

AN EXECUTIVE'S GUIDE TO REAL-WORLD AI

Lessons from the Front Lines of Business

Sponsored by



SPONSOR PERSPECTIVE



JIM WHITEHURST
PRESIDENT AND CHIEF EXECUTIVE OFFICER
RED HAT

Artificial intelligence will likely transform every aspect of society, business, and industry over the coming decade. AI could impact everything from customers to employees to operations, making it imperative that organizations begin understanding their place in the AI era. If you find the prospect of all this change overwhelming, you are not alone.

The hype around AI and machine learning and the speed at which they are evolving exacerbate the feeling that if you don't get started now, you'll be left behind. CIOs and business leaders can't afford to be paralyzed by these fears.

The good news, according to this new research from Harvard Business Review Analytic Services, is that you don't need to be a large, tech-forward enterprise to win the AI race. You just need to be smart about where to start. CIOs positioning their companies to take advantage of AI now are building a foundation of data, technology, and talent that their competitors are not. They know that, like many overhyped technologies before it, AI is just a tool—one that will work only if you know where and how to wield it.

That's what makes this new report so remarkable. Going beyond hype and hypotheticals, you'll read real examples from more than a dozen leading CIOs, chief data officers, and business leaders who are applying AI to their domains in both big and small ways.

I encourage readers to use this report, sponsored by Red Hat as part of The Enterprisers Project, as a practical guide to start your journey. The stories demonstrate that AI is within reach of all businesses. We plan to continue exploring the unique ways CIOs are transforming their organizations with AI and machine learning at The Enterprisers Project, our online community of IT leaders. We invite you to join us there to continue the conversation.

AN EXECUTIVE'S GUIDE TO REAL-WORLD AI

Lessons from the Front Lines of Business

Artificial intelligence, machine learning, and intelligent automation are commanding huge mind share among business and technology leaders today. There's a lot of hype fueling that interest. Technology and consulting companies are pumping millions of dollars into marketing campaigns that make it seem that if you don't have a talking robot on your road map, you will quickly become irrelevant. Business leaders fear their organizations will be left behind with no chance of recovery as their competitors become radically more efficient and effective. On the flip side, articles abound about the potential negative impact on jobs and the risks and unintended consequences (legal, regulatory, reputational, financial) of turning more decisions and actions over to machines. Hype in tech is nothing new. What's different this time is the degree to which reasonable and knowledgeable people believe that there is, indeed, a real urgency to get going with AI now. Dan Vesset, group vice president, analytics and information management, at market research firm IDC, warns that "if you're not starting to invest, there's the real possibility of being left behind forever."

MIT research fellow and data science expert Tom Davenport and AI pundit Vikram Mahidhar explain the reason executives shouldn't delay in their *Harvard Business Review* article titled "Why Companies That Wait to Adopt AI May Never Catch Up."¹ It's because a number of foundational pieces must be in place to be successful with AI. These include talent, which is in short supply; having the right data infrastructure as well as sufficient quantity and quality of data to train your models; deciding how AI will be governed; and managing change in the organization, among other things.

But where do you start if you're not a large, tech-forward enterprise? In this paper, we share the insights and real-world experiences of over a dozen leading CIOs, chief data officers, and AI experts. We explore how they identify the best opportunities for AI in their organizations; the necessary building blocks of data, technology, and talent; and the various types of risk involved and how to mitigate them.

SECTIONS

- 1 WHERE TO START: DATA HOTSPOTS, REPETITIVE TASKS
- 2 AI REWRITES THE RULES ON REFINING DATA
- 3 AI SKILLS: HOW TO DEVELOP TALENT
- 4 MITIGATE RISKS: HUMAN INVOLVEMENT AND EXPLAINABLE AI
- 5 PREPARING FOR AI'S WIDESPREAD IMPACT ON JOBS
- 6 CONCLUSION

Artificial intelligence is the science and engineering of making intelligent machines. This includes intelligence that is programmed and rules-based, as well as more advanced techniques such as machine learning, deep learning, and neural networks.

5 KEYS TO AI SUCCESS IN YOUR ENTERPRISE

Here's how to move past the hype and lay the real foundations for AI.



1

WHERE TO START: DATA HOTSPOTS, REPETITIVE TASKS

Nothing creates a sense of urgency like reports that competitors are realizing double-digit improvements in key metrics through their use of AI. “If you believe every company is a data company and every company gets leverage out of being a [data] model company, then the first-mover advantage that we’ve seen in the tech space starts to apply much more broadly,” says Bill Mayo, CIO of the Broad Institute, a biomedical and genomic research center in Cambridge, Mass. “If a competitor can leverage data to make some hundredfold improvement for the customer, they can easily pull away from, and ultimately dominate, the market in no time flat.”

