Professional Perspective

AI Regulation Part 1: What You Need to Know to Stay Ahead of the Curve

Peter J. Schildkraut, Arnold & Porter Kaye Scholer LLP

Bloomberg Law

AI Regulation Part 1: What You Need to Know to Stay Ahead of the Curve

Contributed by Peter J. Schildkraut, Arnold & Porter Kaye Scholer LLP

Artificial intelligence is all around us. AI powers Alexa, Google Assistant, Siri, and other digital assistants. AI makes sense of our natural language searches to deliver, we hope, the optimal results. When we chat with a company representative on a website, we often are chatting with AI, at least initially. AI has defeated the human world champions of chess and Go. AI is advancing diagnostic medicine, driving cars, and evaluating all types of risks.

As AI becomes more common, more powerful, and more influential in our societies and our economies, governments are noticing. When Google CEO Sundar Pichai publicly proclaims "there is no question in my mind that artificial intelligence needs to be regulated," the questions are when and how—not whether—this will happen.

Indeed, certain aspects of AI already are regulated, and the pace of regulatory developments is accelerating. This pair of articles tells you what you need to know—and what steps your company can take—to keep ahead of this curve.

I. What Is AI?

Before discussing the regulation of AI, let's review what AI is and how the leading type works.

Experts broadly conceive of two versions of AI:

1) narrow, meaning it can perform one particular function; and

2) general, meaning it can perform any task and adapt to any situation.

These risks require attention, however. Working from probabilities, not certainty, your company's AI will make mistakes. A big one will attract unwelcome scrutiny from regulators, the media, and the plaintiffs' bar.

All existing AI is narrow. General AI (sometimes known as "artificial general intelligence" or AGI) can perform any task and adapt to any situation. AGI would be as flexible as human intelligence and, theoretically, could improve itself until it far surpasses our capabilities.

For now, AGI remains in the realm of science fiction, and authorities disagree whether AGI is even possible. While serious people do ponder regulating AGI—in case someone creates it—current regulatory initiatives focus on narrow AI.

Machine Learning

Machine learning has enabled the recent explosion of AI applications. As one group explains, "Machine learning systems learn from past data by identifying patterns and correlations within it."

Whereas traditional software (and some other types of AI) run particular inputs through a preprogrammed model or a set of rules and reach a defined result – akin to 2 + 2 = 4a machine learning system builds its own model from the data it is trained upon. The system then can apply the model to make predictions about new data.

According to trade association CompTIA, algorithms are "now probabilistic…. In other words, we are asking computers to make a guess."

For example, in a technology-assisted document review, lawyers will code a small sample of the document collection as responsive or not. The system will identify patterns and correlations distinguishing the sample documents that were coded "responsive" from those coded "not responsive." It then can predict whether any new document is responsive and measure the model's confidence in its prediction.

For validation, the lawyers will review the predictions for another sample of documents, and the system will refine its model with the lawyers' corrections. The process will iterate until the lawyers are satisfied with the model's accuracy. At that point, the lawyers can use the system to code the entire document collection for responsiveness with whatever human quality control they desire.

The quality of the training data set matters greatly. The machine learning system assumes the accuracy of the training data.

In the document review example, if the lawyers incorrectly code every email written by a salesperson as responsive, they will bias the model towards predicting that every sales team email is responsive.

The biased higher probability is not a certainty, however. Other things about an email might overcome the bias. For instance, the lawyers may have coded every email about medical appointments as nonresponsive. As a result, the model still might predict a salesperson's email about a medical appointment is nonresponsive.

II. Why Regulate AI?

Several characteristics of AI drive the calls for regulation. These include accuracy, bias, power, and market failures.

Accuracy

As noted in the above document review example, AI predictions are sometimes inaccurate. Coding every email written by a salesperson as responsive, for example, can lead to faulty results.

Poorly performing AI could underestimate a person's fitness for a job or qualifications for a loan. AI that misperceives environmental conditions could result in the crash of a car. In short, AI could harm individuals in ways that humans already do, but without human intervention.

While the harms may be same, AI's infliction of those harms is novel. In response, legislators may create new causes of action to address harms caused by AI, and regulators may propose prophylactic measures to prevent certain uses of AI.

