Help us choose our Top 10 Topics for 2018
www.deloitte.com/lu/InsideRisk2018g

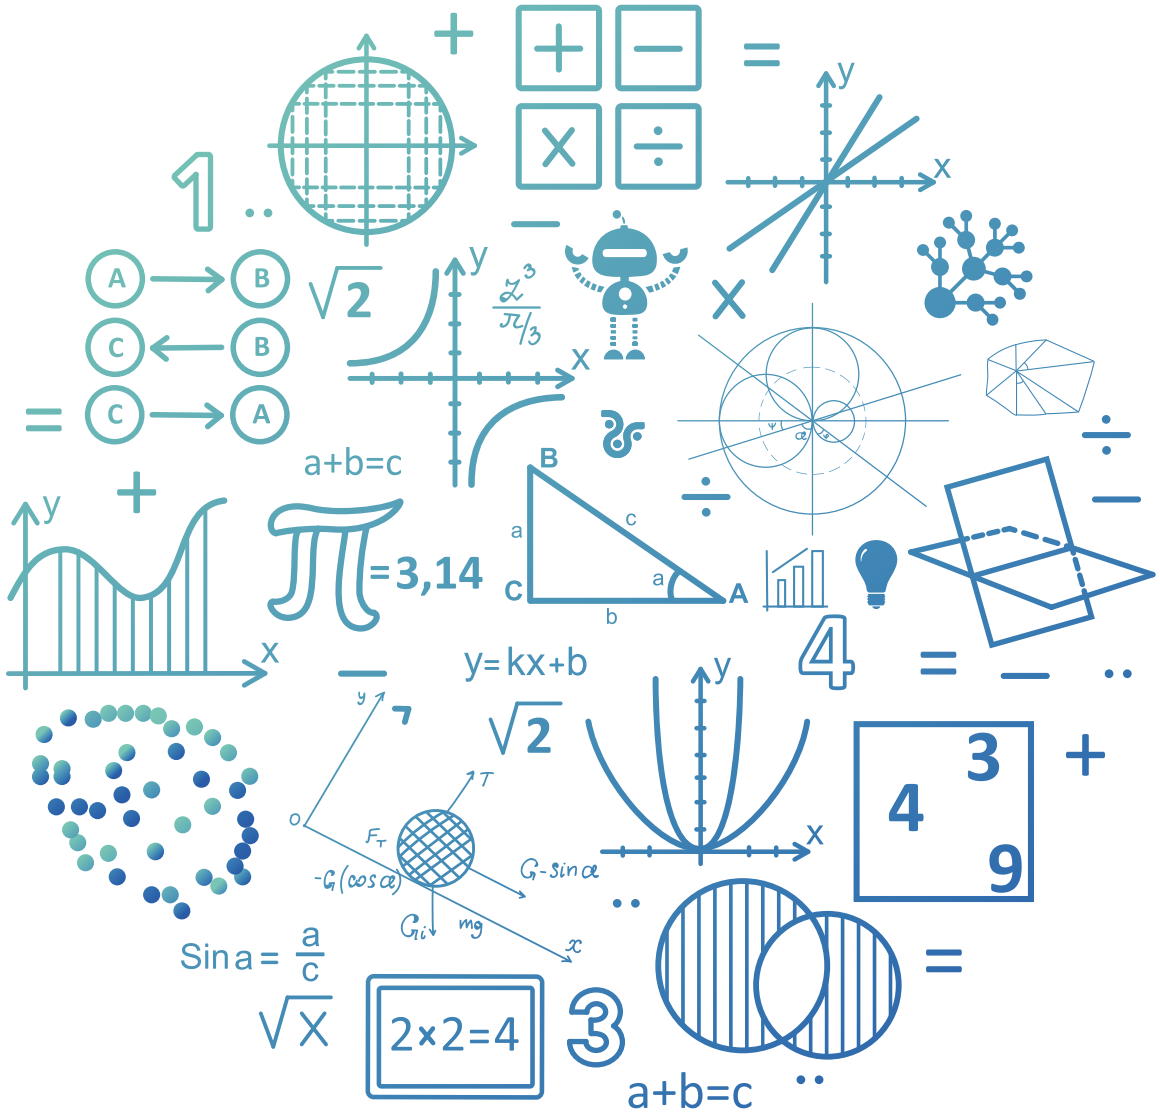

The algorithmic revolution is here

Nancy Albinson
Managing Director
Deloitte Risk and
Financial Advisory
Deloitte US