Given the profound impact that AI can have on an organization’s customers, employees, operations, and cost structure, identifying the right opportunities and approach is unquestionably a C-level issue.

Gurmeet Singh, EVP and chief digital officer at global convenience chain 7-Eleven, looks for two necessary conditions to decide whether something is ripe for AI: “If you have a lot of data, and you are making lots of daily decisions with that data, then the impact of AI is huge,” he says. Framing it that way led 7-Eleven to pursue opportunities in both customer engagement and inventory management. [SEE “7-ELEVEN” SIDEBAR, PAGE 9](#)

The executive team at Caesars Entertainment looked for opportunities across three main areas, according to CIO Les Ottolenghi. These were “what would be good for cost savings and automation; what would be best in customer engagement; and how could we create a new business service or a new business model?” Ottolenghi says. Caesars began with customer engagement for two reasons; It was where the company could see the biggest measurable payoff, and it was where relevant technologies were already available. [SEE “CAESARS” SIDEBAR, PAGE 5](#)

Bryson Koehler, chief technology officer at Equifax, suggests that companies looking to increase efficiency and effectiveness in their internal operations start with repetitive tasks. “There are always 10, 20, 30 highly manual repetitive tasks that end up with an error rate because we have human judgment in play,” says Koehler, who joined Equifax in mid-2018 from IBM, where

he served as CTO for IBM Watson and IBM Cloud. “If you get the data right, those things are pretty straightforward.” Solving a few problems with basic machine learning—a mature technology that has been around for years—creates momentum for doing more.

Automation of IT operations has so far been the most common application for AI in corporations, according to Deloitte’s “State of AI in the Enterprise” research report.² That’s where CIOs, who often take the lead on AI implementations, have the greatest span of control. By taking friction out of the business and just making things work better, CIOs demonstrate the power of AI in

IN FOCUS: IT AUTOMATION

Adobe’s Self-healing Framework

The problem: Identifying and fixing operational issues and outages across Adobe’s systems were taking too long, causing downtime and frustration in the business. A human was required to determine which alerts to react to.

How AI was used: An AI engine built on a knowledge base looks for abnormalities in IT operations. If something goes wrong, the AI applies a programmatic solution it gets from the knowledge base. Adobe CIO Cynthia Stoddard refers to this as “healing as a service.”

Critical element: In addition to detecting, diagnosing, and remediating problems, the system has a built-in learning loop to avoid creating the same problem again in the future.

Results: The time to solve a problem went from an average of 30 minutes to a little over one minute. In 10 months, this saved about 330 hours. There is a dramatic improvement in user experience with the elimination of downtime.

a way that has an immediate and direct impact on everyone in the business. SEE “ADOBE” SIDEBAR, PAGE 2 Now AI is moving quickly into other business processes, as stories about the positive results of AI investments find their way into C-suites and boardrooms and more technology vendors build AI into existing transactional systems, such as customer relationship management (CRM) and enterprise resource planning (ERP) systems.

Similarly, while AI is primarily used at the task level today, it won't stay there for long, according to IDC's Vesset. IDC looks at applications of AI across “a hierarchy of tasks, activities, processes, and systems.” He says, “AI technology in the form of machine learning and deep learning analysis techniques has become quite good at task-level automation.”

Vesset used the example of auto insurance claims processing to illustrate that point. Some insurance companies now allow customers to take a photo or video of damage from an accident and send it in to be processed as part of their claim instead of the company sending out a human adjuster to assess the damage in person. An AI agent processes these incoming images and assigns a general dollar value to the claim before it goes to an agent. “The whole process is not fully automated, but there are tasks and activities within the process that are,” says Vesset. “We see quite a bit of task-level automation but very little evidence of process-level automation” where the entire claims process, for example, is handled by AI.

CIOs at a roundtable discussion in Boston in the early spring of 2019 said they believed that AI adoption and evolution will happen quickly. As benefits accrue and the technology matures, they said, it will not be unusual to see whole processes managed by AI within 12 to 18 months.


That said, more automation should not be a goal in itself. “There are plenty of use cases where human interaction is not only desirable but also advantageous, regardless of the ability of AI technology to automate,” Vesset says. Each enterprise needs to decide for itself what the desired balance of human and machine interaction should be in a given business process.

Before any of this can happen, however, most organizations have a lot of work to do. As 7-Eleven's Singh puts it, “Is your company ready to meet AI?”

Answering that question starts with assessing whether the organization has the data it needs in sufficient quantity and quality for AI. Even if it has amassed a lot of data over the years, where is it stored, in what format, and in what condition?

