MITSloan
Management Review

Using AI to Enhance Business Operations

How organizations can improve processes and capture value through enterprise cognitive computing.

SUMMER 2019
ISSUE

Monideepa Tarafdar
Cynthia M. Beath
Jeanne W. Ross

Vol. 60, No. 4 • Reprint #60402 • https://mitsmr.com/2F3b9vb

Using AI to Enhance Business Operations

How organizations can improve processes and capture value through enterprise cognitive computing.

BY MONIDEEPA TARAFDAR, CYNTHIA M. BEATH, AND JEANNE W. ROSS

Artificial intelligence invariably conjures up visions of self-driving vehicles, obliging personal assistants, and intelligent robots. But AI's effect on how companies operate is no less transformational than its impact on such products.

Enterprise cognitive computing — the use of AI to enhance business operations — involves embedding algorithms into applications that support organizational processes.[1] ECC applications can automate repetitive, formulaic tasks and, in doing so, deliver orders-of-magnitude improvements in the speed of information analysis and in the reliability and accuracy of outputs. For example, ECC call center applications can answer customer calls within 5 seconds on a 24-7-365 basis, accurately address their issues on the first call 90% of the time, and transfer complex issues to employees, with less than half of the customers knowing that they are interacting with a machine.[2] The power of ECC applications stems from their ability to reduce search time and process more data to inform decisions. That's how they enhance productivity and free employees to perform higher-level work — specifically, work that requires human adaptability and creativity. Ultimately, ECC applications can enhance operational excellence, customer satisfaction, and employee experience.[3]

