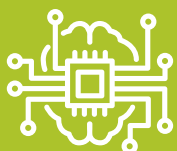C/M/S/

Law . Tax

Integrity and transparency

Ethical considerations for the evolving AI landscape

Artificial
Intelligence

'Artificial Intelligence (AI) is one of the most transformative forces of our time, and is bound to alter the fabric of society,' the European Commission's High Level Expert Group on Artificial Intelligence stated in its December 2018 report *Draft Ethics guidelines for trustworthy AI.*

The report outlines the bounteous ethical questions associated with AI. While AI's potential to alter societies and economies is indisputable, the fears and anxieties linked to it are equally forceful. At a roundtable dinner hosted by CMS in the City of London on 20 March, industry experts, think tank founders, and CMS representatives assessed the current and future AI environment, including the ethical and regulatory boundaries that could be put in place to ensure that these technologies deliver virtuous outcomes.

Amongst a number of key concerns, the roundtable experts evaluated the potential for AI to weaken solidarity and community ties, to be exploited by bad actors and whether regulation could inhibit innovation.

Clive Gringras
Partner
T +44 20 7067 3189
E clive.gringras@cms-cmno.com

Ian Stevens
Partner
T +44 20 7367 2597
E ian.stevens@cms-cmno.com

Dan Tench
Partner
T +44 20 7067 3518
E dan.tench@cms-cmno.com

Rachel Free
Of Counsel (Patent Attorney)
T +44 20 7067 3286
E rachel.free@cms-cmno.com

Brave new worlds

Ever since Alan Turing suggested that a computer could become capable of thought in his 1950 paper *Computing Machinery and Intelligence*, there have been inevitable questions around the ethics and regulation of AI. The legendary mathematician and computer scientist's revelations largely simmered in the background for some 60 years until the use of AI began to proliferate across industry and into the public consciousness over the last decade.

For years AI was seen as a fiction, whether written or on screen. It was portrayed as harmful or at least capable of harm; the sentient computer HAL in Stanley Kubrick's 1968 film *2001: A Space Odyssey* springs to mind. Or in the animated dystopian Pixar film WALL E, computer and robotic technology becomes so powerful that humans become physically and mentally lazy, no longer needing to perform everyday tasks for themselves.

Over the last decade, fiction has given way to actuality, though AI is still very much in the early stages of its evolution. In large part, it is employed to perform routine and volume tasks to heighten rapidity and precision; human interaction and supervision is typically woven into its processes.

From an ethical standpoint, AI has achieved quick wins. CMS powers its document review with systems such as Kira, and others in the legal sector are also embracing AI to achieve higher levels of speed, accuracy and cost efficiency.

But as AI becomes more sophisticated and powerful it has the potential for even greater good and, of course, harm. In this instance, legislators, regulators and the public are conscious that it must remain a benevolent force.

Over the last decade, fiction has given way to actuality
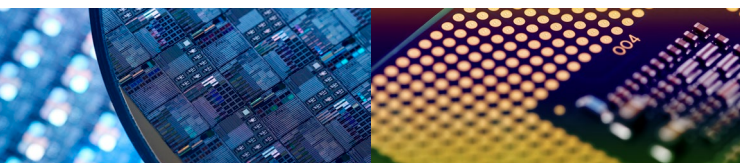
Reassuring the sceptics

From the introduction of the locomotive, to the advent of nuclear power and now to the influx of AI technologies, these ground-breaking moments have all struck at the heart of public anxiety. Current concerns are not without foundation, though perhaps overinflated, but the reality is that AI innovation is making waves in the economy. Matthew Price, Business Development Manager at Oracle Global Startup Ecosystem, says that a high proportion of the emerging businesses that his business unit supports are AI driven.

For many this will be a welcome phenomenon, but fears remain over machines taking over from humans. In March, reflecting on the recent crashes of two Boeing 737 MAX 8 aircraft, a Guardian article *Ethiopian flight 302: second new Boeing 737 to crash in four months* pointed to the 'uncertain interface between human and artificial intelligence'. It went on to suggest that insurers have deep concerns about the co-existence of autonomous cars and human drivers on our roads.

While many highlight the importance of the human guardianship of AI and automated technologies, this does not eradicate an individual's innate fallibility. Dan Tench, a litigation partner at CMS, suggests that too often there is an expectation that AI should achieve complete perfection. It is lambasted for errors, when human mistakes are understood or even forgiven; AI's higher levels of accuracy, efficiency and speed are frequently forgotten.

