Deloitte Review

Issue 21 | July 2017

Complimentary article reprint

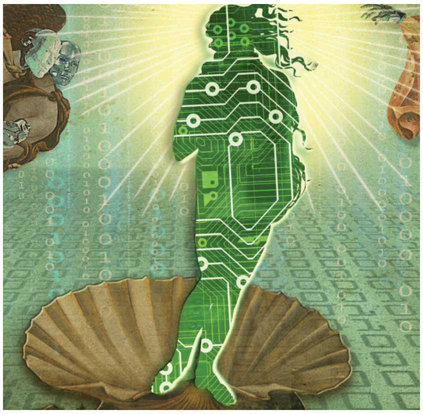

The rise of cognitive work (re)design

Applying cognitive tools to knowledge-based work

By Thomas H. Davenport
Illustration by Anthony Freda

The rise of cognitive work (re)design

Applying cognitive tools to knowledge-based work

By Thomas H. Davenport
Illustration by Anthony Freda

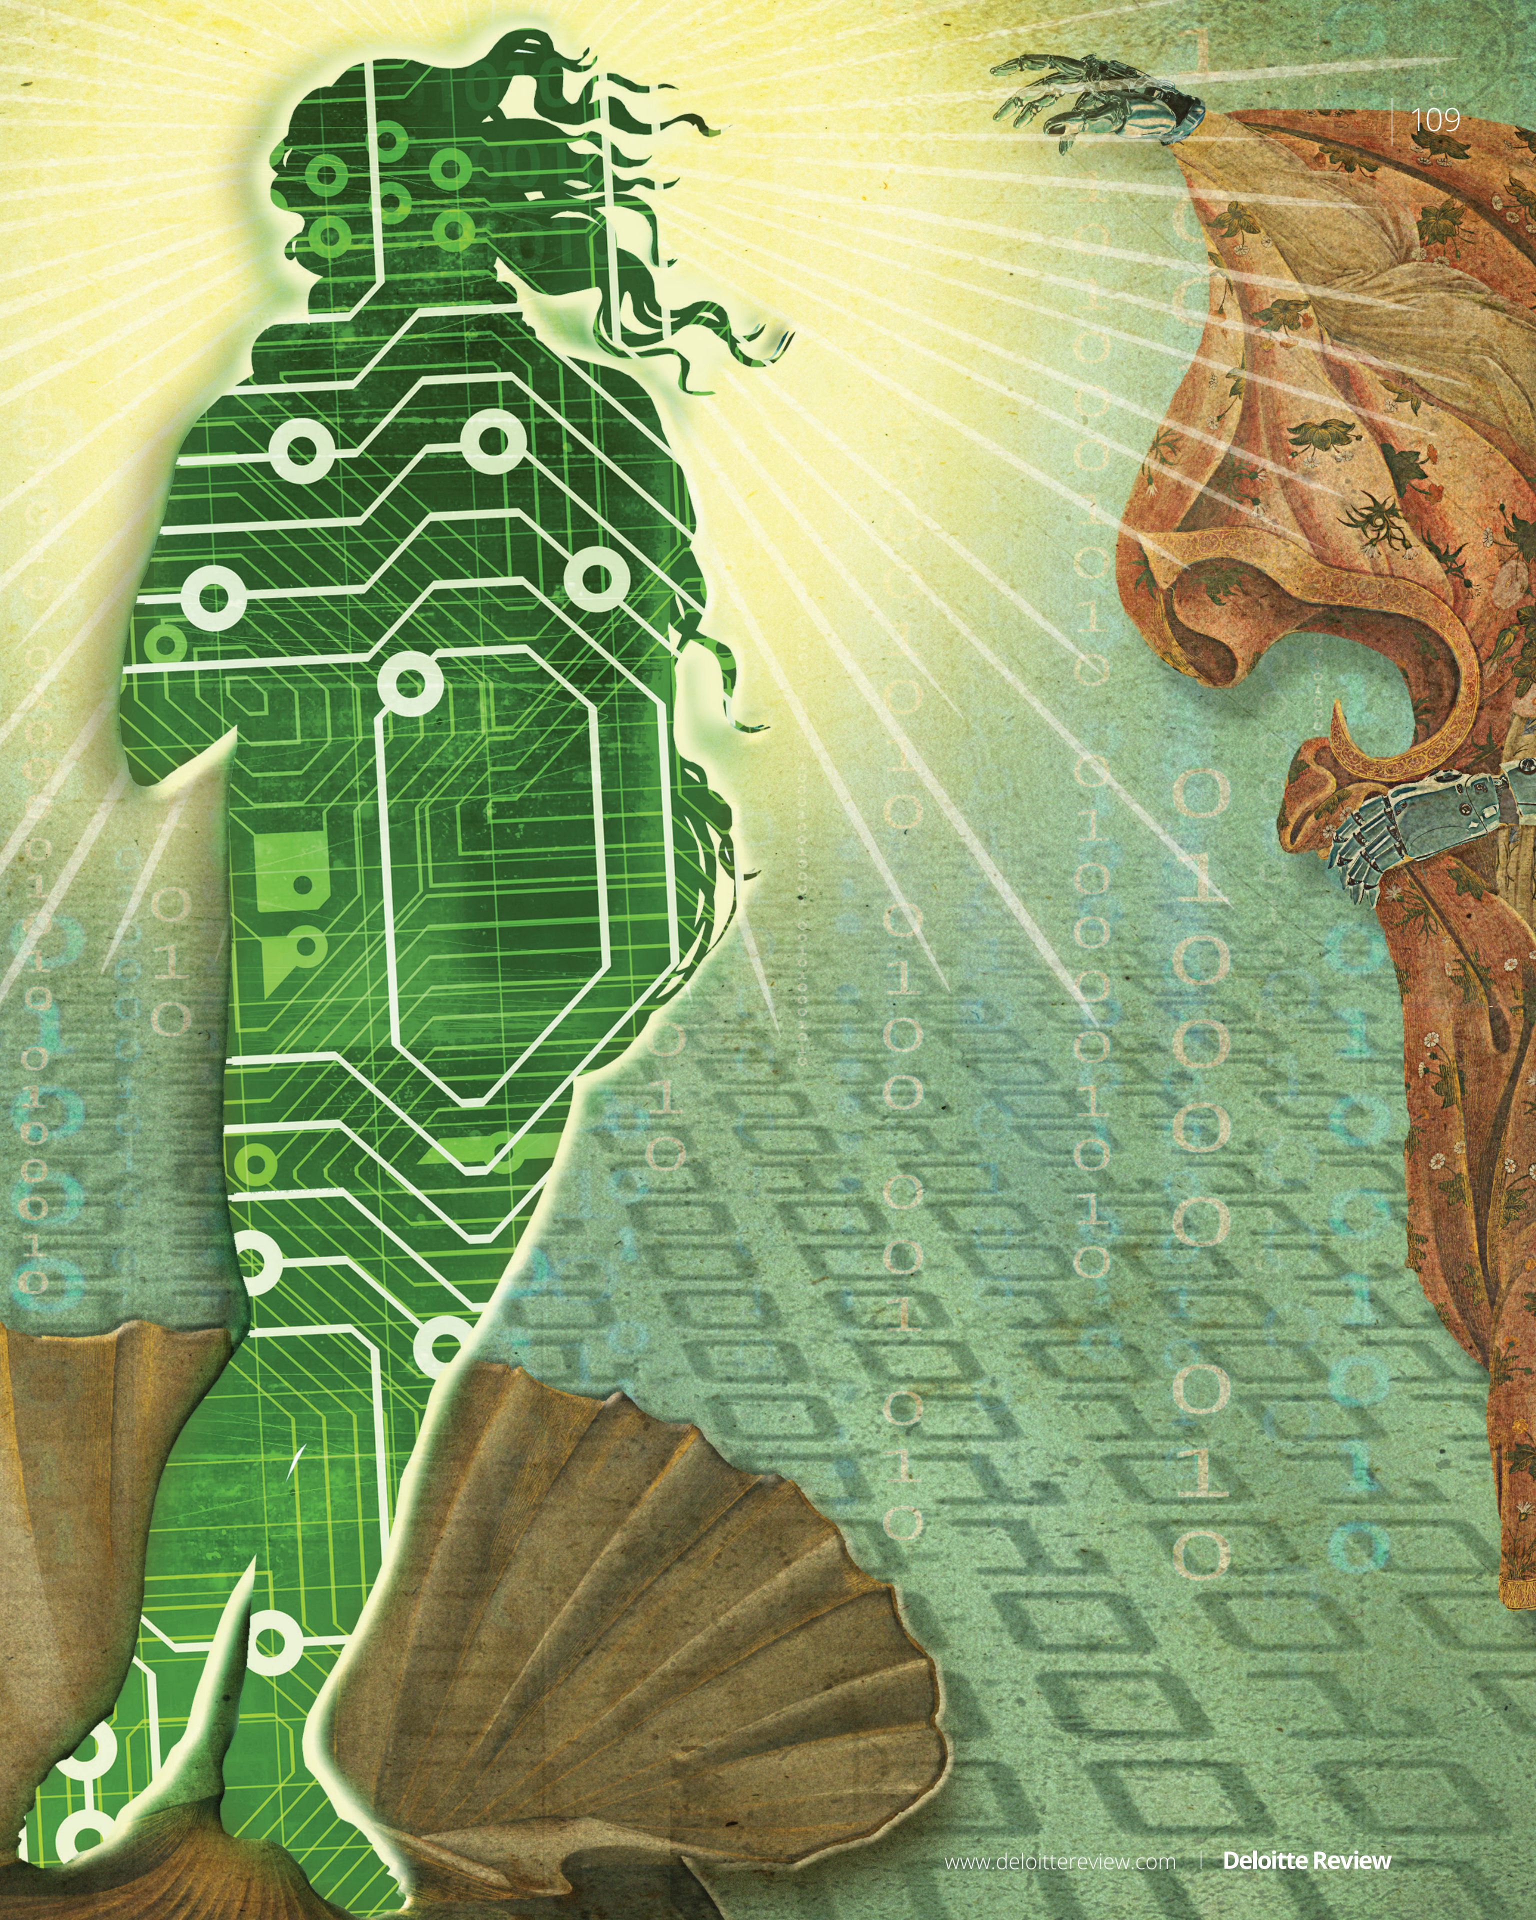

BUSINESS PROCESS CHANGE FOR THE COGNITIVE ERA

NEW automation and cognitive technologies present a unique opportunity to redesign knowledge-based work, but they likely won't do so without a concerted effort to redesign work processes around their capabilities. In order to achieve the productivity and effectiveness benefits that these technologies offer, companies may need to adopt, or readopt, techniques from a variety of systematic approaches to process improvement and change. This time, however, they may want to take a synthetic approach to process change that is consistent with the unique capabilities of cognitive technology.

THE REBIRTH OF REENGINEERING?

IN the early 1990s, one of the most important management trends was "business process reengineering" (BPR).[1] This set of ideas, which encouraged order-of-magnitude improvement in broad business processes, was widely advanced in best-selling books, and led to considerable activity among consulting firms. The primary drivers of the BPR movement were need to substantially improve productivity (in part because of a perceived threat from Japanese competitors) and a powerful new set of information technologies, such as enterprise resource planning (ERP) systems, direct connections between customers and suppliers, and the then-nascent Internet. BPR may have been the only process change approach that specifically addressed information technology as an enabler of innovation and improvement.

Some of the same opportunities and threats appear to be present today. Productivity growth in the United States has slowed for several years,[2] and some prominent economists have proclaimed that information technologies have never fueled the productivity improvements of which they might be capable.[3] As for threats, established firms' primary perceived risks no longer come from large Japanese competitors, but from nimble start-ups in regions like Silicon Valley.

