

EVIDENCE GIVING MEETING #5

All Party Parliamentary Group on Artificial Intelligence
APPG AI

11 September 2017
5:30 – 7:00 pm
House Of Commons, Committee Room 4A

@APPG_AI
www.appg-ai.org

Main Focus

- What is the social impact of AI?
- AI and new cultural systems (human and machines working together)
- The image (sensational image) vs. the reality (what happens in real life) of AI.
- New forms of organisational structures, corporate governance and wealth in an AI focused economy (forms of living, machines and people working together – collaboration, co-creation).

Thought Leaders

- **Krishna Sood** – Technology Lawyer, *Microsoft*
- **Miles Brundage** – AI Policy Research Fellow, Oxford Future for Humanity Institute, *University of Oxford*
- **Dr Joanna Bryson** - Reader at the Department of Computer Science, *University of Bath*, and Affiliate at the Center for Information Technology Policy at *Princeton University*
- **Dr Stephen Cave** – Executive Director of the Leverhulme Centre for the Future of Intelligence, *University of Cambridge*
- **Dr Kate Devlin** – Senior Lecturer, Department of Computing, *Goldsmiths, University of London*
- **Dr Julian Huppert** – Director of the Intellectual Forum, Jesus College, *University of Cambridge*
- **Rodolfo Rosini** – Co-Founder and CEO, *Weave.Ai*
- **Dr Sandra Wachter** – Postdoctoral Researcher in Data Ethics and Algorithms, *Oxford Internet Institute*

Krishna Sood – Technology Lawyer, Microsoft

- Krishna is a media, entertainment and technology lawyer with over twelve years of industry experience. She currently works for Microsoft as a Senior Attorney supporting the Skype team.
- Microsoft has adapted the below principles AI technology should follow:
 - Be designed to assist humanity
 - Be transparent
 - Maximize efficiency, without sacrificing human dignity
 - Respect privacy
 - Address algorithmic accountability
 - Guard against biases/stereotypes
- The UK government should aim to:
 - Build a robust computer regime, with free movement of data
 - Fully implement the General Data Protection Regulation
 - Continue to collaborate in initiatives such as Partnership on AI to discuss social impacts
 - Invest in research and development
 - Promote skills from primary school level to higher education,

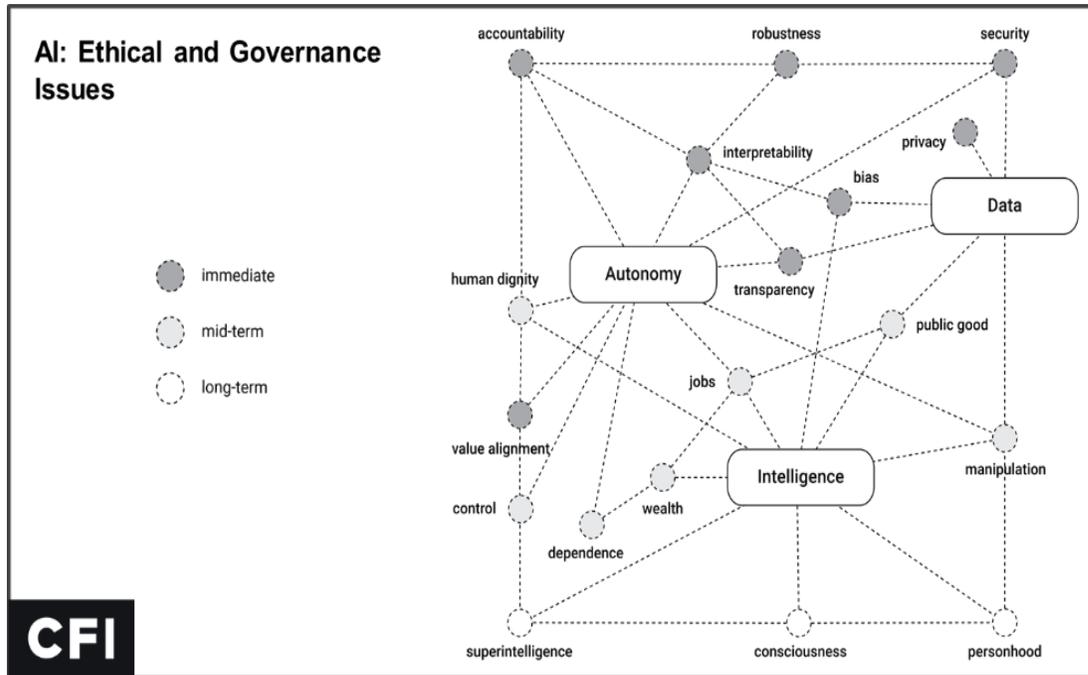
Miles Brundage – AI Policy Research Fellow, Oxford Future for Humanity Institute, University of Oxford