Dilip Krishna
Managing Director
Deloitte Risk and
Financial Advisory
Deloitte US

Yang Chu
Senior Manager
Deloitte Risk and
Financial Advisory
Deloitte US

The rise of advanced data analytics and cognitive technologies has led to an explosion in the use of algorithms across a range of purposes, industries, and business functions. Decisions that have a profound impact on individuals are being influenced by these algorithms—including what information individuals are exposed to, what jobs they're offered, whether their loan applications are approved, what medical treatment their doctors recommend, and even their treatment in the judicial system. At the same time, we're seeing a sharp rise in machine-to-machine interactions that are based in the Internet of Things (IoT) and powered by algorithms.

What's more, dramatically increasing complexity is fundamentally turning algorithms into inscrutable black boxes of decision-making. An aura of objectivity and infallibility may be ascribed to algorithms. However, these black boxes are vulnerable to risks, such as accidental or intentional biases, errors, and fraud, thus raising the question of how to "trust" algorithmic systems.

Embracing this complexity and establishing mechanisms to manage the associated risks will go a long way toward effectively harnessing the power of algorithms—and the upside is significant. Algorithms can be used to achieve desired business goals, accelerate long-term performance, and create differentiation in the marketplace. Organizations that adapt a risk-aware mindset will have an opportunity to use algorithms to lead in the marketplace, better navigate the regulatory environment, and disrupt their industries through innovation.

When algorithms go wrong

From Silicon Valley to the industrial heartland, the use of data-driven insights powered by algorithms is skyrocketing. Growth in sensor-generated data and advancements in data analytics and cognitive technologies have been the biggest drivers of this change, enabling businesses to produce rich insights to guide strategic, operational, and financial decisions. Business spending on cognitive technologies has been growing rapidly, and it's expected to continue at a five-year compound annual growth rate of 55 percent to nearly US$47 billion by 2020, paving the way for an even broader use of machine learning-based algorithms.[1] Going forward, these algorithms will be powering many of the IoT-based smart applications across sectors.

While such a change is transformative and impressive, cases of algorithms going wrong or being misused have also increased significantly. Some recent examples include:

- In the 2016 US elections, social media algorithms were cited for shaping and swaying public opinion by creating opinion echo chambers and failing to clamp down on fake news.
- During the 2016 Brexit referendum, algorithms were blamed for the flash crash of the British pound by six percent in a matter of two minutes.[2]
- Investigations have found that the algorithm used by criminal justice systems across the United States to predict recidivism rates is biased against certain racial classes.[3]
- Researchers have found erroneous statistical assumptions and bugs in functional magnetic-resonance imaging (fMRI) technology, which raised questions about the validity of many brain studies.[4]
- In several instances, employees have manipulated algorithms to suppress negative results of product safety and quality testing.
- Users have manipulated some artificial intelligence-powered tools to make offensive and inflammatory comments.
- According to a recent study, online ads for high-paying jobs were shown more often to men than to women.

1. "Worldwide Cognitive Systems and Artificial Intelligence Revenues Forecast to Surge Past $47 Billion in 2020, According to New IDC Spending Guide," Press release, IDC Research, Inc., October 26, 2016, http://www.idc.com/getdoc.jsp?containerId=prUS41878616.
2. Netty Idayu Ismail and Lukanyo Mnyanda," Flash Crash of the Pound Baffles Traders With Algorithms Being Blamed," Bloomberg, December 7, 2016, https://www.bloomberg.com/news/articles/2016-10-06/pound-plunges-6-1-percent-in-biggest-drop-since-brexit-result.
3. Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, "Machine Bias," Propublica, May 23, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.
4. "When science goes wrong (I) Computer says: oops," The Economist, July 16, 2016, http://www.economist.com/news/science-and-technology/21702166-two-studies-one-neuroscience-and-one-paleoclimatology-cast-doubt.

Increasing complexity, lack of transparency around algorithm design, inappropriate use of algorithms, and weak governance are specific reasons why algorithms are subject to such risks as biases, errors, and malicious acts. These risks, in turn, make it difficult to trust algorithms' decision choices and create concerns around their accuracy.

Many traditional checks and balances are designed for managing "conventional risks" where algorithm-based decisions aren't significantly involved, but these checks and balances aren't sufficient for managing risks associated with today's algorithm-based decision-making systems. This is due to the complexity, unpredictability, and proprietary nature of algorithms, as well as the lack of standards in this space.