On the technology front, perhaps the most disruptive collection of tools is found in cognitive technologies, the contemporary term for artificial intelligence. This group of technologies, which includes deep and machine learning, natural language processing (NLP) and generation, robotic process automation (RPA), and older tools based on rule and recommendation engines, is currently capturing substantial attention as a source of business and workforce disruption. Perhaps, as in the earlier generation of process reengineering, this generation of technologies can become a driver of work transformation. Also, as in the 1990s, the desired transformation won't take place with technology alone.

It may be time, then, for a renaissance of BPR—this time with a specific focus on cognitive technologies as an enabler of process change, and with a more synthetic approach to process

In addition to focusing on broad, cross-functional processes and radical improvements to them, BPR also differs from other process-focused improvement approaches in that it has a strong focus on information technology.

change methods. The marriage seems a good match. Cognitive technologies need a set of management structures and best implementation practices to yield the benefits of which they are capable. BPR could use some updating to accommodate contemporary technologies, and an injection of new change techniques could make it a more effective methodology.

Most importantly, immediate opportunities for business improvement from cognitive technologies are likely not being realized because complementary process changes aren't being designed and implemented. At one large bank, for example, NLP technology was used to extract payment terms from a large volume of vendor contracts. The terms were then compared to the amounts actually paid by the bank in a large number of invoices (from which the payment amounts had also been extracted with a different set of cognitive tools). The automated analysis identified tens of millions of dollars in contract/invoice mismatches, most of the value of which would accrue to the bank. But the value couldn't be captured until the bank redesigned its processes to review the mismatches and approach vendors to negotiate recovery of inaccurate payments.

Another opportunity for cognitive work redesign may be in the thousands of projects underway today involving RPA.[4] This technology makes it relatively easy to automate structured digital tasks that involve interaction with multiple information systems. But perhaps because of the ease of automating these tasks, very few organizations undertake a systematic effort to redesign the processes and underlying tasks before automating them. While RPA typically leads to substantial gains in efficiency, a process reengineering initiative might reveal substantially greater opportunities for efficiency and effectiveness.

THE POWER OF PROCESS

WHILE other approaches to organizational structure—primarily including business functions such as marketing, finance, and supply chain—may be more familiar, business processes can bring a powerful perspective on monitoring and improving work. Process thinking is at the core of not only business process reengineering, but also Total Quality Management, Six Sigma, and Lean.

Processes are structured sets of activities to accomplish a work-related objective. They can be

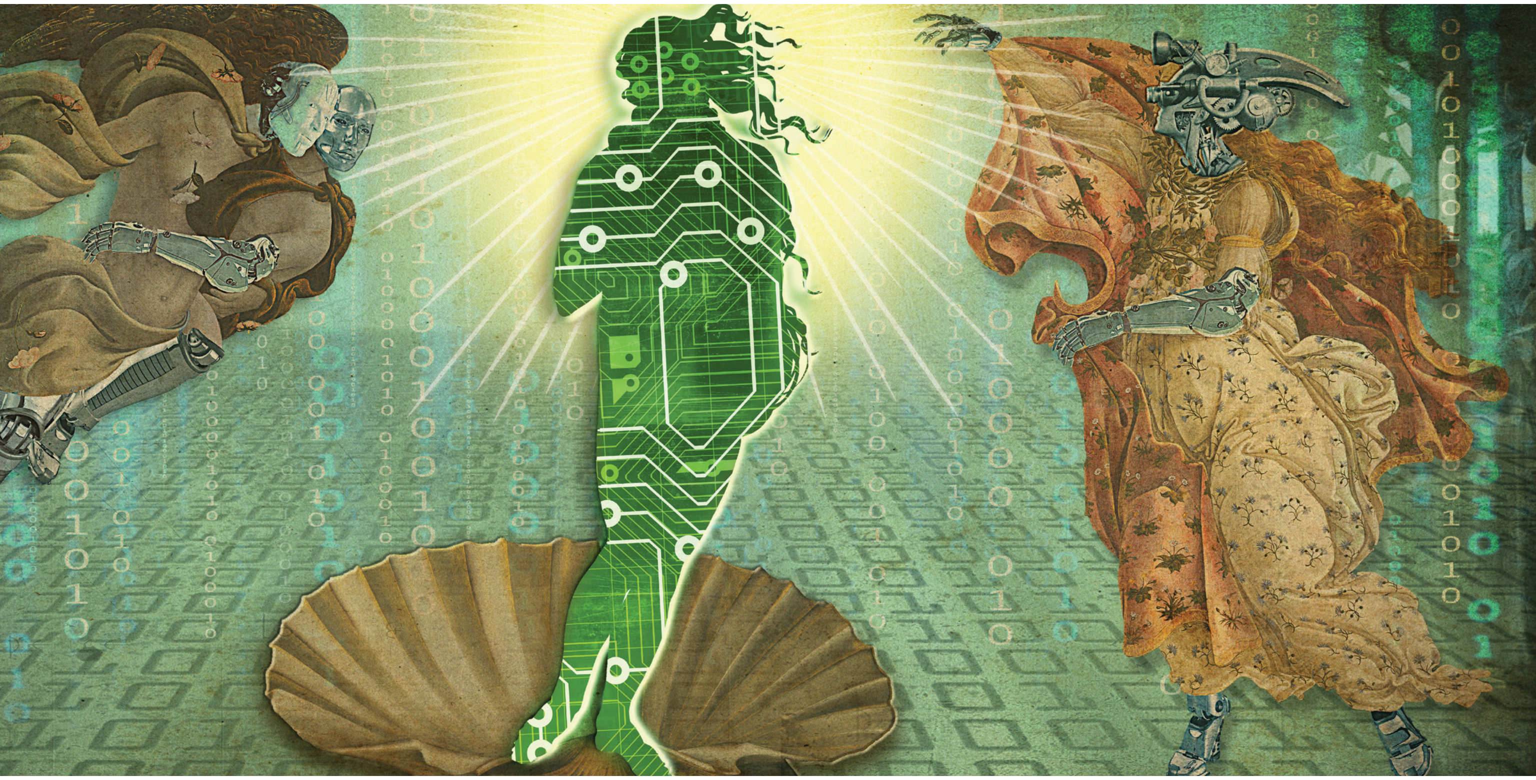

broad, cross-functional processes that encompass many activities ("order to cash" or "procure to pay") or small processes that involve only a few activities ("certify vendor"). BPR was intended to create radical improvements in broad processes, the idea being that radical change required taking on many activities at once and that only broad process improvements would be visible and beneficial to customers. Six Sigma and Lean tend to focus on smaller processes with the idea of making many incremental improvements to them.