- Miles is AI Policy Research Fellow with the Strategic AI Research Center at Oxford Future for Humanity Institute of the University of Oxford. His research focuses on developing methods for rigorous analysis of AI development scenarios and appropriate policy responses.
- Technological progression is exponential → We cannot truly anticipate how fast technology will progress and in what path
 - The UK government should be careful in basing their policies on a specific time frame
- AI near-term implications:
 - The job market will be disrupted
 - Potentially dangerous security implications linked to hacking and autonomous weapons
- AI long-term implications:
 - Safety concerns
 - Advanced AI will add economic pressures. Developing countries are likely to not follow the same safety standards in their attempts to climb up the technology ladder.
- UK has the potential to lead in technological developments in AI and governance of AI

Dr Joanna Bryson - Reader at the Department of Computer Science, University of Bath, and Affiliate at the Center for Information Technology Policy at Princeton University

- Joanna is a Reader at the Department of Computer Science, University of Bath, and Affiliate at the Center for Information Technology Policy at Princeton University. Her primary research interest is using Artificial Intelligence to understand Natural Intelligence.
- Computation is a physical process – not an abstraction like Math
 - It takes time, space and energy
 - → omniscience is not possible
- What sets AI apart from Natural Intelligence is that Artefacts are made deliberately, by humans
 - We shouldn't over-personify humans
 - Intention and responsibility are concepts invented by humans → declaring artefacts responsible would abrogate human responsibility and facilitate corrupt/dangerous human behavior
- We should be careful to not encourage irresponsible corporate behavior
 - We don't want to reward with reduced liability, obfuscated AI
 - We don't want to encourage the pursuit of ill-defined attributes like 'consciousness' in order to avoid tax or legal liability
- Key recommendations to Parliament:
 - Don't separate the core problem of human rights, cybersecurity, biases, etc.
 - Apply standard auditing procedures to regulate ecosystem. We don't need to fully understand every detail before starting to govern it
 - Enable the on-demand and routine auditing of AI and algorithmic systems.
 - Regulate behavior, not necessarily code, and certainly not 'weights'

Dr Stephen Cave – Executive Director of the Leverhulme Centre for the Future of Intelligence, University of Cambridge



- Stephen is Executive Director of the Leverhulme Centre for the Future of Intelligence (CFI) at the University of Cambridge. He is also a Senior Research Associate in the Faculty of Philosophy and a Fellow of Hughes Hall.
- Ethical and governance issues in the sphere of AI often overlap and are interlinked
- The same issues are cross-cutting between the three categories: data, intelligence, and autonomy
- Holistic solutions need to consider this interplay
- Some ethical/governance issues posed by AI are related to those posed by other technologies but others are genuinely new
 - Issues arising from AI systems increasing autonomy
 - Accountability when decisions are made by machines
 - Over-dependence and de-skilling
 - Impact of automated decision-making on a range of human rights
- Recommendations for the UK government:
 - Need to create a well-informed infrastructure to govern AI
 - Need to create an advisory body to govern AI. The Royal Society recently called for government to establish a data governance body. However, as many of these issues are closely related to issues with AI, Stephen suggested an advisory body that takes into account both AI and data.

Dr Kate Devlin – Senior Lecturer, Department of Computing, Goldsmiths, University of London

- Kate Devlin is a Senior Lecturer in the Department of Computing at Goldsmiths, University of London. Coming from an Arts and Humanities background (as an archaeologist) and now working in STEM, Kate has a demonstrable track record of combining diverse fields and methods of research under the umbrella of Cognition, Computation and Culture.
- AI trends:
 - Increasing trends of humanizing technology. People are becoming increasingly dependent on technology.
 - Impact is affecting every industry and sector, including those that are traditionally known as taboo
- The sex tech industry is worth \$30 billion globally. AI has had humongous impacts in the industry with introduction of technologies such as sex robots.
- Technologies such as sex robots are controversial because they raise ethical questions
 - Can we use AI to combat loneliness?
- Recommendations for the UK government:
 - Invest in further research to understand AI's impact in all industries – including those that are not comfortable to discuss
 - Base policies on well-curated evidence

Dr Julian Huppert – Director of the Intellectual Forum, Jesus College, University of Cambridge

- Julian is Research Councils UK Academic Fellow in Computational Biology at the University of Cambridge. Dr Huppert is also founding Director of the Intellectual Forum, a research centre focused on studying problems of public interest such as the future of work, health and social care. He is a former Member of Parliament for Cambridge and University Lecturer in Physics, where he worked on unusual structures of DNA.
- Focus on three main implications of AI:
 - Power
 - Government should be concerned on the overconcentration of power
 - Promote a landscape where competition exists, open standards are followed, and open interactions are norm
 - Governance
 - Companies should adapt new approaches to corporate governance
 - DeepMind Health example – set up a committee that Julian is a member of to produce an annual report on the company’s practices. This model is publicized and the stakeholders are accountable for their actions.
 - Jobs
 - AI will impact the future of jobs. Some jobs will be lost and others created.
 - Challenges us to rethink the purpose of a human being. Perhaps individuals will have more time to do what is meaningful in life?
- Regulators should not frame conversations as simple trade-offs. There is a limiting curve in which you can have more of the ‘good’ option and less of the ‘bad.’ For example, it doesn’t have to be a trade-off between health benefits and giving away your rights to privacy.