Applications of AI causing societal harm—like sex or race discrimination—are of particular concern, as are certain uses of facial recognition technology. A study of 189 facial recognition systems found that minorities were falsely identified much more frequently than Caucasians, and false positives were higher for women than for men. Privacy questions aside, the use of facial recognition by law enforcement to identify criminal suspects makes these differences particularly troubling.

Bias

Bias can arise in various ways. Like all of us, algorithm creators have cultural blind spots, potentially causing them to miss opportunities to correct the disparate impacts of a given algorithm. Moreover, AI predictions are only as good as the underlying training data, which often reflect the society from which they are collected, biases included. To prevent societal biases from infecting AI, developers and operators—and their advisors—must recognize them and determine how to adjust the training data.

Even when AI makes accurate and unbiased predictions, however, the results can be troubling. A targeted advertising AI system that displays job opportunities where their return per placement is highest may inadvertently bypass female candidates if there is higher demand (thus a higher price) to advertise to women. See Ajay Agrawal et al., *Prediction Machines: The Simple Economics of Artificial Intelligence* 196 (2018).

Power

In a 2018 MIT-Harvard class, Joi Ito, then the Director of the MIT Media Lab, relates being told that "machines will be able to win at any game against humans pretty soon." Ito observes:

"A lot of things are games. Markets are like games. Voting can be like games. War can be like games. So, if you could imagine a tool that could win at any game, who controlled it and how it is controlled has a lot of bearing on where the world goes."

It is easy to see why the public might demand regulation of this power.

Market Failures

For all their power, though, AI-enabled goods and services cannot escape market forces and market failures.

Consider programming a self-driving automobile to pass cyclists in the face of oncoming traffic.

The car's occupants will be safer if the car travels closer to the cyclist and farther away from oncoming traffic. The cyclist will arguably be safer if the car moves closer to the oncoming traffic and farther away from the cyclist.

Nobody wants to buy an autonomous vehicle programmed to protect others at its occupants' peril, but everyone wants other people's autonomous vehicles to minimize total traffic casualties.

This is a classic collective action problem where regulation can improve the market outcome.

III. Application of Familiar Regulatory Regimes to AI

While comprehensive regulation of AI is likely for all the reasons discussed above, AI already is regulated in certain economic sectors and activities.

Generally Applicable Law

"AI did it" is, by and large, not an affirmative defense. If something is unlawful for a human or non-AI technology, it probably is illegal for AI. For instance:

- Antidiscrimination laws like Title VII of the U.S. Civil Rights Act of 1964, 42 U.S.C. § 2000e-2(k)(1)(A)(i), the UK Equality Act 2010, c. 15, and the U.S. Equal Credit Opportunity Act, 15 U.S.C. §§ 1691-1691f, prohibit disparate impacts on protected classes, absent specified exceptions, including those caused by AI.
- The U.S. Fair Credit Reporting Act requires certain disclosures to potential employees, tenants, borrowers, and others regarding credit or background checks and further disclosures if the report will lead to an adverse action. 15 U.S.C. § 1681b. Credit and background checks that rely on AI are just as regulated as those that do not.
- The U.S. Securities and Exchange Commission has enforced the Investment Advisers Act of 1940, 15 U.S.C. §§ 80b-1 to 80b-18c, against automated portfolio management services. See, e.g., *In re Hedgeable, Inc.*, Investment Advisers Act Release No. 5087, 2018 WL 6722757 (Dec. 21, 2018). The agency also has warned so-called "robo-advisers" regarding compliance with the Investment Company Act of 1940. 15 U.S.C. §§ 80a-1 to 80a-64.
- There should be little doubt the U.S. Food & Drug Administration will enforce its good manufacturing practices regulations, 21 C.F.R. pts. 210-212, on AI-controlled pharmaceutical production processes.
- Claims about AI applications must not deceive lest they run afoul of consumer protection laws like Section 5 of the U.S. Federal Trade Commission Act, 15 U.S.C. § 45. See, e.g., *In re Everalbum, Inc.*, File No. 1923172, 2021 WL 118892 (FTC Jan. 11, 2021).

