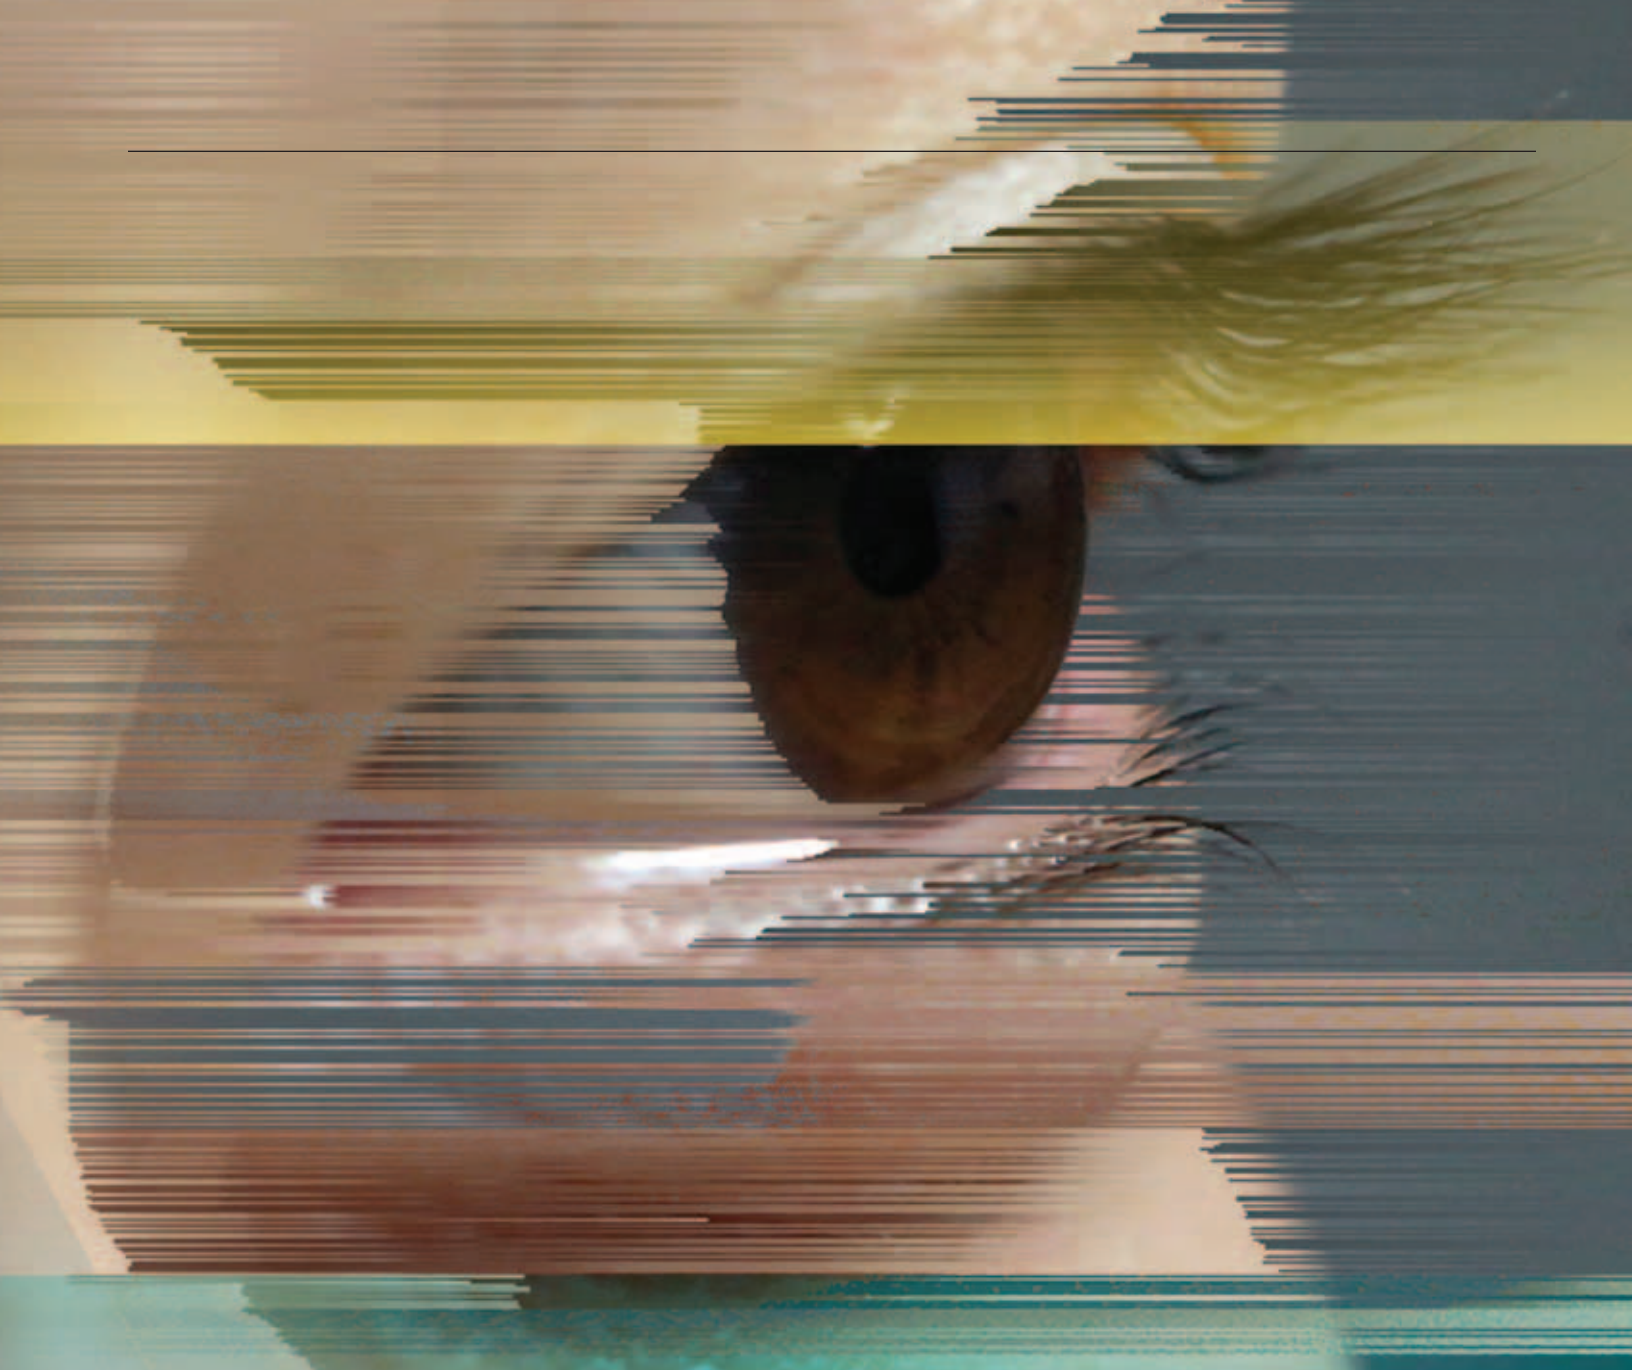


Artificial Intelligence for the Real World

Don't start
with moon shots.

BY THOMAS H. DAVENPORT AND RAJEEV RONANKI

IMAGES BY JAMES WHEATON AND ANDREW NGUYEN

A close-up photograph of a person's face, focusing on the eye area. A medical device, possibly a microscope or a specialized camera, is positioned over the eye, capturing a detailed view of the eye's surface. The background is blurred, showing the person's skin and hair. The image is overlaid with a green and blue gradient at the bottom.

