Hardwiring trust

Data ethics and the application of AI

Summary report of a major
Onward and Atos partnership

ONWARD

Atos

About Onward

Onward is an independent, not-for-profit think tank registered in England and Wales (no. 11326052) whose mission is to create new ideas for the next generation of centre-right thinkers and leaders. We are a new kind of campaigning think tank, focused on developing policies that work, are politically possible, and which have the support of an effective campaign to make them a reality.

About Atos UK and Ireland

Atos is a global leader in digital transformation with over 110,000 employees in 73 countries and annual revenue of over € 11 billion. European number one in Cloud, Cybersecurity and High-Performance Computing, the Group provides end-to-end Orchestrated Hybrid Cloud, Big Data, Business Applications and Digital Workplace solutions.

About this report

This is a summary report of a major Onward and Atos partnership series of events, held in Spring/Summer 2019. Over two roundtable events in the Houses of Parliament, we brought together leading parliamentarians, industry professionals, start-ups, regulators and experts to explore the implications of AI and machine learning, and how policymakers should best respond.

The events were held under the Chatham House Rule, so all comments below are anonymised and the discussion summarised.

For more information, or to explore how you might partner with Onward in future, please contact office@ukonward.com.

Contents

Foreword

If software was eating the world in 2011, artificial intelligence and machine learning are now devouring it. According to the latest figures, AI now attracts $27 billion of venture capital[1] and 30,000 published patent applications a year,[2] up 400% and 3000% respectively compared to a decade ago. The explosion of AI in every aspect of the economy and government has profound implications for society, the economy and government.

To consider the implications of the widespread deployment and adoption of AI, Onward convened a major series of events in partnership with Atos. We were primarily interested in three questions. First, transparency and awareness – can we "look inside the black box" to ensure understanding and eliminate bias? Second, benefits for citizens – how can we ensure the positive aspects of AI are not lost and the tangible benefits understood? Third, protection and regulation – how can we protect citizens without harming innovation?

The resulting conclusions were striking in that they reveal how much we still need to develop. Companies have yet to find a way of mitigating the public's fear of AI. Regulators are still fumbling to find the right frameworks for a set of technologies that defies comparison to existing markets. It is still not clear who is responsible – and crucially accountable – if AI goes wrong. That is partly because the AI debate is a discussion that policymakers are still not having with the urgency it deserves. I hope that this report provides both food for thought and a spur to action for that future policy debate.

Will Tanner
Director, Onward

1 Hogarth, A. (2019), State of AI 2019

2 IPO (2019), Artificial Intelligence: A worldwide overview of AI patents and patenting by the UK AI sector

Introduction from Atos UK and Ireland

We are at a critical time, where the future power of AI is still to be fully defined, but we already know that it will shape advancements in every field, from science and medicine to space exploration and energy, and will result in a remodelling of our society. This is why our politicians must appreciate the size of the change and the opportunity, think about the long-term future, and start to establish the policy frameworks and new public bodies which will work in tandem with those who will design and deliver our future world.

The recent advances in AI are thanks to the immense technological power now at our disposal, taking it out of the realm of research and making it a viable accelerator in all kinds of ways. Our AI and Ethics roundtables have explored this change, how the public can be better informed on the vast potential, and the ethical questions for society. Innovation is by its nature intensely collaborative. By working together with customers and partners, digital leaders and policymakers, we want to help ensure the benefits of the technological revolution are shared by everyone in society. We are at a new dawn that is about much more than technology – it is about advancing human performance and progress to positively transform our world and address our most important challenges both now and in the future.

Kulveer Ranger
Senior Vice-President, UK & Ireland

Trusting AI

How do we know what's in the black box?

21st May, Houses of Parliament

Trust is hard won and easily lost. This is especially true of emerging technologies, as we saw in the early 2000s with the introduction of genetically modified foods in particular. The public are understandably nervous about innovation, especially when stoked by cultural fears, and government and industry can struggle to maintain the positive narrative.