Executives must consider whether the organization has the right technology infrastructure and tools, along with the necessary skills, to process that data. “Investments in building a data platform must go hand in hand with data science efforts,” says Naveen Singla, vice president, data science center of excellence at Bayer Crop Science. “If your data are not in order, there's no point hiring a bunch of data scientists to start working with that.”

People now refer to data as the new oil. For Dallas-based 7-Eleven, that's an apt metaphor. “We are generating petabytes worth of data on a daily basis,” says Singh. “We're sitting on a crude oil field, and we need to build the refinery to generate insights and value to become a customer-obsessed, digitally enabled company.”



THE MANY FLAVORS OF AI

Machine learning

A way of achieving AI without complex programming, rules, and decision trees. Instead, with machine learning, data—often in very large amounts—is fed into an algorithm so the algorithm can train itself and learn.

Deep learning

A form of machine learning but with more layers. Deep artificial neural networks are a set of algorithms that have set new records in accuracy for many important problems, such as image recognition, sound recognition, recommender systems, etc.³

Computer vision

A field of artificial intelligence that trains computers to interpret and understand the visual world. Using digital images and deep learning models, machines can accurately identify and classify objects—and then react to what they “see.”⁴

Natural language processing (NLP)

A subset of AI that helps computers understand, interpret, and manipulate human language.

Robotic process automation (RPA)

The automation of basic tasks using business logic and rules. RPA “bots” are used for highly repetitive tasks such as those used to process transactions, manipulate data, respond to queries, and communicate with other systems. Some question whether RPA qualifies as AI. Under the above definition, it does.

Explainable AI

Machine learning techniques that make it possible for human users to understand, appropriately trust, and effectively manage AI. Various organizations, including the Defense Advanced Research Projects Agency, or DARPA, are working on this.⁵

AIOps

A term coined by Gartner to refer to technology platforms that automate and enhance IT operations by (1) using analytics and machine learning to analyze big data collected from various IT operations tools and devices in order to (2) automatically spot and react to issues in real time.⁶

Companies that have been undergoing a digital transformation, moving to a platform or microservices architecture, have an advantage.

2 AI REWRITES THE RULES ON REFINING DATA

Beth Israel Deaconess Medical Center in Boston is unusual for a hospital in that it built its own electronic health record (EHR) system—back in 1992—and it hasn't deleted any of its 11 petabytes of patient identified data since. This includes “every single OR [operating room] time, start and finish, and every procedure, comorbidities, and medications for every patient,” says CIO John Halamka.

Halamka knew Beth Israel Deaconess had enough of the right data to use AI to schedule its operating rooms more efficiently, but first the data needed work. “We took our raw transactional data, moved it to the cloud, and then curated it in several ways,”

Halamka says. “You need to understand the provenance of the data—who entered it, what purpose it was intended for, and how the very nature of the data collection changed over time—as you create these machine learning models.”

Half of the medical record is structured data, and half is unstructured text. Halamka's team used NLP to provide enough structure to make the unstructured data usable by the machine learning model.

The cloud is a critical enabler of AI because of its ability to process large volumes of data on demand. This capability was key when financial services firm Capital One began innovating with AI for customer engagement, according to Margaret Mayer, vice president of software engineering for conversational AI platforms at Capital One.

“Building and training your models typically require a set of GPU [graphics processing unit] resources that you can use on demand,” says Mayer, who sets the technology strategy for Eno, Capital One's banking chatbot. [SEE “CAPITAL ONE” SIDEBAR, PAGE 10](#) “Eno started out as an innovation project, and we didn't need a lot of infrastructure—much like a startup. But then when we realized this was going to be successful, the cloud enabled us to quickly scale, and as we add components, we're able to scale and provision our resources very easily and have all the security that we need.”

Companies that have been undergoing a digital transformation, moving to a platform or microservices architecture, have an advantage. That shift has been taking place at both 7-Eleven and Caesars Entertainment. Over the past three years, Caesars has moved all of its major systems and business functions to “a cloud-based platform architecture that allows us to unlock data,” says Caesars Entertainment's Ottolenghi. “This has enabled us to experiment with AI in an effective manner. Managing the complexity of the data rationalization, the models that you create, and then the learning that is done ... requires all of that to happen first.”

The use of application programming interfaces (APIs) and microservices makes it possible for organizations to leverage all the data at the company's disposal. At Bayer Crop Science,

IN FOCUS: OPERATIONAL EFFICIENCY

Scheduling ORs at Beth Israel Deaconess Medical Center

The problem: Beth Israel Deaconess was running out of OR capacity. Adding more ORs would be expensive both to build and to run.

