and medium-sized enterprises which are often subject to the whims and service provisions of big tech companies, the White Paper framework treats the private sector as a kind of monolith. It applies a pro-innovation regulatory rhetoric uniformly across an ecosystem characterised by power imbalances that

[17] **Equality Act (2010).** https://www.legislation.gov.uk/ukpga/2010/15/contents

[18] **Human Rights Act (1998)**. https://www.legislation.gov.uk/ukpga/1998/42/contents

[19] Central Digital & Data Office, **'Piloting the UK algorithmic transparency standard'** (June, 2022). https://cdei.blog.gov.uk/2022/06/01/piloting-the-national-algorithmic-transparency-standard/

[20] **Algorithmic Accountability Act (2022**). United States. https://www.congress.gov/bill/117th-congress/house-bill/6580/text

[21] **Public Authority Algorithm Bill.** https://bills.parliament.uk/bills/3074

should be of great regulatory and statutory significance.

This avoidance of thinking seriously and soberly about symmetries and access control and impacts, that are rooted in existing political economic structures leads the framework into a blind spot, where large scale private organisations control data, compute, and software infrastructure as well as goods and services that substantially affect the public interest. These organisations should face higher degrees of legal and regulatory intervention and control.

Crucially this would involve among other things reforming the incentivisation structures that currently drive AI innovation at scale, by keeping in check the market prerogatives that currently shape the move 'fast and break things' attitudes of firms, whose disproportionate size and financial power enable them to act with virtual impunity.

As generative AI tools including foundation models such as **ChatGPT** become more pervasive and technologists seek more assertively to monetise them, this this need becomes ever more urgent. Rectifying the problem entails the establishment of proportionate regulatory processes that recognise AI as the critical infrastructure that it is fast becoming, and restrain such financial impulses accordingly, for instance by holding producers of the system and large base models behind AI generated content legally responsible for their outputs and impacts.

Closing Remarks

To close I want to direct our attention to a useful mechanism for future making and AI governance and regulation. An awareness of the need to respond to power imbalances in market originated and matured in the English common law tradition under the auspices of public utility thinking.

Over the course of modern times regulatory frameworks that have responded to rapid technological change have been anchored in this common law tradition. Industry actors affected with the public interest and whose production of goods and services led to market consolidation of the sort that limited ordinary accountability mechanisms, have been subject to special restrictions such as the duty to serve all commerce, to demand reasonable prices, to ensure public benefit universally even handedly and fairly.

We are at a point in the evolution of AI innovation where more differentiated considerations such as this should guide our regulatory and legislative viewpoints, especially in a rising age of generative AI that is driven by very few private corporations which are motivated only to further consolidate and centralise their economic power.

Sulabh Soral, Chief AI Officer, Deloitte

Introductory Remarks

This is a really interesting topic, a lot of the thoughts shared today are the thoughts most organisations have, including my employer. My personal opinion is based off, not being a regulatory expert in any sense and not understanding law, but as a technologist and a practicing Machine Learning engineer. I've been in this field developing tools for enterprises for the last 20 years.

AI White Paper

When I look at the Pro Innovation Approach for the White Paper, I really think that's the right approach. To support this argument, when I talk to my clients in the industry, when they ask me what's the most fundamental the basic definition of AI, I say “AI is a software tool that can mimic humans”. As a data scientist I can shoot many holes in this definition myself, and I’m sure the academic world will shoot bigger holes in it, but fundamentally we are talking about a software.

If you are able to distil your thinking to that level for a bit, then what we're really talking about is AI's abilities and its focus in helping to scale human intellect, productivity and innovation for our benefit that's what we're trying to do. It’s the bad outcomes that we're guarding against, and it's not the technology itself. That is why I like the context-We approach because we're not trying to control the technology, what we're talking about is trying to sort of think about the use cases and the context in which AI, and specifically now generative AI, is applied.

One of the other things that I really feel, is that the UK is the right place to have this pro-innovation approach. This is because if you have a pro-innovation approach which is context driven, you already need a very good regulatory framework: a set of regulators as well as rule

of law. That's what the UK is known for. We have that infrastructure for a context driven approach. That's why I think in my mind the UK could lead it. If we take a step back, we think about it I think it's going to be very hard to get Governments on the same platform anyway. If you think about the EU they're already going at their own pace, they're thinking about it in a certain way, so is China and also the US.