In 2013, the MD Anderson Cancer Center launched a “moon shot” project: diagnose and recommend treatment plans for certain forms of cancer using IBM’s Watson cognitive system. But in 2017, the project was put on hold after costs topped \$62 million—and the system had yet to be used on patients. At the same time, the cancer center’s IT group was experimenting with using cognitive

technologies to do much less ambitious jobs, such as making hotel and restaurant recommendations for patients' families, determining which patients needed help paying bills, and addressing staff IT problems. The results of these projects have been much more promising: The new systems have contributed to increased patient satisfaction, improved financial performance, and a decline in time spent on tedious data entry by the hospital's care managers. Despite the setback on the moon shot, MD Anderson remains committed to using cognitive technology—that is, next-generation artificial intelligence—to enhance cancer treatment, and is currently developing a variety of new projects at its center of competency for cognitive computing.

The contrast between the two approaches is relevant to anyone planning AI initiatives. Our survey of 250 executives who are familiar with their companies' use of cognitive technology shows that three-quarters of them believe that AI will substantially transform their companies within three years. However, our study of 152 projects in almost as many companies also reveals that highly ambitious moon shots are less likely to be successful than “low-hanging fruit” projects that enhance business processes. This shouldn't be surprising—such has been the case with the great majority of new technologies that companies have adopted in the past. But the hype surrounding artificial intelligence has been especially powerful, and some organizations have been seduced by it.

In this article, we'll look at the various categories of AI being employed and provide a framework for how companies should begin to build up their cognitive capabilities in the next several years to achieve their business objectives.

IN BRIEF

THE PROBLEM

Cognitive technologies are increasingly being used to solve business problems, but many of the most ambitious AI projects encounter setbacks or fail.

THE APPROACH

Companies should take an incremental rather than a transformative approach and focus on augmenting rather than replacing human capabilities.

THE PROCESS

To get the most out of AI, firms must understand which technologies perform what types of tasks, create a prioritized portfolio of projects based on business needs, and develop plans to scale up across the company.

THREE TYPES OF AI

It is useful for companies to look at AI through the lens of business capabilities rather than technologies. Broadly speaking, AI can support three important business needs: automating business processes, gaining insight through data analysis, and engaging with customers and employees. (See the exhibit “Cognitive Projects by Type.”)

Process automation. Of the 152 projects we studied, the most common type was the automation of digital and physical tasks—typically back-office administrative and financial activities—using robotic process automation technologies. RPA is more advanced than earlier business-process automation tools, because the “robots” (that is, code on a server) act like a human inputting and consuming information from multiple IT systems. Tasks include:

- transferring data from e-mail and call center systems into systems of record—for example, updating customer files with address changes or service additions;

- replacing lost credit or ATM cards, reaching into multiple systems to update records and handle customer communications;
- reconciling failures to charge for services across billing systems by extracting information from multiple document types; and
- “reading” legal and contractual documents to extract provisions using natural language processing.

RPA is the least expensive and easiest to implement of the cognitive technologies we'll discuss here, and typically brings a quick and high return on investment. (It's also the least “smart” in the sense that these applications aren't programmed to learn and improve, though developers are slowly adding more intelligence and learning capability.) It is particularly well suited to working across multiple back-end systems.


At NASA, cost pressures led the agency to launch four RPA pilots in accounts payable and receivable, IT spending, and human resources—all managed by a shared services center. The four projects worked well—in the HR application, for example, 86% of transactions were completed without human intervention—and are being rolled out across the organization. NASA is now implementing more RPA bots, some with higher levels of intelligence. As Jim Walker, project leader for the shared services organization notes, “So far it's not rocket science.”

One might imagine that robotic process automation would quickly put people out of work. But across the 71 RPA projects we reviewed (47% of the total), replacing administrative employees was neither the primary objective nor a common outcome. Only a few projects led to reductions in head count, and in most cases, the tasks in question had already been shifted to outsourced workers. As technology improves, robotic automation projects are likely to lead to some job losses in the future, particularly in the offshore business-process outsourcing industry. If you can outsource a task, you can probably automate it.