Moreover, AI tools are (or should not be) not left to their own devices. Where AI is deployed within an appropriate control framework, an individual or individuals are still, in large part, required to ensure that the technology continues to address what it is supposed to and that it doesn't deviate from defined parameters.

In autonomous weapons systems, a human is still required to pull the trigger or press the button. Complex networks of AI tools operating completely autonomously without any human intervention or supervision remains at present a pipe dream.

Ten key questions for the AI landscape

1. Humanity – Should we be concerned that AI could cause social fractures?
2. Employment – Could AI result in lower employment levels?
3. Regulation - Will further regulation empower or inhibit innovation?
4. Supervision – How should the relationship between human and machine be defined?
5. Integrity – How should we ensure that AI remains a force for good?
6. Confidence – How can trust in AI be enhanced with the public and authorities?
7. Transparency – should the public and authorities be given a fuller understanding of how an AI system works?
8. Freedom – How far can we push data freedoms to enable more powerful AI?
9. Accountability – Who is accountable for any harm that might be directly caused by AI?
10. Prejudice – How do we ensure that AI is neutral and free from bias?

The case for regulatory parameters

Debate over whether AI deserves its own defined regulatory and ethical boundaries has been heated, including at our roundtable discussion. A valid argument is that AI is already properly regulated given that it depends on the data it is fed. In the sense that the General Data Protection Regulation (GDPR) and other regulations already deliver limitations to the use of and processing of personal data, then should the use of AI require further rules and directives? GDPR already restricts automated decision making without human involvement. It also allows for consumers to challenge conclusions where automated tools are involved. But of course, GDPR addresses only *personal data* and not the many other petabytes of non-personal data that are fuelling the learning of AI systems.

Perhaps because privacy laws go only so far, the European Union (EU) has taken a proactive line in addressing the AI issue at a higher level than one merely addressing personal data. The European Commission's April 2019 report *Ethics guidelines for trustworthy AI states*:

'Trustworthy AI has three components: (1) it should be lawful, ensuring compliance with all applicable laws and regulations, (2) it should be ethical, ensuring adherence to ethical principles and values and (3) it should be robust, both from a technical and social perspective since to ensure that, even with good intentions, AI systems do not cause any unintentional harm.'

Many might claim that strong regulation would provide the platform for AI businesses to flourish in the same way that a strong regulatory framework has enabled London to blossom as a key global financial centre.

There is also the contention that regulation prevents big tech from racing ahead of smaller players, resulting in a lack of competition and less innovation. The other side of the argument is that regulation actually favours innovation from the largest players because these big businesses are able to absorb additional compliance costs, unlike many start-ups and growth companies.

One of the roundtable participants explained that regulators and legislators are understandably struggling to develop the necessary market insight and knowledge that comes naturally to those that operate within the industry: "Everyone agrees that regulation is good until that regulation starts to inhibit technology and profit. In the market, there's probably a sense that people writing the rules don't necessarily understand the technology as well as the market does. No company says it is anti-regulation until it inhibits what it does."

A view was expressed that regulation can have an important role to play in earning the trust of consumers and customers. For example, if European regulators establish defined parameters for AI technologies to operate in, it could put European AI developers at a distinct advantage in the global environment. "One angle is to build the brand of AI that consumers trust. If people think companies in Europe are building AI in an ethical way then people will put their money there and it will provide the gold standard for market."

However developers outside that regulatory regime may be able to experiment more widely, learn lessons, develop new solutions faster as they will not be constrained in the same way.

" Many might claim that strong regulation would provide the platform for AI businesses to flourish

Is AI-specific regulation required?

Use of technology, including AI, is already directly or indirectly regulated in some sectors, for example through the obligation to manage operational risk and treat customers fairly in the financial services sector.

The introduction of GDPR in 2018 is thought by some to put European businesses at an advantage. It has the potential to create European champions by building trust amongst consumers.

Dan Tench (CMS) says that regulators face especially tough ethical questions. Using the example of someone who is seeking a loan, he poses the idea that the market can simply apply the usual metrics on whether that person has the ability to make the necessary repayments or regulators can intervene to ensure market access for those that might traditionally struggle to access finance. "It is debatable whether that is for the public good," he comments.

But is AI-specific regulation necessary? Clive Gringras, Head of Technology, Media and Telecommunications at CMS, questions whether we should really be concerned with the data that is harnessed by AI, rather than the AI-technology itself. He suggests that data sets are already tightly regulated.