Yet for the Government to fulfil its ambitions for the UK to lead the world in AI development, it will need citizens on-side. The first event in the series discussed how we can foster greater trust and understanding of artificial intelligence, robotics and machine learning, in order to mitigate public concerns and extract the maximum benefit for citizens and users.

The discussion started from the beginning. We discussed what we mean by artificial intelligence, and how much can we reasonably expect citizens to want – or be able – to understand it. While participants had mixed views on many issues, there was universal agreement that trust was a prerequisite for further deployment, not an optional extra that could be disregarded.

The ensuing discussion particularly considered how to develop a clear and accommodating regulatory framework that encourages innovation but gives due weight to ethical concerns. It was through that framework – as with stem cell research and other well-managed technological transitions – that attendees saw the greatest opportunity for growing public trust.

Key points of the discussion

1. Artificial intelligence is poorly defined and misunderstood, even in policy debates. There was widespread agreement that the preceding question for trust in AI is definitional: people mean different computing technologies and algorithmic functions. Politicians struggle because they do not have a simple explanation to present to constituents and there is no commonly accepted "plain English" definition. This is a challenge.

2. Public confidence in AI is weak, and defined by popular culture. Related to definitional concerns, participants were unanimous in their perception that AI was viewed negatively by the public. It was noted that "technology is often seen as something that is foisted upon people, rather than a tool to improve their lives" and what comes to mind with AI is not a Netflix algorithm or a personalised medical programme, but the image of The Terminator. The public fear the unknown, precisely because AI has not been defined in positive terms.

3. This will only change if industry and government proactively seek to understand and respond to public concerns. There was a perception that much of AI research was focused on narrow technological or governance debates and not enough on how AI can benefit public welfare and safety. This led to AI being understood and accepted at the expert periphery, but out of sight for the majority. Most participants agreed on the necessity of demystifying AI and debunking myths, and that this should be led primarily from industry.

   Importantly, attendees agreed that industry and politicians should focus on the material benefits they want AI to help achieve, for example faster and better medical treatment or lower pollution in cities. After all, if we cannot describe what we want from AI, how can we persuade people to be supportive of its development and use?

4. The contents of the 'Black Box' are less important than how it is used. There was a widespread sense that, while they needed to trust AI, most citizens are not interested in – or capable of – understanding the technical mechanics of machine learning algorithms. They are interested in how they interact with AI products and services in their daily lives, the services they offer and the protections they guarantee. It was observed, by way of analogy, that the vast majority of people do not want or need to know how jet engines function before boarding an aeroplane; they trust that the technology works, what it is used for, and that it benefits them.

The public – and politicians – also need to know who is responsible. Just as with jet engines, the system must deliver some level of assurance that an individual or company will bear responsibility for the consequences of any failure of the technology. At the moment, there is no such accountability with AI.

5. Regulators must be proactive with AI and related technologies, not passive. There was complete agreement that regulating AI is an enormous challenge, crossing multiple sectors, technologies and incalculable use cases. Generally, attendees argued that regulators and the industry should be forward-looking on ethical debates but not let them damage the UK's reputation for being friendly to tech firms.

   A number of attendees believed there was greater scope to explore ethical questions in the design and build phase, with regulators baking in harder standards at this early stage before firms bring a product to market. There was some concern of a stand-off between firms that are wary of developing a product with which regulators might take issue and regulators waiting for new products to be developed before taking a stance. This would be the worst of all worlds, and lead to stagnation.

6. Regulation needs to be supportive and inclusive of smaller firms. A number of start-ups around the table argued for any standards and regulations to accommodate the concerns of smaller firms. Too often, the conversation is dominated by a few tech giants, which can harm innovation and reinforce incumbent advantage. It was commented that an essential question in the medical and pharmaceutical industry is how a particular treatment helps patients. Similarly, looking at tangible benefits to people's lives should guide regulation of AI.