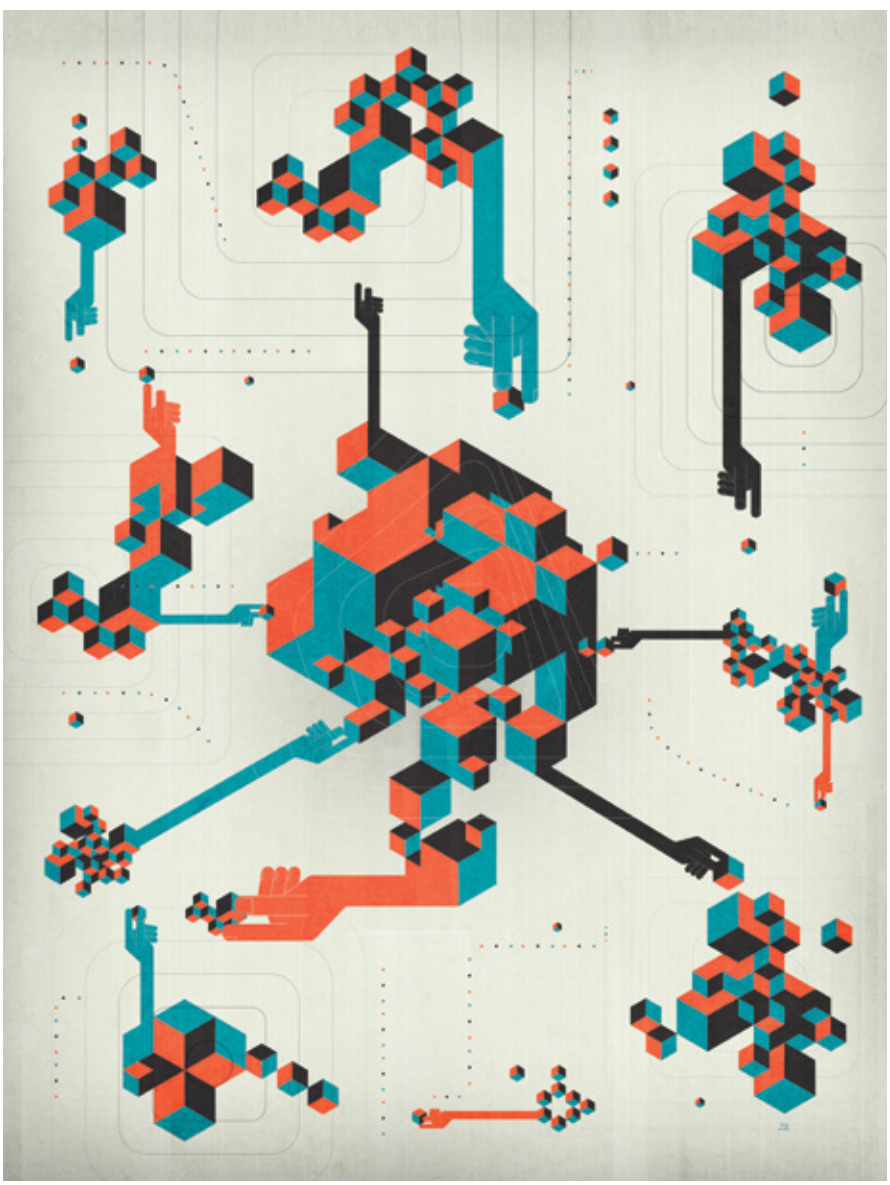

JOHN HERSEY/THEISPOT.COM

THE LEADING QUESTION
How can companies develop their ability to use AI to transform business operations?

- **CEOs recognize the potential of AI to improve operations and capture value but are struggling to realize its promise.**
- **Business domain proficiency — a deep understanding of the tasks, workflows, and logic of existing processes — establishes the essential link between data science and business value.**
- **ECC applications must be managed throughout their life cycles because ever-changing conditions ensure that AI algorithms become a less accurate reflection of reality over time.**

ECC applications come in many flavors. For instance, in addition to call center applications, they include banking applications for processing loan requests and identifying potential fraud, legal applications for identifying relevant case precedents, investment applications for developing buy/sell predictions and recommendations, manufacturing applications for scheduling equipment maintenance, and pharmaceutical R&D applications for predicting the success of drugs under development.

Not surprisingly, most business and technology leaders are optimistic about ECC's value-creating potential. In a 2017 survey of 3,000 senior executives across industries, company sizes, and countries, 63% said that ECC applications would have a large effect on their organization's offerings within five years.[4] However, the actual rate of adoption is low, and benefits have proved elusive for most organizations. In 2017, when we conducted our own survey of senior executives at 106 companies, half of the respondents reported that their company had no ECC applications in place. Moreover, only half of the respondents whose companies had applications believed they had produced measurable business outcomes. Other studies report similar results.[5]

This suggests that generating value from ECC applications is not easy — and that reality has caught many business leaders off guard. Indeed, we found that some of the excitement around ECC resulted from unrealistic expectations about the powers of "intelligent machines." In addition, we observed that many companies that hoped to benefit from ECC but failed to do so had not developed the necessary organizational capabilities. To help address that problem, we undertook a program of research aimed at identifying the foundations of ECC competence. We found five capabilities and four practices that companies need to splice the ECC gene into their organization's DNA.

Five Crucial Capabilities

We found that companies that successfully create value (that is, radically improve business processes to reduce costs and/or generate new revenues) using ECC applications possess five capabilities: data science competence, business domain proficiency, enterprise architecture expertise, an operational IT backbone, and digital inquisitiveness.

Data science competence. Data science competence encompasses a wide range of skills essential to ECC. It involves ensuring the availability and usefulness of massive amounts of data: collecting, cleaning, curating, tagging, and analyzing internal and external data from multiple sources. Such competence also entails identifying and describing relationships between data, as well as developing AI algorithms that have learned from data how to identify patterns and probabilities.

Top-notch data scientists have extensive knowledge in areas such as natural language processing, statistical inference, knowledge representation, and learning algorithms. Wipro, the Indian IT services company, includes these areas among the pillars of its data science expertise. Its data scientists deploy their skills and a variety of tools to create AI algorithms that can be inserted into enterprise applications.

For organizations that cannot develop the talent internally, obtaining data science competence is expensive and can require multiple hires from, for example, software development companies, technology consulting companies, AI startups, or university graduate programs in related fields. At a financial services company we studied — we call it OneBankAssure — the CEO hired a new direct report who was a technically accomplished data science academic and consultant. This person, in turn, hired the 20 data scientists who became the core ECC development team. Companies that are serious about ECC spend the money to hire the right data science talent. To raise the money, one pharmaceutical company we studied reduced its operational IT costs (by eliminating duplication in systems and standardizing processes across its business units) and redirected the savings to the acquisition of data science skills.

