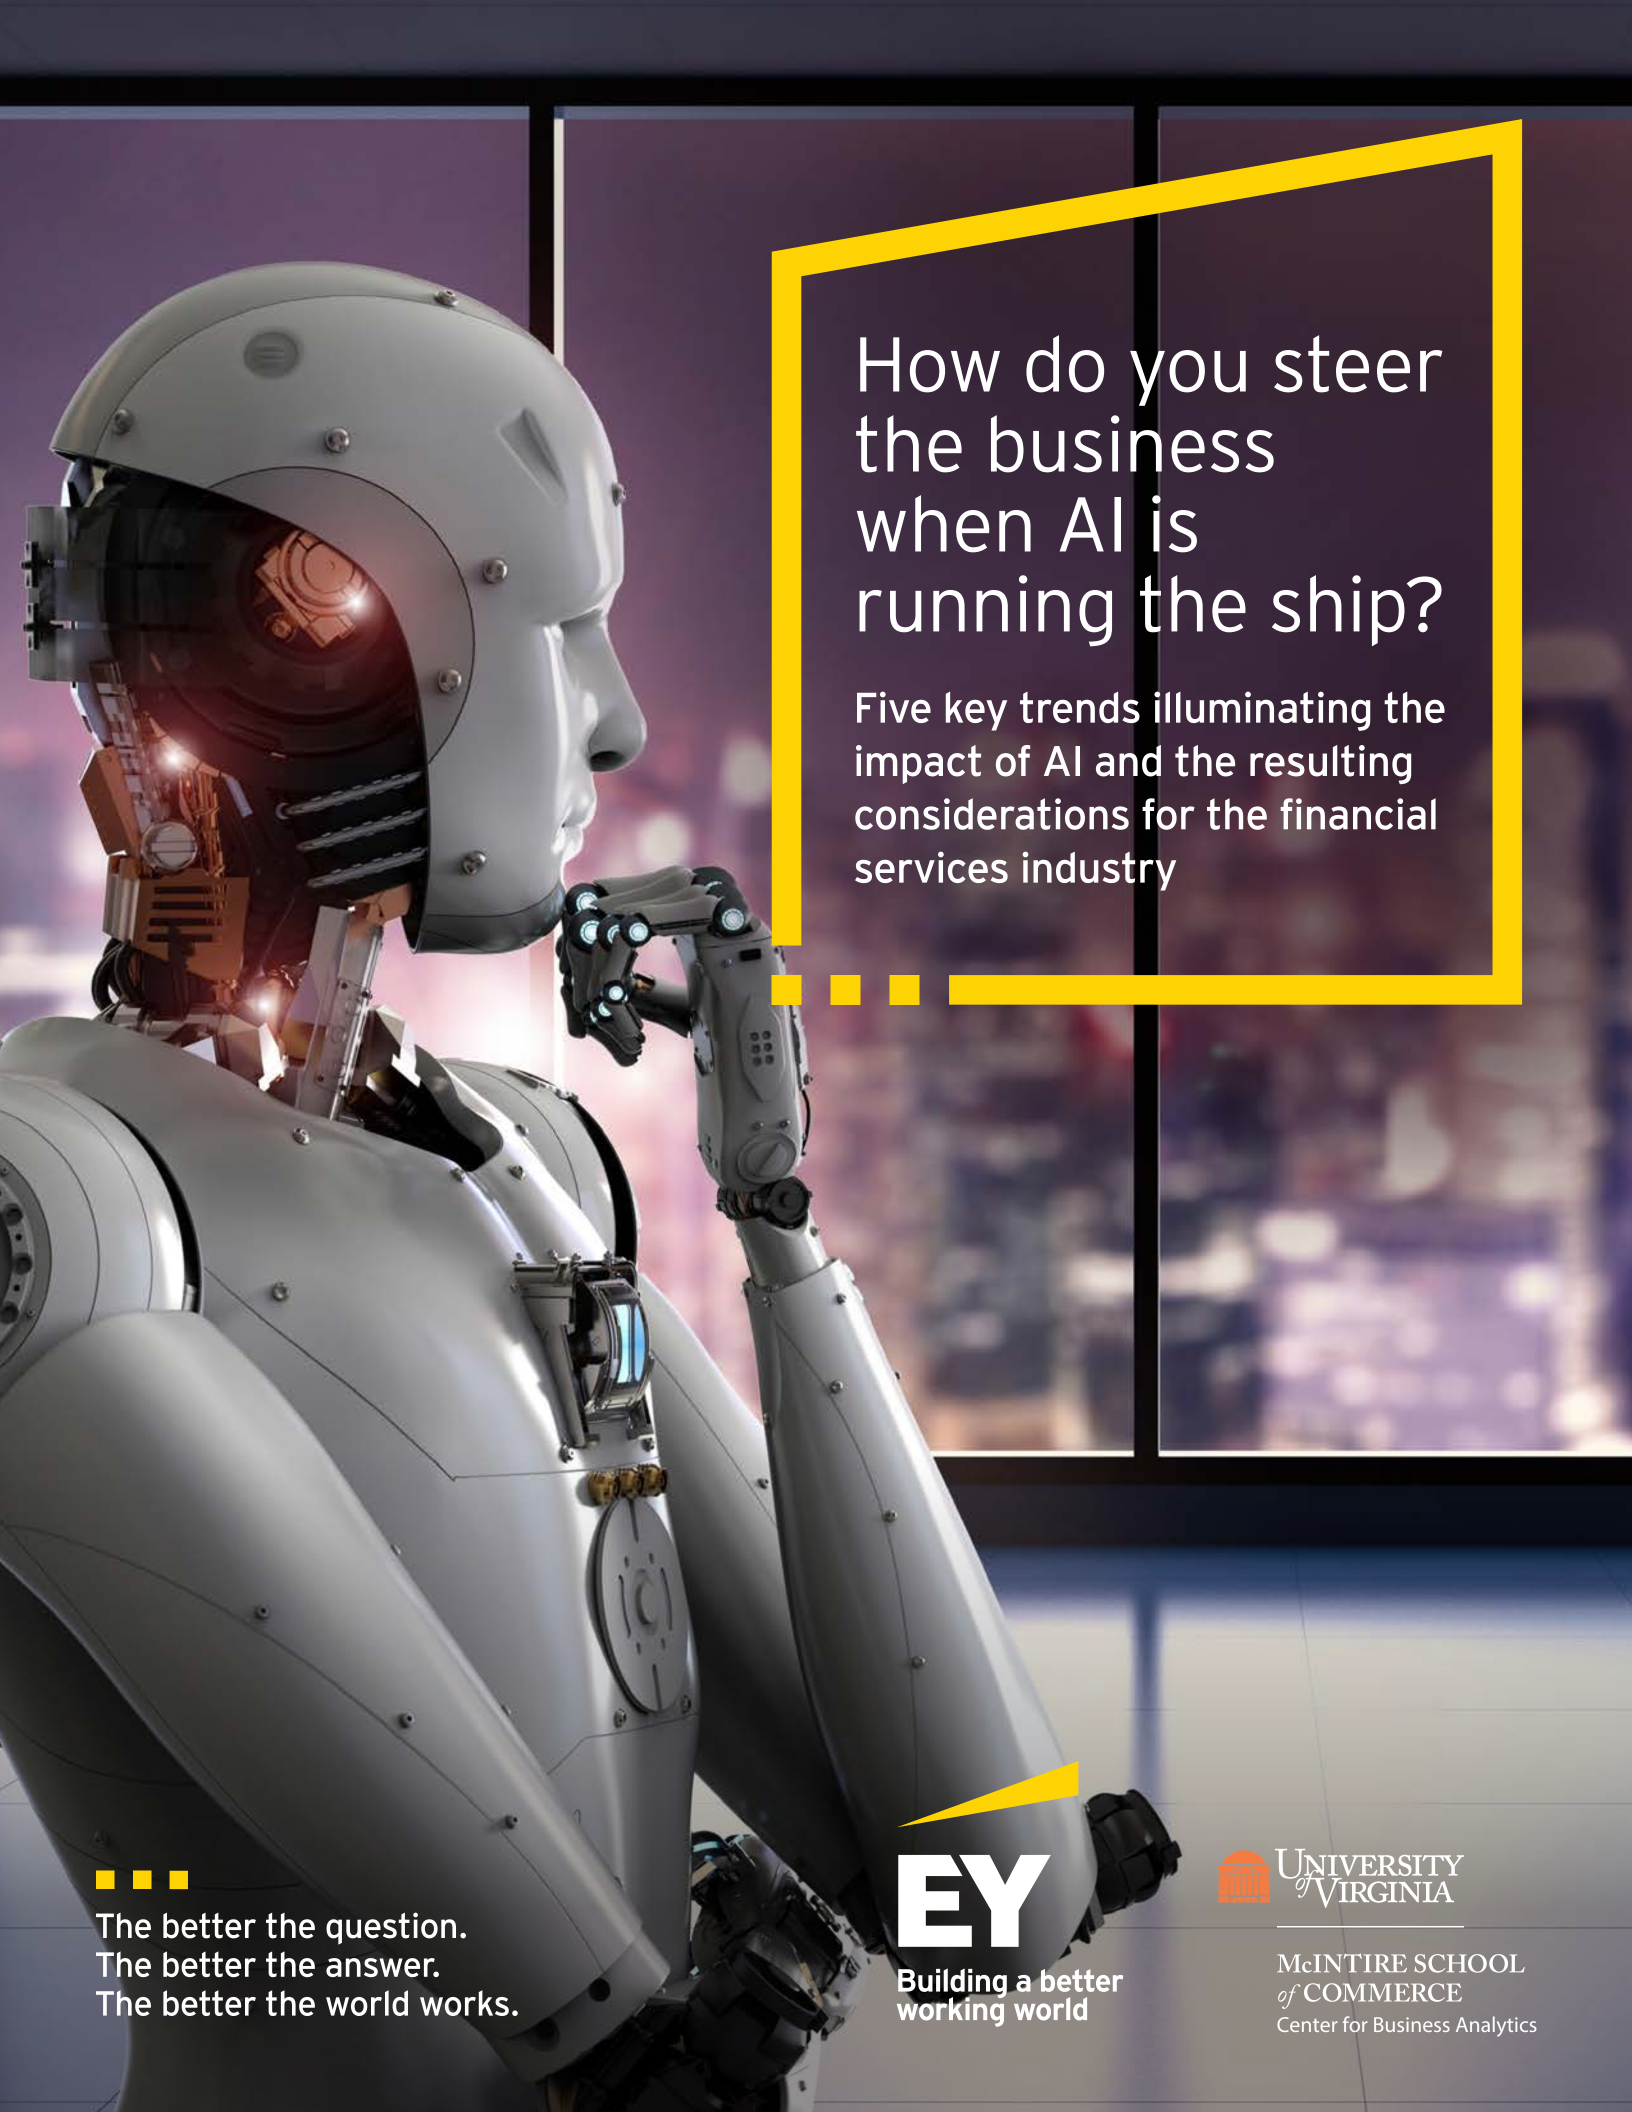

How do you steer the business when AI is running the ship?

Five key trends illuminating the impact of AI and the resulting considerations for the financial services industry

The better the question.
The better the answer.
The better the world works.

