

Data and the Digital Economy – Auditing AI and Algorithms

February 2021



Overview

Last year saw an accelerated growth in the use of data and data driven technologies, such as AI, algorithms, and machine learning, in response to the Covid-19 pandemic. But even before 2020 there was already a widespread use of these technologies. For example, the content we see on social media platforms or watch on streaming services, decisions made on extending credit, the selection of candidates in recruitment processes and the conduct of due diligence can all be carried out by AI and algorithms. It is vital that the use of these technologies is managed effectively, ethically, securely and in accordance with regulatory requirements.

Auditing AI and algorithms – how to best prepare for the future: webinar

In a webinar hosted by Travers Smith with techUK on 11 February 2021 we examined the factors that businesses should consider when deploying algorithms and AI, the legal and other risks they face and how the need to check the use of AI and algorithms could lead to the development of an effective audit and assurance framework.

The webinar was chaired by Nicky Morgan, former Digital Secretary, now a consultant at Travers Smith and our speakers were Travers Smith IP and Technology Partner, James Longster, Head of Legal Technology at Travers Smith, Shawn Curran, Maria Axente who is the Responsible AI and AI for Good Lead at PwC and Alister Pearson, Senior Policy Officer at the Information Commissioner's Office.

This briefing sets out the broad points that were made and a [recording](#) of the whole webinar is available on our website.

The regulation of AI and algorithms

As James Longster said, the laws relevant to AI and algorithms are many and varied although there is not, at this stage, a law of algorithms per se. James advised that organisations need to look at the context of how algorithms are being used to identify the relevant laws for any legal assurance or audit programme. He focused on two main areas – the laws and regulations relating to equalities and the laws around the processing of personal data.

In relation to equalities laws which, in the UK, primarily means the Equality Act 2010, the key point is that people are protected from discrimination caused by technology in the same way as any other form of discrimination. The discrimination is likely to arise from the data used to train the algorithm. As James said, "If you put bad data in, you are going to get bad decisions."

The Equality Act does not apply just in workplaces and applies, for example, in the provision of services to customers. This is an area likely to see an increased use of technology solutions which brings with it a risk of contravening section 29 (2) of the 2010 Act.

In terms of managing the risk of using such technologies James recommends that it will become increasingly important for an organisation to be able to demonstrate the steps, from a legal perspective, that have been taken at the design and, if relevant, learning phase of the technologies. This will ensure that if questions about the use of and decisions made by the algorithm are raised later then the work carried out earlier can be used to assist in the assurance or audit process.

On the issue of the use of personal data in accordance with UK or EU GDPR, the advice is that organisations should go back to basics to assess the requisite lawful basis on which they are relying to process personal data - e.g. with consent or in the performance of a contract or using the legitimate interests basis. In the context of the use of AI, the organisation needs a lawful basis on which to process not only the personal data which is used by AI to make a decision, but also any personal data which is used for the training of that AI. The lawful basis for each type of processing might not necessarily be the same one. James also pointed out that if an organisation is processing 'special category data', it won't be able to rely solely on legitimate interest as its lawful basis for such processing.



The processing of personal data also needs to be fair and transparent. So, for example, will the AI be processing personal data in a way that an individual would not reasonably expect? If the processing breaches the Equality Act then it is likely to fall foul of the fairness principle in data protection law as well. In relation to the transparency requirement, an organisation's privacy information notice should reflect what the AI is actually doing, including whether its use results in a solely automated decision-making process. James recommended that a data protection impact assessment is a good way of keeping a record of data processing activities, the risks associated with such activities and the steps taken to mitigate them, and documenting the decision-making process, and can ultimately form the basis of any subsequent assurance or audit process.

James also touched on the need for organisations to consider how their suppliers are using AI and algorithms, particularly where they are sub-contracted to make decisions on behalf of an organisation. He recommended that organisations get contractual comfort around the dataset that a supplier has used to train its AI, and where data portability is an important commercial objective for an organisation, who owns the data, not only in case the supplier relationship changes later but also in order to justify how decisions are made where challenged, eg under the Equality Act. James also made the point that the code that sits behind these technologies is the same as any other proprietary software and can be protected by copyright, so it can also be a valuable asset for any business.