These risks have the potential to cascade across an organization and negatively affect its reputation, revenues, business operations, and even regulatory compliance. That's why it's important for organizations to understand and proactively manage the risks presented by algorithms to fully capture the algorithms' value and drive marketplace differentiation.

What are algorithmic risks?

Algorithmic risks arise from the use of data analytics and cognitive technology-based software algorithms in various automated and semi-automated decision-making environments. Figure 1 provides a framework for understanding the different areas that are vulnerable to such risks and the underlying factors causing them.

- *Input data* is vulnerable to risks, such as biases in the data used for training; incomplete, outdated, or irrelevant data; insufficiently large and diverse sample size; inappropriate data collection techniques; and a mismatch between the data used for training the algorithm and the actual input data during operations.
- *Algorithm design* is vulnerable to risks, such as biased logic, flawed assumptions or judgments, inappropriate modeling techniques, coding errors, and identifying spurious patterns in the training data.
- *Output decisions* are vulnerable to risks, such as incorrect interpretation of the output, inappropriate use of the output, and disregard of the underlying assumptions.

These risks can be caused by several underlying factors:

Human biases: Cognitive biases of model developers or users can result in flawed output. In addition, lack of governance and misalignment between the organization's values and individual employees' behavior can yield unintended outcomes.
Example: Developers provide biased historical data to train an image recognition algorithm, resulting in the algorithm being unable to correctly recognize minorities.

Technical flaws: Lack of technical rigor or conceptual soundness in the development, training, testing, or validation of the algorithm can lead to an incorrect output.
Example: Bugs in trading algorithms drive erratic trading of shares and sudden fluctuations in prices, resulting in millions of dollars in losses in a matter of minutes.

Usage flaws: Flaws in the implementation of an algorithm, its integration with operations, or its use by end users can lead to inappropriate decision making.
Example: Drivers over-rely on driver assistance features in modern cars, believing them to be capable of completely autonomous operation, which can result in traffic accidents.

Security flaws: Internal or external threat actors can gain access to input data, algorithm design, or its output, and manipulate them to introduce deliberately flawed outcomes.
Example: By intentionally feeding incorrect data into a self-learning facial recognition algorithm, attackers are able to impersonate victims via biometric authentication systems.

It's important for organizations to understand and proactively manage the risks presented by algorithms.