EY
Building a better working world

University of Virginia
McIntire School of Commerce
Center for Business Analytics

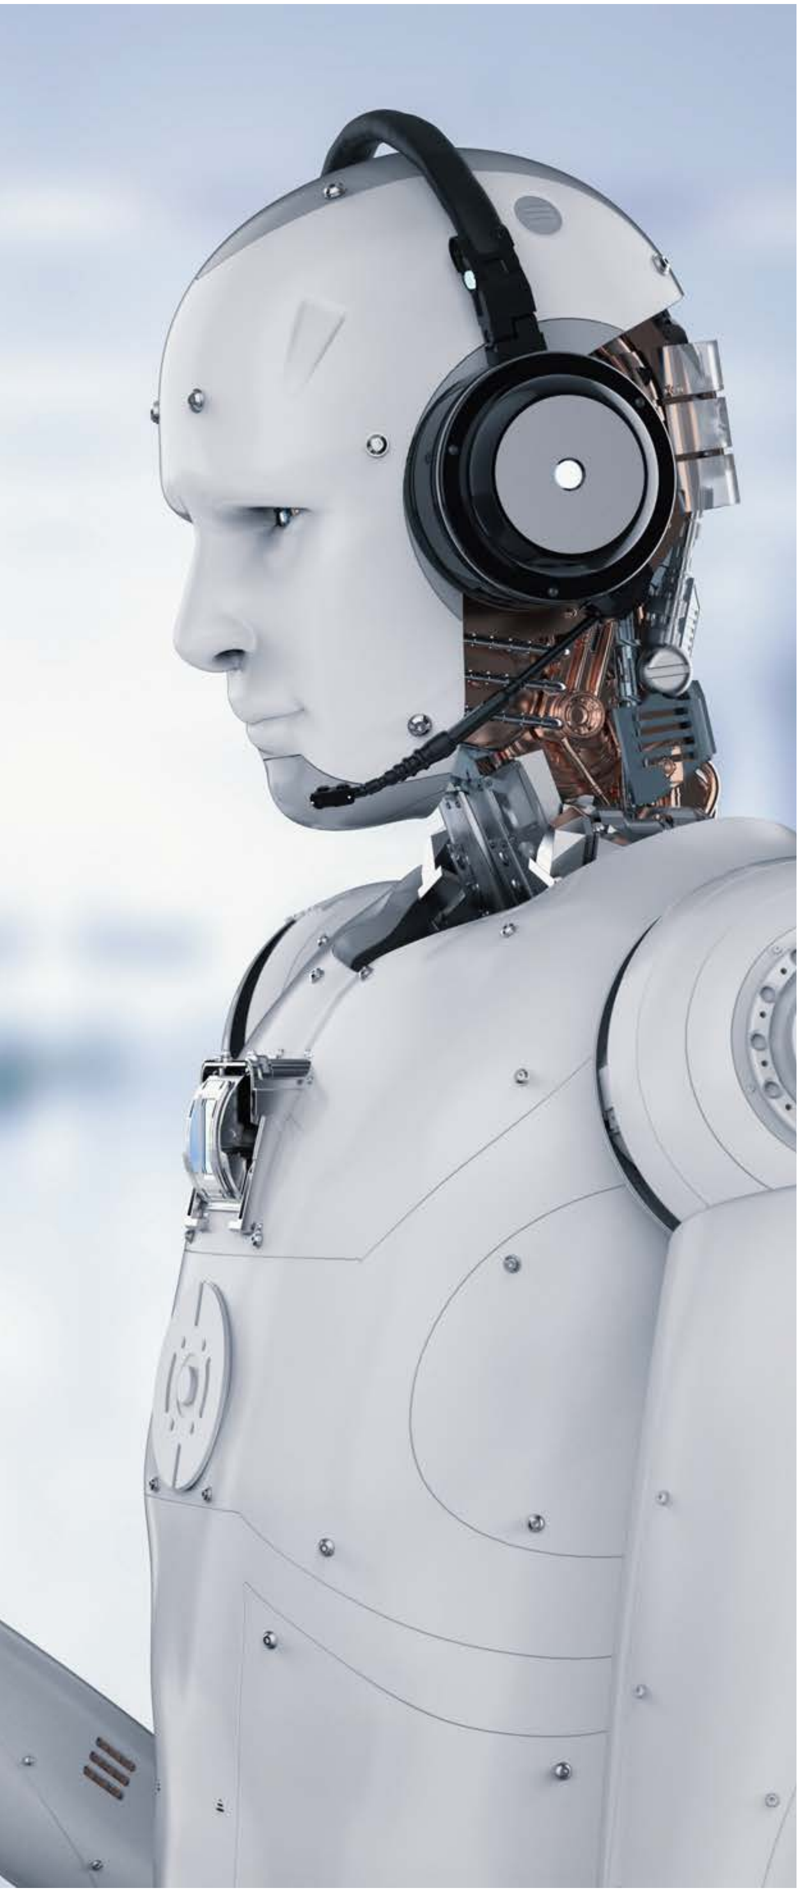

Introduction

Companies across many industries are exploring the use of artificial intelligence (AI) to enhance their processes and operations. The emergence and rapid growth of AI capabilities have resulted in a number of challenges for many firms, including financial services institutions — driving the need for organizations to evaluate the technology's impact on their existing processes and adapt to the changing environment.

Based on research conducted by the Center for Business Analytics in the McIntire School of Commerce at the University of Virginia (UVA), several key trends highlight some of the issues driving the need for organizations to evaluate the potential impact of the growing use of AI:

1 Machine learning (ML) models are being used for a wider array of macro- and micro-level prediction tasks[1]

"Datafication" has made rich new forms of data more available and accessible, thereby making myriad predictive analytics applications at both macro and micro levels of granularity more tangible than ever. For instance, new forms of ML-based predictive models are being developed to anticipate human behavior, social dynamics, economic activity, operational efficiency, financial market trends and security-related events at the individual, group, organizational and industry levels. These differences in prediction scope present an array of exciting possibilities for organizations but also introduce risk because AI is being used to tackle increasingly complex and challenging tasks.

---

1 Donald E. Brown, Ahmed Abbasi and Raymond Y.K. Lau, "Predictive Analytics," *IEEE Intelligent Systems*, 2015, Vol. 30, Issue 2, pp. 6–8; Donald E. Brown, Ahmed Abbasi and Raymond Y.K. Lau, "Predictive Analytics: Predictive Modeling at the Micro Level," *IEEE Intelligent Systems*, 2015, Vol. 30, Issue 3, pp. 6–8; Ahmed Abbasi, Raymond Y.K. Lau and Donald E. Brown, "Predicting behavior," *IEEE Intelligent Systems*, 2015, Vol. 30, Issue 3, pp. 35–43.