Positive AI

How does AI make lives better?

3rd July, Houses of Parliament

This morning, you probably browsed the news, checked your emails, planned your route to work and paid for the journey – all using artificial intelligence. From personalised news aggregation and spam filters to route planners and fraud detection, machine learning is being used to immeasurably improve the consumer experience. Yet most of us will never even notice.

The question for policymakers and industry is how to maximise these benefits. There are clear opportunities to use AI to positively impact citizens through improved education, healthcare and tax policy through existing public sector data assets. For companies, there are countless opportunities to improve existing products and develop new ones to benefit consumers and businesses.

This discussion considered the positive benefits of AI, and the main barriers to further adoption. Participants agreed that citizens remain wary of artificial intelligence and, as we found in the first roundtable, unaware of many of the benefits it could bring to their lives.

It was also agreed that we must ensure the benefits accrue to everyone. Co-production, developing in collaboration with the end users and intended beneficiaries, is essential for developing products ethically and effectively.

Key points of the discussion

1. There was a debate about why perceptions of AI were negative. While everyone agreed that perception of AI was increasingly negative, perhaps informed by science fiction, people disagreed as to why this might be. Some suggested that people feel a sense of a lack of control over new technologies. The internet, for example, seems dominated by the commercial interests of a few big companies. Losing control over innovation and the use of technology in daily life can lead to feelings of insecurity. Separately, some participants thought that this may be a function of AI's relative market immaturity. Concerns about technological development are hardly new. As the use of AI becomes increasingly ubiquitous, fear and suspicion may decline.

   One other problem that was raised on this issue is that public engagement is often only for short, defined projects. There is too little sustained long-term engagement and no single point of authority for the public to go to find information.

2. AI practitioners and policymakers should focus on the everyday benefits of AI, not the edge cases. There was widespread agreement that the public will be more receptive to innovation that delivers direct benefits to them, rather than to distant others or in pursuit of exotic goals. Participants discussed whether AI could be visibly used to tackle run-of-the-mill concerns such as bin collection or pothole detection, in order to popularise the benefits of AI. It was also noted that greater uptake among SMEs was essential to bring the commercial benefits of AI to a wider group of people – and that while large firms often adopt AI, they do not do enough to diffuse it to their supply chains or client networks.

3. Public trust needs to be earned, not assumed. There was unanimous agreement on the need for greater responsibility around AI. When technology fails, where do you go? Who resolves a problem if the system does not work? Where there are crises, should there be someone to respond and reassure users and the general public? As in our first roundtable, participants agreed that at the moment this does not exist.

   One of the repercussions of this is that well-meaning organisations shy away from using AI out of fear or lack of understanding. For example, there is widespread fear amongst charities and civil society groups concerning how they can use personal data for fundraising, leading to low uptake of AI in the sector. In failing to put in place a reasonable framework for AI, we risk deterring institutions from collecting and using data for good causes or beneficial aims.

It was argued that professionalisation would necessarily include discussions about ethics and best practice – as is the case for doctors, lawyers and teachers. Such a move would help to instill a greater sense of social responsibility.

4. AI systems should be developed with concern for – and input from – end users. It was argued that we cannot have a situation where, as one participant suggested, a firm develops a solution that they think the NHS needs without involving patients or practitioners in the product development, or thoroughly understanding the impact that those technologies would have when deployed or scaled up in a medical setting. This puts a much greater onus than currently exists on engineers to look at the wider context rather than provide a one-off solution for a particular issue.

   Finland's smart cities were referenced as an exemplar. The model is one of co-production and co-creation. Residents are actively involved in the planning and projects, and backers must demonstrate tangible benefits for the community. In the health sector, this would mean that developers should consult with doctors and patients and place their solutions in the context of wider problems. The group generally agreed that a holistic and inclusive approach to designing solutions is the most effective.