Figure 1: Framework for understanding algorithmic risks

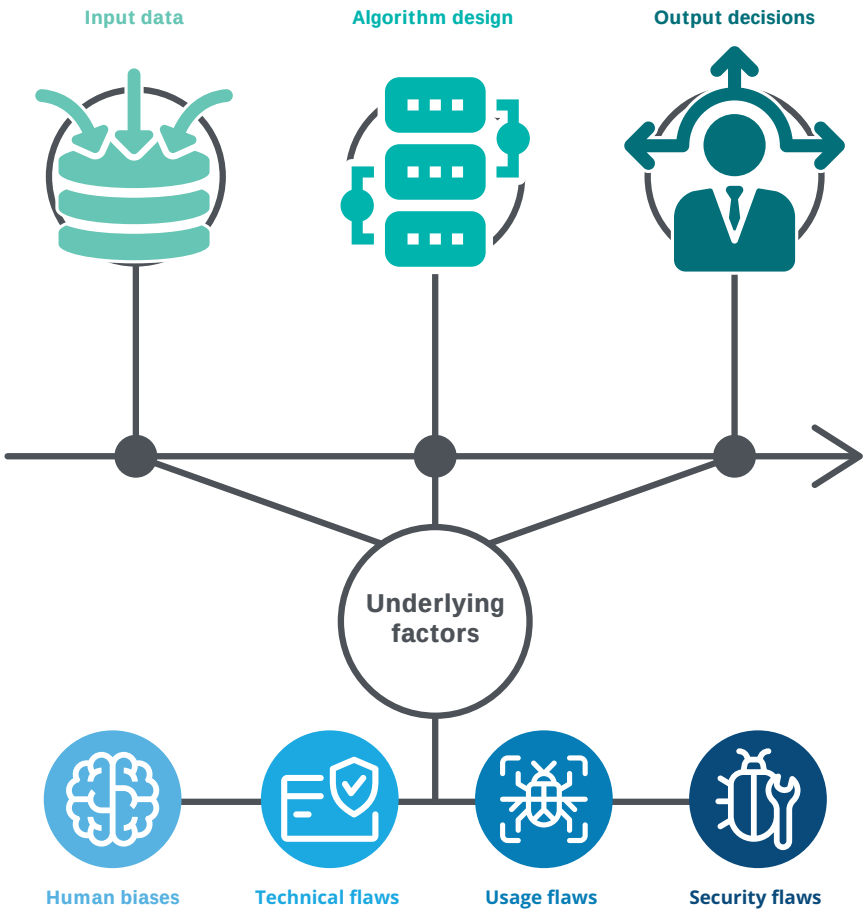

Why are algorithmic risks gaining prominence today?

While algorithms have been in use for many years, the need to critically evaluate them for biases, lack of technical rigor, usage flaws, and security vulnerabilities has grown significantly in recent times. This growing prominence of algorithmic risks can be attributed to the following factors:

Algorithms are becoming pervasive

With the increasing adoption of advanced data analytics and machine learning technology, algorithm use is becoming more prevalent and integral to business processes across industries and functions. It's also becoming a source of competitive advantage. One study predicts that 47 percent of jobs will be automated by 2033.[5] Figure 2 highlights some prominent business use cases of algorithms.

These use cases are expected to significantly expand in the near future, given the tremendous growth in IoT-enabled systems. These systems can lead to the development and proliferation of new algorithms for connecting IoT devices and enabling smart applications.

Machine learning techniques are evolving

Improvements in computational power coupled with the availability of large volumes of training data—data used to train algorithms—are driving advancements in machine learning. Neural networks are becoming an increasingly popular way of implementing machine learning. Techniques such as deep learning are being used for tasks like computer vision and speech recognition.

These advances in machine learning techniques are enabling the creation of algorithms that have better predictive capabilities but are significantly more complex.

Algorithms are becoming more powerful

Not only are algorithms becoming more pervasive, but the power and responsibility entrusted to them is increasing as well. Due to advancements in deep learning techniques, algorithms are becoming better at prediction and making complex decisions. Today, algorithms are being used to help make many important decisions, such as detecting crime and assigning punishment, deciding investment of millions of dollars, and saving the lives of patients.

5 Carl Benedikt Frey and Michael Osborne, "The Future of Employment," Oxford Martin Programme on Technology and Employment, September 17, 2013, http://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf.

Algorithms are becoming more opaque

Algorithms run in the background and often function as black boxes. With their internal workings and functioning largely hidden from developers and end users, monitoring algorithms can be difficult. Many new machine learning techniques, such as deep learning, are so opaque that it's practically impossible to understand what they deduce from training data and how they reach their conclusions—thus making it hard to judge their correctness. This difficulty in understanding algorithmic decisions, coupled with the unpredictability and continuously evolving nature of algorithms, makes inspecting them a challenge.

Algorithms are becoming targets of hacking

Machine learning algorithms are exhibiting vulnerabilities that can be exploited by hackers. A common vulnerability is the data used to train algorithms. Manipulating that training data as it's presented to the algorithms results in skewed algorithms that produce erroneous output, which in turn leads to unintended actions and decisions. In addition, attackers are also tampering with the actual live data to which the algorithms are applied. A recent report revealed that cyber criminals are making close to US$5 million per day by tricking ad purchasing algorithms with fraudulent ad click data, which is generated by bots rather than humans.[6]