International Co-operation

Secondly, we need to think about what suits us, however we also need to have international cooperation which the Prime Minister and the Government are currently doing.

What are those use cases, what are the parts where we definitely need to have convergence with the international community? For example, if it is very customer centric, so if you're using AI to do something that impacts the customer, obviously we need to have international convergence on that because it's not only our customers, as a lot of our companies are actually doing business in Europe or China, so it's customers in other jurisdictions.

Think about it another way, if you're using AI to find the next best material, the strongest material, which generative AI is being used right now by material sciences for carbon nanotubes and diamond nanotubes as a co-pilot, what level of convergence do we need? It is something that our economy can use as a co-pilot, doing the next generation of research and put us at the front. AI is a tool, and its context is what basically drives the risks that we're talking about here, and we need to differentiate what that context is. Is it for scientific research, is it for something else?

Regulation

One of the things I really worry about that actually could have a profound impact is that the marginal cost of knowledge is going down, the marginal cost of knowledge production is going down, but we know from history that when we introduce technologies like this, the human productivity goes up. What does that mean? It means that if we don't have a growth-oriented mindset in the application of AI, this can be deflationary. How do we regulate such that we don't stop growth? This is going to be the most important thing for populations.

If we do hard regulation, it's the big players who practice regulatory capture, so we need to have two things:

1. Context – this is a very important principle.
2. Think not only about the risks but actually the opportunity loss that can happen to this economy if we do not enable AI to fulfil its purpose.

Concluding Remarks

I've been talking to the experts in at **Deloitte**, from our risk advisory practice and regulated

practice[22], and what they say is that in spirit it's good, but in practicality there are some things lacking. The primary thing that they really talk about is the regulatory uncertainty in the White Paper, who's going to coordinate it, how they're going to coordinate. All the contenders have different charters, they have different KPIs to manage. If you do not have a statutory obligation, they're not going to be able to coordinate on that.

[22] **Deloitte Risk Advisory Practice.** https://www2.deloitte.com/uk/en/services/risk-advisory.html

Roger Taylor, Advisor to Responsible AI Programme, Accenture

Introductory Remarks

I agree with the need for growth, and I agree with the objectives set out in this White Paper. However, I do not think it's going to work. I do not think this is a credible plan to deliver what the White Paper aims to deliver.

Just to be clear, I agree with the principles. I agree that we need an agile and pro-innovation and iterative approach, that is right. I agree we do not want to rush to legislation when we don't really know what we're talking about, that is the worst thing we could do. We need to be context specific, and, in that sense, it makes sense to allow sectoral regulators to take the lead. We need to not come down like a hammer on start-up businesses but to recognise where the real risk it. These are all good things.

I think we can forgive it for the fact that it hasn't really dealt with you know the big questions about foundation models, in which the White Paper came out at the wrong time.

However, given all of that I don't think it is going to work. I want to pull out a few examples of why I don't think it's going to work, in line with custom and practice these days I have put them under three headings: money, mechanisms, and motivation.

Money

Money is the point about the scale of this problem: how much do we actually need to do? If you take the description of the central functions, I think there's an excellent description. I think it captures it quite well. I think that what it doesn't recognise, is that given the number of context specific use cases, where regulators are going to be tasked to come up with both how you

implement this stuff in terms of fairness, which is relatively well understood, or something like transparency or interpretability, which are much less well understood, that they're all going to do this in a way that is coherent across the country, and furthermore is interoperable with international regulation. This is a management nightmare and an incredibly difficult problem to solve.

I do think the interoperability thing is crucial because if anyone has a sense that you're going to have to do a bit different in Britain, then they're just not going to bother, and they're not going to have time for that. You could still have lower standards or higher standards, but they have to be interoperable. The G7 is good on this. I don't think the scale the problem is recognised on the national scale and the notion that this would be well performed by a group of generalist civil servants within the ministry struck me as being hopelessly optimistic, and just isn't going to work. It needs real expertise brought to it.

