ico.
Information Commissioner's Office

Select Committee on Artificial Intelligence – Submission from the Information Commissioner

Contents

Executive Summary

1. The Information Commissioner's Office (ICO) is responsible for promoting and enforcing the Data Protection Act 1998 (DPA) and welcomes the opportunity to respond to the Committee's Call for Evidence.

2. Our interest in Artificial Intelligence (AI) technology lies in the processing of personal data. The automated processing of personal data without appropriate checks has been a privacy concern for many years. Successive data protection laws including the current DPA, and the EU General Data Protection Regulation coming into effect in May 2018 require organisations who process personal data to comply with a number of important principles. These principles help to mitigate privacy risks and provide certain rights to individuals.

3. The rapidly increasing use of AI, although a type of automated processing, presents its own unique risks. Whereas 'traditional' processing involves a human-being making decisions as to how and for what purpose data is processed, AI enabled processing involves a computer making these decisions with little or no human oversight. There are questions to be answered as to whether a computer can show the same level of empathy and reasonableness in making often significant decisions about individuals.

4. People have mistrust in the use of AI technology. We believe the key elements for preparing the public are: transparency- providing individuals with information about the implications and likely outcomes from the use of AI; control – ensuring a significant element of human oversight and intervention through knowledgeable, appropriately senior, dedicated staff; and effective regulatory oversight – organisations taking a number of compliance steps including regular reviews and privacy impact assessments.

5. The use of AI raises ethics as well as privacy concerns. Data protection law, especially new requirements contained in the soon to be implemented General Data Protection Regulation, go a long way in tackling these concerns. Ultimately, data protection is about the relationship between those that process personal data and the people whose data is being processed. If those who use AI do so fairly then many of these concerns about its use will be addressed. This will be to the benefit of impacted individuals and society as a whole.

Introduction

6. The Information Commissioner's Office (ICO) has responsibility in the UK for promoting and enforcing the Data Protection Act 1998 (DPA), the Freedom of Information Act 2000 (FOIA), the Environmental Information Regulations 2004 (EIR), the Privacy and Electronic Communications Regulations 2003, as amended (PECR), and the eIDAS Regulations (2016). We also deal with complaints under the Re-use of Public Sector Information Regulations 2015 (RPSI) and the INSPIRE Regulations 2009. We are independent of Government and uphold information rights in the public interest, promoting openness by public bodies and data privacy for individuals. We do this by providing guidance to individuals and organisations, solving problems where we can, and taking appropriate action where the law is broken. We welcome the opportunity to respond to your call for evidence and are grateful for your consideration of this submission.

7. The Information Commissioner's interest in artificial intelligence (AI) lies primarily where its use involves the processing of personal data. Personal data is defined in the DPA as "data which relate to a living individual who can be identified from those data, or from those data and other information which is in the possession of, or is likely to come into the possession of, the data controller."[1]

8. The processing of personal data by automated means has always posed a privacy risk for individuals. Since the UK's first Data Protection Act in 1984, successive legislation has required organisations to take certain steps to mitigate these risks whilst giving individuals' specific rights over their own personal data.

9. AI, although a type of automated processing, creates its own unique risks that are potentially even more intrusive to individuals' privacy. Unlike other forms of automated processing, AI programs don't linearly analyse data in the way they were originally programmed. Instead they learn from the data they have already analysed in order to respond intelligently to new data and adapt their outputs accordingly. This brings the possibility of AI-enabled technologies making significant decisions about people, with little or no human oversight. This evidence makes clear that data protection rules have become more relevant than ever, and if applied effectively can help to protect individuals, mitigate risk and to allow society to reap the benefits of AI technology.