Business domain proficiency. Domain proficiency is needed to understand the tasks, workflows, and logic of existing business processes, as well as to imagine how ECC applications could improve them. As many organizations have learned the hard way, it's possible — even easy — to develop an elegant AI algorithm that uses massive amounts of data to learn how to predict or categorize something but doesn't improve the business. Having the right technical skills isn't enough. Domain proficiency links data science competence to business value.

ABOUT THE RESEARCH

The research activities on which this article is based were undertaken between January 2016 and December 2017, and covered companies across industries in North America, Europe, Asia, and Australia. We interviewed senior executives in IT and innovation units in 33 companies, as well as industry and technical experts in eight enterprise cognitive computing developer/vendor organizations, regarding ECC uptake in a range of organizations and industries, and ECC challenges and opportunities. We studied 51 ECC use cases (37% deployed; 48% in ideation or design stages; and 15% abandoned prior to development). We surveyed senior IT and technology leaders in 106 companies about ECC applications in place, application development and management issues, and outcomes. Finally, we researched and prepared three in-depth case studies on three organizations, for which we interviewed 35 people: C-level officers; functional leaders in IT, marketing, sales, and strategy; and data science and domain/process experts.[i]

For example, the ability of data scientists to effectively curate, tag, and analyze data depends on a clear understanding of the relationships among the data from a process and business point of view. Domain proficiency provides clarity around those relationships, referred to as ontologies. Data ontologies can become quite complex and even counterintuitive. Here's how a domain expert at a pharmaceutical company described some of the complexities he encountered capturing the data ontologies needed to support the company's research on diabetes: "A big part of diabetes is being overweight. Should there be an obesity dimension in our ontology of diabetes? Or is diabetes an attribute of obesity? Oh, and people who are overweight often have joint replacement issues. If they're overweight and their joints hurt and they have diabetes, the incidence of depression is very high, and dealing with depression is an important part of generating outcomes. Do I train the algorithm on depression?"

Domain proficiency is also important for creating the business rules that shape how the outputs from the algorithm are handled by the ECC application. For example, an ECC application that helps banks predict which customers are most likely to repay loans on time must include business rules for how the algorithm's prediction will be applied, such as: Will some loans be granted automatically? If so, under what conditions? With whom will the predictions be shared? Under what circumstances can a prediction be overridden?

For any given ECC application, domain proficiency is needed in all the functional areas that have a bearing on — or are stakeholders in — the operations of the focal process. For example, a team at a U.S. bank that developed an ECC application to detect financial fraud needed proficiency not only in fraud identification and prevention but also in the related areas of regulatory compliance and banking law.

People with domain proficiency have deep process knowledge. They may be process owners, although they are often people with a regular hands-on role. Some companies seek to hire data scientists with domain expertise. Indeed, such individuals can partner well with business domain experts, but they cannot substitute for them when an ECC application is being developed. That's because they usually lack enterprise-specific knowledge about processes, policies, and practices currently in play.

Enterprise architecture expertise. Implementations of enterprise systems have a history of disappointing leaders who underestimated the organizational changes needed to capture their value. Too many leaders are reliving this disappointment with ECC applications. ECC applications do not deliver value by simply processing data and delivering outputs. They deliver value when the organization changes its behavior — that is, when it changes processes, policies, and practices — to gain and apply the insights from those outputs. Experts in enterprise architecture design the new organization needed to create business value from ECC applications, and they help manage the transition from the old organization to the new one.

The most ambitious ECC applications usually affect several, often fundamentally different business processes. In such cases, enterprise architects are needed to orchestrate the redesign of the systems, processes, and roles across organizational units. The more ambitious the ECC application, the more likely it will require far-reaching organizational changes.