It is the same when you look at it in terms of the sector. For example, there's a suggestion that in recruitment the **Equalities and Human Rights Commission**[23] working with the **Information Commissioner's Office**[24] and others, might put out some guidance to recruitment companies using AI as to what fairness might look like, and so on. Again, the difficulty of the task of addressing what is the problem with fairness is on a scale that is simply not recognised in this document.

Mechanisms

I'm going to illustrate that by talking about the next thing which is mechanisms. The one thing that works with regulation, the reason that laws and regulations work, is that people are pretty clear which sign of the line they stand. If it is unclear to you, which side of the line you stand, it is just a waste of everybody's time, it puts friction into the of economy, it creates mass confusion, it is bizarre. We've seen it with GDPR, with ideas like anonymity or automated decision making which have been put into law before and moved just so fast that it created a high degree of chaos.

The problem is we've got these principles, but what does it mean in practice? How do I know I'm on the right side of this? It is a difficult problem. It's a difficult problem partly because there are standards. We've got £24 million going into the **AI Standards Hub**[25]which is helpful, and there's lots of work going on around the world on standards.

The hardest bit which is judgement. Within this you're going to generate a bunch of data from a system according to certain standards, you're going to have certain obligations to look at this and make certain decisions, and then you've got to know if your view on this is within the reasonable toleration. That's what you need to know to know whether you're on your right side of the law. Nobody knows the answer to these questions, there's nothing in this White Paper

[23] **Equality & Human Rights Commission**. https://www.equalityhumanrights.com/en

[24] **Information Commissioner's Office.** https://ico.org.uk/

[25] **AI Standards Hub.** https://aistandardshub.org/

that properly suggests that anybody is going to be better informed about that by the time that is implemented. So that is a huge problem.

Motivation

The last thing I want to talk about is motivation because I do not think you can solve this problem of what do these rules mean, and how I know I'm doing this right, without industry being very involved. There is academia, but we have to involve industry because they are actually doing this stuff and they are live with the problem. Now the problem for industry, and I think talking about sandboxing is excellent and it's certainly helpful, but industry is the only who's really at risk here. For Government this is not a risk. For academia it is not a risk it's just a growth opportunity, for industry it is a real risk.

The problem is probably explained quite well with sandboxing because a sandbox is sound and tidy but broadly speaking, you're going to open all your secrets to the regulator and pray they kind of go yeah that's good, as opposed to what are you doing, stop it now. It's just not a very good bet.

We need to think quite a lot about how we can create a safe mechanism for industry to engage in this, and for industry really want this. This is a nightmare if you are in industry, and this was described quite well recently, with someone saying that everyone wants to be in the middle of the pack, everyone is a bit scared, and they don’t want anyone to notice them as it were, and they just don't want to have been able to make just something to worry about.

I do think we need to think very imaginatively about how to get around this and I don't know the answer but the sorts of things where, for example, Government puts pressure on industries to get together and come up with a self-regulating proposal at least as a starting point. If you're talking about the evaluation of large language models, they’ve all got different approaches, standards are only just being developed, but they need to go into a room, Government needs to go outside, put its fingers in its ears and then hear the proposals, and try and make that process happen.

Concluding Remarks

There are some problems here. The last point I want to mention is it just not that this proposal is not going to be effective, it's worse than useless. If you put an obligation on a regulator to do something without the resources, they need to do it properly, they will do it in slightly half-arsed way. What that will create is a burden on industry to deal with this lack of clarity in regulations which will just generate a huge amount of work that doesn't really get very far. Worse than that, it will create a degree of uncertainty which, the more responsible organisations will retreat from, with fear, and the less responsible organisations will have the path cleared – you’re actively skewing the tables to produce worse outcomes.

Fran Bennett, Interim Director, Ada Lovelace Institute

Introductory Remarks

At **Ada Lovelace Institute** we see the White Paper as an important milestone on the UK's journey towards comprehensive AI regulation. As a UK-based research institute we're pleased to see our government engage with this incredibly difficult regulatory challenge. The Summit announced last week confirmed the UK's intention to lead the world in AI governance, but credible global leadership needs to be underpinned by a trustworthy regulatory framework here in the UK.