[1] Data Protection Act 1998, S1 (1)

Historical Context

10. The use of automated data processing without appropriate checks and balances has been a privacy concern for many years. The OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (1980) represent one of the earliest international data protection instruments. The Guidelines' Explanatory Memorandum states:

> "As far as the legal problems of automatic data processing (ADP) are concerned, the protection of privacy and individual liberties constitutes perhaps the most widely debated aspect. Among the reasons for such widespread concern are the ubiquitous use of computers for the processing of personal data, vastly expanded possibilities of storing, comparing, linking, selecting and accessing personal data, and the combination of computers and telecommunications technology which may place personal data simultaneously at the disposal of thousands of users at geographically dispersed locations and enables the pooling of data and the creation of complex national and international data networks."

11. Subsequently, the current European Data Protection Directive (95/46/EC), adopted in October 1995 - which still forms the basis of the UK's current data protection law - states in its second Recital:

> "Whereas data-processing systems are designed to serve man; whereas they must, whatever the nationality or residence of natural persons, respect their fundamental rights and freedoms, notably the right to privacy, and contribute to economic and social progress, trade expansion and the well-being of individuals."

It's worth pointing out that the OECD Guidelines and the Directive 95/46/EC were drafted during an era of stand-alone computers and basic telephony systems with very limited functionality. The internet was not widely used in business or for personal use. Most of the technology companies whose services we are so dependent on today had not yet been founded. The roots of social media, wide-spread data-sharing, Big Data and artificial intelligence (AI) were just forming. Therefore we believe that both the OECD Guidelines and the Directive were highly prescient, and were right to acknowledge the threats as well as the opportunities of information technology. At the time of drafting, to many, the risk of mankind serving technology, and not vice versa, must have seemed the stuff of dystopian science fiction. However, the use of AI has the potential to bring this risk closer to home.

12. DP laws have recently been modernised to tackle the challenges of technology in the twenty-first century. The EU General Data Protection Regulation was passed in 2016 and will come into effect in May 2018. A UK Data Protection Bill will also be introduced to cover national implementing measures and areas where member states are allowed to derogate. The legislation increases individuals' rights – for example rights in relation to profiling and introduces new concepts such as data protection by design and data protection impact assessments.

The Pace of Technological Change

13. Question one. What is the current state of artificial intelligence and what factors have contributed to this? How is it likely to develop over the next 5, 10 and 20 years? What factors, technical or societal, will accelerate or hinder this development?

14. There will others who are directly involved in technical development who may be better placed to comment on the technical factors affecting development in this area or how it is likely to develop over the next few years. However, we are aware of a general increase in the adoption of AI technology and how swiftly this is becoming a mainstream technology, with a wide range of potential uses in both the public and private sectors. The volume and range of datasets available, increases in computing power and online storage are rapidly driving forward these advances. The Information Commissioner published a report on the implications of AI for data protection earlier this year[2]. AI will also feature as a priority area in the Commissioner's new Technology Strategy, which will be published later in 2017.

15. A lack of public trust could be a factor that hinders the take-up of AI, particularly in personal data processing contexts. ICO research conducted in 2016 found that only one in four UK adults trust businesses with their personal information.[3] Trust may be even more lacking when it comes to the use of AI and automated processing more generally. There could be a point at which public suspicion - arising from a lack of control and understanding – undermines trust and inhibits the take-up and development of new services – particularly digital ones.

16. The lack of public trust could be compounded by the perception that decisions based on AI are opaque at best and have unfair or otherwise undesirable consequences for people. Questions arise such as what criteria is the computer using to carry out certain actions? How will I

[2] https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf
[3] https://ico.org.uk/media/about-the-ico/documents/1624382/ico-annual-track-2016.pptx p.10

know if the 'computer is wrong'? And what can I do about it? A useful automated processing albeit not AI example is that of the Border Systems Programme use for security purposes. There are public perceptions that the system is setup to target individuals on the grounds of race or religion, where in fact this is a misunderstanding. A lack of information as to how the system works likely contributes to this mistrust.

17. The responsible use of AI for the electronic delivery of government services is very important – of course we want the public to provide accurate data and to take up willingly the new services that technology facilitates. This depends on the transparency, control and oversight that we will elaborate on later in our evidence.