In addition to focusing on broad, cross-functional processes and radical improvements to them, BPR also differs from other process-focused improvement approaches in that it has a strong focus on information technology. Again, the rationale is that IT has the ability to enable dramatically new ways of working, which is one way to achieve radical improvements in a process. This was the first—and perhaps still the only—business improvement method to couple the power of technology and structured ways of looking at work.

In practice, the technology most likely to support BPR initiatives in the 1990s was ERP systems, which became popular at the same time. The breadth of these systems and their inherent process orientation made them a good fit for reengineering. However, the difficulty and expense of implementing ERP systems and the challenges of adapting them to fit customized business processes probably contributed to the high failure rate of BPR projects—estimated at between 50 and 70 percent, though never with any rigorous attempts at classification of success and failure.[5]

Cognitive technologies are almost always narrower in their scope of application than ERP. Hence, reengineering methods may need to be modified to some extent to accommodate the fact that cognitive technologies automate or support tasks, not entire processes. Perhaps a synthesis of reengineering methods and Lean or Six Sigma approaches—which can also be relatively narrow in their focus—would be appropriate. Such a blend could couple a broad process innovation vision using cognitive capabilities with a set of shorter-term improvements in specific tasks.

COGNITIVE TECHNOLOGIES AND THEIR IMPACT ON PROCESS TASKS

COGNITIVE technologies have in common the ability to perform tasks with some degree of autonomy that previously only humans could perform. They differ, however, in the types of tasks for which they were intended.[6] Four types of tasks that can be commonly addressed by cognitive technologies include analyzing numbers, analyzing text and images, performing digital tasks, and performing physical tasks.

Analyzing numbers. A key aspect of some cognitive technologies—most forms of statistical machine learning, for example—involves analyzing numbers in structured formats. If any statistical analysis is to be used in a cognitive system, at some point, all forms of data must be converted into structured number formats.

Early numerical analysis was primarily for human decision support, requiring skilled users to direct their use. Now, however, they can run on their own in an automated or semi-automated fashion. Simple machine learning methods can bring different variables into and out of the model to try to create the best fit to the data and the best set of predictions. More complex machine learning models can learn from labeled data and determine strategies in complex business situations, including fraud detection and personalized marketing.

Analyzing words and images. It's always been the province of human beings to read or listen to words and view images, and determine their meaning and significance—a key aspect of human cognition. But now there are a wide variety of tools that are beginning to do just that. Words are increasingly "understood"—counted, classified, interpreted, predicted, and so on—through technologies such as machine learning, natural language processing, neural networks, deep learning, and so forth. Some of the same technologies—deep learning in particular—are being used to analyze and identify images.

Your smartphone can perform many of these tasks. But the analysis of words and images on a large scale comprises a different category of capability. One such application involves translating large volumes of text across languages. Another is to answer questions as a human would. A third is to make sense of language in a way that can either summarize it or generate new passages. A fourth common application, which is mentioned above, is to use linguistic understanding to extract relevant information from documents such as contracts and invoices. This relatively prosaic task is often quite useful in administrative processes.

Image identification and classification is the other key activity in this category. "Machine vision" has existed for many years, but today, many companies are interested in more sensitive and accurate vision tasks: recognizing faces, classifying photos on the Internet, or assessing the collision damage to a car. This sort of automated vision requires more sophisticated tools to match particular patterns of pixels to a recognizable image.[7] Our eyes and brains are great at this, but computers are just beginning to get good at it. Machine learning and "deep learning" neural networks seem to be the most promising technology for this application.

"Deep learning" neural network approaches are particularly well-suited to analyzing data in multiple dimensions (x and y location coordinates; color; intensity; and, in videos, time). The "deep" refers not to the profundity of the learning, but rather to a hierarchy of dimensions in the data. It's this technology that is letting engineers identify photos of cats on the Internet. Perhaps in the near future, smart machines could watch video taken by drones and security cameras and determine whether something bad is happening.

When implemented broadly across an organization, the benefits of RPA can add up quickly.

The most capable systems in this task category are capable of "learning" in that their decisions get better with more data, and they "remember" previously ingested information. IBM's Watson, for example, can be fed more and more documents as they become available over time; that's what makes it well suited for keeping track of cancer research, for example. Other systems in this category can get better at their cognitive task by having more data for training purposes. As more documents that have been translated from Urdu to Hindi become available to Google Translate, for example, it should get better with its machine translations across those languages.

Performing digital tasks. One of the more pragmatic roles for cognitive technology over the past few years has been to automate administrative tasks and decisions. Companies typically have thousands of such tasks and decisions to perform, and it was realized early on that if they could be expressed in a formal logic, they could be automated.

In order to make this possible, a couple of technical capabilities were necessary. One was the expression of the decision logic itself; this came to be known as "business rules." Rules can bring precision, consistency, speed, and computer-driven efficiency to operations. They can be embedded in any sort of computer program, but they are much easier to manage and modify when they are incorporated into a "rules engine," for which there are a variety of vendors.

In addition to business rules, administrative task automation also needed technologies that could move a case or task through the series of steps required to complete it. In the early days of business rules, that technology was "workflow" (also known as "business process management," "case management," or an "orchestration engine"; the most recent version is "complex event processing," or CEP). Regardless of the name, its role was to move a case or project through a series of information-oriented tasks to completion.

