



Technology and Analytics



Stop Tinkering with AI

It's time to go all in. **by Thomas H. Davenport and Nitin Mittal**

Stop Tinkering with AI

It's time to go all in. **by Thomas H. Davenport and Nitin Mittal**

From the Magazine (January–February 2023) / Reprint R2301J



Carolyn Doucette

If you ask someone to name a company that's putting artificial intelligence at the center of its business, you'll probably hear a predictable list of technology powerhouses: Alphabet (Google), Meta (Facebook), Amazon, Microsoft, Tencent, and Alibaba. But at legacy organizations in other industries many leaders feel that it's beyond the capabilities of their companies to transform themselves using AI. Because this technology is relatively new, however, no company was powered by AI a decade ago, so all those that have been successful had to accomplish the same fundamental tasks: They put people in charge of creating the AI; they rounded up the required data, talent, and

monetary investments; and they moved as aggressively as possible to build capabilities.

Easier said than done? Yes. At many organizations AI initiatives are too small and too tentative; they never get to the only step that can add economic value—deploying a model on a large scale. In a 2019 survey conducted by *MIT Sloan Management Review* and Boston Consulting Group, seven out of 10 companies reported that their AI efforts had had minimal or no impact. The same survey showed that among the 90% of companies that had made some investment in AI, fewer than 40% had achieved business gains over the previous three years. That's not surprising: A pilot program or an experiment can take you only so far.

In our research over the past several years we have identified 30 companies and government agencies (which are not always known for their technological savvy) that have gone all in on AI—and reaped the benefits. Many of the companies compete in industries such as banking, retail, and consumer products. Having studied their journeys, we've identified 10 actions those 30 organizations took to become successful AI adopters.

To get substantial value from AI, your organization must fundamentally rethink the way that humans and machines interact in work environments. You should focus on applications that will change how employees perform and how customers interact with your company. You should consider systematically deploying AI across every key function and operation to support new processes and data-driven decision-making. Likewise, AI should drive new product and service offerings and business models. In other words, the technology should eventually transform every aspect of your business.

Each of the 10 undertakings we list in this article will bring your business closer to transformation—but to fully achieve it, you must avoid piecemeal efforts and attack all 10 tasks. The accompanying examples detail how some organizations succeeded. Your business may choose to handle the tasks differently or to approach them in a different order.

[1]

Know What You Want to Accomplish

Ambitious companies have a specific sense of how they mean to apply AI. They want to be more financially successful, of course, but identifying and developing transformational AI requires a clearer objective. Some businesses begin using the technology to improve process speed, reduce operating costs, or become better marketers. No matter what your reason is for harnessing AI, we recommend identifying one well-defined, overarching objective and making it a guiding principle for your adoption.



Carolyn Doucette's digitally altered photographs investigate the human presence in nature.

When Deloitte's audit and assurance practice began developing Omnia, a proprietary AI platform, in 2014, the guiding principle was to improve service quality globally. Creating a global tool in that field isn't as simple as translating data into multiple languages. Important differences exist in how countries regulate data, including standards for privacy, audit processes, and risk management.

A significant part of auditing a company is gathering financial and operational data in a format that can be easily analyzed. Because data structures differ between companies, extracting relevant data and loading it onto an auditing platform can be labor-intensive. Although Omnia was piloted with a U.S. client, the goal of making it a global tool created several unique challenges at the outset, such as developing a single data model that would work across clients and regions.

Envisioning Omnia as a global tool before it had been created allowed Deloitte's developers to focus on standardizing information from

different companies in different countries—a huge undertaking that would have been even more challenging later in the development process.

[2]

Work with an Ecosystem of Partners

Building Omnia required the audit and assurance practice to monitor technology start-ups around the world to find solutions that fit Deloitte’s needs. Without those partners, Deloitte would have had to develop the technologies in-house, which may have been possible, but at a much higher cost and on a much slower timeline. A company needs strong partnerships to succeed with AI.

Deloitte worked with Kira Systems, a Canada-based start-up with software that extracts contract terms from legal documents. Deloitte’s auditors have historically had to read through many contracts and perform this task manually, but now Kira’s natural-language-processing technology automatically identifies and extracts the key terms. Another partner, Signal AI, built a platform that analyzes publicly available financial data to identify potential risk factors in a client’s business. A recent addition to the Omnia platform is Trustworthy AI, a module developed in partnership with Chatterbox Labs, which evaluates AI models for bias.