The solution: They created greater capacity in existing OR utilization.

How AI is used: Machine learning using data from a million patients—including OR times of the past, procedures done, and patients' disease, gender, age, comorbidities, medications, etc.—determines how much OR time is needed for any given patient.

Who was involved: Beth Israel Deaconess's 30-member cross-functional IT steering committee defined the problem and desired outcome. Internal staff cleaned and curated the data; a technology partner developed the model and provided the AI framework.

Results: This freed up 30% of OR capacity, making better use of surgical teams' time and avoiding the cost of expansion.

“we have every type of data storage under the sun,” says Jim Swanson, Bayer Crop Science’s chief information officer and head of digital transformation. “We have relational databases as well as open NoSQL databases, graph databases, and open source technologies like HBase and Cassandra.” APIs and microservices provide access to that data, which “allows it to be democratized and accessible, while simultaneously ensuring data security.”

Companies that want to tap into the ongoing stream of unstructured data that comes from customer behavior, operational systems, internet of things (IoT) sensors, and more are moving beyond traditional ways of storing data and building data lakes and data ingestion services. “This is not, oh, let’s create a data warehouse,” says Equifax’s Koehler. “Operational systems don’t move onto a data warehouse, but operational systems can move onto a data lake. You need to get that right first.”

Taking care of the back end is just half the battle, according to Akshay Kumar, senior vice president and chief data officer at Discover. “It’s not just about creating a data lake and moving data out to the cloud,” he asserts. “You have to actually make it accessible and easy to use.” APIs and microservices come into play here as well. [SEE “DISCOVER” SIDEBAR, PAGE 12](#)

Discover has built a platform that “gives users much more capacity and capability without [their] having to learn thousands of new things,” Kumar says. “We are making the experience of doing analytics much easier for the data scientists, the data engineers, the technologists, the business user—making it easy for them to consume analytics.” Users have access to all the analytics they need through a webpage. “Everything they want is here.”

With the data and data infrastructure in place, CIOs can finally turn to the AI tools themselves. These are becoming more readily available, with open source machine learning frameworks like TensorFlow and Caffe. “The tools that allow you to train models are becoming easier to use,” says Kumar. “Every cloud player out there has their own variant of AI tools. As long as you have the underlying data, over time the problem becomes not a PhD-level problem—it’s probably still a graduate degree-level problem—but there is more ease of use in the space.”

Some companies will choose to rely more heavily than others on outside resources. For instance, Beth Israel Deaconess has collaborative agreements with Amazon, Google, IBM, and MIT “to bring their tools and expertise to our data and domain expertise in health care,” Halamka notes. This arrangement is beneficial to all parties as the cloud vendors seek access to new sources of high-quality data in different domains.

Many small and mid-sized organizations will tap into AI primarily through their software providers, with AI embedded in their transactional systems and processes. For example, some vendors are building AI-based predictive lead scoring into their CRM systems, according to Davenport. “It’s very easy,” he says. “It’s got the data that you need already embedded in it. It’s got the interface that your people are familiar with. And it’s going to make your salespeople more productive.” The same could apply to hiring systems, supply chain management, and more.



The cloud is a critical enabler of AI because of its ability to process large volumes of data on demand.

IN FOCUS: CUSTOMER EXPERIENCE

Caesars Entertainment’s Virtual Concierge

The opportunity: Increase and enhance guest engagement at select Caesars resort properties.

How AI was used: Caesars implemented a commercially available 24-hour virtual concierge chatbot system to connect with guests via SMS text. The system uses AI to make recommendations and book dining, entertainment, and spa experiences; facilitate housekeeping and maintenance requests; and more.

Time to get up and running: The development for the pilot was completed in two weeks; setup, training, and deployment took six weeks. After the success in pilot properties, Caesars integrated with the vendor using the industry-standard integration methodology (REST APIs). This implementation took about eight weeks of development effort. The design of the integration makes it possible for the company to deploy this solution easily to new properties with minimal effort in two to four hours.

The outcome/results: Improves guest experiences by providing more relevant information and services, reduces call volume to front desk, reduces wait times to check out, identifies unhappy guests in real time for in-house guest recovery, and improves net promoter scores.



AI SKILLS: HOW TO DEVELOP TALENT

The skills you need on your AI team will depend on how deeply you plan to get into the AI development work itself. Are you addressing PhD-level data science problems—building AI frameworks, algorithms, and models—or graduate degree-level problems?

“The good part is the cloud providers are offering so many tools for free that AI is getting democratized,” says 7-Eleven’s Singh.