Organization design and change issues can surface for seemingly small-scale applications. One medical drug distributor failed to recoup its investment in an ECC application that could accurately predict whether an online customer's insurance would cover a claim 90% of the time because the accounts payable department balked at making costly process changes required to support the application. If an enterprise architect had been engaged in the project from the outset, this loss might have been avoided.

The organizational changes needed to unlock the potential of an ECC application can be complex and intertwined. Enterprise architects are familiar with the organizational roadblocks that drive up costs or limit impact. At Wipro, enterprise architecture expertise helped smooth the way for a new help desk ECC application by first merging the company's existing help desk applications, reducing the types of fault tickets from 3,000 to 2,200, and eliminating redundancies in support tasks. By simplifying and standardizing the help desk process prior to the development of the ECC application, the company reduced and simplified the work of getting the data needed to train the AI algorithm, developing it, and ultimately automating the process, thus unlocking additional value.

Enterprise architects also recognize when ECC applications require changes in employees' jobs. They may see the need for upskilling, re-skilling, or the creation of entirely new roles. When a seemingly simple sales-lead-generating ECC application required its agents to do more cold calling and make more targeted pitches, OneBankAssure's enterprise architects designed a new coaching role to help agents that proved essential to generating benefits from the application.

Given the breadth of skills that enterprise architects draw on, this expertise can be difficult to develop. It often resides in people who are steeped in organization design and change management such as business leaders with experience managing technology-driven transformations or other reorganizations. Human resource professionals with exposure to a broad range of organizational roles can be a good source of architectural expertise in role design and redesign, as well as skills training. IT professionals with exposure to many different business processes, who can help streamline processes and establish the proper division of work between ECC applications and employees, also can be tapped.

Operational IT backbone. A company's existing technology and data foundation — its operational IT backbone — and the people responsible for it support the development and running of ECC applications. They supply the IT capabilities needed to store and access critical data, integrate ECC applications with other applications, provide reliable operations, and ensure privacy and security.

As noted earlier, for an AI algorithm to learn from data, a company must make available massive amounts of high-quality data that is cleaned and tagged. The lack of high-quality data is the most pernicious and least anticipated obstacle in the development of AI algorithms. OneBankAssure overcame this obstacle and accelerated its adoption of ECC by separating responsibilities for developing AI algorithms from responsibilities for providing the data. Since the IT unit already maintained the underlying operational infrastructure and good-quality operational data, it was able to support the algorithm developers by providing them with access to a data lake containing operational and external data. Responsibility for structuring the data for developing algorithms rested with the data scientists.

Almost no new enterprise application can operate in isolation from other enterprise applications. ECC is no exception. If an application is not

Enterprise architects recognize when ECC applications require changes in employees' jobs. They may see the need for upskilling, re-skilling, or the creation of entirely new roles.

Often, users of ECC applications must apply human judgment to predictions made by algorithms. They need to possess digital inquisitiveness — the inclination to question and evaluate the data before them.

properly integrated, it will be hard to use and possibly ignored. That's why the IT unit at OneBankAssure embedded the company's new sales lead system into its customer relationship manager system, which was part of its operational IT backbone. The CRM linked up-to-date contact information and customer history data to the sales leads. It also provided a set of processes within which the ECC sales leads could be seamlessly presented to users. Being part of the IT backbone also meant that the sales lead system would be scalable, reliable, and secure.

The existing IT staff is the logical source for operational IT backbone expertise. At OneBankAssure, the IT function set the standards for ECC plug-ins and adapted applications to the company's production environment by refactoring and retesting the code. It also managed disaster recovery and security for installed ECC applications.

Digital inquisitiveness. The AI algorithms in ECC applications do not produce definitive answers. Rather, they produce predictions based on probabilities: the probability that a customer will buy a product, that a patient has a disease, that a loan will be repaid. Often, application users must consider these predictions and apply human judgment to arrive at decisions about how and where to promote offerings, what treatments to prescribe, or what loans to approve. To do this effectively, they need to possess digital inquisitiveness — a habitual inclination to question and evaluate the data before them. They must use that skill to better understand the options provided by ECC applications and continually improve outcomes.