We've set out three tests which we believe any such framework should pass:

1. Coverage of the framework.
2. Capability of those tasked with regulating.
3. Criticality of regulation coming into effect.

Coverage of the Framework

The principles the White Paper sets out describe outcomes that regulators should aim for when AI is used within their domains. For example, the **Financial Conduct Authority**[26] is expected to ensure an AI system used to assess a loan application is appropriately robust, fair and transparent, and that any decision it makes can be contested. However, many parts of the UK economy are only partially regulated, including sensitive areas like recruitment and employment.

If an AI system is used to sort candidates for a job interview, for example, there is no domain-

[26] **Financial Conduct Authority.** https://www.fca.org.uk/

specific regulator that would be responsible for ensuring that system is safe, fair or robust, or that an individual would have any recourse to redress if the system failed. That individual might be able to fall back on equalities law and data protection law, both of which apply across the economy.

But both of these areas of law are weakly enforced and the **Data Protection and Digital Information Bill**[27] waters down protections related to automated decision making, and in some cases existing law is completely inadequate.

An independent review we commissioned led by Matthew Ryder KC[28] demonstrated the insufficiency of the current legal framework to govern biometric data, which underlies many AI systems. So, for all these reasons, it's unclear whether the UK's current regulatory patchwork will evenly cover the risks of AI in different parts of the economy.

Capability of those tasked with Regulating

The powers and resourcing of UK regulators vary considerably, which will affect their capability to implement the AI principles.

For example, if a regulator wants to assess the robustness of an AI system, it will likely need to carry out a technical audit, and not every regulator will have the powers needed to require companies to open their systems up. Even more significantly most regulators – save for those focused on safety – look downstream at where AI is deployed and used if they look at AI at all.

With the advent of foundation models, we're going to need greater governance capability earlier in the value chain, where AI is developed, and tools to manage how liability gets distributed across all the players. A few regulators already boast significant digital expertise and will be set up to do this well already, but the majority simply aren't.

Regardless of whether regulatory resources are delivered centrally or on a distributed basis, it will be important to ensure that regulators – and policy capacity within government itself – are properly resourced.

Criticality of regulation coming into effect

The Government envisions at least a year for the first iteration of the new AI framework to be set up. In ordinary times, that would be a reasonable timeline, but there are urgent risks associated with AI use today, particularly arising from the growing capabilities of general-

[27] **Data Protection & Digital Information Bill.** https://bills.parliament.uk/bills/3322

[28] Ada Lovelace Institute **'Independent legal review of the governance of biometric data in England and Wales'** (June, 2022). https://www.adalovelaceinstitute.org/project/ryder-review-biometrics/

purpose AI.

These systems are being integrated into the economy and our everyday lives at pace, from search engines to productivity software, and through broadly available APIs on which businesses can build their own services. These systems have associated risks, such as 'hallucination' – their tendency to confidently assert untrue statements – or enabling bad actors to produce deliberate misinformation at scale. Unless Government intervention is timely, the unchecked distribution of these systems will only make the regulatory challenge harder to solve when we get to it.

Concluding Remarks

Coverage, capability and criticality. These three tests are at the heart of Ada's response to the White Paper.

We'll continue to work with Government – as well as parliamentarians from across both Houses – to ensure that these tests are met, and that the UK boasts regulation to match its global leadership aspirations.

Daniel Hulme, CEO, Satalia

Introductory Remarks

I wear two hats: an academic hat, I have been researching AI for the past 25 years, I run a masters programme, with 100 students applying emerging technologies to solving problems, in a bunch of different industries. Additionally, I'm currently entrepreneur in residence in **UCL**[29].

I started my company, **Satalia**, during my PhD, which was acquired by **WPP**[30], 18 months ago. My company has been building AI solutions for some of the biggest companies for the past 15 years, and I take on the Chief AI Officer role across **WPP**, and I also invest in new and emerging technologies which is great fun.

White Paper

First of all, I was really heartened to see that the definition of AI in the White Paper, links to the concept of adaptation. It is probably inspired by Sternberg and Salter's definition of intelligence which is called 'directed, adapted behaviour'. For the most part, everything we do in industry isn't adaptive. I built a career helping organisations understand what AI is and isn't, but everything we do doesn't link to that definition of AI.