Cognitive insight. The second most common type of project in our study (38% of the total) used algorithms to detect patterns in vast volumes of data and interpret their meaning. Think of it as “analytics on steroids.” These machine-learning applications are being used to:

- predict what a particular customer is likely to buy;
- identify credit fraud in real time and detect insurance claims fraud;
- analyze warranty data to identify safety or quality problems in automobiles and other manufactured products;
- automate personalized targeting of digital ads; and
- provide insurers with more-accurate and detailed actuarial modeling.

Cognitive insights provided by machine learning differ from those available from traditional analytics



COGNITIVE PROJECTS BY TYPE

We studied 152 cognitive technology projects and found that they fell into three categories.

ROBOTICS &
COGNITIVE
AUTOMATION

71

COGNITIVE
INSIGHT

57

COGNITIVE
ENGAGEMENT

24

in three ways: They are usually much more data-intensive and detailed, the models typically are trained on some part of the data set, and the models get better—that is, their ability to use new data to make predictions or put things into categories improves over time.


Versions of machine learning (deep learning, in particular, which attempts to mimic the activity in the human brain in order to recognize patterns) can perform feats such as recognizing images and speech. Machine learning can also make available new data for better analytics. While the activity of data curation has historically been quite labor-intensive, now machine learning can identify probabilistic matches—data that is likely to be associated with the same person or company but that appears in slightly different

formats—across databases. GE has used this technology to integrate supplier data and has saved \$80 million in its first year by eliminating redundancies and negotiating contracts that were previously managed at the business unit level. Similarly, a large bank used this technology to extract data on terms from supplier contracts and match it with invoice numbers, identifying tens of millions of dollars in products and services not supplied. Deloitte’s audit practice is using cognitive insight to extract terms from contracts, which enables an audit to address a much higher proportion of documents, often 100%, without human auditors’ having to painstakingly read through them.

Cognitive insight applications are typically used to improve performance on jobs only machines can do—tasks such as programmatic ad buying that involve

THE BUSINESS BENEFITS OF AI

We surveyed 250 executives who were familiar with their companies' use of cognitive technologies to learn about their goals for AI initiatives. More than half said their primary goal was to make existing products better. Reducing head count was mentioned by only 22%.



SOURCE DELOITTE 2017

such high-speed data crunching and automation that they've long been beyond human ability—so they're not generally a threat to human jobs.

Cognitive engagement. Projects that engage employees and customers using natural language processing chatbots, intelligent agents, and machine learning were the least common type in our study (accounting for 16% of the total). This category includes:

- intelligent agents that offer 24/7 customer service addressing a broad and growing array of issues from password requests to technical support questions—all in the customer's natural language;

- internal sites for answering employee questions on topics including IT, employee benefits, and HR policy;
- product and service recommendation systems for retailers that increase personalization, engagement, and sales—typically including rich language or images; and
- health treatment recommendation systems that help providers create customized care plans that take into account individual patients' health status and previous treatments.

The companies in our study tended to use cognitive engagement technologies more to interact with employees than with customers. That may change as firms become more comfortable turning customer interactions over to machines. Vanguard, for example, is piloting an intelligent agent that helps its customer service staff answer frequently asked questions. The plan is to eventually allow customers to engage with the cognitive agent directly, rather than with the human customer-service agents. SEBank, in Sweden, and the medical technology giant Becton, Dickinson, in the United States, are using the lifelike intelligent-agent avatar Amelia to serve as an internal employee help desk for IT support. SEBank has recently made Amelia available to customers on a limited basis in order to test its performance and customer response.

Companies tend to take a conservative approach to customer-facing cognitive engagement technologies largely because of their immaturity. Facebook, for example, found that its Messenger chatbots couldn't answer 70% of customer requests without human intervention. As a result, Facebook and several other firms are restricting bot-based interfaces to certain topic domains or conversation types.

