



# Building the Unbiased and Continually Self-Improving Machine

## Authors

**Dinanath Kholkar**

Vice President and Global Head of Analytics & Insights,  
Tata Consultancy Services

Large enterprises that take a Machine First™ approach to digital transformation are using Artificial Intelligence (AI) to automate both manual and knowledge work. While the technology will supplant some workers, we believe the most effective implementations will help people do their work better, whether that means a robot performing repetitive tasks in a factory, freeing up workers for more creative tasks, or an AI system illuminating traffic patterns in a retail store so that salespeople can better serve shoppers.

The benefits of using AI and machine learning are piling up fast. These intelligent systems can help companies react faster to fleeting revenue opportunities, such as identifying customer needs (at the moment when customers experience them) because AI-based machines are tracking those needs at a volume and pace that's beyond the capability of humans. They can pinpoint organizational bottlenecks, such as suboptimal manufacturing processes and delivery routes

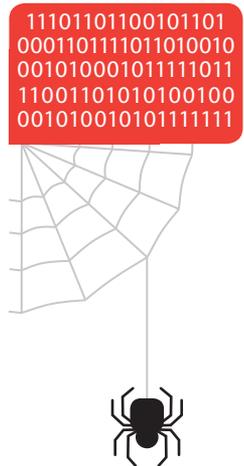
to make them more efficient, reducing waste and costs. They can help organizations make better hiring decisions by logging the qualities of a firm's most productive people and using them to screen potential employees. They can automate processes, from identifying suspicious financial transactions with greater accuracy (and in less time) than humans, to predicting when a machine might fail so it can be fixed before it brings a line to a grinding halt. Without human intervention, they can ensure that corporate purchases comply with an organization's procurement policies.

With this wide array of opportunities, a clear challenge presents itself: to be effective, all these systems rely on data that must be continually refreshed and as complete as possible. Dated or incomplete data yields results that are meaningless, or worse, lead to errors.

The answer to this challenge is for enterprises to build unbiased and self-improving machines that continuously take in more data, from more sources.

## The Problems of Outdated and Incomplete Data

Increasingly, businesses are trusting their AI-based systems, but no matter how smart they are, they can't do the work for which they are designed if the data they are ingesting is flawed. There are two ways that using AI in automated systems can create more problems than benefits: first, relying on data sources that are outdated or not refreshed frequently enough, and second, relying on incomplete data.



Data that is outdated or not refreshed constantly can lead to biased recommendations or incorrectly automated actions. It may prompt a system to consider the wrong things or fail to consider the most recent things. Data needs to be updated to reflect changes in a marketplace, or changes in the makeup of a company's customer base.

## Data that is outdated or not refreshed constantly can lead to biased recommendations or incorrectly automated actions.

The state of Michigan's child welfare system, for example, suffered from poor data quality (including data entry errors) which led to mistakes in tracking the status of neglected children relying on the government for protection and care.<sup>36</sup>

Relying on incomplete data can send a package to an old address or direct a delivery truck to a recently closed road. It can lead to hiring the wrong person for a critical job or eliminate a whole pool of candidates due to unconsciously biased hiring parameters.

In both cases, outdated or incomplete data can lead to failure. If leaders rely on AI systems to make superior decisions, faulty data can derail their organization quickly.

<sup>36</sup> "Report: Flawed state software program could hurt Michigan children," Lansing State Journal, March 13, 2019, accessed at: <https://www.lansingstatejournal.com/story/news/2019/03/13/report-shows-flaws-michigan-child-welfare-computer-system-misacwis/3149656002/>

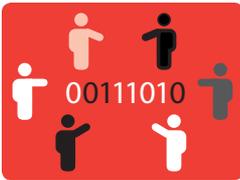
## Building Unbiased Machines

The work of creating unbiased systems differs from traditional software quality practices. Those practices focused on looking for errors in software code and then fixing the bugs. ***With AI, a company's quality measurements must shift from examining lines of software code to examining the quality of the data and the algorithms they use to make meaning of that data.***

Successful implementations must include three quality-control steps:



**1 Ensuring that data sets are complete and well understood.** This requires a company to set up a mature data management capability, review the data inputs for the AI system, and vet the sources of that data for accuracy and completeness.