I was also very encouraged that instead of using the technologies and definitions as a lens for looking at AI, what the paper was trying to do was look at applications. It listed lots of examples of applications using AI.

However, I think where it failed was when trying to categorise and provide framework for the

[29] **UCL Centre for AI.** https://www.ucl.ac.uk/ai-centre/ucl-centre-artificial-intelligence
[30] **WPP**. https://www.wpp.com/

applications of AI. I think there are six applications:

1. **Task automation.** Task automation gets a bad rap in the AI community. It uses very simple algorithms to automate human tasks, but the fact is it drives a huge amount of value in organisations.
2. **Content generation**. Of course, **ChatGPT** and generative AI has accelerated this.
3. **Human representation:** replacing human beings by things that look and behave exactly like a human being.
4. **Insight extraction.** I guess you would call Machine Learning, using these technologies to extract insights from data.
5. **Complex decision making.** This would have been called operations research many years ago.
6. **Human augmentation:** cybernetics and exoskeletons.

Those are the six categories that span almost everything we do in AI, for industry. Each one of those different six applications raise different safety, security, robustness, even sustainability questions. So that's the lens that we use in the **WPP** when thinking about the deploying these technologies and safely and securely.

I think we need to acknowledge that there are already a lot of standards, and principles, and guidelines set out that aid the development of safe and secure software. Really the only difference with AI is compared to traditional software is that it's opaque, for the most part – it is really difficult to see how it is making its decisions. Hence, we need technologies and tools like explainibility, transparency and auditability to make sure that we can understand how these technologies are making their decisions. I don't think it's any more complicated than that. These are software tools that are essentially opaque, that we need to de-opaque to make transparent.

Principles in the White Paper

In the White Paper it talks about five principles. This flattens the complexity of deploying AI by categorising into these 5 principles. Governance was one of the principles that actually sits across what I have just mentioned.

Flattening removes complexity. I think as human beings, we like to think of things in numbers of 5 and when we do flatten things, we lose complexity, and things become nebular. My concern is that we have tried to flatten the complexity of AI into these 5 principles that make it very difficult to actually have a meaningful debate about.

AI Regulation

I would like to close by talking about regulation and AI. I think it is a bit of a mute question, when considering 'what I are we talking about, and what are we trying to regulate against?'. What I failed to see in the paper was a real understanding of the concerns, of the impact these technologies would have on society. I think that we need to understand, what is it we're trying

to mitigate, what risks we're trying to mitigate against and what regulation can help us make sure that we're deploying these systems without too much risk. Risks include post-truth worlds, environmental impacts, surveillance capitalism and surveillance states.

I think to summarise the White Paper it is a good attempt, but the Government must try harder.

5. Institute of Directors Evidence

*Prior to the evidence session, the **Institute of Directors**[31] submitted the below evidence.*

Overview

The IoD's expert advisory group on Science, Innovation and Technology is a community of IoD members with specialist expertise relating to the application of science and technology in the business context. The following is their feedback in relation to the Government's White Paper

What is included, and what is missing in the UK Government's AI White Paper?

The Government White Paper proposes no new specialist regulatory body but recommends the use of existing regulators. Also, no new laws are planned to regulate the technology. Both of these are the exact opposite actions of our nearest neighbours and trading partners the EU, who have moved faster than the UK with its EU AI Act which is a significant piece of legislation that uses a risk-based methodology that has large financial penalties for non-compliance.

The UK's "pro-innovation" approach by contrast is very light touch and the use of many existing regulators, in a patchwork of bodies, will lead to a confused picture for any organisations wanting to comply or use best practice. This puts an unfair and ambiguous burden on any business creating or using AI. Additionally, any efforts for centralised monitoring, as proposed in the paper, will be challenging and costly.

Further consideration towards what is included:

- The CDEI assurance framework, which guides towards elements of consistency in commercial operationalisation, is welcome. However, the companies cited in the CDEI framework are predominantly not British, mostly American, and funded from outside the UK. This seems contrary to the mission to foster British AI innovation.
- The government appears to be putting the mechanics in place to enable them to potentially act to mitigate harms through a 'Central AI Risk function', that is if 'actual' evidence of harms occurs which must be fed through into the sector regulators from industry. Otherwise 'hypothetical' harms are largely disregarded as irrelevant. This approach is contrary to any global principles set out in standards or otherwise which guide risk management. Any 'risk' function should be regulated by an independent body rather than sitting within the political structure of a current Government.