[3]

Master Analytics

Most successful AI adopters had significant analytics initiatives underway before they moved headlong into artificial intelligence. Although any form of machine learning may include other technologies

that are not based on analytics, such as autonomous actions, robotics, and the metaverse, it has analytics at its core, which is why mastering analytics is crucial to AI adoption.

But what exactly does “mastering analytics” mean? In this context it requires a commitment to using data and analytics for most decisions, which means changing the way you deal with customers, embedding AI in products and services, and conducting many tasks—even entire business processes—in a more automated and intelligent fashion. And to transform their businesses with AI, companies must increasingly have unique or proprietary data: If all their competitors have the same data, they will all have similar machine-learning models and similar outcomes.

**To get substantial value from AI,
your company must fundamentally
rethink the way humans and machines
interact in work environments.**

Seagate Technology, the world’s largest disk-drive manufacturer, has tremendous amounts of sensor data in its factories and has been using it extensively over the past five years to improve the quality and efficiency of its manufacturing processes. One focus of this effort has been automating the visual inspection of silicon wafers, from which disk-drive heads are made, and the tools that manufacture them. Multiple microscope images are taken from various tool sets throughout wafer fabrication. Using data provided by the images, Seagate’s Minnesota factory created an automated system that allows machines to find and classify wafer defects directly. Other image-classification models detect out-of-focus electron microscopes in the monitoring tools to determine whether defects actually exist. Since these models were first deployed,

in late 2017, their use has grown extensively across the company's wafer factories in the United States and Northern Ireland, saving millions of dollars in inspection labor costs and scrap prevention. Visual inspection accuracy, at 50% several years ago, now exceeds 90%.

Data is the foundation of machine-learning success, and models can't make accurate predictions without large quantities of good data. It's fair to say that the single biggest obstacle for most organizations in scaling up AI systems is acquiring, cleaning, and integrating the right data. It's also important to actively pursue new sources of data for new AI initiatives—something we'll discuss later in this article.

[4]

Create a Modular, Flexible IT Architecture

You'll need a way to easily deploy data, analytics, and automation across your enterprise applications. That requires a technology infrastructure that can communicate and understand data from other IT environments, both inside and outside your company. Software in a traditional data center is typically designed to communicate only with software from the same data center. Integrating it with software from outside that infrastructure can be time-consuming and expensive.

A flexible IT architecture makes it easier to automate complex processes, such as Deloitte's extraction of key terms from legal documents. If you can't develop such an architecture on your own (few small and midsize businesses can), you may have to partner with a company such as Microsoft Azure, Amazon Web Services (AWS), or Google Cloud.

Capital One, which has been recognized for decades as an analytics powerhouse, has used analytics to understand consumer spending

patterns, reduce credit risk, and improve customer service. (Disclosure: One of us, Tom, has been a paid speaker for Capital One.) In 2011 Capital One made a strategic decision to reinvent and modernize its culture, operating processes, and core technology infrastructure. The transformation involved moving to an agile model for delivering software, building a large-scale engineering organization, and hiring thousands of people for digital roles. It also inspired the company to move its data to the cloud.



Carolyn Doucette

Capital One built its cloud architecture in partnership with AWS. But before the move to the cloud, Capital One’s executives had to reimagine the future of banking. The digital channels to which customers were migrating, such as the bank’s website and mobile app, produced substantially more data than in-person interactions did, giving the bank an opportunity to better understand how customers interacted with it. Shifting to the cloud made strategic sense partly because it would drive

down the costs of data storage. In 1960 storing one gigabyte cost \$2 million, according to data from USC's Marshall School of Business. That cost dropped to \$200,000 in the 1980s, \$7.70 by the early 2000s, and—thanks to cloud storage—as low as 2 cents by 2017.

The bank determined that AWS could provide software-driven, massively scalable, instantly available data storage and computing power in the cloud at a much lower cost than storing data on the premises. Innovative new machine-learning tools and platforms were also available on AWS. It no longer made sense for Capital One's IT organization to build and manage infrastructure solutions for all this data. Instead it began to focus on developing software and business capabilities. Today Capital One analyzes an endless stream of data from web and mobile transactions, ATMs, and card transactions in real time to meet customer needs and prevent fraud. By 2020 the bank had closed its last data center and moved all its applications and data to the AWS cloud.