2 The complexity of ML models is greater than ever[2]

The most commonly used ML methods across multiple industries have evolved over the past 30 years, partly fueled by exponential growth in computing power and availability of training data. The state of the art has moved from simpler, highly interpretable linear model architectures to elaborate architectures that are increasingly complex and require less manual data engineering. The shift toward relatively more sophisticated, non-linear (often metaphorically described as "black box") models necessitates new governance processes.

3 With increased data monetization, questions of fairness and bias warrant careful consideration[3]

Bias can manifest itself not only during the data collection and preparation process but earlier on during the problem framing process and can propagate throughout the various stages of the ML process. The abundance of available data sources and organizational capability to process massive volumes of data, coupled with widespread adoption and advancement of ML, can lead to the development of insensitive and unfair models. Addressing potential bias through the use of ML should be an essential component of every organization's governance and risk management practices.

[2] Ahmed Abbasi, Suprateek Sarker and Roger H.L. Chiang, "Big Data Research in Information Systems: Toward an Inclusive Research Agenda," *Journal of the Association for Information Systems*, 2015, Vol. 17, Issue 2.

[3] Ahmed Abbasi, Jingjing Li, Gari Clifford and Herman Taylor, "Make 'Fairness by Design' Part of Machine Learning," *Harvard Business Review*, August 1, 2018, hbr.org/2018/08/make-fairness-by-design-part-of-machine-learning.