At Discover, “we are looking for people who know enough data science and enough about business that they are able to solve practical business problems for our customers,” Kumar notes.

The blend of domain expertise and data science will become increasingly valuable in all organizations. For instance, as Bayer Crop Science shifts research for its seed breeding programs from the field to computer modeling, “you move from hiring all PhD

plant breeders to hiring data scientists that have relevant domain knowledge,” says CIO Swanson. [SEE “BAYER CROP SCIENCE,” PAGE 11](#)

Indeed, for most businesses, building an AI capability “is not about having 20 PhDs sitting at the table,” says Kumar. “I’ve done that, and I can tell you that doesn’t work.” Instead, in addition to one or two PhDs, the AI team at Discover includes technologists, data engineers, data scientists, and people who understand the business and the company’s customers.

Similarly, the team responsible for Eno at Capital One includes people from the product group, technology, and design, according to Mayer. “So, the team that’s building the platform and the algorithms, that’s designing the answers and determining what features should go into Eno, this is a very diverse-looking team,” she explains. As a result, they’re more likely to ask questions such as “Are we looking too narrowly here? Did we solve for all the problems? Are we creating an error somewhere?”

Even organizations that get most of their AI through their acquired software services or partnerships will need some AI talent on staff—if only to know what questions to ask, what to look for in a provider, and how to test their claims.

To get AI up and running, some large organizations have set up AI centers of excellence (COEs). Both Bayer Crop Science and Capital One describe theirs as a hybrid model. At Bayer Crop Science, the COE serves as a nucleus of expertise in data science and engineering, but “our goal is to distribute this capability into the business as much as possible,” says Swanson. “We look across the enterprise, and where we don’t have the skill set, we embed people so we start to grow the capability in areas like supply chain, commercial, and R&D.”

Capital One’s Center for Machine Learning provides expertise and helps product managers “figure out whether or not the problems they have might be addressed with a machine learning model,” said Mayer. At the same time, data science and machine learning capabilities are distributed throughout the company as well.

Software company Adobe’s Robotic COE is a collaborative effort across IT and the lines of business to explore opportunities for robotic process automation. It started out as a lab environment where people could come together to try out different ideas and compare experiences. “As it evolved into production mode, where we’re ready to put it in and actually use it day to day, that’s when we started to have discussions of center of excellence,” says Adobe CIO Cynthia Stoddard. Incorporating both business process knowledge and technology knowledge, the team started by solving problems in finance and procurement and then taking those learnings out to other areas of the company.

Recognizing the increasing importance of AI to its business, Adobe recently launched a six-month technical AI and machine learning training and certification program for its 5,000 engineers. Engineers must master various technical AI topics in order to complete the course, which is self-paced, with both online courses and in-person sessions. Managers can track learning progress, and senior computer scientists within Adobe act as tutors. The goal is to have all current engineers complete the training by early 2020. An AI training course for non technical roles will be offered to employees later this year.

IN FOCUS: COMPLIANCE


Conquering Export/Import Complexity at Raytheon

The problem being solved: Export regulations for a defense company are complex. With thousands of export/import agreements, each with its own unique configuration of requirements, the number of variables that need to be managed can quickly move beyond a human’s ability to track every detail.

How AI was used: All the characteristics of export/import agreements, along with data from past violations, are fed into a decision tree-based predictive model/algorithm. The output of the model is capable of predicting and flagging noncompliant agreements ahead of time by specific risk scores versus the traditional approach of investigation and analysis after an issue occurs. Raytheon uses 80% of the data to train the model and 20% to validate the model’s statistical strength. The model is run quarterly, and the export/import team uses it to predict where the company is at greatest risk of a violation. This helps ensure that new agreements comply with all the relevant requirements.

Time to get up and running: 12 weeks (two iterative cycles)

The outcome/results: Decrease in overall year-over-year violation trends, new effective audit assessment approach, and ability to generate early alert focus categories.



RECOGNIZING THE INCREASING IMPORTANCE OF AI TO ITS BUSINESS, ADOBE RECENTLY LAUNCHED A SIX-MONTH TECHNICAL AI AND MACHINE LEARNING TRAINING AND CERTIFICATION PROGRAM FOR ITS 5,000 ENGINEERS.



4

MITIGATE RISKS: HUMAN INVOLVEMENT AND EXPLAINABLE AI

The risks associated with AI range from the immediate and known to the distant and unknown. AI is both evolutionary and revolutionary. The risks appear to break down into the following areas:

- **Strategic/financial:** Will a new AI-based product or business model pay off? Should we turn a multimillion-dollar decision over to an AI system?
- **Reputational:** How much automation do we want in our customer processes? What’s the impact if things go wrong?
- **Legal/regulatory:** Does our AI comply, and can we prove how decisions and actions happened?
- **Ethical:** Have we defined our standards, and are we building our AI in a way that will align with them? How are we managing bias?
- **Societal:** When do we reach a tipping point in automation and the resultant impact on jobs? What could or should we do about that?