Our research suggests that cognitive engagement apps are not currently threatening customer service or sales rep jobs. In most of the projects we studied, the goal was not to reduce head count but to handle growing numbers of employee and customer interactions without adding staff. Some organizations were planning to hand over routine communications to machines, while transitioning customer-support personnel to more-complex activities such as handling customer issues that escalate, conducting extended unstructured dialogues, or reaching out to customers before they call in with problems.

As companies become more familiar with cognitive tools, they are experimenting with projects that combine elements from all three categories to reap the benefits of AI. An Italian insurer, for example, developed a "cognitive help desk" within its IT organization. The system engages with employees using deep-learning technology (part of the cognitive insights category) to search frequently asked questions and answers, previously resolved cases, and documentation to come up with solutions to employees' problems. It uses a smart-routing capability (business process automation) to forward the most complex problems to

human representatives, and it uses natural language processing to support user requests in Italian.

Despite their rapidly expanding experience with cognitive tools, however, companies face significant obstacles in development and implementation. On the basis of our research, we've developed a four-step framework for integrating AI technologies that can help companies achieve their objectives, whether the projects are moon shoots or business-process enhancements.

1. UNDERSTANDING THE TECHNOLOGIES

Before embarking on an AI initiative, companies must understand which technologies perform what types of tasks, and the strengths and limitations of each. Rule-based expert systems and robotic process automation, for example, are transparent in how they do their work, but neither is capable of learning and improving. Deep learning, on the other hand, is great at learning from large volumes of labeled data, but it's almost impossible to understand how it creates the models it does. This "black box" issue can be problematic in highly regulated industries such as financial services, in which regulators insist on knowing why decisions are made in a certain way.

We encountered several organizations that wasted time and money pursuing the wrong technology for the job at hand. But if they're armed with a good understanding of the different technologies, companies are better positioned to determine which might best address specific needs, which vendors to work with, and how quickly a system can be implemented. Acquiring this understanding requires ongoing research and education, usually within IT or an innovation group.


In particular, companies will need to leverage the capabilities of key employees, such as data scientists, who have the statistical and big-data skills necessary to learn the nuts and bolts of these technologies. A main success factor is your people's willingness to learn. Some will leap at the opportunity, while others will want to stick with tools they're familiar with. Strive to have a high percentage of the former.

If you don't have data science or analytics capabilities in-house, you'll probably have to build an ecosystem of external service providers in the near term. If you expect to be implementing longer-term AI projects, you will want to recruit expert in-house talent. Either way, having the right capabilities is essential to progress.

Given the scarcity of cognitive technology talent, most organizations should establish a pool of resources—perhaps in a centralized function such as IT or strategy—and make experts available to high-priority projects throughout the organization. As needs and talent proliferate, it may make sense to dedicate groups to particular business functions or units, but even then a central coordinating function can be useful in managing projects and careers.

THE CHALLENGES OF AI

Executives in our survey identified several factors that can stall or derail AI initiatives, ranging from integration issues to scarcity of talent.



SOURCE DELOITTE 2017

2. CREATING A PORTFOLIO OF PROJECTS

The next step in launching an AI program is to systematically evaluate needs and capabilities and then develop a prioritized portfolio of projects. In the companies we studied, this was usually done in workshops or through small consulting engagements. We recommend that companies conduct assessments in three broad areas.

Identifying the opportunities. The first assessment determines which areas of the business could benefit most from cognitive applications. Typically, they are parts of the company where "knowledge"—insight derived from data analysis or a collection of texts—is at a premium but for some reason is not available.

- **Bottlenecks.** In some cases, the lack of cognitive insights is caused by a bottleneck in the flow of information; knowledge exists in the organization, but it is not optimally distributed. That's often the case in health care, for example, where knowledge tends to be siloed within practices, departments, or academic medical centers.

- **Scaling challenges.** In other cases, knowledge exists, but the process for using it takes too long or is expensive to scale. Such is often the case with knowledge developed by financial advisers. That's why many investment and wealth management firms now offer AI-supported "robo-advice" capabilities that provide clients with cost-effective guidance for routine financial issues.