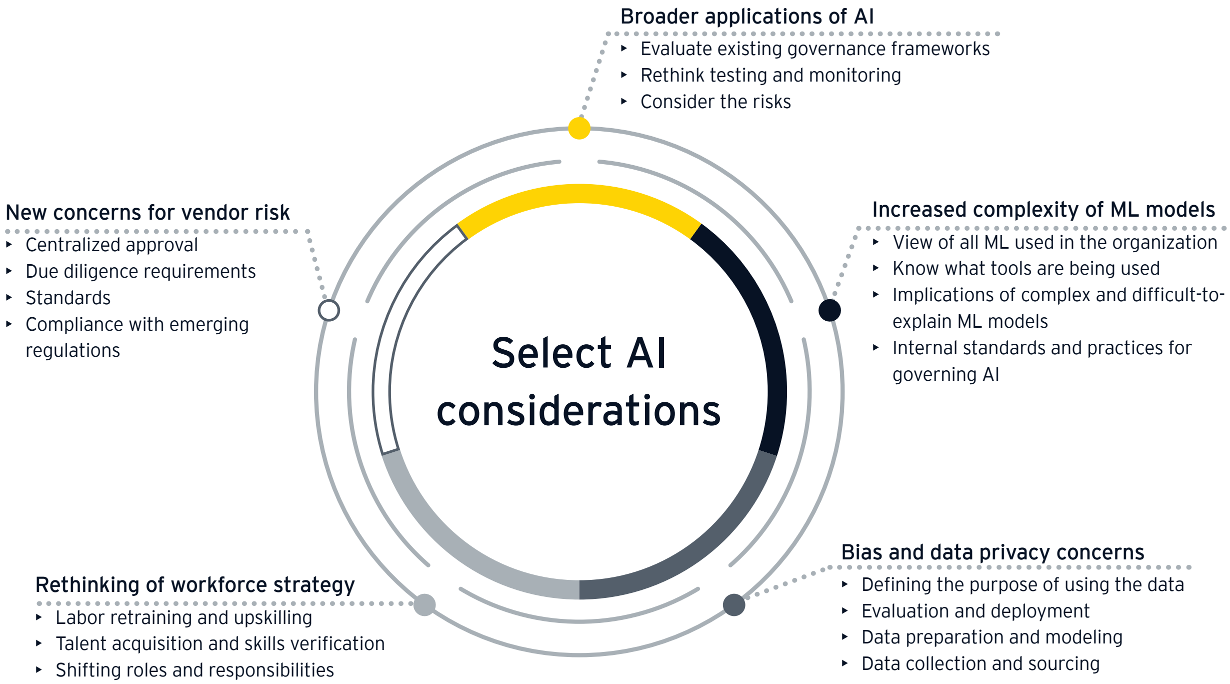

4 The proliferation of AI-enabled automation and augmentation requires an appropriate rethinking of workforce strategy[4]

Process automation and augmentation have generally been well-received in a variety of industries with perceived workforce benefits, including increased productivity, better quality of work, lower stress and an ability to focus on more interesting tasks. However, concerns about job security and obsolescence remain fairly prevalent. Moreover, workers in non-decision-making roles view the risks and rewards of increased automation differently than managers tasked with ultimately making such decisions. Organizations interested in gaining the most benefit from AI-enabled automation should consider appropriate workforce strategies.

5 The gold rush of AI vendors and offerings entails fresh vendor management guidelines[5]

The opacity and complexity of modern ML solutions, coupled with greater use of third-party and open-source offerings, present a new twist on the age-old "buyer beware" challenge. Predictive AI solutions in many contexts, including natural language processing (NLP), cybersecurity and fraud detection, may not perform as intended because of nuances in the application context and need for intelligent interpretation of model results. This performance expectancy gap underscores the need for a new class of vendor and technology management guidelines better geared toward AI applications.

Given these examples of key trends in the AI space, companies must consider how incorporating AI may impact different areas in their organization. Below are some examples of how these trends apply in the financial services industry, and specific considerations are highlighted for organizations to keep in mind as they expand their AI capabilities.

4 Ipsos Public Affairs and the Center for Business Analytics in the McIntire School of Commerce at the University of Virginia, "The Automation Divide," Ipsos, 2017, pp. 1–8.

5 David Zimbra, Ahmed Abbasi, Daniel Zeng and Hsinchun Chen, "The State-of-the-Art in Twitter Sentiment Analysis: A Review and Benchmark Evaluation," *ACM Transactions on Management Information Systems,* 2018, Vol. 9, Issue 2, p. 5.

Widespread and varied uses of machine learning

In alignment with UVA's research, financial services firms are exploring a wide variety of use cases for ML prediction tasks. The following are examples of macro- and micro-level prediction tasks that financial institutions are examining:

- Customer analytics applications (e.g., proactive customer engagement, experience personalization, targeted marketing)
- Call center analytics, such as identification and classification of customer complaints
- Data extraction from unstructured or complex document types
- Unstructured document classification
- Compliance use cases such as anti-money laundering (AML), know your customer (KYC), compliance testing automation
- Cybersecurity and defense against cyber attacks

Although ML enables financial services institutions to tackle challenges using new and enhanced methods and creates the potential for greater accuracy and efficiency, it also can introduce risk into the process.

After being trained with sufficient data, ML models have the potential to evaluate large quantities of data and perform predictive analytics that can help financial services organizations enhance existing business processes. For example, many banks currently have large teams that manually or semi-manually perform AML and KYC processes such as monitoring and investigating potentially fraudulent or suspicious activity. ML models can potentially process a lot of the same information and predict or identify which data may fall into the category of suspicious activity. Similarly, ML can help organizations address customer complaints by flagging information from customer interactions (such as call transcription and other customer touch points), traditionally a manual task. With such classifications, humans can determine what actions need to be taken or employ further automation to trigger actions.

