CORPORATE COMPLIANCE INSIGHTS
THE PREMIER GLOBAL NEWS SOURCE FOR COMPLIANCE, ETHICS, AUDIT & RISK

Adapting Model Governance to AI/ML Systems for Financial Institutions

by Reginald Goeke, Stephanie Monaco and Leslie Cruz ·
May 14, 2020

Firms need not throw away long-established governance principles to adapt to AI/ML. Rather, as Mayer Brown's Reginald Goeke, Stephanie Monaco and Leslie Cruz share, companies can modify their existing model risk management systems to address issues raised by AI/ML tools.

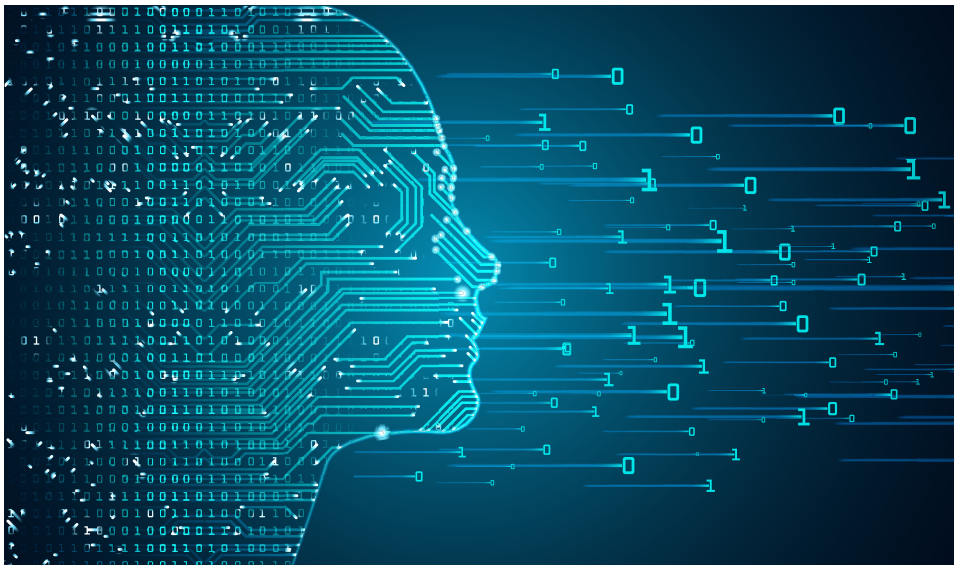

As the financial services industry increasingly turns to the use of artificial intelligence (AI) and machine learning (ML) for a variety of key functions — from underwriting to fraud or risk detection — many institutions grapple with establishing a governance model to manage the new risks posed by such systems. Because AI/ML systems pose risks similar to those posed by quantitative model generally, most financial institutions begin with their existing model risk management frameworks. Indeed, Federal Reserve Board Governor Lael Brainard has suggested that existing supervisory guardrails — including Federal Reserve Guidance SR 11-7 regarding "Guidance on Model Risk Management" and SR 13-19 regarding vendor risk management — are good places to start when implementing AI governance frameworks.

Nonetheless, the unique aspects of AI/ML models strain many elements of traditional model risk management systems and add risks not typically addressed through such systems. AI/ML models typically rely on massive amounts of data (some of which may be unstructured), placing increased importance on data validation and controls. AI/ML models may perform like "black boxes," and therefore can be more difficult to explain, test or validate than their more traditional counterparts. They may rely more on empirical data correlations and be less grounded in fundamental theory, posing a risk that decisions may be premised on impermissible or unintended bias. They also often evolve more rapidly than traditional quantitative models, sometimes changing on a continuous or automated basis. Furthermore, as AI/ML models develop in sophistication and speed, some businesses may deploy such models to make decisions without any human input or review, removing a key control process often used with traditional models. As financial services companies – from banks to asset managers to insurance companies – begin adopting AI/ML systems, they will need to tailor their existing policies to address these unique issues. Although the solution adopted by each company will be unique to its industry and needs, this article outlines the key areas for model risk management changes to address AI/ML challenges.