In the pharmaceutical industry, Pfizer is tackling the scaling problem by using IBM's Watson to accelerate the laborious process of drug-discovery research in immuno-oncology, an emerging approach to cancer treatment that uses the body's immune system to help fight cancer. Immuno-oncology drugs can take up to 12 years to bring to market. By combining a sweeping literature review with Pfizer's own data, such as lab reports, Watson is helping researchers to surface relationships and find hidden patterns that should speed the identification of new drug targets, combination therapies for study, and patient selection strategies for this new class of drugs.

- **Inadequate firepower.** Finally, a company may collect more data than its existing human or computer firepower can adequately analyze and apply. For example, a company may have massive amounts of data on consumers' digital behavior but lack insight about what it means or how it can be strategically applied. To address this, companies are using machine learning to support tasks such as programmatic buying of personalized digital ads or, in the case of Cisco Systems and IBM, to create tens of thousands of "propensity models" for determining which customers are likely to buy which products.

Determining the use cases. The second area of assessment evaluates the use cases in which cognitive applications would generate substantial value and contribute to business success. Start by asking key questions such as: How critical to your overall strategy is addressing the targeted problem? How difficult would it be to implement the proposed AI solution—both technically and organizationally? Would the benefits from launching the application be worth the effort? Next, prioritize the use cases according to which offer the most short- and long-term value, and which might ultimately be integrated into a broader platform or suite of cognitive capabilities to create competitive advantage.

Selecting the technology. The third area to assess examines whether the AI tools being considered for each use case are truly up to the task. Chatbots and intelligent agents, for example, may frustrate some companies because most of them can't yet match human problem solving beyond simple scripted cases (though they are improving rapidly). Other technologies, like robotic process automation that can streamline simple processes such as invoicing, may in fact slow down more-complex production systems. And

while deep learning visual recognition systems can recognize images in photos and videos, they require lots of labeled data and may be unable to make sense of a complex visual field.

In time, cognitive technologies will transform how companies do business. Today, however, it's wiser to take incremental steps with the currently available technology while planning for transformational change in the not-too-distant future. You may ultimately want to turn customer interactions over to bots, for example, but for now it's probably more feasible—and sensible—to automate your internal IT help desk as a step toward the ultimate goal.

3. LAUNCHING PILOTS

Because the gap between current and desired AI capabilities is not always obvious, companies should create pilot projects for cognitive applications before rolling them out across the entire enterprise.

Proof-of-concept pilots are particularly suited to initiatives that have high potential business value or allow the organization to test different technologies at the same time. Take special care to avoid "injections" of projects by senior executives who have been influenced by technology vendors. Just because executives and boards of directors may feel pressure to "do something cognitive" doesn't mean you should bypass the rigorous piloting process. Injected projects often fail, which can significantly set back the organization's AI program.

If your firm plans to launch several pilots, consider creating a cognitive center of excellence or similar structure to manage them. This approach helps build the needed technology skills and capabilities within the organization, while also helping to move small pilots into broader applications that will have a greater impact. Pfizer has more than 60 projects across the company that employ some form of cognitive technology; many are pilots, and some are now in production.

At Becton, Dickinson, a "global automation" function within the IT organization oversees a number of cognitive technology pilots that use intelligent digital agents and RPA (some work is done in partnership with the company's Global Shared Services organization). The global automation group uses end-to-end process maps to guide implementation and identify automation opportunities. The group also uses graphical "heat maps" that indicate the organizational activities most amenable to AI interventions. The company has successfully implemented intelligent agents in IT support processes, but as yet is not ready to support large-scale enterprise processes, like order-to-cash. The health insurer Anthem has developed a similar centralized AI function that it calls the Cognitive Capability Office.

Business-process redesign. As cognitive technology projects are developed, think through how

workflows might be redesigned, focusing specifically on the division of labor between humans and the AI. In some cognitive projects, 80% of decisions will be made by machines and 20% will be made by humans; others will have the opposite ratio. Systematic redesign of workflows is necessary to ensure that humans and machines augment each other's strengths and compensate for weaknesses.