Business leaders manage these risks through a combination of means. IDC has developed a framework to help people think through the degree to which insights, decisions, and actions are made by humans or AI (see graphic below) and about the implications of that division of labor for opportunities and risk.

On the financial risk front, IDC’s Vessel uses the example of a large industrial company using computer vision and machine learning for predictive maintenance. “They’re looking for cracks in their equipment. [The AI] makes it much faster and more efficient to examine the equipment.” The company doesn’t let the AI systems decide whether or not to stop operations or replace the equipment, “because if they’re wrong, it’s potentially tens of millions of dollars in consequences.”

This is one of the reasons a lot of AI and automation are still happening at the task level (find and report signs of stress), with humans making the final decision and initiating the action. Yet it’s easy to see automation progressing quickly from tasks to activities to processes, moving from the left to the right of IDC’s model as results from the AI process get better and more reliable and are better understood.

GETTING TO ACTION WITH AI

Levels of AI-based automation

	 Human Led	 Human Led, Machine Supported	 Machine Led, Human Supported	 Machine Led, Human Governed	 Machine Controlled
1 Who produces insights?	Human analyzes and produces insights using limited technology	Human analyzes and produces insights using a portfolio of tools	Machine analyzes and produces insights with human review	Machine analyzes and produces insights without human review	Machine analyzes and produces insights
2 Who decides and how?	Human decides based on experience and rules	Human decides based on optimized machine prescriptions	Human decides based on machine prescriptions constrained by all factors	Machine decides within a framework of human governance	Machine decides
3 Who acts based on decisions?	Human acts or executes	Human acts or executes	Human acts or executes with machine oversight	Machine acts or executes with human oversight	Machine acts or executes

SOURCE: IDC, 2019

“People get nervous when they don’t understand what the machine is doing,” says Equifax’s Koehler. “That can be anything from your Nest thermostat making a decision around when to turn on or off, or it could be more complex, like your car in auto-drive mode. People like to understand what’s happening.”

Thus the growing interest in “explainable AI.” While neural networks and deep learning hold a lot of promise, in the past it has been hard to provide logical explanations for decisions that the neural network made.

Explainability is particularly important in regulated industries like financial services, where a company working directly with consumers must be able to show how a decision to approve or deny credit, for example, was reached and ensure that it meets regulatory standards. So Equifax has developed a patented approach to “help provide the reasons, the factors, the weightings, and the insight that go into the output of a credit decision made by the financial institutions we work with,” Koehler says. “It basically becomes an explainable output,” making it possible to use these types of models in a highly regulated environment. Put simply, Koehler explains, “it allows the end user, whether that’s a data scientist or a loan officer, to actually see how those weights were applied and impacted the decision their institution has made.”

Early experiments with AI have revealed serious issues with introduced bias—including negative bias based on race or gender, says Tess Posner, CEO of AI4ALL, a nonprofit organization working to increase diversity and inclusion in AI. This is something that needs to be addressed before turning more decisions over to machines.

“As AI becomes ubiquitous, delegating decision making to AI can affect access to loans, jobs, housing—things that have major implications for people’s lives,” Posner says. There are a number of ways companies can and should address this, beginning with ensuring that the teams working on AI projects are themselves diverse. “Having diversity on the team makes it more likely that issues are addressed,” says Posner. “A more diverse set of people are likely to see gaps and ask questions” that a group made up of the same gender, background, education, race, or age might miss.

Diversity on the team is important but not sufficient. To further mitigate and manage the risk of bias, leading organizations provide cognitive bias training for team members to help employees identify bias and how to speak up when bias appears. Additionally, Adobe trains employees in how to avoid bias in the creation of the algorithms in their AI, machine learning, and deep learning processes. This practice includes infusing a diverse data set when training models on applications like image or video tagging.

In addition to addressing diversity, experts encourage companies to institute peer code reviews on their algorithms, run automated checks on their models, and regularly review or audit findings produced by the AI. These steps will help mitigate the risk of unintended consequences.

“People get nervous when they don’t understand what the machine is doing,” says Equifax’s Koehler.

IN FOCUS: INVENTORY MANAGEMENT

Reducing “Out of Stocks” at 7-Eleven

The challenge: Running out of stock can cost retailers millions of dollars in lost sales. For a global convenience store chain, managing what to order each day at the local level is extremely complex, involving millions of decisions. Each store is different—an item that is a heavy seller in one store may not be in others, so it’s important to estimate product demand by store.