Over the past couple of decades, business rules, workflow, and CEP technologies have been used to support a wide variety of administrative tasks, from insurance policy approvals to IT operations to high-speed trading. While these tools can be somewhat inflexible and don't generally learn over time, they have provided a lot of value to organizations. In insurance, for example, they are widely used for policy underwriting and approvals. Their adaptation to a changing business environment has been aided by the relative ease of modifying rules; in many cases, this can be done by a business user. Some rule-based systems are still being implemented for smaller logic-based decisions that require a definite answer versus a probabilistic one.

More recently, companies have begun to employ RPA for digital tasks.[8] Contrary to its name, this technology does not involve actual robots; it makes use of workflow and business rules technology to perform digital tasks. It can automate highly repetitive and transactional tasks, and is usually easily configured and modified by business users. It typically interfaces with multiple information systems as if it were a human user; this is called "presentation layer" integration. RPA technology doesn't learn or improve its performance without human modification, but some vendors are beginning to claim some learning capabilities.

Examples of service industries and processes in which this technology is popular include banking (for example, for back-office customer service tasks, such as replacing a lost ATM card), insurance (process claims and payments), information technology (monitoring system error messages and fixing simple problems), and supply chain management (processing invoices and responding to routine requests from customers and suppliers).

There are substantial benefits from this type of automation, even though it is one of the less exotic forms of cognitive technology. The performance gains can approach 30 or 40 percent improvement in the cost and time to perform a process.[9]

When implemented broadly across an organization, the benefits of RPA can add up quickly. A case study of its application at Telefonica Ó2—the second-largest mobile carrier in the United Kingdom—found that, as of April 2015, the company had automated over 160 process areas involving between 400,000 and 500,000 transactions.[10] Each of the process areas employed a software "robot." The overall return on investment of this technology was between 650 and 800 percent. That's a better payoff than most companies achieved from most other approaches to process improvement.

Performing physical tasks. Physical task automation, of course, is the realm of robots. Though humans love to refer to all automation technologies as robots, the classic usage of the term is "a machine resembling a human being and able to replicate certain human movements and functions automatically."[11] In 2015, more than 250,000 robots were installed in industrial processes across a variety of manufacturing industries.[12]

Robots seem to be evolving in several directions. Some robots are designed from the beginning to provide human support. They include robotic surgery, remotely piloted drone aircraft, and "telecommand" mining machinery. Surgical robots, for example, are driven by human surgeons, but provide them with "superpowers" like better vision, straighter cutting and sutures, and reliable execution of repeated motions. Historically, robots that replaced humans required a high level of programming to do repetitive tasks. They had to be segregated from humans because their movements could be dangerous to us. A new type of robots, however—often called "collaborative robots"—can work alongside humans; they move slowly and stop when they touch anything. These opportunities for human-robot collaboration could be designed into the process, perhaps with some iteration over time as organizations become more familiar with collaborative robots.

Some robots are already somewhat autonomous once programmed, but they are quite limited in their flexibility and their ability to respond to unexpected conditions. More intelligent robots would be able to, for example, look around the proximate area if a part isn't found in the expected location. As robots develop more intelligence, better machine vision, and greater ability to make decisions, they could become a combination of other types of cognitive technologies, but with the added ability to transform the physical environment. IBM Watson software, for example, has been transplanted into several types of robots. FANUC, a Japanese company that is one of the world's largest robot makers, acquired a Japanese deep learning software company, and hopes to make its robots more autonomous using the learning capabilities. As a news article put it,

> Preferred Networks' expertise should allow FANUC's customers to link their robots in new ways. It should also enable the machines to automatically recognize problems and learn to avoid them, or find workarounds in conjunction with other machines.[13]

Similar capabilities are likely to emerge for the "mobile robots" known as autonomous vehicles. Gill Pratt, a Defense Advanced Research Projects Agency (DARPA) program manager who later became head of the Toyota Research Institute, wrote in 2015 that a major change in vehicle intelligence will take place when their intelligence is primarily in the cloud and when vehicles can learn from each other's experiences.[14] These developments suggest that autonomy and awareness are long-term destinations for devices that perform physical tasks, and that the worlds of artificially intelligent software and robots are converging.

Some processes, of course, may involve multiple types of tasks. Tasks may be combined or transformed in applications; some text and images, for example, are converted into numbers for analysis. A customer service application may involve speech recognition, image processing, and machine learning predictions of what is most likely to satisfy the customer. Such combinations are increasingly common with business applications of cognitive technology.

It's important to note that in all of these areas there are still important roles for humans to play. As I've argued (with my co-author Julia Kirby) in a recent book,[15] the most likely future of many processes involves smart humans working alongside smart machines. While there is some possibility of job loss from full automation, most processes can benefit from human oversight, and machines still need some guidance. A redesign effort can determine the tasks within a process for which humans are best suited, and those that can be done primarily by machines.

It's also important to remember that cognitive technologies perform tasks, not jobs or entire processes. It seems that whatever the task, a smart machine can be created to perform it. But a human worker within a business process can typically perform a variety of tasks. Not until we reach the age of "general artificial intelligence" or "the singularity" will this situation change. This suggests that cognitive work redesign efforts within companies should focus on how specific tasks that are supported with cognitive tools fit within broader processes. This is also a better method for thinking about how humans can be redeployed to activities and tasks within processes that make the best use of their capabilities.