The development of this capability requires a broad-based effort. A number of companies we studied instituted mechanisms to cultivate digital inquisitiveness. OneBankAssure's corporate university delivered a training program that introduced executives to the idea of using data effectively in making decisions. One exercise incorporated a strategy game, in which participants vied to develop the highest-value ECC solution to a business problem. At different phases of the game, they had to deal with poor-quality data (in fact, real company data), build decision trees, teach an algorithm to detect patterns, and develop a model to solve a problem. Wipro created an e-learning platform on which employees were able to take courses to understand what AI was, how it could be used in business processes, and how to work effectively in ECC-enabled processes. The company also trained hundreds of domain experts to act as AI champions throughout the organization.

Four Key Practices

Developing the five capabilities equips organizations to derive value from ECC applications, but then companies must apply those capabilities. We've found that four practices in particular help them do that, creating the conditions for a given application — and its underlying AI algorithm — to deliver on its promise.

Develop clear, realistic use cases. A use case provides a clear definition of what an ECC application will do and illustrates how its AI algorithms will enhance the execution and outcomes of a business process or set of processes. It shows how work will be divided between an application and a user. In doing so, a use case establishes the need for process changes and provides initial insights into any new capabilities users will need (as well as any skills that will no longer be needed). A well-designed use case also facilitates the estimation of the costs and benefits of the ECC application.

Consider an ECC application in a call center: Its use case might include a simple version of an AI algorithm that matches customer queries to resolutions. It would show what the algorithm would do and what automated resolutions the

> Because of their business focus, domain experts are often more persuasive ECC champions than are data scientists and IT professionals, who may be perceived as overly enamored with AI.

application could provide to customers. It would also show that some queries would be passed along to call center representatives. Required work and capability changes for call center employees could be inferred as well. All that information would allow a domain expert to roughly gauge the challenges of adoption and estimate intended benefits in terms of reduced response time, reduced labor, fewer follow-up calls, greater customer satisfaction, or a combination of outcomes.

Developing an ECC use case that is grounded in reality is a team activity. It is primarily the responsibility of domain experts and data scientists, who specify how an AI algorithm will enhance organizational outcomes and what data is needed to create it. But enterprise architects weigh in, too, identifying any new structures, roles, and systems required by a proposed ECC application, especially those affected indirectly by the new application. IT experts assess the need for integration with other applications and identify any additional IT support the application might require.

Properly developed use cases can help companies avoid sloppy or ill-considered ECC implementations that waste resources and may limit enthusiasm for — and effective implementation of — ECC. In fact, if the use cases for early ECC applications highlight quick wins for high-profile issues, they can be a powerful driver of organizational uptake of ECC. Bench scientists at one pharmaceutical company suggested developing an ECC application that could mine patent data for a specific disease knowing that if it was successful, the application itself would serve as a use case for similar applications for other diseases — and it did. Sometimes the algorithms themselves can be substantially reused. Wipro developed a use case for new-customer verification in the financial services sector, in which the AI algorithm automated the extraction and interpretation of information from customers' financial documents. This use case gave rise to an ECC application in the engineering sector that extracted and interpreted information from digitized blueprints.

Manage ECC application learning. AI algorithms in products such as smartphones use the data they process to improve themselves without human intervention. In contrast, ECC applications have a much more complex feedback loop. Business conditions and demands change constantly. As a result, the data used to create an AI algorithm becomes a less accurate reflection of reality over time — the algorithm *drifts*. It thus becomes necessary to manage the learning of the ECC applications throughout their life cycles.

Algorithm drift may occur quickly, as in predicting the sales of fashion apparel, or slowly, as in predicting the presence of a disease. To manage drift and keep ECC applications up-to-date, companies usually rely on a combination of IT backbone capabilities, data science competence, and domain proficiency. They build reporting mechanisms into ECC applications that generate alerts if the business results derived from the application's outputs are no longer aligned with the organization's goals, the algorithm's recommendations aren't within preestablished error ranges, or the application isn't running properly.