Figure 2. Algorithm use across business functions

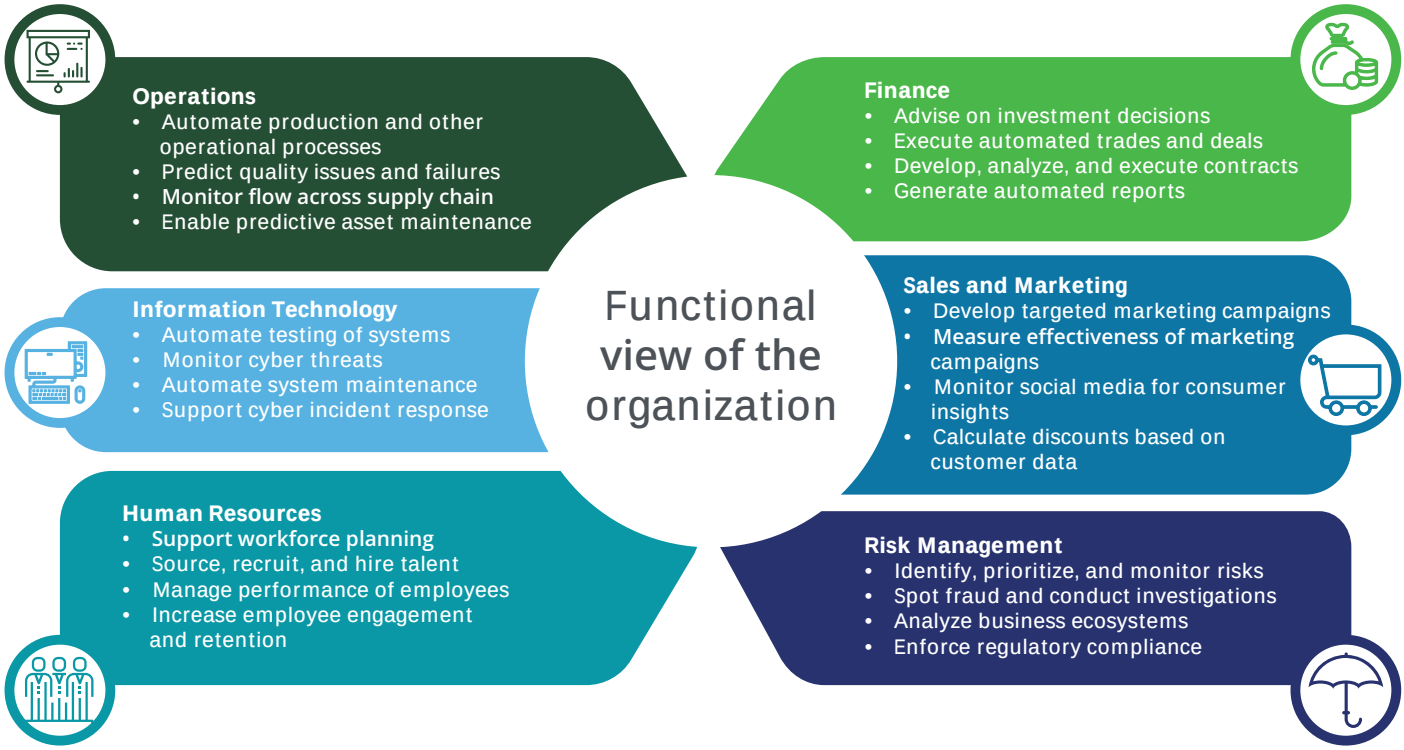

6. Thomas Fox-Brewster, " 'Biggest Ad Fraud Ever': Hackers Make $5M A Day By Faking 300M Video Views," Forbes, December 20, 2016, https://www.forbes.com/sites/thomasbrewster/2016/12/20/methbot-biggest-ad-fraud-busted/#4324f6c74899.

What do algorithmic risks mean for your organization?

As noted previously, data analytics and cognitive technology-based algorithms are increasingly becoming integral to many business processes, and organizations are investing heavily in them. Nevertheless, if the issues highlighted in this report aren't adequately managed, the investments may not yield the anticipated benefits. Worse yet, they may subject organizations to unanticipated risks.

The immediate fallouts of these algorithmic risks can include inappropriate and potentially illegal decisions relating to:

- Finance, such as inaccurate financial reporting resulting in regulatory penalties and shareholder backlash, as well as taking on unanticipated market risks beyond the organization's risk appetite.
- Sales and marketing, such as discrimination against certain groups of customers in product pricing, product offerings, and ratings.
- Operations, such as credit offers, access to health care and education, and product safety and quality.
- Risk management, such as not detecting significant risks.
- Information technology, such as inadequate business continuity planning and undetected cyber threats.
- Human resources, such as discrimination in hiring and performance management practices.

Algorithms operate at faster speeds in fully automated environments, and they become increasingly volatile as algorithms interact with other algorithms or social media platforms. Therefore, algorithmic risks can quickly get out of hand.