COGNITIVE PROCESSES: REDESIGN NEEDED

IN general, cognitive technologies fit best where there is a substantial amount of knowledge needed to make the process effective. Given that cognitive technologies create (from data) and apply knowledge, there are business process contexts for which they are particularly suited. These have historically been processes like product development, health care delivery, and decision making around capital investments, mergers and acquisitions, and strategy. The attributes of likely candidates include the following types of situations:

A knowledge bottleneck—knowledge is unevenly distributed but broadly needed. Knowledge bottlenecks exist where there is substantial knowledge available in one part of a process, but a shortage of it in another. Medical diagnosis and treatment is a classic example. In cancer care, for example, there is substantial knowledge available in academic cancer centers but much less available to the average general medical practitioner—particularly someone in a remote area. A cognitive system can capture the knowledge of the expert (albeit with difficulty, as early results of cognitive cancer treatment systems suggest) and make it available much more broadly. Sofie, the cognitive system for veterinarians from vendor LifeLearn, is a similar solution to a knowledge bottleneck that is particularly severe, given the broad range of animal species for which veterinarians are expected to provide knowledgeable care. Sofie extracts knowledge from the medical literature on animal health and makes it broadly available to veterinarians.[16]

Knowledge is too expensive. In some processes, the requisite knowledge may be available, but is too expensive—perhaps because the knowledge is scarce, or its practitioners are well compensated. The expense could limit the breadth of its application. For example, providers of investment advice have typically charged a fee of 1 percent of invested assets or more. Many less well-off investors don't want to spend that much. And cognitive technologies are now supporting "robo-advisors" that charge less; for a $35,000 portfolio, for example, several robo-advisors charge between 0 and 0.38 percent.[17] College education, widely viewed as too expensive for many students, may also benefit in the future from cognitive technologies such as adaptive learning.

Too much data or knowledge for the human brain to master. There are also processes in which we have little choice about employing cognitive technologies, simply because there is too much data and analysis in the process for the human brain to master.[18] In automated digital advertising (also known as "programmatic buying"), for example, a set of complex calculations (including cost comparisons, auction bidding, and personalization to the user) must take place within approximately 200 milliseconds so that an ad can be served on a publisher's website.[19] No human brain could make such calculations in that time frame. The soaring amount of knowledge about cancer treatment has also been cited as a rationale for cognitive diagnosis and treatment approaches to the disease.

Need for high decision quality and consistency. Typical applications of previous generations of cognitive technology (rule-based systems in particular) included automated underwriting systems in insurance and automated consumer credit issuance systems in banking. These are high-volume processes in which an ongoing high level of performance is critical.[20] Rule-based systems are not as capable as more modern cognitive technologies, but there are contemporary technologies that can support these decision-quality and consistency objectives.

Regulatory requirements. Regulators may require a certain approach to decision making or to descriptions of decisions. While regulators do not require companies to use cognitive technologies, these tools may be helpful in achieving regulatory compliance. For example, some firms are creating anti-money laundering "suspicious activity reports" with automated text generation technologies. Having machines do these relatively structured tasks can free human knowledge workers to perform more value-adding roles.

Companies are just beginning to seize on the work redesign idea for cognitive technologies.

Virtually all business processes require data and information to function, and some data-intensive processes may also be suitable for improvement through data-derived knowledge; that is, analytics. In the traditionally transactional process of order management, for example, customer orders might be treated differently based on their lifetime value predictions. Sales processes could be redesigned around the likelihood of converting a lead to a sale or to assess a customer's propensity to buy. These types of models, as they become more detailed and granular, often require machine learning rather than traditional analytics.

COGNITIVE WORK REDESIGN AT VANGUARD

TO see how these concepts can be put into practice, let's look at how the Vanguard Group approached using cognitive technologies in one of its client-facing activities.

In 2015, the Vanguard Group, an investment management company that manages over $4 trillion in assets,[21] announced a new service for semi-automated investment advice called Personal Advisor Services.[22] The three-year project involved product, technology, and process design, as well as the redesign of the role of investment advisor at the company. This discussion focuses primarily on the work process design and role changes, but it's also important to mention that the advising product that Vanguard chose to offer was relatively straightforward. That made it ideal for a cognitive-based intervention, given the relatively early stage of those technologies. Investment advice is, of course, a knowledge-based offering, so cognitive technologies are appropriate for supporting its delivery.

The Personal Advisor Services product primarily involves basic investment analysis and retirement planning, although it can also address college planning and saving for a home. As with most of Vanguard's business, its analysis largely involves index funds and exchange-traded funds (ETFs) as the investment vehicles it chooses from. These are relatively simple investments, and Vanguard already possessed recommendations for what types of funds were appropriate for different investors' ages and risk preferences.

The goal of the new offering was to have an intelligent system take over many of the tasks of investment advising, including constructing a customized portfolio, rebalancing portfolios over time, tax loss harvesting, and tax-efficient investment selection (figure 1). The system took over some tasks from advisors, including acquiring basic information from customers and presenting financial status information to them. This was sometimes considered tedious for human advisors anyway.

The new process required customers to input more information about themselves, and to

Figure 1. The role of cognitive technologies in delivering Vanguard Personal Advisor Services

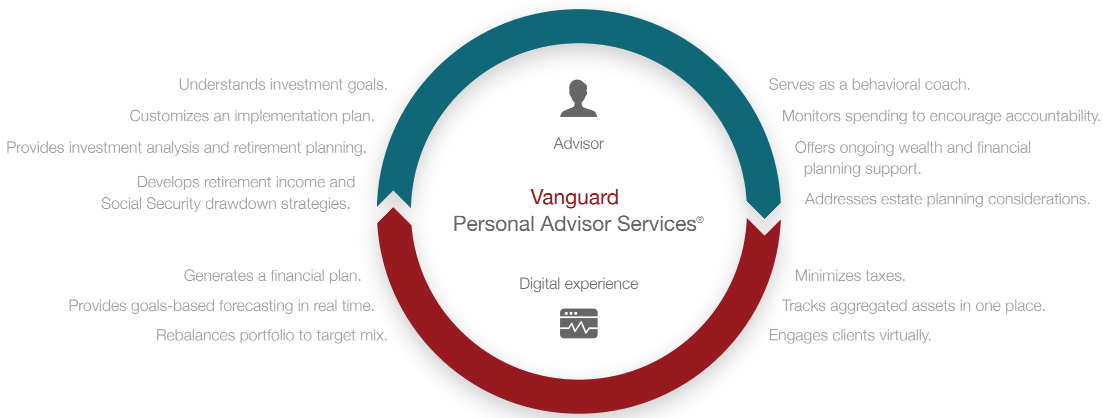

Note: The "digital experience" includes, but is not limited to, an intelligent system.
Source: Vanguard Group, 2017.

furnish information about non-Vanguard assets to their advisor or directly to the system. It made somewhat complex information (for example, about Monte Carlo simulations of how long a portfolio would last in retirement) available to customers, and gave them the ability to override actions that the automated system planned.