When deviations occur, AI algorithms need to be retrained and ECC applications relaunched. Domain experts and data scientists need to work together to identify, access, clean, tag, and architect new sources of data to improve the accuracy of AI algorithms and the utility of ECC applications. In addition, as the performance of the algorithm is better understood or as users become more proficient with the application, new business rules or processes that can enhance the value of the application may be required.

At OneBankAssure, domain experts and data scientists identified new external sources of data that could

help identify productive sales leads, so they retrained their AI algorithm. They also learned that agent experience affected sales success, so they developed more elaborate rules to govern how the ECC application presented leads. The new data and business rules led to a richer, more complex ECC application that OneBankAssure continues to enhance.

Cocreate throughout the application life cycle. A data scientist or business domain expert cannot develop and sustain an ECC application in isolation. Interviewees in companies that effectively exploited AI repeatedly told us that they had, at first, badly underestimated the intense level of interdisciplinary cocreation needed to achieve success with ECC. They said they began to make progress only when they realized that ECC applications require people from disparate specialties and disciplines to work as a single team, not just during initial development and implementation but also in ongoing development throughout the application life cycle.

One reason cocreation is important for ECC is because business experts do not yet understand what AI can and can't do. During the development of ECC, close and sustained collaborative relationships across diverse areas of expertise can ameliorate this problem. At OneBankAssure, as a matter of hard-learned policy, every ECC application is created by a team of process owners and users with domain expertise, enterprise architects, and data scientists, with added assistance from the IT function. There are few handoffs within the team. No team member ever works completely alone, and in the end, no one team member is responsible for success or failure. The interaction of team members results in a shared vocabulary about the business need and potential solutions, enabling them to better visualize and make sense of how people will actually use the application.

During implementation, owners of the IT operational backbone get involved not with a single handoff but rather by working with the ECC application team to cocreate a solution for integrating the application with the backbone upon production. After implementation, responsibility for maintaining and sustaining ECC applications continues to be highly interdependent in nature, as described above.

Think "cognitive." Companies that successfully develop and use ECC applications champion the uptake of AI and create positive buzz and excitement around its use. They encourage employees to generate ideas for new ECC applications that can improve their own work.

The employee response to ECC varies widely. Some people do not see the potential of ECC at first. Others have exaggerated expectations, thinking that ECC applications will automatically solve difficult business problems. Still others do not trust AI and see risks to ECC-enabled business processes, such as rogue behavior in AI algorithms and capability or job losses.

Domain experts who have seen what AI can do are the best stewards of realistic and credible conversations about ECC within their companies. Because of their business focus, they are more likely to be able to create a positive buzz around ECC than are data scientists and IT professionals, who

LEARNING THROUGH DOING

As companies apply their enterprise cognitive computing capabilities through the four key practices, they're also enriching their capabilities. Practices are, after all, opportunities to practice.

The pharmaceutical company we studied offers a good example. Recognizing that data science and ECC applications would become increasingly important to curing and preventing disease, the company hired data scientists to conduct workshops that would help senior staff (mainly business domain experts and enterprise architects) imagine the possibilities. They worked with business leaders to identify information-processing bottlenecks that created backlogs in drug discovery, clinical trials, manufacturing, and commercialization. The bottlenecks highlight opportunities for AI applications that could solve problems for small groups of analysts and decision makers in the organization.

These early efforts generated incremental business value, but the business leaders were far more focused on building capabilities than on building game-changing applications. They carefully chose use cases to meet the needs of people who naturally think "cognitive" and then engaged all the needed expertise — data scientists, domain experts, and IT specialists — to cocreate and manage the applications. In those pockets of the company, people deepened their understanding of organizational impacts and developed the capabilities to identify and pursue more ambitious ECC applications. What's more, the gains they made in efficiency and productivity inspired others in the company to seek out their own use cases and build their own capabilities. Creating this virtuous cycle of continuous organizational learning has mitigated the risks of the company's AI investments and positioned the company to make ECC a competitive advantage.

may be perceived as overly enamored with AI. Indeed, at Wipro, domain experts were enlisted as AI champions — conducting "walkabouts" in their various departments, evangelizing ECC, and listening to ideas put forth by their colleagues.