Financial markets have already experienced significant instability because of algorithms. The most high-profile instance was the flash crash of 2010, which sent the Dow Jones Industrial Average on a 1,000-point slide.[7]

Algorithmic risks can also carry broader and long-term implications for an organization, such as:

- Reputational risks: The use of algorithms can significantly increase an organization's exposure to reputation risks. This is particularly true if the various stakeholders believe that the workings of the algorithm aren't aligned to the ethics and values of the organization, or if the algorithms are designed to covertly manipulate consumers, regulators, or employees.
- Financial risks: Errors or vulnerabilities in algorithms, especially those used for financial and strategic decision-making, can result in significant revenue loss for organizations and negatively affect the integrity of their financial reporting.
- Operational risks: As algorithms are used to automate supply chain and other operational areas, errors can result in significant operational disruptions.
- Regulatory risks: Algorithms making decisions that violate the law, circumvent existing rules and regulations, or discriminate against certain groups of people can expose organizations to regulatory and legal actions.
- Technology risks: The wide-scale use of advanced algorithms can open up new points of vulnerability for IT infrastructure.
- Strategic risks: With algorithms being used increasingly as sources for strategic decision-making, errors or vulnerabilities within them can put an organization at a competitive disadvantage.

7. Silla Brush, Tom Schoenberg, and Suzi Ring, "How a Mystery Trader With an Algorithm May Have Caused the Flash Crash," Bloomberg, April 22, 2015, https://www.bloomberg.com/news/articles/2015-04-22/mystery-trader-armed-with-algorithms-rewrites-flash-crash-story.

It's important for organizations to evaluate their use of algorithms in high-risk and high-impact situations.

What's different about managing algorithmic risks?

With the growing urgency of algorithmic risk management, it's important to note that conventional risk management approaches may not be effective for that purpose. Instead, organizations should rethink and reengineer some of their existing risk management processes due to the inherent nature of algorithms and how they're used within organizations. For example, algorithmic risk management can't be a periodic point-in-time exercise. It requires continuous monitoring of algorithms, perhaps through the use of other algorithms. Three factors differentiate algorithmic risk management from traditional risk management:

Algorithms are proprietary

Algorithms are typically based on proprietary data, models, and techniques. They're considered trade secrets and sources of competitive advantage. As a result, organizations are typically unwilling to share data, source code, or the internal workings of their algorithms. This makes it difficult for regulatory agencies and outside watchdog groups to monitor them.

Algorithms are complex, unpredictable, and difficult to explain

Even if organizations were to share their algorithm codes, understanding them may be difficult because of their inherent complexity. Many of today's algorithms are based on machine learning and other advanced technologies. They evolve over time based on input data. In many cases, even the teams that develop them might not be able to predict or explain their behaviors.

Machine learning algorithms can even develop their own languages to communicate with each other. This is an area with both tremendous potential and risk, given the anticipated growth in IoT and machine-to-machine communications.

There's a lack of standards and regulations

In financial services, model validation has become very important over the past few years, and there are widely accepted standards such as SR 11-7: Guidance on Model Risk Management. However, these standards have limitations when applied to complex machine learning techniques such as deep learning. Currently, no widely accepted cross-industry standards exist to govern many types of machine learning algorithms, including processes around data collection, training, and algorithm design. As a result, there's a lack of consistent business controls for development, implementation, and use of algorithms. Developers frequently use their experience and theoretical knowledge to make these decisions without management oversight, leading to variations in processes and the increased likelihood of errors.

In addition, regulations in this space are still evolving and apply to only a limited set of algorithms, such as those relating to capital management and stress testing in the banking sector. While there have been some attempts to broadly regulate the use of algorithms (especially in Europe), there's still a lack of clarity about, and many unanswered questions around, how these regulations will be implemented.
This lack of standards and regulations makes it difficult to drive accountability and fairness in the use of algorithms.

Figure 3. A framework for algorithmic risk management

Is your organization ready to manage algorithmic risks?

The rapid proliferation of powerful algorithms in many facets of business is in full swing and is likely to grow unabated for years to come. The use of intelligent algorithms offers a wide range of potential benefits to organizations, from innovative products to improved customer experience, to strategic planning, to operational efficiency, and even to risk management. Yet as this article has discussed, some of those benefits could be diminished by inherent risks associated with the design, implementation, and use of algorithms—risks that are also likely to increase unless organizations invest effectively in algorithmic risk management capabilities.

It's not a journey that organizations must take alone. The growing awareness of algorithmic risks among researchers, consumer advocacy groups, lawmakers, regulators, and other stakeholders should contribute to a growing body of knowledge about algorithmic risks and, over time, risk management standards. In the meantime, it's important for organizations to evaluate their use of algorithms in high-risk and high-impact situations and implement leading practices to manage those risks intelligently so algorithms can be harnessed for competitive advantage.