5. There was agreement that one of the most positive arenas for AI deployment was in education, but disagreement on the limits. It was noted that machine learning heralds the possibility of tailored education for each pupil in and out of a classroom setting. However it was countered that reliance on algorithmic teaching methods would necessarily lead to children being taught to think in ways that are understandable to machines. How is a child that is educated in that way going to compete with AI when their education has been optimised according to how a machine thinks?

   There was broad agreement, however, that AI meant that schools should increasingly focus on skills that are more resistant to automation, such as creativity and critical thinking, and invest in human skills rather than training children to think like machines or algorithms. It was agreed that there is a greater need for people to be able to find knowledge and apply it situationally rather than memorising or replicating facts and figures.

Conclusions

Even as artificial intelligence sweeps through the economy and society, it is clear that much of the infrastructure and institutions needed to maintain trust and legitimacy are still being put in place.

This is not for lack of action: the AI Sector Deal, establishment of the Turing Institute and Centre for Data Ethics and Innovation, and the investment in PhD places and research funding all exemplify the Government's determination for the UK to lead the world in this technology. But there is more to do.

This series has exposed a number of key areas for policymakers and practitioners to focus their energies.

1. Communication

First, it is clear that there is a core communications challenge. Government ministers and engineers alike must take seriously the need to bring the public with them, or risk a backlash which could set the industry back years, as it did with GM. That should be built around the everyday benefits of AI to citizens, not to science, and communicated clearly and transparently.

2. Accountability

Second, people must be accountable. At the moment, it is unclear who bears the costs when AI goes wrong, and which regulator is responsible for meting out judgement. Regularising the framework within which AI systems are developed and deployed is an urgent task. It seems both industry and regulators are ready to engage.

3. Accessibility

Third, we must ensure that small companies and start-ups are part of this conversation. Too often policy can be made for the biggest common denominator, not the start-up without a seat at the table and with potentially as much trust risk attached to their product.

There is no doubt that AI can be a force for good. Ensuring it is able to be is the bigger challenge.

Thank you

Onward and Atos would like to thank everyone who made this series possible. We were thrilled to have incredibly high levels of attendance at our roundtables, from senior parliamentarians and civil servants to leading civil society organisations and industry groups to startups and global technology companies.

The varied and nuanced discussions were testament to the diversity of experience and depth of expertise involved and we look forward to working with all of these partners, and more, as we develop this important strand of work. Onward was established to come up with new ideas for the next generation and few policy debates embody that mission as keenly as the future of technology and data.

Support Onward

Onward is an independent, not-for-profit thinktank. We rely on the generous support of individuals and trusts, as well as partnerships with charities and businesses, to support our leading programme of research and events.

Individual and Trust donations

Onward's core programme of research is funded by individual and trust donations. If you are an individual or represent a philanthropic trust and would like to support our ongoing research agenda, we would love to hear from you. We do not accept corporate funding for research reports.

Please contact us on office@ukonward.com if you would like to donate by cheque or you can donate electronically using Onward's account details below.

UK Onward Thinktank Ltd
Not-for-profit company no. 11326052

Bank: Natwest
Account number: 21328412
Sort code: 50–10–05

Please note that Onward retains copyright and full editorial control over any written research it produces, irrespective of funding.

Partnerships with companies and charities

Alongside our research, we want to work closely with with charities, trusts and businesses, large and small, to further the political debate and bring the expertise and ideas from outside government to bear on the policymaking process.

If you would like to partner with Onward as a sponsor of one of our roundtable events or conferences, or to join our Business Leaders Network to engage further with our work, please get in touch at office@ukonward.com.

A commitment to transparency

Onward is committed to transparency and will publish the names of individuals and organisations who give us more than £5,000 of support each year. This is published on our website twice a year.

ONWARD

www.ukonward.com

office@ukonward.com
4 Millbank, London, SW1P 3JQ

Atos

www.atos.net

River Ouest, 80 quai Voltaire,
95877 Bezons cedex, France