**2 Employing experts to validate the use of data by AI algorithms.** Skilled people trained in specific domains need to ensure that systems using AI (or machine learning or natural language processing) are producing high-quality outputs. For example, that can mean having a customer experience expert review the results of a system that produces automated responses to customer requests.

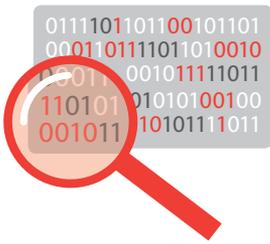


**3 Adding new data.** To improve, a successful AI system requires fresh data on an ongoing basis. Updating the data continuously makes for more accurate outputs. Additionally, more data sources give the AI algorithm more evidence from which to draw insights, improving the quality of its work. Therefore, it is important to take advantage of new technologies—such as visually enabled systems that can interpret text and images, as well as augmented and virtual reality systems that can replicate physical environments—that can supply new data to AI systems.

When combined, these efforts will increase the probability that an automated system will provide accurate results. They will become self-improving, which guards against the risk of biased outcomes based on incorrect or incomplete data.

## How Unbiased Machines Reduce Risk

This article looks at two areas where we have worked to create and sustain unbiased and self-improving machines: systems that automate operations monitoring and systems that automate labor-intensive, regulated processes. In both areas, machines can detect anomalies in patterns of data more quickly and accurately than humans.



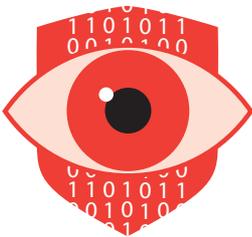
### **Operations monitoring.**

Systems that monitor operations collect data from a range of sources, analyze it to detect irregularities, and automatically determine the next-best action based on those signals, whether the system performs the next action itself or alerts a person. The data the system digests can arrive in the form of text, voice, images, video, transaction streams (as in finance or corporate purchases), unstructured social media data, emails, or online chats. With AI-enabled systems, the more data, and the greater its variety, the more accurate the system, and the less potential for bias. This makes it essential for companies using these systems to seek out new forms of relevant data to add to the analysis.

*Where do enterprises deploy such self-improving systems to monitor their operations?* Retail banks use them to mitigate the risk of fraud by assessing the flow of credit card transactions. Oil and gas companies, as well as transportation firms, have long automated preventive maintenance by monitoring the wear of equipment to

predict when it's time for repair. Manufacturing plant operators can run simulations based on data about plant conditions to optimize operations, such as finding the best, most efficient ways to run a blast furnace. Corporate procurement can automate the auditing of purchasing histories to detect anomalies and ensure compliance with policies.

Each of these settings lend themselves to the continuous addition of new data to analyze patterns of activity and automatically determine the next action. A credit alert is automatically sent to a consumer when the system identifies a never-before-seen purchase. An oil company dispatches a repair crew to replace a pump before it fails and disrupts operations. Procurement automatically generates a monthly list of purchases that cost too much, shedding light on previously dark corporate expenditures.



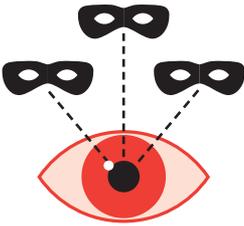
### **Security surveillance.**

Historically, security professionals have relied on manual efforts to track threats. They scan documents, electronic messages, video, and satellite images. A bank might collect employees' electronic and recorded voice communications with customers and other external parties. It's still a prevalent practice: firms assign staff to sift through inputs from various channels, manually, as best they can. But it is difficult, time-consuming work. Vast quantities of 'junk data' that on first blush may signal an important event turn out to be a false positive.

Now, AI and machine-learning applied to monitoring systems makes the work easier at both government agencies and private-sector enterprises.