The most likely sources of ideas for new ECC applications are people with domain proficiency or data science competence (or both). At OneBankAssure, operational managers spent several months in discussions with data science professionals to envision how their business might be affected by AI in the future, to develop ideas for new ECC applications, and to draft road maps for how their ideas could be developed and commercialized.

Proactive data science leaders also can be effective idea generators. At a pharmaceutical company we studied, one ECC project got its start at a lunch in which a business leader told a data scientist about a business problem, and the data scientist proposed a simple solution leveraging an already developed AI algorithm. In another company, the head of the data science unit organized seminars for functional and business leaders to identify areas in which ECC applications could best serve them.

The digital inquisitiveness of the entire workforce should be harnessed, too. Wipro, for example, crowdsources ideas from employees. It encourages them to envision and suggest new ECC applications, evaluating the ideas for their potential contribution to top-line growth, bottom-line profits, customer satisfaction, or employee satisfaction.

BUSINESS APPLICATIONS OF AI may not create the same buzz as a self-driving car, but they can generate handsome returns — dramatic improvements in performance, profitability, revenues, and customer satisfaction. By cultivating the five capabilities and applying the four practices described in this article, business leaders can splice the ECC gene into their organizational DNA and set themselves up to reap those rewards.

It's a virtuous cycle: The capabilities enable employees to execute the practices, and the practices themselves exercise and strengthen the capabilities. This cycle helps companies become evermore adept at developing and using ECC applications that improve operations and create business value.

Monideepa Tarafdar *is a professor of information systems and codirector of the Centre for Technological Futures at Lancaster University Management School in the United Kingdom. Her research focuses on how digital technologies transform organizations and societies.* ***Cynthia M. Beath*** *is professor emerita at the University of Texas McCombs School of Business in Austin. Her research focuses on how organizations get value from investments in IT.* ***Jeanne W. Ross*** *is a principal research scientist at MIT Sloan's Center for Information Systems Research in Cambridge, Massachusetts. She writes a quarterly column for* MIT SMR *about digital business management issues. Comment on this article at http://sloanreview.mit.edu/x/60402.*

REFERENCES

1. ECC applications are distinct from other kinds of enterprise software in that AI tools, rather than human deduction, are used to figure out what logic will optimize business outcomes. AI software tools apply computational and analytical techniques, such as neural network analysis, machine learning, and Bayesian statistics, to large sets of structured and unstructured data to create AI algorithms that will classify, cluster, predict, and match patterns. These algorithms become part of the logic of the ECC application.

2. J. Bughin and E. Hazan, "Five Management Strategies for Getting the Most From AI," MIT Sloan Management Review, Sept. 19, 2017, www.sloanreview.mit.edu.

3. M. Tarafdar, C.M. Beath, and J.W. Ross, "Enterprise Cognitive Computing Applications: Opportunities and Challenges," IEEE IT Professional 19, no. 4 (August 2017): 21-27.

4. S. Ransbotham, D. Kiron, P. Gerbert, et al., "Reshaping Business With Artificial Intelligence," MIT Sloan Management Review research report, Sept. 6, 2017.

5. S. Norton, "Machine Learning at Scale Remains Elusive for Many Firms," The Wall Street Journal, April 27, 2018; J. Bughin and E. Hazan, "Five Management Strategies"; and Ransbotham, et al., "Reshaping Business."

i. C.M. Beath, M. Tarafdar, and J.W. Ross, "OneBankAssure: Customer Intimacy Through Machine Learning," working paper, MIT Center for Information Systems Research, Cambridge, MA, March 12, 2018; M. Tarafdar and C.M. Beath, "Wipro Limited: Developing a Cognitive DNA," working paper, MIT Center for Information Systems Research, Cambridge, MA, April 27, 2018; and J.W. Ross, K. Moloney, and C.M. Beath, "Pharmco: Becoming a Data-Science Driven Company," working paper, MIT Center for Information Systems Research, Cambridge, MA, Feb. 21, 2019.

Reprint 60402.