The longer it takes for broad regulation to arrive, the more we should expect government enforcers to apply their existing powers in new ways. Just as the FTC has used its Section 5 authority in the absence of a federal privacy statute, see, e.g., Facebook, Inc., Docket No. C-4365, 2020 WL 2197924 (FTC Apr. 27, 2020), that agency is turning its attention to AI abuses.

Express Regulation of AI

In addition, the General Data Protection Regulation restricts the development and use of AI in connection with individuals and their personal data. Regulation (EU) 2016/679, 2016 O.J. (L 119) 1.

U.S. states, too, are adopting laws expressly regulating AI.

- As amended by the California Privacy Rights Act of 2020, the California Consumer Privacy Act of 2018 contains requirements like the GDPR's, although the restrictions on automatic decision-making are left to be fleshed out by regulations. See Cal. Civ. Code §§ 1798.100-1798.199.100.
- The Virginia Consumer Data Protection Act does as well. See 2021 Va. Acts ch. 36.
- Another California law targets intentionally deceitful uses of chatbots masquerading as real people in certain situations. Cal. Bus. & Prof. Code §§ 17940-17943.
- Depending on the technology involved, an Illinois statute may govern private entities' use of facial recognition technology. See 740 Ill. Comp. Stat. 14/1; *In re Facebook Biometric Info. Privacy Litig.*, No. 15-cv-03747-JD, 2021 BL 69600, at *6 (N.D. Cal. Feb. 26, 2021).

- Illinois also regulates employer use of AI for screening video interviews. 820 Ill. Comp. Stat. 42/5.

IV. Development of New Regulatory Regimes for AI

Because existing rules don't address all concerns about AI, policymakers worldwide are considering new regulatory regimes.

The advanced democracies appear to have a loose consensus that the degree of regulation should correlate with an application's risk. For instance, a music-recommendation algorithm has much lower stakes than does AI diagnosing disease. The Organization for Economic Cooperation and Development (OECD) and the G20 have adopted principles embodying this high-level consensus.

OECD/G20 Principles for Responsible Stewardship of Trustworthy AI

- AI should benefit people and the planet by driving inclusive growth, sustainable development and well-being.
- AI systems should be designed in a way that respects the rule of law, human rights, democratic values, and diversity, and they should include appropriate safeguards–for example, enabling human intervention where necessary–to ensure a fair and just society.
- There should be transparency and responsible disclosure around AI systems to ensure that people understand AI-based outcomes and can challenge them.
- AI systems must function in a robust, secure, and safe way throughout their lifecycles, and potential risks should be continually assessed and managed.
- Organizations and individuals developing, deploying, or operating AI systems should be held accountable for their proper functioning in line with the above principles.

The international consensus, however, seems likely to fray when it comes to regulating particular risks. The U.S. has a culture of permission-less innovation, waiting for harms to emerge and be well-understood before regulating. European governments tend to be quicker to invoke the precautionary principle: innovations should be regulated until they are proven safe. Yet, even in the U.S., some warn against repeating what they see as the mistake of not regulating the internet until it was too late to prevent various harms.

U.S. Steps Toward Regulating AI

The Trump Administration sought to "promote a light-touch regulatory approach," to AI. Last year, the U.S. Office of Management and Budget published guidance for agencies consistent with this light-touch approach. According to OMB:

- "The appropriate regulatory or non-regulatory response to privacy and other risks must necessarily depend on the nature of the risk presented and the tools available to mitigate those risks." In particular, "[i]t is not necessary to mitigate every foreseeable risk …. [A] risk-based approach should be used to determine which risks are acceptable and which risks present the possibility of unacceptable harm, or harm that has expected costs greater than expected benefits."

- Instead of "[r]igid, design-based regulations" with technical prescriptions, agencies should consider sector-specific policy guidance or frameworks, pilot programs and experiments, and voluntary consensus standards and frameworks.

- Agencies should evaluate training data quality, assess protection of privacy and cybersecurity, promote nondiscrimination and general fairness in AI application outcomes, both absolutely and compared to existing processes, and consider what constitutes adequate disclosure and transparency about using AI.

