

The algorithmic revolution

What are algorithms?

Algorithms are processes or sequences of instructions used to analyze data, solve problems and perform tasks. For example, when you make an online purchase, algorithms will commonly record your purchase and develop recommendations for other things you may want to buy from the online retailer.

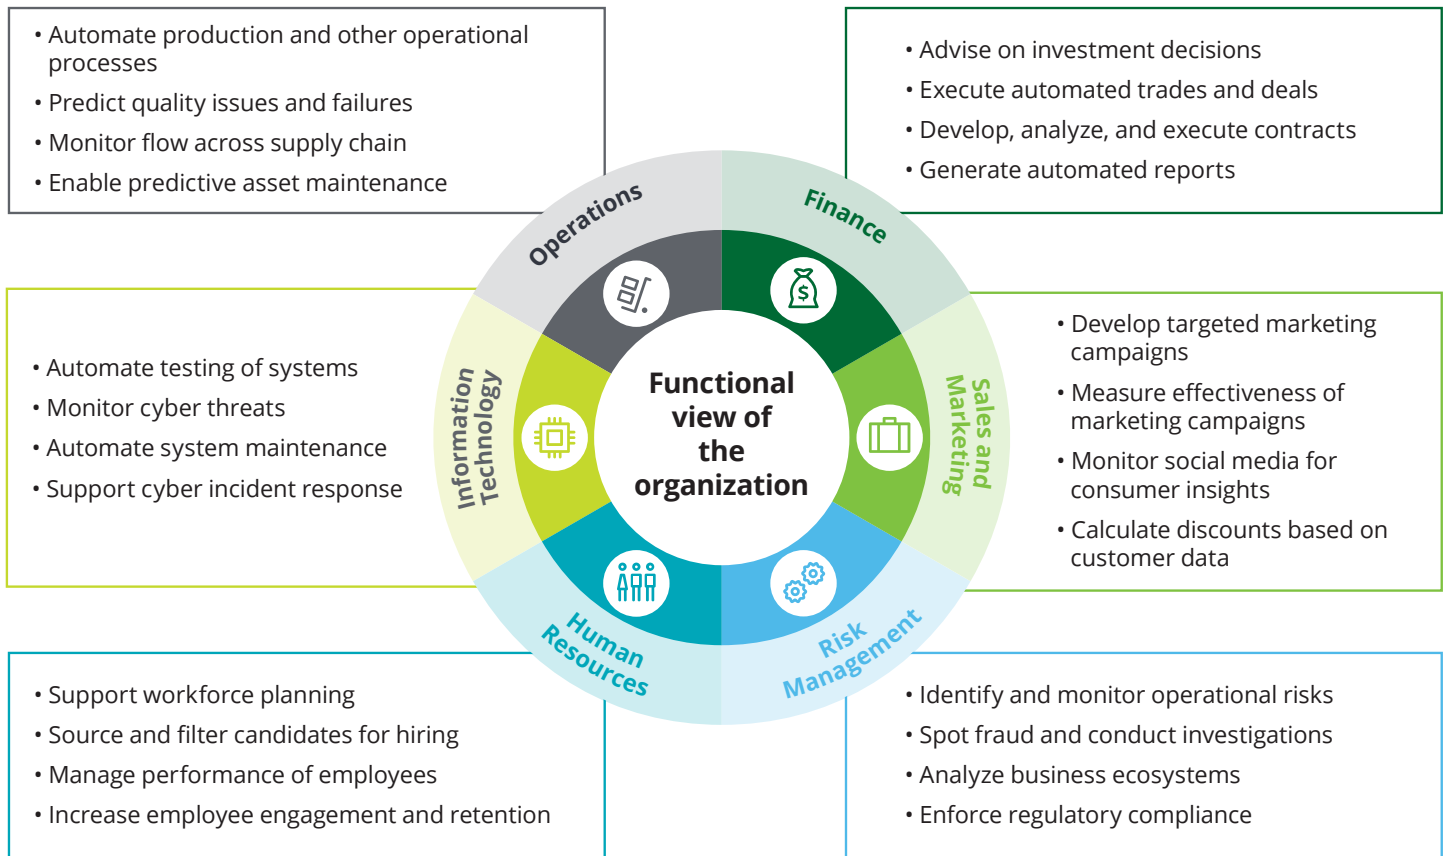
Initially, algorithms were programmed by people to do their jobs; however,

“self-learning” algorithms are increasingly replacing pre-programmed algorithms. A self-learning algorithm is able to expand, without human intervention, the range of tasks it can perform based upon the information it receives and processes.

What do algorithms do?

