How To Comply With New UK Guidance On AI Privacy

By **William Long and Francesca Blythe** (November 9, 2020)

William Long

Francesca Blythe

As artificial intelligence plays an increasingly important role in nearly all aspects of everyday life, from national security, law enforcement, business, health care, education, culture, personal and property protection, and social networking, the U.K.'s Information Commissioner's Office appears to be positioning itself as a regulatory front-runner in this regard, with the aim of placing the U.K. at the front of emerging digital technologies.

The ICO has identified AI as one of its top three strategic priorities and enabling good practice in AI as one of its regulatory priorities. In turn, the ICO has published a series of blog posts since March 2019 and guidance documents focusing on data protection-compliant AI.

Most recently, on July 30, the ICO announced the publication of its guidance on AI and data protection as part of the development of its wider AI framework. The AI guidance builds upon and should be read in conjunction with the ICO's guidance on explaining decisions made with AI, which was published in May in collaboration with the Alan Turing Institute.

The AI guidance, which is focused on machine-learning-based AI, acknowledges the significant benefits that AI technologies can provide but also recognizes that, without appropriate regulation, the use of AI technology can pose risks to the rights and freedoms of individuals. The AI guidance is two-fold in that it provides (1) auditing tools and procedures that the ICO will use in AI audits and investigations; and (2) practical guidance for organizations on how to use AI in accordance with the data protection principles.

Companies developing and/or using AI technologies will need to closely review the AI guidance and implement policies and procedures to address the ICO's recommendations — which are largely focused on the concepts of accountability and data protection by design and by default. Helpfully, in doing so the ICO does not expect companies to adopt a zero-tolerance approach to risks, and, instead, companies should ensure that these risks are identified, managed and mitigated.

In practice, this will most likely be documented in the form of a data-protection impact assessment, or DPIA, which according to the ICO, should be carried out in the vast majority of cases and which must equate to more than a tick- box exercise.

The ICO, in turn, proposes that companies maintain two versions of the DPIA — one that presents a thorough technical description for specialist audiences, and another that provides a high-level description of the processing for individuals to address the transparency requirement of fair and lawful processing and/or customers to satisfy their due diligence process.

Interestingly, while the draft AI guidance included an obligation to consult affected individuals unless there was a good reason not to — something which in practice would be virtually impossible for companies with no direct interaction with the individuals — the final version of the AI guidance includes a helpful caveat of where the company can demonstrate

that such consultation would "compromise commercial confidentiality, undermine security, or be disproportionate or impracticable." This is indeed something that companies should be considering and documenting in the DPIA.

Ultimately, to the extent companies utilizing AI technologies do not already have a process in place for carrying out DPIAs, they should consider developing and implementing such a process.

The process should ensure that the need to carry out a DPIA is triggered early on in the project, i.e., predeployment, and it should involve all necessary personnel, i.e., it should not be delegated to data scientists or engineering teams — in particular, to ensure that where tradeoffs and made, i.e., where a balance needs to be struck between competing interests of privacy and other rights and interests, these are approved at an appropriate level within the company.

The DPIA should be regarded as a central assessment for any AI project and should identify and assess these tradeoffs, which, according to the ICO, will inevitably be triggered by the use of AI technologies.

Further, the ICO requires that the processes in relation to the tradeoffs should be documented to an auditable standard and be captured within the DPIA. Unfortunately, the level of detail expected is not clear and the examples of the types of tradeoffs the ICO anticipates and how these can be managed in line with the data protection requirements are fairly limited.

With regard to the designation of a party as a controller or processor in the context of the AI supply chain, the AI guidance acknowledges the complexity of assigning these roles when AI systems involve a number of organizations and states that the ICO intends to further explore this in consultation with relevant stakeholders when it reviews its cloud computing guidance in 2021.

However, this in turn means that for the time-being, and irrespective of the draft guidance recently published by the European Data Protection Board on controllers and processors, there remains little clarity in this regard — in particular, for example, in the context of the role of a developer post-deployment.

A similar potential area of concern for a number of companies, in particular those that cease to have any control over the product or data post-deployment, are the various references in the AI guidance to ongoing monitoring, i.e., in relation to statistical accuracy, security measures and data subject rights.

While we agree that statistical accuracy and the appropriate measures to evaluate it should be considered as part of the obligation to implement data protection by design and by default throughout the AI life cycle, in practice the requirement for ongoing monitoring in this regard and, indeed, in relation to the security measures implemented, should not apply to a manufacturer of a product after the product has been deployed. However, the AI guidance does not provide details as to whether the liability of a manufacturer or developer should continue post-deployment.

More generally, in terms of security, the risk-based approach to compliance and the acknowledgement by the ICO that there is no one-size-fits-all approach is welcome. The ICO also acknowledges the unique risks presented from a security perspective in the context of AI in particular, as common practices about how to process personal data securely in the

context of AI are still under development.

However, this means that companies involved in the AI supply chain will need to constantly work to expand their knowledge in this regard to ensure the products they develop and/or use promote privacy and data protection. This approach will however, also serve to meet the continued obligation to maintain data protection by design and by default throughout the life cycle of the data-processing activities and assist in ensuring the measures implemented are fit for purpose.

A further action for companies developing AI technologies is the need to implement risk-management practices designed to ensure that data minimization is considered from the design phase. Likewise, where AI systems are purchased from or operated by third parties, these considerations should form part of the due diligence process.

In particular, the AI guidance considers the data minimization principle in relation to supervised machine learning systems at various stages in the AI life cycle and provides examples of privacy enhancing methods which can be used to minimize personal data processed e.g., perturbation, use of synthetic data and federated learning.

Finally, the ICO addresses data subject rights and acknowledges that the implementation of effective mechanisms for individuals to exercise their rights will be more challenging in the context of AI. Nonetheless, the ICO confirms that the rights apply at all stages of the AI life cycle. In particular, the ICO dedicates an entire subsection to individual rights relating to solely automated decisions with legal or similar effect, i.e., Article 22 of the European Union's General Data Protection Regulation, as this is particularly relevant in the context of AI systems.

The draft guidance discusses the concepts of automation bias and lack of interpretability — both identified as factors that, according to the ICO, could potentially cause a system to be considered solely automated — and states that organizations should: (1) consider the system requirements necessary to support a meaningful human review from the design phase of the AI system; (2) implement a process for individuals to exercise their data protection rights which is simple and user friendly; and (3) record all decisions made by an AI system as part of the organization's accountability obligations.

In practice, this means companies will likely need to adapt their existing data subject rights policies and processes to ensure they are fit for purpose also in the context of AI. A process will also need to be implemented to ensure that the decisions referred to in (3) are adequately recorded.

To conclude, this is a quickly developing area and one in which regulation will need to adapt accordingly. The ICO has acknowledged that the AI guidance will "continue to evolve … to keep pace with [AI]," and, in turn, the ICO will continue to seek feedback on the AI guidance and engage with experts.

In particular, we understand that the ICO is developing a toolkit designed to provide further practical support to organizations auditing the compliance of their own AI technologies, which will be published in due course.

William Long is a partner and a global co-leader of the privacy and cybersecurity practice at Sidley Austin LLP.

Francesca Blythe is a senior associate at the firm.