[31] **Institute of Directors.** https://www.iod.com/

- The reference to standards which also underpin the EU Regulation is welcome, this makes it less costly for businesses to develop/provide/use AI across borders. The standards are unlikely to become mandatory in the UK under the current White Paper proposal, which could result in an inconsistent approach by business leaders, or inertia due to lack of uniformity and legal certainty.

What is missing:

- Independent regulation and legal clarity around the development, provision and deployment of potentially harmful AI technologies. We would like to see a technology which is cited by world AI leaders as holding the potential to threaten humanity, to have a dedicated AI Regulator in the UK.
- Consistency in the use of language to define AI. The UK Government has created yet another definition of AI which is not consistent with the EU AI Act or any internationally recognised bodies who have been working on AI policy such as the OECD.
- The principles-based approach is vague and lacks any means through which organisations may operationalise the recommendations. It follows some of the guidance laid out in the OECD AI Principles but does not go far enough, also missing any reference to sustainability and the environment.
- A means by which a buyer of AI technologies can assess and compare/contrast the merits or risks associated with the adoption of one AI technology over another. The driver for UK business, the only driver, one could argue, could end up being linked to whether the AI decisions are used in the EU or not. AI systems which are developed and hosted outside the EU, but the decisions are used in the EU are subject to the EU AI Act.
- A definition of what 'pro-innovation' actually means.
- Recognition that AI technologies and their applications, both now and in the future, will not be constrained by international, regional or national borders. Understanding the impact of the future ubiquitous availability of AI technology enablers and systems across the world is a prerequisite to developing successful regulation and increasing innovation.

What is required for more effective and efficient adoption of AI in the UK?

Aside from a thorough, credible and workable piece of legislation underpinned by an AI regulator, there is a requirement for the upskilling of boards.

It is clear that AI needs to be on the board agenda and considered seriously as part of the G in ESG (Environmental, Social and Governance) and the CSR (Corporate Social Responsibility) requirements. AI should not be confined within the realms of IT, although the CIO may take responsibility for implementation and management.

- An IoD members' survey revealed in 2022 that 80% of boards did not have a

process in place to audit their use of AI. They said they did not know what questions to ask.

- Board governance of AI at the project inception is important because boards do not wish to unravel ethical issues later which may have a reputation and cost impact.
- Research reveals a gap between board governance and the use of AI in their businesses.
- Over 86% of businesses already use some form of AI without the board being aware of this.
- AI applications in business are expanding so AI opportunities and threats need to be understood.
- AI can amplify existing bias in human decisions and safeguards are needed to prevent this and its impact on perpetuating bias in existing culture.
- AI governance may be rooted in the core ethical values of the business.
- AI risk and governance model requires a framework which boards can use as a blueprint.

The UK government needs to give clarity to businesses and organisations in a short timeframe. AI has the potential to touch every UK citizens' life via the workplace, home or as a consumer.

Consequently, a specialist body of experts should be formed to regulate what is arguably the 5th industrial revolution. The Government's announcement that it will invest £100m into an AI taskforce which examines the British potential for the development of its own generative AI seems too little too late.

The support of British companies developing, providing or using AI needs to be a priority. The UK has not shown a great deal of success in this area. for example, Deepmind, a British start-up was acquired by Google in 2014 and since, it has been reported that Google wrote off nearly $1.5 billion of debt, seeing the value of the innovation not just to humanity but its potential across the entire Alphabet enterprise too.

Deepmind is to be consolidated into Google's research products with revenues reported to be almost entirely generated from the Alphabet internal customers. It is also reported to be central to Google competing with ChatGPT.

Deepmind is cited in the White Paper as a 'British' success because it is based in the UK which as a statement, to some, lacks credibility. If Deepmind had been supported by the British Government and not sold to Google, the ability of the British AI community to develop a home-grown foundational model would be vastly ahead of where it is now.