How AI was used: 7-Eleven evaluated a number of machine and deep learning techniques to estimate demand. To better explain to store operators how the model works, the company has operationalized statistical machine learning techniques. It uses historical sales, promotions, out-of-stock conditions, weather, holidays, weekends/weekdays, and store demographics to estimate demand by product and store.

For the initial prototype, 7-Eleven processed about a billion transactions (which will be 100 times that once fully rolled out) to train the models and uses cloud services for big data processing and predictions. To predict demand for items daily requires looking at 30-40 million combinations.

Ordering AI is a feature that is embedded in the ordering screen used by franchisees to replenish inventory. The goal is to automate inventory ordering across all the key items in a store. Store operators have full control over the “demand” ordering. AI will suggest what to order, and they can change it if they have additional information that is specific to that location.

The results: Double-digit improvements in avoiding out-of-stock situations.

5

PREPARING FOR AI'S WIDESPREAD IMPACT ON JOBS



As AI changes the nature of work, it will inevitably affect the skills and competencies needed by employees throughout the business. At a conference on the future of the workforce in April 2019, IBM CEO Ginni Rometty said, “I expect AI to change 100% of jobs within the next five to 10 years.” As an example of what those changes might look like, research firm Gartner predicts that 80% of project management tasks that would typically be handled by a person today will be eliminated by AI by 2030, spanning traditional project management functions such as data collection, tracking, and reporting.

Executives talk about “upskilling” the people whose repetitive work is being automated through AI, whether that’s a database administrator in IT, a customer service rep in a call center, or someone in procurement who has spent time creating contracts and purchase orders. However, given the profound changes underway, it’s remarkable that more companies aren’t investing in the training and change efforts to facilitate this.

“I’ve been very disappointed that HR organizations aren’t helping to think about this,” says MIT’s Davenport. “[They should be asking,] what’s this going to do to jobs for us? How are we going to retrain people? And what’s our strategy for dealing with AI? You just don’t hear much about it at all.”

It may be that with AI most commonly used to automate tasks to make existing processes more efficient and effective, executives see the impact so far as minimal. Salespeople still call on clients when all is said and done, even if the information that’s informing their decisions is better and is shaped by AI.

Adobe is one company that is tackling this workplace reorientation now. “It’s change management,” says Stoddard, “so people understand they are getting a virtual workforce buddy to amplify their capabilities and effectiveness; we’re not eliminating a job. We’ve opened up capacity for them to do other activities. So it’s important to help them move up the scale.” Within IT, instead of doing what Stoddard refers to as toil-manual tasks, “we’ve up-leveled the people to write engineering scripts in order to feed the AI engine. So you can show that it’s kind of a career-advancing skill activity as well.”

IN FOCUS: CUSTOMER EXPERIENCE

Capital One’s Banking Assistant

The opportunity: Customers who are increasingly mobile want to be able to easily ask and get answers to questions about things like their bank balance or when their bill is due by using their mobile or digital channels, including SMS text.

The challenge: Banking terminology is very specific, and Capital One wanted customers to be able to ask their questions using whatever terms make sense to them. But that meant being able to interpret the millions of different ways customers ask questions.

How AI was used: A deep learning model using natural language processing was trained on hundreds of thousands of customer conversations and chat logs to be able to match customer queries to the bank’s framework—e.g., “What’s my credit card bill?” equals “return customer account balance.” Customers text these queries to “Eno,” the bank’s chatbot.

The specific technologies: Models were built using a combination of convolutional neural networks and long short-term memory networks with both unsupervised learning—to build out the language model and vocabulary—and supervised learning (to build the intent model). A large-scale cloud-native distributed platform manages customer conversations. APIs access customer data, such as account balances, and get it back to the customer.

The outcome: Customers can SMS text questions to the bank using their own terminology, and Eno will translate. There are over 2,000 ways customers can ask about their balance alone.

Added value: Eno is used for two-way fraud alerts, asking customers whether they made a particular charge, and it’s able to interpret all the different ways customers answer, including with a thumbs up or down.

We all should be thinking about the long-term implications of where this is heading.

As AI changes the nature of work, it will inevitably affect the skills and competencies needed by employees throughout the business.

6

CONCLUSION

AI is becoming pervasive, as these current examples illustrate. But getting started with AI requires building a foundation of data, technology, and talent that many organizations still lack. Companies at the forefront will have an advantage that others may never be able to overcome.