18. Recent UK research[4] found that 55% of UK consumers find AI 'creepy'. However, as we explain later in our evidence, there are generally ways of mitigating the risks and of keeping this mistrust at bay. However, it is possible that some uses of AI will always be unacceptable. Should an individual's innocence or criminality ever be automatically inferred using solely automated / AI means? [5] If so, where should the legal and ethical limits to the deployment of such technology lie?

19. Question two. Is the current level of excitement which surrounds artificial intelligence warranted?

20. The Information Commissioner believes that the current level of excitement is warranted but needs to be tempered by caution and a comprehensive assessment of the risks and benefits. There are certainly significant benefits to the use of AI but there are also data protection implications. We discuss this in detail in our research paper.[6]

21. The use of AI presents some novel challenges for data protection safeguards. A classical paradigm in data protection is where a data controller (the person, usually an organisation, who decides the purpose for and manner in which the data is processed) processes information about an individual for a particular purpose – for example to work out a housing benefit claim. A human being will work out what information is necessary to process the claim, where it should come from and how it should be analysed to produce the right result. The technology used will be essentially inert and will only process the

[4] https://www.research-live.com/article/news/half-of-uk-consumers-find-artificial-intelligence-creepy/id/5024372

[5] https://arxiv.org/pdf/1611.04135v1.pdf

[6] https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

information in the way it is programmed to. Once the claim is processed, a human being will deal with any ensuing disputes or queries, hopefully using the very human principles of reasonableness, fairness and with perhaps an element of empathy.

22. An AI-enabled scenario can be different to the above scenario in several key respects. Whilst issues of data controller responsibility and purpose might be the same, decisions over the sources of the data and the methods used to analyse it could be taken by the AI-enabled devices themselves. This in turn has implications in terms of compliance with other data protection rules, such as transparency, fairness, necessity, relevance and adequacy. Although AI is reportedly becoming more intelligent there are also issues over whether a machine could really display the reasonableness and empathy that can be needed to deal with individuals. This illustrates the importance of human supervision and intervention when AI is in use – we discuss this in greater detail below.

Impact on Society

23. Question three. How can the general public best be prepared for more widespread use of artificial intelligence?

In this question, you may wish to address issues such as the impact on everyday life, jobs, education and retraining needs, which skills will be most in demand, and the potential need for more significant social policy changes. You may also wish to address issues such as the impact on democracy, cyber security, privacy, and data ownership.

24. It is clear that the use of AI is increasing and is being used to make decisions that can have significant impact on people. Examples of this include the use of AI in internet counter-terrorism surveillance, offender management, credit referencing and on-line dispute resolution. It seems likely that the use of AI to analyse personal and non-personal information will continue to expand into more areas and to have a more significant impact on peoples' lives.

25. In our view the key elements for preparing the public for the more widespread use of AI are transparency, control and effective regulatory oversight. We consider each of these elements in turn.

Transparency

26. It is a basic and crucial requirement of data protection law that – in normal circumstances – people should be aware of such matters as who is collecting their information, how it will be used and whether it will be disclosed to a third party. This information is usually communicated to the public through an organisation's privacy notice. Even where a data processing operation involves the use of AI it should still be possible to provide this basic privacy information. However, current data protection law also contains provisions intended to protect individuals against the potentially negative impact of automated decision-making, including the use of AI.

27. The General Data Protection Regulation (GDPR), which will be implemented in the UK in May 2018, places more emphasis than the current law on automated decision making, when used for purposes such as profiling an individual – for example to target behavioural advertising.

28. In terms of transparency, in certain circumstances, the legal requirement under the GDPR will be for individuals to be made aware that automated decision making is taking place, to be provided with *meaningful information about the logic involved*, as well as the significance and the envisaged consequences of the data processing. This is where the use of AI as part of a personal data processing activity poses real challenges in terms of transparency and intelligibility to the public; what is meaningful information about AI? The problem is that the 'math' behind the algorithms used in AI would only be understandable to a limited number of experts and it would be very difficult for the vast majority of members of the public to challenge an AI-supported automatic decision on the grounds that its outcome is unwarranted, unfair or otherwise detrimental.

