MARCH 2023

PRIVACY AND DATA SECURITY
CLIENT ALERT

SHOOK
HARDY & BACON

Tips for Developing a Responsible AI Framework

Stated simply, artificial intelligence (AI) involves applying a set of instructions to a data set to solve a specific problem in an automated way. AI is common on the internet—think recommendations for feeds, personal assistants and search functions—and is becoming increasingly common in business—think network security functions, fraud detection, credit/risk rating, employee and applicant screening, consumer behavior prediction and medical diagnosis. But as AI development continues and its application becomes increasingly enmeshed in day-to-day operations and transactions, it is important for companies to develop a Responsible AI Framework. This governance framework serves multiple purposes including (i) educating business teams on the impact of using AI in assisting or replacing human decision-making; (ii) complying with evolving legal requirements; and (iii) avoiding unintended consequences like discrimination against, or adverse effects on, individuals. Below are a few issues to consider when developing an AI framework, which are intended to help issue-spot and guide discussions on incorporating AI into your business operations.

Accountability

A good framework starts with policies and procedures to guide business teams on the responsible use of AI. Such policies should address issues like:

- *Purpose specification*. What problem is the company trying to solve through use of AI? Is the AI system narrowly tailored to address the specific problem? Are there simpler or less technologically intensive means to solve the same problem?

SUBSCRIBE

ARCHIVE

Shook, Hardy & Bacon understands that companies face challenges securing information in an increasingly electronic world.

Shook guides its clients through an ever-changing patchwork of data security and data privacy laws and regulations, and helps its clients manage litigation and other risks associated with maintaining and using electronic information.

To learn more about Shook's Privacy and Data Security capabilities, please visit shb.com or contact:

Camila Tobón
Partner
303.285.5318
ctobon@shb.com

- *AI system oversight.* Developing the AI system is not the end of the story. Instead, the system must be continually checked to ensure that it is working as intended. There should also be human oversight over the decisions made by the AI system. This helps ensure that the AI system is not achieving any unintended consequences.

- *Recordkeeping.* It is important in the development of an AI system to keep records of the intended purpose, how the system will achieve the purpose, what data was used to train the algorithm, and the methodologies adopted to train, build and test the system. At its most basic, this can be a decision log tracking the different development and implementation decisions made along the way. But more robust records will contain details—and examples—of input and output.

- *Compliance with applicable laws.* While the legal landscape regulating AI is sparse, there are a few jurisdictions that either have passed or are considering specific requirements. For example, New York City recently enacted a law requiring notice of use of AI in hiring decisions as well as annual audits to assess potential bias in the use of such AI. Furthermore, the EU is considering a draft AI Regulation, which would take a risk-based approach to regulating AI systems, and the American Data Privacy and Protection Act—introduced in Congress last year—would require impact assessments for certain uses of AI. Being aware of current AI initiatives will help companies prepare for any proposed or pending legal requirements.

Transparency and Explainability

There are two aims here. First, letting individuals know that their data is being used with an AI system. This can be done through a general or specific privacy notice. Second, being able to explain how the AI system works. This involves explaining the logic behind the AI system, what data is input into the system, and how the system achieves the intended output.

Fairness

The impact of any AI system must be fair to the individuals affected by it. Such impacts include detrimental, discriminatory, biased, unexpected or misleading outcomes. For example, if an AI system is used to pre-screen job applicants, users must ensure that the system does not favor certain populations or disfavor others. Furthermore, where an AI system has specific legal or material consequences on individuals (e.g., credit decisions or job performance determinations) there must be some oversight to ensure that the system is treating all individuals fairly.

Privacy and Security

Aside from the transparency issue identified above, comprehensive privacy laws (like those in California, Colorado, Connecticut, Utah and Virginia as well as the EU and elsewhere) restrict profiling and automated decision-making activities. In addition, these laws provide individuals rights with respect to their information, including access, correction and deletion. These rights must be respected in any AI system lifecycle.

As for security, not only do you need safeguards in place to protect data from unauthorized use or disclosure, you also need them to ensure the integrity and reliability of the AI system itself. An outside actor tampering with the algorithm or its underlying dataset can lead not only to inaccurate results, but also to potentially harmful ones that may have material consequences on individuals (like denying someone's application for a loan or a job).

SHB.COM

ABOUT | CONTACT | SERVICES | LOCATIONS | CAREERS | PRIVACY

Unsubscribe | Forward to a Colleague | Privacy Notice