Benevolent AI, a British 'golden child' has just announced that it will be cutting 180 jobs to allegedly extend its 'runway'. Based on desk research, the organisation appears to have consumed $270m in investment to produce c £4m revenue in 2021.

Commercialisation does not appear to be a British strength. The UK Government should be encouraging the 'long game' by supporting investors in game changing technology which may in turn lead them to become more 'patient'.

What is required for more effective and efficient adoption of AI in the UK?

- The UK government should accept that it must regulate AI and do so very soon.
- Enhance R&D Tax Credits for software companies.
- Deliver timely, clear legislation regulated by an expert body.
- AI legislation should focus on the outcomes aimed at achieving and managing the associated risks. This risk-based outcomes approach is more adaptable and flexible in adjusting to the continuously evolving AI sector.
- Create a moat around the UK AI Governance space and favour British companies who are funded by British backers.
- Recognise the leading bodies in the UK which drive good governance in the UK and engage effectively, e.g., the IoD which provides world-leading chartered qualifications for Directors.
- Increase the public's awareness and understanding of AI and its applications.

6. Speaker Bios

Prof. David Leslie, Director of Ethics & Responsible Innovation Research, Alan Turing Institute

David Leslie is the Director of Ethics and Responsible Innovation Research at The Alan Turing Institute and Professor of Ethics, Technology and Society at Queen Mary University of London. He previously taught at Princeton's University Center for Human Values, where he also participated in the UCHV's 2017-2018 research collaboration with Princeton's Center for Information Technology Policy on "Technology Ethics, Political Philosophy and Human Values: Ethical Dilemmas in AI Governance." Prior to teaching at Princeton, David held academic appointments at Yale's programme in Ethics, Politics and Economics and at Harvard's Committee on Degrees in Social Studies, where he received over a dozen teaching awards including the 2014 Stanley Hoffman Prize for Teaching Excellence. He was also a 2017-2018 Mellon-Sawyer Fellow in Technology and the Humanities at Boston University and a 2018-2019 Fellow at MIT's Dalai Lama Center for Ethics and Transformative Values.

David is the author of the UK Government's official guidance on the responsible design and implementation of AI systems in the public sector, Understanding artificial intelligence ethics and safety (2019) and a principal co-author of Explaining decisions made with AI (2020), a co-badged guidance on AI explainibility published by the Information Commissioner's Office and The Alan Turing Institute. After serving as an elected member of the Bureau of the Council of Europe's (Coe) Ad Hoc Committee on Artificial Intelligence (CAHAI) (2021-2022), he was appointed Specialist Advisor to the CoE's Committee on AI where he has led the writing of the zero draft of its Human Rights, Democracy and the Rule of Law Impact Assessment for AI,

which will accompany its forthcoming AI Convention. In his advisory role with the CoE, David led his team in writing a Primer to support the meaningful participation of citizens and civil society organisations in the stakeholder outreach of the CAHAI's Feasibility Study, published March 2021 and translated into French and Dutch. He and his team were then asked by the CoE to carry out research into the Impact Assessment Instrument to be submitted to the Council of Ministers as part of the CoE legal framework on AI and published the 335-page Human Rights, Democracy, and the Rule of Law Assurance Framework for AI Systems submitted in September 2021. As part of his international work, he also serves on UNESCO's High-Level Expert Group steering the implementation of its Recommendation on the Ethics of Artificial Intelligence, a first-of-its-kind document adopted by the 193 Member States of the Organisation.

Anthony Walker, Deputy CEO, TechUK

Antony Walker is deputy CEO of TechUK, which he played a lead role in launching in November 2013.

Antony is a member of the senior leadership team and has overall responsibility for TechUK's policy work. Prior to his appointment in July 2012 Antony was chief executive of the Broadband Stakeholder Group (BSG), the UK's independent advisory group on broadband policy. Antony was closely involved in the development of broadband policy development in the UK since the BSG was established in 2001 and authored several major reports to government. He also led the development of the UK's world leading Open Internet Code of Practice that addresses the issue of net neutrality in the UK. Prior to setting up the BSG, Antony spent six years working in Brussels for the American Chamber of Commerce following and writing about telecoms issues and as a consultant working on EU social affairs and environmental issues. Antony is a graduate of Aberdeen University and KU Leuven and is also a Policy Fellow Alumni of the Centre for Science and Policy at Cambridge University.