Algorithms are now often an integral part of daily life for individuals and enterprises alike. However, the rise of advanced data analytics and cognitive technologies is far broader—in fact, it 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. The below chart illustrates how algorithms can work in an enterprise. ➔



Algorithmic risk

Algorithms gone wrong

As indicated in the graphic, algorithms can increase performance in an enterprise in many ways, by automating some existing processes and tackling new activities previously not feasible using manual processes. However, algorithms can (and do) go wrong and can have serious adverse effects when they do. In the example above—where an online purchase generates algorithmic recommendations for additional purchases—a wrong or offensive recommendation could cause the customer to avoid the retailer in the future. Multiply that across a class of customers and there is the potential for a business meltdown.

The implications of “algorithms gone wrong” for the community or society at large can be far broader. Some examples:

- Researchers found erroneous statistical assumptions and bugs in functional magnetic-resonance imaging technology, which raised questions about the validity of certain brain studies.
- Employees of a manufacturer were accused of installing hidden software that suppressed negative results of product testing.
- A bank was fined more than \$100 million for deceitful use of algorithms on its trading platform to increase its profits.
- Users manipulated artificial intelligence tools to make inflammatory comments.

Why do algorithms go wrong?

There are several causes of algorithmic risk, as illustrated below.

- *Input data* is vulnerable to risks such as biases in the data; incomplete, outdated, or irrelevant data; inappropriate sample size or data collection techniques; and mismatches between the data used for developing the algorithm and 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 overfitting of algorithm to 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:

- *biases* of developers or users, or misalignment between values and individual behavior;
- *technical flaws* arising from a lack of technical rigor or conceptual soundness in the development, training, testing, or validation of the algorithm;
- *usage flaws* in the implementation of an algorithm, its integration with operations, or its use by end users; and
- *security flaws* that permit internal or external parties to access input data, algorithm design, or its output and manipulate them to introduce deliberately flawed outcomes.

A growing concern

Concerns about algorithm risk have gained, and continue to gain, prominence. First, increased use of and reliance upon powerful algorithms across industries and processes have made users, and even the general public, more aware of algorithms, including their vulnerabilities. These vulnerabilities are likely to increase in the near term, given the rapid growth of the “Internet of Things” and other technological advances, coupled with the fact that the technologies associated with algorithms—including “self-learning” algorithms—continue to evolve.

In addition, the ever-increasing awareness of cybersecurity issues has generated concern about algorithms, which are no less susceptible to hacking than other forms of technology.

Also, algorithmic risk can become insidious in that it may not occur in “obvious” places. For example, an algorithm that decides which customers can use automated processing and which must use manual processing can be problematic if the time involved in manual processing turns out to affect valued customers.

Finally, many algorithms are opaque. They may function as “black boxes” that run in the background, and their internal workings are hidden from developers and users. As such, they can be difficult to monitor or audit, and their flaws may not be known or knowable until it is too late—i.e., after their conclusions are reached and acted upon. ➔



Impacts on the enterprise

The type and nature of the algorithmic risks to a particular enterprise will depend upon its nature and size, its industry or field of endeavor, and other factors that make it unique. However, some key examples of risk are as follows:

Area of impact	Illustrative risks	Impact
Finance	Inaccurate financial reporting, incorrect monitoring of data, incorrect metrics used in risk analysis, financial and strategic decision making	Regulatory issues, breaches of loan covenants, operational problems, shareholder discontent or activism, reputation loss
Sales and marketing	Discrimination against certain customers in pricing, product offerings, and ratings	Customer dissatisfaction, revenue loss, regulatory issues, reputation loss
Operations	Product safety and quality, supply chain problems, disruption of normal operations	Business disruption, health and safety impact, revenue loss, regulatory issues, reputation loss
Risk management	Missing detection of significant risks	Business disruption, regulatory issues, shareholder discontent or activism, reputation loss
Information technology	Cyber vulnerabilities, inadequate business continuity planning	Business disruption, regulatory issues, shareholder discontent or activism, reputation loss
Human resources	Discrimination in hiring or performance management	Regulatory or litigation issues, shareholder and customer discontent, reputation loss

Overseeing algorithmic risk

For the board to engage in effective oversight of algorithmic risk, it is advisable to understand the challenges that algorithmic risk poses, and how management is handling these challenges.

Managing algorithmic risk

Algorithmic risk differs from traditional technology risks in some key aspects, which makes it harder to manage, and harder for the board to exercise oversight as to its management.

Algorithmic risk can differ from other types of risk because algorithms can be:

- proprietary;
- complex, unpredictable, and difficult to explain; and
- are not subject to widely accepted cross-industry standards and regulations.

In addition, monitoring algorithmic risk can differ from more “conventional” risk management in that it involves reviewing ongoing data quality in subtle ways. For example, in self-learning algorithms, the data being used by the algorithm should be evaluated as much as the algorithm itself, since the data determine how the algorithm will “learn”.