The Information Commissioner's Office (ICO) view on processing personal data in the context of using AI and algorithms

As Alister Pearson from the ICO made clear, the ICO's interest in the use of AI stems from the fact that in a vast proportion of cases where AI and algorithms are being used to make a decision, personal data is being processed. As the data protection regulator, the ICO has a responsibility to regulate how personal data is being processed in that context.

In early 2019 the ICO started to develop an [AI auditing framework](#) to provide a clear methodology on the auditing of AI systems, how the ICO thinks data collection and processing applies to the use of AI, and to support the ICO's investigations and assurance teams with their work. The guidance was published in July 2020 and the ICO will shortly be publishing an AI Risk Toolkit.

Alister made the point that, for the ICO, there is no agreed universal definition of AI and there is no definition of AI in data protection law – so their approach is to "refer to it as an umbrella term for a range of technologies and approaches that often attempt to mimic human thought to solve complex tasks." As far as the ICO is concerned "...what is important is to have a basic understanding of what AI is and what it can do, and also what risks it can create and exacerbate [in relation to] individual rights."

The July 2020 guidance on AI and data protection is designed to provide organisations with what the ICO considers to be best practice when it comes to developing, designing, and deploying AI systems that process personal data. The other element is an interpretation of the law which may apply to AI systems. The ICO holds that the law requires organisations in the vast majority of cases, where they intend to use AI systems to process personal data, to carry out a data protection impact assessment because such use is likely to result in high-risk processing.

The ICO has taken a risk-based approach in the guidance which means that organisations should assess the risks and "implement appropriate and proportionate controls to address those risks". The guidance is not a step-by-step guide for organisations to follow in order to secure 100% compliance. It really means that they need to look at the particular context that the AI system is being used in, assess the risks involved and consider how to mitigate those risks and implement appropriate and proportionate controls. The guidance also talks about the new risks to individual rights from AI and some possible measures that organisations can take to address those risks.

The three key areas for organisations processing personal data in the context of using algorithms to focus on are "documentation, transparency and understanding". This requires organisations to document the decisions made about risks and proportionate controls which can then be used if any third party auditors, including the ICO, later ask how an organisation is complying with the applicable law. Transparency means being transparent about how a person's personal data is being used and the likely impact on them from that use. Finally, the ICO is keen that organisations should really understand the AI they have in their hands, be able to say what capabilities their system has, what they want the AI system to do, and also to understand what the risks are to individual rights and how to address them.

The ICO's AI risk toolkit is designed for risk practitioners and it will contain risk statements produced from the guidance. These risk statements are divided into the different domains of data protection law ie. risk statements associated with lawfulness, fairness, transparency, security and all of the data protection principles. Both the guidance and the toolkit aim to help organisations to assess why their particular use of AI might create or increase each of these risks. There will also be a list of suggested practical steps that organisations can take to address the risks that are presented, and a work plan that they can use. It is likely that a beta version of the toolkit will be published within the next few months. The ICO is interested in hearing people's views about the AI risk toolkit and invites individuals to email AI@ico.org.uk to find out more.

Why does the way businesses use AI matter?

The use of AI is often referred to in conjunction with words such as 'responsible' or 'ethical'. Maria Axente from PwC pointed out that AI is a "a major disruptive force that promises a lot of benefits, especially in a business context, [but] it comes with a set of risks and we don't talk about risks of AI enough."

Her argument is that "in order for us to really grasp the full benefits and the risks of AI, we need to have a top-down approach, not only a bottom-up approach and this is why responsible AI looks at issues in the context of an end to end governance programme, not just the development of the product in isolation. Appropriate proactive risk management embedded throughout the lifecycle, not just the one-off exercise at the end or the beginning of the project." Maria also argued that the "...integrating of core ethical principles in the same manner is absolutely critical."

She referenced a PwC survey in late 2020 which asked several thousand clients, based both internationally and in the UK how they felt about AI, what was the rate of adoption and what did they think about issues such as governance and ethics. The survey found that the pandemic had significantly accelerated the adoption of AI.