To be sure, many companies have already migrated data and applications to the cloud (or they originated there). Those that haven't will have a harder time becoming aggressive AI adopters.

[5]

Integrate AI into Existing Workflows

Inflexible business processes can be as limiting as inflexible IT architectures are. The companies described in this article took pains to integrate AI in the daily workflows of employees and customers. To do this at your organization, determine which of your workflows are ripe for AI speed and intelligence and begin integrating AI into them as soon as possible. Avoid trying to cram it into workflows that wouldn't benefit from machine speed and scale, such as seldom-used business

processes that neither involve nor generate enormous amounts of data and repetition.

Workflow integration requires an even more specific plan of attack than does task 1, “Know what you want to accomplish.” Say you have determined that you want to improve customer service. But integrating AI in existing customer-service workflows requires acute on-the-ground knowledge of those processes that few C-suite executives have. Line employees, however, have an ideal perspective for determining which processes can benefit from artificial intelligence and how the processes can be specifically improved.

Some branches of the U.S. government identified specific tasks and workflows that were ideal for AI speed and scale. NASA, for example, launched pilot projects in accounts payable and receivable, IT spending, and human resources. (As a result of the HR project, 86% of its HR transactions were completed without human intervention.) The Social Security Administration has used AI and machine learning in its adjudication work to address challenges from heavy caseloads and to ensure accuracy and consistency in decision-making. At the start of the Covid-19 pandemic the Department of Veterans Affairs implemented AI chatbots to field questions, to help determine the severity of confirmed cases, and to find potential locations for patient admission. The Transportation Security Lab at the Department of Homeland Security Science and Technology Directorate is exploring ways to incorporate AI and machine learning in the TSA screening process to improve passenger and bag scanning. The Internal Revenue Service is using AI to test which combinations of formal notices are most likely to induce a taxpayer who owes money to send a check.

[6]

Build Solutions Across the Organization

Once you've internally tested and mastered AI across a specific workflow, you'll want to become more aggressive in deploying it throughout the organization. Rather than designing one algorithmic model for one process, your goal should be to find a unified approach that can be replicated across the company.

Cleveland Clinic has "AI popping up all over the place," according to Chris Donovan, its executive director of enterprise analytics and information management. His group facilitates worker-led efforts to develop and deploy AI while also providing executive-led governance approaches. The effort thus far has been driven by a cross-organizational community of practice anchored in the enterprise analytics, IT, and ethics departments.

The greatest challenge leaders face is creating a culture that emphasizes data-driven decisions and actions and is enthusiastic about AI's potential to transform the business.

Like most organizations that are beginning aggressive AI transformations, the clinic faces a huge challenge involving data and analytics. According to Donovan, hospitals have much less data than organizations in other industries, and it is less likely to be clean and well structured. Cleveland's data, he says, has quality issues, is captured poorly, is entered in different ways, and involves different definitions across the institution. Even a common metric such as blood pressure can be taken while the patient is standing, sitting, or supine—

typically with varying outcomes—and is recorded in a variety of ways. Knowledge of each practice’s data structures is required to interpret the BP data accurately. Rather than leave data preparation to each practice within the clinic for each individual data set, Donovan’s group makes it a part of every AI project and works to provide useful data sets to all AI projects.

Cleveland Clinic also uses AI to assess risk in the population health area, where it has built a predictive model that helps prioritize the use of scarce resources to deliver care to the patients most in need. The predictive risk score is now its primary method for determining who gets a phone call to check in. A diabetic patient who has difficulty managing the disease, for example, would get a high-risk score. The clinic built another model to identify patients who are at risk for a disease but have no history or symptoms of it. It’s used to proactively schedule them for preventive care. CC is also working to identify patients with problematic living or working conditions that affect their health; they may need a social worker as much as a physician, or a bus pass to get to medical appointments.