As banks implement AI solutions for these types of tasks, there may be greater pressure on existing governance and controls frameworks, which are not designed to cover new AI techniques and the associated risk being introduced. As a result, financial services organizations may want to consider:

- **Evaluating existing governance frameworks:** How does incorporating AI into your organization impact your existing governance frameworks? How can existing frameworks be leveraged and adapted to keep pace with emerging technologies and their application to business processes?
- **Rethinking testing and monitoring under the lens of AI:** Does the AI application involve a feedback loop or automated creation and ingestion of training data? Do existing monitoring practices provide sufficient coverage for reviewing and validating model output, or should the cadence or scope be adjusted? What is the definition of success, and how do you know if the ML model is performing correctly or as expected? Whose responsibility is it?
- **Analyzing the risks:** How does the use of AI and ML impact both existing and evolving areas of risk, such as organizational, model, data, third-party, technology, compliance and legal, and business process risk? How can this risk be mitigated and/or proactively addressed?

These are only a few key considerations. Below, we will highlight other areas in response to the four additional trends from UVA's research, some of which overlap with the key points we have addressed.

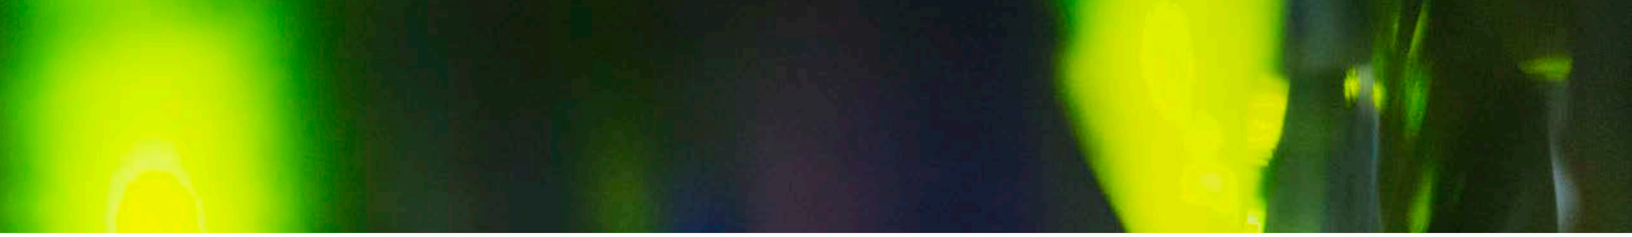

Increased complexity of machine learning models

The significant advancements in computing power and ML techniques have resulted in higher complexity in terms of how some machine learning models work. Advanced techniques may lead to ML models that are not readily explainable and outputs that cannot be easily interpreted or attributed to the variables driving them.

Although many financial services firms have well-established model risk management practices for stress-testing and other models leveraged in their organizations, there are additional considerations to keep in mind as firms repurpose and apply existing model risk management practices to ML models. The transparency, "explainability" and interpretability of machine learning models being used in financial services (as well as in other industries) are critical for sound decision-making, risk management and compliance with regulatory requirements. Without the transparency of ML models, it is difficult for stakeholders to trust the solution; this creates challenges for successfully using the technology for business and everyday processes. EY professionals have written about the value of trust when it comes to AI.

What makes an ML model difficult to understand? The reasons may include that:

- The model is applied to solve high-dimensional or nonlinear problems, which makes it difficult for humans to visualize, and therefore easily interpret.
- The model training process involves high-dimensional, voluminous or unstructured data.
- The linkage between inputs and outputs of the model is difficult to determine; for example, inputs may be correlated, or the input-output mapping may involve nonlinear transformations.

In simpler terms, if we don't understand how the model is structured or comes to its decisions, or if we don't understand the data being used to train the model well enough, we run into trouble regarding transparency and explainability.

Although there are tangible technical approaches that financial services institutions can use to increase the transparency of ML models being leveraged (outside of the scope of this article), robust governance is also at the core of the issue. The following are questions worth answering for your organization so you can successfully tackle model transparency:

- Do I know and have a comprehensive view of everyone who is building and/or leveraging machine learning in my organization?
- What tools are my teams using?
- Do we have sufficient internal standards and practices to control risk, and have we evaluated them for governing AI specifically?
- What are the potential implications of complex and opaque machine learning models for our business?

For financial services organizations, it is critical to have clear answers to these questions while exploring methods for developing explainable ML models that can help firms better understand model outcomes, decisions, factors behind a model's success or failure, and situations in which ML models can be trusted.

Risk of bias and data privacy concerns

There have been numerous high-profile examples of AI failures in the news, where large technology companies and other organizations attempted to create ML models that resulted in unintended consequences and sometimes alarming incidents of bias. Examples include controversial chatbot conversations, biased recruiting tools and discriminatory advertising.