For advisors, the new work process required them to undertake some new roles. Several of them were actively involved in the product and process design. The primary description of their new role was to be an "investing coach," able to answer investor questions, encourage healthy financial behaviors, and be, in Vanguard's words, "an emotional circuit breaker" to keep investors on plan.[23] Advisors were encouraged to learn about behavioral finance to perform these roles effectively. To keep costs down and preserve face-to-face contact with investors, advisors were encouraged to employ videoconferencing technology for occasional meetings.

The business goals for the new offering were to further lower the cost of advice and to make customized advice available to investors with lower assets. Both goals were met by the new offering. Vanguard lowered its own fees for ongoing asset management advice to 30 basis points, substantially less than the industry average of around 1 percent. Minimum asset requirements for customized portfolios and advice was reduced from $500,000 to $50,000.[24] And Vanguard has accumulated assets under management rapidly in the program—they are now over $65 billion.[25]

HOW WOULD COGNITIVE WORK REDESIGN WORK?

COMPANIES are just beginning to seize on the work redesign idea for cognitive technologies. Thus far, many have "paved the cow path" by automating the basic existing work process, particularly with RPA technology. Simply automating existing workflows can be a fast way to get to implementation and return on investment, but it can miss an opportunity for substantial improvement in the process.

In essence, work redesign is an instance of "design thinking," which has largely been developed since the first generation of reengineering. Design thinking can involve the design of products, strategies, facilities, and work processes. At least one cognitive technology expert—Manoj Saxena, the chairman of Cognitive Scale, and former general manager of IBM Watson—argues that design thinking is a useful method for harnessing cognitive technology.[26] It seems likely that some components of design thinking could be added to a synthetic approach to cognitive process redesign. Some of the principles of design thinking that can be applied in this context include:

Understand customer (end user) needs. In processes, the customer is the person or unit that receives the output of the process. That may (and often should) be an external customer if the process is defined broadly; it may also be an internal customer. In either case, cognitive process designers should interview and spend time with customers to understand their met and unmet needs, the job that the process is performing for them, and how a cognitive technology solution might make it better. The customers may not understand the capabilities of cognitive technology, so process designers may have to translate customer needs into cognitive capabilities.

Work collaboratively, and include people who perform the process. Reengineering had difficulties in part because it didn't involve people who performed the process to be redesigned. There is a "practice" dimension of work processes that involves workarounds, extraordinary steps to meet customer needs, and departures from official procedure.[27] Involving those who do the work not only helps capture the practice dimension, but can also facilitate buy-in once the process has been designed. This can be particularly important for knowledge workers, who may not be interested in being told how to do their jobs.[28] Other participants in the process might include process design experts, cognitive technology experts, and customers or their representatives.

Design iteratively and experimentally. To test a new process design in action, it's important to create prototypes and pilots to assess different aspects of the design. Scale-up can happen later. If possible, consider breaking the design effort into stages in which different aspects can be piloted or experimented with over time. Try to accomplish something visible each week. In short, this is an "agile" approach to cognitive work redesign. Neither business process reengineering nor large cognitive projects have historically been particularly agile, so this is a departure from the norm.

Keep the cognitive enablers in mind. A key principle of design thinking is to connect technology possibilities with customer needs. In order to do that, the team doing the cognitive process design project should have a high level of familiarity with the capabilities of cognitive technology, key cognitive technology families, common use cases, and so forth. Some examples of these capabilities and use cases include image and speech recognition, creating more granular and personalized marketing models, or automating back-office digital tasks. A cognitive expert on the team could educate other team members on this.

Consider multiple alternatives. One danger in a design exercise is often converging too rapidly on a particular design or technology. It is often more valuable to think of a portfolio of technologies and process innovations that can be tested against the needs of the process and its customers. Since cognitive technology includes a variety of technology types, this should be easy to do.

Start with easy and relatively inexpensive problems. Typical design thinking may not advise starting with simple, inexpensive business problems, but that can be good advice for cognitive work redesign. Cognitive "moonshots" have often proven to be very expensive, at least in the early days of this technology. "Picking low-hanging fruit" appears to be a more successful strategy for cognitive technology for now. For example, in advertising, cognitive technology (machine learning in particular) has been quite successful with digital ads, which are inexpensive. The cost of a bad algorithm is quite low. In television advertising, however, ads can be very expensive—and the industry is probably wise to rely largely on human decision making at this point.

It will probably also be useful to employ at least some of the typical tools used in reengineering and other process-centric methods—such as understanding and measuring the current process and laying out the steps and flows of the "to be" process—in a quick, agile fashion. In addition, it's important to describe the specific "division of labor" between humans and machines at different steps within the process. One call center company, for example, determined that only humans were able to deal with the breadth of call topics from customers calling in for service. So it employs humans for the initial triage of calls, and then connects customers to one of more than a thousand "bots" to handle detailed questions. Another company—a financial asset management and brokerage firm—chose the opposite approach, designing the bot to handle first-line questions and deploying humans to address detailed questions on particular topics. There's no one right answer to this sort of question—only a solution that fits your situation and strategy.

In the future, we expect to see more efforts to use cognitive technologies to redesign key aspects of work. Companies will likely use a blend of participative, iterative methods to incorporate these powerful cognitive tools to capture, apply, and distribute knowledge more effectively within their enterprises. Through these synthetic methods, they can determine the right "division of labor" among smart humans and smart machines. Those who use process-based thinking can be more likely to achieve their business goals, please their customers, and get returns on their investments.