“At some point, we’re going to hit an inflection point,” says IDC’s Vesset, “and if you’re really better at using AI, you will be in a completely different league from your competitors.”

Participants at the Boston CIO roundtable posited that the transition to a more automated, AI-based economy will happen much more quickly than economic disruptions of the past. “In five years, it’s going to be impactful enough that society will have to adjust the way we think about our economy, our societal practices, and our culture,” said participant Kevin Neifert, CIO at defense company Raytheon. [SEE “RAYTHEON” SIDEBAR, PAGE 6.](#)

Whatever the actual time frame, figuring out how an organization is going to employ AI and what that means for its customers, employees, and shareholders should be at the top of every executive’s short-term priority list. At the same time, we all should be thinking about the long-term implications of where this is heading. The decisions that business and technology leaders make today will determine that future—not just for their own organization but for society as a whole.

IN FOCUS: RESEARCH & DEVELOPMENT

Bayer Crop Science’s Seed Breeding Program

The problem being solved: Until recently, the method used to develop better seeds primarily was for plant breeders and agronomists to work with physically planted seeds, see how they performed through the growing season, study their properties, and manually cross-breed them to produce the best germplasm, which would then be subject to further testing. It could take multiple growing cycles and a number of years to get to a new seed variety with desirable characteristics.

How AI was used: Computer simulations use AI to predict the best phenotype (yield, a good stalk, disease susceptibility, etc.) from genomic data, running multiple billions of scenarios before even a single seed goes into the ground. Machine learning models are used to predict the phenotypes from the genomic data, leveraging more than 15 years of data collected from the plant breeding programs at Bayer (previously Monsanto). The model development and deployment are done on the company’s internally developed data science platform, science-at-scale, with a cloud service providing the infrastructure and computational backend.

Time to get up and running: The model development, deployment, technology, and process changes were parts of a multi-year journey, starting with unlocking data from source systems, connecting them across the different process steps, building models, and validating them with experiments in the field progressing to large-scale deployment across the breeding programs for corn. The Bayer team continuously improves the models by employing technologies to collect more and better data and enhancing the underlying machine learning models.

The outcome/results: This shaved one year off the seed breeding R&D pipeline. Fewer growing fields are needed, and they are able to get better products into the marketplace faster.

Figuring out how an organization is going to employ AI and what that means for its customers, employees, and shareholders should be at the top of every executive's short-term priority list.

IN FOCUS: CREDIT MANAGEMENT

Improving Customer's Financial Health at Discover

The opportunity: Identify when a customer might be getting into financial trouble, and address that before the account goes into delinquency and collections. Detecting patterns up front lets Discover be more proactive to engage customers earlier and steer them into a financial program that better meets their needs, lifestyle, and spending behaviors, according to CIO Glenn Schneider. The company is also using this capability to improve credit decisions for new accounts.

How AI is used: AI identifies patterns based on transaction data, payments, and other events, including engagement or lack of engagement with the company, and external data. Given the nature of the predictions in this highly regulated industry, Discover ensures a degree of explainability so that at any step they can demonstrate how a particular outcome was achieved.

Customers first: Discover is also using AI to understand and respond to customers' engagement preferences, such as text versus phone call or if they want to speak with an agent only under certain circumstances.

The outcome/results: Time to identify and act on an issue improved significantly; on-track payments increased by more than 20%, enabling customers to improve their credit standing. Discover also saw a 40% decrease in outbound calls to customers for late payments, and overall, they improved the customer experience through better communication and customized programs.

ENDNOTES

- 1 Why Companies That Wait to Adopt AI May Never Catch Up, by Vikram Mahidhar and Thomas H. Davenport, *Harvard Business Review* <https://hbr.org/2018/12/why-companies-that-wait-to-adopt-ai-may-never-catch-up>
- 2 State of AI in the Enterprise, Deloitte, October 2018 <https://www2.deloitte.com/insights/us/en/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>
- 3 Artificial Intelligence (AI) vs. Machine Learning vs. Deep Learning <https://skymind.ai/wiki/ai-vs-machine-learning-vs-deep-learning>
- 4 Computer Vision: What it is and why it matters https://www.sas.com/en_us/insights/analytics/computer-vision.html
- 5 Explainable Artificial Intelligence (XAI) <https://www.darpa.mil/program/explainable-artificial-intelligence>
- 6 What Is AIOps? Artificial Intelligence for IT Operations Explained <https://www.bmc.com/blogs/what-is-aiops/>



**Harvard
Business
Review**

ANALYTIC SERVICES

hbr.org/hbr-analytic-services



CONTACT US

hbranalyticsservices@hbr.org

Copyright © 2019 Harvard Business School Publishing.

MC212860419