Although many of the big news stories have involved well-known technology companies, there are similar risks for financial services firms. Published articles, books and research have explored the use of AI in developing models for expanding access to credit, which can benefit some groups of consumers more than others. The EY organization has also published a piece highlighting some of the risks. How do we ensure that ML models or other AI solutions do not systematically disadvantage specific individuals or groups? How do we protect the way personal data is used?

Financial services firms must consider how bias could manifest in any one of the stages of the ML implementation process: during data collection, data preparation, modeling, evaluation and/or deployment. These considerations are closely tied to recently resurging concerns regarding data privacy: who owns people's personal data and the insights derived from it, and how do we effectively control and authorize the use of it?

We are seeing different trends in countries across the world. For example, China enables companies and the government to maintain significantly greater control over individual personal data, and this data is used in different ways — including being monetized to drive development of additional consumer goods and products. Data monetization is also increasingly a trend in the rest of the world. However, in comparison, many European countries have established that individuals own the data they may or may not provide to companies and platforms and, in some cases, also own the insights derived from the use of their data. As a result, the European Union responded by developing regulations and legislation to protect against what may be considered improper use of data and provide general guidance on data protection and privacy as well as the exporting of personal data outside the EU. The US is also exploring this topic but has not yet passed comprehensive legislation to address it, though individual pieces of legislation have been passed for specific situations.

Given the varying approaches around the world to legislate data insight and action ownership, financial services firms and other organizations may want to consider:

- **Defining the purpose of using the data:** What is the problem statement? Is the designed problem framework fair and free from bias? What are you trying to solve with your ML model? Does the use of the data comply with existing regulatory guidance? Be specific, and then evaluate whether the problem you are seeking to address may run the risk of leading to biased results through direct or indirect correlations.
- **Data collection and sourcing:** How was the data collected for modeling purposes? What is the composition and distribution of the population the data was collected from, and is it representative of the population the model is being built for? Are there any potential inherent biases to this type of data, or idoes it reflect existing biases and prejudices?
- **Data preparation and modeling:** How was the data standardized, normalized, or otherwise modified or prepared for modeling purposes? Are there any potential risks with the approaches used? Which attributes or features are being selected, and how is the data being used in the model?
- **Evaluation and deployment:** How are the model results and outcomes being evaluated and tested for bias? How are the results being used to impact actions or decisions, and how is this being monitored?

Rethinking workforce strategy as a result of AI

Many financial services firms are developing proofs of concept using ML, natural language processing and other AI techniques to test specific use cases for augmenting business processes. Although most of the industry's organizations are still in the early exploration stages, the use of AI — especially in combination with other tools and techniques, such as robotic process automation and optical character recognition — creates significant implications for financial organizations and their people:

- **Labor retraining and upskilling:** The introduction of AI in combination with other automation can reduce or eliminate many manual tasks, thereby changing the nature of work for the employees previously handling those tasks. Organizations must develop and implement strategies for upskilling or retraining employees whose jobs are affected by AI and help shift the nature of their work from manual and repetitive tasks to more thoughtful and insightful ones. How do we develop appropriate training programs for these types of situations? How do we promote internal mobility in the organization to shift and retain existing employees?
- **Talent acquisition and skills verification:** Because the incorporation of AI is changing the nature of the roles needed, financial services firms are beginning to reassess the positions they need to fill and recruit for. There is a need for AI-focused resources, ranging from software engineers and data scientists to AI managers, enablers and strategists. Simultaneously, there is increased and widespread access to online learning tools, open-source software, software engineering boot camps, certifications and other ways of obtaining competency in all related areas as opposed to more traditional undergraduate and graduate education programs. Firms must consider how to effectively assess and verify that prospective candidates have the required skills for the open positions. Additionally, as AI technology continues to evolve rapidly, how does the organization keep its programs up to date and in pace with these changes? How do we enable sustainable talent in the organization and create a shared awareness of sound and ethical practices?
- **Shifting roles and responsibilities:** Incorporating AI also often requires an increased need for different teams or functions to collaborate more closely. Updated roles and responsibilities must enable this fluid communication between different teams, such as business units and technology teams. Some responsibilities may shift from one group to another, and these changes need to be communicated effectively. In other cases, the lines may appear more blurred about which team owns certain responsibilities. For example, for a use-case-specific ML model, who owns the model and model results? Who is responsible for continuously monitoring model performance and timely maintenance of developed models? Is it the technology team that created it, or is it the business unit that specified the requirements? If individual business units begin to assemble internal AI teams, how do these teams interface and split responsibilities with the organization's existing technology function? These are questions that must be considered as an organization re-evaluates existing roles and responsibilities and adapts them to a future state AI-enabled environment.

New considerations for vendor risk

Vendor and third-party risk is an established risk for which there is existing regulatory guidance, and many financial services institutions must conduct third-party risk management to be compliant. However, the increasingly rapid emergence of vendors in the AI space has implications that warrant additional consideration for organizations.