Model governance systems typically set out relevant policies and procedures and identify the roles and responsibilities for those involved in the process. The policies and procedures will govern the identification, development, documentation, testing, validation, implementation, modification, use and retirement of models, including any models (or aspects of them) that are acquired from third parties. As discussed below, when expanding those policies to include AI/ML systems, companies should consider the unique risks those systems pose at each step of the model process, including the data and inputs, the model itself and the outputs and usage of the model. The model governance roles and responsibilities typically include those responsible for (1) developing the models, (2) implementing an appropriate model control environment and (3) ensuring that the governance system is followed. As also discussed below, companies should consider potential

revisions to those roles and responsibilities to address the unique risks posed by AI/ML models.

Policies and Procedures Should Address Unique Risks of AI/ML Models

Many of the challenges companies face in adopting model risk management policies — including defining what qualifies as a "model," determining the appropriate level of documentation and adopting appropriate rules governing third-party models — apply equally in the context of AI/ML models. In addition to those issues, AI/ML models pose unique risks at each step of the model life cycle, including (1) model inputs and data, (2) model development and operation and (3) model output and use. Effective policies and procedures should anticipate those unique challenges.

Model Inputs/Data
Effective model governance incorporates a rigorous assessment of data quality, reliability, security and relevance; ensures the data is fit for purpose and representative of the data for which the model will be used; and ensures adequate documentation around the data and its sources. Because AI/ML models use massive amounts of data, sometimes in unstructured format, those models can challenge existing data quality review processes. Companies may increasingly adopt new data governance systems to govern such topics as data standardization and data quality, data reliability, data privacy and security, metadata management (e.g., maintaining data about the data) and data life cycle and architecture, among other items.

Further, while some models may digest massive data sets, they may not specify the data actually used by the model. Model governance should ensure that model owners identify the data that is available to any model and should ensure that such data is appropriate to be used in the decisions made by the model (e.g., do not violate privacy, regulatory or other legal requirements), that the data used has a consistent meaning across the database and over time and, if the company has obligations to identify the bases for its decisions, that the company can trace the source of the data if needed. Given the dynamic nature of data, companies should develop policies regarding the retention and overwriting of data, particularly with respect to data used to train AI/ML models, and should ensure that the frequency of data testing is consistent with the frequency of data acquisition.

Model Development and Operation
A key tenet of model risk management is that the model owners should understand the fundamental theory or logic underlying a model's application. AI/ML models, however, often operate by identifying new empirical relationships or correlations within data, which may be unintuitive or inconsistent with existing theory. The model governance system should ensure that model users can explain and support the theoretical bases for model decisions. In some cases, this can be done by using tools such as global variable importance charts, partial dependence charts and decision trees to help to understand and explain the factors driving model output.

A second critical element of model risk management is the existence and use of validators, independent of the model developers, who can ensure the model is fit for its stated purpose, can identify its assumptions and limits and can ensure that the model is stable and robust. For AI/ML models, which can often appear as "black box" models, such independent validation can be a challenge. The adoption of the tools noted above may help. In addition, companies may consider conducting more in-depth analyses of the model decisions and outputs, including back testing the model against multiple sets of data to ensure that the model results are stable and robust.

Model governance policies should also ensure that changes to the model are controlled, identified, documented and tested and validated as necessary. Yet AI/ML models may develop and change over time without human prompting or intervention. Companies may consider conducting frequent testing of the model to evaluate the extent to which key factors (or the weights assigned to them) change over time through the model's learning process. Companies may also consider conducting more frequent tests of the outputs (e.g., portfolio of loans, assets or insured risks created using the model) to determine the extent of any model drift over time.

Model Output and Use
Traditionally, the outputs of many models served as inputs to, or in conjunction with, human decision-making. With the inclusion of humans in the process, models both had an additional control function and an important feedback mechanism, because the frequency of user overrides provided important data regarding the reliability of model decisions. One of the benefits of AI/ML models is the ability to vastly accelerate decision-making processes (e.g., for credit or insurance underwriting), and in those uses there may be no opportunity for manual overrides. In those instances, companies should consider the use of automated controls built into models (e.g., setting maximum debt to income ratios for loans) to serve as guardrails and to constrain the scope of decisions that AI/ML systems can take without user input. Further, companies should also consider more frequent testing or audits of model-driven decisions to evaluate whether the model is performing as expected.