It remains to be seen whether the Biden Administration will share this light touch. In recent decades, Democrats, like Republicans, largely have supported permission-less innovation to encourage new technologies.

However, members of both parties have begun pressing for greater regulation of at least some internet companies and applications, arguing they have amassed too much economic, social, and political power. Shaped in part by this experience

and spurred by concerns that algorithms are reinforcing biases in U.S. society, various Democrats have called for regulation of AI even at this early stage. For example, according to its sponsors, the proposed Algorithmic Accountability Act, S. 1108, 116th Cong. (2019); H.R. 2231, 116th Cong. (2019), would:

- Authorize FTC regulations requiring companies to conduct impact assessments of highly sensitive automated decision systems, both new and existing.
- "Require companies to assess their use of automated decision systems, including training data, for impacts on accuracy, fairness, bias, discrimination, privacy[,] and security."
- Compel self-evaluation of protections for consumers' personal information.
- Mandate correction of problems companies discover through impact assessments.

Except for data brokers, smaller companies would be exempt to these rules. The proposed Consumer Online Privacy Rights Act similarly would require algorithmic decision-making impact assessments. S. 2968, 116th Cong. § 108(b) (2019). With recent changes in Washington and greater public sensitivity to systemic discrimination, such legislation may gain traction and even become law.

European Steps Toward Regulating AI

While the U.S. has yet to embrace widespread regulation of AI, Europe is plowing ahead. To supplement the GDPR, the EU is moving towards regulating high-risk AI applications.

In April 2021, the European Commission released proposed legislation regulating AI. The proposal classifies AI systems as high-risk or not based upon intended use. If adopted, the legislation would require all high-risk AI systems to meet standards for compliance programs and risk management; human oversight; documentation, disclosure and explainability; robustness, accuracy, and cybersecurity; and record retention. High-risk systems would have to demonstrate compliance through conformity assessments before introduction into the European market. Some AI systems—high-risk or not—would have to meet transparency standards.

Certain uses of AI would be prohibited altogether. Most current uses of AI would remain unregulated. In the transportation sector, high-risk AI safety components, products, or systems covered by eight existing legislative acts would be exempt from the proposal, although those acts would be amended to take the proposal's requirements for high-risk systems into account.

The Commission simultaneously proposed updated legislation regulating machinery, which, among other changes, would address the integration of AI into machinery, consistent with the AI proposal. Interested parties may submit comments on the AI proposal during the consultation period ending July 6, 2021. The European Parliament and Council of the European Union will then consider the Commission's proposal in light of those comments.

Previously, the European Parliament had adopted its own recommendations for proposed legislation. See Framework of Ethical Aspects of Artificial Intelligence, Robotics and Related Technologies, P9_TA(2020)0275.

Parliament's recommendations would not ban specific AI practices as the Commission has proposed, but it would define "high-risk" somewhat more expansively than the Commission. Parliament further suggests requiring high-risk AI not to "interfere in elections or contribute to the dissemination of disinformation" or cause certain other social harms.

Parliament's recommendations also propose that all conformity assessments be performed by approved third parties, whereas the Commission would allow providers of certain types of high-risk AI to assess themselves. Moreover, Parliament has included an individual right to redress for violations and whistleblower protections, which the Commission did not.

Any of these proposals could find their way into the final legislation that ultimately is adopted.

In addition, Parliament has recommended legislation amending the civil liability regime for AI systems. See Civil Liability Regime for Artificial Intelligence, P9_TA(2020)0276, annex B. An earlier Commission report had indicated such changes might be necessary, so proposed legislation may be forthcoming.