[7]

Create an AI Governance and Leadership Structure

Putting someone in charge of determining how artificial intelligence is deployed throughout the organization makes transformation easier. The best leaders are aware of what AI can do in general, what it can do for their companies, and what implications it might have for strategies, business models, processes, and people. But the greatest challenge leaders face is creating a culture that emphasizes data-driven decisions and actions and makes employees enthusiastic about AI’s potential to improve the business. In the absence of that kind of culture, even if a few AI advocates are scattered around the organization, they won’t get

the resources they need to build great applications, and they won't be able to hire great people. And if AI applications are built, the business won't make effective use of them.

What kind of leader can foster the right culture? First, it helps to have a CEO or another C-level executive who is familiar with information technology leading the initiative. Although someone with no technical knowledge can lead AI efforts at your company, that person would have to learn a lot, and quickly. Second, it's important that the leader work on multiple fronts. The specific initiatives in which he or she chooses to get involved will vary by organization, but participation by a senior executive is particularly important to signaling interest in the technology, establishing a culture of data-driven decisions, prompting innovation across the business, and motivating employees to adopt new skills. Third, leaders hold the power of the purse. Exploring, developing, and deploying AI is expensive. Leaders must invest—or persuade others to invest—enough to enable all levels of adoption.

Having a single AI leader helps, but ultimately commitment to this work must go deep into the organization. If upper, middle, and even frontline managers are only paying lip service to the idea of transforming with AI, things will move slowly, and the organization will most likely revert to old habits. We've seen some highly committed leaders build AI-focused companies with multiple initiatives. But their successors weren't believers, so the focus on AI lapsed.

[8]

Develop and Staff Centers of Excellence

Most heads of AI and analytics still spend a lot of time evangelizing to other managers about the value and purpose of the technology. Decision-makers from all business units should ensure that AI projects

get sufficient funding and time, and they should also implement AI in their own work. It's important to educate that group on how AI functions, when it's appropriate, and what a major commitment to it involves. For the great majority of companies it's still early days for this upskilling and reskilling work, and not every employee needs to be trained in AI. But some clearly do, and probably the more the better. Each company referenced in this article realized that if it was to be successful, it needed considerable talent and training in AI, data engineering, and data science.

Decision-makers from all business units should ensure that AI projects get sufficient funding and time, and they should also implement AI in their own work.

When Piyush Gupta joined DBS Bank as CEO, in 2009, it was Singapore's lowest-rated bank for customer service. Gupta has invested heavily in AI experimentation—about \$300 million a year over the past few years—and has given business units and functions the flexibility to hire data scientists to see what they can accomplish. The bank's head of HR, who had no technical background, created a small working group to identify and pilot AI applications, including JIM—the Job Intelligence Maestro—a model that predicts personnel attrition and helps the bank recruit the most-qualified employees. DBS used it to hire many of the 1,000 data scientists and data engineers who work at the organization today.

DBS now has twice as many engineers as bankers, Gupta says. They work on emerging technologies such as blockchain and asset-backed tokens as well as on AI projects. And the bank's culture has greatly improved. *Euromoney* named DBS the world's best bank for each of the four years from 2018 to 2021, and its capital positions and credit

ratings are now among the highest in the Asia-Pacific region. In 2019 *Harvard Business Review* named Gupta the 89th best-performing CEO in the world.

[9]

Invest Continually

Choosing to be aggressive with AI is not a decision leaders make lightly. That move will have a major influence on the company for decades and for large enterprises may ultimately involve hundreds of millions or billions of dollars. Every successful AI adopter we studied told us that's the cost of committing to ambitious AI adoption at the enterprise level. At first such resource commitments may be scary for organizations. But after seeing the benefits they received from early projects, the AI-powered companies we investigated found it much easier to spend on AI-oriented data, technologies, and people.

CCC Intelligent Solutions, for example, has spent and expects to continue spending more than \$100 million a year on AI and data. (Disclosure: Tom has been a paid speaker for CCC.) The company was founded in 1980 as Certified Collateral Corporation. It was originally created to provide car valuation information to insurers. If you've had a car accident requiring substantial repair work, you've probably benefited from CCC's data, ecosystem, and AI-based decision-making. Over its 40-plus years CCC has evolved to collect and manage more and more data, to establish more and more relationships with parties in the automobile insurance industry, and to make more and more decisions with data, analytics, and, eventually, AI. For the past 23 years the company has been led by Githesh Ramamurthy, who was previously its chief technology officer. CCC has enjoyed solid growth and is approaching \$700 million in annual revenues.