A second policy question many companies face is the extent to which the output of AI/ML models should be tested to identify any unintended bias or discriminatory impact. Even where AI/ML models do not have access to traits, such as race, gender or other prohibited categories, model decisions might still result in disparate impact on certain groups based on the use of data correlated to such traits. Companies should consider the extent to which they will review the model's input variables or use proxy data (such as the BISG method used by many regulators) to test model outputs for potential disparate impact, and they should develop a framework for evaluating the results of such analysis in light of the model's

potential benefits. The answers to those questions will likely differ by company and will depend on a number of variables, including the access to appropriate data (or sufficient proxy data), the extent to which alternative model designs can reduce the disparate impact, the product at issue and associated legal liability rules and the expectations of regulators, among other factors.

Relatedly, in many circumstances, financial institutions must explain to customers or regulators the bases for specific decisions generated by a model. Depending on the model, it may be difficult to pinpoint the specific variables that resulted in the model's decision. Companies subject to such rules might seek to incorporate software and tools that can help provide such explanations of the model output, including tools such as "Shapley values" and "LIME" tools.

Oversight Structure

An effective model governance system should have an appropriate oversight structure that (1) establishes roles and responsibilities of model developers, model validators and model users; (2) identifies those responsible for establishing model risk controls and those responsible for ensuring compliance with the model risk management policy; and (3) establishes senior-level responsibility for overall model risk management policy and reporting. In many organizations, model oversight often incorporates model risk committees, which typically include members from the compliance and/or risk management function, legal, the relevant business users of the model and certain technical experts. In addition to developing effective controls for the development and use of AI models, such groups can also facilitate early discussions of compliance, regulatory or legal issues before developers invest substantial time developing the models.

In the context of AI/ML models, companies should consider the inclusion of members of any data governance team (or any cybersecurity or privacy team) to ensure that data privacy, security and reliability factors are adequately considered. Further, in addition to including technical programming expertise, companies adopting AI/ML models that make decision through data mining and statistical extrapolation might consider supplementing the committee with relevant data analytics expertise.

In addition to typical model governance committees, some have suggested that companies using AI/ML models establish a group to address "AI ethics." Such ethical decisions might include determining appropriate trade-offs between developing increased predictive lift from models versus the potential for increased disparate impact. The AI ethics team might also set out policy regarding the circumstances in which the company will conduct a disparate impact analysis on AI/ML model outputs. Further, the AI ethics team can evaluate the potential privacy and security implications of various AI/ML models to determine when the company might deploy such models. In some companies, the AI ethics team may become the final arbiter of which AI applications the company will proceed to develop. Particularly with respect to larger financial institutions with multiple business lines, such a senior-level policy group might help to ensure that such decisions are consistent across all businesses.

Conclusion

While AI/ML models may pose a challenge to existing model governance principles and frameworks, companies can adapt their existing model governance policies to address those AI/ML models. Companies should consider: (1) closely aligning their model governance and data management and security systems, (2) updating their model governance policies to address the unique risks posed by AI/ML models, (3) adopting policies to supplement testing of model outputs, including predictive disparate impact testing and (4) supplementing their model governance committees to help address the ethical trade-offs associated with AI/ML models.

Reginald Goeke *is co-leader of Mayer Brown's litigation and commercial litigation groups. He has substantial experience representing clients with complex modeling-related investigations, litigations and assisting them in addressing model governance issues.*

Stephanie Monaco *is a partner of Mayer Brown's Corporate & Securities practice. She advises investment management firms, investment companies and hedge funds across a broad range of investment management needs.*

Leslie Cruz *serves as counsel and is a member of Mayer Brown's Corporate & Securities practice. She focuses her practice on representing registered investment companies, investment advisers and other financial institutions engaged in market, financial or investment management activities.*

MAYER | BROWN