**Government agencies** can use machines to analyze ever-growing data sources that provide a bigger picture from which to detect unusual patterns. As technology improves—for example, satellite systems that can identify objects with greater precision from greater distances than ever before—the ability to detect anomalies, even minuscule ones, improves. The results can bring greater clarity to risk assessments as the systems and analytics surfaces risks with greater precision in less time.



**Banks** can use AI to identify potentially illegal transactions, mitigating the risk that they are enabling fraud and money laundering, violating government sanctions, or unwittingly facilitating financing for terrorists. The system analyzes transaction data as well as data from conversations between traders and external parties conducted via voice, email, and electronic chat. It automatically identifies correlations that emerge between transactions assessed as risky and the parties involved. Then it can offer analyses to security experts for further evaluation and investigation, saving time by having machines carry out more tasks than people could. These intelligent systems also reduce the number of false-positives, helping experts focus on truly risky transactions without having to employ armies of investigators.

## How Automated Systems Can Improve Processes in Regulated Industries



The same principles—strong data management, the continuous inclusion

of additional data sources, and AI-enabled automation—can be of great worth in the heavily regulated and resource-intensive pharmaceutical industry. Automating essential aspects of pharmacovigilance, a labor-intensive process that requires monitoring the effects of approved drugs on patients and showing all the results to regulators, demonstrates how a commitment to building unbiased and self-improving machines can pay off.

Traditionally, pharmaceutical firms employ teams of 250 to 300 people, including doctors and pharmacists, trained to interpret data in a variety of formats about patients' experiences with medications, especially adverse reactions. The process must also consider health events that may or may not be related to the drug being monitored.

Much of this work is manual: collecting documents about patient reactions and complaints, records of hospital visits and doctor examinations, and scanning social media sites for patient posts alerting friends about their experiences. Then the work of evaluating the data begins, with a team of doctors assigned to assess complaints to determine if the issues are relevant to the company's drug. Finally, the company gathers the findings in reports to submit to regulators.

With AI, pharmaceutical firms can automate this difficult process. First, the system receives data files, both structured and unstructured: emails, medical records from doctors' offices and hospitals, and data from social media feeds. Sometimes, the data is contained in partially completed forms, or fragments of reports. All of these sources serve as 'training data' for the system.

Next, the system performs a triage process to determine the relevance and severity of an identified medical issue: Is it related to the drug? How could

it be related? How serious is it? The system makes its analysis using another set of data: information about the monitored drug, such as its interactions with other drugs and other conditions, white papers, legal cases, and any other relevant information.

The automated pharmacovigilance system is not static. It continues to gather more data to sharpen the insights it generates according to rules developed by experts. These rules enhance accuracy, reduce the chance of error, and provide transparency into the system's workings. That is critical. A misidentified health facility cited as the source of a patient reaction, for example, could lead a regulator to question the accuracy of the company's entire report. Rules programmed into the system can identify a potential error, signal for more scrutiny, estimate the probability that it is an error, and alert trained specialists (such as a physician) to

scrutinize the results. This step is also required to show regulators how the system achieved its results.

Such an automated process can save millions of dollars, shrinking hours of manual pharmacovigilance work to minutes.

This AI-enabled automation approach will have applications in other regulated industries, from banks checking loan activity for compliance with policies to insurance firms examining claim payouts for potential fraud risks. In all these cases, the systems must do the work while also being able to produce transparent reports about how they did it that can be examined by regulators.

What connects all these cases is the requirement that organizations provide an ever-larger set of data sources so their machines can continue to improve their accuracy while avoiding the potential for bias caused by missing or incomplete data. This takes strong data management capabilities. It takes

expertise in developing and managing systems that use AI, machine learning, and natural language processing. It requires a commitment to an ongoing search for new sources of data.

Many companies are adopting automation techniques to improve their operations. The technology is readily available, and opportunities for their use look promising. However, the winners will be those companies that combine a strong data management foundation with a Machine First™ approach to automating processes where AI systems can do the work of many in less time and with better results. These organizations will take advantage of the opportunities to gain benefits that will grow over time with continually self-improving systems.

The winners will be those companies that combine a strong data management foundation with a Machine First approach.