In *Perspectives on Issues in AI Governance*, a paper published by Google in 2019, the tech giant states: 'Harnessed appropriately, we believe AI can deliver great benefits for economies and society, and support decision-making which is fairer, safer and more inclusive and informed. But such promise will not be realized without great care and effort, which includes consideration of how its development and usage should be governed, and what degree of legal and ethical oversight — by whom, and when — is needed.'

The paper recognises the unique position that AI has in the annals of technological innovation. Its potential economic, societal and existential effects are nothing less than transformative. So compared to other technologies, it is not surprising that AI is discretely treated from an ethical and regulatory standpoint. "What is different about AI?" asks Clive Gringras (CMS). "It's the potential for harm with loss of control" answers Tom Marshall, Legal Counsel at Samsung.

The loss of control might come in various guises. Several roundtable participants recognised the potential for unintentional cartel activity driven by the use of automated algorithms. If businesses in a certain industry use similar algorithms to identify the market price or respond to market changes, this could result in artificially high price points for consumers. Where specific algorithms communicate with other industry players through AI, there is the potential for cartel behaviour irrespective of the intent by the company or human that operates the technology. Where a network of sophisticated machines communicate over market practice and pricing, there is the possibility for distorted behaviour or a deviation from free market economics. This is a topic that CMS commented on in Kabir Garyali's thoughtful article, "*Is the competition regime ready to take on the AI decision maker*".

Freeing the data landscape

The issue of data freedoms is at the heart of the AI debate. Smart cities or intelligent transport systems, for example, depend on efficient sharing of data so that automated technologies and AI tools can perform to expectations.

For AI to function in the interests of public good, there are some who argue it would benefit most from open source environments where data is leveraged for a positive outcome. Sharing data for the purposes of curing cancer would be an obvious example, but finding a balance between discretion and expediency is not straightforward. One of the roundtable participants explained that "good actors" will "move the needle". Where the public and authorities recognise AI's benefits, it will enable greater freedoms.

"There needs to be a natural ceiling where rights and access can be granted," says Ian Stevens, a partner in CMS's technology team. "Governments need to find a way to allow access to better data sets, but recognise that these can be used for harm as well as good."

The Open Data Institute (ODI) announced in November 2018 that a series of pilot projects supported by the government's Office for Artificial Intelligence would be launched in the UK to assess the potential positive effects of data trusts. The ODI defines a data trust as 'a legal structure that provides independent stewardship of data'.

Following initial research, the ODI's Chief Executive Jeni Tennison said in April 2019: "We only unlock the full value of data when it gets used, so we really need to find good ways to share data more widely without putting people at risk." Business Secretary Greg Clark added: "Access to data is pushing forward huge technological change that can benefit our economy, provide better services to consumers and lead to the creation of new highly skilled jobs."

❝

Governments need to find a way to allow access to better data sets, but recognise that these can be used for harm as well as good

Assigning accountability and liability

The use of AI-specific technology provides clear benefits in the tracing of decisions and processes. The lack of transparency that comes with human thought and unrecorded conversations is eradicated through the use of machines. Audit trails should be much clearer. Yet, what happens when decision making processes become fully automated? Humans taking the ultimate or final decision provides a level of comfort that the public expects or is used to. Commercial air passengers still expect a pilot to observe the controls and data provided by the aircraft, even if much of the flight is automated. In this instance, the pilot and airline is still typically accountable if an incident occurs, although the recent crashes involving Boeing 737 MAX 8 aircraft have raised questions about whether manufacturers themselves could be held liable.

In the game of man versus machine, there are natural arguments in favour of the machine. A fully synthesised tool should in theory be free of prejudice and bias. An AI tool should make unemotional decisions based on hard data. Yet it still depends on the data sets that are presented to it, a process that is largely managed by humans. AI tools might simply reflect the biases of their developers or architects. They might be influenced by the data they are fed or draw the wrong conclusions from data sets. It might then be argued that AI tools will simply continue historical prejudices and partiality.

Could we then use AI to detect bias? Ashton East is General Counsel at Behavox, the AI-driven platform that addresses behaviours in the workplace. He points out that AI can be used to detect human bias. He also notes that if AI is used to detect bias in other AI then a situation arises which is like opening an "endless Pandora's box".

Dan Tench (CMS) warns that AI could be very effective in reinforcing bias. He suggests that AI could negatively influence the recruitment process by identifying that a high proportion of recent recruits were educated at a particular school or university. Ethical processes and regulation should be put in place to ensure these glitches do not creep in. Ashton East (Behavox) also spoke about the dangers of AI reinforcing bias when he explained that "you don't want AI used to implement prejudice and you need ethics to prevent that." He comments, "you remain liable for discrimination regardless of how you arrived at that result."