Organizations may lack visibility into these new vendors' products, resulting in the same challenges regarding lack of model transparency that we discussed. Additionally, a plethora of open-source software is available to developers online, for potential use internally at organizations. Open-source software tools are constantly and rapidly changing as well and may not have been assessed through existing third-party risk management processes.

Furthermore, vendor risk is amplified by the concerns around bias and data privacy that we addressed. Financial services organizations must evaluate their vendors to determine compliance with data protection regulations; if a vendor employs poor data protection standards or is negligent in some way, whether the vendor is aware of it or not, this could pose risk and potential liability for organizations using the vendor's services.

All of these considerations also apply to financial services firms partnering with or acquiring FinTech companies, which may leverage AI-based AI-based applications. In addition to assessing open-source tools and the technology stack being leveraged by vendors, firms should also assess the AI-based technology that partners or acquired companies use for the same reasons.

As a result, firms must revisit existing third-party risk management practices to confirm that they are addressing AI-specific considerations. Key considerations for mitigating third-party risk with AI and analytics vendors may include:

- **Centralized approval:** setting up a structure requiring centralized approval of vendor platforms and open-source libraries and establishing libraries of pre-vetted vendor and open-source solutions for data, infrastructure, code and other needs
- **Standards:** developing or enhancing existing standards detailing specific requirements for use of public open-source libraries and verifying that these standards meet firm and model risk management vendor management guidelines
- **Due diligence requirements:** developing requirements for conducting due diligence on third-party vendors and confirming sufficient transparency, especially given potential fourth-party risk or vendors' own use of open-source tools and software; similarly, developing due diligence requirements for potential partnerships, alliances and acquisition targets
- **Compliance with emerging regulations:** emerging regulatory guidance and legislation presenting further implications and complexities for third-party risk management and thereby potentially requiring evaluation of each vendor's own data protection standards

For more information on additional applications of AI and trends we are seeing in the financial services industry, see our other thought leadership pieces.

AI considerations: case study examples

To help illustrate some of the considerations we detailed, here are two case studies for reference and discussion.

Case study 1: Enhanced model explainability for auto insurance ML model

Description: An auto insurance client is seeking to develop and test approaches for explaining a deep learning model built to evaluate and classify images of car damage based on severity. The goal is to test multiple feature importance approaches on the deep learning model to develop an explanation for the model's classification results.

Challenge: The auto insurance firm's deep learning model can ingest an image of a car accident or auto damage sent by customers to the insurance company, then analyze and evaluate the severity of the car damage to classify images into different categories; this helps with claim calculations and improves accuracy in premium pricing. However, the auto insurance firm needs to generate sufficient textual explanations on why the deep learning techniques and methods arrived at their conclusions.

Approach: Several technical approaches can be used to measure or help identify feature importance in the model by decomposing the overall prediction to the most relevant features, using techniques such as Layer-wise Relevance Propagation (LRP) and Local Interpretable Model-agnostic Explanations (LIME). LRP methods help to identify the relative importance of pixels that contributed to the overall damage assessment. Sensitivity analysis is used to cross-validate the results. These methods help provide reasoning into how the model comes to its classification decision; this justification is collected along with the claim calculation to then provide a textual explanation to the claim adjuster and eventually the customer.

Case study 2: Operating model redesign for bank's shared services team

Description: A financial institution's internal shared services team seeks to increase the efficiency of its operating model and its ability to extract data from multiple structured and unstructured document types by incorporating ML and NLP technologies and additional vendor platforms into its existing processes.

Challenge: The shared services team receives information in various structured and unstructured data formats, such as paper documents, PDFs and emails across multiple channels. The team is responsible for extracting information from these documents and sending extracted data to the appropriate recipients. The team already uses tools to help extract data from structured documents, but there are additional use cases that could be better performed if supplemented with more vendor platforms and AI technologies. The team needs help with updating their operating model accordingly to align to the target state vision and developing a proof of concept for a more complex use case.

Approach: The bank's shared services team obtains support to develop a target state operating model, identify key considerations as a result of incorporating AI, and develop a proof of concept using ML/NLP for a specific use case or document type. The ML/NLP model offers robust classification and extraction results for identifying the document type and key attributes located in it. Because of the new risks and aspects of the target operating model, the key considerations identified address the updated structure for roles and responsibilities between the shared services team and the technology or data science teams. The considerations also address model risk management and model life cycle components and other important areas as part of the new operating model.

Authors

Yang S. Shim
Digital Enterprise Transformation Leader,
EY Americas Financial Services Advisory
Ernst & Young LLP

Ahmed Abbasi, PhD.
Director, McIntire Center
for Business Analytics
University of Virginia

Chandra Devarkonda
Managing Director, Financial
Services Advisory
Ernst & Young LLP

Sameer Gupta
AI and Advanced Analytics Leader,
Financial Services Advisory
Ernst & Young LLP

Tom Reilly
Senior Manager, Financial
Services Advisory
Ernst & Young LLP

Rita Kirzhner
Senior Consultant, Financial
Services Advisory
Ernst & Young LLP

EYG no. 003989-19Gbl
CSG No. 1906-3181145
ED None.