CCC's machine-learning models are based on more than a trillion dollars' worth of historical claims, billions of historical images, and other data on automobile parts, repair shops, collision injuries, and regulations. It also has gathered more than 50 billion miles' worth of historical data through telematics and sensors in vehicles. It provides data—and, increasingly, decisions—to an extensive ecosystem of some 300 insurers, 26,000 repair facilities, 3,500 parts suppliers, and all major automobile-original-equipment manufacturers. CCC's goal is to link those diverse organizations in a seamless ecosystem to process claims quickly. All those transactions take place in the cloud, where CCC's systems have been based since 2003. They connect 30,000 companies and 500,000 individual users and process \$100 billion worth of commercial transactions annually. As you can imagine, reaching this point has been expensive and time-consuming.

[10]

Always Seek New Sources of Data

Gathering data is typically not a problem for large companies, but AI strategies are driven in large part by whatever data can be assembled. More data is good. More accurate data is great. More accurate, structured data that can be applied to AI models immediately is ideal. Integrating data from client systems was perhaps the most challenging component of Deloitte's AI journey. Capital One always had strong data, but it needed a way to store and make use of it within a flexible IT architecture. CCC began accumulating data with its first business model and was therefore well prepared for a shift to an AI-based model. But CCC's transition from a data-oriented business to an AI-oriented one was solidified when it learned how to use a tremendous trove of data that hadn't existed five years earlier.

When you think of data, don't assume that it's just words and numbers. For CCC, vehicle images represent data that can be applied to several critical processes. CCC had accumulated billions of images over its history, but they were taken by adjusters at the site of vehicle damage or by repair shops. Those photos required professional cameras with special graphics cards to store and send the images.

Around 2012 CCC executives noticed that amateur cameras were getting better at a rapid pace and were being incorporated into smartphones. They envisioned a future in which the owners of damaged vehicles would be able to take their own photos for insurance estimates and send them directly from their phones. The executives expected that with no need for professional photographers and cameras, the process would be quicker and more cost-effective. They engaged several professors at leading universities to explore the capability. Meanwhile, CCC's executives began to read about a new AI approach to image analysis—deep-learning neural networks—that with enough training data could sometimes equal or surpass human analysis.

CCC assembled a pool of talented data scientists who learned how to map photos onto the structure of various vehicles and to annotate or label the photos for training. By mid-2021 the system was ready for deployment, and USAA signed on as one of its first customers. The virtuous circle of more data, better models, more business, and more data is what makes CCC's application of smartphone imagery so powerful. New data will continue to flow in to the company, and it will be used to improve estimate predictions and other functions. That will help CCC clients make better decisions, which will most likely bring CCC more business and more data.

• • •

We believe that companies with the most aggressive AI adoption, the best integration with strategy and operations, and the best implementation will achieve the greatest business value. Knowing what the leading adopters are doing can help others as they attempt to assess technology's potential to transform their business. Your organization can take the 10 actions outlined here to move in the same direction.

We also believe that AI—applied strategically and in large doses—will be critical to the success of almost every business in the future. Data is increasing at a rapid pace, and that's not going to change. AI is a means of making sense of data at scale and of ensuring smart decisions throughout an organization. That's not going to change either. Artificial intelligence is here to stay. Companies that apply it vigorously will dominate their industries over the next several decades.

Editor's note: Thomas H. Davenport and Nitin Mittal are the authors of [All-in on AI: How Smart Companies Win Big with Artificial Intelligence](#) (Harvard Business Review Press, 2023), from which this article is adapted.

A version of this article appeared in the [January–February 2023](#) issue of Harvard Business Review.



Thomas H. Davenport is the President's Distinguished Professor of Information Technology and Management at Babson College, a visiting scholar at the MIT Initiative on the Digital Economy, and a senior adviser to Deloitte's AI practice. He is a coauthor of [All-in on AI: How Smart Companies Win Big with Artificial Intelligence](#) (Harvard Business Review Press, 2023).



Nitin Mittal is a principal at Deloitte Consulting, the leader of its analytics and cognitive offering, and a coleader of Deloitte's AI strategic growth offering. He is a coauthor of [All-in on AI: How Smart Companies Win Big with Artificial Intelligence](#) (Harvard Business Review Press, 2023).