29. Individuals not being able to challenge such decisions would mean that – as the use of AI develops – there is the possibility of a widening rift of understanding between the public and the organisations that are using AI to make decisions about them. It may be more realistic for individuals to be provided with information about the implications and possible outcomes of the AI, rather than detail of the algorithm itself. In reality, it could be very difficult for members of the public to exercise their legal rights – and to be protected from the possible excesses of AI - without some form of expert mediation – we discuss the form that this might take in our comments on regulation below. Transparency will remain important but must be complemented by other effective safeguards.

Control

30. The second main right that individuals enjoy in respect of automated decision making, including the use of AI, is the right not to be subjected to a solely automated decision making process if the decision has 'legal' or a 'similarly significant' effect on an individual. The relevant provisions in the GDPR are complex, but in certain circumstances the individual also has a right to have an automated decision subjected to human scrutiny, to express his or her point of view and to contest the decision.

31. It is important to be aware, however, that it seems likely that organisations such as large e-commerce sites that use AI for purchaser – vendor dispute resolution will deal with a very large number of cases. It could be a challenge therefore for companies like this to offer complainants a second decision, taken using human intervention. There are also issues around how these – and other companies make individuals aware that AI is being used. Clearly individuals cannot use their 'automated decision making' rights unless they know automated decision making – possibly involving AI – is in use.

32. The rights in data protection law are potentially very powerful in respect of AI-based decision making. They would mean, for example, that if a Credit Reference Agency (CRA) recommends that a credit grantor turn-down an application for a loan – based on an automated decision – then the person applying for credit would be able to contact the CRA if he or she considers the decision to be unfair, and the CRA would have to ask a 'real person' to re-assess the factors that were used to make the original decision; of course the outcome might be the same. However, the point is that the law recognises the risks that AI and automated decision-making can pose and gives people a 'human defence' against this.

33. It is worth noting that the ICO currently receives very few complaints about AI or automated decision-making – this suggests that on the whole these technologies are being used responsibly and with reasonable outcomes for individuals. (Or, on the other hand this could be the result of a lack of public awareness.) However, we expect complaint numbers and volumes of queries to rise as the use of AI becomes more prevalent and moves into potentially more controversial uses of data – for example using social media data to predict an individual's credit score.

34. We believe that the element of human intervention addressed above is particularly important in the specific context of AI. A unique aspect of AI is that algorithms can 'teach themselves' and develop, based on their 'experience' performing a particular task. This can of

course have positive social consequences – for example an algorithm used to select particular travellers for counter-terrorism checks at airports – could become more accurate in the light of experience, leading to fewer false-positives and minimising collateral privacy damage. However, there is a danger that as technology 'makes its own rules' the results of its use could deviate from intended outcomes.

35. It is very important that organisations using AI applications review periodically the consequences of their use on the individuals whose data they are analysing and ensure the processing activity has not deviated from its intended purpose and is not having unintended consequences. Ensuring that organisations have rigorous processes underpinned by knowledgeable dedicated staff, including data protection officers with the correct level resources and organisational influence, will also be important.

Regulation and organisational accountability

36. As we have explained above, current data protection law and the GDPR both contain features that are highly relevant to the use of AI in personal data processing contexts. Appropriate use of AI for the processing of personal data depends on going through a series of compliance steps. We have included a list of such steps as an annex to this submission. Such checks, however, should not be viewed as a 'tick box' process and should be considered as comprehensively as possible likely through the use of ongoing privacy impact assessments.

37. As noted above, the requirement for organisations to be transparent about the uses of AI will have limitations for individuals. This therefore highlights the importance of ensuring organisations are accountable for their use of AI to data protection authorities. The concept of algorithmic accountability is important - organisations will need to provide evidence of how they have assessed and audited the impact and effects of the AI they have deployed. This may require the development of automated tools and new audit methodologies.