The investment firm Vanguard, for example, has a new “Personal Advisor Services” (PAS) offering, which combines automated investment advice with guidance from human advisers. In the new system, cognitive technology is used to perform many of the traditional tasks of investment advising, including constructing a customized portfolio, rebalancing portfolios over time, tax loss harvesting, and tax-efficient investment selection. Vanguard's human advisers serve as “investing coaches,” tasked with answering investor questions, encouraging healthy financial behaviors, and being, in Vanguard's words, “emotional circuit breakers” to keep investors on plan. Advisers are encouraged to learn about behavioral finance to perform these roles effectively. The PAS approach has quickly gathered more than \$80 billion in assets under management, costs are lower than those for purely human-based advising, and customer satisfaction is high. (See the exhibit “One Company's Division of Labor.”)

Vanguard understood the importance of work redesign when implementing PAS, but many companies simply “pave the cow path” by automating existing work processes, particularly when using RPA technology. By automating established workflows, companies can quickly implement projects and achieve ROI—but they forgo the opportunity to take full advantage of AI capabilities and substantively improve the process.

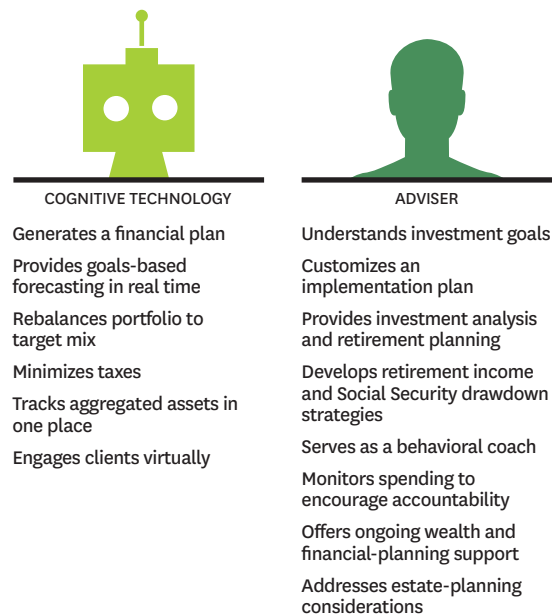
Cognitive work redesign efforts often benefit from applying design-thinking principles: understanding customer or end-user needs, involving employees whose work will be restructured, treating designs as experimental “first drafts,” considering multiple alternatives, and explicitly considering cognitive technology capabilities in the design process. Most cognitive projects are also suited to iterative, agile approaches to development.

4. SCALING UP

Many organizations have successfully launched cognitive pilots, but they haven't had as much success rolling them out organization-wide. To achieve their goals, companies need detailed plans for scaling up, which requires collaboration between technology experts and owners of the business process being automated. Because cognitive technologies typically support individual tasks rather than entire processes, scale-up almost always requires integration with

ONE COMPANY'S DIVISION OF LABOR

Vanguard, the investment services firm, uses cognitive technology to provide customers with investment advice at a lower cost. Its Personal Advisor Services system automates many traditional tasks of investment advising, while human advisers take on higher-value activities. Here's how Vanguard redesigned its work processes to get the most from the new system.



SOURCE VANGUARD GROUP

FURTHER READING

“Big Idea: The Business of Artificial Intelligence”

by Erik Brynjolfsson and Andrew McAfee
HBR.org/ai

“Inside Facebook’s AI Workshop”

by Scott Berinato
HBR.org/ai

“AI Can Be a Troublesome Teammate”

by Kurt Gray
HBR.org/ai

existing systems and processes. Indeed, in our survey, executives reported that such integration was the greatest challenge they faced in AI initiatives.

