Centre for
Information
Policy
Leadership
Hunton Andrews Kurth

The Centre for Information Policy Leadership at Hunton Andrews Kurth LLP

Artificial Intelligence and Data Protection: Delivering Sustainable AI Accountability in Practice

Project Proposal

February 13, 2018

“Artificial Intelligence”

Significant advances in the analytical capacity of modern computers are increasingly challenging data protection laws and norms. Those advances are often described by the term “artificial intelligence” (or “AI”) a term that describes the broad goal of empowering “computer systems to perform tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”[1] This one term encompasses a variety of technical innovations, each of which may present distinct challenges to data protection tools.

Most AI in use today involves computer systems that perform discrete tasks—playing games, recognizing images, verifying identity—by identifying patterns in large amounts of data. The mathematical concept dates back to the 1950s but has found real-world applications in recent years, thanks to advances in processing power and the vast amounts of digital data recently available for analysis. As a result, AI almost always is associated with “big data.” However, recent developments such as the recent use of AI to defeat CAPTCHA and Google’s AlphaGoZeros that taught itself to play Go at the championship level, both occurred with minimal training data, so big data may not always be linked with AI.

These are all examples of “narrow” AI, which is AI designed to perform one task or set of tasks. It is still pretty complicated and includes tools such as “boosted decision trees,” that allow an algorithm to change the weighting it gives to each data point, and “random forests,” that average together many thousands of randomly generated decision trees. As the *New York Times* recently noted, even narrow AI tools can be “bafflingly opaque” and “evade understanding because they involve an avalanche of statistical probability.”[2] This is an obvious challenge both for building confidence in the technologies and for compliance with data protection laws.

Researchers are increasingly seeking to develop “autonomous” AI or even Artificial General Intelligence (called “AGI”), which are programs capable of altering themselves or even of mimicking human intelligence. This is the field of “machine learning,” which Stanford University professor Andrew Ng defines as “the science of getting computers to act without being explicitly programmed.”[3] Machines with true intelligence are the future of AI, and will be even more difficult to explain, predict, or build compliance programs around.

[1] English Oxford Living Dictionaries, “Artificial Intelligence,” at https://en.oxforddictionaries.com/definition/artificial_intelligence.

[2] Cliff Kuang, Can A.I. Be Taught to Explain Itself?, *New York Times Magazine* (Nov. 21, 2017)., at https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html?_r=0.

[3] https://www.coursera.org/learn/machine-learning.

The Challenge for Data Protection

With this broad understanding of AI—encompassing, but not necessarily requiring big data, and extending from today's smart machines to increasingly autonomous and nimble computers of the future—it is easy to see how data protection laws and norms might be challenged. For example, beginning with the OECD's Guidelines on the Protection of Privacy and Transborder Flows of Personal Data adopted in 1980, modern data protection law has recognized both Purpose Specification and Use Limitation principles—namely, that personal data should be collected for specified purposes and then used only for those or similar purposes. Similarly, the Openness and Individual Participation principles require that data processing be transparent and that individuals have access to personal data about them, as well as to information about how those data are used. Most national laws reflect these principles, and they are especially prevalent in the EU Data Protection Directive and the new EU General Data Protection Regulation.

The challenge, of course, is how to comply with these requirements when data are being used for unforeseen, unpredictable purposes, by advanced computational machines that are not always understood by their own programmers and will be increasingly programmed only by other computers. As Georgetown professor Paul Ohm has stressed, when a program "thrives on surprising correlations and produces inferences and predictions that defy human understanding…. [h]ow can you provide notice about the unpredictable and unexplainable?"[4]

Implicit in the OECD Guidelines, and made explicit in EU General Data Protection Regulation, is another widely shared principle: Data Minimization. "Personal data must be adequate, relevant and limited to what is necessary in relation to the purposes for which those data are processed."[5] Yet how does one know in advance "what is necessary" in a world of "surprising correlations" and computer-generated discoveries?

The challenges to data protection presented by AI are frequently remarked on, but usually addressed in policy settings only at a surface level. The result has been a fair amount of hand-wringing and assertions about the need to achieve the extraordinary advances AI makes possible while still complying with all applicable data protection laws. Often, this sounds like a call to do the impossible or face the threat of regulatory consequences. There is an urgent need for a more nuanced, detailed understanding, especially by regulators, of the opportunities presented by AI, and of potential challenges and practical ways of addressing them, in terms of both legal compliance and the ethical issues that AI may raise.

CIPL's Project

CIPL proposes a new project focusing on AI, its reliance on big data, and the evolution towards autonomous AI. The project will:

1. Describe clearly the wide range of technological innovations encompassed by "AI," both today and in the foreseeable future, including examples of the ways in which they are being deployed in specific sectors and the benefits that result. This would include the continuing and intensive

[4] Paul Ohm, "Changing the Rules: General Principles for Data Use and Analysis," in Julia Lane, Victoria Stodden, Stefan Bender, Helen Nissenbaum, eds., *Privacy, Big Data, and the Public Good* 100 (Cambridge 2014).
[5] Rec.39; Art.5(1)(c).

processing of innovation, the need for training data, and the ways in which AI is being used to facilitate privacy.

2. Address in precise, specific terms the opportunities and challenges presented by these innovations to data protection laws and norms.
3. Address in precise, specific terms the opportunities and challenges presented by these innovations to norms about the ethical use of personal data and other societal issues.
4. Describe practical steps for addressing today's challenges and those on the horizon, including best practices already in use by leading companies; efforts to create user-centric designs that facilitate trust, transparency, and control while also delivering a frictionless, enjoyable experience; innovative applications of existing legal concepts; the role of accountability; and proposals for new approaches.
5. Provide a frank acknowledgment of issues that cannot be resolved within existing laws and regulations and of the limits of what we know about AI and its future.

The project will proceed in phases. The first will be working with member companies to develop a white paper that builds on CIPL's 2017 discussion paper, *Data Privacy Accountability for Artificial Intelligence and Machine Learning*. This version of the white paper will focus on objectives 1 through 3 above. Given the breadth of the topic and the issues, the white paper will likely focus on specific sectors.

The second phase will involve expanding the discussion to include regulators and policymakers, with a goal of engaging and informing them and refining the white paper, to address all five objectives above. The third phase will involve socializing the white paper, especially with regulators and policy makers, with an eye towards identifying topics that require further development or addressing AI in other sectors, perhaps as separate documents.

The goals of CIPL's project are to: (1) engage with policy makers and regulators about AI and related innovations so that they develop a more practical and detailed understanding of the technologies, the legal and ethical challenges the technologies present for data protection, the benefits of those technologies broadly and for data protection, the practical approaches that companies are using and developing for addressing the challenges presented by AI, and the limits of those approaches; (2) facilitate the deployment of AI in ways that builds trust among regulators, policymakers, and the public; (3) enhance information-sharing among member companies about AI and related tools; and (4) build on CIPL's prior work on accountability and establish CIPL and participating member companies as thought-leaders on the critical issues surrounding AI.

Given the speed of innovation, the project is necessarily both iterative and forward-looking—it will evolve over time and it will focus on tools to the challenges AI presents for data protection that work today and in the future. The project will seek opportunities to work collaboratively with regulatory, policymakers, and others. And the work product will be designed to be comprehensible, precise, and practical, avoiding the sweeping generalizations and often vague generalizations that have characterized much of the policy work in this field to date.