38. The Commissioner recognises that she will need to recruit more technical experts to audit and investigate issues related to AI. It will also be important that the market provides more services that audit AI – this also fits with the concept of 'certification' in GDPR - where the Commissioner will be able to accredit expert third parties to provide data protection certification that demonstrates compliance with the law.

39. The Information Commissioner recently completed an investigation into the trial of a service provided by the AI company, Google

Deepmind, to the Royal Free Hospital[7]. She concluded that the Hospital breached the Data Protection Act and required an undertaking to be signed to address non-compliance. The findings highlighted the importance of transparency, rigorous privacy impact assessment, robust contractual arrangements to prevent the re-use of patient data and verifying processes in practice using audits.

40. Despite robust data protection compliance, the law only takes us so far. We believe that it can be highly challenging to apply certain data protection concepts such as fairness and relevance to advanced AI applications. For example, empathic computing involves the use of AI to examine an individuals' on-line behaviour. It considers the vocabulary individuals use, the way they input type and the pictures they look at longest in order to assess that individual's mood and deliver content accordingly. This certainly involves the processing of personal data and therefore engages data protection law. However, whilst the pure data protection compliance aspects of using AI in empathic computing and other contexts can be addressed using the compliance steps outlined in the annex, the use of AI raises wider ethical issues of significant public interest.

41. Data protection law deals well with data processing activities – including those using AI – when the information being processed is about individuals and has an effect on those individuals. However, the broader social effects of technology, including AI, go beyond this. The creation of a data ethics advisory body may be a means to help ensure the public is engaged in these ethical issues. It would act to monitor the effects of technology on society, engaging with the public and providing advice to existing regulators to help ensure that the balance between the power of technology – and those controlling it - and wider societal concerns including the rights of individuals is struck correctly. The Information Commissioner is keen to ensure the right solutions are in place and is working with government to help with its consideration of the issue. It is important to ensure that any new advisory body would complement the existing work of the Information Commissioner and other regulators rather than seek to replace existing functions.

42. A data ethics advisory body's role should involve identifying data-related problems that existing regulators may not be able to counter, because they are unaware of them or because the problem falls outside their area of statutory competence. It could detect areas where the societal advantage of data use (personal or non-personal) is not being gained because, perhaps, of a misunderstanding or lack of relevant law. In such cases, a data ethics advisory body could invite

[7] https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2017/07/royal-free-google-deepmind-trial-failed-to-comply-with-data-protection-law/

the appropriate regulator – or regulators - to provide clarification or make recommendations for law reform.

43. Question four. Who in society is gaining the most from the development and use of artificial intelligence and data? Who is gaining the least? How can potential disparities be mitigated?

44. This question is not applicable to the Information Commissioner.

Public Perception

45. Question five. Should efforts be made to improve the public's understanding of, and engagement with, artificial intelligence? If so, how?

46. The situation with AI is much the same as with other technologies. Most people would probably be unable to explain what the main components of a computer do or how coding works. Nonetheless, they may be able to use a computer and understand the consequences of their digital activity.

47. Ideally members of the public would understand what artificial intelligence is and how it affects them. However, we need to be realistic about the public's ability to understand in detail how the technology works. Perhaps it would be better to focus on the effect of the technology – in terms of benefits and detriments – and to ensure that there is an effective regulatory system which does have the necessary technical understanding in place.

48. As we have explained elsewhere in our evidence, even though the 'math' may be difficult for non-experts to understand, it ought to still be possible to explain the purpose(s) for which peoples' data is being processed, who is doing the processing and the consequences of this. If we focus on the consequences of AI, rather than on the way it works, then it is possible to bring about public understanding and to allow individuals to exercise their rights.

49. The ICO has produced a new code of practice on privacy notices[8]. This code stresses the need to communicate with the public in clear, accessible ways. The guidance in the code of practice is as applicable to data processing carried out using AI as it is to more conventional forms of data processing.