Maria said that four in 10 organizations now have a strategy that looks at proactively identifying and managing risk across AI operations. While there is no consistent approach to embedding ethics, many organisations in the UK (almost 40–45%) have either ethical principles or boards that are aware of the key ethical considerations that apply in their own business.

However, the survey also revealed that "...the top consideration among the ethical principles was not fairness or privacy. It was in fact reliability and security, which means that businesses are still, to a certain extent, concerned about the quality of the output delivered by AI while embedding it in current processes that are by no means ready to embrace that disruption." In Maria's view the key ethical principles are fairness, privacy of data, accountability and transparency.

Maria also argued that all of the C-suite in a business are critical in assessing how that business handles AI and it is a competitive advantage to have executive leadership who understand how to ensure AI is used ethically and responsibly. Handling AI well is not only important from a safeguarding and brand reputation standpoint, but also positively impacts on improving business performance.

The challenge is to bring together the conversations on governance of AI, behaving ethically and managing AI risk. Maria pointed to the maturity of conversations about the use of AI in financial services businesses and, in particular, their understanding that every part of the business from compliance to frontline business to HR have to work together on deploying AI successfully and responsibly. As Maria said, the use of AI "triggers a profound organisational change."

On the question of audit Maria believes that "...we need to be very specific about what we audit for and what we audit against. I think having that clarity and also support from the legislature will give AI audit or assurance ...a boost".

It is relevant to ask whether new laws mandating the audit of AI and the use of algorithms are likely. Given that AI is difficult to even define it seems unlikely AI audit laws are on the horizon – instead much will be left to regulators such as the ICO and the FCA as well as overarching governments' strategies. However, the ethical, responsible and transparent use of AI and algorithms is definitely on the Westminster agenda. Although the rules on the use of data by AI don't differ between the public and private sectors it seems likely that the bar on transparency may be set higher in the public sector, particularly given the large datasets the public sector has which can contain very detailed personal information such as medical records.

A practical example of the use of AI

Our webinar also included an overview given by the Head of Legal Technology at Travers Smith, Shawn Curran, on how the firm has used AI to assist with work for clients.

As Shawn explained, two years ago Travers Smith looked at the market for "what we call contract review systems which are essentially AI systems that help their clients and their lawyers assess risk across contract data." Shawn and his team were unable to find the right product so decided to build the firm's own AI capability. The team's extensive work on the labelling and categorising of any data (not just personal data) for the purpose of helping to train AI tools, led to the development of a data labelling platform called Etatonna.

Shawn refers to the platform as "the value capture layer between lawyer and AI." For any other businesses it can be viewed as the value capture layer between the domain experts in the underlying machine learning model. This is because it has capacity to store training and input data in its original form, albeit organised in such a way that you can understand, identify, find and segregate the data that is used to train the AI tool. This makes it much easier to remove and ringfence the data as necessary.

The Etatonna platform allows a rules-based approach to be developed (used for example in reviewing force majeure clauses where key words will be repeated every time) but crucially its ability to label and organise the underlying data, could allow for example, for an organisation to cut through vast amounts of datasets in order to identify and delete a data subject's personal data if they so requested, and for the underlying model to regenerate.

Travers Smith has open sourced the Etatonna platform under the GNU General Public license. Please contact shawn.curran@traverssmith.com if you would like to know more.

Key contacts



Nicky Morgan
Consultant, Technology
Sector Group

nicky.morgan@traverssmith.com
+44 (0) 20 7295 3279



James Longster
Partner, Commercial, IP &
Technology

james.longster@traverssmith.com
+44 (0) 20 7295 3496



Shawn Curran
Head of Legal Technology

shawn.curran@traverssmith.com
+44 (0) 20 7295 3381

The information contained in this booklet is based on information available at the time of publication. Travers Smith LLP has made every effort to ensure the accuracy of the information in this booklet.

However, readers should always obtain professional advice before deciding to take any action (or not, as the case may be) in relation to information contained in this booklet. Travers Smith LLP is a limited liability partnership registered in England and Wales under number OC 336962 and is regulated by the Solicitors Regulation Authority. The word "partner" is used to refer to a member of Travers Smith LLP. A list of the members of Travers Smith LLP is open to inspection at our registered office and principal place of business: 10 Snow Hill, London EC1A 2AL.