Rodolfo Rosini – Co-Founder and CEO, Weave.AI

- Rodolfo is the Co-Founder and CEO of Weave.AI, which is aiming to transform the mobile experience using AI, contextual search and deeplinking, learning from users and offering them information and services. Rodolfo has taken three start-ups from inception to market, recruited their management teams, raised VC funding for each and successfully sold one.
- The concept of exponential growth is not new
- Technology acceleration will dramatically change existing social concepts
 - Notions such as privacy might cease to exist
- Applied AI is changing the world and is a force multiplier of other technologies
- We are living in an inflection point in which UK has the opportunity to lead
 - China has announced AI research as priority
 - US had their first Congressional hearing on AI
 - Russia noted that whoever will lead the world in AI will govern the world
- Recommends that the UK national strategy:
 - Addresses the lack of entrepreneurship and encourages graduates to start their own businesses
 - Confronts job disruption and prepares those whose jobs will be automated
 - Rethinks current migration policies
 - Reforms the education system to prepare citizens to benefit from AI

Dr Sandra Wachter – Postdoctoral Researcher in Data Ethics and Algorithms, Oxford Internet Institute

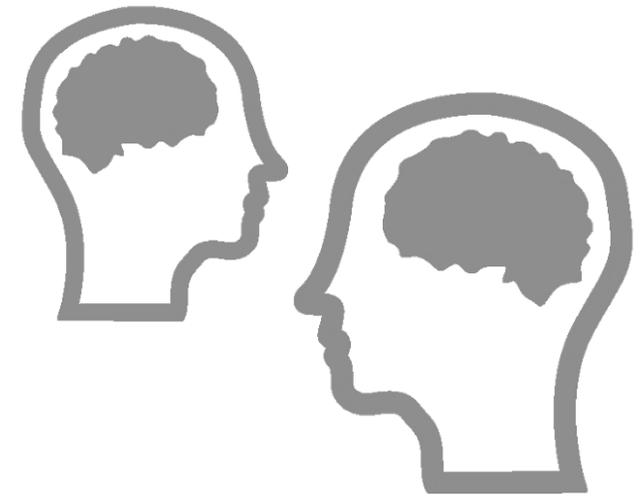
- Sandra is a lawyer and Postdoctoral Researcher in Data Ethics and Algorithms at the Oxford Internet Institute. Sandra is a member of the Ethics and Philosophy of Information Research Cluster and the Digital Ethics Lab. She is also a Turing Research Fellow at the Alan Turing Institute in London, a member of the Law Committee of the IEEE and serves as a policy advisor for governments and NGO's around the world on regulatory and ethical questions concerning emerging technologies.
- AI is increasingly being used to make decisions across the world.
 - Pro – AI has the potential of making decisions that are more fair, transparent, and accountable than humans.
 - Con- AI is opaque and hard to understand. AI is created by humans and there are embedded biases build within them.
- AI is fueled by data and there is much debate about data ethics and governance.
- GDPR is expected to come to force in March 2018. This regulation is good because it gives citizens the right to contest a decision made about him/her. Its flaws are:
 - The regulation applies only for decisions that are made through a fully automatic process
 - The right to explanation is not legally binding
- Recommendations for the UK
 - UK should lead the world in closing the accountability gap in the GDPR
 - Restate the regulation to include the right to contest a decision predominantly based on automated making
 - Make the right to explanation legally binding
- We want a future in which AI and humans can work side by side. The legal frameworks should reflect this.

Questions and Answers

- 1. Where should we study the future of work?
 - There should be a cross-disciplinary approach to researching what the future job market will look like. Job disruption will eventually affects all professions.
 - Academia, industry, and government have a shared responsibility to explore the topic
- 2. What are the biggest AI threats for society?
 - Implications that are worrying for society include: job disruption, privacy, accountability, biases, inequality, etc.
 - These have always been threats in society. AI simply puts them in the debate and urges stakeholders to address them.

Key Takeaways

- AI can have a social purpose. AI can be used to make us better humans.
- Society can reap many benefits from AI technologies but we also have to discuss the risks.
 - Issues such as transparency, job disruption, privacy, algorithmic biases, etc. are interlinked
 - These issues are not new or caused by AI technology. They existed in our society and AI has brought them into our agendas.
- Solution requires a holistic approach in which all stakeholders work together.
- A key question to consider is: *even if we can make something, should we?*
- The UK can lead in this revolution and shape how AI will be governed moving forward





BIG INNOVATION CENTRE

Secretariat

All Party Parliamentary Group on Artificial Intelligence (APPG AI)

Big Innovation Centre

appg@biginnovationcentre.com

Ergon House, Horseferry Road
Westminster, London SW1P 2AL

www.appg-ai.org

www.biginnovationcentre.com