[8] https://ico.org.uk/for-organisations/guide-to-data-protection/privacy-notices-transparency-and-control/

Industry

50. Question six. What are the key sectors that stand to benefit from the development and use of artificial intelligence? Which sectors do not?

In this question, you may also wish to address why some sectors stand to benefit over others, and what barriers there are for any sector looking to use artificial intelligence.

51. This question is not applicable to the Information Commissioner.

52. Question seven. How can the data-based monopolies of some large corporations, and the 'winner-takes-all' economies associated with them, be addressed? How can data be managed and safeguarded to ensure it contributes to the public good and a well-functioning economy?

53. Please refer to our evidence above.

Ethics

54. Question eight. What are the ethical implications of the development and use of artificial intelligence? How can any negative implications be resolved?

In this question, you may wish to address issues such as privacy, consent, safety, diversity and the impact on democracy.

55. We have already addressed most of these issues earlier in our evidence. However, we would like to clarify that data protection law is framed in terms of the relationship between data controllers (organisations) and data subjects (individuals). However, it could also be seen as being about the relationship between individuals (acting on behalf of organisations) who make decisions about information use and individuals who are affected by those decisions. In that sense, data protection can be seen as a branch of ethics. We believe that many ethical issues and questions relating to societal norms will be addressed provided that the relationship between organisations and the people whose data they analyse, whether or not using AI, is a fair one.

56. Utilising modern data protection regulatory concepts will also be important. These include: ensuring technological capabilities are used

in a proactive way to safeguard privacy- privacy by design; the impacts are understood and addressed at the outset- privacy impact assessments; and that organisations take proactive responsibility once processing is underway- accountability.

57. Other risks, such as diversity, also highlight the importance of organisations undertaking privacy impact assessments and broader ethical impact assessments before commencing the implementation of AI. Recent research highlights the risks that AI can pose for gender and ethnicity issues[9].

58. The Information Commissioner recognises the importance of applied research that considers the risks of AI but also looks for innovative privacy enhancing solutions that can make a real difference to the public. Her recent Grants Programme encouraged applications in relation to AI. 119 applications have been received for the programme and the grants awarded will be announced before the end of the year[10].

Consent

59. We would like to add a comment about consent. The role of consent is often misunderstood – it can be seen both as a cure-all and as a legal requirement of data protection law. For the reasons we have already discussed, there can be real problems in expecting people to consent to their data being processed by AI systems. Many people will not know what AI is or the implications of its usage, and in data protection law consent has to be fully informed to be valid. This means that there may be significant problems in legitimising the use of AI on the basis of individuals' consent.

60. Individuals can suffer from 'consent fatigue' – as may be the case with repeated 'cookie consents'. Many individuals might prefer the services and systems they use to do what they expect them to and to use their personal data fairly and responsibly, with benign and predictable outcomes, but not to be repeatedly asked for their consent. Another problem is that if the use of AI is occurring on the basis of consent, then it would likely have to cease if consent is withdrawn or found to have not been given in the first place. This could lead to the scenario of organisations offering AI and non-AI enabled services, something that it could be infeasible to deliver in practice.

61. On a legal point, data protection law is sometimes portrayed as requiring individuals' consent in order to process their personal data.

[9] http://www.sciencemag.org/news/2017/04/even-artificial-intelligence-can-acquire-biases-against-race-and-gender

[10] https://ico.org.uk/about-the-ico/what-we-do/grants-programme/

This is not the case. The law usually provides a number of bases for processing personal data, consent is just one. Organisations can process personal data, including the use of AI provided the activity is legitimate and does not have a detrimental effect on people. If this is the case and the compliance issues we have discussed earlier in our evidence are addressed properly, then organisations should be able to go ahead with the processing without the individual's consent. It is important to be clear, however, that this is not the case with regards to the processing of 'sensitive' personal data – for example, data relating to the health, racial or ethnic origin, political opinions or sexual orientation of individuals. Here, consent or another appropriate basis will need to be used.