A greater emphasis on ethical behaviours should not be seen as a burden. Birgitte Andersen, Chief Executive Officer & CoCreator of the Big Innovation Centre, says that good governance provides the basis for more external investment and trust.

> In the game of man versus machine, there are natural arguments in favour of the machine. A fully synthesised tool should in theory be free of the prejudice and bias. An AI tool should make unemotional decisions based on hard data.

The effects on humanity

While the good that is stemming from or will branch from AI is irrefutable, it still presents a wealth of further ethical questions. AI is clarifying points and enabling decisions to be made much quicker. It is elevating services to a personalised level, but it could have negative wider impacts. "If we build our AI correctly it can actually be less biased and more ethical than human decision making, as long as we train it accurately and use the right data sources. The power of AI is in our hands, and we can chose how to use it, and I believe this is what we should focus on - empowering businesses with the right ethical underpinnings to develop AIs that have a positive impact on business, the economy and society." comments Marja Verbon, COO at Jump.Work, a talent platform that uses behavioural data to personalise job searches for professionals.

If lives become so personalised then society has the potential to become fragmented. People lose the sense of commonality. Rachel Free, a patent attorney at CMS, who holds an MSc in Artificial Intelligence from Edinburgh University, talks of the social ties that could be broken by AI. "I used to go to the baby clinic, where I met all the other mums and we were treated in the same way," she says. "We all felt a sense of community and solidarity, but that could change with AI creating more individualised services where you might only see the midwife when AI predicts it is needed."

Already technology has had the effect of fewer shoppers on the streets and a scarcity of visitors to cinemas and other entertainment locations. Home deliveries and the advent of media streaming have dissuaded people from venturing from their dwellings. "If everything is built around the individual, then you lose social capital," comments Birgitte Andersen (Big Innovation Centre).

For this and other potentially negative outcomes, Abhijit Akerkar, Head of Applied Sciences, Business Integration at Lloyds Banking Group, says that the interface between humans and AI is especially important. He points to the use of AI in healthcare where doctors will still make the ultimate diagnosis or decision. AI is simply there to augment and assist conclusions and evaluations. Moreover, one of the roundtable participants identifies the key ethical question: "All businesses run on trust, but doctors are educated in ethics from the beginning and computer scientists are not."

Where AI is allowed a certain level of autonomy, it should still come under regular checks and supervision. Ian Stevens (CMS) recognises that this is fundamental to ensure that AI does not deliver negative effects: "If you deploy AI in the loan environment, for example, you will need to carry out checks every now and then to ensure that it is performing as it is expected and that it is still making appropriate decisions." CMS Senior Associate in Banking and Finance, Gurminder Muker, explained the legal and fiscal implications of AI in the loan market in his thoughtful article, *"Big Data, AI and the life of a corporate loan"*.

> **Already technology has had the effect of fewer shoppers on the streets and a scarcity of visitors to cinemas and other entertainment locations. Home deliveries and the advent of media streaming have dissuaded people from venturing from their dwellings.**

Building transparency into the AI landscape

For those without an intimate understanding of learning technologies and AI, it remains a considerable challenge to win trust. Greater transparency could earn more faith and confidence in AI technologies. Abhijit Akerkar (Lloyds) believes this isn't just about public disclosure, but also about giving confidence to employees to trust the recommendations made by machine learning models. It enables employees to have a clearer understanding of when AI tools are working as they should and when they are not.

But Ashton East (Behavox) says that software architects themselves typically struggle to demystify the technology. "If you ask a software developer to explain how it works, they will eventually give up and tell you it's magic," he jokes. This is because documenting how an AI process works is resource intensive and often only understood by AI experts.

The problem with earning more trust from regulators and the public, is that the AI world looks so opaque from the outside. "What does transparency of a really complex algorithm or process give you?" asks Ian Stevens (CMS). "Because the average person on the street might still not understand."

The issue of transparency leads the market down a difficult path. It becomes a delicate ethical question where innovation and transparency are uneasy bedfellows. "If you have developed an impressive AI product, should you then have to reveal your trade secrets?" asks Clive Gringras (CMS). His assertion is just one of many that regulators and legislators will continue to consider as the world absorbs the transformative effects of AI in all its forms.

> The issue of transparency leads the market down a difficult path. It becomes a delicate ethical question where innovation and transparency are uneasy bedfellows.