However, companies can often effectively manage algorithmic risk by developing and adopting new approaches built on “conventional” enterprise risk management approaches. These approaches consist of the following:

- developing an algorithmic risk strategy and governance structure, including policies, risk assessments, training, compliance, and complaint procedures;
- preparing a strong inventory of key algorithms that have been tested and “risk-rated” to enable a focus on algorithms that pose the greatest risk and potential impact;
- developing processes and approaches, aligned with the governance structure, to address the entire algorithm life cycle; and
- establishing processes for assessing the algorithm process—testing data inputs, workings, and outputs, monitoring results, and seeking independent reviews of algorithms. ➔



The board's role

Companies using algorithms need boards that understand the unique challenges associated with overseeing algorithmic risk. While those challenges can be formidable, the board should consider an approach similar to those used in other areas:

- develop a knowledge base;
- work with management to establish an acceptable level of risk associated with the use of algorithms;
- determine specific areas (if any) of risk focus; and
- determine the cadence for periodic reviews of algorithmic risk.

Developing a knowledge base

Developing a knowledge base can begin by asking questions, particularly given that directors may not be aware of or familiar with algorithms and their use within the organization.



Questions for boards to consider:

1. Where and how are algorithms used in our organization?
2. What are the potential impacts if the algorithms go wrong?
3. Are we aware of any algorithms that have functioned improperly? Have we received any complaints about them from any of our constituencies—customers, suppliers, employees, communities, etc.?
4. If so, what kinds of problems have those improper functions created, and how have they been addressed or resolved?
5. What monitoring systems are in place to give us indications of problems with our algorithms?
6. Who oversees our use of algorithms and related risks?
7. What processes do we have in place to monitor and test our algorithms, including data inputs, workings, and outputs?
8. Are our algorithms independently reviewed? By whom? How often?
9. How secure are our key algorithms from cyber-theft or hacking?
10. Has management developed an inventory of tested and “risk-rated” algorithms so that we can rely upon them and focus on other algorithms that may pose greater risks?

Establish a risk tolerance level

Once these questions are answered—along with any other questions arising from the answers—the board should work with management to determine a level of risk tolerance or “risk appetite” associated with algorithms. The discussions with management should result in an understanding of a level of algorithmic risk that will enable the use of algorithms without exposing the company to excessive exposure.

Areas of risk focus

In considering appropriate levels of algorithmic risk, the board and management may determine that certain areas of risk merit specific focus. Algorithmic risk also needs to be considered in a wide range of scenarios, from new product launches to acquisitions.

Risk review cadence

Once a risk tolerance level has been established and determined whether any areas merit particular focus, the board should determine the cadence of algorithmic risk review. A board might determine that it should receive reports on risk—including algorithmic risk—at every meeting or on some other periodic basis, with a deeper dive on certain types of risk at specific meetings.

Also, as with any risk that can be significant, boards need to be satisfied that management has a “crisis” plan in place in the event of an algorithm gone wrong. An algorithmic problem can have a rapid and expansive impact that calls for having a plan in place covering all necessary members of the team and the actions to be taken to address the crisis.

Conclusion

The rapid proliferation of powerful algorithms in every facet of business is in full swing and is likely to grow for years to come as artificial intelligence technologies improve and gain wider adoption. 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. Some benefits could be diminished or completely negated by risks associated with the use of algorithms—risks that are also likely to grow unless organizations develop processes to address algorithmic risk, including an appropriate level of board oversight. ➔



Authors



Nancy Albinson

Managing Director

Deloitte Risk and Financial Advisory
Deloitte & Touche LLP
nalbinson@deloitte.com



Dilip Krishna

Managing Director

Deloitte Risk and Financial Advisory
Deloitte & Touche LLP
dkrishna@deloitte.com



Bob Lamm

Independent Senior Advisor

Center for Board Effectiveness
Deloitte LLP
rlamm@deloitte.com



Yang Chu

Senior Manager

Deloitte Risk and Financial Advisory
Deloitte & Touche LLP
yangchu@deloitte.com

Contact us



Deborah DeHaas

**Vice Chairman, Chief Inclusion Officer,
and National Managing Partner**


Center for Board Effectiveness
Deloitte
ddehaas@deloitte.com



Henry Phillips

**Vice Chairman and
National Managing Partner**

Center for Board Effectiveness
Deloitte & Touche LLP
henryphillips@deloitte.com



Maureen Bujno

Managing Director

Center for Board Effectiveness
Deloitte LLP
mbunjo@deloitte.com



Debbie McCormack

Managing Director

Center for Board Effectiveness
Deloitte LLP
dmccormack@deloitte.com



Krista Parsons

Managing Director

Center for Board Effectiveness
Deloitte & Touche LLP
kparsons@deloitte.com

About this publication

This publication contains general information only and is not a substitute for professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. The authors shall not be responsible for any loss sustained by any person who relies on this communication.

About the Center for Board Effectiveness

The Center for Board Effectiveness helps directors deliver value to the organizations they serve through a portfolio of high quality, innovative experiences throughout their tenure as board members. Whether an individual is aspiring to board participation or a veteran of many board experiences, the Center's programs enable them to contribute effectively and provide focus in the areas of governance and audit, strategy, risk, innovation, compensation and succession.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.