Centre for Information Policy Leadership
Hunton Andrews Kurth LLP

Centre for Information Policy Leadership
Hard Data Protection Issues Roundtable
London, March 12, 2019

Issues

1. **Fairness**: Data protection regimes around the world consider fairness, either as an explicit requirement or an overarching principle of data protection. However, fairness requirements are not clearly defined, especially in the context of AI.

    a. What is "fairness" in the context of AI?

    b. Is there any limit to which factors should be considered when assessing the "fairness" of AI? Should organisations and data protection regulators consider factors outside of data protection such as job displacement or wealth disparity?

    c. Who has the responsibility for determining "fairness"?

    d. How can organisations determine, monitor, and document fairness in the development and deployment of AI? Is it a matter of transparency, of ethics boards, or something else? Are there specific procedural elements that are necessary for a thoughtful, valid determination of "fairness" (or "unfairness")?

2. **Transparency**: Transparency is also a requirement of most data protection regimes. However, AI can make transparency more difficult because of the technology's ability to make previously unrecognized correlations, inferences, and results.

    a. What does "transparency" mean in the context of AI? Does it require transparency about purposes and/or uses of personal data, algorithms, consequences, alternative outcomes, or other elements of AI?

    b. What purposes should "transparency" serve in the context of AI?

    c. How can organisations provide appropriate and effective transparency for individuals with regard to the processing of personal data by AI applications?

    d. What tools are available to help address transparency challenges relating to the complexity of algorithms, the changing nature of algorithms and data, and the need to protect trade secrets?

    e. Are there effective alternatives to "transparency"?

3. **Use limitations**: Many data protection laws require organisations to articulate the purpose of collection or processing, and then limit processing to only those specified (or compatible) purposes. Use limitation and purpose specification pose a particular challenge when

advancements in AI may create new and valuable uses for existing data or yield unforeseen results.

a. If the results of processing are different than the expectations, what options do organisations have to be able to use the results?

b. To what extent should regulators consider the potential transactional burden created for organisations and the potential privacy burden created for individuals by the need to return to individuals to obtain new consent for an originally unanticipated use?

c. How should "compatible" be defined?

d. Should it matter that transactional burdens and use limitations may hinder the development of AI for public beneficial uses?

e. Should organisations and regulators consider whether unanticipated uses of personal data pose little to no risk of harm to the individual? How might one appropriately determine that?

4. **Data minimisation**: The OECD principles and all EU laws require organisations to limit data collection to what is necessary for the organisation's stated purpose. If interpreted narrowly, this can limit the effectiveness of AI by eliminating protected variables from algorithms or by requiring the destruction of no longer necessary data. Similarly, heightened consent requirements and other processing limitations for "sensitive" or "special" categories of data also have the effect of eliminating relevant and often necessary variables from data sets. A lack of data may render the algorithm inaccurate or biased, while the absence of protected variables may make it more difficult to identify and attribute bias.

a. With AI, it is almost unpredictable what variables are necessary or relevant for an algorithm, but it is likely that the more data used in the algorithmic development phase, the more accurate the results will be. How then can organisations comply with limitations of relevance and necessity for processing? If everything is necessary for the intended purpose, there is no practical effect. However, without additional data, the AI may be less accurate. How can organisations and regulators navigate these two realities?

b. Is it permissible to demonstrate data minimisation with respect to an AI system by proactively articulating and documenting the need to collect and process data as well as the expectations or goals of the processing?

c. How broadly or narrowly may a purpose be described to satisfy the legal requirements of data minimisation?

d. To what extent can a public social benefit be the stated purpose for processing?

e. If protected variables cannot be collected and bias becomes more difficult to detect, how should organisations proceed with respect to collection and processing?

f. In addition to use limitations, regulations often set limitations on the retention of data to only as long as necessary. Because AI uses data for training, development, and deployment, should those retention limits be extended due to the potential future use?

5. **Automated Decision-Making**: Restrictions or prohibitions on automated decision-making are driven by concerns of algorithmic bias, incorrect decisions based on inaccurate or incomplete data, or general unfairness to individuals. These limitations are often prescribed to decisions that will have a legal or similarly significant impact. Automation of certain decisions can be beneficial to both organisations and individuals, so navigating the tensions between automated decisions and data protection is especially important.

   a. What impacts to individuals should be considered legally significant?

   b. Even if the impact is legally significant, should automated decisions be permitted if organisations provide meaningful, effective, and expedient options for review and redress? Must human intervention always occur prior to automated decision-making?

   c. If an automated decision does not provide a negative consequence to individuals (for example – if the algorithm can pre-approve an individual, but everyone else goes through an individual human review), is this a legally significant impact?

   d. Should an individual be able to consent to automated decision-making in certain circumstances? If so, what are those circumstances?

How might the large and expanding volume of automated decisions be addressed meaningfully?