62. Question nine. In what situations is a relative lack of transparency in artificial intelligence systems (so-called 'black boxing') acceptable? When should it not be permissible?

63. As we have already explained, transparency is one of the basic requirements of data protection law. However, there are exemptions from this, for example where providing too much information about how a system operates would prejudice the purposes of law enforcement, or where providing information about the logic involved in decision-taking constitutes a trade secret. The rules and norms here are well-established and apply equally to AI and non-AI processing of personal data.

64. Regardless as to whether an exemption does or does not apply, there is still a concern with 'blackboxing' in terms of accountability. An issue with the 'black-box' is that no-one understands how an AI system got from input to output. Where there is zero transparency how can the processing be demonstrably compliant with data protection laws? We discuss potential methods to approach algorithmic transparency in our research paper.[11]

The Role of the Government

65. Question 10. What role should the Government take in the development and use of artificial intelligence in the United Kingdom? Should artificial intelligence be regulated? If so, how?

66. Government must recognise that there are unique features of AI that mean that it presents risks as well as opportunities. We do not

[11] https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf

think that AI should be regulated as a discrete topic. We should look more at the purpose and effects of its use, rather than the technology itself. As we have already explained many uses of AI are already subject to regulation through data protection and other laws. However, as discussed above, given the complexity of the regulatory landscape and the fact that AI straddles several areas of regulatory responsibility, we do think there is a case for some form of ethics advisory body to take a holistic view, providing advice to existing regulators so that the best protection is offered to individuals and to society as a whole.

67. We should not underestimate the potential consequences of AI for individuals, ones that can be irreversible. This is why it is so important that organisations deploying AI-enabled systems have a clear set of compliance rules so they can design and deploy AI systems properly, with proper respect for the individuals whose data they may be processing. We have explained earlier on in our evidence how data protection legislation provides appropriate safeguards where personal data is involved.

Learning from Others

68. Question 11. What lessons can be learnt from other countries or international organisations (e.g. the European Union, the World Economic Forum) in their policy approach to artificial intelligence?

69. The International Conference of Data Protection and Privacy Commissioners, the forum for the world's data protection and privacy authorities of which the ICO is member, focussed specifically on the topic of AI as part of its 38th gathering in 2016. The fact that this theme was chosen for the conference demonstrates the significantly increased level of global attention that AI devices have attracted in the last two to three years. There is consensus across data protection and privacy commissioners that we are only just beginning to understand the challenges that AI brings to data protection. The Information Commissioner will continue to work with her international counterparts in furthering the understanding of these challenges and proposing potential solutions.

Annex- Basic Compliance Steps for the Responsible use of AI

70. These are the basic steps that should be taken when implementing an AI-enabled data processing system. They are based on the premise that the use of AI is going to become more prevalent and that organisations need to understand the rules needed to deploy it responsibly.

1) Initial assessment of the need for the data processing; what are you trying to achieve – e.g. detect fraudulent benefit claims and why is AI based processing a necessary and proportionate response to achieving this? (Commissioning a form of privacy impact assessment may help with this and to identify necessary conventional data protection safeguards if personal data is involved).
2) If the decision is taken to use AI, specify a range of data inputs (i.e. data sources and data items) as well as limits on algorithmic self-improvement.
3) Test the system, ideally using synthetic data or, if this is not possible, a small sample of live data (in accordance with appropriate safeguards) and assess the results – e.g. is benefit fraud being detected accurately?
4) If the system is intended to go live, ensure that during the design and testing phases, transparency procedures for informing the public of general privacy information but also of the use of automated decision making / artificial intelligence are developed.
5) Carry out regular audits to ensure that the system is working in the expected manner - i.e. that the correct data items are being utilised and that they are being analysed in accordance with design parameters.
6) Put systems in place for the periodic review of outcomes – is the system continuing to achieve its intended objectives? If not, modify the system or deploy a better one.
7) Ensure there are procedures in place for dealing with queries and complaints from the public, including means of re-taking a decision with an element of human intervention, and for delivering all relevant individuals’ rights.