Thomas H. Davenport *is the President's Distinguished Professor of Information Technology and Management at Babson College, the co-founder of the International Institute for Analytics, a Fellow of the MIT Center for Digital Business, and an independent senior advisor to Deloitte Analytics.*

Endnotes

1. Thomas H. Davenport, *Process Innovation: Reengineering Work through Information Technology* (Harvard Business Review Press, 1993); Michael Hammer and James Champy, *Reengineering the Corporation: A Manifesto for Business Revolution* (HarperCollins, 1993).

2. Martin Neil Baily and Nicholas Montalbano, “Why is US productivity growth so slow? Possible explanations and policy responses,” Brookings Institution, September 1, 2016, https://www.brookings.edu/research/why-is-us-productivity-growth-so-slow-possible-explanations-and-policy-responses/.

3. Robert J. Gordon, *The Rise and Fall of American Growth: The US Standard of Living Since the Civil War* (Princeton University Press, 2016).

4. For more background on RPA, see David Schatsky, Craig Muraskin, and Kaushik Iyengar, *Robotic process automation: A path to the cognitive enterprise*, Deloitte University Press, September 14, 2016, https://dupress.deloitte.com/dup-us-en/focus/signals-for-strategists/cognitive-enterprise-robotic-process-automation.html.

5. Michael Hammer and James A. Champy, *Reengineering the Corporation*, revised and updated edition (HarperCollins, 2003), p. 221.

6. An alternative task-type classification for cognitive technologies is found in David Schatsky and Jeff Schwartz, “Redesigning Work in an Era of Cognitive Technologies,” *Deloitte Review* 17, July 2015, https://dupress.deloitte.com/dup-us-en/deloitte-review/issue-17/work-redesign-and-cognitive-technology.html.

7. It is possible to imagine the use of cognitive technologies and “machine vision” to enhance the vision-picking capabilities, enabled by augmented reality, being explored at some warehouses and distribution centers. See Alan Taliaferro, Charles-Andre Guenette, Ankit Agarwal, and Mathilde Pochon, *Industry 4.0 and distribution centers: Transforming distribution operations through innovation*, Deloitte University Press, September 12, 2016, https://dupress.deloitte.com/dup-us-en/focus/industry-4-0/warehousing-distributed-center-operations.html.

8. Leslie P. Willcocks and Mary C. Lacity, *Service Automation: Robots and the Future of Work* (Steve Brookes Publishing, 2016).

9. Mary C. Lacity and Leslie P. Willcocks, “A New Approach to Automating Services,” *MIT Sloan Management Review*, fall 2016, http://sloanreview.mit.edu/article/a-new-approach-to-automating-services/.

10. Mary Lacity, Leslie Willcocks, and Andrew Craig, “Robotic process automation at Telefónica O2,” London School of Economics case study, April 2015, http://www.umsl.edu/~lacitym/TelefonicaOUWP022015FINAL.pdf.

11. Oxford English Dictionary, s.v. “robot,” https://en.oxforddictionaries.com/definition/robot, accessed May 5, 2017.

12. International Federation of Robotics, “Executive Summary World Robotics 2016 Industrial robots,” https://ifr.org/img/uploads/Executive_Summary_WR_Industrial_Robots_20161.pdf, accessed May 5, 2017.

13. Nikkei Asian Review, “Fanuc to build brainier robots with startup,” June 11, 2015, http://asia.nikkei.com/Business/Companies/Fanuc-to-build-brainier-robots-with-startup.

14. Gill A. Pratt, “Is a Cambrian explosion coming for robotics?,” *Journal of Economic Perspectives* 29, no. 3, summer 2015: pp. 51–60.

15. Thomas H. Davenport and Julia Kirby, *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines* (New York: Harper Business, 2016).

16. LifeLearn, “LifeLearn Sofie: Improving the health of patients and practices,” http://www.lifelearn.com/innovations/lifelearn-sofie/, accessed April 18, 2017.

17. Larry Ludwig, “The true costs of robo-advisors: What are the annual fees?,” InvestorJunkie, February 26, 2017, https://investorjunkie.com/42668/true-costs-robo-advisors/.

18. One could envision cognitive technologies being used in conjunction with other technologies, such as augmented reality, that can make large amounts of data more accessible and understandable. See Joe Mariani, Brenna Sniderman, and Cary Harr, “More real than reality: Transforming work through augmented reality,” *Deloitte Review* 21, Deloitte University Press, July 2017.

19. WhatRunsWhere.com, “Media buying 101: What is programmatic buying?,” http://blog.whatrunswhere.com/media-buying-101-programmatic-buying/.

20. Automating high-volume processes can free up humans to focus on other tasks, helping to prevent a "scarcity mind-set" that can compromise decision making and cognitive processing capabilities. See Kelly Monahan, Mark Cotteleer, and Jen Fisher, *Does scarcity make you dumb? A behavioral understanding of how scarcity diminishes our decision making and control*, Deloitte University Press, July 19, 2016, https://dupress.deloitte.com/dup-us-en/focus/behavioral-economics/scarcity-mind-set-improving-decision-making.html.

21. The Vanguard Group, "Fast facts about Vanguard," https://about.vanguard.com/who-we-are/fast-facts/, accessed May 4, 2017.

22. Alessandra Malito, "Vanguard officially launches its robo adviser, drops minimum investment to $50,000," *Investment News*, May 5, 2015, http://www.investmentnews.com/article/20150505/FREE/150509967/vanguard-officially-launches-its-robo-adviser-drops-minimum.

23. Karin Risi (managing director, retail investor group, Vanguard Group), interview with the author, 2015.

24. Malito, "Vanguard officially launches its robo adviser."

25. Personal communication from Vanguard spokesperson, May 5, 2017.

26. Manoj Saxena, "What's still missing from the AI revolution," *Fast Co. Design*, February 8, 2017, https://www.fastcodesign.com/3068005/whats-still-missing-from-the-ai-revolution.

27. John Seely Brown and Paul Duguid, "Organizational learning and communities of practice: Towards a unified view of working, learning, and innovation," *Organization Science* 2, no. 1, 1991: pp. 40–57.

28. Thomas H. Davenport, *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers* (Harvard Business School Press, 2003).