Selected European Commission Proposals

Prohibited Uses

- Harmful distortion of human behavior through subliminal techniques or exploitation of age or disability
- Real-time remote biometric identification systems in public places for law enforcement (limited exceptions)
- Governmental social scoring of individuals leading to discriminatory treatment across contexts or disproportionate to behavior

High-Risk Uses

- Many types of safety components and products
- Remote biometric identification and categorization of people
- Admission, assignment, and assessment of students
- Recruitment and other employment decisions
- Evaluation of creditworthiness and credit scoring (limited exception)
- Emergency services dispatch
- Law enforcement use for individual risk assessments, credibility determinations, emotion detection, identification of "deep fakes," evidentiary reliability evaluations, predictive policing, profiling of individuals, and crime analytics
- Immigration determinations
- Judicial decision-making
- Other uses the Commission later designates as high-risk

Requirements for High-Risk AI Systems

- Compliance (Providers)
  - Quality management system (compliance program), including regularly updated prescribed risk management system reflecting state of the art.
  - Pre-market conformity assessment (certifications valid for up to five years absent substantial modifications) and post-market monitoring with immediate correction and reporting requirements upon reason to consider system noncompliant or risky to health, safety, or fundamental rights.
  - Registration in EU database.
- Compliance (Others)
  - Third party that sells or installs under own brand, alters intended purpose, substantially modifies system, or incorporates system into product treated as provider.
  - Importers and distributors, among other obligations, must verify upstream compliance, not degrade compliance, and report if system risky to health, safety, or fundamental rights; distributors have correction duty upon reason to consider system noncompliant.
  - Professional users must operate consistent with provider instructions, monitor operations, input only relevant data, and assess data protection impact.

Selected European Commission Proposals

- Human Oversight—Humans with necessary competence, training, and authority must oversee operation and be able to:
  - Stop operation.
  - Disregard, override, or reverse the output.
- Documentation, Disclosure, and Explainability—To enable users to understand and control operation and to facilitate governmental oversight, providers must supply:
  - Concise, complete, correct, clear, relevant, accessible, and comprehensible instructions describing:
    - Characteristics, capabilities, and limitations of performance, including foreseeable unintended outcomes and other risks.
    - Human oversight and related technical safety measures.
    - Expected lifetime and necessary maintenance and care.
  - Detailed prescription of continuously updated technical documentation covering (in part):
    - Descriptions of system and development process, including compliance tradeoffs.
    - Monitoring, functioning, and control, including system's risks and mitigation.
    - Provider's risk management system.
- Robustness, Accuracy, and Cybersecurity—High-risk AI systems must:
  - Perform consistently at appropriate levels throughout their lifecycles, notwithstanding attempts at manipulation of training or operation or unauthorized alteration.
  - Meet training, validation, and testing data quality requirements.
- Penalties for Violations—Up to greater of:
  - €30 million.
  - 6% of global annual revenue.
- Retention of Records and Data
  - Automatic logging of operations, ensuring traceability.
  - Ten years:
    - Technical documentation.
    - Documentation of quality management system.
    - Certain conformity records.

Requirements for Certain AI Systems

- Transparency—For high- and low-risk systems, if applicable:
  - System must inform people they are interacting with an AI system unless obvious.
  - Individuals exposed to emotion recognition or (unless for law enforcement) biometric categorization system must be notified.
  - "Deep fakes" must be identified (qualified law enforcement and expressive, artistic, and scientific freedom exceptions).

For enforcement and other AI governance tasks, both the Commission and Parliament would look to assorted agencies. Industry regulators at both the EU and member state levels would continue to implement their mandates. New national authorities would fill gaps, coordinated by a new European Artificial Intelligence Board. How this mix of regulators gels, or not, will significantly influence how burdensome the contemplated regime becomes.

Meanwhile, the U.K. may be forging its own path. The Information Commissioner's Office (ICO) advises developers and users of AI on their obligations under the U.K. GDPR and other legislation. When warranted, the ICO can impose significant monetary penalties. The new Centre for Data Ethics and Innovation (CDEI) advises the government on AI regulation. In that capacity, CDEI recently called for clarification of how existing statutes and regulations apply to algorithmic bias. However, CDEI apparently won't be a regulator itself, leaving that task with the ICO and sector-specific agencies.

Whatever results from the debates in the EU and the U.K., exporters of AI-enabled products and services into Europe will likely be covered and should plan for compliance.

Conclusion

Like the technology, the legal landscape for AI is changing rapidly and requires monitoring. The second article in this pair will discuss what your company can do now to keep ahead of the curve.

With assistance from Darrel Pae, Katerina Kostaridi, and Elliot S. Rosenwald, Arnold & Porter Kaye Scholer LLP