Roger Taylor, Advisor to Responsible AI Programme, Accenture

Roger Taylor has worked as an entrepreneur, a regulator and a writer. He was inaugural chair of the Centre for Data Ethics and Innovation from 2018-2021, advising the UK government on AI and data policy. Before that he was chair of Ofqual, the qualifications regulator and National Advisor on Intelligence to the Care Quality Commission. He co-founded Dr Foster which pioneered the use of public data to provide independent ratings of healthcare. He worked as a correspondent for the Financial Times in the UK and the US. He has written two books: God Bless the NHS (Faber & Faber 2014) and Transparency and the Open Society (Policy Press 2016).

Fran Bennett, Interim Director, Ada Lovelace Institute

Francine Bennett is Interim Director of the Ada Lovelace Institute, covering Director Carly Kind's maternity leave in 2023, and she has been a founder member of the Institute's Board since 2019.

Prior to joining Ada, Francine was VP of Data at Healx, a biotech company which uses AI to find treatments for rare diseases; cofounder of Mastodon C, a data science consultancy that supports local and central government to realise the social and business potential of their data; and was a founding trustee of DataKind UK, which provides data science support to UK charities.

She is also currently a member of the Gambling Commission's Digital Advisory Board and the Advisory Council of the British Library.

Daniel Hulme, CEO, Satalia

Daniel Hulme (PhD) is a globally recognised expert in Artificial Intelligence (AI) and investor in Emerging Technologies. He's the CEO of Satalia, an award-winning company that provides AI products and solutions for global companies such as Tesco and PwC. Satalia joined the world's largest marketing company in 2021, WPP, where Daniel is now the Chief AI Officer; helping define, identify, curate and promote AI capability and new opportunities for the benefit of the wider group and society.

Daniel has over 20 years academic experience with AI. Having received a Masters and Doctorate in AI at UCL, Daniel was previously Director of UCL's Applied AI MSc (Business Analytics), where he is now UCL's Computer Science Entrepreneur-in-Residence and a lecturer at LSE's Marshall Institute, focused on using AI to solve business and social problems. Daniel is also an Impact Board Member of St Andrew's University Computer Science department.

Passionate about how technology can be used to govern organisations and bring positive social impact, Daniel is a popular keynote speaker specialising in the topics of AI, ethics, metaverse, emerging technology, innovation, decentralisation and organisational design. He is a serial speaker for Google and TEDx, holds an international Kauffman Global Entrepreneur Scholarship, and is a faculty member of Singularity University.

Daniel is a contributor to numerous books, podcasts and articles on AI and the future of work. His mission is to create a world where everyone has the freedom to spawn and contribute to innovations, and have those innovations become free to everyone. He has advisory and executive positions across companies and governments, and actively promotes purposeful entrepreneurship and technology innovation across the globe.

Sulabh Soral, Chief AI Officer, Deloitte

Sulabh is Chief AI Officer at Deloitte Consulting & Innovator with 20 years product innovation, data science and technology experience.

He leads the Deloitte AI Institute in the UK, bringing together the best of global AI thinking and helping advance human-machine collaboration in the Age of With.

Sulabh advises Global C-level executives on AI and exponential technology led transformations, particularly in the financial services industry. He works with a fantastic team of highly skilled experts spanning Fintech, AI theory, applied data science, ML engineering and innovation/design thinking, and has led significant greenfield and brownfield transformations.

Some of the projects Sulabh has led include - AI driven UW in commercial insurance, AI powered Digital Banking solutions, satellite and imagery-based systems for risk mitigation, intelligent fleet solutions based on telematics and auto quotation in motor insurance.

7. Contact

APPG AI Secretariat

Big Innovation Centre
14-16 Dowgate Hill
London EC4R 2SU
United Kingdom

info@biginnovationcentre.com
www.biginnovationcentre.com

appg@biginnovationcentre.com
https://bicpavilion.com/about/appg-artificial-intelligence

www.biginnovationcentre.com