Companies should begin the scaling-up process by considering whether the required integration is even possible or feasible. If the application depends on special technology that is difficult to source, for example, that will limit scale-up. Make sure your business process owners discuss scaling considerations with the IT organization before or during the pilot phase: An end run around IT is unlikely to be successful, even for relatively simple technologies like RPA.

The health insurer Anthem, for example, is taking on the development of cognitive technologies as part of a major modernization of its existing systems. Rather than bolting new cognitive apps onto legacy technology, Anthem is using a holistic approach that maximizes the value being generated by the cognitive applications, reduces the overall cost of development and integration, and creates a halo effect on legacy systems. The company is also redesigning processes at the same time to, as CIO Tom Miller puts it, “use cognitive to move us to the next level.”

In scaling up, companies may face substantial change-management challenges. At one U.S. apparel retail chain, for example, the pilot project at a small subset of stores used machine learning for online product recommendations, predictions for optimal inventory and rapid replenishment models, and—most difficult of all—merchandising. Buyers, used to ordering product on the basis of their intuition, felt threatened and made comments like “If you’re going to trust this, what do you need me for?” After the pilot, the buyers went as a group to the chief merchandising officer and requested that the program be killed. The executive pointed out that the results were positive and warranted expanding the project. He assured the buyers that, freed of certain merchandising tasks, they could take on more high-value work that humans can still do better than machines, such as understanding younger customers’ desires and determining apparel manufacturers’ future plans. At the same time, he acknowledged that the merchandisers needed to be educated about a new way of working.

If scale-up is to achieve the desired results, firms must also focus on improving productivity. Many, for example, plan to grow their way into productivity—adding customers and transactions without adding staff. Companies that cite head count reduction as the primary justification for the AI investment should ideally plan to realize that goal over time through attrition or from the elimination of outsourcing.

THE FUTURE COGNITIVE COMPANY

Our survey and interviews suggest that managers experienced with cognitive technology are bullish

on its prospects. Although the early successes are relatively modest, we anticipate that these technologies will eventually transform work. We believe that companies that are adopting AI in moderation now—and have aggressive implementation plans for the future—will find themselves as well positioned to reap benefits as those that embraced analytics early on.

Through the application of AI, information-intensive domains such as marketing, health care, financial services, education, and professional services could become simultaneously more valuable and less expensive to society. Business drudgery in every industry and function—overseeing routine transactions, repeatedly answering the same questions, and extracting data from endless documents—could become the province of machines, freeing up human workers to be more productive and creative. Cognitive technologies are also a catalyst for making other data-intensive technologies succeed, including autonomous vehicles, the Internet of Things, and mobile and multi-channel consumer technologies.

The great fear about cognitive technologies is that they will put masses of people out of work. Of course, some job loss is likely as smart machines take over certain tasks traditionally done by humans. However, we believe that most workers have little to fear at this point. Cognitive systems perform tasks, not entire jobs. The human job losses we’ve seen were primarily due to attrition of workers who were not replaced or through automation of outsourced work. Most cognitive tasks currently being performed augment human activity, perform a narrow task within a much broader job, or do work that wasn’t done by humans in the first place, such as big-data analytics.

Most managers with whom we discuss the issue of job loss are committed to an augmentation strategy—that is, integrating human and machine work, rather than replacing humans entirely. In our survey, only 22% of executives indicated that they considered reducing head count as a primary benefit of AI.

We believe that every large company should be exploring cognitive technologies. There will be some bumps in the road, and there is no room for complacency on issues of workforce displacement and the ethics of smart machines. But with the right planning and development, cognitive technology could usher in a golden age of productivity, work satisfaction, and prosperity. 🍷

HBR Reprint R1801H



THOMAS H. DAVENPORT is the President’s Distinguished Professor of Information Technology and Management at Babson College, a research fellow at the MIT Initiative on the Digital Economy, and a senior adviser at Deloitte Analytics. **RAJEEV RONANKI** is a principal at Deloitte Consulting, where he leads the cognitive computing and health care innovation practices. Some of the companies mentioned in this article are Deloitte clients.